

# Quartus II Handbook Volume 1: Design and Synthesis



Subscribe



Send Feedback

**QII5V1**  
2014.12.15

101 Innovation Drive  
San Jose, CA 95134  
[www.altera.com](http://www.altera.com)



2014.12.15

QII5V1



Subscribe



Send Feedback

The Quartus II software organizes and manages the elements of your design within a *project*. The project encapsulates information about your design hierarchy, libraries, constraints, and project settings. Click **File > New Project Wizard** to quickly create a new project and specify basic project settings

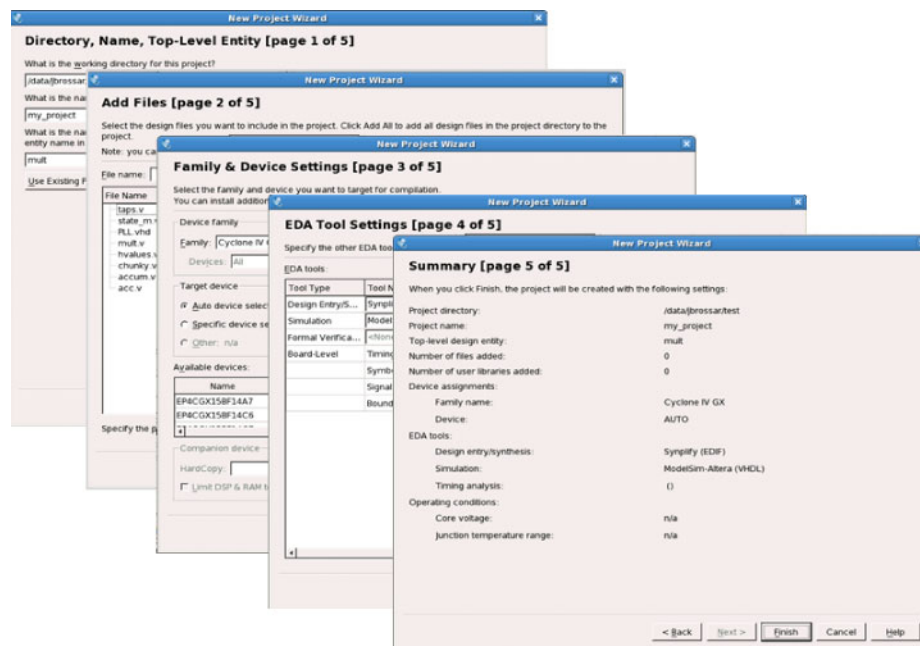
When you open a project, a unified GUI displays integrated project information. The Project Navigator allows you to view and edit the elements of your project. The Messages window lists important information about project processing.

You can save multiple revisions of your project to experiment with settings that achieve your design goals. Quartus II projects support team-based, distributed work flows and a scripting interface.

## Quick Start

To quickly create a project and specify basic settings, click **File > New Project Wizard**.

Figure 1-1: New Project Wizard



© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



## Understanding Quartus II Projects

A single Quartus II Project File (**.qpf**) represents each project. The text-based **.qpf** references the Quartus II Settings File (**.qsf**), that lists all project files and stores project and entity settings. When you make project changes in the GUI, these text files automatically store the changes. The GUI helps to manage:

- Design, EDA, IP core, and Qsys system files
- Project settings and constraint files
- Project archive and migration files

**Table 1-1: Quartus II Project Files**

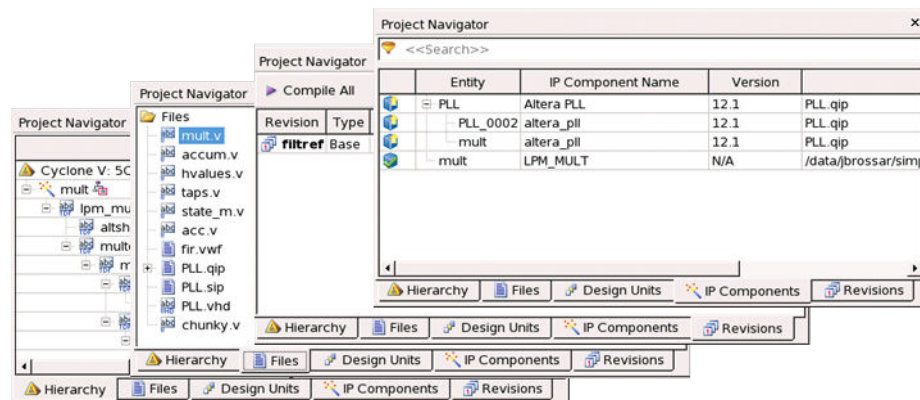
File Type	Contains	To Edit	Format
Project file	Project and revision name	<b>File &gt; New Project Wizard</b>	Quartus II Project File ( <b>.qpf</b> )
Project settings	Lists design files, entity settings, target device, synthesis directives, placement constraints	<b>Assignments &gt; Settings</b>	Quartus II Settings File ( <b>.qsf</b> )
Project database	Compilation results	<b>Project &gt; Export Database</b>	Exported Partition File ( <b>.qxp</b> )
Timing constraints	Clock properties, exceptions, setup/hold	<b>Tools &gt; TimeQuest Timing Analyzer</b>	Synopsys Design Constraints File ( <b>.sdc</b> )
Logic design files	RTL and other design source files	<b>File &gt; New</b>	All supported HDL files
Program- ming files	Device programming image and information	<b>Tools &gt; Programmer</b>	SRAM Object File ( <b>.sof</b> ) Programmer Object File ( <b>.pof</b> )
Project library	Project and global library information	<b>Tools &gt; Options &gt; Libraries</b>	<b>.qsf</b> (project) <b>quartus2.ini</b> (global)
IP core files	IP core logic, synthesis, and simulation information	<b>Tools &gt; IP Catalog</b>	All supported HDL files Quartus II IP File ( <b>.qip</b> )
Qsys system files	Qsys system and IP core files	<b>Tools &gt; Qsys</b>	Qsys System File ( <b>.qsys</b> )
EDA tool files	Generated for third-party EDA tools	<b>Tools &gt; Options &gt; EDA Tool Options</b>	Verilog Output File ( <b>.vo</b> ) VHDL Output File ( <b>.vho</b> ) Verilog Quartus Mapping File ( <b>.vqm</b> )
Archive files	Complete project as single compressed file	<b>Project &gt; Archive Project</b>	Quartus II Archive File ( <b>.qar</b> )

## Viewing Basic Project Information

View basic information about your project in the Project Navigator, Report panel, and Messages window. View project elements in the Project Navigator ( **View > Utility Windows > Project Navigator**). The Project Navigator displays key project information, including design files, IP components, and revisions of your project. Use the Project Navigator to:

- View and modify the design hierarchy (**right-click > Set as Top-Level Entity**)
- Set the project revision (**right-click > Set Current Revision**)
- View and update logic design files and constraint files (**right-click > Open**)
- Update IP component version information (**right-click > Upgrade IP Component**)

Figure 1-2: Project Navigator Hierarchy, Files, Revisions, and IP



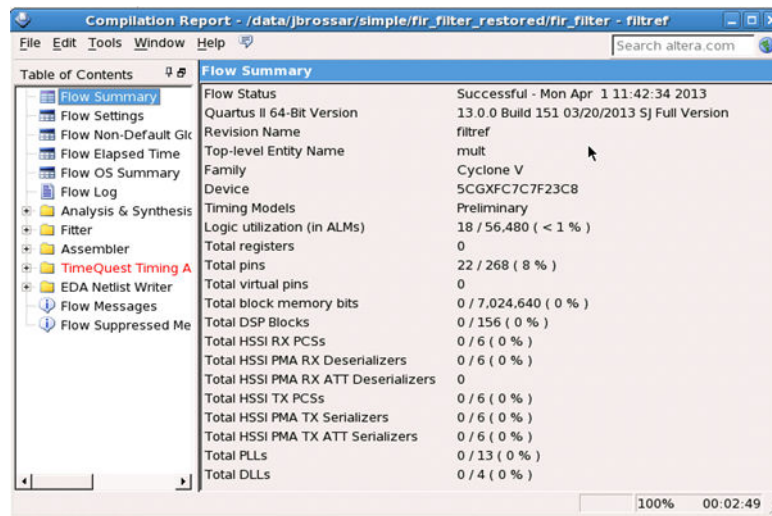
## Viewing Project Reports

The Report panel (**Processing > Compilation Report**) displays detailed reports after project processing, including the following:

- Analysis & Synthesis reports
- Fitter reports
- Timing analysis reports
- Power analysis reports
- Signal integrity reports

Analyze the detailed project information in these reports to determine correct implementation. Right-click report data to locate and edit the source in project files.

Figure 1-3: Report Panel



### Related Information

#### List of Compilation Reports

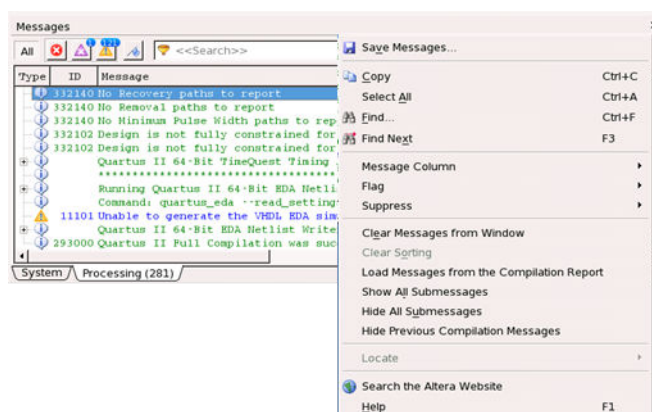
## Viewing Project Messages

The Messages window (**View > Utility Windows > Messages**) displays information, warning, and error messages about Quartus II processes. Right-click messages to locate the source or get message help.

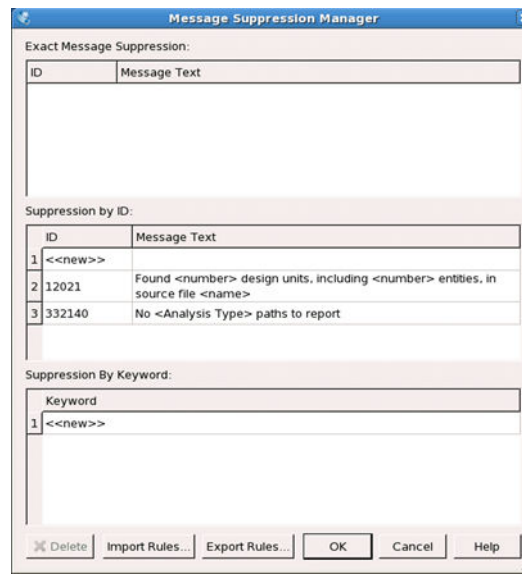
- **Processing** tab—displays messages from the most recent process
- **System** tab—displays messages unrelated to design processing
- **Search**—locates specific messages

Messages are written to `stdout` when you use command-line executables.

Figure 1-4: Messages Window



You can suppress display of unimportant messages so they do not obscure valid messages.

**Figure 1-5: Message Suppression by Message ID Number**

## Suppressing Messages

To suppress messages, right-click a message and choose any of the following:

- **Suppress Message**—suppresses all messages matching exact text
- **Suppress Messages with Matching ID**—suppresses all messages matching the message ID number, ignoring variables
- **Suppress Messages with Matching Keyword**—suppresses all messages matching keyword or hierarchy

## Message Suppression Guidelines

- You cannot suppress error or Altera legal agreement messages.
- Suppressing a message also suppresses any submessages.
- Message suppression is revision-specific. Derivative revisions inherit any suppression.
- You cannot edit messages or suppression rules during compilation.

## Managing Project Settings

The New Project Wizard helps you initially assign basic project settings. Optimizing project settings enables the Compiler to generate programming files that meet or exceed your specifications.

The **.qsf** stores each revision's project settings.

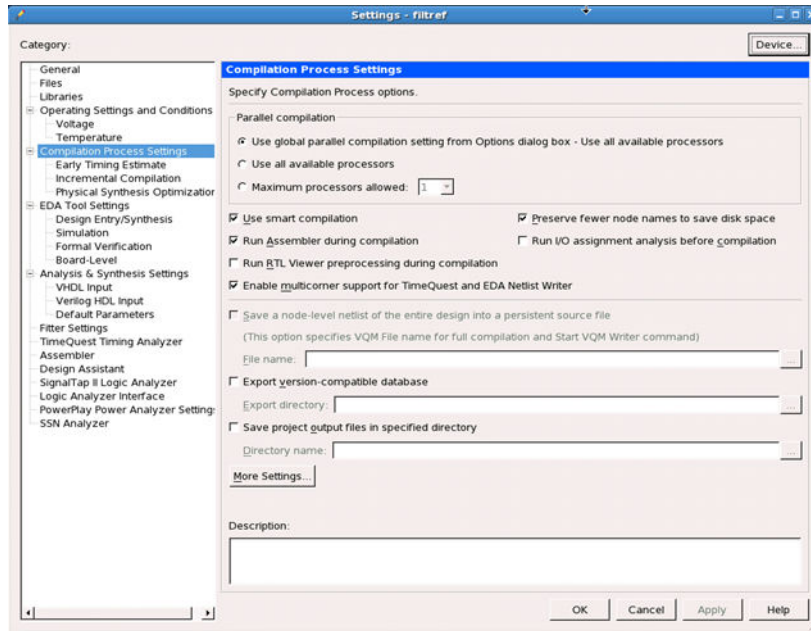
Click **Assignments > Settings** to access global project settings, including:

- Project files list
- Synthesis directives and constraints
- Logic options and compiler effort levels
- Placement constraints

- Timing constraint files
- Operating temperature limits and conditions
- File generation for other EDA tools
- Target device (click **Assignments** > **Device**)

The Quartus II Default Settings File (<revision name>\_assignment\_defaults.qdf) stores initial settings and constraints for each new project revision.

Figure 1-6: Settings Dialog Box for Global Project Settings



The Assignment Editor (**Tools** > **Assignment Editor**) provides a spreadsheet-like interface for assigning all instance-specific settings and constraints.

Figure 1-7: Assignment Editor Spreadsheet

The screenshot shows the 'Assignment Editor' window with a spreadsheet-like interface. The table has columns for 'From', 'To', 'Assignment Name', 'Value', 'Enabled', 'Entity', and 'Comment'. The 'From' column contains 'ab' and '<<new>>'. The 'To' column contains various instance names like 'follow', 'yn\_out[7]', 'yn\_out[6]', 'yn\_out[5]', 'yn\_out[4]', 'yn\_out[3]', 'yn\_out[2]', 'yn\_out[1]', 'yn\_out[0]', 'yvalid', and '\*PLL...\*'. The 'Assignment Name' column contains 'Current Strength', 'PLL Compensation Mode', 'PLL Automatic Self-Reset', and 'PLL Bandwidth Preset'. The 'Value' column contains 'Minimum Current', 'Normal', 'Off', and 'Auto'. The 'Enabled' column contains 'Yes' for all entries. The 'Entity' column contains 'mult' for all entries. The 'Comment' column is empty. The status bar at the bottom shows '0%' and '00:00:00'.

From	To	Assignment Name	Value	Enabled	Entity	Comment
ab	follow	Current Strength	Minimum Current	Yes	mult	
	yn_out[7]	Current Strength	Minimum Current	Yes	mult	
	yn_out[6]	Current Strength	Minimum Current	Yes	mult	
	yn_out[5]	Current Strength	Minimum Current	Yes	mult	
	yn_out[4]	Current Strength	Minimum Current	Yes	mult	
	yn_out[3]	Current Strength	Minimum Current	Yes	mult	
	yn_out[2]	Current Strength	Minimum Current	Yes	mult	
	yn_out[1]	Current Strength	Minimum Current	Yes	mult	
	yn_out[0]	Current Strength	Minimum Current	Yes	mult	
	yvalid	Current Strength	Minimum Current	Yes	mult	
	*PLL...*	PLL Compensation Mode	Normal	Yes	mult	
	*PLL...*	PLL Automatic Self-Reset	Off	Yes	mult	
	*PLL...*	PLL Bandwidth Preset	Auto	Yes	mult	
<<new>>	<<new>>	<<new>>				

## Optimizing Project Settings

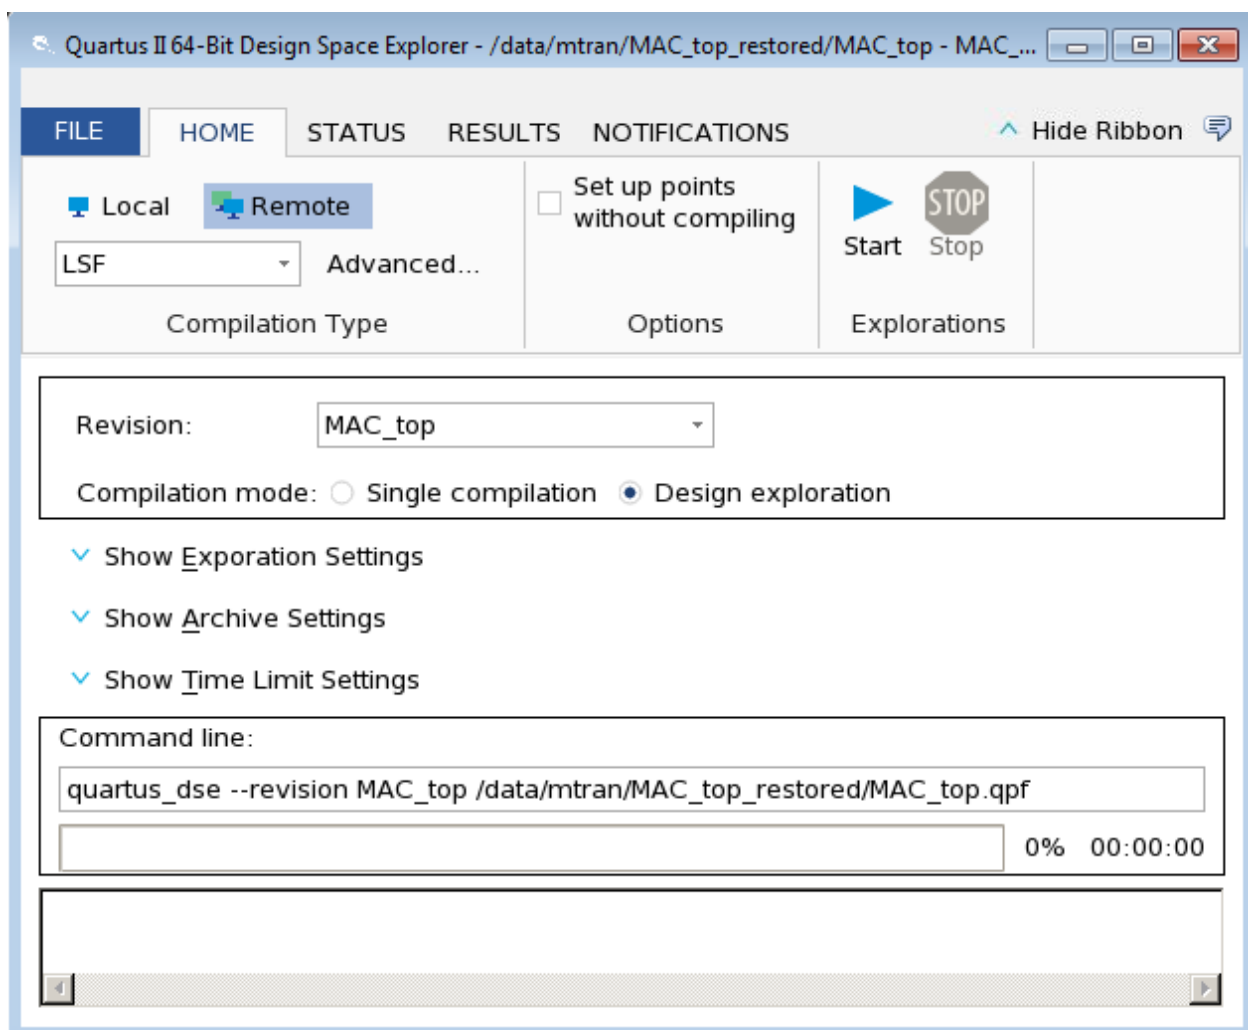
Optimize project settings to meet your design goals. The Quartus II Design Space Explorer II iteratively compiles your project with various setting combinations to find the optimal setting for your goals. Alternatively, you can create a project revision or project copy to manually compare various project settings and design combinations.

### Optimizing with Design Space Explorer

Use Design Space Explorer II (**Tools > Launch Design Space Explorer**) to find optimal project settings for resource, performance, or power optimization goals. Design Space Explorer II (DSE) processes your design using various setting and constraint combinations, and reports the best settings for your design.

DSE II attempts multiple seeds to identify one meeting your requirements. DSE II can run different compilations on multiple computers in parallel to streamline timing closure.

Figure 1-8: Design Space Explorer II



### Optimizing with Project Revisions

You can save multiple, named project revisions within your Quartus II project (**Project > Revisions**).



Each revision captures a unique set of project settings and constraints, but does not capture any logic design file changes. Use revisions to experiment with different settings while preserving the original. You can compare revisions to determine the best combination, or optimize different revisions for various applications. Use revisions for the following:

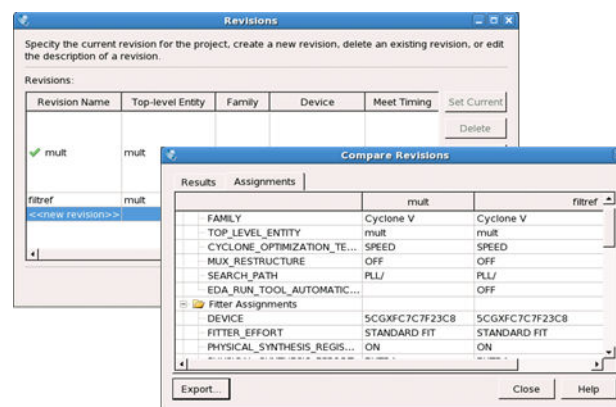
- Create a unique revision to optimize a design for different criteria, such as by area in one revision and by  $f_{MAX}$  in another revision.
- When you create a new revision the default Quartus II settings initially apply.
- Create a revision of a revision to experiment with settings and constraints. The child revision includes all the assignments and settings of the parent revision.

You create, delete, specify current, and compare revisions in the **Revisions** dialog box. Each time you create a new project revision, the Quartus II software creates a new **.qsf** using the revision name.

To compare each revision's synthesis, fitting, and timing analysis results side-by-side, click **Project > Revisions** and then click **Compare**.

In addition to viewing the compilation results of each revision, you can also compare the assignments for each revision. This comparison reveals how different optimization options affect your design.

**Figure 1-9: Comparing Project Revisions**



## Copying Your Project

Click **Project > Copy Project** to create a separate copy of your project, rather than just a revision within the same project.

The project copy includes all design files, **.qsf(s)**, and project revisions. Use this technique to optimize project copies for different applications. For example, optimize one project to interface with a 32-bit data bus, and optimize a project copy to interface with a 64-bit data bus.


## Managing Logic Design Files

The Quartus II software helps you create and manage the logic design files in your project. Logic design files contain the logic that implements your design. When you add a logic design file to the project, the Compiler automatically compiles that file as part of the project. The Compiler synthesizes your logic design files to generate programming files for your target device.

The Quartus II software includes full-featured schematic and text editors, as well as HDL templates to accelerate your design work. The Quartus II software supports VHDL Design Files (**.vhd**), Verilog HDL Design Files (**.v**), SystemVerilog (**.sv**) and schematic Block Design Files (**.bdf**). The Quartus II software also supports Verilog Quartus Mapping (**.vqm**) design files generated by other design entry and synthesis tools. In addition, you can combine your logic design files with Altera and third-party IP core design files, including combining components into a Qsys system (**.qsys**).

The New Project Wizard prompts you to identify logic design files. Add or remove project files by clicking **Project > Add/Remove Files in Project**. View the project's logic design files in the Project Navigator.

Figure 1-10: Design and IP Files in Project Navigator



Right-click files in the Project Navigator to:

- **Open** and edit the file
- **Remove File from Project**
- **Set as Top-Level Entity** for the project revision
- **Create a Symbol File for Current File** for display in schematic editors
- Edit file **Properties**

## Including Design Libraries

You can include design files libraries in your project. Specify libraries for a single project, or for all Quartus II projects. The **.qsf** stores project library information.

The **quartus2.ini** file stores global library information.

### Related Information

[Design Library Migration Guidelines](#) on page 1-38

## Specifying Design Libraries

To specify project libraries from the GUI:

1. Click **Assignment > Settings**.
2. Click **Libraries** and specify the **Project Library name** or **Global Library name**. Alternatively, you can specify project libraries with `SEARCH_PATH` in the **.qsf**, and global libraries in the **quartus2.ini** file.

### Related Information

- [Recommended Design Practices](#) on page 11-1

- [Recommended HDL Coding Styles](#) on page 12-1

## Managing Timing Constraints

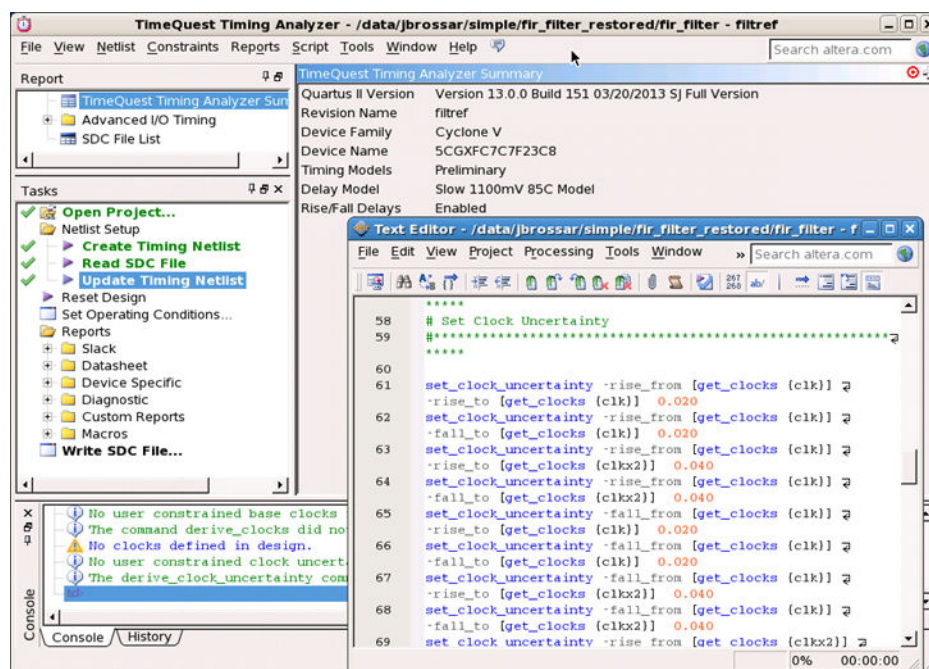
View basic information about your project in the Project Navigator, Report panel, and Messages window.

Apply appropriate timing constraints to correctly optimize fitting and analyze timing for your design. The Fitter optimizes the placement of logic in the device to meet your specified timing and routing constraints.

Specify timing constraints in the TimeQuest Timing Analyzer (**Tools > TimeQuest Timing Analyzer**), or in an **.sdc** file. Specify constraints for clock characteristics, timing exceptions, and external signal setup and hold times before running analysis. TimeQuest reports the detailed information about the performance of your design compared with constraints in the Compilation Report panel.

Save the constraints you specify in the GUI in an industry-standard Synopsys Design Constraints File (**.sdc**). You can subsequently edit the text-based **.sdc** file directly.

Figure 1-11: TimeQuest Timing Analyzer and SDC Syntax Example



### Related Information

[Quartus II TimeQuest Timing Analyzer](#)

## Introduction to Altera IP Cores

Altera® and strategic IP partners offer a broad portfolio of off-the-shelf, configurable IP cores optimized for Altera devices. The Altera Complete Design Suite (ACDS) installation includes the Altera IP library.

The OpenCore and OpenCore Plus IP evaluation features enable fast acquisition, evaluation, and hardware testing of Altera IP cores.

You can integrate optimized and verified IP cores into your design to shorten design cycles and maximize performance. The Quartus II® software also supports IP cores from other sources. Use the IP Catalog to efficiently parameterize and generate a custom IP variation for instantiation in your design.

The Altera IP library includes the following IP core types:

- Basic functions
- DSP functions
- Interface protocols
- Memory interfaces and controllers
- Processors and peripherals

**Note:** The IP Catalog (**Tools > IP Catalog**) and parameter editor replace the MegaWizard™ Plug-In Manager for IP selection and parameterization, beginning in Quartus II software version 14.0. Use the IP Catalog and parameter editor to locate and parameterize Altera and other supported IP cores.

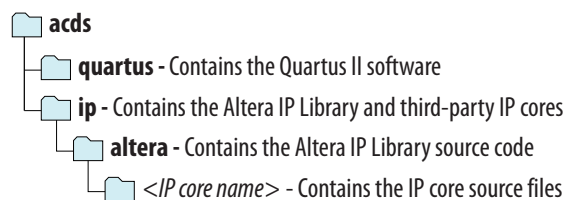
#### Related Information

- [IP User Guide Documentation](#)
- [Altera IP Release Notes](#)

## Installing and Licensing IP Cores

The Altera IP Library provides many useful IP core functions for production use without purchasing an additional license. You can evaluate any Altera IP core in simulation and compilation in the Quartus® II software using the OpenCore® evaluation feature. Some Altera IP cores, such as MegaCore® functions, require that you purchase a separate license for production use. You can use the OpenCore Plus feature to evaluate IP that requires purchase of an additional license until you are satisfied with the functionality and performance. After you purchase a license, visit the Self Service Licensing Center to obtain a license number for any Altera product.

Figure 1-12: IP Core Installation Path



**Note:** The default IP installation directory on Windows is `<drive>:\altera\<version number>`; on Linux it is `<home directory>/altera/ <version number>`.

#### Related Information

- [Altera Licensing Site](#)
- [Altera Software Installation and Licensing Manual](#)

## OpenCore Plus IP Evaluation

Altera's free OpenCore Plus feature allows you to evaluate licensed MegaCore IP cores in simulation and hardware before purchase. You need only purchase a license for MegaCore IP cores if you decide to take your design to production. OpenCore Plus supports the following evaluations:

- Simulate the behavior of a licensed IP core in your system.
- Verify the functionality, size, and speed of the IP core quickly and easily.
- Generate time-limited device programming files for designs that include IP cores.
- Program a device with your IP core and verify your design in hardware

OpenCore Plus evaluation supports the following two operation modes:

- Untethered—run the design containing the licensed IP for a limited time.
- Tethered—run the design containing the licensed IP for a longer time or indefinitely. This requires a connection between your board and the host computer.

**Note:** All IP cores that use OpenCore Plus time out simultaneously when any IP core in the design times out.

## IP Catalog and Parameter Editor

The Quartus II IP Catalog (**Tools > IP Catalog**) and parameter editor help you easily customize and integrate IP cores into your project. You can use the IP Catalog and parameter editor to select, customize, and generate files representing your custom IP variation.

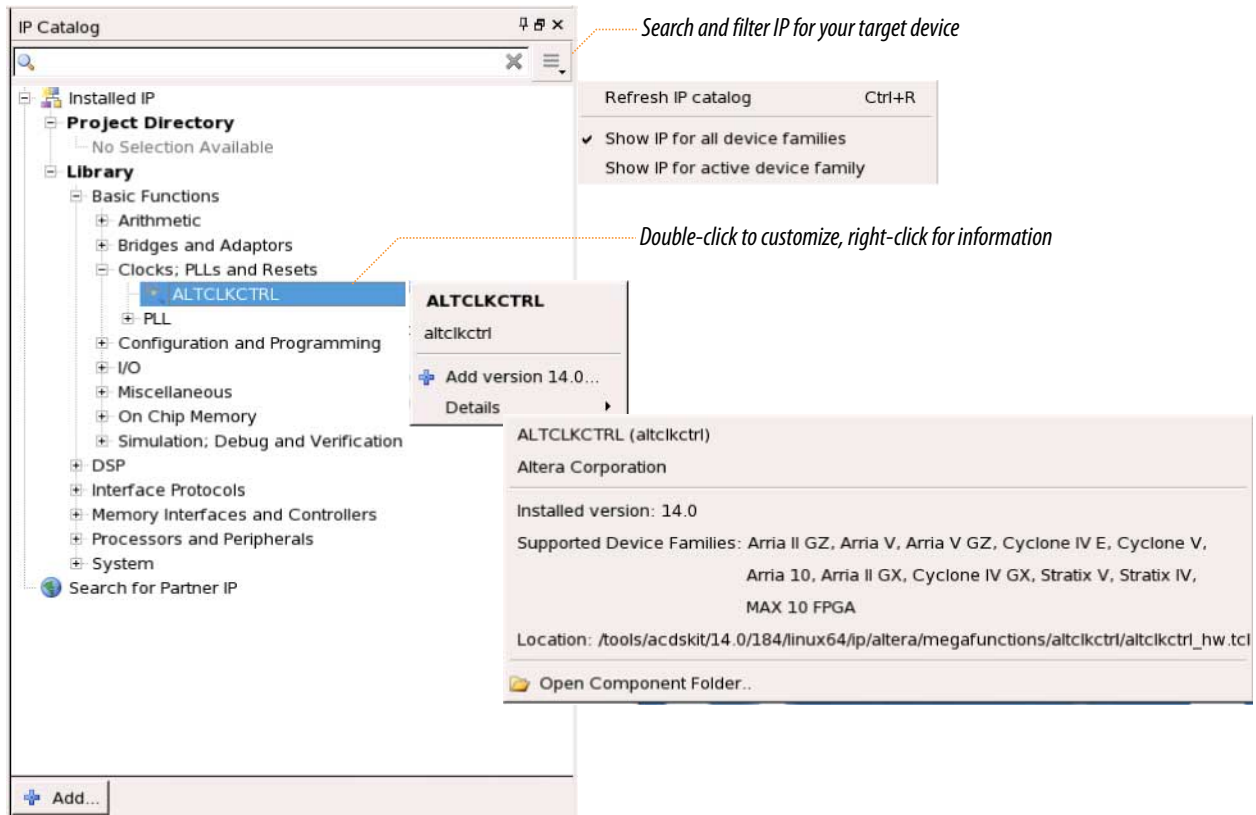
**Note:** The IP Catalog (**Tools > IP Catalog**) and parameter editor replace the MegaWizard™ Plug-In Manager for IP selection and parameterization, beginning in Quartus II software version 14.0. Use the IP Catalog and parameter editor to locate and parameterize Altera IP cores.

The IP Catalog lists IP cores available for your design. Double-click any IP core to launch the parameter editor and generate files representing your IP variation. The parameter editor prompts you to specify an IP variation name, optional ports, and output file generation options. The parameter editor generates a top-level Qsys system file (**.qsys**) or Quartus II IP file (**.qip**) representing the IP core in your project. You can also parameterize an IP variation without an open project.

Use the following features to help you quickly locate and select an IP core:

- Filter IP Catalog to **Show IP for active device family** or **Show IP for all device families**.
- Search to locate any full or partial IP core name in IP Catalog. Click **Search for Partner IP**, to access partner IP information on the Altera website.
- Right-click an IP core name in IP Catalog to display details about supported devices, open the IP core's installation folder, and/or view links to documentation.

Figure 1-13: Quartus II IP Catalog



**Note:** The IP Catalog is also available in Qsys (**View > IP Catalog**). The Qsys IP Catalog includes exclusive system interconnect, video and image processing, and other system-level IP that are not available in the Quartus II IP Catalog. For more information about using the Qsys IP Catalog, refer to *Creating a System with Qsys* in the *Quartus II Handbook*.

**Related Information**

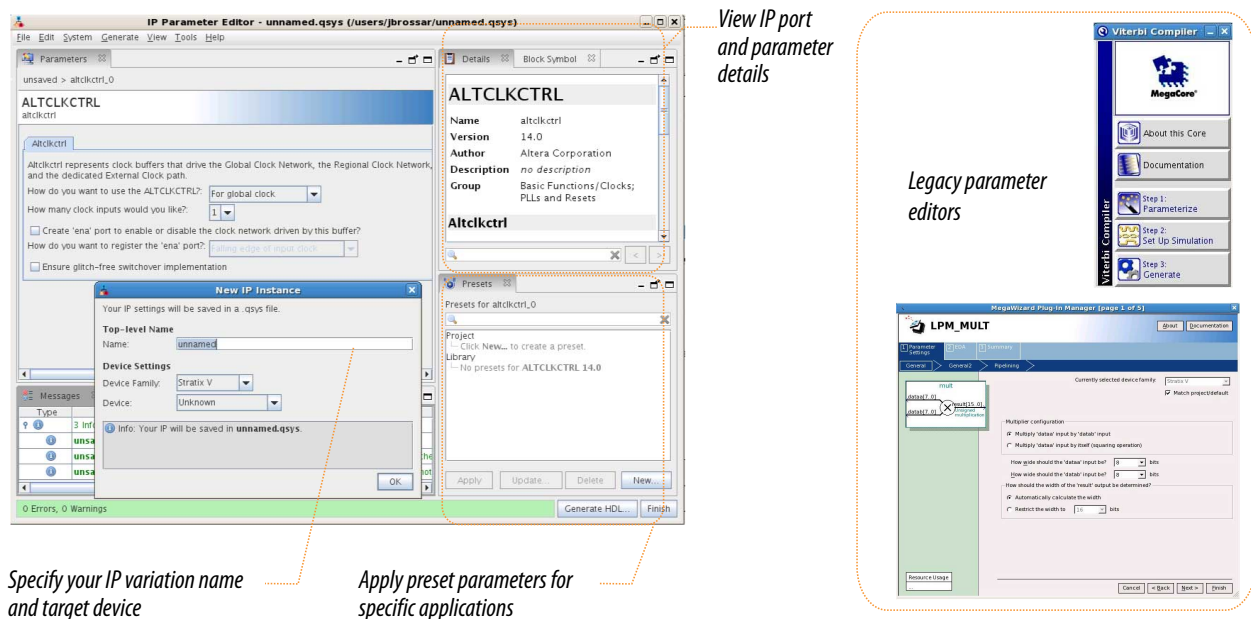
[Creating a System With Qsys](#) on page 5-1

**Using the Parameter Editor**

The parameter editor helps you to configure IP core ports, parameters, and output file generation options.

- Use preset settings in the parameter editor (where provided) to instantly apply preset parameter values for specific applications.
- View port and parameter descriptions, and links to documentation.
- Generate testbench systems or example designs (where provided).

Figure 1-14: IP Parameter Editors



## Adding IP Cores to IP Catalog

The IP Catalog automatically displays Altera IP cores found in the project directory, in the Altera installation directory, and in the defined IP search path. The IP Catalog can include Altera-provided IP components, third-party IP components, custom IP components that you provide, and previously generated Qsys systems.

You can use the **IP Search Path** option (**Tools > Options**) to include custom and third-party IP components in the IP Catalog. The IP Catalog displays all IP cores in the IP search path. The Quartus software searches the directories listed in the IP search path for the following IP core files:

- Component Description File (**\_hw.tcl**)—Defines a single IP core.
- IP Index File (**.ipx**)—Each **.ipx** file indexes a collection of available IP cores, or a reference to other directories to search. In general, **.ipx** files facilitate faster searches.

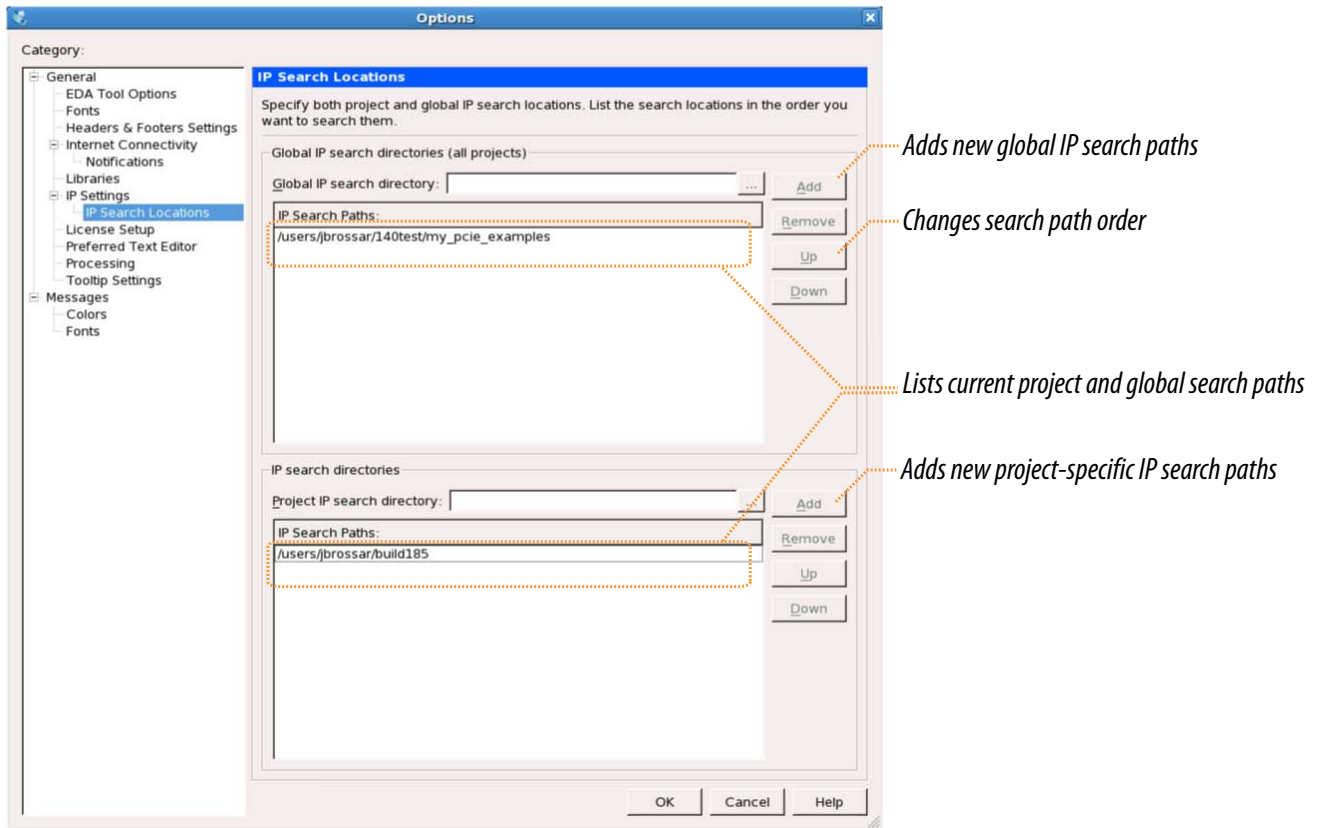
The Quartus software searches some directories recursively and other directories only to a specific depth. When the search is recursive, the search stops at any directory that contains an **\_hw.tcl** or **.ipx** file.

In the following list of search locations, a recursive descent is annotated by \*\*. A single \* signifies any file.

Table 1-2: IP Search Locations

Location	Description
<b>PROJECT_DIR/*</b>	Finds IP components and index files in the Quartus project directory.
<b>PROJECT_DIR/ip/**/*</b>	Finds IP components and index files in any subdirectory of the <b>/ip</b> subdirectory of the Quartus project directory.

Figure 1-15: Specifying IP Search Locations



If the Quartus software recognizes two IP cores with the same name, the following search path precedence rules determine the resolution of files:

1. Project directory.
2. Project database directory.
3. Project IP search path specified in **IP Search Locations**, or with the `SEARCH_PATH` assignment in the Quartus Settings File (`.qsf`) for the current project revision.
4. Global IP search path specified in **IP Search Locations**, or with the `SEARCH_PATH` assignment in the `quartus2.ini` file.
5. Quartus software libraries directory, such as `<Quartus Installation>\libraries`.

**Note:** If you add a component to the search path, you must refresh your system by clicking **File > Refresh** to update the IP Catalog.

## General IP Core Settings

You can use the following settings to control how the Quartus software manages IP cores in your project.



Table 1-3: IP Core General Setting Locations

Setting Location	Description
<b>Tools &gt; Options &gt; IP Settings</b> Or <b>Assignments &gt; Settings &gt; IP Settings</b> (only enabled with open project)	<ul style="list-style-type: none"> <li>Specify your <b>IP generation HDL preference</b>. The parameter editor generates IP files in your preferred HDL by default.</li> <li>Increase <b>Maximum Qsys memory usage size</b> if you experience slow processing for large systems, or if Qsys reports an Out of Memory error.</li> <li>Specify whether to <b>Automatically add Quartus II IP files</b> to all projects. Disable this option to control addition of IP files manually. You may want to experiment with IP before adding to a project.</li> <li><b>Use the IP Generation Policy</b> setting to control when synthesis files are regenerated for each IP variation. Typically you <b>Always regenerate synthesis files for IP cores</b> after making changes to an IP variation.</li> </ul>
<b>Tools &gt; Options &gt; IP Catalog Search Locations</b> Or <b>Assignments &gt; Settings &gt; IP Catalog Search Locations</b>	<ul style="list-style-type: none"> <li>Specify project and global IP search locations. The Quartus II software searches for IP cores in the project directory, in the Altera installation directory, and in the IP search path.</li> </ul>
<b>Assignments &gt; Settings &gt; Simulation</b>	<ul style="list-style-type: none"> <li><b>NativeLink Settings</b> allow you to automatically compile testbenches for supported simulators. You can also specify a script to compile the testbench, and a script to set up the simulation.</li> </ul>

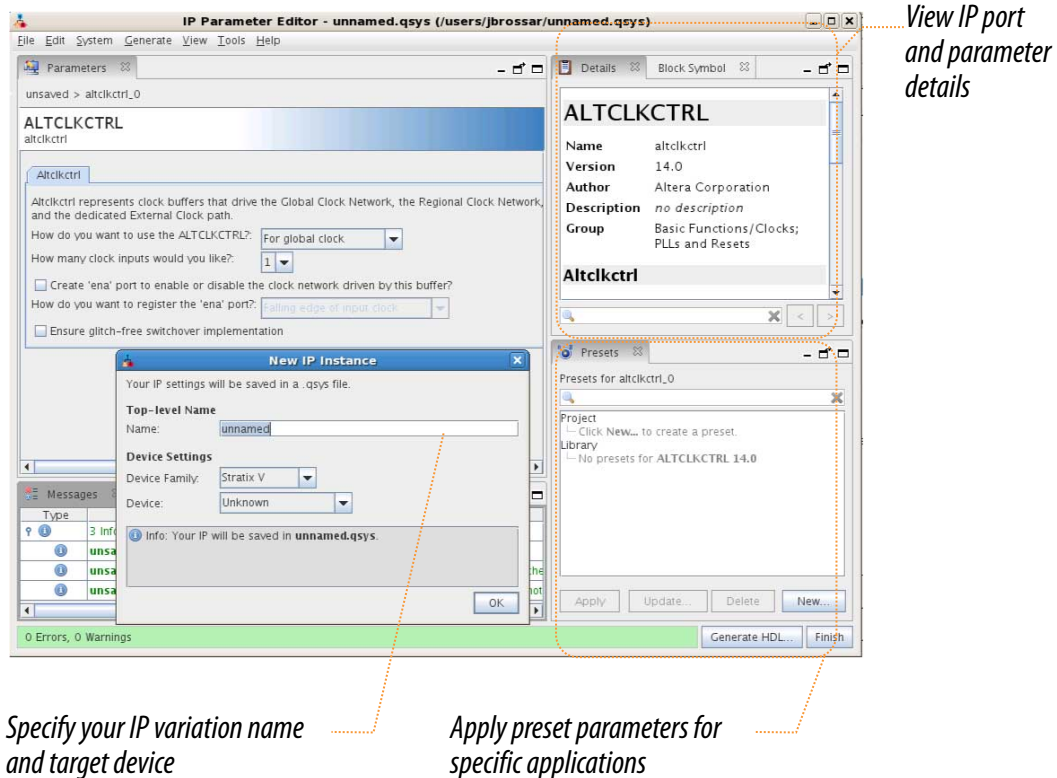
## Specifying IP Core Parameters and Options

The parameter editor GUI allows you to quickly configure your custom IP variation. Use the following steps to specify IP core options and parameters in the Quartus II software. Refer to *Specifying IP Core Parameters and Options (Legacy Parameter Editors)* for configuration of IP cores using the legacy parameter editor.

- In the IP Catalog (**Tools > IP Catalog**), locate and double-click the name of the IP core to customize. The parameter editor appears.
- Specify a top-level name for your custom IP variation. The parameter editor saves the IP variation settings in a file named `<your_ip>.qsys`. Click **OK**.
- Specify the parameters and options for your IP variation in the parameter editor, including one or more of the following. Refer to your IP core user guide for information about specific IP core parameters.
  - Optionally select preset parameter values if provided for your IP core. Presets specify initial parameter values for specific applications.
  - Specify parameters defining the IP core functionality, port configurations, and device-specific features.
  - Specify options for processing the IP core files in other EDA tools.
- Click **Generate HDL**, the **Generation** dialog box appears.

5. Specify output file generation options, and then click **Generate**. The IP variation files generate according to your specifications.
6. To generate a simulation testbench, click **Generate > Generate Testbench System**.
7. To generate an HDL instantiation template that you can copy and paste into your text editor, click **Generate > HDL Example**.
8. Click **Finish**. The parameter editor adds the top-level **.qsys** file to the current project automatically. If you are prompted to manually add the **.qsys** file to the project, click **Project > Add/Remove Files in Project** to add the file.
9. After generating and instantiating your IP variation, make appropriate pin assignments to connect ports.

Figure 1-16: IP Parameter Editor



## Files Generated for Altera IP Cores

The Quartus software generates the following IP core output file structure.

Figure 1-17: IP Core Generated Files

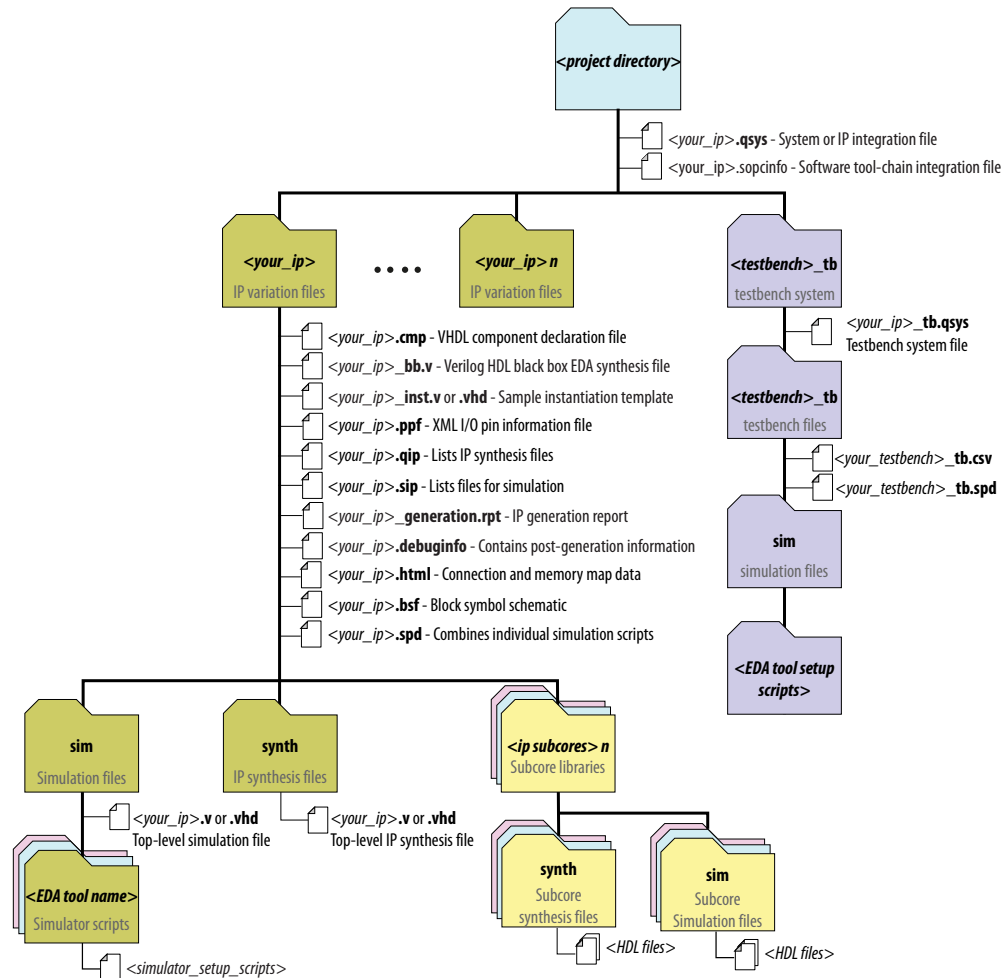


Table 1-4: IP Core Generated Files

File Name	Description
<b>&lt;my_ip&gt;.qsys</b>	The Qsys system or top-level IP variation file. <my_ip> is the name that you give your IP variation.
<b>&lt;system&gt;.sopcinfo</b>	<p>Describes the connections and IP component parameterizations in your Qsys system. You can parse its contents to get requirements when you develop software drivers for IP components.</p> <p>Downstream tools such as the Nios II tool chain use this file. The <b>.sopcinfo</b> file and the <b>system.h</b> file generated for the Nios II tool chain include address map information for each slave relative to each master that accesses the slave. Different masters may have a different address map to access a particular slave component.</p>

File Name	Description
<b>&lt;my_ip&gt;.cmp</b>	The VHDL Component Declaration ( <b>.cmp</b> ) file is a text file that contains local generic and port definitions that you can use in VHDL design files.
<b>&lt;my_ip&gt;.html</b>	A report that contains connection information, a memory map showing the address of each slave with respect to each master to which it is connected, and parameter assignments.
<b>&lt;my_ip&gt;_generation.rpt</b>	IP or Qsys generation log file. A summary of the messages during IP generation.
<b>&lt;my_ip&gt;.debuginfo</b>	Contains post-generation information. Used to pass System Console and Bus Analyzer Toolkit information about the Qsys interconnect. The Bus Analysis Toolkit uses this file to identify debug components in the Qsys interconnect.
<b>&lt;my_ip&gt;.qip</b>	Contains all the required information about the IP component to integrate and compile the IP component in the Quartus software.
<b>&lt;my_ip&gt;.csv</b>	Contains information about the upgrade status of the IP component.
<b>&lt;my_ip&gt;.bsf</b>	A Block Symbol File ( <b>.bsf</b> ) representation of the IP variation for use in Quartus Block Diagram Files ( <b>.bdf</b> ).
<b>&lt;my_ip&gt;.spd</b>	Required input file for <code>ip-make-simscript</code> to generate simulation scripts for supported simulators. The <b>.spd</b> file contains a list of files generated for simulation, along with information about memories that you can initialize.
<b>&lt;my_ip&gt;.ppf</b>	The Pin Planner File ( <b>.ppf</b> ) stores the port and node assignments for IP components created for use with the Pin Planner.
<b>&lt;my_ip&gt;_bb.v</b>	You can use the Verilog black-box ( <b>_bb.v</b> ) file as an empty module declaration for use as a black box.
<b>&lt;my_ip&gt;.sip</b>	Contains information required for NativeLink simulation of IP components. You must add the <b>.sip</b> file to your Quartus project.
<b>&lt;my_ip&gt;_inst.v or _inst.vhd</b>	HDL example instantiation template. You can copy and paste the contents of this file into your HDL file to instantiate the IP variation.
<b>&lt;my_ip&gt;.regmap</b>	If IP contains register information, <b>.regmap</b> file generates. The <b>.regmap</b> file describes the register map information of master and slave interfaces. This file complements the <b>.sopcinfo</b> file by providing more detailed register information about the system. This enables register display views and user customizable statistics in the System Console.

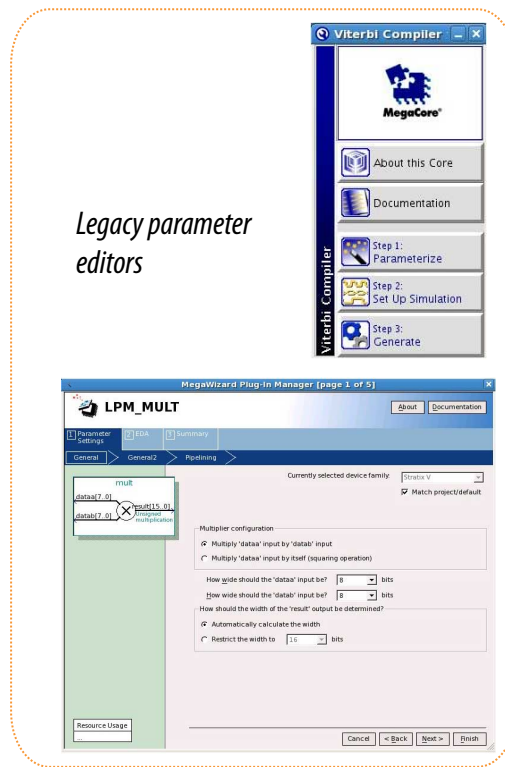
File Name	Description
<my_ip>.svd	Allows HPS System Debug tools to view the register maps of peripherals connected to HPS within a Qsys system.  During synthesis, the .svd files for slave interfaces visible to System Console masters are stored in the .sof file in the debug section. System Console reads this section, which Qsys can query for register map information. For system slaves, Qsys can access the registers by name.
<my_ip>.v or <my_ip>.vhd	HDL files that instantiate each submodule or child IP core for synthesis or simulation.
mentor/	Contains a ModelSim® script <b>msim_setup.tcl</b> to set up and run a simulation.
aldec/	Contains a Riviera-PRO script <b>rivierapro_setup.tcl</b> to setup and run a simulation.
/synopsys/vcs /synopsys/vcsmx	Contains a shell script <b>vcs_setup.sh</b> to set up and run a VCS® simulation.  Contains a shell script <b>vcsmx_setup.sh</b> and <b>synopsys_sim.setup</b> file to set up and run a VCS MX® simulation.
/cadence	Contains a shell script <b>ncsim_setup.sh</b> and other setup files to set up and run an NCSIM simulation.
/submodules	Contains HDL files for the IP core submodule.
<child IP cores>/	For each generated child IP core directory, Qsys generates <b>/synth</b> and <b>/sim</b> sub-directories.

## Specifying IP Core Parameters and Options (Legacy Parameter Editors)

Some IP cores use a legacy version of the parameter editor for configuration and generation. Use the following steps to configure and generate an IP variation using a legacy parameter editor.

**Note:** The legacy parameter editor generates a different output file structure than the latest parameter editor. Refer to *Specifying IP Core Parameters and Options* for configuration of IP cores that use the latest parameter editor.

Figure 1-18: Legacy Parameter Editors



1. In the IP Catalog (**Tools > IP Catalog**), locate and double-click the name of the IP core to customize. The parameter editor appears.
2. Specify a top-level name and output HDL file type for your IP variation. This name identifies the IP core variation files in your project. Click **OK**.
3. Specify the parameters and options for your IP variation in the parameter editor. Refer to your IP core user guide for information about specific IP core parameters.
4. Click **Finish** or **Generate** (depending on the parameter editor version). The parameter editor generates the files for your IP variation according to your specifications. Click **Exit** if prompted when generation is complete. The parameter editor adds the top-level **.qip** file to the current project automatically.

**Note:** To manually add an IP variation generated with legacy parameter editor to a project, click **Project > Add/Remove Files in Project** and add the IP variation **.qip** file.

### Files Generated for Altera IP Cores (Legacy Parameter Editors)

The Quartus software generates one of the following output file structures for Altera IP cores that use a legacy parameter editor.

Figure 1-19: IP Core Generated Files (Legacy Parameter Editor)



**Note:** To manually add an IP variation to a Quartus project, click **Project > Add/Remove Files in Project** and add only the IP variation .qip or .qsys file, but not both, to the project. Do not manually add the top-level HDL file to the project.

## Generate a Qsys System or IP Variation with qsys-generate

You can use the `qsys-generate` utility to generate RTL for your Qsys system, an IP core variation, or simulation models and scripts. You can create testbench systems for testing your Qsys system in a simulator using bus functional models (BFMs). Output from the `qsys-generate` command is the same as when generating using the Qsys GUI.

**Table 1-5: qsys-generate Command-Line Options**

Option	Usage	Description
<code>&lt;1st arg file&gt;</code>	Required	The name of the <b>.qsys</b> system file to generate.
<code>--synthesis=&lt;VERILOG VHDL&gt;</code>	Optional	Creates synthesis HDL files that Qsys uses to compile the system in a Quartus II project. You must specify the preferred generation language for the top-level RTL file for the generated Qsys system.
<code>--block-symbol-file</code>	Optional	Creates a Block Symbol File ( <b>.bsf</b> ) for the Qsys system.
<code>--simulation=&lt;VERILOG VHDL&gt;</code>	Optional	Creates a simulation model for the Qsys system. The simulation model contains generated HDL files for the simulator, and may include simulation-only features. You must specify the preferred simulation language.
<code>--testbench=&lt;SIMPLE STANDARD&gt;</code>	Optional	Creates a testbench system that instantiates the original system, adding bus functional models (BFMs) to drive the top-level interfaces. When you generate the system, the BFMs interact with the system in the simulator.
<code>--testbench-simulation=&lt;VERILOG VHDL&gt;</code>	Optional	After you create the testbench system, you can create a simulation model for the testbench system.
<code>--search-path=&lt;value&gt;</code>	Optional	If you omit this command, Qsys uses a standard default path. If you provide this command, Qsys searches a comma-separated list of paths. To include the standard path in your replacement, use "\$", for example, "/extra/dir,\$".
<code>--jvm-max-heap-size=&lt;value&gt;</code>	Optional	The maximum memory size that Qsys uses for allocations when running <code>qsys-generate</code> . You specify the value as <code>&lt;size&gt; &lt;unit&gt;</code> , where <code>unit</code> is <code>m</code> (or <code>M</code> ) for multiples of megabytes or <code>g</code> (or <code>G</code> ) for multiples of gigabytes. The default value is 512m.



Option	Usage	Description
<code>--family=&lt;value&gt;</code>	Optional	Specifies the device family.
<code>--part=&lt;value&gt;</code>	Optional	Specifies the device part number. If set, this option overrides the <code>--family</code> option.
<code>--allow-mixed-language-simulation</code>	Optional	Enables a mixed language simulation model generation. If true, if a preferred simulation language is set, Qsys uses a <code>fileset</code> of the component for the simulation model generation. When false, which is the default, Qsys uses the language specified with <code>--file-set=&lt;value&gt;</code> for all components for simulation model generation.

## Modifying an IP Variation

You can easily modify the parameters of any Altera IP core variation in the parameter editor to match your design requirements. Use any of the following methods to modify an IP variation in the parameter editor.

Table 1-6: Modifying an IP Variation

Menu Command	Action
<b>File &gt; Open</b>	Select the top-level HDL ( <code>.v</code> , or <code>.vhd</code> ) IP variation file to launch the parameter editor and modify the IP variation. Regenerate the IP variation to implement your changes.
<b>View &gt; Utility Windows &gt; Project Navigator &gt; IP Components</b>	Double-click the IP variation to launch the parameter editor and modify the IP variation. Regenerate the IP variation to implement your changes.
<b>Project &gt; Upgrade IP Components</b>	Select the IP variation and click <b>Upgrade in Editor</b> to launch the parameter editor and modify the IP variation. Regenerate the IP variation to implement your changes.

## Upgrading IP Cores

IP core variants generated with a previous version of the Quartus II software may require upgrading before use in the current version of the Quartus II software. Click **Project > Upgrade IP Components** to identify and upgrade IP core variants.

The **Upgrade IP Components** dialog box provides instructions when IP upgrade is required, optional, or unsupported for specific IP cores in your design. You must upgrade IP cores that require it before you can compile the IP variation in the current version of the Quartus II software. Many Altera IP cores support automatic upgrade.

The upgrade process renames and preserves the existing variation file (`.v`, `.sv`, or `.vhd`) as `<my_variant>_BAK.v`, `.sv`, `.vhd` in the project directory.

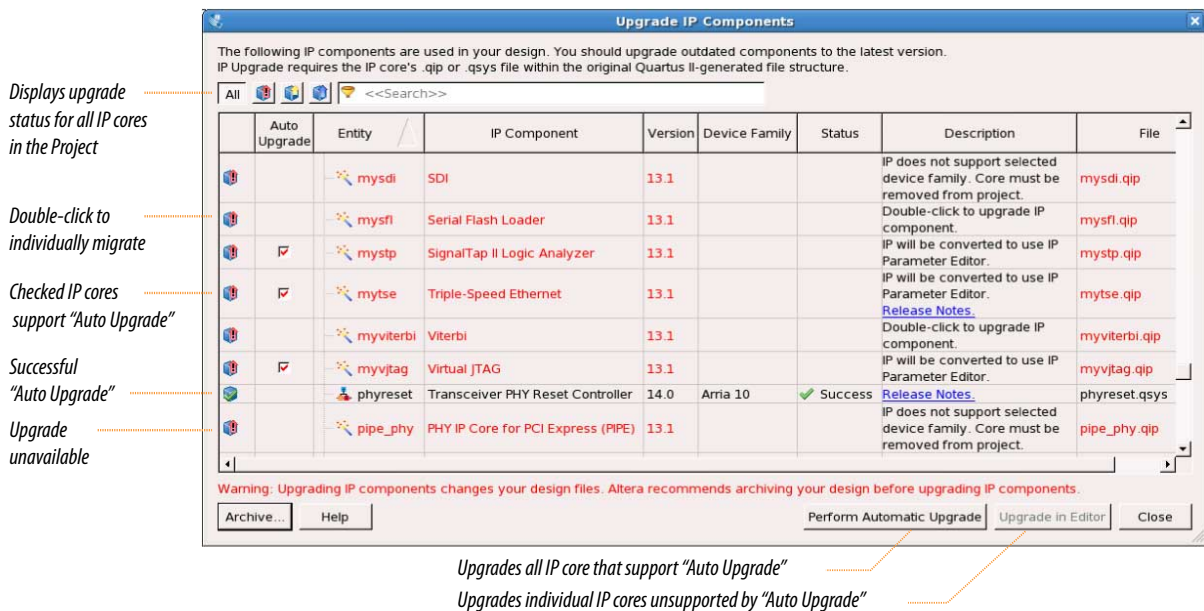
**Table 1-7: IP Core Upgrade Status**

IP Core Status	Corrective Action
Required Upgrade IP Components	You must upgrade the IP variation before compiling in the current version of the Quartus II software.
Optional Upgrade IP Components	Upgrade is optional for this IP variation in the current version of the Quartus II software. You can upgrade this IP variation to take advantage of the latest development of this IP core. Alternatively you can retain previous IP core characteristics by declining to upgrade.
Upgrade Unsupported	Upgrade of the IP variation is not supported in the current version of the Quartus II software due to IP core end of life or incompatibility with the current version of the Quartus II software. You are prompted to replace the obsolete IP core with a current equivalent IP core from the IP Catalog.

**Before you begin**

- Archive the Quartus II project containing outdated IP cores in the original version of the Quartus II software: Click **Project > Archive Project** to save the project in your previous version of the Quartus II software. This archive preserves your original design source and project files.
  - Restore the archived project in the latest version of the Quartus II software: Click **Project > Restore Archived Project**. Click **OK** if prompted to change to a supported device or overwrite the project database. File paths in the archive must be relative to the project directory. File paths in the archive must reference the IP variation **.v** or **.vhd** file or **.qsys** file (not the **.qip** file).
1. In the latest version of the Quartus II software, open the Quartus II project containing an outdated IP core variation. The **Upgrade IP Components** dialog automatically displays the status of IP cores in your project, along with instructions for upgrading each core. Click **Project > Upgrade IP Components** to access this dialog box manually.
  2. To simultaneously upgrade all IP cores that support automatic upgrade, click **Perform Automatic Upgrade**. The **Status** and **Version** columns update when upgrade is complete. Example designs provided with any Altera IP core regenerate automatically whenever you upgrade the IP core.

Figure 1-20: Upgrading IP Cores



### Example 1-1: Upgrading IP Cores at the Command Line

You can upgrade IP cores that support auto upgrade at the command line. IP cores that do not support automatic upgrade do not support command line upgrade.

- To upgrade a single IP core that supports auto-upgrade, type the following command:

```
quartus_sh -ip_upgrade -variation_files <my_ip_filepath/my_ip>.<hdl>
<qii_project>
```

Example:

```
quartus_sh -ip_upgrade -variation_files mega/pll25.v hps_testx
```

- To simultaneously upgrade multiple IP cores that support auto-upgrade, type the following command:

```
quartus_sh -ip_upgrade -variation_files "<my_ip_filepath/my_ip1>.<hdl>;
<my_ip_filepath/my_ip2>.<hdl>" <qii_project>
```

Example:

```
quartus_sh -ip_upgrade -variation_files "mega/pll_tx2.v;mega/pll3.v"
hps_testx
```

**Note:** IP cores older than Quartus II software version 12.0 do not support upgrade. Altera verifies that the current version of the Quartus II software compiles the previous version of each IP core. The *Altera IP Release Notes* reports any verification exceptions for Altera IP cores. Altera does not verify compilation for IP cores older than the previous two releases.

**Related Information**[Altera IP Release Notes](#)

## Migrating IP Cores to a Different Device

IP migration allows you to target the latest device families with IP originally generated for a different device. Some Altera IP cores require individual migration to upgrade. The **Upgrade IP Components** dialog box prompts you to double-click IP cores that require individual migration.

1. To display IP cores requiring migration, click **Project > Upgrade IP Components**. The **Description** field prompts you to double-click IP cores that require individual migration.
2. Double-click the IP core name, and then click **OK** after reading the information panel. The parameter editor appears showing the original IP core parameters.
3. For the **Currently selected device family**, turn off **Match project/default**, and then select the new target device family.
4. Click **Finish**, and then click **Finish** again to migrate the IP variation using best-effort mapping to new parameters and settings. Click **OK** if you are prompted that the IP core is unsupported for the current device. A new parameter editor opens displaying best-effort mapped parameters.
5. Click **Generate HDL**, and then confirm the **Synthesis** and **Simulation** file options. Verilog is the parameter editor default HDL for synthesis files. If your original IP core was generated for VHDL, select **VHDL** to retain the original output HDL format.
6. To regenerate the new IP variation for the new target device, click **Generate**. When generation is complete, click **Close**.
7. Click **Finish** to complete migration of the IP core. Click **OK** if you are prompted to overwrite IP core files. The **Device Family** column displays the migrated device support. The migration process replaces `<my_ip>.qip` with the `<my_ip>.qsys` top-level IP file in your project.

**Note:** If migration does not replace `<my_ip>.qip` with `<my_ip>.qsys`, click **Project > Add/Remove Files in Project** to replace the file in your project.

8. Review the latest parameters in the parameter editor or generated HDL for correctness. IP migration may change ports, parameters, or functionality of the IP core. During migration, the IP core's HDL generates into a library that is different from the original output location of the IP core. Update any assignments that reference outdated locations. If your upgraded IP core is represented by a symbol in a supporting Block Design File schematic, replace the symbol with the newly generated `<my_ip>.bsf` after migration.

**Note:** The migration process may change the IP variation interface, parameters, and functionality. This may require you to change your design or to re-parameterize your variant after the **Upgrade IP Components** dialog box indicates that migration is complete. The **Description** field identifies IP cores that require design or parameter changes.

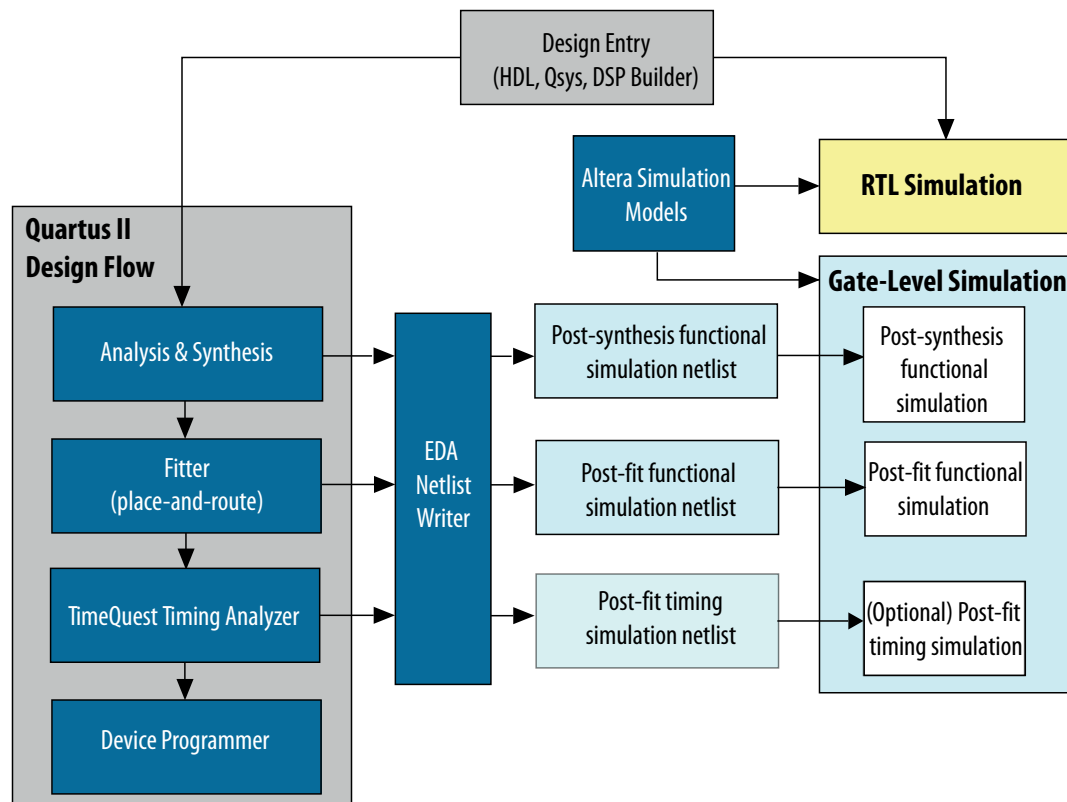
**Related Information**[Altera IP Release Notes](#)

## Simulating Altera IP Cores in other EDA Tools

The Quartus II software supports RTL and gate-level design simulation of Altera IP cores in supported EDA simulators. Simulation involves setting up your simulator working environment, compiling simulation model libraries, and running your simulation.

You can use the functional simulation model and the testbench or example design generated with your IP core for simulation. The functional simulation model and testbench files are generated in a project subdirectory. This directory may also include scripts to compile and run the testbench. For a complete list of models or libraries required to simulate your IP core, refer to the scripts generated with the testbench. You can use the Quartus II NativeLink feature to automatically generate simulation files and scripts. NativeLink launches your preferred simulator from within the Quartus II software.

**Figure 1-21: Simulation in Quartus II Design Flow**



**Note:** Post-fit timing simulation is not supported for 28nm and later device architectures. Altera IP supports a variety of simulation models, including simulation-specific IP functional simulation models and encrypted RTL models, and plain text RTL models. These are all cycle-accurate models. The models support fast functional simulation of your IP core instance using industry-standard VHDL or Verilog HDL simulators. For some cores, only the plain text RTL model is generated, and you can simulate that model. Use the simulation models only for simulation and not for synthesis or any other purposes. Using these models for synthesis creates a nonfunctional design.

#### Related Information

#### [Simulating Altera Designs](#)

### Generating Simulation Scripts

You can automatically generate simulation scripts to set up supported simulators. These scripts compile the required device libraries and system design files in the correct order, and then elaborate or load the top-level design for simulation. You can also use scripts to modify the top-level simulation environment,

independent of IP simulation files that are replaced during regeneration. You can modify the scripts to set up supported simulators.

Use the NativeLink feature to generate simulation scripts to automate simulation steps. You can reuse these generated files and simulation scripts in a custom simulation flow. NativeLink optionally generates scripts for your simulator in the project subdirectory.

1. Click **Assignments > Settings**.
2. Under **EDA Tool Settings**, click **Simulation**.
3. Select the **Tool name** of your simulator.
4. Click **More NativeLink Settings**.
5. Turn on **Generate third-party EDA tool command scripts without running the EDA tool**.

**Table 1-8: NativeLink Generated Scripts for RTL Simulation**

Simulator(s)	Simulation File	Use
Mentor Graphics ModelSim QuestaSim	<code>/simulation/modelsim/&lt;my_ip&gt;.do</code>	Source directly with your simulator.
Aldec Riviera Pro	<code>/simulation/modelsim/&lt;my_ip&gt;.do</code>	Source directly with your simulator.
Synopsys VCS	<code>/simulation/modelsim/&lt;revision name&gt;_ _&lt;rtl or gate&gt;.vcs</code>	Add your testbench file name to this options file to pass the file to VCS using the <code>-file</code> option. If you specify a testbench file to NativeLink, NativeLink generates an <code>.sh</code> script that runs VCS.
Synopsys VCS MX	<code>/simulation/scsim/&lt;revision name&gt;_ vcsmx_&lt;rtl or gate&gt;_&lt;verilog or vhdl&gt; .tcl</code>	Run this script at the command line using the command: <code>quartus_sh -t &lt;script&gt;</code> Any testbench you specify with NativeLink is included in this script.
Cadence Incisive (NC SIM)	<code>/simulation/ncsim/&lt;revision name&gt;_ ncsim_&lt;rtl or gate&gt;_&lt;verilog or vhdl&gt; .tcl</code>	Run this script at the command line using the command: <code>quartus_sh -t &lt;script&gt;</code> . Any testbench you specify with NativeLink is included in this script.

You can use the following script variables:

- `TOP_LEVEL_NAME`—The top-level entity of your simulation is often a testbench that instantiates your design, and then your design instantiates IP cores and/or Qsys systems. Set the value of `TOP_LEVEL_NAME` to the top-level entity.
- `QSYS_SIMDIR`—Specifies the top-level directory containing the simulation files.
- Other variables control the compilation, elaboration, and simulation process.

## Generating Custom Simulation Scripts with ip-make-simscript

Use the `ip-make-simscript` utility to generate simulation command scripts for multiple IP cores or Qsys systems. Specify all Simulation Package Descriptor files (`.spd`), each of which lists the required simulation files for the corresponding IP core or Qsys system. The IP parameter editor generates the `.spd` files.

`ip-make-simscript` compiles IP simulation models into various simulation libraries. Use the `compile-to-work` option to compile all simulation files into a single work library. Use this option only if you require a simplified library structure.

When you specify multiple `.spd` files, the `ip-make-simscript` utility generates a single simulation script containing all required simulation information. The default value of `TOP_LEVEL_NAME` is the `TOP_LEVEL_NAME` defined in the IP core or Qsys `.spd` file.

Set appropriate variables in the script, or edit the variable assignment directly in the script. If the simulation script is a Tcl file that is sourced in the simulator, set the variables before sourcing the script. If the simulation script is a shell script, pass in the variables as command-line arguments to the shell script.

- To run `ip-make-simscript`, type the following at the command prompt:

```
<Quartus installation path>\quartus\sopc_builder\bin\ip-make-simscript
```

**Table 1-9: ip-make-simscript Examples**

Option	Description	Status
<code>--spd=&lt;file&gt;</code>	Describes the list of compiled files and memory model hierarchy. If your design includes multiple IP cores or Qsys systems that include <code>.spd</code> files, use this option for each file. For example:  <code>ip-make-simscript --spd=ip1.spd --spd=ip2.spd</code>	Required
<code>--output-directory=&lt;directory&gt;</code>	Specifies the location of output files. If unspecified, the default setting is the directory from which <code>ip-make-simscript</code> is run.	Optional
<code>--compile-to-work</code>	Compiles all design files to the default work library. Use this option only if you encounter problems managing your simulation with multiple libraries.	Optional
<code>--use-relative-paths</code>	Uses relative paths whenever possible.	Optional

## Synthesizing Altera IP Cores in Other EDA Tools

You can use supported EDA tools to synthesize a design that includes Altera IP cores. When you generate the IP core synthesis files for use with third-party EDA synthesis tools, you can optionally create an area and timing estimation netlist. To enable generation, turn on **Create timing and resource estimates for third-party EDA synthesis tools** when customizing your IP variation.

The area and timing estimation netlist describes the IP core connectivity and architecture, but does not include details about the true functionality. This information enables certain third-party synthesis tools to

better report area and timing estimates. In addition, synthesis tools can use the timing information to achieve timing-driven optimizations and improve the quality of results.

The Quartus II software generates the `<variant name>_syn.v` netlist file in Verilog HDL format regardless of the output file format you specify. If you use this netlist for synthesis, you must include the IP core wrapper file `<variant name>.v` or `<variant name>.vhd` in your Quartus II project.

#### Related Information

- [Quartus II Integrated Synthesis](#) on page 16-1

## Instantiating IP Cores in HDL

You can instantiate an IP core directly in your HDL code by calling the IP core name and declaring its parameters, in the same manner as any other module, component, or subdesign. When instantiating an IP core in VHDL, you must include the associated libraries.

### Example Top-Level Verilog HDL Module

Verilog HDL ALTFP\_MULT in Top-Level Module with One Input Connected to Multiplexer.

```
module MF_top (a, b, sel, datab, clock, result);
    input [31:0] a, b, datab;
    input clock, sel;
    output [31:0] result;
    wire [31:0] wire_dataa;

    assign wire_dataa = (sel)? a : b;
    altfp_mult inst1
    (.dataa(wire_dataa), .datab(datab), .clock(clock), .result(result));

    defparam
        inst1.pipeline = 11,
        inst1.width_exp = 8,
        inst1.width_man = 23,
        inst1.exception_handling = "no";
endmodule
```

### Example Top-Level VHDL Module

VHDL ALTFP\_MULT in Top-Level Module with One Input Connected to Multiplexer.

```
library ieee;
use ieee.std_logic_1164.all;
library altera_mf;
use altera_mf.altera_mf_components.all;

entity MF_top is
    port (clock, sel : in std_logic;
          a, b, datab : in std_logic_vector(31 downto 0);
          result : out std_logic_vector(31 downto 0));
end entity;

architecture arch_MF_top of MF_top is
    signal wire_dataa : std_logic_vector(31 downto 0);
begin

    wire_dataa <= a when (sel = '1') else b;

    inst1 : altfp_mult
        generic map (
            pipeline => 11,
```



```

width_exp => 8,
width_man => 23,
exception_handling => "no")
port map (
  dataa => wire_dataa,
  datab => datab,
  clock => clock,
  result => result);
end arch_MF_top;

```

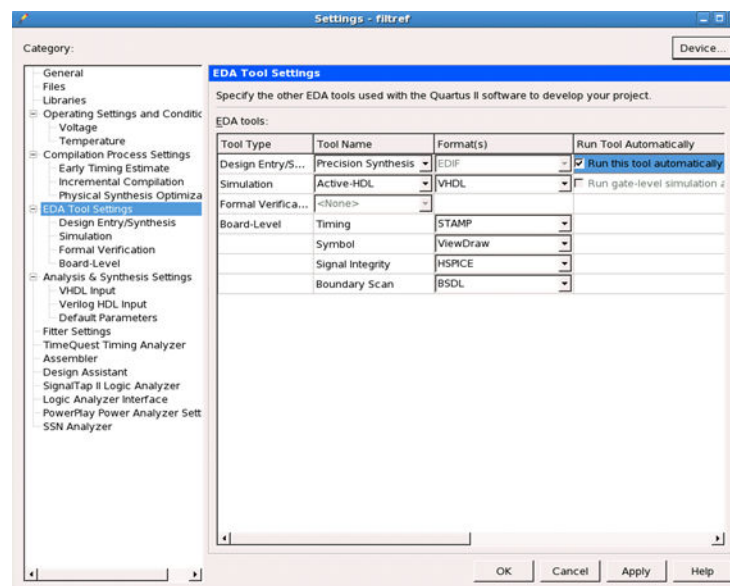
## Integrating Other EDA Tools

You can integrate supported EDA design entry, synthesis, simulation, physical synthesis, and formal verification tools into the Quartus II design flow. The Quartus II software supports netlist files from other EDA design entry and synthesis tools. The Quartus II software optionally generates various files for use in other EDA tools.

The Quartus II software manages EDA tool files and provides the following integration capabilities:

- Automatically generate files for synthesis and simulation and automatically launch other EDA tools (**Assignments > Settings > EDA Tool Settings > NativeLink Settings**).
- Compile all RTL and gate-level simulation model libraries for your device, simulator, and design language automatically (**Tools > Launch Simulation Library Compiler**).
- Include files (.edf, .vqm) generated by other EDA design entry or synthesis tools in your project as synthesized design files (**Project > Add/Remove File from Project**).
- Automatically generate optional files for board-level verification (**Assignments > Settings > EDA Tool Settings**).

Figure 1-22: EDA Tool Settings



### Related Information

[Mentor Graphics Precision Synthesis SupportGraphics](#) on page 18-1

## Simulating Altera Designs

# Managing Team-based Projects

The Quartus II software supports multiple designers, design iterations, and platforms. You can use the following techniques to preserve and track project changes in a team-based environment. These techniques may also be helpful for individual designers.

### Related Information

- [Preserving Compilation Results](#) on page 1-33
- [Archiving Projects](#) on page 1-35
- [Using External Revision Control](#) on page 1-36
- [Migrating Projects Across Operating Systems](#) on page 1-37

## Preserving Compilation Results

The Quartus II software maintains a database of compilation results for each project revision. The databases files store results of incremental or full compilation. Do not edit these files directly. However, you can use the database files in the following ways:

- Preserve compilation results for migration to a new version of the Quartus II software. Export a post-synthesis or post-fit, version-compatible database (**Project > Export Database**), and then import it into a newer version of the Quartus II software (**Project > Import Database**), or into another project.
- Optimize and lock down the compilation results for individual blocks. Export the post-synthesis or post-fit netlist as a Quartus II Exported Partition File (.qxp) (**Project > Export Design Partition**). You can then import the partition as a new project design file.
- Purge the content of the project database (**Project > Clean Project**) to remove unwanted previous compilation results at any time.

## Factors Affecting Compilation Results

Changes to any of the following factors can impact compilation results:

- Project Files—project settings (.qsf), design files, and timing constraints (.sdc).
- Hardware—CPU architecture, not including hard disk or memory size differences. Windows XP x32 results are not identical to Windows XP x64 results. Linux x86 results is not identical to Linux x86\_64.
- Quartus II Software Version—including build number and installed patches. Click **Help > About** to obtain this information.
- Operating System—Windows or Linux operating system, excluding version updates. For example, Windows XP, Windows Vista, and Windows 7 results are identical. Similarly, Linux RHEL, CentOS 4, and CentOS 5 results are identical.

### Related Information

- [Quartus II Incremental Compilation for Hierarchical and Team-Based Design](#) on page 3-1
- [Design Planning for Partial Reconfiguration](#) on page 4-1

The Partial Reconfiguration (PR) feature in the Quartus II software allows you to reconfigure a portion of the FPGA dynamically, while the remainder of the device continues to operate. The Quartus II software supports the PR feature for the Altera® Stratix® V device family.

## Migrating Results Across Quartus II Software Versions

View basic information about your project in the Project Navigator, Report panel, and Messages window.

To preserve compilation results for migration to a later version of the Quartus II software, export a version-compatible database file, and then import it into the later version of the Quartus II software. A few device families do not support version-compatible database generation, as indicated by project messages.

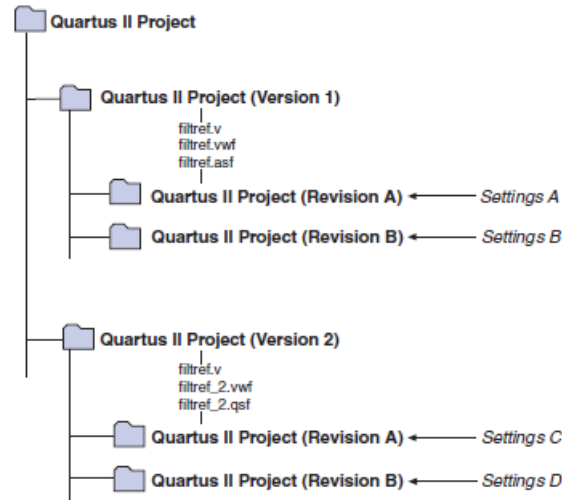
### Exporting and Importing the Results Database

To save the compilation results in a version-compatible format for migration to a later version of the Quartus II software, follow these steps:

1. Open the project for migration in the original version of the Quartus II software.
2. Generate the project database and netlist with one of the following:
  - Click **Processing > Start > Start Analysis & Synthesis** to generate a post-synthesis netlist.
  - Click **Processing > Start Compilation** to generate a post-fit netlist.
3. Click **Project > Export Database** and specify the **Export directory**.
4. In a later version of the Quartus II software, click **New Project Wizard** and create a new project with the same top-level design entity name as the migrated project.
5. Click **Project > Import Database** and select the **<project directory> /export\_db/exported database directory**. The Quartus II software opens the compiled project and displays compilation results.

**Note:** You can turn on **Assignments > Settings > Compilation Process Settings > Export version-compatible database** if you want to always export the database following compilation.

Figure 1-23: Quartus II Version-Compatible Database Structure



### Cleaning the Project Database

To clean the project database and remove all prior compilation results, follow these steps:

1. Click **Project > Clean Project**.
2. Select **All revisions** to remove the databases for all revisions of the current project, or specify a **Revision name** to remove only that revision's database.
3. Click **OK**. A message indicates when the database is clean.

## Archiving Projects

You can save the elements of a project in a single, compressed Quartus II Archive File (. **qar**) by clicking **Project > Archive Project**.

The **.qar** captures logic design, project, and settings files required to restore the project.

Use this technique to share projects between designers, or to transfer your project to a new version of the Quartus II software, or to Altera support. You can optionally add compilation results, Qsys system files, and third-party EDA tool files to the archive. If you restore the archive in a different version of the Quartus II software, you must include the original **.qdf** in the archive to preserve original compilation results.

## Manually Adding Files To Archives

To manually add files to an archive:

1. Click **Project > Archive Project** and specify the archive file name.
2. Click **Advanced**.
3. Select the **File set** for archive or select **Custom**. Turn on **File subsets** for archive.
4. Click **Add** and select Qsys system or EDA tool files. Click **OK**.
5. Click **Archive**.

## Archiving Compilation Results

You can include compilation results in a project archive to avoid recompilation and preserve original results in the restored project. To archive compilation results, export the post-synthesis or post-fit version compatible database and include this file in the archive.

1. Export the project database.
2. Click **Project > Archive Project** and specify the archive file name.
3. Click **Advanced**.
4. Under **File subsets**, turn on **Version-compatible database files** and click **OK**.
5. Click **Archive**.

To restore an archive containing a version-compatible database, follow these steps:

1. Click **Project > Restore Archived Project**.
2. Select the archive name and destination folder and click **OK**.
3. After restoring the archived project, click **Project > Import Database** and import the version-compatible database.

### Related Information

[Exporting and Importing the Results Database](#) on page 1-34

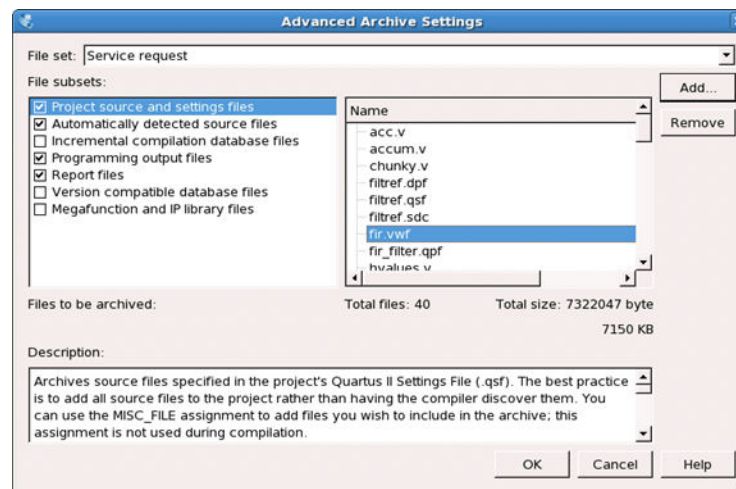
## Archiving Projects for Altera Service Requests

When archiving projects for an Altera service request, include all of the following file types for proper debugging by Altera Support:

To quickly identify and include appropriate archive files for an Altera service request:

1. Click **Project > Archive Project** and specify the archive file name.
2. Click **Advanced**.
3. In **File set**, select **Service Request** to include files for Altera Support.
  - Project source and setting files (**.v**, **.vhd**, **.vqm**, **.qsf**, **.sdc**, **.qip**, **.qpf**, **.cmp**, **.sip**)
  - Automatically detected source files (various)
  - Programming output files (**.jdi**, **.sof**, **.pof**)
  - Report files (**.rpt**, **.pin**, **.summary**, **.smsg**)
  - Qsys system and IP files (**.qsys**, **.qip**)
4. Click **OK**, and then click **Archive**.

**Figure 1-24: Archiving Project for Service Request**



## Using External Revision Control

Your project may involve different team members with distributed responsibilities, such as sub-module design, device and system integration, simulation, and timing closure. In such cases, it may be useful to track and protect file revisions in an external revision control system.

While Quartus II project revisions preserve various project setting and constraint combinations, external revision control systems can also track and merge RTL source code, simulation testbenches, and build scripts. External revision control supports design file version experimentation through branching and merging different versions of source code from multiple designers. Refer to your external revision control documentation for setup information.

### Files to Include In External Revision Control

Include the following Quartus project file types in external revision control systems:

- Logic design files (**.v**, **.vhd**, **.bdf**, **edf**, **.vqm**)
- Timing constraint files (**.sdc**)
- Quartus project settings and constraints (**.qdf**, **.qpf**, **.qsf**)
- IP files (**.v**, **.sv**, **.vhd**, **.qip**, **.sip**, **.qsys**)
- Qsys-generated files (**.qsys**, **.qip**, **.sip**)
- EDA tool files (**.vo**, **.who**)

You can generate or modify these files manually if you use a scripted design flow. If you use an external source code control system, you can check-in project files anytime you modify assignments and settings in the Quartus software.

## Migrating Projects Across Operating Systems


Consider the following cross-platform issues when moving your project from one operating system to another (for example, from Windows to Linux).

### Migrating Design Files and Libraries

Consider the following file naming differences when migrating projects across operating systems:

- Use appropriate case for your platform in file path references.
- Use a character set common to both platforms.
- Do not change the forward-slash (/) and back-slash (\) path separators in the **.qsf**. The Quartus II software automatically changes all back-slash (\) path separators to forward-slashes (/) in the **.qsf**.
- Observe the target platform's file name length limit.
- Use underscore instead of spaces in file and directory names.
- Change library absolute path references to relative paths in the **.qsf**.
- Ensure that any external project library exists in the new platform's file system.
- Specify file and directory paths as relative to the project directory. For example, for a project titled **foo\_design**, specify the source files as: **top.v**, **foo\_folder /foo1.v**, **foo\_folder /foo2.v**, and **foo\_folder / bar\_folder /bar1.vhdl**.
- Ensure that all the subdirectories are in the same hierarchical structure and relative path as in the original platform.

Figure 1-25: All Inclusive Project Directory Structure




### Use Relative Paths

Express file paths using relative path notation (**../**).

For example, in the directory structure shown you can specify **top.v** as **../source/top.v** and **foo1.v** as **../source/foo\_folder/foo1.v**.

Figure 1-26: Quartus II Project Directory Separate from Design Files



## Design Library Migration Guidelines

The following guidelines apply to library migration across computing platforms:

1. The project directory takes precedence over the project libraries.
2. For Linux, the Quartus II software creates the file in the **altera.quartus** directory under the *<home>* directory.
3. All library files are relative to the libraries. For example, if you specify the **user\_lib1** directory as a project library and you want to add the **/user\_lib1/foo1.v** file to the library, you can specify the **foo1.v** file in the **.qsf** as **foo1.v**. The Quartus II software includes files in specified libraries.
4. If the directory is outside of the project directory, an absolute path is created by default. Change the absolute path to a relative path before migration.
5. When copying projects that include libraries, you must either copy your project library files along with the project directory or ensure that your project library files exist in the target platform.
  - On Windows, the Quartus II software searches for the **quartus2.ini** file in the following directories and order:
    - USERPROFILE, for example, **C:\Documents and Settings\ <user name>**
    - Directory specified by the TMP environmental variable
    - Directory specified by the TEMP environmental variable
    - Root directory, for example, **C:\**

## Scripting API

You can use command-line executables or scripts to execute project commands, rather than using the GUI. The following commands are available for scripting project management.

## Scripting Project Settings

You can use a Tcl script to specify settings and constraints, rather than using the GUI. This can be helpful if you have many settings and wish to track them in a single file or spreadsheet for iterative comparison.

The **.qsf** supports only a limited subset of Tcl commands. Therefore, pass settings and constraints using a Tcl script:

1. Create a text file with the extension **.tcl** that contains your assignments in Tcl format.
2. Source the Tcl script file by adding the following line to the **.qsf**: `set_global_assignment -name SOURCE_TCL_SCR IPT_FILE <file name>`.

## Project Revision Commands

Use the following commands for scripting project revisions.

[Create Revision Command](#) on page 1-39

[Set Current Revision Command](#) on page 1-39

[Get Project Revisions Command](#) on page 1-39

[Delete Revision Command](#) on page 1-39

### Create Revision Command

```
create_revision <name> -based_on <revision_name> -copy_results -set_current
```

Option	Description
based_on (optional)	Specifies the revision name on which the new revision bases its settings.
copy_results	Copies the results from the "based_on" revision.
set_current (optional)	Sets the new revision as the current revision.

### Set Current Revision Command

The **-force** option enables you to open the revision that you specify under revision name and overwrite the compilation database if the database version is incompatible.

```
set_current_revision -force <revision name>
```

### Get Project Revisions Command

```
get_project_revisions <project_name>
```

### Delete Revision Command

```
delete_revision <revision name>
```

## Project Archive Commands

You can use Tcl commands and the `quartus_sh` executable to create and manage archives of a Quartus project.



## Creating a Project Archive

in a Tcl script or from a Tcl prompt, you can use the following command to create a Quartus archive:

```
project_archive <name>.qar
```

You can specify the following other options:

- `-all_revisions` - Includes all revisions of the current project in the archive.
- `-auto_common_directory` - Preserves original project directory structure in archive
- `-common_directory /<name>` - Preserves original project directory structure in specified subdirectory
- `-include_libraries` - Includes libraries in archive
- `-include_outputs` - Includes output files in archive
- `-use_file_set <file_set>` Includes specified fileset in archive
- `-version_compatible_database` - Includes version-compatible database files in archive

**Note:** Version-compatible databases are not available for some device families. If you require the database files to reproduce the compilation results in the same Quartus software version, use the `-use_file_set full_db` option to archive the complete database.

## Restoring an Archived Project

Use the following Tcl command to restore a Quartus project:

```
project_restore <name>.qar -destination restored -overwrite
```

This example restores to a destination directory named "restored".

## Project Database Commands

Use the following commands for managing Quartus project databases:

[Import and Export Version-Compatible Databases](#) on page 1-40

[Import and Export Version-Compatible Databases from a Flow Package](#) on page 1-40

[Generate Version-Compatible Database After Compilation](#) on page 1-41

[quartus\\_cdb and quartus\\_sh Executables to Manage Version-Compatible Databases](#) on page 1-41

## Import and Export Version-Compatible Databases

Use the following commands to import or export a version-compatible database:

- `import_database <directory>`
- `export_database <directory>`

## Import and Export Version-Compatible Databases from a Flow Package

The following are Tcl commands from the `flow` package to import or export version-compatible databases. If you use the `flow` package, you must specify the database directory variable name. `flow` and `database_manager` packages contain commands to manage version-compatible databases.

- `set_global_assignment -name VER_COMPATIBLE_DB_DIR <directory>`
- `execute_flow -flow export_database`
- `execute_flow -flow import_database`

## Generate Version-Compatible Database After Compilation

Use the following commands to generate a version-compatible database after compilation:

- `set_global_assignment -name AUTO_EXPORT_VER_COMPATIBLE_DB ON`
- `set_global_assignment -name VER_COMPATIBLE_DB_DIR <directory>`

## quartus\_cdb and quartus\_sh Executables to Manage Version-Compatible Databases

Use the following commands to manage version-compatible databases:

- `quartus_cdb <project> -c <revision> --export_database=<directory>`
- `quartus_cdb <project> -c <revision> --import_database=<directory>`
- `quartus_sh -flow export_database <project> -c \ <revision>`
- `quartus_sh -flow import_database <project> -c \ <revision>`

## Project Library Commands

Use the following commands to script project library changes.

[Specify Project Libraries With SEARCH\\_PATH Assignment](#) on page 1-41

[Report Specified Project Libraries Commands](#) on page 1-41

### Related Information

- [Tcl Scripting](#)
- [CommandLine Scripting](#)
- [Quartus Settings File Manual](#)

## Specify Project Libraries With SEARCH\_PATH Assignment

In Tcl, use commands in the `::quartus::project` package to specify project libraries, and the `set_global_assignment` command.

Use the following commands to script project library changes:

- `set_global_assignment -name SEARCH_PATH "../other_dir/library1"`
- `set_global_assignment -name SEARCH_PATH "../other_dir/library2"`
- `set_global_assignment -name SEARCH_PATH "../other_dir/library3"`

## Report Specified Project Libraries Commands

To report any project libraries specified for a project and any global libraries specified for the current installation of the Quartus software, use the `get_global_assignment` and `get_user_option` Tcl commands.

Use the following commands to report specified project libraries:

- `get_global_assignment -name SEARCH_PATH`
- `get_user_option -name SEARCH_PATH`

## Document Revision History

Table 1-10: Document Revision History

Date	Version	Changes
2014.12.15	14.1.0	<ul style="list-style-type: none"> <li>Updated content for DSE II GUI and optimizations.</li> <li>Added information about new <b>Assignments &gt; Settings &gt; IP Settings</b> that control frequency of synthesis file regeneration and automatic addition of IP files to the project.</li> </ul>
2014.08.18	14.0a10.0	<ul style="list-style-type: none"> <li>Added information about specifying parameters for IP cores targeting Arria 10 devices.</li> <li>Added information about the latest IP output for Quartus II version 14.0a10 targeting Arria 10 devices.</li> <li>Added information about individual migration of IP cores to the latest devices.</li> <li>Added information about editing existing IP variations.</li> </ul>
2014.06.30	14.0.0	<ul style="list-style-type: none"> <li>Replaced MegaWizard Plug-In Manager information with IP Catalog.</li> <li>Added standard information about upgrading IP cores.</li> <li>Added standard installation and licensing information.</li> <li>Removed outdated device support level information. IP core device support is now available in IP Catalog and parameter editor.</li> </ul>
November 2013	13.1.0	<ul style="list-style-type: none"> <li>Conversion to DITA format</li> </ul>
May 2013	13.0.0	<ul style="list-style-type: none"> <li>Overhaul for improved usability and updated information.</li> </ul>
June 2012	12.0.0	<ul style="list-style-type: none"> <li>Removed survey link.</li> <li>Updated information about <code>VERILOG_INCLUDE_FILE</code>.</li> </ul>
November 2011	10.1.1	Template update.
December 2010	10.1.0	<ul style="list-style-type: none"> <li>Changed to new document template.</li> <li>Removed Figure 4-1, Figure 4-6, Table 4-2.</li> <li>Moved “Hiding Messages” to Help.</li> <li>Removed references about the <code>set_user_option</code> command.</li> <li>Removed Classic Timing Analyzer references.</li> </ul>

### Related Information

[Quartus II Handbook Archive](#)

2014.06.30

QII5V1



Subscribe



Send Feedback

## Design Planning with the Quartus II Software

Because of the significant increase in FPGA device densities, designs are complex and can sometimes involve multiple designers. System architects must also resolve design issues when integrating design blocks. However, you can solve potential problems early in the design cycle by following the design planning considerations in this chapter. This chapter provides only an introduction to various design planning features in the Quartus® II software.

Before reading the design planning guidelines discussed in this chapter, consider your design priorities. More device features, density, or performance requirements can increase system cost. Signal integrity and board issues can impact I/O pin locations. Power, timing performance, and area utilization all affect each other, and compilation time is affected when optimizing these priorities.

The Quartus II software optimizes designs for the best overall results; however, you can change the settings to better optimize one aspect of your design, such as power utilization. Certain tools or debugging options can lead to restrictions in your design flow. Your design priorities help you choose the tools, features, and methodologies to use for your design.

After you select a device family, to check if additional guidelines are available, refer to the design guidelines section of the appropriate device handbook.

## Creating Design Specifications

Before you create your design logic or complete your system design, create detailed design specifications that define the system, specify the I/O interfaces for the FPGA, identify the different clock domains, and include a block diagram of basic design functions.

In addition, creating a test plan helps you to design for verification and ease of manufacture. For example, you might need to validate interfaces incorporated in your design. To perform any built-in self-test functions to drive interfaces, you can use a UART interface with a Nios® II processor inside the FPGA device.

If more than one designer works on your design, you must consider a common design directory structure or source control system to make design integration easier. Consider whether you want to standardize on an interface protocol for each design block.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



**Related Information**

- [Planning for On-Chip Debugging Tools](#) on page 2-8  
For guidelines related to analyzing and debugging the device after it is in the system.
- [Planning for Hierarchical and Team-Based Design](#) on page 2-11  
For more suggestions on team-based designs.
- [Using Qsys and Standard Interfaces in System Design](#) on page 2-2  
For improved reusability and ease of integration.

## Selecting Intellectual Property

Altera and its third-party intellectual property (IP) partners offer a large selection of standardized IP cores optimized for Altera devices. The IP you select often affects system design, especially if the FPGA interfaces with other devices in the system. Consider which I/O interfaces or other blocks in your system design are implemented using IP cores, and plan to incorporate these cores in your FPGA design.

The OpenCore Plus feature, which is available for many IP cores, allows you to program the FPGA to verify your design in the hardware before you purchase the IP license. The evaluation supports the following modes:

- Untethered—the design runs for a limited time.
- Tethered—the design requires an Altera serial JTAG cable connected between the JTAG port on your board and a host computer running the Quartus II Programmer for the duration of the hardware evaluation period.

**Related Information****[Intellectual Property](#)**

For descriptions of available IP cores.

## Using Qsys and Standard Interfaces in System Design

You can use the Quartus II Qsys system integration tool to create your design with fast and easy system-level integration. With Qsys, you can specify system components in a GUI and generate the required interconnect logic automatically, along with adapters for clock crossing and width differences.

Because system design tools change the design entry methodology, you must plan to start developing your design within the tool. Ensure all design blocks use appropriate standard interfaces from the beginning of the design cycle so that you do not need to make changes later.

Qsys components use Avalon<sup>®</sup> standard interfaces for the physical connection of components, and you can connect any logical device (either on-chip or off-chip) that has an Avalon interface. The Avalon Memory-Mapped interface allows a component to use an address mapped read or write protocol that enables flexible topologies for connecting master components to any slave components. The Avalon Streaming interface enables point-to-point connections between streaming components that send and receive data using a high-speed, unidirectional system interconnect between source and sink ports.

In addition to enabling the use of a system integration tool such as Qsys, using standard interfaces ensures compatibility between design blocks from different design teams or vendors. Standard interfaces simplify the interface logic to each design block and enable individual team members to test their individual design blocks against the specification for the interface protocol to ease system integration.

### Related Information

- [System Design with Qsys](#)  
For more information about using Qsys to improve your productivity.
- [SOPC Builder User Guide](#)  
For more information about SOPC Builder.

## Device Selection

The device you choose affects board specification and layout. This section provides guidelines in the device selection process.

Choose the device family that best suits your design requirements. Families differ in cost, performance, logic and memory density, I/O density, power utilization, and packaging. You must also consider feature requirements, such as I/O standards support, high-speed transceivers, global or regional clock networks, and the number of phase-locked loops (PLLs) available in the device.

Each device family also has a device handbook, including a data sheet, which documents device features in detail. You can also see a summary of the resources for each device in the **Device** dialog box in the Quartus II software.

Carefully study the device density requirements for your design. Devices with more logic resources and higher I/O counts can implement larger and more complex designs, but at a higher cost. Smaller devices use lower static power. Select a device larger than what your design requires if you want to add more logic later in the design cycle to upgrade or expand your design, and reserve logic and memory for on-chip debugging. Consider requirements for types of dedicated logic blocks, such as memory blocks of different sizes, or digital signal processing (DSP) blocks to implement certain arithmetic functions.

If you have older designs that target an Altera device, you can use their resources as an estimate for your design. Compile existing designs in the Quartus II software with the **Auto device selected by the Fitter** option in the **Settings** dialog box. Review the resource utilization to learn which device density fits your design. Consider coding style, device architecture, and the optimization options used in the Quartus II software, which can significantly affect the resource utilization and timing performance of your design.

### Related Information

- [Planning for On-Chip Debugging Tools](#) on page 2-8  
For information about on-chip debugging.
- [Altera Product Selector](#)  
For You can refer to the Altera website to help you choose your device.
- [Selector Guides](#)  
You can review important features of each device family in the refer to the Altera website.
- [Devices and Adapters](#)  
For a list of device selection guides.
- [IP and Megafunctions](#)  
For information on how to obtain resource utilization estimates for certain configurations of Altera's IP, refer to the user guides for Altera megafunctions and IP MegaCores on the literature page of the Altera website.

## Device Migration Planning

Determine whether you want to migrate your design to another device density to allow flexibility when your design nears completion. You may want to target a smaller (and less expensive) device and then

move to a larger device if necessary to meet your design requirements. Other designers may prototype their design in a larger device to reduce optimization time and achieve timing closure more quickly, and then migrate to a smaller device after prototyping. If you want the flexibility to migrate your design, you must specify these migration options in the Quartus II software at the beginning of your design cycle.

Selecting a migration device impacts pin placement because some pins may serve different functions in different device densities or package sizes. If you make pin assignments in the Quartus II software, the Pin Migration View in the Pin Planner highlights pins that change function between your migration devices.

#### Related Information

- [Early Pin Planning and I/O Analysis](#) on page 2-5
- [Specifying Devices for Device Migration](#)

For more information about specifying the target migration devices in Quartus II Help.

## Planning for Device Programming or Configuration

Planning how to program or configure the device in your system allows system and board designers to determine what companion devices, if any, your system requires. Your board layout also depends on the type of programming or configuration method you plan to use for programmable devices. Many programming options require a JTAG interface to connect to the devices, so you might have to set up a JTAG chain on the board. Additionally, the Quartus II software uses the settings for the configuration scheme, configuration device, and configuration device voltage to enable the appropriate dual purpose pins as regular I/O pins after you complete configuration. The Quartus II software performs voltage compatibility checks of those pins during I/O assignment analysis and compilation of your design. You can use the **Configuration** tab of the **Device and Pin Options** dialog box to select your configuration scheme.

The device family handbooks describe the configuration options available for a device family. For information about programming CPLD devices, refer to your device data sheet or handbook.

#### Related Information

##### [Configuration Handbook](#)

For more details about configuration options.

## Estimating Power

You can use the Quartus II power estimation and analysis tools to provide information to PCB board and system designers. Power consumption in FPGA devices depends on the design logic, which can make planning difficult. You can estimate power before you create any source code, or when you have a preliminary version of the design source code, and then perform the most accurate analysis with the PowerPlay Power Analyzer when you complete your design.

You must accurately estimate device power consumption to develop an appropriate power budget and to design the power supplies, voltage regulators, heat sink, and cooling system. Power estimation and analysis helps you satisfy two important planning requirements:

- **Thermal**—ensure that the cooling solution is sufficient to dissipate the heat generated by the device. The computed junction temperature must fall within normal device specifications.
- **Power supply**—ensure that the power supplies provide adequate current to support device operation.

The PowerPlay Early Power Estimator (EPE) spreadsheet allows you to estimate power utilization for your design.

You can manually enter data into the EPE spreadsheet, or use the Quartus II software to generate device resource information for your design.

To manually enter data into the EPE spreadsheet, enter the device resources, operating frequency, toggle rates, and other parameters for your design. If you do not have an existing design, estimate the number of device resources used in your design, and then enter the data into the EPE spreadsheet manually.

If you have an existing design or a partially completed design, you can use the Quartus II software to generate the PowerPlay Early Power Estimator File (.txt, .csv) to assist you in completing the PowerPlay EPE spreadsheet.

The PowerPlay EPE spreadsheet includes the Import Data macro that parses the information in the PowerPlay EPE File and transfers the information into the spreadsheet. If you do not want to use the macro, you can manually transfer the data into the EPE spreadsheet. For example, after importing the PowerPlay EPE File information into the PowerPlay EPE spreadsheet, you can add device resource information. If the existing Quartus II project represents only a portion of your full design, manually enter the additional device resources you use in the final design.

Estimating power consumption early in the design cycle allows planning of power budgets and avoids unexpected results when designing the PCB.

When you complete your design, perform a complete power analysis to check the power consumption more accurately. The PowerPlay Power Analyzer tool in the Quartus II software provides an accurate estimation of power, ensuring that thermal and supply limitations are met.

#### Related Information

- [PowerPlay Power Analysis](#)  
For more information about power estimation and analysis.
- [Performing an Early Power Estimate Using the PowerPlay Early Power Estimator](#)  
For more information about generating the PowerPlay EPE File, refer to Quartus II Help.
- [PowerPlay Early Power Estimator and Power Analyzer](#)  
The PowerPlay EPE spreadsheets for each supported device family are available on the Altera website.

## Early Pin Planning and I/O Analysis

In many design environments, FPGA designers want to plan the top-level FPGA I/O pins early to help board designers begin the PCB design and layout. The I/O capabilities and board layout guidelines of the FPGA device influence pin locations and other types of assignments. If the board design team specifies an FPGA pin-out, the pin locations must be verified in the FPGA placement and routing software to avoid board design changes.

You can create a preliminary pin-out for an Altera FPGA with the Quartus II Pin Planner before you develop the source code, based on standard I/O interfaces (such as memory and bus interfaces) and any other I/O requirements for your system. The Quartus II I/O Assignment Analysis checks that the pin locations and assignments are supported in the target FPGA architecture. You can then use I/O Assignment Analysis to validate I/O-related assignments that you create or modify throughout the design process. When you compile your design in the Quartus II software, I/O Assignment Analysis runs automatically in the Fitter to validate that the assignments meet all the device requirements and generates error messages.



Early in the design process, before creating the source code, the system architect has information about the standard I/O interfaces (such as memory and bus interfaces), the IP cores in your design, and any other I/O-related assignments defined by system requirements. You can use this information with the **Early Pin Planning** feature in the Pin Planner to specify details about the design I/O interfaces. You can then create a top-level design file that includes all I/O information.

The Pin Planner interfaces with the IP core parameter editor, which allows you to create or import custom IP cores that use I/O interfaces. You can configure how to connect the functions and cores to each other by specifying matching node names for selected ports. You can create other I/O-related assignments for these interfaces or other design I/O pins in the Pin Planner, as described in this section. The Pin Planner creates virtual pin assignments for internal nodes, so internal nodes are not assigned to device pins during compilation. After analysis and synthesis of the newly generated top-level wrapper file, use the generated netlist to perform I/O Analysis with the **Start I/O Assignment Analysis** command.

You can use the I/O analysis results to change pin assignments or IP parameters even before you create your design, and repeat the checking process until the I/O interface meets your design requirements and passes the pin checks in the Quartus II software. When you complete initial pin planning, you can create a revision based on the Quartus II-generated netlist. You can then use the generated netlist to develop the top-level design file for your design, or disregard the generated netlist and use the generated Quartus II Settings File (.qsf) with your design.

During this early pin planning, after you have generated a top-level design file, or when you have developed your design source code, you can assign pin locations and assignments with the Pin Planner.

With the Pin Planner, you can identify I/O banks, voltage reference (VREF) groups, and differential pin pairings to help you through the I/O planning process. If you selected a migration device, the **Pin Migration View** highlights the pins that have changed functions in the migration device when compared to the currently selected device. Selecting the pins in the Device Migration view cross-probes to the rest of the Pin Planner, so that you can use device migration information when planning your pin assignments. You can also configure board trace models of selected pins for use in “board-aware” signal integrity reports generated with the **Enable Advanced I/O Timing** option. This option ensures that you get accurate I/O timing analysis. You can use a Microsoft Excel spreadsheet to start the I/O planning process if you normally use a spreadsheet in your design flow, and you can export a Comma-Separated Value File (.csv) containing your I/O assignments for spreadsheet use when you assign all pins.

When you complete your pin planning, you can pass pin location information to PCB designers. The Pin Planner is tightly integrated with certain PCB design EDA tools, and can read pin location changes from these tools to check suggested changes. Your pin assignments must match between the Quartus II software and your schematic and board layout tools to ensure the FPGA works correctly on the board, especially if you must make changes to the pin-out. The system architect uses the Quartus II software to pass pin information to team members designing individual logic blocks, allowing them to achieve better timing closure when they compile their design.

Start FPGA planning before you complete the HDL for your design to improve the confidence in early board layouts, reduce the chance of error, and improve the overall time to market of the design. When you complete your design, use the Fitter reports for the final sign-off of pin assignments. After compilation, the Quartus II software generates the Pin-Out File (.pin), and you can use this file to verify that each pin is correctly connected in board schematics.

#### Related Information

- [Device Migration Planning](#) on page 2-3

- **Set Up Top-Level Design File Window (Edit Menu)**  
For more information about setting up the nodes in your design, refer to Quartus II Help.
- **I/O Management**  
For more information about I/O assignment and analysis.
- **Mentor Graphics PCB Design Tools Support**
- **Cadence PCB Design Tools Support**  
For more information about passing I/O information between the Quartus II software and third-party EDA tools.

## Simultaneous Switching Noise Analysis

Simultaneous switching noise (SSN) is a noise voltage inducted onto a victim I/O pin of a device due to the switching behavior of other aggressor I/O pins in the device.

Altera provides tools for SSN analysis and estimation, including SSN characterization reports, an Early SSN Estimator (ESE) spreadsheet tool, and the SSN Analyzer in the Quartus II software. SSN often leads to the degradation of signal integrity by causing signal distortion, thereby reducing the noise margin of a system. You must address SSN with estimation early in your system design, to minimize later board design changes. When your design is complete, verify your board design by performing a complete SSN analysis of your FPGA in the Quartus II software.

### Related Information

- **Altera's Signal Integrity Center**  
For more information and device support for the ESE spreadsheet tool on the Altera website.
- **Simultaneous Switching Noise (SSN)**  
For more information about the SSN Analyzer.

## Selecting Third-Party EDA Tools

Your complete FPGA design flow may include third-party EDA tools in addition to the Quartus II software. Determine which tools you want to use with the Quartus II software to ensure that they are supported and set up properly, and that you are aware of their capabilities.

## Synthesis Tool

The Quartus II software includes integrated synthesis that supports Verilog HDL, VHDL, Altera Hardware Description Language (AHDL), and schematic design entry.

You can also use supported standard third-party EDA synthesis tools to synthesize your Verilog HDL or VHDL design, and then compile the resulting output netlist file in the Quartus II software. Different synthesis tools may give different results for each design. To determine the best tool for your application, you can experiment by synthesizing typical designs for your application and coding style. Perform placement and routing in the Quartus II software to get accurate timing analysis and logic utilization results.

The synthesis tool you choose may allow you to create a Quartus II project and pass constraints, such as the EDA tool setting, device selection, and timing requirements that you specified in your synthesis project. You can save time when setting up your Quartus II project for placement and routing.

Tool vendors frequently add new features, fix tool issues, and enhance performance for Altera devices, you must use the most recent version of third-party synthesis tools.

### Related Information

- [Incremental Compilation with Design Partitions](#) on page 2-12  
For more information on how to use incremental compilation with design partitions.
- [Volume 1: Design and Synthesis](#)  
For more information about synthesis tool flows.
- [Quartus II Software and Device Support Release Notes](#)  
For more information about the supported versions of synthesis tools in your version of the Quartus II Help.

## Simulation Tool

Altera provides the Mentor Graphics ModelSim<sup>®</sup>-Altera Starter Edition with the Quartus II software. You can also purchase the ModelSim-Altera Edition or a full license of the ModelSim software to support large designs and achieve faster simulation performance. The Quartus II software can generate both functional and timing netlist files for ModelSim and other third-party simulators.

Use the simulator version that your Quartus II software version supports for best results. You must also use the model libraries provided with your Quartus II software version. Libraries can change between versions, which might cause a mismatch with your simulation netlist.

### Related Information

- [Quartus II Software and Device Support Release Notes](#)  
For a list of the version of each simulation tool that is supported with a given version of the Quartus II software.
- [Simulation](#)  
For more information about simulation tool flows.

## Formal Verification Tools

Consider whether the Quartus II software supports the formal verification tool that you want to use, and whether the flow impacts your design and compilation stages of your design.

Using a formal verification tool can impact performance results because performing formal verification requires turning off certain logic optimizations, such as register retiming, and forces you to preserve hierarchy blocks, which can restrict optimization. Formal verification treats memory blocks as black boxes. Therefore, you must keep memory in a separate hierarchy block so other logic does not get incorporated into the black box for verification. If formal verification is important to your design, plan for limitations and restrictions at the beginning of the design cycle rather than make changes later.

### Related Information

#### [Volume 3: Verification](#)

For more information about formal verification flows and the supported tools

## Planning for On-Chip Debugging Tools

In-system debugging tools offer different advantages and trade-offs. A particular debugging tool may work better for different systems and designers.

You must evaluate on-chip debugging tools early in your design process, to ensure that your system board, Quartus II project, and design can support the appropriate tools. You can reduce debugging time and avoid making changes to accommodate your preferred debugging tools later.

If you intend to use any of these tools, you may have to plan for the tools when developing your system board, Quartus II project, and design. Consider the following debugging requirements when you plan your design:

- JTAG connections—required to perform in-system debugging with JTAG tools. Plan your system and board with JTAG ports that are available for debugging.
- Additional logic resources—required to implement JTAG hub logic. If you set up the appropriate tool early in your design cycle, you can include these device resources in your early resource estimations to ensure that you do not overload the device with logic.
- Reserve device memory—required if your tool uses device memory to capture data during system operation. To ensure that you have enough memory resources to take advantage of this debugging technique, consider reserving device memory to use during debugging.
- Reserve I/O pins—required if you use the Logic Analyzer Interface (LAI) or SignalProbe tools, which require I/O pins for debugging. If you reserve I/O pins for debugging, you do not have to later change your design or board. The LAI can multiplex signals with design I/O pins if required. Ensure that your board supports a debugging mode, in which debugging signals do not affect system operation.
- Instantiate an IP core in your HDL code—required if your debugging tool uses an Altera IP core.
- Instantiate the SignalTap II Logic Analyzer IP core—required if you want to manually connect the SignalTap II Logic Analyzer to nodes in your design and ensure that the tapped node names do not change during synthesis. You can add the analyzer as a separate design partition for incremental compilation to minimize recompilation times.

**Table 2-1: Factors to Consider When Using Debugging Tools During Design Planning Stages**

Design Planning Factor	SignalTap II Logic Analyzer	System Console	In-System Memory Content Editor	Logic Analyzer Interface (LAI)	SignalProbe	In-System Sources and Probes	Virtual JTAG IP Core
JTAG connections	Yes	Yes	Yes	Yes	—	Yes	Yes
Additional logic resources	—	Yes	—	—	—	—	Yes
Reserve device memory	Yes	Yes	—	—	—	—	—
Reserve I/O pins	—	—	—	Yes	Yes	—	—
Instantiate IP core in your HDL code	—	—	—	—	—	Yes	Yes

**Related Information**

- [System Debugging Tools Overview](#)  
For an overview of debugging tools that can help you decide which tools to use.
- [Design Debugging Using the SignalTap II Logic Analyzer](#)  
For more information on using the SignalTap II Logic Analyzer.

## Design Practices and HDL Coding Styles

When you develop complex FPGA designs, design practices and coding styles have an enormous impact on the timing performance, logic utilization, and system reliability of your device.

## Design Recommendations

Use synchronous design practices to consistently meet your design goals. Problems with asynchronous design techniques include reliance on propagation delays in a device, incomplete timing analysis, and possible glitches.

In a synchronous design, a clock signal triggers all events. When you meet all register timing requirements, a synchronous design behaves in a predictable and reliable manner for all process, voltage, and temperature (PVT) conditions. You can easily target synchronous designs to different device families or speed grades.

Clock signals have a large effect on the timing accuracy, performance, and reliability of your design. Problems with clock signals can cause functional and timing problems in your design. Use dedicated clock pins and clock routing for best results, and if you have PLLs in your target device, use the PLLs for clock inversion, multiplication, and division. For clock multiplexing and gating, use the dedicated clock control block or PLL clock switchover feature instead of combinational logic, if these features are available in your device. If you must use internally-generated clock signals, register the output of any combinational logic used as a clock signal to reduce glitches.

The Design Assistant in the Quartus II software is a design-rule checking tool that enables you to verify design issues. The Design Assistant checks your design for adherence to Altera-recommended design guidelines. You can also use third-party lint tools to check your coding style.

Consider the architecture of the device you choose so that you can use specific features in your design. For example, the control signals should use the dedicated control signals in the device architecture. Sometimes, you might need to limit the number of different control signals used in your design to achieve the best results.

### Related Information

- [About the Design Assistant](#)

For more information about running the Design Assistant, refer to Quartus II Help.

- [Recommended Design Practices](#) on page 11-1

For more information about design recommendations and using the Design Assistant.

- [www.sunburst-design.com](http://www.sunburst-design.com)

You can also refer to industry papers for more information about multiple clock design. For a good analysis, refer to *Synthesis and Scripting Techniques for Designing Multi-Asynchronous Clock Designs* under **Papers**.

## Recommended HDL Coding Styles

HDL coding styles can have a significant effect on the quality of results for programmable logic designs.

If you design memory and DSP functions, you must understand the target architecture of your device so you can use the dedicated logic block sizes and configurations. Follow the coding guidelines for inferring megafunctions and targeting dedicated device hardware, such as memory and DSP blocks.

### Related Information

- [Recommended HDL Coding Styles](#) on page 12-1

For HDL coding examples and recommendations, refer to the Recommended HDL Coding Styles chapter in volume 1 of the Quartus II Handbook. For additional tool-specific guidelines

## Managing Metastability

Metastability problems can occur in digital design when a signal is transferred between circuitry in unrelated or asynchronous clock domains, because the designer cannot guarantee that the signal meets the setup and hold time requirements during the signal transfer.

Designers commonly use a synchronization chain to minimize the occurrence of metastable events. Ensure that your design accounts for synchronization between any asynchronous clock domains. Consider using a synchronizer chain of more than two registers for high-frequency clocks and frequently-toggling data signals to reduce the chance of a metastability failure.

You can use the Quartus II software to analyze the average mean time between failures (MTBF) due to metastability when a design synchronizes asynchronous signals, and optimize your design to improve the metastability MTBF. The MTBF due to metastability is an estimate of the average time between instances when metastability could cause a design failure. A high MTBF (such as hundreds or thousands of years between metastability failures) indicates a more robust design. Determine an acceptable target MTBF given the context of your entire system and the fact that MTBF calculations are statistical estimates.

The Quartus II software can help you determine whether you have enough synchronization registers in your design to produce a high enough MTBF at your clock and data frequencies.

### Related Information

- [Managing Metastability with the Quartus II Software](#) on page 13-1  
For information about metastability analysis, reporting, and optimization features in the Quartus II software

## Planning for Hierarchical and Team-Based Design

To create a hierarchical design so that you can use compilation-time savings and performance preservation with the Quartus II software incremental compilation feature, plan for an incremental compilation flow from the beginning of your design cycle. The following subsections describe the flat compilation flow, in which the design hierarchy is flattened without design partitions, and then the incremental compilation flow that uses design partitions. Incremental compilation flows offer several advantages, but require more design planning to ensure effective results. The last subsections discuss planning an incremental compilation flow, planning design partitions, and optionally creating a design floorplan.

### Related Information

- [Quartus II Incremental Compilation for Hierarchical and Team-Based Design](#) on page 3-1  
For information about using the incremental compilation flow methodology in the Quartus II software.

## Flat Compilation Flow with No Design Partitions

In the flat compilation flow with no design partitions in the Quartus II software, the Quartus II software compiles the entire design in a “flat” netlist.

Your source code can have hierarchy, but the Quartus II software flattens your design during compilation and synthesizes all the design source code and fits in the target device whenever the software recompiles your design after any change in your design. By processing the entire design, the software performs all available logic and placement optimizations on the entire design to improve area and performance. You

can use debugging tools in an incremental design flow, such as the SignalTap II Logic Analyzer, but you do not specify any design partitions to preserve design hierarchy during compilation.

The flat compilation flow is easy to use; you do not have to plan any design partitions. However, because the Quartus II software recompiles the entire design whenever you change your design, compilation times can be slow for large devices. Additionally, you may find that the results for one part of the design change when you change a different part of your design. You run **Rapid Recompile** to preserve portions of previous placement and routing in subsequent compilations. This option can reduce your compilation time in a flat or partitioned design when you make small changes to your design.

## Incremental Compilation with Design Partitions

In an incremental compilation flow, the system architect splits a large design into partitions. When hierarchical design partitions are well chosen and placed in the device floorplan, you can speed up your design compilation time while maintaining the quality of results.

Incremental compilation preserves the compilation results and performance of unchanged partitions in the design, greatly reducing design iteration time by focusing new compilations on changed design partitions only. Incremental compilation then merges new compilation results with the previous compilation results from unchanged design partitions. Additionally, you can target optimization techniques, such as physical synthesis, to specific design partitions while leaving other partitions unchanged. You can also use empty partitions to indicate that parts of your design are incomplete or missing, while you compile the rest of your design.

Third-party IP designers can also export logic blocks to be integrated into the top-level design. Team members can work on partitions independently, which can simplify the design process and reduce compilation time. With exported partitions, the system architect must provide guidance to designers or IP providers to ensure that each partition uses the appropriate device resources. Because the designs may be developed independently, each designer has no information about the overall design or how their partition connects with other partitions. This lack of information can lead to problems during system integration. The top-level project information, including pin locations, physical constraints, and timing requirements, must be communicated to the designers of lower-level partitions before they start their design.

The system architect plans design partitions at the top level and allows third-party designs to access the top-level project framework. By designing in a copy of the top-level project (or by checking out the project files in a source control environment), the designers of the lower-level block have full information about the entire project, which helps to ensure optimal results.

When you plan your design code and hierarchy, ensure that each design entity is created in a separate file so that the entities remain independent when you make source code changes in the file. If you use a third-party synthesis tool, create separate Verilog Quartus Mapping or EDIF netlists for each design partition in your synthesis tool. You may have to create separate projects in your synthesis tool, so that the tool synthesizes each partition separately and generates separate output netlist files. The netlists are then considered the source files for incremental compilation.

### Related Information

- [Quartus II Incremental Compilation for Hierarchical and Team-Based Design](#) on page 3-1  
For more information about support for Quartus II incremental compilation.

## Planning Design Partitions and Floorplan Location Assignments

Partitioning a design for an FPGA requires planning to ensure optimal results when you integrate the partitions. Following Altera's recommendations for creating design partitions should improve the overall quality of results. For example, registering partition I/O boundaries keeps critical timing paths inside one

partition that can be optimized independently. When you specify the design partitions, you can use the Incremental Compilation Advisor to ensure that partitions meet Altera's recommendations.

If you have timing-critical partitions that are changing through the design flow, or partitions exported from another Quartus II project, you can create design floorplan assignments to constrain the placement of the affected partitions. Good partition and floorplan design helps partitions meet top-level design requirements when integrated with the rest of your design, reducing time you spend integrating and verifying the timing of the top-level design.

#### Related Information

[Best Practices for Incremental Compilation Partitions and Floorplan Assignments](#) on page 14-1  
For detailed guidelines about creating design partitions and organizing your source code, as well as information about when and how to create floorplan assignments .

#### [Analyzing and Optimizing the Design Floorplan](#)

For more information about creating floorplan assignments in the Chip Planner .

## Fast Synthesis and Early Timing Estimation

You save time when you find design issues early in the design cycle rather than in the final timing closure stages. When the first version of the design source code is complete, you might want to perform a quick compilation to create a kind of silicon virtual prototype (SVP) that you can use to perform timing analysis.

If you synthesize with the Quartus II software, you can choose to perform a **Fast** synthesis, which reduces the compilation time, but may give reduced quality of results.

Regardless of your compilation flow, you can run an early timing estimate to perform a quick placement and routing, and a timing analysis of your design. The software chooses a device automatically if required, places any LogicLock regions to create a floorplan, finds a quick initial placement for all the design logic, and provides a useful estimate of the final design performance. If you have entered timing constraints, timing analysis reports on these constraints.

If you design individual design blocks or partitions separately, you can use the Fast synthesis and early timing estimate features as you develop your design. Any issues highlighted in the lower-level design blocks are communicated to the system architect. Resolving these issues might require allocating additional device resources to the individual partition, or changing the timing budget of the partition.

Expert designers can also use fast synthesis and early timing estimation to prototype the entire design. Incomplete partitions are marked as empty in an incremental compilation flow, while the rest of the design is compiled to get an early timing estimate and detect any problems with design integration.

#### Related Information

- [Synthesis Effort logic option](#)  
For more information about Fast synthesis, refer to Quartus II Help.
- [Running a Timing Analysis](#)  
For more information about how to run an early timing estimate, refer to Quartus II Help.



## Document Revision History

Table 2-2: Document Revision History

Date	Version	Changes
2014.06.30	14.0.0	Updated format.
November 2013	13.1.0	Removed HardCopy device information.
November, 2012	12.1.0	Update for changes to early pin planning feature
June 2012	12.0.0	Editorial update.
November 2011	11.0.1	Template update.
May 2011	11.0.0	<ul style="list-style-type: none"> <li>Added link to System Design with Qsys in “Creating Design Specifications” on page 1–2</li> <li>Updated “Simultaneous Switching Noise Analysis” on page 1–8</li> <li>Updated “Planning for On-Chip Debugging Tools” on page 1–10</li> <li>Removed information from “Planning Design Partitions and Floorplan Location Assignments” on page 1–15</li> </ul>
December 2010	10.1.0	<ul style="list-style-type: none"> <li>Changed to new document template</li> <li>Updated “System Design and Standard Interfaces” on page 1–3 to include information about the Qsys system integration tool</li> <li>Added link to the Altera Product Selector in “Device Selection” on page 1–3</li> <li>Converted information into new table (Table 1–1) in “Planning for On-Chip Debugging Options” on page 1–10</li> <li>Simplified description of incremental compilation usages in “Incremental Compilation with Design Partitions” on page 1–14</li> <li>Added information about the Rapid Recompile option in “Flat Compilation Flow with No Design Partitions” on page 1–14</li> <li>Removed details and linked to Quartus II Help in “Fast Synthesis and Early Timing Estimation” on page 1–16</li> </ul>

Date	Version	Changes
July 2010	10.0.0	<ul style="list-style-type: none"> <li>• Added new section “System Design” on page 1–3</li> <li>• Removed details about debugging tools from “Planning for On-Chip Debugging Options” on page 1–10 and referred to other handbook chapters for more information</li> <li>• Updated information on recommended design flows in “Incremental Compilation with Design Partitions” on page 1–14 and removed “Single-Project Versus Multiple-Project Incremental Flows” heading</li> <li>• Merged the “Planning Design Partitions” section with the “Creating a Design Floorplan” section. Changed heading title to “Planning Design Partitions and Floorplan Location Assignments” on page 1–15</li> <li>• Removed “Creating a Design Floorplan” section</li> <li>• Removed “Referenced Documents” section</li> <li>• Minor updates throughout chapter</li> </ul>
November 2009	9.1.0	<ul style="list-style-type: none"> <li>• Added details to “Creating Design Specifications” on page 1–2</li> <li>• Added details to “Intellectual Property Selection” on page 1–2</li> <li>• Updated information on “Device Selection” on page 1–3</li> <li>• Added reference to “Device Migration Planning” on page 1–4</li> <li>• Removed information from “Planning for Device Programming or Configuration” on page 1–4</li> <li>• Added details to “Early Power Estimation” on page 1–5</li> <li>• Updated information on “Early Pin Planning and I/O Analysis” on page 1–6</li> <li>• Updated information on “Creating a Top-Level Design File for I/O Analysis” on page 1–8</li> <li>• Added new “Simultaneous Switching Noise Analysis” section</li> <li>• Updated information on “Synthesis Tools” on page 1–9</li> <li>• Updated information on “Simulation Tools” on page 1–9</li> <li>• Updated information on “Planning for On-Chip Debugging Options” on page 1–10</li> <li>• Added new “Managing Metastability” section</li> <li>• Changed heading title “Top-Down Versus Bottom-Up Incremental Flows” to “Single-Project Versus Multiple-Project Incremental Flows”</li> <li>• Updated information on “Creating a Design Floorplan” on page 1–18</li> <li>• Removed information from “Fast Synthesis and Early Timing Estimation” on page 1–18</li> </ul>
March 2009	9.0.0	<ul style="list-style-type: none"> <li>• No change to content</li> </ul>
November 2008	8.1.0	<ul style="list-style-type: none"> <li>• Changed to 8-1/2 x 11 page size. No change to content.</li> </ul>

Date	Version	Changes
May 2008	8.0.0	<ul style="list-style-type: none"><li>• Organization changes</li><li>• Added “Creating Design Specifications” section</li><li>• Added reference to new details in the In-System Design Debugging section of volume 3</li><li>• Added more details to the “Design Practices and HDL Coding Styles” section</li><li>• Added references to the new Best Practices for Incremental Compilation and Floorplan Assignments chapter</li><li>• Added reference to the Quartus II Language Templates</li></ul>

**Related Information**[Quartus II Handbook Archive](#)

For previous versions of the Quartus II Handbook

2014.12.15

QII5V1



Subscribe



Send Feedback

## About Quartus II Incremental Compilation

This manual provides information and design scenarios to help you partition your design to take advantage of the Quartus® II incremental compilation feature.

The ability to iterate rapidly through FPGA design and debugging stages is critical. The Quartus II software introduced the FPGA industry's first true incremental design and compilation flow, with the following benefits:

- Preserves the results and performance for unchanged logic in your design as you make changes elsewhere.
- Reduces design iteration time by an average of 75% for small changes in large designs, so that you can perform more design iterations per day and achieve timing closure efficiently.
- Facilitates modular hierarchical and team-based design flows, as well as design reuse and intellectual property (IP) delivery.

Quartus II incremental compilation supports the Arria®, Stratix®, and Cyclone® series of devices.

## Deciding Whether to Use an Incremental Compilation Flow

The Quartus II incremental compilation feature enhances the standard Quartus II design flow by allowing you to preserve satisfactory compilation results and performance of unchanged blocks of your design.

### Flat Compilation Flow with No Design Partitions

In the flat compilation flow with no design partitions, all the source code is processed and mapped during the Analysis and Synthesis stage, and placed and routed during the Fitter stage whenever the design is recompiled after a change in any part of the design. One reason for this behavior is to ensure optimal push-button quality of results. By processing the entire design, the Compiler can perform global optimizations to improve area and performance.

You can use a flat compilation flow for small designs, such as designs in CPLD devices or low-density FPGA devices, when the timing requirements are met easily with a single compilation. A flat design is satisfactory when compilation time and preserving results for timing closure are not concerns.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



**Related Information**[Reducing Compilation Time documentation](#)**Incremental Capabilities Available When A Design Has No Partitions**

The Quartus II software has incremental compilation features available even when you do not partition your design, including Smart Compilation, Rapid Recompile, and incremental debugging. These features work in either an incremental or flat compilation flow.

**With Smart Compilation**

In any Quartus II compilation flow, you can use Smart Compilation to allow the Compiler to determine which compilation stages are required, based on the changes made to the design since the last smart compilation, and then skip any stages that are not required.

For example, when Smart Compilation is turned on, the Compiler skips the Analysis and Synthesis stage if all the design source files are unchanged. When Smart Compilation is turned on, if you make any changes to the logic of a design, the Compiler does not skip any compilation stage. You can turn on Smart Compilation on the **Compilation Process Settings** page of the **Setting** dialog box.

**Note:** Arria 10 devices do not support the smart compilation feature.

**Related Information**[Smart Compilation online help](#)**With Rapid Recompile**

The Quartus II software also includes a Rapid Recompile feature that instructs the Compiler to reuse the compatible compilation results if most of the design has not changed since the last compilation. This feature reduces compilation times for small and isolated design changes. You do not have control over which parts of the design are recompiled using this option; the Compiler determines which parts of the design must be recompiled. The Rapid Recompile feature preserves performance and can save compilation time by reducing the amount of changed logic that must be recompiled.

**Related Information**[About Rapid Recompile online help](#)**With SignalTap II Logic Analyzer**

During the debugging stage of the design cycle, you can add the SignalTap<sup>®</sup> II Logic Analyzer to your design, even if the design does not have partitions. To preserve the compilation netlist for the entire design, instruct the software to reuse the compilation results for the automatically-created "Top" partition that contains the entire design.

**Related Information**[About SignalTap II Logic Analyzer online help](#)**Incremental Compilation Flow With Design Partitions**

In the standard incremental compilation design flow, the top-level design is divided into design partitions, which can be compiled and optimized together in the top-level Quartus II project. You can preserve fitting results and performance for completed partitions while other parts of the design are changing, which reduces the compilation times for each design iteration.

If you use the incremental compilation feature at any point in your design flow, it is easier to accommodate the guidelines for partitioning a design and creating a floorplan if you start planning for incremental compilation at the beginning of your design cycle.

Incremental compilation is recommended for large designs and high resource densities when preserving results is important to achieve timing closure. The incremental compilation feature also facilitates team-based design flows that allow designers to create and optimize design blocks independently, when necessary.

To take advantage of incremental compilation, start by splitting your design along any of its hierarchical boundaries into design blocks to be compiled incrementally, and set each block as a design partition. The Quartus II software synthesizes each individual hierarchical design partition separately, and then merges the partitions into a complete netlist for subsequent stages of the compilation flow. When recompiling your design, you can use source code, post-synthesis results, or post-fitting results to preserve satisfactory results for each partition.

In a team-based environment, part of your design may be incomplete, or it may have been developed by another designer or IP provider. In this scenario, you can add the completed partitions to the design incrementally. Alternatively, other designers or IP providers can develop and optimize partitions independently and the project lead can later integrate the partitions into the top-level design.

**Related Information**

- [Team-Based Design Flows and IP Delivery](#) on page 3-6
- [Incremental Compilation Summary](#) on page 3-7
- [Best Practices for Incremental Compilation Partitions and Floorplan Assignments documentation](#) on page 14-1

**Impact of Using Incremental Compilation with Design Partitions**

**Table 3-1: Impact Summary of Using Incremental Compilation**

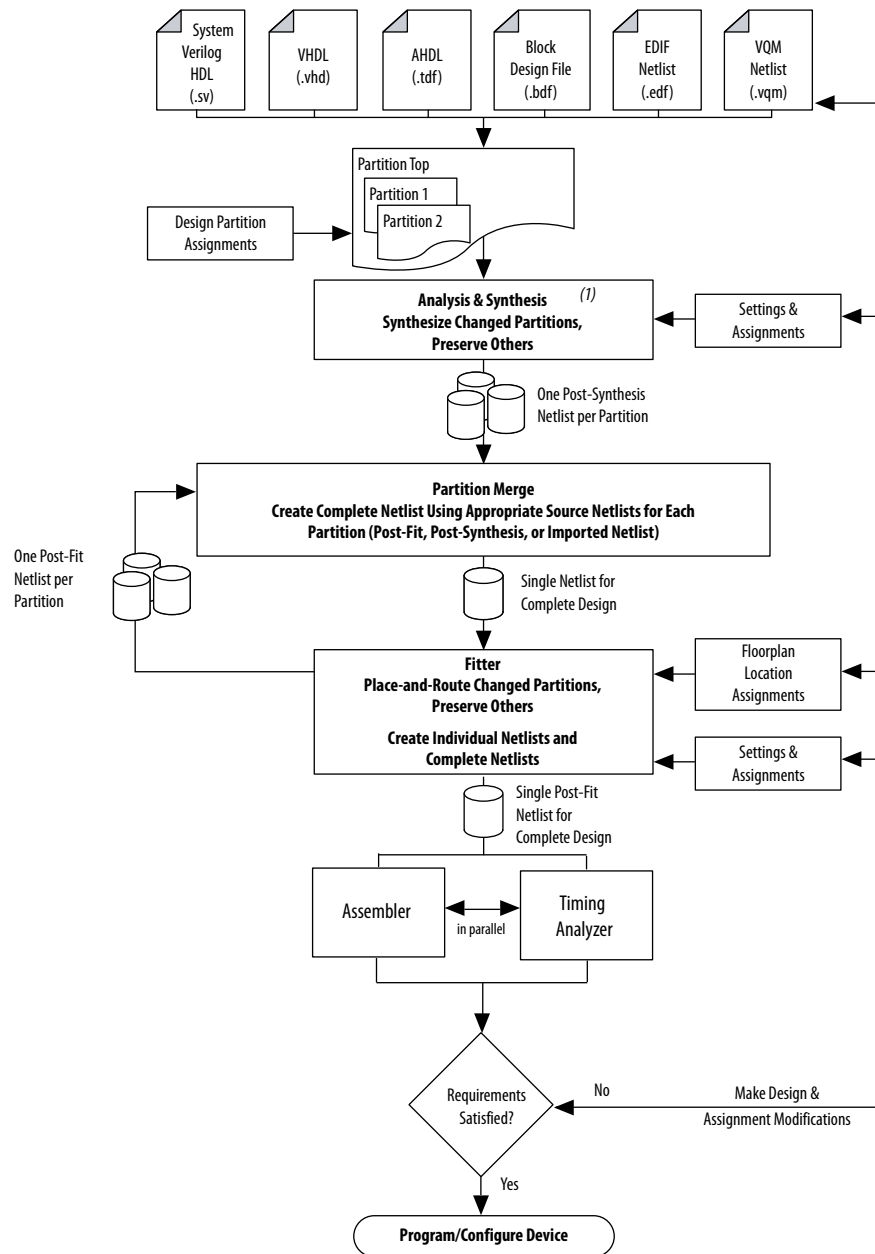
Characteristic	Impact of Incremental Compilation with Design Partitions
Compilation Time Savings	Typically saves an average of 75% of compilation time for small design changes in large designs when post-fit netlists are preserved; there are savings in both Quartus II Integrated Synthesis and the Fitter. <sup>(1)</sup>
Performance Preservation	Excellent performance preservation when timing critical paths are contained within a partition, because you can preserve post-fitting information for unchanged partitions.
Node Name Preservation	Preserves post-fitting node names for unchanged partitions.

<sup>(1)</sup> Quartus II incremental compilation does not reduce processing time for the early "pre-fitter" operations, such as determining pin locations and clock routing, so the feature cannot reduce compilation time if runtime is dominated by those operations.

Characteristic	Impact of Incremental Compilation with Design Partitions
Area Changes	The area (logic resource utilization) might increase because cross-boundary optimizations are limited, and placement and register packing are restricted.
$f_{MAX}$ Changes	The design's maximum frequency might be reduced because cross-boundary optimizations are limited. If the design is partitioned and the floorplan location assignments are created appropriately, there might be no negative impact on $f_{MAX}$ .

## Quartus II Design Stages for Incremental Compilation

Figure 3-1: Design Stages for Incremental Compilation



**Note:** When you use EDIF or VQM netlists created by third-party EDA synthesis tools, Analysis and Synthesis creates the design database, but logic synthesis and technology mapping are performed only for black boxes.

### Analysis and Synthesis Stage

The figure above shows a top-level partition and two lower-level partitions. If any part of the design changes, Analysis and Synthesis processes the changed partitions and keeps the existing netlists for the



unchanged partitions. After completion of Analysis and Synthesis, there is one post-synthesis netlist for each partition.

### Partition Merge Stage

The Partition Merge step creates a single, complete netlist that consists of post-synthesis netlists, post-fit netlists, and netlists exported from other Quartus II projects, depending on the netlist type that you specify for each partition.

### Fitter Stage

The Fitter then processes the merged netlist, preserves the placement and routing of unchanged partitions, and refits only those partitions that have changed. The Fitter generates the complete netlist for use in future stages of the compilation flow, including timing analysis and programming file generation, which can take place in parallel if more than one processor is enabled for use in the Quartus II software. The Fitter also generates individual netlists for each partition so that the Partition Merge stage can use the post-fit netlist to preserve the placement and routing of a partition, if specified, for future compilations.

### How to Compare Incremental Compilation Results with Flat Design Results

If you define partitions, but want to check your compilation results without partitions in a “what if” scenario, you can direct the Compiler to ignore all partitions assignments in your project and compile the design as a “flat” netlist. When you turn on the **Ignore partitions assignments during compilation** option on the **Incremental Compilation** page, the Quartus II software disables all design partition assignments in your project and runs a full compilation ignoring all partition boundaries and netlists. Turning off the **Ignore partitions assignments during compilation** option restores all partition assignments and netlists for subsequent compilations.

#### Related Information

- [Incremental Compilation Page online help](#)
- [Design Partition Properties Dialog Box online help](#)

## Team-Based Design Flows and IP Delivery

The Quartus II software supports various design flows to enable team-based design and third-party IP delivery. A top-level design can include one or more partitions that are designed or optimized by different designers or IP providers, as well as partitions that will be developed as part of a standard incremental methodology.

### With a Single Quartus II Project

In a team-based environment, part of your design may be incomplete because it is being developed elsewhere. The project lead or system architect can create empty placeholders in the top-level design for partitions that are not yet complete. Designers or IP providers can create and verify HDL code separately, and then the project lead later integrates the code into the single top-level Quartus II project. In this scenario, you can add the completed partitions to the design incrementally, however, the design flow allows all design optimization to occur in the top-level design for easiest design integration. Altera recommends using a single Quartus II project whenever possible because using multiple projects can add significant up-front and debugging time to the development cycle.

### With Multiple Quartus II Projects

Alternatively, partition designers can design their partition in a copy of the top-level design or in a separate Quartus II project. Designers export their completed partition as either a post-synthesis netlist or

optimized placed and routed netlist, or both, along with assignments such as LogicLock™ regions, as appropriate. The project lead then integrates each design block as a design partition into the top-level design. Altera recommends that designers export and reuse post-synthesis netlists, unless optimized post-fit results are required in the top-level design, to simplify design optimization.

### Additional Planning Needed

Teams with a bottom-up design approach often want to optimize placement and routing of design partitions independently and may want to create separate Quartus II projects for each partition. However, optimizing design partitions in separate Quartus II projects, and then later integrating the results into a top-level design, can have the following potential drawbacks that require careful planning:

- Achieving timing closure for the full design may be more difficult if you compile partitions independently without information about other partitions in the design. This problem may be avoided by careful timing budgeting and special design rules, such as always registering the ports at the module boundaries.
- Resource budgeting and allocation may be required to avoid resource conflicts and overuse. Creating a floorplan with LogicLock regions is recommended when design partitions are developed independently in separate Quartus II projects.
- Maintaining consistency of assignments and timing constraints can be more difficult if there are separate Quartus II projects. The project lead must ensure that the top-level design and the separate projects are consistent in their assignments.

### Collaboration on a Team-Based Design

A unique challenge of team-based design and IP delivery for FPGAs is the fact that the partitions being developed independently must share a common set of resources. To minimize issues that might arise from sharing a common set of resources, you can design partitions within a single Quartus II project or a copy of the top-level design. A common project ensures that designers have a consistent view of the top-level project framework.

For timing-critical partitions being developed and optimized by another designer, it is important that each designer has complete information about the top-level design in order to maintain timing closure during integration, and to obtain the best results. When you want to integrate partitions from separate Quartus II projects, the project lead can perform most of the design planning, and then pass the top-level design constraints to the partition designers. Preferably, partition designers can obtain a copy of the top-level design by checking out the required files from a source control system. Alternatively, the project lead can provide a copy of the top-level project framework, or pass design information using Quartus II-generated design partition scripts. In the case that a third-party designer has no information about the top-level design, developers can export their partition from an independent project if required.

#### Related Information

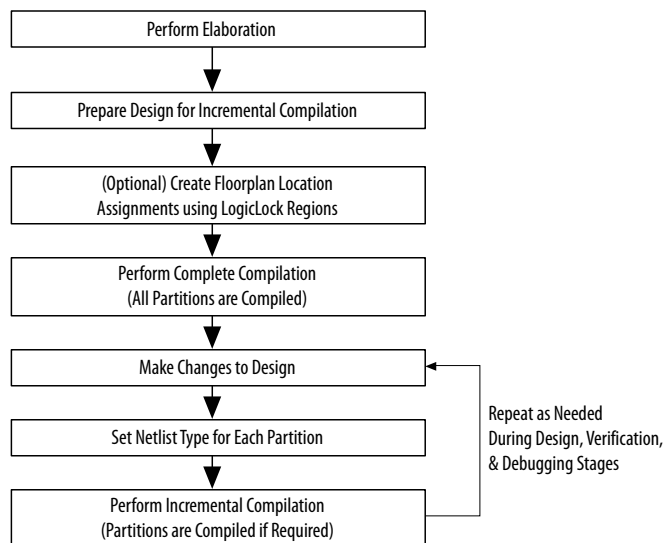
- [Exporting Design Partitions from Separate Quartus II Projects](#) on page 3-31
- [Project Management— Making the Top-Level Design Available to Other Designers](#) on page 3-33

## Incremental Compilation Summary

### Incremental Compilation Single Quartus II Project Flow

The figure illustrates the incremental compilation design flow when all partitions are contained in one top-level design.

Figure 3-2: Top-Down Design Flow



## Steps for Incremental Compilation

For an interactive introduction to implementing an incremental compilation design flow, refer to the **Getting Started Tutorial** on the Help menu in the Quartus II software.

### Related Information

[Using the Incremental Compilation Design Flow online help](#)

## Preparing a Design for Incremental Compilation

1. Elaborate your design, or run any compilation flow (such as a full compilation) that includes the elaboration step. Elaboration is the part of the synthesis process that identifies your design's hierarchy.
2. Designate specific instances in the design hierarchy as design partitions.
3. If required for your design flow, create a floorplan with LogicLock regions location assignments for timing-critical partitions that change with future compilations. Assigning a partition to a physical region on the device can help maintain quality of results and avoid conflicts in certain situations.

### Related Information

- [Creating Design Partitions](#) on page 3-9
- [Creating a Design Floorplan With LogicLock Regions](#) on page 3-47

## Compiling a Design Using Incremental Compilation

The first compilation after making partition assignments is a full compilation, and prepares the design for subsequent incremental compilations. In subsequent compilations of your design, you can preserve satisfactory compilation results and performance of unchanged partitions with the **Netlist Type** setting in the Design Partitions window. The **Netlist Type** setting determines which type of netlist or source file the Partition Merge stage uses in the next incremental compilation. You can choose the Source File, Post-Synthesis netlist, or Post-Fit netlist.

#### Related Information

[Specifying the Level of Results Preservation for Subsequent Compilations](#) on page 3-25

## Creating Design Partitions

There are several ways to designate a design instance as a design partition.

#### Related Information

[Deciding Which Design Blocks Should Be Design Partitions](#) on page 3-20

### Creating Design Partitions in the Project Navigator

You can right-click an instance in the list under the **Hierarchy** tab in the Project Navigator and use the sub-menu to create and delete design partitions.

#### Related Information

[Creating Design Partitions online help](#)

### Creating Design Partitions in the Design Partitions Window

The Design Partitions window, available from the Assignments menu, allows you to create, delete, and merge partitions, and is the main window for setting the netlist type to specify the level of results preservation for each partition on subsequent compilations.

The Design Partitions window also lists recommendations at the bottom of the window with links to the Incremental Compilation Advisor, where you can view additional recommendations about partitions. The **Color** column indicates the color of each partition as it appears in the Design Partition Planner and Chip Planner.

You can right-click a partition in the window to perform various common tasks, such as viewing property information about a partition, including the time and date of the compilation netlists and the partition statistics.

When you create a partition, the Quartus II software automatically generates a name based on the instance name and hierarchy path. You can edit the partition name in the Design Partitions Window so that you avoid referring to them by their hierarchy path, which can sometimes be long. This is especially useful when using command-line commands or assignments, or when you merge partitions to give the partition a meaningful name. Partition names can be from 1 to 1024 characters in length and must be unique. The name can consist of alphanumeric characters and the pipe ( | ), colon ( : ), and underscore ( \_ ) characters.

#### Related Information

- [Netlist Type for Design Partitions](#) on page 3-25
- [Creating Design Partitions online help](#)

### Creating Design Partitions With the Design Partition Planner

The Design Partition Planner allows you to view design connectivity and hierarchy, and can assist you in creating effective design partitions that follow Altera's guidelines.

The Design Partition Planner displays a visual representation of design connectivity and hierarchy, as well as partitions and entity relationships. You can explore the connectivity between entities in the design, evaluate existing partitions with respect to connectivity between entities, and try new partitioning schemes in "what if" scenarios.

When you extract design blocks from the top-level design and drag them into the Design Partition Planner, connection bundles are drawn between entities, showing the number of connections existing between pairs of entities. In the Design Partition Planner, you can then set extracted design blocks as design partitions.

The Design Partition Planner also has an **Auto-Partition** feature that creates partitions based on the size and connectivity of the hierarchical design blocks.

#### Related Information

[Using the Design Partition Planner online help](#)

[Best Practices for Incremental Compilation Partitions and Floorplan Assignments documentation](#) on page 14-1

## Creating Design Partitions With Tcl Scripting

You can also create partitions with Tcl scripting commands.

#### Related Information

[Scripting Support](#) on page 3-55

## Automatically-Generated Partitions

The Compiler creates some partitions automatically as part of the compilation process, which appear in some post-compilation reports. For example, the `sld_hub` partition is created for tools that use JTAG hub connections, such as the SignalTap II Logic Analyzer. The `hard_block` partition is created to contain certain "hard" or dedicated logic blocks in the device that are implemented in a separate partition so that they can be shared throughout the design.

## Common Design Scenarios Using Incremental Compilation

#### Related Information

[Steps for Incremental Compilation](#) on page 3-8

## Reducing Compilation Time When Changing Source Files for One Partition

Scenario background: You set up your design to include partitions for several of the major design blocks, and now you have just performed a lengthy compilation of the entire design. An error is found in the HDL source file for one partition and it is being fixed. Because the design is currently meeting timing requirements, and the fix is not expected to affect timing performance, it makes sense to compile only the affected partition and preserve the rest of the design.

Use the flow in this example to update the source file in one partition without having to recompile the other parts of the design. To reduce the compilation time, instruct the software to reuse the post-fit netlists for the unchanged partitions. This flow also preserves the performance of these blocks, which reduces additional timing closure efforts.

Perform the following steps to update a single source file:

1. Apply and save the fix to the HDL source file.
2. On the Assignments menu, open the **Design Partitions window**.
3. Change the netlist type of each partition, including the top-level entity, to **Post-Fit** to preserve as much as possible for the next compilation.

- The Quartus II software recompiles partitions by default when changes are detected in a source file. You can refer to the Partition Dependent Files table in the Analysis and Synthesis report to determine which partitions were recompiled. If you change an assignment but do not change the logic in a source file, you can set the netlist type to **Source File** for that partition to instruct the software to recompile the partition's source design files and its assignments.
4. Click **Start Compilation** to incrementally compile the fixed HDL code. This compilation should take much less time than the initial full compilation.
  5. Simulate the design to ensure that the error is fixed, and use the TimeQuest Timing Analyzer report to ensure that timing results have not degraded.

#### Related Information

[List of Compilation and Simulation Reports online help](#)

## Optimizing a Timing-Critical Partition

Scenario background: You have just performed a lengthy full compilation of a design that consists of multiple partitions. The TimeQuest Timing Analyzer reports that the clock timing requirement is not met, and you have to optimize one particular partition. You want to try optimization techniques such as raising the Placement Effort Multiplier, enabling Physical Synthesis, and running Design Space Explorer II. Because these techniques all involve significant compilation time, you should apply them to only the partition in question.

Use the flow in this example to optimize the results of one partition when the other partitions in the design have already met their requirements. You can use this flow iteratively to lock down the performance of one partition, and then move on to optimization of another partition.

Perform the following steps to preserve the results for partitions that meet their timing requirements, and to recompile a timing-critical partition with new optimization settings:

1. Open the **Design Partitions** window.
2. For the partition in question, set the netlist type to **Source File**.
  - If you change a setting that affects only the Fitter, you can save additional compilation time by setting the netlist type to **Post-Synthesis** to reuse the synthesis results and refit the partition.
3. For the remaining partitions (including the top-level entity), set the netlist type to **Post-Fit**.
  - You can optionally set the **Fitter Preservation Level** on the **Advanced** tab in the **Design Partitions Properties** dialog box to **Placement** to allow for the most flexibility during routing.
4. Apply the desired optimization settings.
5. Click **Start Compilation** to perform incremental compilation on the design with the new settings. During this compilation, the Partition Merge stage automatically merges the critical partition's new synthesis netlist with the post-fit netlists of the remaining partitions. The Fitter then refits only the required partition. Because the effort is reduced as compared to the initial full compilation, the compilation time is also reduced.

To use Design Space Explorer II, perform the following steps:

1. Repeat steps 1–3 of the previous procedure.
2. Save the project and run Design Space Explorer II.

## Adding Design Logic Incrementally or Working With an Incomplete Design

Scenario background: You have one or more partitions that are known to be timing-critical in your full design. You want to focus on developing and optimizing this subset of the design first, before adding the rest of the design logic.

Use this flow to compile a timing-critical partition or partitions in isolation, optionally with extra optimizations turned on. After timing closure is achieved for the critical logic, you can preserve its content and placement and compile the remaining partitions with normal or reduced optimization levels. For example, you may want to compile an IP block that comes with instructions to perform optimization before you incorporate the rest of your custom logic.

To implement this design flow, perform the following steps:

1. Partition the design and create floorplan location assignments. For best results, ensure that the top-level design includes the entire project framework, even if some parts of the design are incomplete and are represented by an empty wrapper file.
2. For the partitions to be compiled first, in the Design Partitions window, set the netlist type to **Source File**.
3. For the remaining partitions, set the netlist type to **Empty**.
4. To compile with the desired optimizations turned on, click **Start Compilation**.
5. Check the Timing Analyzer reports to ensure that timing requirements are met. If so, proceed to step 6. Otherwise, repeat steps 4 and 5 until the requirements are met.
6. In the Design Partitions window, set the netlist type to **Post-Fit** for the first partitions. You can set the **Fitter Preservation Level** on the **Advanced** tab in the **Design Partitions Properties** dialog box to **Placement** to allow more flexibility during routing if exact placement and routing preservation is not required.
7. Change the netlist type from **Empty** to **Source File** for the remaining partitions, and ensure that the completed source files are added to the project.
8. Set the appropriate level of optimizations and compile the design. Changing the optimizations at this point does not affect any fitted partitions, because each partition has its netlist type set to **Post-Fit**.
9. Check the Timing Analyzer reports to ensure that timing requirements are met. If not, make design or option changes and repeat step 8 and step 9 until the requirements are met.

The flow in this example is similar to design flows in which a module is implemented separately and is later merged into the top-level. Generally, optimization in this flow works only if each critical path is contained within a single partition. Ensure that if there are any partitions representing a design file that is missing from the project, you create a placeholder wrapper file to define the port interface.

#### Related Information

- [Designing in a Team-Based Environment](#) on page 3-41
- [Deciding Which Design Blocks Should Be Design Partitions](#) on page 3-20
- [Empty Partitions](#) on page 3-33

## Debugging Incrementally With the SignalTap II Logic Analyzer

Scenario background: Your design is not functioning as expected, and you want to debug the design using the SignalTap II Logic Analyzer. To maintain reduced compilation times and to ensure that you do not negatively affect the current version of your design, you want to preserve the synthesis and fitting results and add the SignalTap II Logic Analyzer to your design without recompiling the source code.

Use this flow to reduce compilation times when you add the logic analyzer to debug your design, or when you want to modify the configuration of the SignalTap II File without modifying your design logic or its placement.

It is not necessary to create design partitions in order to use the SignalTap II incremental compilation feature. The SignalTap II Logic Analyzer acts as its own separate design partition.

Perform the following steps to use the SignalTap II Logic Analyzer in an incremental compilation flow:

1. Open the Design Partitions window.
2. Set the netlist type to **Post-fit** for all partitions to preserve their placement.
  - The netlist type for the top-level partition defaults to **Source File**, so be sure to change this “Top” partition in addition to any design partitions that you have created.
3. If you have not already compiled the design with the current set of partitions, perform a full compilation. If the design has already been compiled with the current set of partitions, the design is ready to add the SignalTap II Logic Analyzer.
4. Set up your SignalTap II File using the **post-fitting** filter in the **Node Finder** to add signals for logic analysis. This allows the Fitter to add the SignalTap II logic to the post-fit netlist without modifying the design results.

To add signals from the pre-synthesis netlist, set the partition’s netlist type to **Source File** and use the **presynthesis** filter in the **Node Finder**. This allows the software to resynthesize the partition and to tap directly to the pre-synthesis node names that you choose. In this case, the partition is resynthesized and refit, so the placement is typically different from previous fitting results.

#### Related Information

[Design Debugging Using the SignalTap II Embedded Logic Analyzer documentation](#)

## Functional Safety IP Implementation

In functional safety designs, recertification is required when logic is modified in safety or standard areas of the design. Recertification is required because the FPGA programming file has changed. You can reduce the amount of required recertification if you use the functional safety separation flow in the Quartus II software. By partitioning your safety IP (SIP) from standard logic, you ensure that the safety critical areas of the design remain the same when the standard areas in your design are modified. The safety-critical areas remain the same at the bit level.

The functional safety separation flow supports only Cyclone IV and Cyclone V device families.

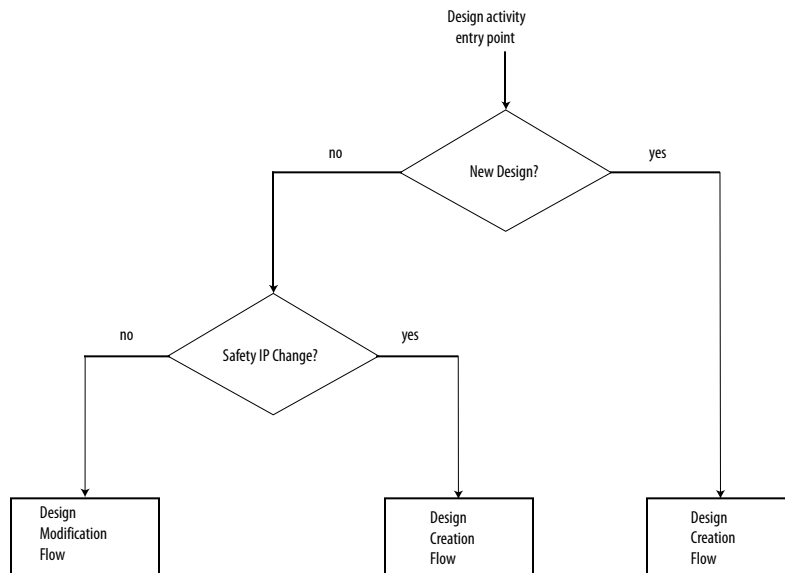
### Software Tool Impact on Safety

The Quartus II software can partition your design into safety partitions and standard partitions, but the Quartus II software does not perform any online safety-related functionality. The Quartus II software generates a bitstream that performs the safety functions. For the purpose of compliance with a functional safety standard, the Quartus II software should be considered as an offline support tool.

### Functional Safety Separation Flow

The functional safety separation flow consists of two separate work flows. The design creation flow and the design modification flow both use incremental compilation, but the two flows have different use-case scenarios.



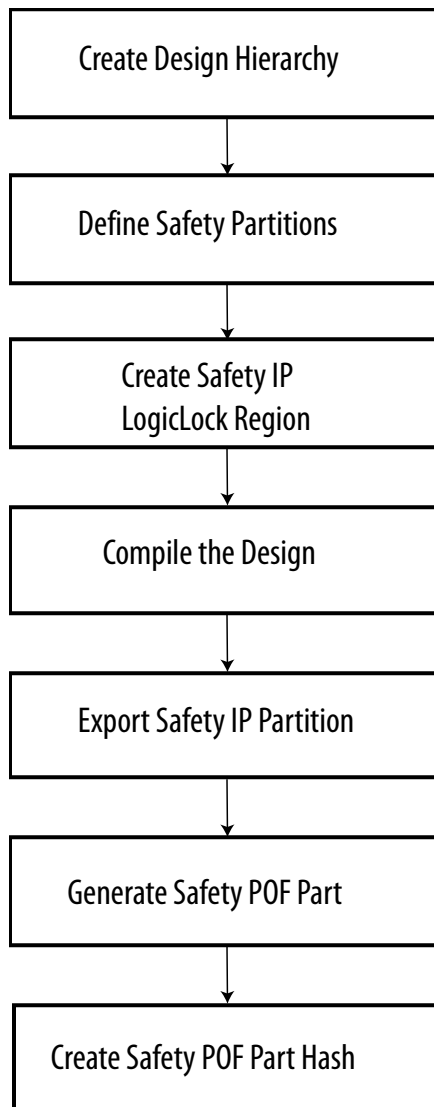
**Figure 3-3: Functional Safety Separation Flow**

### Design Creation Flow

The design creation flow describes the necessary steps for initial design creation in a way that allows you to modify your design. Some of the steps are architectural constraints and the remaining steps are steps that you need to perform in the Quartus II software. Use the design creation flow for the first pass certification of your product.

When you make modifications to the safety IP in your design, you must use the design creation flow.

Figure 3-4: Design Creation Flow



The design creation flow becomes active when you have a valid safety IP partition in your Quartus II project and that safety IP partition does not have place and route data from a previous compile. In the design creation flow, the Assembler generates a Partial Settings Mask (**.psm**) file for each safety IP partition. Each **.psm** file contains a list of programming bits for its respective safety IP partition.

The Quartus II software determines whether to use the design creation flow or design modification flow on a per partition basis. It is possible to have multiple safety IP partitions in a design where some are running the design creation flow and others are running the design modification flow.

To reset the complete design to the design creation flow, remove the previous place and route data by cleaning the project (removing the dbs). Alternatively, use the partition import flow, to selectively reset the design. You can remove the netlists for the imported safety IP partitions individually using the **Design Partitions** window.

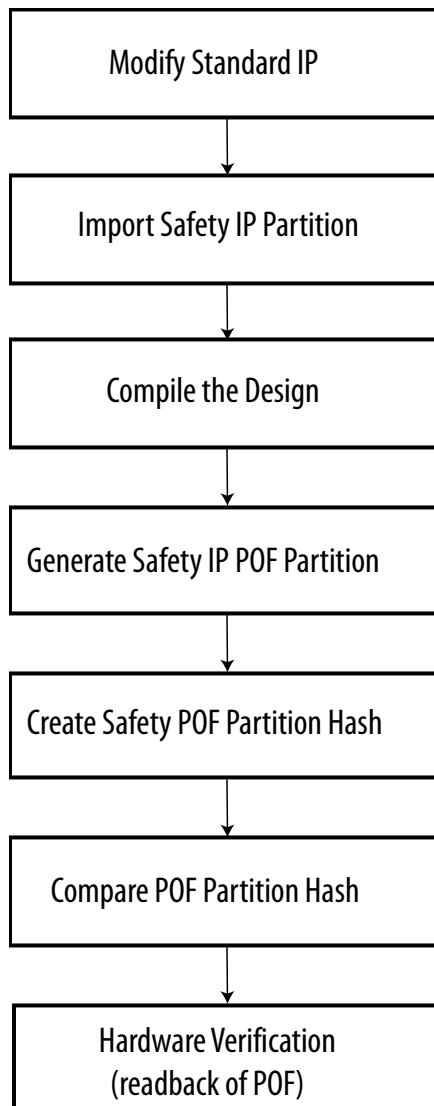
**Related Information**

- [Exporting and Importing Your Safety IP](#) on page 3-19
- [Design Partitions Window online help](#)

**Design Modification Flow**

The design modification flow describes the necessary steps to make modifications to the standard IP in your design. This flow ensures that the previously compiled safety IP that the project uses remains unchanged when you change or compile standard IP.

Use the design modification flow only after you qualify your design in the design creation flow.

**Figure 3-5: Design Modification Flow**

When the design modification flow is active for a safety IP partition, the Fitter runs in Strict Preservation mode for that partition. The Assembler performs run-time checks that compare the Partial Settings Mask information matches the **.psm** file generated in the design creation flow. If the Assembler detects a

mismatch, a "Bad Mask!" or "ASM\_STRICT\_PRESERVATION\_BITS\_UTILITY::compare\_masked\_byte\_array failed" internal error message is shown. If you see either error message while compiling your design, contact [Altera support](#) for assistance.

When a change is made to any HDL source file that belongs to a safety IP, the default behavior of the Quartus II software is to resynthesize and perform a clean place and route for that partition, which then activates the design creation flow for that partition. To change this default behavior and keep the design modification flow active, do the following:

- Use the partition export/import flow.

or

- Use the Design Partitions window to modify the design partition properties and turn on **Ignore changes in source files and strictly use the specified netlist, if available**.

The Fitter applies the same design flow to all partitions that belong to the same safety IP. If more than one safety IP is used in the design, the Fitter may evoke different flows for different safety IPs.

**Note:** If your safety IP is a sub-block in a Qsys system, every time you regenerate HDL for the Qsys system, the timestamp for the safety IP HDL changes. This results in resynthesis of the safety IP, unless the default behavior (described above) is changed.

#### Related Information

- [Exporting and Importing Your Safety IP](#) on page 3-19
- [Design Partitions Window online help](#)

## How to Turn On the Functional Safety Separation Flow

Every safety-related IP component in your design should be implemented in a partition(s) so the safety IPs are protected from recompilation. The global assignment `PARTITION_ENABLE_STRICT_PRESERVATION` is used to identify safety IP in your design.

```
set_global_assignment -name PARTITION_ENABLE_STRICT_PRESERVATION <ON/OFF> -  
section_id <partition_name>
```

When this global assignment is designated as ON for a partition, the partition is protected from recompilation, exported as a safety IP, and included in the safety IP POF mask. Specifying the value as ON for any partition turns on the functional safety separation flow.

When this global assignment is designated as OFF, the partition is considered as standard IP or as not having a `PARTITION_ENABLE_STRICT_PRESERVATION` assignment at all. Logic that is not assigned to a partition is considered as part of the top partition and treated as standard logic.

**Note:** Only partitions and I/O pins can be assigned to SIP.

A partition assigned to safety IP can contain safety logic only. If the parent partition is assigned to a safety IP, then all the child partitions for this parent partition are considered as part of the safety IP. If you do not explicitly specify a child partition as a safety IP, a critical warning notifies you that the child partition is treated as part of a safety IP.

A design can contain several safety IPs. All the partitions containing logic that implements a single safety IP function should belong with the same top-level parent partition.

You can also turn on the functional safety separation flow from the **Design Partition Properties** dialog box. Click the **Advanced** tab and turn on **Allow partition to be strictly preserved for safety**.

When the functional safety separation flow is active, you can view which partitions in your design have the Strict Preservation property turned on. The **Design Partitions** window displays a on or off value for safety IP in your design (in the **Strict Preservation** column).

#### Related Information

- [Design Partition Properties Dialog box online help](#)
- [Design Partitions Window online help](#)

### Preservation of Device Resources

The preservation of the partition's netlist atoms and the atoms placement and routing, in the design modification flow, is done by setting the netlist type to **Post-fit** with the Fitter preservation level set to **Placement and Routing Preserved**.

### Preservation of Placement in the Device with LogicLock

In order to fix the safety IP logic into specific areas of the device, you should define LogicLock regions. By using preserved LogicLock regions, device placement is reserved for the safety IP to prevent standard logic from being placed into the unused resources of the safety IP region. You establish a fixed size and origin to ensure location preservation. You need to use LogicLock to ensure a valid safety IP POF mask is generated when you turn on the functional safety separation flow. The POF comparison tool for functional safety can check that the safety region is unchanged between compiles. A LogicLock region assigned to a safety IP can only contain safety IP logic.

### Assigning I/O Pins

You can use a global assignment to specify that a pin is assigned to a safety IP.

```
set_instance_assignment -name ENABLE_STRICT_PRESERVATION ON/OFF -to <hpath> -
section_id <region_name>
```

- *<hpath>* refers to an I/O pin (pad).
- *<region\_name>* refers to the top-level safety IP partition name.

A value of ON indicates that the pin is a safety pin that should be preserved along with the safety IP. A value of OFF indicates that the pin that connects up to the safety IP, should be treated as a standard pin, and is not preserved along with the safety IP.

All the pins that connect up to a safety IP should have an explicit assignment.

An error is reported if a pin that connects up the safety IP does not have an assignment or a pin does not connect up to the specified *<region\_name>*.

If an IO\_REG group contains a pin that is assigned to a safety IP, then all the pins in the IO\_REG group are reserved for this safety IP. All pins in the IO\_REG group need to be assigned to the same safety IP and none of the pins in the group can be assigned to standard signals.

## General Guidelines for Implementation

- An internal clock source, such as a PLL, should be implemented in a safe partition.
- An I/O pin driving the external clock should be indicated as a safety pin.
- To export a safety IP containing several partitions, the top-level partition for the safety IP should be exported. A safety IP containing several partitions is flattened and converted into a single partition during export. This hierarchical safety IP is flattened to ensure bit-level settings are preserved.
- Hard blocks implemented in a safe partition need to stay with the safe partition.

## Reports for Safety IP

When you have the functional safety separation flow turned on, the Quartus II software displays safety IP and standard IP information in the Fitter report.

### Fitter Report

The Fitter report includes information for each safety IP and the respective partition and I/O usage. The report contains the following information:

- Safety IP name defined as the name of the top-level safety IP partition
- Effective design flow for the safety IP
- Names of all partitions that belong to the safety IP
- Number of safety/standard inputs to the safety IP
- Number of safety/standard outputs to the safety IP
- LogicLock region names along with size and locations for the regions
- I/O pins used for the respective safety IP in your design
- Safety-related error messages

## SIP Partial Bitstream Generation

The Programmer generates a bitstream file containing only the bits for a safety IP. This partial preserved bitstream (**.ppb**) file is for the safety IP region mask. The command lines to generate the partial bitstream file are the following:

- `quartus_cpf --genppb safel.psm design.sof safel.rbf.ppb`
- `quartus_cpf -c safel.psm safel.rbf.ppb`

The **.ppb** file is generated in two steps.

1. Generation of partial SOF.
2. Generation of **.ppb** file using the partial SOF.

The **.psm** file, **.ppb** file, and MD5 hash signature (**.md5.sign**) file created during partial bitstream generation should be archived for use in future design modification flow compiles.

## Exporting and Importing Your Safety IP

### Safety IP Partition Export

After you have successfully compiled the safety IP(s) in the Quartus II software, save the safety partition place and route information for use in any subsequent design modification flow. Saving the partition information allows the safety IP to be imported to a clean Quartus II project where no previous compilation results have been removed (even if the version of the Quartus II software being used is newer than the Quartus II software version with which the safety IP was originally compiled). Use the **Design Partitions** window to export the design partition. Verify that only the post-fit netlist and export routing options are

turned on when you generate the **.qxp** file for each safety IP. The **.qxp** files should be archived along with the partial bitstream files for use in later design modification flow compiles.

### Safety IP Partition Import

You can import a previously exported safety IP partition into your Quartus II project. There are two use-cases for this.

- (Optional) Import into the original project to ensure that any potential source code changes do not trigger the design creation flow unintentionally.
- Import into a new or clean project where you want to use the design modification flow for the safety IP. As the exported partition is independent of your Quartus II software version, you can import the **.qxp** into a future Quartus II software release.

To import a previously exported design partition, use the **Design Partitions** window and import the **.qxp**.

#### Related Information

- [Export Design Partition online help](#)
- [Import Design Partition online help](#)

### POF Comparison Tool for Verification

There is a separate safe/standard partitioning verification tool that is licensed to safety users. Along with the **.ppb** file, a **.md5.sign** file is generated. The MD5 hash signature can be used for verification. For more detailed verification, the POF comparison tool should be used. This POF comparison tool is available in the Altera Functional Safety Data Package.

## Deciding Which Design Blocks Should Be Design Partitions

The incremental compilation design flow requires more planning than flat compilations. For example, you might have to structure your source code or design hierarchy to ensure that logic is grouped correctly for optimization.

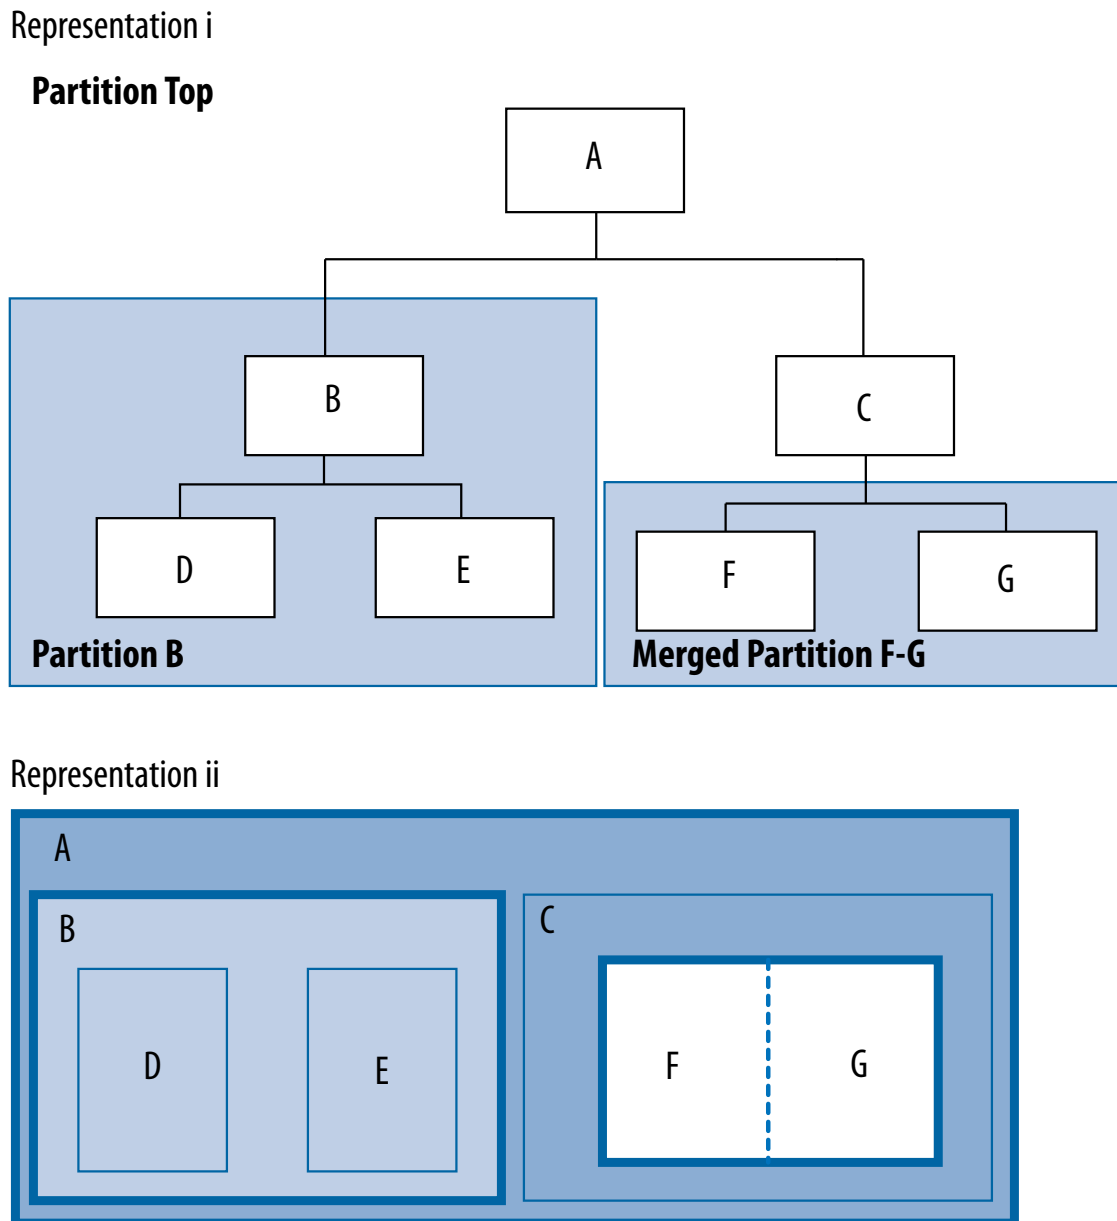
It is a common design practice to create modular or hierarchical designs in which you develop each design entity separately, and then instantiate them in a higher-level entity, forming a complete design. The Quartus II software does not automatically consider each design entity or instance to be a design partition for incremental compilation; instead, you must designate one or more design hierarchies below the top-level project as a design partition. Creating partitions might prevent the Compiler from performing optimizations across partition boundaries. However, this allows for separate synthesis and placement for each partition, making incremental compilation possible.

Partitions must have the same boundaries as hierarchical blocks in the design because a partition cannot be a portion of the logic within a hierarchical entity. You can merge partitions that have the same immediate parent partition to create a single partition that includes more than one hierarchical entity in the design. When you declare a partition, every hierarchical instance within that partition becomes part of the same partition. You can create new partitions for hierarchical instances within an existing partition, in which case the instances within the new partition are no longer included in the higher-level partition, as described in the following example.

In the figure below, a complete design is made up of instances **A, B, C, D, E, F,** and **G**. The shaded boxes in Representation i indicate design partitions in a “tree” representation of the hierarchy. In Representation ii, the lower-level instances are represented inside the higher-level instances, and the partitions are

illustrated with different colored shading. The top-level partition, called “Top”, automatically contains the top-level entity in the design, and contains any logic not defined as part of another partition. The design file for the top level may be just a wrapper for the hierarchical instances below it, or it may contain its own logic. In this example, partition **B** contains the logic in instances **B**, **D**, and **E**. Entities **F** and **G** were first identified as separate partitions, and then merged together to create a partition **F-G**. The partition for the top-level entity **A**, called “Top”, includes the logic in one of its lower-level instances, **C**, because **C** was not defined as part of any other partition.

Figure 3-6: Partitions in a Hierarchical Design



You can create partition assignments to any design instance. The instance can be defined in HDL or schematic design, or come from a third-party synthesis tool as a VQM or EDIF netlist instance.



To take advantage of incremental compilation when source files change, create separate design files for each partition. If you define two different entities as separate partitions but they are in the same design file, you cannot maintain incremental compilation because the software would have to recompile both partitions if you changed either entity in the design file. Similarly, if two partitions rely on the same lower-level entity definition, changes in that lower-level affect both partitions.

The remainder of this section provides information to help you choose which design blocks you should assign as partitions.

## Impact of Design Partitions on Design Optimization

The boundaries of your design partitions can impact the design's quality of results. Creating partitions might prevent the Compiler from performing logic optimizations across partition boundaries, which allows the software to synthesize and place each partition separately in an incremental flow. Therefore, consider partitioning guidelines to help reduce the effect of partition boundaries.

Whenever possible, register all inputs and outputs of each partition. This helps avoid any delay penalty on signals that cross partition boundaries and keeps each register-to-register timing path within one partition for optimization. In addition, minimize the number of paths that cross partition boundaries. If there are timing-critical paths that cross partition boundaries, rework the partitions to avoid these inter-partition paths. Including as many of the timing-critical connections as possible inside a partition allows you to effectively apply optimizations to that partition to improve timing, while leaving the rest of the design unchanged.

Avoid constant partition inputs and outputs. You can also merge two or more partitions to allow cross-boundary optimizations for paths that cross between the partitions, as long as the partitions have the same parent partition. Merging related logic from different hierarchy blocks into one partition can be useful if you cannot change the design hierarchy to accommodate partition assignments.

If critical timing paths cross partition boundaries, you can perform timing budgeting and make timing assignments to constrain the logic in each partition so that the entire timing path meets its requirements. In addition, because each partition is optimized independently during synthesis, you may have to perform resource allocation to ensure that each partition uses an appropriate number of device resources. If design partitions are compiled in separate Quartus II projects, there may be conflicts related to global routing resources for clock signals when the design is integrated into the top-level design. You can use the Global Signal logic option to specify which clocks should use global or regional routing, use the ALTCLK\_CTRL IP core to instantiate a clock control block and connect it appropriately in both the partitions being developed in separate Quartus II projects, or find the compiler-generated clock control node in your design and make clock control location assignments in the Assignment Editor.

## Turning On Supported Cross-boundary Optimizations

You can improve the optimizations performed between design partitions by turning on supported cross-boundary optimizations. These optimizations are turned on a per partition basis and you can select the optimizations as individual assignments. This allows the cross-boundary optimization feature to give you more control over the optimizations that work best for your design. You can turn on the cross-boundary optimizations for your design partitions on the **Advanced** tab of the **Design Partition Properties** dialog box. Once you change the optimization settings, the Quartus II software recompiles your partition from source automatically. Cross-boundary optimizations include the following: propagate constants, propagate inversions on partition inputs, merge inputs fed by a common source, merge electrically equivalent bidirectional pins, absorb internal paths, and remove logic connected to dangling outputs.

Cross-boundary optimizations are implemented top-down from the parent partition into the child partition, but not vice-versa. Also, cross-boundary optimizations cannot be enabled for partitions that allow multiple personas (partial reconfiguration partitions).

#### Related Information

[Design Partition Properties Dialog Box online help](#)

[Best Practices for Incremental Compilation Partitions and Floorplan Assignments documentation](#) on page 14-1

## Design Partition Assignments Compared to Physical Placement Assignments

Design partitions for incremental compilation are logical partitions, which is different from physical placement assignments in the device floorplan. A logical design partition does not refer to a physical area of the device and does not directly control the placement of instances. A logical design partition sets up a virtual boundary between design hierarchies so that each is compiled separately, preventing logical optimizations from occurring between them. When the software compiles the design source code, the logic in each partition can be placed anywhere in the device unless you make additional placement assignments.

If you preserve the compilation results using a Post-Fit netlist, it is not necessary for you to back-annotate or make any location assignments for specific logic nodes. You should not use the incremental compilation and logic placement back-annotation features in the same Quartus II project. The incremental compilation feature does not use placement “assignments” to preserve placement results; it simply reuses the netlist database that includes the placement information.

You can assign design partitions to physical regions in the device floorplan using LogicLock region assignments. In the Quartus II software, LogicLock regions are used to constrain blocks of a design to a particular region of the device. Altera recommends using LogicLock regions for timing-critical design blocks that will change in subsequent compilations, or to improve the quality of results and avoid placement conflicts in some cases.

#### Related Information

[Creating a Design Floorplan With LogicLock Regions](#) on page 3-47

[Best Practices for Incremental Compilation Partitions and Floorplan Assignments documentation](#) on page 14-1

## Using Partitions With Third-Party Synthesis Tools

If you are using a third-party synthesis tool, set up your tool to create a separate VQM or EDIF netlist for each hierarchical partition. In the Quartus II software, assign the top-level entity from each netlist to be a design partition. The VQM or EDIF netlist file is treated as the source file for the partition in the Quartus II software.

### Synopsys Synplify Pro/Premier and Mentor Graphics Precision RTL Plus

The Synplify Pro and Synplify Premier software include the MultiPoint synthesis feature to perform incremental synthesis for each design block assigned as a Compile Point in the user interface or a script. The Precision RTL Plus software includes an incremental synthesis feature that performs block-based synthesis based on Partition assignments in the source HDL code. These features provide automated block-based incremental synthesis flows and create different output netlist files for each block when set up for an Altera device.

Using incremental synthesis within your synthesis tool ensures that only those sections of a design that have been updated are resynthesized when the design is compiled, reducing synthesis run time and preserving the results for the unchanged blocks. You can change and resynthesize one section of a design without affecting other sections of the design.

#### Related Information

- [Synopsys Synplify Support documentation](#) on page 17-1
- [Mentor Graphics Precision Synthesis Support documentation](#) on page 18-1

## Other Synthesis Tools

You can also partition your design and create different netlist files manually with the basic Synplify software (non-Pro/Premier), the basic Precision RTL software (non-Plus), or any other supported synthesis tool by creating a separate project or implementation for each partition, including the top level. Set up each higher-level project to instantiate the lower-level VQM/EDIF netlists as black boxes. Synplify, Precision, and most synthesis tools automatically treat a design block as a black box if the logic definition is missing from the project. Each tool also includes options or attributes to specify that the design block should be treated as a black box, which you can use to avoid warnings about the missing logic.

## Assessing Partition Quality

The Quartus II software provides various tools to assess the quality of your assigned design partitions. You can take advantage of these tools to assess your partition quality, and use the information to improve your design or assignments as required to achieve the best results.

### Partition Statistics Reports

After compilation, you can view statistics about design partitions in the Partition Merge Partition Statistics report, and on the **Statistics** tab in the **Design Partitions Properties** dialog box.

The Partition Merge Partition Statistics report lists statistics about each partition. The statistics for each partition (each row in the table) include the number of logic cells it contains, as well as the number of input and output pins it contains, and how many are registered or unconnected.

You can also view post-compilation statistics about the resource usage and port connections for a particular partition on the **Statistics** tab in the **Design Partition Properties** dialog box.

#### Related Information

- [Best Practices for Incremental Compilation Partitions and Floorplan Assignments documentation](#) on page 14-1

### Partition Timing Reports

You can generate a Partition Timing Overview report and a Partition Timing Details report by clicking **Report Partitions** in the Tasks pane in the TimeQuest Timing Analyzer, or using the `report_partitions` Tcl command.

The Partition Timing Overview report shows the total number of failing paths for each partition and the worst-case slack for any path involving the partition.

The Partition Timing Details report shows the number of failing partition-to-partition paths and worst-case slack for partition-to-partition paths, to provide a more detailed breakdown of where the critical paths in the design are located with respect to design partitions.

## Incremental Compilation Advisor

You can use the Incremental Compilation Advisor to check that your design follows Altera's recommendations for creating design partitions and floorplan location assignments.

Recommendations are split into **General Recommendations**, **Timing Recommendations**, and **Team-Based Design Recommendations** that apply to design flows in which partitions are compiled independently in separate Quartus II projects before being integrated into the top-level design. Each recommendation provides an explanation, describes the effect of the recommendation, and provides the action required to make a suggested change. In some cases, there is a link to the appropriate Quartus II settings page where you can make a suggested change to assignments or settings. For some items, if your design does not follow the recommendation, the **Check Recommendations** operation creates a table that lists any nodes or paths in your design that could be improved. The relevant timing-independent recommendations for the design are also listed in the Design Partitions window and the LogicLock Regions window.

To verify that your design follows the recommendations, go to the **Timing Independent Recommendations** page or the **Timing Dependent Recommendations** page, and then click **Check Recommendations**. For large designs, these operations can take a few minutes.

After you perform a check operation, symbols appear next to each recommendation to indicate whether the design or project setting follows the recommendations, or if some or all of the design or project settings do not follow the recommendations. Following these recommendations is not mandatory to use the incremental compilation feature. The recommendations are most important to ensure good results for timing-critical partitions.

For some items in the Advisor, if your design does not follow the recommendation, the **Check Recommendations** operation lists any parts of the design that could be improved. For example, if not all of the partition I/O ports follow the **Register All Non-Global Ports** recommendation, the advisor displays a list of unregistered ports with the partition name and the node name associated with the port.

When the advisor provides a list of nodes, you can right-click a node, and then click **Locate** to cross-probe to other Quartus II features, such as the RTL Viewer, Chip Planner, or the design source code in the text editor.

**Note:** Opening a new TimeQuest report resets the Incremental Compilation Advisor results, so you must rerun the Check Recommendations process.

## Specifying the Level of Results Preservation for Subsequent Compilations

The netlist type of each design partition allows you to specify the level of results preservation. The netlist type determines which type of netlist or source file the Partition Merge stage uses in the next incremental compilation.

When you choose to preserve a post-fit compilation netlist, the default level of Fitter preservation is the highest degree of placement and routing preservation supported by the device family. The advanced Fitter Preservation Level setting allows you to specify the amount of information that you want to preserve from the post-fit netlist file.

### Netlist Type for Design Partitions

Before starting a new compilation, ensure that the appropriate netlist type is set for each partition to preserve the desired level of compilation results. The table below describes the settings for the netlist type,

explains the behavior of the Quartus II software for each setting, and provides guidance on when to use each setting.

**Table 3-2: Partition Netlist Type Settings**

Netlist Type	Quartus II Software Behavior for Partition During Compilation
Source File	<p>Always compiles the partition using the associated design source file(s). <sup>(2)</sup></p> <p>Use this netlist type to recompile a partition from the source code using new synthesis or Fitter settings.</p>
Post-Synthesis	<p>Preserves post-synthesis results for the partition and reuses the post-synthesis netlist when the following conditions are true:</p> <ul style="list-style-type: none"> <li>• A post-synthesis netlist is available from a previous synthesis.</li> <li>• No change that initiates an automatic resynthesis has been made to the partition since the previous synthesis. <sup>(3)</sup></li> </ul> <p>Compiles the partition from the source files if resynthesis is initiated or if a post-synthesis netlist is not available. <sup>(2)</sup></p> <p>Use this netlist type to preserve the synthesis results unless you make design changes, but allow the Fitter to refit the partition using any new Fitter settings.</p>
Post-Fit	<p>Preserves post-fit results for the partition and reuses the post-fit netlist when the following conditions are true:</p> <ul style="list-style-type: none"> <li>• A post-fit netlist is available from a previous fitting.</li> <li>• No change that initiates an automatic resynthesis has been made to the partition since the previous fitting. <sup>(3)</sup></li> </ul> <p>When a post-fit netlist is not available, the software reuses the post-synthesis netlist if it is available, or otherwise compiles from the source files. Compiles the partition from the source files if resynthesis is initiated. <sup>(2)</sup></p> <p>The Fitter Preservation Level specifies what level of information is preserved from the post-fit netlist.</p> <p>Assignment changes, such as Fitter optimization settings, do not cause a partition set to Post-Fit to recompile.</p>

<sup>(2)</sup> If you use Rapid Recompile, the Quartus II software might not recompile the entire partition from the source code as described in this table; it will reuse compatible results if there have been only small changes to the logic in the partition.

<sup>(3)</sup> You can turn on the **Ignore changes in source files and strictly use the specified netlist, if available** option on the **Advanced** tab in the **Design Partitions Properties** dialog box to specify whether the Compiler should ignore source file changes when deciding whether to recompile the partition.

Netlist Type	Quartus II Software Behavior for Partition During Compilation
Empty	<p>Uses an empty placeholder netlist for the partition. The partition's port interface information is required during Analysis and Synthesis to connect the partition correctly to other logic and partitions in the design, and peripheral nodes in the source file including pins and PLLs are preserved to help connect the empty partition to the rest of the design and preserve timing of any lower-level non-empty partitions within empty partitions. If the source file is not available, you can create a wrapper file that defines the design block and specifies the input, output, and bidirectional ports. In Verilog HDL: a module declaration, and in VHDL: an entity and architecture declaration.</p> <p>You can use this netlist type to skip the compilation of a partition that is incomplete or missing from the top-level design. You can also set an empty partition if you want to compile only some partitions in the design, such as to optimize the placement of a timing-critical block such as an IP core before incorporating other design logic, or if the compilation time is large for one partition and you want to exclude it.</p> <p>If the project database includes a previously generated post-synthesis or post-fit netlist for an unchanged Empty partition, you can set the netlist type from <b>Empty</b> directly to <b>Post-Synthesis</b> or <b>Post-Fit</b> and the software reuses the previous netlist information without recompiling from the source files.</p>

**Related Information**

- [What Changes Initiate the Automatic Resynthesis of a Partition?](#) on page 3-29
- [Fitter Preservation Level for Design Partitions](#) on page 3-27
- [Incremental Capabilities Available When A Design Has No Partitions](#) on page 3-2

**Fitter Preservation Level for Design Partitions**

The default Fitter Preservation Level for partitions with a **Post-Fit** netlist type is the highest level of preservation available for the target device family and provides the most compilation time reduction.

You can change the advanced Fitter Preservation Level setting to provide more flexibility in the Fitter during placement and routing. You can set the Fitter Preservation Level on the **Advanced** tab in the **Design Partitions Properties** dialog box.

Table 3-3: Fitter Preservation Level Settings

Fitter Preservation Level	Quartus II Behavior for Partition During Compilation
Placement and Routing	<p>Preserves the design partition's netlist atoms and their placement and routing.</p> <p>This setting reduces compilation times compared to Placement only, but provides less flexibility to the router to make changes if there are changes in other parts of the design.</p> <p>By default, the Fitter preserves the usage of high-speed programmable power tiles contained within the selected partition, for devices that support high-speed and low-power tiles. You can turn off the <b>Preserve high-speed tiles when preserving placement and routing</b> option on the <b>Advanced</b> tab in the <b>Design Partitions Properties</b> dialog box.</p>
Placement	<p>Preserves the netlist atoms and their placement in the design partition. Reroutes the design partition and does not preserve high-speed power tile usage.</p>
Netlist Only	<p>Preserves the netlist atoms of the design partition, but replaces and reroutes the design partition. A post-fit netlist with the atoms preserved can be different than the Post-Synthesis netlist because it contains Fitter optimizations; for example, Physical Synthesis changes made during a previous Fitting.</p> <p>You can use this setting to:</p> <ul style="list-style-type: none"> <li>• Preserve Fitter optimizations but allow the software to perform placement and routing again.</li> <li>• Reapply certain Fitter optimizations that would otherwise be impossible when the placement is locked down.</li> <li>• Resolve resource conflicts between two imported partitions.</li> </ul>

**Related Information**

[Setting the Netlist Type and Fitter Preservation Level for Design Partitions online help](#)

**Where Are the Netlist Databases Saved?**

The incremental compilation database folder (`\incremental_db`) includes all the netlist information from previous compilations. To avoid unnecessary recompilations, these database files must not be altered or deleted.

If you archive or reproduce the project in another location, you can use a Quartus II Archive File (.qar). Include the incremental compilation database files to preserve post-synthesis or post-fit compilation results.

To manually create a project archive that preserves compilation results without keeping the incremental compilation database, you can keep all source and settings files, and create and save a Quartus II Settings File (.qxp) for each partition in the design that will be integrated into the top-level design.

#### Related Information

- [Using Incremental Compilation With Quartus II Archive Files](#) on page 3-50
- [Exporting Design Partitions from Separate Quartus II Projects](#) on page 3-31

## Deleting Netlists

You can choose to abandon all levels of results preservation and remove all netlists that exist for a particular partition with the **Delete Netlists** command in the Design Partitions window. When you delete netlists for a partition, the partition is compiled using the associated design source file(s) in the next compilation. Resetting the netlist type for a partition to **Source** would have the same effect, though the netlists would not be permanently deleted and would be available for use in subsequent compilations. For an imported partition, the **Delete Netlists** command also optionally allows you to remove the imported .qxp.

## What Changes Initiate the Automatic Resynthesis of a Partition?

A partition is synthesized from its source files if there is no post-synthesis netlist available from a previous synthesis, or if the netlist type is set to **Source File**. Additionally, certain changes to a partition initiate an automatic resynthesis of the partition when the netlist type is **Post Synthesis** or **Post-Fit**. The software resynthesizes the partition in these cases to ensure that the design description matches the post-place-and-route programming files.

The following list explains the changes that initiate a partition's automatic resynthesis when the netlist type is set to **Post-Synthesis** or **Post-Fit**:

- The device family setting has changed.
- Any dependent source design file has changed.
- The partition boundary was changed by an addition, removal, or change to the port boundaries of a partition (for example, a new partition has been defined for a lower-level instance within this partition).
- A dependent source file was compiled into a different library (so it has a different `-library` argument).
- A dependent source file was added or removed; that is, the partition depends on a different set of source files.



- The partition's root instance has a different entity binding. In VHDL, an instance may be bound to a specific entity and architecture. If the target entity or architecture changes, it triggers resynthesis.
- The partition has different parameters on its root hierarchy or on an internal AHDL hierarchy (AHDL automatically inherits parameters from its parent hierarchies). This occurs if you modified the parameters on the hierarchy directly, or if you modified them indirectly by changing the parameters in a parent design hierarchy.
- You have moved the project and compiled database between a Windows and Linux system. Due to the differences in the way new line feeds are handled between the operating systems, the internal checksum algorithm may detect a design file change in this case.

The software reuses the post-synthesis results but re-fits the design if you change the device setting within the same device family. The software reuses the post-fitting netlist if you change only the device speed grade.

Synthesis and Fitter assignments, such as optimization settings, timing assignments, or Fitter location assignments including pin assignments, do not trigger automatic recompilation in the incremental compilation flow. To recompile a partition with new assignments, change the netlist type for that partition to one of the following:

- **Source File** to recompile with all new settings
- **Post-Synthesis** to recompile using existing synthesis results but new Fitter settings
- **Post-Fit** with the **Fitter Preservation Level** set to **Placement** to rerun routing using existing placement results, but new routing settings (such as delay chain settings)

You can use the LogicLock Origin location assignment to change or fine-tune the previous Fitter results from a Post-Fit netlist.

#### Related Information

[Changing Partition Placement with LogicLock Changes](#) on page 3-48

## Resynthesis Due to Source Code Changes

The Quartus II software uses an internal checksum algorithm to determine whether the contents of a source file have changed. Source files are the design description files used to create the design, and include Memory Initialization Files (.mif) as well as .qxp from exported partitions. When design files in a partition have dependencies on other files, changing one file may initiate an automatic recompilation of another file. The Partition Dependent Files table in the Analysis and Synthesis report lists the design files that contribute to each design partition. You can use this table to determine which partitions are recompiled when a specific file is changed.

For example, if a design has file **A.v** that contains entity **A**, **B.v** that contains entity **B**, and **C.v** that contains entity **C**, then the Partition Dependent Files table for the partition containing entity **A** lists file **A.v**, the table for the partition containing entity **B** lists file **B.v**, and the table for the partition containing entity **C** lists file **C.v**. Any dependencies are transitive, so if file **A.v** depends on **B.v**, and **B.v** depends on **C.v**, the entities in file **A.v** depend on files **B.v** and **C.v**. In this case, files **B.v** and **C.v** are listed in the report table as dependent files for the partition containing entity **A**.

**Note:** If you use Rapid Recompile, the Quartus II software might not recompile the entire partition from the source code as described in this section; it will reuse compatible results if there have been only small changes to the logic in the partition.

If you define module parameters in a higher-level module, the Quartus II software checks the parameter values when determining which partitions require resynthesis. If you change a parameter in a higher-level

module that affects a lower-level module, the lower-level module is resynthesized. Parameter dependencies are tracked separately from source file dependencies; therefore, parameter definitions are not listed in the **Partition Dependent Files** list.

If a design contains common files, such as an **includes.v** file that is referenced in each entity by the command `include includes.v`, all partitions are dependent on this file. A change to **includes.v** causes the entire design to be recompiled. The VHDL statement `use work.all` also typically results in unnecessary recom compilations, because it makes all entities in the work library visible in the current entity, which results in the current entity being dependent on all other entities in the design.

To avoid this type of problem, ensure that files common to all entities, such as a common include file, contain only the set of information that is truly common to all entities. Remove `use work.all` statements in your VHDL file or replace them by including only the specific design units needed for each entity.

#### Related Information

[Incremental Capabilities Available When A Design Has No Partitions](#) on page 3-2

### Forcing Use of the Compilation Netlist When a Partition has Changed

Forcing the use of a post-compilation netlist when the contents of a source file has changed is recommended only for advanced users who understand when a partition must be recompiled. You might use this assignment, for example, if you are making source code changes but do not want to recompile the partition until you finish debugging a different partition, or if you are adding simple comments to the source file but you know the design logic itself is not being changed and you want to keep the previous compilation results.

To force the Fitter to use a previously generated netlist even when there are changes to the source files, right-click the partition in the Design Partitions window and then click **Design Partition Properties**. On the **Advanced** tab, turn on the **Ignore changes in source files and strictly use the specified netlist, if available** option.

Turning on this option can result in the generation of a functionally incorrect netlist when source design files change, because source file updates will not be recompiled. Use caution when setting this option.

### Exporting Design Partitions from Separate Quartus II Projects

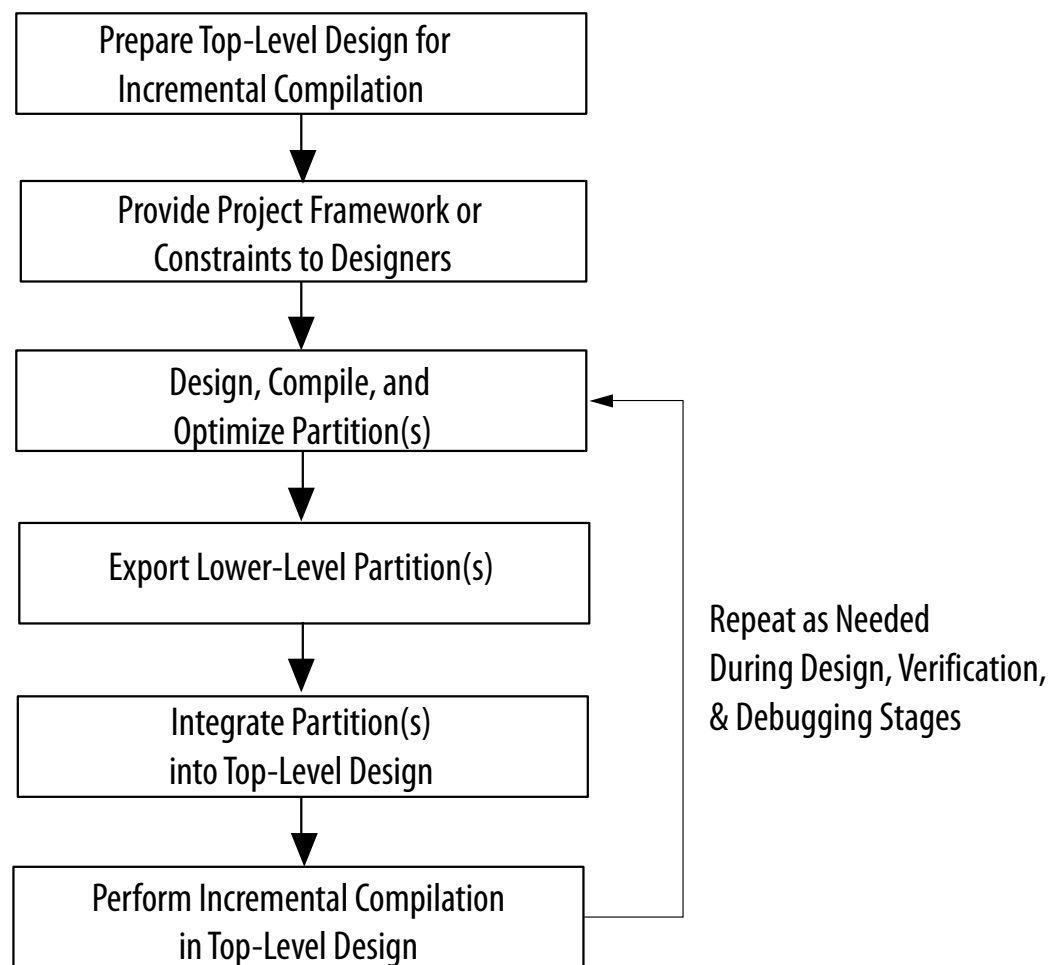
Partitions that are developed by other designers or team members in the same company or third-party IP providers can be exported as design partitions to a Quartus II Exported Partition File (**.qxp**), and then integrated into a top-level design. A **.qxp** is a binary file that contains compilation results describing the exported design partition and includes a post-synthesis netlist, a post-fit netlist, or both, and a set of assignments, sometimes including LogicLock placement constraints. The **.qxp** does not contain the source design files from the original Quartus II project.

To enable team-based development and third-party IP delivery, you can design and optimize partitions in separate copies of the top-level Quartus II project framework, or even in isolation. If the designers have access to the top-level project framework through a source control system, they can access project files as read-only and develop their partition within the source control system. If designers do not have access to a source control system, the project lead can provide the designer with a copy of the top-level project framework to use as they develop their partitions. The project lead also has the option to generate design partition scripts to manage resource and timing budgets in the top-level design when partitions are developed outside the top-level project framework.

The exported compilation results of completed partitions are given to the project lead, preferably using a source control system, who is then responsible for integrating them into the top-level design to obtain a fully functional design. This type of design flow is required only if partition designers want to optimize their placement and routing independently, and pass their design to the project lead to reuse placement and routing results. Otherwise, a project lead can integrate source HDL from several designers in a single Quartus II project, and use the standard incremental compilation flow described previously.

The figure below illustrates the team-based incremental compilation design flow using a methodology in which partitions are compiled in separate Quartus II projects before being integrated into the top-level design. This flow can be used when partitions are developed by other designers or IP providers.

**Figure 3-7: Team-Based Incremental Compilation Design Flow**



**Note:** You cannot export or import partitions that have been merged.

#### Related Information

- [Deciding Which Design Blocks Should Be Design Partitions](#) on page 3-20

- [Incremental Compilation Restrictions](#) on page 3-49

## Preparing the Top-Level Design

To prepare your design to incorporate exported partitions, first create the top-level project framework of the design to define the hierarchy for the subdesigns that will be implemented by other team members, designers, or IP providers.

In the top-level design, create project-wide settings, for example, device selection, global assignments for clocks and device I/O ports, and any global signal constraints to specify which signals can use global routing resources.

Next, create the appropriate design partition assignments and set the netlist type for each design partition that will be developed in a separate Quartus II project to **Empty**. It may be necessary to constrain the location of partitions with LogicLock region assignments if they are timing-critical and are expected to change in future compilations, or if the designer or IP provider wants to place and route their design partition independently, to avoid location conflicts.

Finally, provide the top-level project framework to the partition designers, preferably through a source control system.

### Related Information

[Creating a Design Floorplan With LogicLock Regions](#) on page 3-47

## Empty Partitions

You can use a design flow in which some partitions are set to an **Empty** netlist type to develop pieces of the design separately, and then integrate them into the top-level design at a later time. In a team-based design environment, you can set the netlist type to **Empty** for partitions in your design that will be developed by other designers or IP providers. The **Empty** setting directs the Compiler to skip the compilation of a partition and use an empty placeholder netlist for the partition.

When a netlist type is set to **Empty**, peripheral nodes including pins and PLLs are preserved and all other logic is removed. The peripheral nodes including pins help connect the empty partition to the design, and the PLLs help preserve timing of non-empty partitions within empty partitions.

When you set a design partition to **Empty**, a design file is required during Analysis and Synthesis to specify the port interface information so that it can connect the partition correctly to other logic and partitions in the design. If a partition is exported from another project, the **.qxp** contains this information. If there is no **.qxp** or design file to represent the design entity, you must create a wrapper file that defines the design block and specifies the input, output, and bidirectional ports. For example, in Verilog HDL, you should include a module declaration, and in VHDL, you should include an entity and architecture declaration.

## Project Management— Making the Top-Level Design Available to Other Designers

In team-based incremental compilation flows, whenever possible, all designers or IP providers should work within the same top-level project framework. Using the same project framework among team members ensures that designers have the settings and constraints needed for their partition, and makes timing closure easier when integrating the partitions into the top-level design. If other designers do not have access to the top-level project framework, the Quartus II software provides tools for passing project information to partition designers.

## Distributing the Top-Level Quartus II Project

There are several methods that the project lead can use to distribute the “skeleton” or top-level project framework to other partition designers or IP providers.

- If partition designers have access to the top-level project framework, the project will already include all the settings and constraints needed for the design. This framework should include PLLs and other interface logic if this information is important to optimize partitions.
  - If designers are part of the same design environment, they can check out the required project files from the same source control system. This is the recommended way to share a set of project files.
  - Otherwise, the project lead can provide a copy of the top-level project framework so that each design develops their partition within the same project framework.
- If a partition designer does not have access to the top-level project framework, the project lead can give the partition designer a Tcl script or other documentation to create the separate Quartus II project and all the assignments from the top-level design.

If the partition designers provide the project lead with a post-synthesis **.qxp** and fitting is performed in the top-level design, integrating the design partitions should be quite easy. If you plan to develop a partition in a separate Quartus II project and integrate the optimized post-fitting results into the top-level design, use the following guidelines to improve the integration process:

- Ensure that a LogicLock region constrains the partition placement and uses only the resources allocated by the project lead.
- Ensure that you know which clocks should be allocated to global routing resources so that there are no resource conflicts in the top-level design.
  - Set the Global Signal assignment to **On** for the high fan-out signals that should be routed on global routing lines.
  - To avoid other signals being placed on global routing lines, turn off **Auto Global Clock and Auto Global Register Controls** under **More Settings** on the Fitter page in the **Settings** dialog box. Alternatively, you can set the Global Signal assignment to **Off** for signals that should not be placed on global routing lines.

Placement for LABs depends on whether the inputs to the logic cells within the LAB use a global clock. You may encounter problems if signals do not use global lines in the partition, but use global routing in the top-level design.

- Use the Virtual Pin assignment to indicate pins of a partition that do not drive pins in the top-level design. This is critical when a partition has more output ports than the number of pins available in the target device. Using virtual pins also helps optimize cross-partition paths for a complete design by enabling you to provide more information about the partition ports, such as location and timing assignments.
- When partitions are compiled independently without any information about each other, you might need to provide more information about the timing paths that may be affected by other partitions in the top-level design. You can apply location assignments for each pin to indicate the port location after incorporation in the top-level design. You can also apply timing assignments to the I/O ports of the partition to perform timing budgeting.

### Related Information

- [Best Practices for Incremental Compilation Partitions and Floorplan Assignments documentation](#) on page 14-1

## Generating Design Partition Scripts

If IP providers or designers on a team want to optimize their design blocks independently and do not have access to a shared project framework, the project lead must perform some or all of the following tasks to ensure successful integration of the design blocks:

- Determine which assignments should be propagated from the top-level design to the partitions. This requires detailed knowledge of which assignments are required to set up low-level designs.
- Communicate the top-level assignments to the partitions. This requires detailed knowledge of Tcl or other scripting languages to efficiently communicate project constraints.
- Determine appropriate timing and location assignments that help overcome the limitations of team-based design. This requires examination of the logic in the partitions to determine appropriate timing constraints.
- Perform final timing closure and resource conflict avoidance in the top-level design. Because the partitions have no information about each other, meeting constraints at the lower levels does not guarantee they are met when integrated at the top-level. It then becomes the project lead's responsibility to resolve the issues, even though information about the partition implementation may not be available.

Design partition scripts automate the process of transferring the top-level project framework to partition designers in a flow where each design block is developed in separate Quartus II projects before being integrated into the top-level design. If the project lead cannot provide each designer with a copy of the top-level project framework, the Quartus II software provides an interface for managing resources and timing budgets in the top-level design. Design partition scripts make it easier for partition designers to implement the instructions from the project lead, and avoid conflicts between projects when integrating the partitions into the top-level design. This flow also helps to reduce the need to further optimize the designs after integration.

You can use options in the **Generate Design Partition Scripts** dialog box to choose which types of assignments you want to pass down and create in the partitions being developed in separate Quartus II projects.

### Related Information

- [Enabling Designers on a Team to Optimize Independently](#) on page 3-42
- [Generating Design Partition Scripts for Project Management online help](#)
- [Generate Design Partition Scripts Dialog Box online help](#)

## Exporting Partitions

When partition designers achieve the design requirements in their separate Quartus II projects, each designer can export their design as a partition so it can be integrated into the top-level design by the project lead. The **Export Design Partition** dialog box, available from the Project menu, allows designers to export a design partition to a Quartus II Exported Partition File (**.qxp**) with a post-synthesis netlist, a post-fit netlist, or both. The project lead then adds the **.qxp** to the top-level design to integrate the partition.

A designer developing a timing-critical partition or who wants to optimize their partition on their own would opt to export their completed partition with a post-fit netlist, allowing for the partition to more reliably meet timing requirements after integration. In this case, you must ensure that resources are allocated appropriately to avoid conflicts. If the placement and routing optimization can be performed in

the top-level design, exporting a post-synthesis netlist allows the most flexibility in the top-level design and avoids potential placement or routing conflicts with other partitions.

When designing the partition logic to be exported into another project, you can add logic around the design block to be exported as a design partition. You can instantiate additional design components for the Quartus II project so that it matches the top-level design environment, especially in cases where you do not have access to the full top-level design project. For example, you can include a top-level PLL in the project, outside of the partition to be exported, so that you can optimize the design with information about the frequency multipliers, phase shifts, compensation delays, and any other PLL parameters. The software then captures timing and resource requirements more accurately while ensuring that the timing analysis in the partition is complete and accurate. You can export the partition for the top-level design without any auxiliary components that are instantiated outside the partition being exported.

If your design team uses makefiles and design partition scripts, the project lead can use the **make** command with the **master\_makefile** command created by the scripts to export the partitions and create **.qxp** files. When a partition has been compiled and is ready to be integrated into the top-level design, you can export the partition with option on the **Export Design Partition** dialog box, available from the Project menu.

#### Related Information

[Using a Team-Based Incremental Compilation Design Flow online help](#)

## Viewing the Contents of a Quartus II Exported Partition File (.qxp)

The QXP report allows you to view a summary of the contents in a **.qxp** when you open the file in the Quartus II software. The **.qxp** is a binary file that contains compilation results so the file cannot be read in a text editor. The QXP report opens in the main Quartus II window and contains summary information including a list of the I/O ports, resource usage summary, and a list of the assignments used for the exported partition.

## Integrating Partitions into the Top-Level Design

To integrate a partition developed in a separate Quartus II project into the top-level design, you can simply add the **.qxp** as a source file in your top-level design (just like a Verilog or VHDL source file). You can also use the **Import Design Partition** dialog box to import the partition.

The **.qxp** contains the design block exported from the partition and has the same name as the partition. When you instantiate the design block into a top-level design and include the **.qxp** as a source file, the software adds the exported netlist to the database for the top-level design. The **.qxp** port names are case sensitive if the original HDL of the partition was case sensitive.

When you use a **.qxp** as a source file in this way, you can choose whether you want the **.qxp** to be a partition in the top-level design. If you do not designate the **.qxp** instance as a partition, the software reuses just the post-synthesis compilation results from the **.qxp**, removes unconnected ports and unused logic just like a regular source file, and then performs placement and routing.

If you assigned the **.qxp** instance as a partition, you can set the netlist type in the Design Partitions Window to choose the level of results to preserve from the **.qxp**. To preserve the placement and routing results from the exported partition, set the netlist type to **Post-Fit** for the **.qxp** partition in the top-level design. If you assign the instance as a design partition, the partition boundary is preserved.

#### Related Information

[Impact of Design Partitions on Design Optimization](#) on page 3-22

## Integrating Assignments from the .qxp

The Quartus II software filters assignments from **.qxp** files to include appropriate assignments in the top-level design. The assignments in the **.qxp** are treated like assignments made in an HDL source file, and are not listed in the Quartus II Settings File (**.qsf**) for the top-level design. Most assignments from the **.qxp** can be overridden by assignments in the top-level design.

### Design Partition Assignments Within the Exported Partition

Design partition assignments defined within a separate Quartus II project are not added to the top-level design. All logic under the exported partition in the project hierarchy is treated as single instance in the **.qxp**.

### Synopsys Design Constraint Files for the Quartus II TimeQuest Timing Analyzer

Timing assignments made for the Quartus II TimeQuest analyzer in a Synopsys Design Constraint File (**.sdc**) in the lower-level partition project are not added to the top-level design. Ensure that the top-level design includes all of the timing requirements for the entire project.

#### Related Information

- [Best Practices for Incremental Compilation Partitions and Floorplan Assignments documentation](#) on page 14-1

### Global Assignments

The project lead should make all global project-wide assignments in the top-level design. Global assignments from the exported partition's project are not added to the top-level design. When it is possible for a particular constraint, the global assignment is converted to an instance-specific assignment for the exported design partition.

### LogicLock Region Assignments

The project lead typically creates LogicLock region assignments in the top-level design for any lower-level partition designs where designer or IP providers plan to export post-fit information to be used in the top-level design, to help avoid placement conflicts between partitions. When you use the **.qxp** as a source file, LogicLock constraints from the exported partition are applied in the top-level design, but will not appear in your **.qsf** file or LogicLock Regions window for you to view or edit. The LogicLock region itself is not required to constrain the partition placement in the top-level design if the netlist type is set to **Post-Fit**, because the netlist contains all the placement information.

## Integrating Encrypted IP Cores from .qxp Files

Proper license information is required to compile encrypted IP cores. If an IP core is exported as a **.qxp** from another Quartus II project, the top-level designer instantiating the **.qxp** must have the correct license. The software requires a full license to generate an unrestricted programming file. If you do not have a license, but the IP in the **.qxp** was compiled with OpenCore Plus hardware evaluation support, you can generate an evaluation programming file without a license. If the IP supports OpenCore simulation only, you can fully compile the design and generate a simulation netlist, but you cannot create programming files unless you have a full license.

## Advanced Importing Options

You can use advanced options in the **Import Design Partition** dialog box to integrate a partition developed in a separate Quartus II project into the top-level design. The import process adds more control than using the **.qxp** as a source file, and is useful only in the following circumstances:



- **If you want LogicLock regions in your top-level design (.qsf)**—If you have regions in your partitions that are not also in the top-level design, the regions will be added to your **.qsf** during the import process.
- **If you want different settings or placement for different instantiations of the same entity**—You can control the setting import process with the advanced import options, and specify different settings for different instances of the same **.qxp** design block.

When you use the **Import Design Partition** dialog box to integrate a partition into the top-level design, the import process sets the partition's netlist type to **Imported** in the Design Partitions window.

After you compile the entire design, if you make changes to the place-and-route results (such as movement of an imported LogicLock region), use the **Post-Fit** netlist type on subsequent compilations. To discard an imported netlist and recompile from source code, you can compile the partition with the netlist type set to **Source File** and be sure to include the relevant source code in the top-level design. The import process sets the partition's Fitter Preservation Level to the setting with the highest degree of preservation supported by the imported netlist. For example, if a post-fit netlist is imported with placement information, the Fitter Preservation Level is set to **Placement**, but you can change it to the **Netlist Only** value.

When you import a partition from a **.qxp**, the **.qxp** itself is not part of the top-level design because the netlists from the file have been imported into the project database. Therefore if a new version of a **.qxp** is exported, the top-level designer must perform another import of the **.qxp**.

When you import a partition into a top-level design with the **Import Design Partition** dialog box, the software imports relevant assignments from the partition into the top-level design. If required, you can change the way some assignments are imported, as described in the following subsections.

#### Related Information

- [Netlist Type for Design Partitions](#) on page 3-25
- [Fitter Preservation Level for Design Partitions](#) on page 3-27

### Importing LogicLock Assignments

LogicLock regions are set to a fixed size when imported. If you instantiate multiple instances of a subdesign in the top-level design, the imported LogicLock regions are set to a Floating location. Otherwise, they are set to a Fixed location. You can change the location of LogicLock regions after they are imported, or change them to a Floating location to allow the software to place each region but keep the relative locations of nodes within the region wherever possible. To preserve changes made to a partition after compilation, use the **Post-Fit** netlist type.

The LogicLock Member State assignment is set to **Locked** to signify that it is a preserved region.

LogicLock back-annotation and node location data is not imported because the **.qxp** contains all of the relevant placement information. Altera strongly recommends that you do not add to or delete members from an imported LogicLock region.

#### Related Information

[Changing Partition Placement with LogicLock Changes](#) on page 3-48

### Advanced Import Settings

The **Advanced Import Settings** dialog box allows you to disable assignment import and specify additional options that control how assignments and regions are integrated when importing a partition into a top-level design, including how to resolve assignment conflicts.

### Related Information

[Advanced Import Settings Dialog Box online help](#)

## Team-Based Design Optimization and Third-Party IP Delivery Scenarios

### Using an Exported Partition to Send to a Design Without Including Source Files

Scenario background: A designer wants to produce a design block and needs to send out their design, but to preserve their IP, they prefer to send a synthesized netlist instead of providing the HDL source code to the recipient. You can use this flow to implement a black box.

Use this flow to package a full design as a single source file to send to an end customer or another design location.

As the sender in this scenario perform the following steps to export a design block:

1. Provide the device family name to the recipient. If you send placement information with the synthesized netlist, also provide the exact device selection so they can set up their project to match.
2. Create a black box wrapper file that defines the port interface for the design block and provide it to the recipient for instantiating the block as an empty partition in the top-level design.
3. Create a Quartus II project for the design block, and complete the design.
4. Export the level of hierarchy into a single **.qxp**. Following a successful compilation of the project, you can generate a **.qxp** from the GUI, the command-line, or with Tcl commands, as described in the following:
  - If you are using the Quartus II GUI, use the **Export Design Partition** dialog box.
  - If you are using command-line executables, run **quartus\_cdb** with the `--incremental_compilation_export` option.
  - If you are using Tcl commands, use the following command: `execute_flow -incremental_compilation_export`.
5. Select the option to include just the **Post-synthesis netlist** if you do not have to send placement information. If the recipient wants to reproduce your exact Fitter results, you can select the **Post-fitting netlist** option, and optionally enable **Export routing**.
6. If a partition contains sub-partitions, then the sub-partitions are automatically flattened and merged into the partition netlist before exporting. You can change this behavior and preserve the sub-partition hierarchy by turning off the **Flatten sub-partitions** option on the **Export Design Partition** dialog box. Optionally, you can use the `-dont_flatten` sub-option for the `export_partition` Tcl command.
7. Provide the **.qxp** to the recipient. Note that you do not have to send any of your design source code.

As the recipient in this example, first create a Quartus II project for your top-level design and ensure that your project targets the same device (or at least the same device family if the **.qxp** does not include placement information), as specified by the IP designer sending the design block. Instantiate the design block using the port information provided, and then incorporate the design block into a top-level design.

Add the **.qxp** from the IP designer as a source file in your Quartus II project to replace any empty wrapper file. If you want to use just the post-synthesis information, you can choose whether you want the file to be a partition in the top-level design. To use the post-fit information from the **.qxp**, assign the instance as a design partition and set the netlist type to **Post-Fit**.

### Related Information

- [Creating Design Partitions](#) on page 3-9
- [Netlist Type for Design Partitions](#) on page 3-25

## Creating Precompiled Design Blocks (or Hard-Wired Macros) for Reuse

Scenario background: An IP provider wants to produce and sell an IP core for a component to be used in higher-level systems. The IP provider wants to optimize the placement of their block for maximum performance in a specific Altera device and then deliver the placement information to their end customer. To preserve their IP, they also prefer to send a compiled netlist instead of providing the HDL source code to their customer.

Use this design flow to create a precompiled IP block (sometimes known as a hard-wired macro) that can be instantiated in a top-level design. This flow provides the ability to export a design block with post-synthesis or placement (and, optionally, routing) information and to import any number of copies of this pre-compiled block into another design.

The customer first specifies which Altera device is being used for this project and provides the design specifications.

As the IP provider in this example, perform the following steps to export a preplaced IP core (or hard macro):

1. Create a black box wrapper file that defines the port interface for the IP core and provide the file to the customer to instantiate as an empty partition in the top-level design.
2. Create a Quartus II project for the IP core.
3. Create a LogicLock region for the design hierarchy to be exported.

Using a LogicLock region for the IP core allows the customer to create an empty placeholder region to reserve space for the IP in the design floorplan and ensures that there are no conflicts with the top-level design logic. Reserved space also helps ensure the IP core does not affect the timing performance of other logic in the top-level design. Additionally, with a LogicLock region, you can preserve placement either absolutely or relative to the origin of the associated region. This is important when a **.qxp** is imported for multiple partition hierarchies in the same project, because in this case, the location of at least one instance in the top-level design does not match the location used by the IP provider.

4. If required, add any logic (such as PLLs or other logic defined in the customer's top-level design) around the design hierarchy to be exported. If you do so, create a design partition for the design hierarchy that will be exported as an IP core.
5. Optimize the design and close timing to meet the design specifications.
6. Export the level of hierarchy for the IP core into a single **.qxp**.
7. Provide the **.qxp** to the customer. Note that you do not have to send any of your design source code to the customer; the design netlist and placement and routing information is contained within the **.qxp**.

### Related Information

- [Creating Design Partitions](#) on page 3-55
- [Netlist Type for Design Partitions](#) on page 3-25
- [Changing Partition Placement with LogicLock Changes](#) on page 3-48

## Incorporate IP Core

As the customer in this example, incorporate the IP core in your design by performing the following steps:

1. Create a Quartus II project for the top-level design that targets the same device and instantiate a copy or multiple copies of the IP core. Use a black box wrapper file to define the port interface of the IP core.
2. Perform Analysis and Elaboration to identify the design hierarchy.
3. Create a design partition for each instance of the IP core with the netlist type set to **Empty**.
4. You can now continue work on your part of the design and accept the IP core from the IP provider when it is ready.
5. Include the **.qxp** from the IP provider in your project to replace the empty wrapper-file for the IP instance. Or, if you are importing multiple copies of the design block and want to import relative placement, follow these additional steps:
  - a. Use the **Import** command to select each appropriate partition hierarchy. You can import a **.qxp** from the GUI, the command-line, or with Tcl commands:
    - If you are using the Quartus II GUI, use the **Import Design Partition** command.
    - If you are using command-line executables, run **quartus\_cdb** with the `incremental_compilation_import` option.
    - If you are using Tcl commands, use the following command:`execute_flow - incremental_compilation_import`.
  - b. When you have multiple instances of the IP block, you can set the imported LogicLock regions to floating, or move them to a new location, and the software preserves the relative placement for each of the imported modules (relative to the origin of the LogicLock region). Routing information is preserved whenever possible.

**Note:** The Fitter ignores relative placement assignments if the LogicLock region's location in the top-level design is not compatible with the locations exported in the **.qxp**.
6. You can control the level of results preservation with the **Netlist Type** setting.

If the IP provider did not define a LogicLock region in the exported partition, the software preserves absolute placement locations and this leads to placement conflicts if the partition is imported for more than one instance

## Designing in a Team-Based Environment

Scenario background: A project includes several lower-level design blocks that are developed separately by different designers and instantiated exactly once in the top-level design.

This scenario describes how to use incremental compilation in a team-based design environment where each designer has access to the top-level project framework, but wants to optimize their design in a separate Quartus II project before integrating their design block into the top-level design.

As the project lead in this scenario, perform the following steps to prepare the top-level design:

1. Create a new Quartus II project to ultimately contain the full implementation of the entire design and include a "skeleton" or framework of the design that defines the hierarchy for the subdesigns implemented by separate designers. The top-level design implements the top-level entity in the design

and instantiates wrapper files that represent each subdesign by defining only the port interfaces, but not the implementation.

2. Make project-wide settings. Select the device, make global assignments such as device I/O ports, define the top-level timing constraints, and make any global signal allocation constraints to specify which signals can use global routing resources.
3. Make design partition assignments for each subdesign and set the netlist type for each design partition to be imported to **Empty** in the Design Partitions window.
4. Create LogicLock regions to create a design floorplan for each of the partitions that will be developed separately. This floorplan should consider the connectivity between partitions and estimates of the size of each partition based on any initial implementation numbers and knowledge of the design specifications.
5. Provide the top-level project framework to partition designers using one of the following procedures:
  - Allow access to the full project for all designers through a source control system. Each designer can check out the projects files as read-only and work on their blocks independently. This design flow provides each designer with the most information about the full design, which helps avoid resource conflicts and makes design integration easy.
  - Provide a copy of the top-level Quartus II project framework for each designer. You can use the **Copy Project** command on the Project menu or create a project archive.

## Exporting Your Partition

As the designer of a lower-level design block in this scenario, design and optimize your partition in your copy of the top-level design, and then follow these steps when you have achieved the desired compilation results:

1. On the Project menu, click **Export Design Partition**.
2. In the **Export Design Partition** dialog box, choose the netlist(s) to export. You can export a Post-synthesis netlist if placement or performance preservation is not required, to provide the most flexibility for the Fitter in the top-level design. Select Post-fit netlist to preserve the placement and performance of the lower-level design block, and turn on **Export routing** to include the routing information, if required. One **.qxp** can include both post-synthesis and post-fitting netlists.
3. Provide the **.qxp** to the project lead.

## Integrating Your Partitions

Finally, as the project lead in this scenario, perform these steps to integrate the **.qxp** files received from designers of each partition:

1. Add the **.qxp** as a source file in the Quartus II project, to replace any empty wrapper file for the previously **Empty** partition.
2. Change the netlist type for the partition from **Empty** to the required level of results preservation.

## Enabling Designers on a Team to Optimize Independently

Scenario background: A project consists of several lower-level design blocks that are developed separately by different designers who do not have access to a shared top-level project framework. This scenario is similar to creating precompiled design blocks for reuse, but assumes that there are several design blocks being developed independently (instead of just one IP block), and the project lead can provide some information about the design to the individual designers. If the designers have shared access to the top-level design, use the instructions for designing in a team-based environment.

This scenario assumes that there are several design blocks being developed independently (instead of just one IP block), and the project lead can provide some information about the design to the individual designers.

This scenario describes how to use incremental compilation in a team-based design environment where designers or IP developers want to fully optimize the placement and routing of their design independently in a separate Quartus II project before sending the design to the project lead. This design flow requires more planning and careful resource allocation because design blocks are developed independently.

#### Related Information

- [Creating Precompiled Design Blocks \(or Hard-Wired Macros\) for Reuse](#) on page 3-40
- [Designing in a Team-Based Environment](#) on page 3-41

## Preparing Your Top-level Design

As the project lead in this scenario, perform the following steps to prepare the top-level design:

1. Create a new Quartus II project to ultimately contain the full implementation of the entire design and include a “skeleton” or framework of the design that defines the hierarchy for the subdesigns implemented by separate designers. The top-level design implements the top-level entity in the design and instantiates wrapper files that represent each subdesign by defining only the port interfaces but not the implementation.
2. Make project-wide settings. Select the device, make global assignments such as device I/O ports, define the top-level timing constraints, and make any global signal constraints to specify which signals can use global routing resources.
3. Make design partition assignments for each subdesign and set the netlist type for each design partition to be imported to **Empty** in the Design Partitions window.
4. Create LogicLock regions. This floorplan should consider the connectivity between partitions and estimates of the size of each partition based on any initial implementation numbers and knowledge of the design specifications.
5. Provide the constraints from the top-level design to partition designers using one of the following procedures.
  - Use design partition scripts to pass constraints and generate separate Quartus II projects. On the Project menu, use the **Generate Design Partition Scripts** command, or run the script generator from a Tcl or command prompt. Make changes to the default script options as required for your project. Altera recommends that you pass all the default constraints, including LogicLock regions, for all partitions and virtual pin location assignments. If partitions have not already been created by the other designers, use the partition script to set up the projects so that you can easily take advantage of makefiles. Provide each partition designer with the Tcl file to create their project with the appropriate constraints. If you are using makefiles, provide the makefile for each partition.
  - Use documentation or manually-created scripts to pass all constraints and assignments to each partition designer.

## Exporting Your Design

As the designer of a lower-level design block in this scenario, perform the appropriate set of steps to successfully export your design, whether the design team is using makefiles or exporting and importing the design manually.

If you are using makefiles with the design partition scripts, perform the following steps:

1. Use the **make** command and the makefile provided by the project lead to create a Quartus II project with all design constraints, and compile the project.
2. The information about which source file should be associated with which partition is not available to the software automatically, so you must specify this information in the makefile. You must specify the dependencies before the software rebuilds the project after the initial call to the makefile.
3. When you have achieved the desired compilation results and the design is ready to be imported into the top-level design, the project lead can use the `master_makefile` command to export this partition and create a **.qxp**, and then import it into the top-level design.

## Exporting Without Makefiles

If you are not using makefiles, perform the following steps:

1. If you are using design partition scripts, source the Tcl script provided by the Project Lead to create a project with the required settings:
  - To source the Tcl script in the Quartus II software, on the Tools menu, click **Utility Windows** to open the Tcl console. Navigate to the script's directory, and type the following command: `source <filename>`.
  - To source the Tcl script at the system command prompt, type the following command:  
`quartus_cdb -t <filename>.tcl`
2. If you are not using design partition scripts, create a new Quartus II project for the subdesign, and then apply the following settings and constraints to ensure successful integration:
  - Make LogicLock region assignments and global assignments (including clock settings) as specified by the project lead.
  - Make Virtual Pin assignments for ports which represent connections to core logic instead of external device pins in the top-level design.
  - Make floorplan location assignments to the Virtual Pins so they are placed in their corresponding regions as determined by the top-level design. This provides the Fitter with more information about the timing constraints between modules. Alternatively, you can apply timing I/O constraints to the paths that connect to virtual pins.
3. Proceed to compile and optimize the design as needed.
4. When you have achieved the desired compilation results, on the Project menu, click **Export Design Partition**.
5. In the **Export Design Partition** dialog box, choose the netlist(s) to export. You can export a Post-synthesis netlist instead if placement or performance preservation is not required, to provide the most flexibility for the Fitter in the top-level design. Select **Post-fit** to preserve the placement and performance of the lower-level design block, and turn on Export routing to include the routing information, if required. One **.qxp** can include both post-synthesis and post-fitting netlists.
6. Provide the **.qxp** to the project lead.

## Importing Your Design

Finally, as the project lead in this scenario, perform the appropriate set of steps to import the **.qxp** files received from designers of each partition.

If you are using makefiles with the design partition scripts, perform the following steps:

1. Use the `master_makefile` command to export each partition and create `.qxp` files, and then import them into the top-level design.
2. The software does not have all the information about which source files should be associated with which partition, so you must specify this information in the makefile. The software cannot rebuild the project if source files change unless you specify the dependencies.

### Importing Without Makefiles

If you are not using makefiles, perform the following steps:

1. Add the `.qxp` as a source file in the Quartus II project, to replace any empty wrapper file for the previously Empty partition.
2. Change the netlist type for the partition from Empty to the required level of results preservation.

### Resolving Assignment Conflicts During Integration

When integrating lower-level design blocks, the project lead may notice some assignment conflicts. This can occur, for example, if the lower-level design block designers changed their LogicLock regions to account for additional logic or placement constraints, or if the designers applied I/O port timing constraints that differ from constraints added to the top-level design by the project lead. The project lead can address these conflicts by explicitly importing the partitions into the top-level design, and using options in the **Advanced Import Settings** dialog box. After the project lead obtains the `.qxp` for each lower-level design block from the other designers, use the **Import Design Partition** command on the Project menu and specify the partition in the top-level design that is represented by the lower-level design block `.qxp`. Repeat this import process for each partition in the design. After you have imported each partition once, you can select all the design partitions and use the **Reimport using latest import files at previous locations** option to import all the files from their previous locations at one time. To address assignment conflicts, the project lead can take one or both of the following actions:

- Allow new assignments to be imported
- Allow existing assignments to be replaced or updated

When LogicLock region assignment conflicts occur, the project lead may take one of the following actions:

- Allow the imported region to replace the existing region
- Allow the imported region to update the existing region
- Skip assignment import for regions with conflicts

If the placement of different lower-level design blocks conflict, the project lead can also set the partition's **Fitter Preservation Level** to **Netlist Only**, which allows the software to re-perform placement and routing with the imported netlist.

### Importing a Partition to be Instantiated Multiple Times

In this variation of the design scenario, one of the lower-level design blocks is instantiated more than once in the top-level design. The designer of the lower-level design block may want to compile and optimize the entity once under a partition, and then import the results as multiple partitions in the top-level design.

If you import multiple instances of a lower-level design block into the top-level design, the imported LogicLock regions are automatically set to Floating status.



If you resolve conflicts manually, you can use the import options and manual LogicLock assignments to specify the placement of each instance in the top-level design.

## Performing Design Iterations With Lower-Level Partitions

Scenario background: A project consists of several lower-level subdesigns that have been exported from separate Quartus II projects and imported into the top-level design. In this example, integration at the top level has failed because the timing requirements are not met. The timing requirements might have been met in each individual lower-level project, but critical inter-partition paths in the top-level design are causing timing requirements to fail.

After trying various optimizations in the top-level design, the project lead determines that the design cannot meet the timing requirements given the current partition placements that were imported. The project lead decides to pass additional information to the lower-level partitions to improve the placement.

Use this flow if you re-optimize partitions exported from separate Quartus II projects by incorporating additional constraints from the integrated top-level design.

### Providing the Complete Top-Level Project Framework

The best way to provide top-level design information to designers of lower-level partitions is to provide the complete top-level project framework using the following steps:

1. For all partitions other than the one(s) being optimized by a designer(s) in a separate Quartus II project(s), set the netlist type to **Post-Fit**.
2. Make the top-level design directory available in a shared source control system, if possible. Otherwise, copy the entire top-level design project directory (including database files), or create a project archive including the post-compilation database.
3. Provide each partition designer with a checked-out version or copy of the top-level design.
4. The partition designers recompile their designs within the new project framework that includes the rest of the design's placement and routing information as well top-level resource allocations and assignments, and optimize as needed.
5. When the results are satisfactory and the timing requirements are met, export the updated partition as a **.qxp**.

### Providing Information About the Top-Level Framework

If this design flow is not possible, you can generate partition-specific scripts for individual designs to provide information about the top-level project framework with these steps:

1. In the top-level design, on the Project menu, click **Generate Design Partition Scripts**, or launch the script generator from Tcl or the command line.
2. If lower-level projects have already been created for each partition, you can turn off the **Create lower-level project if one does not exist** option.
3. Make additional changes to the default script options, as necessary. Altera recommends that you pass all the default constraints, including LogicLock regions, for all partitions and virtual pin location assignments. Altera also recommends that you add a maximum delay timing constraint for the virtual I/O connections in each partition.
4. The Quartus II software generates Tcl scripts for all partitions, but in this scenario, you would focus on the partitions that make up the cross-partition critical paths. The following assignments are important in the script:
  - Virtual pin assignments for module pins not connected to device I/O ports in the top-level design.

- Location constraints for the virtual pins that reflect the initial top-level placement of the pin's source or destination. These help make the lower-level placement "aware" of its surroundings in the top-level design, leading to a greater chance of timing closure during integration at the top level.
  - `INPUT_MAX_DELAY` and `OUTPUT_MAX_DELAY` timing constraints on the paths to and from the I/O pins of the partition. These constrain the pins to optimize the timing paths to and from the pins.
5. The partition designers source the file provided by the project lead.
- To source the Tcl script from the Quartus II GUI, on the Tools menu, click **Utility Windows** and open the Tcl console. Navigate to the script's directory, and type the following command:

```
source <filename>
```

- To source the Tcl script at the system command prompt, type the following command:

```
quartus_cdb -t <filename>.tcl
```

6. The partition designers recompile their designs with the new project information or assignments and optimize as needed. When the results are satisfactory and the timing requirements are met, export the updated partition as a **.qxp**.

The project lead obtains the updated **.qxp** files from the partition designers and adds them to the top-level design. When a new **.qxp** is added to the files list, the software will detect the change in the "source file" and use the new **.qxp** results during the next compilation. If the project uses the advanced import flow, the project lead must perform another import of the new **.qxp**.

You can now analyze the design to determine whether the timing requirements have been achieved. Because the partitions were compiled with more information about connectivity at the top level, it is more likely that the inter-partition paths have improved placement which helps to meet the timing requirements.

## Creating a Design Floorplan With LogicLock Regions

A floorplan represents the layout of the physical resources on the device. Creating a design floorplan, or floorplanning, describe the process of mapping the logical design hierarchy onto physical regions in the device floorplan. After you have partitioned the design, you can create floorplan location assignments for the design to improve the quality of results when using the incremental compilation design flow. Creating a design floorplan is not a requirement to use an incremental compilation flow, but it is recommended in certain cases. Floorplan location planning can be important for a design that uses incremental compilation for the following reasons:

- To avoid resource conflicts between partitions, predominantly when partitions are imported from another Quartus II project
- To ensure a good quality of results when recompiling individual timing-critical partitions

Design floorplan assignments prevent the situation in which the Fitter must place a partition in an area of the device where most resources are already used by other partitions. A physical region assignment provides a reasonable region to re-place logic after a change, so the Fitter does not have to scatter logic throughout the available space in the device.

Floorplan assignments are not required for non-critical partitions compiled as part of the top-level design. The logic for partitions that are not timing-critical (such as simple top-level glue logic) can be placed anywhere in the device on each recompilation, if that is best for your design.

The simplest way to create a floorplan for a partitioned design is to create one LogicLock region per partition (including the top-level partition). If you have a compilation result for a partitioned design with no LogicLock regions, you can use the Chip Planner with the Design Partition Planner to view the partition placement in the device floorplan. You can draw regions in the floorplan that match the general location and size of the logic in each partition. Or, initially, you can set each region with the default settings of **Auto** size and **Floating** location to allow the Quartus II software to determine the preliminary size and location for the regions. Then, after compilation, use the Fitter-determined size and origin location as a starting point for your design floorplan. Check the quality of results obtained for your floorplan location assignments and make changes to the regions as needed. Alternatively, you can perform synthesis, and then set the regions to the required size based on resource estimates. In this case, use your knowledge of the connections between partitions to place the regions in the floorplan.

Once you have created an initial floorplan, you can refine the region using tools in the Quartus II software. You can also use advanced techniques such as creating non-rectangular regions by merging LogicLock regions.

You can use the Incremental Compilation Advisor to check that your LogicLock regions meet Altera's guidelines.

#### Related Information

[Incremental Compilation Advisor](#) on page 3-25

[Best Practices for Incremental Compilation Partitions and Floorplan Assignments documentation](#) on page 14-1

## Creating and Manipulating LogicLock Regions

Options in the **LogicLock Regions Properties** dialog box, available from the Assignments menu, allow you to enter specific sizing and location requirements for a region. You can also view and refine the size and location of LogicLock regions in the Quartus II Chip Planner. You can select a region in the graphical interface in the Chip Planner and use handles to move or resize the region.

Options in the **Layer Settings** panel in the Chip Planner allow you to create, delete, and modify tasks to determine which objects, including LogicLock regions and design partitions, to display in the Chip Planner.

#### Related Information

[Creating and Manipulating LogicLock Regions online help](#)

## Changing Partition Placement with LogicLock Changes

When a partition is assigned to a LogicLock region as part of a design floorplan, you can modify the placement of a post-fit partition by moving the LogicLock region. Most assignment changes do not initiate a recompilation of a partition if the netlist type specifies that Fitter results should be preserved. For example, changing a pin assignment does not initiate a recompilation; therefore, the design does not use the new pin assignment unless you change the netlist type to **Post Synthesis** or **Source File**.

Similarly, if a partition's placement is preserved, and the partition is assigned to a LogicLock region, the Fitter always reuses the corresponding LogicLock region size specified in the post-fit netlist. That is, changes to the LogicLock **Size** setting do not initiate refitting if a partition's placement is preserved with the **Post-Fit** netlist type, or with **.qxp** that includes post-fit information.

However, you can use the LogicLock **Origin** location assignment to change or fine-tune the previous Fitter results. When you change the **Origin** setting for a region, the Fitter can move the region in the following manner, depending upon how the placement is preserved for that region's members:

- When you set a new region Origin, the Fitter uses the new origin and replaces the logic, preserving the relative placement of the member logic.
- When you set the region Origin to **Floating**, the following conditions apply:
  - If the region's member placement is preserved with an imported partition, the Fitter chooses a new Origin and re-places the logic, preserving the relative placement of the member logic within the region.
  - If the region's member placement is preserved with a **Post-Fit** netlist type, the Fitter does not change the Origin location, and reuses the previous placement results.

#### Related Information

[What Changes Initiate the Automatic Resynthesis of a Partition?](#) on page 3-29

## Taking Advantage of the Early Timing Estimator

When creating a floorplan you can take advantage of the Early Timing Estimator to enable quick compilations of the design while creating assignments. The Early Timing Estimator feature provides a timing estimate for a design without having to run a full compilation. You can use the Chip Planner to view the "placement estimate" created by this feature, identify critical paths by locating from the timing analyzer reports, and, if necessary, add or modify floorplan constraints. You can then rerun the Early Timing Estimator to quickly assess the impact of any floorplan location assignments or logic changes, enabling rapid iterations on design variants to help you find the best solution. This faster placement has an impact on the quality of results. If getting the best quality of results is important in a given design iteration, perform a full compilation with the Fitter instead of using the Early Timing Estimate feature.

## Incremental Compilation Restrictions

### When Timing Performance May Not Be Preserved Exactly

Timing performance might change slightly in a partition with placement and routing preserved when other partitions are incorporated or re-placed and routed. Timing changes are due to changes in parasitic loading or crosstalk introduced by the other (changed) partitions. These timing changes are very small, typically less than 30 ps on a timing path. Additional fan-out on routing lines when partitions are added can also degrade timing performance.

To ensure that a partition continues to meet its timing requirements when other partitions change, a very small timing margin might be required. The Fitter automatically works to achieve such margin when compiling any design, so you do not need to take any action.

### When Placement and Routing May Not Be Preserved Exactly

The Fitter may have to refit affected nodes if the two nodes are assigned to the same location, due to imported netlists or empty partitions set to re-use a previous post-fit netlist. There are two cases in which routing information cannot be preserved exactly. First, when multiple partitions are imported, there might be routing conflicts because two lower-level blocks could be using the same routing wire, even if the floorplan assignments of the lower-level blocks do not overlap. These routing conflicts are automatically resolved by the Quartus II Fitter re-routing on the affected nets. Second, if an imported LogicLock region

is moved in the top-level design, the relative placement of the nodes is preserved but the routing cannot be preserved, because the routing connectivity is not perfectly uniform throughout a device.

## Using Incremental Compilation With Quartus II Archive Files

The post-synthesis and post-fitting netlist information for each design partition is stored in the project database, the **incremental\_db** directory. When you archive a project, the database information is not included in the archive unless you include the compilation database in the **.qar** file.

If you want to re-use post-synthesis or post-fitting results, include the database files in the **Archive Project** dialog box so compilation results are preserved. Click **Advanced**, and choose a file set that includes the compilation database, or turn on **Incremental compilation database files** to create a Custom file set.

When you include the database, the file size of the **.qar** archive file may be significantly larger than an archive without the database.

The netlist information for imported partitions is already saved in the corresponding **.qxp**. Imported **.qxp** files are automatically saved in a subdirectory called **imported\_partitions**, so you do not need to archive the project database to keep the results for imported partitions. When you restore a project archive, the partition is automatically reimported from the **.qxp** in this directory if it is available.

For new device families with advanced support, a version-compatible database might not be available. In this case, the archive will not include the compilation database. If you require the database files to reproduce the compilation results in the same Quartus II version, you can use the following command-line option to archive a full database:

```
quartus_sh --archive -use_file_set full_db [-revision <revision name>]<project name>
```

## Formal Verification Support

You cannot use design partitions for incremental compilation if you are creating a netlist for a formal verification tool.

## SignalProbe Pins and Engineering Change Orders

ECO and SignalProbe changes are performed only during ECO and SignalProbe compilations. Other compilation flows do not preserve these netlist changes.

When incremental compilation is turned on and your design contains one or more design partitions, partition boundaries are ignored while making ECO changes and SignalProbe signal settings. However, the presence of ECO and/or SignalProbe changes does not affect partition boundaries for incremental compilation. During subsequent compilations, ECO and SignalProbe changes are not preserved regardless of the **Netlist Type** or **Fitter Preservation Level** settings. To recover ECO changes and SignalProbe signals, you must use the Change Manager to re-apply the ECOs after compilation.

For partitions developed independently in separate Quartus II projects, the exported netlist includes all currently saved ECO changes and SignalProbe signals. If you make any ECO or SignalProbe changes that affect the interface to the lower-level partition, the software issues a warning message during the export process that this netlist does not work in the top-level design without modifying the top-level HDL code to reflect the lower-level change. After integrating the **.qxp** partition into the top-level design, the ECO changes will not appear in the Change Manager.

### Related Information

- [Quick Design Debugging Using SignalProbe documentation](#)
- [Engineering Change Management with the Chip Planner documentation](#)

## SignalTap II Logic Analyzer in Exported Partitions

You can use the SignalTap II Embedded Logic Analyzer in any project that you can compile and program into an Altera device.

When incremental compilation is turned on, debugging logic is added to your design incrementally and you can tap post-fitting nodes and modify triggers and configuration without recompiling the full design. Use the appropriate filter in the Node Finder to find your node names. Use **SignalTap II: post-fitting** if the netlist type is Post-Fit to incrementally tap node names in the post-fit netlist database. Use **SignalTap II: pre-synthesis** if the netlist type is **Source File** to make connections to the source file (pre-synthesis) node names when you synthesize the partition from the source code.

If incremental compilation is turned off, the debugging logic is added to the design during Analysis and Elaboration, and you cannot tap post-fitting nodes or modify debug settings without fully compiling the design.

For design partitions that are being developed independently in separate Quartus II projects and contain the logic analyzer, when you export the partition, the Quartus II software automatically removes the SignalTap II logic analyzer and related SLD\_HUB logic. You can tap any nodes in a Quartus II project, including nodes within **.qxp** partitions. Therefore, you can use the logic analyzer within the full top-level design to tap signals from the **.qxp** partition.

You can also instantiate the SignalTap II IP core directly in your lower-level design (instead of using an **.stp** file) and export the entire design to the top level to include the logic analyzer in the top-level design.

### Related Information

[Design Debugging Using the SignalTap II Embedded Logic Analyzer documentation](#)

## External Logic Analyzer Interface in Exported Partitions

You can use the Logic Analyzer Interface in any project that you can compile and program into an Altera device. You cannot export a partition that uses the Logic Analyzer Interface. You must disable the Logic Analyzer Interface feature and recompile the design before you export the design as a partition.

### Related Information

[In-System Debugging Using External Logic Analyzers documentation](#)

## Assignments Made in HDL Source Code in Exported Partitions

Assignments made with I/O primitives or the `altera_attribute` HDL synthesis attribute in lower-level partitions are passed to the top-level design, but do not appear in the top-level **.qsf** file or Assignment Editor. These assignments are considered part of the source netlist files. You can override assignments made in these source files by changing the value with an assignment in the top-level design.

## Design Partition Script Limitations

### Related Information

[Generating Design Partition Scripts](#) on page 3-35

### Warnings About Extra Clocks Due to Design Partition Scripts

The generated scripts include applicable clock information for all clock signals in the top-level design. Some of those clocks may not exist in the lower-level projects, so you may see warning messages related to clocks that do not exist in the project. You can ignore these warnings or edit your constraints so the messages are not generated.

### Synopsys Design Constraint Files for the TimeQuest Timing Analyzer in Design Partition Scripts

After you have compiled a design using TimeQuest constraints, and the timing assignments option is turned on in the scripts, a separate Tcl script is generated to create an .sdc file for each lower-level project. This script includes only clock constraints and minimum and maximum delay settings for the TimeQuest Timing Analyzer.

**Note:** PLL settings and timing exceptions are not passed to lower-level designs in the scripts.

### Related Information

- [Best Practices for Incremental Compilation Partitions and Floorplan Assignments documentation](#) on page 14-1

### Wildcard Support in Design Partition Scripts

When applying constraints with wildcards, note that wildcards are not analyzed across hierarchical boundaries. For example, an assignment could be made to these nodes: `Top|A:inst|B:inst|*`, where A and B are lower-level partitions, and hierarchy B is a child of A, that is B is instantiated in hierarchy A. This assignment is applied to modules A, B, and all children instances of B. However, the assignment `Top|A:inst|B:inst*` is applied to hierarchy A, but is not applied to the B instances because the single level of hierarchy represented by `B:inst*` is not expanded into multiple levels of hierarchy. To avoid this issue, ensure that you apply the wildcard to the hierarchical boundary if it should represent multiple levels of hierarchy.

When using the wildcard to represent a level of hierarchy, only single wildcards are supported. This means assignments such as `Top|A:inst|*|B:inst|*` are not supported. The Quartus II software issues a warning in these cases.

### Derived Clocks and PLLs in Design Partition Scripts

If a clock in the top level is not directly connected to a pin of a lower-level partition, the lower-level partition does not receive assignments and constraints from the top-level pin in the design partition scripts.

This issue is of particular importance for clock pins that require timing constraints and clock group settings. Problems can occur if your design uses logic or inversion to derive a new clock from a clock input pin. Make appropriate timing assignments in your lower-level Quartus II project to ensure that clocks are not unconstrained.

If the lower-level design uses the top-level project framework from the project lead, the design will have all the required information about the clock and PLL settings. Otherwise, if you use a PLL in your top-level design and connect it to lower-level partitions, the lower-level partitions do not have information

about the multiplication or phase shift factors in the PLL. Make appropriate timing assignments in your lower-level Quartus II project to ensure that clocks are not unconstrained or constrained with the incorrect frequency. Alternatively, you can manually duplicate the top-level derived clock logic or PLL in the lower-level design file to ensure that you have the correct multiplication or phase-shift factors, compensation delays and other PLL parameters for complete and accurate timing analysis. Create a design partition for the rest of the lower-level design logic for export to the top level. When the lower-level design is complete, export only the partition that contains the relevant logic.

### Pin Assignments for GXB and LVDS Blocks in Design Partition Scripts

Pin assignments for high-speed GXB transceivers and hard LVDS blocks are not written in the scripts. You must add the pin assignments for these hard IP blocks in the lower-level projects manually.

### Virtual Pin Timing Assignments in Design Partition Scripts

Design partition scripts use `INPUT_MAX_DELAY` and `OUTPUT_MAX_DELAY` assignments to specify inter-partition delays associated with input and output pins, which would not otherwise be visible to the project. These assignments require that the software specify the clock domain for the assignment and set this clock domain to `" * "`.

This clock domain assignment means that there may be some paths constrained and reported by the timing analysis engine that are not required.

To restrict which clock domains are included in these assignments, edit the generated scripts or change the assignments in your lower-level Quartus II project. In addition, because there is no known clock associated with the delay assignments, the software assumes the worst-case skew, which makes the paths seem more timing critical than they are in the top-level design. To make the paths appear less timing-critical, lower the delay values from the scripts. If required, enter negative numbers for input and output delay values.

### Top-Level Ports that Feed Multiple Lower-Level Pins in Design Partition Scripts

When a single top-level I/O port drives multiple pins on a lower-level module, it unnecessarily restricts the quality of the synthesis and placement at the lower-level. This occurs because in the lower-level design, the software must maintain the hierarchical boundary and cannot use any information about pins being logically equivalent at the top level. In addition, because I/O constraints are passed from the top-level pin to each of the children, it is possible to have more pins in the lower level than at the top level. These pins use top-level I/O constraints and placement options that might make them impossible to place at the lower level. The software avoids this situation whenever possible, but it is best to avoid this design practice to avoid these potential problems. Restructure your design so that the single I/O port feeds the design partition boundary and the single connection is split into multiple signals within the lower-level partition.

### Restrictions on IP Core Partitions

The Quartus II software does not support partitions for IP core instantiations. If you use the parameter editor to customize an IP core variation, the IP core generated wrapper file instantiates the IP core. You can create a partition for the IP core generated wrapper file.

The Quartus II software does not support creating a partition for inferred IP cores (that is, where the software infers an IP core to implement logic in your design). If you have a module or entity for the logic that is inferred, you can create a partition for that hierarchy level in the design.



The Quartus II software does not support creating a partition for any Quartus II internal hierarchy that is dynamically generated during compilation to implement the contents of an IP core.

## Register Packing and Partition Boundaries

The Quartus II software performs register packing during compilation automatically. However, when incremental compilation is enabled, logic in different partitions cannot be packed together because partition boundaries might prevent cross-boundary optimization. This restriction applies to all types of register packing, including I/O cells, DSP blocks, sequential logic, and unrelated logic. Similarly, logic from two partitions cannot be packed into the same ALM.

## I/O Register Packing

Cross-partition register packing of I/O registers is allowed in certain cases where your input and output pins exist in the top-level hierarchy (and the Top partition), but the corresponding I/O registers exist in other partitions.

The following specific circumstances are required for input pin cross-partition register packing:

- The input pin feeds exactly one register.
- The path between the input pin and register includes only input ports of partitions that have one fan-out each.

The following specific circumstances are required for output register cross-partition register packing:

- The register feeds exactly one output pin.
- The output pin is fed by only one signal.
- The path between the register and output pin includes only output ports of partitions that have one fan-out each.

Output pins with an output enable signal cannot be packed into the device I/O cell if the output enable logic is part of a different partition from the output register. To allow register packing for output pins with an output enable signal, structure your HDL code or design partition assignments so that the register and tri-state logic are defined in the same partition.

Bidirectional pins are handled in the same way as output pins with an output enable signal. If the registers that need to be packed are in the same partition as the tri-state logic, you can perform register packing.

The restrictions on tri-state logic exist because the I/O atom (device primitive) is created as part of the partition that contains tri-state logic. If an I/O register and its tri-state logic are contained in the same partition, the register can always be packed with tri-state logic into the I/O atom. The same cross-partition register packing restrictions also apply to I/O atoms for input and output pins. The I/O atom must feed the I/O pin directly with exactly one signal. The path between the I/O atom and the I/O pin must include only ports of partitions that have one fan-out each.

### Related Information

- [Best Practices for Incremental Compilation Partitions and Floorplan Assignments documentation](#) on page 14-1

## Scripting Support

You can run procedures and make settings described in this chapter in a Tcl script or at a command-line prompt.

### Tcl Scripting and Command-Line Examples

The `::quartus::incremental_compilation` Tcl package contains a set of functions for manipulating design partitions and settings related to the incremental compilation feature.

#### Related Information

- [::quartus::incremental\\_compilation online help](#)
- [Quartus II Software Scripting Support website](#)  
Scripting support information, design examples, and training
- [Tcl Scripting documentation](#)
- [Command-Line Scripting documentation](#)

### Creating Design Partitions

To create a design partition to a specified hierarchy name, use the following command:

#### Example 3-1: Create Design Partition

```
create_partition [-h | -help] [-long_help] -contents
<hierarchy name> -partition <partition name>
```

**Table 3-4: Tcl Script Command: `create_partition`**

Argument	Description
<code>-h   -help</code>	Short help
<code>-long_help</code>	Long help with examples and possible return values
<code>-contents &lt;hierarchy name&gt;</code>	Partition contents (hierarchy assigned to Partition)
<code>-partition &lt;partition name&gt;</code>	Partition name

### Enabling or Disabling Design Partition Assignments During Compilation

To direct the Quartus II Compiler to enable or disable design partition assignments during compilation, use the following command:

#### Example 3-2: Enable or Disable Partition Assignments During Compilation

```
set_global_assignment -name IGNORE_PARTITIONS <value>
```

**Table 3-5: Tcl Script Command: `set_global_assignment`**

Value	Description
OFF	The Quartus II software recognizes the design partitions assignments set in the current Quartus II project and recompiles the partition in subsequent compilations depending on their netlist status.
ON	The Quartus II software does not recognize design partitions assignments set in the current Quartus II project and performs a compilation without regard to partition boundaries or netlists.

## Setting the Netlist Type

To set the partition netlist type, use the following command:

### Example 3-3: Set Partition Netlist Type

```
set_global_assignment -name PARTITION_NETLIST_TYPE <value>
-section_id <partition name>
```

**Note:** The `PARTITION_NETLIST_TYPE` command accepts the following values: `SOURCE`, `POST_SYNTH`, `POST_FIT`, and `EMPTY`.

## Setting the Fitter Preservation Level for a Post-fit or Imported Netlist

To set the Fitter Preservation Level for a post-fit or imported netlist, use the following command:

### Example 3-4: Set Fitter Preservation Level

```
set_global_assignment -name PARTITION_FITTER_PRESERVATION_LEVEL
<value> -section_id <partition name>
```

**Note:** The `PARTITION_FITTER_PRESERVATION` command accepts the following values: `NETLIST_ONLY`, `PLACEMENT`, and `PLACEMENT_AND_ROUTING`.

## Preserving High-Speed Optimization

To preserve high-speed optimization for tiles contained within the selected partition, use the following command:

### Example 3-5: Preserve High-Speed Optimization

```
set_global_assignment -name PARTITION_PRESERVE_HIGH_SPEED_TILES_ON
```

## Specifying the Software Should Use the Specified Netlist and Ignore Source File Changes

To specify that the software should use the specified netlist and ignore source file changes, even if the source file has changed since the netlist was created, use the following command:

### Example 3-6: Specify Netlist and Ignore Source File Changes

```
set_global_assignment -name PARTITION_IGNORE_SOURCE_FILE_CHANGES ON  
-section_id "<partition name>"
```

## Reducing Opening a Project, Creating Design Partitions, and Performing an Initial Compilation

Scenario background: You open a project called `AB_project`, set up two design partitions, entities A and B, and then perform an initial full compilation.

### Example 3-7: Set Up and Compile `AB_project`

```
set project AB_project  
  
load_package incremental_compilation  
load_package flow  
project_open $project  
  
# Ensure that design partition assignments are not ignored  
set_global_assignment -name IGNORE_PARTITIONS \ OFF  
  
# Set up the partitions  
create_partition -contents A -name "Partition_A"  
create_partition -contents B -name "Partition_B"  
  
# Set the netlist types to post-fit for subsequent  
# compilations (all partitions are compiled during the  
# initial compilation since there are no post-fit netlists)  
set_partition -partition "Partition_A" -netlist_type POST_FIT  
set_partition -partition "Partition_B" -netlist_type POST_FIT  
  
# Run initial compilation  
export_assignments  
execute_flow -full_compile  
  
project_close
```

## Optimizing the Placement for a Timing-Critical Partition

Scenario background: You have run the initial compilation shown in the example script below. You would like to apply Fitter optimizations, such as physical synthesis, only to partition A. No changes have been made to the HDL files. To ensure the previous compilation result for partition B is preserved, and to ensure that Fitter optimizations are applied to the post-synthesis netlist of partition A, set the netlist type of B to **Post-Fit** (which was already done in the initial compilation, but is repeated here for safety), and the netlist type of A to **Post-Synthesis**, as shown in the following example:

### Example 3-8: Fitter Optimization for `AB_project`

```
set project AB_project  
  
load_package flow
```

```

load_package incremental_compilation
load_package project
project_open $project

# Turn on Physical Synthesis Optimization
set_high_effort_fmax_optimization_assignments

# For A, set the netlist type to post-synthesis
set_partition -partition "Partition_A" -netlist_type POST_SYNTH

# For B, set the netlist type to post-fit
set_partition -partition "Partition_B" -netlist_type POST_FIT

# Also set Top to post-fit
set_partition -partition "Top" -netlist_type POST_FIT

# Run incremental compilation
export_assignments
execute_flow -full_compile

project_close

```

## Generating Design Partition Scripts

To generate design partition scripts, use the following script:

### Example 3-9: Generate Partition Script

```

# load required package
load_package database_manager

# name and open the project
set project <project_path/project_name>
project_open $project

# generate the design partition scripts
generate_bottom_up_scripts <options>

#close project
project_close

```

#### Related Information

[Generate Design Partition Scripts Dialog Box online help](#)

## Exporting a Partition

To open a project and load the `::quartus::incremental_compilation` package before you use the Tcl commands to export a partition to a **.qxp** that contains both a post-synthesis and post-fit netlist, with routing, use the following script:

### Example 3-10: Export .qxp

```

# load required package
load_package incremental_compilation

# open project
project_open <project name>

# export partition to the .qxp and set preservation level

```

```
export_partition -partition <partition name>
-qxp <.qxp file name> -<options>

#close project
project_close
```

## Importing a Partition into the Top-Level Design

To import a .qxp into a top-level design, use the following script:

### Example 3-11: Import .qxp into Top-Level Design

```
# load required packages
load_package incremental_compilation
load_package project
load_package flow

# open project
project_open <project name>

#import partition
import_partition -partition <partition name> -qxp <.qxp file>
-<options>

#close project
project_close
```

## Makefiles

For an example of how to use incremental compilation with a `makefile` as part of the team-based incremental compilation design flow, refer to the **read\_me.txt** file that accompanies the `incr_comp` example located in the `/qdesigns/incr_comp_makefile` subdirectory.

When using a team-based incremental compilation design flow, the **Generate Design Partition Scripts** dialog box can write makefiles that automatically export lower-level design partitions and import them into the top-level design whenever design files change.

### Related Information

[Generate Design Partition Scripts Dialog Box online help](#)

## Document Revision History

Table 3-6: Document Revision History

Date	Version	Changes
2014.12.15	14.1.0	<ul style="list-style-type: none"> <li>Updated location of Fitter Settings, Analysis &amp; Synthesis Settings, and Physical Optimization Settings to Compiler Settings.</li> <li>Updated DSE II content.</li> </ul>
2014.08.18	14.0a10.0	Added restriction about smart compilation in Arria 10 devices.

Date	Version	Changes
June 2014	14.0.0	<ul style="list-style-type: none"> <li>• Dita conversion.</li> <li>• Replaced MegaWizard Plug-In Manager content with IP Catalog and Parameter Editor content.</li> <li>• Revised functional safety section. Added export and import sections.</li> </ul>
November 2013	13.1.0	Removed HardCopy device information. Revised information about Rapid Recompile. Added information about functional safety. Added information about flattening sub-partition hierarchies.
November 2012	12.1.0	Added Turning On Supported Cross-boundary Optimizations.
June 2012	12.0.0	Removed survey link.
November 2011	11.0.1	Template update.
May 2011	11.0.0	<ul style="list-style-type: none"> <li>• Updated “Tcl Scripting and Command-Line Examples”.</li> </ul>
December 2010	10.1.0	<ul style="list-style-type: none"> <li>• Changed to new document template.</li> <li>• Reorganized Tcl scripting section. Added description for new feature: <b>Ignore partitions assignments during compilation</b> option.</li> <li>• Reorganized “Incremental Compilation Summary” section.</li> </ul>
July 2010	10.0.0	<ul style="list-style-type: none"> <li>• Removed the explanation of the “bottom-up design flow” where designers work completely independently, and replaced with Altera’s recommendations for team-based environments where partitions are developed in the same top-level project framework, plus an explanation of the bottom-up process for including independent partitions from third-party IP designers.</li> <li>• Expanded the Merge command explanation to explain how it now accommodates cross-partition boundary optimizations.</li> <li>• Restructured Altera recommendations for when to use a floorplan.</li> <li>• Added “Viewing the Contents of a Quartus II Exported Partition File (.qxp)” section.</li> <li>• Reorganized chapter to make design flow scenarios more visible; integrated into various sections rather than at the end of the chapter.</li> </ul>

Date	Version	Changes
October 2009	9.1.0	<ul style="list-style-type: none"> <li>Redefined the bottom-up design flow as team-based and reorganized previous design flow examples to include steps on how to pass top-level design information to lower-level designers.</li> <li>Moved SDC Constraints from Lower-Level Partitions section to the <i>Best Practices for Incremental Compilation Partitions and Floorplan Assignments</i> chapter in volume 1 of the <i>Quartus II Handbook</i>.</li> <li>Reorganized the “Conclusion” section.</li> <li>Removed HardCopy APEX and HardCopy Stratix Devices section.</li> </ul>
March 2009	9.0.0	<ul style="list-style-type: none"> <li>Split up netlist types table</li> <li>Moved “Team-Based Incremental Compilation Design Flow” into the “Including or Integrating partitions into the Top-Level Design” section.</li> <li>Added new section “Including or Integrating Partitions into the Top-Level Design”.</li> <li>Removed “Exporting a Lower-Level Partition that Uses a JTAG Feature” restriction</li> <li>Other edits throughout chapter</li> </ul>
November 2008	8.1.0	<ul style="list-style-type: none"> <li>Added new section “Importing SDC Constraints from Lower-Level Partitions” on page 2–44</li> <li>Removed the Incremental Synthesis Only option</li> <li>Removed section “OpenCore Plus Feature for MegaCore Functions in Bottom-Up Flows”</li> <li>Removed section “Compilation Time with Physical Synthesis Optimizations”</li> <li>Added information about using a <b>.qxp</b> as a source design file without importing</li> <li>Reorganized several sections</li> <li>Updated Figure 2–10</li> </ul>

**Related Information**

[Quartus II Handbook Archive](#)



2014.12.15

QII5V1



Subscribe



Send Feedback

The Partial Reconfiguration (PR) feature in the Quartus II software allows you to reconfigure a portion of the FPGA dynamically, while the remainder of the device continues to operate. The Quartus II software supports the PR feature for the Altera® Stratix® V device family.

This chapter assumes a basic knowledge of Altera's FPGA design flow, incremental compilation, and LogicLock™ region features available in the Quartus II software. It also assumes knowledge of the internal FPGA resources such as logic array blocks (LABs), memory logic array blocks (MLABs), memory types (RAM and ROM), DSP blocks, clock networks.

**Note:** For assistance with support for partial reconfiguration with the Arria® V or Cyclone® V device families, file a service request at *mySupport* using the link below.

## Related Information

- [Terminology](#) on page 4-1
- [An Example of a Partial Reconfiguration Design](#) on page 4-4
- [Partial Reconfiguration Design Flow](#) on page 4-7
- [Implementation Details for Partial Reconfiguration](#) on page 4-18
- [Partial Reconfiguration with an External Host](#) on page 4-25
- [Partial Reconfiguration with an Internal Host](#) on page 4-27
- [Partial Reconfiguration Project Management](#) on page 4-28
- [Programming Files for a Partial Reconfiguration Project](#) on page 4-30
- [Partial Reconfiguration Known Limitations](#) on page 4-37
- [mySupport](#)

## Terminology

The following terms are commonly used in this chapter.

**project:** A Quartus II project contains the design files, settings, and constraints files required for the compilation of your design.

**revision:** In the Quartus II software, a revision is a set of assignments and settings for one version of your design. A Quartus II project can have several revisions, and each revision has its own set of assignments and settings. A revision helps you to organize several versions of your design into a single project.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



**incremental compilation:** This is a feature of the Quartus II software that allows you to preserve results of previous compilations of unchanged parts of the design, while changing the implementation of the parts of your design that you have modified since your previous compilation of the project. The key benefits include timing preservation and compile time reduction by only compiling the logic that has changed.

**partition:** You can partition your design along logical hierarchical boundaries. Each design partition is independently synthesized and then merged into a complete netlist for further stages of compilation. With the Quartus II incremental compilation flow, you can preserve results of unchanged partitions at specific preservation levels. For example, you can set the preservation levels at post-synthesis or post-fit, for iterative compilations in which some part of the design is changed. A partition is only a logical partition of the design, and does not necessarily refer to a physical location on the device. However, you may associate a partition with a specific area of the FPGA by using a floorplan assignment.

For more information on design partitions, refer to the *Best Practices for Incremental Compilation Partitions and Floorplan Assignments* chapter in the *Quartus II Handbook*.

**LogicLock region:** A LogicLock region constrains the placement of logic in your design. You can associate a design partition with a LogicLock region to constrain the placement of the logic in the partition to a specific physical area of the FPGA.

For more information about LogicLock regions, refer to the *Analyzing and Optimizing the Design Floorplan with the Chip Planner* chapter in the *Quartus II Handbook*.

**PR project:** Any Quartus II design project that uses the PR feature.

**PR region:** A design partition with an associated contiguous LogicLock region in a PR project. A PR project can have one or more PR regions that can be partially reconfigured independently. A PR region may also be referred to as a PR partition.

**static region:** The region outside of all the PR regions in a PR project that cannot be reprogrammed with partial reconfiguration (unless you reprogram the entire FPGA). This region is called the static region, or fixed region.

**persona:** A PR region has multiple implementations. Each implementation is called a persona. PR regions can have multiple personas. In contrast, static regions have a single implementation or persona.

**PR control block:** Dedicated block in the FPGA that processes the PR requests, handshake protocols, and verifies the CRC.

#### Related Information

[Best Practices for Incremental Compilation Partitions and Floorplan Assignments](#) on page 14-1

[Analyzing and Optimizing the Design Floorplan with the Chip Planner](#)

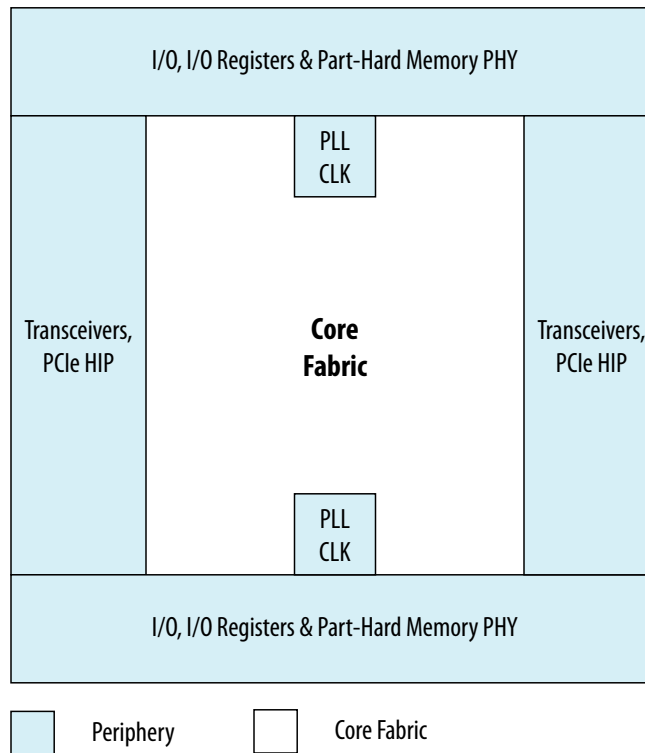
## Determining Resources for Partial Reconfiguration

You can use partial reconfiguration to configure only the resources such as LABs, embedded memory blocks, and DSP blocks in the FPGA core fabric that are controlled by configuration RAM (CRAM).

The functions in the periphery, such as GPIOs or I/O Registers, are controlled by I/O configuration bits and therefore cannot be partially reconfigured. Clock multiplexers for `GCLK` and `QCLK` are also not partially reconfigurable because they are controlled by I/O periphery bits.

### Figure 4-1: Partially Reconfigurable Resources

These are the types of resource blocks in a Stratix V device.



**Table 4-1: Reconfiguration Modes of the FPGA Resource Block**

The following table describes the reconfiguration type supported by each FPGA resource block, which are shown in the figure.

Hardware Resource Block	Reconfiguration Mode
Logic Block	Partial Reconfiguration
Digital Signal Processing	Partial Reconfiguration
Memory Block	Partial Reconfiguration
Transceivers	Dynamic Reconfiguration ALTGX_Reconfig
PLL	Dynamic Reconfiguration ALTGX_Reconfig
Core Routing	Partial Reconfiguration
Clock Networks	Clock network sources cannot be changed, but a PLL driving a clock network can be dynamically reconfigured

Hardware Resource Block	Reconfiguration Mode
I/O Blocks and Other Periphery	Not supported

The transceivers and PLLs in Altera FPGAs can be reconfigured using dynamic reconfiguration. For more information on dynamic reconfiguration, refer to the *Dynamic Reconfiguration in Stratix V Devices* chapter in the *Stratix V Handbook*.

#### Related Information

[Dynamic Reconfiguration in Stratix V Devices](#)

## An Example of a Partial Reconfiguration Design

A PR design is divided into two parts. The static region where the design logic does not change, and one or more PR regions.


Each PR region can have different design personas, that change with partial reconfiguration.

PR Region A has three personas associated with it; A1, A2, and A3. PR Region B has two personas; B1 and B2. Each persona for the two PR regions can implement different application specific logic, and using partial reconfiguration, the persona for each PR region can be modified without interrupting the operation of the device in the static or other PR region.

When a region can access more than one persona, you must create control logic to swap between personas for a PR region.

**Figure 4-2: Partial Reconfiguration Project Structure**

The following figure shows the top-level of a PR design, which includes a static region and two PR regions.



## Partial Reconfiguration Modes

When you implement a design on an Altera FPGA device, your design implementation is controlled by bits stored in CRAM inside the FPGA.

You can use partial reconfiguration in the SCRUB mode or the AND/OR mode. The mode you select affects your PR flow in ways detailed later in this chapter.

The CRAM bits control individual LABs, MLABs, M20K memory blocks, DSP blocks, and routing multiplexers in a design. The CRAM bits are organized into a frame structure representing vertical areas that correspond to specific locations on the FPGA. If you change a design and reconfigure the FPGA in a non-PR flow, the process reloads all the CRAM bits to a new functionality.

Configuration bitstreams used in a non-PR flow are different than those used in a PR flow. In addition to standard data and CRC check bits, configuration bitstreams for partial reconfiguration also include instructions that direct the PR control block to process the data for partial reconfiguration.

The configuration bitstream written into the CRAM is organized into configuration frames. If a LAB column passes through multiple PR regions, those regions share some programming frames.

## SCRUB Mode

In the SCRUB mode, the unchanging CRAM bits from the static region are "scrubbed" back to their original values. They are neither erased nor reset.

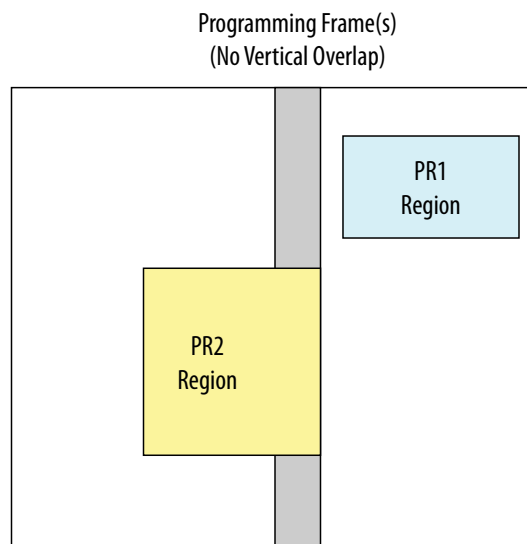
The static regions controlled by the CRAM bits from the same programming frame as the PR region continue to operate. All the CRAM bits corresponding to a PR region are overwritten with new data, regardless of what was previously contained in the region.

The SCRUB mode of partial reconfiguration involves re-writing all the bits in an entire LAB column of the CRAM, including bits controlling any PR regions above or below the region being reconfigured. As a result, it is not currently possible to correctly determine the bits associated with a PR region above or below the region being reconfigured, because those bits could have already been reconfigured and changed to an unknown value. This restriction does not apply to static bits above or below the PR region, since those bits never change and you can rewrite them with the same value as the current state of the configuration bit. You cannot use the SCRUB mode when two PR regions have a vertically overlapping column in the device.

The advantage of using the SCRUB mode is that the programming file size is much smaller than the AND/OR mode.

**Figure 4-3: SCRUB Mode**

This is the floorplan of a FPGA using SCRUB mode, with two PR regions, whose columns do not overlap.



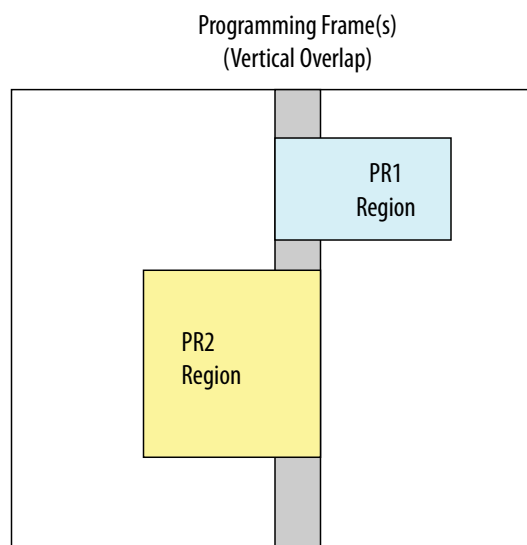
## AND/OR Mode

The AND/OR mode refers to how the bits are rewritten. Partial reconfiguration with AND/OR uses a two-pass method.

Simplistically, this can be compared to bits being ANDed with a MASK, and ORed with new values, allowing multiple PR regions to vertically overlap a single column. In the first pass, all the bits in the CRAM frame for a column passing through a PR region are ANDed with 0's while those outside the PR region are ANDed with 1's. After the first pass, all the CRAM bits corresponding to the PR region are reset without modifying the static region. In the second pass for each CRAM frame, new data is ORed with the current value of 0 inside the PR region, and in the static region, the bits are ORed with 0's so they remain unchanged. The programming file size of a PR region using the AND/OR mode could be twice the programming file size of the same PR region using SCRUB mode.

**Figure 4-4: AND/OR Mode**

This is the floorplan of a FPGA using AND/OR mode, with two PR regions, with columns that overlap.



**Note:** If you have overlapping PR regions in your design, you must use AND/OR mode to program all PR regions, including PR regions with no overlap. The Quartus II software will not permit the use of SCRUB mode when there are overlapping regions. If none of your regions overlap, you can use AND/OR, SCRUB, or a mixture of both.

## Programming File Sizes for a Partial Reconfiguration Project

The programming file size for a partial reconfiguration is proportional to the area of the PR region.

A partial reconfiguration programming bitstream for AND/OR mode makes two passes on the PR region; the first pass clears all relevant bits, and the second pass sets the necessary bits. Due to this two-pass sequence, the size of a partial bitstream can be larger than a full FPGA programming bitstream depending on the size of the PR region.

When using the AND/OR mode for partial reconfiguration, the formula which describes the approximate file size within ten percent is:

$$\text{PR bitstream size} = ((\text{Size of region in the horizontal direction}) / (\text{full horizontal dimension of the part})) * 2 * (\text{size of full bitstream})$$

The way the Fitter reserves routing for partial reconfiguration increases the effective size for small PR regions from a bitstream perspective. PR bitstream sizes in designs with a single small PR region will not match the file size computed by this equation.

**Note:** The PR bitstream size is approximately half of the size computed above when using SCRUB mode.

You can control expansion of the routing regions by adding the following two assignments to your Quartus II Settings file (**.qsf**):

```
set_global_assignment -name LL_ROUTING_REGION Expanded -section_id <region name>  
set_global_assignment -name LL_ROUTING_REGION_EXPANSION_SIZE 0 -section_id <region  
name>
```

Adding these to your **.qsf** disables expansion and minimizes the bitstream size.

## Partial Reconfiguration Design Flow

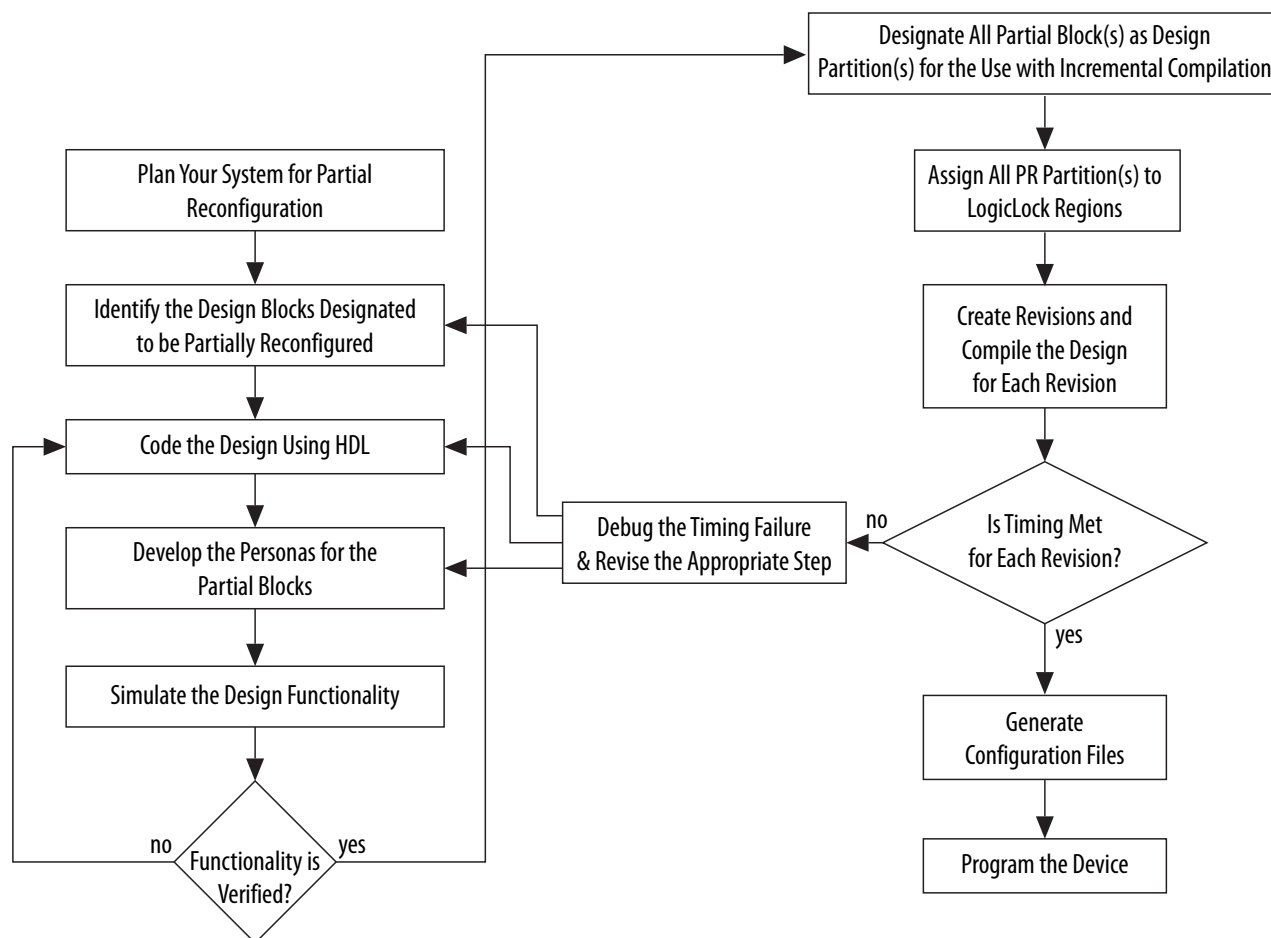
Partial reconfiguration is based on the revision feature in the Quartus II software. Your initial design is the base revision, where you define the boundaries of the static region and reconfigurable regions on the FPGA. From the base revision, you create multiple revisions, which contain the static region and describe the differences in the reconfigurable regions.

Two types of revisions are specific to partial reconfiguration: reconfigurable and aggregate. Both import the persona for the static region from the base revision. A reconfigurable revision generates personas for PR regions. An aggregate revision is used to combine personas from multiple reconfigurable revisions to create a complete design suitable for timing analysis.

The design flow for partial reconfiguration also utilizes the Quartus II incremental compilation flow. To take advantage of incremental compilation for partial reconfiguration, you must organize your design into logical and physical partitions for synthesis and fitting. For the PR flow, these partitions are treated as PR regions that must also have associated LogicLock assignments.

Revisions make use of personas, which are subsidiary archives describing the characteristics of both static and reconfigurable regions, that contain unique logic which implements a specific set of functions to reconfigure a PR region of the FPGA. Partial reconfiguration uses personas to pass this logic from one revision to another.

Figure 4-5: Partial Reconfiguration Design Flow



The PR design flow requires more initial planning than a standard design flow. Planning requires setting up the design logic for partitioning, and determining placement assignments to create a floorplan. Well-planned partitions can help improve design area utilization and performance, and make timing closure easier. You should also decide whether your system requires partial reconfiguration to originate from the FPGA pins or internally, and which mode you are using; the AND/OR mode or the SCRUB mode, because this influences some of the planning steps described in this section.

You must structure your source code or design hierarchy to ensure that logic is grouped correctly for optimization. Implementing the correct logic grouping early in the design cycle is more efficient than restructuring the code later. The PR flow requires you to be more rigorous about following good design practices. The guidelines for creating partitions for incremental compilation also include creating partitions for partial reconfiguration.



Use the following best practice guidelines for designing in the PR flow, which are described in detail in this section:

- Determining resources for partial reconfiguration
- Partitioning the design for partial reconfiguration
- Creating incremental compilation partitions for partial reconfiguration
- Instantiating the PR controller in the design
- Creating wrapper logic for PR regions
- Creating freeze logic for PR regions
- Planning clocks and other global signals for the PR design
- Creating floorplan assignments for the PR design

## Design Partitions for Partial Reconfiguration

You must create design partitions for each PR region that you want to partially reconfigure. Optionally, you can also create partitions for the static parts of the design for timing preservation and/or for reducing compilation time.

There is no limit on the number of independent partitions or PR regions you can create in your design. You can designate any partition as a PR partition by enabling that feature in the LogicLock Regions window in the Quartus II software.

Partial reconfiguration regions do not support the following IP blocks that require a connection to the JTAG controller:

- In-System Memory Content EditorI
- In-System Signals & Probes
- Virtual JTAG
- Nios II with debug module
- SignalTap II tap or trigger sources

**Note:** PR partitions can contain only core resources, they cannot contain I/O or periphery elements.

## Incremental Compilation Partitions for Partial Reconfiguration

Use the following best practices guidelines when creating partitions for PR regions in your design:

- Register all partition boundaries; register all inputs and outputs of each partition when possible. This practice prevents any delay penalties on signals that cross partition boundaries and keeps each register-to-register timing path within one partition for optimization.
- Minimize the number of paths that cross partition boundaries.
- Minimize the timing-critical paths passing in or out of PR regions. If there are timing-critical paths that cross PR region boundaries, rework the PR regions to avoid these paths.
- The Quartus II software can optimize some types of paths between design partitions for non-PR designs. However, for PR designs, such inter-partition paths are strictly not optimized.

For more information about incremental compilation, refer to the following chapter in the *Quartus II Handbook*.

### Related Information

- [Quartus II Incremental Compilation for Hierarchical and Team-Based Design](#) on page 3-1

## Partial Reconfiguration Controller Instantiation in the Design

You must instantiate the Stratix V PR control block and the Stratix V CRC block in your design in order to use the PR feature in Stratix V devices. You may find that adding the PR control block and CRC block at the top level of the design offers the most convenience.

For example, in a design named `Core_Top`, all the logic is contained under the `Core_Top` module hierarchy. Create a wrapper (`Chip_Top`) at the top-level of the hierarchy that instantiates this `Core_Top` module, the Stratix V PR control block, and the Stratix V CRC check modules.

If you are performing partial reconfiguration from pins, then the required pins should be on the I/O list for the top-level (`Chip_Top`) of the project, as shown in the code in the following examples. If you are performing partial reconfiguration from within the core, you may choose another configuration scheme, such as Active Serial, to transmit the reconfiguration data into the core, and then assemble it to 16-bit wide data inside the FPGA within your logic. In such cases, the PR pins are not part of the FPGA I/O.

**Note:** Verilog HDL does not require a component declaration. You can instantiate the PR control block as shown in the following example.

### Component Declaration of the PR Control Block and CRC Block in VHDL

This code sample has the component declaration in VHDL, showing the ports of the Stratix V PR control block and the Stratix V CRC block. In the following example, the PR function is performed from within the core (code located in `Core_Top`) and you must add additional ports to `Core_Top` to connect to both components.

```
-- The Stratix V control block interface

component stratixv_prblock is
    port(
        corectl: in STD_LOGIC ;
        prrequest: in STD_LOGIC ;
        data: in STD_LOGIC_VECTOR(15 downto 0);
        error: out STD_LOGIC ;
        ready: out STD_LOGIC ;
        done: out STD_LOGIC
    ) ;
end component ;

-- The Stratix V CRC block for diagnosing CRC errors

component stratixv_crcblock is
    port(
        shiftnld: in STD_LOGIC ;
        clk: in STD_LOGIC ;
        crcerror: out STD_LOGIC
    ) ;
end component ;
```

The following rules apply when connecting the PR control block to the rest of your design:

- The `corectl` signal must be set to '1' (when using partial reconfiguration from core) or to '0' (when using partial reconfiguration from pins).
- The `corectl` signal has to match the **Enable PR pins** option setting in the Device and Pin Options dialog box on the Setting page; if you have turned on **Enable PR pins**, then the `corectl` signal on the PR control block instantiation must be toggled to '0'.
- When performing partial reconfiguration from pins the Quartus II software automatically assigns the PR unassigned pins. If you so choose, you can make pin assignments to all the dedicated PR pins in **Pin Planner** or **Assignment Editor**.
- When performing partial reconfiguration from core, you can connect the `prblock` signals to either core logic or I/O pins, excluding the dedicated programming pin such as `DCLK`.

## Instantiating the PR Control Block and CRC Block in VHDL

This code example instantiates a PR control block in VHDL, inside your top-level project, `Chip_Top`:

```
module Chip_Top (
  //User I/O signals (excluding PR related signals)
  ..
  ..
  //PR interface & configuration signals
    pr_request,
    pr_ready,
    pr_done,
    crc_error,
    dclk,
    pr_data,
    init_done
);
//user I/O signal declaration
..
..
//PR interface and configuration signals declaration
input pr_request;
output pr_ready;
output pr_done;
output crc_error;
input dclk;
input [15:0] pr_data;
output init_done

// Following shows the connectivity within the Chip_Top module
Core_Top : Core_Top
port_map (
  ..
  ..
);

m_pr : stratixv_prblock
port map(
  clk => dclk,
  corectl => '0', //1 - when using PR from inside
                //0 - for PR from pins; You must also enable
                // the appropriate option in Quartus II settings
  prrequest => pr_request,
  data      => pr_data,
  error     => pr_error,
  ready     => pr_ready,
  done      => pr_done
);
m_crc : stratixv_crcblock
port map(
  shiftlnld=> '1', //If you want to read the EMR register when
```

```

        clk=> dummy_clk,          //error occurs, refer to AN539 for the
                                //connectivity for this signal. If you only want
                                //to detect CRC errors, but plan to take no
                                //further action, you can tie the shiftnld
                                //signal to logical high.
    crcerror    => crc_error
    );

```

For more information on port connectivity for reading the Error Message Register (EMR), refer to the following application note.

#### Related Information

[AN539: Test Methodology of Error Detection and Recovery using CRC in Altera FPGA Devices](#)

### Instantiating the PR Control Block and CRC Block in Verilog HDL

The following example instantiates a PR control block in Verilog HDL, inside your top-level project, Chip\_Top:

```

    module Chip_Top (
    //User I/O signals (excluding PR related signals)
    ..
    ..
    //PR interface & configuration signals
        pr_request,
        pr_ready,
        pr_done,
        crc_error,
        dclk,
        pr_data,
        init_done
    );
    //user I/O signal declaration
    ..
    ..
    //PR interface and configuration signals declaration
        input pr_request;
        output pr_ready;
        output pr_done;
        output crc_error;
        input dclk;
        input [15:0] pr_data;
        output init_done

    // Following shows the connectivity within the Chip_Top module
    Core_Top : Core_Top
    port_map (
        ..
        ..
    );

    m_pr : stratixv_prblock
        //set corectl to '1' when using PR from inside
        //set corectl to '0' for PR from pins. You must also enable
        // the appropriate option in Quartus II settings.
    port map(
        clk => dclk,
        corectl=> '0',
        prrequest=> pr_request,
        data=> pr_data,
        error=> pr_error,
        ready=> pr_ready,
        done=> pr_done
    );

```

```

m_crc : stratixv_crcblock
    //If you want to read the EMR register when an error occurs, refer to AN539 for
    the
    //connectivity for this signal. If you only want to detect CRC errors, but plan
    to take no
    //further action, you can tie the shiftnld signal to logical high.
    port map(
        shiftnld=> '1',
        clk=> dummy_clk,
        crcerror=> crc_error
    );

```

For more information on port connectivity for reading the Error Message Register (EMR), refer to the following application note.

#### Related Information

[AN539: Test Methodology of Error Detection and Recovery using CRC in Altera FPGA Devices](#)


## Wrapper Logic for PR Regions

Each persona of a PR region must implement the same input and output boundary ports. These ports act as the boundary between static and reconfigurable logic.

Implementing the same boundary ports ensures that all ports of a PR region remain stationary regardless of the underlying persona, so that the routing from the static logic does not change with different PR persona implementations.

**Figure 4-6: Wire-LUTs at PR Region Boundary**

The Quartus II software automatically instantiates a wire-LUT for each port of the PR region to lock down the same location for all instances of the PR persona.



If one persona of your PR region has a different number of ports than others, then you must create a wrapper so that the static region always communicates with this wrapper. In this wrapper, you can create dummy ports to ensure that all of the PR personas of a PR region have the same connection to the static region.

The sample code below each create two personas; `persona_1` and `persona_2` are different functions of one PR region. Note that one persona has a few dummy ports. The first example creates partial reconfiguration wrapper logic in Verilog HDL:

```

// Partial Reconfiguration Wrapper in Verilog HDL
module persona_1
(
    input reset,
    input [2:0] a,
    input [2:0] b,
    input [2:0] c,
    output [3:0] p,

```

```

        output [7:0] q
    );
    reg [3:0] p, q;
    always@(a or b) begin
        p = a + b ;
    end

    always@(a or b or c or p)begin
        q = (p*a - b*c )
    end
endmodule

module persona_2
(
    input reset,
    input [2:0] a,
    input [2:0] b,
    input [2:0] c, //never used in this persona
    output [3:0] p,
    output [7:0] q //never assigned in this persona
);
    reg [3:0] p, q;
    always@(a or b) begin
        p = a * b;
    // note q is not assigned value in this persona
    end
endmodule

```

The following example creates partial reconfiguration wrapper logic in VHDL.

```

-- Partial Reconfiguration Wrapper in VHDL
entity persona_1 is
    port( a:in STD_LOGIC_VECTOR (2 downto 0);
          b:in STD_LOGIC_VECTOR (2 downto 0);
          c:in STD_LOGIC_VECTOR (2 downto 0);
          p: out STD_LOGIC_VECTOR (3 downto 0);
          q: out STD_LOGIC_VECTOR (7 downto 0));
end persona_1;

architecture synth of persona_1 is
    begin
        process(a,b)
            begin
                p <= a + b;
            end process;

        process (a, b, c, p)
            begin
                q <= (p*a - b*c);
            end process;
    end synth;

entity persona_2 is
    port( a:in STD_LOGIC_VECTOR (2 downto 0);
          b:in STD_LOGIC_VECTOR (2 downto 0);
          c:in STD_LOGIC_VECTOR (2 downto 0); --never used in this persona
          p:out STD_LOGIC_VECTOR (3 downto 0);
          q:out STD_LOGIC_VECTOR (7 downto 0)); --never used in this persona
end persona_2;

architecture synth of persona_2 is
    begin
        process(a, b)
            begin

```

```

        p <= a *b; --note q is not assigned a value in this persona
    end process;
end synth;


```

## Freeze Logic for PR Regions

When you use partial reconfiguration, you must freeze all non-global inputs of a PR region except global clocks. Locally routed signals are not considered global signals, and must also be frozen during partial reconfiguration. Freezing refers to driving a '1' on those PR region inputs. When you start a partial reconfiguration process, the chip is in user mode, with the device still running.

Freezing all non-global inputs for the PR region ensures there is no contention between current values that may result in unexpected behavior of the design after partial reconfiguration is complete. Global signals going into the PR region should not be frozen to high. The Quartus II software freezes the outputs from the PR region; therefore the logic outside of the PR region is not affected.

**Figure 4-7: Freezing at PR Region Boundary**



During partial reconfiguration, the static region logic should not depend on the outputs from PR regions to be at a specific logic level for the continued operation of the static region.

The easiest way to control the inputs to PR regions is by creating a wrapper around the PR region in RTL. In addition to freezing all inputs high, you can also drive the outputs from the PR block to a specific value, if required by your design. For example, if the output drives a signal that is active high, then your wrapper could freeze the output to GND.

The following example implements a freeze wrapper in Verilog HDL, on a module named `pr_module`.

```

module freeze_wrapper
(
    input reset,
    input freeze, //PR process active, generated by user logic
    input clk1, //global clock signal
    input clk2, // non-global clock signal
    input [3:0] control_mode,
    input [3:0] framer_ctl,
    output [15:0] data_out
);
    wire [3:0]control_mode_wr, framer_ctl_wr;

```

```

wire clk2_to_wr;
//instantiate pr_module
pr_module pr_module
(
  .reset (reset), //input
  .clk1 (clk1), //input, global clock
  .clk2 (clk2_to_wr), // input, non-global clock
  .control_mode (control_mode_wr), //input
  .framer_ctl (framer_ctl_wr), //input
  .pr_module_out (data_out)// collection of outputs from pr_module
);

// Freeze all inputs

assign control_mode_wr = freeze ? 4'hF: control_mode;
assign framer_ctl_wr = freeze ? 4'hF: framer_ctl;
assign clk2_to_wr = freeze ? 1'b1 : clk2;

endmodule

```

The following example implements a freeze wrapper in VHDL, on a module named `pr_module`.

```

entity freeze_wrapper is
port( reset:in STD_LOGIC;
      freeze:in STD_LOGIC;
      clk1: in STD_LOGIC; --global signal
      clk2: in STD_LOGIC; --non-global signal
      control_mode: in STD_LOGIC_VECTOR (3 downto 0);
      framer_ctl: in STD_LOGIC_VECTOR (3 downto 0);
      data_out: out STD_LOGIC_VECTOR (15 downto 0));
end freeze_wrapper;

architecture behv of freeze_wrapper is

  component pr_module
  port(reset:in STD_LOGIC;
        clk1:in STD_LOGIC;
        clk2:in STD_LOGIC;
        control_mode:in STD_LOGIC_VECTOR (3 downto 0);
        framer_ctl:in STD_LOGIC_VECTOR (3 downto 0);
        pr_module_out:out STD_LOGIC_VECTOR (15 downto 0));
  end component

  signal control_mode_wr: in STD_LOGIC_VECTOR (3 downto 0);
  signal framer_ctl_wr : in STD_LOGIC_VECTOR (3 downto 0);
  signal clk2_to_wr : STD_LOGIC;
  signal data_out_temp : STD_LOGIC_VECTOR (15 downto 0);
  signal logic_high : STD_LOGIC_VECTOR (3 downto 0):="1111";

begin

  data_out(15 downto 0) <= data_out_temp(15 downto 0);

  m_pr_module: pr_module
  port map (
    reset => reset,
    clk1 => clk1,
    clk2 => clk2_to_wr,
    control_mode =>control_mode_wr,
    framer_ctl => framer_ctl_wr,
    pr_module_out => data_out_temp);

  -- freeze all inputs

  control_mode_wr <= logic_high when (freeze ='1') else control_mode;

```



```

framer_ctl_wr <= logic_high when (freeze = '1') else framer_ctl;

clk2_to_wr <= logic_high(0) when (freeze = '1') else clk2;

end architecture;

```

## Clocks and Other Global Signals for a PR Design

For non-PR designs, the Quartus II software automatically promotes high fan-out signals onto dedicated clocks or other forms of global signals during the pre-fitter stage of design compilation using a process called global promotion. For PR designs, however, automatic global promotion is disabled by default for PR regions, and you must assign the global clock resources necessary for PR partitions. Clock resources can be assigned by making Global Signal assignments in the Quartus II Assignment Editor, or by adding Clock Control Block (altclkctrl) Megafunction blocks in the design that drive the desired global signals.

There are 16 global clock networks in a Stratix V device. However, only six unique clocks can drive a row clock region limiting you to a maximum of six global signals in each PR region. The Quartus II software must ensure that any global clock can feed every location in the PR region.

The limit of six global signals to a PR region includes the `GCLK`, `QCLK` and `PCLKs` used inside of the PR region. Make QSF assignments for global signals in your project's Quartus II Settings File (`.qsf`), based on the clocking requirements for your design. In designs with multiple clocks that are external to the PR region, it may be beneficial to align the PR region boundaries to be within the global clock boundary (such as `QCLK` or `PCLK`).

If your PR region requires more than six global signals, modify the region architecture to reduce the number of global signals within this to six or fewer. For example, you can split a PR region into multiple regions, each of which uses only a subset of the clock domains, so that each region does not use more than six.

Every instance of a PR region that uses the global signals (for example, `PCLK`, `QCLK`, `GCLK`, `ACLK`) must use a global signal for that input.

Global signals can only be used to route certain secondary signals into a PR region and the restrictions for each block are listed in the following table. Data signals and other secondary signals not listed in the table, such as synchronous clears and clock enables are not supported.

**Table 4-2: Supported Signal Types for Driving Clock Networks in a PR Region**

Block Types	Supported Signals for Global/Periphery/Quadrant Clock Networks
LAB	Clock, ACLR
RAM	Clock, ACLR, Write Enable(WE), Read Enable(RE)
DSP	Clock, ACLR

**Note:** PR regions are allowed to contain output ports that are used outside of the PR region as global signals.

- If a global signal feeds both static and reconfigurable logic, the restrictions in the table also apply to destinations in the static region. For example, the same global signal cannot be used as an `SCLR` in the static region and an `ACLR` in the PR region.
- A global signal used for a PR region should only feed core blocks inside and outside the PR region. In particular you should not use a clock source for a PR region and additionally connect the signal to an I/O register on the top or bottom of the device. Doing so may cause the Assembler to give an error because it is unable to create valid programming mask files.

## Floorplan Assignments for PR Designs

You must create a LogicLock region so the interface of the PR region with the static region is the same for any persona you implement. If different personas of a PR region have different area requirements, you must make a LogicLock region assignment that contains enough resources to fit the largest persona for the region. The static regions in your project do not necessarily require a floorplan, but depending on any other design requirement, you may choose to create a floorplan for a specific static region. If you create multiple PR regions, and are using SCRUB mode, make sure you have one column or row of static region between each PR region.

There is no minimum or maximum size for the LogicLock region assigned for a PR region. Because wire-LUTs are added on the periphery of a PR region by the Quartus II software, the LogicLock region for a PR region must be slightly larger than an equivalent non-PR region. Make sure the PR regions include only the resources that can be partially reconfigured; LogicLock regions for PR can only contain only LABs, DSPs, and RAM blocks. When creating multiple PR regions, make sure there is at least one static region column between each PR region. When multiple PR regions are present in a design, the shape and alignment of the region determines whether you use the SCRUB or AND/OR PR mode.

You can use the default **Auto size** and **Floating location** LogicLock region properties to estimate the preliminary size and location for the PR region.

You can also define regions in the floorplan that match the general location and size of the logic in each partition. You may choose to create a LogicLock region assignment that is non-rectangular, depending on the design requirements, but disjoint LogicLock regions are not allowed for PR regions.

After compilation, use the Fitter-determined size and origin location as a starting point for your design floorplan. Check the quality of results obtained for your floorplan location assignments and make changes to the regions as needed.

Alternatively, you can perform Analysis and Synthesis, and then set the regions to the required size based on resource estimates. In this case, use your knowledge of the connections between partitions to place the regions in the floorplan.

For more information on making design partitions and using an incremental design flow, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Floorplan Design* chapter in the *Quartus II Handbook*. For more design guidelines to ensure good quality of results, and suggestions on making design floorplan assignments with LogicLock regions, refer to the *Best Practices for Incremental Compilation Partitions and Floorplan Floorplan Assignments* chapter in the *Quartus II Handbook*.

### Related Information

- [Quartus II Incremental Compilation for Hierarchical and Team-Based Floorplan](#) on page 3-1
- [Best Practices for Incremental Compilation Partitions and Floorplan](#) on page 14-1

## Implementation Details for Partial Reconfiguration

This section describes implementation details that help you create your PR design.

## Partial Reconfiguration Pins

Partial reconfiguration can be performed through external pins or from inside the core of the FPGA.

When using PR from pins, some of the I/O pins are dedicated for implementing partial reconfiguration functionality. If you perform partial reconfiguration from pins, then you must use the passive parallel with 16 data bits (FPPx16) configuration mode. All dual-purpose pins should also be specified to **Use as regular I/O**.

To enable partial reconfiguration from pins in the Quartus II software, perform the following steps:

1. From the Assignments menu, click **Device**, then click **Device and Pin Options**.
2. In the **Device and Pin Options** dialog box, select **General** in the **Category** list and turn on **Enable PR pins** from the **Options** list.
3. Click **Configuration** in the **Category** list and select **Passive Parallel x16** from the **Configuration scheme** list.
4. Click **Dual-Purpose Pins** in the **Category** list and verify that all pins are set to **Use as regular I/O** rather than **Use as input tri-stated**.
5. Click **OK**, or continue to modify other settings in the **Device and Pin Options** dialog box.
6. Click **OK**.

**Note:** You can enable open drain on PR pins from the **Device and Pin Options** dialog box in the **Settings** page of the Quartus II software.

**Table 4-3: Partial Reconfiguration Dedicated Pins Description**

Pin Name	Pin Type	Pin Description
PR_REQUEST	Input	Dedicated input when <b>Enable PR pins</b> is turned on; otherwise, available as user I/O.  Logic high on pin indicates the PR host is requesting partial reconfiguration.
PR_READY	Output	Dedicated output when <b>Enable PR pins</b> is turned on; otherwise, available as user I/O.  Logic high on this pin indicates the Stratix V control block is ready to begin partial reconfiguration.
PR_DONE	Output	Dedicated output when <b>Enable PR pins</b> is turned on; otherwise, available as user I/O.  Logic high on this pin indicates that partial reconfiguration is complete.

Pin Name	Pin Type	Pin Description
PR_ERROR	Output	Dedicated output when <b>Enable PR pins</b> is turned on; otherwise, available as user I/O.  Logic high on this pin indicates the device has encountered an error during partial reconfiguration.
DATA[15:0]	Input	Dedicated input when <b>Enable PR pins</b> is turned on; otherwise available as user I/O.  These pins provide connectivity for PR_DATA when <b>Enable PR pins</b> is turned on.
DCLK	Bidirectional	Dedicated input when <b>Enable PR pins</b> is turned on; PR_DATA is sent synchronous to this clock.  This is a dedicated programming pin, and is not available as user I/O even if <b>Enable PR pins</b> is turned off.

For more information on different configuration modes for Stratix V devices, and specifically about FPPx16 mode, refer to the *Configuration, Design Security, and Remote System Upgrades in Stratix V Devices* chapter of the *Stratix V Handbook*.

#### Related Information

[Configuration, Design Security, and Remote System Upgrades in Stratix V Devices](#)

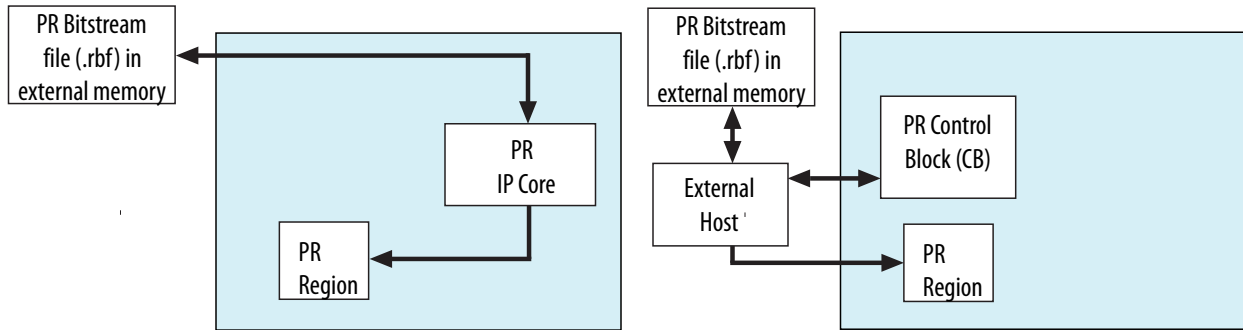
## Interface with the PR Control Block through a PR Host

You communicate between your PR control IP and the PR Control Block (CB) via control signals, while executing partial reconfiguration.

You can communicate with the PR control block via an internal host which communicates with the CB through internal control signals. You can also use an external host with handshake signals accessed via external pins. The internal PR host can be user logic or a Nios® II processor.

**Figure 4-8: Managing Partial Reconfiguration with an Internal or External Host**

The figure shows how these blocks should be connected to the PR control block (CB). In your system, you will have either the External Host or the Internal Host, but not both.



The PR mode is independent of the full chip programming mode. For example, you can configure the full chip using a JTAG download cable, or other supported configuration modes. When configuring PR regions, you must use the FPPx16 interface to the PR control block whether you choose to partially reconfigure the chip from an external or internal host.

When using an external host, you must implement the control logic for managing system aspects of partial reconfiguration on an external device. By using an internal host, you can implement all of your logic necessary for partial reconfiguration in the FPGA, therefore external devices are not required to support partial reconfiguration. When using an internal host, you can use any interface to load the PR bitstream data to the FPGA, for example, from a serial or a parallel flash device, and then format the PR bitstream data to fit the FPPx16 interface on the PR Control Block.

To use the external host for your design, turn on the **Enable PR Pins** option in the **Device and Pin Options** dialog box in the Quartus II software when you compile your design. If this setting is turned off, then you must use an internal host. Also, you must tie the `corectl` port on the PR control block instance in the top-level of the design to the appropriate level for the selected mode.

#### Related Information

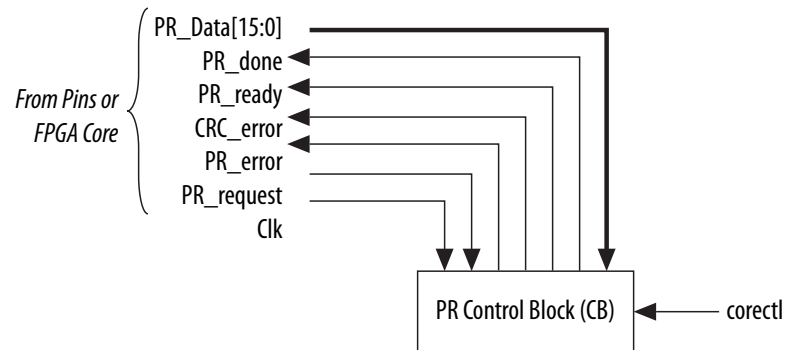
[Partial Reconfiguration Pins](#) on page 4-19  
[Partial Reconfiguration Dedicated Pins Table](#)

## PR Control Signals Interface

The Quartus II Programmer allows you to generate the different bit-streams necessary for full chip configuration and for partial reconfiguration. The programming bit-stream for partial reconfiguration contains the instructions (opcodes) as well as the configuration bits, necessary for reconfiguring each of the partial regions. When using an external host, the interface ports on the control block are mapped to FPGA pins. When using an internal host, these signals are within the core of the FPGA.

### Figure 4-9: Partial Reconfiguration Interface Signals

These handshaking control signals are used for partial reconfiguration.



- **PR\_DATA:** The configuration bitstream is sent on `PR_DATA[ 15:0 ]`, synchronous to the `Clk`.
- **PR\_DONE:** Sent from CB to control logic indicating the PR process is complete.
- **PR\_READY:** Sent from CB to control logic indicating the CB is ready to accept PR data from the control logic.
- **CRC\_ERROR:** The `CRC_Error` generated from the device's CRC block, is used to determine whether to partially reconfigure a region again, when encountering a `CRC_Error`.
- **PR\_ERROR:** Sent from CB to control logic indicating an error during partial reconfiguration.
- **PR\_REQUEST:** Sent from your control logic to CB indicating readiness to begin the PR process.
- **corectl:** Determines whether partial reconfiguration is performed internally or through pins.

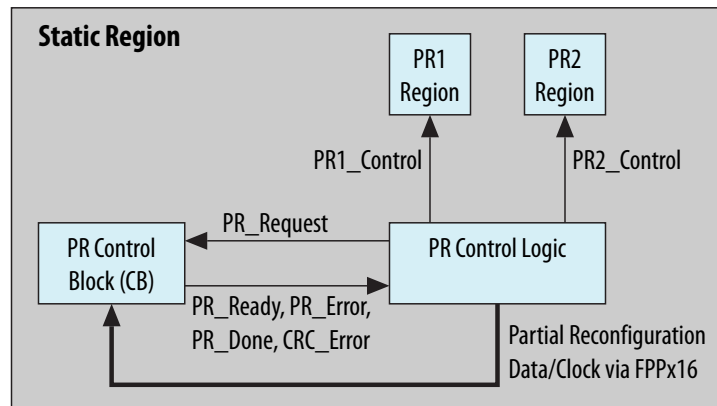
## Reconfiguring a PR Region

The figure below shows a system in which your PR Control logic is implemented inside the FPGA. However, this section is also applicable for partial reconfiguration with an external host.

The PR control block (CB) represents the Stratix V PR controller inside the FPGA. PR1 and PR2 are two PR regions in a user design. In addition to the four control signals (`PR_REQUEST`, `PR_READY`, `PR_DONE`, `PR_ERROR`) and the data/clock signals interfacing with the PR control block, your PR Control IP should also send a control signal (`PR_CONTROL`) to each PR region. This signal implements the freezing and unfreezing of the PR Interface signals. This is necessary to avoid contention on the FPGA routing fabric.

**Figure 4-10: Example of a PR System with Two PR Regions**

Implementation of PR Control logic in the FPGA.



After the FPGA device has been configured with a full chip configuration at least once, the `INIT_DONE` signal is released, and the signal is asserted high due to the external resistor on this pin. The `INIT_DONE` signal must be assigned to a pin to monitor it externally. When a full chip configuration is complete, and the device is in user mode, the following steps describe the PR sequence:

1. Begin a partial reconfiguration process from your PR Control logic, which initiates the PR process for one or more of the PR regions (asserting `PR1_Control` or `PR2_Control` in the figure). The wrapper HDL described earlier freezes (pulls high) all non-global inputs of the PR region before the PR process.
2. Send `PR_REQUEST` signal from your control logic to the PR Control Block (CB). If your design uses an external controller, monitor `INIT_DONE` to verify that the chip is in user mode before asserting the `PR_REQUEST` signal. The CB initializes itself to accept the PR data and clock stream. After that, the CB asserts a `PR_READY` signal to indicate it can accept PR data. Exactly four clock cycles must occur before sending the PR data to make sure the PR process progresses correctly. Data and clock signals are sent to the PR control block to partially reconfigure the PR region interface.
  - If there are multiple PR personas for the PR region, your PR Control IP must determine the programming file data for partial reconfiguration.
  - When there are multiple PR regions in the design, then the same PR control IP determines which regions require reconfiguration based on system requirements.
  - At the end of the PR process, the PR control block asserts a `PR_DONE` signal and de-asserts the `PR_READY` signal.
  - If you want to suspend sending data, you can implement logic to pause the clock at any point.
3. Your PR control logic must de-assert the `PR_REQUEST` signal within eight clock cycles after the `PR_DONE` signal goes high. If your logic does not de-assert the `PR_REQUEST` signal within eight clock cycles, a new PR cycle starts.
4. If your design includes additional PR regions, repeat steps 2 – 3 for each region. Otherwise, proceed to step 5.
5. Your PR Control logic de-asserts the `PR_CONTROL` signal(s) to the PR region. The freeze wrapper releases all input signals of the PR region, thus the PR region is ready for normal user operation.
6. You must perform a reset cycle to the PR region to bring all logic in the region to a known state. After partial reconfiguration is complete for a PR region, the states in which the logic in the region come up is unknown.

The PR event is now complete, and you can resume operation of the FPGA with the newly configured PR region.

At any time after the start of a partial reconfiguration cycle, the PR host can suspend sending the PR\_DATA, but the host must suspend sending the PR\_CLK at the same time. If the PR\_CLK is suspended after a PR process, there must be at least 20 clock cycles after the PR\_DONE or PR\_ERROR signal is asserted to prevent incorrect behavior.

For an overview of different reset schemes in Altera devices, refer to the *Recommended Design Practices* chapter in the *Quartus II Handbook*.

#### Related Information

[Partial Reconfiguration Cycle Waveform](#) on page 4-24

For more information on clock requirements for partial reconfiguration.

[Recommended Design Practices](#) on page 11-1

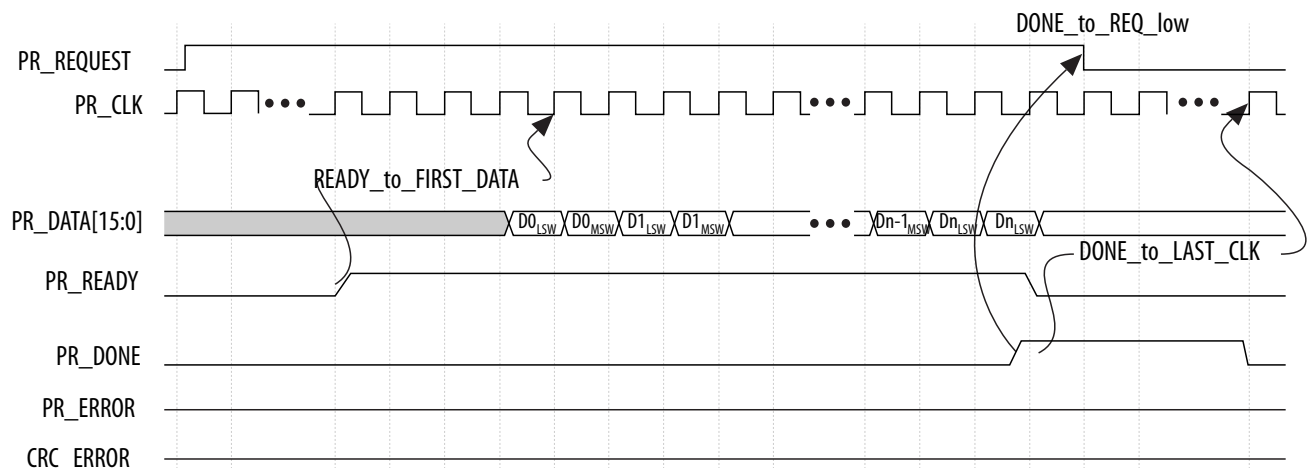
## Partial Reconfiguration Cycle Waveform

The PR host initiates the PR request, transfers the data to the FPGA device when it is ready, and monitors the PR process for any errors or until it is done.

A PR cycle is initiated by the host (internal or external) by asserting the PR\_REQUEST signal high. When the FPGA device is ready to begin partial reconfiguration, it responds by asserting the PR\_READY signal high. The PR host responds by sending configuration data on DATA [15:0]. The data is sent synchronous to PR\_CLK. When the FPGA device receives all PR data successfully, it asserts the PR\_DONE high, and de-asserts PR\_READY to indicate the completion of the PR cycle.

**Figure 4-11: Partial Reconfiguration Timing Diagram**

The partial reconfiguration cycle waveform with a hand-shaking protocol.



If there is an error encountered during partial reconfiguration, the FPGA device asserts the PR\_ERROR signal high and de-asserts the PR\_READY signal low.

The PR host must continuously monitor the PR\_DONE and PR\_ERROR signals status. Whenever either of these two signals are asserted, the host must de-assert PR\_REQUEST within eight PR\_CLK cycles. As a response to PR\_ERROR error, the host can optionally request another partial reconfiguration or perform a full FPGA configuration.

To prevent incorrect behavior, the PR\_CLK signal must be active a minimum of twenty clock cycles after PR\_DONE or PR\_ERROR signal is asserted high. Once PR\_DONE is asserted, PR\_REQUEST must be de-asserted



within eight clock cycles. `PR_DONE` is de-asserted by the device within twenty `PR_CLK` cycles. The host can assert `PR_REQUEST` again after the 20 clocks after `PR_DONE` is de-asserted.

**Table 4-4: Partial Reconfiguration Clock Requirements**

Signal timing requirements for partial reconfiguration.

Timing Parameters	Value (clock cycles)
<code>PR_READY</code> to first data	4 (exact)
<code>PR_ERROR</code> to last clock	20 (minimum)
<code>PR_DONE</code> to last clock	20 (minimum)
<code>DONE_to_REQ_low</code>	8 (maximum)
Compressed <code>PR_READY</code> to first data	4 (exact)
Encrypted <code>PR_READY</code> to first data (when using double PR)	8 (exact)
Encrypted and Compressed <code>PR_READY</code> to first data (when using double PR)	12 (exact)

At any time during partial reconfiguration, to pause sending `PR_DATA`, the PR host can stop toggling `PR_CLK`. The clock can be stopped either high or low.

At any time during partial reconfiguration, the PR host can terminate the process by de-asserting the PR request. A partially completed PR process results in a PR error. You can have the PR host restart the PR process after a failed process by sending out a new PR request 20 cycles later.

If you terminate a PR process before completion, and follow it up with a full FPGA configuration by asserting `nConfig`, then you must toggle `PR_CLK` for an additional 20 clock cycles prior to asserting `nConfig` to flush the `PR_CONTROL_BLOCK` and avoid lock up.

During these steps, the PR control block might assert a `PR_ERROR` or a `CRC_ERROR` signal to indicate that there was an error during the partial reconfiguration process. Assertion of `PR_ERROR` indicates that the PR bitstream data was corrupt, and the assertion of `CRC_ERROR` indicates a CRAM CRC error either during or after completion of PR process. If the `PR_ERROR` or `CRC_ERROR` signals are asserted, you must plan whether to reconfigure the PR region or reconfigure the whole FPGA, or leave it unconfigured.

**Important:** The `PR_CLK` signal has different a nominal maximum frequency for each device. Most Stratix V devices have a nominal maximum frequency of at least 62.5 MHz.

## Partial Reconfiguration with an External Host

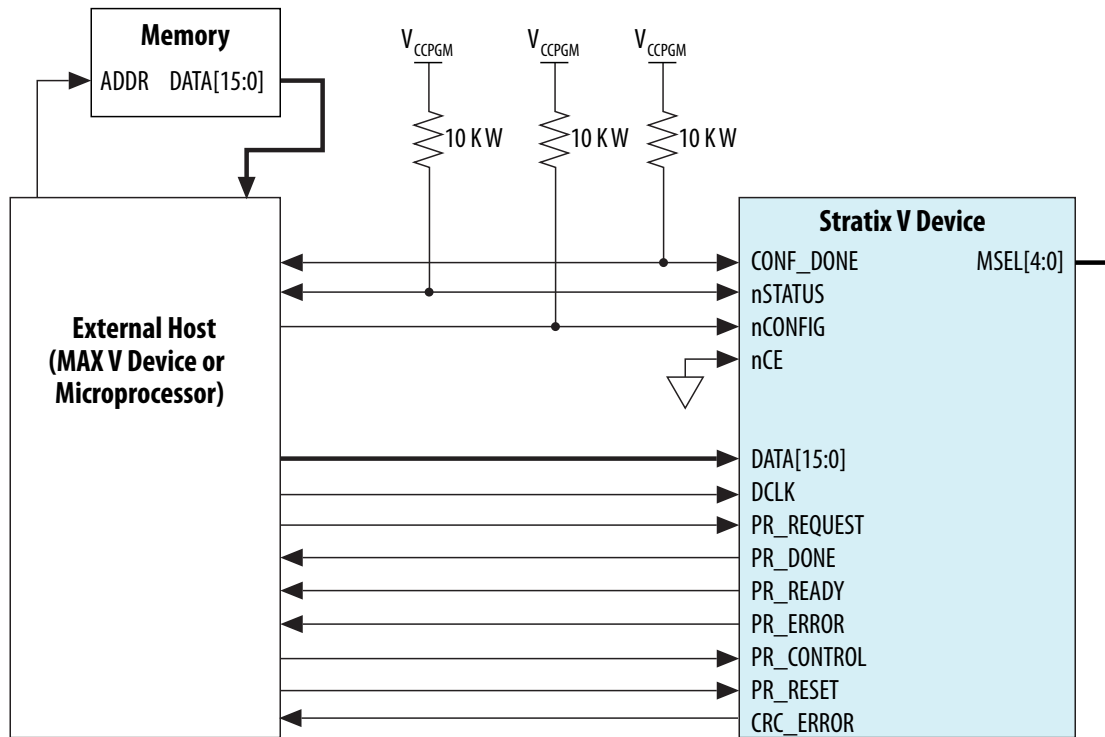
For partial reconfiguration using an external host, you must set the `MSEL [4:0]` pins for FPPx16 configuration scheme.

You can use a microcontroller, another FPGA, or a CPLD such as a MAX V device, to implement the configuration and PR controller. In this setup, the Stratix V device configures in FPPx16 mode during power-up. Alternatively, you can use a JTAG interface to configure the Stratix V device.

At any time during user-mode, the external host can initiate partial reconfiguration and monitor the status using the external PR dedicated pins: `PR_REQUEST`, `PR_READY`, `PR_DONE`, and `PR_ERROR`. In this mode, the external host must respond appropriately to the hand-shaking signals for a successful partial reconfiguration. This includes acquiring the data from the flash memory and loading it into the Stratix V device on `DATA[15:0]`.

**Figure 4-12: Connecting to an External Host**

The connection setup for partial reconfiguration with an external host in the FPPx16 configuration scheme.

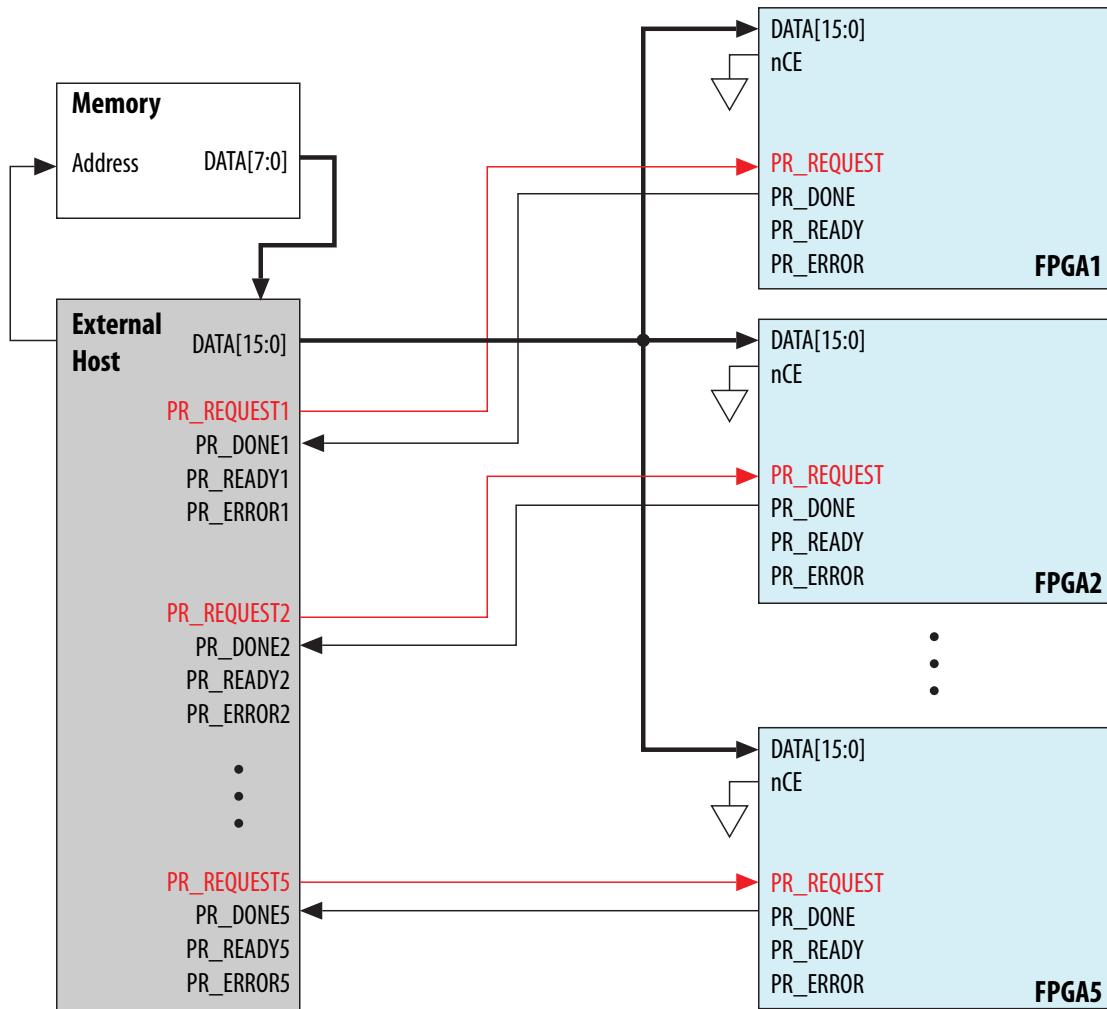


## Using an External Host with Multiple Devices

You must design the external host to accommodate the arbitration scheme that is required for your system, as well as the partial reconfiguration interface requirement for each device.

**Figure 4-13: Connecting Multiple FPGAs to an External Host**

An example of an external host controlling multiple Stratix V devices on a board.



## Partial Reconfiguration with an Internal Host

The PR internal host is a piece of soft logic implemented in the FPGA that you must design to accommodate the hand-shaking protocol with the PR control block.

For example, PR programming bitstream(s) stored in an external flash device can be routed through the regular I/Os of the FPGA device, or received through the high speed transceiver channel (PCI Express, SRIO or Gigabit Ethernet), and can be stored in on-chip memory such as MLABs or M20K blocks, for processing by the internal logic. This data must be formatted into the 16 bit wide data so that it can be transmitted to the PR control block by the internal IP, because the PR control block can only accept PR data via its FPPx16 interface.

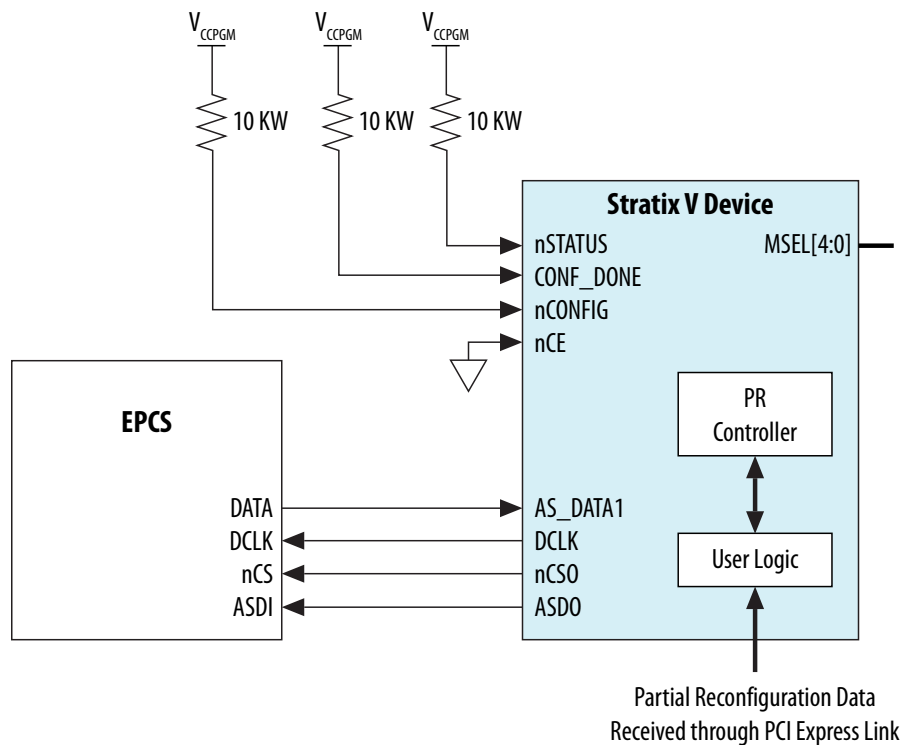
The PR dedicated pins (PR\_REQUEST, PR\_READY, PR\_DONE, and PR\_ERROR) can be used as regular I/Os when performing partial reconfiguration with an internal host. For the full FPGA configuration upon power-up, you can set the MSEL[4:0] pins to match the configuration scheme, for example, AS, PS,

FPPx8, FPPx16, or FPPx32. Alternatively, you can use the JTAG interface to configure the FPGA device. At any time during user-mode, you can initiate partial reconfiguration through the FPGA core fabric using the PR internal host.

In the following figure, the programming bitstream for partial reconfiguration is received through the PCI Express link, and your logic converts the data to the FPPx16 mode.

**Figure 4-14: Connecting to an Internal Host**

An example of the configuration setup when performing partial reconfiguration using the internal host.



## Partial Reconfiguration Project Management

When compiling your PR project, you must create a base revision, and one or more reconfigurable revisions. The project revision you start out is termed the base revision.

### Create Reconfigurable Revisions

To create a reconfigurable revision, use the **Revisions** tab of the **Project Navigator** window in the Quartus II software. When you create a reconfigurable revision, the Quartus II software adds the required assignments to associate the reconfigurable revision with the base revision of the PR project. You can add the necessary files to each revision with the **Add/Remove Files** option in the **Project** option under the **Project** menu in the Quartus II software. With this step, you can associate the right implementation files for each revision of the PR project.

**Important:** You must use the **Revisions** tab of the **Project Navigator** window in the Quartus II software when creating revisions for partial reconfiguration. Revisions created using **Project > Revisions** cannot be reconfigured.

## Compiling Reconfigurable Revisions

Altera recommends that you use the largest persona of the PR region for the base compilation so that the Quartus II software can automatically budget sufficient routing.

Here are the typical steps involved in a PR design flow.

1. Compile the base revision with the largest persona for each PR region.
2. Create reconfigurable revisions for other personas of the PR regions by right-clicking in the **Revisions** tab in the Project Navigator.
3. Compile your reconfigurable revisions.
4. Analyze timing on each reconfigurable revision to make sure the design performs correctly to specifications.
5. Create aggregate revisions as needed.
6. Create programming files.

For more information on compiling a partial reconfiguration project, refer to *Performing Partial Reconfiguration* in Quartus II Help.

### Related Information

#### [Performing Partial Reconfiguration](#)

## Timing Closure for a Partial Reconfiguration Project

As with any other FPGA design project, simulate the functionality of various PR personas to make sure they perform to your system specifications. You must also make sure there are no timing violations in the implementation of any of the personas for every PR region in your design project.

In the Quartus II software, this process is manual, and you must run multiple timing analyses, on the base, reconfigurable, and aggregate revisions. The different timing requirements for each PR persona can be met by using different SDC constraints for each of the personas.

The interface between the partial and static partitions remains identical for each reconfigurable and aggregate revision in the PR flow. If all the interface signals between the static and the PR regions are registered, and there are no timing violations within the static region as well as within the PR regions, the reconfigurable and aggregate revisions should not have any timing violations.

However, you should perform timing analysis on the reconfigurable and aggregate revisions, in case you have any unregistered signals on the interface between partial reconfiguration and static regions.

## Bitstream Compression and Encryption for PR Designs

You can choose to independently compress and encrypt the base bitstream as well as the PR bitstream for your PR project using options available in the Quartus II software.

When you choose to compress the bitstreams, you can compress the base and PR programming bitstreams independently, based on your design requirements. However, if you want to encrypt only the base image, you can choose whether or not to encrypt the PR images.

- When you want to encrypt the bitstreams, you can encrypt the PR images only when the base image is encrypted.
- The Encryption Key Programming (.ekp) file generated when encrypting the base image must be used for encrypting PR bitstream.
- When you compress the bitstream, you must present each `PR_DATA[15:0]` word for exactly four clock cycles.

**Table 4-5: Partial Reconfiguration Clock Requirements for Bitstream Compression**

Timing Parameters	Value (clock cycles)
PR_READY to first data	4 (exact)
PR_ERROR to last clock	80 (minimum)
PR_DONE to last clock	80 (minimum)
DONE_to_REQ_low	8 (maximum)

**Related Information**

- [Enable Partial Reconfiguration Bitstream Decompression when Configuring Base Design SOF file in JTAG mode](#) on page 4-35
- [Enable Bitstream Decryption Option](#) on page 4-36
- [Generate PR Programming Files with the Convert Programming Files Dialog Box](#) on page 4-33

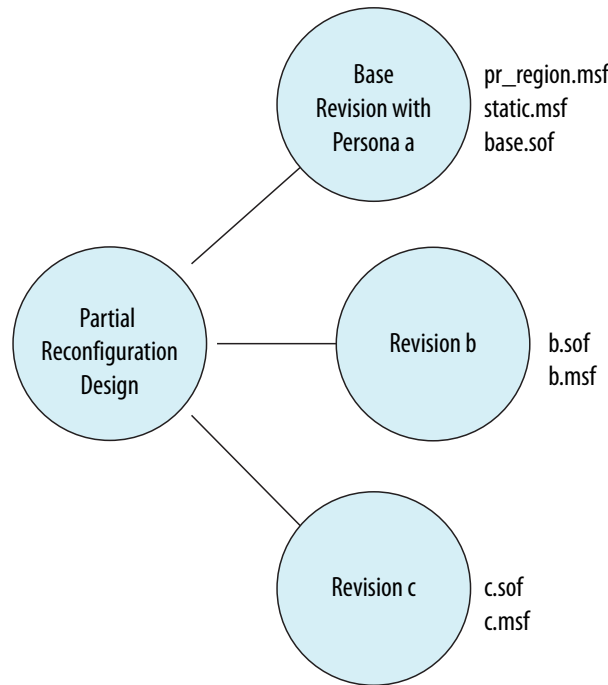
## Programming Files for a Partial Reconfiguration Project

You must generate PR bitstream(s) based on the designs and send them to the control block for partial reconfiguration.

Compile the PR project, including the base revision and at least one reconfigurable revision before generating the PR bitstreams. The Quartus II Programmer generates PR bitstreams. This generated bitstream can be sent to the PR ports on the control block for partial reconfiguration.

**Figure 4-15: PR Project with Three Revisions**

Consider a partial reconfiguration design that has three revisions and one PR region, a base revision with persona a, one PR revision with persona b, and a second PR revision with persona c.




When these individual revisions are compiled in the Quartus II software, the assembler produces Masked SRAM Object Files (**.msf**) and the SRAM Object Files (**.sof**) for each revision. The **.sof** files are created as before (for non-PR designs). Additionally, **.msf** files are created specifically for partial reconfiguration, one for each revision. The **pr\_region.msf** file is the one of interest for generating the PR bitstream. It contains the mask bits for the PR region. Similarly, the **static.msf** file has the mask bits for the static region. The **.sof** files have the information on how to configure the static region as well as the corresponding PR region. The **pr\_region.msf** file is used to mask out the static region so that the bitstream can be computed for the PR region. The default file name of the pr region **.msf** corresponds to the LogicLock region name, unless the name is not alphanumeric. In the case of a non-alphanumeric region name, the **.msf** file is named after the location of the lower left most coordinate of the region.

**Note:** Altera recommends naming all LogicLock regions to enhance documenting your design.

**Figure 4-16: Generation of Partial-Masked SRAM Object Files (.pmsf)**

You can convert files in the Convert Programming Files window or run the `quartus_cpf -p` command to process the `pr_region.msf` and `.sof` files to generate the Partial-Masked SRAM Object File (`.pmsf`).




The `.msf` file helps determine the PR region from each of the `.sof` files during the PR bitstream computation.

Once all the `.pmsf` files are created, process the PR bitstreams by running the `quartus_cpf -o` command to produce the raw binary `.rbf` files for reconfiguration.

If one wishes to partially reconfigure the PR region with persona a, use the `a.rbf` bitstream file, and so on for the other personas.

**Figure 4-17: Generating PR Bitstreams**

This figure shows how three bitstreams can be created to partially reconfigure the region with persona a, persona b, or persona c as desired.



In the Quartus II software, the Convert Programming Files window supports the generation of the required programming bitstreams. When using the `quartus_cpf` from the command line, the following options for generating the programming files are read from an option text file, for example, `option.txt`.

- If you want to use SCRUB mode, before generating the bitstreams create an option text file, with the following line:  
`use_scrub=on`
- If you have initialized M20K blocks in the PR region (ROM/Initialized RAM), then add the following line in the option text file, before generating the bitstreams:  
`write_block_memory_contents=on`
- If you want to compress the programming bitstream files, add the following line in the option text file. This option is available when converting base `.sof` to any supported programming file types, such as `.rbf`, `.pof` and JTAG Indirect Configuration File (`.jic`).

`bitstream_compression=on`



### Related Information

[Generate PR Programming Files with the Convert Programming Files Dialog Box](#) on page 4-33

## Generating Required Programming Files

1. Generate **.sof** and **.msf** files (part of a full compilation of the base and PR revisions).
2. Generate a Partial-Masked SRAM Object File (**.pmsf**) using the following commands:

```
quartus_cpf -p <pr_revision>.msf <pr_revision>.sof <new_filename>.pmsf
```

for example:

```
quartus_cpf -p x7y48.msف switchPRBS.sof x7y48_new.pmsf
```

3. Convert the **.pmsf** file for every PR region in your design to **.rbf** file format. The **.rbf** format is used to store the bitstream in an external flash memory. This command should be run in the same directory where the files are located:

```
quartus_cpf -o scrub.txt -c <pr_revision >.pmsf <pr_revision>.rbf
```

for example:

```
quartus_cpf -o scrub.txt -c x7y48_new.pmsf x7y48.rbf
```

When you do not have an option text file such as **scrub.txt**, the files generated would be for AND/OR mode of PR, rather than SCRUB mode.

## Generate PR Programming Files with the Convert Programming Files Dialog Box

In the Quartus II software, the flow to generate PR programming files is supported in the Convert Programming Files dialog box. You can specify how the Quartus II software processes file types such as **.msf**, **.pmsf**, and **.sof** to create **.rbf** and merged **.msf** and **.pmsf** files.

You can create

- A **.pmsf** output file, from **.msf** and **.sof** input files
- A **.rbf** output file from a **.pmsf** input file
- A merged **.msf** file from two or more **.msf** input files
- A merged **.pmsf** file from two or more **.pmsf** input files

Convert Programming Files dialog box also allows you to enable the option bit for bitstream decompression during partial reconfiguration, when converting the base **.sof** (full design **.sof**) to any supported file type.

For additional details, refer to the *Quartus II Programmer* chapter in the *Quartus II Handbook*.

### Related Information

[Quartus II Programmer](#)

## Generating a .pmsf File from a .msf and .sof Input File

Perform the following steps in the Quartus II software to generate the **.pmsf** file in the **Convert Programming Files** dialog box.

1. Open the **Convert Programming Files** dialog box.
2. Specify the programming file type as **Partial-Masked SRAM Object File (.pmsf)**.
3. Specify the output file name.
4. Select input files to convert (only a single **.msf** and **.sof** file are allowed). Click **Add**.
5. Click **Generate** to generate the **.pmsf** file.

## Generating a .rbf File from a .pmsf Input File

Perform the following steps in the Quartus II software to generate the partial reconfiguration **.rbf** file in the **Convert Programming Files** dialog box.

1. From the File menu, click **Convert Programming Files**.
2. Specify the programming file type as **Raw Binary File for Partial Reconfiguration (.rbf)**.
3. Specify the output file name.
4. Select input file to convert. Only a single **.pmsf** input file is allowed. Click **Add**.
5. Select the new **.pmsf** and click **Properties**.
6. Turn the **Compression**, **Enable SCRUB mode**, **Write memory contents**, and **Generate encrypted bitstream** options on or off depending on the requirements of your design. Click **Generate** to generate the **.rbf** file for partial reconfiguration.
  - **Compression:** Enables compression on the PR bitstream.
  - **Enable SCRUB mode:** Default is based on AND/OR mode. This option is valid only when your design does not contain vertically overlapped PR masks. The **.rbf** generation fails otherwise.
  - **Write memory contents:** Turn this on when you have a **.mif** that was used during compilation. Otherwise, turning this option on forces you to use double PR in AND/OR mode.
  - **Generate encrypted bitstream:** If this option is enabled, you must specify the Encrypted Key Programming (**.ekp**) file, which generated when converting a base **.sof** to an encrypted bitstream. The same **.ekp** must be used to encrypt the PR bitstream.

When you turn on **Compression**, you must present each `PR_DATA[15:0]` word for exactly four clock cycles.

Turn on the **Write memory contents** option only if you are using AND/OR mode and have M20K blocks in your PR design that need to be initialized. When you check this box, you must to perform double PR for regions with initialized M20K blocks.

### Related Information

- [Initializing M20K Blocks with a Double PR Cycle](#) on page 4-42
- [Initializing M20K Blocks with a Double PR Cycle](#) on page 4-42

## Create a Merged .msf File from Multiple .msf Files

You can merge two or more **.msf** files in the **Convert Programming Files** window.

1. Open the **Convert Programming Files** window.
2. Specify the programming file type as **Merged Mask Settings File (.msf)**.
3. Specify the output file name.
4. Select **MSF Data** in the **Input files to convert** window.
5. Click **Add File** to add input files. You must specify two or more files for merging.
6. Click **Generate** to generate the merged file.

To merge two or more **.msf** files from the command line, type:

```
quartus_cpf --merge_msf=<number of merged files> <msf_input_file_1>  
<msf_input_file_2> <msf_input_file_etc> <msf_output_file>
```

For example, to merge two **.msf** files, type:

```
quartus_cpf --merge_msf=<2> <msf_input_file_1> <msf_input_file_2>  
<msf_output_file>
```

## Generating a Merged .pmsf File from Multiple .pmsf Files

You can merge two or more **.pmsf** files in the **Convert Programming Files** window.

1. Open the **Convert Programming Files** window.
2. Specify the programming file type as **Merged Partial-Mask SRAM Object File (.pmsf)**.
3. Specify the output file name.
4. Select **PMSF Data** in the **Input files to convert** window.
5. Click **Add File** to add input files. You must specify two or more files for merging.
6. Click **Generate** to generate the merged file.

To merge two or more **.pmsf** files from the command line, type:

```
quartus_cpf --merge_pmsf=<number of merged files> <pmsf_input_file_1>  
<pmsf_input_file_2> <pmsf_input_file_etc> <pmsf_output_file>
```

For example, to merge two **.pmsf** files, type:

```
quartus_cpf --merge_pmsf=<2> <pmsf_input_file_1> <pmsf_input_file_2>  
<pmsf_output_file>
```

The merge operation checks for any bit conflict on the input files, and the operation fails with error message if a bit conflict is detected. In most cases, a successful file merge operation indicates input files do not have any bit conflict.

## Enable Partial Reconfiguration Bitstream Decompression when Configuring Base Design SOF file in JTAG mode

In the Quartus II software, the **Convert Programming Files** window provides the option in the **.sof** file properties to enable bitstream decompression during partial reconfiguration.

This option is available when converting base **.sof** to any supported programming file types, such as **.rbf**, **.pof**, and **.jic**.

In order to view this option, the base **.sof** must be targeted on Stratix V devices in the **.sof File Properties**. This option must be turned on if you turned on the **Compression** option during **.pmsf** to **.rbf** file generation.

### Enable Bitstream Decryption Option

The **Convert Programming Files** window provides the option in the **.sof** file properties to enable bitstream decryption during partial reconfiguration.

This option is available when converting base **.sof** to any supported programming file types, such as **.rbf**, **.pof**, and **.jic**.

The base **.sof** must have partial reconfiguration enabled and the base **.sof** generated from a design that has a PR Control Block instantiated, to view this option in the **.sof File Properties**. This option must be turned on if you want to turn on the Generate encrypted bitstream option during **.pmsf** to **.rbf** file generation.

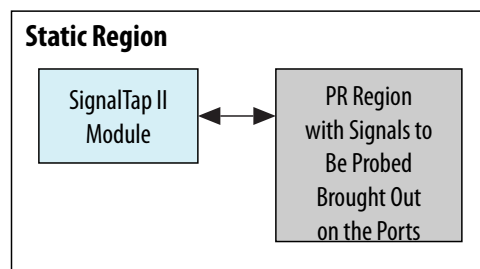
## On-Chip Debug for PR Designs

You cannot instantiate a SignalTap II block inside a PR region. If you must monitor signals within a PR region for debug purposes, bring those signals to the ports of the PR region.

The Quartus II software does not support the Incremental SignalTap feature for PR designs. After you instantiate the SignalTap II block inside the static region, you must recompile your design. When you recompile your design, the static region may have a modified implementation and you must also recompile your PR revisions. If you modify an existing SignalTap II instance you must also recompile your entire design; base revision and reconfigurable revisions.

**Figure 4-18: Using SignalTap II with a PR Design**

You can instantiate the SignalTap II block in the static region of the design and probe the signals you want to monitor.



You can use other on-chip debug features in the Quartus II software, such as the In-System Sources and Probes or SignalProbe, to debug a PR design. As in the case of SignalTap, In-System Sources and Probes can only be instantiated within the static region of a PR design. If you have to probe any signal inside the PR region, you must bring those signals to the ports of the PR region in order to monitor them within the static region of the design.

## Partial Reconfiguration Known Limitations

There are restrictions that derive from hardware limitations in specific Stratix V devices.

The restrictions in the following sections apply only if your design uses M20K blocks as RAMs or ROMs in your PR project.

### Memory Blocks Initialization Requirement for PR Designs

For a non-PR design, the power up value for the contents of a M20K RAM or a MLAB RAM are all set at zero. However, at the end of performing a partial reconfiguration, the contents of a M20K or MLAB memory block are unknown. You must intentionally initialize the contents of all the memory to zero, if required by the functionality of the design, and not rely upon the power on values.

### M20K RAM Blocks in PR Designs

When your PR design uses M20K RAM blocks in Stratix V devices, there are some restrictions which limit how you utilize the respective memory blocks as ROMs or as RAMs with initial content.

#### Related Information

[Implementing Memories with Initialized Content](#) on page 4-40

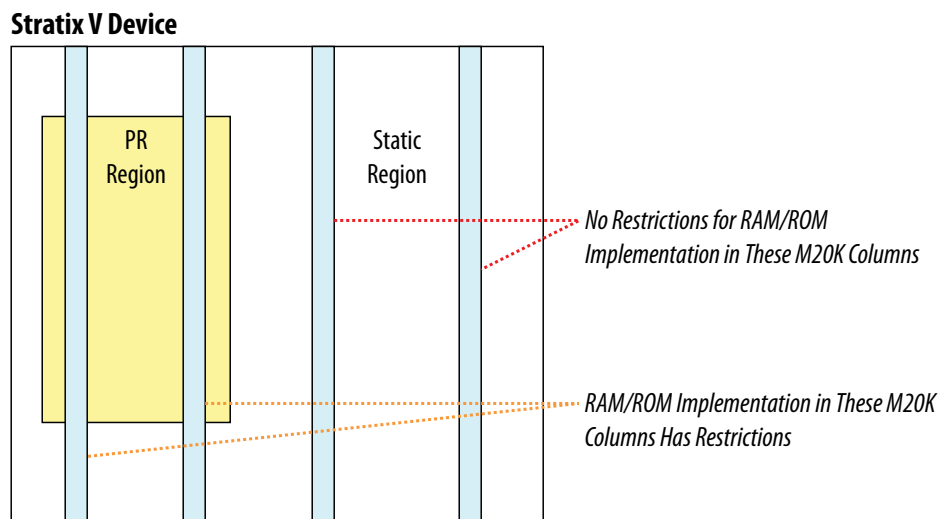
If your design requires initialized memory content either as a ROM or a RAM inside a PR region, you must follow these guidelines.

### Limitations When Using Stratix V Production Devices

These workarounds allow your design to use M20K blocks with PR.

**Figure 4-19: Limitations for Using M20Ks in PR Regions**

If you implement a M20K block in your PR region as a ROM or a RAM with initialized content, when the PR region is reconfigured, any data read from the memory blocks in static regions in columns that cross the PR region is incorrect.




If the functionality of the static region depends on any data read out from M20K RAMs in the static region, the design will malfunction.

Use one of the following workarounds, which are applicable to both AND/OR and SCRUB modes of partial reconfiguration:

- Do not use ROMs or RAMs with initialized content inside PR regions.
- If this is not possible for your design, you can program the memory content for M20K blocks with a `.mif` using the suggested workarounds.
- Make sure your PR region extends vertically all the way through the device, in such a way that the M20K column lies entirely inside a PR region.

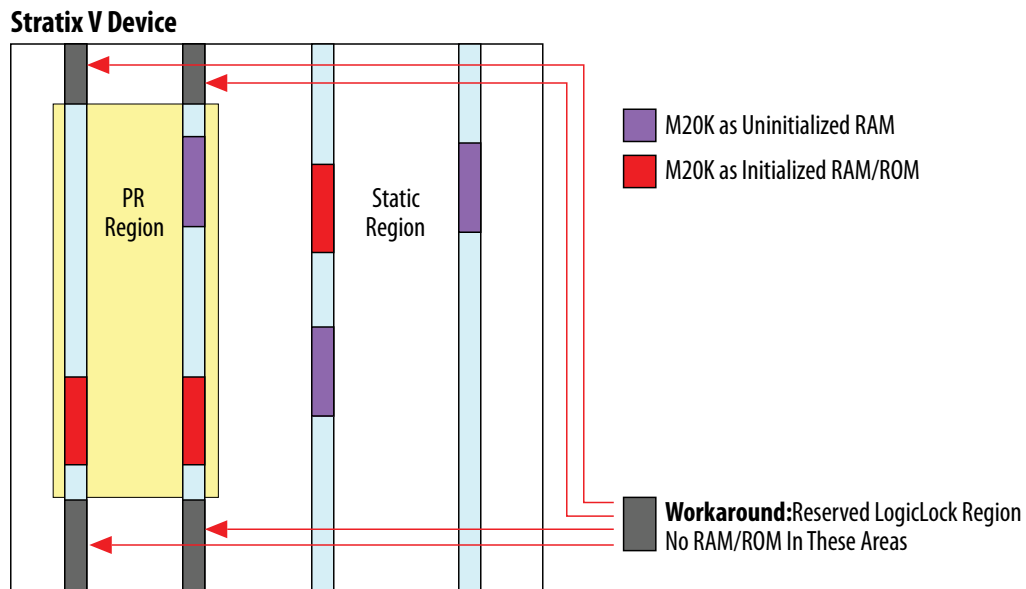
**Figure 4-20: Workaround for Using M20Ks in PR Regions**

This figure shows the LogicLock region extended as a rectangle reducing the area available for the static region. However, you can create non-rectangular LogicLock regions to allocate the resources required for the partition more optimally. If saving area is a concern, extend the LogicLock region to include M20K columns entirely.



**Figure 4-21: Alternative Workaround for Using M20Ks in PR Region**

Using Reserved LogicLock Regions, block all the M20K columns that are not inside a PR region, but that are in columns above or below a PR region. In this case, you may choose to under-utilize M20K resources, in order to gain ROM functionality within the PR region.



For more information including a list of the Stratix V production devices, refer to the *Errata Sheet for Stratix V Devices*.

**Related Information**

[Errata Sheet for Stratix V Devices](#)

**MLAB Blocks in PR designs**

Stratix V devices include dual-purpose blocks called MLABs, which can be used to implement RAMs or LABs for user logic.

This section describes the restrictions while using MLAB blocks (sometimes also referred to as LUT-RAM) in Stratix V devices for your PR designs.

If your design uses MLABS as LUT RAM, you must use all available MLAB bits within the region.

**Table 4-6: RAM Implementation Restrictions Summary**

The following table shows a summary of the LUT-RAM Restrictions.

PR Mode	Type of memory in PR region	Stratix V Production
SCRUB mode	LUT RAM (no initial content)	OK
	LUT ROM and LUT RAM with your initial content	OK

PR Mode	Type of memory in PR region	Stratix V Production
AND/OR mode	LUT RAM (no initial content)	<p><i>While design is running:</i> Write 1s to all locations before partial reconfiguration</p> <p><i>At compile time:</i> Explicitly initialize all memory locations in each new persona to 1 via initialization file (.mif).</p>
	LUT ROM and LUT RAM with your initial content	No

If your design does not use any MLAB blocks as RAMs, the following discussion does not apply. The restrictions listed below are the result of hardware limitations in specific devices.

### Limitations with Stratix V Production Devices

*When using SCRUB mode:*

- LUT-RAMs without initialized content, LUT-RAMs with initialized content, and LUT-ROMs can be implemented in MLABs within PR regions without any restriction.

*When using AND/OR mode:*

- LUT-RAMs with initialized content or LUT-ROMs cannot be implemented in a PR region.
- LUT-RAMs without initialized content in MLABs inside PR regions are supported with the following restrictions.
- MLAB blocks contain 640 bits of memory. The LUT RAMs in PR regions in your design must occupy all MLAB bits, you should not use partial MLABs.
- You must include control logic in your design with which you can write to all MLAB locations used inside PR region.
- Using this control logic, write '1' at each MLAB RAM bit location in the PR region before starting the PR process. This is to work around a false EDCRC error during partial reconfiguration.
- You must also specify a .mif that sets all MLAB RAM bits to '1' immediately after PR is complete.
- ROMs cannot be implemented in MLABs (LUT-ROMs).
- There are no restrictions to using MLABs in the static region of your PR design.

For more information, refer to the following documents in the *Stratix V Handbook*:

#### Related Information

[Errata Sheet for Stratix V Devices](#)

## Implementing Memories with Initialized Content

If your Stratix V PR design implements ROMs, RAMs with initialization, or ROMs within the PR regions, using either M20K blocks or LUT-RAMs, then you must follow the following design guidelines to determine what is applicable in your case.



**Table 4-7: Implementing Memory with Initialized Content in PR Designs**


Mode		Production Devices	
		AND/OR	SCRUB
LUT-RAM without initialization	<b>Suggested Method</b>	<p><i>While design is running:</i> Write '1' to all locations before partial reconfiguration.</p> <p><i>At compile time:</i> Explicitly initialize all memory locations in each new persona to '1' via initialization file (.mif)</p> <p>Make sure no spurious write on PR entry <sup>(4)</sup></p>	No special method required
	<b>Without Suggested Method</b>	CRC Error	No special method required
LUT-RAM with initialization	<b>Suggested Method</b>	Not supported	Make sure no spurious write on PR exit <sup>(4)</sup>
	<b>Without Suggested Method</b>		Incorrect results
M20K without initialization	<b>Suggested Method</b>	No special method required	
	<b>Without Suggested Method</b>	No special method required	
M20K with initialization	<b>Suggested Method</b>	<p>Use double PR cycle <sup>(5)</sup></p> <p>Make sure no spurious write on PR exit <sup>(4)</sup></p>	No special method required
	<b>Without Suggested Method</b>	Incorrect results	No special method required

<sup>(4)</sup> Use the circuit shown in the **M20K/LUTRAM** figure to create clock enable logic to safely exit partial reconfiguration without spurious writes.

<sup>(5)</sup> Double partial reconfiguration is described in *Initializing M20K Blocks with a Double PR Cycle*

**Figure 4-22: M20K/LUTRAM**

To avoid spurious writes during PR entry and exit, implement the following clock enable circuit in the same PR region as the RAM.



The circuit depends on an active-high clear signal from the static region. Before entering PR, freeze this signal in the same manner as all PR inputs. Your host control logic should de-assert the clear signal as the final step in the PR process.

#### Related Information


[Initializing M20K Blocks with a Double PR Cycle](#) on page 4-42

## Initializing M20K Blocks with a Double PR Cycle

When a PR region in your PR design contains an initialized M20K block and is reconfigured via AND/OR mode, your host logic must complete a double PR cycle, instead of a single PR cycle.

**Figure 4-23: Next PR Request Assertion During Double PR Cycle**

This figure displays the second phase of a double PR cycle, where the host logic must issue another PR\_REQUEST signal after exactly seven clock cycles after the PR\_DONE signal is asserted.



If the PR encryption feature (without compression) is enabled, the host logic must issue another PR\_REQUEST signal exactly six clock cycles after PR\_DONE is asserted.

If the PR compression feature is enabled (with or without encryption), the host logic must issue another PR\_REQUEST signal exactly two clock cycles after PR\_DONE is asserted. The FPGA responds with PR\_READY signal to the second PR\_REQUEST signal assertion.

The PR host must continue sending PR\_DATA signal exactly four clock cycles after the PR\_READY signal, just as in the first PR cycle. The data on PR\_DATA pins can be don't care between the first PR\_DONE signal and until four clock cycles after the PR\_READY signal is asserted for the second PR cycle.

The host must continue sending a PR\_DATA signal for the second PR cycle, until it receives the PR\_DONE signal for the second request, similar to the first PR cycle. After the PR\_DONE signal is asserted for the second time, the host should de-assert the PR\_REQUEST signal and continue with other operations needed for region bring up, such as issuing a reset to bring the region to a known state.

### Double PR with Compressed Programming Bitstream

You can use bitstream compression along with PR designs that also require memory initialization for M20K blocks.

For a compressed bitstream requiring a double PR cycle, the PR host must stop sending the PR\_DATA signal in the bitstream as soon as the first PR\_DONE is asserted. The PR host must resume sending the PR\_DATA signal immediately after the second PR\_READY signal is asserted.

## Document Revision History

Table 4-8: Document Revision History

Date	Version	Changes
2015.12.15	14.1.0	Minor revisions to some topics to resolve design refinements: <ul style="list-style-type: none"> <li>• Implementing Memories with Initialized Content</li> <li>• Instantiating the PR Control Block and CRC Block in Verilog HDL</li> <li>• Partial Reconfiguration Pins</li> </ul>
June 2014	14.0.0	Minor updates to "Programming File Sizes for a Partial Reconfiguration Project" and code samples in "Freeze Logic for PR Regions" sections.
November 2013	13.1.0	Added support for merging multiple .msf and .pmsf files. Added support for PR Megafunction. Updated for revisions on timing requirements.
May 2013	13.0.0	Added support for encrypted bitstreams. Updated support for double PR.
November 2012	12.1.0	Initial release.

### Related Information

#### [Quartus II Handbook Archive](#)

For previous versions of the *Quartus II Handbook*.

2014.12.15

QIISV1



Subscribe



Send Feedback

Qsys is a system integration tool included as part of the Quartus II software. Qsys captures system-level hardware designs at a high level of abstraction and simplifies the task of defining and integrating customized IP components. These components include verification IP cores, and other design modules. Qsys facilitates design reuse by packaging and integrating your custom IP components with Altera Altera and third-party IP components. Qsys automatically creates interconnect logic from the high-level connectivity that you specify, thereby eliminating the error-prone and time-consuming task of writing HDL to specify system-level connections.

Qsys is more powerful if you design your custom IP components using standard interfaces. By using standard interfaces, your custom IP components inter-operate with the Altera IP components in the IP Catalog. In addition, you can take advantage of bus functional models (BFMs), monitors, and other verification IP to verify your system.

Qsys supports Avalon<sup>®</sup>, AMBA<sup>®</sup> AXI3<sup>™</sup> (version 1.0), AMBA AXI4<sup>™</sup> (version 2.0), AMBA AXI4-Lite<sup>™</sup> (version 2.0), AMBA AXI4-Stream (version 1.0), and AMBA APB<sup>™</sup>3 (version 1.0) interface specifications.

Qsys provides the following advantages:

- Simplifies the process of customizing and integrating IP components into systems
- Generates an IP core variation for use in your Quartus II software projects
- Supports up to 64-bit addressing
- Supports modular system design
- Supports visualization of systems
- Supports optimization of interconnect and pipelining within the system
- Supports auto-adaptation of different data widths and burst characteristics
- Supports inter-operation between standard protocols, such as Avalon and AXI
- Fully integrated with the Quartus II software

**Note:** For information on how to define and generate single IP cores for use in your Quartus II software projects, refer to *Introduction to Altera IP Cores* and *Managing Quartus II Projects*.

#### Related Information

- [Introduction to Altera IP Cores](#)
- [Managing Quartus II Projects](#) on page 1-1
- [Avalon Interface Specifications](#)
- [AMBA Protocol Specifications](#)

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



## IP Component Interface Support in Qsys

IP components can have any number of interfaces in any combination. Each interface represents a set of signals that you can connect within a Qsys system, or export outside of a Qsys system.

Qsys IP components can include the following interface types:

**Table 5-1: IP Component Interface Types**

Interface Type	Description
Memory-Mapped	Connects memory-referencing master devices with slave memory devices. Master devices may be processors and DMAs, while slave memory devices may be RAMs, ROMs, and control registers. Data transfers between master and slave may be unidirectional (read only or write only), or bi-directional (read or write).
Streaming	Connects Avalon Streaming (Avalon-ST) sources and sinks that stream unidirectional data, as well as high-bandwidth, low-latency IP components. Streaming creates datapaths for unidirectional traffic, including multichannel streams, packets, and DSP data. The Avalon-ST interconnect is flexible and can implement on-chip interfaces for industry standard telecommunications and data communications cores, such as Ethernet, Interlaken, and video. You can define bus widths, packets, and error conditions.
Interrupts	Connects interrupt senders to interrupt receivers. Qsys supports individual, single-bit interrupt requests (IRQs). In the event that multiple senders assert their IRQs simultaneously, the receiver logic (typically under software control) determines which IRQ has highest priority, then responds appropriately
Clocks	Connects clock output interfaces with clock input interfaces. Clock outputs can fan-out without the use of a bridge. A bridge is required only when a clock from an external (exported) source connects internally to more than one source.
Resets	Connects reset sources with reset input interfaces. If your system requires a particular positive-edge or negative-edge synchronized reset, Qsys inserts a reset controller to create the appropriate reset signal. If you design a system with multiple reset inputs, the reset controller ORs all reset inputs and generates a single reset output.
Conduits	Connects point-to-point conduit interfaces, or represent signals that are exported from the Qsys system. Qsys uses conduits for component I/O signals that are not part of any supported standard interface. You can connect two conduits directly within a Qsys system as a point-to-point connection, or conduit interfaces can be exported and brought to the top-level of the system as top-level system I/O. You can use conduits to connect to external devices, for example external DDR SDRAM memory, and to FPGA logic defined outside of the Qsys system.

## Create a Qsys System

Click **Tools** > **Qsys** in the Quartus II software to open Qsys. A **.qsys** or **.qip** file represents your Qsys system in your Quartus II software project.

### Related Information

- [Creating Qsys Components](#) on page 6-1
- [Component Interface Tcl Reference](#) on page 9-1

## Start a New Project or Open a Recent Project in Qsys

1. To start a new Qsys project, save the default system that appears when you open Qsys (**File** > **Save**), or click **File** > **New System**, and then save your new project.  
Qsys saves the new project in the Quartus II project directory. To alternatively save your Qsys project in a different directory, click **File** > **Save As**.
2. To open a recent Qsys project, click **File** > **Open** to browse for the project, or locate a recent project with the **File** > **Recent Projects** command.
3. To revert the project currently open in Qsys to the saved version, click the first item in the **Recent Projects** list.

**Note:** You can edit the directory path information in the **recent\_projects.ini** file to reflect a new location for items that appear in the Recent Projects list.

## Specify the Target Device

In Qsys, the **Device Family** tab allows you to select the device family and device for your Qsys system. IP components, parameters, and output options that appear with your Qsys system vary according to the device type. Qsys saves device settings in the **.qsys** file.

When you generate your Qsys system, the generated output is for the Qsys-selected device, which may be different than your Quartus II project device settings.

The Quartus II software always uses the device specified in the Quartus II project settings, even if you have already generated your Qsys system with the Qsys-selected device.

**Note:** Qsys generates a warning message if the Qsys device family and device do not match the Quartus II project settings. Also, when you open Qsys from within the Quartus II software, the Quartus II project device settings apply.

## Add IP Components to your Qsys System

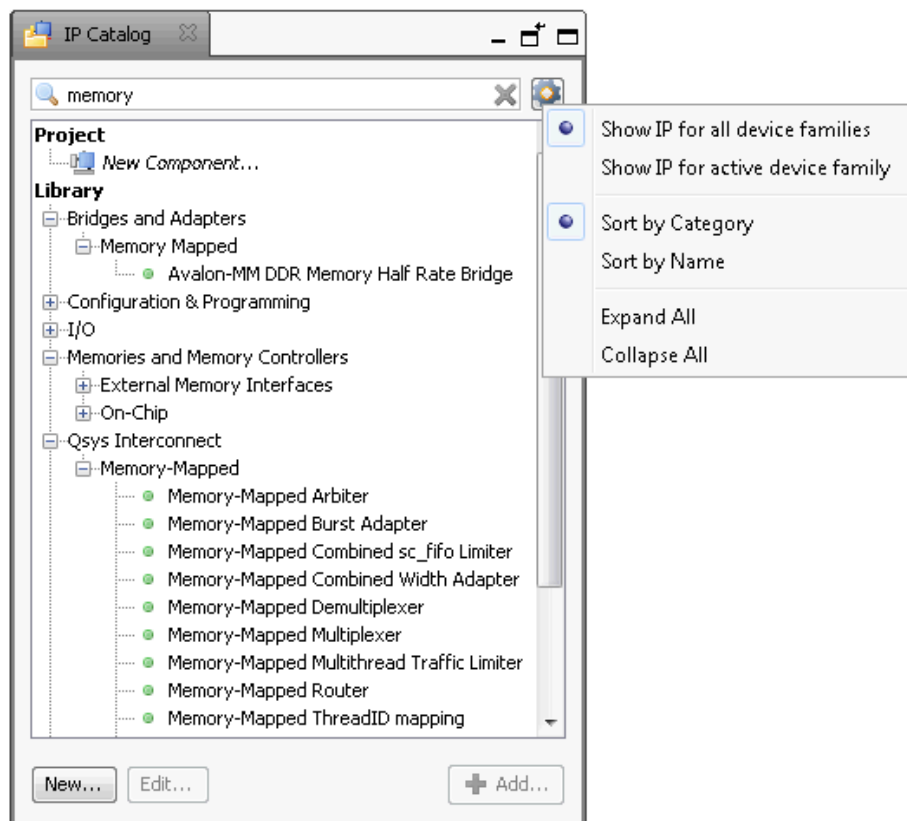
In Qsys, the IP Catalog displays IP components available for your target device. Double-click any component to launch the parameter editor. The parameter editor allows you to create a custom IP variation of the selected component. A Qsys system can contain a single instance of an IP component, or multiple, individually parameterized variations of the same IP component.

Some components have preset settings, which allow you to instantly apply preset parameter values appropriate for a specific application. When you complete your customization and click **Finish**, the component displays in the **System Contents** tab.

Right-click any IP component name in IP Catalog to display details about device support, installation location, version, and links to documentation. The IP Catalog maintains multiple versions of a component.

To locate a specific type of component, type some or all of the component's name in the IP Catalog search box. For example, you can type **memory** to locate memory-mapped IP components, or **axi** to locate AXI IP. You can also filter the IP Catalog display with options on the right-click menu.

Figure 5-1: IP Catalog



## Connect IP Components in Your Qsys System

The **System Contents** tab is the primary interface that you use to connect and configure components.

You connect interfaces of compatible types and opposite directions. For example, you can connect a memory-mapped master interface to a slave interface, and an interrupt sender interface to an interrupt receiver interface. You can connect any interfaces exported from a Qsys system within a parent Qsys system.

Possible connections between interfaces appear as gray lines and open circles. To make a connection, click the open circle at the intersection of the interfaces. When you make a connection, Qsys draws the connection line in black and fills the connection circle. Clicking a filled-in circle removes a connection.

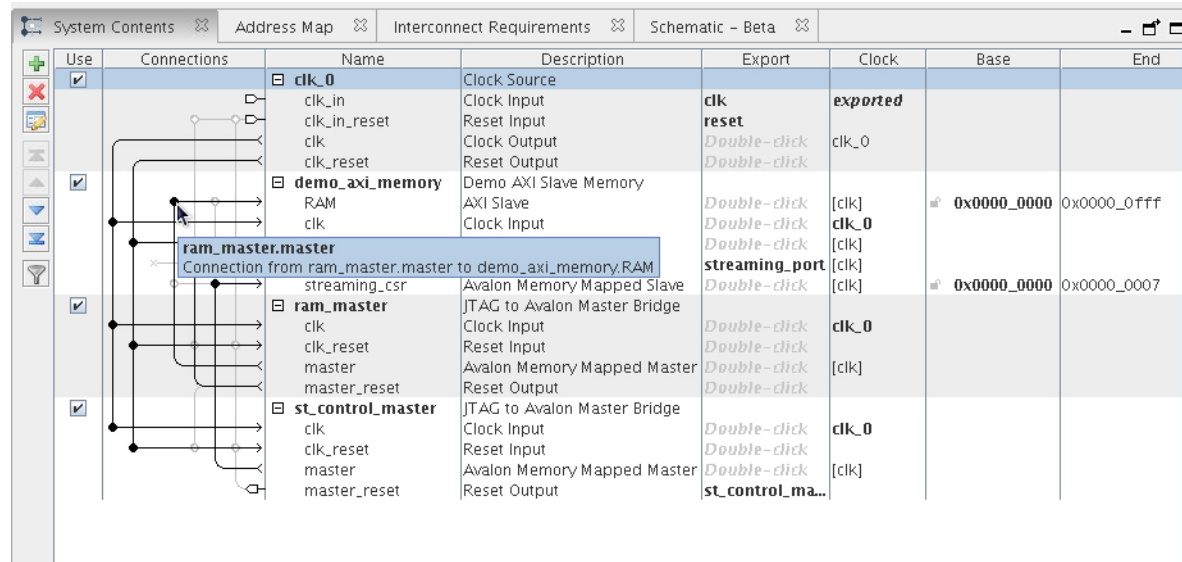
Qsys interconnect connects interface signals during system generation.



The **Connections** tab (**View > Connections**) shows a list of current and possible connections for selected instances or interfaces in the **Hierarchy** or **System Contents** tabs. You can add and remove connections by clicking the check box for each connection. Reporting columns provide information about each connection. For example, **Clock Crossing**, **Data Width**, and **Burst** columns provide information after system generation when the Qsys interconnect adds adapters that could possibly result in slower  $f_{MAX}$ , or larger area.

Figure 5-2: Connections Column in the System Contents Tab

When you finish adding connections, you can deselect **Allow Connection Editing** in the right-click menu. This option sets the **Connections** column to read-only and hides the possible connections.



**Related Information**

[Connecting Components](#)

**Create Connections Between Masters and Slaves**

The **Address Map** tab provides the address range that each memory-mapped master must use to connect to each slave in your system.

Qsys shows the slaves on the left, the masters across the top, and the address span of the connection in each cell. If there is no connection between a master and a slave, the table cell is empty.

You can design a system where two masters access a slave at different addresses. If you use this feature, Qsys labels the **Base** and **End** address columns in the **System Contents** tab as "mixed" rather than providing the address range.

Follow these steps to change or create a connection between master and slave IP components:

1. In Qsys, click or open the **Address Map** tab.
2. Locate the table cell that represents the connection between the master and slave component pair.
3. Either type in a base address, or update the current base address in the cell.

**Note:** The base address of a slave component must be a multiple of the address span of the component. This restriction is a requirement of the Qsys interconnect. The result is an efficient address decoding logic, which allows Qsys to achieve the best possible  $f_{MAX}$ .

## View Your Qsys System

Qsys allows you to change the display of your system to match your design development. Each tab on View menu allows you to view your design with a unique perspective. Multiple tabs open in your workspace allow you to focus on a selected element in your system under different perspectives.

Qsys GUI supports global selection and edit. When you make a selection or apply an edit in the **Hierarchy** tab, Qsys updates all other open tabs to reflect your action. For example, when you select `cpu_0` in the **Hierarchy** tab, Qsys updates the **Parameters** tab to show the parameters for `cpu_0`.

By default, when you open Qsys, the IP Catalog and **Hierarchy** tab display to the left of the main frame. The **System Contents**, **Address Map**, **Interconnect Requirements**, and **Device Family** tabs display in the main frame.

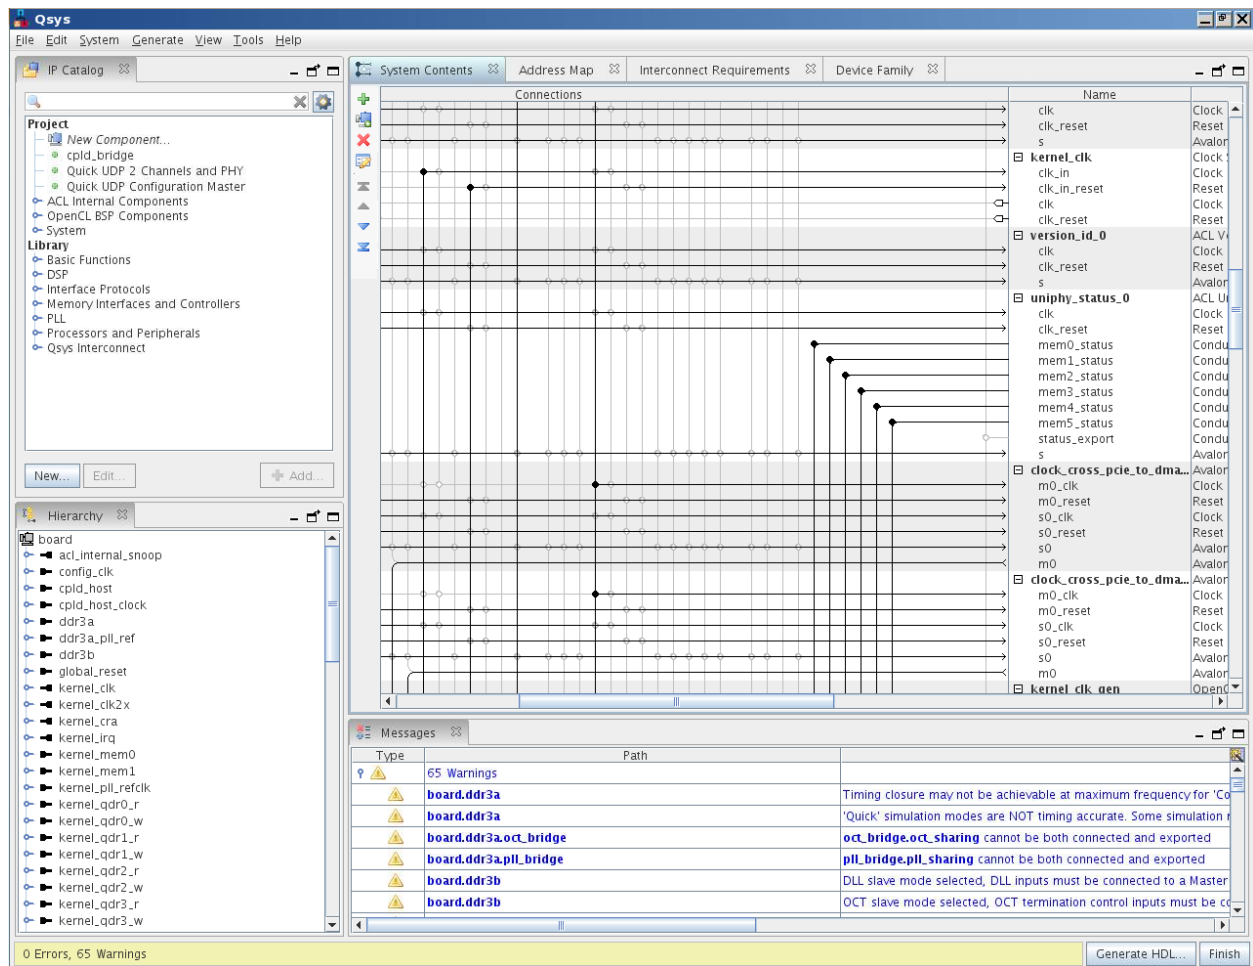
The **Messages** tab displays in the lower portion of Qsys. Double-clicking a message in the **Messages** tab changes focus to the associated element in the relevant tab to facilitate debugging. When the **Messages** tab is closed or not open in your workspace, error and warning message counts continue to display in the status bar of the Qsys window.

You can dock tabs in the main frame as a group, or individually by clicking the tab control in the upper-right corner of the main frame. You can arrange your workspace by dragging and dropping, and then grouping tabs in an order appropriate to your design development, or close or dock tabs that you are not using. Tool tips on the upper-right corner of the tab describe possible workspace arrangements, for example, restoring or disconnecting a tab to or from your workspace. When you save your system, Qsys also saves the current workspace configuration. When you re-open a saved system, Qsys restores the last saved workspace.

The **Reset to System Layout** command on the View menu restores the workspace to its default configuration for Qsys system design. The **Reset to IP Layout** command restores the workspace to its default configuration for defining and generating single IP cores.

**Note:** Qsys contains some tabs which are not documented and appear on the View menu as "Beta". The purpose in presenting these tabs is to allow designers to explore their usefulness in Qsys system development.

Figure 5-3: View Your Qsys System



## Manage Qsys Window Views with Layouts

Qsys Layout controls what tabs are open in your Qsys design window. When you create a Qsys window configuration that you want to keep, Qsys allows you to save that configuration as a custom layout. By default, Qsys contains a layout suitable for Qsys system design, as well as a layout that allows you to easily define and generate single IP cores for use in your Quartus II software projects.

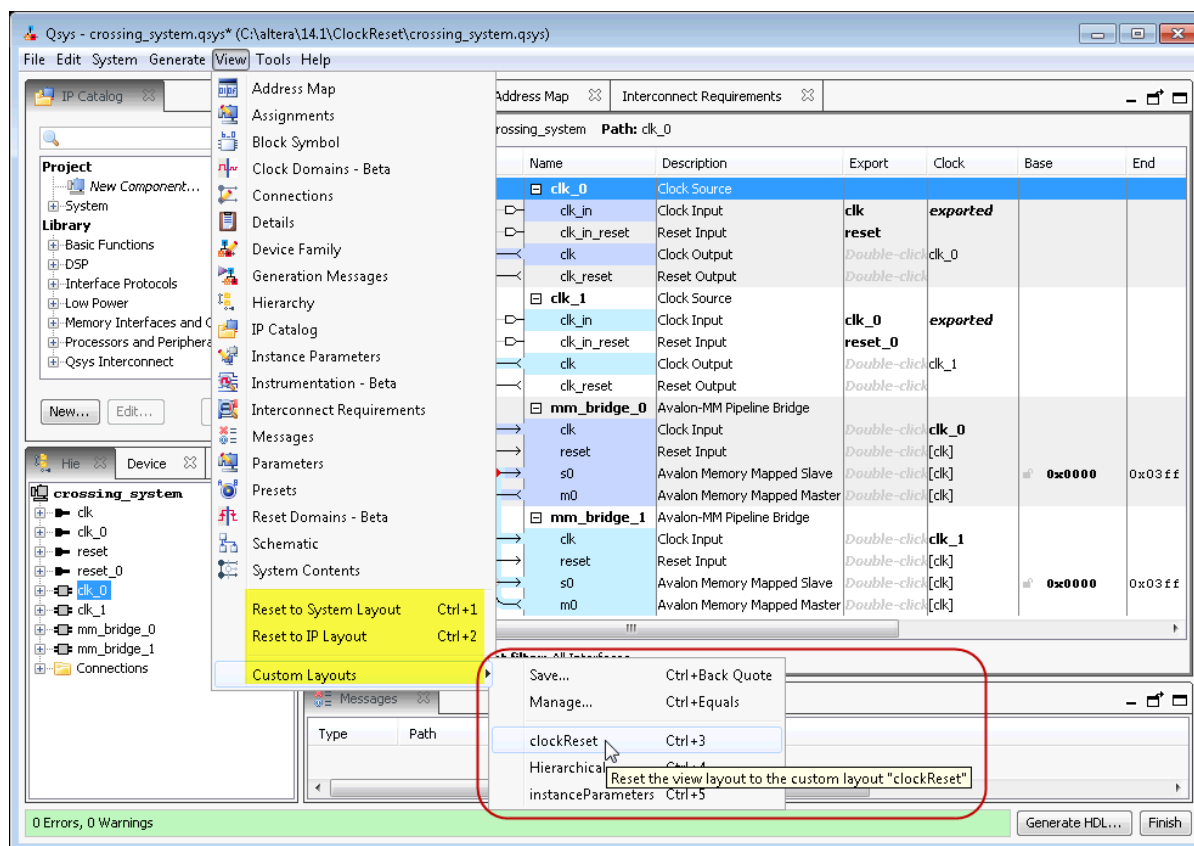
1. To configure your Qsys window with a layout suitable for Qsys system design, click **View > Reset to System Layout**.  
The **System Contents**, **Address Map**, **Interconnect Requirements**, and **Messages** tabs open in the main pane, and the **IP Catalog** and **Hierarchy** tabs along the left pane.
2. To configure your Qsys window with a layout suitable for single IP core design, click **View > Reset to IP Layout**.  
The **Parameters** and **Messages** tabs open in the main pane, and the **Details**, **Block Symbol** and **Presets** tabs along the right pane.
3. To save your current Qsys window configuration as a custom layout, click **View > Custom Layouts > Save**.

Qsys saves your custom layout in your project directory, and adds the layout to the custom layouts list, and the **layouts.ini** file. The **layouts.ini** file controls the order in which the layouts appear in the list.

- To reset your Qsys window configuration to a previously saved configuration, click **View > Custom Layouts**, and then select the custom layout in the list.

The Qsys windows opens with your previously saved Qsys window configuration.

**Figure 5-4: Save Your Qsys Window Views and Layouts**




- To manage your saved custom layouts, click **View > Custom Layouts**.

The **Manage Custom Layouts** dialog box opens and allows you to apply a variety of functions that facilitate custom layout management, including the ability to import or export a layout from or to a different directory.

**Figure 5-5: Manage Custom Layouts**

The shortcut, **Ctrl-3**, for example, allows you to quickly change your Qsys window view with a quick keystroke.

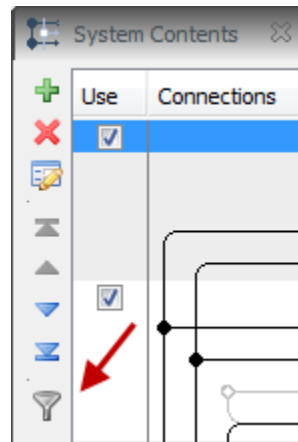


### Filter the Display of the System Contents tab

You can use the **Filters** dialog box in the to filter the display of your system by interface type, instance name, or by using custom tags.

For example, in the **System Contents** tab, you can show only instances that include memory-mapped interfaces, instances that are connected to a particular Nios II processor, or temporarily hide clock and reset interfaces to simplify the display.

Figure 5-6: Filter Icon in the System Contents Tab



### Related Information

#### [Filters Dialog Box](#)

### Display Details About a Component or Parameter

The **Details** tab provides information for a selected component or parameter. Qsys updates the information in the **Details** tab as you select different components.

As you click through the parameters for a component in the parameter editor, Qsys displays the description of the parameter in the **Details** tab. To return to the complete description for the component, click the header in the **Parameters** tab.

### Display a Graphical Representation of a Component

In the **Block Symbol** tab, Qsys displays a graphical representation of the element that you select in the **Hierarchy** or **System Contents** tabs. You can view the selected component's port interfaces and signals. The **Show signals** option allows you to turn on or off signal graphics.

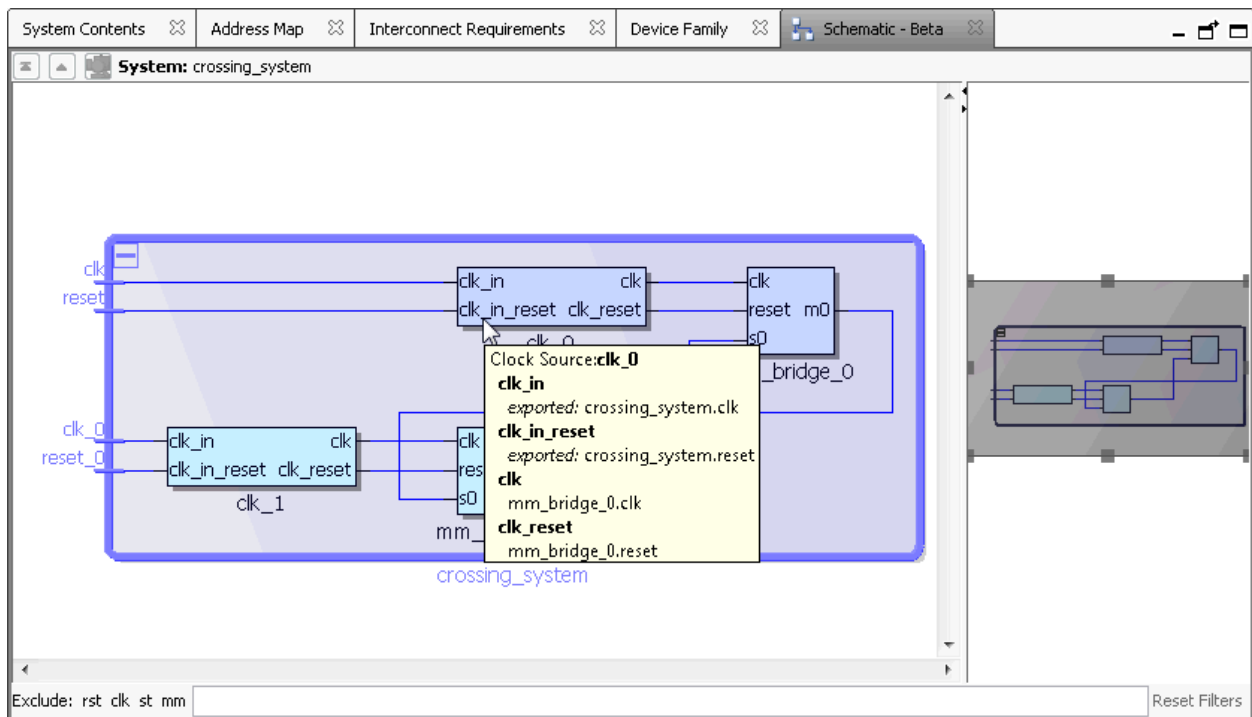
The **Block Symbol** tab appears by default in the parameter editor when you add a component to your system. When the **Block Symbol** tab is open in your workspace, it reflects changes that you make in other tabs.

### View a Schematic of Your Qsys System

The **Schematic** tab displays a schematic representation of your Qsys system. Tab controls allow you to zoom into a component or connection, or to obtain tooltip details for your selection. You can use the image handles in the right panel to resize the schematic image.

If your selection is a subsystem, use the Hierarchy tool to navigate to the parent subsystem, move up one level, or to drill into the currently open subsystem.

Figure 5-7: Qsys Schematic Tab




**Related Information**

[Edit a Qsys Subsystem](#) on page 5-23

**View Assignments and Connections in Your Qsys System**

On the **Assignments** tab (**View > Assignments**), you can view assignments for a module or element that you select in the **System Contents** tab. The **Connections** tab displays a lists of connections in your Qsys system. On the **Connections** tab (**View > Connections**), you can choose to connect or un-connect a module in your system, and then view the results in the **System Contents** tab.

Figure 5-8: Assignments and Connections tabs in Qsys



## Navigate Your Qsys System

The **Hierarchy** tab is a full system hierarchical navigator that expands the Qsys system contents to show all elements in your system, including, subsystems, components, interfaces, signals, and connections.

You can use the **Hierarchy** tab to browse, connect, parameterize IP, and drive changes in other open tabs. Expanding each interface in the **Hierarchy** tab allows you to view sub-components, associated elements, and signals for the interface. You can focus on a particular area of your system by coordinating selections in the **Hierarchy** tab with other open tabs in your workspace.

Navigating your system using the **Hierarchy** tab in conjunction with relevant tabs is useful during the debugging phase because you can contain and focus your debugging efforts to a single element in your system.

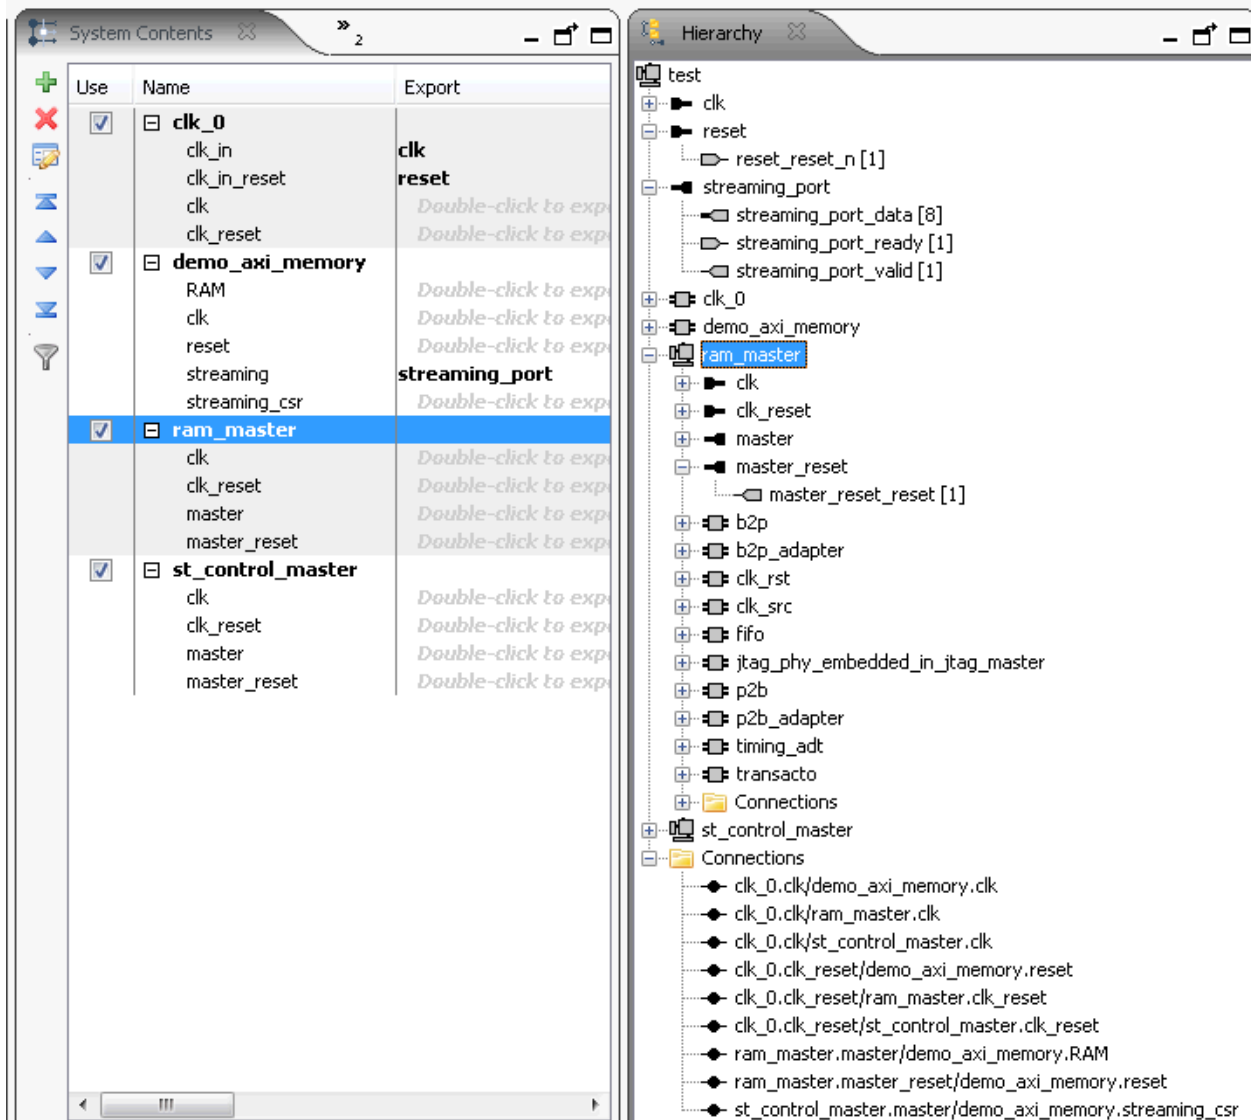
The **Hierarchy** tab provides the following information and functionality:

- Connections between signals.
- Names of signals in exported interfaces.
- Right-click menu to connect, edit, add, remove, or duplicate elements in the hierarchy.
- Internal connections of Qsys subsystems that are included as IP components. In contrast, the **System Contents** tab displays only the exported interfaces of Qsys subsystems.



**Figure 5-9: Expanding System Contents in the Hierarchy Tab**

The **Hierarchy** tab displays a unique icon for each element in the system. Context sensitivity between tabs facilitates design development and debugging. For example, when you select an element in the **Hierarchy** tab, Qsys selects the same element in other open tabs. This allows you to interact with your system in more detail. In the example below, the `ram_master` selection appears selected in both the **System Contents** and **Hierarchy** tabs.



**Related Information**

[Create and Manage Hierarchical Qsys Systems](#) on page 5-21

**Specify IP Component Parameters**

The **Parameters** tab allows you to specify parameters that define the IP component's functionality.

When you add a component to your system, or when you double-click a component in an open tab, the parameter editor appears. Many IP components in the IP Catalog have parameters that you can configure.

If you create your own IP components, use the Hardware Component Description File (`_hw.tcl`) to specify configurable parameters.

With the **Parameters** tab open, when you select an element in the **Hierarchy** tab, Qsys shows the same element in the **Parameters** tab. You can then make changes to the parameters that appear in the parameter editor, including changing the name for top-level instance that appears in the **System Contents** tab. Changes that you make in the **Parameters** tab affect your entire system and appear dynamically in other open tabs in your workspace.

In the parameter editor, the **Documentation** button provides information about a component's parameters, including the version.

At the top of the parameter editor, Qsys shows the hierarchical path for the component and its elements. This feature is useful when you navigate deep within your system with the **Hierarchy** tab. When you select an interface in the **Hierarchy** tab, the **Parameters** tab also allows you to review the timing for that interface. Qsys displays the read and write waveforms at the bottom of the **Parameters** tab.

### Configure Your IP Component with a Pre-Defined Set of Parameters

The **Presets** tab allows you to apply a pre-defined set of parameter values to your IP component. The **Presets** tab opens the preset editor and allows you to create, modify, and save custom component parameter values as a preset file. Not all IP components have preset files.

When you add a new component to your system, if there are preset values available for the component, the preset editor appears in the parameter editor window and lists preset files that you can apply the component. The name of each preset file describes a particular protocol and contains the required parameter values for that protocol.

You can search for text to filter the **Presets** list. For example, if you add the **DDR3 SDRAM Controller with UniPHY** component to your system, and type `1g micron 256` in the search box, the **Presets** list shows only those parameter values that apply to the 1g micron 256 filter request. Presets whose parameter values match the current parameter settings appear in bold.

If the available preset files do not meet the requirements of your design, you can create a new preset file. In the presets editor, click **New** to open the **New Preset** dialog box. Specify the new custom preset name, component version, description of the preset, and which parameters to include in the preset. You can also specify where you want to save the preset file. If the file location that you specify is not already in the IP search path, Qsys adds the location of the new preset file to the IP search path.

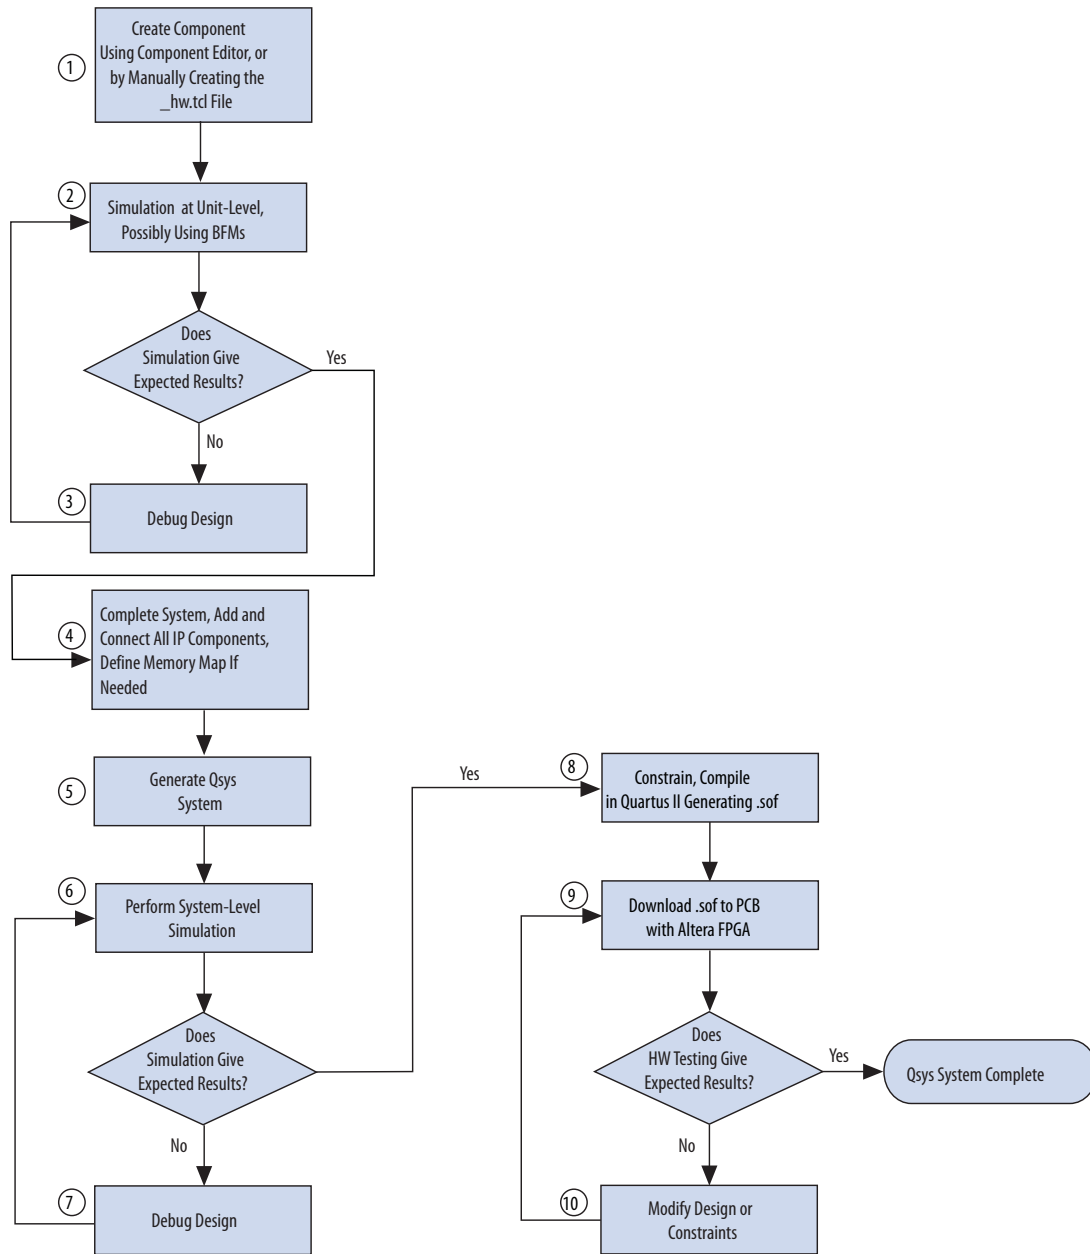
In the preset editor, click **Update** to update parameter values for a custom preset. The **Update Preset** dialog box displays the default values, which you can edit, and the current value, which are static. Click **Delete** to remove a custom preset from the **Presets** list.

When you access the presets editor by clicking **View > Presets**, you can apply the available presets to the currently selected component.

## Create a Custom IP Component (\_hw.tcl)

Figure 5-10: Qsys System Design Flow

The Qsys system design flow describes how to create a custom IP component using the Qsys Component Editor. You can optionally manually create a `_hw.tcl` file. The flow shows the simulation your custom IP, and at what point you can integrate it with other IP components to create a Qsys system and complete Quartus II project.



**Note:** For information on how to define and generate single IP cores for use in your Quartus II software projects, refer to *Introduction to Altera IP Cores*.

**Related Information**

[Creating Qsys Components](#) on page 6-1

[Introduction to Altera IP Cores](#)

[Managing Quartus II Projects](#) on page 1-1

## Add IP Components to the Qsys IP Catalog

The Qsys IP Catalog lists IP components that you can use in your Qsys systems. IP components can include Altera-provided, third-party, and custom IP components that you create on your own.

Because Qsys supports hierarchical design, previously created Qsys systems also appear in the IP Catalog. Additionally, instance parameters set on a Qsys system, can be parameterized within a parent Qsys system.

Altera and third-party developers provide ready-to-use IP components, which are installed with the Quartus II software and appear in the IP Catalog. The Qsys IP Catalog includes the following IP component types:

- Microprocessors, such as the Nios<sup>®</sup> II processor
- DSP IP components, such as the Reed Solomon Decoder II
- Interface protocols, such as the IP Compiler for PCI Express
- Memory controllers, such as the RLDRAM II Controller with UniPHY
- Avalon Streaming (Avalon-ST) IP components, such as the Avalon-ST Multiplexer
- Qsys Interconnect IP components
- Verification IP (VIP) Bus Functional Models (BFMs)

You can use the **IP Search Path** option (**Tools > Options**) to include custom and third-party IP components in the IP Catalog. The IP Catalog displays all IP cores that are found in the IP search path.

Qsys searches the directories listed in the IP search path for the following file types:

- **\*\_hw.tcl**—Defines a single IP component.
- **\*.ipx**—Each IP Index File (**.ipx**) indexes a collection of available IP components, or a reference to other directories to search. In general, **.ipx** files facilitate faster startup for Qsys because Qsys does not traverse directories below the **.ipx** file. Additionally, the **.ipx** file can contain information about the IP component, which eliminates the need for Qsys to read the **\_hw.tcl** file to understand the basic attributes of the IP core.

Qsys recursively searches directories specified in the IP search path until it finds a **\_hw.tcl** file. When Qsys finds a **\_hw.tcl** file in a directory, Qsys does not search subdirectories for additional **\_hw.tcl** files.

When you specify an IP search location in Qsys, a recursive descent is annotated by **\*\***. A single **\*** signifies a match against any file within the specified directory. However, even if a recursive descent is specified, if Qsys finds a **\_hw.tcl** or **.ipx** file, it does not search any subdirectories beyond that level.

**Note:** When you add a component to the search path, you must refresh your system by clicking **File > Refresh** to update the IP Catalog.

Table 5-2: IP Search Locations

Location	Description
<b>PROJECT_DIR/*</b>	Finds IP components and index files in the Quartus II project directory.
<b>PROJECT_DIR/ip/**/*</b>	Finds IP components and index files in any subdirectory of the <b>/ip</b> subdirectory of the Quartus project directory.

## Save IP Components to the Default Search Directory

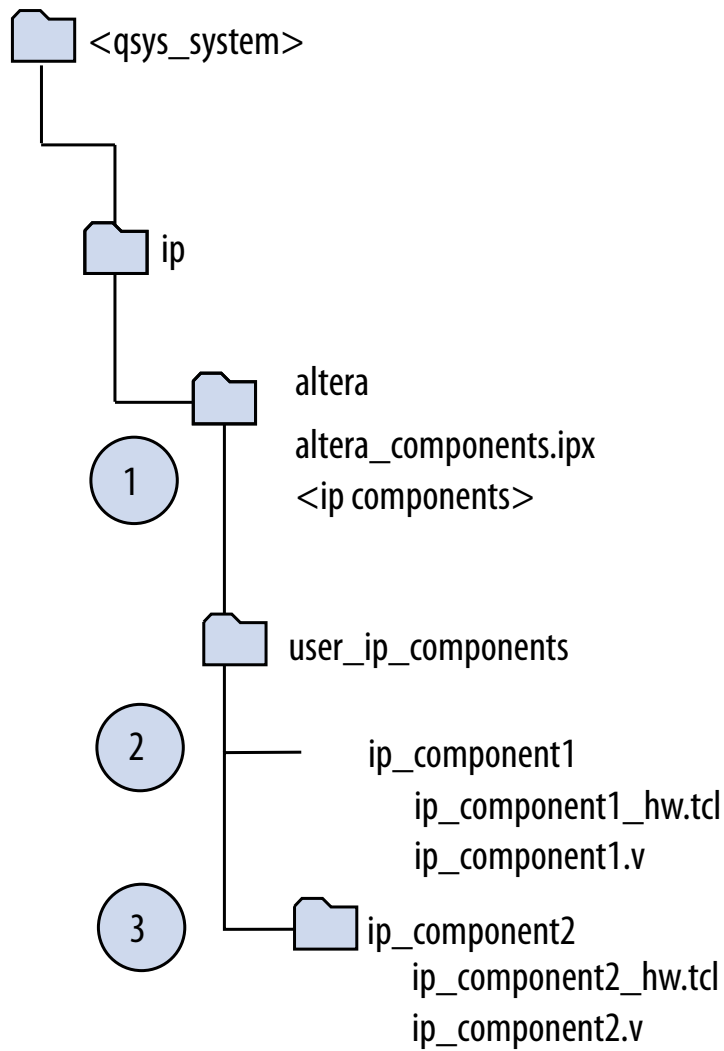
The simplest method to add an IP component to the Qsys IP Catalog is to copy the IP component files into an IP search path location.

When you create a component in the Qsys Component Editor, you can save the IP component in your Qsys project directory, or in the **/ip** subdirectory of your Qsys project directory. These methods are useful if you want to associate your IP component with a specific Qsys project.

You can also use the global IP search path in the Quartus II software (**Tools > Options > IP Catalog Search Locations**) to modify the IP search path across all Quartus II projects, Qsys systems, and Altera command line tools.

Figure 5-11: User IP Component Search

In this example, `user_ip_components` is a user-created directory for custom components.



Qsys performs the IP component search algorithm to locate `.ipx` and `_hw.tcl` files, as follows:

1. Qsys recursively searches the `<qsys_system>/ip` directory by default. The recursive search stops when Qsys discovers an `.ipx` file.

Because of this traversal, if changes to the `ip_component1_hw.tcl` file are made, but the `.ipx` file is not rebuilt using `ip-make-ipx`, those changes are not reflected in Qsys because the `.ipx` file contains old information about the `ip_component1` directory.

2. If the `.ipx` file is removed, Qsys discovers both the `ip_component1` and `ip_component2` directories, which contain `_hw.tcl` files.

**Note:** If you save your `_hw.tcl` file in the `<qsys_system>/ip` directory, Qsys finds your `_hw.tcl` file and will not search subdirectories adjacent to the `_hw.tcl` file. For more information about the IP search path, refer to *Introduction to Altera IP Cores*.

## Related Information

### [Introduction to Altera IP Cores](#)

## Set up the IP Index File (.ipx) to Search for IP Components

An IP Index File (.ipx) is an XML file that contains a search path that Qsys uses to search for IP components that are available for a Qsys system. You can use the `ip-make-ipx` command to create an .ipx file for a directory tree, which can reduce the startup time for Qsys.

You can specify a search path in a `user_components.ipx` file located anywhere in a customizable search path either in Qsys (**Tools > Options**) or the Quartus II software (**Tools > Options > IP Catalog Search Locations**). This method of discovering IP components allows you to add a location that is independent of the default search path. The `user_components.ipx` file directs Qsys to the location of each IP component or directory to search.

A `<path>` element in the .ipx file specifies a directory where multiple IP components may be found. A `<component>` entry specifies the path to a single component. A `<path>` element can use wildcards in its definition. An asterisk matches any file name. If you use an asterisk as a directory name, it matches any number of subdirectories.

### Example 5-1: Path Element in an .ipx File

```
<library>
  <path path="...<user directory>" />
  <path path="...<user directory>" />
  ...
  <component ... file="...<user directory>" />
  ...
</library>
```

A `<component>` element in an .ipx file contains several attributes to define a component. If you provide the required details for each component in an .ipx file, the startup time for Qsys is less than if Qsys must discover the files in a directory. The example below shows two `<component>` elements. Note that the paths for file names are specified relative to the .ipx file.

### Example 5-2: Component Element in an .ipx File

```
<library>
  <component
    name="A Qsys Component"
    displayName="Qsys FIR Filter Component"
    version="2.1"
    file="./components/qsys_filters/fir_hw.tcl"
  />
  <component
    name="rgb2cmyk_component"
    displayName="RGB2CMYK Converter(Color Conversion Category!)"
    version="0.9"
    file="./components/qsys_converters/color/rgb2cmyk_hw.tcl"
  />
</library>
```

**Note:** You can verify that IP components are available with the `ip-catalog` command.

## Integrate Third-Party IP Components into the Qsys IP Catalog

You can use IP components created by Altera partners in your Qsys systems. These IP components have interfaces that are supported by Qsys, such as Avalon-MM or AXI. Additionally, some include timing and placement constraints, software drivers, simulation models, and reference designs.

To locate supported third-party IP components on Altera's web page, navigate to the *Intellectual Property & Reference Designs* page, type `Qsys Certified` in the **Search** box, select **IP Core & Reference Designs**, and then press **Enter**.

Refer to Altera's *Intellectual Property & Reference Designs* page for more information.

### Related Information

[Intellectual Property & Reference Designs](#)

## Upgrade Outdated IP Components

When you open a Qsys system that contains outdated IP components, Qsys automatically attempts to upgrade the IP components if it cannot locate the requested version. IP components that Qsys successfully updates appear in the **Upgrade IP Cores** dialog box with a green checkmark. Most Qsys IP components support automatic upgrade.

You can include a path to older IP components in the IP Search Path, which Qsys uses even if updated versions are available. However, older versions of IP components may not work in newer version of Qsys.

**Note:** If your Qsys system includes an IP component(s) outside of the project directory, the directory of the `.qsys` file, or other any other directory location, you must add these dependency paths to the Qsys IP Search Path (**Tools > Options**).

1. With your Qsys system open, click **System > Upgrade IP Cores**.  
Only IP Components that are associated with the open Qsys system, and that do not support automatic upgrade appear in **Upgrade IP Cores** dialog box.
2. In the **Upgrade IP Cores** dialog box, click one or multiple IP components, and then click **Upgrade**.  
A green checkmark appears for the IP components that Qsys successfully upgrades.
3. Generate your Qsys system.

**Note:** Qsys supports command-line upgrade for IP components with the following command:

```
qsys-generate --upgrade-ip-cores <qsys_file>
```

The `<qsys_file>` variable accepts a path to the `.qsys` file so that you are not constrained to running this command in the same directory as the `.qsys` file. Qsys reports the start and finish of the command-line upgrade, but does not name the particular IP component(s) upgraded.

For device migration information, refer to *Introduction to Altera IP Cores*.

### Related Information

[Add IP Components to the Qsys IP Catalog](#) on page 5-16

[Introduction to Altera IP Cores](#)



## Create and Manage Hierarchical Qsys Systems

Qsys supports hierarchical system design. You can add any Qsys system as a subsystem in another Qsys system. Qsys hierarchical system design allows you to create, explore and edit hierarchies dynamically within a single instance of the Qsys editor. Qsys generates the complete hierarchy during the top-level system's generation.

**Note:** You can explore parameterizable Qsys systems and `_hw.tcl` files, but you cannot edit their elements.

Your Qsys systems appear in the IP Catalog under the System category under Project. You can reuse systems across multiple designs. In a team-based hierarchical design flow, you can divide large designs into subsystems and have team members develop subsystems simultaneously.

### Related Information

[Navigate Your Qsys System](#) on page 5-12


## Add a Subsystem to Your Qsys Design

You can create a child subsystem or nest subsystems at any level in the hierarchy. Qsys adds a subsystem to the system you are currently editing. This can be the top-level system, or a subsystem.

To create or nest subsystems in your Qsys design, use the following methods within the **System Contents** tab:

- Right-click command: **Add a new subsystem to the current system.**
- Left panel icon.
- **CTRL+SHIFT+N.**

Figure 5-12: Add a Subsystem to Your Qsys Design



## Drill into a Qsys Subsystem to Explore its Contents

The ability to drill into a system provides visibility into its elements and connections. When you drill into an instance, you open the system it instantiates for editing.

You can drill into a subsystem with the following commands:

- Double-click a system in the **Hierarchy** tab.
- Right-click a system in the **Hierarchy**, **System Contents**, or **Schematic** tabs, and then select **Drill into subsystem**.
- CTRL+SHIFT+D in the **System Contents** tab.

**Note:** You can only drill into **.qsys** files, not parameterizable Qsys systems or **\_hw.tcl** files.

The **Hierarchy** tab is rooted at the top-level and drives global selection. You can manage a hierarchical Qsys system that you build across multiple Qsys files, and view and edit their interconnected paths and address maps simultaneously. As an example, you can select a path to a subsystem in the **Hierarchy** tab, and then drill deeper into the subsystem in the **System Contents** or **Schematic** tabs. You could also select a subsystem in the **System Contents** tab, and then drill into the selected subsystem in the **Hierarchy** tab.

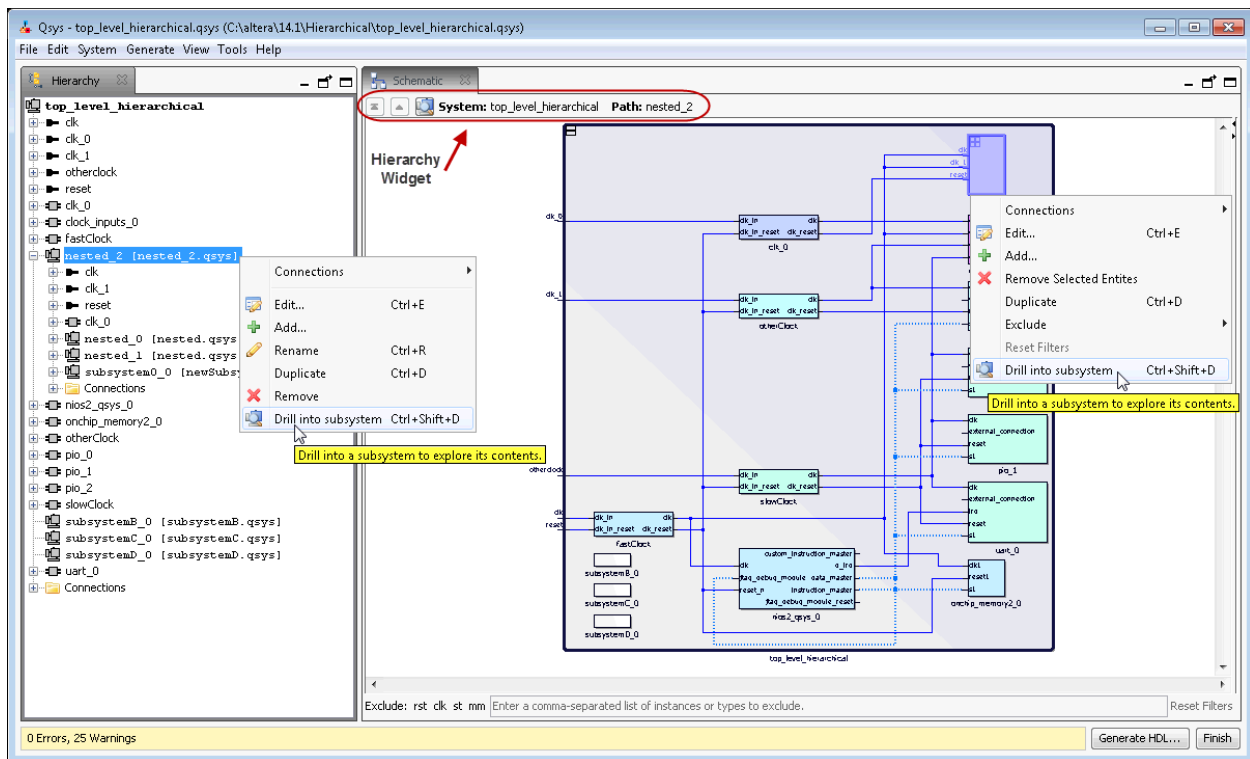
Views that manage system-level editing, for example, the **System Contents** and **Schematic** tabs, contain the hierarchy widget, which allows you to efficiently navigate your subsystems. The hierarchy widget also displays the name of the current selection, and its path in the context of the system or subsystem.

The hierarchy widget contains the following controls and information:

- **Top**—Navigates to the project-level **.qsys** file that contains the subsystem.
- **Up**—Navigates up one level from the current selection.
- **Drill Into**—Allows you to drill into an editable system.
- **System**—Displays the hierarchical location of the system you are currently editing.
- **Path**—Displays the relative path to the current selection.

**Note:** In the **System Contents** tab, you can use CTRL+SHIFT+U to navigate up one level, and CTRL+SHIFT+D to drill into a system.


Figure 5-13: Drill into a Qsys System to Explore its Contents



## Edit a Qsys Subsystem

You can double-click a Qsys subsystem in the **Hierarchy** tab to edit its contents in any tab. When you make a change, open tabs refresh their content to reflect your edit. You can change the level of a subsystem, or push it into another subsystem with commands in the **System Contents** tab.

**Note:** To edit a **.qsys** file, the file must be writeable and reside outside of the ACDS installation directory. You cannot edit systems that you create from composed **\_hw.tcl** files, or systems that define instance parameters.



1. In the **System Contents** or **Schematic** tabs, use the hierarchy widget to navigate to the top-level system, up one level, or down one level (drill into a system).

All tabs refresh and display the requested hierarchy level.

2. To edit a system, double-click the system in the **Hierarchy** tab. You can also drill into the system with the Hierarchy tool or right-click commands, which are available in the **Hierarchy**, **Schematic**, **System Contents** tabs.

The system is open and available for edit in all Qsys views. A system currently open for edit appears as bold in the **Hierarchy** tab.

3. In the **System Contents** tab, you can rename any element, add, remove, or duplicate connections, and export interfaces, as appropriate.

Changes to a subsystem affect all instances. Qsys identifies unsaved changes to a subsystem with an asterisk next to the subsystem in the **Hierarchy** tab.

#### Related Information

[View a Schematic of Your Qsys System](#) on page 5-10

## Change the Hierarchy Level of a Qsys Component

You can push selected components down into their own subsystem, which can simplify your top-level system view. Similarly, you can pull a component up out of a subsystem to perhaps share it between two unique subsystems. Hierarchical-level management facilitates system optimization and can reduce complex connectivity in your subsystems. When you make a change, open tabs refresh their content to reflect your edit.

1. In the **System Contents** tab, to group multiple components that perhaps share a system-level component, select the components, right-click, and then select **Push down into new subsystem**. Qsys pushes the components into their own subsystem and re-establishes the exported signals and connectivity in the new location.
2. In the **System Contents** tab, to pull a component up out of a subsystem, select the component, and then click **Pull up**. Qsys pulls the component up out of the subsystem and re-establishes the exported signals and connectivity in the new location.

## Save New Qsys Subsystem

When you save a subsystem to your Qsys design, Qsys confirms the new subsystem(s) in the **Confirm New System Filenames** dialog box. The **Confirm New System Filenames** dialog box appears when you save your Qsys design. Qsys uses the name that you give a subsystem as **.qsys** filename, and saves the subsystems in the project's ip directory.

1. Click **File > Save** to save your Qsys design.
2. In the **Confirm New System Filenames** dialog box, click **OK** to accept the subsystem file names.

**Note:** If you have not yet saved your top-level system, or multiple subsystems, you can type a name, and then press **Enter**, to move to the next un-named system.

3. In the **Confirm New System Filenames** dialog box, to edit the name of a subsystem, click the subsystem, and then type the new name.
4. To cancel the save process, click **Cancel** in the **Confirm New System Filenames** dialog box.

## Create an IP Component Based on a Qsys System

The **Export System as hw.tcl Component** command on the Qsys File menu allows you to save the system currently open in Qsys as an **\_hw.tcl** file in project directory. The saved system displays as a new component under the **System** category in the IP Catalog.

## Hierarchical System Using Instance Parameters Example

This example illustrates how you can use instance parameters to control the implementation of an On-Chip Memory component, `onchip_memory_0` when instantiated into a higher-level Qsys system.

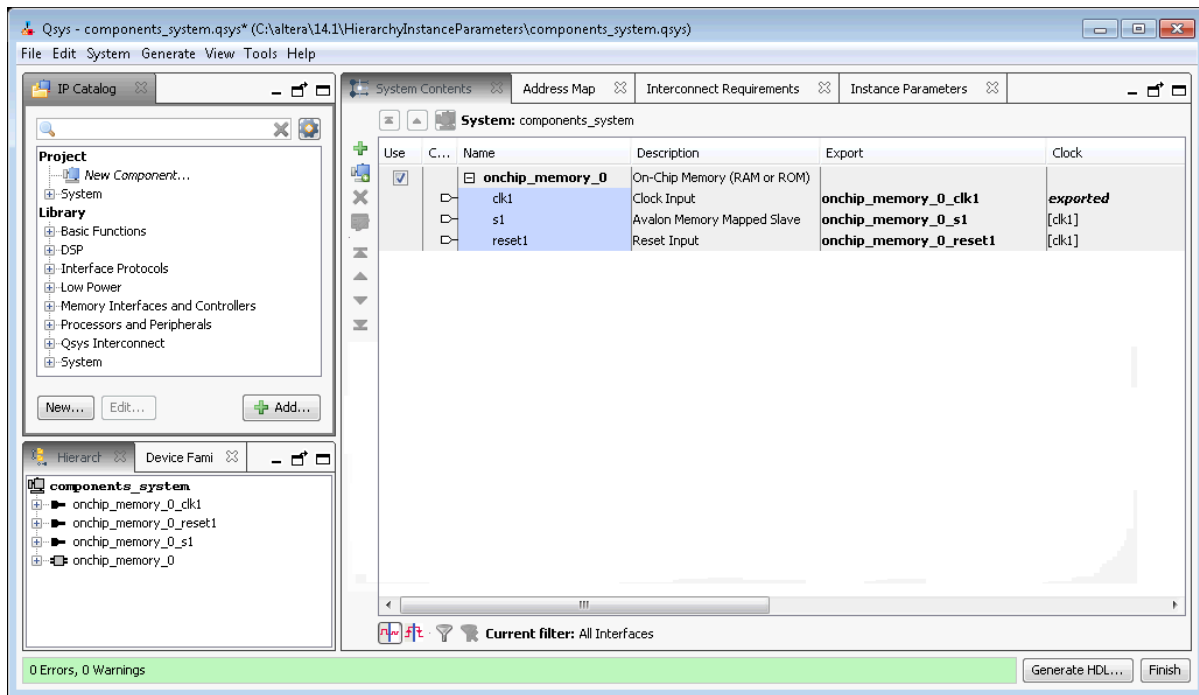
Follow the steps below to create a system that contains an On-Chip Memory IP component with instance parameters, and the instantiating higher-level Qsys system. With your completed system, you can vary the values of the instance parameters to review their effect within the On-Chip Memory component.

### Create the Component System

This procedure creates a Qsys system to use as subsystem in a higher-level system as part of a hierarchical instance parameter example.

1. In Qsys, click **File > New System**.
2. Right-click `clk_0`, and then click **Remove**.
3. In the IP Catalog search box, type **on-chip** to locate the On-Chip Memory (RAM or ROM) component.
4. Double-click to add the On-Chip Memory component to your system.  
The parameter editor opens. When you click **Finish**, Qsys adds the component to your system.
5. Rename the On-Chip Memory component to `onchip_memory_0`.
6. In the **System Contents** tab, for the `clk1` element (`onchip_memory_0`), double-click the **Export** column.
7. In the **System Contents** tab, for the `s1` element (`onchip_memory_0`), double-click the **Export** column.
8. In the **System Contents** tab, for the `reset1` element (`onchip_memory_0`), double-click the **Export** column.
9. Click **File > Save** to save your Qsys system as **components\_system.qsys**.

**Figure 5-14: On-Chip Memory Component System and Instance Parameters (components\_system.qsys)**

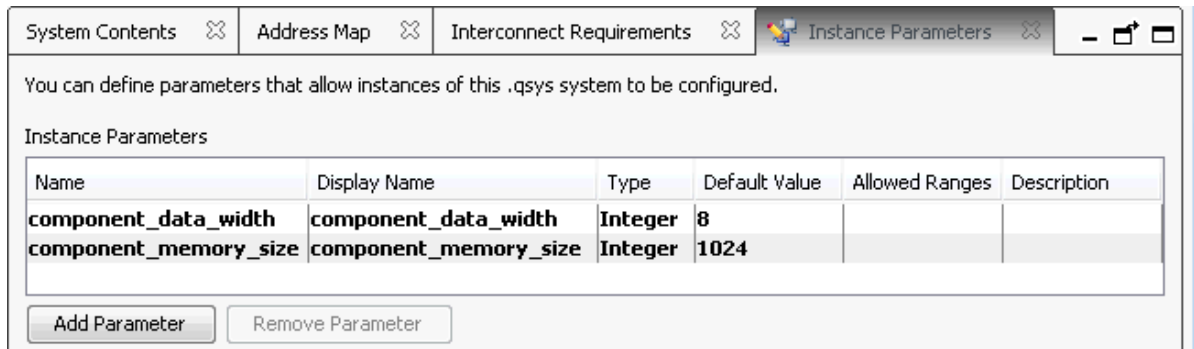


## Add Qsys Instance Parameters

The **Instance Parameters** tab allows you to define parameters to control the implementation of a subsystem component. Each column in the **Instance Parameters** table defines a property of the parameter. This procedure creates instance parameters in a Qsys system to use as subsystem in a higher-level system as part of a hierarchical instance parameter example.

1. In the **components\_system.qsys** system, click **View > Instance Parameters**.
2. Click **Add Parameter**.
3. In the **Name** and **Display Name** columns, rename the `new_parameter_0` parameter to `component_data_width`.
4. For `component_data_width`, select **Integer** for **Type**, and 8 as the **Default Value**.
5. Click **Add Parameter**.
6. In the **Name** and **Display Name** columns, rename the `new_parameter_0` parameter to `component_memory_size`.
7. For `component_size`, select **Integer** for **Type**, and 1024 as the **Default Value**.

Figure 5-15: Qsys Instance Parameters Tab



- In the **Instance Script** section, type the commands that control how the parameters are passed to an instance from the higher-level system. For example, in the script below, the `onchip_memory_0` instance receives its `dataWidth` and `memorySize` parameter values from the instance parameters that you define.

Figure 5-16: Instance Script

```
Instance Script
# request a specific version of the scripting API
package require -exact qsys 14.1

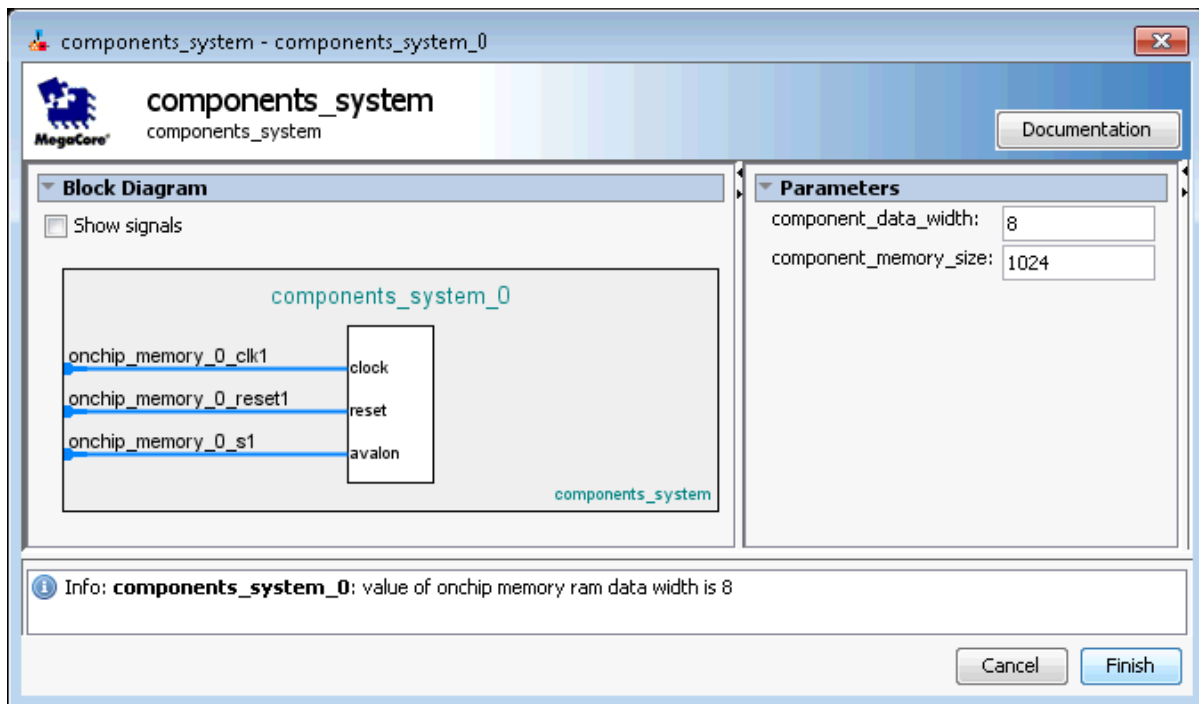
# Set the name of the procedure to manipulate parameters
set_module_property COMPOSITION_CALLBACK compose

proc compose () {
    # manipulate parameters in here
    set_instance_parameter_value onchip_memory_0 dataWidth [ get_parameter_value component_data_width ]
    set_instance_parameter_value onchip_memory_0 memorySize [ get_parameter_value component_memory_size ]

    set value [get_instance_parameter_value onchip_memory_0 dataWidth]
    send_message info "value of onchip memory ram data width is $value "
}
```

- Click **Preview Instance** to view the parameter editor GUI. **Preview Instance** allows you to see how an instance of a system appears when you use it in another system.

Figure 5-17: Preview Your Instance in the Parameter Editor



10. Click **File > Save**.

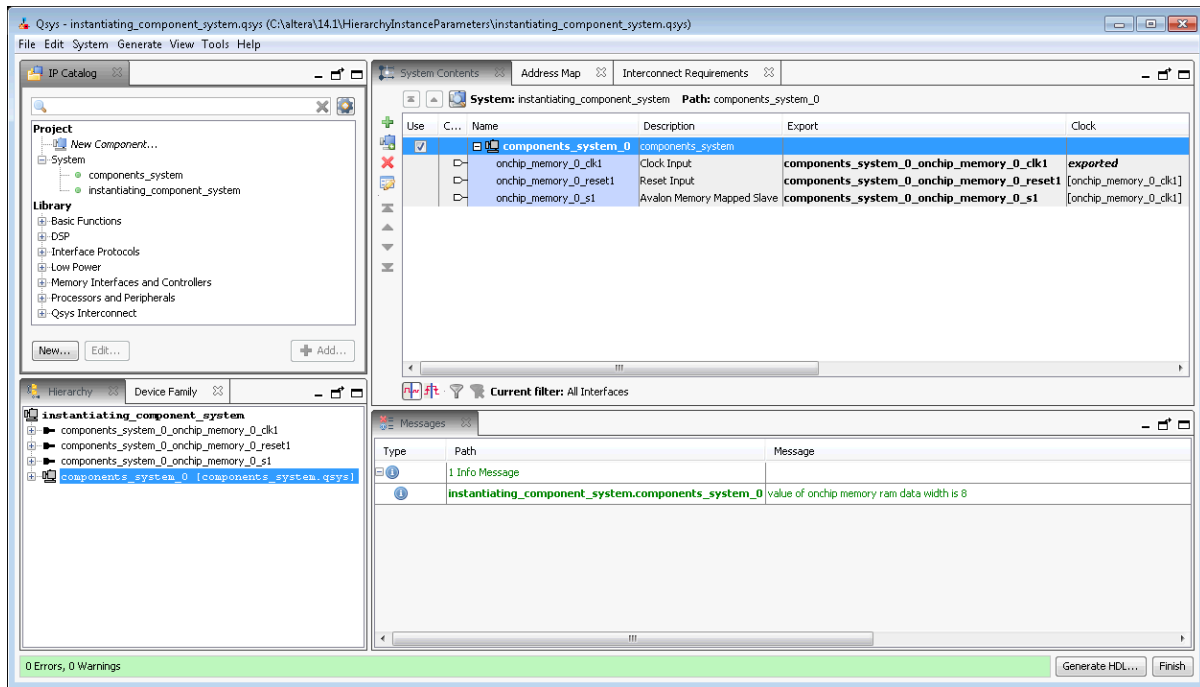
### Create a Qsys Instantiating Component System

This procedure creates a Qsys system to use as a higher-level system as part of a hierarchical instance parameter example.

1. In Qsys, click **File > New System**.
2. Right-click `clk_0`, and then click **Remove**.
3. In the IP Catalog, under **System**, double-click **components\_system**.
4. In the **Systems Contents** tab, for each element under **system\_0**, double-click the **Export** column.
5. Click **File > Save** to save your Qsys as **instantiating\_component\_system.qsys**.



Figure 5-18: Instantiating Component System (instantiating\_component\_system.qsys)

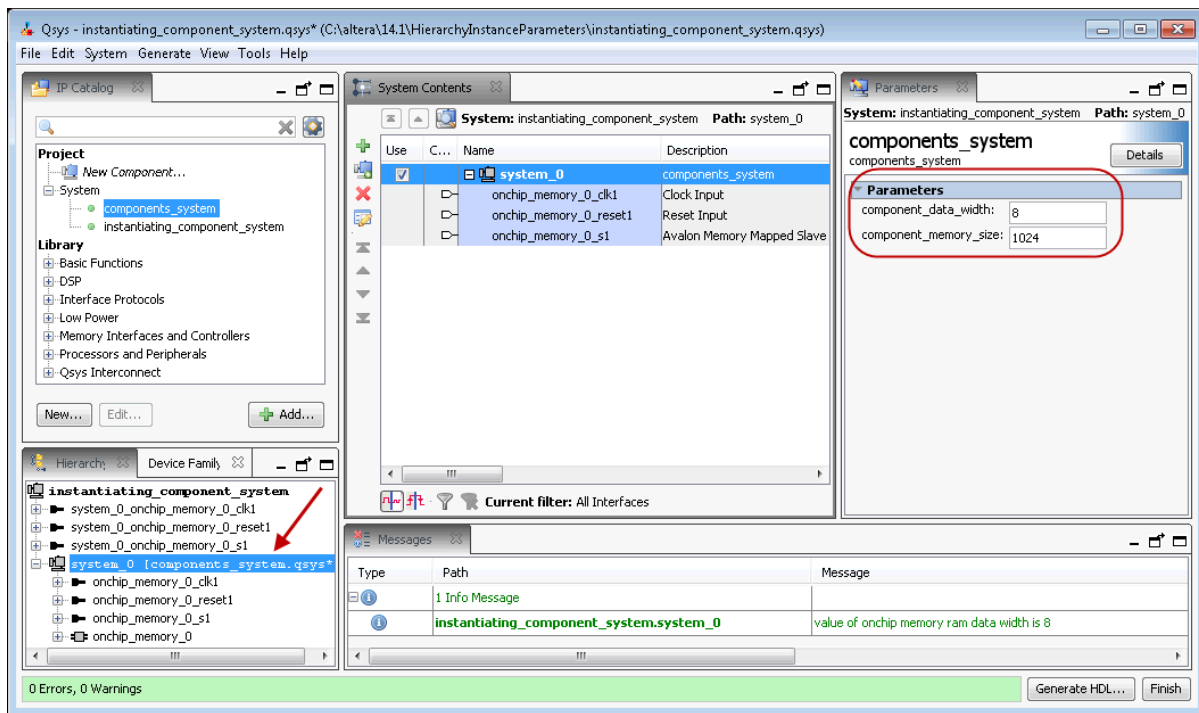


## Apply Instance Parameters at a Higher-Level Qsys System and Pass the Parameters to the Instantiated Lower-Level System

This procedure shows you how to use instance parameters to control the implementation of an On-Chip Memory component as part of a hierarchical instance parameter example.

1. In the **instantiating\_component\_system.qsys** system, in the **Hierarchy** tab, click and expand **system\_0 (components\_system.qsys)**.
2. Click **View > Parameters**.

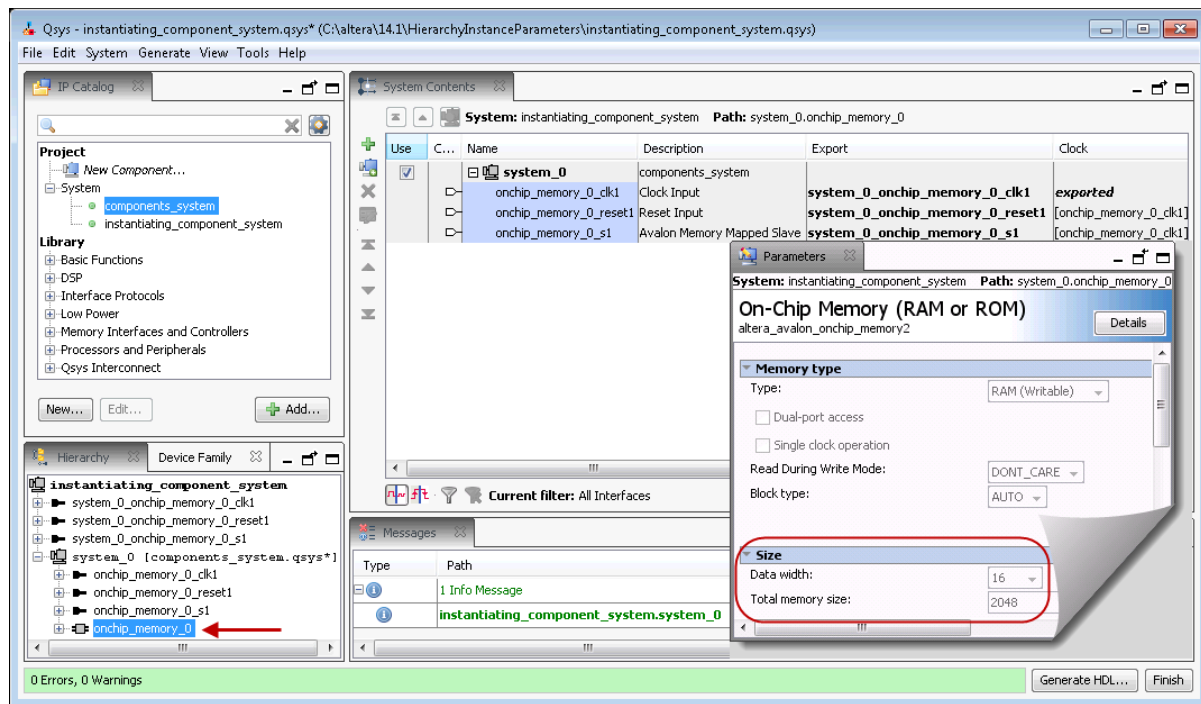
The instance parameters that you defined in **components\_system.qsys** display in the Parameter Editor.



3. In the **Parameters** tab, change the value of **component\_data\_width** to 16, and **component\_memory\_size** to 2048.
4. In the **Hierarchy** tab, under **system\_0** (**components\_system.qsys**), click **onchip\_memory\_0**. When you select **onchip\_memory\_0**, the new parameter values for **Data width** and **Total memory size** are displayed.

The entries for the On-Chip Memory component appear grayed out at this lower-level of hierarchy.

Figure 5-20: Vary the Values for Your Instance Parameters



## View and Filter Clock and Reset Domains in Your Qsys System

The Qsys clock and reset domains tabs allow you to see clock domains and reset domains in your Qsys system. Qsys determines clock and reset domains by the associated clocks and resets, which are displayed in tooltips for each interface in your system. You can filter your system to display particular components or interfaces within a selected clock or reset domain. The clock and reset domain tabs also provide quick access to performance bottlenecks by indicating connection points where Qsys automatically inserts clock crossing adapters and reset synchronizers during system generation. With these tools, you can more easily create optimal connections between interfaces.

Click **View > Clock Domains**, or **View > Reset Domains** to open the respective tabs in your workspace. The domain tools display as a tree with the current system at the root. You can select each clock or reset domain in the list to view associated interfaces.

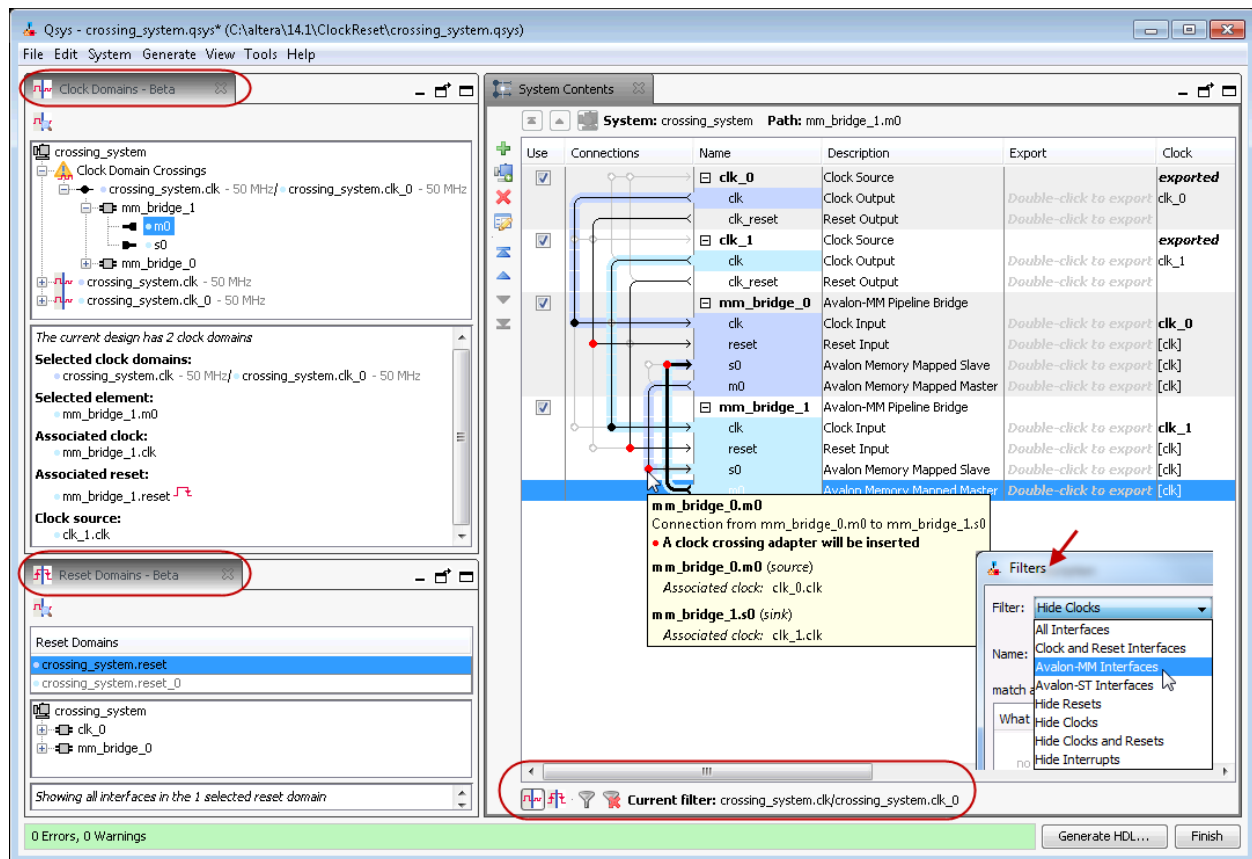
When you select an element in the **Clock Domains** tab, the corresponding selection appears in the **System Contents** tab. You can select single or multiple interface(s) and module(s). Mouse over tooltips in the **System Contents** tab to provide detailed information for all elements and connections. Colors that appear for the clocks and resets in the domain tools correspond to the colors in the **System Contents** and **Schematic** tabs.

Clock and reset control tools at the bottom on the **System Contents** tab allow you to toggle between highlighting clock or reset domains. You can further filter your view with options in the **Filters** dialog box, which is accessible by clicking the filter icon at the bottom of the **System Contents** tab. In the **Filters** dialog box, you can choose to view a single interface, or to hide clock, reset, or interrupt interfaces.

Clock and reset domain tools respond to global selection and edits, and help to provide answers to the following system design questions:

- How many clock and reset domains do you have in your Qsys system?
- What interfaces and modules does each clock or reset domain contain?
- Where do clock or reset crossings occur?
- At what connection points does Qsys automatically insert clock or reset adapters?
- Where do you have to manually insert a clock or reset adapter?

Figure 5-21: Qsys Clock and Reset Domains




## View Clock Domains in Your Qsys System

With the **Clock Domains** tab, you can filter the **System Contents** tab to display a single clock domain, or multiple clock domains. You can further filter your view with selections in the **Filters** dialog box. When you select an element in the **Clock Domains** tab, the corresponding selection appears highlighted in the **System Contents** tab.

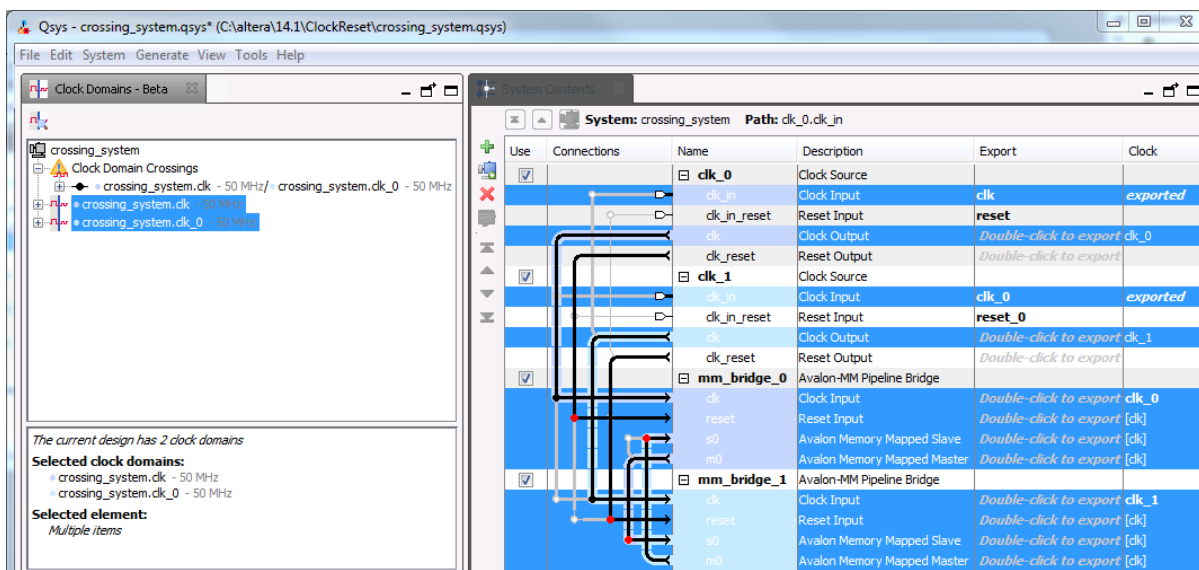
1. To view clock domain interfaces and their connections in your Qsys system, click **View > Clock Domains** to open the Clock Domains tab.
2. To enable and disable highlighting of the clock domains in the **System Contents** tab, click the clock control tool at the bottom of the **System Contents** tab.

Figure 5-22: Clock Control Tool



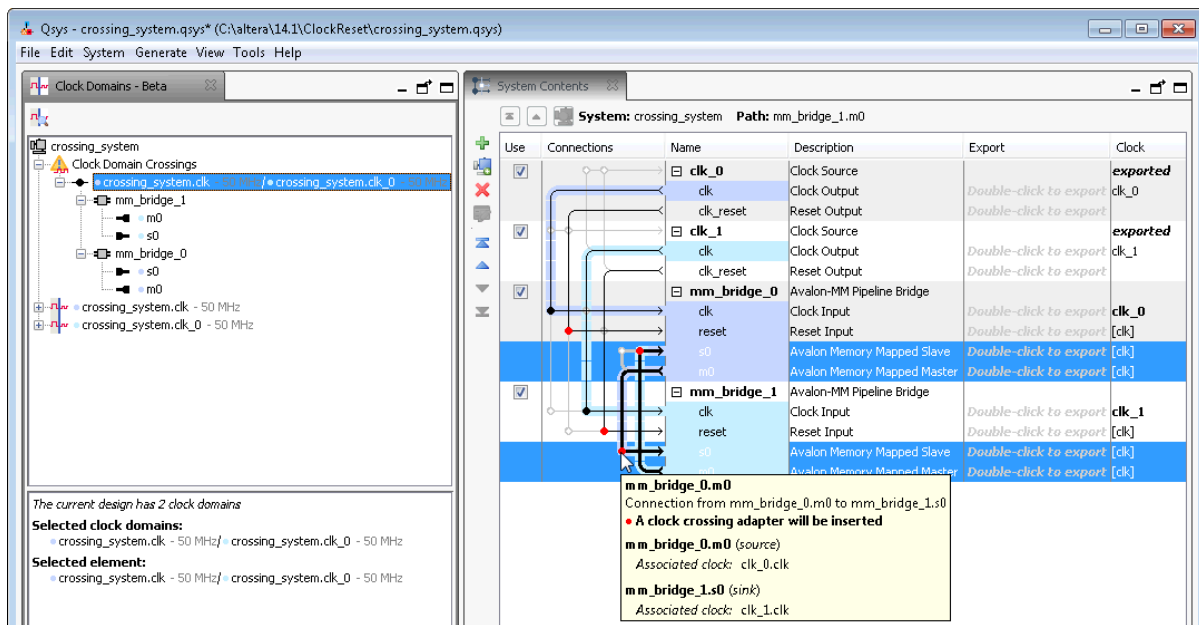
- To view a single clock domain, or multiple clock domains and their modules and connections, click the clock name(s) in the **Clock Domains** tab.  
The modules for the selected clock domain(s) and their connections appear highlighted in the **System Contents** tab. Detailed information for the current selection appears in the clock domain details pane. Red dots in the **Connections** column indicate auto insertions by Qsys during system generation, for example, a reset synchronizer or clock crossing adapter.

Figure 5-23: Clock Domains



- To view interfaces that cross clock domains, expand the **Clock Domain Crossings** icon in the **Clock Domains** tab, and select each element to view its details in the **System Contents** tab.  
Qsys lists the interfaces that cross clock domain under **Clock Domain Crossings**. As you click through the elements, detailed information appears in the clock domain details pane. Qsys also highlights the selection in the **System Contents** tab.  
If a connection crosses a clock domain, the connection circle appears as a red dot in the **System Contents** tab. Mouse over tooltips at the red dot connections provide details about the connection, as well as what adapter type Qsys automatically inserts during system generation.

Figure 5-24: Clock Domain Crossings




## View Reset Domains in Your Qsys System

With the **Reset Domains** tab, you can filter the **System Contents** tab to display a single reset domain, or multiple reset domains. When you select an element in the **Reset Domains** tab, the corresponding selection appears in the **System Contents** tab.

1. To view reset domain interfaces and their connections in your Qsys system, click **View > Reset Domains** to open the **Reset Domains** tab.
2. To show reset domains in the **System Contents** tab, click the reset control tool at the bottom of the **System Contents** tab.

Figure 5-25: Reset Control Tool



3. To view a single reset domain, or multiple reset domains and their modules and connections, click the reset name(s) in the **Reset Domain** tab.

Qsys displays your selection according to the following rules:

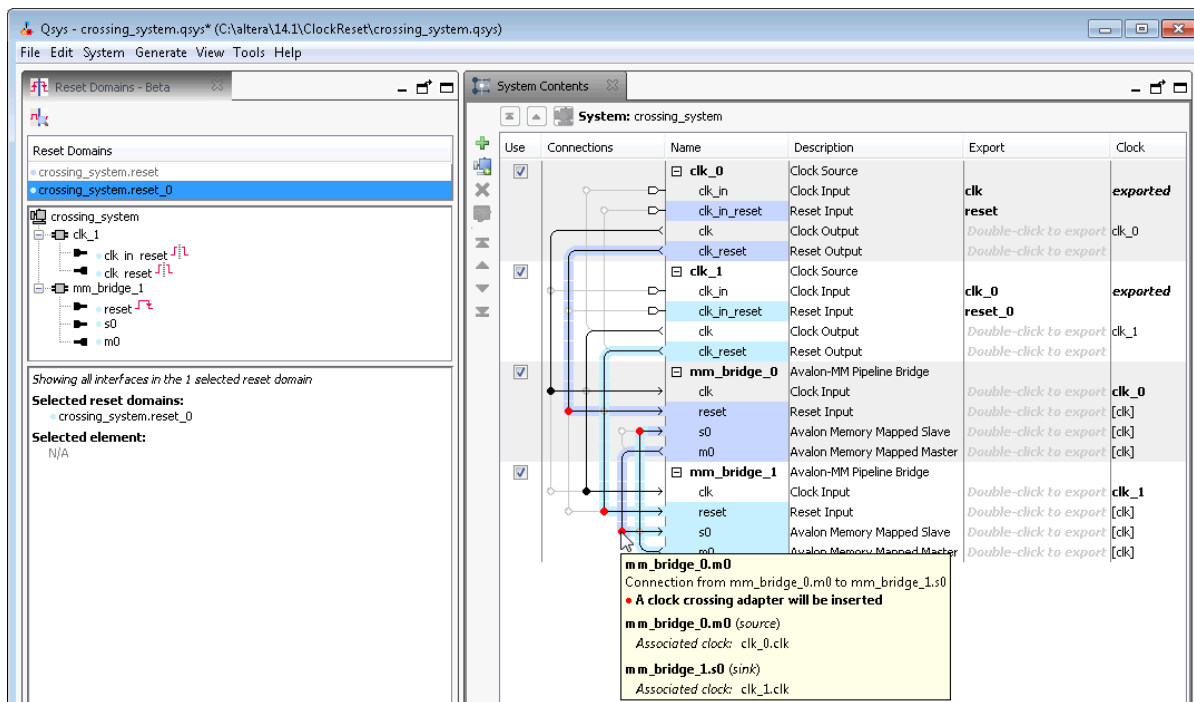
- When you select multiple reset domains, the **System Contents** tab shows interfaces and modules in both reset domains.
- When you select a single reset domain, the other reset domain(s) are grayed out, unless the two domains have interfaces in common.
- Reset interfaces appear black when connected to multiple reset domains.
- Reset interfaces appear gray when they are not connected to all of the selected reset domains.
- If an interface is contained in multiple reset domains, the interface is grayed out.

Detailed information for your selection appears in the reset domain details pane.

**Note:** Red dots in the **Connections** column between reset sinks and sources indicate auto insertions by Qsys during system generation, for example, a reset synchronizer. Qsys decides when to display a red dot with the following protocol, and ends the decision process at first match.

- Multiple resets fan into a common sink.
- Reset inputs are associated with different clock domains.
- Reset inputs have different synchronicity.

Figure 5-26: Reset Domains

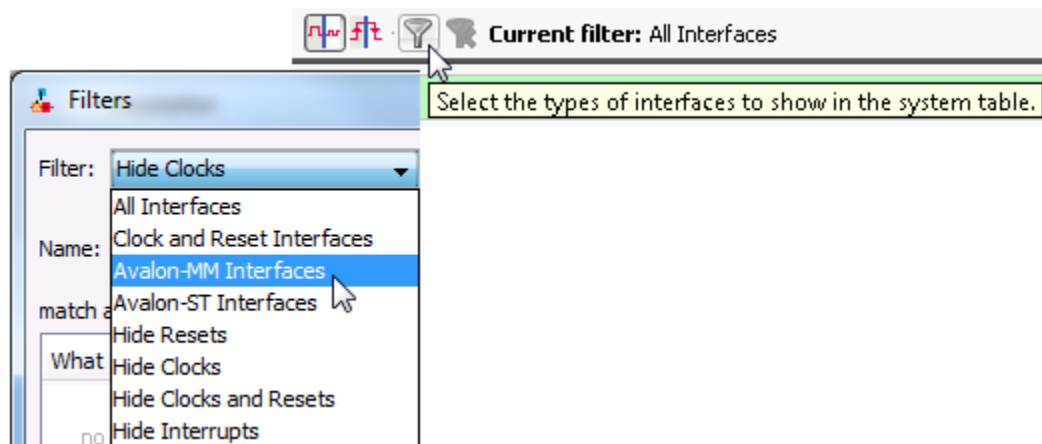


## Filter Qsys Clock and Reset Domains in the System Contents Tab

You can filter the display of your Qsys clock and reset domains in the **System Contents** tab.


1. To filter the display in the **System Contents** tab to view only a particular interface and its connections, or to choose to hide clock, reset, or interrupt interfaces, click the **Filters** icon in the clock and reset control tool to open the **Filters** dialog box.  
The selected interfaces appear in the **System Contents** tab.

Figure 5-27: Filters Dialog Box



- To clear all clock and reset filters in the **System Contents** tab and show all interfaces, click the **Filters** icon with the red "x" in the clock and reset control tool.

Figure 5-28: Show All Interfaces



## Define Qsys Instance Parameters

You can use instance parameters to test the functionality of your Qsys system when you use another system as a sub-component. A higher-level Qsys system can assign values to instance parameters, and then test those values in the lower-level system.

The **Instance Script** on the **Instance Parameters** tab defines how the specified values for the instance parameters affect the sub-components in your Qsys system. The instance script allows you to create queries for the instance parameters and set the values of the parameters for the lower-level system components.

When you click **Preview Instance**, Qsys creates a preview of the current Qsys system with the specified parameters and instance script and opens the parameter editor. This command allows you to see how an instance of a system appears when you use it in another system. The preview instance does not affect your saved system.

To use instance parameters, the IP components or subsystems in your Qsys system must have parameters that can be set when they are instantiated in a higher-level system.

If you create hierarchical Qsys systems, each Qsys system in the hierarchy can include instance parameters to pass parameter values through multiple levels of hierarchy.

### Related Information

[Working with Instance Parameters in Qsys](#)



## Create an Instance Parameter Script in Qsys

The first command in an instance parameter script must specify the Tcl command version. The version command ensures the Tcl commands behave identically in future versions of the tool. Use the following command to specify the version of the Tcl commands, where *<version>* is the Quartus II software version number, such as 14.1:

```
package require -exact qsys <version>
```

To use Tcl commands that work with instance parameters in the instance script, you must specify the commands within a Tcl composition callback. In the instance script, you specify the name for the composition callback with the following command:

```
set_module_property COMPOSITION_CALLBACK <name of callback procedure>
```

Specify the appropriate Tcl commands inside the Tcl procedure with the following syntax:

```
proc <name of procedure defined in previous command> {}  
{#Tcl commands to query and set parameters go here}
```

### Example 5-3: Instance Script Example

In this example, an instance script uses the `pio_width` parameter to set the `width` parameter of a parallel I/O (PIO) component. The script combines the `get_parameter_value` and `set_instance_parameter_value` commands using brackets.

```
# Request a specific version of the scripting API  
package require -exact qsys 13.1  
  
# Set the name of the procedure to manipulate parameters:  
set_module_property COMPOSITION_CALLBACK compose  
  
proc compose {} {  
  
# Get the pio_width parameter value from this Qsys system and  
# pass the value to the width parameter of the pio_0 instance  
  
set_instance_parameter_value pio_0 width \  
[get_parameter_value pio_width]  
}
```

### Related Information

- [Component Interface Tcl Reference](#) on page 9-1

## Supported Tcl Commands for Qsys Instance Parameter Scripts

You can use standard Tcl commands to manipulate parameters in the script, such as the `set` command to create variables, or the `expr` command for mathematical manipulation of the parameter values. Instance scripts also use Tcl commands to query the parameters of a Qsys system, or to set the values of the parameters of the sub-IP-components instantiated in the system.

### `get_instance_parameter_value`

#### Description

Returns the value of a parameter in a child instance.

#### Usage

```
get_instance_parameter_value <instance> <parameter>
```

#### Returns

The value of the parameter.

#### Arguments

**instance**

The name of the child instance.

**parameter**

The name of the parameter in the instance.

#### Example

```
get_instance_parameter_value pixel_converter input_DPI
```



## get\_instance\_parameters

### Description

Returns the names of all parameters on a child instance that can be manipulated by the parent. It omits parameters that are derived and those that have the `SYSTEM_INFO` parameter property set.

### Usage

```
get_instance_parameters <instance>
```

### Returns

A list of parameters in the instance.

### Arguments

**instance**

The name of the child instance.

### Example

```
get_instance_parameters instance
```

## `get_parameter_value`

### Description

Returns the current value of a parameter defined previously with the `add_parameter` command.

### Usage

```
get_parameter_value <parameter>
```

### Returns

The value of the parameter.

### Arguments

#### **parameter**

The name of the parameter whose value is being retrieved.

### Example

```
get_parameter_value fifo_width
```

## get\_parameters

### Description

Returns the names of all the parameters in the component.

### Usage

```
get_parameters
```

### Returns

A list of parameter names.

### Arguments

No arguments.

### Example

```
get_parameters
```

## send\_message

### Description

Sends a message to the user of the component. The message text is normally interpreted as HTML. The `<b>` element can be used to provide emphasis. If you do not want the message text to be interpreted as HTML, then pass a list like `{ Info Text }` as the message level

### Usage

```
send_message <level> <message>
```

### Returns

No return value.

### Arguments

#### level

The following message levels are supported:

- `ERROR`--Provides an error message.
- `WARNING`--Provides a warning message.
- `INFO`--Provides an informational message.
- `PROGRESS`--Provides a progress message.
- `DEBUG`--Provides a debug message when debug mode is enabled.

#### message

The text of the message.

### Example

```
send_message ERROR "The system is down!"  
send_message { Info Text } "The system is up!"
```

## set\_instance\_parameter\_value

### Description

Sets the value of a parameter for a child instance. Derived parameters and `SYSTEM_INFO` parameters for the child instance may not be set using this command.

### Usage

```
set_instance_parameter_value <instance> <parameter> <value>
```

### Returns

No return value.

### Arguments

**instance**

The name of the child instance.

**parameter**

The name of the parameter.

**value**

The new parameter value.

### Example

```
set_instance_parameter_value uart_0 baudRate 9600
```

## set\_module\_property

### Description

Used to specify the Tcl procedure invoked to evaluate changes in Qsys system instance parameters.

### Usage

```
set_module_property <property> <value>
```

### Returns

No return value.

### Arguments

#### property

The name of the property. Refer to [Module Properties](#).

#### value

The new value of the property.

### Example

```
set_module_property COMPOSITION_CALLBACK "my_composition_callback"
```

## Specify Qsys Interconnect Requirements

The **Interconnect Requirements** tab allows you to apply system-wide, `$system`, and interface interconnect requirements for IP components in your system. Options in the **Setting** column vary depending on what you select in the **Identifier** column

Table 5-3: Specifying System-Wide Interconnect Requirements

Option	Description
<b>Limit interconnect pipeline stages to</b>	Specifies the maximum number of pipeline stages that Qsys may insert in each command and response path to increase the $f_{MAX}$ at the expense of additional latency. You can specify between 0–4 pipeline stages, where 0 means that the interconnect has a combinational data path. Choosing 3 or 4 pipeline stages may significantly increase the logic utilization of the system. This setting is specific for each Qsys system or subsystem, meaning that each subsystem can have a different setting. Additional latency is added once on the command path, and once on the response path. You can manually adjust this setting in the <b>Memory-Mapped Interconnect</b> tab. Access this tab by clicking <b>Show System With Qsys Interconnect command</b> on the System menu.



Option	Description
<p><b>Clock crossing adapter type</b></p>	<p>Specifies the default implementation for automatically inserted clock crossing adapters:</p> <ul style="list-style-type: none"> <li>• <b>Handshake</b>—This adapter uses a simple hand-shaking protocol to propagate transfer control signals and responses across the clock boundary. This methodology uses fewer hardware resources because each transfer is safely propagated to the target domain before the next transfer can begin. The <b>Handshake</b> adapter is appropriate for systems with low throughput requirements.</li> <li>• <b>FIFO</b>—This adapter uses dual-clock FIFOs for synchronization. The latency of the FIFO-based adapter is a couple of clock cycles more than the handshaking clock crossing component. However, the FIFO-based adapter can sustain higher throughput because it supports multiple transactions at any given time. FIFO-based clock crossing adapters require more resources. The <b>FIFO</b> adapter is appropriate for memory-mapped transfers requiring high throughput across clock domains.</li> <li>• <b>Auto</b>—If you select <b>Auto</b>, Qsys specifies the <b>FIFO</b> adapter for bursting links, and the <b>Handshake</b> adapter for all other links.</li> </ul>

**Table 5-4: Specifying \$system Interconnect Requirements**

You can apply the following interconnect requirements when you select `$system` as the **Identifier** in the **Interconnect Requirements** tab, in the **All Requirements** table. `$system` is the current system that is open in Qsys.

Option	Description
<p><b>Limit interconnect pipeline stages to</b></p>	<p>Refer to <i>Specifying System-Wide Interconnect Requirements</i>.</p>
<p><b>Clock crossing adapter type</b></p>	<p>Refer to <i>Specifying System-Wide Interconnect Requirements</i>.</p>
<p><b>Automate default slave insertion</b></p>	<p>Specifies whether you want Qsys to automatically insert a default slave for undefined memory region accesses during system generation.</p>
<p><b>Enable instrumentation</b></p>	<p>When you set this option to <b>TRUE</b>, Qsys enables debug instrumentation in the Qsys interconnect, which then monitors interconnect performance in the system console.</p>

Option	Description
<b>Burst Adapter Implementation</b>	<p>Allows you to choose the converter type that Qsys applies to each burst.</p> <ul style="list-style-type: none"> <li>• <b>Generic converter (slower, lower area)</b>—Default. Controls all burst conversions with a single converter that is able to adapt incoming burst types. This results in an adapter that has lower <math>f_{max}</math>, but smaller area.</li> <li>• <b>Per-burst-type converter (faster, higher area)</b>—Controls incoming bursts with a particular converter, depending on the burst type. This results in an adapter that has higher <math>f_{max}</math>, but higher area. This setting is useful when you have AXI masters or slaves and you want a higher <math>f_{max}</math>.</li> </ul>

**Table 5-5: Specifying Interface Interconnect Requirements**

You can apply the following interconnect requirements when you select a component interface as the **Identifier** in the **Interconnect Requirements** tab, in the **All Requirements** table.

Option	Value	Description
<b>Security</b>	<ul style="list-style-type: none"> <li>• Non-secure</li> <li>• Secure</li> <li>• Secure ranges</li> <li>• TrustZone-aware</li> </ul>	<p>After you establish connections between the masters and slaves, allows you to set the security options, as needed, for each master and slave in your system.</p> <p><b>Note:</b> You can also set these values in the <b>Security</b> column in the <b>System Contents</b> tab.</p>
<b>Secure address ranges</b>	Accepts valid address range.	Allows you to type in any valid address range.

For more information about HPS, refer to the *Cyclone V Device Handbook* in volume 3 of the *Hard Processor System Technical Reference Manual*.

## Manage Qsys System Security

TrustZone is the security extension of the ARM<sup>®</sup>-based architecture. It includes secure and non-secure transactions designations, and a protocol for processing between the designations. TrustZone security support is a part of the Qsys interconnect.

The AXI `AXIPROT` protection signal specifies a secure or non-secure transaction. When an AXI master sends a command, the `AXIPROT` signal specifies whether the command is secure or non-secure. When an AXI slave receives a command, the `AXIPROT` signal determines whether the command is secure or non-secure. Determining the security of a transaction while sending or receiving a transaction is a run-time protocol.

The Avalon specification does not include a protection signal as part of its specification. When an Avalon master sends a command, it has no embedded security and Qsys recognizes the command as non-secure.

When an Avalon slave receives a command, it also has no embedded security, and the slave always accepts the command and responds.

AXI masters and slaves can be TrustZone-aware. All other master and slave interfaces, such as Avalon-MM interfaces, are non-TrustZone-aware. You can set compile-time security support for all components (except AXI masters, including AXI3, AXI4, and AXI4-Lite) in the **Security** column in the **System Contents** tab, or in the **Interconnect Requirements** tab under the **Identifier** column for the master or slave interface. To begin creating a secure system, you must first add masters and slaves to your system, and the connections between them. After you establish connections between the masters and slaves, you can then set the security options, as needed

An example of when you may need to specify compile-time security support is when an Avalon master needs to communicate with a secure AXI slave, and you can specify whether the connection point is secure or non-secure. You can specify a compile-time secure address ranges for a memory slave if an interface-level security setting is not sufficient.

**Related Information**

- [Qsys Interconnect](#) on page 7-1
- [Qsys System Design Components](#) on page 10-1

## Configure Qsys Security Settings Between Interfaces

The AXI `axprot` signal specifies a transaction as secure or non-secure at runtime when a master sends a transaction. Qsys identifies AXI master interfaces as TrustZone-aware. You can configure AXI slaves as Trustzone-aware, secure, non-secure, or secure ranges.

**Table 5-6: Compile-Time Security Options**

For non-TrustZone-aware components, compile-time security support options are available in Qsys on the **System Contents** tab, or on the **Interconnect Requirements** tab.

Compile-Time Security Options	Description
<b>Non-secure</b>	Master sends only non-secure transactions, and the slave receives any transaction, secure or non-secure.
<b>Secure</b>	Master sends only secure transactions, and the slave receives only secure transactions.
<b>Secure ranges</b>	Applies to only the slave interface. The specified address ranges within the slave's address span are secure, all other address ranges are not. The format is a comma-separated list of inclusive-low and inclusive-high addresses, for example, <code>0x0:0xffff, 0x2000:0x20ff</code> .

After setting compile-time security options for non-TrustZone-aware master and slave interfaces, you must identify those masters that require a default slave before generation. To designate a slave interface as the default slave, turn on **Default Slave** in the **System Contents** tab. A master can have only one default slave.

**Note:** The **Security** and **Default Slave** columns in the **System Contents** tab are hidden by default. Right-click the **System Contents** header to select which columns you want to display.

The following are descriptions of security support for master and slave interfaces. These description can guide you in your design decisions when you want to create secure systems that have mixed secure and non-TrustZone-aware components:

- All AXI, AXI4, and AXI4-Lite masters are TrustZone-aware.
- You can set AXI, AXI4, and AXI4-Lite slaves as Trust-Zone-aware, secure, non-secure, or secure range ranges.
- You can set non-AXI master interfaces as secure or non-secure.
- You can set non-AXI slave interfaces as secure, non-secure, or secure address ranges.

## Specify a Default Slave in a Qsys System

If a master issues "per-access" or "not allowed" transactions, your design must contain a default slave. Per-access refers to the ability of a TrustZone-aware master to allow or disallow access or transactions. A transaction that violates security is rerouted to the default slave and subsequently responds to the master with an error. You can designate any slave as the default slave.

You can share a default slave between multiple masters. You should have one default slave for each interconnect domain. An interconnect domain is a group of connected memory-mapped masters and slaves that share the same interconnect. The `altera_axi_default_slave` component includes the required TrustZone features.

You can achieve an optimized secure system by partitioning your design and carefully designating secure or non-secure address maps to maintain reliable data. Avoid a design where, under the same hierarchy, a non-secure master initiates transactions to a secure slave resulting in unsuccessful transfers.

**Table 5-7: Secure and Non-Secure Access Between Master, Slave, and Memory Components**

Transaction Type	TrustZone-aware Master	Non-TrustZone-aware Master Secure	Non-TrustZone-aware Master Non-Secure
TrustZone-aware slave/memory	OK	OK	OK
Non-TrustZone-aware slave (secure)	Per-access	OK	Not allowed
Non-TrustZone-aware slave (non-secure)	OK	OK	OK
Non-TrustZone-aware memory (secure region)	Per-access	OK	Not allowed
Non-TrustZone-aware memory (non-secure region)	OK	OK	OK

## Access Undefined Memory Regions

When a transaction from a master targets a memory region that is not specified in the slave memory map, it is known as an "access to an undefined memory region." To ensure predictable response behavior when this occurs, you must add a default slave to your design. Qsys then routes undefined memory region accesses to the default slave, which terminates the transaction with an error response.

You can designate any memory-mapped slave as a default slave. Altera recommends that you have only one default slave for each interconnect domain in your system. Accessing undefined memory regions can occur in the following cases:

- When there are gaps within the accessible memory map region that are within the addressable range of slaves, but are not mapped.
- Accesses by a master to a region that does not belong to any slaves that is mapped to the master.
- When a non-secured transaction is accessing a secured slave. This applies to only slaves that are secured at compilation time.
- When a read-only slave is accessed with a write command, or a write-only slave is accessed with a read command.

To designate a slave as the default slave, for the selected component, turn on **Default Slave** in the **Systems Content** tab.

**Note:** If you do not specify the default slave, Qsys automatically assigns the slave at the lowest address within the memory map for the master that issues the request as the default slave.

### Related Information

- [Qsys System Design Components](#) on page 10-1

## Integrate a Qsys System with the Quartus II Software

To integrate a Qsys system with your Quartus II project, you must add either the Qsys System File (**.qsys**) or the Quartus II IP File (**.qip**), but never both to your Quartus II project. Qsys creates the **.qsys** file when you save your Qsys system, and produces the **.qip** file when you generate your Qsys system. Both the **.qsys** and **.qip** files contain the information necessary for compiling your Qsys system within a Quartus II project.

You can choose to include the **.qsys** file automatically in your Quartus II project when you generate your Qsys system by turning on the **Automatically add Quartus II IP files to all projects** option in the Quartus II software (**Tools > Options > IP Settings**). If this option is turned off, the Quartus II software asks you if you want to include the **.qsys** file in your Quartus II project after you exit Qsys.

If you want file generation to occur as part of the Quartus II software's compilation, you should include the **.qsys** file in your Quartus II project. If you want to manually control file generation outside of the Quartus II software, you should include the **.qip** file in your Quartus II project.

**Note:** The Quartus II software generates an error message during compilation if you add both the **.qsys** and **.qip** files to your Quartus II project.

## Does Quartus II Overwrite Qsys-Generated Files During Compilation?

Qsys supports standard and legacy device generation. Standard device generation refers to generating files for the Arria 10 device, and later device families. Legacy device generation refers to generating files for device families prior to the release of the Arria 10 device, including Max 10 devices.

When you integrate your Qsys system with the Quartus II software, if a **.qsys** file is included as a source file, Qsys generates standard device files under **<system>/** next to the location of the **.qsys** file. For legacy devices, if a **.qsys** file is included as a source file, Qsys generates HDL files in the Quartus II project directory under **/db/ip**.

For standard devices, Qsys-generated files are only overwritten during Quartus II compilation if the **.qip** file is removed or missing. For legacy devices, each time you compile your Quartus II project with a **.qsys** file, the Qsys-generated files are overwritten. Therefore, you should not edit Qsys-generated HDL in the **/db/ip** directory; any edits made to these files are lost and never used as input to the Quartus HDL synthesis engine.

### Related Information

- [Add IP Components to the Qsys IP Catalog](#) on page 5-16
- [Generate a Qsys System](#) on page 5-53
- [Qsys Synthesis Standard and Legacy Device Output Directories](#) on page 5-56
- [Qsys Simulation Standard and Legacy Device Output Directories](#) on page 5-57
- [Introduction to Altera IP Cores](#)
- [Implementing and Parameterizing Memory IP](#)

## Integrate a Qsys System and the Quartus II Software With the .qsys File

Use the following steps to integrate your Qsys system and your Quartus II project using the **.qsys** file:

1. In Qsys, create and save a Qsys system.
2. To automatically include the **.qsys** file in the your Quartus II project during compilation, in the Quartus II software, select **Tools > Options > IP Settings**, and turn on **Automatically add Quartus II IP files to all projects**.
3. When the **Automatically add Quartus II IP files to all projects** option is not checked, when you exit Qsys, the Quartus II software displays a dialog box asking whether you want to add the **.qsys** file to your Quartus II project. Click **Yes** to add the **.qsys** file to your Quartus II project.
4. In the Quartus II software, select **Processing > Start Compilation**.

## Integrate a Qsys System and the Quartus II Software With the .qip File

Use the following steps to integrate your Qsys system and your Quartus II project using the **.qip** file:

1. In Qsys, create and save a Qsys system.
2. In Qsys, click **Generate HDL**.
3. In the Quartus II software, select **Assignments > Settings > Files**.
4. On the **Files** page, use the controls to locate your **.qip** file, and then add it to your Quartus II project.
5. In the Quartus II software, select **Processing > Start Compilation**.

## Manage IP Settings in the Quartus II Software

To specify the following IP Settings in the Quartus II software, click **Tools > Option > IP Settings**:

**Table 5-8: IP Settings**

Setting	Description
<b>Maximum Qsys memory usage</b>	Allows you to increase memory usage for Qsys if you experience slow processing for large systems, or if Qsys reports an <b>Out of Memory</b> error.
<b>IP generation HDL preference</b>	The Quartus II software uses this setting when the <b>.qsys</b> file appears in the <b>Files</b> list for the current project in the <b>Settings</b> dialog box and you run Analysis & Synthesis. Qsys uses this setting when you generate HDL files.
<b>Automatically add Quartus II IP files to all projects</b>	The Quartus II software uses this setting when you create an IP core file variation with options in the Quartus II IP Catalog and parameter editor. When turned on, the Quartus II software adds the IP variation files to the project currently open.
<b>IP Catalog Search Locations</b>	<p>The Quartus II software uses the settings that you specify for global and project search paths under <b>IP Search Locations</b>, in addition to the <b>IP Search Path</b> in Qsys (<b>Tools &gt; Options</b>), to populate the Quartus II software IP Catalog.</p> <p>Qsys uses the settings that you specify for global search paths under <b>IP Search Locations</b> to populate the Qsys IP Catalog, which appears in Qsys (<b>Tools &gt; Options</b>). Qsys uses the project search path settings to populate the Qsys IP Catalog when you open Qsys from within the Quartus II software (<b>Tools &gt; Qsys</b>), but not when you open Qsys from the command-line.</p>

**Note:** You can also access **IP Settings** by clicking **Assignments > Settings > IP Settings**. This access is available only when you have a Quartus II project open. This allows you access to **IP Settings** when you want to create IP cores independent of a Quartus II project. Settings that you apply or create in either location are shared.

### Opening Qsys with Additional Memory

If your Qsys system requires more than the 512 megabytes of default memory, you can increase the amount of memory either in the Quartus II software **Options** dialog box, or at the command-line.

- When you open Qsys from within the Quartus II software, you can increase memory for your Qsys system, by clicking **Tools > Options > IP Settings**, and then selecting the appropriate amount of memory with the **Maximum Qsys memory usage** option.
- When you open Qsys from the command-line, you can add an option to increase the memory. For example, the following `qsys-edit` command allows you to open Qsys with 1 gigabytes of memory.

```
qsys-edit --jvm-max-heap-size=1g
```

## Set Qsys Clock Constraints

Many IP components include Synopsys Design Constraint (**.sdc**) files that provide timing constraints. Generated **.sdc** files are included in your Quartus II project with the generated **.qip** file. For your top-level clocks and PLLs, you must provide clock and timing constraints in SDC format to direct synthesis and fitting to optimize the design appropriately, and to evaluate performance against timing constraints.

You can specify a base clock assignment for each clock input in the TimeQuest GUI or with the `create_clock` command, and then use the `derive_pll_clocks` command to define the PLL clock output frequencies and phase shifts for all PLLs in the Quartus II project using the **.sdc** file.

**Figure 5-29: Single Clock Input Signal**

For the case of a single clock input signal called `clk`, and one PLL with a single output, you can use the following commands in your Synopsys Design Constraint (**.sdc**) file:

```
create_clock -name master_clk -period 20 [get_ports {clk}]
derive_pll_clocks
```

Use	Connections	Name	Description	Export	Clock	Base
<input checked="" type="checkbox"/>		clk_0	Clock Source			
		clk_in	Clock Input	clk	exported	
		clk_in_reset	Reset Input	reset		
		clk	Clock Output	Double-click to export	clk_0	
		clk_reset	Reset Output	Double-click to export		
<input checked="" type="checkbox"/>		master_0	JTAG to Avalon Master Bridge			
		clk	Clock Input	Double-click to export	clk_0	
		clk_reset	Reset Input	Double-click to export		
		master	Avalon Memory Mapped Master	Double-click to export	[clk]	
		master_reset	Reset Output	Double-click to export		
<input checked="" type="checkbox"/>		my_pll	Avalon ALTPLL			
		inclk_interface	Clock Input	Double-click to export	clk_0	
		inclk_interface_reset	Reset Input	Double-click to export	[inclk_interfa...	
		pll_slave	Avalon Memory Mapped Slave	Double-click to export	[inclk_interfa...	0x0000_0000
		c0	Clock Output	Double-click to export	my_pll_c0	
		areset_conduit	Conduit	Double-click to export		
		locked_conduit	Conduit	Double-click to export		
		phasedone_conduit	Conduit	Double-click to export		
<input checked="" type="checkbox"/>		pio_0	PIO (Parallel I/O)			
		clk	Clock Input	Double-click to export	my_pll_c0	
		reset	Reset Input	Double-click to export	[clk]	
		s1	Avalon Memory Mapped Slave	Double-click to export	[clk]	0x0000_0020
		external_connection	Conduit	Double-click to export		



## Related Information

### [The Quartus II TimeQuest Timing Analyzer](#)

## Generate a Qsys System

In Qsys, you can choose options for generation of synthesis, simulation and testbench files for your Qsys system.

Qsys system generation creates the interconnect between IP components and generates synthesis and simulation HDL files. You can generate a testbench system that adds Bus Functional Models (BFMs) that interact with your system in a simulator.

When you make changes to a system, Qsys gives you the option to exit without generating. If you choose to generate your system before you exit, the **Generation** dialog box opens and allows you to select generation options.

The **Generate HDL** button in the lower-right of the Qsys window allows you to quickly generate synthesis and simulation files for your system.

## Related Information

- [Avalon Verification IP Suite User Guide](#)
- [Mentor Verification IP \(VIP\) Altera Edition \(AE\)](#)

## Set the Generation ID

The **Generation Id** parameter is a unique integer value that is set to a timestamp during Qsys system generation. System tools, such as NIOS II or HPS (Hard Processor System) use the **Generation ID** to ensure software-build compatibility with your Qsys system.

To set the **Generation Id** parameter, select the top-level system in the **Hierarchy** tab, and then locating the parameter in the open **Parameters** tab.

## Generate Files for Synthesis and Simulation

The Quartus II software uses Qsys-generated synthesis HDL files during compilation.

In Qsys, you can generate simulation HDL files (**Generate > Generate HDL**), which can include simulation-only features targeted towards your simulator. You can generate simulation files as Verilog, VHDL, or as a mixed-language simulation for use in your simulation environment.

**Note:** For a list of Altera-supported simulators, refer to *Simulating Altera Designs*.

Qsys supports standard and legacy device generation. Standard device generation refers to generating files for the Arria 10 device, and later device families. Legacy device generation refers to generating files for device families prior to the release of the Arria 10 device, including Max 10 devices.

The **Output Directory** option applies to both synthesis and simulation generation. By default, the path of the generation output directory is fixed relative to the **.qsys** file. You can change the default directory in the **Generation** dialog box for legacy devices. For standard devices, the generation directory is fixed to the Qsys project directory.

**Note:** If you need to change top-level I/O pin or instance names, create a top-level HDL file that instantiates the Qsys system. The Qsys-generated output is then instantiated in your design without changes to the Qsys-generated output files.

The following options in the **Generation** dialog box (**Generate > Generate HDL**) allow you to generate synthesis and simulation files:

Option	Description
<b>Create HDL design files for synthesis</b>	Generates Verilog HDL or VHDL design files for the system's top-level definition and child instances for the selected target language. Synthesis file generation is optional.
<b>Create timing and resource estimates for third-party EDA synthesis tools</b>	Generates a non-functional Verilog Design File (.v) for use by some third-party EDA synthesis tools. Estimates timing and resource usage for your IP component. The generated netlist file name is <code>&lt;your_ip_component_name&gt;_syn.v</code> .
<b>Create Block Symbol File (.bsf)</b>	Allows you to optionally create a (.bsf) file to use in a schematic Block Diagram File (.bdf).
<b>Create simulation model</b>	Allows you to optionally generate Verilog HDL or VHDL simulation model files, and simulation scripts.
<b>Allow mixed-language simulation</b>	Generates a simulation model that contains both Verilog and VHDL as specified by the individual IP cores. Using this option, each IP core produces their HDL using it's native implementation, which results in simulation HDL that is easier to understand and faster to simulate. You must have a simulator that supports mixed language simulation.  When turned off, Qsys generates all simulation files in the selected simulation model language.

#### Related Information

[Simulating Altera Designs](#)

## Files Generated for Qsys IP Components

Qsys generates the following files for your Qsys system that are used during synthesis and simulation.

**Table 5-9: Qsys-Generated IP Component Files**

File Name	Description
<code>&lt;system&gt;.qsys</code>	The Qsys system file. <code>&lt;system&gt;</code> is the name that you give your system.

File Name	Description
<system>.sopcinfo	<p>Describes the connections and IP component parameterizations in your Qsys system. You can parse its contents to get requirements when you develop software drivers for IP components.</p> <p>Downstream tools such as the Nios II tool chain use this file. The <b>.sopcinfo</b> file and the <b>system.h</b> file generated for the Nios II tool chain include address map information for each slave relative to each master that accesses the slave. Different masters may have a different address map to access a particular slave component.</p>
<system>.cmp	The VHDL Component Declaration ( <b>.cmp</b> ) file is a text file that contains local generic and port definitions that you can use in VHDL design files.
<system>.html	A system report that contains connection information, a memory map showing the address of each slave with respect to each master to which it is connected, and parameter assignments.
<system>_generation.rpt	Qsys generation log file. A summary of the messages that Qsys issues during system generation.
<system>.debuginfo	Contains post-generation information. Used to pass System Console and Bus Analyzer Toolkit information about the Qsys interconnect. The Bus Analysis Toolkit uses this file to identify debug components in the Qsys interconnect.
<system>.qip	Contains all the required information about the IP component or system to integrate and compile the component in the Quartus II software.
<system>.bsf	A Block Symbol File ( <b>.bsf</b> ) representation of the top-level Qsys system for use in Quartus II Block Diagram Files ( <b>.bdf</b> ).
<system>.spd	Input file for <code>ip-make-simscript</code> to generate simulation scripts for supported simulators. The <b>.spd</b> file contains a list of files generated for simulation, including memory initialization files.
<system>.ppf	The Pin Planner File ( <b>.ppf</b> ) stores the port and node assignments for IP components created for use with the Pin Planner.
<system>_bb.v	You can use the Verilog black-box ( <b>_bb.v</b> ) file as an empty module declaration for use as a black box.
<system>.sip	Contains information required for NativeLink simulation of IP components. You must add the <b>.sip</b> file to your Quartus II project.
<system>_inst.v or _inst.vhd	HDL example instantiation template. You can copy and paste the contents of this file into your HDL file to instantiate a Qsys system.

File Name	Description
<b>&lt;system&gt;.regmap</b>	If IP in the system contain register information, Qsys generates a <b>.regmap</b> file. The <b>.regmap</b> file describes the register map information of master and slave interfaces. This file complements the <b>.sopcinfo</b> file by providing more detailed register information about the system. This enables register display views and user customizable statistics in the System Console.
<b>&lt;system&gt;.svd</b>	Allows HPS System Debug tools to view the register maps of peripherals connected to HPS within a Qsys system.  During synthesis, the <b>.svd</b> files for slave interfaces visible to System Console masters are stored in the <b>.sof</b> file in the debug section. System Console reads this section, which Qsys can query for register map information. For system slaves, Qsys can access the registers by name.
<b>&lt;system&gt;.v</b> or <b>&lt;system&gt;.vhd</b>	HDL files that instantiate each submodule or child IP core in the system for synthesis or simulation.
<b>mentor/</b>	Contains a ModelSim script <b>msim_setup.tcl</b> to set up and run a simulation.
<b>aldec/</b>	Contains a Riviera-PRO script <b>rivierapro_setup.tcl</b> to setup and run a simulation.
<b>/synopsys/vcs</b> <b>/synopsys/vcsmx</b>	Contains a shell script <b>vcs_setup.sh</b> to set up and run a VCS simulation.  Contains a shell script <b>vcsmx_setup.sh</b> and <b>synopsys_sim.setup</b> file to set up and run a VCS MX simulation.
<b>/cadence</b>	Contains a shell script <b>ncsim_setup.sh</b> and other setup files to set up and run an NCSIM simulation.
<b>/submodules</b>	Contains HDL files for the submodule of the Qsys system.
<b>&lt;child IP cores&gt;/</b>	For each generated child IP core directory, Qsys generates <b>/synth</b> and <b>/sim</b> subdirectories.

#### Related Information

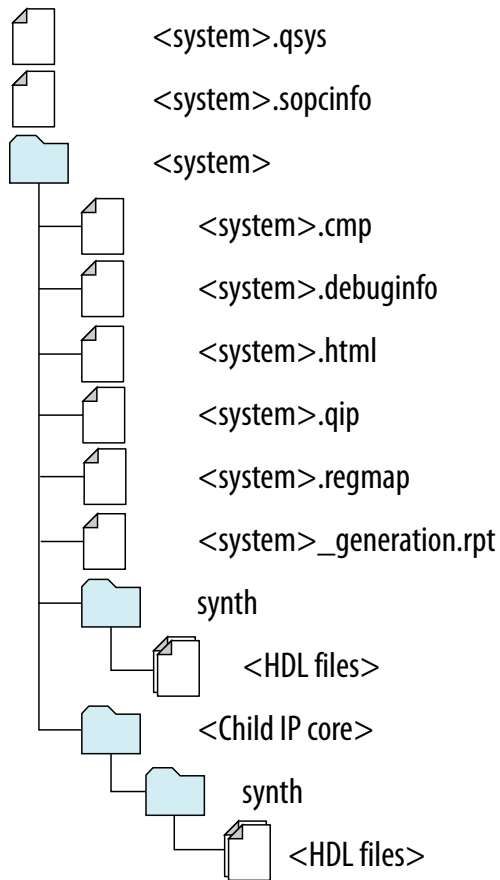
- [Qsys Synthesis Standard and Legacy Device Output Directories](#) on page 5-56
- [Qsys Simulation Standard and Legacy Device Output Directories](#) on page 5-57

### Qsys Synthesis Standard and Legacy Device Output Directories

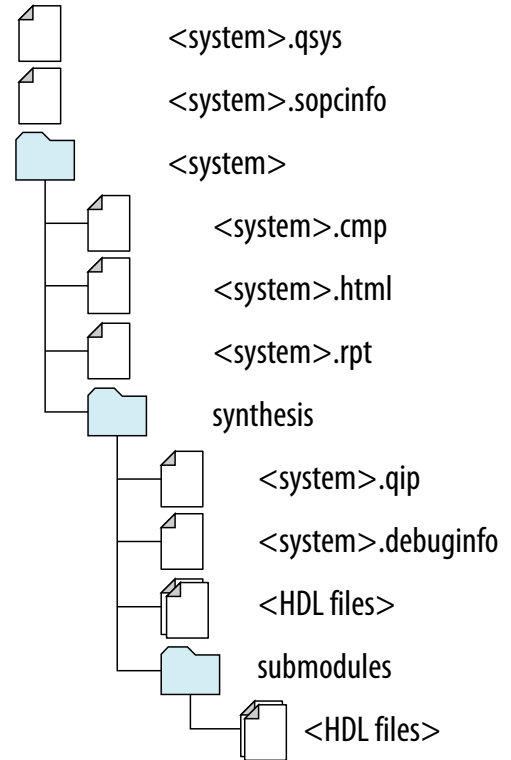
The **/synth** or **/synthesis** directories contain the Qsys-generated files that the Quartus II software uses to synthesize your design.

Figure 5-30: Qsys Synthesis Output Directories

Standard Directory Structure



Legacy Directory Structure



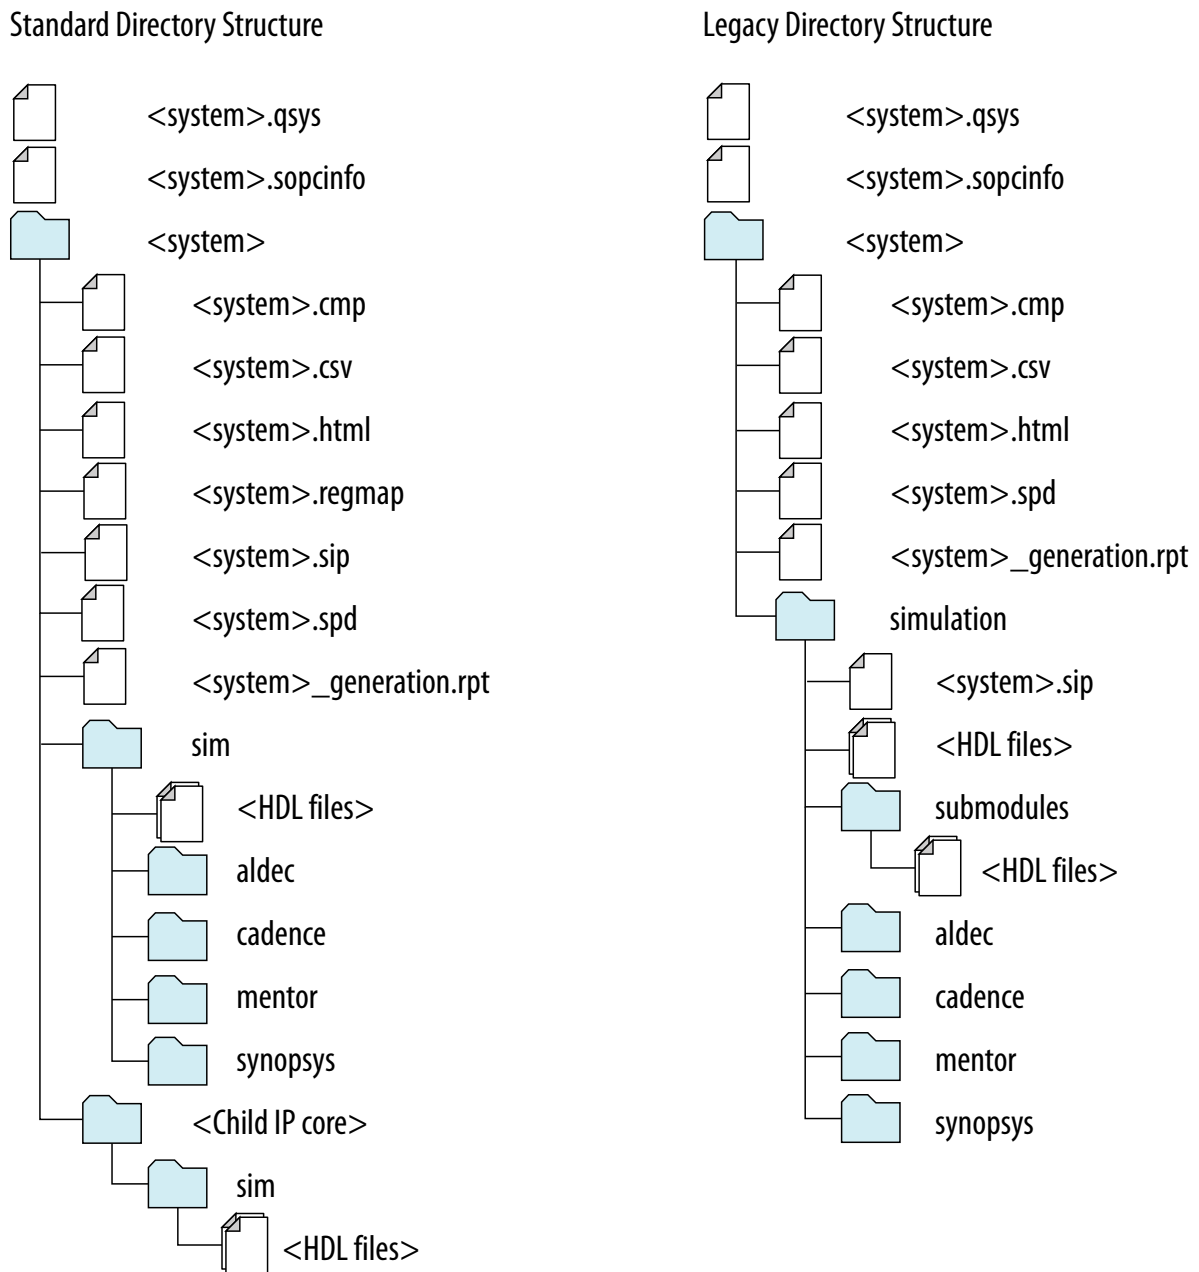
Related Information

[Files Generated for Qsys IP Components](#) on page 5-54

Qsys Simulation Standard and Legacy Device Output Directories

The `/sim` and `/simulation` directories contain the Qsys-generated output files to simulate your Qsys system.

Figure 5-31: Qsys Simulation Output Directories

**Related Information**

[Files Generated for Qsys IP Components](#) on page 5-54

**Generate Files for a Testbench Qsys System**

Qsys testbench is a new system that instantiates the current Qsys system by adding BFM's to drive the top-level interfaces. BFM's interact with the system in the simulator. You can use options in the **Generation** dialog box (**Generate > Generate Testbench System**) to generate a testbench Qsys system.

You can generate a standard or simple testbench system with BFM or Mentor Verification IP (for AXI3/AXI4) IP components that drive the external interfaces of your system. Qsys generates a Verilog HDL or VHDL simulation model for the testbench system to use in your simulation tool. You should first generate a testbench system, and then modify the testbench system in Qsys before generating its simulation model. In most cases, you should select only one of the simulation model options.

By default, the path of the generation output directory is fixed relative to the **.qsys** file. You can change the default directory in the **Generation** dialog box for legacy devices. For standard devices, the generation directory is fixed to the Qsys project directory.

The following options are available for generating a Qsys testbench system:

Option	Description
<b>Create testbench Qsys system</b>	<ul style="list-style-type: none"> <li>• <b>Standard, BFMs for standard Qsys Interconnect</b>—Creates a testbench Qsys system with BFM IP components attached to exported Avalon and AXI3/AXI4 interfaces. Includes any simulation partner modules specified by IP components in the system. The testbench generator supports AXI interfaces and can connect AXI3/AXI4 interfaces to Mentor Graphics AXI3/AXI4 master/slave BFMs. However, BFMs support address widths only up to 32-bits.</li> <li>• <b>Simple, BFMs for clocks and resets</b>—Creates a testbench Qsys system with BFM IP components driving only clock and reset interfaces. Includes any simulation partner modules specified by IP components in the system.</li> </ul>
<b>Create testbench simulation model</b>	Creates Verilog HDL or VHDL simulation model files and simulation scripts for the testbench Qsys system currently open in your workspace. Use this option if you do not need to modify the Qsys-generated testbench before running the simulation.
<b>Allow mixed-language simulation</b>	Generates a mixed-language simulation model. If you have a mixed-language simulator license, generating for mixed-language simulation can shorten the generation time, and produce files that can simulate faster. When turned off, all simulation files are generated in the selected simulation model language.

## Files Generated for Qsys Testbench

Table 5-10: Qsys-Generated Testbench Files

File Name or Directory Name	Description
<system>_tb.qsys	The Qsys testbench system.

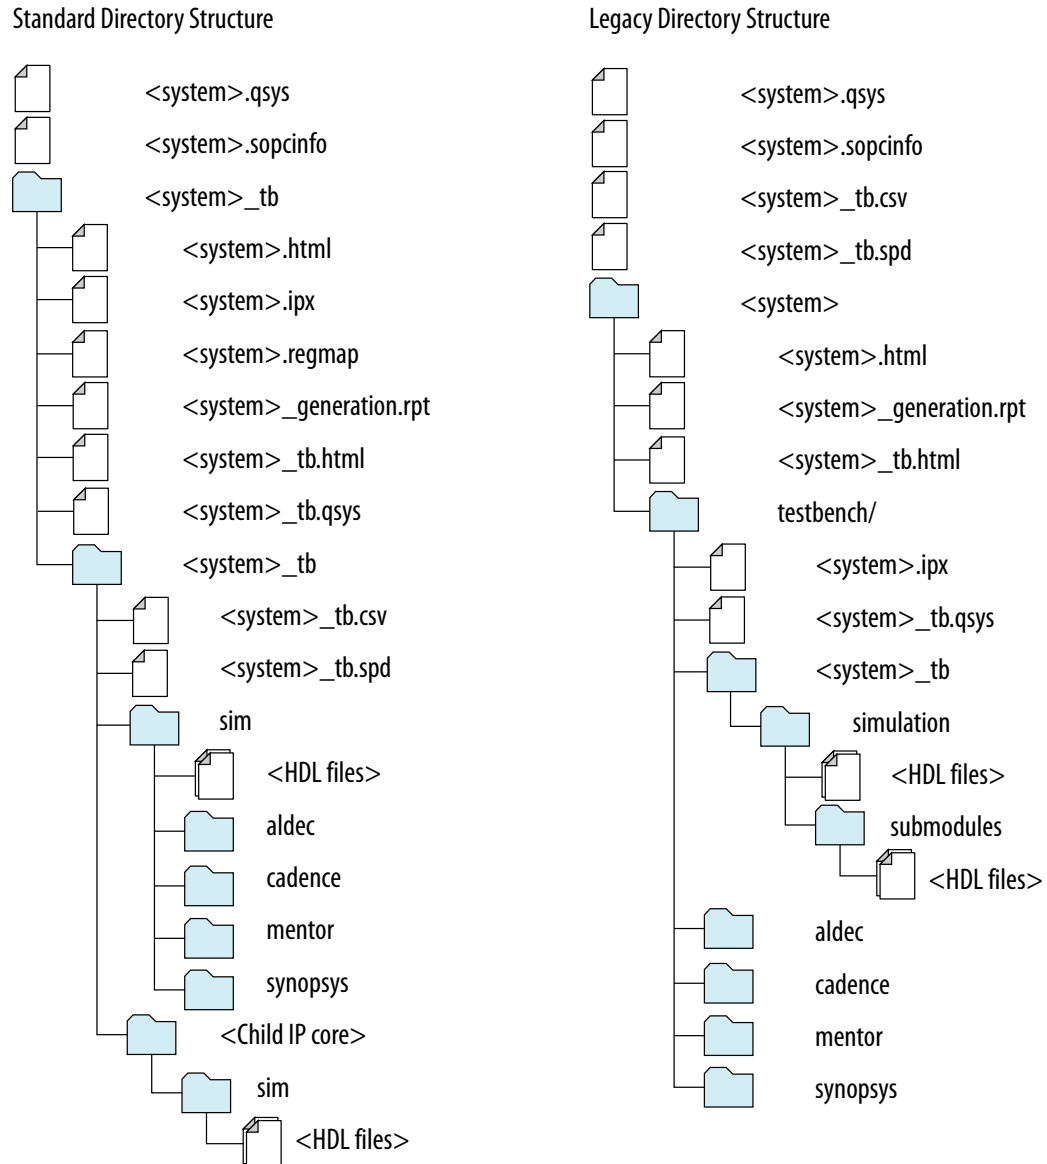
File Name or Directory Name	Description
<p><b>&lt;system&gt;_tb.v</b></p> <p>or</p> <p><b>&lt;system&gt;_tb.vhd</b></p>	The top-level testbench file that connects BFM's to the top-level interfaces of <b>&lt;system&gt;_tb.qsys</b> .
<b>&lt;system&gt;_tb.spd</b>	Required input file for <code>ip-make-simscript</code> to generate simulation scripts for supported simulators. The <b>.spd</b> file contains a list of files generated for simulation and information about memory that you can initialize.
<p><b>&lt;system&gt;.html</b></p> <p>and</p> <p><b>&lt;system&gt;_tb.html</b></p>	A system report that contains connection information, a memory map showing the address of each slave with respect to each master to which it is connected, and parameter assignments.
<b>&lt;system&gt;_generation.rpt</b>	Qsys generation log file. A summary of the messages that Qsys issues during testbench system generation.
<b>&lt;system&gt;.ipx</b>	The IP Index File ( <b>.ipx</b> ) lists the available IP components, or a reference to other directories to search for IP components.
<b>&lt;system&gt;.svd</b>	Allows HPS System Debug tools to view the register maps of peripherals connected to HPS within a Qsys system.  Similarly, during synthesis the <b>.svd</b> files for slave interfaces visible to System Console masters are stored in the <b>.sof</b> file in the debug section. System Console reads this section, which Qsys can query for register map information. For system slaves, Qsys can access the registers by name.
<b>mentor/</b>	Contains a ModelSim script <b>msim_setup.tcl</b> to set up and run a simulation
<b>aldec/</b>	Contains a Riviera-PRO script <b>rivierapro_setup.tcl</b> to setup and run a simulation.
<p><b>/synopsys/vcs</b></p> <p><b>/synopsys/vcsmx</b></p>	<p>Contains a shell script <b>vcs_setup.sh</b> to set up and run a VCS simulation.</p> <p>Contains a shell script <b>vcsmx_setup.sh</b> and <b>synopsys_sim.setup</b> file to set up and run a VCS MX simulation.</p>
<b>/cadence</b>	Contains a shell script <b>ncsim_setup.sh</b> and other setup files to set up and run an NCSIM simulation.
<b>/submodules</b>	Contains HDL files for the submodule of the Qsys testbench system.
<b>&lt;child IP cores&gt;/</b>	For each generated child IP core directory, Qsys testbench generates <b>/synth</b> and <b>/sim</b> subdirectories.



## Qsys Testbench Simulation Standard and Legacy Device Output Directories

The `/sim` and `/simulation` directories contain the Qsys-generated output files to simulate your Qsys testbench system.

Figure 5-32: Qsys Simulation Testbench Directory Structure



### Generate and Modify a Qsys Testbench System

You can use the following steps to create a Qsys testbench system of your Qsys system.

1. Create a Qsys system.
2. Generate a testbench system in the Qsys **Generation** dialog box (**Generate** > **Generate Testbench System**).
3. Open the testbench system in Qsys. Make changes to the BFM, as needed, such as changing the instance names and **VHDL ID** value. For example, you can modify the **VHDL ID** value in the **Altera Avalon Interrupt Source** IP component.
4. If you modify a BFM, re-generate the simulation model for the testbench system.
5. Create a custom test program for the BFM.
6. Compile and load the Qsys system and testbench into your simulator, and then run the simulation.

## Simulation Scripts

Qsys generates simulation scripts to set up the simulation environment for Mentor Graphics Modelsim<sup>®</sup> and Questasim<sup>®</sup>, Synopsys VCS and VCS MX, Cadence Incisive Enterprise Simulator<sup>®</sup> (NCSIM), and the Aldec Riviera-PRO<sup>®</sup> Simulator.

You can use the scripts to compile the required device libraries and system design files in the correct order and elaborate or load the top-level system for simulation.

**Table 5-11: Simulation Script Variables**

The simulation scripts provide variables that allow flexibility in your simulation environment.

Variable	Description
TOP_LEVEL_NAME	If the testbench Qsys system is not the top-level instance in your simulation environment because you instantiate the Qsys testbench within your own top-level simulation file, set the TOP_LEVEL_NAME variable to the top-level hierarchy name.
QSYS_SIMDIR	If the simulation files generated by Qsys are not in the simulation working directory, use the QSYS_SIMDIR variable to specify the directory location of the Qsys simulation files.
QUARTUS_INSTALL_DIR	Points to the Quartus installation directory that contains the device family library.

### Example 5-4: Top-Level Simulation HDL File for a Testbench System

The example below shows the `pattern_generator_tb` generated for a Qsys system called `pattern_generator`. The `top.sv` file defines the top-level module that instantiates the `pattern_generator_tb` simulation model, as well as a custom SystemVerilog test program with BFM transactions, called `test_program`.

```
module top();
  pattern_generator_tb tb();
  test_program pgm();
endmodule
```

**Note:** The VHDL version of the Altera Tristate Conduit BFM is not supported in Synopsys VCS, NCSim, and Riviera-PRO in the Quartus II software version 14.0. These simulators do not support the VHDL protected type, which is used to implement the BFM. For a workaround, use a simulator that supports the VHDL protected type.

### Related Information

- [ModelSim-Altera software, Mentor Graphics ModelSim support](#)
- [Synopsys VCS and VCS MX support](#)
- [Cadence Incisive Enterprise Simulator \(IES\) support](#)
- [Aldec Active-HDL and Rivera-PRO support](#)

## Simulating Software Running on a Nios II Processor

To simulate the software in a system driven by a Nios II processor, generate the simulation model for the Qsys testbench system with the following steps:

1. In the **Generation** dialog box (**Generate > Generate Testbench System**), select **Simple, BFM's for clocks and resets**.
2. For the **Create testbench simulation model** option select **Verilog** or **VHDL**.
3. Click **Generate**.
4. Open the **Nios II Software Build Tools for Eclipse**.
5. Set up an application project and board support package (BSP) for the `<system>` **.sopcinfo** file.
6. To simulate, right-click the application project in Eclipse, and then click **Run as > Nios II ModelSim**. Sets up the ModelSim simulation environment, and compiles and loads the Nios II software simulation.
7. To run the simulation in ModelSim, type `run -all` in the ModelSim transcript window.
8. Set the ModelSim settings and select the Qsys Testbench Simulation Package Descriptor (**.spd**) file, `<system > _tb.spd`. The **.spd** file is generated with the testbench simulation model for Nios II designs and specifies the files required for Nios II simulation.

### Related Information

- [Getting Started with the Graphical User Interface \(Nios II\)](#)
- [Getting Started from the Command-Line \(Nios II\)](#)

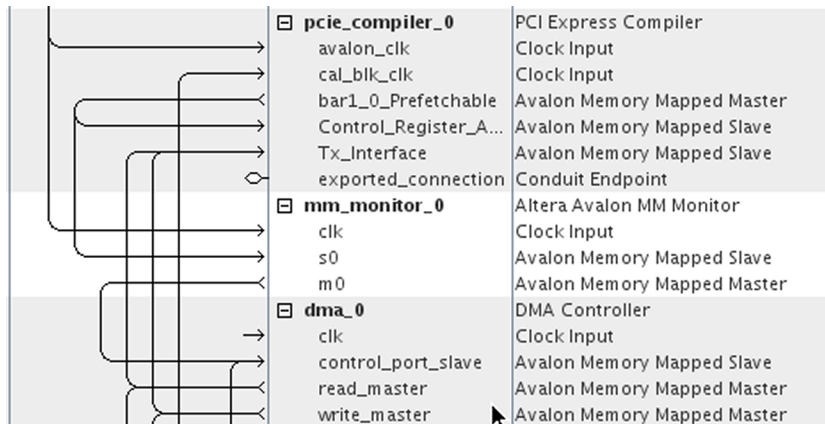
## Add Assertion Monitors for Simulation

You can add monitors to Avalon-MM, AXI, and Avalon-ST interfaces in your system to verify protocol and test coverage with a simulator that supports SystemVerilog assertions.

**Note:** Modelsim Altera Edition does not support SystemVerilog assertions. If you want to use assertion monitors, you must use a supported third-party simulators such as Mentor Questasim, Synopsys VCS, or Cadence Incisive. For more information, refer to *Introduction to Altera IP Cores*.

**Figure 5-33: Inserting an Avalon-MM Monitor Between an Avalon-MM Master and Slave Interface**

This example demonstrates the use of a monitor with an Avalon-MM monitor between the `pcie_compiler_bar1_0_Prefetchable` Avalon-MM master interface, and the `dma_0_control_port_slave` Avalon-MM slave interface.



Similarly, you can insert an Avalon-ST monitor between Avalon-ST source and sink interfaces.

#### Related Information

[Introduction to Altera IP Cores](#)

## CMSIS Support for the HPS IP Component

Qsys systems that contain an HPS IP component generate a System View Description (`.svd`) file that lists peripherals connected to the ARM processor.

The `.svd` (or CMSIS-SVD) file format is an XML schema specified as part of the Cortex Microcontroller Software Interface Standard (CMSIS) provided by ARM. The `.svd` file allows HPS system debug tools (such as the DS-5 Debugger) to view the register maps of peripherals connected to HPS in a Qsys system.

#### Related Information

[Component Interface Tcl Reference](#) on page 9-1

[CMSIS - Cortex Microcontroller Software](#)

## Explore and Manage Qsys Interconnect

The System with Qsys Interconnect window allows you to see the contents of the Qsys interconnect before you generate your system. In this display of your system, you can review a graphical representation of the generated interconnect. Qsys converts connections between interfaces to interconnect logic during system generation.

You access the System with Qsys Interconnect window by clicking **Show System With Qsys Interconnect** command on the System menu.

The System with Qsys Interconnect window has the following tabs:

- **System Contents**—Displays the original instances in your system, as well as the inserted interconnect instances. Connections between interfaces are replaced by connections to interconnect where applicable.
- **Hierarchy**—Displays a system hierarchical navigator, expanding the system contents to show modules, interfaces, signals, contents of subsystems, and connections.
- **Parameters**—Displays the parameters for the selected element in the **Hierarchy** tab.
- **Memory-Mapped Interconnect**—Allows you to select a memory-mapped interconnect module and view its internal command and response networks. You can also insert pipeline stages to achieve timing closure.

The **System Contents**, **Hierarchy**, and **Parameters** tabs are read-only. Edits that you apply on the **Memory-Mapped Interconnect** tab are automatically reflected on the **Interconnect Requirements** tab.

The **Memory-Mapped Interconnect** tab in the System with Qsys Interconnect window displays a graphical representation of command and response data paths in your system. Data paths allow you precise control over pipelining in the interconnect. Qsys displays separate figures for the command and response data paths. You can access the data paths by clicking their respective tabs in the **Memory-Mapped Interconnect** tab.

Each node element in a figure represents either a master or slave that communicates over the interconnect, or an interconnect sub-module. Each edge is an abstraction of connectivity between elements, and its direction represents the flow of the commands or responses.

Click **Highlight Mode (Path, Successors, Predecessors)** to identify edges and data paths between modules. Turn on **Show Pipeline Locations** to add greyed-out registers on edges where pipelining is allowed in the interconnect.

**Note:** You must select more than one module to highlight a path.

## Manually Controlling Pipelining in the Qsys Interconnect

The **Memory-Mapped Interconnect** tab allows you to manipulate pipeline connections in the Qsys interconnect. You access the **Memory-Mapped Interconnect** tab by clicking the **Show System With Qsys Interconnect** command on the System menu.

**Note:** To increase interconnect frequency, you should first try increasing the value of the **Limit interconnect pipeline stages to** option on the **Interconnect Requirements** tab. You should only consider manually pipelining the interconnect if changes to this option do not improve frequency, and you have tried all other options to achieve timing closure, including the use of a bridge. Manually pipelining the interconnect should only be applied to complete systems.

1. In the **Interconnect Requirements** tab, first try increasing the value of the **Limit interconnect pipeline stages to** option until it no longer gives significant improvements in frequency, or until it causes unacceptable effects on other parts of the system.
2. In the Quartus II software, compile your design and run timing analysis.
3. Using the timing report, identify the critical path through the interconnect and determine the approximate mid-point. The following is an example of a timing report:

```
2.800 0.000 cpu_instruction_master|out_shifter[63]|q
3.004 0.204 mm_domain_0|addr_router_001|Equal5~0|datac
3.246 0.242 mm_domain_0|addr_router_001|Equal5~0|combout
3.346 0.100 mm_domain_0|addr_router_001|Equal5~1|dataa
3.685 0.339 mm_domain_0|addr_router_001|Equal5~1|combout
```

```
4.153 0.468 mm_domain_0|addr_router_001|src_channel[5]~0|datad
4.373 0.220 mm_domain_0|addr_router_001|src_channel[5]~0|combout
```

4. In Qsys, click **System** > **Show System With Qsys Interconnect**.
5. In the **Memory-Mapped Interconnect** tab, select the interconnect module that contains the critical path. You can determine the name of the module from the hierarchical node names in the timing report.
6. Click **Show Pipelinable Locations**. Qsys display all possible pipeline locations in the interconnect. Right-click the possible pipeline location to insert or remove a pipeline stage.
7. Locate the possible pipeline location that is closest to the mid-point of the critical path. The names of the blocks in the memory-mapped interconnect tab correspond to the module instance names in the timing report.
8. Right-click the location where you want to insert a pipeline, and then click **Insert Pipeline**.
9. Regenerate the Qsys system, recompile the design, and then rerun timing analysis. If necessary, repeat the manual pipelining process again until timing requirements are met.

Manual pipelining has the following limitations:

- If you make changes to your original system's connectivity after manually pipelining an interconnect, your inserted pipelines may become invalid. Qsys displays warning messages when you generate your system if invalid pipeline stages are detected. You can remove invalid pipeline stages with the **Remove Stale Pipelines** option in the **Memory-Mapped Interconnect** tab. Altera recommends that you do not make changes to the system's connectivity after manual pipeline insertion.
- Review manually-inserted pipelines when upgrading to newer versions of Qsys. Manually-inserted pipelines in one version of Qsys may not be valid in a future version.

#### Related Information

[Specify Qsys \\$system Interconnect Requirements](#)

[Qsys System Design Components](#) on page 10-1

## Implement Performance Monitoring

You use the Qsys **Instrumentation** tab in to set up real-time performance monitoring using throughput metrics such as read and write transfers. The **Add debug instrumentation to the Qsys Interconnect** option allows you to interact with the Bus Analyzer Toolkit, which you can access on the Tools menu in the Quartus II software.

Qsys supports performance monitoring for only Avalon-MM interfaces. In your Qsys system, you can monitor the performance of no less than three, and no greater than 15 components at one time. The performance monitoring feature works with Quartus II software devices 13.1 and newer.

**Note:** For more information about the Bus Analyzer Toolkit and the Qsys **Instrumentation** tab, refer to the **Bus Analyzer Toolkit** page.

#### Related Information

[Bus Analyzer Toolkit](#)

## Qsys 64-Bit Addressing Support

Qsys interconnect supports up to 64-bit addressing for all Qsys interfaces and IP components, with a range of: 0x0000 0000 0000 0000 to 0xFFFF FFFF FFFF FFFF, inclusive.

Address parameters appear in the **Base** and **End** columns in the **System Contents** tab, on the **Address Map** tab, in the parameter editor, and in validation messages. Qsys displays as many digits as needed in order to display the top-most set bit, for example, 12 hex digits for a 48-bit address.

A Qsys system can have multiple 64-bit masters, with each master having its own address space. You can share slaves between masters and masters can map slaves to different addresses. For example, one master can interact with slave 0 at base address 0000\_0000\_0000, and another master can see the same slave at base address c000\_000\_000.

Quartus II debugging tools provide access to the state of an addressable system via the Avalon-MM interconnect. These are also 64-bit compatible, and process within a 64-bit address space, including a JTAG to Avalon master bridge.

For more information on design practices when using slaves with large address spans, refer to *Address Span Extender* in volume 1 of the *Quartus II Handbook*.

### Related Information

- [Qsys System Design Components](#) on page 10-1

## Support for Avalon-MM Non-Power of Two Data Widths

Qsys requires that you connect all multi-point Avalon-MM connections to interfaces with data widths that are equal to powers of two.

Qsys issues a validation error if an Avalon-MM master or slave interface on a multi-point connection is parameterized with a non-power of two data width.

**Note:** Avalon-MM point-to-point connections between an Avalon-MM master and an Avalon-MM slave are an exception and may set their data widths to a non-power of two.

## View the Qsys HDL Example

Click **Generate > HDL Example** to generate a template for the top-level HDL definition of your Qsys system in either Verilog HDL or VHDL. The HDL template displays the IP component declaration.

You can copy and paste the example into a top-level HDL file that instantiates the Qsys system, if the Qsys system is not the top-level module in your Quartus II project.

## Qsys System Example Designs

Click the **Example Design** button in the parameter editor to generate an example design.

If there are multiple example designs for an IP component, then there is a button for each example in the parameter editor. When you click the **Example Design** button, the **Select Example Design Directory** dialog box appears, where you can select the directory to save the example design.

The **Example Design** button does not appear in the parameter editor if there is no example. For some IP components, you can click **Generate > Example Designs** to access an example design.

The following Qsys system example designs demonstrate various design features and flows that you can replicate in your Qsys system.

#### Related Information

- [NIOS II Qsys Example Design](#)
- [PCI Express Avalon-ST Qsys Example Design](#)
- [Triple Speed Ethernet Qsys Example Design](#)

## Qsys Command-Line Utilities

You can perform many of the functions available in the Qsys GUI from the command-line with the `qsys-edit`, `qsys-generate`, and `qsys-script` utilities.

You run Qsys command-line executables from the Quartus II installation directory:

```
<Quartus II installation directory>\quartus\sopc_builder\bin
```

You can use `qsys-generate` to generate a Qsys system or IP variation outside of the Qsys GUI. You can use `qsys-script` to create, manipulate or manage a Qsys system with command-line scripting.

For command-line help listing all options for these executables, type the following command:

```
<Quartus II installation directory>\quartus\sopc_builder\bin\<executable name> --help
```

### Example 5-5: Qsys Command-Line Scripting Example

```
qsys-script --script=my_script.tcl \  
--system-file=fancy.qsys my_script.tcl contains:  
package require -exact qsys 13.1  
# get all instance names in the system and print one by one  
set instances [ get_instances ]  
foreach instance $instances {  
    send_message Info "$instance"  
}
```

**Note:** You must add `$QUARTUS_ROOTDIR/sopc_builder/bin/` to the `PATH` variable to access command-line utilities. Once you add this `PATH` variable, you can launch the utilities from any directory location.

#### Related Information

- [Working with Instance Parameters in Qsys](#)
- [Altera Wiki Qsys Scripts](#)

## Run the Qsys Editor with `qsys-edit`

You can use the `qsys-edit` utility to run the Qsys editor from the command-line.



You can use the following options with the `qsys-edit` utility:

**Table 5-12: qsys-edit Command-Line Options**

Option	Usage	Description
<i>1st arg file</i>	Optional	The name of the <code>.qsys</code> system or <code>.qvar</code> variation file to edit.
<code>--search-path[=&lt;value&gt;]</code>	Optional	If omitted, Qsys uses a standard default path. If provided, Qsys searches a comma-separated list of paths. To include the standard path in your replacement, use "\$", for example: <code>/extra/dir.\$</code> .
<code>--project-directory=&lt;directory&gt;</code>	Optional	Allows you to find IP components in certain locations relative to the project, if any. By default, the current directory is: ' '. To exclude any project directory, use ' '.
<code>--new-component-type=&lt;value&gt;</code>	Optional	Allows you to specify the kind of instance that is parameterized in a variation.
<code>--debug</code>	Optional	Enables debugging features and output.
<code>--host-controller</code>	Optional	Launches the application with an XML host controller interface on standard input/output.
<code>--jvm-max-heap-size=&lt;value&gt;</code>	Optional	The maximum memory size Qsys uses for allocations when running <code>qsys-edit</code> . You specify this value as <code>&lt;size&gt;&lt;unit&gt;</code> , where unit is <code>m</code> (or <code>M</code> ) for multiples of megabytes, or <code>g</code> (or <code>G</code> ) for multiples of gigabytes. The default value is 512m.
<code>--help</code>	Optional	Display help for <code>qsys-edit</code> .

## Generate a Qsys System or IP Variation with `qsys-generate`

You can use the `qsys-generate` utility to generate RTL for your Qsys system, an IP core variation, or simulation models and scripts. You can create testbench systems for testing your Qsys system in a simulator using bus functional models (BFMs). Output from the `qsys-generate` command is the same as when generating using the Qsys GUI.

**Table 5-13: qsys-generate Command-Line Options**

Option	Usage	Description
<i>&lt;1st arg file&gt;</i>	Required	The name of the <code>.qsys</code> system file to generate.

Option	Usage	Description
<code>--synthesis=&lt;VERILOG VHDL&gt;</code>	Optional	Creates synthesis HDL files that Qsys uses to compile the system in a Quartus II project. You must specify the preferred generation language for the top-level RTL file for the generated Qsys system.
<code>--block-symbol-file</code>	Optional	Creates a Block Symbol File ( <b>.bsf</b> ) for the Qsys system.
<code>--simulation=&lt;VERILOG VHDL&gt;</code>	Optional	Creates a simulation model for the Qsys system. The simulation model contains generated HDL files for the simulator, and may include simulation-only features. You must specify the preferred simulation language.
<code>--testbench=&lt;SIMPLE STANDARD&gt;</code>	Optional	Creates a testbench system that instantiates the original system, adding bus functional models (BFMs) to drive the top-level interfaces. When you generate the system, the BFMs interact with the system in the simulator.
<code>--testbench-simulation=&lt;VERILOG VHDL&gt;</code>	Optional	After you create the testbench system, you can create a simulation model for the testbench system.
<code>--search-path=&lt;value&gt;</code>	Optional	If you omit this command, Qsys uses a standard default path. If you provide this command, Qsys searches a comma-separated list of paths. To include the standard path in your replacement, use "\$", for example, "/extra/dir,\$".
<code>--jvm-max-heap-size=&lt;value&gt;</code>	Optional	The maximum memory size that Qsys uses for allocations when running <code>qsys-generate</code> . You specify the value as <code>&lt;size&gt; &lt;unit&gt;</code> , where <code>unit</code> is <code>m</code> (or <code>M</code> ) for multiples of megabytes or <code>g</code> (or <code>G</code> ) for multiples of gigabytes. The default value is 512m.
<code>--family=&lt;value&gt;</code>	Optional	Specifies the device family.
<code>--part=&lt;value&gt;</code>	Optional	Specifies the device part number. If set, this option overrides the <code>--family</code> option.

Option	Usage	Description
<code>--allow-mixed-language-simulation</code>	Optional	Enables a mixed language simulation model generation. If true, if a preferred simulation language is set, Qsys uses a <code>fileset</code> of the component for the simulation model generation. When false, which is the default, Qsys uses the language specified with <code>--file-set=&lt;value&gt;</code> for all components for simulation model generation.

## Display Available IP Components with ip-catalog

The `ip-catalog` command displays a list of available IP components relative to the current Quartus II project directory. Use the following format for the `ip-catalog` command:

```
ip-catalog
    [--project-dir=<directory>]
    [--name=<value>]
    [--verbose]
    [--xml]
    [--help]
```

**Table 5-14: ip-catalog Command-Line Options**

Option	Usage	Description
<code>--project-dir= &lt;directory&gt;</code>	Optional	Finds IP components relative to the Quartus II project directory. By default, Qsys uses '.' as the current directory. To exclude a project directory, leave the value empty.
<code>--name=&lt;value&gt;</code>	Optional	Provides a pattern to filter the names of the IP components found. To show all IP components, use a '*' or '.'. By default, Qsys shows all IP components. The argument is not case sensitive.
<code>--verbose</code>	Optional	Reports the progress of the command.
<code>--xml</code>	Optional	Generates the output in XML format, in place of colon-delimited format.
<code>--help</code>	Optional	Displays help for the <code>ip-catalog</code> command.

## Create an .ipx File with ip-make-ipx

The `ip-make-ipx` command creates an **.ipx** file and is a convenient way to include a collection of IP components from an arbitrary directory. You can edit the **.ipx** file to disable visibility of one or more IP components in the IP Catalog. Use the following format for the `ip-make-ipx` command:

```
ip-make-ipx
  [--source-directory=<directory>]
  [--output=<file>]
  [--relative-vars=<value>]
  [--thorough-descent]
  [--message-before=<value>]
  [--message-after=<value>]
  [--help]
```

**Table 5-15: ip-make-ipx Command-Line Options**

Option	Usage	Description
<code>--source-directory=&lt;directory&gt;</code>	Optional	Specifies the root directory for IP component files. The default directory is <code>.*</code> . You can provide a comma-separated list of directories.
<code>--output=&lt;file&gt;</code>	Optional	Specifies the name of the file to generate. The default name is <code>/component.ipx</code> .
<code>--relative-vars=&lt;value&gt;</code>	Optional	Causes the output file to include references relative to the specified variable(s) where possible. You can specify multiple variables as a comma-separated list.
<code>--thorough-descent</code>	Optional	If set, a component or <b>.ipx</b> file in a directory does not stop Qsys from searching subdirectories.
<code>--message-before=&lt;value&gt;</code>	Optional	Prints a message: <code>stdout</code> when indexing begins.
<code>--message-after=&lt;value&gt;</code>	Optional	Sends a message: <code>stdout</code> when indexing is done.
<code>--help</code>	Optional	Displays help for the <code>ip-make-ipx</code> command.

## Generate a Qsys System with qsys-script

You can use the `qsys-script` utility to create and manipulate a Qsys system with Tcl scripting commands.

**Note:** You must provide a package version for the `qsys-script`. If you do not specify the `--package-version=<value>` command, you must then provide a Tcl script and request the system scripting API directly with the `package require -exact qsys <version>` command.

**Table 5-16: qsys-script Command-Line Options**

Option	Usage	Description
<code>--system-file=&lt;file&gt;</code>	Optional	Specifies the path to a <b>.qsys</b> file. Qsys loads the system before running scripting commands.
<code>--script=&lt;file&gt;</code>	Optional	A file that contains Tcl scripting commands that you can use to create or manipulate Qsys systems. If you specify both <code>--cmd</code> and <code>--script</code> , Qsys runs the <code>--cmd</code> commands before the script specified by <code>--script</code> .
<code>--cmd=&lt;value&gt;</code>	Optional	A string that contains Tcl scripting commands that you can use to create or manipulate a Qsys system. If you specify both <code>--cmd</code> and <code>--script</code> , Qsys runs the <code>--cmd</code> commands before the script specified by <code>--script</code> .
<code>--package-version=&lt;value&gt;</code>	Optional	Specifies which Tcl API scripting version to use and determines the functionality and behavior of the Tcl commands. The Quartus II software supports Tcl API scripting commands. If you do not specify the version on the command-line, your script must request the scripting API directly with the <code>package require -exact qsys &lt;version&gt;</code> command.
<code>--search-path=&lt;value&gt;</code>	Optional	If you omit this command, a Qsys uses a standard default path. If you provide this command, Qsys searches a comma-separated list of paths. To include the standard path in your replacement, use "\$", for example, <code>/&lt;directory path&gt;/dir,\$</code> . Separate multiple directory references with a comma.
<code>--jvm-max-heap-size=&lt;value&gt;</code>	Optional	The maximum memory size that the <code>qsys-script</code> tool uses. You specify this value as <code>&lt;size&gt;&lt;unit&gt;</code> , where unit is <code>m</code> (or <code>M</code> ) for multiples of megabytes, or <code>g</code> (or <code>G</code> ) for multiples of gigabytes.
<code>--help</code>	Optional	Displays help for the <code>qsys-script</code> utility.

## Qsys Scripting Command Reference

The following are Qsys scripting commands:

- [add\\_connection](#) on page 5-77
- [add\\_instance](#) on page 5-78
- [add\\_interface](#) on page 5-79
- [apply\\_preset](#) on page 5-80
- [auto\\_assign\\_base\\_addresses](#) on page 5-81
- [auto\\_assign\\_system\\_base\\_addresses](#) on page 5-82
- [auto\\_assign\\_irqs](#) on page 5-83
- [auto\\_connect](#) on page 5-84
- [create\\_system](#) on page 5-85
- [export\\_hw\\_tcl](#) on page 5-86
- [get\\_composed\\_connection\\_parameter\\_value](#) on page 5-87
- [get\\_composed\\_connection\\_parameters](#) on page 5-88
- [get\\_composed\\_connections](#) on page 5-89
- [get\\_composed\\_instance\\_assignment](#) on page 5-90
- [get\\_composed\\_instance\\_assignments](#) on page 5-91
- [get\\_composed\\_instance\\_parameter\\_value](#) on page 5-92
- [get\\_composed\\_instance\\_parameters](#) on page 5-93
- [get\\_composed\\_instances](#) on page 5-94
- [get\\_connection\\_parameter\\_property](#) on page 5-95
- [get\\_connection\\_parameter\\_value](#) on page 5-96
- [get\\_connection\\_parameters](#) on page 5-97
- [get\\_connection\\_properties](#) on page 5-98
- [get\\_connection\\_property](#) on page 5-99
- [get\\_connections](#) on page 5-100
- [get\\_instance\\_assignment](#) on page 5-101
- [get\\_instance\\_assignments](#) on page 5-102
- [get\\_instance\\_documentation\\_links](#) on page 5-103
- [get\\_instance\\_interface\\_assignment](#) on page 5-104
- [get\\_instance\\_interface\\_assignments](#) on page 5-105
- [get\\_instance\\_interface\\_parameter\\_property](#) on page 5-106
- [get\\_instance\\_interface\\_parameter\\_value](#) on page 5-107
- [get\\_instance\\_interface\\_parameters](#) on page 5-108

[get\\_instance\\_interface\\_port\\_property](#) on page 5-109

[get\\_instance\\_interface\\_ports](#) on page 5-110

[get\\_instance\\_interface\\_properties](#) on page 5-111

[get\\_instance\\_interface\\_property](#) on page 5-112

[get\\_instance\\_interfaces](#) on page 5-113

[get\\_instance\\_parameter\\_property](#) on page 5-114

[get\\_instance\\_parameter\\_value](#) on page 5-115

[get\\_instance\\_parameters](#) on page 5-116

[get\\_instance\\_port\\_property](#) on page 5-117

[get\\_instance\\_properties](#) on page 5-118

[get\\_instance\\_property](#) on page 5-119

[get\\_instances](#) on page 5-120

[get\\_interconnect\\_requirement](#) on page 5-121

[get\\_interconnect\\_requirements](#) on page 5-122

[get\\_interface\\_port\\_property](#) on page 5-123

[get\\_interface\\_ports](#) on page 5-124

[get\\_interface\\_properties](#) on page 5-125

[get\\_interface\\_property](#) on page 5-126

[get\\_interfaces](#) on page 5-127

[get\\_module\\_properties](#) on page 5-128

[get\\_module\\_property](#) on page 5-129

[get\\_parameter\\_properties](#) on page 5-130

[get\\_port\\_properties](#) on page 5-131

[get\\_project\\_properties](#) on page 5-132

[get\\_project\\_property](#) on page 5-133

[load\\_system](#) on page 5-134

[lock\\_avalon\\_base\\_address](#) on page 5-135

[remove\\_connection](#) on page 5-136

[remove\\_dangling\\_connections](#) on page 5-137

[remove\\_instance](#) on page 5-138

[remove\\_interface](#) on page 5-139

[save\\_system](#) on page 5-140

[send\\_message](#) on page 5-141

[set\\_connection\\_parameter\\_value](#) on page 5-142

[set\\_instance\\_parameter\\_value](#) on page 5-143

[set\\_instance\\_property](#) on page 5-144

[set\\_interconnect\\_requirement](#) on page 5-145

[set\\_interface\\_property](#) on page 5-146

[set\\_module\\_property](#) on page 5-147

[set\\_project\\_property](#) on page 5-148

[set\\_use\\_testbench\\_naming\\_pattern](#) on page 5-149

[set\\_validation\\_property](#) on page 5-150

[unlock\\_avalon\\_base\\_address](#) on page 5-151

[validate\\_connection](#) on page 5-152

[validate\\_instance](#) on page 5-153

[validate\\_instance\\_interface](#) on page 5-154

[validate\\_system](#) on page 5-155



## add\_connection

### Description

Connects the named interfaces using an appropriate connection type. Both interface names consist of a child instance name, followed by the name of an interface provided by that module. For example, `mux0.out` is the interface named on the instance named `mux0`. Be careful to connect the start to the end, and not the reverse.

### Usage

```
add_connection <start> [<end>]
```

### Returns

No return value.

### Arguments

#### start

The start interface that is connected, in `<instance_name>.<interface_name>` format. If the `end` argument is omitted, the connection must be of the form `<instance1>.<interface>/<instance2>.<interface>`.

#### end (optional)

The end interface that is connected, `<instance_name>.<interface_name>`.

### Example

```
add_connection dma.read_master sdram.s1
```

### Related Information

- [get\\_connection\\_parameter\\_value](#) on page 5-96
- [get\\_connection\\_property](#) on page 5-99
- [get\\_connections](#) on page 5-100
- [remove\\_connection](#) on page 5-136
- [set\\_connection\\_parameter\\_value](#) on page 5-142

## add\_instance

### Description

Adds an instance of a component, referred to as a *child* or *child instance*, to the system.

### Usage

```
add_instance <name> <type> [<version>]
```

### Returns

No return value.

### Arguments

**name**

Specifies a unique local name that you can use to manipulate the instance. Qsys uses this name in the generated HDL to identify the instance.

**type**

Refers to a kind of instance available in the IP Catalog, for example `altera_avalon_uart`.

**version (optional)**

The required version of the specified instance type. If no version is specified, Qsys uses the latest version.

### Example

```
add_instance uart_0 altera_avalon_uart
```

### Related Information

- [get\\_instance\\_parameter\\_value](#) on page 5-115
- [get\\_instance\\_property](#) on page 5-119
- [get\\_instances](#) on page 5-120
- [remove\\_instance](#) on page 5-138
- [set\\_instance\\_parameter\\_value](#) on page 5-143

## add\_interface

### Description

Adds an interface to your system, which Qsys uses to export an interface from within the system. You specify the exported internal interface with `set_interface_property <interface> EXPORT_OF instance.interface`.

### Usage

```
add_interface <name> <type> <direction>.
```

### Returns

No return value.

### Arguments

**name**

The name of the interface that Qsys exports from the system.

**type**

The type of interface.

**direction**

The interface direction.

### Example

```
add_interface my_export conduit end  
set_interface_property my_export EXPORT_OF uart_0.external_connection
```

### Related Information

- [get\\_interface\\_ports](#) on page 5-124
- [get\\_interface\\_properties](#) on page 5-125
- [get\\_interface\\_property](#) on page 5-126
- [set\\_interface\\_property](#) on page 5-146

## `apply_preset`

### Description

Applies the settings in a preset to the specified instance.

### Usage

```
apply_preset <instance> <preset_name>
```

### Returns

No return value.

### Arguments

**instance**

The name of the instance.

**preset\_name**

The name of the preset.

### Example

```
apply_preset cpu_0 "Custom Debug Settings"
```

## auto\_assign\_base\_addresses

### Description

Assigns base addresses to all memory mapped interfaces on an instance in the system. Instance interfaces that are locked with `lock_avalon_base_address` keep their addresses during address auto-assignment.

### Usage

```
auto_assign_base_addresses <instance>
```

### Returns

No return value.

### Arguments

**instance**

The name of the instance with memory mapped interfaces.

### Example

```
auto_assign_base_addresses sdram
```

### Related Information

- [auto\\_assign\\_system\\_base\\_addresses](#) on page 5-82
- [lock\\_avalon\\_base\\_address](#) on page 5-135
- [unlock\\_avalon\\_base\\_address](#) on page 5-151

## auto\_assign\_system\_base\_addresses

### Description

Assigns legal base addresses to all memory mapped interfaces on all instances in the system. Instance interfaces that are locked with `lock_avalon_base_address` keep their addresses during address auto-assignment.

### Usage

```
auto_assign_system_base_addresses
```

### Returns

No return value.

### Arguments

No arguments.

### Example

```
auto_assign_system_base_addresses
```

### Related Information

- [auto\\_assign\\_base\\_addresses](#) on page 5-81
- [lock\\_avalon\\_base\\_address](#) on page 5-135
- [unlock\\_avalon\\_base\\_address](#) on page 5-151



## auto\_assign\_irqs

### Description

Assigns interrupt numbers to all connected interrupt senders on an instance in the system.

### Usage

```
auto_assign_irqs <instance>
```

### Returns

No return value.

### Arguments

**instance**

The name of the instance with an interrupt sender.

### Example

```
auto_assign_irqs uart_0
```

## auto\_connect

### Description

Creates connections from an instance or instance interface to matching interfaces in other instances in the system. For example, Avalon-MM slaves connect to Avalon-MM masters.

### Usage

```
auto_connect <element>
```

### Returns

No return value.

### Arguments

#### element

The name of the instance interface, or the name of an instance.

### Example

```
auto_connect sdram  
auto_connect uart_0.s1
```

### Related Information

[add\\_connection](#) on page 5-77



## create\_system

### Description

Replaces the current system with a new system with the specified name.

### Usage

```
create_system [<name>]
```

### Returns

No return value.

### Arguments

#### **name (optional)**

The name of the new system.

### Example

```
create_system my_new_system_name
```

### Related Information

- [load\\_system](#) on page 5-134
- [save\\_system](#) on page 5-140

## export\_hw\_tcl

### Description

Allows you to save the currently open system as an **\_hw.tcl** file in the project directory. The saved systems appears under the **System** category in the IP Category.

### Usage

```
export_hw_tcl
```

### Returns

No return value.

### Arguments

No arguments

### Example

```
export_hw_tcl
```

## get\_composed\_connection\_parameter\_value

### Description

Returns the value of a parameter in a connection in a child instance containing a subsystem.

### Usage

```
get_composed_connection_parameter_value <instance> <child_connection> <parameter>
```

### Returns

The parameter value.

### Arguments

**instance**

The child instance that contains a subsystem

**child\_connection**

The name of the connection in the subsystem

**parameter**

The name of the parameter to query on the connection.

### Example

```
get_composed_connection_parameter_value subsystem_0 cpu.data_master/memory.s0  
baseAddress
```

### Related Information

- [get\\_composed\\_connection\\_parameters](#) on page 5-88
- [get\\_composed\\_connections](#) on page 5-89

## get\_composed\_connection\_parameters

### Description

Returns a list of all connections in the subsystem, for an instance that contains a subsystem.

### Usage

```
get_composed_connection_parameters <instance> <child_connection>
```

### Returns

A list of parameter names.

### Arguments

**instance**

The child instance containing a subsystem.

**child\_connection**

The name of the connection in the subsystem.

### Example

```
get_composed_connection_parameters subsystem_0 cpu.data_master/memory.s0
```

### Related Information

- [get\\_composed\\_connection\\_parameter\\_value](#) on page 5-87
- [get\\_composed\\_connections](#) on page 5-89



## get\_composed\_connections

### Description

For an instance that contains a subsystem of the Qsys system, returns a list of all connections found in a subsystem.

### Usage

```
get_composed_connections <instance>
```

### Returns

A list of connection names in the subsystem. These connection names are not qualified with the instance name.

### Arguments

**instance**

The child instance containing a subsystem.

### Example

```
get_composed_connections subsystem_0
```

### Related Information

- [get\\_composed\\_connection\\_parameter\\_value](#) on page 5-87
- [get\\_composed\\_connection\\_parameters](#) on page 5-88

## get\_composed\_instance\_assignment

### Description

For an instance that contains a subsystem of the Qsys system, returns the value of an assignment found on the instance in the subsystem.

### Usage

```
get_composed_instance_assignment <instance> <child_instance> <assignment>
```

### Returns

The value of the assignment.

### Arguments

**instance**

The child instance containing a subsystem.

**child\_instance**

The name of a child instance found in the subsystem.

**assignment**

The assignment key.

### Example

```
get_composed_instance_assignment subsystem_0 video_0 "embeddedsw.CMacro.colorSpace"
```

### Related Information

- [get\\_composed\\_instance\\_assignments](#) on page 5-91
- [get\\_composed\\_instances](#) on page 5-94



## get\_composed\_instance\_assignments

### Description

For an instance that contains a subsystem of the Qsys system, returns a list of assignments found on the instance in the subsystem.

### Usage

```
get_composed_instance_assignments <instance> <child_instance>
```

### Returns

A list of assignment names.

### Arguments

**instance**

The child instance containing a subsystem.

**child\_instance**

The name of a child instance found in the subsystem.

### Example

```
get_composed_instance_assignments subsystem_0 cpu
```

### Related Information

- [get\\_composed\\_instance\\_assignment](#) on page 5-90
- [get\\_composed\\_instances](#) on page 5-94

## get\_composed\_instance\_parameter\_value

### Description

For an instance that contains a subsystem of the Qsys system, returns the value of a parameters found on the instance in the subsystem.

### Usage

```
get_composed_instance_parameter_value <instance> <child_instance> <parameter>
```

### Returns

The value of a parameter on the instance in the subsystem.

### Arguments

**instance**

The child instance containing a subsystem.

**child\_instance**

The name of a child instance found in the subsystem.

**parameter**

The name of the parameter to query on the instance in the subsystem.

### Example

```
get_composed_instance_parameter_value subsystem_0 cpu DATA_WIDTH
```

### Related Information

- [get\\_composed\\_instance\\_parameters](#) on page 5-93
- [get\\_composed\\_instances](#) on page 5-94





## get\_composed\_instance\_parameters

### Description

For an instance that contains a subsystem of the Qsys system, returns a list of parameters found on the instance in the subsystem.

### Usage

```
get_composed_instance_parameters <instance> <child_instance>
```

### Returns

A list of parameter names.

### Arguments

**instance**

The child instance containing a subsystem.

**child\_instance**

The name of a child instance found in the subsystem.

### Example

```
get_composed_instance_parameters subsystem_0 cpu
```

### Related Information

- [get\\_composed\\_instance\\_parameter\\_value](#) on page 5-92
- [get\\_composed\\_instances](#) on page 5-94

## get\_composed\_instances

### Description

For an instance that contains a subsystem of the Qsys system, returns a list of child instances found in the subsystem.

### Usage

```
get_composed_instances <instance>
```

### Returns

A list of instance names found in the subsystem.

### Arguments

**instance**

The child instance containing a subsystem.

### Example

```
get_composed_instances subsystem_0
```

### Related Information

- [get\\_composed\\_instance\\_assignment](#) on page 5-90
- [get\\_composed\\_instance\\_assignments](#) on page 5-91
- [get\\_composed\\_instance\\_parameter\\_value](#) on page 5-92
- [get\\_composed\\_instance\\_parameters](#) on page 5-93



## get\_connection\_parameter\_property

### Description

Returns the value of a property on a parameter in a connection. Parameter properties are metadata about how Qsys uses the parameter.

### Usage

```
get_connection_parameter_property <connection> <parameter> <property>
```

### Returns

The value of the parameter property.

### Arguments

**connection**

The connection to query.

**parameter**

The name of the parameter.

**property**

The property of the connection. Refer to *Parameter Properties*.

### Example

```
get_connection_parameter_property cpu.data_master/dma0.csr baseAddress UNITS
```

### Related Information

- [get\\_connection\\_parameter\\_value](#) on page 5-96
- [get\\_connection\\_property](#) on page 5-99
- [get\\_connections](#) on page 5-100
- [get\\_parameter\\_properties](#) on page 5-130
- [Parameter Properties](#) on page 5-165

## get\_connection\_parameter\_value

### Description

Returns the value of a parameter on the connection. Parameters represent aspects of the connection that you can modify, such as the base address for an Avalon-MM connection.

### Usage

```
get_connection_parameter_value <connection> <parameter>
```

### Returns

The value of the parameter.

### Arguments

**connection**

The connection to query.

**parameter**

The name of the parameter.

### Example

```
get_connection_parameter_value cpu.data_master/dma0.csr baseAddress
```

### Related Information

- [get\\_connection\\_parameters](#) on page 5-97
- [get\\_connections](#) on page 5-100
- [set\\_connection\\_parameter\\_value](#) on page 5-142



## get\_connection\_parameters

### Description

Returns a list of parameters found on a connection.

### Usage

```
get_connection_parameters <connection>
```

### Returns

A list of parameter names.

### Arguments

#### connection

The connection to query.

### Example

```
get_connection_parameters cpu.data_master/dma0.csr
```

### Related Information

- [get\\_connection\\_parameter\\_property](#) on page 5-95
- [get\\_connection\\_parameter\\_value](#) on page 5-96
- [get\\_connection\\_property](#) on page 5-99

## get\_connection\_properties

### Description

Returns a list of properties found on a connection.

### Usage

```
get_connection_properties
```

### Returns

A list of connection properties.

### Arguments

No arguments.

### Example

```
get_connection_properties
```

### Related Information

- [get\\_connection\\_property](#) on page 5-99
- [get\\_connections](#) on page 5-100

## get\_connection\_property

### Description

Returns the value of a property found on a connection. Properties represent aspects of the connection that you can modify, such as the type of connection.

### Usage

```
get_connection_property <connection> <property>
```

### Returns

The value of a connection property.

### Arguments

#### connection

The connection to query.

#### property

The name of the connection. property. Refer to *Connection Properties*.

### Example

```
get_connection_property cpu.data_master/dma0.csr TYPE
```

### Related Information

- [get\\_connection\\_properties](#) on page 5-98
- [Connection Properties](#) on page 5-157

## get\_connections

### Description

Returns a list of all connections in the system if no element is specified. If you specify a child instance, for example `cpu`, Qsys returns all connections to any interface on the instance. If specify an interface on a child instance, for example `cpu.instruction_master`, Qsys returns all connections to that interface.

### Usage

```
get_connections [<element>]
```

### Returns

A list of connections.

### Arguments

#### **element (optional)**

The name of a child instance, or the qualified name of an interface on a child instance.

### Example

```
get_connections
get_connections cpu
get_connections cpu.instruction_master
```

### Related Information

- [add\\_connection](#) on page 5-77
- [remove\\_connection](#) on page 5-136



## get\_instance\_assignment

### Description

Returns the value of an assignment on a child instance. Qsys uses assignments to transfer information about hardware to embedded software tools and applications.

### Usage

```
get_instance_assignment <instance> <assignment>
```

### Returns

The value of the specified assignment.

### Arguments

**instance**

The name of the child instance.

**assignment**

The assignment key to query.

### Example

```
get_instance_assignment video_0 embeddedsw.CMacro.colorSpace
```

### Related Information

[get\\_instance\\_assignments](#) on page 5-102

## get\_instance\_assignments

### Description

Returns a list of assignment keys for any assignments defined for the instance.

### Usage

```
get_instance_assignments <instance>
```

### Returns

A list of assignment keys.

### Arguments

**instance**

The name of the child instance.

### Example

```
get_instance_assignments sdram
```

### Related Information

[get\\_instance\\_assignment](#) on page 5-101

[get\\_instance\\_assignment](#) on page 5-101

## get\_instance\_documentation\_links

### Description

Returns a list of all documentation links provided by an instance.

### Usage

```
get_instance_documentation_links <instance>
```

### Returns

A list of documentation links.

### Arguments

**instance**

The name of the child instance.

### Example

```
get_instance_documentation_links cpu_0
```

### Notes

The list of documentation links includes titles and URLs for the links. For instance, a component with a single data sheet link may return:

```
{Data Sheet} {http://url/to/data/sheet}
```

## get\_instance\_interface\_assignment

### Description

Returns the value of an assignment on an interface of a child instance. Qsys uses assignments to transfer information about hardware to embedded software tools and applications.

### Usage

```
get_instance_interface_assignment <instance> <interface> <assignment>
```

### Returns

The value of the specified assignment.

### Arguments

**instance**

The name of the child instance.

**interface**

The name of an interface on the child instance.

**assignment**

The assignment key to query.

### Example

```
get_instance_interface_assignment sdram s1 embeddedsw.configuration.isFlash
```

### Related Information

[get\\_instance\\_interface\\_assignments](#) on page 5-105



## get\_instance\_interface\_assignments

### Description

Returns a list of assignment keys for any assignments defined for an interface of a child instance.

### Usage

```
get_instance_interface_assignments <instance> <interface>
```

### Returns

A list of assignment keys.

### Arguments

**instance**

The name of the child instance.

**interface**

The name of an interface on the child instance.

### Example

```
get_instance_interface_assignments sdram s1
```

### Related Information

[get\\_instance\\_interface\\_assignment](#) on page 5-104

## get\_instance\_interface\_parameter\_property

### Description

Returns the value of a property on a parameter in an interface of a child instance. Parameter properties are metadata about how Qsys uses the parameter.

### Usage

```
get_instance_interface_parameter_property <instance> <interface> <parameter>  
<property>
```

### Returns

The value of the parameter property.

### Arguments

**instance**

The name of the child instance.

**interface**

The name of an interface on the child instance.

**parameter**

The name of the parameter on the interface.

**property**

The name of the property on the parameter. Refer to *Parameter Properties*.

### Example

```
get_instance_interface_parameter_property uart_0 s0 setupTime ENABLED
```

### Related Information

- [get\\_instance\\_interface\\_parameters](#) on page 5-108
- [get\\_instance\\_interfaces](#) on page 5-113
- [get\\_parameter\\_properties](#) on page 5-130
- [Parameter Properties](#) on page 5-165

## get\_instance\_interface\_parameter\_value

### Description

Returns the value of a parameter of an interface in a child instance.

### Usage

```
get_instance_interface_parameter_value <instance> <interface> <parameter>
```

### Returns

The value of the parameter.

### Arguments

**instance**

The name of the child instance.

**interface**

The name of an interface on the child instance.

**parameter**

The name of the parameter on the interface.

### Example

```
get_instance_interface_parameter_value uart_0 s0 setupTime
```

### Related Information

- [get\\_instance\\_interface\\_parameters](#) on page 5-108
- [get\\_instance\\_interfaces](#) on page 5-113

## get\_instance\_interface\_parameters

### Description

Returns a list of parameters for an interface in a child instance.

### Usage

```
get_instance_interface_parameters <instance> <interface>
```

### Returns

A list of parameter names for parameters in the interface.

### Arguments

**instance**

The name of the child instance.

**interface**

The name of an interface on the uart\_0 s0.

### Example

```
get_instance_interface_parameters instance interface
```

### Related Information

- [get\\_instance\\_interface\\_parameter\\_value](#) on page 5-107
- [get\\_instance\\_interfaces](#) on page 5-113



## get\_instance\_interface\_port\_property

### Description

Returns the value of a property of a port found in the interface of a child instance.

### Usage

```
get_instance_interface_port_property <instance> <interface> <port> <property>
```

### Returns

The value of the port property.

### Arguments

**instance**

The name of the child instance.

**interface**

The name of an interface on the child instance.

**port**

The name of the port in the interface.

**property**

The name of the property of the port. Refer to *Port Properties*.

### Example

```
get_instance_interface_port_property uart_0 exports tx WIDTH
```

### Related Information

- [get\\_instance\\_interface\\_ports](#) on page 5-110
- [get\\_port\\_properties](#) on page 5-131
- [Port Properties](#) on page 5-170

## get\_instance\_interface\_ports

### Description

Returns a list of ports found in an interface of a child instance.

### Usage

```
get_instance_interface_ports <instance> <interface>
```

### Returns

A list of port names found in the interface.

### Arguments

**instance**

The name of the child instance.

**interface**

The name of an interface on the child instance.

### Example

```
get_instance_interface_ports uart_0 s0
```

### Related Information

- [get\\_instance\\_interface\\_port\\_property](#) on page 5-109
- [get\\_instance\\_interfaces](#) on page 5-113

## get\_instance\_interface\_properties

### Description

Returns a list of properties that can be queried for an interface in a child instance.

### Usage

```
get_instance_interface_properties
```

### Returns

A list of property names.

### Arguments

No arguments.

### Example

```
get_instance_interface_properties
```

### Related Information

- [get\\_instance\\_interface\\_property](#) on page 5-112
- [get\\_instance\\_interfaces](#) on page 5-113

## get\_instance\_interface\_property

### Description

Returns the value of a property for an interface in a child instance.

### Usage

```
get_instance_interface_property <instance> <interface> <property>
```

### Returns

The value of the property.

### Arguments

**instance**

The name of the child instance.

**interface**

The name of an interface on the child instance.

**property**

The name of the property of the interface. Refer to *Element Properties*.

### Example

```
get_instance_interface_property uart_0 s0 DESCRIPTION
```

### Related Information

- [get\\_instance\\_interface\\_properties](#) on page 5-111
- [get\\_instance\\_interfaces](#) on page 5-113
- [Element Properties](#) on page 5-160



## get\_instance\_interfaces

### Description

Returns a list of interfaces found in a child instance

### Usage

```
get_instance_interfaces <instance>
```

### Returns

A list of interface names.

### Arguments

**instance**

The name of the child instance.

### Example

```
get_instance_interfaces uart_0
```

### Related Information

- [get\\_instance\\_interface\\_ports](#) on page 5-110
- [get\\_instance\\_interface\\_properties](#) on page 5-111
- [get\\_instance\\_interface\\_property](#) on page 5-112

## get\_instance\_parameter\_property

### Description

Returns the value of a property on a parameter in a child instance. Parameter properties are metadata about how Qsys uses the parameter.

### Usage

```
get_instance_parameter_property <instance> <parameter> <property>
```

### Returns

The value of the parameter property.

### Arguments

**instance**

The name of the child instance.

**parameter**

The name of the parameter in the instance.

**property**

The name of the property of the parameter. Refer to *Parameter Properties*.

### Example

```
get_instance_parameter_property uart_0 baudRate ENABLED
```

### Related Information

- [get\\_instance\\_parameters](#) on page 5-116
- [get\\_parameter\\_properties](#) on page 5-130
- [Parameter Properties](#) on page 5-165



## get\_instance\_parameter\_value

### Description

Returns the value of a parameter in a child instance.

### Usage

```
get_instance_parameter_value <instance> <parameter>
```

### Returns

The value of the parameter.

### Arguments

**instance**

The name of the child instance.

**parameter**

The name of the parameter in the instance.

### Example

```
get_instance_parameter_value uart_0 baudRate
```

### Related Information

- [get\\_instance\\_parameters](#) on page 5-116
- [set\\_instance\\_parameter\\_value](#) on page 5-143

## get\_instance\_parameters

### Description

Returns a list of parameters in a child instance.

### Usage

```
get_instance_parameters <instance>
```

### Returns

A list of parameters in the instance.

### Arguments

**instance**

The name of the child instance.

### Example

```
get_instance_parameters uart_0
```

### Related Information

- [get\\_instance\\_parameter\\_property](#) on page 5-114
- [get\\_instance\\_interface\\_parameter\\_value](#) on page 5-107
- [set\\_instance\\_parameter\\_value](#) on page 5-143





## get\_instance\_port\_property

### Description

Returns the value of a property of a port contained by an interface in a child instance.

### Usage

```
get_instance_port_property <instance> <port> <property>
```

### Returns

The value of the property for the port.

### Arguments

**instance**

The name of the child instance.

**port**

The name of a port in one of the interfaces on the child instance.

**property**

The name of a property found on the port. Refer to *Port Properties*.

### Example

```
get_instance_port_property uart_0 tx WIDTH
```

### Related Information

- [get\\_instance\\_interface\\_ports](#) on page 5-110
- [get\\_port\\_properties](#) on page 5-131
- [Port Properties](#) on page 5-170

## get\_instance\_properties

### Description

Returns a list of properties for a child instance.

### Usage

```
get_instance_properties
```

### Returns

A list of property names for the child instance.

### Arguments

No arguments.

### Example

```
get_instance_properties
```

### Related Information

[get\\_instance\\_property](#) on page 5-119

## get\_instance\_property

### Description

Returns the value of a property for a child instance.

### Usage

```
get_instance_property <instance> <property>
```

### Returns

The value of the property.

### Arguments

**instance**

The name of the child instance.

**property**

The name of a property found on the instance. Refer to *Element Properties*.

### Example

```
get_instance_property uart_0 ENABLED
```

### Related Information

- [get\\_instance\\_properties](#) on page 5-118
- [Element Properties](#) on page 5-160

## `get_instances`

### Description

Returns a list of the instance names for all child instances in the system.

### Usage

```
get_instances
```

### Returns

A list of child instance names.

### Arguments

No arguments.

### Example

```
get_instances
```

### Related Information

- [add\\_instance](#) on page 5-78
- [remove\\_instance](#) on page 5-138

## get\_interconnect\_requirement

### Description

Returns the value of an interconnect requirement for a system or interface on a child instance.

### Usage

```
get_interconnect_requirement <element_id> <requirement>
```

### Returns

The value of the interconnect requirement.

### Arguments

#### element\_id

{*\$system*} for the system, or the qualified name of the interface of an instance, in *<instance>.<interface>* format. In Tcl, the system identifier is escaped, for example, {*\$system*}.

#### requirement

The name of the requirement.

### Example

```
get_interconnect_requirement {$system} qsys_mm.maxAdditionalLatency
```

## get\_interconnect\_requirements

### Description

Returns a list of all interconnect requirements in the system.

### Usage

```
get_interconnect_requirements
```

### Returns

A flattened list of interconnect requirements. Every sequence of three elements in the list corresponds to one interconnect requirement. The first element in the sequence is the element identifier. The second element is the requirement name. The third element is the value. You can loop over the returned list with a `foreach` loop, for example:

```
foreach { element_id name value } $requirement_list { loop_body
}
```

### Arguments

No arguments.

### Example

```
get_interconnect_requirements
```

## get\_interface\_port\_property

### Description

Returns the value of a property of a port contained by one of the top-level exported interfaces

### Usage

```
get_interface_port_property <interface> <port> <property>
```

### Returns

The value of the property.

### Arguments

**interface**

The name of a top-level interface on the system.

**port**

The name of a port found in the interface.

**property**

The name of a property found on the port. Refer to *Port Properties*.

### Example

```
get_interface_port_property uart_exports tx DIRECTION
```

### Related Information

- [get\\_interface\\_ports](#) on page 5-124
- [get\\_port\\_properties](#) on page 5-131
- [Port Properties](#) on page 5-170

## `get_interface_ports`

### Description

Returns the names of all of the ports that have been added to a given interface.

### Usage

```
get_interface_ports <interface>
```

### Returns

A list of port names.

### Arguments

#### **interface**

The name of a top-level interface on the system.

### Example

```
get_interface_ports export_clk_out
```

### Related Information

- [get\\_interface\\_port\\_property](#) on page 5-123
- [get\\_interfaces](#) on page 5-127



## get\_interface\_properties

### Description

Returns the names of all the available interface properties common to all interface types.

### Usage

```
get_interface_properties
```

### Returns

A list of interface properties.

### Arguments

No arguments.

### Example

```
get_interface_properties
```

### Related Information

- [get\\_interface\\_property](#) on page 5-126
- [set\\_interface\\_property](#) on page 5-146

## get\_interface\_property

### Description

Returns the value of a single interface property from the specified interface.

### Usage

```
get_interface_property <interface> <property>
```

### Returns

The property value.

### Arguments

#### interface

The name of a top-level interface on the system.

#### property

The name of the property. Refer to *Interface Properties*.

### Example

```
get_interface_property export_clk_out EXPORT_OF
```

### Related Information

- [get\\_interface\\_properties](#) on page 5-125
- [set\\_interface\\_property](#) on page 5-146
- [Interface Properties](#) on page 5-162



## get\_interfaces

### Description

Returns a list of top-level interfaces in the system.

### Usage

```
get_interfaces
```

### Returns

A list of the top-level interfaces exported from the system.

### Arguments

No arguments.

### Example

```
get_interfaces
```

### Related Information

- [add\\_interface](#) on page 5-79
- [get\\_interface\\_ports](#) on page 5-124
- [get\\_interface\\_property](#) on page 5-126
- [remove\\_interface](#) on page 5-139
- [set\\_interface\\_property](#) on page 5-146

## get\_module\_properties

### Description

Returns the properties that you can manage for top-level module of the Qsys system.

### Usage

```
get_module_properties
```

### Returns

A list of property names.

### Arguments

No arguments.

### Example

```
get_module_properties
```

### Related Information

- [get\\_module\\_property](#) on page 5-129
- [set\\_module\\_property](#) on page 5-147

## get\_module\_property

### Description

Returns the value of a top-level system property.

### Usage

```
get_module_property <property>
```

### Returns

The value of the property.

### Arguments

**property**

The name of the property to query. Refer to *Module Properties*.

### Example

```
get_module_property NAME
```

### Related Information

- [get\\_module\\_properties](#) on page 5-128
- [set\\_module\\_property](#) on page 5-147

## get\_parameter\_properties

### Description

Returns a list of properties that you can query for any parameters, for example parameters on instances, interfaces, instance interfaces, and connections.

### Usage

```
get_parameter_properties
```

### Returns

A list of parameter properties.

### Arguments

No arguments.

### Example

```
get_parameter_properties
```

### Related Information

- [get\\_connection\\_parameter\\_property](#) on page 5-95
- [get\\_instance\\_interface\\_parameter\\_property](#) on page 5-106
- [get\\_instance\\_parameter\\_property](#) on page 5-114



## get\_port\_properties

### Description

Returns a list of properties that you can query for ports.

### Usage

```
get_port_properties
```

### Returns

A list of port properties.

### Arguments

No arguments.

### Example

```
get_port_properties
```

### Related Information

- [get\\_instance\\_interface\\_port\\_property](#) on page 5-109
- [get\\_instance\\_interface\\_ports](#) on page 5-110
- [get\\_instance\\_port\\_property](#) on page 5-117
- [get\\_interface\\_port\\_property](#) on page 5-123
- [get\\_interface\\_ports](#) on page 5-124

## get\_project\_properties

### Description

Returns a list of properties that you can query for properties pertaining to the Quartus II project.

### Usage

```
get_project_properties
```

### Returns

A list of project properties.

### Arguments

No arguments

### Example

```
get_project_properties
```

### Related Information

- [get\\_project\\_property](#) on page 5-133
- [set\\_project\\_property](#) on page 5-148



## get\_project\_property

### Description

Returns the value of a Quartus II project property. Not all Quartus II project properties are available.

### Usage

```
get_project_property <property>
```

### Returns

The value of the property.

### Arguments

**property**

The name of the project property. Refer to *Project properties*.

### Example

```
get_project_property DEVICE_FAMILY
```

### Related Information

- [get\\_module\\_properties](#) on page 5-128
- [get\\_module\\_property](#) on page 5-129
- [set\\_module\\_property](#) on page 5-147

## load\_system

### Description

Loads a Qsys system from a file, and uses the system as the current system for scripting commands.

### Usage

```
load_system <file>
```

### Returns

No return value.

### Arguments

#### file

The path to a **.qsys** file.

### Example

```
load_system example.qsys
```

### Related Information

- [create\\_system](#) on page 5-85
- [save\\_system](#) on page 5-140

## lock\_avalon\_base\_address

### Description

Prevents the memory-mapped base address from being changed for connections to the specified interface on an instance when Qsys runs the `auto_assign_base_addresses` or `auto_assign_system_base_addresses` commands.

### Usage

```
lock_avalon_base_address <instance.interface>
```

### Returns

No return value.

### Arguments

#### **instance.interface**

The qualified name of the interface of an instance, in `<instance>.<interface>` format.

### Example

```
lock_avalon_base_address sdram.s1
```

### Related Information

- [auto\\_assign\\_base\\_addresses](#) on page 5-81
- [auto\\_assign\\_system\\_base\\_addresses](#) on page 5-82
- [unlock\\_avalon\\_base\\_address](#) on page 5-151

## remove\_connection

### Description

This command removes a connection from the system.

### Usage

```
remove_connection <connection>
```

### Returns

no return value

### Arguments

#### connection

The name of the connection to remove

### Example

```
remove_connection cpu.data_master/sdram.s0
```

### Related Information

- [add\\_connection](#) on page 5-77
- [get\\_connections](#) on page 5-100

## remove\_dangling\_connections

### Description

Removes connections where both end points of the connection no longer exist in the system.

### Usage

```
remove_dangling_connections
```

### Returns

No return value.

### Arguments

No arguments.

### Example

```
remove_dangling_connections
```

## `remove_instance`

### Description

Removes a child instance from the system.

### Usage

```
remove_instance <instance>
```

### Returns

No return value.

### Arguments

#### **instance**

The name of the child instance to remove.

### Example

```
remove_instance cpu
```

### Related Information

- [add\\_instance](#) on page 5-78
- [get\\_instances](#) on page 5-120

## remove\_interface

### Description

Removes an exported top-level interface from the system.

### Usage

```
remove_interface <interface>
```

### Returns

No return value.

### Arguments

#### **interface**

The name of the exported top-level interface.

### Example

```
remove_interface clk_out
```

### Related Information

- [add\\_interface](#) on page 5-79
- [get\\_interfaces](#) on page 5-127

## save\_system

### Description

Saves the current system to the named file. If you do not specify the file, Qsys saves the system to the same file that was opened with the `load_system` command. You can specify the file as an absolute or relative path. Relative paths are relative to directory of the most recently loaded system, or relative to the working directory if no systems are loaded.

### Usage

```
save_system <file>
```

### Returns

No return value.

### Arguments

**file**

If available, the path of the `.qsys` file to save.

### Example

```
save_system
```

```
save_system file.qsys
```



## send\_message

### Description

Sends a message to the user of the script. The message text is normally interpreted as HTML. You can use the `<b>` element to provide emphasis.

### Usage

```
send_message <level> <message>
```

### Returns

No return value.

### Arguments

#### level

The following message levels are supported:

- `ERROR`--Provides an error message.
- `WARNING`--Provides a warning message.
- `INFO`--Provides an informational message.
- `PROGRESS`--Provides a progress message.
- `DEBUG`--Provides a debug message when debug mode is enabled. Refer to *Message Levels Properties*.

#### message

The text of the message.

### Example

```
send_message ERROR "The system is down!"
```

#### Related Information

[Message Levels Properties](#) on page 5-163

## set\_connection\_parameter\_value

### Description

Sets the value of a parameter for a connection.

### Usage

```
set_connection_parameter_value <connection> <parameter> <value>
```

### Returns

No return value.

### Arguments

**connection**

The name of the connection.

**parameter**

The name of the parameter.

**value**

The new parameter value.

### Example

```
set_connection_parameter_value cpu.data_master/dma0.csr baseAddress "0x000a0000"
```

### Related Information

- [get\\_connection\\_parameter\\_value](#) on page 5-96
- [get\\_connection\\_parameters](#) on page 5-97



## set\_instance\_parameter\_value

### Description

Set the value of a parameter for a child instance. You cannot set derived parameters and `SYSTEM_INFO` parameters for the child instance with this command.

### Usage

```
set_instance_parameter_value <instance> <parameter> <value>
```

### Returns

No return value.

### Arguments

**instance**

The name of the child instance.

**parameter**

The name of the parameter.

**value**

The new parameter value.

### Example

```
set_instance_parameter_value uart_0 baudRate 9600
```

### Related Information

- [get\\_instance\\_parameter\\_value](#) on page 5-115
- [get\\_instance\\_parameter\\_property](#) on page 5-114

## set\_instance\_property

### Description

Sets the value of a property of a child instance. Most instance properties are read-only and can only be set by the instance itself. The primary use for this command is to update the `ENABLED` parameter, which includes or excludes a child instance when generating Qsys interconnect.

### Usage

```
set_instance_property <instance> <property> <value>
```

### Returns

No return value.

### Arguments

**instance**

The name of the child instance.

**property**

The name of the property. Refer to *Instance Properties*.

**value**

The new property value.

### Example

```
set_instance_property cpu ENABLED false
```

### Related Information

- [get\\_instance\\_parameters](#) on page 5-116
- [get\\_instance\\_property](#) on page 5-119
- [Instance Properties](#) on page 5-161



## set\_interconnect\_requirement

### Description

Sets the value of an interconnect requirement for a system or an interface on a child instance.

### Usage

```
set_interconnect_requirement <element_id> <requirement> <value>
```

### Returns

No return value.

### Arguments

#### element\_id

{*\$system*} for the system, or qualified name of the interface of an instance, in *<instance>.<interface>* format. In Tcl, the system identifier is escaped, for example, {*\$system*}.

#### requirement

The name of the requirement.

#### value

The new requirement value.

### Example

```
set_interconnect_requirement {$system} qsys_mm.clockCrossingAdapter HANDSHAKE
```

## set\_interface\_property

### Description

Sets the value of a property on an exported top-level interface. You use this command to set the `EXPORT_OF` property to specify which interface of a child instance is exported via this top-level interface.

### Usage

```
set_interface_property <interface> <property> <value>
```

### Returns

No return value.

### Arguments

**interface**

The name of an exported top-level interface.

**property**

The name of the property. Refer to *Interface Properties*.

**value**

The new property value.

### Example

```
set_interface_property clk_out EXPORT_OF clk.clk_out
```

### Related Information

- [add\\_interface](#) on page 5-79
- [get\\_interface\\_properties](#) on page 5-125
- [get\\_interface\\_property](#) on page 5-126
- [Interface Properties](#) on page 5-162



## set\_module\_property

### Description

Sets the value of a system property, such as the name of the system using the `NAME` property.

### Usage

```
set_module_property <property> <value>
```

### Returns

No return value.

### Arguments

**property**

The name of the property. Refer to *Module Properties*.

**value**

The new property value.

### Example

```
set_module_property NAME "new_system_name"
```

### Related Information

- [get\\_module\\_properties](#) on page 5-128
- [get\\_module\\_property](#) on page 5-129
- [Module Properties](#) on page 5-164

## set\_project\_property

### Description

Sets the value of a project property, such as the device family.

### Usage

```
set_project_property <property> <value>
```

### Returns

No return value.

### Arguments

**property**

The name of the property. Refer to *Project Properties*.

**value**

The new property value.

### Example

```
set_project_property DEVICE_FAMILY "Cyclone IV GX"
```

### Related Information

- [get\\_project\\_properties](#) on page 5-132
- [set\\_project\\_property](#) on page 5-148
- [Project Properties](#) on page 5-171
- [get\\_project\\_properties](#) on page 5-132
- [get\\_project\\_property](#) on page 5-133
- [Project Properties](#) on page 5-171





## set\_use\_testbench\_naming\_pattern

### Description

Use this command to create testbench systems so that the generated file names for the test system match the system's original generated file names. Without setting this, the generated file names for the test system receive the top-level testbench system name.

### Usage

```
set_use_testbench_naming_pattern <value>
```

### Returns

No return value.

### Arguments

**value**

True or false.

### Example

```
set_use_testbench_naming_pattern true
```

### Notes

Use this command only to create testbench systems.

## set\_validation\_property

### Description

Sets a property that affects how and when validation is run. To disable system validation after each scripting command, set `AUTOMATIC_VALIDATION` to `False`.

### Usage

```
set_validation_property <property> <value>
```

### Returns

No return value.

### Arguments

**property**

The name of the property. Refer to *Validation Properties*.

**value**

The new property value.

### Example

```
set_validation_property AUTOMATIC_VALIDATION false
```

### Related Information

- [validate\\_system](#) on page 5-155
- [Validation Properties](#) on page 5-175

## unlock\_avalon\_base\_address

### Description

Allows the memory-mapped base address to change for connections to the specified interface on an instance when Qsys runs the `auto_assign_base_addresses` or `auto_assign_system_base_addresses` commands.

### Usage

```
unlock_avalon_base_address <instance.interface>
```

### Returns

No return value.

### Arguments

#### **instance.interface**

The qualified name of the interface of an instance, in `<instance>.<interface>` format.

### Example

```
unlock_avalon_base_address sdram.s1
```

### Related Information

- [auto\\_assign\\_base\\_addresses](#) on page 5-81
- [auto\\_assign\\_system\\_base\\_addresses](#) on page 5-82
- [lock\\_avalon\\_base\\_address](#) on page 5-135

## `validate_connection`

### Description

Validates the specified connection and returns validation messages.

### Usage

```
validate_connection <connection>
```

### Returns

A list of messages produced during validation.

### Arguments

#### **connection**

The name of the connection to validate.

### Example

```
validate_connection cpu.data_master/sdram.s1
```

### Related Information

- [validate\\_instance](#) on page 5-153
- [validate\\_instance\\_interface](#) on page 5-154
- [validate\\_system](#) on page 5-155

## validate\_instance

### Description

Validates the specified child instance and returns validation messages.

### Usage

```
validate_instance <instance>
```

### Returns

A list of messages produced during validation.

### Arguments

#### instance

The name of the child instance to validate.

### Example

```
validate_instance cpu
```

### Related Information

- [validate\\_connection](#) on page 5-152
- [validate\\_instance\\_interface](#) on page 5-154
- [validate\\_system](#) on page 5-155

## validate\_instance\_interface

### Description

Validates an interface on a child instance and returns validation messages.

### Usage

```
validate_instance_interface <instance> <interface>
```

### Returns

A list of messages produced during validation.

### Arguments

**instance**

The name of a child instance.

**interface**

The name of the interface on the child instance to validate.

### Example

```
validate_instance_interface cpu data_master
```

### Related Information

- [validate\\_connection](#) on page 5-152
- [validate\\_instance](#) on page 5-153
- [validate\\_system](#) on page 5-155

## validate\_system

### Description

Validates the system and returns validation messages.

### Usage

```
validate_system
```

### Returns

A list of validation messages produced during validation.

### Arguments

No arguments.

### Example

```
validate_system
```

### Related Information

- [validate\\_connection](#) on page 5-152
- [validate\\_instance](#) on page 5-153
- [validate\\_instance\\_interface](#) on page 5-154

## Qsys Scripting Property Reference

Interface properties work differently for `_hw.tcl` scripting than with qsys scripting. In `_hw.tcl`, interfaces do not distinguish between properties and parameters. In qsys scripting, properties and parameters are unique.

[Connection Properties](#) on page 5-157

[Design Environment Type Properties](#) on page 5-158

[Direction Properties](#) on page 5-159

[Element Properties](#) on page 5-160

[Instance Properties](#) on page 5-161

[Interface Properties](#) on page 5-162

[Message Levels Properties](#) on page 5-163

[Module Properties](#) on page 5-164

[Parameter Properties](#) on page 5-165

[Parameter Status Properties](#) on page 5-168

[Parameter Type Properties](#) on page 5-169

[Port Properties](#) on page 5-170

[Project Properties](#) on page 5-171

[System Info Type Properties](#) on page 5-172

[Units Properties](#) on page 5-174

[Validation Properties](#) on page 5-175



## Connection Properties

Type	Name	Description
string	END	The end interface of the connection.
string	NAME	The name of the connection.
string	START	The start interface of the connection.
String	TYPE	The type of the connection.

## Design Environment Type Properties

### Description

IP cores use the design environment to identify what sort of interfaces are most appropriate to connect in the parent system.

Name	Description
NATIVE	The design environment supports native IP interfaces.
QSYS	The design environment supports standard Qsys interfaces.

## Direction Properties

Name	Description
BIDIR	The direction for a bidirectional signal.
INOUT	The direction for an input signal.
OUTPUT	The direction for an output signal.

## Element Properties

### Description

Element properties are, with the exception of `ENABLED` and `NAME`, read-only properties of the types of instances, interfaces, and connections. These read-only properties represent metadata that does not vary between copies of the same type. `ENABLED` and `NAME` properties are specific to particular instances, interfaces, or connections.

Type	Name	Description
String	<code>AUTHOR</code>	The author of the component or interface.
Boolean	<code>AUTO_EXPORT</code>	Indicates whether unconnected interfaces on the instance are automatically exported.
String	<code>CLASS_NAME</code>	The type of the instance, interface or connection, for example, <code>altera_nios2</code> or <code>avalon_slave</code> .
String	<code>DESCRIPTION</code>	The description of the instance, interface or connection type.
String	<code>DISPLAY_NAME</code>	The display name for referencing the type of instance, interface or connection.
Boolean	<code>EDITABLE</code>	Indicates whether you can edit the component in the Qsys Component Editor.
Boolean	<code>ENABLED</code>	Indicates whether the instance is turned on.
String	<code>GROUP</code>	The IP Catalog category.
Boolean	<code>INTERNAL</code>	Hides internal IP components or sub-components from the IP Catalog..
String	<code>NAME</code>	The name of the instance, interface or connection.
String	<code>VERSION</code>	The version number of the instance, interface or connection, for example, 14.0.

## Instance Properties

Type	Name	Description
String	AUTO_EXPORT	Indicates whether unconnected interfaces on the instance are automatically exported.
Boolean	ENABLED	If true, this instance is included in the generated system. if false, it is not included.
String	NAME	The name of the system, which is used as the name of the top-level module in the generated HDL.

## Interface Properties

Type	Name	Description
String	EXPORT_OF	Indicates which interface of a child instance to export through the top-level interface. Before using this command, you must create the top-level interface using the <code>add_interface</code> command. You must use the format: <code>&lt;instanceName.interfaceName&gt;</code> . For example:  <pre>set_interface_property CSC_input EXPORT_OF my_colorSpace-Converter.input_port</pre>

## Message Levels Properties

Name	Description
COMPONENT_INFO	Reports an informational message only during component editing.
DEBUG	Provides messages when debug mode is turned on.
ERROR	Provides an error message.
INFO	Provides an informational message.
PROGRESS	Reports progress during generation.
TODOERROR	Provides an error message that indicates the system is incomplete.
WARNING	Provides a warning message.

**Module Properties**

Type	Name	Description
String	GENERATION_ID	The generation ID for the system.
String	NAME	The name of the instance.



## Parameter Properties

Type	Name	Description
Boolean	AFFECTS_ELABORATION	Set AFFECTS_ELABORATION to false for parameters that do not affect the external interface of the module. An example of a parameter that does not affect the external interface is <code>isNonVolatileStorage</code> . An example of a parameter that does affect the external interface is <code>width</code> . When the value of a parameter changes and AFFECTS_ELABORATION is false, the elaboration phase does not repeat and improves performance. When AFFECTS_ELABORATION is set to true, the default value, Qsys re-analyzes the HDL file to determine the port widths and configuration each time a parameter changes.
Boolean	AFFECTS_GENERATION	The default value of AFFECTS_GENERATION is false if you provide a top-level HDL module. The default value is true if you provide a fileset callback. Set AFFECTS_GENERATION to false if the value of a parameter does not change the results of fileset generation.
Boolean	AFFECTS_VALIDATION	The AFFECTS_VALIDATION property determines whether a parameter's value sets derived parameters, and whether the value affects validation messages. When set to false, this may improve response time in the parameter editor when the value changes.
String[]	ALLOWED_RANGES	Indicates the range or ranges of the parameter. For integers, Each range is a single value, or a range of values defined by a start and end value, and delimited by a colon, for example, <code>11:15</code> . This property also specifies the legal values and description strings for integers, for example, <code>{0:None 1:Monophonic 2:Stereo 4:Quadrophonic}</code> , where 0, 1, 2, and 4 are the legal values. You can assign description strings in the parameter editor for string variables. For example, <pre>ALLOWED_RANGES { "dev1:Cyclone IV GX" "dev2:Stratix V GT" }</pre>
String	DEFAULT_VALUE	The default value.
Boolean	DERIVED	When <code>True</code> , indicates that the parameter value is set by the component and cannot be set by the user. Derived parameters are not saved as part of an instance's parameter values. The default value is <code>False</code> .
String	DESCRIPTION	A short user-visible description of the parameter, suitable for a tooltip description in the parameter editor.
String[]	DISPLAY_HINT	Provides a hint about how to display a property.

Type	Name	Description
		<ul style="list-style-type: none"> <li><code>boolean</code>--For integer parameters whose value are 0 or 1. The parameter displays as an option that you can turn on or off.</li> <li><code>radio</code>--Displays a parameter with a list of values as radio buttons.</li> <li><code>hexadecimal</code>--For integer parameters, displays and interprets the value as a hexadecimal number, for example: <code>0x00000010</code> instead of 16.</li> <li><code>fixed_size</code>--For <code>string_list</code> and <code>integer_list</code> parameters, the <code>fixed_size DISPLAY_HINT</code> eliminates the <b>Add</b> and <b>Remove</b> buttons from tables.</li> </ul>
String	<code>DISPLAY_NAME</code>	The GUI label that appears to the left of this parameter.
String	<code>DISPLAY_UNITS</code>	The GUI label that appears to the right of the parameter.
Boolean	<code>ENABLED</code>	When <code>False</code> , the parameter is turned off. It displays in the parameter editor but is greyed out, indicating that you cannot edit this parameter.
String	<code>GROUP</code>	Controls the layout of parameters in the GUI.
Boolean	<code>HDL_PARAMETER</code>	When <code>True</code> , Qsys passes the parameter to the HDL component description. The default value is <code>False</code> .
String	<code>LONG_DESCRIPTION</code>	A user-visible description of the parameter. Similar to <code>DESCRIPTION</code> , but allows a more detailed explanation.
String	<code>NEW_INSTANCE_VALUE</code>	Changes the default value of a parameter without affecting older components that do not explicitly set a parameter value, and use the <code>DEFAULT_VALUE</code> property. Older instances continue to use <code>DEFAULT_VALUE</code> for the parameter and new instances use the value assigned by <code>NEW_INSTANCE_VALUE</code> .
String[]	<code>SYSTEM_INFO</code>	Allows you to assign information about the instantiating system to a parameter that you define. <code>SYSTEM_INFO</code> requires an argument specifying the type of information for example, <p style="margin-left: 40px;"><code>SYSTEM_INFO &lt;info-type&gt;</code></p>
String	<code>SYSTEM_INFO_ARG</code>	Defines an argument to pass to <code>SYSTEM_INFO</code> . For example, the name of a reset interface.
(various)	<code>SYSTEM_INFO_TYPE</code>	Specifies the types of system information that you can query. Refer to <i>System Info Type Properties</i> .
(various)	<code>TYPE</code>	Specifies the type of the parameter. Refer to <i>Parameter Type Properties</i> .
(various)	<code>UNITS</code>	Sets the units of the parameter. Refer to <i>Units Properties</i> .
Boolean	<code>VISIBLE</code>	Indicates whether or not to display the parameter in the parameter editor.

Type	Name	Description
String	WIDTH	Indicates the width of the logic vector for the <code>STD_LOGIC_VECTOR</code> parameter.

#### Related Information

- [System Info Type Properties](#) on page 5-172
- [Parameter Type Properties](#) on page 5-169
- [Units Properties](#) on page 5-174

**Parameter Status Properties**

Type	Name	Description
Boolean	ACTIVE	Indicates that this parameter is an active parameter.
Boolean	DEPRECATED	Indicates that this parameter exists only for backwards compatibility, and may not have any effect.
Boolean	EXPERIMENTAL	Indicates that this parameter is experimental and not exposed in the design flow.

## Parameter Type Properties

Name	Description
BOOLEAN	A boolean parameter set to <code>true</code> or <code>false</code> .
FLOAT	A signed 32-bit floating point parameter. (Not supported for HDL parameters.)
INTEGER	A signed 32-bit integer parameter.
INTEGER_LIST	A parameter that contains a list of 32-bit integers. (Not supported for HDL parameters.)
LONG	A signed 64-bit integer parameter. (Not supported for HDL parameters.)
NATURAL	A 32-bit number that contains values 0 to 2147483647 (0x7fffffff).
POSITIVE	A 32-bit number that contains values 1 to 2147483647 (0x7fffffff).
STD_LOGIC	A single bit parameter set to 0 or 1.
STD_LOGIC_VECTOR	An arbitrary-width number. The parameter property <code>WIDTH</code> determines the size of the logic vector.
STRING	A string parameter.
STRING_LIST	A parameter that contains a list of strings. (Not supported for HDL parameters.)

**Port Properties**

Type	Name	Description
(various)	DIRECTION	The direction of the signal. Refer to <i>Direction Properties</i> .
String	ROLE	The type of the signal. Each interface type defines a set of interface types for its ports.
Integer	WIDTH	The width of the signal in bits.

## Project Properties

Type	Name	Description
String	DEVICE	The device part number in the Quartus II project that contains the Qsys system.
String	DEVICE_FAMILY	The device family name in the Quartus II project that contains the Qsys system.

## System Info Type Properties

Type	Name	Description
String	ADDRESS_MAP	An XML-formatted string that describes the address map for the interface specified in the <code>SYSTEM_INFO</code> parameter property.
Integer	ADDRESS_WIDTH	The number of address bits that Qsys requires to address memory-mapped slaves connected to the specified memory-mapped master in this instance.
String	AVALON_SPEC	The version of the Qsys interconnect. Refer to <i>Avalon Interface Specifications</i> .
Integer	CLOCK_DOMAIN	An integer that represents the clock domain for the interface specified in the <code>SYSTEM_INFO</code> parameter property. If this instance has interfaces on multiple clock domains, you can use this property to determine which interfaces are on each clock domain. The absolute value of the integer is arbitrary.
Long, Integer	CLOCK_RATE	The rate of the clock connected to the clock input specified in the <code>SYSTEM_INFO</code> parameter property. If zero, the clock rate is currently unknown.
String	CLOCK_RESET_INFO	The name of this instance's primary clock or reset sink interface. You use this property to determine the reset sink for global reset when you use SOPC Builder interconnect that conforms to <i>Avalon Interface Specifications</i> .
String	CUSTOM_INSTRUCTION_SLAVES	Provides slave information, including the name, base address, address span, and clock cycle type.
String	DESIGN_ENVIRONMENT	A string that identifies the current design environment. Refer to <i>Design Environment Type Properties</i> .
String	DEVICE	The device part number of the selected device.
String	DEVICE_FAMILY	The family name of the selected device.
String	DEVICE_FEATURES	A list of key/value pairs delimited by spaces that indicate whether a device feature is available in the selected device family. The format of the list is suitable for passing to the <code>array</code> command. The keys are device features. The values are 1 if the feature is present, and 0 if the feature is absent.
String	DEVICE_SPEEDGRADE	The speed grade of the selected device.
Integer	GENERATION_ID	A integer that stores a hash of the generation time that Qsys uses as a unique ID for a generation run.



Type	Name	Description
BigInteger, Long	INTERRUPTS_USED	A mask indicating which bits of an interrupt receiver are connected to interrupt senders. The interrupt receiver is specified in the system info argument.
Integer	MAX_SLAVE_DATA_WIDTH	The data width of the widest slave connected to the specified memory-mapped master.
String, Boolean, Integer	QUARTUS_INI	The value of the <b>quartus.ini</b> setting specified in the system info argument.
Integer	RESET_DOMAIN	An integer representing the reset domain for the interface specified in the <code>SYSTEM_INFO</code> parameter property. If this instance has interfaces on multiple reset domains, you can use this property to determine which interfaces are on each reset domain. The absolute value of the integer is arbitrary.
String	TRISTATECONDUIT_INFO	An XML description of the tri-state conduit masters connected to a tri-state conduit slave. The slave is specified as the <code>SYSTEM_INFO</code> parameter property. The value contains information about the slave, connected master instance and interface names, and signal names, directions, and widths.
String	TRISTATECONDUIT_MASTERS	The names of the instance's interfaces that are tri-state conduit slaves.
String	UNIQUE_ID	A string guaranteed to be unique to this instance.

#### Related Information

- [Design Environment Type Properties](#) on page 5-158
- [Avalon Interface Specifications](#)
- [Qsys Interconnect](#) on page 7-1

**Units Properties**

Name	Description
ADDRESS	A memory-mapped address.
BITS	Memory size in bits.
BITSPERSECOND	Rate in bits per second.
BYTES	Memory size in bytes.
CYCLES	A latency or count in clock cycles.
GIGABITSPERSECOND	Rate in gigabits per second.
GIGABYTES	Memory size in gigabytes.
GIGAHERTZ	Frequency in GHz.
HERTZ	Frequency in Hz.
KILOBITSPERSECOND	Rate in kilobits per second.
KILOBYTES	Memory size in kilobytes.
KILOHERTZ	Frequency in kHz.
MEGABITSPERSECOND	Rate, in megabits per second.
MEGABYTES	Memory size in megabytes.
MEGAHERTZ	Frequency in MHz.
MICROSECONDS	Time in microseconds.
MILLISECONDS	Time in milliseconds.
NANOSECONDS	Time in nanoseconds.
NONE	Unspecified units.
PERCENT	A percentage.
PICOSECONDS	Time in picoseconds.
SECONDS	Time in seconds.



## Validation Properties

Type	Name	Description
Boolean	<code>AUTOMATIC_VALIDATION</code>	When <code>true</code> , Qsys runs system validation and elaboration after each scripting command. When <code>false</code> , Qsys runs system validation with validation scripting commands. Some queries affected by system elaboration may be incorrect if automatic validation is turned off. You can disable validation to make a system script run faster.

## Document Revision History

The table below indicates edits made to the *Creating a System With Qsys* content since its creation.

**Table 5-17: Document Revision History**

Date	Version	Changes
December 2014	14.1.0	<ul style="list-style-type: none"> <li>• Create and Manage Hierarchical Qsys Systems.</li> <li>• Schematic tab.</li> <li>• View and Filter Clock and Reset Domains.</li> <li>• <b>File &gt; Recent Projects</b> menu item.</li> <li>• Updated example: Hierarchical System Using Instance Parameters</li> </ul>
August 2014	14.0a10.0	<ul style="list-style-type: none"> <li>• Added distinction between legacy and standard device generation.</li> <li>• Updated: <i>Upgrading Outdated IP Components</i>.</li> <li>• Updated: <i>Generating a Qsys System</i>.</li> <li>• Updated: <i>Integrating a Qsys System with the Quartus II Software</i>.</li> <li>• Added screen shot: <i>Displaying Your Qsys System</i>.</li> </ul>
June 2014	14.0.0	<ul style="list-style-type: none"> <li>• Added tab descriptions: Details, Connections.</li> <li>• Added <i>Managing IP Settings in the Quartus II Software</i>.</li> <li>• Added <i>Upgrading Outdated IP Components</i>.</li> <li>• Added <i>Support for Avalon-MM Non-Power of Two Data Widths</i>.</li> </ul>
November 2013	13.1.0	<ul style="list-style-type: none"> <li>• Added <i>Integrating with the .qsys File</i>.</li> <li>• Added <i>Using the Hierarchy Tab</i>.</li> <li>• Added <i>Managing Interconnect Requirements</i>.</li> <li>• Added <i>Viewing Qsys Interconnect</i>.</li> </ul>
May 2013	13.0.0	<ul style="list-style-type: none"> <li>• Added AMBA APB support.</li> <li>• Added qsys-generate utility.</li> <li>• Added VHDL BFM ID support.</li> <li>• Added <i>Creating Secure Systems (TrustZones)</i>.</li> <li>• Added <i>CMSIS Support for Qsys Systems With An HPS Component</i>.</li> <li>• Added VHDL language support options.</li> </ul>
November 2012	12.1.0	<ul style="list-style-type: none"> <li>• Added AMBA AXI4 support.</li> </ul>

Date	Version	Changes
June 2012	12.0.0	<ul style="list-style-type: none"> <li>• Added AMBA AX3I support.</li> <li>• Added Preset Editor updates.</li> <li>• Added command-line utilities, and scripts.</li> </ul>
November 2011	11.1.0	<ul style="list-style-type: none"> <li>• Added Synopsys VCS and VCS MX Simulation Shell Script.</li> <li>• Added Cadence Incisive Enterprise (NCSIM) Simulation Shell Script.</li> <li>• Added <i>Using Instance Parameters and Example Hierarchical System Using Parameters</i>.</li> </ul>
May 2011	11.0.0	<ul style="list-style-type: none"> <li>• Added simulation support in Verilog HDL and VHDL.</li> <li>• Added testbench generation support.</li> <li>• Updated simulation and file generation sections.</li> </ul>
December 2010	10.1.0	Initial release.

**Related Information**

[Quartus II Handbook Archive](#)

2014.06.30

QII5V1



Subscribe



Send Feedback

In order to describe and package IP components for use in a Qsys system, you must create a Hardware Component Definition File (`_hw.tcl`) which will describes your component, its interfaces and HDL files. Qsys provides the Component Editor to help you create a simple `_hw.tcl` file.

The **Demo AXI Memory** example on the **Qsys Design Examples** page of the Altera web site provides the full code examples that appear in the following topics.

Qsys supports Avalon, AMBA AXI3 (version 1.0), AMBA AXI4 (version 2.0), AMBA AXI4-Lite (version 2.0), AMBA AXI4-Stream (version 1.0), and AMBA APB3 (version 1.0) interface specifications.

#### Related Information

- [Avalon Interface Specifications](#)
- [AMBA Protocol Specifications](#)
- [Demo AXI Memory Example](#)

## Qsys Components

A Qsys component includes the following elements:

- Information about the component type, such as name, version, and author.
- HDL description of the component's hardware, including SystemVerilog, Verilog HDL, or VHDL files
- Constraint files (Synopsys Design Constraints File (`.sdc`) and/or Quartus II IP File (`.qip`)) that define the component for synthesis and simulation.
- A component's interfaces, including I/O signals.
- The parameters that configure the operation of the component.

## IP Component Interface Support in Qsys

IP components can have any number of interfaces in any combination. Each interface represents a set of signals that you can connect within a Qsys system, or export outside of a Qsys system.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



Qsys IP components can include the following interface types:

**Table 6-1: IP Component Interface Types**

Interface Type	Description
Memory-Mapped	Connects memory-referencing master devices with slave memory devices. Master devices may be processors and DMAs, while slave memory devices may be RAMs, ROMs, and control registers. Data transfers between master and slave may be unidirectional (read only or write only), or bi-directional (read or write).
Streaming	Connects Avalon Streaming (Avalon-ST) sources and sinks that stream unidirectional data, as well as high-bandwidth, low-latency IP components. Streaming creates datapaths for unidirectional traffic, including multichannel streams, packets, and DSP data. The Avalon-ST interconnect is flexible and can implement on-chip interfaces for industry standard telecommunications and data communications cores, such as Ethernet, Interlaken, and video. You can define bus widths, packets, and error conditions.
Interrupts	Connects interrupt senders to interrupt receivers. Qsys supports individual, single-bit interrupt requests (IRQs). In the event that multiple senders assert their IRQs simultaneously, the receiver logic (typically under software control) determines which IRQ has highest priority, then responds appropriately
Clocks	Connects clock output interfaces with clock input interfaces. Clock outputs can fan-out without the use of a bridge. A bridge is required only when a clock from an external (exported) source connects internally to more than one source.
Resets	Connects reset sources with reset input interfaces. If your system requires a particular positive-edge or negative-edge synchronized reset, Qsys inserts a reset controller to create the appropriate reset signal. If you design a system with multiple reset inputs, the reset controller ORs all reset inputs and generates a single reset output.
Conduits	Connects point-to-point conduit interfaces, or represent signals that are exported from the Qsys system. Qsys uses conduits for component I/O signals that are not part of any supported standard interface. You can connect two conduits directly within a Qsys system as a point-to-point connection, or conduit interfaces can be exported and brought to the top-level of the system as top-level system I/O. You can use conduits to connect to external devices, for example external DDR SDRAM memory, and to FPGA logic defined outside of the Qsys system.

## Component Structure

Altera provides components automatically installed with the Quartus II software. You can obtain a list of Qsys-compliant components provided by third-party IP developers on Altera's **Intellectual Property & Reference Designs** page by typing: **qsys certified** in the **Search** box, and then selecting **IP Core & Reference Designs**. Components are also provided with Altera development kits, which are listed on the **All Development Kits** page.

Every component is defined with a `<component_name>_hw.tcl` file, a text file written in the Tcl scripting language that describes the component to Qsys. When you design your own custom component, you can create the `_hw.tcl` file manually, or by using the Qsys Component Editor.

The Component Editor simplifies the process of creating `_hw.tcl` files by creating a file that you can edit outside of the Component Editor to add advanced procedures. When you edit a previously saved `_hw.tcl` file, Qsys automatically backs up the earlier version as `_hw.tcl~`.

You can move component files into a new directory, such as a network location, so that other users can use the component in their systems. The `_hw.tcl` file contains relative paths to the other files, so if you move an `_hw.tcl` file, you should also move all the HDL and other files associated with it.

There are three component types:

- **Static**— Static components always generate the same output, regardless of their parameterization. Components that instantiate static components must have only static children.
- **Generated**— A generated component's fileset callback allows an instance of the component to create unique HDL design files based on the instance's parameter values.
- **Composed**— Composed components are subsystems constructed from instances of other components. You can use a composition callback to manage the subsystem in a composed component.

#### Related Information

- [Creating a Composed Component or Subsystem](#) on page 6-28
- [Adding Component Instances to a Static or Generated Component](#) on page 6-31
- [Intellectual Property & Reference Designs](#)

## Component File Organization

A typical component uses the following directory structure where the names of the directories are not significant:

`<component_directory>/`

- `<hdl>/`— Contains the component HDL design files, for example `.v`, `.sv`, or `.vhd` files that contain the top-level module, along with any required constraint files.
- `<component_name>_hw.tcl`— The component description file.
- `<component_name>_sw.tcl`— The software driver configuration file. This file specifies the paths for the `.c` and `.h` files associated with the component, when required.
- `<software>/`— Contains software drivers or libraries related to the component.

**Note:** Refer to the *Nios II Software Developer's Handbook* for information about writing a device driver or software package suitable for use with the Nios II processor.

#### Related Information

- [Hardware Abstraction Layer Tool Reference \(Nios II Software Developer's Handbook\)](#)
- [Nios II Software Build Tool Reference \(Nios II Software Developer's Handbook\)](#)

## Component Versions

Qsys systems support multiple versions of the same component within the same system; you can create and maintain multiple versions of the same component.



If you have multiple `_hw.tcl` files for components with the same NAME module properties and different VERSION module properties, both versions of the component are available.

If multiple versions of the component are available in the IP Catalog, you can add a specific version of a component by right-clicking the component, and then selecting **Add version** `<version_number>`.

## Upgrading IP Components to the Latest Version

When you open a Qsys design, if Qsys detects IP components that require regeneration, the **Upgrade IP Cores** dialog box appears and allows you to upgrade outdated components.

Components that you must upgrade in order to successfully compile your design appear in red. Status icons indicate whether a component is currently being regenerated, the component is encrypted, or that there is not enough information to determine the status of component. To upgrade a component, in the **Upgrade IP Cores** dialog box, select the component that you want to upgrade, and then click **Upgrade**. The Quartus II software maintains a list of all IP components associated with your design on the **Components** tab in the Project Navigator.

### Related Information

[Upgrade IP Components Dialog Box](#)

## Life Cycle of an IPComponent

When you define a component with the Qsys Component Editor, or a custom `_hw.tcl` file, you specify the information that Qsys requires to instantiate the component in a Qsys system and to generate the appropriate output files for synthesis and simulation.

The following phases describe the process when working with components in Qsys:

- **Discovery**—During the discovery phase, Qsys reads the `_hw.tcl` file to identify information that appears in the IP Catalog, such as the component's name, version, and documentation URLs. Each time you open Qsys, the tool searches for the following file types using the default search locations and entries in the **IP Search Path**:
  - `_hw.tcl` files—Each `_hw.tcl` file defines a single component.
  - IP Index (`.ipx`) files—Each `.ipx` file indexes a collection of available components, or a reference to other directories to search.
- **Static Component Definition**—During the static component definition phase, Qsys reads the `_hw.tcl` file to identify static parameter declarations, interface properties, interface signals, and HDL files that define the component. At this stage of the life cycle, the component interfaces may be only partially defined.
- **Parameterization**—During the parameterization phase, after an instance of the component is added to a Qsys system, the user of the component specifies parameters with the component's parameter editor.
- **Validation**—During the validation phase, Qsys validates the values of each instance's parameters against the allowed ranges specified for each parameter. You can use callback procedures that run during the validation phase to provide validation messages. For example, if there are dependencies between parameters where only certain combinations of values are supported, you can report errors for the unsupported values.

- **Elaboration**—During the elaboration phase, Qsys queries the component for its interface information. Elaboration is triggered when an instance of a component is added to a system, when its parameters are changed, or when a system property changes. You can use callback procedures that run during the elaboration phase to dynamically control interfaces, signals, and HDL files based on the values of parameters. For example, interfaces defined with static declarations can be enabled or disabled during elaboration. When elaboration is complete, the component's interfaces and design logic must be completely defined.
- **Composition**—During the composition phase, a component can manipulate the instances in the component's subsystem. The `_hw.tcl` file uses a callback procedure to provide parameterization and connectivity of sub-components.
- **Generation**—During the generation phase, Qsys generates synthesis or simulation files for each component in the system into the appropriate output directories, as well as any additional files that support associated tools.

## Creating Qsys Components in the Component Editor

The Qsys Component Editor, accessed by clicking **New Component** in the IP Catalog, allows you to create and package a component for use in Qsys. When you use the Component Editor to define a component, the Component Editor writes the information to the `_hw.tcl` file.

The Component Editor allows you to perform the following tasks:

- Specify component's identifying information, such as name, version, author, etc.
- Specify the SystemVerilog, Verilog HDL, or VHDL files, and constraint files that define the component for synthesis and simulation.
- Create an HDL template for a component by first defining its parameters, signals, and interfaces.
- Associate and define signals for a component's interfaces.
- Set parameters on interfaces, which specify characteristics.
- Specify relationships between interfaces.
- Declare parameters that alter the component structure or functionality.

If the component is HDL-based, you must define the parameters and signals in the HDL file, and cannot add or remove them in the Component Editor. If you have not yet created the top-level HDL file, you declare the parameters and signals in the Component Editor, and they are then included in the HDL template file that Qsys creates.

In a Qsys system, the interfaces of a component are connected within the system, or exported as top-level signals from the system.

If you are creating the component using an existing HDL file, the order in which the tabs appear in the Component Editor reflects the recommended design flow for component development. You can use the **Prev** and **Next** buttons at the bottom of the Component Editor window to guide you through the tabs.

If the component is not based on an existing HDL file, enter the parameters, signals, and interfaces first, and then return to the **Files** tab to create the top-level HDL file template. When you click **Finish**, Qsys creates the component `_hw.tcl` file with the details provided on the Component Editor tabs.

After the component is saved, it is available in the IP Catalog.

If you require features in the component that are not supported by the Component Editor, such as callback procedures, you can use the Component Editor to create the `_hw.tcl` file, and then manually edit

the file to complete the component definition. Subsequent topics document the `_hw.tcl` commands that are generated by the Component Editor, as well as some of the advanced features that you can add with your own `_hw.tcl` commands.

**Note:** By default, custom component do not have registered outputs, even if they are exported out of the Qsys system. For a custom component, if you want to export the signals, you must add the registered outputs.

#### Related Information

- [Component Interface Tcl Reference](#) on page 9-1

## Saving a Component and Creating an `_hw.tcl` File

You save a component by clicking **Finish** in the Component Editor. The Component Editor saves the component to a file with the file name `<component_name>_hw.tcl`.

Altera recommends that you save `_hw.tcl` files and their associated files in an `ip/ <class-name>` directory within your Quartus II project directory. You can also publish component information for use by software, such as a C compiler and a board support package (BSP) generator.

Refer to *Creating a System with Qsys* for information on how to search for and add components to the IP Catalog for use in your designs.

#### Related Information

[Publishing Component Information to Embedded Software \(Nios II Software Developer's Handbook\)](#)

[Creating a System with Qsys](#) on page 5-1

## Editing a Component with the Component Editor

In Qsys, you make changes to a component by right-clicking the component in the System Contents view, and then clicking **Edit**. After making changes, click **Finish** to save the changes to the `_hw.tcl` file. You can open the `_hw.tcl` file in a text editor to view the hardware Tcl for the component. If you edit the `_hw.tcl` file to customize the component with advanced features, you cannot use the Component Editor to make further changes without over-writing your customized file.

You cannot use the Component Editor to edit components installed with the Quartus II software, such as Altera-provided components. If you edit the HDL for a component and change the interface to the top-level module, you must edit the component to reflect the changes you made to the HDL.

#### Related Information

[Creating Qsys Components](#)

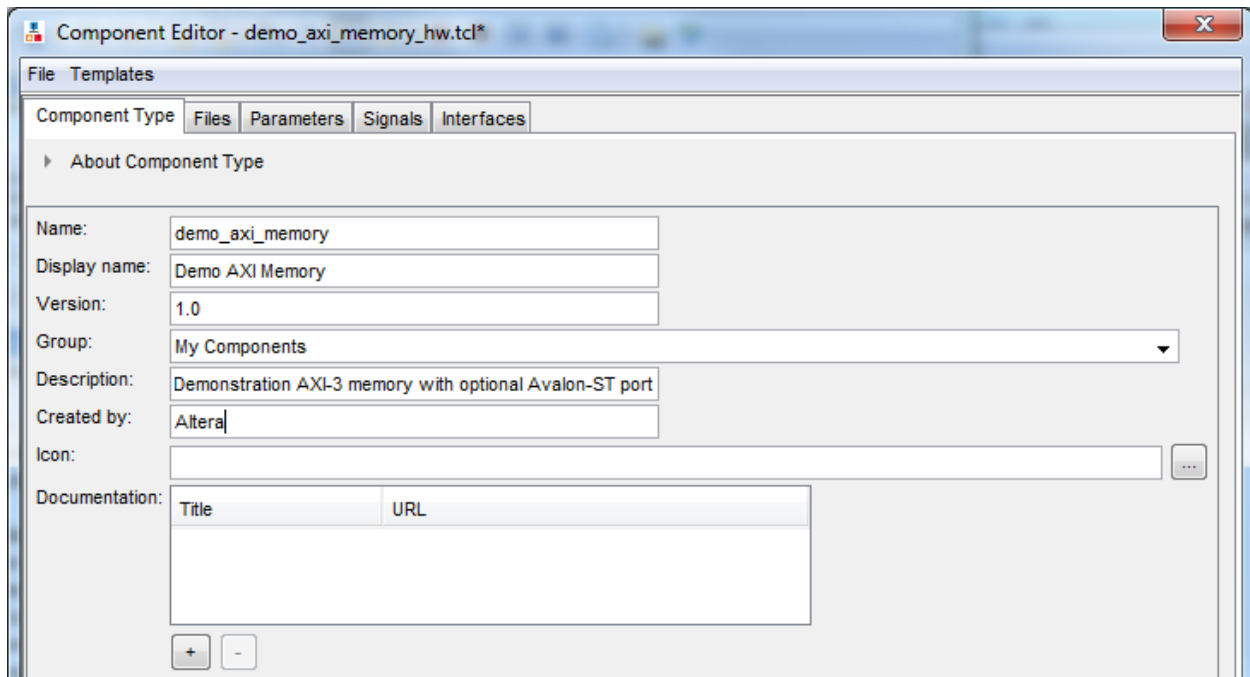
## Specifying Basic Component Information

The **Component Type** tab in the Component Editor allows you to specify the following information about the component:

- **Name**—Specifies the name used in the `_hw.tcl` filename, as well as in the top-level module name when you create a synthesis wrapper file for a non HDL-based component.
- **Display name**—Identifies the component in the parameter editor, which you use to configure and instance of the component, and also appears in the IP Catalog under **Project** and on the **System Contents** tab.
- **Version**—Specifies the version number of the component.
- **Group**—Represents the category of the component in the list of available components in the IP Catalog. You can select an existing group from the list, or define a new group by typing a name in the **Group** box. Separating entries in the **Group** box with a slash defines a subcategory. For example, if you type **Memories and Memory Controllers/On-Chip**, the component appears in the IP Catalog under the **On-Chip** group, which is a subcategory of the **Memories and Memory Controllers** group. If you save the component in the project directory, the component appears in the IP Catalog in the group you specified under **Project**. Alternatively, if you save the component in the Quartus II installation directory, the component appears in the specified group under **IP Catalog**.
- **Description**—Allows you to describe the component. This description appears when the user views the component details.
- **Created By**—Allows you to specify the author of the component.
- **Icon**—Allows you to enter the relative path to an icon file (`.gif`, `.jpg`, or `.png` format) that represents the component and appears as the header in the parameter editor for the component. The default image is the Altera MegaCore function icon.
- **Documentation**—Allows you to add links to documentation for the component, and appears when you right-click the component in the IP Catalog, and then select **Details**.
  - To specify an Internet file, begin your path with `http://`, for example: `http://mydomain.com/datasheets/my_memory_controller.html`.
  - To specify a file in the file system, begin your path with `file:///` for Linux, and `file://` for Windows; for example (Windows): `file:///company_server/datasheets my_memory_controller.pdf`.

**Figure 6-1: Component Type Tab in the Component Editor**

The **Display name**, **Group**, **Description**, **Created By**, **Icon**, and **Documentation** entries are optional.



When you use the Component Editor to create a component, it writes this basic component information in the **\_hw.tcl file**. The example below shows the component hardware Tcl code related to the entries for the **Component Type** tab in figure above. The `package require` command specifies the Quartus II software version that Qsys uses to create the **\_hw.tcl** file, and ensures compatibility with this version of the Qsys API in future ACDS releases.

#### Example 6-1: **\_hw.tcl** Created from Entries in the Component Type Tab

The component defines its basic information with various module properties using the `set_module_property` command. For example, `set_module_property NAME` specifies the name of the component, while `set_module_property VERSION` allows you to specify the version of the component. When you apply a version to the **\_hw.tcl** file, it allows the file to behave exactly the same way in future releases of the Quartus II software.

```
# request TCL package from ACDS 14.0
package require -exact qsys 14.0

# demo_axi_memory

set_module_property DESCRIPTION \
"Demo AXI-3 memory with optional Avalon-ST port"

set_module_property NAME demo_axi_memory
set_module_property VERSION 1.0
set_module_property GROUP "My Components"
set_module_property AUTHOR Altera
set_module_property DISPLAY_NAME "Demo AXI Memory"
```

### Related Information

- [Component Interface Tcl Reference](#) on page 9-1

## Specifying Files for Synthesis and Simulation

The **Files** tab in the Component Editor allows you to specify files for synthesis and simulation. If you already have HDL code that describes the Qsys component that you want to create, you can specify the files on the **Files** tab. If you have not yet created the HDL code that describes the component, but you have identified the signals and parameters that you want in the component, you can use the **Files** tab to create a top-level HDL template file. The Component Editor generates the appropriate `_hw.tcl` commands to specify the files. You can also write your own `hw.tcl` file with the same commands, if you are not using the Component Editor.

A component uses filesets to specify the different sets of files that can be generated for an instance of the component. The supported fileset types are: `QUARTUS_SYNTH`, for synthesis and compilation in the Quartus II software, `SIM_VERILOG`, for Verilog HDL simulation, and `SIM_VHDL`, for VHDL simulation.

In a `_hw.tcl` file, you add a fileset to a component with the `add_fileset` command. You then list specific files with the `add_fileset_file` command, which adds the specified files to the most recently declared fileset. The `add_fileset_property` command allows you to add properties such as `TOP_LEVEL`, which specifies the top-level HDL module for the component.

You can populate a fileset with a fixed list of files, add different files based on a parameter value, or even generate an HDL file with a custom HDL generator function outside of the `_hw.tcl` file.

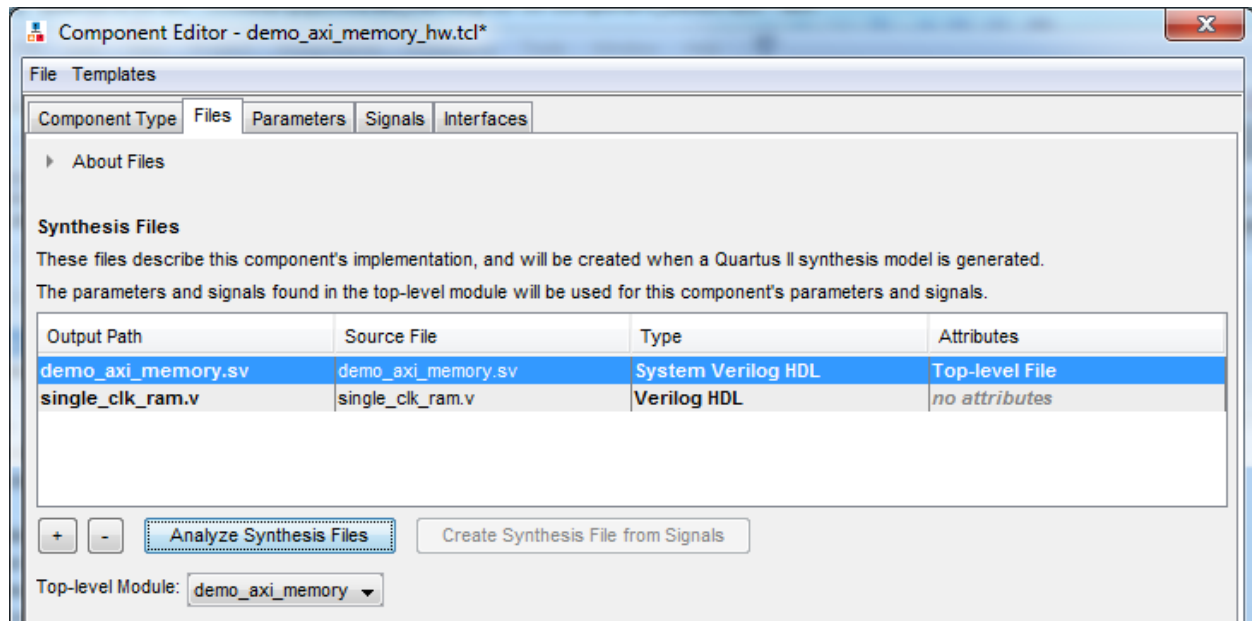
## Specifying HDL Files for Synthesis

In the Component Editor, you can add HDL files and other support files that should be included when this component is created to the list of **Synthesis Files** by clicking **+**, and then selecting the files in the **Open** dialog box.

A component must specify an HDL file as the top-level file, which contains the top-level module. The **Synthesis Files** list may also include supporting HDL files, such as timing constraints, or other files required to successfully synthesize and compile in the Quartus II software. The synthesis files for a component are copied to the generation output directory during Qsys system generation.

**Figure 6-2: Using HDL Files to Define a Component**

In the **Synthesis Files** section on the **Files** tab, the **demo\_axi\_memory.sv** file is selected as the top-level file for the component.



## Creating a New HDL File for Synthesis

If you do not already have an HDL implementation of the component, you can use the Component Editor to define the component, and then create a simple top-level synthesis file containing the signals and parameters for the component. You can then edit this HDL file to add the logic that directs the component's behavior.

To begin, you first specify the information about the component on the **Parameters**, **Signals**, and **Interfaces** tabs. Then, you return to the **Files** tab to create an HDL file by clicking **Create Synthesis File from Signals**. The Component Editor creates an HDL file from the specified parameters and signals.

## Analyzing Synthesis Files

After the top-level HDL file is specified in the Component Editor, click **Analyze Synthesis Files** to analyze the parameters and signals in the top-level, and then select the top-level module from the **Top Level Module** list. If there is a single module or entity in the HDL file, Qsys automatically populates the **Top-level Module** list.

Once analysis is complete and the top-level module is selected, the parameters and signals found in the top-level module are used as the parameters and signals for the component, and you can view them on the **Parameters** and **Signals** tabs. The Component Editor may report errors or warnings at this stage, because the signals and interfaces are not yet fully defined.

**Note:** At this stage in the Component Editor flow, you cannot add or remove parameters or signals created from a specified HDL file without editing the HDL file itself.

The synthesis files are added to a fileset with the name `QUARTUS_SYNTH` and type `QUARTUS_SYNTH` in the `_hw.tcl` file created by the Component Editor. The top-level module is used to specify the `TOP_LEVEL`

fileset property. Each synthesis file is individually added to the fileset. If the source files are saved in a different directory from the working directory where the Component Editor is launched and the `_hw.tcl` is located, you can use standard fixed or relative path notation to identify the file location for the `PATH` variable.

### Example 6-2: `_hw.tcl` Created from Entries in the Files tab in the Synthesis Files Section

```
# file sets

add_fileset QUARTUS_SYNTH QUARTUS_SYNTH "" ""
set_fileset_property QUARTUS_SYNTH TOP_LEVEL demo_axi_memory

add_fileset_file demo_axi_memory.sv
SYSTEM_VERILOG PATH demo_axi_memory.sv

add_fileset_file single_clk_ram.v VERILOG PATH single_clk_ram.v
```

#### Related Information

[Specifying HDL Files for Synthesis](#) on page 6-9

[Component Interface Tcl Reference](#) on page 9-1

## Naming HDL Signals for Automatic Interface and Type Recognition

If you create the component's top-level HDL file before using the Component Editor, the Component Editor recognizes the interface and signal types based on the signal names in the source HDL file. This auto-recognition feature eliminates the task of manually assigning each interface and signal type in the Component Editor.

To enable this auto-recognition feature, you must create signal names using the following naming convention:

*<interface type prefix>\_<interface name>\_<signal type>*

Specifying an interface name with *<interface name>* is optional if you have only one interface of each type in the component definition. For interfaces with only one signal, such as clock and reset inputs, the *<interface type prefix>* is also optional.

**Table 6-2: Interface Type Prefixes for Automatic Signal Recognition**

When the Component Editor recognizes a valid prefix and signal type for a signal, it automatically assigns an interface and signal type to the signal based on the naming convention. If no interface name is specified for a signal, you can choose an interface name on the **Interfaces** tab in the Component Editor.

Interface Prefix	Interface Type
asi	Avalon-ST sink (input)
aso	Avalon-ST source (output)
avm	Avalon-MM master
avs	Avalon-MM slave



Interface Prefix	Interface Type
axm	AXI master
axs	AXI slave
apm	APB master
aps	APB slave
coe	Conduit
csi	Clock Sink (input)
cso	Clock Source (output)
inr	Interrupt receiver
ins	Interrupt sender
ncm	Nios II custom instruction master
ncs	Nios II custom instruction slave
rsi	Reset sink (input)
rso	Reset source (output)
tcm	Avalon-TC master
tcs	Avalon-TC slave

Refer to the *Avalon Interface Specifications* or the *AMBA Protocol Specification* for the signal types available for each interface type.

#### Related Information

- [Avalon Interface Specifications](#)
- [AMBA Protocol Specification](#)

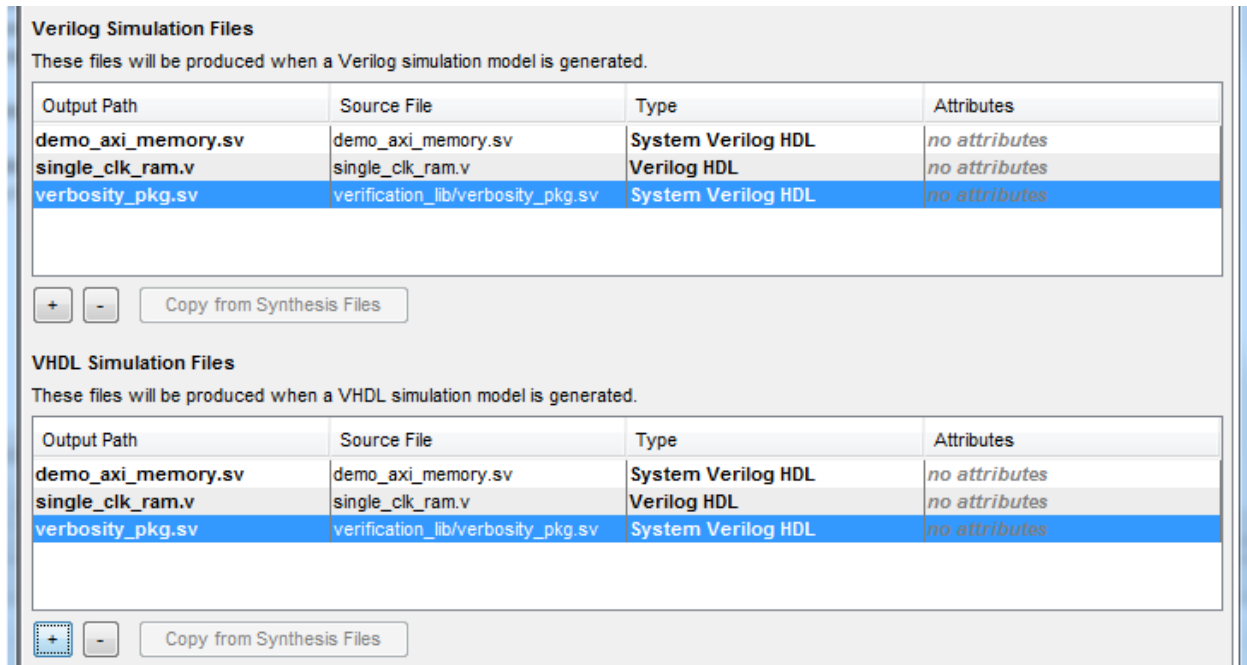
## Specifying Files for Simulation

To support Qsys system generation for simulation, a component must specify the VHDL or Verilog simulation files. Simulation files are generated when a user adds the component to a Qsys system and chooses to generate Verilog or VHDL simulation files. In most cases, these files are the same as the synthesis files. If there are simulation-specific HDL files or simulation models, you can use them in addition to, or in place of the synthesis files. To use your synthesis files as your simulation files in the Component Editor, on the **Files** tab, click **Copy From Synthesis Files** to copy the list of synthesis files to the **Verilog Simulation Files** or **VHDL Simulation Files** lists.

You specify the simulation files in a similar way as the synthesis files with the fileset commands in a `_hw.tcl` file. The code example below shows `SIM_VERILOG` and `SIM_VHDL` filesets for Verilog and VHDL simulation output files. In this example, the same Verilog files are used for both Verilog and VHDL outputs, and there is one additional System Verilog file added. This method works for designers of Verilog IP to support users who want to generate a VHDL top-level simulation file when they have a mixed-language simulation tool and license that can read the Verilog output for the component.

**Note:** The order that you add files to the fileset determines the order of compilation. For VHDL filesets with VHDL files, you must add the files bottom-up, adding the top-level file last.

**Figure 6-3: Specifying the Simulation Output Files on the Files Tab**



**Example 6-3: `_hw.tcl` Created from Entries in the Files tab in the Simulation Files Section**

```
add_filesset SIM_VERILOG SIM_VERILOG "" ""
set_filesset_property SIM_VERILOG TOP_LEVEL demo_axi_memory
add_filesset_file single_clk_ram.v VERILOG PATH single_clk_ram.v

add_filesset_file verbosity_pkg.sv SYSTEM_VERILOG PATH \
verification_lib/verbosity_pkg.sv

add_filesset_file demo_axi_memory.sv SYSTEM_VERILOG PATH \
demo_axi_memory.sv

add_filesset SIM_VHDL SIM_VHDL "" ""
set_filesset_property SIM_VHDL TOP_LEVEL demo_axi_memory
set_filesset_property SIM_VHDL ENABLE_RELATIVE_INCLUDE_PATHS false

add_filesset_file demo_axi_memory.sv SYSTEM_VERILOG PATH \
demo_axi_memory.sv

add_filesset_file single_clk_ram.v VERILOG PATH single_clk_ram.v
```

```
add_fileset_file verbosity_pkg.sv SYSTEM_VERILOG PATH \
verification_lib/verbosity_pkg.sv
```

#### Related Information

- [Component Interface Tcl Reference](#) on page 9-1

## Including Internal Register Map Description in the .svd for Slave Interfaces Connected to an HPS Component

Qsys supports the ability for IP component designers to specify register map information on their slave interfaces. This allows components with slave interfaces that are connected to an HPS component to include their internal register description in the generated .svd file.

To specify their internal register map, the IP component designer must write and generate their own .svd file and attach it to the slave interface using the following command:

```
set_interface_property <slave interface> CMSIS_SVD_FILE <file path>
```

The CMSIS\_SVD\_VARIABLES interface property allows for variable substitution inside the .svd file. You can dynamically modify the character data of the .svd file by using the CMSIS\_SVD\_VARIABLES property.

### Example 6-4: Setting the CMSIS\_SVD\_VARIABLES Interface Property

For example, if you set the CMSIS\_SVD\_VARIABLES in the `_hw tcl` file, then in the .svd file if there is a variable `{width}` that describes the element `<size>${width}</size>`, it is replaced by `<size>23</size>` during generation of the .svd file. Note that substitution works only within character data (the data enclosed by `<element>...</element>`) and not on element attributes.

```
set_interface_property <interface name> \
CMSIS_SVD_VARIABLES "{width} {23}"
```

#### Related Information

[Component Interface Tcl Reference](#) on page 9-1

[CMSIS - Cortex Microcontroller Software](#)

## Specifying Component Parameters

Components can include parameterized HDL, which allows users of the component flexibility in meeting their system requirements. For example, a component may have a configurable memory size or data width, where one HDL implementation can be used in many different systems, each with unique parameters values.

The **Parameters** tab in the Component Editor allows you specify the parameters that are used to configure instances of the component in a Qsys system. You can specify various properties for each parameter that describe how the parameter is displayed and used. You can also specify a range of allowed values that are checked during the Validation phase. The **Parameters** table displays the HDL parameters that are declared in the top-level HDL module. If you have not yet created the top-level HDL file, the parameters that you create on the **Parameters** tab are included in the top-level synthesis file template created from the **Files** tab.

When the component includes HDL files, the parameters match those defined in the top-level module, and you cannot add or remove them on the **Parameters** tab. To add or remove the parameters, edit your HDL source, and then re-analyze the file.

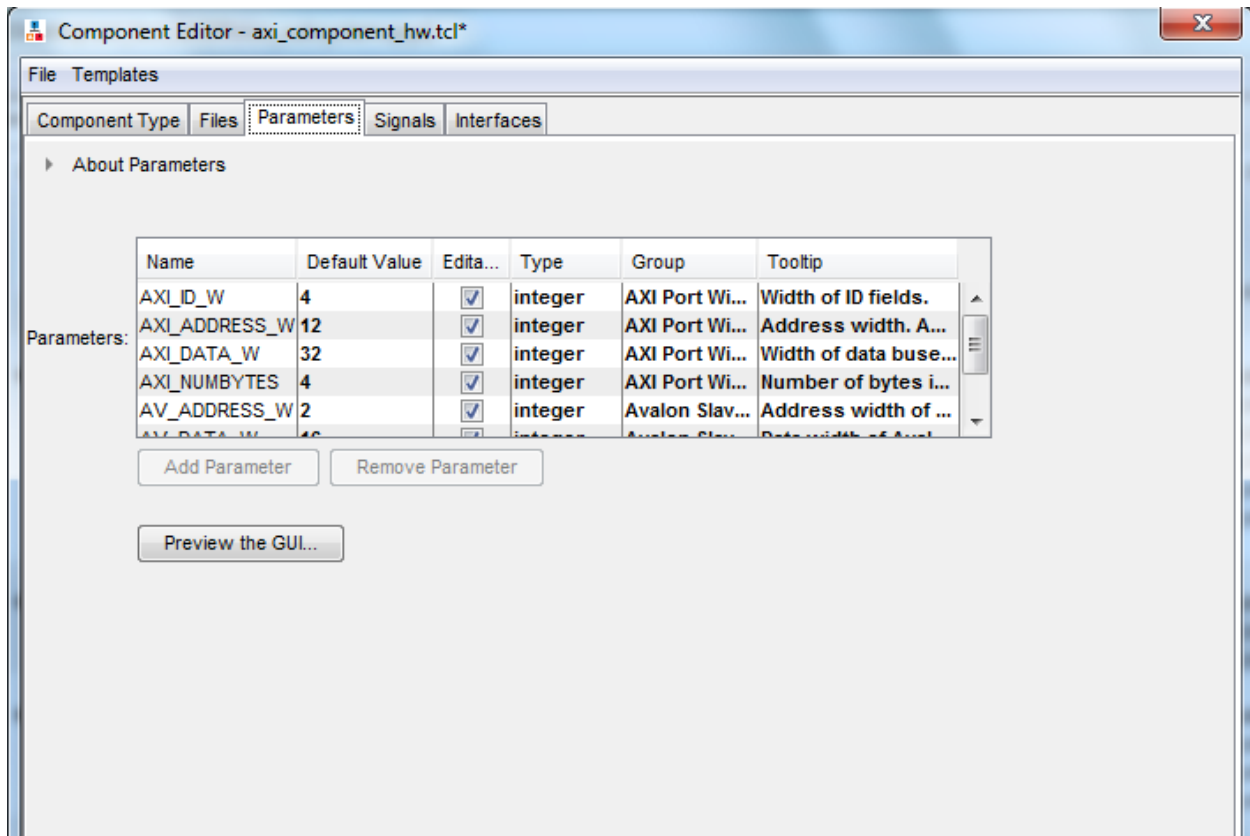
If you used the Component Editor to create a top-level template HDL file for synthesis, you can remove the newly-created file from the **Synthesis Files** list on the **Files** tab, make your parameter changes, and then re-analyze the top-level synthesis file.

You can use the **Parameters** table to specify the following information about each parameter:

- **Name**—Specifies the name of the parameter.
- **Default Value**—Sets the default value used in new instances of the component.
- **Editable**—Specifies whether or not the user can edit the parameter value.
- **Type**—Defines the parameter type as string, integer, boolean, std\_logic, logic vector, natural, or positive.
- **Group**—Allows you to group parameters in parameter editor.
- **Tooltip**—Allows you to add a description of the parameter that appears when the user of the component points to the parameter in the parameter editor.

Figure 6-4: Parameters Tab in the Components Editor

On the **Parameters** tab, you can click **Preview the GUI** at any time to see how the declared parameters appear in the parameter editor. Parameters with their default values appear with checks in the **Editable** column, indicating that users of this component are allowed to modify the parameter value. Editable parameters cannot contain computed expressions. You can group parameters under a common heading or section in the parameter editor with the **Group** column, and a tooltip helps users of the component understand the function of the parameter. Various parameter properties allow you to customize the component's parameter editor, such as using radio buttons for parameter selections, or displaying an image.



Example 6-5: \_hw.tcl Created from Entries in the Parameters Tab

In this example, the first `add_parameter` command includes commonly-specified properties. The `set_parameter_property` command specifies each property individually. The **Tooltip** column on the **Parameters** tab maps to the `DESCRIPTION` property, and there is an additional unused `UNITS` property created in the code. The `HDL_PARAMETER` property specifies that the value of the parameter is specified in the HDL instance wrapper when creating instances of the component. The **Group** column in the **Parameters** tab maps to the display items section with the `add_display_item` commands.

**Note:** If a parameter  $\langle n \rangle$  defines the width of a signal, the signal width must follow the format:  $\langle n-1 \rangle:0$ .

```
#
# parameters
#
add_parameter AXI_ID_W INTEGER 4 "Width of ID fields"
set_parameter_property AXI_ID_W DEFAULT_VALUE 4
set_parameter_property AXI_ID_W DISPLAY_NAME AXI_ID_W
set_parameter_property AXI_ID_W TYPE INTEGER
set_parameter_property AXI_ID_W UNITS None
set_parameter_property AXI_ID_W DESCRIPTION "Width of ID fields"
set_parameter_property AXI_ID_W HDL_PARAMETER true
add_parameter AXI_ADDRESS_W INTEGER 12
set_parameter_property AXI_ADDRESS_W DEFAULT_VALUE 12

add_parameter AXI_DATA_W INTEGER 32
...
#
# display items
#
add_display_item "AXI Port Widths" AXI_ID_W PARAMETER ""
```

**Note:** If an AXI slave's ID bit width is smaller than required for your system, the AXI slave response may not reach all AXI masters. The formula of an AXI slave ID bit width is calculated as follows:

$$\text{maximum\_master\_id\_width\_in\_the\_interconnect} + \log_2(\text{number\_of\_masters\_in\_the\_same\_interconnect})$$

For example, if an AXI slave connects to three AXI masters and the maximum AXI master ID length of the three masters is 5 bits, then the AXI slave ID is 7 bits, and is calculated as follows:

$$5 \text{ bits} + 2 \text{ bits} (\log_2(3 \text{ masters})) = 7$$

#### Related Information

- [Component Interface Tcl Reference](#) on page 9-1

## Allowed Ranges Parameter Property

In a component's `hw.tcl` file, you can specify valid ranges for parameters. In Qsys, validation checks each parameter value against the `ALLOWED_RANGES` property. If the values specified are outside of the allowed ranges, Qsys displays an error message. Specifying choices for the allowed values enables users of the component to choose the parameter value from a drop-down list or radio button in the parameter editor GUI instead of entering a value.

The `ALLOWED_RANGES` property is a list of valid ranges, where each range is a single value, or a range of values defined by a start and end value.

Table 6-3: ALLOWED\_RANGES Property Examples

ALLOWED_RANGES	Meaning
{a b c}	a, b, or c
{"No Control" "Single Control" "Dual Controls"}	Unique string values. Quotation marks are required if the strings include spaces
{1 2 4 8 16}	1, 2, 4, 8, or 16
{1:3}	1 through 3, inclusive
{1 2 3 7:10}	1, 2, 3, or 7 through 10 inclusive

**Related Information**

[Declaring Parameters with Custom \\_hw.tcl Commands](#) on page 6-19

**Types of Parameters**

Qsys uses the following parameter types: user parameters, system information parameters, and derived parameters.

[User Parameters](#) on page 6-18

[System Information Parameters](#) on page 6-18

[Derived Parameters](#) on page 6-19

**Related Information**

[Declaring Parameters with Custom \\_hw.tcl Commands](#) on page 6-19

**User Parameters**

User parameters are parameters that users of a component can control, and appear in the parameter editor for instances of the component. User parameters map directly to parameters in the component HDL. For user parameter code examples, such as `AXI_DATA_W` and `ENABLE_STREAM_OUTPUT`, refer to *Declaring Parameters with Custom hw.tcl Commands*.

**System Information Parameters**

A `SYSTEM_INFO` parameter is a parameter whose value is set automatically by the Qsys system. When you define a `SYSTEM_INFO` parameter, you provide an `information` type, and additional arguments.

For example, you can configure a parameter to store the clock frequency driving a clock input for your component. To do this, define the parameter as `SYSTEM_INFO` of type `CLOCK_RATE`:

```
set_parameter_property <param> SYSTEM_INFO CLOCK_RATE
```

You then set the name of the clock interface as the `SYSTEM_INFO_ARG` argument:

```
set_parameter_property <param> SYSTEM_INFO_ARG <clkname>
```

## Derived Parameters

Derived parameter values are calculated from other parameters during the Elaboration phase, and are specified in the `hw.tcl` file with the `DERIVED` property. Derived parameter values are calculated from other parameters during the Elaboration phase, and are specified in the `hw.tcl` file with the `DERIVED` property. For example, you can derive a clock period parameter from a data rate parameter. Derived parameters are sometimes used to perform operations that are difficult to perform in HDL, such as using logarithmic functions to determine the number of address bits that a component requires.

### Related Information

[Declaring Parameters with Custom `\_hw.tcl` Commands](#) on page 6-19

### Parameterized Parameter Widths

Qsys allows a `std_logic_vector` parameter to have a width that is defined by another parameter, similar to derived parameters. The width can be a constant or the name of another parameter.

## Declaring Parameters with Custom `_hw.tcl` Commands

The example below illustrates a custom `_hw.tcl` file, with more advanced parameter commands than those generated when you specify parameters in the Component Editor. Commands include the `ALLOWED_RANGES` property to provide a range of values for the `AXI_ADDRESS_W` (**Address Width**) parameter, and a list of parameter values for the `AXI_DATA_W` (**Data Width**) parameter. This example also shows the parameter `AXI_NUMBYTES` (**Data width in bytes**) parameter; that uses the `DERIVED` property. In addition, these commands illustrate the use of the `GROUP` property, which groups some parameters under a heading in the parameter editor GUI. You use the `ENABLE_STREAM_OUTPUT_GROUP` (**Include Avalon streaming source port**) parameter to enable or disable the optional Avalon-ST interface in this design, and is displayed as a check box in the parameter editor GUI because the parameter is of type `BOOLEAN`. Refer to figure below to see the parameter editor GUI resulting from these `hw.tcl` commands.

### Example 6-6: Parameter Declaration

In this example, the `AXI_NUMBYTES` parameter is derived during the Elaboration phase based on another parameter, instead of being assigned to a specific value. `AXI_NUMBYTES` describes the number of bytes in a word of data. Qsys calculates the `AXI_NUMBYTES` parameter from the `DATA_WIDTH` parameter by dividing by 8. The `_hw.tcl` code defines the `AXI_NUMBYTES` parameter as a derived parameter, since its value is calculated in an elaboration callback procedure. The `AXI_NUMBYTES` parameter value is not editable, because its value is based on another parameter value.

```
add_parameter AXI_ADDRESS_W INTEGER 12

set_parameter_property AXI_ADDRESS_W DISPLAY_NAME \
"AXI Slave Address Width"

set_parameter_property AXI_ADDRESS_W DESCRIPTION \
"Address width."

set_parameter_property AXI_ADDRESS_W UNITS bits
set_parameter_property AXI_ADDRESS_W ALLOWED_RANGES 4:16
set_parameter_property AXI_ADDRESS_W HDL_PARAMETER true

set_parameter_property AXI_ADDRESS_W GROUP \
"AXI Port Widths"

add_parameter AXI_DATA_W INTEGER 32
```



```

set_parameter_property AXI_DATA_W DISPLAY_NAME "Data Width"

set_parameter_property AXI_DATA_W DESCRIPTION \
"Width of data buses."

set_parameter_property AXI_DATA_W UNITS bits

set_parameter_property AXI_DATA_W ALLOWED_RANGES \
{8 16 32 64 128 256 512 1024}

set_parameter_property AXI_DATA_W HDL_PARAMETER true
set_parameter_property AXI_DATA_W GROUP "AXI Port Widths"

add_parameter AXI_NUMBYTES INTEGER 4
set_parameter_property AXI_NUMBYTES DERIVED true

set_parameter_property AXI_NUMBYTES DISPLAY_NAME \
"Data Width in bytes; Data Width/8"

set_parameter_property AXI_NUMBYTES DESCRIPTION \
"Number of bytes in one word"

set_parameter_property AXI_NUMBYTES UNITS bytes
set_parameter_property AXI_NUMBYTES HDL_PARAMETER true
set_parameter_property AXI_NUMBYTES GROUP "AXI Port Widths"

add_parameter ENABLE_STREAM_OUTPUT BOOLEAN true

set_parameter_property ENABLE_STREAM_OUTPUT DISPLAY_NAME \
"Include Avalon Streaming Source Port"

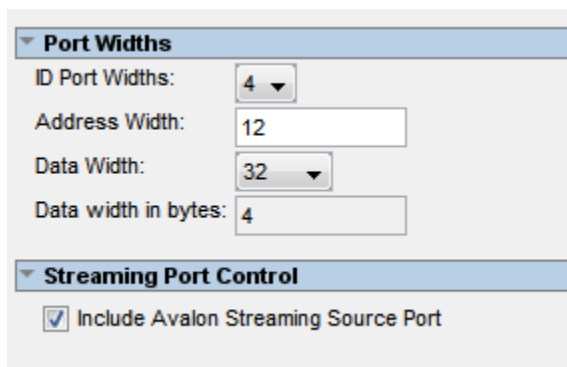
set_parameter_property ENABLE_STREAM_OUTPUT DESCRIPTION \
"Include optional Avalon-ST source (default),\
or hide the interface"

set_parameter_property ENABLE_STREAM_OUTPUT GROUP \
"Streaming Port Control"

...

```

Figure 6-5: Resulting Parameter Editor GUI from Parameter Declarations



#### Related Information

[Controlling Interfaces Dynamically with an Elaboration Callback](#) on page 6-26

[Component Interface Tcl Reference](#) on page 9-1

## Validating Parameter Values with a Validation Callback

You can use a validation callback procedure to validate parameter values with more complex validation operations than the `ALLOWED_RANGES` property allows. You define a validation callback by setting the `VALIDATION_CALLBACK` module property to the name of the Tcl callback procedure that runs during the validation phase. In the validation callback procedure, the current parameter values is queried, and warnings or errors are reported about the component's configuration.

### Example 6-7: Demo AXI Memory Example

If the optional Avalon streaming interface is enabled, then the control registers must be wide enough to hold an AXI RAM address, so the designer can add an error message to ensure that the user enters allowable parameter values.

```
set_module_property VALIDATION_CALLBACK validate
proc validate {} {
  if {
    [get_parameter_value ENABLE_STREAM_OUTPUT ] &&
    ([get_parameter_value AXI_ADDRESS_W] >
     [get_parameter_value AV_DATA_W])
  }
  send_message error "If the optional Avalon streaming port\
is enabled, the AXI Data Width must be equal to or greater\
than the Avalon control port Address Width"
}
```

#### Related Information

[Component Interface Tcl Reference](#) on page 9-1

[Demo AXI Memory Example](#)

## Specifying Interface and Signal Types

The **Signals** tab in the Components Editor allows you to specify the interface and signal type of each signal in the component. When you add HDL files to the **Synthesis Files** table on the **Files** tab, and then click **Analyze Synthesis Files**, the signals on the top-level module appear on the **Signals** tab.

If you have not yet created your top-level HDL file, you can click **Add Signal** to specify each top-level signal in the component. For each signal that you add, you must provide the appropriate values in the **Name**, **Interface**, **Signal Type**, **Width**, and **Direction** columns. You can use the error and warning messages at the bottom of the window to guide your selections. You can edit the signal name by double-clicking the **Name** column, and then typing the new name.

After you have analyzed the component's top-level HDL file on the **Files** tab, you cannot add or remove signals or change the signal names on the **Signals** tab. To change the signals, edit your HDL source, and then click **Generate Synthesis File from Signals**.

If you used the Component Editor to create a top-level template HDL file for synthesis, you can remove the newly-created file from the **Synthesis Files** list on the **Files** tab, make your signal changes, and then re-analyze the top-level synthesis file.

The **Interface** column allows you assign a signal to an interface. Each signal must belong to an interface and be assigned a legal signal type for that interface. To create a new interface of a specific type, select **new** *<interface type>* from the list; this new interface then become available in the list for subsequent signal assignments. You can highlight all of the signals in an interface and then select an Interface from the list to apply the Interface name to each signal in the interface.

You edit the interface name on the **Interface** tab; you cannot edit the interface name on the **Signals** tab.

Figure 6-6: Signals Tab in the Qsys Components Editor

Name	Interface	Signal Type	Width	Direction
clk	clock	clk	1	input
reset_n	reset	reset_n	1	input
axs_awid	altera_axi_slave	awid	AXI_ID...	input
axs_awaddr	altera_axi_slave	awaddr	AXI_AD...	input
axs_awlen	clock	awlen	4	input
axs_awsize	reset	awsize	3	input
axs_awburst	avalon_streaming_source_0	awburst	2	input
axs_awlock	altera_axi_slave	awlock	2	input
axs_awcache	new Avalon Memory Mapped Master...	awcache	4	input
axs_awprot	new Avalon Memory Mapped Slave...	awprot	3	input
axs_awvalid	new Avalon Streaming Source...	awvalid	1	input
axs_awready	new Avalon Streaming Sink...	awready	1	output
axs_wid	new Avalon Memory Mapped Tristate Slave...	wid	AXI_ID...	input
axs_wdata	new AXI Master...	wdata	AXI_DA...	input
axs_wstrb	new AXI Slave...	wstrb	AXI_NU...	input
axs_wlast	new AXI4 Master...	wlast	1	input
axs_wvalid	new AXI4 Slave...	wvalid	1	input
axs_wready	new Clock Output...	wready	1	output
axs_bid	new Clock Input...	bid	AXI_ID...	output
axs_bresp	new Conduit...	bresp	2	output
axs_bvalid	new Interrupt Receiver...	bvalid	1	output
axs_bready	new Interrupt Sender...	bready	1	input
axs_arid	new Custom Instruction Master...	arid	AXI_ID...	input
axs_araddr	new Custom Instruction Slave...	araddr	AXI_AD...	input

#### Related Information

[Adding Interfaces and Managing Interface Settings](#) on page 6-22

[Component Interface Tcl Reference](#) on page 9-1

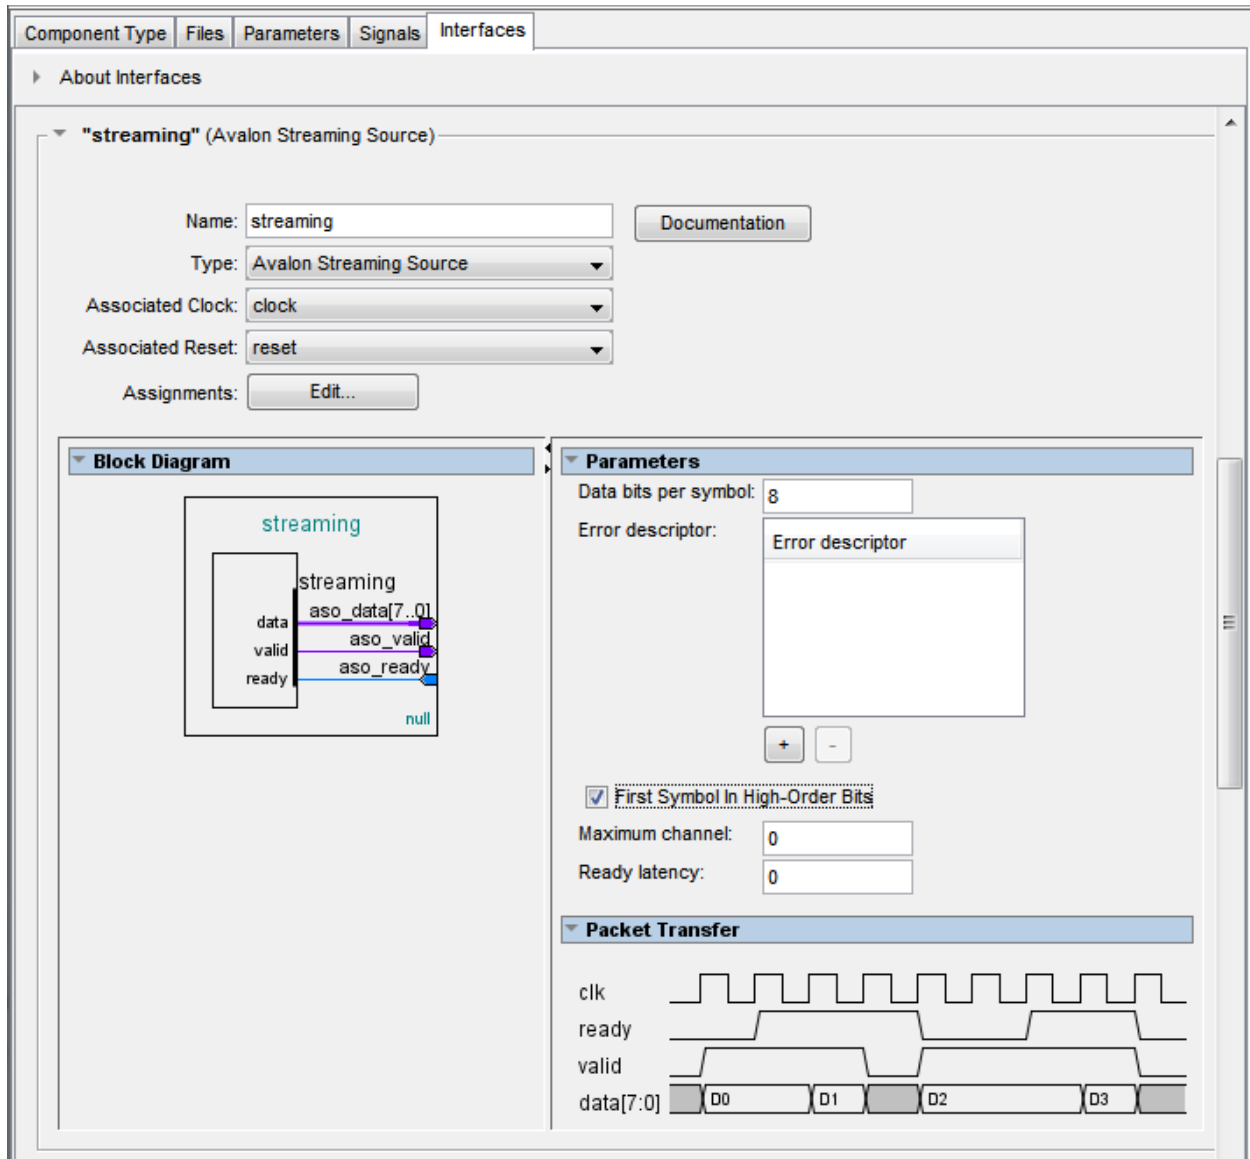
## Adding Interfaces and Managing Interface Settings

The **Interfaces** tab in the Component Editor allows you to manage settings for each interface of the component. The interface name appears on the **Signals** tab, and in the Qsys **System Contents** tab when the component is added to a system.

You can configure the type and properties of each interface. Some interfaces display waveforms that illustrate the timing for the interface. If you update timing parameters, the waveforms automatically update.

You add additional interfaces by clicking **Add Interface**, and then you must specify the signals for the added interface on the **Signals** tab. You can remove interfaces that have no assigned signals by clicking **Remove Interfaces With No Signals**.

Figure 6-7: Avalon Streaming Source Interface on the Interfaces Tab



**Example 6-8: \_hw.tcl Created from Entries in the Interface Tab**

Each interface is created with the `add_interface` command. You specify the properties of each interface with the `set_interface_property` command. The interface's signals are specified with the `add_interface_port` command.

```
#
# connection point clock
#
```

```

add_interface clock clock end
set_interface_property clock clockRate 0
set_interface_property clock ENABLED true

add_interface_port clock clk clk Input 1

#
# connection point reset
#
add_interface reset reset end
set_interface_property reset associatedClock clock
set_interface_property reset synchronousEdges DEASSERT
set_interface_property reset ENABLED true

add_interface_port reset reset_n reset_n Input 1

#
# connection point streaming
#
add_interface streaming avalon_streaming start
set_interface_property streaming associatedClock clock
set_interface_property streaming associatedReset reset
set_interface_property streaming dataBitsPerSymbol 8
set_interface_property streaming errorDescriptor ""
set_interface_property streaming firstSymbolInHighOrderBits true
set_interface_property streaming maxChannel 0
set_interface_property streaming readyLatency 0
set_interface_property streaming ENABLED true

add_interface_port streaming aso_data data Output 8
add_interface_port streaming aso_valid valid Output 1
add_interface_port streaming aso_ready ready Input 1

#
# connection point slave
#
add_interface slave axi end
set_interface_property slave associatedClock clock
set_interface_property slave associatedReset reset
set_interface_property slave readAcceptanceCapability 1
set_interface_property slave writeAcceptanceCapability 1
set_interface_property slave combinedAcceptanceCapability 1
set_interface_property slave readDataReorderingDepth 1
set_interface_property slave ENABLED true

add_interface_port slave axs_awid awid Input AXI_ID_W
...
add_interface_port slave axs_rresp rresp Output 2

```

**Table 6-4: AXI Master and Slave Parameters**

Qsys refers to AXI interface parameters to build AXI interconnect. If these parameter settings are incompatible with the component's HDL behavior, Qsys interconnect and transactions may not work correctly. To prevent unexpected interconnect behavior, you must set the AXI component parameters.

AXI Master Parameters	AXI Slave Parameters
readIssuingCapability	readAcceptanceCapability
writeIssuingCapability	writeAcceptanceCapability
combinedIssuingCapability	combinedAcceptanceCapability

AXI Master Parameters	AXI Slave Parameters
	readDataReorderingDepth

#### Related Information

- [Component Interface Tcl Reference](#) on page 9-1

## Creating Custom \_hw.tcl Interface Settings and Properties

### Example 6-9: Clock, Reset, AXI Slave, and Avalon Streaming Interfaces Using Variables

The clock, reset, AXI slave, and Avalon streaming interfaces use variables for the interface names to make the file easier to read and update. The interface declaration statement includes the name, type, and direction of the interface, as well as the associated clock and reset interfaces. Also in the example below, some of the AXI memory signals use parameters to specify their width.

```

set CLOCK_INTERFACE "clk"
add_interface $CLOCK_INTERFACE clock end
add_interface_port $CLOCK_INTERFACE clk clk Input 1

set RESET_INTERFACE "reset"
add_interface $RESET_INTERFACE reset end
set_interface_property $RESET_INTERFACE associatedClock clk
set_interface_property $RESET_INTERFACE synchronousEdges DEASSERT
add_interface_port reset reset_n reset_n Input 1

set SLAVE_INTERFACE "slave"
add_interface $SLAVE_INTERFACE axi end
set_interface_property $SLAVE_INTERFACE associatedClock "clk"
set_interface_property $SLAVE_INTERFACE associatedReset "reset"

set_interface_property $SLAVE_INTERFACE \
readAcceptanceCapability 1
...
add_interface_port $SLAVE_INTERFACE axs_wdata wdata \
Input AXI_DATA_W

add_interface_port $SLAVE_INTERFACE axs_wstrb wstrb \
Input AXI_NUMBYTES

add_interface_port $SLAVE_INTERFACE axs_wlast wlast Input 1
...
set STREAMING_INTERFACE "streaming"
add_interface $STREAMING_INTERFACE avalon_streaming start
set_interface_property $STREAMING_INTERFACE associatedClock "clk"
...
add_interface_port $STREAMING_INTERFACE aso_data data Output 8
add_interface_port $STREAMING_INTERFACE aso_valid valid Output 1
add_interface_port $STREAMING_INTERFACE aso_ready ready Input 1

```

#### Related Information

- [Component Interface Tcl Reference](#) on page 9-1

## Controlling Interfaces Dynamically with an Elaboration Callback

You can allow user parameters to dynamically control your component's behavior with an elaboration callback procedure during the elaboration phase. Using an elaboration callback allows you to change interface properties, remove interfaces, or add new interfaces as a function of a parameter value. You define an elaboration callback by setting the module property `ELABORATION_CALLBACK` to the name of the Tcl callback procedure that runs during the elaboration phase. In the callback procedure, you can query the parameter values of the component instance, and then change the interfaces accordingly.

### Example 6-10: Avalon-ST Source Interface Optionally Included in a Component Specified with an Elaboration Callback

```
set_module_property ELABORATION_CALLBACK elaborate

proc elaborate {} {

    # Optionally disable the Avalon- ST data output

    if{[ get_parameter_value ENABLE_STREAM_OUTPUT] == "false" }{
        set_port_property aso_data      termination true
        set_port_property aso_valid     termination true
        set_port_property aso_ready    termination true
        set_port_property aso_ready    termination_value 0
    }
    # Calculate the Data Bus Width in bytes

    set bytewidth_var [expr [get_parameter_value AXI_DATA_W]/8]
    set_parameter_value AXI_NUMBYTES $bytewidth_var
}
```

#### Related Information

- [Creating Custom \\_hw.tcl Interface Settings and Properties](#) on page 6-25
- [Validating Parameter Values with a Validation Callback](#) on page 6-21
- [Component Interface Tcl Reference](#) on page 9-1

## Controlling File Generation Dynamically with Parameters and a Fileset Callback

You can use a fileset callback to control which files are created in the output directories during the generation phase based on parameter values, instead of providing a fixed list of files. In a callback procedure, you can query the values of the parameters and use them to generate the appropriate files. To define a fileset callback, you specify a callback procedure name as an argument in the `add_fileset` command. You can use the same fileset callback procedure for all of the filesets, or create separate procedures for synthesis and simulation, or Verilog and VHDL.

### Example 6-11: Fileset Callback Using Parameters to Control Filesets in Two Different Ways

The `RAM_VERSION` parameter chooses between two different source files to control the implementation of a RAM block. For the top-level source file, a custom Tcl routine generates HDL that

optionally includes control and status registers, depending on the value of the `CSR_ENABLED` parameter.

During the generation phase, Qsys creates a top-level Qsys system HDL wrapper module to instantiate the component top-level module, and applies the component's parameters, for any parameter whose parameter property `HDL_PARAMETER` is set to true.

```
#Create synthesis fileset with fileset_callback and set top level
add_fileset my_synthesis_fileset QUARTUS_SYNTH fileset_callback

set_fileset_property my_synthesis_fileset TOP_LEVEL \
demo_axi_memory

# Create Verilog simulation fileset with same fileset_callback
# and set top level
add_fileset my_verilog_sim_fileset SIM_VERILOG fileset_callback

set_fileset_property my_verilog_sim_fileset TOP_LEVEL \
demo_axi_memory

# Add extra file needed for simulation only
add_fileset_file verbosity_pkg.sv SYSTEM_VERILOG PATH \
verification_lib/verbosity_pkg.sv

# Create VHDL simulation fileset (with Verilog files
# for mixed-language VHDL simulation)
add_fileset my_vhdl_sim_fileset SIM_VHDL fileset_callback
set_fileset_property my_vhdl_sim_fileset TOP_LEVEL demo_axi_memory

add_fileset_file verbosity_pkg.sv SYSTEM_VERILOG PATH
verification_lib/verbosity_pkg.sv

# Define parameters required for fileset_callback
add_parameter RAM_VERSION INTEGER 1
set_parameter_property RAM_VERSION ALLOWED_RANGES {1 2}
set_parameter_property RAM_VERSION HDL_PARAMETER false
add_parameter CSR_ENABLED BOOLEAN enable
set_parameter_property CSR_ENABLED HDL_PARAMETER false

# Create Tcl callback procedure to add appropriate files to
# filesets based on parameters
proc fileset_callback { entityName } {
    send_message INFO "Generating top-level entity $entityName"
    set ram [get_parameter_value RAM_VERSION]
    set csr_enabled [get_parameter_value CSR_ENABLED]

    send_message INFO "Generating memory
implementation based on RAM_VERSION $ram    "

    if {$ram == 1} {
        add_fileset_file single_clk_ram1.v VERILOG PATH \
single_clk_ram1.v
    } else {
        add_fileset_file single_clk_ram2.v VERILOG PATH \
single_clk_ram2.v
    }
}

send_message INFO "Generating top-level file for \
CSR_ENABLED $csr_enabled"
```



```
generate_my_custom_hdl $csr_enabled demo_axi_memory_gen.sv

add_fileset_file demo_axi_memory_gen.sv VERILOG PATH \
demo_axi_memory_gen.sv
}
```

### Related Information

[Specifying Files for Synthesis and Simulation](#) on page 6-9

[Component Interface Tcl Reference](#) on page 9-1

## Creating a Composed Component or Subsystem

A composed component is a subsystem containing instances of other components. Unlike an HDL-based component, a composed component's HDL is created by generating HDL for the components in the subsystem, in addition to the Qsys interconnect to connect the subsystem instances.

You can add child instances in a composition callback of the `_hw.tcl` file.

With a composition callback, you can also instantiate and parameterize sub-components as a function of the composed component's parameter values. You define a composition callback by setting the `COMPOSITION_CALLBACK` module property to the name of the composition callback procedures.

A composition callback replaces the validation and elaboration phases. HDL for the subsystem is generated by generating all of the sub-components and the top-level that combines them.


To connect instances of your component, you must define the component's interfaces. Unlike an HDL-based component, a composed component does not directly specify the signals that are exported. Instead, interfaces of submodules are chosen as the external interface, and each internal interface's ports are connected through the exported interface.

Exporting an interface means that you are making the interface visible from the outside of your component, instead of connecting it internally. You can set the `EXPORT_OF` property of the externally visible interface from the main program or the composition callback, to indicate that it is an exported view of the submodule's interface.

Exporting an interface is different than defining an interface. An exported interface is an exact copy of the subcomponent's interface, and you are not allowed to change properties on the exported interface. For example, if the internal interface is a 32-bit or 64-bit master without bursting, then the exported interface is the same. An interface on a subcomponent cannot be exported and also connected within the subsystem.

When you create an exported interface, the properties of the exported interface are copied from the subcomponent's interface without modification. Ports are copied from the subcomponent's interface with only one modification; the names of the exported ports on the composed component are chosen to ensure that they are unique.

Figure 6-8: Top-Level of a Composed Component



**Example 6-12: Composed\_hw.tcl File that Instantiates Two Sub-Components**

Qsys connects the components, and also connects the clocks and resets. Note that clock and reset bridge components are required to allow both sub-components to see common clock and reset inputs.

```

package require -exact qsys 14.0
set_module_property name my_component
set_module_property COMPOSITION_CALLBACK composed_component

proc composed_component {} {
    add_instance clk altera_clock_bridge
    add_instance reset altera_reset_bridge
    add_instance regs my_regs_microcore
    add_instance phy my_phy_microcore

    add_interface clk clock end
    add_interface reset reset end
    add_interface slave avalon slave
    add_interface pins conduit end

    set_interface_property clk EXPORT_OF clk.in_clk
    set_instance_property_value reset synchronous_edges deassert
    set_interface_property reset EXPORT_OF reset.in_reset
    set_interface_property slave EXPORT_OF regs.slave
    set_interface_property pins EXPORT_OF phy.pins
    
```

**Files**

```

add_connection clk.out_clk reset.clk
add_connection clk.out_clk regs.clk
add_connection clk.out_clk phy.clk
add_connection reset.out_reset regs.reset
add_connection reset.out_reset phy.clk_reset
add_connection regs.output phy.input
add_connection phy.output regs.input
}

```

**Related Information**

- [Component Interface Tcl Reference](#) on page 9-1

## Creating a Component With Differing Structural Qsys View and Generated Output Files

There are cases where it may be beneficial to have the structural Qsys system view of a component differ from the generated synthesis output files. The structural composition callback allows you to define a structural hierarchy for a component separately from the generated output files.

One application of this feature is for IP designers who want to send out a placed-and-routed component that represents a Qsys system in order to ensure timing closure for the end-user. In this case, the designer creates a design partition for the Qsys system, and then exports a post-fit Quartus II Exported Partition File (**.qxp**) when satisfied with the placement and routing results.

The designer specifies a **.qxp** file as the generated synthesis output file for the new component. The designer can specify whether to use a simulation output fileset for the custom simulation model file, or to use simulation output files generated from the original Qsys system.

When the end-user adds this component to their Qsys system, the designer wants the end-user to see a structural representation of the component, including lower-level components and the address map of the original Qsys system. This structural view is a logical representation of the component that is used during the elaboration and validation phases in Qsys.

### Example 6-13: Structural Composition Callback and .qxp File as the Generated Output

To specify a structural representation of the component for Qsys, connect components or generate a hardware Tcl description of the Qsys system, and then insert the Tcl commands into a structural composition callback. To invoke the structural composition callback use the command:

```

set_module_property STRUCTURAL_COMPOSITION_CALLBACK structural_hierarchy

package require -exact qsys 14.0
set_module_property name example_structural_composition

set_module_property STRUCTURAL_COMPOSITION_CALLBACK \
structural_hierarchy

add_fileset synthesis_fileset QUARTUS_SYNTH \
synth_callback_procedure

add_fileset simulation_fileset SIM_VERILOG \
sim_callback_procedure

set_fileset_property synthesis_fileset TOP_LEVEL \
my_custom_component

```

```
set_fileset_property simulation_fileset TOP_LEVEL \  
my_custom_component  
  
proc structural_hierarchy {} {  
  
# called during elaboration and validation phase  
# exported ports should be same in structural_hierarchy  
# and generated QXP  
  
# These commands could come from the exported hardware Tcl  
  
    add_interface clk clock sink  
    add_interface reset reset sink  
  
    add_instance clk_0 clock_source  
    set_interface_property clk EXPORT_OF clk_0.clk_in  
    set_interface_property reset EXPORT_OF clk_0.clk_in_reset  
  
    add_instance pll_0 altera_pll  
    # connections and connection parameters  
    add_connection clk_0.clk pll_0.refclk clock  
    add_connection clk_0.clk_reset pll_0.reset reset  
}  
  
proc synth_callback_procedure { entity_name } {  
  
# the QXP should have the same name for ports  
# as exported in structural_hierarchy  
  
    add_fileset_file my_custom_component.qxp QXP PATH \  
    "my_custom_component.qxp"  
}  
  
proc sim_callback_procedure { entity_name } {  
  
# the simulation files should have the same name for ports as  
# exported in structural_hierarchy  
  
    add_fileset_file my_custom_component.v VERILOG PATH \  
    "my_custom_component.v"  
    ...  
    ...  
}
```

### Related Information

[Creating a Composed Component or Subsystem](#) on page 6-28

## Adding Component Instances to a Static or Generated Component

You can create nested components by adding component instances to an existing component. Both static and generated components can create instances of other components. You can add child instances of a component in a `_hw.tcl` using elaboration callback.

With an elaboration callback, you can also instantiate and parameterize sub-components with the `add_hdl_instance` command as a function of the parent component's parameter values.

When you instantiate multiple nested components, you must create a unique variation name for each component with the `add_hdl_instance` command. Prefixing a variation name with the parent

component name prevents conflicts in a system. The variation name can be the same across multiple parent components if the generated parameterization of the nested component is exactly the same.

**Note:** If you do not adhere to the above naming variation guidelines, Qsys validation-time errors occur, which are often difficult to debug.

#### Related Information

- [Static Components](#) on page 6-32
- [Generated Components](#) on page 6-33

## Static Components

Static components always generate the same output, regardless of their parameterization. Components that instantiate static components must have only static children.

A design file that is static between all parameterizations of a component can only instantiate other static design files. Since static IPs always render the same HDL regardless of parameterization, Qsys generates static IPs only once across multiple instantiations, meaning they have the same top-level name set.

### Example 6-14: Typical Usage of the `add_hdl_instance` Command for Static Components

```
package require -exact qsys 14.0

set_module_property name add_hdl_instance_example
add_fileset synth_fileset QUARTUS_SYNTH synth_callback
set_fileset_property synth_fileset TOP_LEVEL basic_static
set_module_property elaboration_callback elab

proc elab {} {
    # Actual API to instantiate an IP Core
    add_hdl_instance emif_instance_name altera_mem_if_ddr3_emif

    # Make sure the parameters are set appropriately
    set_instance_parameter_value emif_instance_name SPEED_GRADE {7}
    ...
}

proc synth_callback { output_name } {
    add_fileset_file "basic_static.v" VERILOG PATH basic_static.v
}
```

### Example 6-15: Top-Level HDL Instance and Wrapper File Created by Qsys

In this example, Qsys generates a wrapper file for the instance name specified in the `_hw.tcl` file.

```
//Top Level Component HDL
module basic_static (input_wire, output_wire, inout_wire);
input [31:0] input_wire;
output [31:0] output_wire;
inout [31:0] inout_wire;

// Instantiation of the instance added via add_hdl_instance
// command. This is an example of how the instance added via
// the add_hdl_instance command can be used
// in the top-level file of the component.

emif_instance_name fixed_name_instantiation_in_top_level(
.pll_ref_clk (input_wire), // pll_ref_clk.clk
```

```
.global_reset_n (input_wire), // global_reset.reset_n
.soft_reset_n (input_wire), // soft_reset.reset_n
...
... );
endmodule

//Wrapper for added HDL instance
// emif_instance_name.v
// Generated using ACDS version 14.0

`timescale 1 ps / 1 ps
module emif_instance_name (
input wire pll_ref_clk, // pll_ref_clk.clk
input wire global_reset_n, // global_reset.reset_n
input wire soft_reset_n, // soft_reset.reset_n
output wire afi_clk, // afi_clk.clk
...
...);
example_addhdlinstance_system
_add_hdl_instance_example_0_emif_instance
_name_emif_instance_name emif_instance_name (

.pll_ref_clk (pll_ref_clk), // pll_ref_clk.clk
.global_reset_n (global_reset_n), // global_reset.reset_n
.soft_reset_n (soft_reset_n), // soft_reset.reset_n
...
...);
endmodule
```

## Generated Components

A generated component's fileset callback allows an instance of the component to create unique HDL design files based on the instance's parameter values. For example, you can write a fileset callback to include a control and status interface based on the value of a parameter. The callback overcomes a limitation of HDL languages, which do not allow runtime parameters.

Generated components change their generation output (HDL) based on their parameterization. If a component is generated, then any component that may instantiate it with multiple parameter sets must also be considered generated, since its HDL changes with its parameterization. This case has an effect that propagates up to the top-level of a design.

Since generated components are generated for each unique parameterized instantiation, when implementing the `add_hdl_instance` command, you cannot use the same fixed name (specified using `instance_name`) for the different variants of the child HDL instances. To facilitate unique naming for the wrapper of each unique parameterized instantiation of child HDL instances, you must use the following command so that Qsys generates a unique name for each wrapper. You can then access this unique wrapper name with a fileset callback so that the instances are instantiated inside the component's top-level HDL.

- To declare auto-generated fixed names for wrappers, use the command:

```
set_instance_property instance_name HDLINSTANCE_USE_GENERATED_NAME \
true
```

**Note:** You can only use this command with a generated component in the global context, or in an elaboration callback.

- To obtain auto-generated fixed name with a fileset callback, use the command:

```
get_instance_property instance_name HDLINSTANCE_GET_GENERATED_NAME
```

**Note:** You can only use this command with a fileset callback. This command returns the value of the auto-generated fixed name, which you can then use to instantiate inside the top-level HDL.

### Example 6-16: Typical Usage of the `add_hdl_instance` Command for Generated Components

Qsys generates a wrapper file for the instance name specified in the `_hw.tcl` file.

```
package require -exact qsys 14.0
set_module_property name generated_toplevel_component
set_module_property ELABORATION_CALLBACK elaborate
add_fileset QUARTUS_SYNTH QUARTUS_SYNTH generate
add_fileset SIM_VERILOG SIM_VERILOG generate
add_fileset SIM_VHDL SIM_VHDL generate

proc elaborate {} {
    # Actual API to instantiate an IP Core
    add_hdl_instance emif_instance_name altera_mem_if_ddr3_emif

    # Make sure the parameters are set appropriately
    set_instance_parameter_value emif_instance_name SPEED_GRADE {7}
    ...
    # instruct Qsys to use auto generated fixed name
    set_instance_property emif_instance_name \
    HDLINSTANCE_USE_GENERATED_NAME 1
}

proc generate { entity_name } {
    # get the autogenerated name for emif_instance_name added
    # via add_hdl_instance

    set autogeneratedfixedname [get_instance_property \
    emif_instance_name HDLINSTANCE_GET_GENERATED_NAME]

    set fileID [open "generated_toplevel_component.v" r]
    set temp ""

    # read the contents of the file

    while {[eof $fileID] != 1} {
        gets $fileID lineInfo

        # replace the top level entity name with the name provided
        # during generation

        regsub -all "substitute_entity_name_here" $lineInfo \
        "${entity_name}" lineInfo

        # replace the autogenerated name for emif_instance_name added
        # via add_hdl_instance
    }
}
```

```

regsub -all "substitute_autogenerated_emifinstancename_here" \
${lineInfo}"${autogeneratedfixedname}" lineInfo \
append temp "${lineInfo}\n"
}

# adding a top level component file

add_fileset_file ${entity_name}.v VERILOG TEXT $temp
}

```

### Example 6-17: Top-Level HDL Instance and Wrapper File Created By Qsys

```

// Top Level Component HDL

module substitute_entity_name_here (input_wire, output_wire,
inout_wire);

input [31:0] input_wire;
output [31:0] output_wire;
inout [31:0] inout_wire;

// Instantiation of the instance added via add_hdl_instance
// command. This is an example of how the instance added
// via add_hdl_instance command can be used
// in the top-level file of the component.

substitute_autogenerated_emifinstancename_here
fixed_name_instantiation_in_top_level (
.pll_ref_clk (input_wire), // pll_ref_clk.clk
.global_reset_n (input_wire), // global_reset.reset_n
.soft_reset_n (input_wire), // soft_reset.reset_n
...
);
endmodule

// Wrapper for added HDL instance
// generated_toplevel_component_0_emif_instance_name.v is the
// auto generated //emif_instance_name
// Generated using ACDS version 13.

`timescale 1 ps / 1 ps
module generated_toplevel_component_0_emif_instance_name (
input wire pll_ref_clk, // pll_ref_clk.clk
input wire global_reset_n, // global_reset.reset_n
input wire soft_reset_n, // soft_reset.reset_n
output wire afi_clk, // afi_clk.clk
...
);
example_addhdlinstance_system_add_hdl_instance_example_0_emif
_instance_name_emif_instance_name emif_instance_name (

.pll_ref_clk (pll_ref_clk), // pll_ref_clk.clk
.global_reset_n (global_reset_n), // global_reset.reset_n
.soft_reset_n (soft_reset_n), // soft_reset.reset_n
...
);
endmodule

```

#### Related Information

[Controlling File Generation Dynamically with Parameters and a Fileset Callback](#) on page 6-26



## Design Guidelines for Adding Component Instances

In order to promote standard and predictable results when generating static and generated components, Altera recommends the following best-practices:

- For two different parameterizations of a component, a component must never generate a file of the same name with different instantiations. The contents of a file of the same name must be identical for every parameterization of the component.
- If a component generates a nested component, it must never instantiate two different parameterizations of the nested component using the same instance name. If the parent component's parameterization affects the parameters of the nested component, the parent component must use a unique instance name for each unique parameterization of the nested component.
- Static components that generate differently based on parameterization have the potential to cause problems in the following cases:
  - Different file names with the same entity names, results in same entity conflicts at compilation-time
  - Different contents with the same file name results in overwriting other instances of the component, and in either file, compile-time conflicts or unexpected behavior.
- Generated components that generate files not based on the output name and that have different content results in either compile-time conflicts, or unexpected behavior.

## Document Revision History

The table below indicates edits made to the *Creating Qsys Components* content since its creation.

**Table 6-5: Document Revision History**

Date	Version	Changes
December 2014	14.1.0	<ul style="list-style-type: none"> <li>• Note added about the order in which to add simulation files: <i>Specifying Files for Simulation</i>.</li> </ul>
November 2013	13.1.0	<ul style="list-style-type: none"> <li>• add_hdl_instance</li> <li>• Added <i>Creating a Component With Differing Structural Qsys View and Generated Output Files</i>.</li> </ul>
May 2013	13.0.0	<ul style="list-style-type: none"> <li>• Consolidated content from other Qsys chapters.</li> <li>• Added <i>Upgrading IP Components to the Latest Version</i>.</li> <li>• Updated for AMBA APB support.</li> </ul>

Date	Version	Changes
November 2012	12.1.0	<ul style="list-style-type: none"> <li>Added AMBA AXI4 support.</li> <li>Added the <b>demo_axi_memory</b> example with screen shots and example <b>_hw.tcl</b> code.</li> </ul>
June 2012	12.0.0	<ul style="list-style-type: none"> <li>Added new tab structure for the Component Editor.</li> <li>Added AMBA AXI3 support.</li> </ul>
November 2011	11.1.0	Template update.
May 2011	11.0.0	<ul style="list-style-type: none"> <li>Removed beta status.</li> <li>Added Avalon Tri-state Conduit (Avalon-TC) interface type.</li> <li>Added many interface templates for Nios custom instructions and Avalon-TC interfaces.</li> </ul>
December 2010	10.1.0	Initial release.

For previous versions of the *Quartus II Handbook*, refer to the *Quartus II Handbook Archive*.

**Related Information**

[Quartus II Handbook Archive](#)

2014.12.15

QI15V1



Subscribe



Send Feedback

Qsys interconnect is a high-bandwidth structure that allows you to connect IP components to other IP components with various interfaces.

Qsys supports Avalon, AMBA AXI3 (version 1.0), AMBA AXI4 (version 2.0), AMBA AXI4-Lite (version 2.0), AMBA AXI4-Stream (version 1.0), and AMBA APB3 (version 1.0) interface specifications.

**Note:** The video, *AMBA AXI and Altera Avalon Interoperation Using Qsys*, describes seamless integration of IP components using the AMBA AXI interface, and the Altera Avalon interface.

#### Related Information

- [Avalon Interface Specifications](#)
- [AMBA Specifications](#)
- [Creating a System with Qsys](#) on page 5-1
- [Creating Qsys Components](#) on page 6-1
- [Qsys System Design Components](#) on page 10-1
- [AMBA AXI and Altera Avalon Interoperation Using Qsys](#)

## Memory-Mapped Interfaces

Qsys supports the implementation of memory-mapped interfaces for Avalon, AXI, and APB protocols.

Qsys interconnect transmits memory-mapped transactions between masters and slaves in packets. The command network transports read and write packets from master interfaces to slave interfaces. The response network transports response packets from slave interfaces to master interfaces.

For each component interface, Qsys interconnect manages memory-mapped transfers and interacts with signals on the connected interface. Master and slave interfaces can implement different signals based on interface parameterizations, and Qsys interconnect provides any necessary adaptation between them. In the path between master and slaves, Qsys interconnect may introduce registers for timing synchronization, finite state machines for event sequencing, or nothing at all, depending on the services required by the interfaces.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered




Qsys interconnect supports the following implementation scenarios:

- Any number of components with master and slave interfaces. The master-to-slave relationship can be one-to-one, one-to-many, many-to-one, or many-to-many.
- Masters and slaves of different data widths.
- Masters and slaves operating in different clock domains.
- IP Components with different interface properties and signals. Qsys adapts the component interfaces so that interfaces with the following differences can be connected:
  - Avalon and AXI interfaces that use active-high and active-low signaling. AXI signals are active high, except for the reset signal.
  - Interfaces with different burst characteristics.
  - Interfaces with different latencies.
  - Interfaces with different data widths.
  - Interfaces with different optional interface signals.

**Note:** AXI3/4 to AXI3/4 interface connections declare a fixed set of signals with variable latency. As a result, there is no need for adapting between active-low and active-high signaling, burst characteristics, different latencies, or port signatures. Some adaptation may be necessary between Avalon interfaces.

**Figure 7-1: Qsys interconnect for an Avalon-MM System with Multiple Masters**

In this example, there are two components mastering the system, a processor and a DMA controller, each with two master interfaces. The masters connect through the Qsys interconnect to several slaves in the Qsys system. The dark blue blocks represent interconnect components. The dark grey boxes indicate items outside of the Qsys system and the Quartus II software design, and show how component interfaces can be exported and connected to external devices.



## Qsys Packet Format

The Qsys packet format supports Avalon, AXI, and APB transactions. Memory-mapped transactions between masters and slaves are encapsulated in Qsys packets. For Avalon systems without AXI or APB interfaces, some fields are ignored or removed.

### Qsys Packet Format

**Table 7-1: Qsys Packet Format for Memory-Mapped Master and Slave Interfaces**

The fields of the Qsys packet format are of variable length to minimize resource usage. However, if the majority of components in a design have a single data width, for example 32-bits, and a single component has a data width of 64-bits, Qsys inserts a width adapter to accommodate 64-bit transfers.

Command	Description
Address	Specifies the byte address for the lowest byte in the current cycle. There are no restrictions on address alignment.
Size	Encodes the run-time size of the transaction.  In conjunction with address, this field describes the segment of the payload that contains valid data for a beat within the packet.
Address Sideband	Carries “address” sideband signals. The interconnect passes this field from master to slave. This field is valid for each beat in a packet, even though it is only produced and consumed by an address cycle.  Up to 8-bit sideband signals are supported for both read and write address channels.
Cache	Carries the AXI cache signals.
Transaction (Exclusive)	Indicates whether the transaction has exclusive access.
Transaction (Posted)	Used to indicate non-posted writes (writes that require responses).
Data	For command packets, carries the data to be written. For read response packets, carries the data that has been read.

Command	Description
Byteenable	<p>Specifies which symbols are valid. AXI can issue or accept any byteenable pattern. For compatibility with Avalon, Altera recommends that you use the following legal values for 32-bit data transactions between Avalon masters and slaves:</p> <ul style="list-style-type: none"> <li>• <b>1111</b>—Writes full 32 bits</li> <li>• <b>0011</b>—Writes lower 2 bytes</li> <li>• <b>1100</b>—Writes upper 2 bytes</li> <li>• <b>0001</b>—Writes byte 0 only</li> <li>• <b>0010</b>—Writes byte 1 only</li> <li>• <b>0100</b>—Writes byte 2 only</li> <li>• <b>1000</b>—Writes byte 3 only</li> </ul>
Source_ID	The ID of the master or slave that initiated the command or response.
Destination_ID	The ID of the master or slave to which the command or response is directed.
Response	Carries the AXI response signals.
Thread ID	Carries the AXI transaction ID values.
Byte count	The number of bytes remaining in the transaction, including this beat. Number of bytes requested by the packet.

Command	Description
Burstwrap	<p>The burstwrap value specifies the wrapping behavior of the current burst. The burstwrap value is of the form <math>2^{&lt;n&gt;} - 1</math>. The following types are defined:</p> <ul style="list-style-type: none"> <li>• Variable wrap—Variable wrap bursts can wrap at any integer power of 2 value. When the burst reaches the wrap boundary, it wraps back to the previous burst boundary so that only the low order bits are used for addressing. For example, a burst starting at address 0x1C, with a burst wrap boundary of 32 bytes and a burst size of 20 bytes, would write to addresses 0x1C, 0x0, 0x4, 0x8, and 0xC.</li> <li>• For a burst wrap boundary of size <math>&lt;m&gt;</math>, <math>\text{Burstwrap} = &lt;m&gt; - 1</math>, or for this case <math>\text{Burstwrap} = (32 - 1) = 31</math> which is <math>2^5 - 1</math>.</li> <li>• For AXI masters, the burstwrap boundary value (m) is based on the different AXBURST: <ul style="list-style-type: none"> <li>• Burstwrap set to all 1's. For example, for a 6-bit burstwrap, burstwrap is 6'b111111.</li> <li>• For WRAP bursts, <math>\text{burstwrap} = \text{AXLEN} * \text{size} - 1</math>.</li> <li>• For FIXED bursts, <math>\text{burstwrap} = \text{size} - 1</math>.</li> <li>• Sequential bursts increment the address for each transfer in the burst. For sequential bursts, the <code>Burstwrap</code> field is set to all 1s. For example, with a 6-bit <code>Burstwrap</code> field, the value for a sequential burst is 6'b111111 or 63, which is <math>2^6 - 1</math>.</li> </ul> </li> </ul> <p>For Avalon masters, Qsys adaptation logic sets a hardwired value for the burstwrap field, according the declared master burst properties. For example, for a master that declares sequential bursting, the burstwrap field is set to ones. Similarly, masters that declare burst have their burstwrap field set to the appropriate constant value.</p> <p>AXI masters choose their burst type at run-time, depending on the value of the <code>AW</code> or <code>ARBURST</code> signal. The interconnect calculates the burstwrap value at run-time for AXI masters.</p>
Protection	<p>Access level protection. When the lowest bit is 0, the packet has normal access. When the lowest bit is 1, the packet has privileged access. For Avalon-MM interfaces, this field maps directly to the privileged access signal, which allows an memory-mapped master to write to an on-chip memory ROM instance. The other bits in this field support AXI secure accesses and uses the same encoding, as described in the AXI specification.</p>
QoS	<p>QoS (Quality of Service Signaling) is a 4-bit field that is part of the AXI4 interface that carries QoS information for the packet from the AXI master to the AXI slave.</p> <p>Transactions from AXI3 and Avalon masters have the default value 4'b0000, that indicates that they are not participating in the QoS scheme. QoS values are dropped for slaves that do not support QoS.</p>



Command	Description
Data sideband	Carries data sideband signals for the packet. On a write command, the data sideband directly maps to <code>WUSER</code> . On a read response, the data sideband directly maps to <code>RUSER</code> . On a write response, the data sideband directly maps to <code>BUSER</code> .

## Transaction Types for Memory-Mapped Interfaces

**Table 7-2: Transaction Types for Memory-Mapped Interfaces**

The table below describes the information that each bit transports in the packet format's transaction field.


Bit	Name	Definition
0	<code>PKT_TRANS_READ</code>	When asserted, indicates a read transaction.
1	<code>PKT_TRANS_COMPRESSED_READ</code>	For read transactions, specifies whether or not the read command can be expressed in a single cycle, that is whether or not it has all <code>byteenables</code> asserted on every cycle.
2	<code>PKT_TRANS_WRITE</code>	When asserted, indicates a write transaction.
3	<code>PKT_TRANS_POSTED</code>	When asserted, no response is required.
4	<code>PKT_TRANS_LOCK</code>	When asserted, indicates arbitration is locked. Applies to write packets.

## Qsys Transformations

The memory-mapped master and slave components connect to network interface modules that encapsulate the transaction in Avalon-ST packets. The memory-mapped interfaces have no information about the encapsulation or the function of the layer transporting the packets. The interfaces operate in accordance with memory-mapped protocol and use the read and write signals and transfers.

**Figure 7-2: Transformation when Generating a System with Memory-Mapped and Slave Components**

Qsys components that implement the blocks appear shaded.



#### Related Information

- [Master Network Interfaces](#) on page 7-11
- [Slave Network Interfaces](#) on page 7-13

## Interconnect Domains

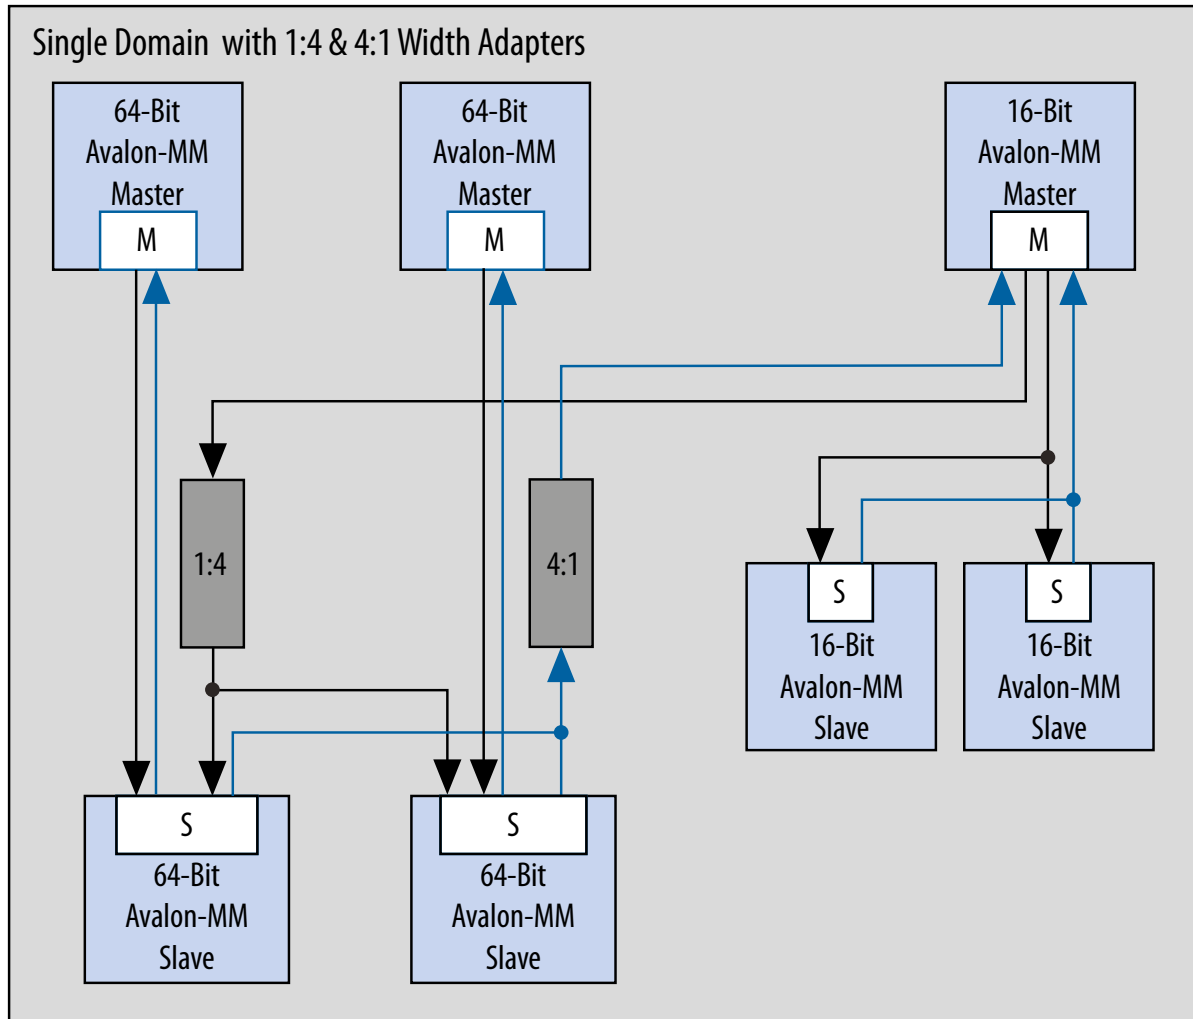
An interconnect domain is a group of connected memory-mapped masters and slaves that share the same interconnect. The components in a single interconnect domain share the same packet format.

### Using One Domain with Width Adaptation

When one of the masters in a system connects to all of the slaves, Qsys creates a single domain with two packet formats: one with 64-bit data, and one with 16-bit data. A width adapter manages accesses between the 16-bit master and 64-bit slaves.

**Figure 7-3: One Domain with 1:4 and 4:1 Width Adapters**

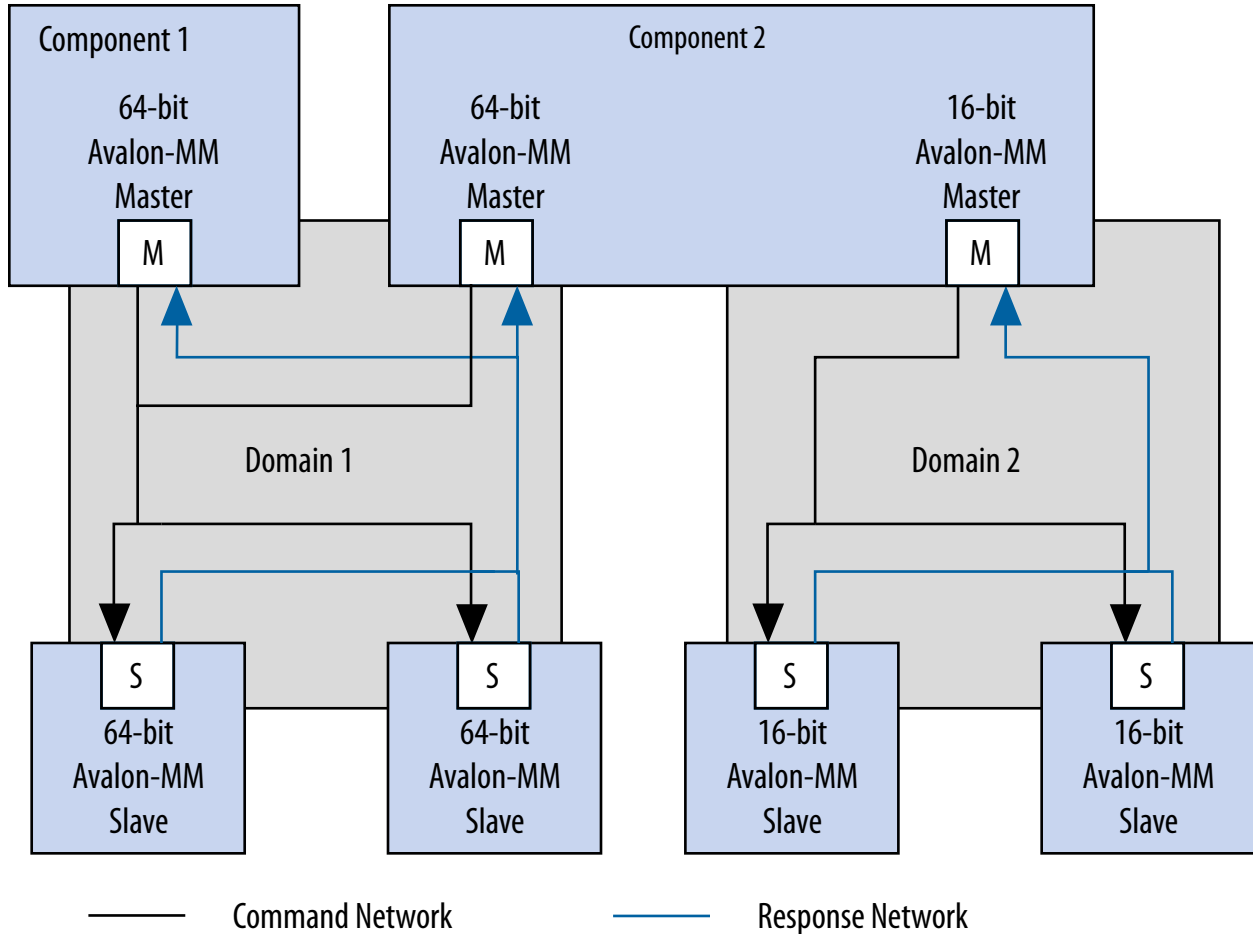
In this system example, there are two 64-bit masters that access two 64-bit slaves. It also includes one 16-bit master, that accesses two 16-bit slaves and two 64-bit slaves. The 16-bit Avalon-MM master connects through a 1:4 adapter, then a 4:1 adapter to reach its 16-bit slaves.



## Using Two Separate Domains

Figure 7-4: Two Separate Domains

In this system example, Qsys uses two separate domains. The first domain includes two 64-bit masters connected to two 64-bit slaves. A second domain includes one 16-bit master connected to two 16-bit slaves. Because the interfaces in Domain 1 and Domain 2 do not share any connections, Qsys can optimize the packet format for the two separate domains. In this example, the first domain uses a 64-bit data width and the second domain uses 16-bit data.

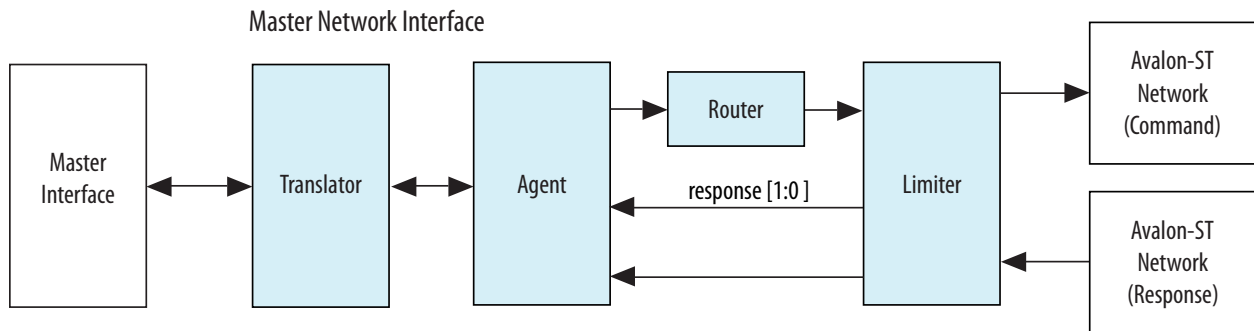


## Master Network Interfaces

**Figure 7-5: Avalon-MM Master Network Interface**

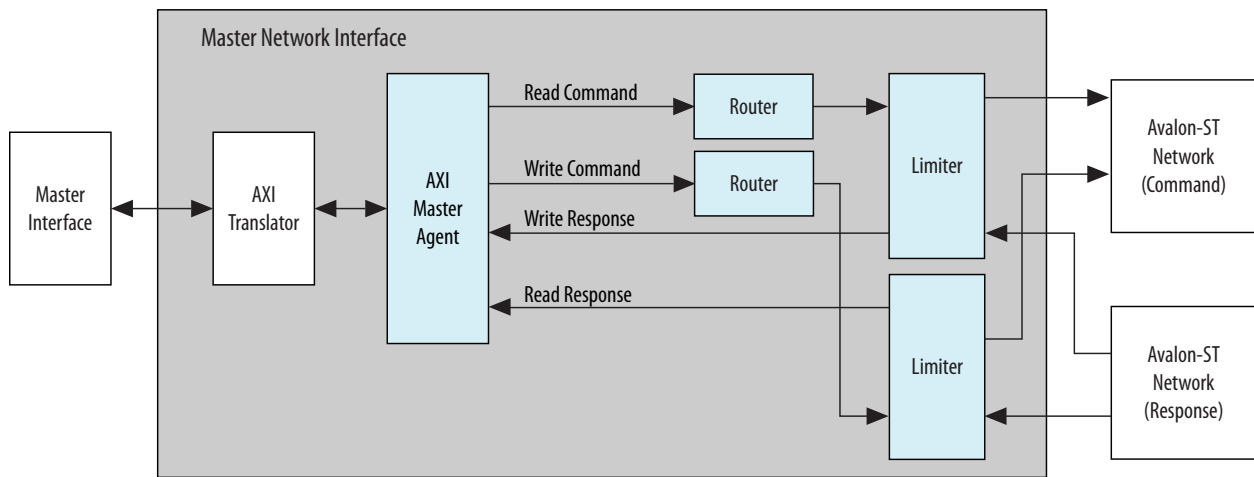
Avalon network interfaces drive default values for the QoS and BUSER, WUSER, and RUSER packet fields in the master agent, and drop the packet fields in the slave agent.

**Note:** The `response` signal from the Limiter to the Agent is optional.



**Figure 7-6: AXI Master Network Interface**

An AXI4 master supports INCR bursts up to 256 beats, QoS signals, and data sideband signals.



**Note:** For a complete definition of the optional read response signal, refer to *Avalon Memory-Mapped Interface Signal Types* in the *Avalon Interface Specifications*.

### Related Information

- [Responses](#) on page 7-24
- [Avalon Interface Specifications](#)
- [Creating a System with Qsys](#) on page 5-1

## Avalon-MM Master Agent

The Avalon-MM Master Agent translates Avalon-MM master transactions into Qsys command packets and translates the Qsys Avalon-MM slave response packets into Avalon-MM responses.

## Avalon-MM Master Translator

The Avalon-MM Master Translator interfaces with an Avalon-MM master component and converts the Avalon-MM master interface to a simpler representation for use in Qsys.

The Avalon-MM Master translator performs the following functions:

- Translates active-low signaling to active-high signaling
- Inserts wait states to prevent an Avalon-MM master from reading invalid data
- Translates word and symbol addresses
- Translates word and symbol burst counts
- Manages re-timing and re-sequencing bursts
- Removes unnecessary address bits

## AXI Master Agent

An AXI Master Agent accepts AXI commands and produces Qsys command packets. It also accepts Qsys response packets and converts those into AXI responses. This component has separate packet channels for read commands, write commands, read responses, and write responses. Avalon master agent drives the `QoS` and `BUSER`, `WUSER`, and `RUSER` packet fields with default values `AXQ0` and `b0000`, respectively.

**Note:** For signal descriptions, refer to *Qsys Packet Format*.

### Related Information

[Qsys Packet Format](#) on page 7-4

## AXI Translator

AXI4 allows some signals to be omitted from interfaces. The translator bridges between these “incomplete” AXI4 interfaces and the “complete” AXI4 interface on the network interfaces.

The AXI translator is inserted for both AXI4 masters and slaves and performs the following functions:

- Matches ID widths between the master and slave in 1x1 systems.
- Drives default values as defined in the *AMBA Protocol Specifications* for missing signals.
- Performs lock transaction bit conversion when an AXI3 master connects to an AXI4 slave in 1x1 systems.

### Related Information

[AMBA Protocol Specifications](#)

## APB Master Agent

An APB master agent accepts APB commands and produces or generates Qsys command packets. It also converts Qsys response packets to APB responses.

## APB Slave Agent

An APB slave agent issues resulting transaction to the APB interface. It also accepts creates Qsys response packets.

## APB Translator

An APB peripheral does not require `pslverr` signals to support additional signals for the APB debug interface.

The APB translator is inserted for both the master and slave and performs the following functions:

- Sets the response value default to `OKAY` if the APB slave does not have a `pslverr` signal.
- Turns on or off additional signals between the APB debug interface, which is used with HPS (Altera SoC's Hard Processor System).

## Memory-Mapped Router

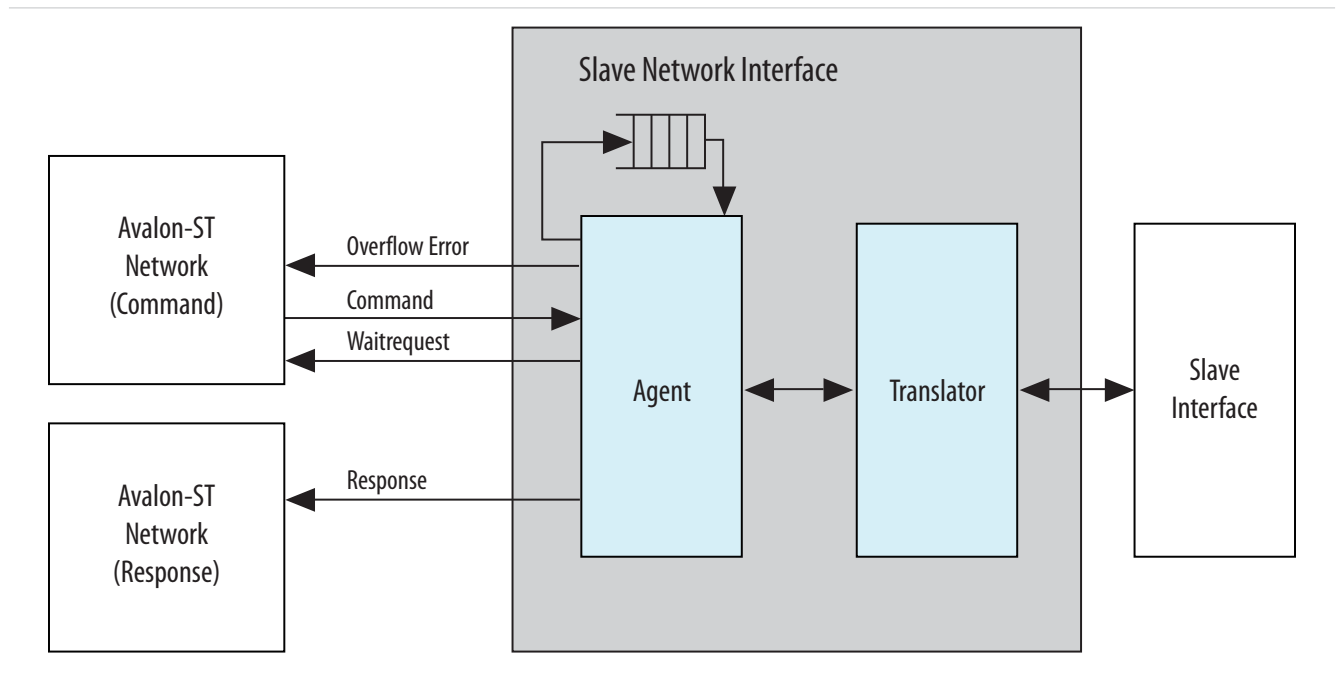
The Memory-Mapped Router routes command packets from the master to the slave, and response packets from the slave to the master. For master command packets, the router uses the address to set the `Destination_ID` and Avalon-ST channel. For the slave response packet, the router uses the `Destination_ID` to set the Avalon-ST channel. The demultiplexers use the Avalon-ST channel to route the packet to the correct destination.

## Memory-Mapped Traffic Limiter

The Memory-Mapped Traffic Limiter ensures the responses arrive in order. It prevents any command from being sent if the response could conflict with the response for a command that has already been issued. By guaranteeing in-order responses, the Traffic Limiter simplifies the response network.

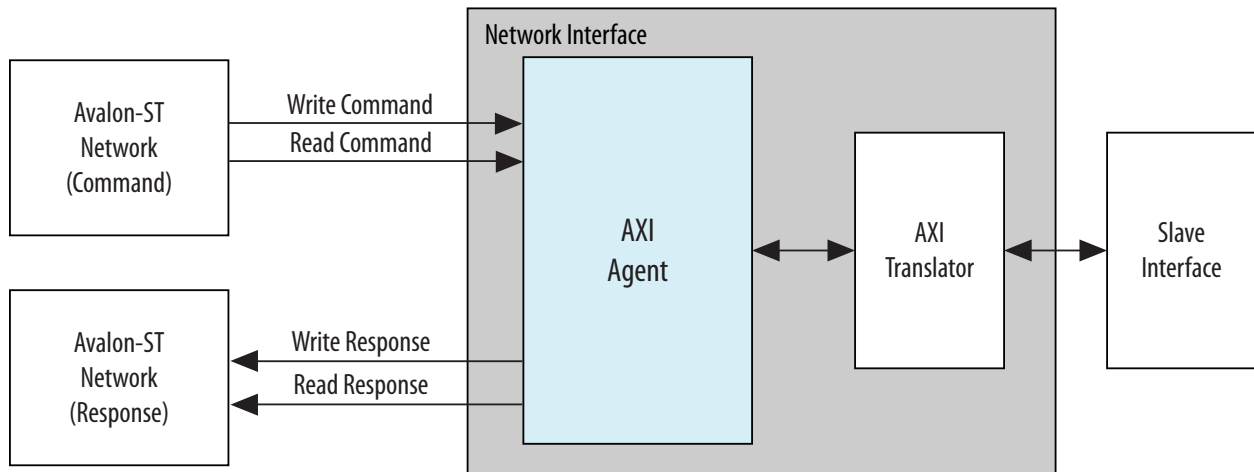
## Slave Network Interfaces

Figure 7-7: Avalon-MM Slave Network Interface



**Figure 7-8: AXI Slave Network Interface**

An AXI4 slave supports up to 256 beat `INCR` bursts, QoS signals, and data sideband signals.

**Avalon-MM Slave Translator**

The Avalon-MM Slave Translator interfaces to an Avalon-MM slave component as the *Avalon-MM Slave Network Interface* figure illustrates. It converts the Avalon-MM slave interface to a simplified representation that the Qsys network can use.

An Avalon-MM Merlin Slave Translator performs the following functions:

- Drives the `beginbursttransfer` and `byteenable` signals.
- Supports Avalon-MM slaves that operate using fixed timing and or slaves that use the `readdatavalid` signal to identify valid data.
- Translates the `read`, `write`, and `chipselct` signals into the representation that the Avalon-ST slave response network uses.
- Converts active low signals to active high signals.
- Translates word and symbol addresses and burstcounts.
- Handles burstcount timing and sequencing.
- Removes unnecessary address bits.

**Related Information**

[Slave Network Interfaces](#) on page 7-13

**AXI Translator**

AXI4 allows some signals to be omitted from interfaces. The translator bridges between these “incomplete” AXI4 interfaces and the “complete” AXI4 interface on the network interfaces.

The AXI translator is inserted for both AXI4 master and slave, and performs the following functions:

- Matches ID widths between master and slave in 1x1 systems.
- Drives default values as defined in the *AMBA Protocol Specifications* for missing signals.
- Performs lock transaction bit conversion when an AXI3 master connects to an AXI4 slave in 1x1 systems.




## Wait State Insertion

Wait states extend the duration of a transfer by one or more cycles. Wait state insertion logic accommodates the timing needs of each slave, and causes the master to wait until the slave can proceed. Qsys interconnect inserts wait states into a transfer when the target slave cannot respond in a single clock cycle, as well as in cases when slave `read` and `write` signals have setup or hold time requirements.

**Figure 7-9: Wait State Insertion Logic for One Master and One Slave**

Wait state insertion logic is a small finite-state machine that translates control signal sequencing between the slave side and the master side. Qsys interconnect can force a master to wait for the wait state needs of a slave. For example, arbitration logic in a multi-master system. Qsys generates wait state insertion logic based on the properties of all slaves in the system.



## Avalon-MM Slave Agent

The Avalon-MM Slave Agent accepts command packets and issues the resulting transactions to the Avalon interface. For pipelined slaves, an Avalon-ST FIFO stores information about pending transactions. The size of this FIFO is the maximum number of pending responses that you specify when creating the slave component. The Avalon-MM Slave Agent also *backpressures* the Avalon-MM master command interface when the FIFO is full if the slave component includes the `waitrequest` signal.

## AXI Slave Agent

An AXI Slave Agent works similar to a master agent in reverse. The AXI slave Agent accepts Qsys command packets to create AXI commands, and accepts AXI responses to create Qsys response packets. This component has separate packet channels for read commands, write commands, read responses, and write responses.

## Arbitration

When multiple masters contend for access to a slave, Qsys automatically inserts arbitration logic, which grants access in fairness-based, round-robin order. You can alternatively choose to designate a slave as a fixed priority arbitration slave, and then manually assign priorities in the Qsys GUI.

## Round-Robin Arbitration

When multiple masters contend for access to a slave, Qsys automatically inserts arbitration logic which grants access in fairness-based, round-robin order.

In a fairness-based arbitration protocol, each master has an integer value of transfer *shares* with respect to a slave. One share represents permission to perform one transfer. The default arbitration scheme is equal

share round-robin that grants equal, sequential access to all requesting masters. You can change the arbitration scheme to weighted round-robin by specifying a relative number of arbitration shares to the masters that access a particular slave. AXI slaves have separate arbitration for their independent read and write channels, and the **Arbitration Shares** setting affects both the read and write arbitration. To display arbitration settings, right-click an instance on the **System Contents** tab, and then click **Show Arbitration Shares**.

Figure 7-10: Arbitration Shares in the Connections Column


Connections	Name	Description
	<input type="checkbox"/> <b>mm_master_bfm_0_avalon</b> clk clk_reset m0	Altera UVM Avalon-MM Master BFM Clock Input Reset Input Avalon Memory Mapped Master
	<input type="checkbox"/> <b>mm_master_bfm_1_axi</b> clk clk_reset altera_axi_master	Altera AXI3 Master Module Clock Input Reset Input AXI Master
	<input type="checkbox"/> <b>mm_master_bfm_2_axi</b> clk clk_reset altera_axi_master	Altera AXI3 Master Module Clock Input Reset Input AXI Master
	<input type="checkbox"/> <b>mm_slave_bfm_0_avalon</b> clk clk_reset s0	Altera UVM Avalon-MM Slave BFM Clock Input Reset Input Avalon Memory Mapped Slave
	<input type="checkbox"/> <b>mm_slave_bfm_1_avalon</b> clk clk_reset s0	Altera UVM Avalon-MM Slave BFM Clock Input Reset Input Avalon Memory Mapped Slave
	<input type="checkbox"/> <b>mm_slave_bfm_2_axi</b> clk clk_reset altera_axi_slave	Altera AXI3 Slave Module Clock Input Reset Input AXI Slave
	<input type="checkbox"/> <b>CLOCK_0</b> clk clk_reset dummy_src dummy_snk	Altera Avalon Clock and Reset Source Clock Output Reset Output Avalon Streaming Source Avalon Streaming Sink

### Fairness-Based Shares

In a fairness-based arbitration scheme, each master-to-slave connection provides a transfer share count. This count is a request for the arbiter to grant a specific number of transfers to this master before giving control to a different master. One share represents permission to perform one transfer.


**Figure 7-11: Arbitration of Continuous Transfer Requests from Two Masters**

Consider a system with two masters connected to a single slave. Master 1 has its arbitration shares set to three, and Master 2 has its arbitration shares set to four. Master 1 and Master 2 continuously attempt to perform back-to-back transfers to the slave. The arbiter grants Master 1 access to the slave for three transfers, and then grants Master 2 access to the slave for four transfers. This cycle repeats indefinitely. The figure below describes the waveform for this scenario.



**Figure 7-12: Arbitration of Two Masters with a Gap in Transfer Requests**

If a master stops requesting transfers before it exhausts its shares, it forfeits all of its remaining shares, and the arbiter grants access to another requesting master. After completing one transfer, Master 2 stops requesting for one clock cycle. As a result, the arbiter grants access back to Master 1, which gets a replenished supply of shares.



### Round-Robin Scheduling

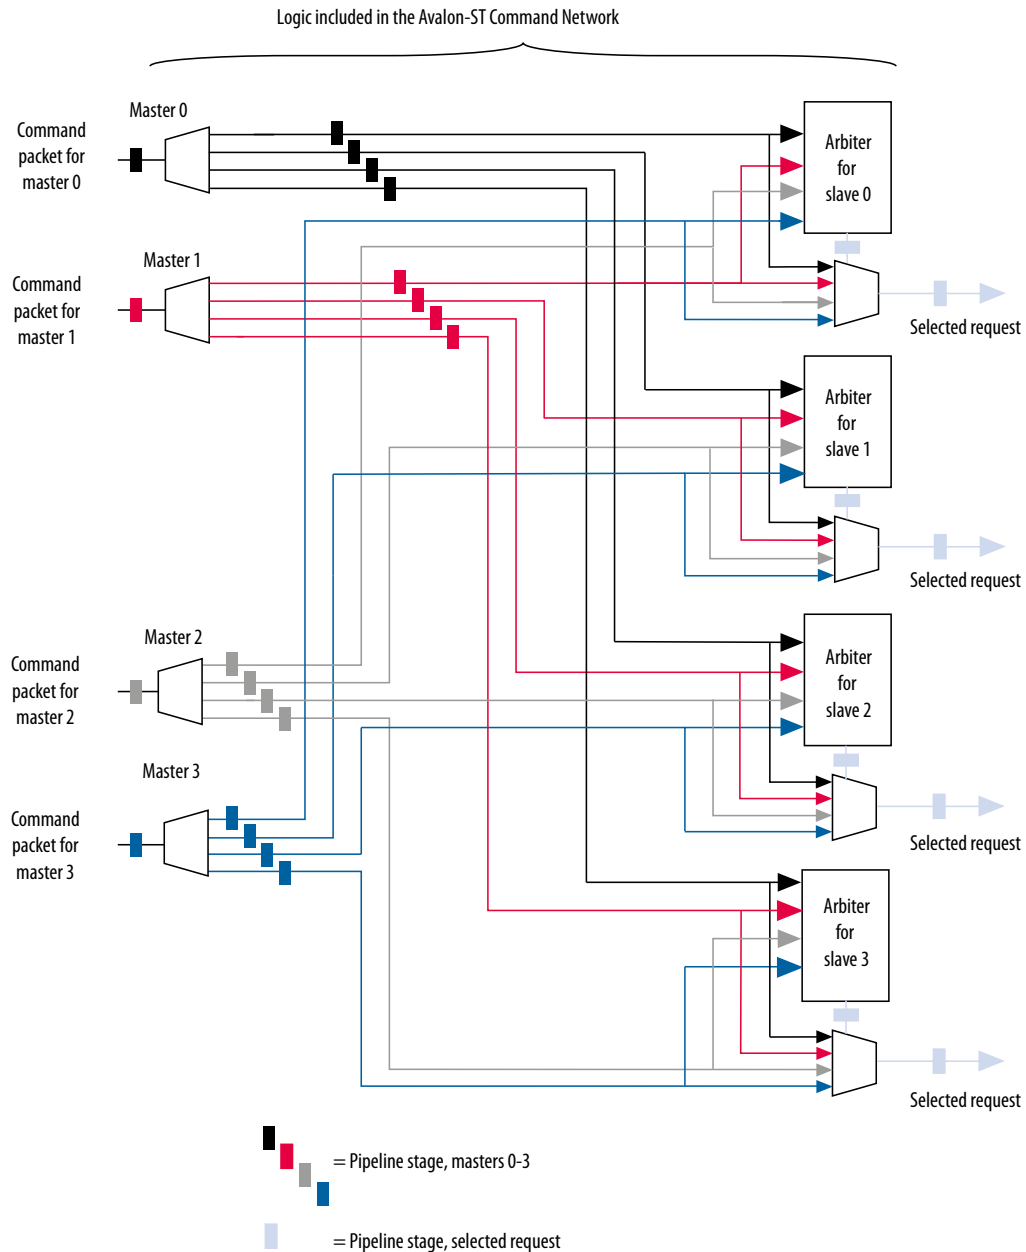
When multiple masters contend for access to a slave, the arbiter grants shares in round-robin order. Qsys includes only requesting masters in the arbitration for each slave transaction.

### Memory-Mapped Arbiter

The input to the Memory-Mapped Arbiter is the command packet for all masters requesting access to a particular slave. The arbiter outputs the channel number for the selected master. This channel number controls the output of a multiplexer that selects the slave device. The figure below illustrates this logic.

**Figure 7-13: Arbitration Logic**

In this example, four Avalon-MM masters connect to four Avalon-MM slaves. In each cycle, an arbiter positioned in front of each Avalon-MM slave selects among the requesting Avalon-MM masters.




**Note:** If you specify a **Limit interconnect pipeline stages** to parameter greater than zero, the output of the Arbiter is registered. Registering this output reduces the amount of combinational logic between the master and the interconnect, increasing the  $f_{MAX}$  of the system.

**Note:** You can use the Memory-Mapped Arbiter for both round-robin and fixed priority arbitration.

## Datapath Multiplexing Logic

Datapath multiplexing logic drives the `writedata` signal from the granted master to the selected slave, and the `readdata` signal from the selected slave back to the requesting master. Qsys generates separate datapath multiplexing logic for every master in the system (`readdata`), and for every slave in the system (`writedata`). Qsys does not generate multiplexing logic if it is not needed.

Figure 7-14: Datapath Multiplexing Logic for One Master and Two Slaves



## Width Adaptation

Qsys width adaptation converts between Avalon memory-mapped master and slaves with different data and byte enable widths, and manages the run-time size requirements of AXI. Width adaptation for AXI to Avalon interfaces is also supported.

### Memory-Mapped Width Adapter

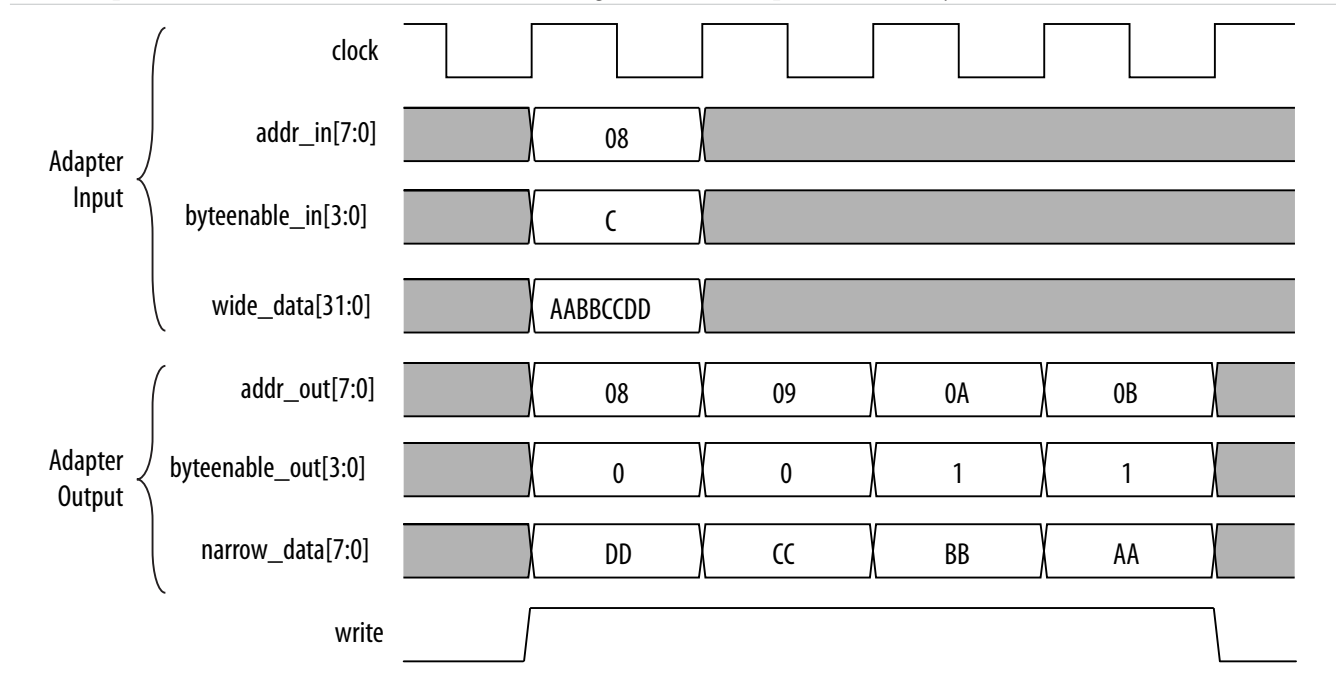
The Memory-Mapped Width Adapter is used in the Avalon-ST domain and operates with information contained in the packet format.

The memory-mapped width adapter accepts packets on its sink interface with one data width and produces output packets on its source interface with a different data width. The ratio of the narrow data width must be a power of two, such as 1:4, 1:8, and 1:16. The ratio of the wider data width to the narrower width must also be a power of two, such as 4:1, 8:1, and 16:1. These output packets may have a different size if the input size exceeds the output data bus width, or if data packing is enabled.

When the width adapter converts from narrow data to wide data, each input beat's data and byte enables are copied to the appropriate segment of the wider output data and byte enables signals.

**Figure 7-15: Width Adapter Timing for a 4:1 Adapter**

This adapter assumes that the field ordering of the input and output packets is the same, with the only difference being the width of the data and accompanying byte enable fields. When the width adapter converts from wide data to narrow data, the narrower data is transmitted over several beats. The first output beat contains the lowest addressed segment of the input data and byte enables.

**AXI Wide-to-Narrow Adaptation**

For all cases of AXI wide-to-narrow adaptation, read data is re-packed to match the original size. Responses are merged, with the following error precedence: `DECERR`, `SLVERR`, `OKAY`, and `EXOKAY`.

**Table 7-3: AXI Wide-to-Narrow Adaptation (Downsizing)**

Burst Type	Behavior
Incrementing Incrementing	<p>If the transaction size is less than or equal to the output width, the burst is unmodified. Otherwise, it is converted to an incrementing burst with a larger length and size equal to the output width.</p> <p>If the resulting burst is unsuitable for the slave, the burst is converted to multiple sequential bursts of the largest allowable lengths. For example, for a 2:1 downsizing ratio, an <code>INCR9</code> burst is converted into <code>INCR16</code> + <code>INCR2</code> bursts. This is true if the maximum burstcount a slave can accept is 16, which is the case for AXI3 slaves. Avalon slaves have a maximum burstcount of 64.</p>

Burst Type	Behavior
Wrapping	<p>If the transaction size is less than or equal to the output width, the burst is unmodified. Otherwise, it is converted to a wrapping burst with a larger length, with a size equal to the output width.</p> <p>If the resulting burst is unsuitable for the slave, the burst is converted to multiple sequential bursts of the largest allowable lengths; respecting wrap boundaries. For example, for a 2:1 downsizing ratio, a <code>WRAP16</code> burst is converted into two or three <code>INCR</code> bursts, depending on the address.</p>
Fixed	<p>If the transaction size is less than or equal to the output width, the burst is unmodified. Otherwise, it is converted into repeated sequential bursts over the same addresses. For example, for a 2:1 downsizing ratio, a <code>FIXED</code> single burst is converted into an <code>INCR2</code> burst.</p>

### AXI Narrow-to-Wide Adaptation

Table 7-4: AXI Narrow-to-Wide Adaptation (Upsizing)

Burst Type	Behavior
Incrementing	The burst (and its response) passes through unmodified. Data and write strobes are placed in the correct output segment.
Wrapping	The burst (and its response) passes through unmodified.
Fixed	The burst (and its response) passes through unmodified.

### Burst Adapter

Qsys interconnect uses the memory-mapped burst adapter to accommodate the burst capabilities of each interface in the system, including interfaces that do not support burst transfers.

The maximum burst length for each interface is a property of the interface and is independent of other interfaces in the system. Therefore, a particular master may be capable of initiating a burst longer than a slave's maximum supported burst length. In this case, the burst adapter translates the large master burst into smaller bursts, or into individual slave transfers if the slave does not support bursting. Until the master completes the burst, arbiter logic prevents other masters from accessing the target slave. For example, if a master initiates a burst of 16 transfers to a slave with maximum burst length of 8, the burst adapter initiates 2 bursts of length 8 to the slave.

Avalon-MM and AXI burst transactions allow a master uninterrupted access to a slave for a specified number of transfers. The master specifies the number of transfers when it initiates the burst. Once a burst begins between a master and slave, arbiter logic is locked until the burst completes. For burst masters, the length of the burst is the number of cycles that the master has access to the slave, and the selected arbitration shares have no effect.

**Note:** AXI masters can issue burst types that Avalon cannot accept, for example, fixed bursts. In this case, the burst adapter converts the fixed burst into a sequence of transactions to the same address.

**Note:** For AXI4 slaves, Qsys allows 256-beat INCR bursts. You must ensure that 256-beat narrow-sized INCR bursts are shortened to 16-beat narrow-sized INCR bursts for AXI3 slaves.

Avalon-MM masters always issue addresses that are aligned to the size of the transfer. However, when Qsys uses a narrow-to-wide width adaptation, the resulting address may be unaligned. For unaligned addresses, the burst adapter issues the maximum sized bursts with appropriate byte enables. This brings the burst-in-progress up to an aligned slave address. Then, it completes the burst on aligned addresses.

The burst adapter supports variable wrap or sequential burst types to accommodate different properties of memory-mapped masters. Some bursting masters can issue more than one burst type.

Burst adaptation is available for Avalon to Avalon, Avalon to AXI, and AXI to Avalon, and AXI to AXI connections. For information about AXI-to-AXI adaptation, refer to *AXI Wide-to-Narrow Adaptation*

**Note:** For AXI4 to AXI3 connections, Qsys follows an AXI4 256 burst length to AXI3 16 burst length.

## Burst Adapter Implementation Options

Qsys automatically inserts burst adapters into your system depending on your master and slave connections, and properties. You can select burst adapter implementation options on the **Interconnect Requirements** tab.

To access the implementation options, you must select the **Burst adapter implementation** setting for the `$system` identifier.

- **Generic converter (slower, lower area)**—Default. Controls all burst conversions with a single converter that is able to adapt incoming burst types. This results in an adapter that has lower  $f_{\max}$ , but smaller area.
- **Per-burst-type converter (faster, higher area)**—Controls incoming bursts with a particular converter, depending on the burst type. This results in an adapter that has higher  $f_{\max}$ , but higher area. This setting is useful when you have AXI masters or slaves and you want a higher  $f_{\max}$ .

**Note:** For more information about the **Interconnect Requirements** tab, refer to *Creating a System with Qsys*.

### Related Information

- [Creating a System with Qsys](#) on page 5-1



## Burst Adaptation: AXI to Avalon

**Table 7-5: Burst Adaptation: AXI to Avalon**

Entries specify the behavior when converting between AXI and Avalon burst types.

Burst Type	Behavior
Incrementing	<p><b>Sequential Slave</b></p> <p>Bursts that exceed <code>slave_max_burst_length</code> are converted to multiple sequential bursts of a length less than or equal to the <code>slave_max_burst_length</code>. Otherwise, the burst is unconverted. For example, for an Avalon slave with a maximum burst length of 4, an <code>INCR7</code> burst is converted to <code>INCR4 + INCR3</code>.</p> <p><b>Wrapping Slave</b></p> <p>Bursts that exceed the <code>slave_max_burst_length</code> are converted to multiple sequential bursts of length less than or equal to the <code>slave_max_burst_length</code>. Bursts that exceed the wrapping boundary are converted to multiple sequential bursts that respect the slave's wrapping boundary.</p>
Wrapping	<p><b>Sequential Slave</b></p> <p>A <code>WRAP</code> burst is converted to multiple sequential bursts. The sequential bursts are less than or equal to the <code>max_burst_length</code> and respect the transaction's wrapping boundary</p> <p><b>Wrapping Slave</b></p> <p>If the <code>WRAP</code> transaction's boundary matches the slave's boundary, then the burst passes through. Otherwise, the burst is converted to sequential bursts that respect both the transaction and slave wrap boundaries.</p>
Fixed	<p>Fixed bursts are converted to sequential bursts of length 1 that repeatedly access the same address.</p>
Narrow	<p>All narrow-sized bursts are broken into multiple bursts of length 1.</p>

## Burst Adaptation: Avalon to AXI

**Table 7-6: Burst Adaptation: Avalon to AXI**

Entries specify the behavior when converting between Avalon and AXI burst types.

Burst Type	Definition
Sequential	<p>Bursts of length greater than 16 are converted to multiple <code>INCR</code> bursts of a length less than or equal to 16. Bursts of length less than or equal to 16 are not converted.</p>

Burst Type	Definition
Wrapping	Only Avalon masters with <code>alwaysBurstMaxBurst = true</code> are supported. The WRAP burst is passed through if the length is less than or equal to 16. Otherwise, it is converted to two or more <code>INCR</code> bursts that respect the transaction's wrap boundary.
<code>GENERIC_CONVERTER</code>	Controls all burst conversions with a single converter that is able to adapt all incoming burst types. This results in an adapter that has smaller area, but lower $f_{Max}$ .

## Responses

Qsys merges write responses if a write is converted (burst adapted) into multiple bursts. Qsys requires read response merging for a downsized (wide-to-narrow width adapted) read.

Qsys merges responses based on the following precedence rule:

DECERR > SLVERR > OKAY > EXOKAY

### Figure 7-16: Mismatched Master and Slave Response Support

The Interconnect sends the slave response transaction back to the master that issues the transaction. When there is a mismatch in response support between master and slave, the interconnect resolves the transaction flow according to the following scenarios.

	Slave with Response	Slave without Response
Master with Response	Interconnect delivers response from the slave to the master.  Response merging or duplication may be necessary for bus sizing.	Interconnect delivers an <code>OKAY</code> response to the master for all reads.
Master without Response	Master ignores responses from the slave.	No need for responses. Master, slave, and interconnect operate without response support.

For the response case where the transaction violates security settings or uses an illegal address, the interconnect routes the transactions to the default slave. For information about Qsys system security and how to specify a default slave, refer to *Creating a System with Qsys*.

**Note:** For Avalon-MM slaves without the `response` signal, there is no way to notify a connected master that a transaction has not completed successfully. As a result, the Qsys interconnect generates an `OKAY` response on behalf of an Avalon-MM slave that does not have the `response` signal.

### Related Information

[Master Network Interfaces](#) on page 7-11

## Qsys Address Decoding

Address decoding logic forwards appropriate addresses to each slave.

Address decoding logic simplifies component design in the following ways:

- The interconnect selects a slave whenever it is being addressed by a master. Slave components do not need to decode the address to determine when they are selected.
- Slave addresses are properly aligned to the slave interface.
- Changing the system memory map does not involve manually editing HDL.

**Figure 7-17: Address Decoding for One Master and Two Slaves**

In this example, Qsys generates separate address decoding logic for each master in a system. The address decoding logic processes the difference between the master address width ( $\langle M \rangle$ ) and the individual slave address widths ( $\langle S \rangle$ ) and ( $\langle T \rangle$ ). The address decoding logic also maps only the necessary master address bits to access words in each slave's address space.





Figure 7-18: Address Decoding Base Settings

Qsys controls the base addresses with the **Base** setting of active components on the **System Contents** tab. The base address of a slave component must be a multiple of the address span of the component. This restriction is part of the Qsys interconnect to allow the address decoding logic to be efficient, and to achieve the best possible  $f_{MAX}$ .



Connections	Name	Description	Export	Clock	Base	End	IRQ
mem_clk	clk_in	Clock Source	mem_clk	exported			
	clk_in_reset	Clock Input	mem_reset				
	clk	Clock Output	Double-click to	mem_clk			
	clk_reset	Reset Output	Double-click to				
cpu_clk	clk_in	Clock Source	cpu_clk	exported			
	clk_in_reset	Reset Input	cpu_reset				
	clk	Clock Output	Double-click to	cpu_clk			
	clk_reset	Reset Output	Double-click to				
cpu	clk	Nios II Processor	Double-click to	cpu_clk			
	reset_n	Reset Input	Double-click to	[clk]			
	data_master	Avalon Memory Mapped Master	Double-click to	[clk]	IRQ 0	IRQ 31	
	instruction_master	Avalon Memory Mapped Master	Double-click to	[clk]			
	jtag_debug_module_reset	Reset Output	Double-click to	cpu_jtag_debug...			
onchip_ram	clk1	On-Chip Memory (RAM or ROM)	Double-click to	cpu_clk			
	s1	Avalon Memory Mapped Slave	Double-click to	[clk1]	0x0001_0000	0x0001_1fff	
	reset1	Reset Input	Double-click to	[clk1]			
jtag_uart	clk	JTAG UART	Double-click to	cpu_clk			
	reset	Reset Input	Double-click to	[clk]			
	avalon_jtag_slave	Avalon Memory Mapped Slave	Double-click to	[clk]	0x0001_2800	0x0001_2807	
pipeline_bridge	clk	Avalon-MM Pipeline Bridge	Double-click to	mem_clk			
	reset	Reset Input	Double-click to	[clk]			
	s0	Avalon Memory Mapped Slave	Double-click to	[clk]	0x0000_0000	0x0000_ffff	
	m0	Avalon Memory Mapped Master	Double-click to	master			

## Avalon Streaming Interfaces

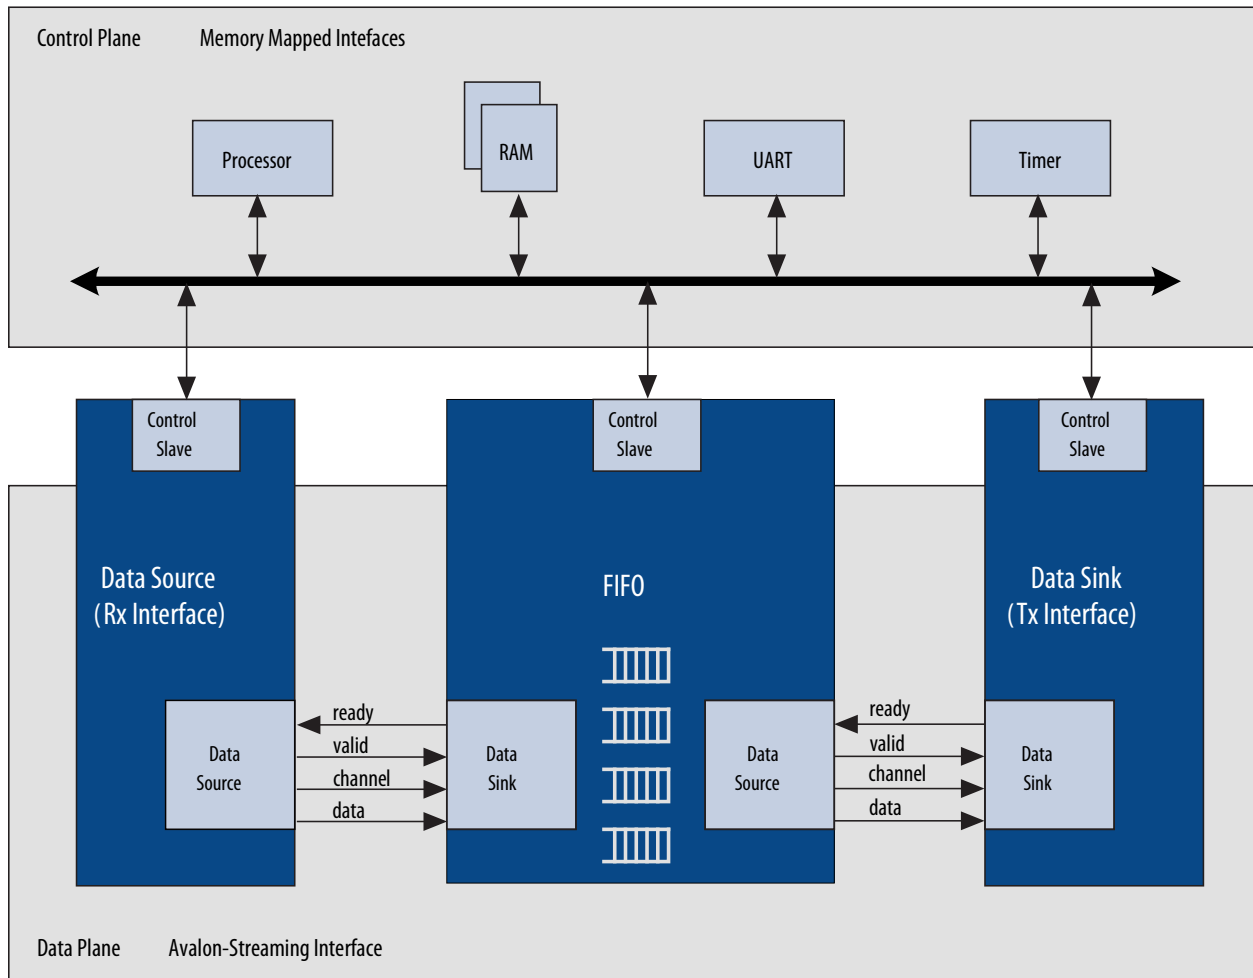
High bandwidth components with streaming data typically use Avalon-ST interfaces for the high throughput datapath. Streaming interfaces can also use memory-mapped connection interfaces to provide an access point for control. In contrast to the memory-mapped interconnect, the Avalon-ST interconnect always creates a point-to-point connection between a single data source and data sink.

**Figure 7-19: Memory-Mapped and Avalon-ST Interfaces**

In this example, there are the following connection pairs:


- Data source in the Rx Interface transfers data to the data sink in the FIFO.
- Data source in the FIFO transfers data to the Tx Interface data sink.

The memory-mapped interface allows a processor to access the data source, FIFO, or data sink to provide system control. If your source and sink interfaces have different formats, for example, a 32-bit source and an 8-bit sink, Qsys automatically inserts the necessary adapters. You can view the adapters on the **System Contents** tab by clicking **System > Show System with Qsys Interconnect**.

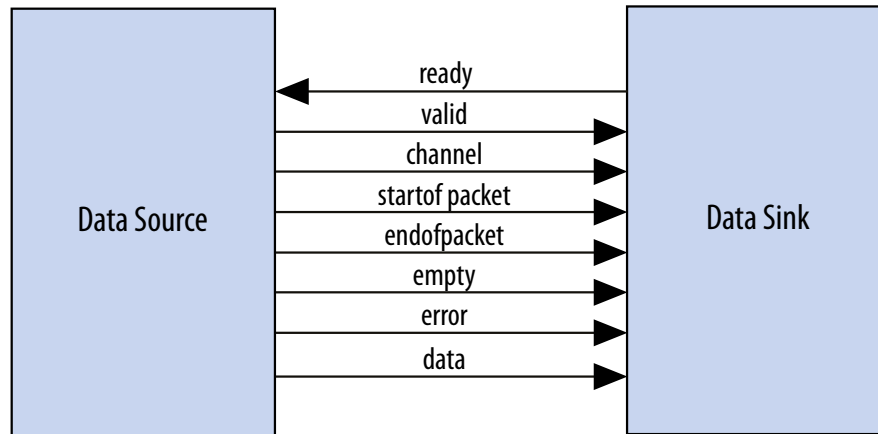


**Figure 7-20: Avalon-ST Connection Between the Source and Sink**

This source-sink pair includes only the `data` signal. The sink must be able to receive data as soon as the source interface comes out of reset.

**Figure 7-21: Signals Indicating the Start and End of Packets, Channel Numbers, Error Conditions, and Backpressure**

All data transfers using Avalon-ST interconnect occur synchronously on the rising edge of the associated clock interface. Throughput and frequency of a system depends on the components and how they are connected.



The IP Catalog includes a number of Avalon-ST components that you can use to create datapaths, including datapaths whose input and output streams have different properties. Generated systems that include memory-mapped master and slave components may also use these Avalon-ST components because Qsys generation creates interconnect with a structure similar to a network topology, as described in *Qsys Transformations*. The following sections introduce the Avalon-ST components.

**Related Information**

- [Avalon-ST Adapters](#) on page 7-28
- [Qsys Transformations](#) on page 7-7
- [Avalon Interface Specification](#)

**Avalon-ST Adapters**

Qsys automatically adds Avalon-ST adapters between two components during system generation when it detects mismatched interfaces. If you connect mismatched Avalon-ST sources and sinks, for example, a

32-bit source and an 8-bit sink, Qsys inserts the appropriate adapter type to connect the mismatched interfaces.

After generation, you can view the inserted adapters with the **Show System With Qsys Interconnect** command in the System menu. For each mismatched source-sink pair, Qsys inserts an Avalon-ST Adapter. The adapter instantiates the necessary adaptation logic as sub-components. You can review the logic for each adapter instantiation in the Hierarchy view by expanding each adapter's source and sink interface and comparing the relevant ports. For example, to determine why a channel adapter is inserted, expand the channel adapter's sink and source interfaces and review the channel port properties for each interface.

You can turn off the auto-inserted adapters feature by adding the `qsys_enable_avalon_streaming_transform=off` command to the **quartus.ini** file. When you turn off the auto-inserted adapters feature, if mismatched interfaces are detected during system generation, Qsys does not insert adapters and reports the mismatched interface with validation error message.

**Note:** The auto-inserted adapters feature does not work for video IP core connections.

## Avalon-ST Adapter

The Avalon-ST adapter combines the logic of the channel, error, data format, and timing adapters. The Avalon-ST adapter provides adaptations between interfaces that have mismatched Avalon-ST endpoints. Based on the source and sink interface parameterizations for the Avalon-ST adapter, Qsys instantiates the necessary adapter logic (channel, error, data format, or timing) as hierarchal sub-components.

### Avalon-ST Adapter Parameters Common to Source and Sink Interfaces

Table 7-7: Avalon-ST Adapter Parameters Common to Source and Sink Interfaces

Parameter Name	Description
<b>Symbol Width</b>	Width of a single symbol in bits.
<b>Use Packet</b>	Indicates whether the source and sink interfaces connected to the adapter's source and sink interfaces include the <code>startofpacket</code> and <code>endofpacket</code> signals, and the optional <code>empty</code> signal.

### Avalon-ST Adapter Upstream Source Interface Parameters

Table 7-8: Avalon-ST Adapter Upstream Source Interface Parameters

Parameter Name	Description
<b>Source Data Width</b>	Controls the data width of the source interface <code>data</code> port.
<b>Source Top Channel</b>	Maximum number of output channels allowed.
<b>Source Channel Port Width</b>	Sets the bit width of the source interface <code>channel</code> port. If set to 0, there is no <code>channel</code> port on the sink interface.
<b>Source Error Port Width</b>	Sets the bit width of the source interface <code>error</code> port. If set to 0, there is no <code>error</code> port on the sink interface.
<b>Source Error Descriptors</b>	A list of strings that describe the error conditions for each bit of the source interface <code>error</code> signal.

Parameter Name	Description
<b>Source Uses Empty Port</b>	Indicates whether the source interface includes the <code>empty</code> port, and whether the sink interface should also include the <code>empty</code> port.
<b>Source Empty Port Width</b>	Indicates the bit width of the source interface <code>empty</code> port, and sets the bit width of the sink interface <code>empty</code> port.
<b>Source Uses Valid Port</b>	Indicates whether the source interface connected to the sink interface uses the <code>valid</code> port, and if set, configures the sink interface to use the <code>valid</code> port.
<b>Source Uses Ready Port</b>	Indicates whether the sink interface uses the <code>ready</code> port, and if set, configures the source interface to use the <code>ready</code> port.
<b>Source Ready Latency</b>	Specifies what ready latency to expect from the source interface connected to the adapter's sink interface.

### Avalon-ST Adapter Downstream Sink Interface Parameters

Table 7-9: Avalon-ST Adapter Downstream Sink Interface Parameters

Parameter Name	Description
<b>Sink Data Width</b>	Indicates the bit width of the <code>data</code> port on the sink interface connected to the source interface.
<b>Sink Top Channel</b>	Maximum number of output channels allowed.
<b>Sink Channel Port Width</b>	Indicates the bit width of the <code>channel</code> port on the sink interface connected the source interface.
<b>Sink Error Port Width</b>	Indicates the bit width of the <code>error</code> port on the sink interface connected to the adapter's source interface. If set to zero, there is no error port on the source interface.
<b>Sink Error Descriptors</b>	A list of strings that describe the error conditions for each bit of the <code>error</code> port on the sink interface connected to the source interface.
<b>Sink Uses Empty Port</b>	Indicates whether the sink interface connected to the source interface uses the <code>empty</code> port, and whether the source interface should also use the <code>empty</code> port.
<b>Sink Empty Port Width</b>	Indicates the bit width of the <code>empty</code> port on the sink interface connected to the source interface, and configures a corresponding <code>empty</code> port on the source interface.
<b>Sink Uses Valid Port</b>	Indicates whether the sink interface connected to the source interface uses the <code>valid</code> port, and if set, configures the source interface to use the <code>valid</code> port.
<b>Sink Uses Ready Port</b>	Indicates whether the <code>ready</code> port on the sink interface is connected to the source interface , and if set, configures the sink interface to use the <code>ready</code> port.



Parameter Name	Description
<b>Sink Ready Latency</b>	Specifies what ready latency to expect from the source interface connected to the sink interface.

## Channel Adapter

The channel adapter provides adaptations between interfaces that have different channel signal widths.

**Table 7-10: Channel Adapter Adaptations**

Condition	Description of Adapter Logic
The source uses channels, but the sink does not.	Qsys gives a warning at generation time. The adapter provides a simulation error and signals an error for data for any channel from the source other than 0.
The sink has channel, but the source does not.	Qsys gives a warning at generation time, and the channel inputs to the sink are all tied to a logical 0.
The source and sink both support channels, and the source's maximum channel number is less than the sink's maximum channel number.	The source's channel is connected to the sink's channel unchanged. If the sink's channel signal has more bits, the higher bits are tied to a logical 0.
The source and sink both support channels, but the source's maximum channel number is greater than the sink's maximum channel number.	The source's channel is connected to the sink's channel unchanged. If the source's channel signal has more bits, the higher bits are left unconnected. Qsys gives a warning that channel information may be lost.  An adapter provides a simulation error message and an error indication if the value of channel from the source is greater than the sink's maximum number of channels. In addition, the <code>valid</code> signal to the sink is deasserted so that the sink never sees data for channels that are out of range.

## Avalon-ST Channel Adapter Input Interface Parameters

**Table 7-11: Avalon-ST Channel Adapter Input Interface Parameters**

Parameter Name	Description
<b>Channel Signal Width (bits)</b>	Width of the input channel signal in bits
<b>Max Channel</b>	Maximum number of input channels allowed.

## Avalon-ST Channel Adapter Output Interface Parameters

Table 7-12: Avalon-ST Channel Adapter Output Interface Parameters

Parameter Name	Description
Channel Signal Width (bits)	Width of the output channel signal in bits.
Max Channel	Maximum number of output channels allowed.

## Avalon-ST Channel Adapter Common to Input and Output Interface Parameters

Table 7-13: Avalon-ST Channel Adapter Common to Input and Output Interface Parameters

Parameter Name	Description
Data Bits Per Symbol	Number of bits for each symbol in a transfer.
Include Packet Support	When the Avalon-ST Channel adapter supports packets, the <code>startofpacket</code> , <code>endofpacket</code> , and optional <code>empty</code> signals are included on its sink and source interfaces.
Include Empty Signal	Indicates whether an <code>empty</code> signal is required.
Data Symbols Per Beat	Number of symbols per transfer.
Support Backpressure with the ready signal	Indicates whether a <code>ready</code> signal is required.
Ready Latency	Specifies the ready latency to expect from the sink connected to the module's source interface.
Error Signal Width (bits)	Bit width of the <code>error</code> signal.
Error Signal Description	A list of strings that describes what each bit of the <code>error</code> signal represents.

### Data Format Adapter


The data format adapter allows you to connect interfaces that have different values for the parameters defining the `data` signal, or interfaces where the source does not use the `empty` signal, but the sink does use the `empty` signal. One of the most common uses of this adapter is to convert data streams of different widths.

**Table 7-14: Data Format Adapter Adaptations**

Condition	Description of Adapter Logic
The source and sink's bits per symbol parameters are different.	The connection cannot be made.
The source and sink have a different number of symbols per beat.	<p>The adapter converts the source's width to the sink's width.</p> <p>If the adaptation is from a wider to a narrower interface, a beat of data at the input corresponds to multiple beats of data at the output. If the input <code>error</code> signal is asserted for a single beat, it is asserted on output for multiple beats.</p> <p>If the adaptation is from a narrow to a wider interface, multiple input beats are required to fill a single output beat, and the output <code>error</code> is the logical OR of the input <code>error</code> signal.</p>
The source uses the <code>empty</code> signal, but the sink does not use the <code>empty</code> signal.	Qsys cannot make the connection.

**Figure 7-22: Avalon Streaming Interconnect with Data Format Adapter**

In this example, the data format adapter allows a connection between a 128-bit output data stream and three 32-bit input data streams.



### Avalon-ST Data Format Adapter Input Interface Parameters

**Table 7-15: Avalon-ST Data Format Adapter Input Interface Parameters**

Parameter Name	Description
<b>Data Symbols Per Beat</b>	Number of symbols per transfer.
<b>Include Empty Signal</b>	Indicates whether an <code>empty</code> signal is required.

### Avalon-ST Data Format Adapter Output Interface Parameters

**Table 7-16: Avalon-ST Data Format Adapter Output Interface Parameters**

Parameter Name	Description
<b>Data Symbols Per Beat</b>	Number of symbols per transfer.
<b>Include Empty Signals</b>	Indicates whether an <code>empty</code> signal is required.

## Avalon-ST Data Format Adapter Common to Input and Output Interface Parameters

Table 7-17: Avalon-ST Data Format Adapter Common to Input and Output Interface Parameters

Parameter Name	Description
<b>Data Bits Per Symbol</b>	Number of bits for each symbol in a transfer.
<b>Include Packet Support</b>	When the Avalon-ST Data Format adapter supports packets, Qsys uses <code>startofpacket</code> , <code>endofpacket</code> , and <code>empty</code> signals.
<b>Channel Signal Width (bits)</b>	Width of the output channel signal in bits.
<b>Max Channel</b>	Maximum number of channels allowed.
<b>Read Latency</b>	Specifies the ready latency to expect from the sink connected to the module's source interface.
<b>Error Signal Width (bits)</b>	Width of the <code>error</code> signal output in bits.
<b>Error Signal Description</b>	A list of strings that describes what each bit of the <code>error</code> signal represents.

### Error Adapter

The error adapter ensures that per-bit-error information provided by the source interface is correctly connected to the sink interface's input error signal. Error conditions that both the source and sink are able to process are connected. If the source has an `error` signal representing an error condition that is not supported by the sink, the signal is left unconnected; the adapter provides a simulation error message and an error indication if the error is asserted. If the sink has an error condition that is not supported by the source, the sink's input error bit corresponding to that condition is set to 0.

**Note:** The output interface error signal descriptor accepts an error set with an `other` descriptor. Qsys assigns the bit-wise `ORing` of all input error bits that are unmatched, to the output interface error bits set with the `other` descriptor.

### Avalon-ST Error Adapter Input Interface Parameters

Table 7-18: Avalon-ST Error Adapter Input Interface Parameters

Parameter Name	Description
<b>Error Signal Width (bits)</b>	The width of the <code>error</code> signal. Valid values are 0–256 bits. Type 0 if the <code>error</code> signal is not used.
<b>Error Signal Description</b>	The description for each of the error bits. If scripting, separate the description fields by commas. For a successful connection, the description strings of the error bits in the source and sink must match and are case sensitive.

## Avalon-ST Error Adapter Output Interface Parameters

Table 7-19: Avalon-ST Error Adapter Output Interface Parameters

Parameter Name	Description
<b>Error Signal Width (bits)</b>	The width of the <code>error</code> signal. Valid values are 0–256 bits. Type 0 if you do not need to send error values.
<b>Error Signal Description</b>	The description for each of the error bits. Separate the description fields by commas. For successful connection, the description of the error bits in the source and sink must match, and are case sensitive.

## Avalon-ST Error Adapter Common to Input and Output Interface Parameters

Table 7-20: Avalon-ST Error Adapter Common to Input and Output Interface Parameters

Parameter Name	Description
<b>Support Backpressure with the ready signal</b>	Turn on this option to add the backpressure functionality to the interface.
<b>Ready Latency</b>	When the <code>ready</code> signal is used, the value for <code>ready_latency</code> indicates the number of cycles between when the ready signal is asserted and when valid data is driven.
<b>Channel Signal Width (bits)</b>	The width of the <code>channel</code> signal. A channel width of 4 allows up to 16 channels. The maximum width of the <code>channel</code> signal is eight bits. Set to 0 if channels are not used.
<b>Max Channel</b>	The maximum number of channels that the interface supports. Valid values are 0–255.
<b>Data Bits Per Symbol</b>	Number of bits per symbol.
<b>Data Symbols Per Beat</b>	Number of symbols per active transfer.
<b>Include Packet Support</b>	Turn on this option if the connected interfaces support a packet protocol, including the <code>startofpacket</code> , <code>endofpacket</code> and <code>empty</code> signals.
<b>Include Empty Signal</b>	Turn this option on if the cycle that includes the <code>endofpacket</code> signal can include empty symbols. This signal is not necessary if the number of symbols per beat is 1.

## Timing Adapter

The timing adapter allows you to connect component interfaces that require a different number of cycles before driving or receiving data. This adapter inserts a FIFO buffer between the source and sink to buffer data or pipeline stages to delay the back pressure signals. You can also use the timing adapter to connect interfaces that support the `ready` signal, and those that do not. The timing adapter treats all signals other than the `ready` and `valid` signals as payload, and simply drives them from the source to the sink.

**Table 7-21: Timing Adapter Adaptations**

Condition	Adaptation
The source has <code>ready</code> , but the sink does not.	In this case, the source can respond to <code>backpressure</code> , but the sink never needs to apply it. The <code>ready</code> input to the source interface is connected directly to logical 1.
The source does not have <code>ready</code> , but the sink does.	The sink may apply <code>backpressure</code> , but the source is unable to respond to it. There is no logic that the adapter can insert that prevents data loss when the source asserts <code>valid</code> but the sink is not ready. The adapter provides simulation time error messages if data is lost. The user is presented with a warning, and the connection is allowed.
The source and sink both support <code>backpressure</code> , but the sink's <code>ready</code> latency is greater than the source's.	The source responds to <code>ready</code> assertion or deassertion faster than the sink requires it. A number of pipeline stages equal to the difference in <code>ready</code> latency are inserted in the <code>ready</code> path from the sink back to the source, causing the source and the sink to see the same cycles as <code>ready</code> cycles.
The source and sink both support <code>backpressure</code> , but the sink's <code>ready</code> latency is less than the source's.	The source cannot respond to <code>ready</code> assertion or deassertion in time to satisfy the sink. A FIFO whose depth is equal to the difference in <code>ready</code> latency is inserted to compensate for the source's inability to respond in time.

### Avalon-ST Timing Adapter Input Interface Parameters

**Table 7-22: Avalon-ST Timing Adapter Input Interface Parameters**

Parameter Name	Description
<b>Support Backpressure with the ready signal</b>	Indicates whether a <code>ready</code> signal is required.
<b>Read Latency</b>	Specifies the <code>ready</code> latency to expect from the sink connected to the module's source interface.

Parameter Name	Description
<b>Include Valid Signal</b>	Indicates whether the sink interface requires a valid signal.

### Avalon-ST Timing Adapter Output Interface Parameters

Table 7-23: Avalon-ST Timing Adapter Output Interface Parameters

Parameter Name	Description
<b>Support Backpressure with the ready signal</b>	Indicates whether a <code>ready</code> signal is required.
<b>Read Latency</b>	Specifies the ready latency to expect from the sink connected to the module's source interface.
<b>Include Valid Signal</b>	Indicates whether the sink interface requires a valid signal.

### Avalon-ST Timing Adapter Common to Input and Output Interface Parameters

Table 7-24: Avalon-ST Timing Adapter Common to Input and Output Interface Parameters

Parameter Name	Description
<b>Data Bits Per Symbol</b>	Number of bits for each symbol in a transfer.
<b>Include Packet Support</b>	Turn this option on if the connected interfaces support a packet protocol, including the <code>startofpacket</code> , <code>endofpacket</code> and <code>empty</code> signals.
<b>Include Empty Signal</b>	Turn this option on if the cycle that includes the <code>endofpacket</code> signal can include empty symbols. This signal is not necessary if the number of symbols per beat is 1.
<b>Data Symbols Per Beat</b>	Number of symbols per active transfer.
<b>Channel Signal Width (bits)</b>	Width of the output channel signal in bits.
<b>Max Channel</b>	Maximum number of output channels allowed.
<b>Error Signal Width (bits)</b>	Width of the output <code>error</code> signal in bits.
<b>Error Signal Description</b>	A list of strings that describes errors.



## Interrupt Interfaces

Using individual requests, the interrupt logic can process up to 32 IRQ inputs connected to each interrupt receiver. With this logic, the interrupt sender connected to interrupt `receiver_0` is the highest priority with sequential receivers being successively lower priority. You can redefine the priority of interrupt senders by instantiating the IRQ mapper component. For more information refer to *IRQ Mapper*.

You can define the interrupt sender interface as asynchronous with no associated clock or reset interfaces. You can also define the interrupt receiver interface as asynchronous with no associated clock or reset interfaces. As a result, the receiver does its own synchronization internally. Qsys does not insert interrupt synchronizers for such receivers.

For clock crossing adaption on interrupts, Qsys inserts a synchronizer, which is clocked with the interrupt end point interface clock when the corresponding starting point interrupt interface has no clock or a different clock (than the end point). Qsys inserts the adapter if there is any kind of mismatch between the start and end points. Qsys does not insert the adapter if the interrupt receiver does not have an associated clock.

### Related Information

[IRQ Mapper](#) on page 7-41

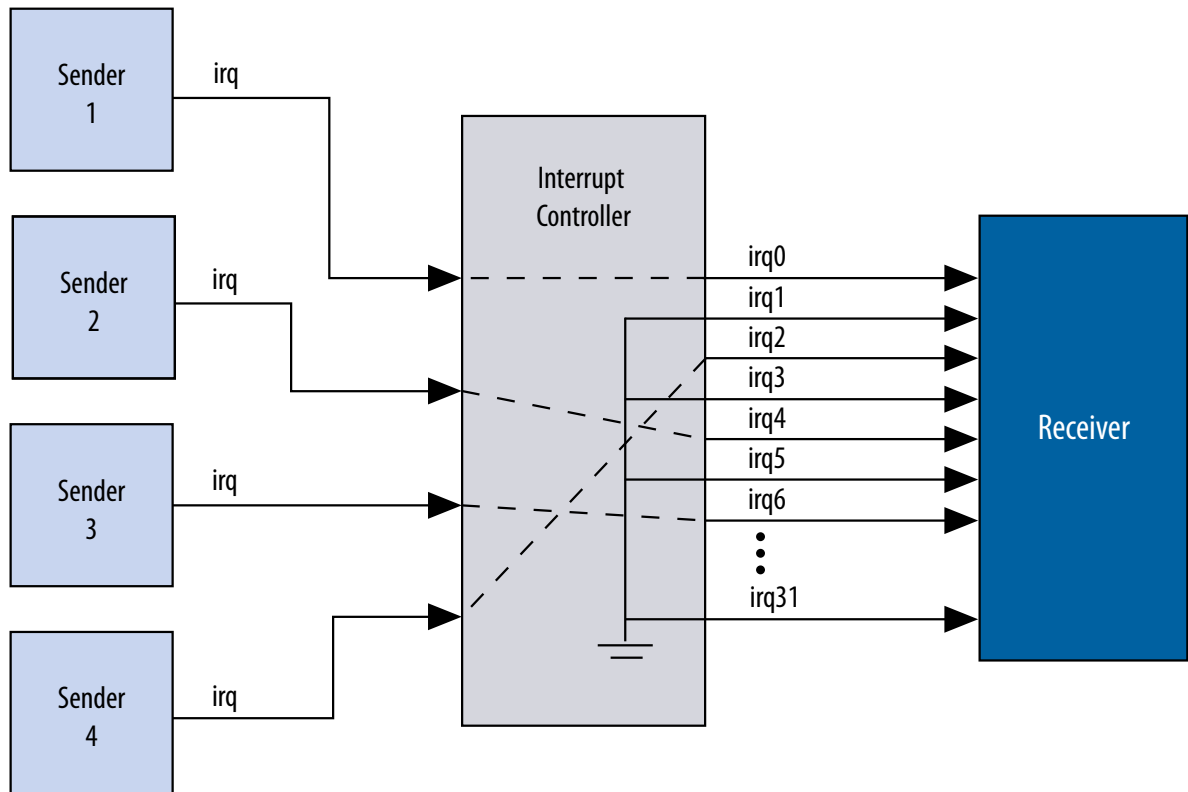
## Individual Requests IRQ Scheme

In the individual requests IRQ scheme, Qsys interconnect passes IRQs directly from the sender to the receiver, without making assumptions about IRQ priority. In the event that multiple senders assert their

IRQs simultaneously, the receiver logic determines which IRQ has highest priority, and then responds appropriately.

### Figure 7-23: Interrupt Controller Mapping IRQs

Using individual requests, the interrupt controller can process up to 32 IRQ inputs. The interrupt controller generates a 32-bit signal `irq[31:0]` to the receiver, and maps slave IRQ signals to the bits of `irq[31:0]`. Any unassigned bits of `irq[31:0]` are disabled.



## Assigning IRQs in Qsys

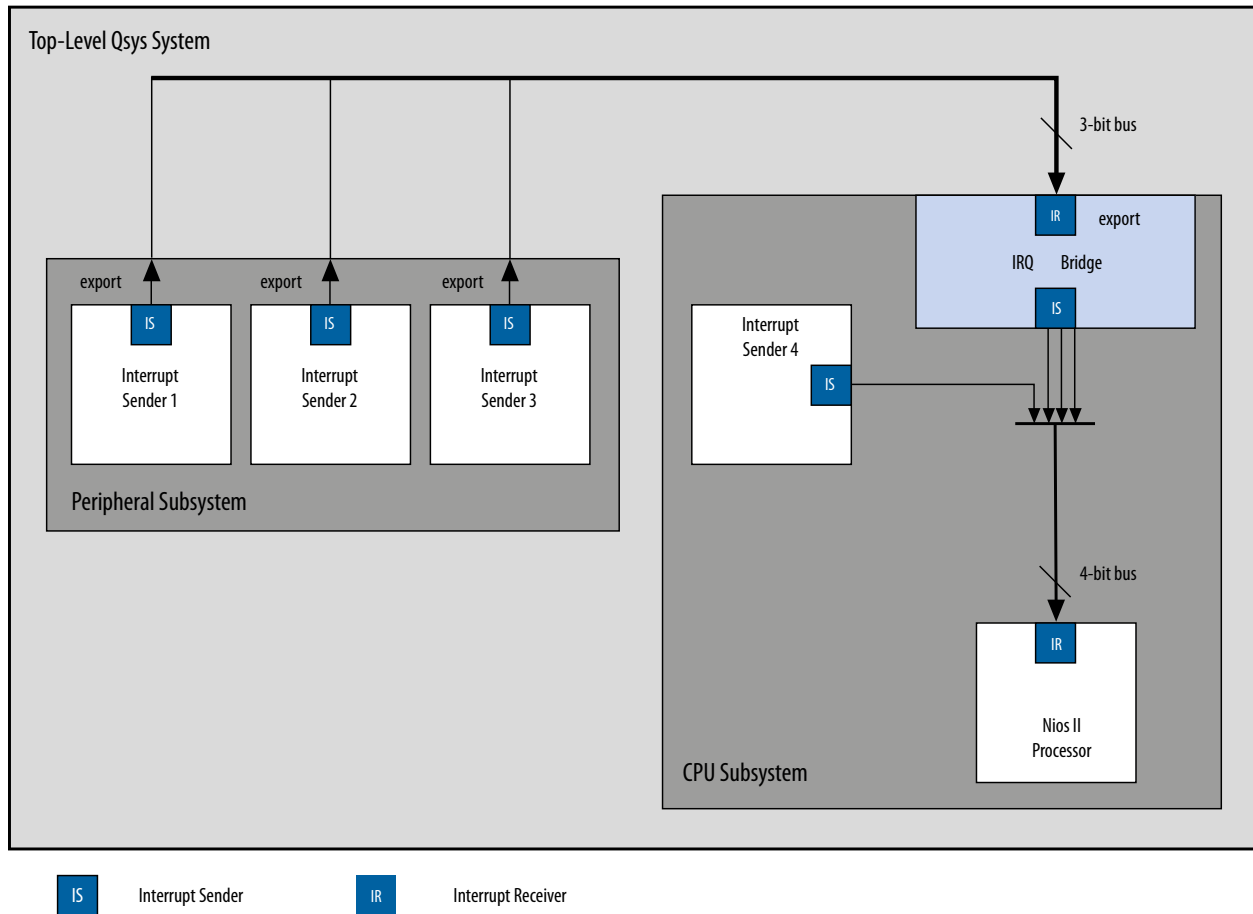
You assign IRQ connections on the **System Contents** tab of Qsys. After adding all components to the system, you connect interrupt senders and receivers. You can use the **IRQ** column to specify an IRQ number with respect to each receiver, or to specify a receiver's IRQ as unconnected. Qsys uses the following three components to implement interrupt handling: IRQ Bridge, IRQ Mapper, and IRQ Clock Crosser.

### IRQ Bridge

The IRQ Bridge allows you to route interrupt wires between Qsys subsystems.

**Figure 7-24: Qsys IRQ Bridge Application**

The peripheral subsystem example below has three interrupt senders that are exported to the to- level of the subsystem. The interrupts are then routed to the CPU subsystem using the IRQ bridge.



**Note:** Nios II BSP tools support the IRQ Bridge. Interrupts connected via an IRQ Bridge appear in the generated **system.h** file. You can use the following properties with the IRQ Bridge, which do not effect Qsys interconnect generation. Qsys uses these properties to generate the correct IRQ information for downstream tools:

- **set\_interface\_property <sender port> bridgesToReceiver <receiver port>**—The <sender port> of the IP generates a signal that is received on the IP's <receiver port>. Sender ports are single bits. Receivers ports can be multiple bits. Qsys requires the `bridgedReceiverOffset` property to identify the <receiver port> bit that the <sender port> sends.
- **set\_interface\_property <sender port> bridgedReceiverOffset <port number>**—Indicates the <port number> of the receiver port that the <sender port> sends.

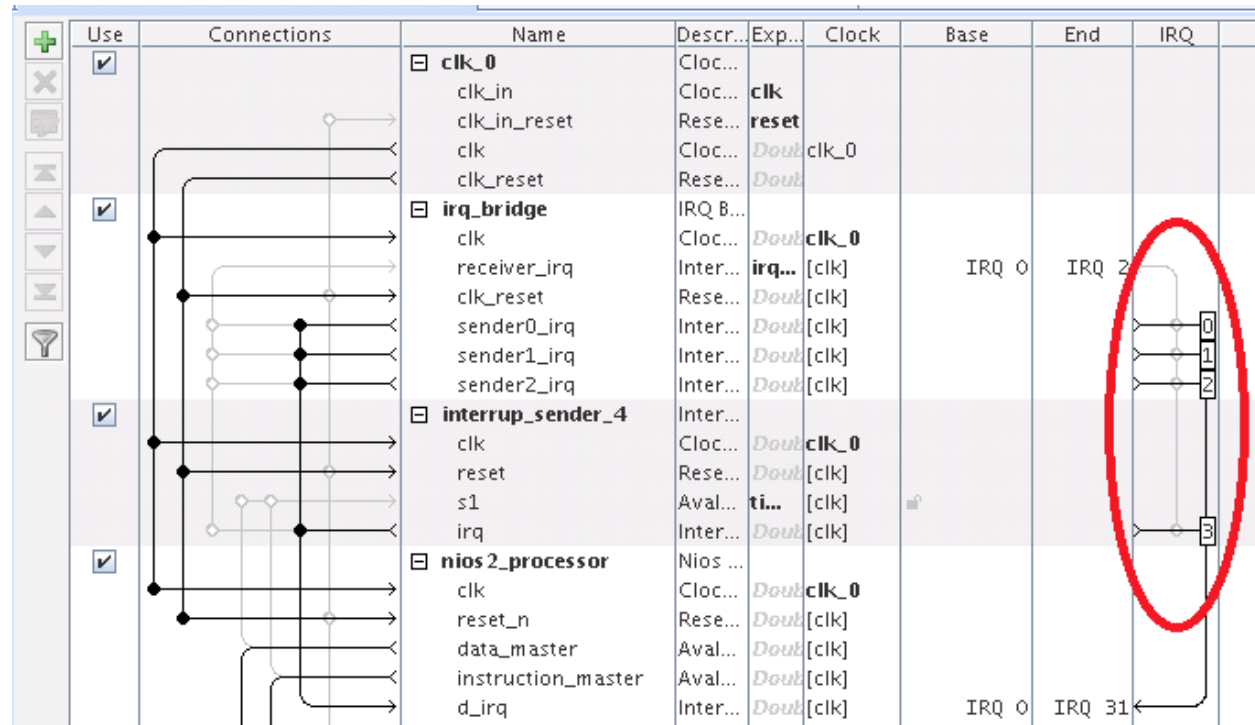
### IRQ Mapper

Qsys inserts the IRQ Mapper automatically during generation. The IRQ Mapper converts individual interrupt wires to a bus, and then maps the appropriate IRQ priority number onto the bus.

By default, the interrupt sender connected to the `receiver0` interface of the IRQ mapper is the highest priority, and sequential receivers are successively lower priority. You can modify the interrupt priority of each IRQ wire by modifying the IRQ priority number in Qsys under the **IRQ** column. The modified priority is reflected in the **IRQ\_MAP** parameter for the auto-inserted IRQ Mapper.

**Figure 7-25: IRQ Column in Qsys**

Circled in the **IRQ** column are the default interrupt priorities allocated for the CPU subsystem.



#### Related Information

[IRQ Bridge](#) on page 7-40

### IRQ Clock Crosser

The IRQ Clock Crosser synchronizes interrupt senders and receivers that are in different clock domains. To use this component, connect the clocks for both the interrupt sender and receiver, and for both the interrupt sender and receiver interfaces. Qsys automatically inserts this component when it is required.

### Clock Interfaces

Clock interfaces define the clocks used by a component. Components can have clock inputs, clock outputs, or both. You can use the **Clock Settings** tab to define external clock sources, for example an oscillator on your board.

The **Clock Source** parameters allows you to set the following options:

- **Clock frequency**—The frequency of the output clock from this clock source.
- **Clock frequency is known**— When turned on, the clock frequency is known. When turned off, the frequency is set from outside the system.

**Note:** If turned off, system generation may fail because the components do not receive the necessary clock information. For best results, turn this option on before system generation.

- **Reset synchronous edges**
  - **None**—The reset is asserted and deasserted asynchronously. You can use this setting if you have internal synchronization circuitry that matches the reset required for the IP in the system.
  - **Both**—The reset is asserted and deasserted synchronously.
  - **Deassert**—The reset is deasserted synchronously and asserted asynchronously.

For more information about synchronous design practices, refer to *Recommended Design Practices*

**Related Information**

- [Recommended Design Practices](#) on page 11-1

## (High Speed Serial Interface) HSSI Clock Interfaces

You can use HSSI Serial Clock and HSSI Bonded Clock interfaces in Qsys to enable high speed serial connectivity between clocks that are used by certain IP protocols.

### HSSI Serial Clock Interface

You can connect the HSSI Serial Clock interface with only similar type of interfaces, for example, you can connect a HSSI Serial Clock Source interface to a HSSI Serial Clock Sink interface.

#### HSSI Serial Clock Source

The HSSI Serial Clock interface includes a source in the **Start** direction.

You can instantiate the HSSI Serial Clock Source interface in the `_hw.tcl` file as:

```
add_interface <name> hssi_serial_clock start
```

You can connect the HSSI Serial Clock Source to multiple HSSI Serial Clock Sinks because the HSSI Serial Clock Source supports multiple fan-outs. This Interface has a single **clk** port role limited to a 1 bit width, and a **clockRate** parameter, which is the frequency of the clock driven by the HSSI Serial Clock Source interface.

An unconnected and unexported HSSI Serial Source is valid and does not generate error messages.

**Table 7-25: HSSI Serial Clock Source Port Roles**

Name	Direction	Width	Description
clk	Output	1 bit	A single bit wide port role, which provides synchronization for internal logic.

**Table 7-26: HSSI Serial Clock Source Parameters**

Name	Type	Default	Derived	Description
clockRate	long	0	No	The frequency of the clock driven by the HSSI Serial Clock Source interface.

**HSSI Serial Clock Sink**

The HSSI Serial Clock interface includes a sink in the **End** direction.

You can instantiate the HSSI Serial Clock Sink interface in the `_hw.tcl` file as:

```
add_interface <name> hssi_serial_clock end
```

You can connect the HSSI Serial Clock Sink interface to a single HSSI Serial Clock Source interface; you cannot connect it to multiple sources. This interface has a single **clk** port role limited to a 1 bit width, and a **clockRate** parameter, which is the frequency of the clock driven by the HSSI Serial Clock Source interface.

An unconnected and unexported HSSI Serial Sink is invalid and generates error messages.

**Table 7-27: HSSI Serial Clock Sink Port Roles**

Name	Direction	Width	Description
clk	Output	1	A single bit wide port role, which provides synchronization for internal logic

**Table 7-28: HSSI Serial Clock Sink Parameters**

Name	Type	Default	Derived	Description
clockRate	long	0	No	The frequency of the clock driven by the HSSI Serial Clock Source interface. When you specify a <b>clockRate</b> greater than 0, then this interface can be driven only at that rate.

**HSSI Serial Clock Connection**

The HSSI Serial Clock Connection defines a connection between a HSSI Serial Clock Source connection point, and a HSSI Serial Clock Sink connection point.

A valid HSSI Serial Clock Connection exists when all of the following criteria are satisfied. If the following criteria are not satisfied, Qsys generates error messages and the connection is prohibited.

- The starting connection point is an HSSI Serial Clock Source with a single port role **clk** and maximum 1 bit in width. The direction of the starting port is **Output**.
- The ending connection point is an HSSI Serial Clock Sink with a single port role **clk**, and maximum 1 bit in width. The direction of the ending port is **Input**.
- If the parameter, **clockRate** of the HSSI Serial Clock Sink is greater than 0, the connection is only valid if the **clockRate** of the HSSI Serial Clock Source is the same as the **clockRate** of the HSSI Serial Clock Sink.

## HSSI Serial Clock Example

### Example 7-1: HSSI Serial Clock Interface Example

You can make connections to declare the HSSI Serial Clock interfaces in the `_hw.tcl`.

```
package require -exact qsys 14.0

set_module_property name hssi_serial_component
set_module_property ELABORATION_CALLBACK elaborate

add_fileset QUARTUS_SYNTH QUARTUS_SYNTH generate
add_fileset SIM_VERILOG SIM_VERILOG generate
add_fileset SIM_VHDL SIM_VHDL generate

set_fileset_property QUARTUS_SYNTH TOP_LEVEL \
"hssi_serial_component"

set_fileset_property SIM_VERILOG TOP_LEVEL "hssi_serial_component"
set_fileset_property SIM_VHDL TOP_LEVEL "hssi_serial_component"

proc elaborate {} {
    # declaring HSSI Serial Clock Source
    add_interface my_clock_start hssi_serial_clock start
    set_interface_property my_clock_start ENABLED true

    add_interface_port my_clock_start hssi_serial_clock_port_out \
    clk Output 1

    # declaring HSSI Serial Clock Sink
    add_interface my_clock_end hssi_serial_clock end
    set_interface_property my_clock_end ENABLED true

    add_interface_port my_clock_end hssi_serial_clock_port_in clk \
    Input 1
}

proc generate { output_name } {
    add_fileset_file hssi_serial_component.v VERILOG PATH \
    "hssi_serial_component.v"
}
```

### Example 7-2: HSSI Serial Clock Instantiated in a Composed Component

If you use the components in a hierarchy, for example, instantiated in a composed component, you can declare the connections as illustrated in this example.

```
add_instance myinst1 hssi_serial_component
add_instance myinst2 hssi_serial_component
# add connection from source of myinst1 to sink of myinst2

add_connection myinst1.my_clock_start myinst2.my_clock_end \
hssi_serial_clock

# adding connection from source of myinst2 to sink of myinst1

add_connection myinst2.my_clock_start myinst2.my_clock_end \
hssi_serial_clock
```

## HSSI Bonded Clock Interface

You can connect the HSSI Bonded Clock interface with only similar type of Interfaces, for example, you can connect a HSSI Bonded Clock Source interface to a HSSI Bonded Clock Sink interface.

### HSSI Bonded Clock Source

The HSSI Bonded Clock interface includes a source in the **Start** direction.

You can instantiate the HSSI Bonded Clock Source interface in the `_hw.tcl` file as:

```
add_interface <name> hssi_bonded_clock start
```

You can connect the HSSI Bonded Clock Source to multiple HSSI Bonded Clock Sinks because the HSSI Serial Clock Source supports multiple fanouts. This Interface has a single **clk** port role limited to a width range of 1 to 1024 bits. The HSSI Bonded Clock Source interface has two parameters: **clockRate** and **serializationFactor**. **clockRate** is the frequency of the clock driven by the HSSI Bonded Clock Source interface, and the **serializationFactor** is the parallel data width that operates the HSSI TX serializer. The serialization factor determines the required frequency and phases of the individual clocks within the HSSI Bonded Clock interface

An unconnected and unexported HSSI Bonded Source is valid and does not generate error messages.

**Table 7-29: HSSI Bonded Clock Source Port Roles**

Name	Direction	Width	Description
clk	Output	1 to 24 bits	A multiple bit wide port role which provides synchronization for internal logic.

**Table 7-30: HSSI Bonded Clock Source Parameters**

Name	Type	Default	Derived	Description
clockRate	long	0	No	The frequency of the clock driven by HSSI Serial Clock Source interface.
serialization	long	0	No	The serialization factor is the parallel data width that operates the HSSI TX serializer. The serialization factor determines the necessary frequency and phases of the individual clocks within the HSSI Bonded Clock interface.

### HSSI Bonded Clock Sink

The HSSI Bonded Clock interface includes a sink in the **End** direction.

You can instantiate the HSSI Bonded Clock Sink interface in the `_hw.tcl` file as:

```
add_interface <name> hssi_bonded_clock end
```

You can connect the HSSI Bonded Clock Sink interface to a single HSSI Bonded Clock Source interface; you cannot connect it to multiple sources. This Interface has a single **clk** port role limited to a width range of 1 to 1024 bits. The HSSI Bonded Clock Source interface has two parameters: **clockRate** and **serializa-**



**tionFactor**. **clockRate** is the frequency of the clock driven by the HSSI Bonded Clock Source interface, and the serialization factor is the parallel data width that operates the HSSI TX serializer. The serialization factor determines the required frequency and phases of the individual clocks within the HSSI Bonded Clock interface

An unconnected and unexported HSSI Bonded Sink is invalid and generates error messages.

**Table 7-31: HSSI Bonded Clock Source Port Roles**

Name	Direction	Width	Description
clk	Output	1 to 24 bits	A multiple bit wide port role which provides synchronization for internal logic.

**Table 7-32: HSSI Bonded Clock Source Parameters**

Name	Type	Default	Derived	Description
clockRate	long	0	No	The frequency of the clock driven byte HSSI Serial Clock Source interface.
serialization	long	0	No	The serialization factor is the parallel data width that operates the HSSI TX serializer. The serialization factor determines the necessary frequency and phases of the individual clocks within the HSSI Bonded Clock interface.

### HSSI Bonded Clock Connection

The HSSI Bonded Clock Connection defines a connection between a HSSI Bonded Clock Source connection point, and a HSSI Bonded Clock Sink connection point.

A valid HSSI Bonded Clock Connection exists when all of the following criteria are satisfied. If the following criteria are not satisfied, Qsys generates error messages and the connection is prohibited.

- The starting connection point is an HSSI Bonded Clock Source with a single port role **clk** with a width range of 1 to 24 bits. The direction of the starting port is **Output**.
- The ending connection point is an HSSI Bonded Clock Sink with a single port role **clk** with a width range of 1 to 24 bits. The direction of the ending port is **Input**.
- The width of the starting connection point **clk** must be the same as the width of the ending connection point.
- If the parameter, **clockRate** of the HSSI Bonded Clock Sink greater than 0, then the connection is only valid if the **clockRate** of the HSSI Bonded Clock Source is same as the **clockRate** of the HSSI Bonded Clock Sink.
- If the parameter, **serializationFactor** of the HSSI Bonded Clock Sink is greater than 0, Qsys generates a warning if the **serializationFactor** of HSSI Bonded Clock Source is not same as the **serializationFactor** of the HSSI Bonded Clock Sink.

## HSSI Bonded Clock Example

### Example 7-3: HSSI Bonded Clock Interface Example

You can make connections to declare the HSSI Bonded Clock interfaces in the `_hw.tcl` file.

```
package require -exact qsys 14.0

set_module_property name hssi_bonded_component
set_module_property ELABORATION_CALLBACK elaborate

add_fileset synthesis QUARTUS_SYNTH generate
add_fileset verilog_simulation SIM_VERILOG generate

set_fileset_property synthesis TOP_LEVEL "hssi_bonded_component"

set_fileset_property verilog_simulation TOP_LEVEL \
"hssi_bonded_component"

proc elaborate {} {
    add_interface my_clock_start hssi_bonded_clock start
    set_interface_property my_clock_start ENABLED true

    add_interface_port my_clock_start hssi_bonded_clock_port_out \
    clk Output 1024

    add_interface my_clock_end hssi_bonded_clock end
    set_interface_property my_clock_end ENABLED true

    add_interface_port my_clock_end hssi_bonded_clock_port_in \
    clk Input 1024
}

proc generate { output_name } {
    add_fileset_file hssi_bonded_component.v VERILOG PATH \
    "hssi_bonded_component.v"
}
```

If you use the components in a hierarchy, for example, instantiated in a composed component, you can declare the connections as illustrated in this example.

### Example 7-4: HSII Bonded Clock Instantiated in a Composed Component

```
add_instance myinst1 hssi_bonded_component
add_instance myinst2 hssi_bonded_component
# add connection from source of myinst1 to sink of myinst2

add_connection myinst1.my_clock_start myinst2.my_clock_end \
hssi_bonded_clock

# adding connection from source of myinst2 to sink of myinst1

add_connection myinst2.my_clock_start myinst2.my_clock_end \
hssi_bonded_clock
```

## Reset Interfaces

Reset interfaces provide both soft and hard reset functionality. Soft reset logic typically re-initializes registers and memories without powering down the device. Hard reset logic initializes the device after power-on. You can define separate reset sources for each clock domain, a single reset source for all clocks, or any combination in between.

You can choose to create a single global reset domain by selecting **Create Global Reset Network** on the System menu. If your design requires more than one reset domain, you can implement your own reset logic and connectivity. The IP Catalog includes a reset controller, reset sequencer, and a reset bridge to implement the reset functionality. You can also design your own reset logic.

**Note:** If you design your own reset circuitry, you must carefully consider situations which may result in system lockup. For example, if an Avalon-MM slave is reset in the middle of a transaction, the Avalon-MM master may lockup.

### Single Global Reset Signal Implemented by Qsys

If you select **Create Global Reset Network** on the System menu, the Qsys interconnect creates a global reset bus. All of the reset requests are ORed together, synchronized to each clock domain, and fed to the reset inputs. The duration of the reset signal is at least one clock period.

The Qsys interconnect inserts the system-wide reset under the following conditions:

- The global reset input to the Qsys system is asserted.
- Any component asserts its `resetrequest` signal.

### Reset Controller

Qsys automatically inserts a reset controller block if the input reset source does not have a reset request, but the connected reset sink requires a reset request.

The Reset Controller has the following parameters that you can specify to customize its behavior:

- **Number of inputs**— Indicates the number of individual reset interfaces the controller ORs to create a signal reset output.
- **Output reset synchronous edges**—Specifies the level of synchronization. You can select one of the following options:
  - **None**—The reset is asserted and deasserted asynchronously. You can use this setting if you have designed internal synchronization circuitry that matches the reset style required for the IP in the system.
  - **Both**—The reset is asserted and deasserted synchronously.
  - **Deassert**—The reset is deasserted synchronously and asserted asynchronously.
- **Synchronization depth**—Specifies the number of register stages the synchronizer uses to eliminate the propagation of metastable events.
- **Reset request**—Enables reset request generation, which is an early signal that is asserted before reset assertion. The reset request is used by blocks that require protection from asynchronous inputs, for example, M20K blocks.

Qsys automatically inserts reset synchronizers under the following conditions:

- More than one reset source is connected to a reset sink
- There is a mismatch between the reset source's synchronous edges and the reset sinks' synchronous edges

## Reset Bridge

The Reset Bridge allows you to use a reset signal in two or more subsystems of your Qsys system. You can connect one reset source to local components, and export one or more to other subsystems, as required.

The Reset Bridge parameters are used to describe the incoming reset and include the following options:

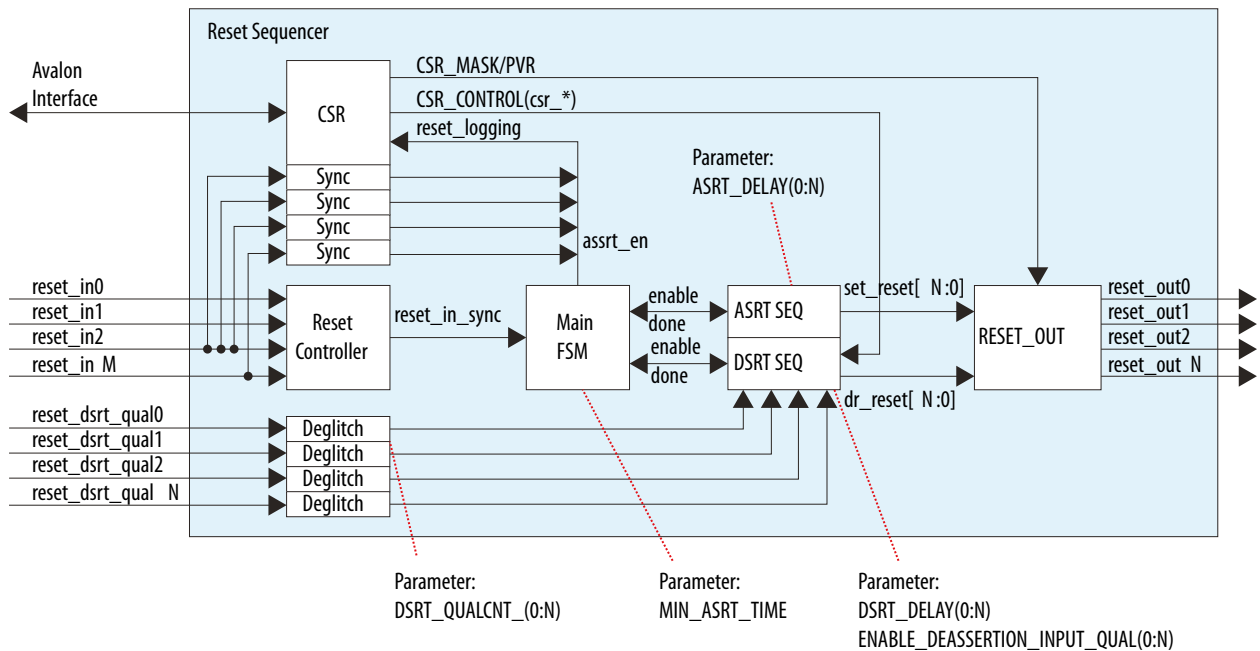
- **Active low reset**—When turned on, reset is asserted low.
- **Synchronous edges**—Specifies the level of synchronization and includes the following options:
  - **None**—The reset is asserted and deasserted asynchronously. Use this setting if you have internal synchronization circuitry.
  - **Both**—The reset is asserted and deasserted synchronously.
  - **Deassert**—The reset is deasserted synchronously, and asserted asynchronously.
- **Number of reset outputs**—The number of reset interfaces that are exported.

**Note:** Qsys supports multiple reset sink connections to a single reset source interface. However, there are situations in composed systems where an internally generated reset must be exported from the composed system in addition to being used to connect internal components. In this situation, you must declare one reset output interface as an export, and use another reset output to connect internal components.

## Reset Sequencer

The Reset Sequencer allows you to control the assertion and de-assertion sequence for Qsys system resets. The Parameter Editor displays the expected assertion and de-assertion sequences based on the current settings. You can connect multiple reset sources to the reset sequencer, and then connect the output of the reset sequencer to components in the system.

Figure 7-26: Elements and Flow of a Reset Sequencer



- **Reset Controller** — Reused reset controller block. It synchronizes the reset inputs into one and feed into the main FSM of the sequencer block.
- **Sync** — Synchronization block (double flip-flop).
- **Deglitch** — Deglitch block. This block waits for a signal to be at a level for X clocks before propagating the input to the output.
- **CSR** — This block contains the CSR Avalon interface and related CSR register and control block in the sequencer.
- **Main FSM** — Main sequencer. This block determines when assertion/deassertion and assertion hold timing occurs.
- **[A/D]SRT SEQ** — Generic sequencer block that sequences out assertion/deassertion of reset from 0:N. The block has multiple counters that saturate upon reaching count.
- **RESET\_OUT** — Controls the end output via:
  - Set/clear from the ASRT\_SEQ/DSRT\_SEQ.
  - Masking/forcing from CSR controls.
  - Remap of numbering (parameterization).

## Reset Sequencer Parameters

Table 7-33: Reset Sequencer Parameters

Parameter	Description
<b>Number of reset outputs</b>	Sets the number of output resets to be sequenced, which is the number of output reset signals defined in the component with a range of 2 to 10.
<b>Number of reset inputs</b>	Sets the number of input reset signals to be sequenced, which is the number of input reset signals defined in the component with a range of 1 to 10.
<b>Minimum reset assertion time</b>	Specifies the minimum assertion cycles between the assertion of the last sequenced reset, and the de-assertion of the first sequenced reset. The range is 0 to 1023.

Parameter	Description
<b>Enable Reset Sequencer CSR</b>	Enables CSR functionality of the Reset Sequencer through an Avalon interface.
<b>reset_out#</b>	Lists the reset output signals. Set the parameters in the other columns for each reset signal in the table.
<b>ASRT Seq#</b>	Determines the order of reset assertion. Enter the values 1, 2, 3, etc. to specify the required non-overlapping assertion order. This value determines the <code>ASRT_REMAP</code> value in the component HDL.
<b>ASRT Cycle#</b>	Number of cycles to wait before assertion of the reset. The value set here corresponds to the <code>ASRT_DELAY</code> value in the component HDL. The range is 0 to 1023.
<b>DSRT Seq#</b>	Determines the reset order of reset de-assertion. Enter the values 1, 2, 3, etc. to specify the required non-overlapping de-assertion order. This value determines the <code>DSRT_REMAP</code> value in the component HDL.
<b>DSRT Cycle#/Deglitch#</b>	Number of cycles to wait before de-asserting or de-glitching the reset. If the <code>USE_DRST_QUAL</code> parameter is set to 0, specifies the number of cycles to wait before de-asserting the reset. If <code>USE_DSRT_QUAL</code> is set to 1, specifies the number of cycles to deglitch the input <code>reset_dsrt_qual</code> signal. This value determines either the <code>DSRT_DELAY</code> , or the <code>DSRT_QUALCNT</code> value in the component HDL, depending on the <code>USE_DSRT_QUAL</code> parameter setting. The range is 0 to 1023.
<b>USE_DSRT_QUAL</b>	If you set <code>USE_DSRT_QUAL</code> to 1, the de-assertion sequence waits for an external input signal for sequence qualification instead of waiting for a fixed delay count. To use a fixed delay count for de-assertion, set this parameter to 0.

## Reset Sequencer Timing Diagrams

Figure 7-27: Basic Sequencing

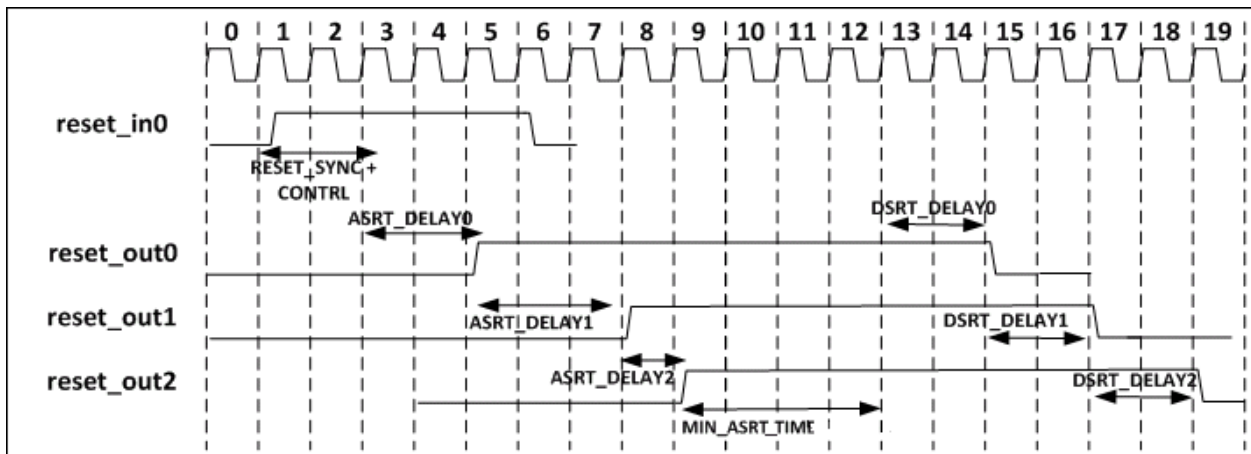
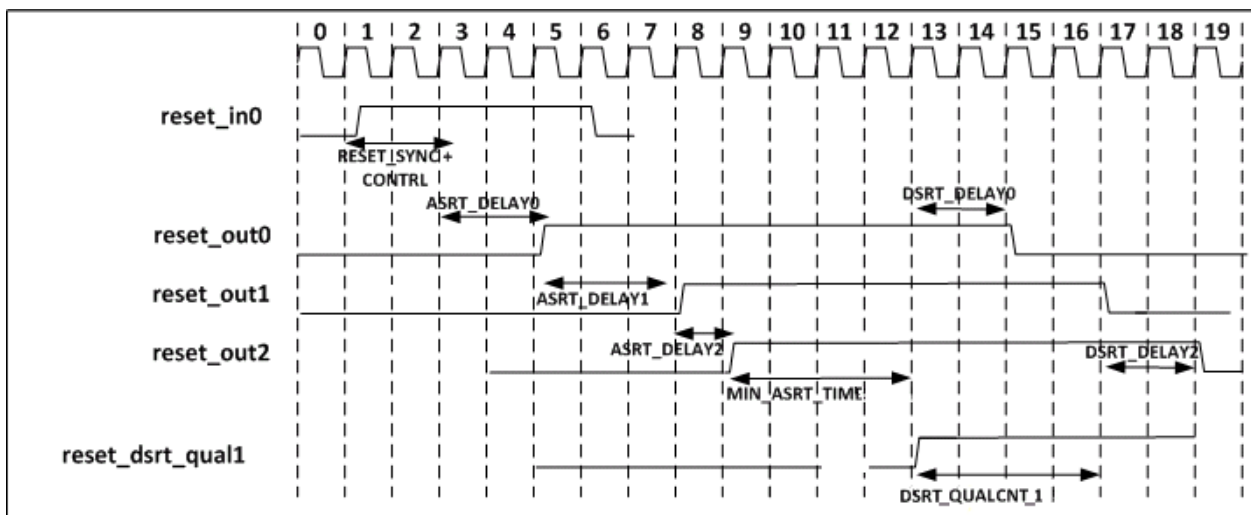


Figure 7-28: Sequencing with USE\_DSRT\_QUAL Set



## Reset Sequencer CSR Registers

The CSR registers on the reset sequencer provide the following functionality:

- **Supports reset logging**
  - Ability to identify which reset is asserted.
  - Ability to determine whether any reset is currently active.
- **Supports software triggered resets**
  - Ability to generate reset by writing to the register.
  - Ability to disable assertion or de-assertion sequence.
- **Supports software sequenced reset**
  - Ability for the software to fully control the assertion/de-assertion sequence by writing to registers and stepping through the sequence.
- **Support reset override**
  - Ability to assert a particular component reset through software.

### Reset Sequencer Status Register Offset 0x00

The **Status** register contains bits that indicate the sources of resets that cause a reset.

You can clear bits by writing 1 to the bit location. The Reset Sequencer ignores writes to bits with a value of 0. If the sequencer is reset (power-on-reset), all bits are cleared, except the power on reset bit.

**Table 7-34: Values for the Status Register at Offset 0x00**

Bit	Attribute	Default	Description
31	RO	0	<b>Reset Active</b> —Indicates that the sequencer is currently active in reset sequence (assertion or de-assertion).
30	RW1C	0	<b>Reset Asserted and waiting for SW to proceed:</b> —Set when there is an active reset assertion, and the next sequence is waiting for the software to proceed. Only valid when the <b>Enable SW sequenced reset entry</b> option is turned on.
29	RW1C	0	<b>Reset De-asserted and waiting for SW to proceed:</b> —Set when there is an active reset de-assertion, and the next sequence is waiting for the software to proceed. Only valid when the <b>Enable SW sequenced reset bring up</b> option is turned on.
28:26	RO	0	Reserved.
25:16	RW1C	0	<b>Reset de-assertion input qualification signal reset_dsrt_qual [9:0] status</b> —Indicates that the reset de-assertion's input signal qualification signal is set. This bit is set on the detection of assertion of the signal.
15:12	RO	0	Reserved.
11	RW1C	0	<b>reset_in9 was triggered</b> —Indicates that <code>reset_in9</code> triggered the reset. Cleared by software by writing 1 to this bit location.



Bit	Attribute	Default	Description
10	RW1C	0	<b>reset_in8 was triggered</b> —Indicates that <code>reset_in8</code> triggered the reset. Cleared by software by writing 1 to this bit location.
9	RW1C	0	<b>reset_in7 was triggered</b> —Indicates that <code>reset_in7</code> triggered the reset. Cleared by software by writing 1 to this bit location.
8	RW1C	0	<b>reset_in6 was triggered</b> —Indicates that <code>reset_in6</code> triggered the reset. Cleared by software by writing 1 to this bit location.
7	RW1C	0	<b>reset_in5 was triggered</b> —Indicates that <code>reset_in5</code> triggered the reset. Cleared by software by writing 1 to this bit location.
6	RW1C	0	<b>reset_in4 was triggered</b> —Indicates that <code>reset_in4</code> triggered the reset. Cleared by software by writing 1 to this bit location.
5	RW1C	0	<b>reset_in3 was triggered</b> —Indicates that <code>reset_in3</code> triggered the reset. Cleared by software by writing 1 to this bit location.
4	RW1C	0	<b>reset_in2 was triggered</b> —Indicates that <code>reset_in2</code> triggered the reset. Cleared by software by writing 1 to this bit location.
3	RW1C	0	<b>reset_in1 was triggered</b> —Indicates that <code>reset_in1</code> triggered the reset. Cleared by software by writing 1 to this bit location.
2	RW1C	0	<b>reset_in0 was triggered</b> —Indicates that <code>reset_in0</code> triggered. Cleared by software by writing 1 to this bit location.
1	RW1C	0	<b>Software triggered reset</b> —Indicates that the software triggered reset is set by the software, and triggering a reset.
0	RW1C	0	<b>Power-On-Reset was triggered</b> —Asserted whenever the reset to the sequencer is triggered. This bit is NOT reset when sequencer is reset. Cleared by software by writing 1 to this bit location.

#### Reset Sequencer Interrupt Enable Register Offset 0x04

The Interrupt Enable register contains the interrupt enable bit that you can use to enable any event triggering the IRQ of the reset sequencer.

**Table 7-35: Values for the Interrupt Enable Register at Offset 0x04**

Bit	Attribute	Default	Description
31	RO	0	Reserved.

Bit	Attribute	Default	Description
30	RW	0	<b>Interrupt on Reset Asserted and waiting for SW to proceed</b> enable. When set, the IRQ is set when the sequencer is waiting for the software to proceed in an assertion sequence.
29	RW	0	<b>Interrupt on Reset De-asserted and waiting for SW to proceed</b> enable. When set, the IRQ is set when the sequencer is waiting for the software to proceed in a de-assertion sequence.
28:26	RO	0	Reserved.
25:16	RW	0	<b>Interrupt on Reset de-assertion input qualification signal reset_dsrt_qual_ [9:0] status</b> — When set, the IRQ is set when the <code>reset_dsrt_qual[9:0]</code> status bit (per bit enable) is set.
15:12	RO	0	Reserved.
11	RW	0	<b>Interrupt on reset_in9 Enable</b> —When set, the IRQ is set when the <code>reset_in9</code> trigger status bit is set.
10	RW	0	<b>Interrupt on reset_in8 Enable</b> —When set, the IRQ is set when the <code>reset_in8</code> trigger status bit is set.
9	RW	0	<b>Interrupt on reset_in7 Enable</b> —When set, the IRQ is set when the <code>reset_in7</code> trigger status bit is set.
8	RW	0	<b>Interrupt on reset_in6 Enable</b> —When set, the IRQ is set when the <code>reset_in6</code> trigger status bit is set.
7	RW	0	<b>Interrupt on reset_in5 Enable</b> —When set, the IRQ is set when the <code>reset_in5</code> trigger status bit is set.
6	RW	0	<b>Interrupt on reset_in4 Enable</b> —When set, the IRQ is set when the <code>reset_in4</code> trigger status bit is set.
5	RW	0	<b>Interrupt on reset_in3 Enable</b> —When set, the IRQ is set when the <code>reset_in3</code> trigger status bit is set.
4	RW	0	<b>Interrupt on reset_in2 Enable</b> —When set, the IRQ is set when the <code>reset_in2</code> trigger status bit is set.
3	RW	0	<b>Interrupt on reset_in1 Enable</b> —When set, the IRQ is set when the <code>reset_in1</code> trigger status bit is set.
2	RW	0	<b>Interrupt on reset_in0 Enable</b> —When set, the IRQ is set when the <code>reset_in0</code> trigger status bit is set.

Bit	Attribute	Default	Description
1	RW	0	<b>Interrupt on Software triggered reset Enable</b> —When set, the IRQ is set when the software triggered reset status bit is set.
0	RW	0	<b>Interrupt on Power-On-Reset Enable</b> —When set, the IRQ is set when the power-on-reset status bit is set.

### Reset Sequencer Control Register Offset 0x08

The Control register contains registers that you can use to control the reset sequencer.

**Table 7-36: Values for the Control Register at Offset 0x08**

Bit	Attribute	Default	Description
31:3	RO	0	Reserved.
2	RW	0	<b>Enable SW sequenced reset entry</b> —Enable a software sequenced reset entry sequence. Timer delays and input qualification are ignored, and only the software can sequence the entry.
1	RW	0	<b>Enable SW sequenced reset bring up</b> —Enable a software sequenced reset bring up sequence. Timer delays and input qualification are ignored, and only the software can sequence the bring up.
0	WO	0	<b>Initiate Reset Sequence</b> —Reset Sequencer writes this bit to 1 a single time in order to trigger the hardware sequenced warm reset. Reset Sequencer verifies that <b>Reset Active</b> is 0 before setting this bit, and always reads the value 0. To monitor this sequence, verify that <b>Reset Active</b> is asserted, and then subsequently de-asserted.

### Reset Sequencer Software Sequenced Reset Entry Control Register Offset 0x0C

You can program the Reset Sequencer Software Sequenced Reset Entry Control register to control the reset entry sequence of the sequencer.

When the corresponding enable bit is set, the sequencer stops when the desired reset asserts, and then sets the **Reset Asserted and waiting for SW to proceed** bit. The Reset Sequencer proceeds only after the **Reset Asserted and waiting for SW to proceed** bit is cleared.

**Table 7-37: Values for the Reset Sequencer Software Sequenced Reset Entry Controls Register at Offset 0x0C**

Bit	Attribute	Default	Description
31:10	RO	0	Reserved.

Bit	Attribute	Default	Description
9:0	RW	3FF	<b>Per-reset SW sequenced reset entry enable</b> —This is a per-bit enable for SW sequenced reset entry. If <code>bitN</code> of this register is set, the sequencer sets the <code>bit30</code> of the status register when a <code>resetN</code> is asserted. It then waits for the <code>bit30</code> of the status register to clear before proceeding with the sequence. By default, all bits are enabled (fully SW sequenced).

#### Reset Sequencer Software Sequenced Reset Bring Up Control Register Offset 0x10

You can program the Software Sequenced Reset Bring Up Control register to control the reset bring up sequence of the sequencer.

When the corresponding enable bit is set, the sequencer stops when the desired reset asserts, and then sets the **Reset De-asserted and waiting for SW to proceed** bit. The Reset Sequencer proceeds only after the **Reset De-asserted and waiting for SW to proceed** bit is cleared..

**Table 7-38: Values for the Reset Sequencer Software Sequenced Bring Up Control Register at Offset 0x10**

Bit	Attribute	Default	Description
31:10	RO	0	Reserved.
9:0	RW	3FF	<b>Per-reset SW sequenced reset entry enable</b> —This is a per-bit enable for SW sequenced reset bring up. If <code>bitN</code> of this register is set, the sequencer sets <code>bit29</code> of the status register when a <code>resetN</code> is asserted. It then waits for the <code>bit29</code> of the status register to clear before proceeding with the sequence. By default, all bits are enabled (fully SW sequenced).

#### Reset Sequencer Software Direct Controlled Resets Offset 0x14

You can write a bit to 1 to assert the `reset_outN` signal, and to 0 to de-assert the `reset_outN` signal.

**Table 7-39: Values for the Software Direct Controlled Resets at Offset 0x14**

Bit	Attribute	Default	Description
31:26	RO	0	Reserved.
25:16	WO	0	<b>Reset Overwrite Trigger Enable</b> —This is a per-bit control trigger bit for the overwrite value to take effect.
15:10	RO	0	Reserved.
9:0	WO	0	<b>reset_outN Reset Overwrite Value</b> —This is a per-bit control of the <code>reset_out</code> bit. The Reset Sequencer can use this to forcefully drive the reset to a specific value. A value of 1 sets the <code>reset_out</code> . A value of 0 clears the <code>reset_out</code> . A write to this register only takes effect if the corresponding trigger bit in this register is set.



### Reset Sequencer Software Reset Masking Offset 0x18

You can write a bit to 1 to assert the `reset_outN` signal, and to 0 to de-assert the `reset_outN` signal.

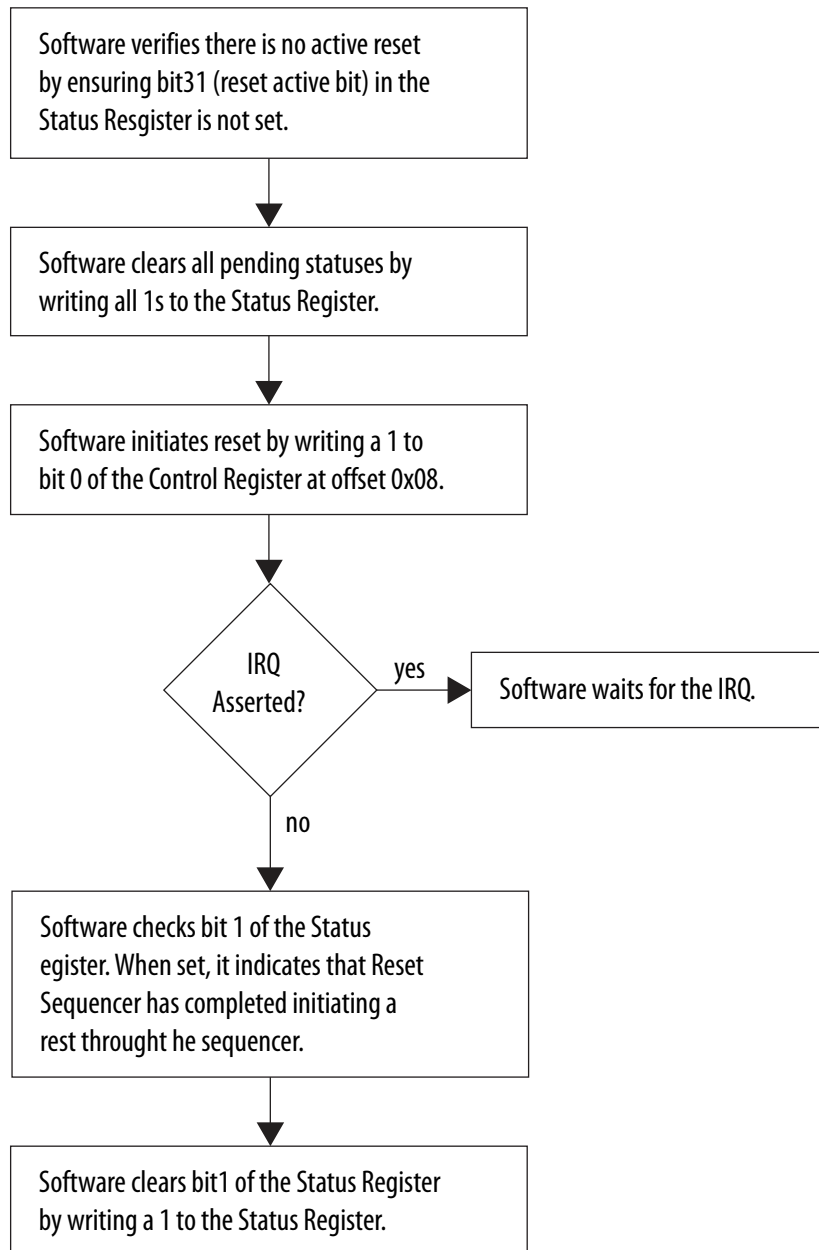
**Table 7-40: Values for the Reset Sequencer Software Reset Masking at Offset 0x18**

Bit	Attribute	Default	Description
31:10	RO	0	Reserved.
9:0	RW	0	<b>reset_outN "Reset Mask Enable"</b> —This is a per-bit control to mask the <code>reset_outN</code> bit. The Software Reset Masking masks the reset bit from being asserted during a reset assertion sequence. If the <code>reset_out</code> is already asserted, it does not de-assert the reset.

## Reset Sequencer Software Flows

### Reset Sequencer (Software-Triggered) Flow

Figure 7-29: Reset Sequencer (Software-Triggered) Flow



## Reset Entry Flow

The following flow sequence occurs for a Reset Entry Flow:

- A reset is triggered either by the software, or when input resets to the Reset Sequencer are asserted.
- The IRQ is asserted, if the IRQ is enabled.
- Software reads the Status register to determine what reset was triggered.


## Reset Bring-Up Flow

The following flow sequence occurs for a Reset Bring-Up Flow:

- When a reset source is de-asserted, or when the reset entry sequence has completed without any more pending resets asserted, the bring-up flow is initiated.
- The IRQ is asserted, if the IRQ is enabled.
- Software reads the Status register to determine what reset was triggered.

## Reset Entry (Software-Sequenced) Flow

Figure 7-30: Reset Entry (Software-Sequenced) Flow



## Reset Bring-Up (Software-Sequenced) Flow

The sequence and flow is similar to the **Reset Entry (SW Sequenced)** flow, though, this flow uses the **reset bring-up** registers/bits in place of the **reset entry** registers/bits.

## Related Information

[Reset Entry \(Software-Sequenced\) Flow](#) on page 7-62



## Conduits

You can use the conduit interface type for interfaces that do not fit any of the other interface types, and to group any arbitrary collection of signals. Like other interface types, you can export or connect conduit interfaces. The *PCI Express-to-Ethernet* example in *Creating a System with Qsys* is an example of using a conduit interface for export. You can declare an associated clock interface for conduit interfaces in the same way as memory-mapped interfaces with the `associatedClock`.

To connect two conduit interfaces inside Qsys, the following conditions must be met:

- The interfaces must match exactly with the same signal roles and widths.
- The interfaces must be the opposite directions.
- Clocked conduit connections must have matching `associatedClocks` on each of their endpoint interfaces.

**Note:** To connect a conduit output to more than one input conduit interface, you can create a custom component. The custom component could have one input that connects to two outputs, and you can use this component between other conduits that you want to connect. For information about the Avalon Conduit interface, refer to the *Avalon Interface Specifications*

### Related Information

#### [Avalon Interface Specifications](#)

[Creating a System with Qsys](#) on page 5-1

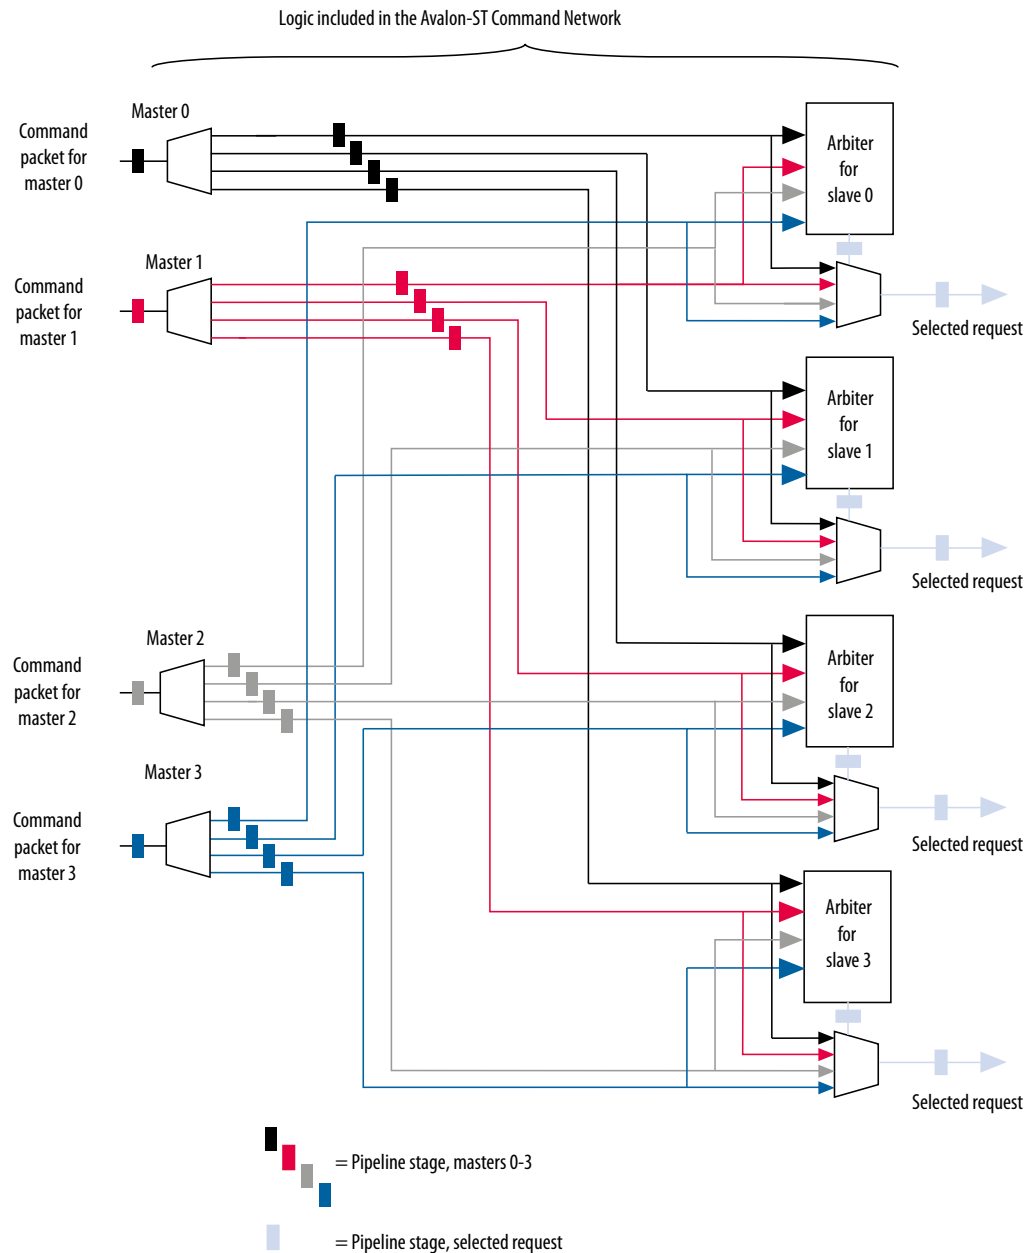
## Interconnect Pipelining

If you set the **Limit interconnect pipeline stages to** parameter to a value greater than 0 on the **Project Settings** tab, Qsys automatically inserts Avalon-ST pipeline stages when you generate your design. The pipeline stages increase the  $f_{MAX}$  of your design by reducing the combinational logic depth. The cost is additional latency and logic.

The insertion of pipeline stages depends upon the existence of certain interconnect components. For example, in a single-slave system, no multiplexer exists; therefore multiplexer pipelining does not occur. In an extreme case, of a single-master to single-slave system, no pipelining occurs, regardless of the value of the **Limit interconnect pipeline stages to** option.

**Figure 7-31: Pipeline Placement in Arbitration Logic**

The example below shows the possible placement of up to four potential pipeline stages, which could be, before the input to the demultiplexer, at the output of the multiplexer, between the arbiter and the multiplexer, and at the outputs of the demultiplexer.



**Note:** For more information about manually inserting and removing pipelines from your system, refer to *Creating a System With Qsys*. Refer to *Optimizing Qsys System Performance* for more information about pipelined Avalon-MM Interfaces.

#### Related Information

- [Creating a System With Qsys](#) on page 5-1

## Manually Controlling Pipelining in the Qsys Interconnect

The **Memory-Mapped Interconnect** tab allows you to manipulate pipeline connections in the Qsys interconnect. You access the **Memory-Mapped Interconnect** tab by clicking the **Show System With Qsys Interconnect** command on the System menu.

**Note:** To increase interconnect frequency, you should first try increasing the value of the **Limit interconnect pipeline stages to** option on the **Interconnect Requirements** tab. You should only consider manually pipelining the interconnect if changes to this option do not improve frequency, and you have tried all other options to achieve timing closure, including the use of a bridge. Manually pipelining the interconnect should only be applied to complete systems.

1. In the **Interconnect Requirements** tab, first try increasing the value of the **Limit interconnect pipeline stages to** option until it no longer gives significant improvements in frequency, or until it causes unacceptable effects on other parts of the system.
2. In the Quartus II software, compile your design and run timing analysis.
3. Using the timing report, identify the critical path through the interconnect and determine the approximate mid-point. The following is an example of a timing report:

```
2.800 0.000 cpu_instruction_master|out_shifter[63]|q
3.004 0.204 mm_domain_0|addr_router_001|Equal5~0|datac
3.246 0.242 mm_domain_0|addr_router_001|Equal5~0|combout
3.346 0.100 mm_domain_0|addr_router_001|Equal5~1|dataa
3.685 0.339 mm_domain_0|addr_router_001|Equal5~1|combout
4.153 0.468 mm_domain_0|addr_router_001|src_channel[5]~0|datad
4.373 0.220 mm_domain_0|addr_router_001|src_channel[5]~0|combout
```

4. In Qsys, click **System > Show System With Qsys Interconnect**.
5. In the **Memory-Mapped Interconnect** tab, select the interconnect module that contains the critical path. You can determine the name of the module from the hierarchical node names in the timing report.
6. Click **Show Pipelinable Locations**. Qsys display all possible pipeline locations in the interconnect. Right-click the possible pipeline location to insert or remove a pipeline stage.
7. Locate the possible pipeline location that is closest to the mid-point of the critical path. The names of the blocks in the memory-mapped interconnect tab correspond to the module instance names in the timing report.
8. Right-click the location where you want to insert a pipeline, and then click **Insert Pipeline**.
9. Regenerate the Qsys system, recompile the design, and then rerun timing analysis. If necessary, repeat the manual pipelining process again until timing requirements are met.

Manual pipelining has the following limitations:

- If you make changes to your original system's connectivity after manually pipelining an interconnect, your inserted pipelines may become invalid. Qsys displays warning messages when you generate your system if invalid pipeline stages are detected. You can remove invalid pipeline stages with the **Remove Stale Pipelines** option in the **Memory-Mapped Interconnect** tab. Altera recommends that you do not make changes to the system's connectivity after manual pipeline insertion.
- Review manually-inserted pipelines when upgrading to newer versions of Qsys. Manually-inserted pipelines in one version of Qsys may not be valid in a future version.

### Related Information

[Specify Qsys \\$system Interconnect Requirements](#)

[Qsys System Design Components](#) on page 10-1

## AMBA 3 AXI Protocol Specification Support (version 1.0)

Qsys allows memory-mapped connections between AXI3 components, AXI3 and AXI4 components, and AXI3 and Avalon interfaces with some unique or exceptional support.

Refer to the *AMBA 3 Protocol Specifications* on the ARM website for more information.

### Related Information

[AMBA 3 Protocol Specifications](#)

## Channels

Qsys 14.0 has the following support and restrictions for AXI3 channels.

### Read and Write Address Channels

All signals are allowed with some limitations.

The following limitations are present in Qsys 14.0:

- Supports 64-bit addressing.
- ID width limited to 18-bits.
- HPS-FPGA master interface has a 12-bit ID.

### Write Data, Write Response, and Read Data Channels

All signals are allowed with some limitations.

The following limitations are present in Qsys 14.0:

- Data widths limited to a maximum of 1024-bits
- Limited to a fixed byte width of 8-bits

### Low Power Channel

Low power extensions are not supported in Qsys, version 14.0.

## Cache Support

AWCACHE and ARCACHE are passed to an AXI slave unmodified.

### Bufferable

Qsys interconnect treats AXI transactions as non-bufferable. All responses must come from the terminal slave.

When connecting to Avalon-MM slaves, since they do not have write responses, the following exceptions apply:

- For Avalon-MM slaves, the write response are generated by the slave agent once the write transaction is accepted by the slave. The following limitation exists for an Avalon bridge:
- For an Avalon bridge, the response is generated before the write reaches the endpoint; users must be aware of this limitation and avoid multiple paths past the bridge to any endpoint slave, or only perform bufferable transactions to an Avalon bridge.



## Cacheable (Modifiable)

Qsys interconnect acknowledges the cacheable (modifiable) attribute of AXI transactions.

It does not change the address, burst length, or burst size of non-modifiable transactions, with the following exceptions:

- Qsys considers a wide transaction to a narrow slave as modifiable because the size requires reduction.
- Qsys may consider AXI read and write transactions as modifiable when the destination is an Avalon slave. The AXI transaction may be split into multiple Avalon transactions if the slave is unable to accept the transaction. This may occur because of burst lengths, narrow sizes, or burst types.

Qsys ignores all other bits, for example, read allocate or write allocate because the interconnect does not perform caching. By default, Qsys considers Avalon master transactions as non-bufferable and non-cacheable, with the allocate bits tied low. Qsys provides compile-time options to control the cache behavior of Avalon transactions on a per-master basis.

## Security Support

TrustZone refers to the security extension of the ARM architecture, which includes the concept of "secure" and "non-secure" transactions, and a protocol for processing between the designations.

The interconnect passes the `AWPROT` and `ARPROT` signals to the endpoint slave without modification. It does not use or modify the `PROT` bits.

Refer to *Creating a System with Qsys* for more information about secure systems and the TrustZone feature.

### Related Information

- [Creating a System with Qsys](#) on page 5-1

## Atomic Accesses

Exclusive accesses are supported for AXI slaves by passing the lock, transaction ID, and response signals from master to slave, with the limitation that slaves that do not reorder responses. Avalon slaves do not support exclusive accesses, and always return `OKAY` as a response. Locked accesses are also not supported.

## Response Signaling

Full response signaling is supported. Avalon slaves always return `OKAY` as a response.

## Ordering Model

Qsys interconnect provides responses in the same order as the commands are issued.

To prevent reordering, for slaves that accept reordering depths greater than 0, Qsys does not transfer the transaction ID from the master, but provides a constant transaction ID of 0. For slaves that do not reorder, Qsys allows the transaction ID to be transferred to the slave. To avoid cyclic dependencies, Qsys supports a single outstanding slave scheme for both reads and writes. Changing the targeted slave before all responses have returned stalls the master, regardless of transaction ID.

## AXI and Avalon Ordering

According to the *AMBA Protocol Specifications*, there is no ordering requirement between reads and writes. However, Avalon has an implicit ordering model that requires transactions from a master to the same slave to be in order. As a result, there is a potential read-after-write risk when Avalon masters

transact to AXI slaves. In response to this potential risk, Avalon interfaces provide a compile-time option to enforce strict order. When turned on, the Avalon interface waits for outstanding write responses before issuing reads.

## Data Buses

Narrow bus transfers are supported. AXI write strobes can have any pattern that is compatible with the address and size information. Altera recommends that transactions to Avalon slaves follow Avalon `byteenable` limitations for maximum compatibility.

**Note:** Byte 0 is always bits [7:0] in the interconnect, following AXI's and Avalon's byte (address) invariance scheme.

## Unaligned Address Commands

Unaligned address commands are commands with addresses that do not conform to the data width of a slave. Since Avalon-MM slaves accept only aligned addresses, Qsys modifies unaligned commands from AXI masters to the correct data width. Qsys must preserve commands issued by AXI masters when passing the commands to AXI slaves.

**Note:** Unaligned transfers are aligned if downsizing occurs. For example, when downsizing to a bus width narrower than that required by the transaction size, `AWSIZE` or `ARSIZE`, the transaction must be modified.

## Avalon and AXI Transaction Support

Qsys 14.0 supports transactions between Avalon and interfaces with some limitations.

### Transaction Cannot Cross 4KB Boundaries

When an Avalon master issues a transaction to an AXI slave, the transaction cannot cross 4KB boundaries. Non-bursting Avalon masters already follow this boundary restriction.

### Handling Read Side Effects

Read side effects can occur when more bytes than necessary are read from the slave, and the unwanted data that are read are later inaccessible on subsequent reads. For write commands, the correct `byteenable` paths are asserted based on the size of the transactions. For read commands, narrow-sized bursts are broken up into multiple non-bursting commands, and each command with the correct `byteenable` paths asserted.

**Note:** Qsys always assumes that the `byteenable` is asserted based on the size of the command, not the address of the command. The following scenarios are examples:

- For a 32-bit AXI master that issues a read command with an unaligned address starting at address `0x01`, and a `burstcount` of 2 to a 32-bit Avalon slave, the starting address is: `0x00`.
- For a 32-bit AXI master that issues a read command with an unaligned address starting at address `0x01`, with 4-bytes to an 8-bit AXI slave, the starting address is: `0x00`.

## AMBA 3 APB Protocol Specification Support (version 1.0)

AMBA APB provides a low-cost interface that is optimized for minimal power consumption and reduced interface complexity. You can use AMBA APB to interface to peripherals which are low-bandwidth and

do not require the high performance of a pipelined bus interface. Signal transitions are sampled at the rising edge of the clock to enable the integration of APB peripherals easily into any design flow.

Qsys allows connections between APB components, and AXI3, AXI4, and Avalon memory-mapped interfaces. The following sections describe unique or exceptional APB support in the Qsys software.

Refer to the *AMBA APB Protocol Specifications* for AXI4 on the ARM website for more information.

#### Related Information

[AMBA APB Protocol Specifications](#)

## Bridges

With APB, you cannot use bridge components that use multiple `PSELx` in Qsys. As a workaround, you can group `PSELx`, and then send the packet to the slave directly.

Altera recommends as an alternative that you instantiate the APB bridge and all the APB slaves in Qsys. You should then connect the slave side of the bridge to any high speed interface and connect the master side of the bridge to the APB slaves. Qsys creates the interconnect on either side of the APB bridge and creates only one `PSEL` signal.

Alternatively, you can connect a bridge to the APB bus outside of Qsys. Use an Avalon/AXI bridge to export the Avalon/AXI master to the top-level, and then connect this Avalon/AXI interface to the slave side of the APB bridge. Alternatively, instantiate the APB bridge in Qsys and export APB master to the top-level, and from there connect to APB bus outside of Qsys.

## Burst Adaptation

APB is a non-bursting interface. Therefore, for any AXI or Avalon master with bursting support, a burst adapter is inserted before the slave interface and the burst transaction is translated into a series of non-bursting transactions before reaching the APB slave.

## Width Adaptation

Qsys allows different data width connections with APB. When connecting a wider master to a narrower APB slave, the width adapter converts the wider transactions to a narrower transaction to fit the APB slave data width. APB does not support Write Strobe. Therefore, when you connect a narrower transaction to a wider APB slave, the slave cannot determine which byte lane to write. In this case, the slave data may be overwritten or corrupted.

## Error Response

Error responses are returned to the master. Qsys performs error mapping if the master is an AXI3 or AXI4 master, for example, `RRESP/BRESP= SLVERR`. For the case when the slave does not use `SLVERR` signal, an `OKAY` response is sent back to master by default.

## AMBA AXI4 Memory-Mapped Interface Support (version 2.0)

Qsys allows memory-mapped connections between AXI4 components, AXI4 and AXI3 components, and AXI4 and Avalon interfaces with some unique or exceptional support.

## Burst Support

Qsys supports `INCR` bursts up to 256 beats. Qsys converts long bursts to multiple bursts in a packet with each burst having a length less than or equal to `MAX_BURST` when going to AXI3 or Avalon slaves.

For narrow-sized transfers, bursts with Avalon slaves as destinations are shortened to multiple non-bursting transactions in order to transmit the correct address to the slaves, since Avalon slaves always perform full-sized `datawidth` transactions.

Bursts with AXI3 slaves as destinations are shortened to multiple bursts, with each burst length less than or equal to 16. Bursts with AXI4 slaves as destinations are not shortened.

## QoS

Qsys routes 4-bit QoS signals (Quality of Service Signaling) on the read and write address channels directly from the master to the slave.

Transactions from AXI3 and Avalon masters have a default value of `4'b0000`, which indicates that the transactions are not part of the QoS flow. QoS values are not used for slaves that do not support QoS.

For Qsys 14.0, there are no programmable QoS registers or compile-time QoS options for a master that overrides its real or default value.

## Regions

For Qsys 14.0, there is no support for the optional regions feature. AXI4 slaves with `AXREGION` signals are allowed. `AXREGION` signals are driven with the default value of `0x0`, and are limited to one entry in a master's address map.

## Write Response Dependency

Write response dependency as specified in the *AMBA Protocol Specifications* for AXI4 is not supported.

### Related Information

[AMBA Protocol Specifications](#)

## AWCACHE and ARCACHE

For AXI4, Qsys meets the requirement for modifiable and non-modifiable transactions. The modifiable bit refers to `ARCACHE[1]` and `AWCACHE[1]`.

## Width Adaptation and Data Packing in Qsys

Data packing applies only to systems where the data width of masters is less than the data width of slaves.

The following rules apply:

- Data packing is supported when masters and slaves are Avalon-MM.
- Data packing is not supported when any master or slave is an AXI3, AXI4, or APB component.

For example, for a read/write command with a 32-bit master connected to a 64-bit slave, and a transaction of 2 burstcounts, Qsys sends 2 separate read/write commands to access the 64-bit data width of the slave. Data packing is only supported if the system does not contain AXI3, AXI4, or APB masters or slaves.



## Ordering Model

Out of order support is not implemented in Qsys, version 14.0. Qsys processes AXI slaves as device non-bufferable memory types.

The following describes the required behavior for the device non-bufferable memory type:

- Write response must be obtained from the final destination.
- Read data must be obtained from the final destination.
- Transaction characteristics must not be modified.
- Reads must not be pre-fetched. Writes must not be merged.
- Non-modifiable read and write transactions.

(`AWCACHE[1] = 0` or `ARCACHE[1] = 0`) from the same ID to the same slave must remain ordered. The interconnect always provides responses in the same order as the commands issued. Slaves that support reordering provide a constant transaction ID to prevent reordering. AXI slaves that do not reorder are provided with transaction IDs, which allows exclusive accesses to be used for such slaves.

## Read and Write Allocate

Read and write allocate does not apply to Qsys interconnect, which does not have caching features, and always receives responses from an endpoint.

## Locked Transactions

Locked transactions are not supported for Qsys, version 14.0.

## Memory Types

For AXI4, Qsys processes transactions as though the endpoint is a device memory type. For device memory types, using non-bufferable transactions to force previous bufferable transactions to finish is irrelevant, because Qsys interconnect always identifies transactions as being non-bufferable.

## Mismatched Attributes

There are rules for how multiple masters issue cache values to a shared memory region. The interconnect meets requirements as long as cache signals are not modified.

## Signals

Qsys supports up to 64-bits for the `BUSER`, `WUSER` and `RUSER` sideband signals. AXI4 allows some signals to be omitted from interfaces by aligning them with the default values as defined in the *AMBA Protocol Specifications* on the ARM website.

### Related Information

[AMBA Protocol Specifications](#)

## AMBA AXI4 Streaming Interface Support (version 1.0)

### Connection Points

Qsys allows you to connect an AXI4 stream interface to another AXI4 stream interface.

The connection is point-to-point without adaptation and must be between an `axi4stream_master` and `axi4stream_slave`. Connected interfaces must have the same port roles and widths.

Non matching master to slave connections, and multiple masters to multiple slaves connections are not supported.

## AXI4 Streaming Connection Point Parameters

Table 7-41: AXI4 Streaming Connection Point Parameters

Name	Type	Description
<code>associatedClock</code>	string	Name of associated clock interface.
<code>associatedReset</code>	string	Name of associated reset interface

## AXI4 Streaming Connection Point Signals

Table 7-42: AXI4 Stream Connection Point Signals

Port Role	Width	Master Direction	Slave Direction	Required
<code>tvalid</code>	1	Output	Input	Yes
<code>tready</code>	1	Input	Output	No
<code>tdata</code> <sup>(6)</sup>	8:4096	Output	Input	No
<code>tstrb</code>	1:512	Output	Input	No
<code>tkeep</code>	1:512	Output	Input	No
<code>tid</code> <sup>(7)</sup>	1:8	Output	Input	No
<code>tdest</code> <sup>(8)</sup>	1:4	Output	Input	No
<code>tuser</code> <sup>(9)</sup>	1:4096	Output	Input	No
<code>tlast</code>	1	Output	Input	No

## Adaptation

AXI4 stream adaptation support is not available. AXI4 stream master and slave interface signals and widths must match.

## AMBA AXI4-Lite Protocol Specification Support (version 2.0)

AXI4-Lite is a sub-set of AMBA AXI4. It is suitable for simpler control register-style interfaces that do not require the full functionality of AXI4.

<sup>(6)</sup> integer in mutiple of bytes

<sup>(7)</sup> maximum 8-bits

<sup>(8)</sup> maximum 4-bits

<sup>(9)</sup> number of bits in multiple of the number of bytes of `tdata`

Qsys 14.0 supports the following AXI4-Lite features:

- Transactions with a burst length of 1.
- Data accesses use the full width of a data bus (32-bit or 64-bit) for data accesses, and no narrow-size transactions.
- Non-modifiable and non-bufferable accesses.
- No exclusive accesses.

## AXI4-Lite Signals

Qsys supports all AXI4-Lite interface signals. All signals are required.

**Table 7-43: AXI4-Lite Signals**

Global	Write Address Channel	Write Data Channel	Write Response Channel	Read Address Channel	Read Data Channel
ACLK	AWVALID	WVALID	BVALID	ARVALID	RVALID
ARESETn	AWREADY	WREADY	BREADY	ARREADY	RREADY
-	AWADDR	WDATA	BRESP	ARADDR	RDATA
-	AWPROT	WSTRB	-	ARPROT	RRESP

## AXI4-Lite Bus Width

AXI4-Lite masters or slaves must have either 32-bit or 64-bit bus widths. Qsys interconnect inserts a width adapter if a master and slave pair have different widths.

## AXI4-Lite Outstanding Transactions

AXI-Lite supports outstanding transactions. The options to control outstanding transactions is set in the parameter editor for the selected component.

## AXI4-Lite IDs

AXI4-Lite does not support IDs. Qsys performs ID reflection inside the slave agent.

## Connections Between AXI3/4 and AXI4-Lite

### AXI4-Lite Slave Requirements

For an AXI4-Lite slave side, the master can be any master interface type, such as an Avalon (with bursting), AXI3, or AXI4. Qsys allows the following connections and inserts adapters, if needed.

- **Burst adapter**—Avalon and AXI3 and AXI4 bursting masters require a burst adapter to shorten the burst length to 1 before sending a transaction to an AXI4-Lite slave.
- Qsys interconnect uses a width adapter for mismatched data widths.
- Qsys interconnect performs ID reflection inside the slave agent.
- An AXI4-Lite slave must have an address width of at least 12-bits.
- AXI4-Lite does not have the `AXSIZE` parameter. Narrow master to a wide AXI4-Lite slave is not supported. For masters that support narrow-sized bursts, for example, AXI3 and AXI4, a burst to an AXI4-Lite slave must have a burst size equal to or greater than the slave's burst size.

### AXI4-Lite Data Packing

Qsys interconnect does not support AXI4-Lite data packing.

### AXI4-Lite Response Merging

When Qsys interconnect merges `SLVERR` and `DECERR`, the error responses are not sticky. The response is based on priority and the master always sees a `DECERR`. When `SLVERR` and `DECERR` are merged, it is based on their priorities, not stickiness. `DECERR` receives priority in this case, even if `SLVERR` returns first.

## Port Roles (Interface Signal Types)

Each interfaces defines a number of signal roles and their behavior. Many signal roles are optional, allowing IP component designers the flexibility to select only the signal roles necessary to implement the required functionality.

### AXI Master Interface Signal Types

Table 7-44: AXI Master Interface Signal Types

Name	Direction	Width
araddr	output	1 - 64
arburst	output	2
arcache	output	4
arid	output	1 - 18
arlen	output	4
arlock	output	2
arprot	output	3
arready	input	1
arsize	output	3
aruser	output	1 - 64
arvalid	output	1
awaddr	output	1 - 64
awburst	output	2

Name	Direction	Width
awcache	output	4
awid	output	1 - 18
awlen	output	4
awlock	output	2
awprot	output	3
awready	input	1
awsize	output	3
awuser	output	1 - 64
awvalid	output	1
bid	input	1 - 18
bready	output	1
bresp	input	2
bvalid	input	1
rdata	input	8, 16, 32, 64, 128, 256, 512, 1024
rid	input	1 - 18
rlast	input	1
rready	output	1
rresp	input	2
rvalid	input	1
wdata	output	8, 16, 32, 64, 128, 256, 512, 1024
wid	output	1 - 18
wlast	output	1
wready	input	1
wstrb	output	1, 2, 4, 8, 16, 32, 64, 128
wvalid	output	1

## AXI Slave Interface Signal Types

Table 7-45: AXI Slave Interface Signal Types

Name	Direction	Width
araddr	input	1 - 64
arburst	input	2

Name	Direction	Width
arcache	input	4
arid	input	1 - 18
arlen	input	4
arlock	input	2
arprot	input	3
arready	output	1
arsize	input	3
aruser	input	1 - 64
arvalid	input	1
awaddr	input	1 - 64
awburst	input	2
awcache	input	4
awid	input	1 - 18
awlen	input	4
awlock	input	2
awprot	input	3
awready	output	1
awsize	input	3
awuser	input	1 - 64
awvalid	input	1
bid	output	1 - 18
breedy	input	1
bresp	output	2
bvalid	output	1
rdata	output	8, 16, 32, 64, 128, 256, 512, 1024
rid	output	1 - 18
rlast	output	1
rready	input	1
rresp	output	2
rvalid	output	1
wdata	input	8, 16, 32, 64, 128, 256, 512, 1024

Name	Direction	Width
wid	input	1 - 18
wlast	input	1
wready	output	1
wstrb	input	1, 2, 4, 8, 16, 32, 64, 128
wvalid	input	1

## AXI4 Master Interface Signal Types

Table 7-46: AXI4 Master Interface Signal Types

Name	Direction	Width
araddr	output	1 - 64
arburst	output	2
arcache	output	4
arid	output	1 - 18
arlen	output	8
arlock	output	1
arprot	output	3
arready	input	1
arregion	output	1 - 4
arsize	output	3
aruser	output	1 - 64
arvalid	output	1
awaddr	output	1 - 64
awburst	output	2
awcache	output	4
awid	output	1 - 18
awlen	output	8
awlock	output	1
awprot	output	3
awqos	output	1 - 4
awready	input	1
awregion	output	1 - 4

Name	Direction	Width
awsize	output	3
awuser	output	1 - 64
awvalid	output	1
bid	input	1 - 18
bready	output	1
bresp	input	2
buser	input	1 - 64
bvalid	input	1
rdata	input	8, 16, 32, 64, 128, 256, 512, 1024
rid	input	1 - 18
rlast	input	1
rready	output	1
rresp	input	2
ruser	input	1 - 64
rvalid	input	1
wdata	output	8, 16, 32, 64, 128, 256, 512, 1024
wid	output	1 - 18
wlast	output	1
wready	input	1
wstrb	output	1, 2, 4, 8, 16, 32, 64, 128
wuser	output	1 - 64
wvalid	output	1

## AXI4 Slave Interface Signal Types

Table 7-47: AXI4 Slave Interface Signal Types

Name	Direction	Width
araddr	input	1 - 64
arburst	input	2
arcache	input	4
arid	input	1 - 18
arlen	input	8



Name	Direction	Width
arlock	input	1
arprot	input	3
argos	input	1 - 4
arready	output	1
arregion	input	1 - 4
arsize	input	3
aruser	input	1 - 64
arvalid	input	1
awaddr	input	1 - 64
awburst	input	2
awcache	input	4
awid	input	1 - 18
awlen	input	8
awlock	input	1
awprot	input	3
awgos	input	1 - 4
awready	output	1
awregion	inout	1 - 4
awsize	input	3
awuser	input	1 - 64
awvalid	input	1
bid	output	1 - 18
bready	input	1
bresp	output	2
bvalid	output	1
rdata	output	8, 16, 32, 64, 128, 256, 512, 1024
rid	output	1 - 18
rlast	output	1
rready	input	1
rresp	output	2
ruser	output	1 - 64

Name	Direction	Width
rvalid	output	1
wdata	input	8, 16, 32, 64, 128, 256, 512, 1024
wlast	input	1
wready	output	1
wstrb	input	1, 2, 4, 8, 16, 32, 64, 128
wuser	input	1 - 64
wvalid	input	1

## AXI4 Stream Master and Slave Interface Signal Types

Table 7-48: AXI4 Stream Master and Slave Interface Signal Types

Name	Width	Master Direction	Slave Direction	Required
tvalid	1	Output	Input	Yes
tready	1	Input	Output	No
tdata	8:4096	Output	Input	No
tstrb	1:512	Output	Input	No
tkeep	1:512	Output	Input	No
tid	1:8	Output	Input	No
tdest	1:4	Output	Input	No
tuser	1	Output	Input	No
tlast	1:4096	Output	Input	No

## APB Interface Signal Types

Table 7-49: APB Interface Signal Types

Name	Width	Direction APB Master	Direction APB Slave	Required
paddr	[1:32]	output	input	yes
psel	[1:16]	output	input	yes
penable	1	output	input	yes
pwrite	1	output	input	yes
pwdata	[1:32]	output	input	yes
prdata	[1:32]	input	output	yes

Name	Width	Direction APB Master	Direction APB Slave	Required
pslverr	1	input	output	no
pready	1	input	output	yes
paddr31	1	output	input	no

## Avalon Memory-Mapped Interface Signal Types

The following table lists the signal roles that constitute the Avalon-MM interface. The signal roles allow you to create masters that use bursts for reads and writes. You can increase the throughput of your system by initiating reads with multiple pipelined slave peripherals. In responding to reads, when a slave peripheral has valid data it asserts `readdatavalid`. The interconnect enables the connection between the master and slave pair.

This specification does not require all signals to exist in an Avalon-MM interface. In fact, there is no one signal that is always required. The minimum requirements are `readdata` for a read-only interface or `writedata` and `write` for a write-only interface.

**Table 7-50: Avalon-MM Signals**

All Avalon signals are active high. Avalon signals that can also be asserted low list `_n` versions of the signal in the **Signal role** column.

Signal Role	Width	Direction	Description
<b>Fundamental Signals</b>			
<code>address</code>	1 - 64	Master → Slave	<p>Masters: By default, the <code>address</code> signal represents a byte address. The value of the address must be aligned to the data width. To write to specific bytes within a data word, the master must use the <code>byteenable</code> signal. Refer to the <code>addressUnits</code> interface property for word addressing.</p> <p>Slaves: By default, the interconnect translates the byte address into a word address in the slave's address space. Each slave access is for a word of data from the perspective of the slave. For example, <code>address= 0</code> selects the first word of the slave. Address 1 selects the second word of the slave. Refer to the <code>addressUnits</code> interface property for byte addressing.</p>

Signal Role	Width	Direction	Description
byteenable byteenable_n	2, 4, 8, 16, 32, 64, 128	Master → Slave	<p>Enables specific byte lane(s) during transfers on ports of width greater than 8 bits. Each bit in <code>byteenable</code> corresponds to a byte in <code>writedata</code> and <code>readdata</code>. The master bit &lt;n&gt; of <code>byteenable</code> indicates whether byte &lt;n&gt; is being written to. During writes, <code>byteenables</code> specify which bytes are being written to. Other bytes should be ignored by the slave. During reads, <code>byteenables</code> indicate which bytes the master is reading. Slaves that simply return <code>readdata</code> with no side effects are free to ignore <code>byteenables</code> during reads. If an interface does not have a <code>byteenable</code> signal, the transfer proceeds as if all <code>byteenables</code> are asserted.</p> <p>When more than one bit of the <code>byteenable</code> signal is asserted, all asserted lanes are adjacent. The number of adjacent lines must be a power of 2. The specified bytes must be aligned on an address boundary for the size of the data. For example, the following values are legal for a 32-bit slave:</p> <ul style="list-style-type: none"> <li>• 1111 writes full 32 bits</li> <li>• 0011 writes lower 2 bytes</li> <li>• 1100 writes upper 2 bytes</li> <li>• 0001 writes byte 0 only</li> <li>• 0010 writes byte 1 only</li> <li>• 0100 writes byte 2 only</li> <li>• 1000 writes byte 3 only</li> </ul> <p>Altera strongly recommends that you use the <code>byteenable</code> signal in components that will be used in systems with different word sizes. Using byte enables avoids unintended side effects in systems that include width adapters.</p>
debugaccess	1	Master → Slave	When asserted, allows internal memories that are normally write-protected to be written. For example, on-chip ROM memories can only be written when <code>debugaccess</code> is asserted.
read read_n	1	Master → Slave	Asserted to indicate a read transfer. If present, <code>readdata</code> is required.
readdata	8,16, 32, 64, 128, 256, 512, 1024	Slave → Master	The <code>readdata</code> driven from the slave to the master in response to a read transfer.

Signal Role	Width	Direction	Description
response (optional)	2	Slave → Master	<p>The <code>response</code> signal indicates the status of a transaction, as follows:</p> <ul style="list-style-type: none"> <li>00: <code>OKAY</code>—Successful response for a transaction.</li> <li>01: <code>RESERVED</code>—Encoding is reserved.</li> <li>10: <code>SLAVEERROR</code>—Error from an endpoint slave. Indicates an unsuccessful transaction.</li> <li>11: <code>DECODEERROR</code>—Indicates attempted access to an undefined location.</li> </ul> <p>Every <code>readdata</code> in a burst receives a read response.</p> <p>To support read response, the interface must be read capable.</p> <p>When a read response returns an error, the slave determines the <code>readdata</code> value, and the master ignores the <code>readdata</code> value.</p> <p>All read transactions must return a response if the component supports read response.</p>
write write_n	1	Master → Slave	Asserted to indicate a <code>write</code> transfer. If present, <code>writedata</code> is required.
writedata	8, 16, 32, 64, 128, 256, 512, 1024	Master → Slave	Data for write transfers. The width must be the same as the width of <code>readdata</code> if both are present.
<b>Wait-State Signals</b>			

Signal Role	Width	Direction	Description
lock	1	Master → Slave	<p>lock ensures that once a master wins arbitration, it maintains access to the slave for multiple transactions. It is asserted coincident with the first <code>read</code> or <code>write</code> of a locked sequence of transactions. It is deasserted on the final transaction of a locked sequence of transactions. lock assertion does not guarantee that arbitration will be won. After the lock-asserting master has been granted, it retains grant until it is deasserted.</p> <p>A master equipped with lock cannot be a burst master. Arbitration priority values for lock-equipped masters are ignored.</p> <p>lock is particularly useful for read-modify-write (RMW) operations. The typical read-modify-write operation includes the following steps:</p> <ol style="list-style-type: none"> <li>1. Master A reads 32-bit data that has multiple bit fields.</li> <li>2. Master A changes one bit field and writes the 32-bit data back.</li> </ol> <p>lock prevents master B from performing a write between Master A's read and write. lock also ensures that master A does not overwrite master B's changes.</p>

Signal Role	Width	Direction	Description
waitrequest waitrequest_n	1	Slave → Master	<p>Asserted by the slave when it is unable to respond to a read or write request. Forces the master to wait until the interconnect is ready to proceed with the transfer. At the start of all transfers, a master initiates the transfer and waits until <code>waitrequest</code> is deasserted. A master must make no assumption about the assertion state of <code>waitrequest</code> when the master is idle: <code>waitrequest</code> may be high or low, depending on system properties.</p> <p>When <code>waitrequest</code> is asserted, master control signals to the slave to remain constant with the exception of <code>beginbursttransfer</code>. For a timing diagram illustrating the <code>beginburst-transfer</code> signal, refer to the figure in <i>Read Bursts</i>.</p> <p>An Avalon-MM slave may assert <code>waitrequest</code> during idle cycles. An Avalon-MM master may initiate a transaction when <code>waitrequest</code> is asserted and wait for that signal to be deasserted. To avoid system lockup, a slave device should assert <code>waitrequest</code> when in reset.</p>

**Pipeline Signals**

readdatavalid readdatavalid_n	1	Slave → Master	<p>Used for variable-latency, pipelined read transfers. When asserted, indicates that the <code>readdata</code> signal contains valid data. A slave with <code>readdatavalid</code> must assert this signal for one cycle for each read access received. There must be at least one cycle of latency between acceptance of the <code>read</code> and assertion of <code>readdatavalid</code>. For a timing diagram illustrating the <code>readdatavalid</code> signal, refer to <i>Pipelined Read Transfer with Variable Latency</i>.</p> <p>A slave may assert <code>readdatavalid</code> to transfer data to the master independently of whether or not the slave is stalling a new command with <code>waitrequest</code>.</p> <p>Required if the master supports pipelined reads. Bursting masters with read functionality must include the <code>readdatavalid</code> signal.</p>
----------------------------------	---	-------------------	---

**Burst Signals**

Signal Role	Width	Direction	Description
burstcount	1 – 11	Master → Slave	<p>Used by bursting masters to indicate the number of transfers in each burst. The value of the maximum <code>burstcount</code> parameter must be a power of 2. A burstcount port of width <code>&lt;n&gt;</code> can encode a max burst of size <math>2^{(&lt;n&gt;-1)}</math>. For example, a 4-bit <code>burstcount</code> signal can support a maximum burst count of 8. The minimum <code>burstcount</code> is 1. The <code>constantBurst</code> property controls the timing of the <code>burstcount</code> signal. Bursting masters with read functionality must include the <code>readdatavalid</code> signal.</p> <p>For bursting masters and slaves, the following restriction applies to the width of the address:</p> $\langle address\_w \rangle \geq \langle burstcount\_w \rangle + \text{floor}(\log_2(\langle symbols\_per\_word\_of\_interface \rangle))$
beginburst-transfer	1	Interconnect → Slave	<p>Asserted for the first cycle of a burst to indicate when a burst transfer is starting. This signal is deasserted after one cycle regardless of the value of <code>waitrequest</code>. The interconnect fabric automatically generates this signal for slaves when requested. For a timing diagram illustrating <code>beginbursttransfer</code>, refer to the figure in <i>Read Bursts</i>.</p> <p><code>beginbursttransfer</code> is optional. A slave can always internally calculate the start of the next write burst transaction by counting data transfers.</p> <p>Altera recommends that you <b>do not</b> use this signal. This signal exists to support legacy memory controllers.</p>

## Avalon Streaming Interface Signals

Each signal in an Avalon-ST source or sink interface corresponds to one Avalon-ST signal role. An Avalon-ST interface may contain only one instance of each signal role. All Avalon-ST signal roles apply to both sources and sinks and have the same meaning for both.

**Table 7-51: Avalon-ST Interface Signals**

In the following table, all signal roles are active high.

Signal Role	Width	Direction	Description
<b>Fundamental Signals</b>			



Signal Role	Width	Direction	Description
channel	1 – 128	Source → Sink	The <code>channel</code> number for data being transferred on the current cycle.  If an interface supports the channel signal, it must also define the <code>maxChannel</code> parameter.
data	1 – 4,096	Source → Sink	The <code>data</code> signal from the source to the sink, typically carries the bulk of the information being transferred.  The contents and format of the <code>data</code> signal is further defined by parameters.
error	1 – 256	Source → Sink	A bit mask used to mark errors affecting the data being transferred in the current cycle. A single bit in <code>error</code> is used for each of the errors recognized by the component, as defined by the <code>errorDescriptor</code> property.
ready	1	Sink → Source	Asserted high to indicate that the sink can accept data. <code>ready</code> is asserted by the sink on cycle $\langle n \rangle$ to mark cycle $\langle n + readyLatency \rangle$ as a ready cycle. The source may only assert <code>valid</code> and transfer data during <code>ready</code> cycles.  Sources without a <code>ready</code> input cannot be backpressured. Sinks without a <code>ready</code> output never need to backpressure.
valid	1	Source → Sink	Asserted by the source to qualify all other source to sink signals. The sink samples data other source-to-sink signals on ready cycles where <code>valid</code> is asserted. All other cycles are ignored.  Sources without a <code>valid</code> output implicitly provide valid data on every cycle that they are not being backpressured. Sinks without a <code>valid</code> input expect valid data on every cycle that they are not backpressuring.

#### Packet Transfer Signals

empty	1 – 8	Source → Sink	Indicates the number of symbols that are empty during cycles that contain the end of a packet. The empty signal is not used on interfaces where there is one symbol per beat. If <code>endofpacket</code> is not asserted, this signal is not interpreted.
endofpacket	1	Source → Sink	Asserted by the source to mark the end of a packet.
startofpacket	1	Source → Sink	Asserted by the source to mark the beginning of a packet.

## Avalon Clock Source Signal Roles

An Avalon Clock source interface drives a clock signal out of a component.

Table 7-52: Clock Source Signal Roles

Signal Role	Width	Direction	Required	Description
clk	1	Output	Yes	An output clock signal.

## Avalon Clock Sink Signal Roles

A clock sink provides a timing reference for other interfaces and internal logic.

Table 7-53: Clock Input Signal Roles

Signal Role	Width	Direction	Required	Description
clk	1	Input	Yes	A clock signal. Provides synchronization for internal logic and for other interfaces.

## Avalon Conduit Signals

Table 7-54: Conduit Signal Roles

Signal Role	Width	Direction	Description
<any>	<n>	In, out, or bidirectional	A conduit interface consists of one or more input, output, or bidirectional signals of arbitrary width. Conduits can have any user-specified role. You can connect compatible Conduit interfaces inside a Qsys system provided the roles and widths match and the directions are opposite.

## Avalon Tristate Conduit Signals

The following table lists the signal defined for the Avalon Tristate Conduit interface. All Avalon-TC signals apply to both masters and slaves and have the same meaning for both

**Table 7-55: Tristate Conduit Interface Signal Roles**

Signal Role	Width	Direction	Required	Description
request	1	Master → Slave	Yes	<p>The meaning of <code>request</code> depends on the state of the <code>grant</code> signal, as the following rules dictate.</p> <p>When <code>request</code> is asserted and <code>grant</code> is deasserted, <code>request</code> is requesting access for the current cycle.</p> <p>When <code>request</code> is asserted and <code>grant</code> is asserted, <code>request</code> is requesting access for the next cycle. Consequently, <code>request</code> should be deasserted on the final cycle of an access.</p> <p>The <code>request</code> is deasserted in the last cycle of a bus access. It can be reasserted immediately following the final cycle of a transfer. This protocol makes both rearbitration and continuous bus access possible if no other masters are requesting access.</p> <p>Once asserted, <code>request</code> must remain asserted until granted. Consequently, the shortest bus access is 2 cycles. Refer to <i>Tristate Conduit Arbitration Timing</i> for an example of arbitration timing.</p>
grant	1	Slave → Master	Yes	<p>When asserted, indicates that a tristate conduit master has been granted access to perform transactions. <code>grant</code> is asserted in response to the <code>request</code> signal. It remains asserted until 1 cycle following the deassertion of <code>request</code>.</p> <p>The design of the Avalon-TC Interface does not allow a default Avalon-TC master to be granted access when no masters are requesting.</p>
<name>_in	1 – 1024	Slave → Master	No	The input signal of a logical tristate signal.
<name>_out	1 – 1024	Master → Slave	No	The output signal of a logical tristate signal.
<name>_outen	1	Master → Slave	No	The output enable for a logical tristate signal.

## Avalon Tri-State Slave Interface Signal Types

Table 7-56: Tri-state Slave Interface Signal Types

Name	Width	Direction	Required	Description
address	1 - 32	input	No	Address lines to the slave port. Specifies a byte offset into the slave's address space.
read read_n	1	input	No	Read-request signal. Not required if the slave port never outputs data.  If present, data must also be used.
write write_n	1	input	No	Write-request signal. Not required if the slave port never receives data from a master.  If present, data must also be present, and <code>writebyteenable</code> cannot be present.
chipselct chipselct_n	1	input	No	When present, the slave port ignores all Avalon-MM signals unless <code>chipselct</code> is asserted. <code>chipselct</code> is always present in combination with read or write
outputenable outputenable_n	1	input	Yes	Output-enable signal. When deasserted, a tri-state slave port must not drive its data lines otherwise data contention may occur.
data	8,16, 32, 64, 128, 256, 512, 1024	bidir	No	Bidirectional data. During write transfers, the FPGA drives the data lines. During read transfers the slave device drives the data lines, and the FPGA captures the data signals and provides them to the master.

Name	Width	Direction	Required	Description
byteenable byteenable_n	2, 4, 8, 16, 32, 64, 128	input	No	<p>Enables specific byte lane(s) during transfers.</p> <p>Each bit in byteenable corresponds to a byte lane in data. During writes, byteenables specify which bytes the master is writing to the slave. During reads, byteenables indicates which bytes the master is reading. Slaves that simply return data with no side effects are free to ignore byteenables during reads.</p> <p>When more than one byte lane is asserted, all asserted lanes are guaranteed to be adjacent. The number of adjacent lines must be a power of 2, and the specified bytes must be aligned on an address boundary for the size of the data. The are legal values for a 32-bit slave:</p> <pre> 1111 writes full 32 bits 0011 writes lower 2 bytes 1100 writes upper 2 bytes 0001 writes byte 0 only 0010 writes byte 1 only 0100 writes byte 2 only 1000 writes byte 3 only </pre>
writebyteenable writebyteenable_n	2,4,8,16, 32, 64,128	input	No	Equivalent to the logical AND of the byteenable and write signals. When used, the write signal is not used.
begintransfer1	1	input	No	Asserted for the first cycle of each transfer.

**Note:** All Avalon signals are active high. Avalon signals that can also be asserted low list both versions in the **Signal Role** column.

## Avalon Interrupt Sender Signal Roles

Table 7-57: Interrupt Sender Signal Roles

Signal Role	Width	Direction	Required	Description
irq irq_n	1	Output	Yes	Interrupt Request. A slave asserts <code>irq</code> when it needs service.

## Avalon Interrupt Receiver Signal Roles

Table 7-58: Interrupt Receiver Signal Roles

Signal Role	Width	Direction	Required	Description
irq	1–32	Input	Yes	<code>irq</code> is an $\langle n \rangle$ -bit vector, where each bit corresponds directly to one IRQ sender. The interrupt sender connected to interrupt receiver 0 is the highest priority with sequential receivers being successively lower priority.

## Document Revision History

The table below indicates edits made to the *Qsys Interconnect* content since its creation.

Table 7-59: Document Revision History

Date	Version	Changes
December 2014	14.1.0	<ul style="list-style-type: none"> <li>Fixed Priority Arbitration.</li> <li>Read error responses, Avalon Memory-Mapped Interface Signal, <i>response</i>.</li> <li>Burst Adapter Implementation Options: Generic converter (slower, lower area), Per-burst-type converter (faster, higher area).</li> </ul>
August 2014	14.0a10.0	<ul style="list-style-type: none"> <li>Updated Qsys Packet Format for Memory-Mapped Master and Slave Interfaces table, <i>Protection</i>.</li> <li>Streaming Interface renamed to Avalon Streaming Interfaces.</li> <li>Added <i>Response Merging</i> under <i>Memory-Mapped Interfaces</i>.</li> </ul>

Date	Version	Changes
June 2014	14.0.0	<ul style="list-style-type: none"> <li>• AXI4-Lite support.</li> <li>• AXI4-Stream support.</li> <li>• Avalon-ST adapter parameters.</li> <li>• IRQ Bridge.</li> <li>• Handling Read Side Effects note added.</li> </ul>
November 2013	13.1.0	<ul style="list-style-type: none"> <li>• HSSI clock support.</li> <li>• Reset Sequencer.</li> <li>• Interconnect pipelining.</li> </ul>
May 2013	13.0.0	<ul style="list-style-type: none"> <li>• AMBA APB support.</li> <li>• Auto-inserted Avalon-ST adapters feature.</li> <li>• Moved Address Span Extender to the <i>Qsys System Design Components</i> chapter.</li> </ul>
November 2012	12.1.0	<ul style="list-style-type: none"> <li>• AMBA AXI4 support.</li> </ul>
June 2012	12.0.0	<ul style="list-style-type: none"> <li>• AMBA AXI3 support.</li> <li>• Avalon-ST adapters.</li> <li>• Address Span Extender.</li> </ul>
November 2011	11.0.1	Template update.
May 2011	11.0.0	Removed beta status.
December 2010	10.1.0	Initial release.

**Related Information**

[Quartus II Handbook Archive](#)

2014.06.30

QI15V1



Subscribe



Send Feedback

You can optimize system interconnect performance for Altera® designs that you create with the Qsys system integration tool.

The foundation of any system is the interconnect logic that connects hardware blocks or components. Creating interconnect logic is prone to errors, is time consuming to write, and is difficult to modify when design requirements change. The Qsys system integration tool addresses these issues and provides an automatically generated and optimized interconnect designed to satisfy your system requirements.

Qsys supports Avalon, AMBA AXI3 (version 1.0), AMBA AXI4 (version 2.0), AMBA AXI4-Lite (version 2.0), AMBA AXI4-Stream (version 1.0), and AMBA APB3 (version 1.0) interface specifications.

**Note:** Recommended Altera practices may improve clock frequency, throughput, logic utilization, or power consumption of your Qsys design. When you design a Qsys system, use your knowledge of your design intent and goals to further optimize system performance beyond the automated optimization available in Qsys.

## Related Information

- [Avalon Interface Specifications](#)
- [AMBA Protocol Specifications](#)
- [Creating a System with Qsys](#) on page 5-1
- [Creating Qsys Components](#) on page 6-1
- [Qsys Interconnect](#) on page 7-1

## Designing with Avalon and AXI Interfaces

Qsys Avalon and AXI interconnect for memory-mapped interfaces is flexible, partial crossbar logic that connects master and slave interfaces.

Avalon Streaming (Avalon-ST) links connect point-to-point, unidirectional interfaces and are typically used in data stream applications. Each a pair of components is connected without any requirement to arbitrate between the data source and sink.

Because Qsys supports multiplexed memory-mapped and streaming connections, you can implement systems that use multiplexed logic for control and streaming for data in a single design.

## Related Information

- [Creating Qsys Components](#) on page 6-1

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



## Designing Streaming Components

When you design streaming component interfaces, you must consider integration and communication for each component in the system. One common consideration is buffering data internally to accommodate latency between components.

For example, if the component's Avalon-ST output or source of streaming data is back-pressured because the ready signal is de-asserted, then the component must back-pressure its input or sink interface to avoid overflow.

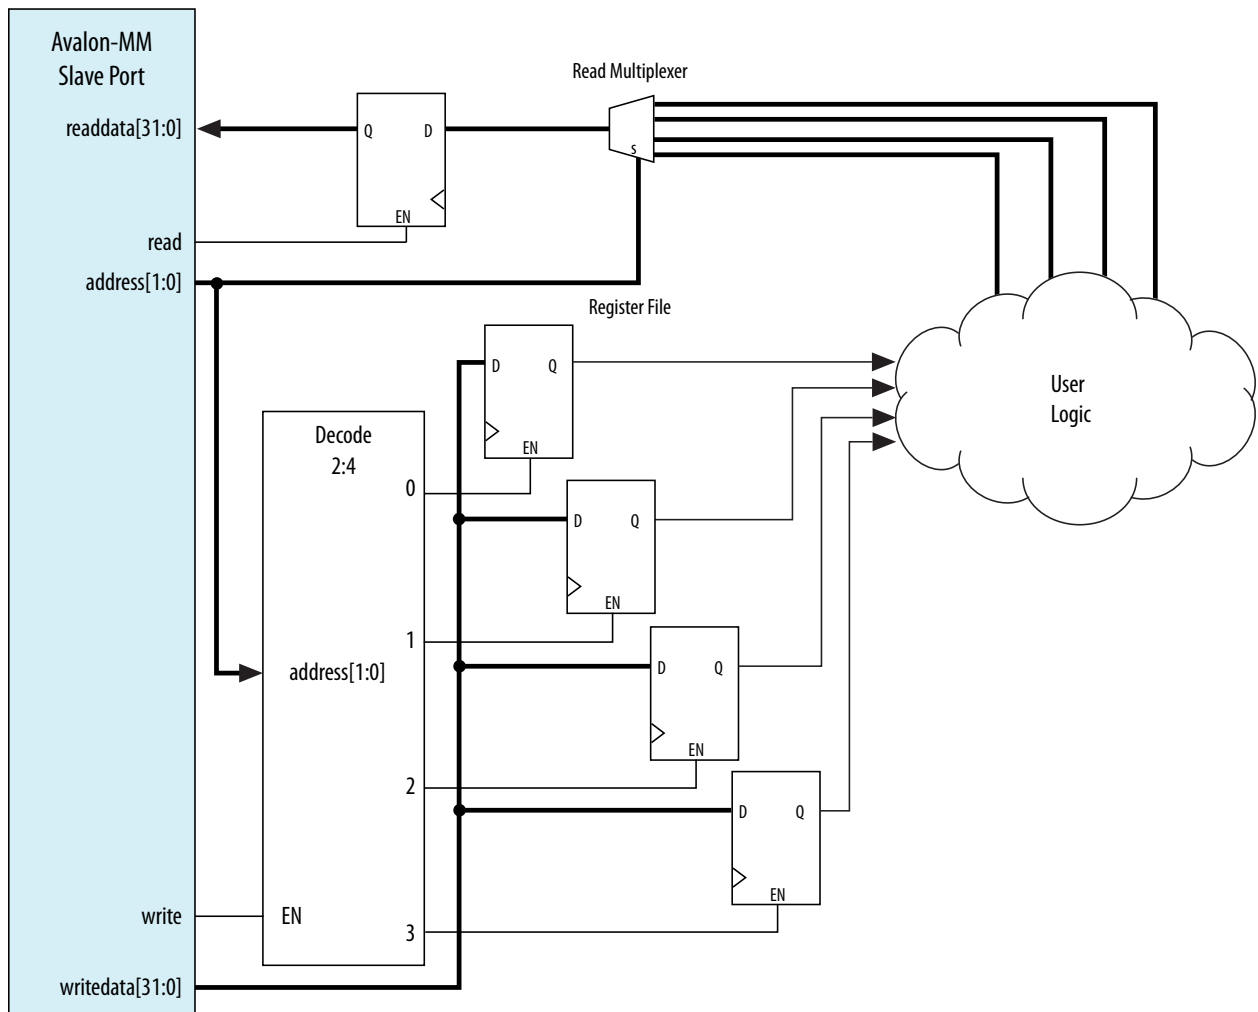
You can use a FIFO to back-pressure internally on the output side of the component so that the input can accept more data even if the output is back-pressured. Then, you can use the FIFO almost full flag to back-pressure the sink interface or input data when the FIFO has only enough space to satisfy the internal latency. You can drive the data valid signal of the output or source interface with the FIFO not empty flag when that data is available.

## Designing Memory-Mapped Components

When designing with memory-mapped components, you can implement any component that contains multiple registers mapped to memory locations, for example, a set of four output registers to support software read back from logic. Components that implement read and write memory-mapped transactions require three main building blocks: an address decoder, a register file, and a read multiplexer.

The decoder enables the appropriate 32-bit or 64-bit register for writes. For reads, the address bits drive the multiplexer selection bits. The read signal registers the data from the multiplexer, adding a pipeline stage so that the component can achieve a higher clock frequency.

Figure 8-1: Control and Status Registers (CSR) in a Slave Component



This slave component has write wait states and one read wait state. Alternatively, if you want high throughput, you may set both the read and write wait states to zero, and then specify a read latency of one, because the component also supports pipelined reads.

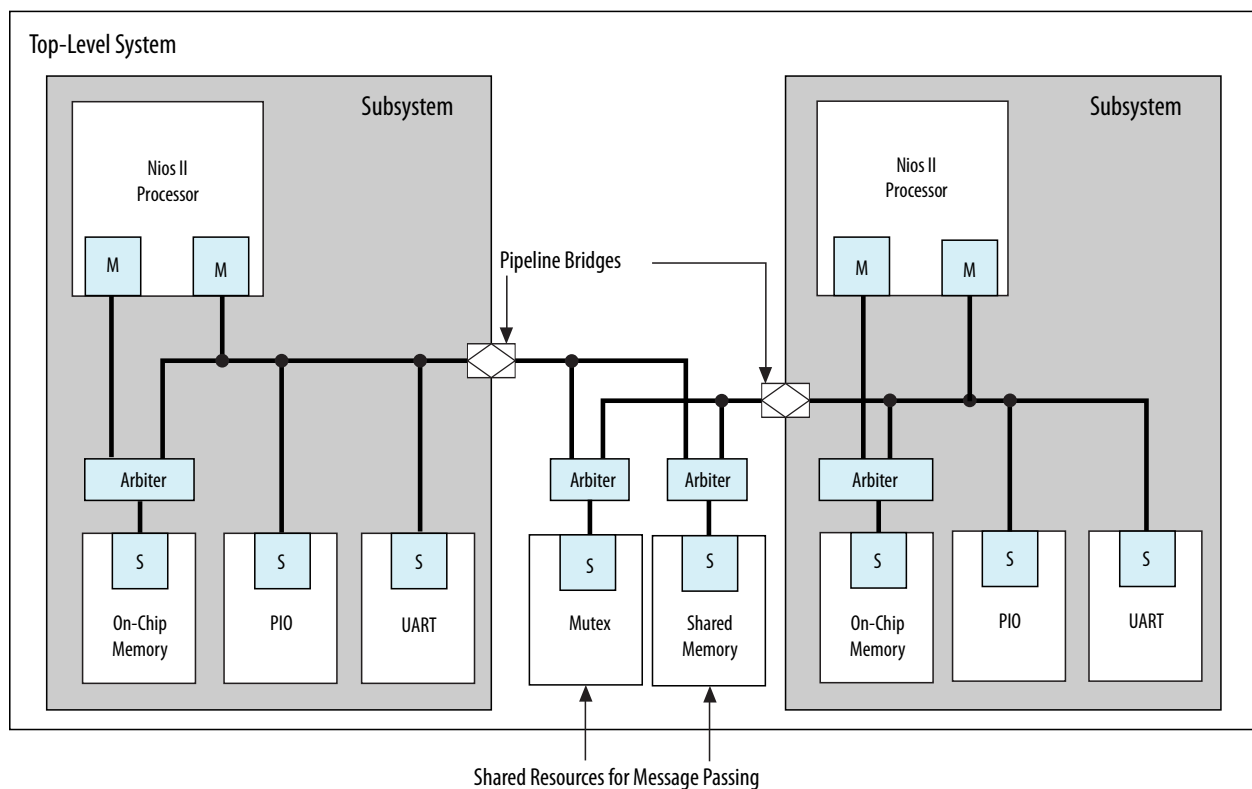
## Using Hierarchy in Systems

You can use hierarchy to sub-divide a system into smaller subsystems that you can then connect in a top-level Qsys system. Additionally, if a design contains one or more identical functional units, the functional unit can be defined as a subsystem and instantiated multiple times within a top-level system.

Hierarchy can simplify verification control of slaves connected to each master in a memory-mapped system. Before you implement subsystems in your design, you should plan the system hierarchical blocks at the top-level, using the following guidelines:

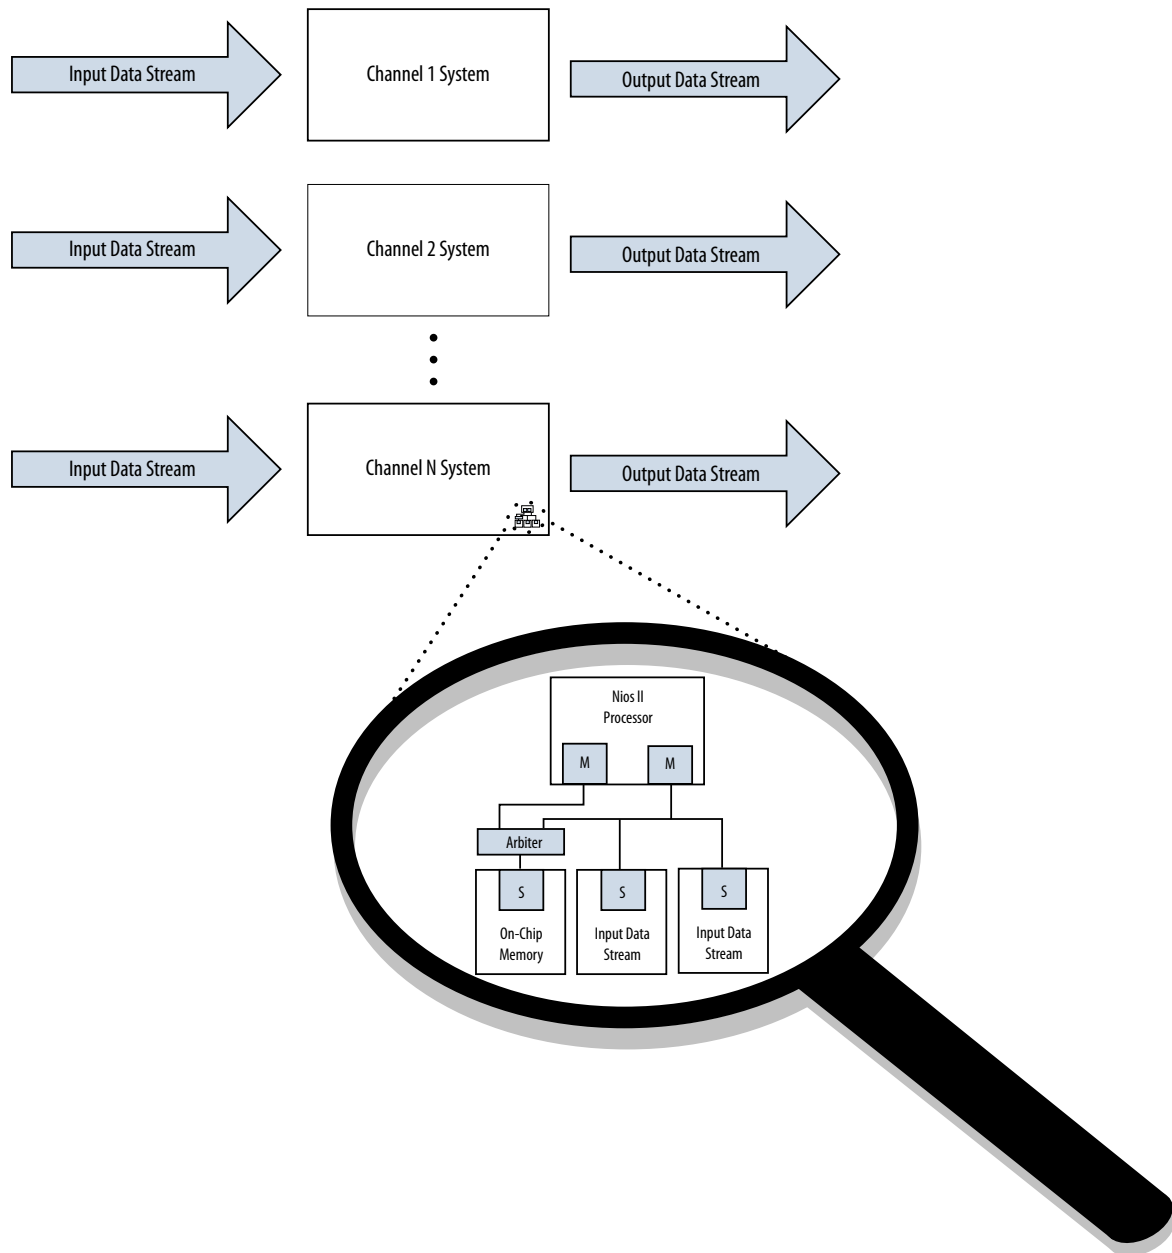
- **Plan shared resources**—Determine the best location for shared resources in the system hierarchy. For example, if two subsystems share resources, add the components that use those resources to a higher-level system for easy access.
- **Plan shared address space between subsystems**—Planning the address space ensures you can set appropriate sizes for bridges between subsystems.
- **Plan how much latency you may need to add to your system**—When you add a pipeline bridge between subsystems, you may add latency to the overall system. You can reduce the added latency by parameterizing the pipeline bridge with zero cycles of latency.

Figure 8-2: Passing Messages Between Subsystems



In this example, two Nios II processor subsystems share resources for message passing. Bridges in each subsystem export the Nios II data master to the top-level system that includes the mutex (mutual exclusion component) and shared memory component (which could be another on-chip RAM, or a controller for an off-chip RAM device).

Figure 8-3: Multi Channel System



You can also design systems that process multiple data channels by instantiating the same subsystem for each channel. This approach is easier to maintain than a larger, non-hierarchical system. Additionally, such systems are easier to scale because you can calculate the required resources as a multiple of the subsystem requirements.

## Using Concurrency in Memory-Mapped Systems

Qsys interconnect uses parallel hardware in FPGAs, which allows you to design concurrency into your system and process transactions simultaneously.

### Implementing Concurrency With Multiple Masters

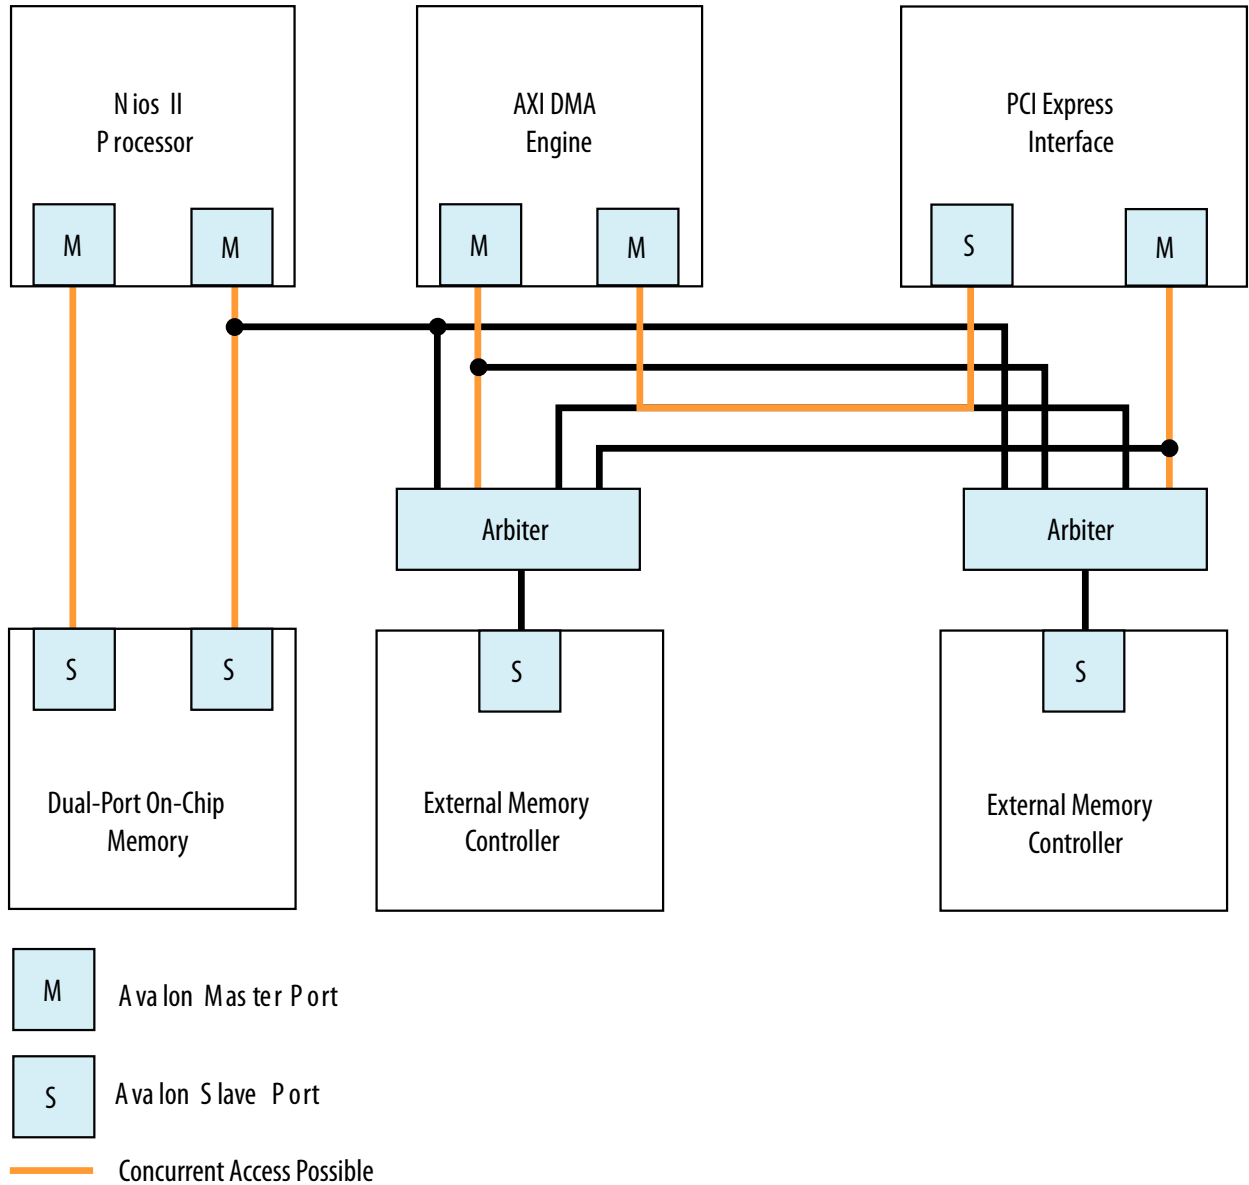
Implementing concurrency requires multiple masters in a Qsys system. Systems that include a processor contain at least two master interfaces because the processors include separate instruction and data masters. You can categorize master components as follows:

- General purpose processors, such as Nios II processors
- DMA (direct memory access) engines
- Communication interfaces, such as PCI Express

Because Qsys generates an interconnect with slave-side arbitration, every master interface in a system can issue transfers concurrently, as long as they are not posting transfers to the same slave. Concurrency is limited by the number of master interfaces sharing any particular slave interface. If a design requires higher data throughput, you can increase the number of master and slave interfaces to increase the number of transfers that occur simultaneously. The example below shows a system with three master interfaces.

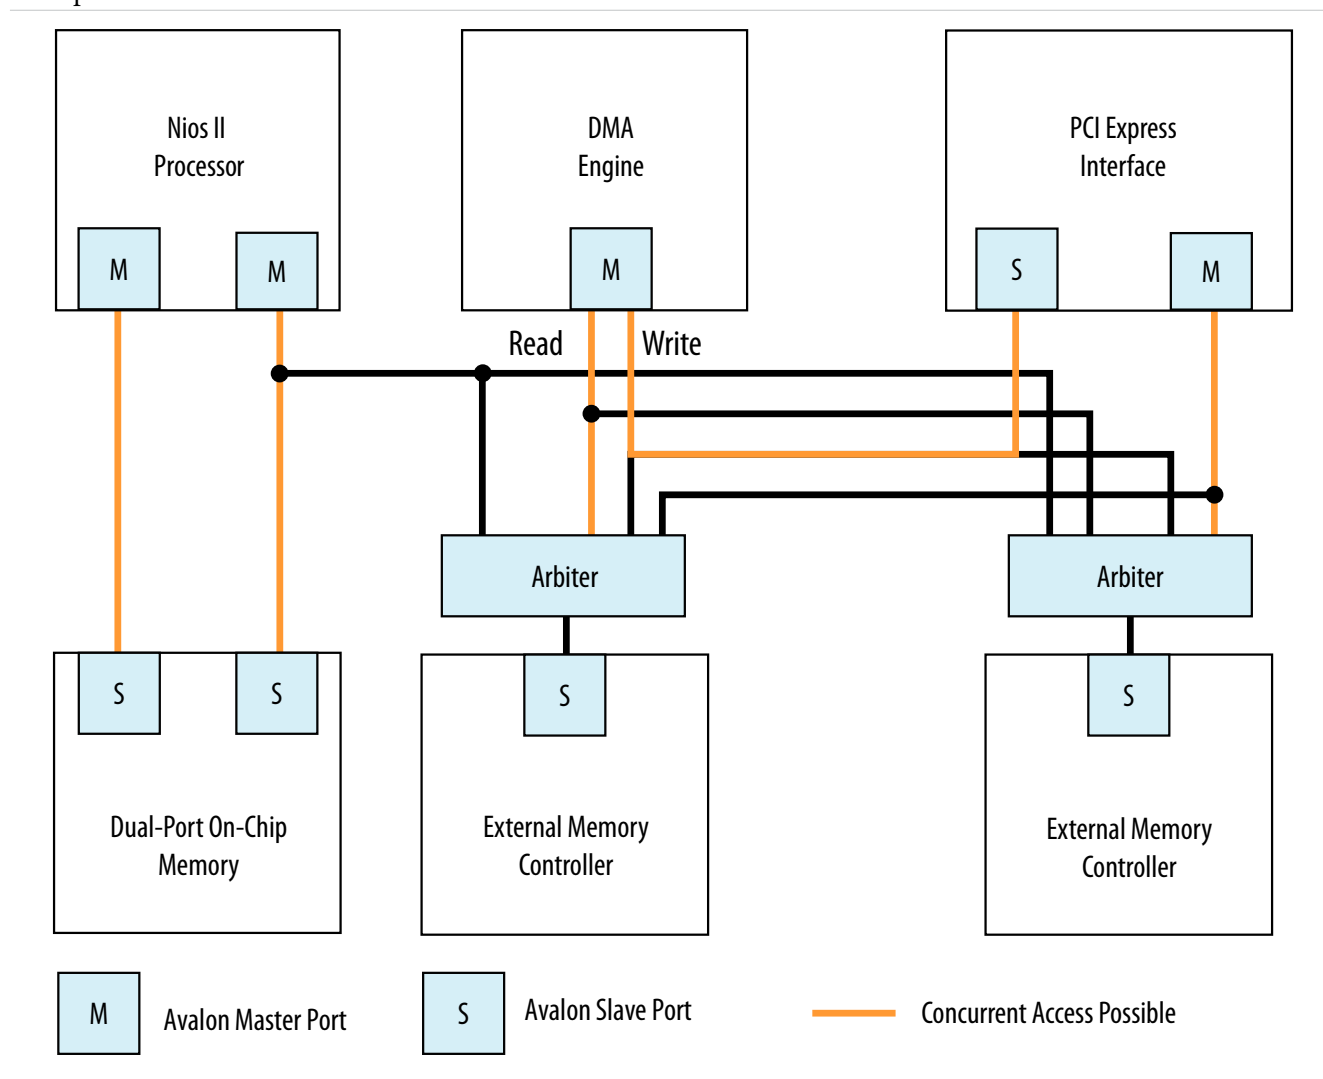
**Figure 8-4: Avalon Multiple Master Parallel Access**

In this Avalon example, the DMA engine operates with Avalon-MM read and write masters. However, an AXI DMA interface typically has only one master, because in the AXI standard, the write and read channels on the master are independent and can process transactions simultaneously. The yellow lines represent active simultaneously connections.



**Figure 8-5: AXI Multiple Master Parallel Access**

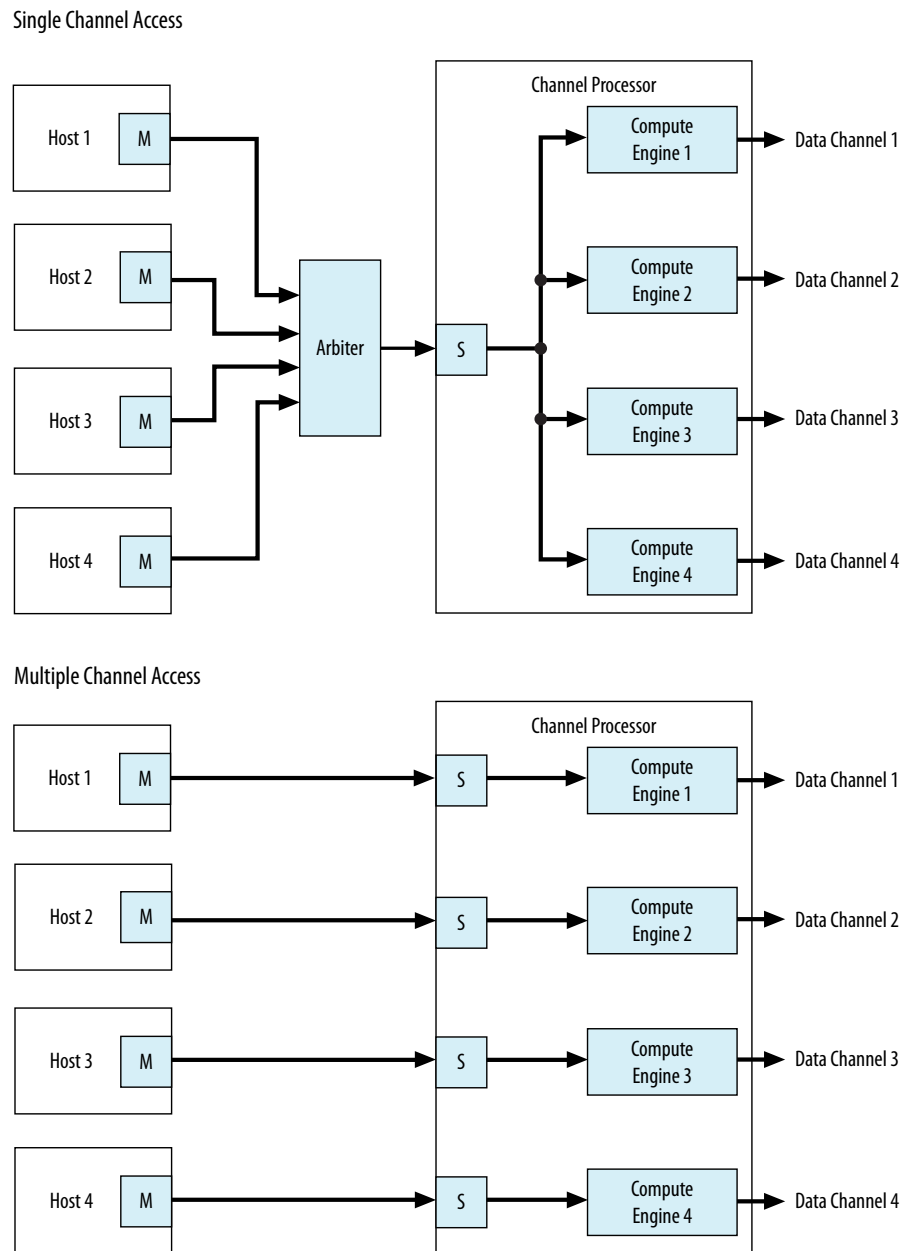
In this example, the DMA engine operates with a single master, because in AXI, the write and read channels on the master are independent and can process transactions simultaneously. There is concurrency between the read and write channels, with the yellow lines representing concurrent data paths.



## Implementing Concurrency With Multiple Slaves

You can create multiple slave interfaces for a particular function to increase concurrency in your design.

Figure 8-6: Single Interface Versus Multiple Interfaces



In this example, there are two channel processing systems. In the first, four hosts must arbitrate for the single slave interface of the channel processor. In the second, each host drives a dedicated slave interface, allowing all master interfaces to simultaneously access the slave interfaces of the component. Arbitration is not necessary when there is a single host and slave interface.

## Implementing Concurrency with DMA Engines

In some systems, you can use DMA engines to increase throughput. You can use a DMA engine to transfer blocks of data between interfaces, which then frees the CPU from doing this task. A DMA engine

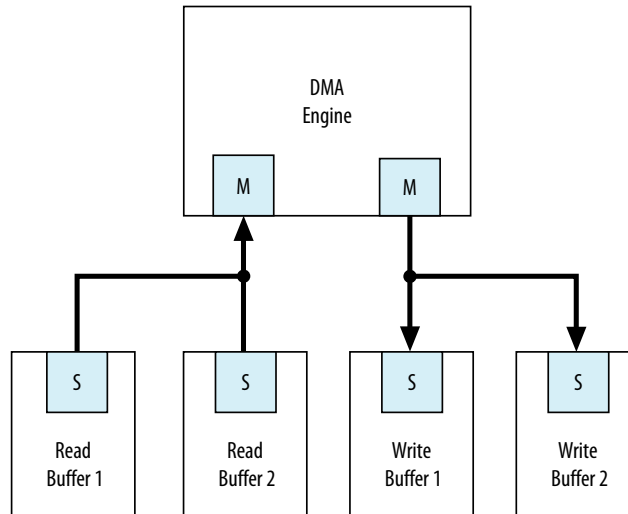


transfers data between a programmed start and end address without intervention, and the data throughput is dictated by the components connected to the DMA. Factors that affect data throughput include data width and clock frequency.

**Figure 8-7: Single or Dual DMA Channels**

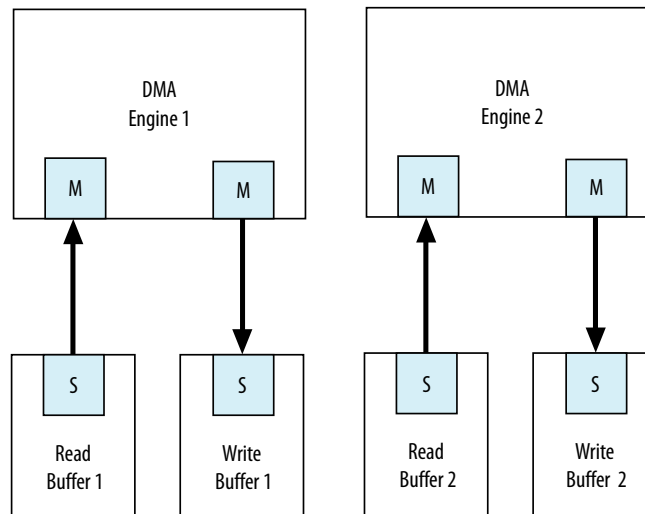
#### Single DMA Channel

Maximum of One Read & One Write Per Clock Cycle



#### Dual DMA Channels

Maximum of two Reads & Two Writes Per Clock Cycle



In this example, the system can sustain more concurrent read and write operations by including more DMA engines. Accesses to the read and write buffers in the top system are split between two DMA engines, as shown in the Dual DMA Channels at the bottom of the figure.

The DMA engine operates with Avalon-MM write and read masters. An AXI DMA typically has only one master, because in AXI, the write and read channels on the master are independent and can process transactions simultaneously.

## Inserting Pipeline Stages to Increase System Frequency

Qsys provides the **Limit interconnect pipeline stages to** option on the **Project Settings** tab to automatically add pipeline stages to the Qsys interconnect when you generate a system.

You can specify between 0 to 4 pipeline stages, where 0 means that the interconnect has a combinational data path. You can specify a unique interconnect pipeline stage value for each subsystem.

Adding pipeline stages may increase the  $f_{MAX}$  of the design by reducing the combinational logic depth, at the cost of additional latency and logic utilization.

The insertion of pipeline stages requires certain interconnect components. For example, in a system with a single slave interface, there is no multiplexer; therefore multiplexer pipelining does not occur. When there is an Avalon or AXI single-master to single-slave system, no pipelining occurs, regardless of the **Limit interconnect pipeline stages to** option.

### Related Information

- [Creating a System with Qsys](#) on page 5-1

## Using Bridges

You can use bridges to increase system frequency, minimize generated Qsys logic, minimize adapter logic, and to structure system topology when you want to control where Qsys adds pipelining. You can also use bridges with arbiters when there is concurrency in the system.

An Avalon bridge has an Avalon-MM slave interface and an Avalon-MM master interface. You can have many components connected to the bridge slave interface, or many components connected to the bridge master interface. You can also have a single component connected to a single bridge slave or master interface.

You can configure the data width of the bridge, which can affect how Qsys generates bus sizing logic in the interconnect. Both interfaces support Avalon-MM pipelined transfers with variable latency, and can also support configurable burst lengths.

Transfers to the bridge slave interface are propagated to the master interface, which connects to components downstream from the bridge. When you need greater control over interconnect pipelining, you can use bridges instead of the **Limit Interconnect Pipeline Stages to** option.

**Note:** You can use Avalon bridges between AXI interfaces, and between Avalon domains. Qsys automatically creates interconnect logic between the AXI and Avalon interfaces, so you do not have to explicitly instantiate bridges between these domains. For more discussion about the benefits and disadvantages of shared and separate domains, refer to the *Qsys Interconnect*.

### Related Information

- [Creating a System with Qsys](#) on page 5-1
- [Qsys Interconnect](#) on page 7-1

## Using Bridges to Increase System Frequency


In Qsys, you can introduce interconnect pipeline stages or pipeline bridges to increase clock frequency in your system. Bridges control the system interconnect topology and allow you to subdivide the interconnect, giving you more control over pipelining and clock crossing functionality.

### Inserting Pipeline Bridges

You can insert an Avalon-MM pipeline bridge to insert registers in the path between the bridges and its master and slaves. If a critical register-to-register delay occurs in the interconnect, a pipeline bridge can help reduce this delay and improve system  $f_{MAX}$ .

The Avalon-MM pipeline bridge component integrates into any Qsys system. The pipeline bridge options can increase logic utilization and read latency. The change in topology may also reduce concurrency if multiple masters arbitrate for the bridge. You can use the Avalon-MM pipeline bridge to control topology without adding a pipeline stage. A pipeline bridge that does not add a pipeline stage is optimal in some latency-sensitive applications. For example, a CPU may benefit from minimal latency when accessing memory.

**Figure 8-8: Avalon-MM Pipeline Bridge**



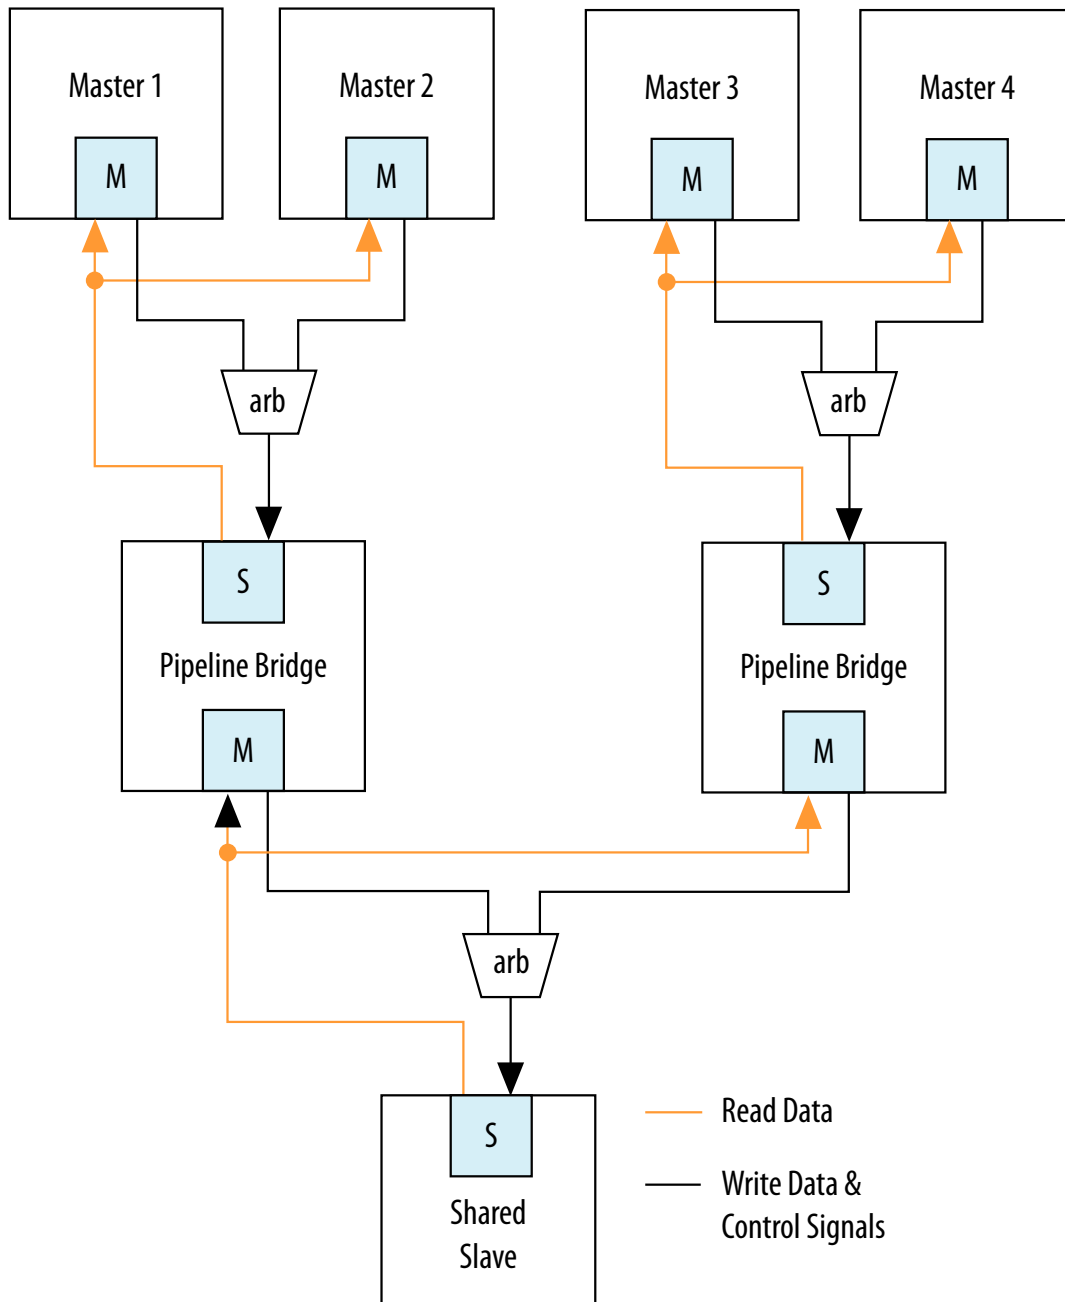
### Implementing Command Pipelining (Master-to-Slave)

When multiple masters share a slave device, you can use command pipelining to improve performance.

The arbitration logic for the slave interface must multiplex the `address`, `writedata`, and `burstcount` signals. The multiplexer width increases proportionally with the number of masters connecting to a single slave interface. The increased multiplexer width may become a timing critical path in the system. If a

single pipeline bridge does not provide enough pipelining, you can instantiate multiple instances of the bridge in a tree structure to increase the pipelining and further reduce the width of the multiplexer at the slave interface.

**Figure 8-9: Tree of Bridges**



**Implementing Response Pipelining (Slave-to-Master)**

When masters connect to multiple slaves that support read transfers, you can use slave-to-master pipelining to improve performance.

The interconnect inserts a multiplexer for every read data path back to the master. As the number of slaves supporting read transfers connecting to the master increases, the width of the read data multiplexer also increases. If the performance increase is insufficient with one bridge, you can use multiple bridges in a tree structure to improve  $f_{MAX}$ .

## Using Clock Crossing Bridges

The clock crossing bridge contains a pair of clock crossing FIFOs, which isolate the master and slave interfaces in separate, asynchronous clock domains. Transfers to the slave interface are propagated to the master interface.

When you use a FIFO clock crossing bridge for the clock domain crossing, you add data buffering. Buffering allows pipelined read masters to post multiple reads to the bridge, even if the slaves downstream from the bridge do not support pipelined transfers.

You can also use a clock crossing bridge to place high and low frequency components in separate clock domains. If you limit the fast clock domain to the portion of your design that requires high performance, you may achieve a higher  $f_{MAX}$  for this portion of the design. For example, the majority of processor peripherals in embedded designs do not need to operate at high frequencies, therefore, you do not need to use a high-frequency clock for these components. When you compile a design with the Quartus II software, compilation may take more time when the clock frequency requirements are difficult to meet because the Fitter needs more time to place registers to achieve the required  $f_{MAX}$ . To reduce the amount of effort that the Fitter uses on low priority and low performance components, you can place these behind a clock crossing bridge operating at a lower frequency, allowing the Fitter to increase the effort placed on the higher priority and higher frequency data paths.

## Using Bridges to Minimize Design Logic

Bridges can reduce interconnect logic by reducing the amount of arbitration and multiplexer logic that Qsys generates. This reduction occurs because bridges limit the number of concurrent transfers that can occur.

## Avoiding Speed Optimizations That Increase Logic

You can add an additional pipeline stage with a pipeline bridge between masters and slaves to reduce the amount of combinational logic between registers, which can increase system performance. If you can increase the  $f_{MAX}$  of your design logic, you may be able to turn off the Quartus II software optimization settings, such as the **Perform register duplication** setting. Register duplication creates duplicate registers in two or more physical locations in the FPGA to reduce register-to-register delays. You may also want to choose **Speed** for the optimization method, which typically results in higher logic utilization due to logic duplication. By making use of the registers or FIFOs available in the bridges, you can increase the design speed and avoid needless logic duplication or speed optimizations, thereby reducing the logic utilization of the design.

## Limiting Concurrency


The amount of logic generated for the interconnect often increases as the system becomes larger because Qsys creates arbitration logic for every slave interface that is shared by multiple master interfaces. Qsys inserts multiplexer logic between master interfaces that connect to multiple slave interfaces if both support read data paths.

Most embedded processor designs contain components that are either incapable of supporting high data throughput, or do not need to be accessed frequently. These components can contain master or slave interfaces. Because the interconnect supports concurrent accesses, you may want to limit concurrency by inserting bridges into the data path to limit the amount of arbitration and multiplexer logic generated.

For example, if a system contains three master and three slave interfaces that are interconnected, Qsys generates three arbiters and three multiplexers for the read data path. If these masters do not require a significant amount of simultaneous throughput, you can reduce the resources that your design consumes by connecting the three masters to a pipeline bridge. The bridge controls the three slave interfaces and reduces the interconnect into a bus structure. Qsys creates one arbitration block between the bridge and the three masters, and a single read data path multiplexer between the bridge and three slaves, and prevents concurrency. This implementation is similar to a standard bus architecture.

You should not use this method for high throughput data paths to ensure that you do not limit overall system performance.

Figure 8-10: Differences Between Systems With and Without a Pipeline Bridge



## Using Bridges to Minimize Adapter Logic

Qsys generates adapter logic for clock crossing, width adaptation, and burst support when there is a mismatch between the clock domains, widths, or bursting capabilities of the master and slave interface pairs.

Qsys creates burst adapters when the maximum burst length of the master is greater than the master burst length of the slave. The adapter logic creates extra logic resources, which can be substantial when your system contains master interfaces connected to many components that do not share the same characteristics. By placing bridges in your design, you can reduce the amount of adapter logic that Qsys generates.

## Determining Effective Placement of Bridges

To determine the effective placement of a bridge, you should initially analyze each master in your system to determine if the connected slave devices support different bursting capabilities or operate in a different clock domain. The maximum burstcount of a component is visible as the `burstcount` signal in the HDL file of the component. The maximum burst length is  $2^{(\text{width}(\text{burstcount} - 1))}$ , therefore, if the `burstcount` width is four bits, the maximum `burstcount` is eight. If no `burstcount` signal is present, the component does not support bursting or has a burst length of 1.

To determine if the system requires a clock crossing adapter between the master and slave interfaces, check the **Clock** column for the master and slave interfaces. If the clock is different for the master and slave interfaces, Qsys inserts a clock crossing adapter between them. To avoid creating multiple adapters, you can place the components containing slave interfaces behind a bridge so that Qsys creates a single adapter. By placing multiple components with the same burst or clock characteristics behind a bridge, you limit concurrency and the number of adapters.

You can also use a bridge to separate AXI and Avalon domains to minimize burst adaptation logic. For example, if there are multiple Avalon slaves that are connected to an AXI master, you can consider inserting a bridge to access the adaptation logic once before the bridge, instead of once per slave. This implementation results in latency, and you would also lose concurrency between reads and writes.

## Changing the Response Buffer Depth

When you use automatic clock-crossing adapters, Qsys determines the required depth of FIFO buffering based on the slave properties. If a slave has a high *Maximum Pending Reads* parameter, the resulting deep response buffer FIFO that Qsys inserts between the master and slave can consume a lot of device resources. To control the response FIFO depth, you can use a clock crossing bridge and manually adjust its FIFO depth to trade off throughput with smaller memory utilization.

For example, if you have masters that cannot saturate the slave, you do not need response buffering. Using a bridge reduces the FIFO memory depth and reduces the **Maximum Pending Reads** available from the slave.

## Considering the Effects of Using Bridges

Before you use pipeline or clock crossing bridges in a design, you should carefully consider their effects. Bridges can have any combination of consequences on your design, which could be positive or negative. Benchmarking your system before and after inserting bridges can help you to determine the impact to the design.

### Increased Latency

Adding a bridge to a design has an effect on the read latency between the master and the slave. Depending on the system requirements and the type of master and slave, this latency increase may or may not be acceptable in your design.

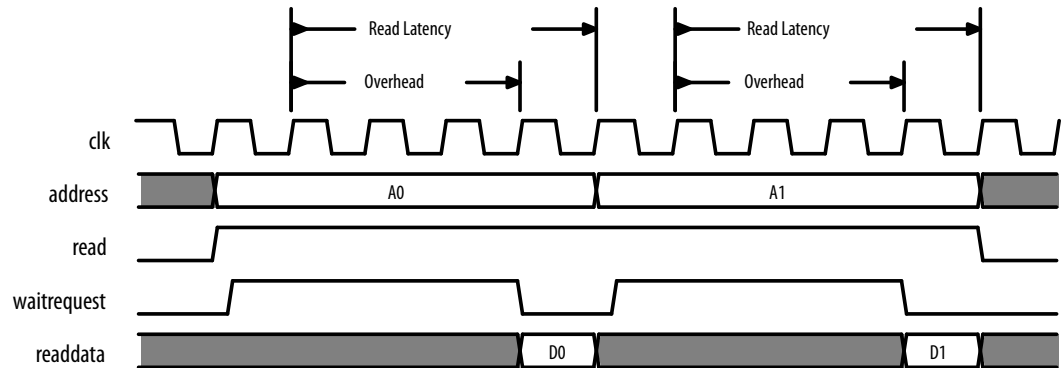
#### Acceptable Latency Increase

For a pipeline bridge, Qsys adds a cycle of latency for each pipeline option that is enabled. The buffering in the clock crossing bridge also adds latency. If you use a pipelined or burst master that posts many read transfers, the increase in latency does not impact performance significantly because the latency increase is very small compared to the length of the data transfer.



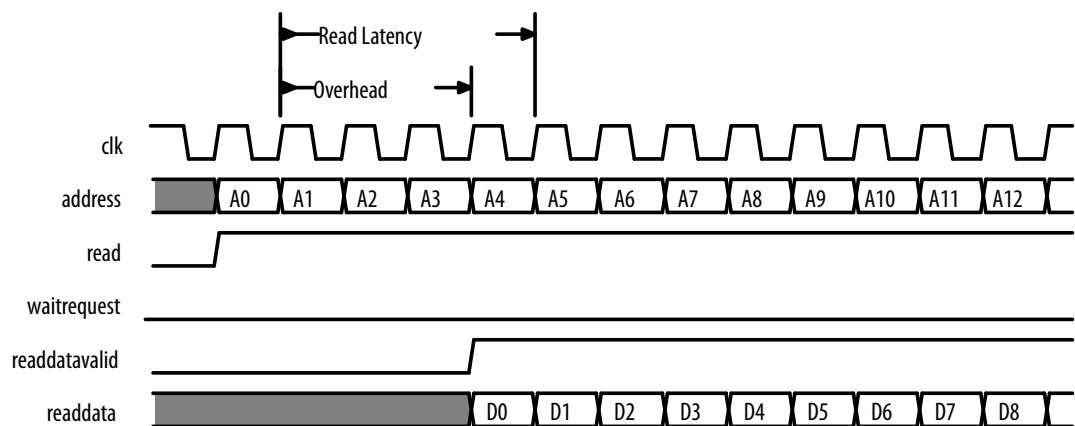
For example, if you use a pipelined read master such as a DMA controller to read data from a component with a fixed read latency of four clock cycles, but only perform a single word transfer, the overhead is three clock cycles out of the total of four. This is true when there is no additional pipeline latency in the interconnect. The read throughput is only 25%.

Figure 8-11: Low-Efficiency Read Transfer



However, if 100 words of data are transferred without interruptions, the overhead is three cycles out of the total of 103 clock cycles. This corresponds to a read efficiency of approximately 97% when there is no additional pipeline latency in the interconnect. Adding a pipeline bridge to this read path adds two extra clock cycles of latency. The transfer requires 105 cycles to complete, corresponding to an efficiency of approximately 94%. Although the efficiency decreased by 3%, adding the bridge may increase the  $f_{MAX}$  by 5%. For example, if the clock frequency can be increased, the overall throughput would improve. As the number of words transferred increases, the efficiency increases to nearly 100%, whether or not a pipeline bridge is present.

Figure 8-12: High Efficiency Read Transfer




### Unacceptable Latency Increase

Processors are sensitive to high latency read times and typically retrieve data for use in calculations that cannot proceed until the data arrives. Before adding a bridge to the data path of a processor instruction or data master, determine whether the clock frequency increase justifies the added latency.


A Nios II processor instruction master has a cache memory with a read latency of four cycles, which is eight sequential words of data return for each read. At 100 MHz, the first read takes 40 ns to complete. Each successive word takes 10 ns so that eight reads complete in 110 ns.

**Figure 8-13: Performance of a Nios II Processor and Memory Operating at 100 MHz**



Adding a clock crossing bridge allows the memory to operate at 125 MHz. However, this increase in frequency is negated by the increase in latency because if the clock crossing bridge adds six clock cycles of latency at 100 MHz, then the memory continues to operate with a read latency of four clock cycles. Consequently, the first read from memory takes 100 ns, and each successive word takes 10 ns because reads arrive at the frequency of the processor, which is 100 MHz. In total, eight reads complete after 170 ns. Although the memory operates at a higher clock frequency, the frequency at which the master operates limits the throughput.

**Figure 8-14: Performance of a Nios II Processor and Eight Reads with Ten Cycles Latency**




## Limited Concurrency

Placing a bridge between multiple master and slave interfaces limits the number of concurrent transfers your system can initiate. This limitation is the same when connecting multiple master interfaces to a single slave interface. The slave interface of the bridge is shared by all the masters and, as a result, Qsys

creates arbitration logic. If the components placed behind a bridge are infrequently accessed, this concurrency limitation may be acceptable.

Bridges can have a negative impact on system performance if you use them inappropriately. For example, if multiple memories are used by several masters, you should not place the memory components behind a bridge. The bridge limits memory performance by preventing concurrent memory accesses. Placing multiple memory components behind a bridge can cause the separate slave interfaces to appear as one large memory to the masters accessing the bridge; all masters must access the same slave interface.

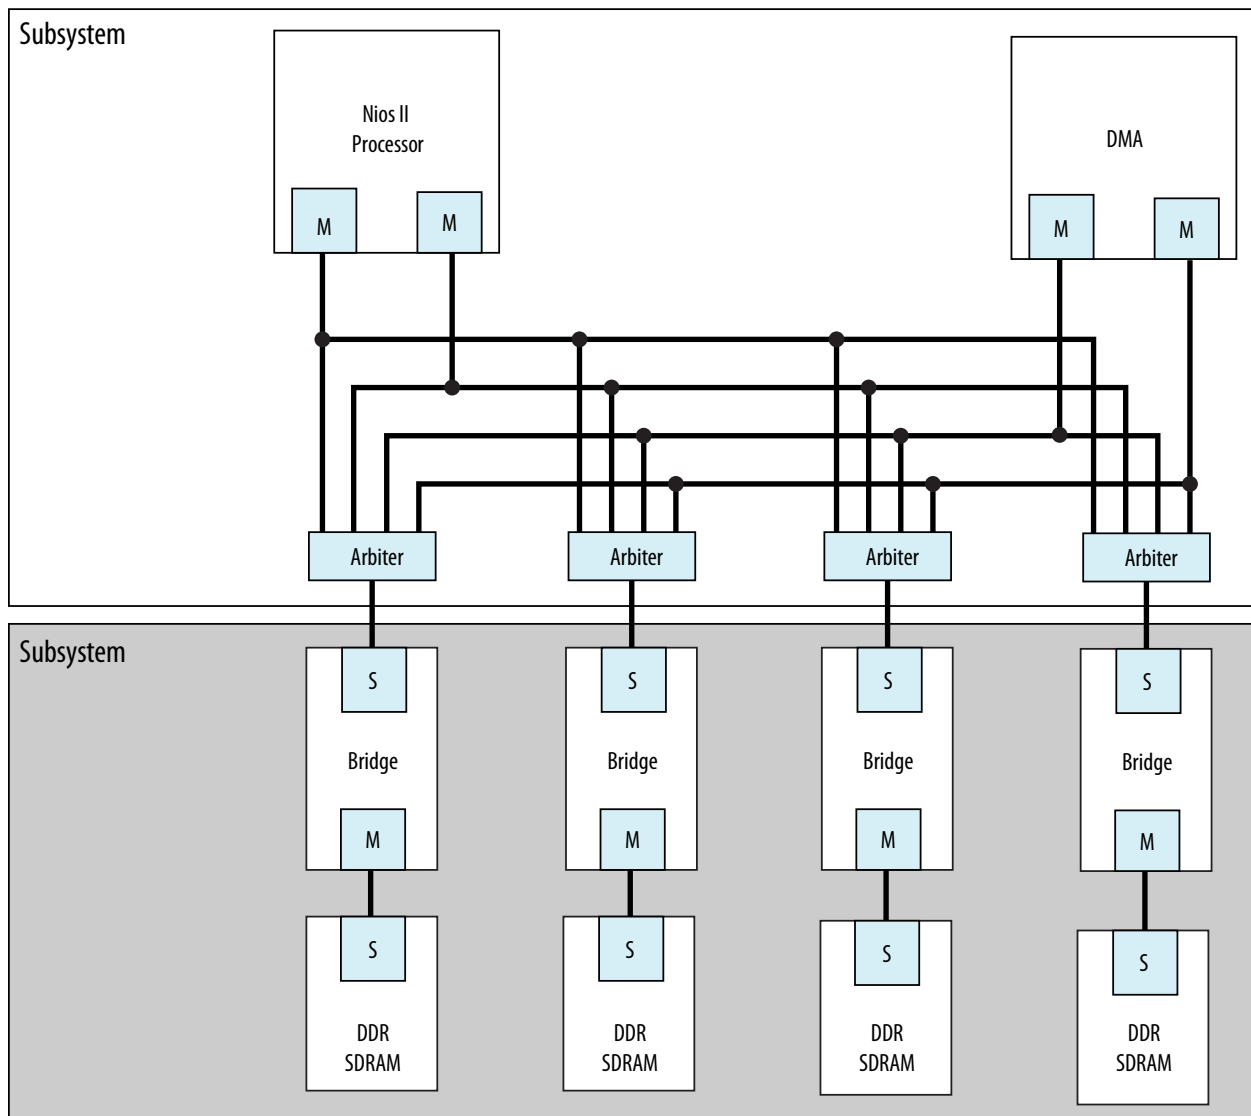
Figure 8-15: Inappropriate Use of a Bridge in a Hierarchical System



A memory subsystem with one bridge that acts as a single slave interface for the Avalon-MM Nios II and DMA masters, which results in a bottleneck architecture. The bridge acts as a bottleneck between the two masters and the memories.

If the  $f_{MAX}$  of your memory interfaces is low and you want to use a pipeline bridge between subsystems, you can place each memory behind its own bridge, which increases the  $f_{MAX}$  of the system without sacrificing concurrency.

Figure 8-16: Efficient Memory Pipelining Without a Bottleneck in a Hierarchical System

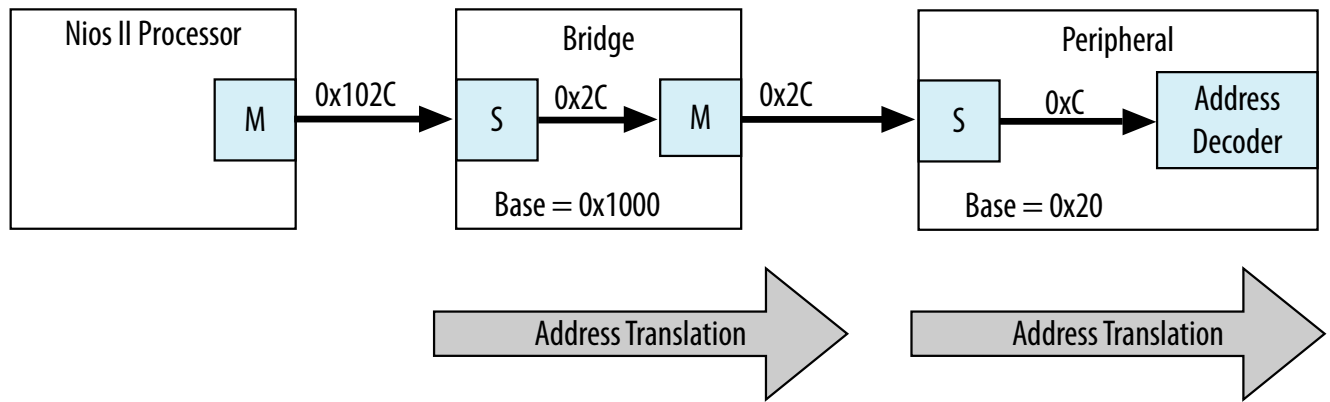


## Address Space Translation

The slave interface of a pipeline or clock crossing bridge has a base address and address span. You can set the base address, or allow Qsys to set it automatically. The address of the slave interface is the base offset address of all the components connected to the bridge. The address of components connected to the bridge is the sum of the base offset and the address of that component.

The master interface of the bridge drives only the address bits that represent the offset from the base address of the bridge slave interface. Any time a master accesses a slave through a bridge, both addresses must be added together, otherwise the transfer fails. The **Address Map** tab displays the addresses of the slaves connected to each master and includes address translations caused by system bridges.

**Figure 8-17: Bridge Address Translation**

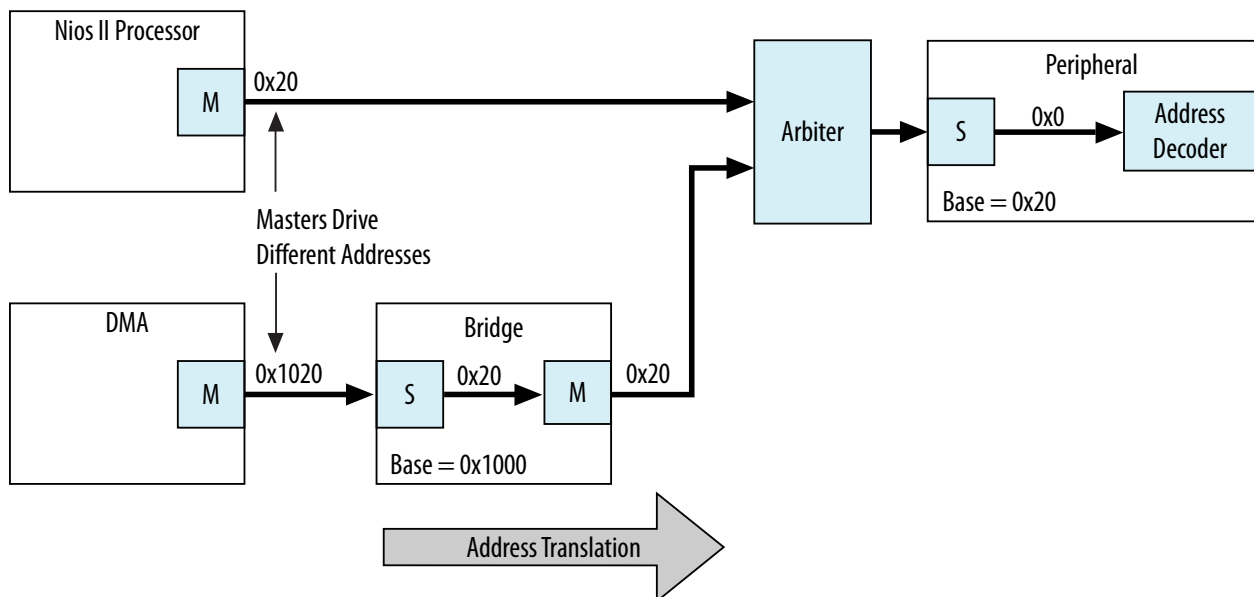


In this example, the Nios II processor connects to a bridge located at base address 0x1000, a slave connects to the bridge master interface at an offset of 0x20, and the processor performs a write transfer to the fourth 32-bit or 64-bit word within the slave. Nios II drives the address 0x102C to interconnect, which is within the address range of the bridge. The bridge master interface drives 0x2C, which is within the address range of the slave, and the transfer completes.

### Address Coherency

To simplify the system design, all masters should access slaves at the same location. In many systems, a processor passes buffer locations to other mastering components, such as a DMA controller. If the processor and DMA controller do not access the slave at the same location, Qsys must compensate for the differences.

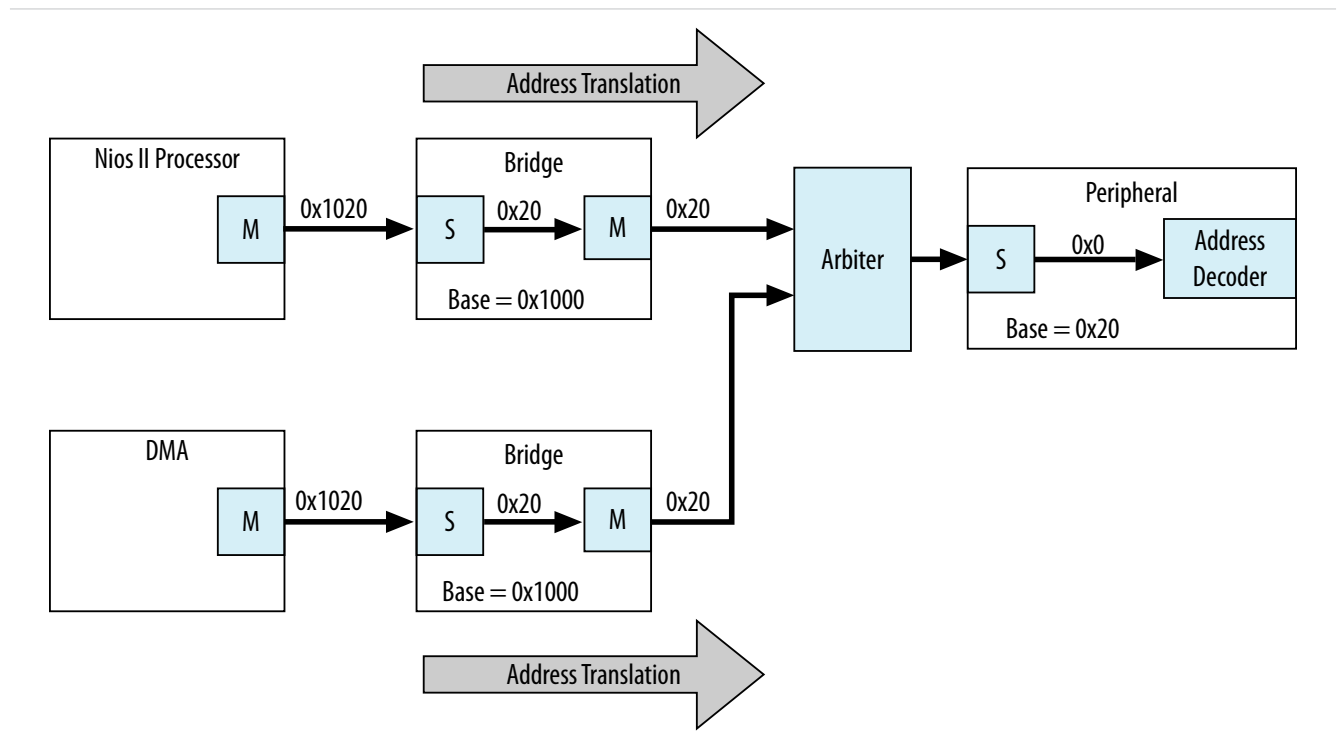
**Figure 8-18: Slaves at Different Addresses and Complicating the System**



A Nios II processor and DMA controller access a slave interface located at address 0x20. The processor connects directly to the slave interface. The DMA controller connects to a pipeline bridge located at address 0x1000, which then connects to the slave interface. Because the DMA controller accesses the pipeline bridge first, it must drive 0x1020 to access the first location of the slave interface. Because the processor accesses the slave from a different location, you must maintain two base addresses for the slave device.

To avoid the requirement for two addresses, you can add an additional bridge to the system, set its base address to 0x1000, and then disable all the pipelining options in the second bridge so that the bridge has minimal impact on system timing and resource utilization. Because this second bridge has the same base address as the original bridge, the processor and DMA controller access the slave interface with the same address range.

Figure 8-19: Address Translation Corrected With Bridge



## Increasing Transfer Throughput

Increasing the transfer efficiency of the master and slave interfaces in your system increases the throughput of your design. Designs with strict cost or power requirements benefit from increasing the transfer efficiency because you can then use less expensive, lower frequency devices. Designs requiring high performance also benefit from increased transfer efficiency because increased efficiency improves the performance of frequency-limited hardware.

Throughput is the number of symbols (such as bytes) of data that Qsys can transfer in a given clock cycle. Read latency is the number of clock cycles between the address and data phase of a transaction. For example, a read latency of two means that the data is valid two cycles after the address is posted. If the

master must wait for one request to finish before the next begins, such as with a processor, then the read latency is very important to the overall throughput.

You can measure throughput and latency in simulation by observing the waveforms, or using the verification IP monitors.

#### Related Information

- [Avalon Verification IP Suite User Guide](#)
- [Mentor® Verification IP Altera Edition AMBA AXI3/4 User Guide](#)

## Using Pipelined Transfers

Pipelined transfers increase the read efficiency by allowing a master to post multiple reads before data from an earlier read returns. Masters that support pipelined transfers post transfers continuously, relying on the `readdatavalid` signal to indicate valid data. Slaves support pipelined transfers by including the `readdatavalid` signal or operating with a fixed read latency.

AXI masters declare how many outstanding writes and reads it can issue with the `writeIssuingCapability` and `readIssuingCapability` parameters. In the same way, a slave can declare how many reads it can accept with the `readAcceptanceCapability` parameter. AXI masters with a read issuing capability greater than one are pipelined in the same way as Avalon masters and the `readdatavalid` signal.

## Using the Maximum Pending Reads Parameter

If you create a custom component with a slave interface supporting variable-latency reads, you must specify the **Maximum Pending Reads** parameter in the Component Editor. Qsys uses this parameter to generate the appropriate interconnect and represent the maximum number of read transfers that your pipelined slave component can process. If the number of reads presented to the slave interface exceeds the **Maximum Pending Reads** parameter, then the slave interface must assert `waitrequest`.

Optimizing the value of the **Maximum Pending Reads** parameter requires an understanding of the latencies of your custom components. This parameter should be based on the component's highest read latency for the various logic paths inside the component. For example, if your pipelined component has two modes, one requiring two clock cycles and the other five, set the **Maximum Pending Reads** parameter to 5 to allow your component to pipeline five transfers, and eliminating dead cycles after the initial five-cycle latency.

You can also determine the correct value for the **Maximum Pending Reads** parameter by monitoring the number of reads that are pending during system simulation or while running the hardware. To use this method, set the parameter to a high value and use a master that issues read requests on every clock. You can use a DMA for this task as long as the data is written to a location that does not frequently assert `waitrequest`. If you implement this method, you can observe your component with a logic analyzer or built-in monitoring hardware.

Choosing the correct value for the **Maximum Pending Reads** parameter of your custom pipelined read component is important. If you underestimate the parameter value, you may cause a master interface to stall with a `waitrequest` until the slave responds to an earlier read request and frees a FIFO position.

The **Maximum Pending Reads** parameter controls the depth of the response FIFO inserted into the interconnect for each master connected to the slave. This FIFO does not use significant hardware resources. Overestimating the **Maximum Pending Reads** parameter results in a slight increase in



hardware utilization. For these reasons, if you are not sure of the optimal value, you should overestimate this value.

If your system includes a bridge, you must set the **Maximum Pending Reads** parameter on the bridge as well. To allow maximum throughput, this value should be equal to or greater than the **Maximum Pending Reads** value for the connected slave that has the highest value. You can limit the maximum pending reads of a slave and reduce the buffer depth by reducing the parameter value on the bridge if the high throughput is not required. If you do not know the **Maximum Pending Reads** value for all the slave components, you can monitor the number of reads that are pending during system simulation while running the hardware. To use this method, set the **Maximum Pending Reads** parameter to a high value and use a master that issues read requests on every clock, such as a DMA. Then, reduce the number of maximum pending reads of the bridge until the bridge reduces the performance of any masters accessing the bridge.

## Arbitration Shares and Bursts

Arbitration shares provide control over the arbitration process. By default, the arbitration algorithm allocates evenly, with all masters receiving one share.

You can adjust the arbitration process by assigning a larger number of shares to the masters that need greater throughput. The larger the arbitration share, the more transfers are allocated to the master to access a slave. The masters gets uninterrupted access to the slave for its number of shares, as long as the master is reading or writing.

If a master cannot post a transfer and other masters are waiting to gain access to a particular slave, the arbiter grants another master access. This mechanism prevents a master from wasting arbitration cycles if it cannot post back-to-back transfers. A bursting transaction contains multiple beats (or words) of data, starting from a single address. Bursts allow a master to maintain access to a slave for more than a single word transfer. If a bursting master posts a write transfer with a burst length of eight, it is guaranteed arbitration for eight write cycles.

You can assign arbitration shares to an Avalon-MM bursting master and AXI masters (which are always considered a bursting master). Each share consists of one burst transaction (such as multi-cycle write), and allows a master to complete a number of bursts before arbitration switches to the next master.

### Related Information

- [Avalon Interface Specifications](#)
- [AMBA Protocol Specification](#)

## Differences Between Arbitration Shares and Bursts

The following three key characteristics distinguish arbitration shares and bursts:

- Arbitration Lock
- Sequential Addressing
- Burst Adapters

### Arbitration Lock

When a master posts a burst transfer, the arbitration is locked for that master; consequently, the bursting master should be capable of sustaining transfers for the duration of the locked period. If, after the fourth

write, the master deasserts the write signal (Avalon-MM write or AXI `wvalid`) for fifty cycles, all other masters continue to wait for access during this stalled period.

To avoid wasted bandwidth, your master designs should wait until a full burst transfer is ready before requesting access to a slave device. Alternatively, you can avoid wasted bandwidth by posting `burstcounts` equal to the amount of data that is ready. For example, if you create a custom bursting write master with a maximum `burstcount` of eight, but only three words of data are ready, you can present a `burstcount` of three. This strategy does not result in optimal use of the system band width if the slave is capable of handling a larger burst; however, this strategy prevents stalling and allows access for other masters in the system.

## Sequential Addressing

An Avalon-MM burst transfer includes a base address and a `burstcount`, which represents the number of words of data that are transferred, starting from the base address and incrementing sequentially. Burst transfers are common for processors, DMAs, and buffer processing accelerators; however, sometimes a master must access non-sequential addresses. Consequently, a bursting master must set the `burstcount` to the number of sequential addresses, and then reset the `burstcount` for the next location.

The arbitration share algorithm has no restrictions on addresses; therefore, your custom master can update the address it presents to the interconnect for every read or write transaction.

## Burst Adapters

Qsys allows you to create systems that mix bursting and non-bursting master and slave interfaces. This design strategy allows you to connect bursting master and slave interfaces that support different maximum burst lengths, with Qsys generating burst adapters when appropriate.

Qsys inserts a burst adapter whenever a master interface burst length exceeds the burst length of the slave interface, or if the master issues a burst type that the slave cannot support. For example, if you connect an AXI master to an Avalon slave, a burst adapter is inserted. Qsys assigns non-bursting masters and slave interfaces a burst length of one. The burst adapter divides long bursts into shorter bursts. As a result, the burst adapter adds logic to the address and `burstcount` paths between the master and slave interfaces.

### Related Information

[Qsys Interconnect](#) on page 7-1

[AMBA Protocol Specification](#)

## Choosing Avalon-MM Interface Types

To avoid inefficient Avalon-MM transfers, custom master or slave interfaces must use the appropriate simple, pipelined, or burst interfaces.

### Simple Avalon-MM Interfaces

Simple interface transfers do not support pipelining or bursting for reads or writes; consequently, their performance is limited. Simple interfaces are appropriate for transfers between masters and infrequently used slave interfaces. In Qsys, the PIO, UART, and Timer include slave interfaces that use simple transfers.

### Pipelined Avalon-MM Interfaces

Pipelined read transfers allow a pipelined master interface to start multiple read transfers in succession without waiting for prior transfers to complete. Pipelined transfers allow master-slave pairs to achieve

higher throughput, even though the slave port may require one or more cycles of latency to return data for each transfer.

In many systems, read throughput becomes inadequate if simple reads are used and pipelined transfers can increase throughput. If you define a component with a fixed read latency, Qsys automatically provides the pipelining logic necessary to support pipelined reads. You can use fixed latency pipelining as the default design starting point for slave interfaces. If your slave interface has a variable latency response time, use the `readdatavalid` signal to indicate when valid data is available. The interconnect implements read response FIFO buffering to handle the maximum number of pending read requests.

To use components that support pipelined read transfers, and to use a pipelined system interconnect efficiently, your system must contain pipelined masters. You can use pipelined masters as the default starting point for new master components. Use the `readdatavalid` signal for these master interfaces.

Because master and slaves sometimes have mismatched pipeline latency, interconnect contains logic to reconcile the differences.

**Table 8-1: Pipeline Latency in a Master-Slave Pair**

Master	Slave	Pipeline Management Logic Structure
No pipeline	No Pipeline	Qsys interconnect does not instantiate logic to handle pipeline latency.
No pipeline	Pipelined with fixed or variable latency	Qsys interconnect forces the master to wait through any slave-side latency cycles. This master-slave pair gains no benefits from pipelining, because the master waits for each transfer to complete before beginning a new transfer. However, while the master is waiting, the slave can accept transfers from a different master.
Pipelined	No pipeline	Qsys interconnect carries out the transfer as if neither master nor slave were pipelined, causing the master to wait until the slave returns data. An example of a non-pipeline slave is an asynchronous off-chip interface.
Pipelined	Pipelined with fixed latency	Qsys interconnect allows the master to capture data at the exact clock cycle when data from the slave is valid, to enable maximum throughput. An example of a fixed latency slave is an on-chip memory.
Pipelined	Pipelined with variable latency	The slave asserts a signal when its <code>readdata</code> is valid, and the master captures the data. The master-slave pair can achieve maximum throughput if the slave has variable latency. Examples of variable latency slaves include SDRAM and FIFO memories.

### Burst Avalon-MM Interfaces

Burst transfers are commonly used for latent memories such as SDRAM and off-chip communication interfaces, such as PCI Express. To use a burst-capable slave interface efficiently, you must connect to a bursting master. Components that require bursting to operate efficiently typically have an overhead penalty associated with short bursts or non-bursting transfers.

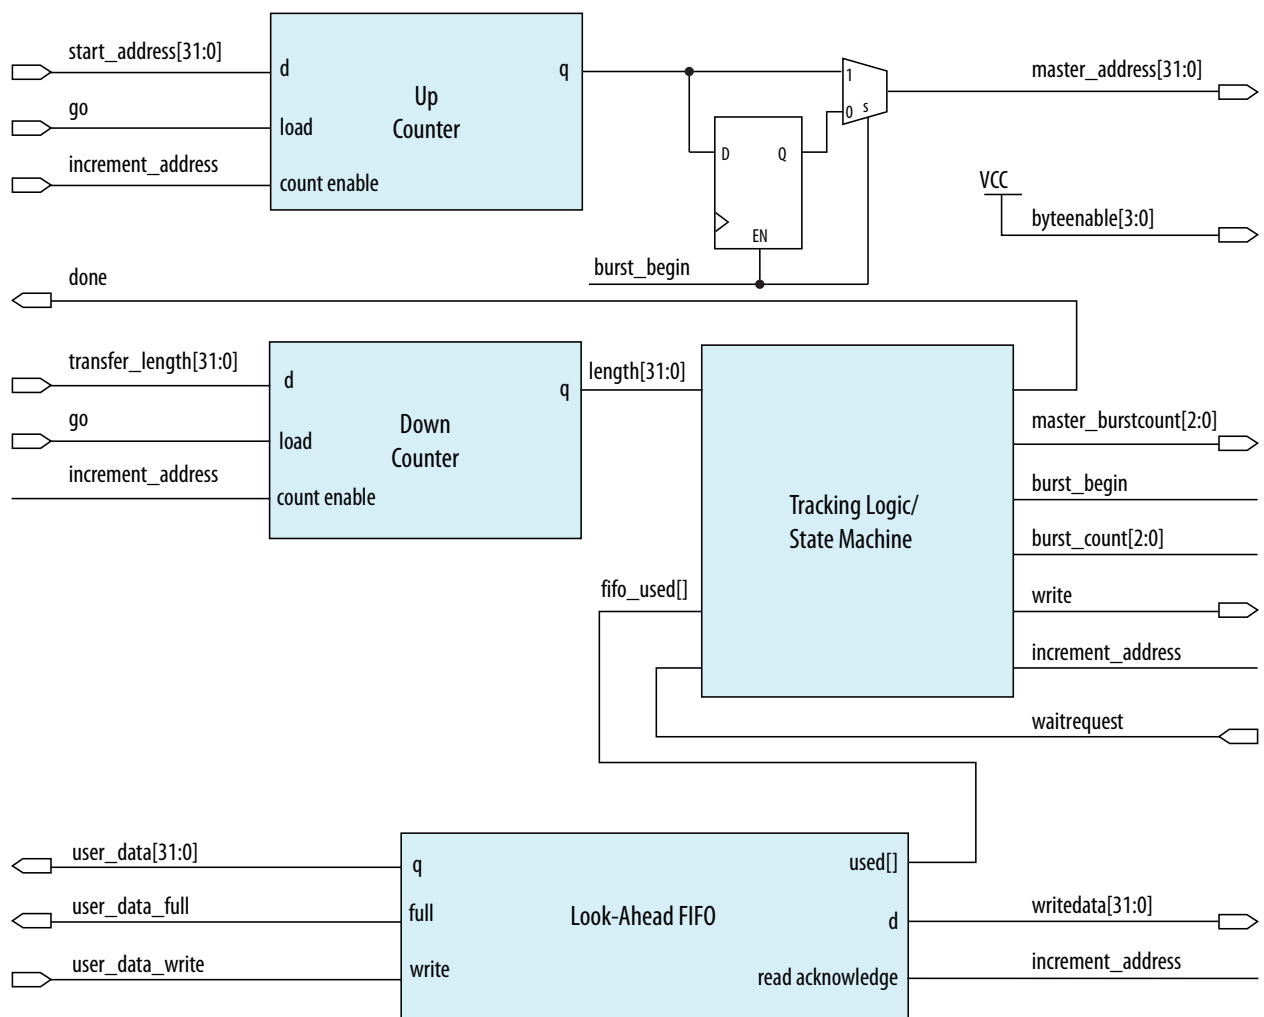
You can use a burst-capable slave interface if you know that your component requires sequential transfers to operate efficiently. Because SDRAM memories incur a penalty when switching banks or rows, performance improves when SDRAM memories are accessed sequentially with bursts.

Architectures that use the same signals to transfer address and data also benefit from bursting. Whenever an address is transferred over shared address and data signals, the throughput of the data transfer is reduced. Because the address phase adds overhead, using large bursts increases the throughput of the connection.

## Avalon-MM Burst Master Example

Figure 8-20: Avalon Bursting Write Master

This example shows the architecture of a bursting write master that receives data from a FIFO and writes the contents to memory. You can use a bursting master as a starting point for your own bursting components, such as custom DMAs, hardware accelerators, or off-chip communication interfaces.



The master performs word accesses and writes to sequential memory locations. When `go` is asserted, the `start_address` and `transfer_length` are registered. On the next clock cycle, the control logic asserts `burst_begin`, which synchronizes the internal control signals in addition to the `master_address` and

`master_burstcount` presented to the interconnect. The timing of these two signals is important because during bursting write transfers, `byteenable`, and `burstcount` must be held constant for the entire burst.

To avoid inefficient writes, the master posts a burst when enough data is buffered in the FIFO. To maximize the burst efficiency, the master should stall only when a slave asserts `waitrequest`. In this example, the FIFO's used signal tracks the number of words of data that are stored in the FIFO and determines when enough data has been buffered.

The `address` register increments after every word transfer, and the `length` register decrements after every word transfer. The address remains constant throughout the burst. Because a transfer is not guaranteed to complete on burst boundaries, additional logic is necessary to recognize the completion of short bursts and complete the transfer.

#### Related Information

- [Avalon Memory-Mapped Master Templates](#)

## Reducing Logic Utilization

You can minimize logic size of Qsys systems. Typically, there is a trade-off between logic utilization and performance. Reducing logic utilization applies to both Avalon and AXI interfaces.

### Minimizing Interconnect Logic to Reduce Logic Utilization

In Qsys, changes to the connections between master and slave reduce the amount of interconnect logic required in the system.

#### Related Information

[Limited Concurrency](#) on page 8-19

### Creating Dedicated Master and Slave Connections to Minimize Interconnect Logic

You can create a system where a master interface connects to a single slave interface. This configuration eliminates address decoding, arbitration, and return data multiplexing, which simplifies the interconnect. Dedicated master-to-slave connections attain the same clock frequencies as Avalon-ST connections.

Typically, these one-to-one connections include an Avalon memory-mapped bridge or hardware accelerator. For example, if you insert a pipeline bridge between a slave and all other master interfaces, the logic between the bridge master and slave interface is reduced to wires. If a hardware accelerator connects only to a dedicated memory, no system interconnect logic is generated between the master and slave pair.

### Removing Unnecessary Connections to Minimize Interconnect Logic

The number of connections between master and slave interfaces affects the  $f_{MAX}$  of your system. Every master interface that you connect to a slave interface increases the width of the multiplexer width. As a multiplexer width increases, so does the logic depth and width that implements the multiplexer in the FPGA. To improve system performance, connect masters and slaves only when necessary.

When you connect a master interface to many slave interfaces, the multiplexer for the read data signal grows. Avalon typically uses a `readdata` signal. AXI read data signals add a response status and last indicator to the read response channel using `rdata`, `rresp`, and `rlast`. Additionally, bridges help control the depth of multiplexers.

**Related Information**

[Implementing Command Pipelining \(Master-to-Slave\)](#) on page 8-12

**Simplifying Address Decode Logic**

If address code logic is in the critical path, you may be able to change the address map to simplify the decode logic. Experiment with different address maps, including a one-hot encoding, to see if results improve.

**Minimizing Arbitration Logic by Consolidating Multiple Interfaces**

As the number of components in a design increases, the amount of logic required to implement the interconnect also increases. The number of arbitration blocks increases for every slave interface that is shared by multiple master interfaces. The width of the read data multiplexer increases as the number of slave interfaces supporting read transfers increases on a per master interface basis. For these reasons, consider implementing multiple blocks of logic as a single interface to reduce interconnect logic utilization.

**Logic Consolidation Trade-Offs**

You should consider the following trade-offs before making modifications to your system or interfaces:

- Consider the impact on concurrency that results when you consolidate components. When a system has four master components and four slave interfaces, it can initiate four concurrent accesses. If you consolidate the four slave interfaces into a single interface, then the four masters must compete for access. Consequently, you should only combine low priority interfaces such as low speed parallel I/O devices if the combination does not impact the performance.
- Determine whether consolidation introduces new decode and multiplexing logic for the slave interface that the interconnect previously included. If an interface contains multiple read and write address locations, the interface already contains the necessary decode and multiplexing logic. When you consolidate interfaces, you typically reuse the decoder and multiplexer blocks already present in one of the original interfaces; however, combining interfaces may simply move the decode and multiplexer logic, rather than eliminate duplication.
- Consider whether consolidating interfaces makes the design complicated. If so, you should not consolidate interfaces.

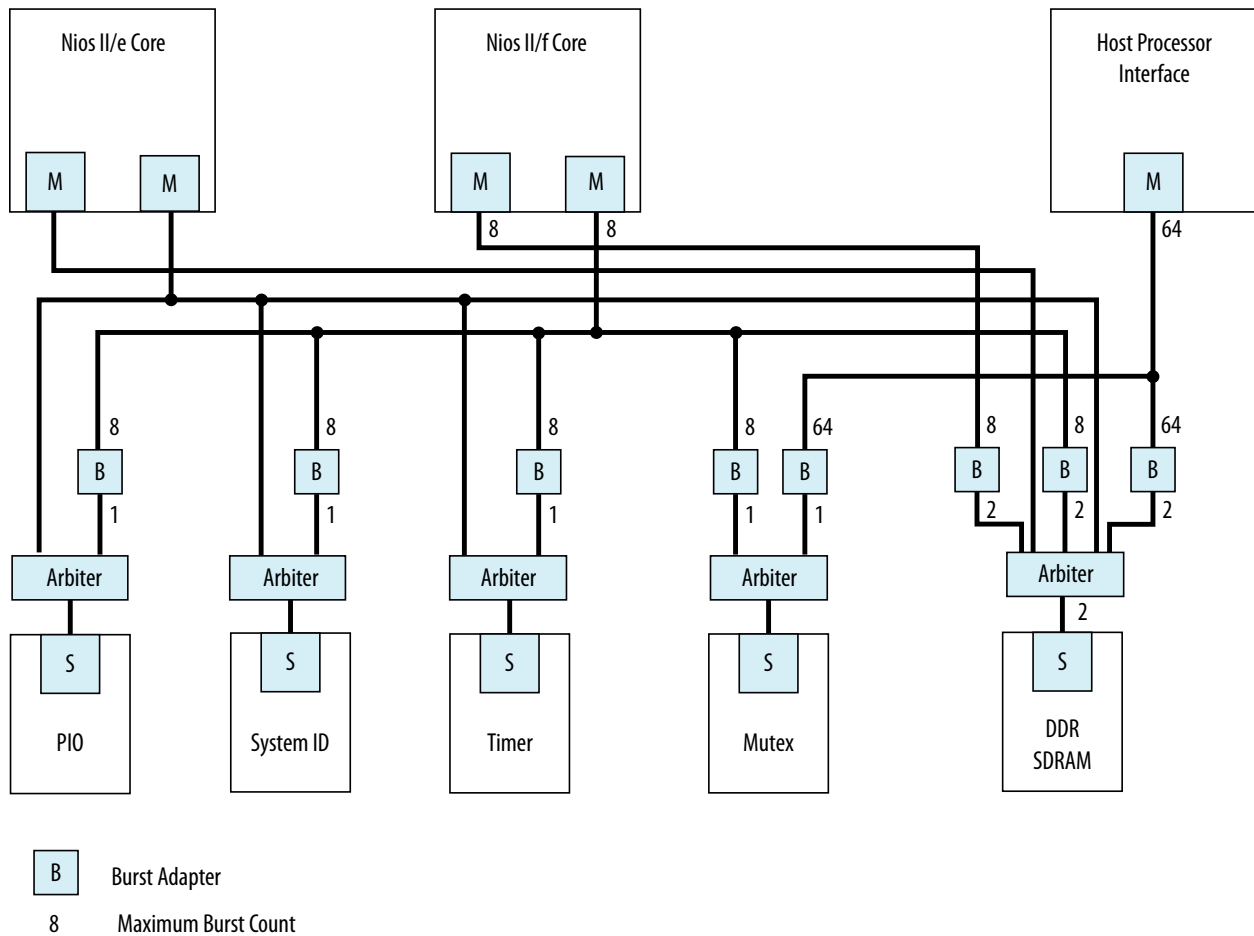
**Related Information**

[Using Concurrency in Memory-Mapped Systems](#) on page 8-6

**Consolidating Interfaces**

The Nios II/e core maintains communication between the Nios II /f core and external processors. The Nios II/f core supports a maximum burst size of eight. The external processor interface supports a maximum burst length of 64. The Nios II/e core does not support bursting. The memory in the system is SDRAM with an Avalon maximum burst length of two.

Figure 8-21: Mixed Bursting System




In this example a system with a mix of components with different burst capabilities with a Nios II/e core, a Nios II/f core, and an external processor, which off-loads some processing tasks to the Nios II/f core.

Qsys automatically inserts burst adapters to compensate for burst length mismatches. The adapters reduce bursts to a single transfer, or the length of two transfers. For the external processor interface connecting to DDR SDRAM, a burst of 64 words is divided into 32 burst transfers, each with a burst length of two. When you generate a system, Qsys inserts burst adapters based on maximum `burstcount` values; consequently, the interconnect logic includes burst adapters between masters and slave pairs that do not require bursting, if the master is capable of bursts.

In this example, Qsys inserts a burst adapter between the Nios II processors and the timer, system ID, and PIO peripherals. These components do not support bursting and the Nios II processor performs a single word read and write accesses to these components.

**Figure 8-22: Mixed Bursting System with Bridges**

To reduce the number of adapters, you can add pipeline bridges. The pipeline bridge, between the Nios II/f core and the peripherals that do not support bursts, eliminates three burst adapters from the previous example. A second pipeline bridge between the Nios II/f core and the DDR SDRAM, with its maximum burst size set to eight, eliminates another burst adapter, as shown below.



## Reducing Logic Utilization With Multiple Clock Domains

You specify clock domains in Qsys on the **System Contents** tab. Clock sources can be driven by external input signals to Qsys, or by PLLs inside Qsys. Clock domains are differentiated based on the name of the clock. You can create multiple asynchronous clocks with the same frequency.




Qsys generates Clock Domain Crossing Logic (CDC) that hides the details of interfacing components operating in different clock domains. The interconnect supports the memory-mapped protocol with each port independently, and therefore masters do not need to incorporate clock adapters in order to interface to slaves on a different domain. Qsys interconnect logic propagates transfers across clock domain boundaries automatically.

Clock-domain adapters provide the following benefits:

- Allows component interfaces to operate at different clock frequencies.
- Eliminates the need to design CDC hardware.
- Allows each memory-mapped port to operate in only one clock domain, which reduces design complexity of components.
- Enables masters to access any slave without communication with the slave clock domain.
- Allows you to focus performance optimization efforts on components that require fast clock speed.

A clock domain adapter consists of two finite state machines (FSM), one in each clock domain, that use a hand-shaking protocol to propagate transfer control signals (`read_request`, `write_request`, and the master `waitrequest` signals) across the clock boundary.

**Figure 8-23: Clock Crossing Adapter**



This example illustrates a clock domain adapter between one master and one slave. The synchronizer blocks use multiple stages of flip flops to eliminate the propagation of meta-stable events on the control signals that enter the handshake FSMs. The CDC logic works with any clock ratio.

The typical sequence of events for a transfer across the CDC logic is as follows:

- The master asserts address, data, and control signals.
- The master handshake FSM captures the control signals and immediately forces the master to wait. The FSM uses only the control signals, not address and data. For example, the master simply holds the address signal constant until the slave side has safely captured it.
- The master handshake FSM initiates a transfer request to the slave handshake FSM.
- The transfer request is synchronized to the slave clock domain.
- The slave handshake FSM processes the request, performing the requested transfer with the slave.
- When the slave transfer completes, the slave handshake FSM sends an acknowledge back to the master handshake FSM. The acknowledge is synchronized back to the master clock domain.
- The master handshake FSM completes the transaction by releasing the master from the wait condition.

Transfers proceed as normal on the slave and the master side, without a special protocol to handle crossing clock domains. From the perspective of a slave, there is nothing different about a transfer initiated by a master in a different clock domain. From the perspective of a master, a transfer across clock domains simply requires extra clock cycles. Similar to other transfer delay cases (for example, arbitration delay or wait states on the slave side), the Qsys forces the master to wait until the transfer terminates. As a result, pipeline master ports do not benefit from pipelining when performing transfers to a different clock domain.

Qsys automatically determines where to insert CDC logic based on the system and the connections between components, and places CDC logic to maintain the highest transfer rate for all components. Qsys evaluates the need for CDC logic for each master and slave pair independently, and generates CDC logic wherever necessary.

#### Related Information

[Avalon Memory-Mapped Design Optimizations](#)

## Duration of Transfers Crossing Clock Domains

CDC logic extends the duration of master transfers across clock domain boundaries. In the worst case, which is for reads, each transfer is extended by five master clock cycles and five slave clock cycles. Assuming the default value of 2 for the master domain synchronizer length and the slave domain synchronizer length, the components of this delay are the following:

- Four additional master clock cycles, due to the master-side clock synchronizer.
- Four additional slave clock cycles, due to the slave-side clock synchronizer.
- One additional clock in each direction, due to potential metastable events as the control signals cross clock domains.

**Note:** Systems that require a higher performance clock should use the Avalon-MM clock crossing bridge instead of the automatically inserted CDC logic. The clock crossing bridge includes a buffering mechanism so that multiple reads and writes can be pipelined. After paying the initial penalty for the first read or write, there is no additional latency penalty for pending reads and writes, increasing throughput by up to four times, at the expense of added logic resources.

## Reducing Power Consumption

Qsys provides various low power design changes that enable you to reduce the power consumption of the interconnect and custom components.

### Reducing Power Consumption With Multiple Clock Domains

When you use multiple clock domains, you should put non-critical logic in the slower clock domain. Qsys automatically reconciles data crossing over asynchronous clock domains by inserting clock crossing logic (handshake or FIFO).

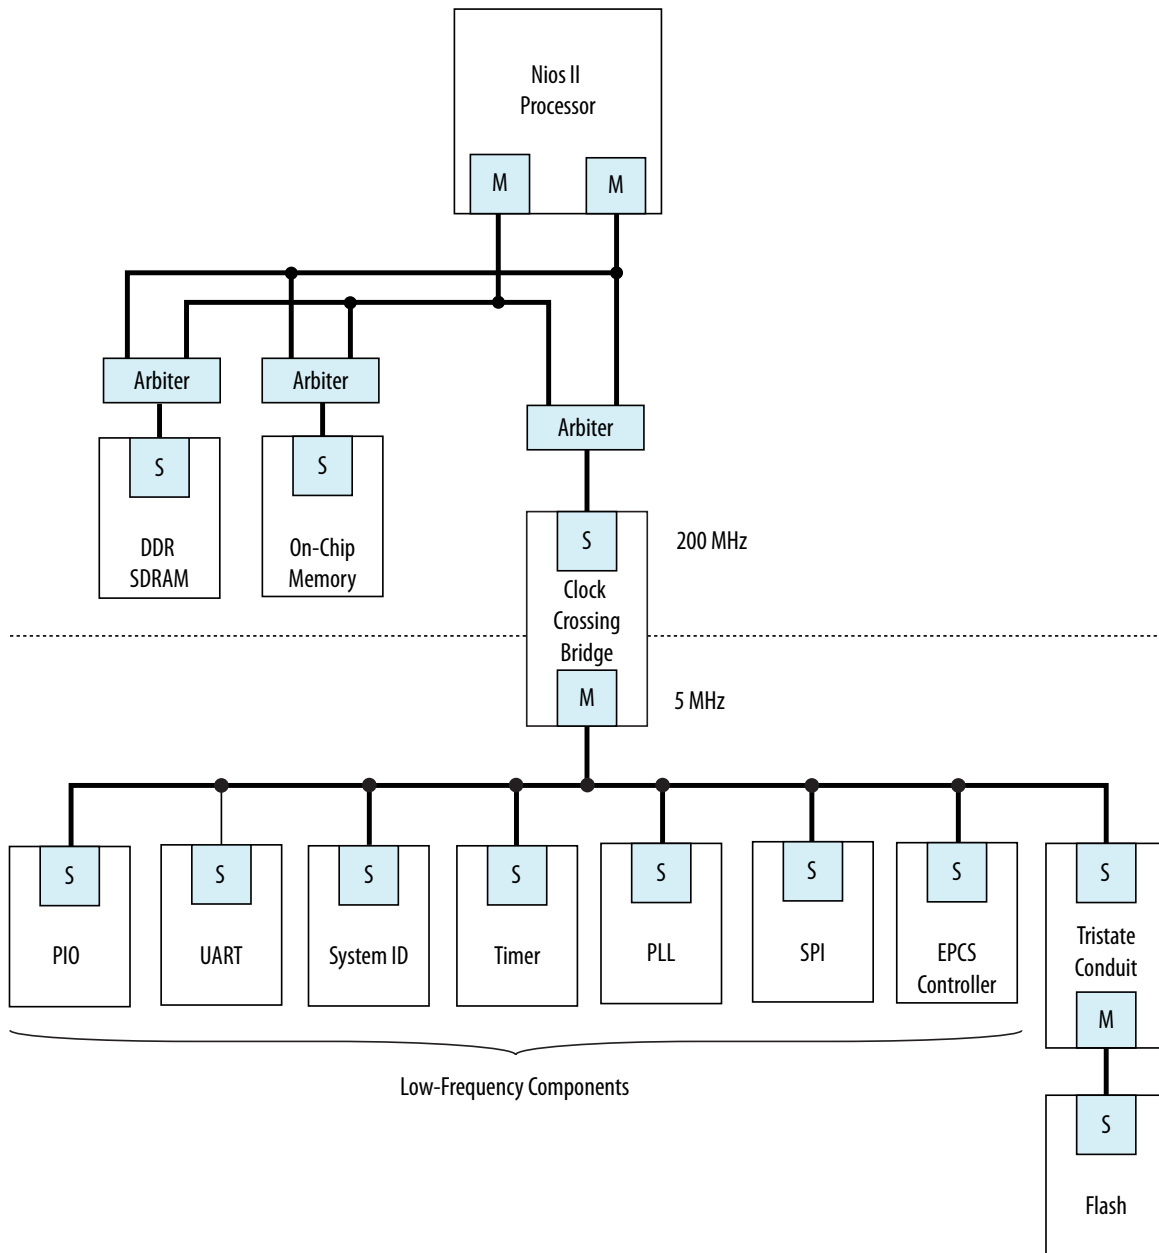
You can use clock crossing in Qsys to reduce the clock frequency of the logic that does not require a high frequency clock, which allows you to reduce power consumption. You can use either handshaking clock crossing bridges or handshaking clock crossing adapters to separate clock domains.

You can use the clock crossing bridge to connect master interfaces operating at a higher frequency to slave interfaces running at a lower frequency. Only connect low throughput or low priority components to a clock crossing bridge that operates at a reduced clock frequency. The following are examples of low throughput or low priority components:

- PIOs
- UARTs (JTAG or RS-232)
- System identification (SysID)
- Timers
- PLL (instantiated within Qsys)
- Serial peripheral interface (SPI)
- EPCS controller
- Tristate bridge and the components connected to the bridge

By reducing the clock frequency of the components connected to the bridge, you reduce the dynamic power consumption of the design. Dynamic power is a function of toggle rates and decreasing the clock frequency decreases the toggle rate.

Figure 8-24: Reducing Power Utilization Using a Bridge to Separate Clock Domains



Qsys automatically inserts clock crossing adapters between master and slave interfaces that operate at different clock frequencies. You can choose the type of clock crossing adapter in the Qsys **Project Settings**

tab. Adapters do not appear in the **Connections** column because you do not insert them. The following clock crossing adapter types are available in Qsys:

- **Handshake**—Uses a simple handshaking protocol to propagate transfer control signals and responses across the clock boundary. This adapter uses fewer hardware resources because each transfer is safely propagated to the target domain before the next transfer begins. The Handshake adapter is appropriate for systems with low throughput requirements.
- **FIFO**—Uses dual-clock FIFOs for synchronization. The latency of the FIFO adapter is approximately two clock cycles more than the handshake clock crossing component, but the FIFO-based adapter can sustain higher throughput because it supports multiple transactions simultaneously. The FIFO adapter requires more resources, and is appropriate for memory-mapped transfers requiring high throughput across clock domains.
- **Auto**—Qsys specifies the appropriate FIFO adapter for bursting links and the Handshake adapter for all other links.

Because the clock crossing bridge uses FIFOs to implement the clock crossing logic, it buffers transfers and data. Clock crossing adapters are not pipelined, so that each transaction is blocking until the transaction completes. Blocking transactions may lower the throughput substantially; consequently, if you want to reduce power consumption without limiting the throughput significantly, you should use the clock crossing bridge or the FIFO clock crossing adapter. However, if the design requires single read transfers, a clock crossing adapter is preferable because the latency is lower.

The clock crossing bridge requires few logic resources other than on-chip memory. The number of on-chip memory blocks used is proportional to the address span, data width, buffering depth, and bursting capabilities of the bridge. The clock crossing adapter does not use on-chip memory and requires a moderate number of logic resources. The address span, data width, and the bursting capabilities of the clock crossing adapter determine the resource utilization of the device.

When you decide to use a clock crossing bridge or clock crossing adapter, you must consider the effects of throughput and memory utilization in the design. If on-chip memory resources are limited, you may be forced to choose the clock crossing adapter. Using the clock crossing bridge to reduce the power of a single component may not justify using more resources. However, if you can place all of the low priority components behind a single clock crossing bridge, you may reduce power consumption in the design.

#### Related Information

#### [Power Optimization](#)

## Reducing Power Consumption by Minimizing Toggle Rates

A Qsys system consumes power whenever logic transitions between on and off states. When the state is held constant between clock edges, no charging or discharging occurs. You can use the following design methodologies to reduce the toggle rates of your design:

- Registering component boundaries
- Using clock enable signals
- Inserting bridges

Qsys interconnect is uniquely combinational when no adapters or bridges are present and there is no interconnect pipelining. When a slave interface is not selected by a master, various signals may toggle and propagate into the component. By registering the boundary of your component at the master or slave interface, you can minimize the toggling of the interconnect and your component. In addition, registering

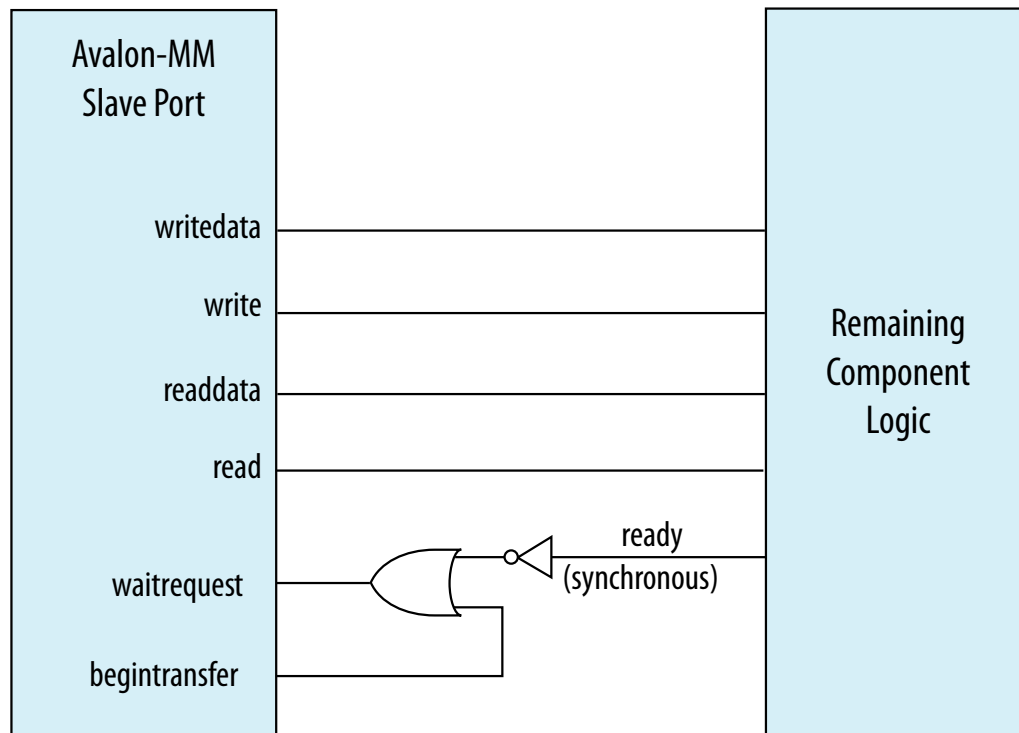
boundaries can improve operating frequency. When you register the signals at the interface level, you must ensure that the component continues to operate within the interface standard specification.

Avalon-MM `waitrequest` is a difficult signal to synchronize when you add registers to your component. The `waitrequest` signal must be asserted during the same clock cycle that a master asserts read or write to in order to prolong the transfer. A master interface can read the `waitrequest` signal too early and post more reads and writes prematurely.

**Note:** There is no direct AXI equivalent for `waitrequest` and `burstcount`, though the *AMBA Protocol Specification* implies that the AXI `ready` signal cannot depend combinatorially on the AXI `valid` signal. Therefore, Qsys typically buffers AXI component boundaries for the `ready` signal.

For slave interfaces, the interconnect manages the `begintransfer` signal, which is asserted during the first clock cycle of any read or write transfer. If the `waitrequest` is one clock cycle late, you can logically OR the `waitrequest` and the `begintransfer` signals to form a new `waitrequest` signal that is properly synchronized. Alternatively, the component can assert `waitrequest` before it is selected, guaranteeing that the `waitrequest` is already asserted during the first clock cycle of a transfer.

Figure 8-25: Variable Latency



### Using Clock Enables

You can use clock enables to hold the logic in a steady state, and the `write` and `read` signals as clock enables for slave components. Even if you add registers to your component boundaries, the interface can potentially toggle without the use of clock enables. You can also use the clock enable to disable combinational portions of the component.

For example, you can use an active high clock enable to mask the inputs into the combinational logic to prevent it from toggling when the component is inactive. Before preventing inactive logic from toggling,

you must determine if the masking causes the circuit to function differently. If masking causes a functional failure, it may be possible to use a register stage to hold the combinational logic constant between clock cycles.

### Inserting Bridges

You can use bridges to reduce toggle rates, if you do not want to modify the component by using boundary registers or clock enables. A bridge acts as a repeater where transfers to the slave interface are repeated on the master interface. If the bridge is not accessed, the components connected to its master interface are also not accessed. The master interface of the bridge remains idle until a master accesses the bridge slave interface.

Bridges can also reduce the toggle rates of signals that are inputs to other master interfaces. These signals are typically `readdata`, `readdatavalid`, and `waitrequest`. Slave interfaces that support read accesses drive the `readdata`, `readdatavalid`, and `waitrequest` signals. A bridge inserts either a register or clock crossing FIFO between the slave interface and the master to reduce the toggle rate of the master input signals.

#### Related Information

- [AMBA Protocol Specification](#)
- [Power Optimization](#)

## Reducing Power Consumption by Disabling Logic

There are typically two types of low power modes: volatile and non-volatile. A volatile low power mode holds the component in a reset state. When the logic is reactivated, the previous operational state is lost. A non-volatile low power mode restores the previous operational state. You can use either software-controlled or hardware-controlled sleep modes to disable a component in order to reduce power consumption.

### Software-Controlled Sleep Mode

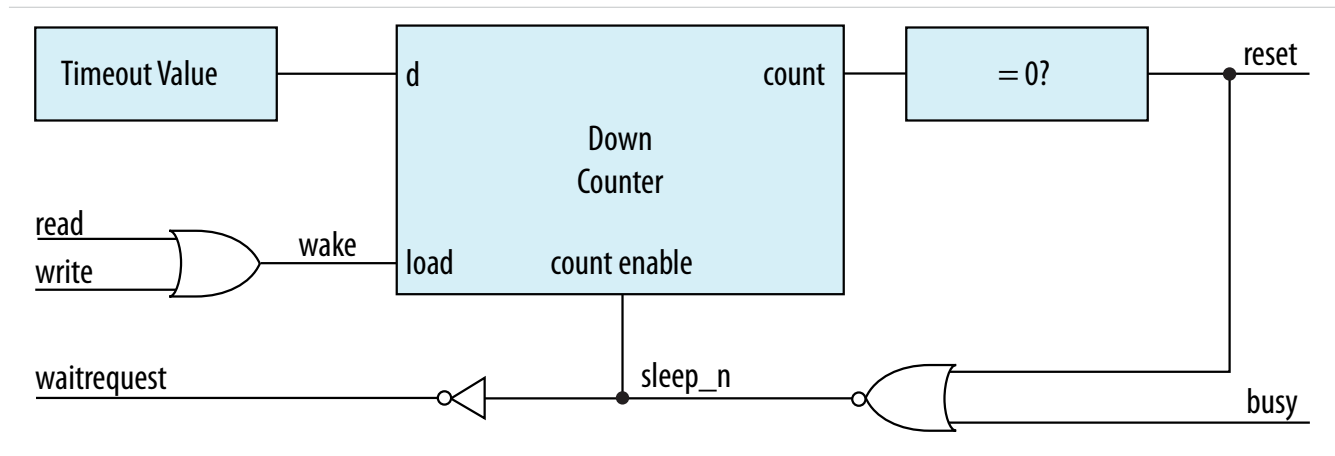
To design a component that supports software-controlled sleep mode, create a single memory-mapped location that enables and disables logic by writing a zero or one. You can use the register's output as a clock enable or reset, depending on whether the component has non-volatile requirements. The slave interface must remain active during sleep mode so that the enable bit is set when the component needs to be activated.

If multiple masters can access a component that supports sleep mode, you can use the mutex core to provide mutually exclusive accesses to your component. You can also build in the logic to re-enable the component on the very first access by any master in your system. If the component requires multiple clock cycles to re-activate, then it must assert a wait request to prolong the transfer as it exits sleep mode.

### Hardware-Controlled Sleep Mode

Alternatively, you can implement a timer in your component that automatically causes the component to enter a sleep mode based on a timeout value specified in clock cycles between read or write accesses. Each access resets the timer to the timeout value. Each cycle with no accesses decrements the timeout value by one. If the counter reaches zero, the hardware enters sleep mode until the next access.

Figure 8-26: Hardware-Controlled Sleep Components



This example provides a schematic for the hardware-controlled sleep mode. If restoring the component to an active state takes a long time, use a long timeout value so that the component is not continuously entering and exiting sleep mode. The slave interface must remain functional while the rest of the component is in sleep mode. When the component exits sleep mode, the component must assert the `waitrequest` signal until it is ready for read or write accesses.

**Related Information**

- [Mutex Core](#)
- [Power Optimization](#)

## Optimizing Qsys System Performance Design Examples

[Avalon Pipelined Read Master Example](#) on page 8-41

[Multiplexer Examples](#) on page 8-43

**Related Information**

[Avalon Interface Specifications](#)

### Avalon Pipelined Read Master Example

For a high throughput system using the Avalon-MM standard, you can design a pipelined read master that allows a system to issue multiple read requests before data returns. Pipelined read masters hide the latency of read operations by posting reads as frequently as every clock cycle. You can use this type of master when the address logic is not dependent on the data returning.

#### Avalon Pipelined Read Master Example Design Requirements

You must carefully design the logic for the control and data paths of pipelined read masters. The control logic must extend a read cycle whenever the `waitrequest` signal is asserted. This logic must also control the master `address`, `byteenable`, and `read` signals. To achieve maximum throughput, pipelined read masters should post reads continuously as long as `waitrequest` is de-asserted. While `read` is asserted, the address presented to the interconnect is stored.



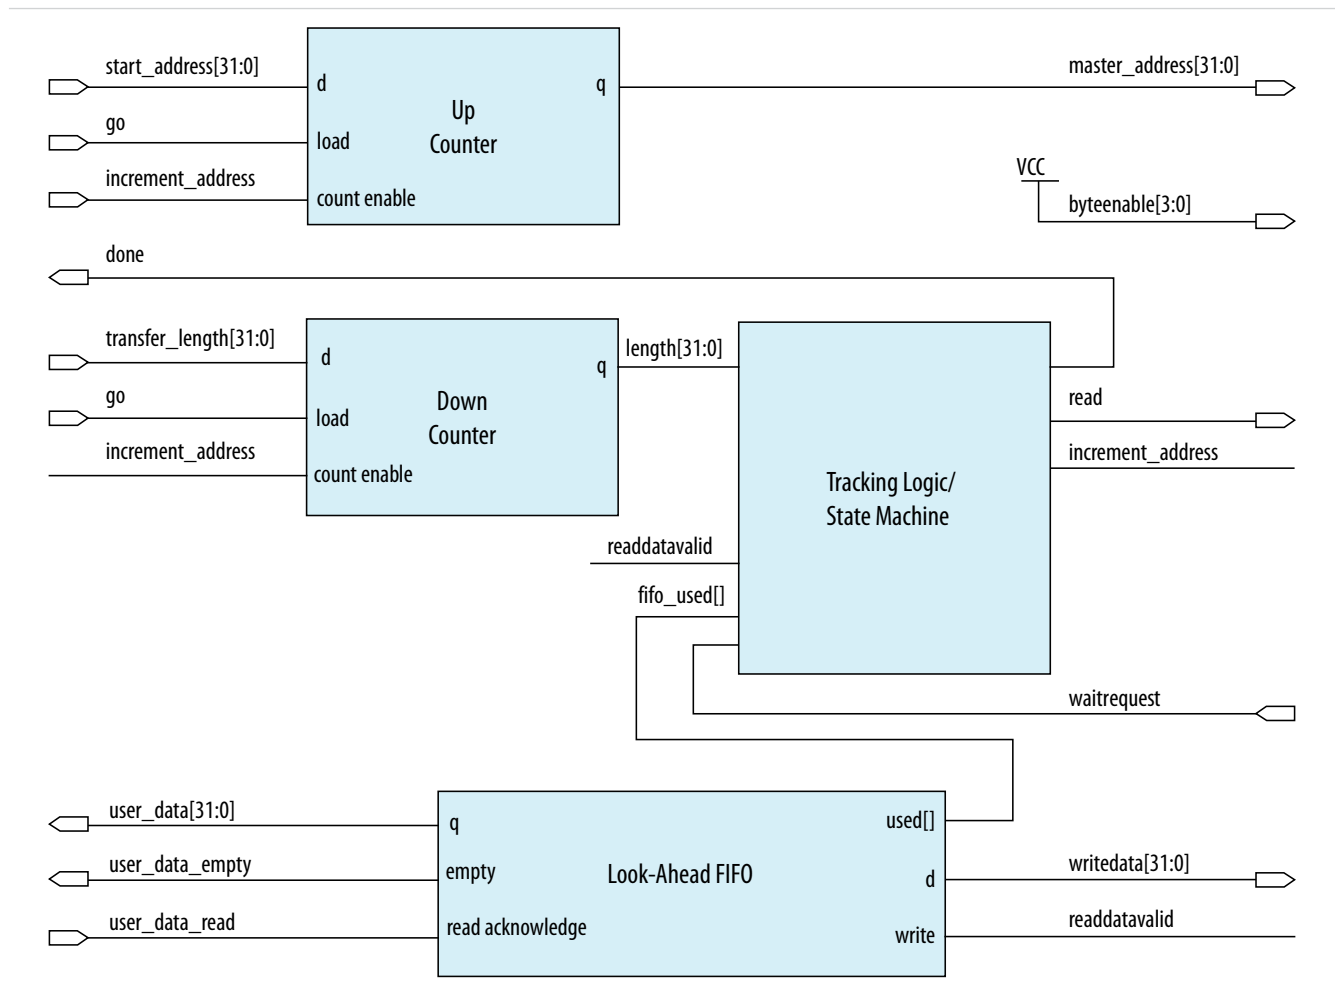
The data path logic includes the `readdata` and `readdatavalid` signals. If your master can accept data on every clock cycle, you can register the data with the `readdatavalid` as an enable bit. If your master cannot process a continuous stream of read data, it must buffer the data in a FIFO. The control logic must stop issuing reads when the FIFO reaches a predetermined fill level to prevent FIFO overflow.

## Expected Throughput Improvement

The throughput improvement that you can achieve with a pipelined read master is typically directly proportional to the pipeline depth of the interconnect and the slave interface. For example, if the total latency is two cycles, you can double the throughput by inserting a pipelined read master, assuming the slave interface also supports pipeline transfers. If either the master or slave does not support pipelined read transfers, then the interconnect asserts `waitrequest` until the transfer completes. You can also gain throughput when there are some cycles of overhead before a read response.

Where reads are not pipelined, the throughput is reduced. When both the master and slave interfaces support pipelined read transfers, data flows in a continuous stream after the initial latency. You can use a pipelined read master that stores data in a FIFO to implement a custom DMA, hardware accelerator, or off-chip communication interface.

**Figure 8-27: Pipelined Read Master**



This example shows a pipelined read master that stores data in a FIFO. The master performs word accesses that are word-aligned and reads from sequential memory addresses. The transfer length is a multiple of the word size.

When the `go` bit is asserted, the master registers the `start_address` and `transfer_length` signals. The master begins issuing reads continuously on the next clock cycle until the length register reaches zero. In this example, the word size is four bytes so that the address always increments by four, and the length decrements by four. The `read` signal remains asserted unless the FIFO fills to a predetermined level. The address register increments and the length register decrements if the length has not reached 0 and a read is posted.

The master posts a read transfer every time the `read` signal is asserted and the `waitrequest` is deasserted. The master issues reads until the entire buffer has been read or `waitrequest` is asserted. An optional tracking block monitors the done bit. When the length register reaches zero, some reads are outstanding. The tracking logic prevents assertion of done until the last read completes, and monitors the number of reads posted to the interconnect so that it does not exceed the space remaining in the `readdata` FIFO. This example includes a counter that verifies that the following conditions are met:

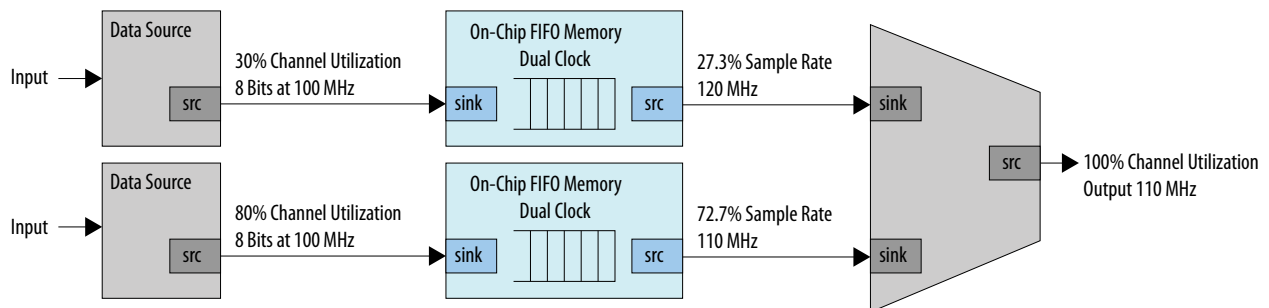
- If a read is posted and `readdatavalid` is deasserted, the counter increments.
- If a read is not posted and `readdatavalid` is asserted, the counter decrements.

When the `length` register and the tracking logic counter reach zero, all the reads have completed and the done bit is asserted. The `done` bit is important if a second master overwrites the memory locations that the pipelined read master accesses. This bit guarantees that the reads have completed before the original data is overwritten.

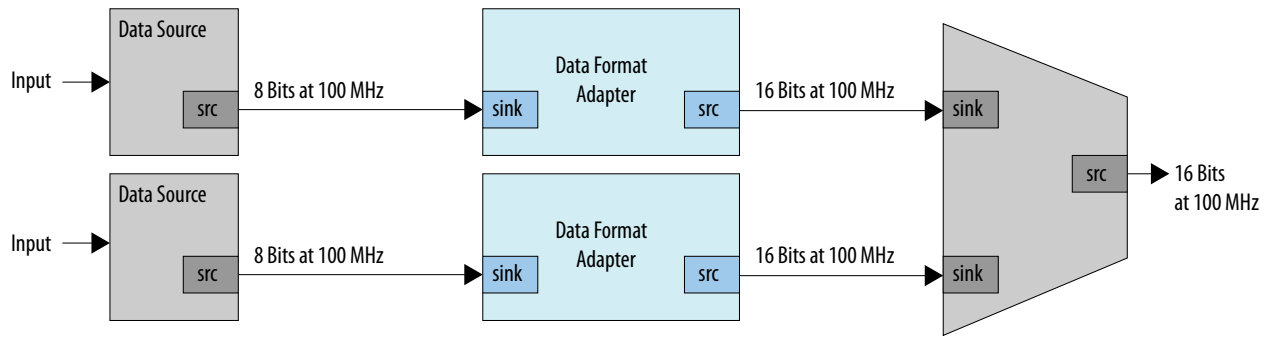
## Multiplexer Examples

You can combine adapters with streaming components to create data paths whose input and output streams have different properties. The following examples demonstrate datapaths in which the output stream exhibits higher performance than the input stream.

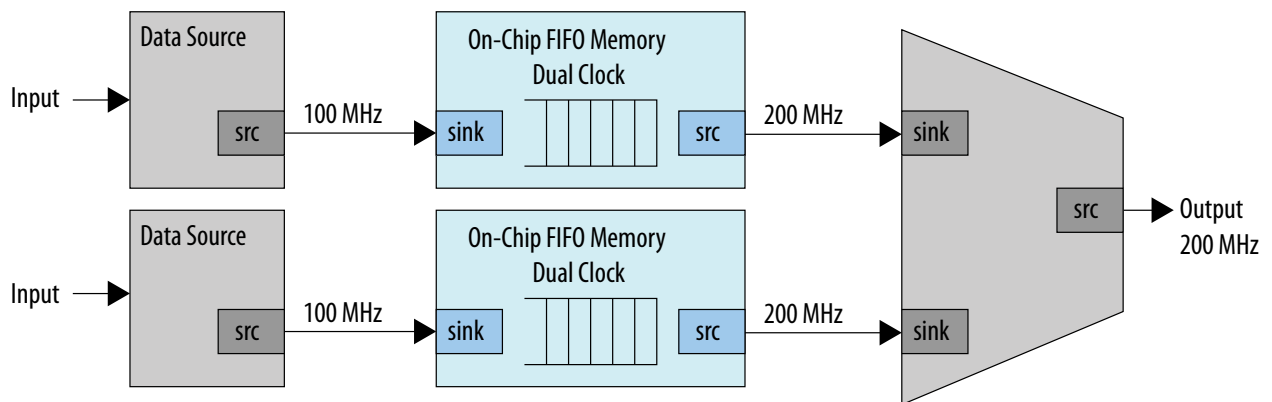
**Figure 8-28: Data Path that Doubles the Clock Frequency**



The diagram below illustrates a data path that uses the dual clock version of the on-chip FIFO memory and Avalon-ST channel multiplexer to merge the 100 MHz input from two streaming data sources into a single 200 MHz streaming output. This example shows an output with double the throughput of each interface with a corresponding doubling of the clock frequency.

**Figure 8-29: Data Path to Double Data Width and Maintain Original Frequency**

The diagram below illustrates a data path that uses the data format adapter and Avalon-ST channel multiplexer to convert two 8-bit inputs running at 100 MHz to a single 16-bit output at 100 MHz.

**Figure 8-30: Data Path to Boost the Clock Frequency**

The diagram below illustrates a data path that uses the dual clock version of the on-chip FIFO memory to boost the frequency of input data from 100 MHz to 110 MHz by sampling two input streams at differential rates. The on-chip FIFO memory has an input clock frequency of 100 MHz, and an output clock frequency of 110 MHz. The channel multiplexer runs at 110 MHz and samples one input stream 27.3 percent of the time, and the second 72.7 percent of the time. You do not need to know what the typical and maximum input channel utilizations are before for this type of design. For example, if the first channel hits 50% utilization, the output stream exceeds 100% utilization.

## Document Revision History

The table below indicates edits made to the *Optimizing Qsys System Performance* content since its creation.

**Table 8-2: Document Revision History**

Date	Version	Changes
May 2013	13.0.0	AMBA APB support.
November 2012	12.1.0	AMBA AXI4 support.
June 2012	12.0.0	AMBA AXI3 support.
November 2011	11.1.0	New document release.

**Related Information**

[Quartus II Handbook Archive](#)

2014.12.15

QI15V1



Subscribe



Send Feedback

Tcl commands allow you to perform a wide range of functions in Qsys. Command descriptions contain the Qsys phases where you can use the command, for example, main program, elaboration, composition, or fileset callback.

Qsys supports Avalon, AMBA AXI3 (version 1.0), AMBA AXI4 (version 2.0), AMBA AXI4-Lite (version 2.0), AMBA AXI4-Stream (version 1.0), and AMBA APB3 (version 1.0) interface specifications.

For more information about procedures for creating IP component **\_hw.tcl** files in the Qsys Component Editor, and supported interface standards, refer to *Creating Qsys Components* and *Qsys Interconnect* in volume 1 of the *Quartus II Handbook*.

If you are developing an IP component to work with the Nios II processor, refer to *Publishing Component Information to Embedded Software* in section 3 of the *Nios II Software Developer's Handbook*, which describes how to publish hardware IP component information for embedded software tools, such as a C compiler and a Board Support Package (BSP) generator.

#### Related Information

- [Avalon Interface Specifications](#)
- [AMBA Protocol Specifications](#)
- [Creating Qsys Components](#) on page 6-1
- [Qsys Interconnect](#) on page 7-1
- [Publishing Component Information to Embedded Software](#)

## Qsys \_hw.tcl Command Reference

To use the current version of the Tcl commands, include the following command at the top of your script:

```
package require -exact qsys <version>
```

## Interfaces and Ports

[add\\_interface](#) on page 9-3

[add\\_interface\\_port](#) on page 9-5

[get\\_interfaces](#) on page 9-7

[get\\_interface\\_assignment](#) on page 9-8

[get\\_interface\\_assignments](#) on page 9-9

[get\\_interface\\_ports](#) on page 9-10

[get\\_interface\\_properties](#) on page 9-11

[get\\_interface\\_property](#) on page 9-12

[get\\_port\\_properties](#) on page 9-13

[get\\_port\\_property](#) on page 9-14

[set\\_interface\\_assignment](#) on page 9-15

[set\\_interface\\_property](#) on page 9-17

[set\\_port\\_property](#) on page 9-18

[set\\_interface\\_upgrade\\_map](#) on page 9-19

## add\_interface

### Description

Adds an interface to your module. An interface represents a collection of related signals that are managed together in the parent system. These signals are implemented in the IP component's HDL, or exported from an interface from a child instance. As the IP component author, you choose the name of the interface.

### Availability

Discovery, Main Program, Elaboration, Composition

### Usage

```
add_interface <name> <type> <direction> [<associated_clock>]
```

### Returns

No returns value.

### Arguments

**name**

A name you choose to identify an interface.

**type**

The type of interface.

**direction**

The interface direction.

**associated\_clock (optional)**

(deprecated) For interfaces requiring associated clocks, use: `set_interface_property <interface> associatedClock <clockInterface>` For interfaces requiring associated resets, use: `set_interface_property <interface> associatedReset <resetInterface>`

### Example

```
add_interface mm_slave avalon slave

add_interface my_export conduit end
set_interface_property my_export EXPORT_OF uart_0.external_connection
```

### Notes

By default, interfaces are enabled. You can set the interface property `ENABLED` to `false` to disable an interface. If an interface is disabled, it is hidden and its ports are automatically terminated to their default values. Active high signals are terminated to 0. Active low signals are terminated to 1.

If the IP component is composed of child instances, the top-level interface is associated with a child instance's interface with `set_interface_property interface EXPORT_OF child_instance.interface`.

The following direction rules apply to Qsys-supported interfaces.

Interface Type	Direction
avalon	master, slave
axi	master, slave
tristate_conduit	master, slave
avalon_streaming	source, sink
interrupt	sender, receiver
conduit	end
clock	source, sink
reset	source, sink
nios_custom_instruction	slave

#### Related Information

- [add\\_interface\\_port](#) on page 9-5
- [get\\_interface\\_assignments](#) on page 9-9
- [get\\_interface\\_properties](#) on page 9-11
- [get\\_interfaces](#) on page 9-7



## add\_interface\_port

### Description

Adds a port to an interface on your module. The name must match the name of a signal on the top-level module in the HDL of your IP component. The port width and direction must be set before the end of the elaboration phase. You can set the port width as follows:

- In the Main program, you can set the port width to a fixed value or a width expression.
- If the port width is set to a fixed value in the Main program, you can update the width in the elaboration callback.

### Availability

Main Program, Elaboration

### Usage

```
add_interface_port <interface> <port> [<signal_type> <direction> <width_expression>]
```

### Returns

### Arguments

#### interface

The name of the interface to which this port belongs.

#### port

The name of the port. This name must match a signal in your top-level HDL for this IP component.

#### signal\_type (optional)

The type of signal for this port, which must be unique. Refer to the *Avalon Interface Specifications* for the signal types available for each interface type.

#### direction (optional)

The direction of the signal. Refer to *Direction Properties*.

#### width\_expression (optional)

The width of the port, in bits. The width may be a fixed value, or a simple arithmetic expression of parameter values.

### Example

#### fixed width:

```
add_interface_port mm_slave s0_rdata readdata output 32
```

#### width expression:

```
add_parameter DATA_WIDTH INTEGER 32  
add_interface_port s0 rdata readdata output "DATA_WIDTH/2"
```

### Related Information

- [add\\_interface](#) on page 9-3
- [get\\_port\\_properties](#) on page 9-13
- [get\\_port\\_property](#) on page 9-14

- [get\\_port\\_property](#) on page 9-14
- [Direction Properties](#) on page 9-99
- [Avalon Interface Specifications](#)

## get\_interfaces

### Description

Returns a list of top-level interfaces.

### Availability

Discovery, Main Program, Edit, Elaboration, Validation, Generation, Composition, Fileset Generation, Parameter Upgrade

### Usage

```
get_interfaces
```

### Returns

A list of the top-level interfaces exported from the system.

### Arguments

No arguments.

### Example

```
get_interfaces
```

### Related Information

[add\\_interface](#) on page 9-3

## `get_interface_assignment`

### Description

Returns the value of the specified assignment for the specified interface

### Availability

Main Program, Elaboration, Validation, Composition

### Usage

```
get_interface_assignment <interface> <assignment>
```

### Returns

The value of the assignment.

### Arguments

#### **interface**

The name of a top-level interface.

#### **assignment**

The name of an assignment.

### Example

```
get_interface_assignment s1 embeddedsw.configuration.isFlash
```

### Related Information

- [add\\_interface](#) on page 9-3
- [get\\_interface\\_assignments](#) on page 9-9
- [get\\_interfaces](#) on page 9-7

## get\_interface\_assignments

### Description

Returns the value of all interface assignments for the specified interface.

### Availability

Main Program, Elaboration, Validation, Composition

### Usage

```
get_interface_assignments <interface>
```

### Returns

A list of assignment keys.

### Arguments

#### interface

The name of the top-level interface whose assignment is being retrieved.

### Example

```
get_interface_assignments s1
```

### Related Information

- [add\\_interface](#) on page 9-3
- [get\\_interface\\_assignment](#) on page 9-8
- [get\\_interfaces](#) on page 9-7

## `get_interface_ports`

### Description

Returns the names of all of the ports that have been added to a given interface. If the interface name is omitted, all ports for all interfaces are returned.

### Availability

Discovery, Main Program, Edit, Elaboration, Validation, Generation, Composition, Fileset Generation, Parameter Upgrade

### Usage

```
get_interface_ports [<interface>]
```

### Returns

A list of port names.

### Arguments

**interface (optional)**

The name of a top-level interface.

### Example

```
get_interface_ports mm_slave
```

### Related Information

- [add\\_interface\\_port](#) on page 9-5
- [get\\_port\\_property](#) on page 9-14
- [set\\_port\\_property](#) on page 9-18

## get\_interface\_properties

### Description

Returns the names of all the interface properties for the specified interface as a space separated list

### Availability

Discovery, Main Program, Edit, Elaboration, Validation, Generation, Composition, Fileset Generation, Parameter Upgrade

### Usage

```
get_interface_properties <interface>
```

### Returns

A list of properties for the interface.

### Arguments

#### interface

The name of an interface.

### Example

```
get_interface_properties interface
```

### Notes

The properties for each interface type are different. Refer to the *Avalon Interface Specifications* for more information about interface properties.

#### Related Information

- [get\\_interface\\_property](#) on page 9-12
- [set\\_interface\\_property](#) on page 9-17
- [Avalon Interface Specifications](#)

## `get_interface_property`

### Description

Returns the value of a single interface property from the specified interface.

### Availability

Discovery, Main Program, Elaboration, Composition, Fileset Generation

### Usage

```
get_interface_property <interface> <property>
```

### Returns

### Arguments

**interface**

The name of an interface.

**property**

The name of the property whose value you want to retrieve. Refer to *Interface Properties*.

### Example

```
get_interface_property mm_slave linewidthBursts
```

### Notes

The properties for each interface type are different. Refer to the *Avalon Interface Specifications* for more information about interface properties.

#### Related Information

- [get\\_interface\\_properties](#) on page 9-11
- [set\\_interface\\_property](#) on page 9-17
- [Avalon Interface Specifications](#)



## get\_port\_properties

### Description

Returns a list of port properties.

### Availability

Discovery, Main Program, Edit, Elaboration, Validation, Generation, Composition, Fileset Generation, Parameter Upgrade

### Usage

```
get_port_properties
```

### Returns

A list of port properties. Refer to *Port Properties*.

### Arguments

No arguments.

### Example

```
get_port_properties
```

### Related Information

- [add\\_interface\\_port](#) on page 9-5
- [get\\_port\\_property](#) on page 9-14
- [set\\_port\\_property](#) on page 9-18
- [Port Properties](#) on page 9-97

## `get_port_property`

### Description

Returns the value of a property for the specified port.

### Availability

Discovery, Main Program, Edit, Elaboration, Validation, Generation, Composition, Fileset Generation, Parameter Upgrade

### Usage

```
get_port_property <port> <property>
```

### Returns

The value of the property.

### Arguments

#### **port**

The name of the port.

#### **property**

The name of a port property. Refer to *Port Properties*.

### Example

```
get_port_property rdata WIDTH_VALUE
```

### Related Information

- [add\\_interface\\_port](#) on page 9-5
- [get\\_port\\_properties](#) on page 9-13
- [set\\_port\\_property](#) on page 9-18
- [Port Properties](#) on page 9-97

## set\_interface\_assignment

### Description

Sets the value of the specified assignment for the specified interface.

### Availability

Main Program, Elaboration, Validation, Composition

### Usage

```
set_interface_assignment <interface> <assignment> [<value>]
```

### Returns

No return value.

### Arguments

#### interface

The name of the top-level interface whose assignment is being set.

#### assignment

The assignment whose value is being set.

#### value (optional)

The new assignment value.

### Example

```
set_interface_assignment s1 embeddedsw.configuration.isFlash 1
```

### Notes

#### Assignments for Nios II Software Build Tools

Interface assignments provide extra data for the Nios II Software Build Tools working with the generated system.

#### Assignments for Qsys Tools

There are several assignments that guide behavior in the Qsys tools.

<code>qsys.ui.export_name:</code>	If present, this interface should always be exported when an instance is added to a Qsys system. The value is the requested name of the exported interface in the parent system.
<code>qsys.ui.connect:</code>	If present, this interface should be auto-connected when an instance is added to a Qsys system. The value is a comma-separated list of other interfaces on the same instance that should be connected with this interface.
<code>ui.blockdiagram.direction:</code>	If present, the direction of this interface in the block diagram is set by the user. The value is either "output" or "input".

**Related Information**

- [add\\_interface](#) on page 9-3
- [get\\_interface\\_assignment](#) on page 9-8
- [get\\_interface\\_assignments](#) on page 9-9

## set\_interface\_property

### Description

Sets the value of a property on an exported top-level interface. You can use this command to set the `EXPORT_OF` property to specify which interface of a child instance is exported via this top-level interface.

### Availability

Main Program, Elaboration, Composition

### Usage

```
set_interface_property <interface> <property> <value>
```

### Returns

No return value.

### Arguments

#### interface

The name of an exported top-level interface.

#### property

The name of the property Refer to *Interface Properties*.

#### value

The new property value.

### Example

```
set_interface_property clk_out EXPORT_OF clk.clk_out  
set_interface_property mm_slave linewidthBursts false
```

### Notes

The properties for each interface type are different. Refer to the *Avalon Interface Specifications* for more information about interface properties.

#### Related Information

- [get\\_interface\\_properties](#) on page 9-11
- [get\\_interface\\_property](#) on page 9-12
- [Interface Properties](#) on page 9-90
- [Avalon Interface Specifications](#)

## set\_port\_property

### Description

Sets a port property.

### Availability

Main Program, Elaboration

### Usage

```
set_port_property <port> <property> [<value>]
```

### Returns

The new value.

### Arguments

**port**

The name of the port.

**property**

One of the supported properties. Refer to *Port Properties*.

**value (optional)**

The value to set.

### Example

```
set_port_property rdata WIDTH 32
```

### Related Information

- [add\\_interface\\_port](#) on page 9-5
- [get\\_port\\_properties](#) on page 9-13
- [set\\_port\\_property](#) on page 9-18

## set\_interface\_upgrade\_map

### Description

Maps the interface name of an older version of an IP core to the interface name for the current IP core. The interface type must be the same between the older and newer versions of the IP cores. This allows system connections and properties to maintain proper functionality. By default, if the older and newer versions of IP core have the same name and type, then Qsys maintains all properties and connections automatically.

### Availability

Parameter Upgrade

### Usage

```
set_interface_upgrade_map { <old_interface_name> <new_interface_name>  
<old_interface_name_2> <new_interface_name_2> ... }
```

### Returns

No return value.

### Arguments

```
{ <old_interface_name> <new_interface_name> }
```

List of mappings between between names of older and newer interfaces.

### Example

```
set_interface_upgrade_map { avalon_master_interface new_avalon_master_interface }
```

## Parameters

[add\\_parameter](#) on page 9-21

[get\\_parameters](#) on page 9-22

[get\\_parameter\\_properties](#) on page 9-23

[get\\_parameter\\_property](#) on page 9-24

[get\\_parameter\\_value](#) on page 9-25

[get\\_string](#) on page 9-26

[load\\_strings](#) on page 9-28

[set\\_parameter\\_property](#) on page 9-29

[set\\_parameter\\_value](#) on page 9-30

[decode\\_address\\_map](#) on page 9-31



## add\_parameter

### Description

Adds a parameter to your IP component.

### Availability

Main Program

### Usage

```
add_parameter <name> <type> [<default_value> <description>]
```

### Returns

### Arguments

**name**

The name of the parameter.

**type**

The data type of the parameter Refer to *Parameter Type Properties*.

**default\_value (optional)**

The initial value of the parameter in a new instance of the IP component.

**description (optional)**

Explains the use of the parameter.

### Example

```
add_parameter seed INTEGER 17 "The seed to use for data generation."
```

### Notes

Most parameter types have a single GUI element for editing the parameter value. `string_list` and `integer_list` parameters are different, because they are edited as tables. A multi-column table can be created by grouping multiple into a single table. To edit multiple list parameters in a single table, the display items for the parameters must be added to a group with a `TABLE` hint:

```
add_parameter coefficients INTEGER_LIST add_parameter positions INTEGER_LIST
add_display_item "" "Table Group" GROUP TABLE add_display_item "Table Group"
coefficients PARAMETER add_display_item "Table Group" positions PARAMETER
```

### Related Information

- [get\\_parameter\\_properties](#) on page 9-23
- [get\\_parameter\\_property](#) on page 9-24
- [get\\_parameter\\_value](#) on page 9-25
- [set\\_parameter\\_property](#) on page 9-29
- [set\\_parameter\\_value](#) on page 9-30
- [Parameter Type Properties](#) on page 9-95

## get\_parameters

### Description

Returns the names of all the parameters in the IP component.

### Availability

Discovery, Main Program, Edit, Elaboration, Validation, Generation, Composition, Fileset Generation, Parameter Upgrade

### Usage

```
get_parameters
```

### Returns

A list of parameter names

### Arguments

No arguments.

### Example

```
get_parameters
```

### Related Information

- [add\\_parameter](#) on page 9-21
- [get\\_parameter\\_property](#) on page 9-24
- [get\\_parameter\\_value](#) on page 9-25
- [get\\_parameters](#) on page 9-22
- [set\\_parameter\\_property](#) on page 9-29

## get\_parameter\_properties

### Description

Returns a list of all the parameter properties as a list of strings. The `get_parameter_property` and `set_parameter_property` commands are used to get and set the values of these properties, respectively.

### Availability

Discovery, Main Program, Edit, Elaboration, Validation, Generation, Composition, Fileset Generation, Parameter Upgrade

### Usage

```
get_parameter_properties
```

### Returns

A list of parameter property names. Refer to *Parameter Properties*.

### Arguments

No arguments.

### Example

```
set property_summary [ get_parameter_properties ]
```

### Related Information

- [add\\_parameter](#) on page 9-21
- [get\\_parameter\\_property](#) on page 9-24
- [get\\_parameter\\_value](#) on page 9-25
- [get\\_parameters](#) on page 9-22
- [set\\_parameter\\_property](#) on page 9-29
- [Parameter Properties](#) on page 9-92

## `get_parameter_property`

### Description

Returns the value of a property of a parameter.

### Availability

Discovery, Main Program, Edit, Elaboration, Validation, Generation, Composition, Fileset Generation, Parameter Upgrade

### Usage

```
get_parameter_property <parameter> <property>
```

### Returns

The value of the property.

### Arguments

#### **parameter**

The name of the parameter whose property value is being retrieved.

#### **property**

The name of the property. Refer to *Parameter Properties*.

### Example

```
set enabled [ get_parameter_property parameter1 ENABLED ]
```

### Related Information

- [add\\_parameter](#) on page 9-21
- [get\\_parameter\\_properties](#) on page 9-23
- [get\\_parameter\\_value](#) on page 9-25
- [get\\_parameters](#) on page 9-22
- [set\\_parameter\\_property](#) on page 9-29
- [set\\_parameter\\_value](#) on page 9-30
- [Parameter Properties](#) on page 9-92

## get\_parameter\_value

### Description

Returns the current value of a parameter defined previously with the `add_parameter` command.

### Availability

Discovery, Edit, Elaboration, Validation, Generation, Composition, Fileset Generation, Parameter Upgrade

### Usage

```
get_parameter_value <parameter>
```

### Returns

The value of the parameter.

### Arguments

#### **parameter**

The name of the parameter whose value is being retrieved.

### Example

```
set width [ get_parameter_value fifo_width ]
```

### Notes

If `AFFECTS_ELABORATION` is `false` for a given parameter, `get_parameter_value` is not available for that parameter from the elaboration callback. If `AFFECTS_GENERATION` is `false` then it is not available from the generation callback.

#### Related Information

- [add\\_parameter](#) on page 9-21
- [get\\_parameter\\_property](#) on page 9-24
- [get\\_parameters](#) on page 9-22
- [set\\_parameter\\_property](#) on page 9-29
- [set\\_parameter\\_value](#) on page 9-30

## get\_string

### Description

Returns the value of an externalized string previously loaded by the `load_strings` command.

### Availability

Discovery, Main Program, Edit, Elaboration, Validation, Generation, Composition, Fileset Generation, Parameter Upgrade

### Usage

```
get_string <identifier>
```

### Returns

The externalized string.

### Arguments

#### identifier

The string identifier.

### Example

```
hw.tcl:  
load_strings test.properties  
set_module_property NAME test  
set_module_property VERSION [get_string VERSION]  
set_module_property DISPLAY_NAME [get_string DISPLAY_NAME]  
add_parameter firepower INTEGER 0 ""  
set_parameter_property firepower DISPLAY_NAME [get_string PARAM_DISPLAY_NAME]  
set_parameter_property firepower TYPE INTEGER  
set_parameter_property firepower DESCRIPTION [get_string PARAM_DESCRIPTION]  
  
test.properties:  
DISPLAY_NAME = Trogdor!  
VERSION = 1.0  
PARAM_DISPLAY_NAME = Firepower  
PARAM_DESCRIPTION = The amount of force to use when breathing fire.
```

### Notes

Use uppercase words separated with underscores to name string identifiers. If you are externalizing module properties, use the module property name for the string identifier:

```
set_module_property DISPLAY_NAME [get_string DISPLAY_NAME]
```

If you are externalizing a parameter property, qualify the parameter property with the parameter name, with uppercase format, if needed:

```
set_parameter_property my_param DISPLAY_NAME [get_string MY_PARAM_DISPLAY_NAME]
```

If you use a string to describe a string format, end the identifier with `_FORMAT`.

```
set formatted_string [ format [ get_string TWO_ARGUMENT_MESSAGE_FORMAT ] "arg1"  
"arg2" ]
```

#### Related Information

[load\\_strings](#) on page 9-28

## load\_strings

### Description

Loads strings from an external `.properties` file.

### Availability

Discovery, Main Program

### Usage

```
load_strings <path>
```

### Returns

No return value.

### Arguments

#### path

The path to the properties file.

### Example

```
hw.tcl:  
load_strings test.properties  
set_module_property NAME test  
set_module_property VERSION [get_string VERSION]  
set_module_property DISPLAY_NAME [get_string DISPLAY_NAME]  
add_parameter firepower INTEGER 0 ""  
set_parameter_property firepower DISPLAY_NAME [get_string PARAM_DISPLAY_NAME]  
set_parameter_property firepower TYPE INTEGER  
set_parameter_property firepower DESCRIPTION [get_string PARAM_DESCRIPTION]  
  
test.properties:  
DISPLAY_NAME = Trogdor!  
VERSION = 1.0  
PARAM_DISPLAY_NAME = Firepower  
PARAM_DESCRIPTION = The amount of force to use when breathing fire.
```

### Notes

Refer to the *Java Properties File* for properties file format. A `.properties` file is a text file with `KEY=value` pairs. For externalized strings, the `KEY` is a string identifier and the `value` is the externalized string.

For example:

```
TROGDOR = A dragon with a big beefy arm
```

### Related Information

- [get\\_string](#) on page 9-26
- [Java Properties File](#)



## set\_parameter\_property

### Description

Sets a single parameter property.

### Availability

Main Program, Edit, Elaboration, Validation, Composition

### Usage

```
set_parameter_property <parameter> <property> <value>
```

### Returns

### Arguments

#### parameter

The name of the parameter that is being set.

#### property

The name of the property. Refer to *Parameter Properties*.

#### value

The new value for the property.

### Example

```
set_parameter_property BAUD_RATE ALLOWED_RANGES {9600 19200 38400}
```

### Related Information

- [add\\_parameter](#) on page 9-21
- [get\\_parameter\\_properties](#) on page 9-23
- [set\\_parameter\\_property](#) on page 9-29
- [Parameter Properties](#) on page 9-92

## set\_parameter\_value

### Description

Sets a parameter value. The value of a derived parameter can be updated by the IP component in the elaboration callback or the edit callback. Any changes to the value of a derived parameter in the edit callback will not be preserved.

### Availability

Edit, Elaboration, Validation, Composition, Parameter Upgrade

### Usage

```
set_parameter_value <parameter> <value>
```

### Returns

No return value.

### Arguments

**parameter**

The name of the parameter that is being set.

**value**

Specifies the new parameter value.

### Example

```
set_parameter_value half_clock_rate [ expr { [ get_parameter_value clock_rate ] /  
2 } ]
```



## decode\_address\_map

### Description

Converts an XML-formatted address map into a list of Tcl lists. Each inner list is in the correct format for conversion to an array. The XML code that describes each slave includes: its name, start address, and end address.

### Availability

Elaboration, Generation, Composition

### Usage

```
decode_address_map <address_map_XML_string>
```

### Returns

No return value.

### Arguments

#### **address\_mapXML\_string**

An XML string that describes the address map of a master.

### Example

In this example, the code describes the address map for the master that accesses the `ext_ssram`, `sys_clk_timer` and `sysid` slaves. The format of the string may differ from the example below; it may have different white space between the elements and include additional attributes or elements. Use the `decode_address_map` command to decode the code that represents a master's address map to ensure that your code works with future versions of the address map.

```
<address-map>
  <slave name='ext_ssram' start='0x01000000' end='0x01200000' />
  <slave name='sys_clk_timer' start='0x02120800' end='0x02120820' />
  <slave name='sysid' start='0x021208B8' end='0x021208C0' />
</address-map>
```

**Note:** Altera recommends that you use the code provided below to enumerate over the IP components within an address map, rather than writing your own parser.

```
set address_map_xml [get_parameter_value my_map_param]
set address_map_dec [decode_address_map $address_map_xml]
foreach i $address_map_dec {
  array set info $i
  send_message info "Connected to slave $info(name)"
}
```

## Display Items

[add\\_display\\_item](#) on page 9-33

[get\\_display\\_items](#) on page 9-35

[get\\_display\\_item\\_properties](#) on page 9-36

[get\\_display\\_item\\_property](#) on page 9-37

[set\\_display\\_item\\_property](#) on page 9-38

## add\_display\_item

### Description

Specifies the following aspects of the IP component display:

- Creates logical groups for a IP component's parameters. For example, to create separate groups for the IP component's timing, size, and simulation parameters. An IP component displays the groups and parameters in the order that you specify the display items in the `_hw.tcl` file.
- Groups a list of parameters to create multi-column tables.
- Specifies an image to provide representation of a parameter or parameter group.
- Creates a button by adding a display item of type `action`. The display item includes the name of the callback to run.

### Availability

Main Program

### Usage

```
add_display_item <parent_group> <id> <type> [<args>]
```

### Returns

### Arguments

**parent\_group**

Specifies the group to which a display item belongs

**id**

The identifier for the display item. If the item being added is a parameter, this is the parameter name. If the item is a group, this is the group name.

**type**

The type of the display item. Refer to *Display Item Kind Properties*.

**args (optional)**

Provides extra information required for display items.

### Example

```
add_display_item "Timing" read_latency PARAMETER
add_display_item "Sounds" speaker_image_id ICON speaker.jpg
```

## Notes

The following examples illustrate further illustrate the use of arguments:

- `add_display_item groupName id icon path-to-image-file`
- `add_display_item groupName parameterName parameter`
- `add_display_item groupName id text "your-text"`

The your-text argument is a block of text that is displayed in the GUI. Some simple HTML formatting is allowed, such as `<b>` and `<i>`, if the text starts with `<html>`.

- `add_display_item parentGroupName childGroupName group [tab]`

The tab is an optional parameter. If present, the group appears in separate tab in the GUI for the instance.

- `add_display_item parentGroupName actionName action buttonClickCallbackProc`

## Related Information

- [get\\_display\\_item\\_properties](#) on page 9-36
- [get\\_display\\_item\\_property](#) on page 9-37
- [get\\_display\\_items](#) on page 9-35
- [set\\_display\\_item\\_property](#) on page 9-38
- [Display Item Kind Properties](#) on page 9-101



## get\_display\_items

### Description

Returns a list of all items to be displayed as part of the parameterization GUI.

### Availability

Discovery, Main Program, Edit, Elaboration, Validation, Generation, Composition, Fileset Generation, Parameter Upgrade

### Usage

```
get_display_items
```

### Returns

List of display item IDs.

### Arguments

No arguments.

### Example

```
get_display_items
```

### Related Information

- [add\\_display\\_item](#) on page 9-33
- [get\\_display\\_item\\_properties](#) on page 9-36
- [get\\_display\\_item\\_property](#) on page 9-37
- [set\\_display\\_item\\_property](#) on page 9-38

## get\_display\_item\_properties

### Description

Returns a list of names of the properties of display items that are part of the parameterization GUI.

### Availability

Main Program

### Usage

```
get_display_item_properties
```

### Returns

A list of display item property names. Refer to *Display Item Properties*.

### Arguments

No arguments.

### Example

```
get_display_item_properties
```

### Related Information

- [add\\_display\\_item](#) on page 9-33
- [get\\_display\\_item\\_property](#) on page 9-37
- [set\\_display\\_item\\_property](#) on page 9-38
- [Display Item Properties](#) on page 9-100



## get\_display\_item\_property

### Description

Returns the value of a specific property of a display item that is part of the parameterization GUI.

### Availability

Main Program, Elaboration, Validation, Composition

### Usage

```
get_display_item_property <display_item> <property>
```

### Returns

The value of a display item property.

### Arguments

#### display\_item

The id of the display item.

#### property

The name of the property. Refer to *Display Item Properties*.

### Example

```
set my_label [get_display_item_property my_action DISPLAY_NAME]
```

### Related Information

- [add\\_display\\_item](#) on page 9-33
- [get\\_display\\_item\\_properties](#) on page 9-36
- [get\\_display\\_items](#) on page 9-35
- [set\\_display\\_item\\_property](#) on page 9-38
- [Display Item Properties](#) on page 9-100

## set\_display\_item\_property

### Description

Sets the value of specific property of a display item that is part of the parameterization GUI.

### Availability

Discovery, Main Program, Edit, Elaboration, Validation, Composition

### Usage

```
set_display_item_property <display_item> <property> <value>
```

### Returns

No return value.

### Arguments

#### display\_item

The name of the display item whose property value is being set.

#### property

The property that is being set. Refer to *Display Item Properties*.

#### value

The value to set.

### Example

```
set_display_item_property my_action DISPLAY_NAME "Click Me"  
set_display_item_property my_action DESCRIPTION "clicking this button runs the  
click_me_callback proc in the hw.tcl file"
```

### Related Information

- [add\\_display\\_item](#) on page 9-33
- [get\\_display\\_item\\_properties](#) on page 9-36
- [get\\_display\\_item\\_property](#) on page 9-37
- [Display Item Properties](#) on page 9-100



## Module Definition

[add\\_documentation\\_link](#) on page 9-40

[get\\_module\\_assignment](#) on page 9-41

[get\\_module\\_assignments](#) on page 9-42

[get\\_module\\_ports](#) on page 9-43

[get\\_module\\_properties](#) on page 9-44

[get\\_module\\_property](#) on page 9-45

[send\\_message](#) on page 9-46

[set\\_module\\_assignment](#) on page 9-47

[set\\_module\\_property](#) on page 9-48

[add\\_hdl\\_instance](#) on page 9-49

[package](#) on page 9-50

## add\_documentation\_link

### Description

Allows you to link to documentation for your IP component.

### Availability

Discovery, Main Program

### Usage

```
add_documentation_link <title> <path>
```

### Returns

No return value.

### Arguments

#### title

The title of the document for use on menus and buttons.

#### path

A path to the IP component documentation, using a syntax that provides the entire URL, not a relative path. For example: `http://www.mydomain.com/my_memory_controller.html` or `file:///datasheet.txt`

### Example

```
add_documentation_link "Avalon Verification IP Suite User Guide" http://  
www.altera.com/literature/ug/ug_avalon_verification_ip.pdf
```

## get\_module\_assignment

### Description

This command returns the value of an assignment. You can use the `get_module_assignment` and `set_module_assignment` and the `get_interface_assignment` and `set_interface_assignment` commands to provide information about the IP component to embedded software tools and applications.

### Availability

Main Program, Elaboration, Validation, Composition

### Usage

```
get_module_assignment <assignment>
```

### Returns

The value of the assignment

### Arguments

#### **assignment**

The name of the assignment whose value is being retrieved

### Example

```
get_module_assignment embeddedsw.CMacro.colorSpace
```

### Related Information

- [get\\_module\\_assignments](#) on page 9-42
- [set\\_module\\_assignment](#) on page 9-47

## `get_module_assignments`

### Description

Returns the names of the module assignments.

### Availability

Main Program, Elaboration, Validation, Composition

### Usage

```
get_module_assignments
```

### Returns

A list of assignment names.

### Arguments

No arguments.

### Example

```
get_module_assignments
```

### Related Information

- [get\\_module\\_assignment](#) on page 9-41
- [set\\_module\\_assignment](#) on page 9-47

## get\_module\_ports

### Description

Returns a list of the names of all the ports which are currently defined.

### Availability

Discovery, Main Program, Edit, Elaboration, Validation, Generation, Composition, Fileset Generation, Parameter Upgrade

### Usage

```
get_module_ports
```

### Returns

A list of port names.

### Arguments

No arguments.

### Example

```
get_module_ports
```

### Related Information

- [add\\_interface](#) on page 9-3
- [add\\_interface\\_port](#) on page 9-5

## get\_module\_properties

### Description

Returns the names of all the module properties as a list of strings. You can use the `get_module_property` and `set_module_property` commands to get and set values of individual properties. The value returned by this command is always the same for a particular version of Qsys

### Availability

Discovery, Main Program, Edit, Elaboration, Validation, Generation, Composition, Fileset Generation, Parameter Upgrade

### Usage

```
get_module_properties
```

### Returns

List of strings. Refer to *Module Properties*.

### Arguments

No arguments.

### Example

```
get_module_properties
```

### Related Information

- [get\\_module\\_property](#) on page 9-45
- [set\\_module\\_property](#) on page 9-48
- [Module Properties](#) on page 9-103





## get\_module\_property

### Description

Returns the value of a single module property.

### Availability

Discovery, Main Program, Edit, Elaboration, Validation, Generation, Composition, Fileset Generation, Parameter Upgrade

### Usage

```
get_module_property <property>
```

### Returns

Various.

### Arguments

#### **property**

The name of the property, Refer to *Module Properties*.

### Example

```
set my_name [ get_module_property NAME ]
```

### Related Information

- [get\\_module\\_properties](#) on page 9-44
- [set\\_module\\_property](#) on page 9-48
- [Module Properties](#) on page 9-103

## send\_message

### Description

Sends a message to the user of the IP component. The message text is normally interpreted as HTML. You can use the `<b>` element to provide emphasis. If you do not want the message text to be interpreted as HTML, then pass a list as the message level, for example, { `Info Text` }.

### Availability

Discovery, Main Program, Edit, Elaboration, Validation, Generation, Composition, Fileset Generation, Parameter Upgrade

### Usage

```
send_message <level> <message>
```

### Returns

No return value .

### Arguments

#### level

The following message levels are supported:

- **ERROR**--Provides an error message. The Qsys system cannot be generated with existing error messages.
- **WARNING**--Provides a warning message.
- **INFO**--Provides an informational message.
- **PROGRESS**--Reports progress during generation.
- **DEBUG**--Provides a debug message when debug mode is enabled.

#### message

The text of the message.

### Example

```
send_message ERROR "The system is down!"  
send_message { Info Text } "The system is up!"
```

## set\_module\_assignment

### Description

Sets the value of the specified assignment.

### Availability

Main Program, Elaboration, Validation, Composition

### Usage

```
set_module_assignment <assignment> [<value>]
```

### Returns

No return value.

### Arguments

#### assignment

The assignment whose value is being set

#### value (optional)

The value of the assignment

### Example

```
set_module_assignment embeddedsw.CMacro.colorSpace CMYK
```

### Related Information

- [get\\_module\\_assignment](#) on page 9-41
- [get\\_module\\_assignments](#) on page 9-42

## set\_module\_property

### Description

Allows you to set the values for module properties.

### Availability

Discovery, Main Program

### Usage

```
set_module_property <property> <value>
```

### Returns

No return value.

### Arguments

**property**

The name of the property. Refer to *Module Properties*.

**value**

The new value of the property.

### Example

```
set_module_property VERSION 10.0
```

### Related Information

- [get\\_module\\_properties](#) on page 9-44
- [get\\_module\\_property](#) on page 9-45
- [Module Properties](#) on page 9-103

## add\_hdl\_instance

### Description

Adds an instance of a predefined module, referred to as a *child* or *child instance*. The HDL entity generated from this instance can be instantiated and connected within this IP component's HDL.

### Availability

Main Program, Elaboration, Composition

### Usage

```
add_hdl_instance <entity_name> <ip_core_type> [<version>]
```

### Returns

The entity name of the added instance.

### Arguments

#### entity\_name

Specifies a unique local name that you can use to manipulate the instance. This name is used in the generated HDL to identify the instance.

#### ip\_core\_type

The type refers to a kind of instance available in the IP Catalog, for example `altera_avalon_uart`.

#### version (optional)

The required version of the specified instance type. If no version is specified, the latest version is used.

### Example

```
add_hdl_instance my_uart altera_avalon_uart
```

### Related Information

- [get\\_instance\\_parameter\\_value](#) on page 9-67
- [get\\_instance\\_parameters](#) on page 9-65
- [get\\_instances](#) on page 9-57
- [set\\_instance\\_parameter\\_value](#) on page 9-70

## package

### Description

Allows you to specify a particular version of the Qsys software to avoid software compatibility issues, and to determine which version of the `_hw.tcl` API to use for the IP component. You must use the `package` command at the beginning of your `_hw.tcl` file.

### Availability

Main Program

### Usage

```
package require -exact qsys <version>
```

### Returns

No return value

### Arguments

#### **version**

The version of Qsys that you require, such as 14.1.

### Example

```
package require -exact qsys 14.1
```

## Composition

[add\\_instance](#) on page 9-52  
[add\\_connection](#) on page 9-52  
[get\\_connections](#) on page 9-54  
[get\\_connection\\_parameters](#) on page 9-55  
[get\\_connection\\_parameter\\_value](#) on page 9-56  
[get\\_instances](#) on page 9-57  
[get\\_instance\\_interfaces](#) on page 9-58  
[get\\_instance\\_interface\\_ports](#) on page 9-59  
[get\\_instance\\_interface\\_properties](#) on page 9-60  
[get\\_instance\\_property](#) on page 9-61  
[set\\_instance\\_property](#) on page 9-62  
[get\\_instance\\_properties](#) on page 9-63  
[get\\_instance\\_interface\\_property](#) on page 9-64  
[get\\_instance\\_parameters](#) on page 9-65  
[get\\_instance\\_parameter\\_property](#) on page 9-66  
[get\\_instance\\_parameter\\_value](#) on page 9-67  
[get\\_instance\\_port\\_property](#) on page 9-68  
[set\\_connection\\_parameter\\_value](#) on page 9-69  
[set\\_instance\\_parameter\\_value](#) on page 9-70

## add\_instance

### Description

Adds an instance of an IP component, referred to as a child or child instance to the subsystem. You can use this command to create IP components that are composed of other IP component instances. The HDL for this subsystem will be generated; no custom HDL will need to be written for the IP component.

### Availability

Main Program, Composition

### Usage

```
add_instance <name> <type> [<version>]
```

### Returns

No return value.

### Arguments

#### name

Specifies a unique local name that you can use to manipulate the instance. This name is used in the generated HDL to identify the instance.

#### type

The type refers to a type available in the IP Catalog, for example `altera_avalon_uart`.

#### version (optional)

The required version of the specified type. If no version is specified, the highest available version is used.

### Example

```
add_instance my_uart altera_avalon_uart  
add_instance my_uart altera_avalon_uart 14.1
```

### Related Information

- [add\\_connection](#) on page 9-52
- [get\\_instance\\_interface\\_property](#) on page 9-64
- [get\\_instance\\_parameter\\_value](#) on page 9-67
- [get\\_instance\\_parameters](#) on page 9-65
- [get\\_instance\\_property](#) on page 9-61
- [get\\_instances](#) on page 9-57
- [set\\_instance\\_parameter\\_value](#) on page 9-70

## add\_connection

### Description

Connects the named interfaces on child instances together using an appropriate connection type. Both interface names consist of a child instance name, followed by the name of an interface provided by that



module. For example, `mux0.out` is the interface named `out` on the instance named `mux0`. Be careful to connect the start to the end, and not the other way around.

## Availability

Main Program, Composition

## Usage

```
add_connection <start> [<end> <kind> <name>]
```

## Returns

The name of the newly added connection in `start.point/end.point` format.

## Arguments

### start

The start interface to be connected, in `<instance_name>.<interface_name>` format.

### end (optional)

The end interface to be connected, `<instance_name>.<interface_name>`.

### kind (optional)

The type of connection, such as `avalon` or `clock`.

### name (optional)

A custom name for the connection. If unspecified, the name will be `<start_instance>.<interface>.<end_instance><interface>`

## Example

```
add_connection dma.read_master sdram.s1 avalon
```

## Related Information

- [add\\_instance](#) on page 9-52
- [get\\_instance\\_interfaces](#) on page 9-58

## `get_connections`

### Description

Returns a list of all connections in the composed subsystem.

### Availability

Main Program, Composition

### Usage

```
get_connections
```

### Returns

A list of connections.

### Arguments

No arguments.

### Example

```
set all_connections [ get_connections ]
```

### Related Information

[add\\_connection](#) on page 9-52

## get\_connection\_parameters

### Description

Returns a list of parameters found on a connection.

### Availability

Main Program, Composition

### Usage

```
get_connection_parameters <connection>
```

### Returns

A list of parameter names

### Arguments

#### **connection**

The connection to query.

### Example

```
get_connection_parameters cpu.data_master/dma0.csr
```

### Related Information

- [add\\_connection](#) on page 9-52
- [get\\_connection\\_parameter\\_value](#) on page 9-56

## get\_connection\_parameter\_value

### Description

Returns the value of a parameter on the connection. Parameters represent aspects of the connection that can be modified once the connection is created, such as the base address for an Avalon Memory Mapped connection.

### Availability

Composition

### Usage

```
get_connection_parameter_value <connection> <parameter>
```

### Returns

The value of the parameter.

### Arguments

**connection**

The connection to query.

**parameter**

The name of the parameter.

### Example

```
get_connection_parameter_value cpu.data_master/dma0.csr baseAddress
```

### Related Information

- [add\\_connection](#) on page 9-52
- [get\\_connection\\_parameters](#) on page 9-55

## get\_instances

### Description

Returns a list of the instance names for all child instances in the system.

### Availability

Main Program, Elaboration, Validation, Composition

### Usage

```
get_instances
```

### Returns

A list of child instance names.

### Arguments

No arguments.

### Example

```
get_instances
```

### Notes

This command can be used with instances created by either `add_instance` or `add_hdl_instance`.

#### Related Information

- [add\\_hdl\\_instance](#) on page 9-49
- [add\\_instance](#) on page 9-52
- [get\\_instance\\_parameter\\_value](#) on page 9-67
- [get\\_instance\\_parameters](#) on page 9-65
- [set\\_instance\\_parameter\\_value](#) on page 9-70

## get\_instance\_interfaces

### Description

Returns a list of interfaces found in a child instance. The list of interfaces can change if the parameterization of the instance changes.

### Availability

Validation, Composition

### Usage

```
get_instance_interfaces <instance>
```

### Returns

A list of interface names.

### Arguments

**instance**

The name of the child instance.

### Example

```
get_instance_interfaces pixel_converter
```

### Related Information

- [add\\_instance](#) on page 9-52
- [get\\_instance\\_interface\\_ports](#) on page 9-59
- [get\\_instance\\_interfaces](#) on page 9-58

## get\_instance\_interface\_ports

### Description

Returns a list of ports found in an interface of a child instance.

### Availability

Validation, Composition, Fileset Generation

### Usage

```
get_instance_interface_ports <instance> <interface>
```

### Returns

A list of port names found in the interface.

### Arguments

#### instance

The name of the child instance.

#### interface

The name of an interface on the child instance.

### Example

```
set port_names [ get_instance_interface_ports cpu data_master ]
```

### Related Information

- [add\\_instance](#) on page 9-52
- [get\\_instance\\_interfaces](#) on page 9-58
- [get\\_instance\\_port\\_property](#) on page 9-68

## `get_instance_interface_properties`

### Description

Returns the names of all of the properties of the specified interface

### Availability

Validation, Composition

### Usage

```
get_instance_interface_properties <instance> <interface>
```

### Returns

List of property names.

### Arguments

#### **instance**

The name of the child instance.

#### **interface**

The name of an interface on the instance.

### Example

```
set properties [ get_instance_interface_properties cpu data_master ]
```

### Related Information

- [add\\_instance](#) on page 9-52
- [get\\_instance\\_interface\\_property](#) on page 9-64
- [get\\_instance\\_interfaces](#) on page 9-58



## get\_instance\_property

### Description

Returns the value of a single instance property.

### Availability

Main Program, Elaboration, Validation, Composition, Fileset Generation

### Usage

```
get_instance_property <instance> <property>
```

### Returns

Various.

### Arguments

#### **instance**

The name of the instance.

#### **property**

The name of the property. Refer to *Instance Properties*.

### Example

```
set my_name [ get_instance_property myinstance NAME ]
```

### Related Information

- [add\\_instance](#) on page 9-52
- [get\\_instance\\_properties](#) on page 9-63
- [set\\_instance\\_property](#) on page 9-62
- [Instance Properties](#) on page 9-91

## set\_instance\_property

### Description

Allows a user to set the properties of a child instance.

### Availability

Main Program, Elaboration, Validation, Composition

### Usage

```
set_instance_property <instance> <property> <value>
```

### Returns

### Arguments

**instance**

The name of the instance.

**property**

The name of the property to set. Refer to *Instance Properties*.

**value**

The new property value.

### Example

```
set_instance_property myinstance SUPPRESS_ALL_WARNINGS true
```

### Related Information

- [add\\_instance](#) on page 9-52
- [get\\_instance\\_properties](#) on page 9-63
- [get\\_instance\\_property](#) on page 9-61
- [Instance Properties](#) on page 9-91

## get\_instance\_properties

### Description

Returns the names of all the instance properties as a list of strings. You can use the `get_instance_property` and `set_instance_property` commands to get and set values of individual properties. The value returned by this command is always the same for a particular version of Qsys

### Availability

Discovery, Main Program, Edit, Elaboration, Validation, Generation, Composition, Fileset Generation, Parameter Upgrade

### Usage

```
get_instance_properties
```

### Returns

List of strings. Refer to *Instance Properties*.

### Arguments

No arguments.

### Example

```
get_instance_properties
```

### Related Information

- [add\\_instance](#) on page 9-52
- [get\\_instance\\_property](#) on page 9-61
- [set\\_instance\\_property](#) on page 9-62
- [Instance Properties](#) on page 9-91

## get\_instance\_interface\_property

### Description

Returns the value of a property for an interface in a child instance.

### Availability

Validation, Composition

### Usage

```
get_instance_interface_property <instance> <interface> <property>
```

### Returns

The value of the property.

### Arguments

**instance**

The name of the child instance.

**interface**

The name of an interface on the child instance.

**property**

The name of the property of the interface.

### Example

```
set value [ get_instance_interface_property cpu data_master setupTime ]
```

### Related Information

- [add\\_instance](#) on page 9-52
- [get\\_instance\\_interfaces](#) on page 9-58



## get\_instance\_parameters

### Description

Returns a list of names of the parameters on a child instance that can be set using `set_instance_parameter_value`. It omits parameters that are derived and those that have the `SYSTEM_INFO` parameter property set.

### Availability

Main Program, Elaboration, Validation, Composition

### Usage

```
get_instance_parameters <instance>
```

### Returns

A list of parameters in the instance.

### Arguments

**instance**

The name of the child instance.

### Example

```
set parameters [ get_instance_parameters instance ]
```

### Notes

You can use this command with instances created by either `add_instance` or `add_hdl_instance`.

#### Related Information

- [add\\_hdl\\_instance](#) on page 9-49
- [add\\_instance](#) on page 9-52
- [get\\_instance\\_parameter\\_value](#) on page 9-67
- [get\\_instances](#) on page 9-57
- [set\\_instance\\_parameter\\_value](#) on page 9-70

## get\_instance\_parameter\_property

### Description

Returns the value of a property on a parameter in a child instance. Parameter properties are metadata about how the parameter will be used by the Qsys tools.

### Availability

Validation, Composition

### Usage

```
get_instance_parameter_property <instance> <parameter> <property>
```

### Returns

The value of the parameter property.

### Arguments

**instance**

The name of the child instance.

**parameter**

The name of the parameter in the instance.

**property**

The name of the property of the parameter. Refer to *Parameter Properties*.

### Example

```
get_instance_parameter_property instance parameter property
```

### Related Information

- [add\\_instance](#) on page 9-52
- [Parameter Properties](#) on page 9-92

## get\_instance\_parameter\_value

### Description

Returns the value of a parameter in a child instance. You cannot use this command to get the value of parameters whose values are derived or those that are defined using the `SYSTEM_INFO` parameter property.

### Availability

Elaboration, Validation, Composition

### Usage

```
get_instance_parameter_value <instance> <parameter>
```

### Returns

The value of the parameter.

### Arguments

**instance**

The name of the child instance.

**parameter**

Specifies the parameter whose value is being retrieved.

### Example

```
set dpi [ get_instance_parameter_value pixel_converter input_DPI ]
```

### Notes

You can use this command with instances created by either `add_instance` or `add_hdl_instance`.

#### Related Information

- [add\\_hdl\\_instance](#) on page 9-49
- [add\\_instance](#) on page 9-52
- [get\\_instance\\_parameters](#) on page 9-65
- [get\\_instances](#) on page 9-57
- [set\\_instance\\_parameter\\_value](#) on page 9-70

## get\_instance\_port\_property

### Description

Returns the value of a property of a port contained by an interface in a child instance.

### Availability

Validation, Composition, Fileset Generation

### Usage

```
get_instance_port_property <instance> <port> <property>
```

### Returns

The value of the property for the port.

### Arguments

#### instance

The name of the child instance.

#### port

The name of a port in one of the interfaces on the child instance.

#### property

The property whose value is being retrieved. Only the following port properties can be queried on ports of child instances: `ROLE`, `DIRECTION`, `WIDTH`, `WIDTH_EXPR` and `VHDL_TYPE`. Refer to *Port Properties*.

### Example

```
get_instance_port_property instance port property
```

### Related Information

- [add\\_instance](#) on page 9-52
- [get\\_instance\\_interface\\_ports](#) on page 9-59
- [Port Properties](#) on page 9-97



## set\_connection\_parameter\_value

### Description

Sets the value of a parameter of the connection. The start and end are each interface names of the format `<instance>.<interface>`. Connection parameters depend on the type of connection, for Avalon-MM they include base addresses and arbitration priorities.

### Availability

Main Program, Composition

### Usage

```
set_connection_parameter_value <connection> <parameter> <value>
```

### Returns

No return value.

### Arguments

#### connection

Specifies the name of the connection as returned by the `add_connectioncommand`. It is of the form `start.point/end.point`.

#### parameter

The name of the parameter.

#### value

The new parameter value.

### Example

```
set_connection_parameter_value cpu.data_master/dma0.csr baseAddress "0x000a0000"
```

### Related Information

- [add\\_connection](#) on page 9-52
- [get\\_connection\\_parameter\\_value](#) on page 9-56

## set\_instance\_parameter\_value

### Description

Sets the value of a parameter for a child instance. Derived parameters and `SYSTEM_INFO` parameters for the child instance can not be set with this command.

### Availability

Main Program, Elaboration, Composition

### Usage

```
set_instance_parameter_value <instance> <parameter> <value>
```

### Returns

Vo return value.

### Arguments

**instance**

Specifies the name of the child instance.

**parameter**

Specifies the parameter that is being set.

**value**

Specifies the new parameter value.

### Example

```
set_instance_parameter_value uart_0 baudRate 9600
```

### Notes

You can use this command with instances created by either `add_instance` or `add_hdl_instance`.

#### Related Information

- [add\\_hdl\\_instance](#) on page 9-49
- [add\\_instance](#) on page 9-52
- [get\\_instance\\_parameter\\_value](#) on page 9-67
- [get\\_instances](#) on page 9-57



## Fileset Generation

[add\\_fileset](#) on page 9-72

[add\\_fileset\\_file](#) on page 9-73

[set\\_fileset\\_property](#) on page 9-74

[get\\_fileset\\_file\\_attribute](#) on page 9-75

[set\\_fileset\\_file\\_attribute](#) on page 9-76

[get\\_fileset\\_properties](#) on page 9-77

[get\\_fileset\\_property](#) on page 9-78

[get\\_fileset\\_sim\\_properties](#) on page 9-79

[set\\_fileset\\_sim\\_properties](#) on page 9-80

[create\\_temp\\_file](#) on page 9-81

## add\_fileset

### Description

Adds a generation fileset for a particular target as specified by the `kind`. Qsys calls the target (`SIM_VHDL`, `SIM_VERILOG`, `QUARTUS_SYNTH`, or `EXAMPLE_DESIGN`) when the specified generation target is requested. You can define multiple filesets for each kind of fileset. Qsys passes a single argument to the specified callback procedure. The value of the argument is a generated name, which you must use in the top-level module or entity declaration of your IP component. To override this generated name, you can set the fileset property `TOP_LEVEL`.

### Availability

Main Program

### Usage

```
add_fileset <name> <kind> [<callback_proc> <display_name>]
```

### Returns

No return value.

### Arguments

**name**

The name of the fileset.

**kind**

The kind of fileset. Refer to *Fileset Properties*.

**callback\_proc (optional)**

A string identifying the name of the callback procedure. If you add files in the global section, you can then specify a blank callback procedure.

**display\_name (optional)**

A display string to identify the fileset.

### Example

```
add_fileset my_synthesis_fileset QUARTUS_SYNTH mySynthCallbackProc "My Synthesis"  
proc mySynthCallbackProc { topLevelName } { ... }
```

### Notes

If using the `TOP_LEVEL` fileset property, all parameterizations of the component must use identical HDL.

#### Related Information

- [add\\_fileset\\_file](#) on page 9-73
- [get\\_fileset\\_property](#) on page 9-78
- [Fileset Properties](#) on page 9-105

## add\_fileset\_file

### Description

Adds a file to the generation directory. You can specify source file locations using either an absolute path, or a path that is relative to the IP component's `_hw.tcl` file.

### Availability

Main Program, Fileset Generation

### Usage

```
add_fileset_file <output_file> <file_type> <file_source> <path_or_contents> [<attributes>]
```

### Returns

No return value.

### Arguments

#### output\_file

Specifies the location to store the file after Qsys generation

#### file\_type

The kind of file. Refer to *File Kind Properties*.

#### file\_source

Specifies whether the file is being added by path, or by file contents. Refer to *File Source Properties*.

#### path\_or\_contents

When the `file_source` is `PATH`, specifies the file to be copied to `output_file`. When the `file_source` is `TEXT`, specifies the text contents to be stored in the file.

#### attributes (optional)

An optional list of file attributes. Typically used to specify that a file is intended for use only in a particular simulator. Refer to *File Attribute Properties*.

### Example

```
add_fileset_file "./implementation/rx_pma.sv" SYSTEM_VERILOG PATH synth_rx_pma.sv
add_fileset_file gui.sv SYSTEM_VERILOG TEXT "Customize your IP core"
```

### Related Information

- [add\\_fileset](#) on page 9-72
- [get\\_fileset\\_file\\_attribute](#) on page 9-75
- [File Kind Properties](#) on page 9-109
- [File Source Properties](#) on page 9-110
- [File Attribute Properties](#) on page 9-108

## set\_fileset\_property

### Description

Allows you to set the properties of a fileset.

### Availability

Main Program, Elaboration, Fileset Generation

### Usage

```
set_fileset_property <fileset> <property> <value>
```

### Returns

No return value.

### Arguments

**fileset**

The name of the fileset.

**property**

The name of the property to set. Refer to *Fileset Properties*.

**value**

The new property value.

### Example

```
set_fileset_property mySynthFileset TOP_LEVEL simple_uart
```

### Notes

When a fileset callback is called, the callback procedure will be passed a single argument. The value of this argument is a generated name which must be used in the top-level module or entity declaration of your IP component. If set, the `TOP_LEVEL` specifies a fixed name for the top-level name of your IP component.

The `TOP_LEVEL` property must be set in the global section. It cannot be set in a fileset callback.

If using the `TOP_LEVEL` fileset property, all parameterizations of the IP component must use identical HDL.

#### Related Information

- [add\\_fileset](#) on page 9-72
- [Fileset Properties](#) on page 9-105

## get\_fileset\_file\_attribute

### Description

Returns the attribute of a fileset file.

### Availability

Main Program, Fileset Generation

### Usage

```
get_fileset_file_attribute <output_file> <attribute>
```

### Returns

Value of the fileset File attribute.

### Arguments

#### **output\_file**

Location of the output file.

#### **attribute**

Specifies the name of the attribute Refer to *File Attribute Properties*.

### Example

```
get_fileset_file_attribute my_file.sv ALDEC_SPECIFIC
```

### Related Information

- [add\\_fileset](#) on page 9-72
- [add\\_fileset\\_file](#) on page 9-73
- [get\\_fileset\\_file\\_attribute](#) on page 9-75
- [File Attribute Properties](#) on page 9-108
- [add\\_fileset](#) on page 9-72
- [add\\_fileset\\_file](#) on page 9-73
- [get\\_fileset\\_file\\_attribute](#) on page 9-75
- [File Attribute Properties](#) on page 9-108

## set\_fileset\_file\_attribute

### Description

Sets the attribute of a fileset file.

### Availability

Main Program, Fileset Generation

### Usage

```
set_fileset_file_attribute <output_file> <attribute> <value>
```

### Returns

The attribute value if it was set.

### Arguments

**output\_file**

Location of the output file.

**attribute**

Specifies the name of the attribute Refer to *File Attribute Properties*.

**value**

Value to set the attribute to.

### Example

```
set_fileset_file_attribute my_file_pkg.sv COMMON_SYSTEMVERILOG_PACKAGE  
my_file_package
```



## get\_fileset\_properties

### Description

Returns a list of properties that can be set on a fileset.

### Availability

Discovery, Main Program, Edit, Elaboration, Validation, Generation, Composition, Fileset Generation, Parameter Upgrade

### Usage

```
get_fileset_properties
```

### Returns

A list of property names. Refer to *Fileset Properties*.

### Arguments

No arguments.

### Example

```
get_fileset_properties
```

### Related Information

- [add\\_fileset](#) on page 9-72
- [get\\_fileset\\_properties](#) on page 9-77
- [set\\_fileset\\_property](#) on page 9-74
- [Fileset Properties](#) on page 9-105

## `get_fileset_property`

### Description

Returns the value of a fileset property for a fileset.

### Availability

Main Program, Elaboration, Fileset Generation

### Usage

```
get_fileset_property <fileset> <property>
```

### Returns

The value of the property.

### Arguments

**fileset**

The name of the fileset.

**property**

The name of the property to query. Refer to *Fileset Properties*.

### Example

```
get_fileset_property fileset property
```

### Related Information

[Fileset Properties](#) on page 9-105

## get\_fileset\_sim\_properties

### Description

Returns simulator properties for a fileset.

### Availability

Main Program, Fileset Generation

### Usage

```
get_fileset_sim_properties <fileset> <platform> <property>
```

### Returns

The fileset simulator properties.

### Arguments

#### fileset

The name of the fileset.

#### platform

The operating system for that applies to the property. Refer to *Operating System Properties*.

#### property

Specifies the name of the property to set. Refer to *Simulator Properties*.

### Example

```
get_fileset_sim_properties my_fileset LINUX64 OPT_CADENCE_64BIT
```

### Related Information

- [add\\_fileset](#) on page 9-72
- [set\\_fileset\\_sim\\_properties](#) on page 9-80
- [Operating System Properties](#) on page 9-117
- [Simulator Properties](#) on page 9-111

## set\_fileset\_sim\_properties

### Description

Sets simulator properties for a given fileset

### Availability

Main Program, Fileset Generation

### Usage

```
set_fileset_sim_properties <fileset> <platform> <property> <value>
```

### Returns

The fileset simulator properties if they were set.

### Arguments

#### fileset

The name of the fileset.

#### platform

The operating system that applies to the property. Refer to *Operating System Properties*.

#### property

Specifies the name of the property to set. Refer to *Simulator Properties*.

#### value

Specifies the value of the property.

### Example

```
set_fileset_sim_properties my_fileset LINUX64 OPT_MENTOR_PLI "{libA} {libB}"
```

### Related Information

- [get\\_fileset\\_sim\\_properties](#) on page 9-79
- [Operating System Properties](#) on page 9-117
- [Simulator Properties](#) on page 9-111



## create\_temp\_file

### Description

Creates a temporary file, which you can use inside the fileset callbacks of a `_hw.tcl` file. This temporary file is included in the generation output if it is added using the `add_fileset_file` command.

### Availability

Fileset Generation

### Usage

```
create_temp_file <path>
```

### Returns

The path to the temporary file.

### Arguments

#### path

The name of the temporary file.

### Example

```
set filelocation [create_temp_file "./hdl/compute_frequency.v" ]  
add_fileset_file compute_frequency.v VERILOG PATH ${filelocation}
```

### Related Information

- [add\\_fileset](#) on page 9-72
- [add\\_fileset\\_file](#) on page 9-73

## Miscellaneous

[check\\_device\\_family\\_equivalence](#) on page 9-83

[get\\_device\\_family\\_displayname](#) on page 9-84

[get\\_qip\\_strings](#) on page 9-85

[set\\_qip\\_strings](#) on page 9-86

[set\\_interconnect\\_requirement](#) on page 9-87

## check\_device\_family\_equivalence

### Description

Returns 1 if the device family is equivalent to one of the families in the device families list, Returns 0 if the device family is not equivalent to any families. This command ignores differences in capitalization and spaces.

### Availability

Discovery, Main Program, Edit, Elaboration, Validation, Composition, Fileset Generation, Parameter Upgrade

### Usage

```
check_device_family_equivalence <device_family> <device_family_list>
```

### Returns

1 if equivalent, 0 if not equivalent.

### Arguments

#### **device\_family**

The device family name that is being checked.

#### **device\_family\_list**

The list of device family names to check against.

### Example

```
check_device_family_equivalence "CYLCONE III LS" { "stratixv" "Cyclone IV"  
"cycloneiiiils" }
```

### Related Information

[get\\_device\\_family\\_displayname](#) on page 9-84

## get\_device\_family\_displayname

### Description

Returns the display name of a given device family.

### Availability

Discovery, Main Program, Edit, Elaboration, Validation, Composition, Fileset Generation, Parameter Upgrade

### Usage

```
get_device_family_displayname <device_family>
```

### Returns

The preferred display name for the device family.

### Arguments

**device\_family**

A device family name.

### Example

```
get_device_family_displayname cycloneiiils ( returns: "Cyclone IV LS" )
```

### Related Information

[check\\_device\\_family\\_equivalence](#) on page 9-83





## get\_qip\_strings

### Description

Returns a Tcl list of QIP strings for the IP component.

### Availability

Discovery, Main Program, Edit, Elaboration, Validation, Composition, Parameter Upgrade

### Usage

```
get_qip_strings
```

### Returns

A Tcl list of qip strings set by this IP component.

### Arguments

No arguments.

### Example

```
set strings [ get_qip_strings ]
```

### Related Information

[set\\_qip\\_strings](#) on page 9-86

## set\_qip\_strings

### Description

Places strings in the Quartus II IP File (**.qip**) file, which Qsys passes to the command as a Tcl list. You add the **.qip** file to your Quartus II project on the **Files** page, in the **Settings** dialog box. Successive calls to `set_qip_strings` are not additive and replace the previously declared value.

### Availability

Discovery, Main Program, Edit, Elaboration, Validation, Composition, Parameter Upgrade

### Usage

```
set_qip_strings <qip_strings>
```

### Returns

The Tcl list which was set.

### Arguments

#### **qip\_strings**

A space-delimited Tcl list.

### Example

```
set_qip_strings {"QIP Entry 1" "QIP Entry 2"}
```

### Notes

You can use the following macros in your QIP strings entry:

- %entityName%** The generated name of the entity replaces this macro when the string is written to the **.qip** file.
- %libraryName%** The compilation library this IP component was compiled into is inserted in place of this macro inside the **.qip** file.
- %instanceName%** The name of the instance is inserted in place of this macro inside the **.qip** file.

### Related Information

[get\\_qip\\_strings](#) on page 9-85

## set\_interconnect\_requirement

### Description

Sets the value of an interconnect requirement for a system or an interface on a child instance.

### Availability

Composition

### Usage

```
set_interconnect_requirement <element_id> <name> <value>
```

### Returns

No return value

### Arguments

#### element\_id

{*\$system*} for system requirements, or qualified name of the interface of an instance, in *<instance>.<interface>* format. Note that the system identifier has to be escaped in TCL.

#### name

The name of the requirement.

#### value

The new requirement value.

### Example

```
set_interconnect_requirement {$system} qsys_mm.maxAdditionalLatency 2
```

## Qsys\_hw.tcl Property Reference

- [Script Language Properties](#) on page 9-89
- [Interface Properties](#) on page 9-90
- [Instance Properties](#) on page 9-91
- [Parameter Properties](#) on page 9-92
- [Parameter Type Properties](#) on page 9-95
- [Parameter Status Properties](#) on page 9-96
- [Port Properties](#) on page 9-97
- [Direction Properties](#) on page 9-99
- [Display Item Properties](#) on page 9-100
- [Display Item Kind Properties](#) on page 9-101
- [Display Hint Properties](#) on page 9-102
- [Module Properties](#) on page 9-103
- [Fileset Properties](#) on page 9-105
- [Fileset Kind Properties](#) on page 9-106
- [Callback Properties](#) on page 9-107
- [File Attribute Properties](#) on page 9-108
- [File Kind Properties](#) on page 9-109
- [File Source Properties](#) on page 9-110
- [Simulator Properties](#) on page 9-111
- [Port VHDL Type Properties](#) on page 9-112
- [System Info Type Properties](#) on page 9-113
- [Design Environment Type Properties](#) on page 9-115
- [Units Properties](#) on page 9-116
- [Operating System Properties](#) on page 9-117
- [Quartus.ini Type Properties](#) on page 9-118



## Script Language Properties

Name	Description
TCL	Implements the script in Tcl.

## Interface Properties

Name	Description
CMSIS_SVD_FILE	Specifies the connection point's associated CMSIS file.
CMSIS_SVD_VARIABLES	Defines the variables inside a .svd file.
ENABLED	Specifies whether or not interface is enabled.
EXPORT_OF	For composed <code>_hwl.tcl</code> files, the <code>EXPORT_OF</code> property indicates which interface of a child instance is to be exported through this interface. Before using this command, you must have created the border interface using <code>add_interface</code> . The interface to be exported is of the form <code>&lt;instanceName.interfaceName&gt;</code> .  Example: <code>set_interface_property CSC_input EXPORT_OF my_colorSpace-Converter.input_port</code>
PORT_NAME_MAP	A map of external port names to internal port names, formatted as a Tcl list. Example: <code>set_interface_property &lt;interface name&gt; PORT_NAME_MAP "&lt;new port name&gt; &lt;old port name&gt; &lt;new port name 2&gt; &lt;old port name 2&gt;"</code>
SVD_ADDRESS_GROUP	Generates a CMSIS SVD file. Masters in the same SVD address group will write register data of their connected slaves into the same SVD file
SVD_ADDRESS_OFFSET	Generates a CMSIS SVD file. Slaves connected to this master will have their base address offset by this amount in the SVD file.

## Instance Properties

Name	Description
HDLINSTANCE_GET_GENERATED_NAME	Qsys uses this property to get the auto-generated fixed name when the instance property HDLINSTANCE_USE_GENERATED_NAME is set to true, and only applies to fileSet callbacks.
HDLINSTANCE_USE_GENERATED_NAME	If true, instances added with the <code>add_hdl_instance</code> command are instructed to use unique auto-generated fixed names based on the parameterization.
SUPPRESS_ALL_INFO_MESSAGES	If true, allows you to suppress all Info messages that originate from a child instance.
SUPPRESS_ALL_WARNINGS	If true, allows you to suppress all warnings that originate from a child instance

## Parameter Properties

Type	Name	Description
Boolean	<code>AFFECTS_ELABORATION</code>	Set <code>AFFECTS_ELABORATION</code> to <code>false</code> for parameters that do not affect the external interface of the module. An example of a parameter that does not affect the external interface is <code>isNonVolatileStorage</code> . An example of a parameter that does affect the external interface is <code>width</code> . When the value of a parameter changes, if that parameter has set <code>AFFECTS_ELABORATION=false</code> , the elaboration phase (calling the callback or hardware analysis) is not repeated, improving performance. Because the default value of <code>AFFECTS_ELABORATION</code> is <code>true</code> , the provided HDL file is normally re-analyzed to determine the new port widths and configuration every time a parameter changes.
Boolean	<code>AFFECTS_GENERATION</code>	The default value of <code>AFFECTS_GENERATION</code> is <code>false</code> if you provide a top-level HDL module; it is <code>true</code> if you provide a fileset callback. Set <code>AFFECTS_GENERATION</code> to <code>false</code> if the value of a parameter does not change the results of fileset generation.
Boolean	<code>AFFECTS_VALIDATION</code>	The <code>AFFECTS_VALIDATION</code> property marks whether a parameter's value is used to set derived parameters, and whether the value affects validation messages. When set to <code>false</code> , this may improve response time in the parameter editor UI when the value is changed.
String[]	<code>ALLOWED_RANGES</code>	Indicates the range or ranges that the parameter value can have. For integers, The <code>ALLOWED_RANGES</code> property is a list of ranges that the parameter can take on, where each range is a single value, or a range of values defined by a start and end value separated by a colon, such as <code>11:15</code> . This property can also specify legal values and display strings for integers, such as <code>{0:None 1:Monophonic 2:Stereo 4:Quadrophonic}</code> meaning 0, 1, 2, and 4 are the legal values. You can also assign display strings to be displayed in the parameter editor for string variables. For example, <code>ALLOWED_RANGES {"dev1:Cyclone IV GX" "dev2:Stratix V GT"}</code> .
String	<code>DEFAULT_VALUE</code>	The default value.
Boolean	<code>DERIVED</code>	When <code>true</code> , indicates that the parameter value can only be set by the IP component, and cannot be set by the user. Derived parameters are not saved as part of an instance's parameter values. The default value is <code>false</code> .
String	<code>DESCRIPTION</code>	A short user-visible description of the parameter, suitable for a tooltip description in the parameter editor.
String[]	<code>DISPLAY_HINT</code>	Provides a hint about how to display a property. The following values are possible:



Type	Name	Description
		<ul style="list-style-type: none"> <li><code>boolean</code>--for <code>integer</code> parameters whose value can be 0 or 1. The parameter displays as an option that you can turn on or off.</li> <li><code>radio</code>--displays a parameter with a list of values as radio buttons instead of a drop-down list.</li> <li><code>hexadecimal</code>--for <code>integer</code> parameters, display and interpret the value as a hexadecimal number, for example: <code>0x00000010</code> instead of 16.</li> <li><code>fixed_size</code>--for <code>string_list</code> and <code>integer_list</code> parameters, the <code>fixed_size</code> <code>DISPLAY_HINT</code> eliminates the <b>add</b> and <b>remove</b> buttons from tables.</li> </ul>
String	<code>DISPLAY_NAME</code>	This is the GUI label that appears to the left of this parameter.
String	<code>DISPLAY_UNITS</code>	This is the GUI label that appears to the right of the parameter.
Boolean	<code>ENABLED</code>	When <code>false</code> , the parameter is disabled, meaning that it is displayed, but greyed out, indicating that it is not editable on the parameter editor.
String	<code>GROUP</code>	Controls the layout of parameters in GUI
Boolean	<code>HDL_PARAMETER</code>	When true, the parameter must be passed to the HDL IP component description. The default value is <code>false</code> .
String	<code>LONG_DESCRIPTION</code>	A user-visible description of the parameter. Similar to <code>DESCRIPTION</code> , but allows for a more detailed explanation.
String	<code>NEW_INSTANCE_VALUE</code>	This property allows you to change the default value of a parameter without affecting older IP components that have did not explicitly set a parameter value, and use the <code>DEFAULT_VALUE</code> property. The practical result is that older instances will continue to use <code>DEFAULT_VALUE</code> for the parameter and new instances will use the value assigned by <code>NEW_INSTANCE_VALUE</code> .
String[]	<code>SYSTEM_INFO</code>	Allows you to assign information about the instantiating system to a parameter that you define. <code>SYSTEM_INFO</code> requires an argument specifying the type of information requested, <code>&lt;info-type&gt;</code> .
String	<code>SYSTEM_INFO_ARG</code>	Defines an argument to be passed to a particular <code>SYSTEM_INFO</code> function, such as the name of a reset interface.
(various)	<code>SYSTEM_INFO_TYPE</code>	Specifies one of the types of system information that can be queried. Refer to <i>System Info Type Properties</i> .
(various)	<code>TYPE</code>	Specifies the type of the parameter. Refer to <i>Parameter Type Properties</i> .
(various)	<code>UNITS</code>	Sets the units of the parameter. Refer to <i>Units Properties</i> .
Boolean	<code>VISIBLE</code>	Indicates whether or not to display the parameter in the parameterization GUI.
String	<code>WIDTH</code>	For a <code>STD_LOGIC_VECTOR</code> parameter, this indicates the width of the logic vector.

**Related Information**

- [System Info Type Properties](#) on page 9-113
- [Parameter Type Properties](#) on page 9-95
- [Units Properties](#) on page 9-116

## Parameter Type Properties

Name	Description
BOOLEAN	A boolean parameter whose value is <code>true</code> or <code>false</code> .
FLOAT	A signed 32-bit floating point parameter. Not supported for HDL parameters.
INTEGER	A signed 32-bit integer parameter.
INTEGER_LIST	A parameter that contains a list of 32-bit integers. Not supported for HDL parameters.
LONG	A signed 64-bit integer parameter. Not supported for HDL parameters.
NATURAL	A 32-bit number that contain values 0 to 2147483647 (0x7fffffff).
POSITIVE	A 32-bit number that contains values 1 to 2147483647 (0x7fffffff).
STD_LOGIC	A single bit parameter whose value can be 1 or 0;
STD_LOGIC_VECTOR	An arbitrary-width number. The parameter property <code>WIDTH</code> determines the size of the logic vector.
STRING	A string parameter.
STRING_LIST	A parameter that contains a list of strings. Not supported for HDL parameters.

## Parameter Status Properties

Type	Name	Description
Boolean	ACTIVE	Indicates the parameter is a regular parameter.
Boolean	DEPRECATED	Indicates the parameter exists only for backwards compatibility, and may not have any effect.
Boolean	EXPERIMENTAL	Indicates the parameter is experimental, and not exposed in the design flow.

## Port Properties

Type	Name	Description
(various)	DIRECTION	The direction of the port from the IP component's perspective. Refer to <i>Direction Properties</i> .
String	DRIVEN_BY	Indicates that this output port is always driven to a constant value or by an input port. If all outputs on an IP component specify a <code>driven_by</code> property, the HDL for the IP component will be generated automatically.
String[]	FRAGMENT_LIST	This property can be used in 2 ways: First you can take a single RTL signal and split it into multiple Qsys signals <code>add_interface_port &lt;interface&gt; foo &lt;role&gt; &lt;direction&gt; &lt;width&gt;</code> <code>add_interface_port &lt;interface&gt; bar &lt;role&gt; &lt;direction&gt; &lt;width&gt;</code> <code>set_port_property foo fragment_list "my_rtl_signal(3:0)"</code> <code>set_port_property bar fragment_list "my_rtl_signal(6:4)"</code> Second you can take multiple RTL signals and combine them into a single Qsys signal <code>add_interface_port &lt;interface&gt; baz &lt;role&gt; &lt;direction&gt; &lt;width&gt;</code> <code>set_port_property baz fragment_list "rtl_signal_1(3:0) rtl_signal_2(3:0)"</code> Note: The listed bits in a port fragment must match the declared width of the Qsys signal.
String	ROLE	Specifies an Avalon signal type such as <code>waitrequest</code> , <code>readdata</code> , or <code>read</code> . For a complete list of signal types, refer to the <i>Avalon Interface Specifications</i> .
Boolean	TERMINATION	When <code>true</code> , instead of connecting the port to the Qsys system, it is left unconnected for <code>output</code> and <code>bidir</code> or set to a fixed value for <code>input</code> . Has no effect for IP components that implement a generation callback instead of using the default wrapper generation.
BigInteger	TERMINATION_VALUE	The constant value to drive an input port.
(various)	VHDL_TYPE	Indicates the type of a VHDL port. The default value, <code>auto</code> , selects <code>std_logic</code> if the width is fixed at 1, and <code>std_logic_vector</code> otherwise. Refer to <i>Port VHDL Type Properties</i> .
String	WIDTH	The width of the port in bits. Cannot be set directly. Any changes must be set through the <code>WIDTH_EXPR</code> property.
String	WIDTH_EXPR	The width expression of a port. The <code>width_value_expr</code> property can be set directly to a numeric value if desired. When <code>get_port_property</code> is used <code>width</code> always returns the current integer width of the port while <code>width_expr</code> always returns the unevaluated width expression.
Integer	WIDTH_VALUE	The width of the port in bits. Cannot be set directly. Any changes must be set through the <code>WIDTH_EXPR</code> property.

### Related Information

- [Direction Properties](#) on page 9-99

- [Port VHDL Type Properties](#) on page 9-112
- [Avalon Interface Specifications](#)

## Direction Properties

Name	Description
Bidir	Direction for a bidirectional signal.
Input	Direction for an input signal.
Output	Direction for an output signal.

## Display Item Properties

Type	Name	Description
String	DESCRIPTION	A description of the display item, which you can use as a tooltip.
String[]	DISPLAY_HINT	A hint that affects how the display item displays in the parameter editor.
String	DISPLAY_NAME	The label for the display item in a the parameter editor.
Boolean	ENABLED	Indicates whether the display item is enabled or disabled.
String	PATH	The path to a file. Only applies to display items of type <code>ICON</code> .
String	TEXT	Text associated with a display item. Only applies to display items of type <code>TEXT</code> .
Boolean	VISIBLE	Indicates whether this display item is visible or not.





## Display Item Kind Properties

Name	Description
ACTION	An action displays as a button in the GUI. When the button is clicked, it calls the callback procedure. The button label is the display item <code>id</code> .
GROUP	A group that is a child of the <code>parent_group</code> group. If the <code>parent_group</code> is an empty string, this is a top-level group.
ICON	A <b>.gif</b> , <b>.jpg</b> , or <b>.png</b> file.
PARAMETER	A parameter in the instance.
TEXT	A block of text.

## Display Hint Properties

Name	Description
BIT_WIDTH	Bit width of a number
BOOLEAN	Integer value either 0 or 1.
COLLAPSED	Indicates whether a group is collapsed when initially displayed.
COLUMNS	Number of columns in text field, for example, "columns:N".
EDITABLE	Indicates whether a list of strings allows free-form text entry (editable combo box).
FILE	Indicates that the string is an optional file path, for example, "file:jpg,png,gif".
FIXED_SIZE	Indicates a fixed size for a table or list.
GROW	if set, the widget can grow when the IP component is resized.
HEXADECIMAL	Indicates that the long integer is hexadecimal.
RADIO	Indicates that the range displays as radio buttons.
ROWS	Number of rows in text field, or visible rows in a table, for example, "rows:N".
SLIDER	Range displays as slider.
TAB	if present for a group, the group displays in a tab
TABLE	if present for a group, the group must contain all list-type parameters, which display collectively in a single table.
TEXT	String is a text field with a limited character set, for example, "text:A-Za-z0-9_".
WIDTH	width of a table column

## Module Properties

Name	Description
ANALYZE_HDL	When set to false, prevents a call to the Quartus II mapper to verify port widths and directions, speeding up generation time at the expense of fewer validation checks. If this property is set to false, invalid port widths and directions are discovered during the Quartus II software compilation. This does not affect IP components using filesets to manage synthesis files.
AUTHOR	The IP component author.
COMPOSITION_CALLBACK	The name of the composition callback. If you define a composition callback, you cannot not define the generation or elaboration callbacks.
DATASHEET_URL	Deprecated. Use <code>add_documentation_link</code> to provide documentation links.
DESCRIPTION	The description of the IP component, such as "This IP component puts the shizzle in the frobnitz."
DISPLAY_NAME	The name to display when referencing the IP component, such as "My Qsys IP Component".
EDITABLE	Indicates whether you can edit the IP component in the Component Editor.
ELABORATION_CALLBACK	The name of the elaboration callback. When set, the IP component's elaboration callback is called to validate and elaborate interfaces for instances of the IP component.
GENERATION_CALLBACK	The name for a custom generation callback.
GROUP	The group in the IP Catalog that includes this IP component.
ICON_PATH	A path to an icon to display in the IP component's parameter editor.
INstantiate_in_System_Module	If true, this IP component is implemented by HDL provided by the IP component. If false, the IP component will create exported interfaces allowing the implementation to be connected in the parent.
INTERNAL	An IP component which is marked as internal does not appear in the IP Catalog. This feature allows you to hide the sub-IP-components of a larger composed IP component.
MODULE_DIRECTORY	The directory in which the <code>hw.tcl</code> file exists.
MODULE_TCL_FILE	The path to the <code>hw.tcl</code> file.
NAME	The name of the IP component, such as <code>my_qsys_component</code> .
OPAQUE_ADDRESS_MAP	For composed IP components created using a <code>_hw.tcl</code> file that include children that are memory-mapped slaves, specifies whether the children's addresses are visible to downstream

Name	Description
PREFERRED_SIMULATION_LANGUAGE	software tools. When <code>true</code> , the children's address are not visible. When <code>false</code> , the children's addresses are visible.
REPORT_HIERARCHY	null
STATIC_TOP_LEVEL_MODULE_NAME	Deprecated.
STRUCTURAL_COMPOSITION_CALLBACK	The name of the structural composition callback. This callback is used to represent the structural hierarchical model of the IP component and the RTL can be generated either with module property <code>COMPOSITION_CALLBACK</code> or by <code>ADD_FILESET</code> with target <code>QUARTUS_SYNTH</code>
SUPPORTED_DEVICE_FAMILIES	A list of device family supported by this IP component.
TOP_LEVEL_HDL_FILE	Deprecated.
TOP_LEVEL_HDL_MODULE	Deprecated.
UPGRADEABLE_FROM	null
VALIDATION_CALLBACK	The name of the validation callback procedure.
VERSION	The IP component's version, such as 10.0.

## Fileset Properties

Name	Description
ENABLE_FILE_OVERWRITE_MODE	null
ENABLE_RELATIVE_INCLUDE_PATHS	If true, HDL files can include other files using relative paths in the fileset.
TOP_LEVEL	The name of the top-level HDL module that the fileset generates. If set, the HDL top level must match the <code>TOP_LEVEL</code> name, and the HDL must not be parameterized. Qsys runs the generate callback one time, regardless of the number of instances in the system.

## Fileset Kind Properties

Name	Description
EXAMPLE_DESIGN	Contains example design files.
QUARTUS_SYNTH	Contains files that Qsys uses for the Quartus II software synthesis.
SIM_VERILOG	Contains files that Qsys uses for Verilog HDL simulation.
SIM_VHDL	Contains files that Qsys uses for VHDL simulation.

## Callback Properties

### Description

This list describes each type of callback. Each command may only be available in some callback contexts.

Name	Description
ACTION	Called when an ACTION display item's action is performed.
COMPOSITION	Called during instance elaboration when the IP component contains a subsystem.
EDITOR	Called when the IP component is controlling the parameterization editor.
ELABORATION	Called to elaborate interfaces and signals after a parameter change. In API 9.1 and later, validation is called before elaboration. In API 9.0 and earlier, elaboration is called before validation.
GENERATE_VERILOG_SIMULATION	Called when the IP component uses a custom generator to generate the Verilog simulation model for an instance.
GENERATE_VHDL_SIMULATION	Called when the IP component uses a custom generator to generate the VHDL simulation model for an instance.
GENERATION	Called when the IP component uses a custom generator to generate the synthesis HDL for an instance.
PARAMETER_UPGRADE	Called when attempting to instantiate an IP component with a newer version than the saved version. This allows the IP component to upgrade parameters between released versions of the component.
STRUCTURAL_COMPOSITION	Called during instance elaboration when an IP component is represented by a structural hierarchical model which may be different from the generated RTL.
VALIDATION	Called to validate parameter ranges and report problems with the parameter values. In API 9.1 and later, validation is called before elaboration. In API 9.0 and earlier, elaboration is called before validation.

## File Attribute Properties

Name	Description
ALDEC_SPECIFIC	Applies to Aldec simulation tools and for simulation filesets only.
CADENCE_SPECIFIC	Applies to Cadence simulation tools and for simulation filesets only.
COMMON_SYSTEMVERILOG_PACKAGE	The name of the common SystemVerilog package. Applies to simulation filesets only.
MENTOR_SPECIFIC	Applies to Mentor simulation tools and for simulation filesets only.
SYNOPTIS_SPECIFIC	Applies to Synopsys simulation tools and for simulation filesets only.
TOP_LEVEL_FILE	Contains the top-level module for the fileset and applies to synthesis filesets only.



## File Kind Properties

Name	Description
DAT	DAT Data
FLI_LIBRARY	FLI Library
HEX	HEX Data
MIF	MIF Data
OTHER	Other
PLI_LIBRARY	PLI Library
QXP	QXP File
SDC	Timing Constraints
SYSTEM_VERILOG	System Verilog HDL
SYSTEM_VERILOG_ENCRYPT	Encrypted System Verilog HDL
SYSTEM_VERILOG_INCLUDE	System Verilog Include
VERILOG	Verilog HDL
VERILOG_ENCRYPT	Encrypted Verilog HDL
VERILOG_INCLUDE	Verilog Include
VHDL	VHDL
VHDL_ENCRYPT	Encrypted VHDL
VPI_LIBRARY	VPI Library

## File Source Properties

Name	Description
PATH	Specifies the original source file and copies to <code>output_file</code> .
TEXT	Specifies an arbitrary text string for the contents of <code>output_file</code> .

## Simulator Properties

Name	Description
ENV_ALDEC_LD_LIBRARY_PATH	LD_LIBRARY_PATH when running riviera-pro
ENV_CADENCE_LD_LIBRARY_PATH	LD_LIBRARY_PATH when running ncsim
ENV_MENTOR_LD_LIBRARY_PATH	LD_LIBRARY_PATH when running modelsim
ENV_SYNOPSYS_LD_LIBRARY_PATH	LD_LIBRARY_PATH when running vcs
OPT_ALDEC_PLI	-pli option for riviera-pro
OPT_CADENCE_64BIT	-64bit option for ncsim
OPT_CADENCE_PLI	-loadpli1 option for ncsim
OPT_CADENCE_SVLIB	-sv_lib option for ncsim
OPT_CADENCE_SVROOT	-sv_root option for ncsim
OPT_MENTOR_64	-64 option for modelsim
OPT_MENTOR_CPPPATH	-cpppath option for modelsim
OPT_MENTOR_LDFLAGS	-ldflags option for modelsim
OPT_MENTOR_PLI	-pli option for modelsim
OPT_SYNOPSYS_ACC	+acc option for vcs
OPT_SYNOPSYS_CPP	-cpp option for vcs
OPT_SYNOPSYS_FULL64	-full64 option for vcs
OPT_SYNOPSYS_LDFLAGS	-LDFLAGS option for vcs
OPT_SYNOPSYS_LLIB	-l option for vcs
OPT_SYNOPSYS_VPI	+vpi option for vcs

## Port VHDL Type Properties

Name	Description
AUTO	The VHDL type of this signal is automatically determined. Single-bit signals are <code>STD_LOGIC</code> ; signals wider than one bit will be <code>STD_LOGIC_VECTOR</code> .
STD_LOGIC	Indicates that the signal is not rendered in VHDL as a <code>STD_LOGIC</code> signal.
STD_LOGIC_VECTOR	Indicates that the signal is rendered in VHDL as a <code>STD_LOGIC_VECTOR</code> signal.



## System Info Type Properties

Type	Name	Description
String	ADDRESS_MAP	An XML-formatted string describing the address map for the interface specified in the system info argument.
Integer	ADDRESS_WIDTH	The number of address bits required to address all memory-mapped slaves connected to the specified memory-mapped master in this instance, using byte addresses.
String	AVALON_SPEC	The version of the interconnect. SOPC Builder interconnect uses Avalon Specification 1.0. Qsys interconnect uses Avalon Specification 2.0.
Integer	CLOCK_DOMAIN	An integer that represents the clock domain for the interface specified in the system info argument. If this instance has interfaces on multiple clock domains, this can be used to determine which interfaces are on each clock domain. The absolute value of the integer is arbitrary.
Long, Integer	CLOCK_RATE	The rate of the clock connected to the clock input specified in the system info argument. If 0, the clock rate is currently unknown.
String	CLOCK_RESET_INFO	The name of this instance's primary clock or reset sink interface. This is used to determine the reset sink to use for global reset when using SOPC interconnect.
String	CUSTOM_INSTRUCTION_SLAVES	Provides custom instruction slave information, including the name, base address, address span, and clock cycle type.
(various)	DESIGN_ENVIRONMENT	A string that identifies the current design environment. Refer to <i>Design Environment Type Properties</i> .
String	DEVICE	The device part number of the currently selected device.
String	DEVICE_FAMILY	The family name of the currently selected device.
String	DEVICE_FEATURES	A list of key/value pairs delineated by spaces indicating whether a particular device feature is available in the currently selected device family. The format of the list is suitable for passing to the Tcl <code>array set</code> command. The keys are device features; the values will be 1 if the feature is present, and 0 if the feature is absent.
String	DEVICE_SPEEDGRADE	The speed grade of the currently selected device.
Integer	GENERATION_ID	An integer that stores a hash of the generation time to be used as a unique ID for a generation run.

Type	Name	Description
BigInteger, Long	INTERRUPTS_USED	A mask indicating which bits of an interrupt receiver are connected to interrupt senders. The interrupt receiver is specified in the system info argument.
Integer	MAX_SLAVE_DATA_WIDTH	The data width of the widest slave connected to the specified memory-mapped master.
String, Boolean, Integer	QUARTUS_INI	The value of the quartus.ini setting specified in the system info argument.
Integer	RESET_DOMAIN	An integer that represents the reset domain for the interface specified in the system info argument. If this instance has interfaces on multiple reset domains, this can be used to determine which interfaces are on each reset domain. The absolute value of the integer is arbitrary.
String	TRISTATECONDUIT_INFO	An XML description of the Avalon Tri-state Conduit masters connected to an Avalon Tri-state Conduit slave. The slave is specified as the system info argument. The value will contain information about the slave, connected master instance and interface names, and signal names, directions and widths.
String	TRISTATECONDUIT_MASTERS	The names of the instance's interfaces that are tri-state conduit slaves.
String	UNIQUE_ID	A string guaranteed to be unique to this instance.

**Related Information**

[Design Environment Type Properties](#) on page 9-115

## Design Environment Type Properties

### Description

A design environment is used by IP to tell what sort of interfaces are most appropriate to connect in the parent system.

Name	Description
NATIVE	Design environment prefers native IP interfaces.
QSYS	Design environment prefers standard Qsys interfaces.

## Units Properties

Name	Description
Address	A memory-mapped address.
Bits	Memory size, in bits.
BitsPerSecond	Rate, in bits per second.
Bytes	Memory size, in bytes.
Cycles	A latency or count, in clock cycles.
GigabitsPerSecond	Rate, in gigabits per second.
Gigabytes	Memory size, in gigabytes.
Gigahertz	Frequency, in GHz.
Hertz	Frequency, in Hz.
KilobitsPerSecond	Rate, in kilobits per second.
Kilobytes	Memory size, in kilobytes.
Kilohertz	Frequency, in kHz.
MegabitsPerSecond	Rate, in megabits per second.
Megabytes	Memory size, in megabytes.
Megahertz	Frequency, in MHz.
Microseconds	Time, in micros.
Milliseconds	Time, in ms.
Nanoseconds	Time, in ns.
None	Unspecified units.
Percent	A percentage.
Picoseconds	Time, in ps.
Seconds	Time, in s.





## Operating System Properties

Name	Description
ALL	All operating systems
LINUX32	Linux 32-bit
LINUX64	Linux 64-bit
WINDOWS32	Windows 32-bit
WINDOWS64	Windows 64-bit

## Quartus.ini Type Properties

Name	Description
ENABLED	Returns 1 if the setting is turned on, otherwise returns 0.
STRING	Returns the string value of the <b>.ini</b> setting.

## Document Revision History

The table below indicates edits made to the *Component Interface Tcl Reference* content since its creation.

Date	Version	Changes
December 2014	14.1.0	<ul style="list-style-type: none"> <li>set_interface_upgrade_map</li> <li>Moved <b>Port Roles (Interface Signal Types)</b> section to <i>Qsys Interconnect</i>.</li> </ul>
November 2013	13.1.0	<ul style="list-style-type: none"> <li>add_hdl_instance</li> </ul>
May 2013	13.0.0	<ul style="list-style-type: none"> <li>Consolidated content from other Qsys chapters.</li> <li>Added AMBA APB support.</li> </ul>
November 2012	12.1.0	<ul style="list-style-type: none"> <li>Added the <b>demo_axi_memory</b> example with screen shots and example <code>_hw.tcl</code> code.</li> </ul>
June 2012	12.0.0	<ul style="list-style-type: none"> <li>Added AMBA AXI3 support.</li> <li>Added: set_display_item_property, set_parameter_property, LONG_DESCRIPTION, and static filesets.</li> </ul>
November 2011	11.1.0	<ul style="list-style-type: none"> <li>Template update.</li> <li>Added: set_qip_strings, get_qip_strings, get_device_family_displayname, check_device_family_equivalence.</li> </ul>
May 2011	11.0.0	<ul style="list-style-type: none"> <li>Revised section describing HDL and composed component implementations.</li> <li>Separated reset and clock interfaces in example.</li> <li>Added: TRISTATECONDUIT_INFO, GENERATION_ID, UNIQUE_ID SYSTEM_INFO.</li> <li>Added: WIDTH and SYSTEM_INFO_ARG parameter properties.</li> <li>Removed the doc_type argument from the add_documentation_link command.</li> <li>Removed: get_instance_parameter_properties</li> <li>Added: add_fileset, add_fileset_file, create_temp_file.</li> <li>Updated Tcl examples to show separate clock and reset interfaces.</li> </ul>
December 2010	10.1.0	<ul style="list-style-type: none"> <li>Initial release.</li> </ul>

### Related Information

[Quartus II Handbook Archive](#)

[Qsys Interconnect](#) on page 7-1

# Qsys System Design Components 10

2014.12.15

QI15V1



Subscribe



Send Feedback

You can use Qsys IP components to create Qsys systems. Qsys interfaces include components appropriate for streaming high-speed data, reading and writing registers and memory, controlling off-chip devices, and transporting data between components.

Qsys supports Avalon, AMBA AXI3 (version 1.0), AMBA AXI4 (version 2.0), AMBA AXI4-Lite (version 2.0), AMBA AXI4-Stream (version 1.0), and AMBA APB3 (version 1.0) interface specifications.

## Related Information

- [Avalon Interface Specifications](#)
- [AMBA Protocol Specifications](#)
- [Creating a System with Qsys](#) on page 5-1
- [Qsys Interconnect](#) on page 7-1
- [Embedded Peripherals IP User Guide](#)

## Bridges

Bridges affect the way Qsys transports data between components. You can insert bridges between masters and slave interfaces to control the topology of a Qsys system, which affects the interconnect that Qsys generates. You can also use bridges to separate components into different clock domains to isolate clock domain crossing logic.

A bridge has one slave interface and one master interface. In Qsys, one or more master interfaces from other components connect to the bridge slave. The bridge master connects to one or more slave interfaces on other components.


© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



**Figure 10-1: Using a Bridge in a Qsys System**

In this example, three masters have logical connections to three slaves, although physically each master connects only to the bridge. Transfers initiated to the slave propagate to the master in the same order in which they are initiated on the slave.

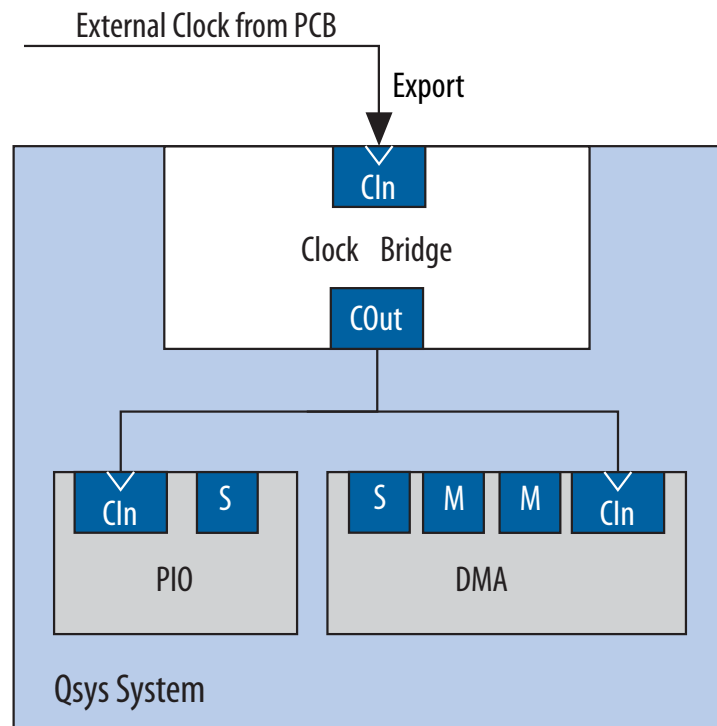


## Clock Bridge

The Clock Bridge allows you to connect a clock source to multiple clock input interfaces. You can use the clock bridge to connect a clock source that is outside the Qsys system. You create the connection through an exported interface, and then connect to multiple clock input interfaces.

Clock outputs have the ability to fan-out without the use of a bridge. You require a bridge only when you want a clock from an exported source to connect internally to more than one source.

Figure 10-2: Clock Bridge



## Avalon-MM Clock Crossing Bridge

The Avalon-MM Clock Crossing Bridge transfers Avalon-MM commands and responses between different clock domains. You can also use the Avalon-MM Clock Crossing Bridge between AXI masters and slaves of different clock domains.

The Avalon-MM Clock Crossing Bridge uses asynchronous FIFOs to implement clock crossing logic. The bridge parameters control the depth of the command and response FIFOs in both the master and slave clock domains. If the number of active reads exceeds the depth of the response FIFO, the Clock Crossing Bridge stops sending reads.

To maintain throughput for high-performance applications, increase the response FIFO depth from the default minimum depth, which is twice the maximum burst size.

**Note:** When you use the FIFO-based clock crossing a Qsys system, the DC FIFO is automatically inserted in the Qsys system. The reset inputs for the DC FIFO are connected to the reset sources for the connected master and slave components on either side of the DC FIFO. With this configuration, both the master side and slave side resets must be asserted at the same time to ensure that the DC FIFO is reset properly.

Alternatively, you can drive both resets from the same reset source to guarantee that the DC FIFO is reset properly.

**Note:** The clock crossing bridge includes appropriate SDC constraints for its internal asynchronous FIFOs. For these SDC constraints to work correctly, you should not set false paths on the pointer crossings in the FIFOs. You should also not split the bridge's clocks into separate clock groups when you declare SDC constraints; this has the same effect as setting false paths.

#### Related Information

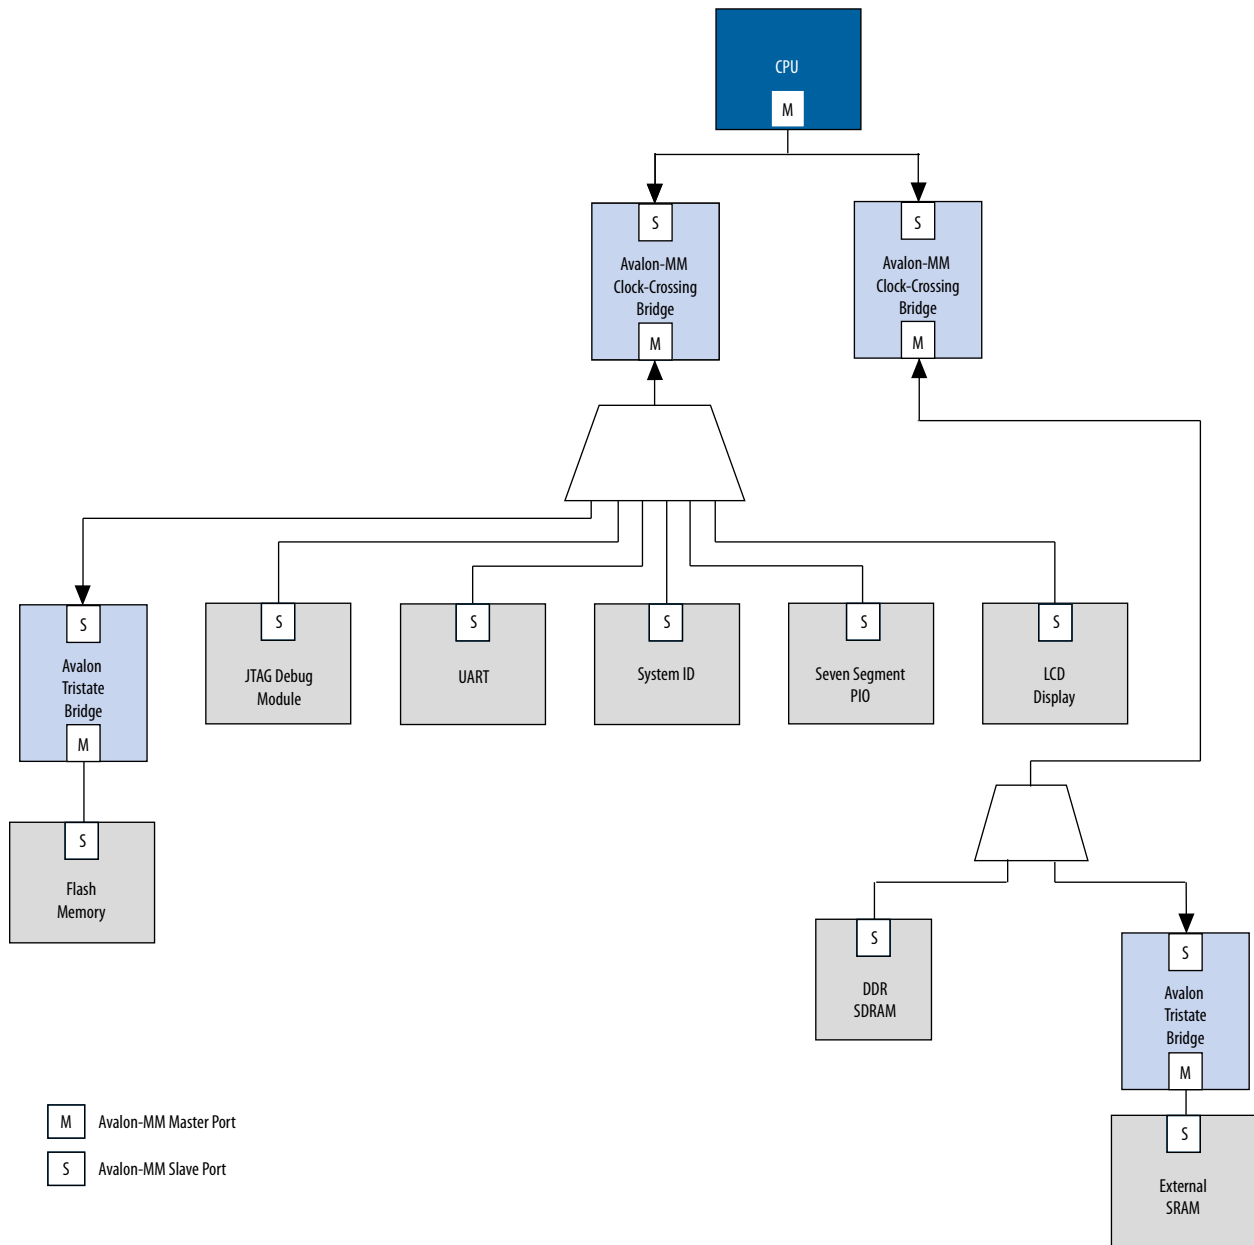
- [Creating a System with Qsys](#) on page 5-1

### Avalon-MM Clock Crossing Bridge Example

In the example shown below, the Avalon-MM Clock Crossing bridges separate slave components into two groups. Low-performance slave components are placed behind a single bridge and are clocked at a low speed. High performance components are placed behind a second bridge and are clocked at a higher speed.

By inserting clock-crossing bridges, you simplify the Qsys interconnect and allow the Quartus II Fitter to optimize paths that require minimal propagation delay.

Figure 10-3: Avalon-MM Clock Crossing Bridge





## Avalon-MM Clock Crossing Bridge Parameters

Table 10-1: Avalon-MM Clock Crossing Bridge Parameters

Parameters	Values	Description
<b>Data width</b>	8, 16, 32, 64, 128, 256, 512, 1024 bits	Determines the data width of the interfaces on the bridge, and affects the size of both FIFOs. For the highest bandwidth, set <b>Data width</b> to be as wide as the widest master that connects to the bridge.
<b>Symbol width</b>	1, 2, 4, 8, 16, 32, 64 (bits)	Number of bits per symbol. For example, byte-oriented interfaces have 8-bit symbols.
<b>Address width</b>	1-32 bits	The address bits needed to address the downstream slaves.
<b>Use automatically-determined address width</b>	-	The minimum bridge address width that is required to address the downstream slaves.
<b>Maximum burst size</b>	1, 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024 bits	Determines the maximum length of bursts that the bridge supports.
<b>Command FIFO depth</b>	2, 4, 8, 16, 32, 64, 128, 256, 512, 1024 2048, 4096, 8192, 16384 bits	Command (master-to-slave) FIFO depth.
<b>Respond FIFO depth</b>	2, 4, 8, 16, 32, 64, 128, 256, 512, 1024 2048, 4096, 8192, 16384 bits	Response (slave-to-master) FIFO depth.
<b>Master clock domain synchronizer depth</b>	2, 3, 4, 5 bits	The number of pipeline stages in the clock crossing logic in the issuing master to target slave direction. Increasing this value leads to a larger meantime between failures (MTBF). You can determine the MTBF for a design by running a TimeQuest timing analysis.

Parameters	Values	Description
<b>Slave clock domain synchronizer depth</b>	2, 3, 4, 5 bits	The number of pipeline stages in the clock crossing logic in the target slave to master direction. Increasing this value leads to a larger meantime between failures (MTBF). You can determine the MTBF for a design by running a TimeQuest timing analysis.


## Avalon-MM Pipeline Bridge

The Avalon-MM Pipeline Bridge inserts a register stage in the Avalon-MM command and response paths. It accepts commands on its Avalon-MM slave port and propagates the commands to its Avalon-MM master port. The pipeline bridge provides separate parameters to turn on pipelining in the command and response networks.

You can use the Avalon-MM bridge to export a single Avalon-MM slave interface to use to control multiple Avalon-MM slave devices. The pipelining feature is optional. You can optionally turn off the pipelining feature of this bridge.

**Figure 10-4: Avalon-MM Pipeline Bridge in a XAUI PHY Transceiver IP Core**

In this example, the bridge transfers commands received on its slave interface to its master port.




Because the slave interface is exported to the pins of the device, having a single slave port, rather than separate ports for each slave device, reduces the pin count of the FPGA.

## Avalon-MM Unaligned Burst Expansion Bridge

The Avalon-MM Unaligned Burst Expansion Bridge aligns read burst transactions from masters connected to its slave interface, to the address space of slaves connected to its master interface. This alignment ensures that all read burst transactions are delivered to the slave as a single transaction.

Figure 10-5: Avalon-MM Unaligned Burst Expansion Bridge



You can use the Avalon Unaligned Burst Expansion Bridge to align read burst transactions from masters that have narrower data widths than the target slaves. Using the bridge for this purpose improves bandwidth utilization for the master-slave pair, and ensures that un-aligned bursts are processed as single transactions rather than multiple transactions.

**Note:** Do not use the Avalon-MM Unaligned Burst Expansion Bridge if any connected slave has read side effects from reading addresses that are exposed to any connected master's address map. This bridge can cause read side effects due to alignment modification to read burst transaction addresses.

**Note:** For Qsys 14.0, the Avalon-MM Unaligned Burst Expansion Bridge does not support VHDL simulation.

#### Related Information

- [Qsys Interconnect](#) on page 7-1

### Using the Avalon-MM Unaligned Burst Expansion Bridge

When a master sends a read burst transaction to a slave, the Avalon-MM Unaligned Burst Expansion Bridge initially determines whether the start address of the read burst transaction is aligned to the slave's memory address space. If the base address is aligned, the bridge does not change the base address. If the base address is not aligned, the bridge aligns the base address to the nearest aligned address that is less than the requested base address.

The Avalon-MM Unaligned Burst Expansion Bridge then determines whether the final word requested by the master is the last word at the slave read burst address. If a single slave address contains multiple words, all of those words must be requested in order for a single read burst transaction to occur.

- If the final word requested by the master is the last word at the slave read burst address, the bridge does not modify the burst length of the read burst command to the slave.
- If the final word requested by the master is not the last word at the slave read burst address, the bridge increases the burst length of the read burst command to the slave. The final word requested by the modified read burst command is then the last word at the slave read burst address.

The bridge stores information about each aligned read burst command that it sends to slaves connected to a master interface. When a read response is received on the master interface, the bridge determines if the base address or burst length of the issued read burst command was altered.

If the bridge alters either the base address or the burst length of the issued read burst command, it receives response words that the master did not request. The bridge suppresses words that it receives from the aligned burst response that are not part of the original read burst command from the master.

## Avalon-MM Unaligned Burst Expansion Bridge Parameters

Figure 10-6: Avalon-MM Unaligned Burst Expansion Bridge Parameter Editor

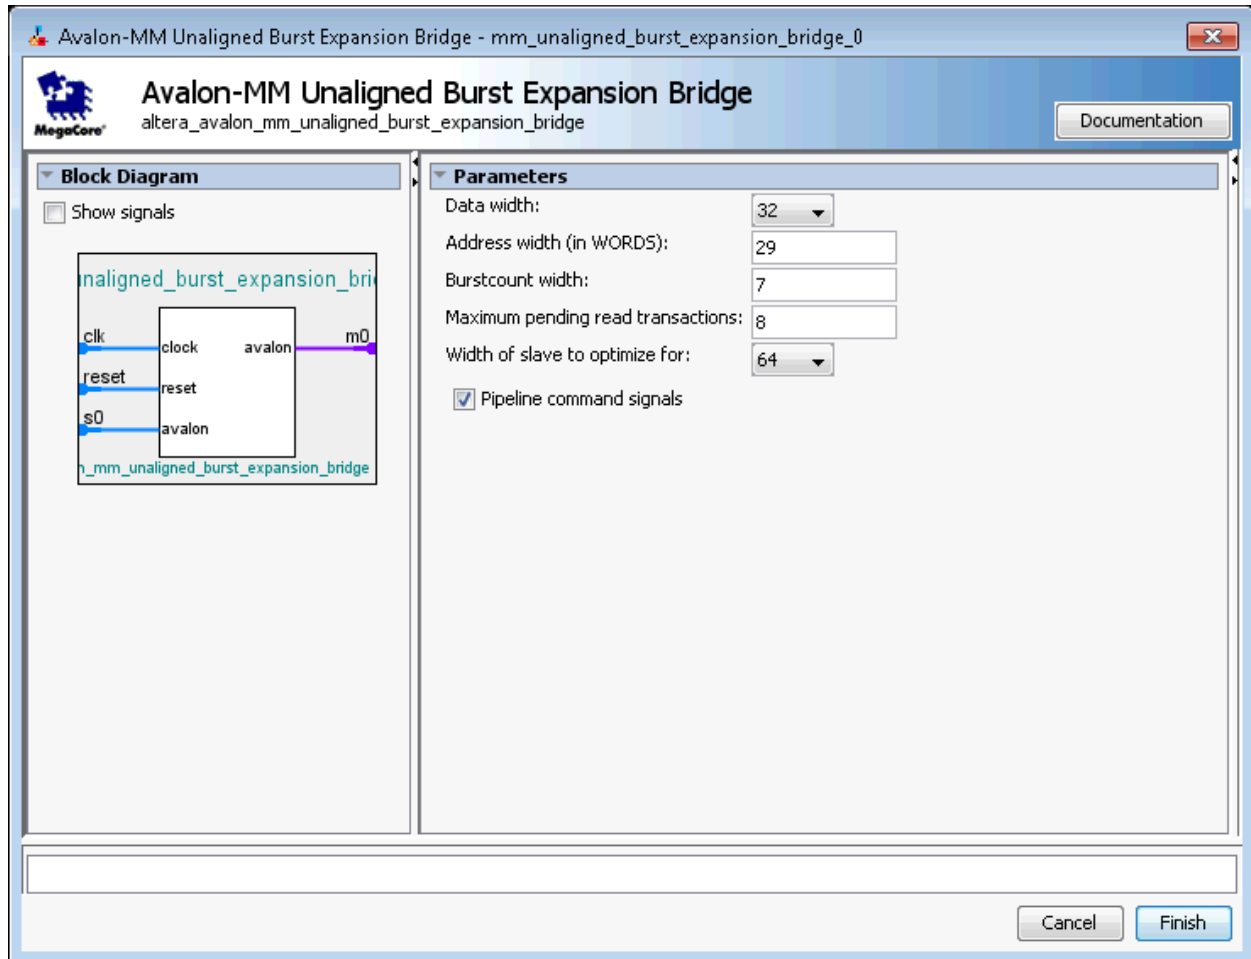


Table 10-2: Avalon-MM Unaligned Burst Expansion Bridge Parameters

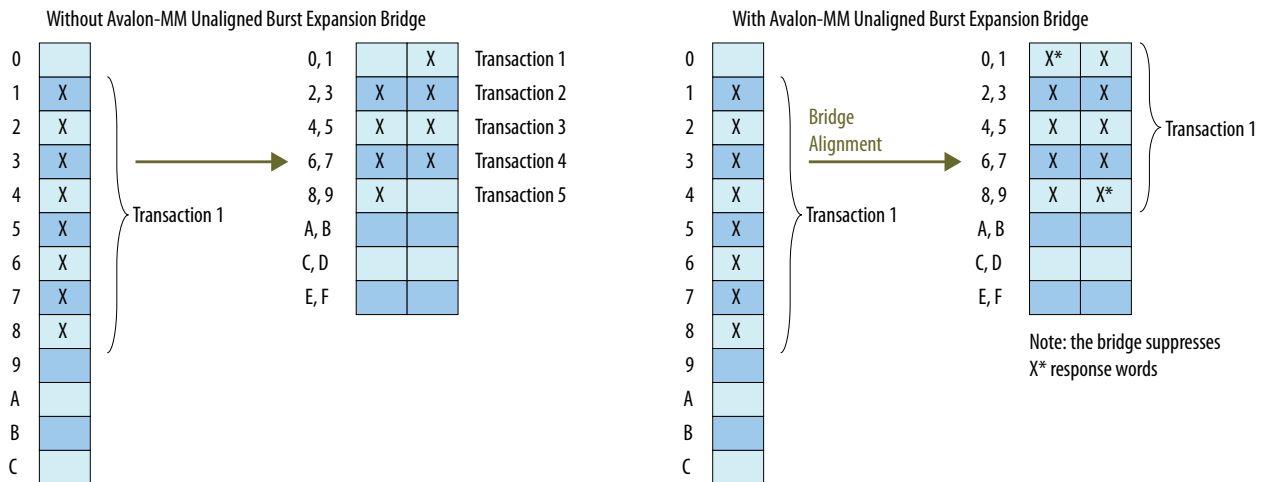
Parameter	Description
<b>Data width</b>	Data width of the master connected to the bridge.
<b>Address width (in WORDS)</b>	The address width of the master connected to the bridge.
<b>Burstcount width</b>	The burstcount signal width of the master connected to the bridge.

Parameter	Description
<b>Maximum pending read transactions</b>	The maximum pending read transactions interface property of the bridge.
<b>Width of slave to optimize for</b>	The data width of the connected slave. Supported values are: 16, 32, 64, 128, 256, 512, 1024, 2048, and 4096 bits. <b>Note:</b> If you connect multiple slaves, all slaves must have the same data width.
<b>Pipeline command signals</b>	When turned on, the command path is pipelined, minimizing the bridge's critical path at the expense of increased logic usage and latency.

### Avalon-MM Unaligned Burst Expansion Bridge Example

Figure 10-7: Unaligned Burst Expansion Bridge

The example below shows an unaligned read burst command from a master that the Avalon-MM Unaligned Burst Expansion Bridge converts to an aligned request for a connected slave, and the suppression of words due to the aligned read burst command. In this example, a 32-bit master requests an 8-beat burst of 32-bit words from a 64-bit slave with a start address that is not 64-bit aligned.



Because the target slave has a 64-bit data width, address 1 is unaligned in the slave's address space. As a result, several smaller burst transactions are needed to request the data associated with the master's read burst command.

With an Avalon-MM Unaligned Burst Expansion Bridge in place, the bridge issues a new read burst command to the target slave beginning at address 0 with burst length 10, which requests data up to the word stored at address 9.

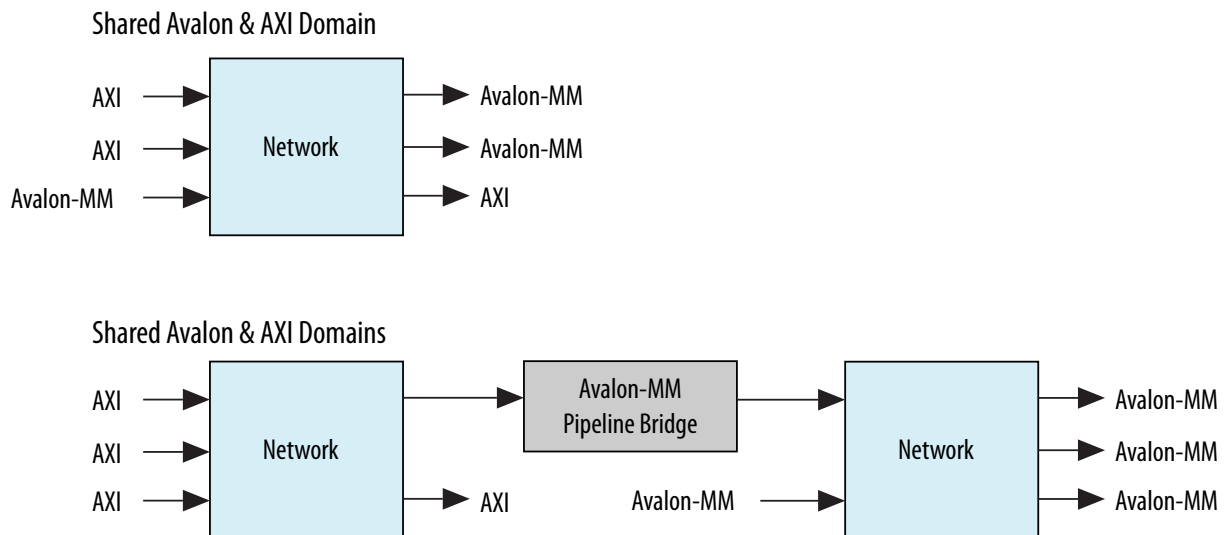
When the bridge receives the word corresponding to address 0, it suppresses it from the master, and then delivers the words corresponding to addresses 1 through 8 to the master. When the bridge receives the word corresponding to address 9, it suppresses that word from the master.

## Bridges Between Avalon and AXI Interfaces

When designing a Qsys system, you can make connections between AXI and Avalon interfaces without the use of explicitly-instantiated bridges; the interconnect provides all necessary bridging logic. However, this does not prevent the use of explicit bridges to separate the AXI and Avalon domains.

**Figure 10-8: Avalon-MM Pipeline Bridge Between Avalon-MM and AXI Domains**

Using an explicit Avalon-MM bridge to separate the AXI and Avalon domains reduces the amount of bridging logic in the interconnect at the expense of concurrency.



### AXI Bridge

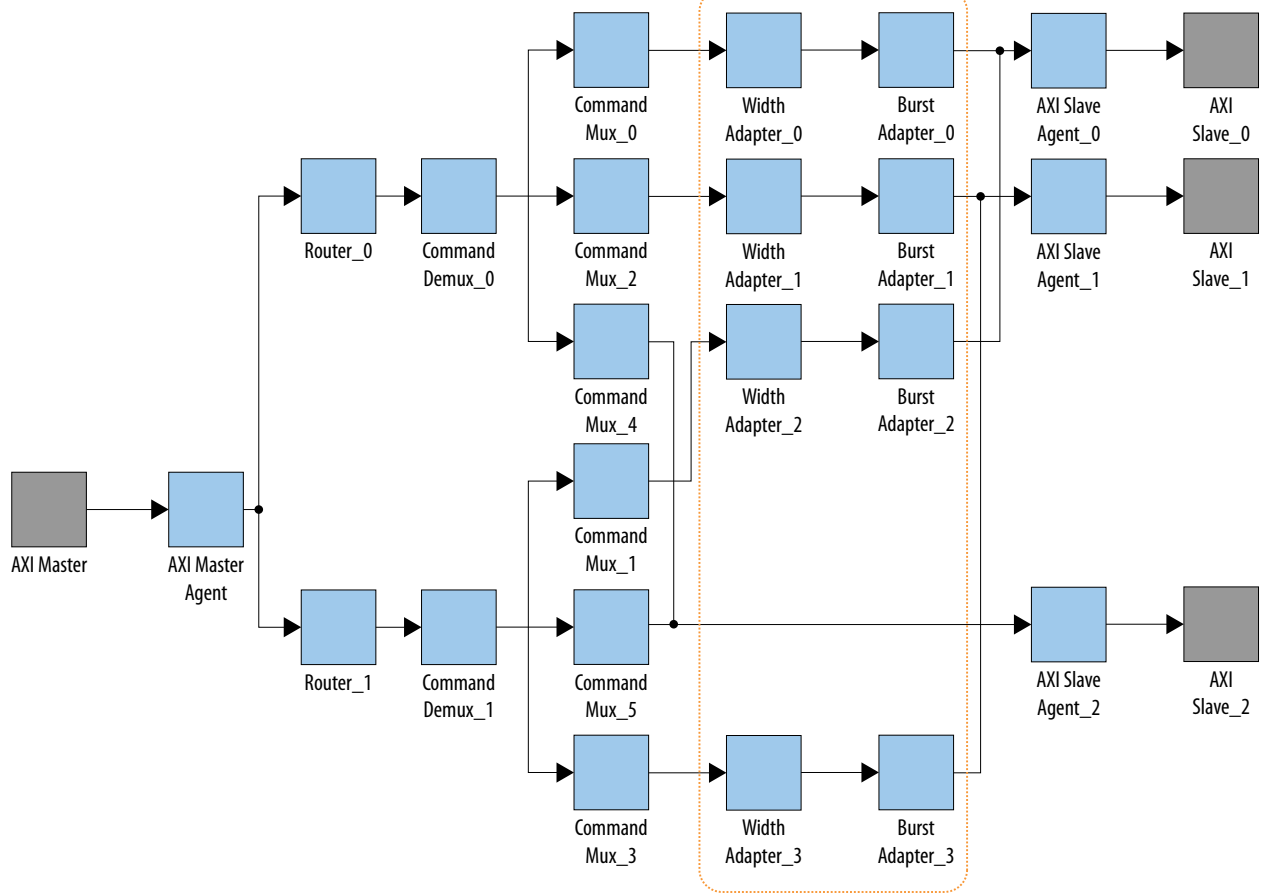
With an AXI bridge, you can influence the placement of resource-intensive components, such as the width and burst adapters. Depending on its use, an AXI bridge may reduce throughput and concurrency, in return for higher  $f_{Max}$  and less logic.

You can use an AXI bridge to group different parts of your Qsys system. Then, other parts of the system connect to the bridge interface instead of to multiple separate master or slave interfaces. You can also use an AXI bridge to export AXI interfaces from Qsys systems.

The example below shows a system with a single AXI master and three AXI slaves. It also has various interconnect components, such as routers, demuxes, and muxes. Two of the slaves have a narrower data width than the master; 16-bit slaves versus a 32-bit master. In this system, Qsys interconnect creates four width adapters and four burst adapters to access the two slaves. In this case, you could improve resource usage by adding an AXI bridge. This would result in Qsys having to add only two width adapters and two burst adapters, one pair for the read channels, and another pair for the write channel.

Figure 10-9: AXI Example Without a Bridge: Adding a Bridge Can Reduce the Number of Adapters

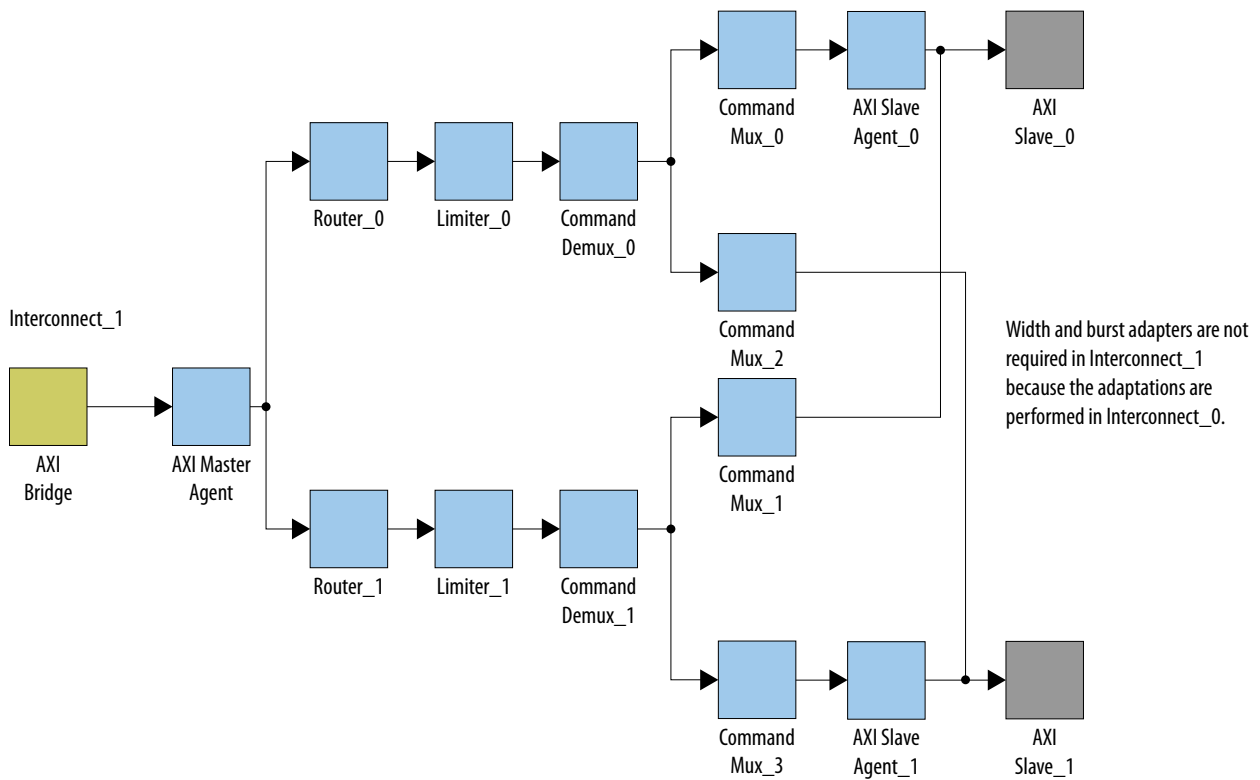
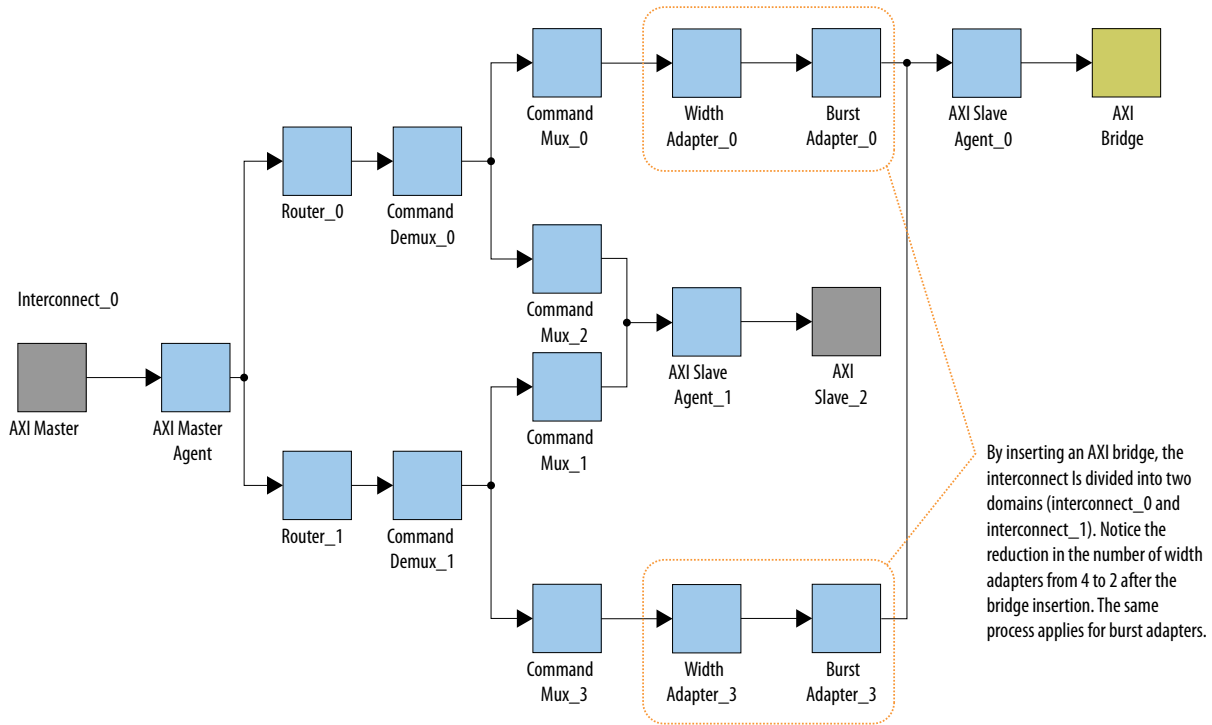
Four width adapters (0 - 3) and four burst adapters (0 - 3) are inserted between the master and slaves for transaction adaptation for the example system.



The example below shows the same system with an AXI bridge component, and the decrease in the number of width and burst adapters. Qsys creates only two width adapters, and two burst adapters, as compared to the four width adapters and four burst adapters in the previous example. The system includes more components, but the overall system performance improves because there are fewer resource-intensive width and burst adapters.



Figure 10-10: Width and Burst Adapters Added to a System With a Bridge



## AXI Bridge Signal Types

Based on parameter selections that you make for the AXI Bridge component, Qsys instantiates either the AXI3 or AXI4 master and slave interfaces into the component.

**Note:** In AXI3, *aw/aruser* accommodates sideband signal usage by hard processor systems (HPS).

**Table 10-3: Sets of Signals for the AXI Bridge Based on the Protocol**

Signal Name	AXI3	AXI4
<i>awid / arid</i>	yes	yes
<i>awaddr / araddr</i>	yes	yes
<i>awlen / arlen</i>	yes (4-bit)	yes (8-bit)
<i>awsizel / arsizel</i>	yes	yes
<i>awburst / arburst</i>	yes	yes
<i>awlock / arlock</i>	yes	yes (1-bit optional)
<i>awcache / arcache</i>	yes (2-bit)	yes (optional)
<i>awprot / arprot</i>	yes	yes
<i>awuser / aruser</i>	yes	yes
<i>awvalid / arvalid</i>	yes	yes
<i>awready / arready</i>	yes	yes
<i>awqos / arqos</i>	no	yes
<i>awregion / arregion</i>	no	yes
<i>wid</i>	yes	no (optional)
<i>wdata / rdata</i>	yes	yes
<i>wstrb</i>	yes	yes
<i>wlast / rvalid</i>	yes	yes
<i>wvalid / rlack</i>	yes	yes
<i>wready / rready</i>	yes	yes
<i>wuser / ruser</i>	no	yes
<i>bid / rid</i>	yes	yes
<i>bresp / rresp</i>	yes	yes (optional)
<i>bvalid</i>	yes	yes
<i>bready</i>	yes	yes

## AXI Bridge Parameters

In the parameter editor, you can customize the parameters for the AXI bridge according to the requirements of your design.

Figure 10-11: AXI Bridge Parameter Editor

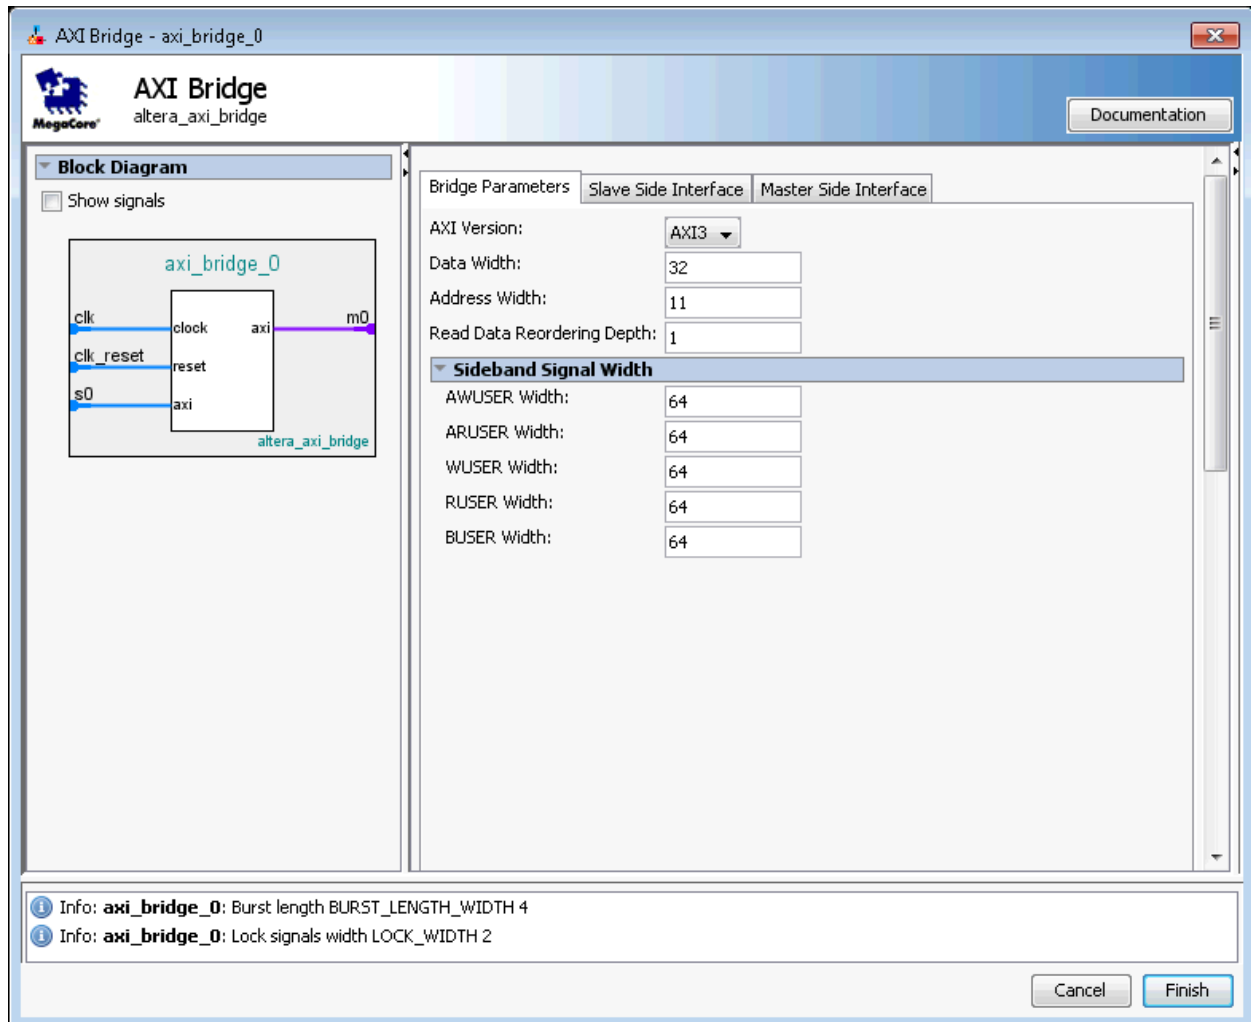


Table 10-4: AXI Bridge Parameters

Parameter	Type	Range	Description
<b>AXI Version</b>	string	AXI3/ AXI4	Specifies the AXI version and signals that Qsys generates for the slave and master interfaces of the bridge.
<b>Data Width</b>	int	8:1024	Controls the width of the data for the master and slave interfaces.
<b>Address Width</b>	int	1-64 bits	Controls the width of the address for the master and slave interfaces.

Parameter	Type	Range	Description
<b>Read Data Reordering Depth</b>	int	1-16	Controls the multithreading feature and out-of-order responses.  If a master issues different thread IDs to different slaves, in order for a slave to view the different thread IDs, you must set the <b>Read Data Reordering Depth</b> to 1.
<b>AWUSER Width</b>	int	1-64 bits	Controls the width of the write address channel sideband signals of the master and slave interfaces.
<b>ARUSER Width</b>	int	1-64 bits	Controls the width of the read address channel sideband signals of the master and slave interfaces.
<b>WUSER Width</b>	int	1-64 bits	Controls the width of the write data channel sideband signals of the master and slave interfaces.
<b>RUSER Width</b>	int	1-16 bits	Controls the width of the read data channel sideband signals of the master and slave interfaces.
<b>BUSER Width</b>	int	1-16 bits	Controls the width of the write response channel sideband signals of the master and slave interfaces.

## AXI Bridge Slave and Master Interface Parameters

Table 10-5: AXI Bridge Slave and Master Interface Parameters

Parameter	Description
<b>ID Width</b>	Controls the width of the thread ID of the master and slave interfaces.
<b>Write/Read/Combined Acceptance Capability</b>	Controls the depth of the FIFO that Qsys needs in the interconnect agents based on the maximum pending commands that the slave interface accepts.

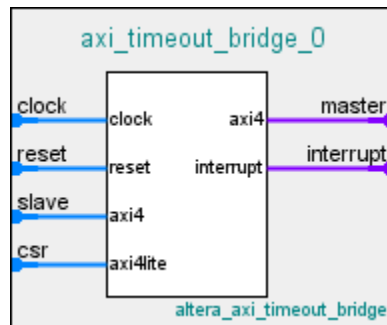
Parameter	Description
<b>Write/Read/Combined Issuing Capability</b>	Controls the depth of the FIFO that Qsys needs in the interconnect agents based on the maximum pending commands that the master interface issues. Issuing capability must follow acceptance capability to avoid unnecessary creation of FIFOs in the bridge.

**Note:** Maximum acceptance/issuing capability is a model-only parameter and does not influence the bridge HDL. The bridge does not backpressure when this limit is reached. Downstream components and/or the interconnect must apply backpressure.

## AXI Timeout Bridge

You can place an AXI Timeout Bridge between a single master and a single slave if you know that the slave may timeout and cause your system to hang. If a slave does not accept a command or respond to a command it accepted, its master can wait indefinitely. The AXI Timeout Bridge allows your system to recover when it hangs, and can also facilitate debugging.

**Figure 10-12: AXI Timeout Bridge**





For a domain with multiple masters and slaves, placement of an AXI Timeout Bridge in your design may be beneficial in the following scenarios:

- To recover from a hang, place the bridge near the slave. If the master attempts to communicate with a slave that hangs, the AXI Timeout Bridge frees the master by generating error responses. The master is then able to communicate with another slave.
- When debugging your system, place the AXI Timeout Bridge near the master. This placement enables you to identify the origin of the burst and to obtain the full address from the master. Additionally, placing an AXI Timeout Bridge near the master enables you to identify the target slave for the burst.

**Note:** If you put the bridge at the slave's side and you have multiple slaves connected to the same master, you do not get the full address.

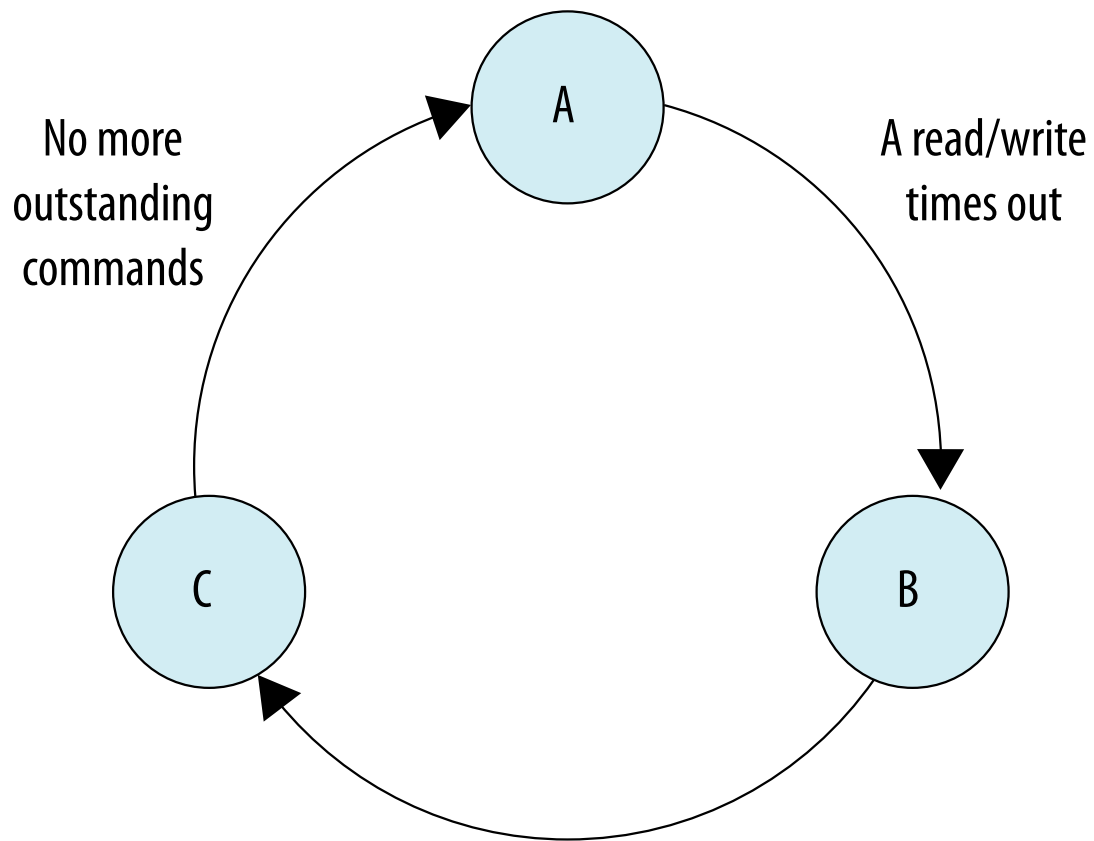
Figure 10-13: AXI Timeout Bridge Placement



### AXI Timeout Bridge Stages

A timeout occurs when the internal timer in the bridge exceeds the specified number of cycles within which a burst must complete from start to end.

Figure 10-14: AXI Bridge Timeout Bridge Stages



The AXI Timeout Bridge is notified that the slave is reset.

- A** Slave is functional - The bridge passes through all bursts.
- B** Slave is unresponsive - The bridge accepts commands and responds (with errors) to commands for the unresponsive slave. Commands are not passed through to the slave at this stage.
- C** Slave is reset - The bridge does not accept new commands, and responds only to commands that are outstanding.

When a timeout occurs, the AXI Timeout Bridge asserts an interrupt and reports the burst that caused the timeout to the CSR. The bridge then generates error responses back to the master on behalf of the unresponsive slave. This stage frees the master and certifies the unresponsive slave as dysfunctional. The AXI Timeout Bridge then accepts subsequent write addresses, write data and read addresses to the dysfunctional slave. The bridge does not accept outstanding write responses and read data from the dysfunctional slave are not passed through to the master. The `awvalid`, `wvalid`, `bready`, `arvalid`, and `rready` ports are "held low" at the master interface of the bridge.

**Note:** After a timeout, `awvalid`, `wvalid` and `arvalid` may be dropped before they are accepted by `awready` at the master interface. While the behavior violates the AXI specification, it occurs only on an interface connected to the slave which has been certified dysfunctional by the AXI Timeout Bridge.

Write channel refers to the AXI write address, data and response channels. Similarly, read channel refers to the AXI read address and data channels. AXI write and read channels are independent of each other. However, when a timeout occurs on either channel, the bridge generates error responses on both channels.

**Table 10-6: Burst Start and End Definitions for the AXI Timeout Bridge**

Channel	Start	End
Write	When an address is issued. First cycle of <code>awvalid</code> , even if data of the same burst is issued before the address (first cycle of <code>wvalid</code> ).	When the response is issued. First cycle of <code>bvalid</code> .
Read	When an address is issued. First cycle of <code>arvalid</code> .	When the last data is issued. First cycle of <code>rvalid</code> and <code>rlast</code> .

The AXI Timeout Bridge has four required interfaces: Master, Slave, Configuration and Status Register (CSR) (AXI4-Lite), and Interrupt. Qsys allows the AXI Timeout bridge to connect to any AXI3, AXI4, or Avalon master or slave interface. Avalon masters must to utilize the bridge's interrupt output to detect a timeout.

The bridge slave interface accepts write addresses, write data, and read addresses, and then generates the `SLVERR` response at the write response and read data channels. You should not expect to use `buser`, `rdata` and `ruser` at this stage of processing.

To resume normal operation, the dysfunctional slave must be reset and the bridge notified of the change in status via the CSR. Once the CSR notifies the bridge that the slave is ready, the bridge does not accept new commands until all outstanding bursts are responded to with an error response.

The CSR has a 4-bit address width, and a 32-bit data width. The CSR reports status and address information when the bridge asserts an interrupt.



Table 10-7: CSR Interrupt Status Information for the AXI Timeout Bridge

Address	Attribute	Name	Description
0x0	write-only	Slave is reset	Write a 1 to notify the AXI Timeout Bridge that the slave is reset and ready. Clears the interrupt.
0x4	read-only	Timed out operation	The operation of the burst that caused the timeout. 1 for a write; 0 for a read.
0x8 through 0xf	read-only	Timed out address	The address of the burst that caused the timeout. For an address width greater than 32-bits, CSR reads addresses 0x8 and 0xc to obtain the complete address.

## AXI Timeout Bridge Parameters

Table 10-8: AXI Timeout Bridge Parameters

Parameter	Description
<b>ID width</b>	The width of <code>awid</code> , <code>bid</code> , <code>arid</code> , or <code>rid</code> .
<b>Address width</b>	The width of <code>awaddr</code> or <code>araddr</code> .
<b>Data width</b>	The width of <code>wdata</code> or <code>rdata</code> .
<b>User width</b>	The width of <code>awuser</code> , <code>wuser</code> , <code>buser</code> , <code>aruser</code> , or <code>ruser</code> .
<b>Maximum number of outstanding writes</b>	The expected maximum number of outstanding writes.
<b>Maximum number of outstanding reads</b>	The expected maximum number of outstanding reads.
<b>Maximum number of cycles</b>	The number of cycles within which a burst must complete.

## Address Span Extender

The Address Span Extender creates a windowed bridge and allows memory-mapped master interfaces to access a larger or smaller address map than the width of their address signals allow. With an address span extender, a restricted master can access a broader address range. The address span extender splits the addressable space into multiple separate windows so that the master can access the appropriate part of the memory through the window.

The address span extender does not limit master and slave widths to a 32-bit and 64-bit configuration. You can use the address span extender for other width configurations. The address span extender supports 1-64 bit address windows.

If a processor can address only 2GB of an address span, and your system contains 4GB of memory, the address span extender can provide two 2GB windows in the 4GB memory address space. This issue sometimes occurs with Altera SoC devices. For example, an HPS subsystem in an SoC device can address only 1GB of an address span within the FPGA using the HPS-to-FPGA bridge. The address span extender enables the SoC device to address all of the address space in the FPGA using multiple 1GB windows.


## CTRL Register Layout

The control registers consist of a 64-bit register for each window. You write the base address that you want for each window to its corresponding control register. For example, if `CTRL_BASE` is the base address of the address span extender's control register, and there are two windows (0 and 1), then window 0's control register starts at `CTRL_BASE`, and window 1's control register starts at `CTRL_BASE + 8` (using byte addresses).

## Calculating the Address Span Extender Slave Address

The diagram below describes how Qsys calculates the slave address. In this example the address span extender is configured with a 28-bit address space for slaves. The lower 26 bits (bits 0 to 25 or `[25:0]`) is the offset into a particular window and originate from the address span extender's data port. The upper 2 bits `[27:26]` originate from the control registers.

Figure 10-15: Address Span Extender



### Using the Address Span Extender

When you implement the address span extender in Qsys, you must know the amount of address space the master uses (the size of the window), the total size of the addressable space (the number of windows), and how much address space (the size of the window) you want a particular slave to occupy in a master's address map.

This component supports 1 to 64 address windows. Qsys requires an assigned number of registers to hold the upper address bits for each window. In the parameter editor, you must select the number of bits in the expanded address map you want to access (**Expanded Master Byte Address Width**), the number of bits you want the master to see (**Slave Word Address Width**), and the number of sub-windows.

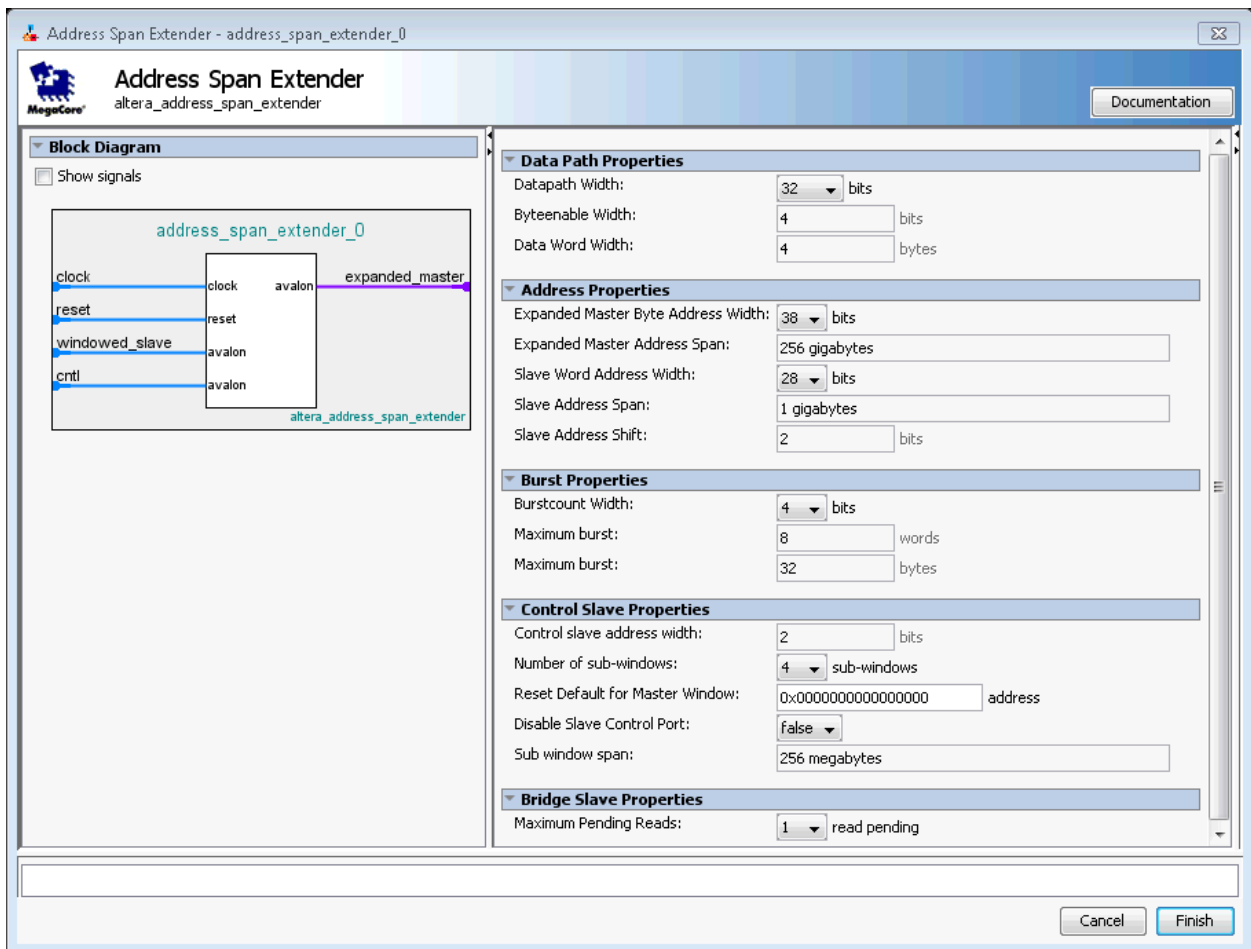
Each sub-window has a 64-bit register set that defines the sub window's upper address, and use only the bits greater than the slave byte address.

- **window 0**—expanded address [63:0]
- **window 1**—expanded address [63:0]

Qsys uses the upper bits of the slave address to pick which window to use. For example, if you specify 4 windows, then Qsys uses the top 2 bits of the slave address to specify window [0, 1, 2, 3]. Therefore having more windows does require the windows to be smaller, for example having 4 windows requires the windows themselves to be 1/4 the size of the slave address space. The total windowed address space is still equal to the original slave address space, but the windows allow access to memory regions in a larger overall address space.

In the parameter editor for the address span extender, you can click **Documentation** to obtain more information about the component.

**Figure 10-16: Address Span Extender Parameter Editor**



### Alternate Options for the Address Span Extender

You can set parameters for the address span extender with an initial fixed address value. Enter an address for the **Reset Default for Master Window** option, and select **True** for the **Disable Slave Control Port** option. This allows the address span extender to function as a fixed, non-programmable component.

Each sub-window is equal in size and stacks sequentially in the windowed slave interface's address space. To control the fixed address bits of a particular sub-window, you can write to the sub-window's register in the register control slave interface. Qsys structures the logic so that Qsys can optimize and remove bits that are not needed.

If **Burstcount Width** is greater than 1, Qsys processes the read burst in a single cycle, and assumes all byteenables are asserted on every cycle.

## NIOS II Support

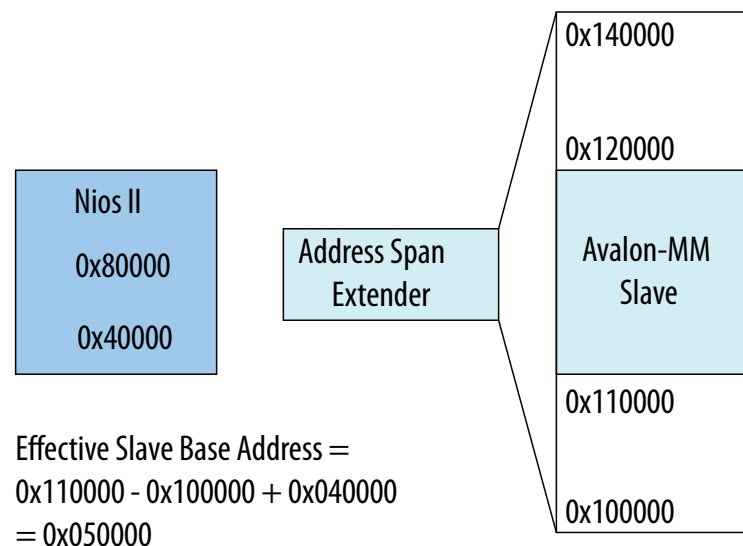
If the address span extender window is fixed, for example, the **Disable Slave Control Port** option is turned on, then the address span extender performs as a bridge. Components on the slave side of the address span extender that are within the window are visible to the NIOS II processor. Components partially within a window appear to NIOS II as if they have a reduced span. For example, a memory partially within a window appears as having a smaller size.

You can also use the address span extender to provide a window for the Nios II processor so that the HPS memory map is visible to NIOS II. In this way it is possible for the Nios II to communicate with HPS peripherals.

In the example below, a NIOS II processor has an address span extender from address 0x40000 to 0x80000. There is a window within the address span extender starting at 0x100000. Within the address span extender's address space there is a slave at base address 0x110000. The slave appears to NIOS II as being at address:

$$0x110000 - 0x100000 + 0x40000 = 0x050000$$

**Figure 10-17: NIOS II Support and the Address Span Extender**



If the address span extender window is dynamic. For example, when the **Disable Slave Control Port** option is turned off, the NIOS II processor is unable to see components on the slave side of the address span extender.

## AXI Default Slave

An AXI Default Slave provides a predictable error response service for master interfaces that send transactions that attempt to access an undefined memory region. This service guarantees an error response, should a master access a memory region that is not decoded to an instantiated slave. The error response service also helps to avoid unpredictable behavior in your system.

The default slave is an AXI3 component and displays in the IP Catalog as either **AXI Default Slave** or **Error Response Slave**.

AXI protocol requires that if the interconnect cannot successfully decode slave access, it must return the `DECERR` error response. Therefore, the default slave is required in AXI systems where the address space is not fully decoded to slave interfaces.

The default slave behaves like any other component in the system and is bound by translation and adaptation interconnect logic. An increase in resource usage may occur when a default slave connects to masters of different data widths, including Avalon or AXI-Lite masters.

You can connect clock, reset, and IRQ signals to a default slave, as well as AXI3 and AXI4 master interfaces without also instantiating a bridge. When you connect a default slave to a master, the default slave accepts cycles sent from the master, and returns the `DECERR` error response. On the AXI interface, the default slave supports only a read and write acceptance of 1, and does not support write data interleaving. The read and write channels are independent, and responses are returned when simultaneously targeted by a read and write cycle.

There is an optional interface on the default slave that supports CSR accesses for debug. CSR registers log the required information when returning an error response. When turned on, this channel acts as an Avalon interface with read and write channels with a fixed latency of 1.

To enable a slave interface as a default slave for a master interface in your system, you must connect the slave to the master in your Qsys system. You specify a default slave for a master by turning on the **Default Slave** column option in the **System Contents** tab. A system can contain more than one default slave. Altera recommends instantiating a separate default slave for each AXI master in your system.

For information about creating secure systems and accessing undefined memory regions, refer to *Creating a System with Qsys* in volume 1 of the *Quartus II Handbook*.

### Related Information

- [Creating a System with Qsys](#) on page 5-1

## AXI Default Slave Parameters

Figure 10-18: AXI Default Slave Parameter Editor




Table 10-9: AXI Default Slave Parameters

Parameter	Value	Description
<b>AXI master ID width</b>	1-8 bits	Determines the master ID width for error logging.
<b>AXI address width</b>	8-64 bits	Determines the address width for error logging. This value also affects the overall address width of the system, and should not exceed the maximum address width required in the system.
<b>AXI data width</b>	32, 64, or 128 bits	Determines the data width for error logging.
<b>Enable CSR Support (for error logging)</b>	On or Off	When turned on, instantiates an Avalon CSR interface for error logging.

Parameter	Value	Description
CSR Error Log Depth	1-16 bits	Depth of the transaction log, for example, the number of transactions the CSR logs for cycles with errors.
Register Avalon CSR inputs	On or Off	When turned on, controls debug access to the CSR interface.

## CSR Registers

When an access violation occurs, and the CSR port is enabled, the AXI Default Slave generates an interrupt and transfers the transaction information into the error log FIFO.

The error log count continues until the  $n^{\text{th}}$  log, where  $n$  is the log depth. When Qsys responds to the interrupt bit, it reads the register until the interrupt bit is no longer valid. The interrupt bit is valid as long as there is a valid bit in FIFO. A cleared interrupt bit is not affected by the FIFO status. When Qsys finishes reading the register, the access violation service is ready to receive new access violation requests. If an access violation occurs when FIFO is full, then an overflow bit is set, indicating more than  $n$  access violations have occurred, and some are not logged.

Qsys exits the access violation service after either the interrupt bit is no longer set, or when it determines that the access violation service has continued for too long.

## CSR Interrupt Status Registers

**Table 10-10: CSR Interrupt Status Registers**

For CSR register maps: Address = Memory Address Base + Offset.

Offset	Bit	Attribute	Default	Description
0x00	31:4	R0	0	Reserved.
	3	RW1C	0	<b>Read Access Violation Interrupt Overflow register</b> Asserted when a read access causes the Interconnect to return a <code>DECERR</code> response, and the buffer log depth is full. Indicates that there is a logging error lost due to an exceeded buffer log depth. Cleared by setting the bit to 1.
	2	RW1C	0	<b>Write Access Violation Interrupt Overflow register</b> Asserted when a write access causes the Interconnect to return a <code>DECERR</code> response, and the buffer log depth is full. Indicates that there is a logging error lost due to an exceeded buffer log depth. Cleared by setting the bit to 1.



Offset	Bit	Attribute	Default	Description
	1	RW1C	0	<b>Read Access Violation Interrupt register</b> Asserted when a read access causes the Interconnect to return a <code>DECERR</code> response. Cleared by setting the bit to 1. <b>Note:</b> Access violation are logged until the bit is cleared.
	0	RW1C	0	<b>Write Access Violation Interrupt register</b> Asserted when a write access causes the Interconnect to return a <code>DECERR</code> response. Cleared by setting the bit to 1. <b>Note:</b> Access violation are logged until the bit is cleared.

### CSR Read Access Violation Log

The CSR read access violation log settings are valid only when an associated read interrupt register is set. This set of registers should be read until the valid bit is cleared.

**Table 10-11: CSR Read Access Violation Log**

Offset	Bit	Attribute	Default	Description
0x100	31:13	R0	0	Reserved.
	12:11	R0	0	Indicates the burst type of the initiating cycle that causes the access violation.
	10:7	R0	0	Indicates the burst length of the initiating cycle that causes the access violation.
	6:4	R0	0	Indicates the burst size of the initiating cycle that causes the access violation.
	3:1	R0	0	Indicates the <code>PROT</code> of the initiating cycle that causes the access violation.
	0	R0	0	Read access violation log for the transaction is valid only when this bit is set. This bit is cleared when the interrupt register is cleared.
0x104	31:0	R0	0	Master ID for the cycle that causes the access violation.
0x108	31:0	R0	0	Read cycle target address for the cycle that causes the access violation (lower 32-bit).

Offset	Bit	Attribute	Default	Description
0x10C	31:0	R0	0	Read cycle target address for the cycle that causes the access violation (upper 32-bit). Valid only if widest address in system is larger than 32 bits.  <b>Note:</b> When this register is read, the current read access violation log is recovered from FIFO.

### CSR Write Access Violation Log

The CSR write access violation log settings are valid only when an associated read interrupt register is set. This set of registers should be read until the valid bit is cleared.

**Table 10-12: CSR Write Access Violation Log**

Offset	Bit	Attribute	Default	Description
0x190	31:13	R0	0	Reserved.
	12:11	R0	0	Indicates the burst type of the initiating cycle that causes the access violation.
	10:7	R0	0	Indicates the burst length of the initiating cycle that causes the access violation.
	6:4	R0	0	Indicates the burst size of the initiating cycle that causes the access violation.
	3:1	R0	0	Indicates the PROT of the initiating cycle that causes the access violation.
	0	R0	0	Write access violation log for the transaction is valid only when this bit is set. This bit is cleared when the interrupt register is cleared.
0x194	31:0	R0	0	Master ID for the cycle that causes the access violation.
0x198	31:0	R0	0	Write target address for the cycle that causes the access violation (lower 32-bit).
0x19C	31:0	R0	0	Write target address for the cycle that causes the access violation (upper 32-bit). Valid only if widest address in system is larger than 32 bits.

Offset	Bit	Attribute	Default	Description
0x1A0	31:0	R0	0	First 32 bits of the write data for the write cycle that causes the access violation.  <b>Note:</b> When this register is read, the current write access violation log is recovered from FIFO, when the data width is 32 bits.
0x1A4	31:0	R0	0	Bits [63:32] of the write data for the write cycle that causes the access violation. Valid only if the data width is greater than 32 -bits.
0x1A8	31:0	R0	0	Bits [95:64] of the write data for the write cycle that causes the access violation. Valid only if the data width is greater than 64 -bits.
0x1AC	31:0	R0	0	The first bits (127:96) of the write data for the write cycle that causes the access violation. Valid only if the data width is greater than 64 -bits.  <b>Note:</b> When this register is read, the current write access violation log is recovered from FIFO.

## Designating a Default Slave in the System Contents Tab

You can designate any slave in your Qsys system as the error response default slave. The designated default slave provides an error response service for masters that attempt access to an undefined memory region.


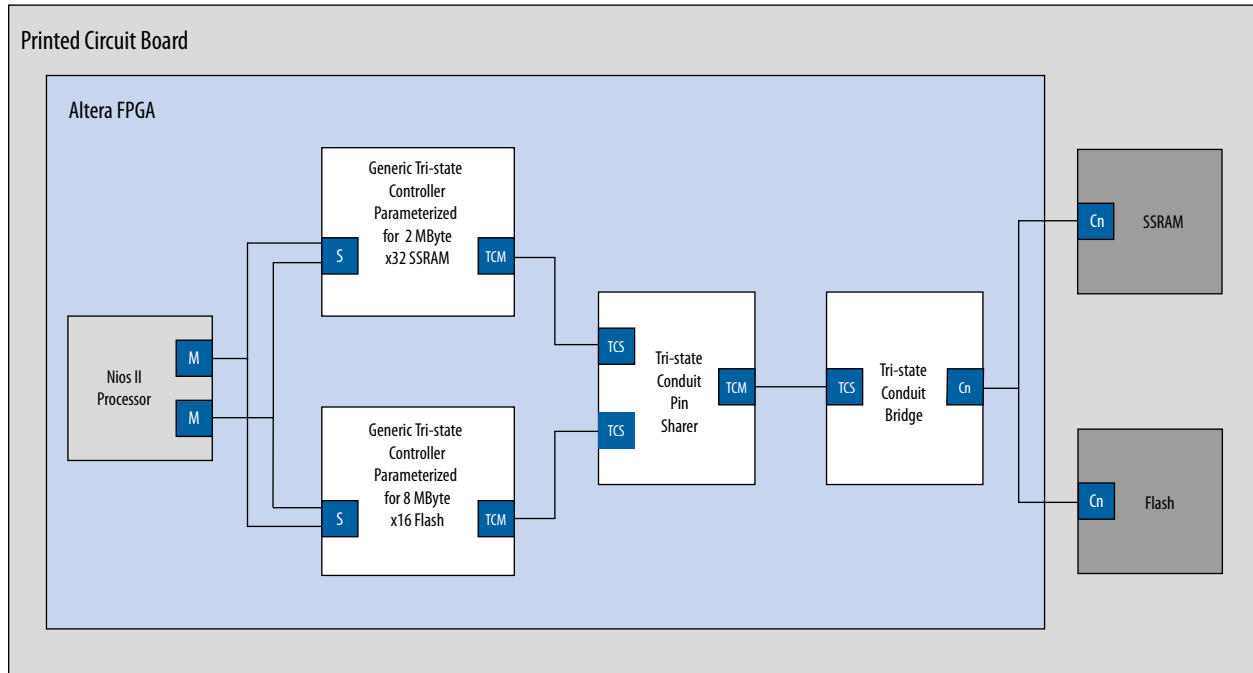
1. In your Qsys system, in the **System Contents** tab, right-click the header and turn on **Show Default Slave Column**.
2. Select the slave that you want to designate as the default slave, and then click the checkbox for the slave in the **Default Slave** column.
3. In the **System Contents** tab, in the **Connections** column, connect the designated default slave to one or more masters.

## Tri-State Components

The tri-state interface type allows you to design Qsys subsystems that connect to tri-state devices on your PCB. You can use tri-state components to implement pin sharing, convert between unidirectional and bidirectional signals, and create tri-state controllers for devices whose interfaces can be described using the tri-state signal types.

**Figure 10-19: Tri-State Conduit System to Control Off-Chip SRAM and Flash Devices**

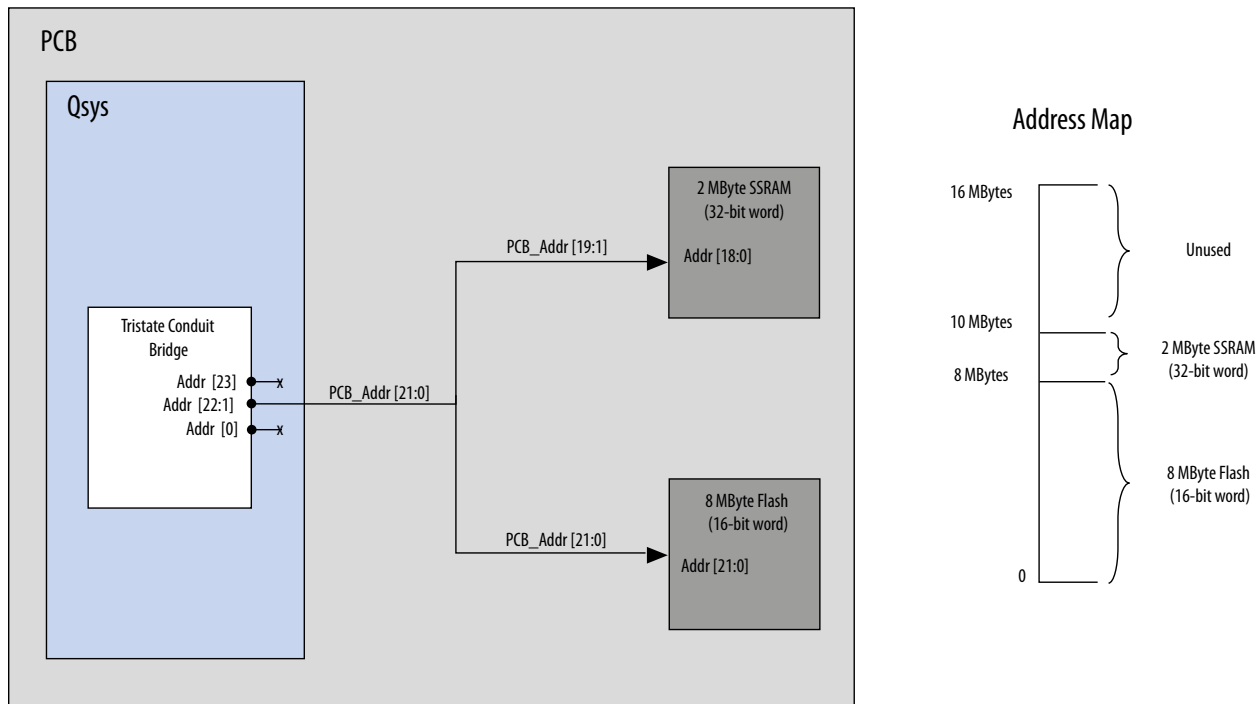
In this example, there are two generic Tri-State Conduit Controllers. The first is customized to control a flash memory. The second is customized to control an off-chip SSRAM. The Tri-State Conduit Pin Sharer multiplexes between these two controllers, and the Tri-State Conduit Bridge converts between an on-chip encoding of tri-state signals and true bidirectional signals. By default, the Tri-State Conduit Pin Sharer and Tri-State Conduit Bridge present byte addresses. Typically, each address location contains more than one byte of data.



**Figure 10-20: Address Connections from Qsys System to PCB**

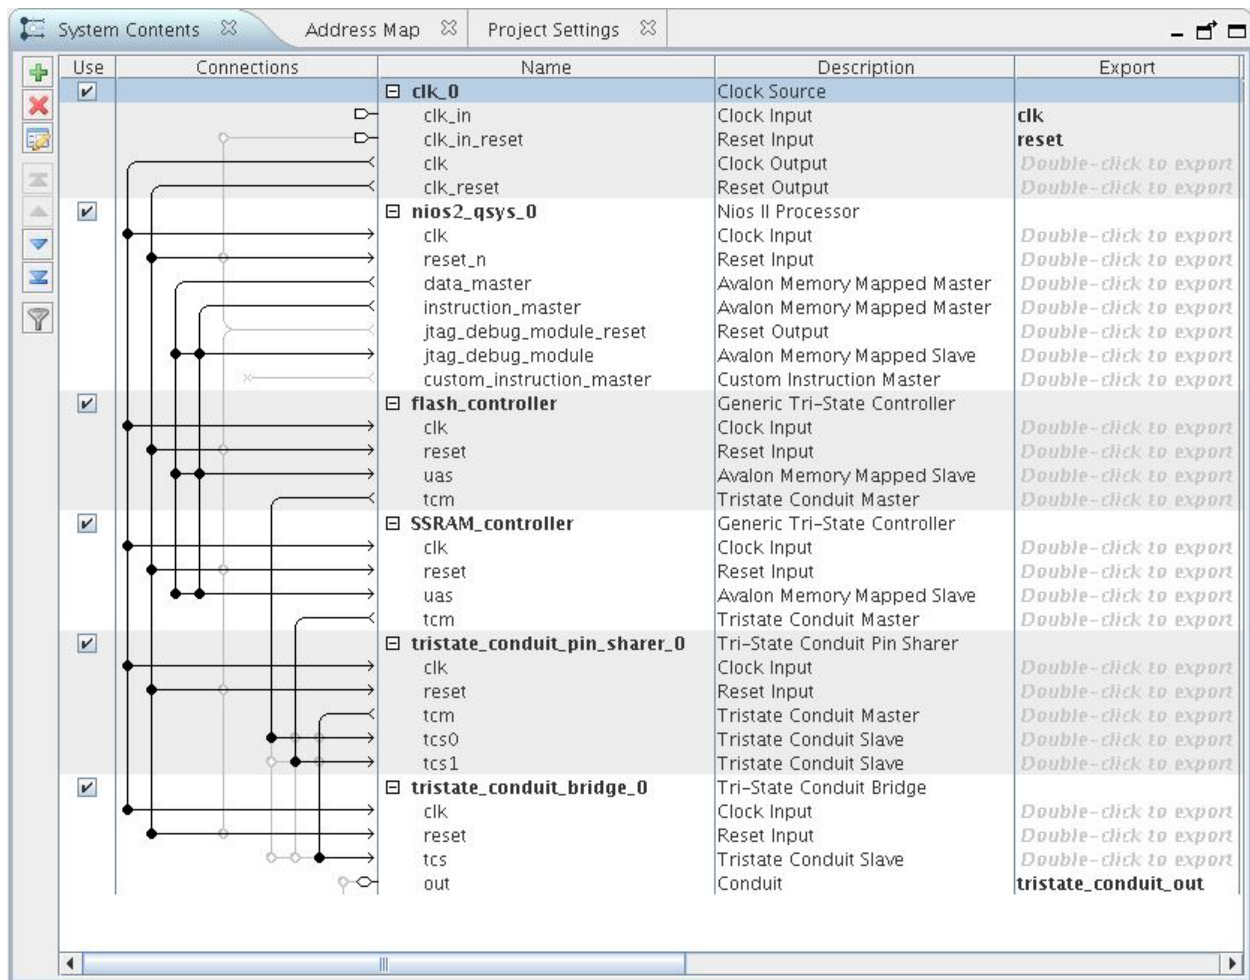
The flash device operates on 16-bit words and must ignore the least-significant bit of the Avalon-MM address, and shows `addr[0]` as not connected. The SSRAM memory operates on 32-bit words and must ignore the two, low-order memory bits. Because neither device requires a byte address, `addr[0]` is not routed on the PCB.

The flash device responds to address range 0 MBytes to 8 MBytes-1. The SSRAM responds to address range 8 MBytes to 10 MBytes-1. The PCB schematic for the PCB connects `addr[21:0]` to `addr[18:0]` of the SSRAM device because the SSRAM responds to 32-bit word address. The 8 MByte flash device accesses 16-bit words; consequently, the schematic does not connect `addr[0]`. The `chipselect` signals select between the two devices.



**Note:** If you create a custom tri-state conduit master with word aligned addresses, the Tri-state Conduit Pin Sharer does not change or align the address signals.

Figure 10-21: Tri-State Conduit System in Qsys



**Related Information**

- [Avalon Interface Specifications](#)
- [Avalon Tri-State Conduit Components User Guide](#)

**Generic Tri-State Controller**

The Generic Tri-State Controller provides a template for a controller. You can customize the tri-state controller with various parameters to reflect the behavior of an off-chip device. The following types of parameters are available for the tri-state controller:

- Width of the address and data signals
- Read and write wait times
- Bus-turnaround time
- Data hold time

**Note:** In calculating delays, the Generic Tri-State Controller chooses the larger of the bus-turnaround time and read latency. Turnaround time is measured from the time that a command is accepted, not from the time that the previous read returned data.

The Generic Tri-State Controller includes the following interfaces:

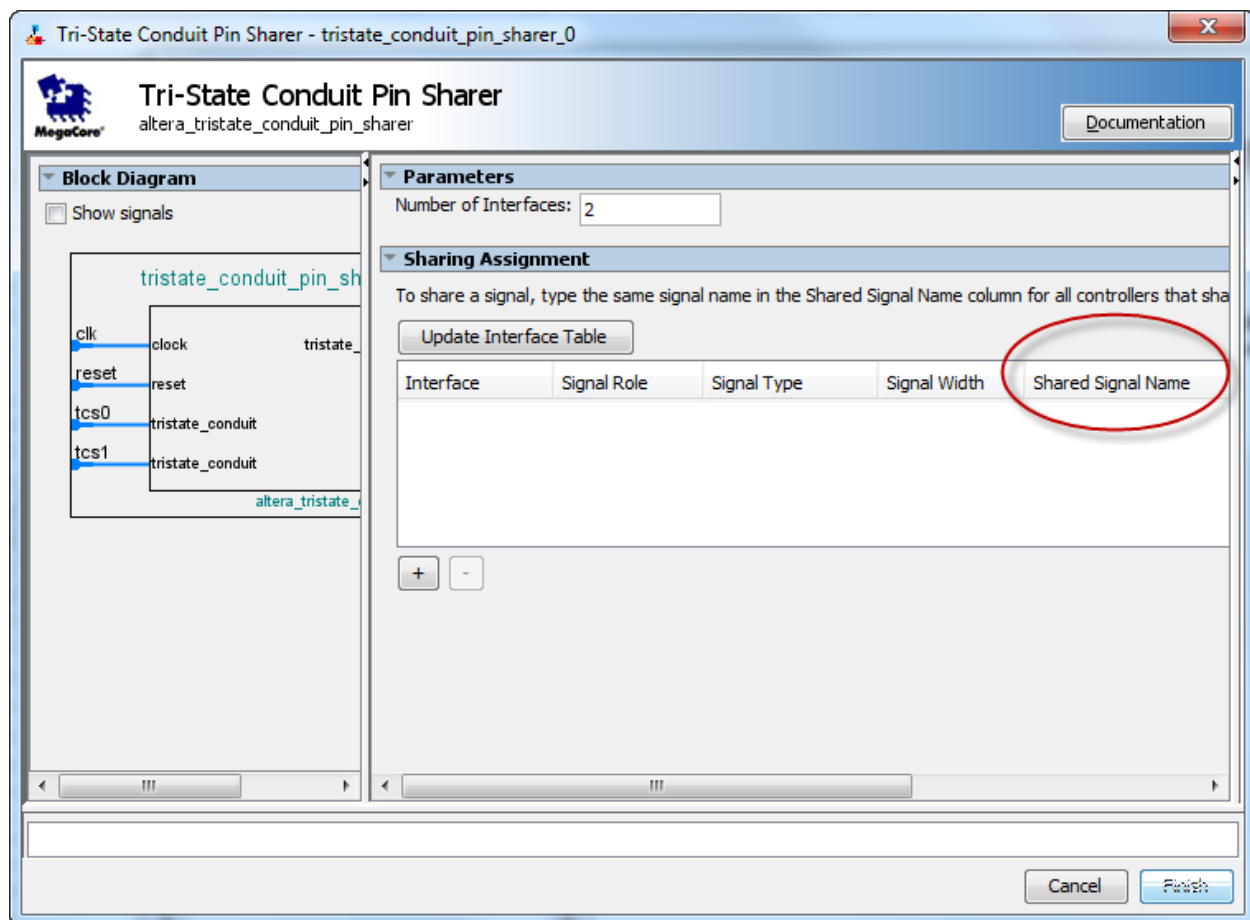
- **Memory-mapped slave interface**—This interface connects to an memory-mapped master, such as a processor.
- **Tristate Conduit Master interface**—Tri-state master interface usually connects to the tri-state conduit slave interface of the tri-state conduit pin sharer.
- **Clock sink**—The component's clock reference. You must connect this interface to a clock source.
- **Reset sink**—This interface connects to a reset source interface.

## Tri-State Conduit Pin Sharer

The Tri-state Conduit Pin Sharer multiplexes between the signals of the connected tri-state controllers. You connect all signals from the tri-state controllers to the Tri-state Conduit Pin Sharer and use the parameter editor to specify the signals that are shared.

**Figure 10-22: Tri-State Conduit Pin Sharer Parameter Editor**

The parameter editor includes a **Shared Signal Name** column. If the widths of shared signals differ, the signals are aligned on their 0<sup>th</sup> bit and the higher-order pins are driven to 0 whenever the smaller signal has control of the bus. Unshared signals always propagate through the pin sharer. The tri-state conduit pin sharer uses the round-robin arbiter to select between tri-state conduit controllers.



**Note:** All tri-state conduit components are connected to a pin sharer must be in the same clock domain.

**Related Information**

[Avalon-ST Round Robin Scheduler](#) on page 10-66

## Tri-State Conduit Bridge

The Tri-State Conduit Bridge instantiates bidirectional signals for each tri-state signal while passing all other signals straight through the component. The Tri-State Conduit Bridge registers all outgoing and incoming signals, which adds two cycles of latency for a read request. You must account for this additional pipelining when designing a custom controller. During reset, all outputs are placed in a high-impedance state. Outputs are enabled in the first clock cycle after reset is deasserted, and the output signals are then bidirectional.

## Test Pattern Generator and Checker Cores

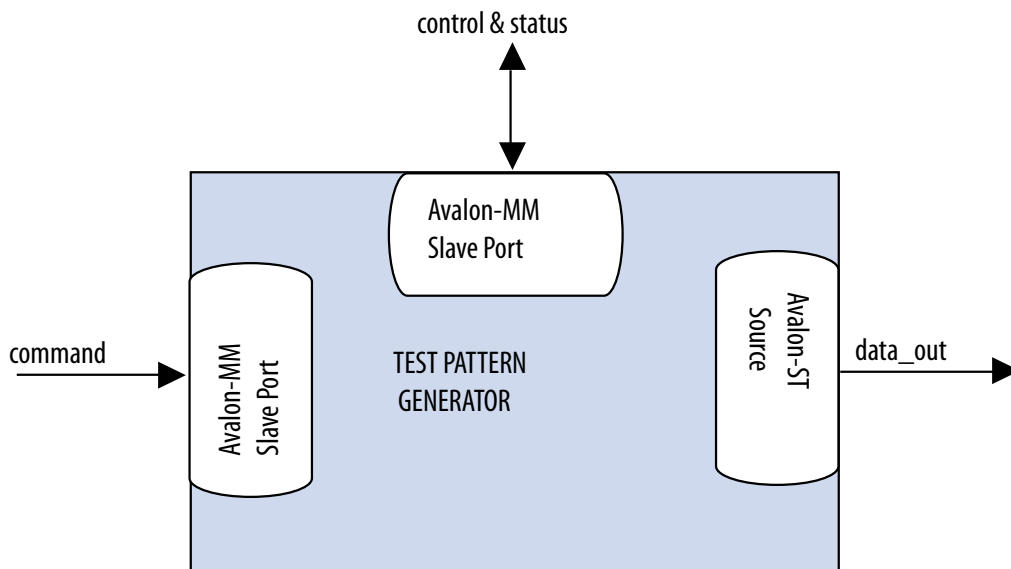
The data generation and monitoring solution for Avalon-ST consists of two components: a test pattern generator core that generates data, and sends it out on an Avalon-ST data interface, and a test pattern checker core that receives the same data and verifies it. Optionally, the data can be formatted as packets, with accompanying `start_of_packet` and `end_of_packet` signals.

The test pattern generator inserts different error conditions, and the test pattern checker reports these error conditions to the control interface, each via an Avalon Memory-Mapped (Avalon-MM) slave. The **Throttle Seed** is the starting value for the throttle control random number generator. Altera recommends a unique value for each instance of the test pattern generator and checker cores in a system.

## Test Pattern Generator

**Figure 10-23: Test Pattern Generator Core**

The test pattern generator core accepts commands to generate data via an Avalon-MM command interface, and drives the generated data to an Avalon-ST data interface. You can parameterize most aspects of the Avalon-ST data interface, such as the number of error bits and data signal width, thus allowing you to test components with different interfaces.





The data pattern is calculated as:  $Symbol\ Value = Symbol\ Position\ in\ Packet \oplus Data\ Error\ Mask$ . Data that is not organized in packets is a single stream with no beginning or end. The test pattern generator has a throttle register that is set via the Avalon-MM control interface. The test pattern generator uses the value of the throttle register in conjunction with a pseudo-random number generator to throttle the data generation rate.

### Test Pattern Generator Command Interface

The command interface for the Test Pattern Generator is a 32-bit Avalon-MM write slave that accepts data generation commands. It is connected to a 16-element deep FIFO, thus allowing a master peripheral to drive a number of commands into the test pattern generator.

The command interface maps to the following registers: `cmd_lo` and `cmd_hi`. The command is pushed into the FIFO when the register `cmd_lo` (address 0) is addressed. When the FIFO is full, the command interface asserts the `waitrequest` signal. You can create errors by writing to the register `cmd_hi` (address 1). The errors are cleared when 0 is written to this register, or its respective fields.

### Test Pattern Generator Control and Status Interface

The control and status interface of the Test Pattern Generator is a 32-bit Avalon-MM slave that allows you to enable or disable the data generation, as well as set the throttle. This interface also provides generation-time information, such as the number of channels and whether or not data packets are supported.

### Test Pattern Generator Output Interface

The output interface of the Test Pattern Generator is an Avalon-ST interface that optionally supports data packets. You can configure the output interface to align with your system requirements. Depending on the incoming stream of commands, the output data may contain interleaved packet fragments for different channels. To keep track of the current symbol's position within each packet, the test pattern generator maintains an internal state for each channel.

You can configure the output interface of the test pattern generator with the following parameters:

- **Number of Channels**—Number of channels that the test pattern generator supports. Valid values are 1 to 256.
- **Data Bits Per Symbol**—Bits per symbol is related to the width of `readdata` and `writedata` signals, which must be a multiple of the bits per symbol.
- **Data Symbols Per Beat**—Number of symbols (words) that are transferred per beat. Valid values are 1 to 256.
- **Include Packet Support**—Indicates whether or not packet transfers are supported. Packet support includes the `startofpacket`, `endofpacket`, and `empty` signals.
- **Error Signal Width (bits)**—Width of the error signal on the output interface. Valid values are 0 to 31. A value of 0 indicates that the error signal is not in use.

**Note:** If you change only bits per symbol, and do not change the data width, errors are generated.

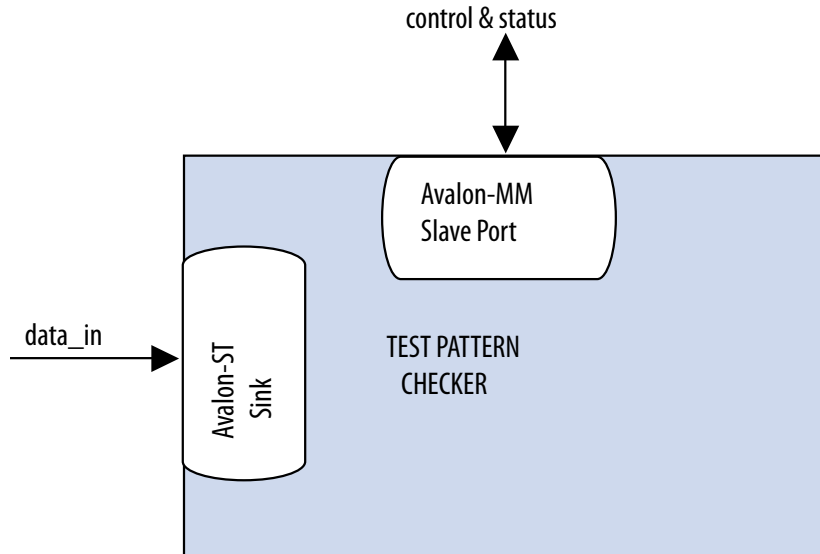
### Test Pattern Generator Functional Parameter

The Test Pattern Generator functional parameter allows you to configure the test pattern generator as a whole system.

## Test Pattern Checker

Figure 10-24: Test Pattern Checker

The test pattern checker core accepts data via an Avalon-ST interface and verifies it against the same predetermined pattern that the test pattern generator uses to produce the data. The test pattern checker core reports any exceptions to the control interface. You can parameterize most aspects of the test pattern checker's Avalon-ST interface such as the number of error bits and the data signal width. This enables the ability to test components with different interfaces. The test pattern checker has a throttle register that is set via the Avalon-MM control interface. The value of the throttle register controls the rate at which data is accepted.



The test pattern checker detects exceptions and reports them to the control interface via a 32-element deep internal FIFO. Possible exceptions are data error, missing start-of-packet (SOP), missing end-of-packet (EOP), and signaled error.

As each exception occurs, an exception descriptor is pushed into the FIFO. If the same exception occurs more than once consecutively, only one exception descriptor is pushed into the FIFO. All exceptions are ignored when the FIFO is full. Exception descriptors are deleted from the FIFO after they are read by the control and status interface.

### Test Pattern Checker Input Interface

The Test Pattern Checker input interface is an Avalon-ST interface that optionally supports data packets. You can configure the input interface to align with your system requirements. Incoming data may contain interleaved packet fragments. To keep track of the current symbol's position, the test pattern checker maintains an internal state for each channel.

### Test Pattern Checker Control and Status Interface

The Test Pattern Checker control and status interface is a 32-bit Avalon-MM slave that allows you to enable or disable data acceptance, as well as set the throttle. This interface provides generation-time information, such as the number of channels and whether the test pattern checker supports data packets. The control and status interface also provides information on the exceptions detected by the test pattern checker. The interface obtains this information by reading from the exception FIFO.

## Test Pattern Checker Functional Parameter

The Test Pattern Checker functional parameter allows you to configure the test pattern checker as a whole system.

### Test Pattern Checker Input Parameters

You can configure the input interface of the test pattern checker using the following parameters:

- **Data Bits Per Symbol**—Bits per symbol is related to the width of `readdata` and `writedata` signals, which must be a multiple of the bits per symbol.
- **Data Symbols Per Beat**—Number of symbols (words) that are transferred per beat. Valid values are 1 to 32.
- **Include Packet Support**—Indicates whether or not data packet transfers are supported. Packet support includes the `startofpacket`, `endofpacket`, and `empty` signals.
- **Number of Channels**—Number of channels that the test pattern checker supports. Valid values are 1 to 256.
- **Error Signal Width (bits)**—Width of the `error` signal on the input interface. Valid values are 0 to 31. A value of 0 indicates that the `error` signal is not in use.

**Note:** If you change only bits per symbol, and do not change the data width, errors are generated.

## Software Programming Model for the Test Pattern Generator and Checker Cores

The HAL system library support, software files, and register maps describe the software programming model for the test pattern generator and checker cores.

### HAL System Library Support

For Nios II processor users, Altera provides HAL system library drivers that allow you to initialize and access the test pattern generator and checker cores. Altera recommends you to use the provided drivers to access the cores instead of accessing the registers directly.

For Nios II IDE users, copy the provided drivers from the following installation folders to your software application directory:

- **<IP installation directory>/ip/sopc\_builder\_ip/altera\_avalon\_data\_source/HAL**
- **<IP installation directory>/ip/sopc\_builder\_ip/altera\_avalon\_data\_sink/HAL**

**Note:** This instruction does not apply if you use the Nios II command-line tools.

### Test Pattern Generator and Test Pattern Checker Core Files

The following files define the low-level access to the hardware, and provide the routines for the HAL device drivers.

**Note:** Do not modify the test pattern generator or test pattern checker core files.

- Test pattern generator core files:
  - **data\_source\_regs.h**—Header file that defines the test pattern generator's register maps.
  - **data\_source\_util.h** , **data\_source\_util.c**—Header and source code for the functions and variables required to integrate the driver into the HAL system library.
- Test pattern checker core files:
  - **data\_sink\_regs.h**—Header file that defines the core's register maps.
  - **data\_sink\_util.h** , **data\_sink\_util.c**—Header and source code for the functions and variables required to integrate the driver into the HAL system library.

## Register Maps for the Test Pattern Generator and Test Pattern Checker Cores

### Test Pattern Generator Control and Status Registers

**Table 10-13: Test Pattern Generator Control and Status Register Map**

Shows the offset for the test pattern generator control and status registers. Each register is 32-bits wide.

Offset	Register Name
base + 0	status
base + 1	control
base + 2	fill

**Table 10-14: Test Pattern Generator Status Register Bits**

Bit(s)	Name	Access	Description
[15:0]	ID	RO	A constant value of 0x64.
[23:16]	NUMCHANNELS	RO	The configured number of channels.
[30:24]	NUMSYMBOLS	RO	The configured number of symbols per beat.
[31]	SUPPORTPACKETS	RO	A value of 1 indicates data packet support.

**Table 10-15: Test Pattern Generator Control Register Bits**

Bit(s)	Name	Access	Description
[0]	ENABLE	RW	Setting this bit to 1 enables the test pattern generator core.
[7:1]	Reserved		

Bit(s)	Name	Access	Description
[16:8]	THROTTLE	RW	Specifies the throttle value which can be between 0–256, inclusively. The test pattern generator uses this value in conjunction with a pseudo-random number generator to throttle the data generation rate.  Setting THROTTLE to 0 stops the test pattern generator core. Setting it to 256 causes the test pattern generator core to run at full throttle. Values between 0–256 result in a data rate proportional to the throttle value.
[17]	SOFT RESET	RW	When this bit is set to 1, all internal counters and statistics are reset. Write 0 to this bit to exit reset.
[31:18]	Reserved		

Table 10-16: Test Pattern Generator Fill Register Bits

Bit(s)	Name	Access	Description
[0]	BUSY	RO	A value of 1 indicates that data transmission is in progress, or that there is at least one command in the command queue.
[6:1]	Reserved		
[15:7]	FILL	RO	The number of commands currently in the command FIFO.
[31:16]	Reserved		

### Test Pattern Generator Command Registers

Table 10-17: Test Pattern Generator Command Register Map

Shows the offset for the command registers. Each register is 32-bits wide.

Offset	Register Name
base + 0	cmd_lo
base + 1	cmd_hi

The `cmd_lo` is pushed into the FIFO only when the `cmd_lo` register is addressed.

**Table 10-18: cmd\_lo Register Bits**

Bit(s)	Name	Access	Description
[15:0]	SIZE	RW	The segment size in symbols. Except for the last segment in a packet, the size of all segments must be a multiple of the configured number of symbols per beat. If this condition is not met, the test pattern generator core inserts additional symbols to the segment to ensure the condition is fulfilled.
[29:16]	CHANNEL	RW	The channel to send the segment on. If the <code>channel</code> signal is less than 14 bits wide, the test pattern generator uses the low order bits of this register to drive the signal.
[30]	SOP	RW	Set this bit to 1 when sending the first segment in a packet. This bit is ignored when data packets are not supported.
[31]	EOP	RW	Set this bit to 1 when sending the last segment in a packet. This bit is ignored when data packets are not supported.

**Table 10-19: cmd\_hi Register Bits**

Bit(s)	Name	Access	Description
[15:0]	SIGNALLED ERROR	RW	Specifies the value to drive the <code>error</code> signal. A non-zero value creates a signalled error.
[23:16]	DATA ERROR	RW	The output data is XORed with the contents of this register to create data errors. To stop creating data errors, set this register to 0.
[24]	SUPPRESS SOP	RW	Set this bit to 1 to suppress the assertion of the <code>startofpacket</code> signal when the first segment in a packet is sent.
[25]	SUPPRESS EOP	RW	Set this bit to 1 to suppress the assertion of the <code>endofpacket</code> signal when the last segment in a packet is sent.

**Test Pattern Checker Control and Status Registers**

**Table 10-20: Test Pattern Checker Control and Status Register Map**

Shows the offset for the control and status registers. Each register is 32 bits wide.

Offset	Register Name
base + 0	status
base + 1	control

Offset	Register Name
base + 2	Reserved
base + 3	
base + 4	
base + 5	exception_descriptor
base + 6	indirect_select
base + 7	indirect_count

Table 10-21: Test Pattern Checker Status Register Bits

Bit(s)	Name	Access	Description
[15:0]	ID	RO	Contains a constant value of 0x65.
[23:16]	NUMCHANNELS	RO	The configured number of channels.
[30:24]	NUMSYMBOLS	RO	The configured number of symbols per beat.
[31]	SUPPORTPACKETS	RO	A value of 1 indicates packet support.

Table 10-22: Test Pattern Checker Control Register Bits

Bit(s)	Name	Access	Description
[0]	ENABLE	RW	Setting this bit to 1 enables the test pattern checker.
[7:1]	Reserved		
[16:8]	THROTTLE	RW	Specifies the throttle value which can be between 0–256, inclusively. Qsys uses this value in conjunction with a pseudo-random number generator to throttle the data generation rate.  Setting THROTTLE to 0 stops the test pattern generator core. Setting it to 256 causes the test pattern generator core to run at full throttle. Values between 0–256 result in a data rate proportional to the throttle value.
[17]	SOFT RESET	RW	When this bit is set to 1, all internal counters and statistics are reset. Write 0 to this bit to exit reset.
[31:18]	Reserved		

If there is no exception, reading the `exception_descriptor` register bit register returns 0.

**Table 10-23: `exception_descriptor` Register Bits**

Bit(s)	Name	Access	Description
[0]	DATA_ERROR	RO	A value of 1 indicates that an error is detected in the incoming data.
[1]	MISSINGSOP	RO	A value of 1 indicates missing start-of-packet.
[2]	MISSINGEOP	RO	A value of 1 indicates missing end-of-packet.
[7:3]	Reserved		
[15:8]	SIGNALLED_ERROR	RO	The value of the <code>error</code> signal.
[23:16]	Reserved		
[31:24]	CHANNEL	RO	The channel on which the exception was detected.

**Table 10-24: `indirect_select` Register Bits**

Bit	Bits Name	Access	Description
[7:0]	INDIRECT_CHANNEL	RW	Specifies the channel number that applies to the <code>INDIRECT_PACKET_COUNT</code> , <code>INDIRECT_SYMBOL_COUNT</code> , and <code>INDIRECT_ERROR_COUNT</code> registers.
[15:8]	Reserved		
[31:16]	INDIRECT_ERROR	RO	The number of data errors that occurred on the channel specified by <code>INDIRECT_CHANNEL</code> .

**Table 10-25: `indirect_count` Register Bits**

Bit	Bits Name	Access	Description
[15:0]	INDIRECT_PACKET_COUNT	RO	The number of data packets received on the channel specified by <code>INDIRECT_CHANNEL</code> .
[31:16]	INDIRECT_SYMBOL_COUNT	RO	The number of symbols received on the channel specified by <code>INDIRECT_CHANNEL</code> .



## Test Pattern Generator API

The following subsections describe application programming interface (API) for the test pattern generator.

**Note:** API functions are currently not available from the interrupt service routine (ISR).

[data\\_source\\_reset\(\)](#) on page 10-46

[data\\_source\\_init\(\)](#) on page 10-47

[data\\_source\\_get\\_id\(\)](#) on page 10-47

[data\\_source\\_get\\_supports\\_packets\(\)](#) on page 10-48

[data\\_source\\_get\\_num\\_channels\(\)](#) on page 10-48

[data\\_source\\_get\\_symbols\\_per\\_cycle\(\)](#) on page 10-48

[data\\_source\\_get\\_enable\(\)](#) on page 10-49

[data\\_source\\_set\\_enable\(\)](#) on page 10-49

[data\\_source\\_get\\_throttle\(\)](#) on page 10-50

[data\\_source\\_set\\_throttle\(\)](#) on page 10-50

[data\\_source\\_is\\_busy\(\)](#) on page 10-51

[data\\_source\\_fill\\_level\(\)](#) on page 10-51

[data\\_source\\_send\\_data\(\)](#) on page 10-51

### data\_source\_reset()

Table 10-26: data\_source\_reset()

Information Type	Description
<b>Prototype</b>	<code>void data_source_reset(alt_u32 base);</code>
<b>Thread-safe</b>	No
<b>Include</b>	<code>&lt;data_source_util.h &gt;</code>
<b>Parameters</b>	<code>base</code> —Base address of the control and status slave.
<b>Returns</b>	<code>void</code>
<b>Description</b>	Resets the test pattern generator core including all internal counters and FIFOs. The control and status registers are not reset by this function.

## data\_source\_init()

Table 10-27: data\_source\_init()

Information Type	Description
<b>Prototype</b>	<code>int data_source_init(alt_u32 base, alt_u32 command_base);</code>
<b>Thread-safe</b>	No
<b>Include</b>	< <i>data_source_util.h</i> >
<b>Parameters</b>	base—Base address of the control and status slave. command_base—Base address of the command slave.
<b>Returns</b>	1—Initialization is successful. 0—Initialization is unsuccessful.
<b>Description</b>	Performs the following operations to initialize the test pattern generator core: <ul style="list-style-type: none"> <li>Resets and disables the test pattern generator core.</li> <li>Sets the maximum throttle.</li> <li>Clears all inserted errors.</li> </ul>

## data\_source\_get\_id()

Table 10-28: data\_source\_get\_id()

Information Type	Description
<b>Prototype</b>	<code>int data_source_get_id(alt_u32 base);</code>
<b>Thread-safe</b>	Yes
<b>Include</b>	< <i>data_source_util.h</i> >
<b>Parameters</b>	base—Base address of the control and status slave.
<b>Returns</b>	Test pattern generator core identifier.
<b>Description</b>	Retrieves the test pattern generator core's identifier.

**data\_source\_get\_supports\_packets()**

Table 10-29: data\_source\_get\_supports\_packets()

Information Type	Description
<b>Prototype</b>	<code>int data_source_init(alt_u32 base);</code>
<b>Thread-safe</b>	Yes
<b>Include</b>	<code>&lt;data_source_util.h &gt;</code>
<b>Parameters</b>	<code>base</code> —Base address of the control and status slave.
<b>Returns</b>	1—Data packets are supported. 0—Data packets are not supported.
<b>Description</b>	Checks if the test pattern generator core supports data packets.

**data\_source\_get\_num\_channels()**

Table 10-30: data\_source\_get\_num\_channels()

Description	Description
<b>Prototype</b>	<code>int data_source_get_num_channels(alt_u32 base);</code>
<b>Thread-safe</b>	Yes
<b>Include</b>	<code>&lt;data_source_util.h &gt;</code>
<b>Parameters</b>	<code>base</code> —Base address of the control and status slave.
<b>Returns</b>	Number of channels supported.
<b>Description</b>	Retrieves the number of channels supported by the test pattern generator core.

**data\_source\_get\_symbols\_per\_cycle()**

Table 10-31: data\_source\_get\_symbols\_per\_cycle()

Description	Description
<b>Prototype</b>	<code>int data_source_get_symbols(alt_u32 base);</code>
<b>Thread-safe</b>	Yes

Description	Description
<b>Include</b>	< <i>data_source_util.h</i> >
<b>Parameters</b>	base—Base address of the control and status slave.
<b>Returns</b>	Number of symbols transferred in a beat.
<b>Description</b>	Retrieves the number of symbols transferred by the test pattern generator core in each beat.

### data\_source\_get\_enable()

Table 10-32: data\_source\_get\_enable()

Information Type	Description
<b>Prototype</b>	<code>int data_source_get_enable(alt_u32 base);</code>
<b>Thread-safe</b>	Yes
<b>Include</b>	< <i>data_source_util.h</i> >
<b>Parameters</b>	base—Base address of the control and status slave.
<b>Returns</b>	Value of the ENABLE bit.
<b>Description</b>	Retrieves the value of the ENABLE bit.

### data\_source\_set\_enable()

Table 10-33: data\_source\_set\_enable()

Information Type	Description
<b>Prototype</b>	<code>void data_source_set_enable(alt_u32 base, alt_u32 value);</code>
<b>Thread-safe</b>	No
<b>Include</b>	< <i>data_source_util.h</i> >
<b>Parameters</b>	base—Base address of the control and status slave. value—ENABLE bit set to the value of this parameter.
<b>Returns</b>	void

Information Type	Description
<b>Description</b>	Enables or disables the test pattern generator core. When disabled, the test pattern generator core stops data transmission but continues to accept commands and stores them in the FIFO

## data\_source\_get\_throttle()

Table 10-34: data\_source\_get\_throttle()

Information Type	Description
<b>Prototype</b>	<code>int data_source_get_throttle(alt_u32 base);</code>
<b>Thread-safe</b>	Yes
<b>Include</b>	< <i>data_source_util.h</i> >
<b>Parameters</b>	<code>base</code> —Base address of the control and status slave.
<b>Returns</b>	Throttle value.
<b>Description</b>	Retrieves the current throttle value.

## data\_source\_set\_throttle()

Table 10-35: data\_source\_set\_throttle()

Information Type	Description
<b>Prototype</b>	<code>void data_source_set_throttle(alt_u32 base, alt_u32 value);</code>
<b>Thread-safe</b>	No
<b>Include</b>	< <i>data_source_util.h</i> >
<b>Parameters</b>	<code>base</code> —Base address of the control and status slave. <code>value</code> —Throttle value.
<b>Returns</b>	void
<b>Description</b>	Sets the throttle value, which can be between 0–256 inclusively. The throttle value, when divided by 256 yields the rate at which the test pattern generator sends data.

## data\_source\_is\_busy()

Table 10-36: data\_source\_is\_busy()

Information Type	Description
<b>Prototype</b>	<code>int data_source_is_busy(alt_u32 base);</code>
<b>Thread-safe</b>	Yes
<b>Include</b>	<code>&lt;data_source_util.h &gt;</code>
<b>Parameters</b>	<code>base</code> —Base address of the control and status slave.
<b>Returns</b>	1—Test pattern generator core is busy. 0—Test pattern generator core is not busy.
<b>Description</b>	Checks if the test pattern generator is busy. The test pattern generator core is busy when it is sending data or has data in the command FIFO to be sent.

## data\_source\_fill\_level()

Table 10-37: data\_source\_fill\_level()

Information Type	Description
<b>Prototype</b>	<code>int data_source_fill_level(alt_u32 base);</code>
<b>Thread-safe</b>	Yes
<b>Include</b>	<code>&lt;data_source_util.h &gt;</code>
<b>Parameters</b>	<code>base</code> —Base address of the control and status slave.
<b>Returns</b>	Number of commands in the command FIFO.
<b>Description</b>	Retrieves the number of commands currently in the command FIFO.

## data\_source\_send\_data()

Table 10-38: data\_source\_send\_data()

Information Type	Description
<b>Prototype</b>	<code>int data_source_send_data(alt_u32 cmd_base, alt_u16 channel, alt_u16 size, alt_u32 flags, alt_u16 error, alt_u8 data_error_mask);</code>

Information Type	Description
<b>Thread-safe</b>	No
<b>Include</b>	< <i>data_source_util.h</i> >
<b>Parameters</b>	<p><code>cmd_base</code>—Base address of the command slave.</p> <p><code>channel</code>—Channel to send the data.</p> <p><code>size</code>—Data size.</p> <p><code>flags</code> —Specifies whether to send or suppress SOP and EOP signals. Valid values are <code>DATA_SOURCE_SEND_SOP</code>, <code>DATA_SOURCE_SEND_EOP</code>, <code>DATA_SOURCE_SEND_SUPPRESS_SOP</code> and <code>DATA_SOURCE_SEND_SUPPRESS_EOP</code>.</p> <p><code>error</code>—Value asserted on the <code>error</code> signal on the output interface.</p> <p><code>data_error_mask</code>—Parameter and the data are XORed together to produce erroneous data.</p>
<b>Returns</b>	Returns 1.
<b>Description</b>	<p>Sends a data fragment to the specified channel. If data packets are supported, applications must ensure consistent usage of SOP and EOP in each channel. Except for the last segment in a packet, the length of each segment is a multiple of the data width.</p> <p>If data packets are not supported, applications must ensure that there are no SOP and EOP indicators in the data. The length of each segment in a packet is a multiple of the data width.</p>

## Test Pattern Checker API

The following subsections describe API for the test pattern checker core. The API functions are currently not available from the ISR.

[data\\_sink\\_reset\(\)](#) on page 10-53

[data\\_sink\\_init\(\)](#) on page 10-53

[data\\_sink\\_get\\_id\(\)](#) on page 10-54

[data\\_sink\\_get\\_supports\\_packets\(\)](#) on page 10-54

[data\\_sink\\_get\\_num\\_channels\(\)](#) on page 10-55

[data\\_sink\\_get\\_symbols\\_per\\_cycle\(\)](#) on page 10-55

[data\\_sink\\_get\\_enable\(\)](#) on page 10-55

[data\\_sink\\_set\\_enable\(\)](#) on page 10-56

[data\\_sink\\_get\\_throttle\(\)](#) on page 10-56

[data\\_sink\\_set\\_throttle\(\)](#) on page 10-57

- [data\\_sink\\_get\\_packet\\_count\(\)](#) on page 10-57
- [data\\_sink\\_get\\_error\\_count\(\)](#) on page 10-58
- [data\\_sink\\_get\\_symbol\\_count\(\)](#) on page 10-58
- [data\\_sink\\_get\\_exception\(\)](#) on page 10-58
- [data\\_sink\\_exception\\_is\\_exception\(\)](#) on page 10-59
- [data\\_sink\\_exception\\_has\\_data\\_error\(\)](#) on page 10-59
- [data\\_sink\\_exception\\_has\\_missing\\_sop\(\)](#) on page 10-60
- [data\\_sink\\_exception\\_has\\_missing\\_eop\(\)](#) on page 10-60
- [data\\_sink\\_exception\\_signalled\\_error\(\)](#) on page 10-61
- [data\\_sink\\_exception\\_channel\(\)](#) on page 10-61

## data\_sink\_reset()

Table 10-39: data\_sink\_reset()

Information Type	Description
<b>Prototype</b>	<code>void data_sink_reset(alt_u32 base);</code>
<b>Thread-safe</b>	No
<b>Include</b>	< <i>data_sink_util.h</i> >
<b>Parameters</b>	<i>base</i> —Base address of the control and status slave.
<b>Returns</b>	void
<b>Description</b>	Resets the test pattern checker core including all internal counters.

## data\_sink\_init()

Table 10-40: data\_sink\_init()

Information Type	Description
<b>Prototype</b>	<code>int data_source_init(alt_u32 base);</code>
<b>Thread-safe</b>	No
<b>Include</b>	< <i>data_sink_util.h</i> >
<b>Parameters</b>	<i>base</i> —Base address of the control and status slave.



Information Type	Description
<b>Returns</b>	1—Initialization is successful. 0—Initialization is unsuccessful.
<b>Description</b>	Performs the following operations to initialize the test pattern checker core: <ul style="list-style-type: none"> <li>Resets and disables the test pattern checker core.</li> <li>Sets the throttle to the maximum value.</li> </ul>

## data\_sink\_get\_id()

Table 10-41: data\_sink\_get\_id()

Information Type	Description
<b>Prototype</b>	<code>int data_sink_get_id(alt_u32 base);</code>
<b>Thread-safe</b>	Yes
<b>Include</b>	<code>&lt;data_sink_util.h &gt;</code>
<b>Parameters</b>	<code>base</code> —Base address of the control and status slave.
<b>Returns</b>	Test pattern checker core identifier.
<b>Description</b>	Retrieves the test pattern checker core's identifier.

## data\_sink\_get\_supports\_packets()

Table 10-42: data\_sink\_get\_supports\_packets()

Information Type	Description
<b>Prototype</b>	<code>int data_sink_init(alt_u32 base);</code>
<b>Thread-safe</b>	Yes
<b>Include</b>	<code>&lt;data_sink_util.h &gt;</code>
<b>Parameters</b>	<code>base</code> —Base address of the control and status slave.
<b>Returns</b>	1—Data packets are supported. 0—Data packets are not supported.
<b>Description</b>	Checks if the test pattern checker core supports data packets.

## data\_sink\_get\_num\_channels()

Table 10-43: data\_sink\_get\_num\_channels()

Information Type	Description
<b>Prototype</b>	<code>int data_sink_get_num_channels(alt_u32 base);</code>
<b>Thread-safe</b>	Yes
<b>Include</b>	< <i>data_sink_util.h</i> >
<b>Parameters</b>	base—Base address of the control and status slave.
<b>Returns</b>	Number of channels supported.
<b>Description</b>	Retrieves the number of channels supported by the test pattern checker core.

## data\_sink\_get\_symbols\_per\_cycle()

Table 10-44: data\_sink\_get\_symbols\_per\_cycle()

Information Type	Description
<b>Prototype</b>	<code>int data_sink_get_symbols(alt_u32 base);</code>
<b>Thread-safe</b>	Yes
<b>Include</b>	< <i>data_sink_util.h</i> >
<b>Parameters</b>	base—Base address of the control and status slave.
<b>Returns</b>	Number of symbols received in a beat.
<b>Description</b>	Retrieves the number of symbols received by the test pattern checker core in each beat.

## data\_sink\_get\_enable()

Table 10-45: data\_sink\_get\_enable()

Information Type	Description
<b>Prototype</b>	<code>int data_sink_get_enable(alt_u32 base);</code>
<b>Thread-safe</b>	Yes

Information Type	Description
<b>Include</b>	<data_sink_util.h >
<b>Parameters</b>	base—Base address of the control and status slave.
<b>Returns</b>	Value of the ENABLE bit.
<b>Description</b>	Retrieves the value of the ENABLE bit.

## data\_sink\_set enable()

Table 10-46: data\_sink\_set enable()

Information Type	Description
<b>Prototype</b>	void data_sink_set_enable(alt_u32 base, alt_u32 value);
<b>Thread-safe</b>	No
<b>Include</b>	<data_sink_util.h >
<b>Parameters</b>	base—Base address of the control and status slave. value—ENABLE bit is set to the value of the parameter.
<b>Returns</b>	void
<b>Description</b>	Enables the test pattern checker core.

## data\_sink\_get\_throttle()

Table 10-47: data\_sink\_get\_throttle()

Information Type	Description
<b>Prototype</b>	int data_sink_get_throttle(alt_u32 base);
<b>Thread-safe</b>	Yes
<b>Include</b>	<data_sink_util.h >
<b>Parameters</b>	base—Base address of the control and status slave.
<b>Returns</b>	Throttle value.
<b>Description</b>	Retrieves the throttle value.

## data\_sink\_set\_throttle()

Table 10-48: data\_sink\_set\_throttle()

Information Type	Description
<b>Prototype</b>	<code>void data_sink_set_throttle(alt_u32 base, alt_u32 value);</code>
<b>Thread-safe</b>	No
<b>Include:</b>	<code>&lt;data_sink_util.h &gt;</code>
<b>Parameters</b>	<p><code>base</code>—Base address of the control and status slave.</p> <p><code>value</code>—Throttle value.</p>
<b>Returns</b>	<code>void</code>
<b>Description</b>	Sets the throttle value, which can be between 0–256 inclusively. The throttle value, when divided by 256 yields the rate at which the test pattern checker receives data.

## data\_sink\_get\_packet\_count()

Table 10-49: data\_sink\_get\_packet\_count()

Information Type	Description
<b>Prototype</b>	<code>int data_sink_get_packet_count(alt_u32 base, alt_u32 channel);</code>
<b>Thread-safe</b>	No
<b>Include</b>	<code>&lt;data_sink_util.h &gt;</code>
<b>Parameters</b>	<p><code>base</code>—Base address of the control and status slave.</p> <p><code>channel</code>—Channel number.</p>
<b>Returns</b>	Number of data packets received on the channel.
<b>Description</b>	Retrieves the number of data packets received on a channel.

**data\_sink\_get\_error\_count()**

Table 10-50: data\_sink\_get\_error\_count()

Information Type	Description
<b>Prototype</b>	<code>int data_sink_get_error_count(alt_u32 base, alt_u32 channel);</code>
<b>Thread-safe</b>	No
<b>Include</b>	<code>&lt;data_sink_util.h &gt;</code>
<b>Parameters</b>	<code>base</code> —Base address of the control and status slave. <code>channel</code> —Channel number.
<b>Returns</b>	Number of errors received on the channel.
<b>Description</b>	Retrieves the number of errors received on a channel.

**data\_sink\_get\_symbol\_count()**

Table 10-51: data\_sink\_get\_symbol\_count()

Information Type	Description
<b>Prototype</b>	<code>int data_sink_get_symbol_count(alt_u32 base, alt_u32 channel);</code>
<b>Thread-safe</b>	No
<b>Include</b>	<code>&lt;data_sink_util.h &gt;</code>
<b>Parameters</b>	<code>base</code> —Base address of the control and status slave. <code>channel</code> —Channel number.
<b>Returns</b>	Number of symbols received on the channel.
<b>Description</b>	Retrieves the number of symbols received on a channel.

**data\_sink\_get\_exception()**

Table 10-52: data\_sink\_get\_exception()

Information Type	Description
<b>Prototype</b>	<code>int data_sink_get_exception(alt_u32 base);</code>

Information Type	Description
<b>Thread-safe</b>	Yes
<b>Include</b>	< <i>data_sink_util.h</i> >
<b>Parameters</b>	<i>base</i> —Base address of the control and status slave.
<b>Returns</b>	First exception descriptor in the exception FIFO. 0—No exception descriptor found in the exception FIFO.
<b>Description</b>	Retrieves the first exception descriptor in the exception FIFO and pops it off the FIFO.

### data\_sink\_exception\_is\_exception()

Table 10-53: data\_sink\_exception\_is\_exception()

Information Type	Description
<b>Prototype</b>	<code>int data_sink_exception_is_exception(int exception);</code>
<b>Thread-safe</b>	Yes
<b>Include</b>	< <i>data_sink_util.h</i> >
<b>Parameters</b>	<i>exception</i> —Exception descriptor
<b>Returns</b>	1—Indicates an exception. 0—No exception.
<b>Description</b>	Checks if an exception descriptor describes a valid exception.

### data\_sink\_exception\_has\_data\_error()

Table 10-54: data\_sink\_exception\_has\_data\_error()

Information Type	Description
<b>Prototype</b>	<code>int data_sink_exception_has_data_error(int exception);</code>
<b>Thread-safe</b>	Yes
<b>Include</b>	< <i>data_sink_util.h</i> >
<b>Parameters</b>	<i>exception</i> —Exception descriptor.

Information Type	Description
<b>Returns</b>	1—Data has errors. 0—No errors.
<b>Description</b>	Checks if an exception indicates erroneous data.

### data\_sink\_exception\_has\_missing\_sop()

Table 10-55: data\_sink\_exception\_has\_missing\_sop()

Information Type	Description
<b>Prototype</b>	<code>int data_sink_exception_has_missing_sop(int exception);</code>
<b>Thread-safe</b>	Yes
<b>Include</b>	<code>&lt;data_sink_util.h &gt;</code>
<b>Parameters</b>	<code>exception</code> —Exception descriptor.
<b>Returns</b>	1—Missing SOP. 0—Other exception types.
<b>Description</b>	Checks if an exception descriptor indicates missing SOP.

### data\_sink\_exception\_has\_missing\_eop()

Table 10-56: data\_sink\_exception\_has\_missing\_eop()

Information Type	Description
<b>Prototype</b>	<code>int data_sink_exception_has_missing_eop(int exception);</code>
<b>Thread-safe</b>	Yes
<b>Include</b>	<code>&lt;data_sink_util.h &gt;</code>
<b>Parameters</b>	<code>exception</code> —Exception descriptor.
<b>Returns</b>	1—Missing EOP. 0—Other exception types.
<b>Description</b>	Checks if an exception descriptor indicates missing EOP.

## data\_sink\_exception\_signalled\_error()

Table 10-57: data\_sink\_exception\_signalled\_error()

Information Type	Description
<b>Prototype</b>	<code>int data_sink_exception_signalled_error(int exception);</code>
<b>Thread-safe</b>	Yes
<b>Include</b>	<code>&lt;data_sink_util.h &gt;</code>
<b>Parameters</b>	<code>exception</code> —Exception descriptor.
<b>Returns</b>	Signal error value.
<b>Description</b>	Retrieves the value of the signaled error from the exception.

## data\_sink\_exception\_channel()

Table 10-58: data\_sink\_exception\_channel()

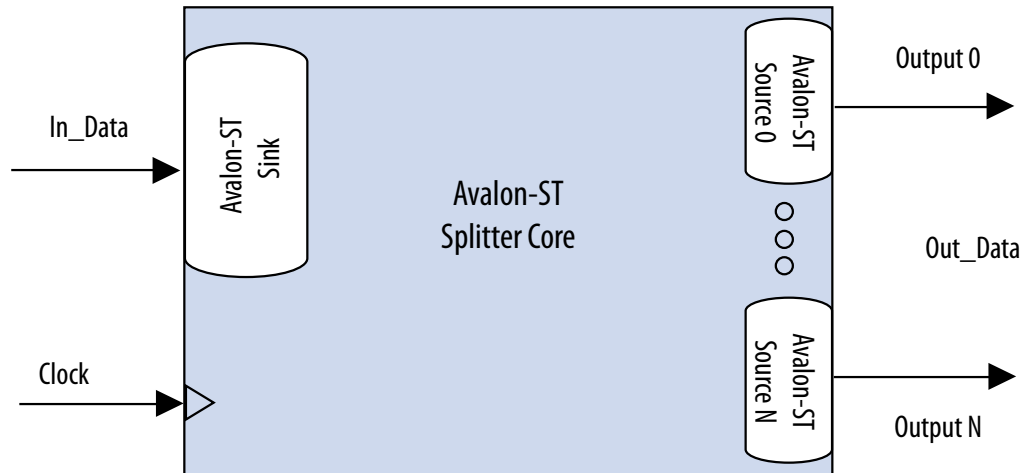
Information Type	Description
<b>Prototype</b>	<code>int data_sink_exception_channel(int exception);</code>
<b>Thread-safe</b>	Yes
<b>Include</b>	<code>&lt;data_sink_util.h &gt;</code>
<b>Parameters</b>	<code>exception</code> —Exception descriptor.
<b>Returns</b>	Channel number on which an exception occurred.
<b>Description</b>	Retrieves the channel number on which an exception occurred.



## Avalon-ST Splitter Core

Figure 10-25: Avalon-ST Splitter Core

The Avalon-ST Splitter Core allows you to replicate transactions from an Avalon-ST source interface to multiple Avalon-ST sink interfaces. This core supports from 1 to 16 outputs.



The Avalon-ST Splitter core copies input signals from the input interface to the corresponding output signals of each output interface without altering the size or functionality. This includes all signals except for the `ready` signal. The core includes a clock signal to determine the Avalon-ST interface and clock domain where the core resides. Because the splitter core does not use the `clock` signal internally, latency is not introduced when using this core.

### Splitter Core Backpressure

The Avalon-ST Splitter core integrates with backpressure by AND-ing the `ready` signals from the output interfaces and sending the result to the input interface. As a result, if an output interface deasserts the `ready` signal, the input interface receives the deasserted `ready` signal, as well. This functionality ensures that backpressure on the output interfaces is propagated to the input interface.

When the **Qualify Valid Out** parameter is set to 1, the `out_valid` signals on all other output interfaces are gated when backpressure is applied from one output interface. In this case, when any output interface deasserts its `ready` signal, the `out_valid` signals on the other output interfaces are also deasserted.

When the **Qualify Valid Out** parameter is set to 0, the output interfaces have a non-gated `out_valid` signal when backpressure is applied. In this case, when an output interface deasserts its `ready` signal, the `out_valid` signals on the other output interfaces are not affected.

Because the logic is combinational, the core introduces no latency.

### Splitter Core Interfaces

The Avalon-ST Splitter core supports streaming data, with optional packet, channel, and error signals. The core propagates backpressure from any output interface to the input interface.

**Table 10-59: Avalon-ST Splitter Core Support**

Feature	Support
Backpressure	Ready latency = 0.
Data Width	Configurable.
Channel	Supported (optional).
Error	Supported (optional).
Packet	Supported (optional).

## Splitter Core Parameters

**Table 10-60: Avalon-ST Splitter Core Parameters**


Parameter	Legal Values	Default Value	Description
<b>Number Of Outputs</b>	1 to 16	2	The number of output interfaces. Qsys supports 1 for some systems where no duplicated output is required.
<b>Qualify Valid Out</b>	0 or 1	1	Determines whether the <code>out_valid</code> signal is gated or non-gated when backpressure is applied.
<b>Data Width</b>	1–512	8	The width of the data on the Avalon-ST data interfaces.
<b>Bits Per Symbol</b>	1–512	8	The number of bits per symbol for the input and output interfaces. For example, byte-oriented interfaces have 8-bit symbols.
<b>Use Packets</b>	0 or 1	0	Indicates whether or not data packet transfers are supported. Packet support includes the <code>startofpacket</code> , <code>endofpacket</code> , and <code>empty</code> signals.
<b>Use Channel</b>	0 or 1	0	The option to enable or disable the channel signal.
<b>Channel Width</b>	0-8	1	The width of the <code>channel</code> signal on the data interfaces. This parameter is disabled when <b>Use Channel</b> is set to 0.

Parameter	Legal Values	Default Value	Description
<b>Max Channels</b>	0-255	1	The maximum number of channels that a data interface can support. This parameter is disabled when <b>Use Channel</b> is set to 0.
<b>Use Error</b>	0 or 1	0	The option to enable or disable the error signal.
<b>Error Width</b>	0–31	1	The width of the <code>error</code> signal on the output interfaces. A value of 0 indicates that the splitter core is not using the <code>error</code> signal. This parameter is disabled when <b>Use Error</b> is set to 0.

## Avalon-ST Delay Core

Figure 10-26: Avalon-ST Delay Core

The Avalon-ST Delay Core provides a solution to delay Avalon-ST transactions by a constant number of clock cycles. This core supports up to 16 clock cycle delays.



The Delay core adds a delay between the input and output interfaces. The core accepts transactions presented on the input interface and reproduces them on the output interface  $N$  cycles later without changing the transaction.

The input interface delays the input signals by a constant  $N$  number of clock cycles to the corresponding output signals of the output interface. The **Number Of Delay Clocks** parameter defines the constant  $N$ , which must be between 0 and 16. The change of the `in_valid` signal is reflected on the `out_valid` signal exactly  $N$  cycles later.

### Delay Core Reset Signal

The Avalon-ST Delay core has a `reset` signal that is synchronous to the `clk` signal. When the core asserts the `reset` signal, the output signals are held at 0. After the `reset` signal is deasserted, the output signals

are held at 0 for  $N$  clock cycles. The delayed values of the input signals are then reflected at the output signals after  $N$  clock cycles.

## Delay Core Interfaces

The Delay core supports streaming data, with optional packet, channel, and error signals. The delay core does not support backpressure.

**Table 10-61: Avalon-ST Delay Core Support**

Feature	Support
Backpressure	Not supported.
Data Width	Configurable.
Channel	Supported (optional).
Error	Supported (optional).
Packet	Supported (optional).

## Delay Core Parameters

**Table 10-62: Avalon-ST Delay Core Parameters**


Parameter	Legal Values	Default Value	Description
<b>Number Of Delay Clocks</b>	0 to 16	1	Specifies the delay the core introduces, in clock cycles. Qsys supports 0 for some systems where no delay is required.
<b>Data Width</b>	1–512	8	The width of the data on the Avalon-ST data interfaces.
<b>Bits Per Symbol</b>	1–512	8	The number of bits per symbol for the input and output interfaces. For example, byte-oriented interfaces have 8-bit symbols.
<b>Use Packets</b>	0 or 1	0	Indicates whether or not data packet transfers are supported. Packet support includes the <code>startofpacket</code> , <code>endofpacket</code> , and <code>empty</code> signals.
<b>Use Channel</b>	0 or 1	0	The option to enable or disable the channel signal.

Parameter	Legal Values	Default Value	Description
<b>Channel Width</b>	0-8	1	The width of the <code>channel</code> signal on the data interfaces. This parameter is disabled when <b>Use Channel</b> is set to 0.
<b>Max Channels</b>	0-255	1	The maximum number of channels that a data interface can support. This parameter is disabled when <b>Use Channel</b> is set to 0.
<b>Use Error</b>	0 or 1	0	The option to enable or disable the error signal.
<b>Error Width</b>	0-31	1	The width of the <code>error</code> signal on the output interfaces. A value of 0 indicates that the error signal is not in use. This parameter is disabled when <b>Use Error</b> is set to 0.

## Avalon-ST Round Robin Scheduler

Figure 10-27: Avalon-ST Round Robin Scheduler

The Avalon-ST Round Robin Scheduler core controls the read operations from a multi-channel Avalon-ST component that buffers data by channels. It reads the almost-full threshold values from the multiple channels in the multi-channel component and issues the read request to the Avalon-ST source according to a round-robin scheduling algorithm.



In a multi-channel component, the component can store data either in the sequence that it comes in (FIFO), or in segments according to the channel. When data is stored in segments according to channels, a scheduler is needed to schedule the read operations.

## Almost-Full Status Interface (Round Robin Scheduler)

The Almost-Full Status interface is an Avalon-ST sink interface that collects the almost-full status from the sink components for the channels in the sequence provided.

**Table 10-63: Avalon-ST Interface Feature Support**

Feature	Property
Backpressure	Not supported
Data Width	Data width = 1; Bits per symbol = 1
Channel	Maximum channel = 32; Channel width = 5
Error	Not supported
Packet	Not supported

## Request Interface (Round Robin Scheduler)

The Request Interface is an Avalon-MM write master interface that requests data from a specific channel. The Avalon-ST Round Robin Scheduler cycles through the channels it supports and schedules data to be read.

## Round Robin Scheduler Operation

If a particular channel is almost full, the Avalon-ST Round Robin Scheduler does not schedule data to be read from that channel in the source component.

The scheduler only requests 1 beat of data from a channel at each transaction. To request 1 beat of data from channel  $n$ , the scheduler writes the value 1 to address  $(4 \times n)$ . For example, if the scheduler is requesting data from channel 3, the scheduler writes 1 to address  $0xc$ . At every clock cycle, the scheduler requests data from the next channel. Therefore, if the scheduler starts requesting from channel 1, at the next clock cycle, it requests from channel 2. The scheduler does not request data from a particular channel if the almost-full status for the channel is asserted. In this case, the scheduler uses one clock cycle without a request transaction.

The Avalon-ST Round Robin Scheduler cannot determine if the requested component is able to service the request transaction. The component asserts `waitrequest` when it cannot accept new requests.

**Table 10-64: Avalon-ST Round Robin Scheduler Ports**

Signal	Direction	Description
<b>Clock and Reset</b>		
clk	In	Clock reference.
reset_n	In	Asynchronous active low reset.

Signal	Direction	Description
<b>Avalon-MM Request Interface</b>		
request_address ( $\log_2$ Max_Channels-1:0)	Out	The write address that indicates which channel has the request.
request_write	Out	Write enable signal.
request_writedata	Out	The amount of data requested from the particular channel. This value is always fixed at 1.
request_waitrequest	In	Wait request signal that pauses the scheduler when the slave cannot accept a new request.

**Avalon-ST Almost-Full Status Interface**

almost_full_valid	In	Indicates that almost_full_channel and almost_full_data are valid.
almost_full_channel (Channel_Width-1:0)	In	Indicates the channel for the current status indication.
almost_full_data ( $\log_2$ Max_Channels-1:0)	In	A 1-bit signal that is asserted high to indicate that the channel indicated by almost_full_channel is almost full.

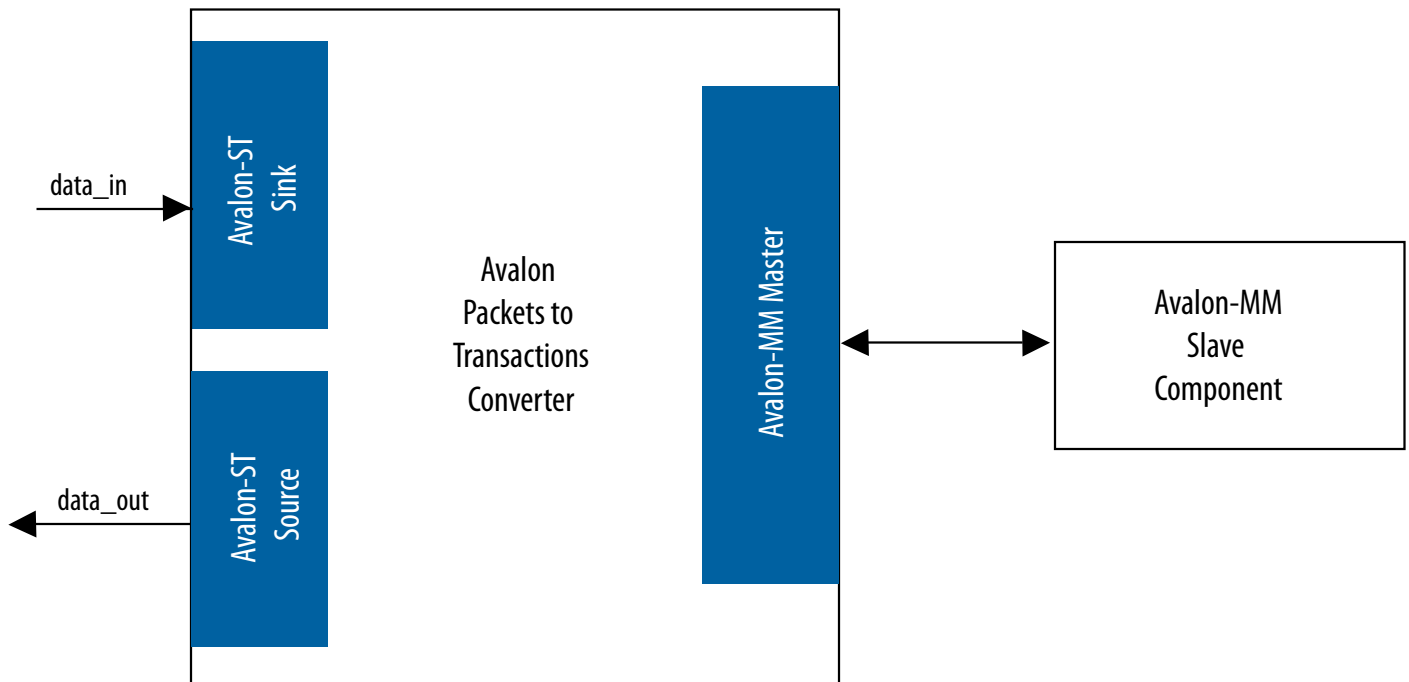
**Round Robin Scheduler Parameters****Table 10-65: Avalon-ST Round Robin Scheduler Parameters**

Parameters	Values	Description
Number of channels	2–32	Specifies the number of channels the Avalon-ST Round Robin Scheduler supports.
Use almost-full status	0–1	Specifies whether the scheduler uses the almost-full interface. If not, the core requests data from the next channel at the next clock cycle.

## Avalon Packets to Transactions Converter

Figure 10-28: Avalon Packets to Transactions Converter Core

The Avalon Packets to Transactions Converter core receives streaming data from upstream components and initiates Avalon-MM transactions. The core then returns Avalon-MM transaction responses to the requesting components.



**Note:** The SPI Slave to Avalon Master Bridge and JTAG to Avalon Master Bridge are examples of the Packets to Transactions Converter core. For more information, refer to the *Avalon Interface Specifications*.

### Related Information

[Avalon Interface Specifications](#)

## Packets to Transactions Converter Interfaces

Table 10-66: Properties of Avalon-ST Interfaces

Feature	Property
Backpressure	Ready latency = 0.
Data Width	Data width = 8 bits; Bits per symbol = 8.
Channel	Not supported.



Feature	Property
Error	Not used.
Packet	Supported.

The Avalon-MM master interface supports read and write transactions. The data width is set to 32 bits, and burst transactions are not supported.

## Packets to Transactions Converter Operation

The Packets to Transactions Converter core receives streams of packets on its Avalon-ST sink interface and initiates Avalon-MM transactions. Upon receiving transaction responses from Avalon-MM slaves, the core transforms the responses to packets and returns them to the requesting components via its Avalon-ST source interface. The core does not report Avalon-ST errors.

### Packets to Transactions Converter Data Packet Formats

A response packet is returned for every write transaction. The core also returns a response packet if a no transaction (`0x7f`) is received. An invalid transaction code is regarded as a no transaction. For read transactions, the core returns the data read.

The Packets to Transactions Converter core expects incoming data streams to be in the formats shown the table below.

**Table 10-67: Data Packet Formats**

Byte	Field	Description
<b>Transaction Packet Format</b>		
0	Transaction code	Type of transaction.
1	Reserved	Reserved for future use.
[3:2]	Size	Transaction size in bytes. For write transactions, the size indicates the size of the <code>data</code> field. For read transactions, the size indicates the total number of bytes to read.
[7:4]	Address	32-bit address for the transaction.
[n:8]	Data	Transaction data; data to be written for write transactions.
<b>Response Packet Format</b>		
0	Transaction code	The transaction code with the most significant bit inverted.
1	Reserved	Reserved for future use.

Byte	Field	Description
[4:2]	Size	Total number of bytes read/written successfully.

**Related Information**

- [Packets to Transactions Converter Interfaces](#) on page 10-69
- [Packets to Transactions Converter Interfaces](#) on page 10-69

**Packets to Transactions Converter Supported Transactions**

**Table 10-68: Packets to Transactions Converter Supported Transactions**

Avalon-MM transactions supported by the Packets to Transactions Converter core.

Transaction Code	Avalon-MM Transaction	Description
0x00	Write, non-incrementing address.	Writes data to the address until the total number of bytes written to the same word address equals to the value specified in the <code>size</code> field.
0x04	Write, incrementing address.	Writes transaction data starting at the current address.
0x10	Read, non-incrementing address.	Reads 32 bits of data from the address until the total number of bytes read from the same address equals to the value specified in the <code>size</code> field.
0x14	Read, incrementing address.	Reads the number of bytes specified in the <code>size</code> parameter starting from the current address.
0x7f	No transaction.	No transaction is initiated. You can use this transaction type for testing purposes. Although no transaction is initiated on the Avalon-MM interface, the core still returns a response packet for this transaction code.

The Packets to Transactions Converter core can process only a single transaction at a time. The `ready` signal on the core's Avalon-ST sink interface is asserted only when the current transaction is completely processed.

No internal buffer is implemented on the data paths. Data received on the Avalon-ST interface is forwarded directly to the Avalon-MM interface and vice-versa. Asserting the `waitrequest` signal on the Avalon-MM interface backpressures the Avalon-ST sink interface. In the opposite direction, if the Avalon-ST source interface is backpressured, the `read` signal on the Avalon-MM interface is not asserted until the backpressure is alleviated. Backpressuring the Avalon-ST source in the middle of a read could result in data loss. In this cases, the core returns the data that is successfully received.

A transaction is considered complete when the core receives an EOP. For write transactions, the actual data size is expected to be the same as the value of the `size` property. Whether or not both values agree, the core always uses the end of packet (EOP) to determine the end of data.

## Packets to Transactions Converter Malformed Packets

The following are examples of malformed packets:

- **Consecutive start of packet (SOP)**—An SOP marks the beginning of a transaction. If an SOP is received in the middle of a transaction, the core drops the current transaction without returning a response packet for the transaction, and initiates a new transaction. This effectively precesses packets without an end of packet (EOP).
- **Unsupported transaction codes**—The core processes unsupported transactions as a no transaction.

## Avalon-ST Streaming Pipeline Stage

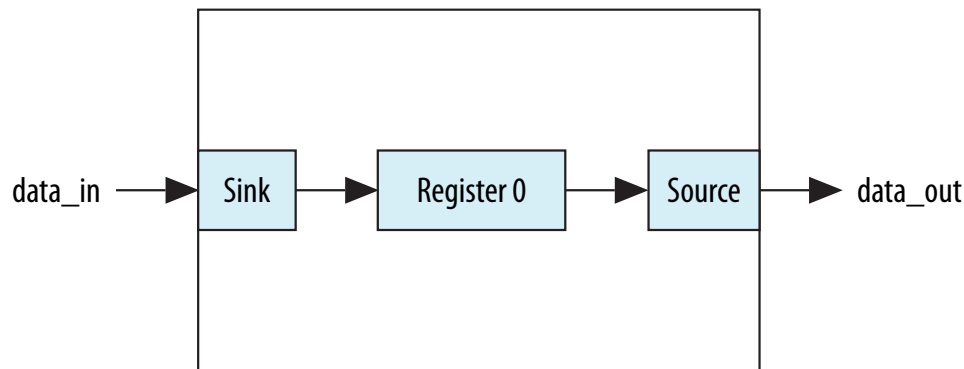
The Avalon-ST pipeline stage receives data from an Avalon-ST source interface, and outputs the data to an Avalon-ST sink interface. In the absence of back pressure, the Avalon-ST pipeline stage source interface outputs data one cycle after receiving the data on its sink interface.

If the pipeline stage receives back pressure on its source interface, it continues to assert its source interface's current data output. While the pipeline stage is receiving back pressure on its source interface and it receives new data on its sink interface, the pipeline stage internally buffers the new data. It then asserts back pressure on its sink interface.

Once the backpressure is deasserted, the pipeline stage's source interface is de-asserted and the pipeline stage asserts internally buffered data (if present). Additionally, the pipeline stage de-asserts back pressure on its sink interface.


**Figure 10-29: Pipeline Stage Simple Register**

If the ready signal is not pipelined, the pipeline stage becomes a simple register.



**Figure 10-30: Pipeline Stage Holding Register**

If the ready signal is pipelined, the pipeline stage must also include a second "holding" register.



## Streaming Channel Multiplexer and Demultiplexer Cores

The Avalon-ST channel multiplexer core receives data from various input interfaces and multiplexes the data into a single output interface, using the optional `channel` signal to indicate the origin of the data. The Avalon-ST channel demultiplexer core receives data from a channelized input interface and drives that data to multiple output interfaces, where the output interface is selected by the input `channel` signal.

The multiplexer and demultiplexer cores can transfer data between interfaces on cores that support unidirectional flow of data. The multiplexer and demultiplexer allow you to create multiplexed or demultiplexed data paths without having to write custom HDL code. The multiplexer includes an Avalon-ST Round Robin Scheduler.

### Related Information

[Avalon-ST Round Robin Scheduler](#) on page 10-66


## Software Programming Model For the Multiplexer and Demultiplexer Components

The multiplexer and demultiplexer components do not have any user-visible control or status registers. Therefore, Qsys cannot control or configure any aspect of the multiplexer or demultiplexer at run-time. The components cannot generate interrupts.

## Avalon-ST Multiplexer

Figure 10-31: Avalon-ST Multiplexer

The Avalon-ST multiplexer takes data from a variety of input data interfaces, and multiplexes the data onto a single output interface. The multiplexer includes a round-robin scheduler that selects from the next input interface that has data. Each input interface has the same width as the output interface, so that the other input interfaces are backpressured when the multiplexer is carrying data from a different input interface.



The multiplexer includes an optional channel signal that enables each input interface to carry channelized data. The output interface channel width is equal to:

$$(\log_2(n-1)) + 1 + w$$

where  $n$  is the number of input interfaces, and  $w$  is the channel width of each input interface. All input interfaces must have the same channel width. These bits are appended to either the most or least significant bits of the output channel signal.

The scheduler processes one input interface at a time, selecting it for transfer. Once an input interface has been selected, data from that input interface is sent until one of the following scenarios occurs:

- The specified number of cycles have elapsed.
- The input interface has no more data to send and the `valid` signal is deasserted on a ready cycle.
- When packets are supported, `endofpacket` is asserted.

### Multiplexer Input Interfaces

Each input interface is an Avalon-ST data interface that optionally supports packets. The input interfaces are identical; they have the same symbol and data widths, error widths, and channel widths.

## Multiplexer Output Interface

The output interface carries the multiplexed data stream with data from the inputs. The symbol, data, and error widths are the same as the input interfaces.

The width of the `channel` signal is the same as the input interfaces, with the addition of the bits needed to indicate the origin of the data.

You can configure the following parameters for the output interface:

- **Data Bits Per Symbol**—The bits per symbol is related to the width of `readdata` and `writedata` signals, which must be a multiple of the bits per symbol.
- **Data Symbols Per Beat**—The number of symbols (words) that are transferred per beat (transfer). Valid values are 1 to 32.
- **Include Packet Support**—Indicates whether or not packet transfers are supported. Packet support includes the `startofpacket`, `endofpacket`, and `empty` signals.
- **Channel Signal Width (bits)**— The number of bits Qsys uses for the channel signal for output interfaces. For example, set this parameter to 1 if you have two input interfaces with no channel, or set this parameter to 2 if you have two input interfaces with a channel width of 1 bit. The input channel can have a width between 0-31 bits.
- **Error Signal Width (bits)**—The width of the `error` signal for input and output interfaces. A value of 0 means the `error` signal is not in use.

**Note:** If you change only bits per symbol, and do not change the data width, errors are generated.

## Multiplexer Parameters


You can configure the following parameters for the multiplexer:

- **Number of Input Ports**—The number of input interfaces that the multiplexer supports. Valid values are 2 to 16.
- **Scheduling Size (Cycles)**—The number of cycles that are sent from a single channel before changing to the next channel.
- **Use Packet Scheduling**—When this parameter is turned on, the multiplexer only switches the selected input interface on packet boundaries. Therefore, packets on the output interface are not interleaved.
- **Use high bits to indicate source port**—When this parameter is turned on, the multiplexer uses the high bits of the output `channel` signal to indicate the origin of the input interface of the data. For example, if the input interfaces have 4-bit channel signals, and the multiplexer has 4 input interfaces, the output interface has a 6-bit channel signal. If this parameter is turned on, bits [5:4] of the output channel signal indicate origin of the input interface of the data, and bits [3:0] are the channel bits that were presented at the input interface.

## Avalon-ST Demultiplexer

Figure 10-32: Avalon-ST Demultiplexer

That Avalon-ST demultiplexer takes data from a channelized input data interface and provides that data to multiple output interfaces, where the output interface selected for a particular transfer is specified by the input `channel` signal.



The data is delivered to the output interfaces in the same order it is received at the input interface, regardless of the value of `channel`, `packet`, `frame`, or any other signal. Each of the output interfaces has the same width as the input interface; each output interface is idle when the demultiplexer is driving data to a different output interface. The demultiplexer uses  $\log_2(\text{num\_output\_interfaces})$  bits of the `channel` signal to select the output for the data; the remainder of the channel bits are forwarded to the appropriate output interface unchanged.

### Demultiplexer Input Interface

Each input interface is an Avalon-ST data interface that optionally supports packets. You can configure the following parameters for the input interface:

- **Data Bits Per Symbol**—The bits per symbol is related to the width of `readdata` and `writedata` signals, which must be a multiple of the bits per symbol.
- **Data Symbols Per Beat**—The number of symbols (words) that are transferred per beat (transfer). Valid values are 1 to 32.
- **Include Packet Support**—Indicates whether or not data packet transfers are supported. Packet support includes the `startofpacket`, `endofpacket`, and `empty` signals.
- **Channel Signal Width (bits)**—The number of bits for the `channel` signal for output interfaces. A value of 0 means that output interfaces do not use the optional `channel` signal.
- **Error Signal Width (bits)**—The width of the `error` signal for input and output interfaces. A value of 0 means the `error` signal is in use.

**Note:** If you change only bits per symbol, and do not change the data width, errors are generated.

### Demultiplexer Output Interface

Each output interface carries data from a subset of channels from the input interface. Each output interface is identical; all have the same symbol and data widths, error widths, and channel widths. The

symbol, data, and error widths are the same as the input interface. The width of the `channel` signal is the same as the input interface, without the bits that the demultiplexer uses to select the output interface.


## Demultiplexer Parameters

You can configure the following parameters for the demultiplexer:

- **Number of Output Ports**—The number of output interfaces that the multiplexer supports. Valid values are 2 to 16.
- **High channel bits select output**—When this option is turned on, the demultiplexing function uses the high bits of the input `channel` signal, and the low order bits are passed to the output. When this option is turned off, the demultiplexing function uses the low order bits, and the high order bits are passed to the output.

Where you place the signals in our design affects the functionality; for example, there is one input interface and two output interfaces. If the low-order bits of the channel signal select the output interfaces, the even channels go to channel 0, and the odd channels go to channel 1. If the high-order bits of the channel signal select the output interface, channels 0 to 7 go to channel 0 and channels 8 to 15 go to channel 1.

Figure 10-33: Select Bits for the Demultiplexer



## Single-Clock and Dual-Clock FIFO Cores

The Avalon-ST Single-Clock and Avalon-ST Dual-Clock FIFO cores are FIFO buffers which operate with a common clock and independent clocks for input and output ports respectively.



Figure 10-34: Avalon-ST Single Clock FIFO Core





Figure 10-35: Avalon-ST Dual Clock FIFO Core



## Interfaces Implemented in FIFO Cores

The following interfaces are implemented in FIFO cores:

[Avalon-ST Data Interface](#) on page 10-79

[Avalon-MM Control and Status Register Interface](#) on page 10-80

[Avalon-ST Status Interface](#) on page 10-80

### Avalon-ST Data Interface

Each FIFO core has an Avalon-ST data sink and source interfaces. The data sink and source interfaces in the dual-clock FIFO core are driven by different clocks.

Table 10-69: Avalon-ST Interfaces Properties

Feature	Property
Backpressure	Ready latency = 0.
Data Width	Configurable.

Feature	Property
Channel	Supported, up to 255 channels.
Error	Configurable.
Packet	Configurable.

## Avalon-MM Control and Status Register Interface

You can configure the single-clock FIFO core to include an optional Avalon-MM interface, and the dual-clock FIFO core to include an Avalon-MM interface in each clock domain. The Avalon-MM interface provides access to 32-bit registers, which allows you to retrieve the FIFO buffer fill level and configure the almost-empty and almost-full thresholds. In the single-clock FIFO core, you can also configure the packet and error handling modes.

## Avalon-ST Status Interface

The single-clock FIFO core has two optional Avalon-ST status source interfaces from which you can obtain the FIFO buffer almost-full and almost empty statuses.

## FIFO Operating Modes

- **Default mode**—The core accepts incoming data on the `in` interface (Avalon-ST data sink) and forwards it to the `out` interface (Avalon-ST data source). The core asserts the `valid` signal on the Avalon-ST source interface to indicate that data is available at the interface.
- **Store and forward mode**—This mode applies only to the single-clock FIFO core. The core asserts the `valid` signal on the `out` interface only when a full packet of data is available at the interface. In this mode, you can also enable the drop-on-error feature by setting the `drop_on_error` register to 1. When this feature is enabled, the core drops all packets received with the `in_error` signal asserted.
- **Cut-through mode**—This mode applies only to the single-clock FIFO core. The core asserts the `valid` signal on the `out` interface to indicate that data is available for consumption when the number of entries specified in the `cut_through_threshold` register are available in the FIFO buffer.

**Note:** To turn on **Cut-through mode**, **Use store and forward** must be set to 0. Turning on **Use store and forward mode** prompts the user to turn on **Use fill level**, and then the CSR appears.

## Fill Level of the FIFO Buffer

You can obtain the fill level of the FIFO buffer via the optional Avalon-MM control and status interface. Turn on the **Use fill level** parameter (**Use sink fill level** and **Use source fill level** in the dual-clock FIFO core) and read the `fill_level` register.

The dual-clock FIFO core has two fill levels, one in each clock domain. Due to the latency of the clock crossing logic, the fill levels reported in the input and output clock domains may be different for any instance. In both cases, the fill level may report badly for the clock domain; that is, the fill level is reported high in the input clock domain, and low in the output clock domain.

The dual-clock FIFO has an output pipeline stage to improve  $f_{MAX}$ . This output stage is accounted for when calculating the output fill level, but not when calculating the input fill level. Therefore, the best measure of the amount of data in the FIFO is by the fill level in the output clock domain. The fill level in

the input clock domain represents the amount of space available in the FIFO (available space = FIFO depth – input fill level).

## Almost-Full and Almost-Empty Thresholds to Prevent Overflow and Underflow

You can use almost-full and almost-empty thresholds as a mechanism to prevent FIFO overflow and underflow. This feature is available only in the single-clock FIFO core. To use the thresholds, turn on the **Use fill level**, **Use almost-full status**, and **Use almost-empty status** parameters. You can access the `almost_full_threshold` and `almost_empty_threshold` registers via the csr interface and set the registers to an optimal value for your application.

You can obtain the almost-full and almost-empty statuses from `almost_full` and `almost_empty` interfaces (Avalon-ST status source). The core asserts the `almost_full` signal when the fill level is equal to or higher than the almost-full threshold. Likewise, the core asserts the `almost_empty` signal when the fill level is equal to or lower than the almost-empty threshold.

## Single-Clock and Dual-Clock FIFO Core Parameters

Table 10-70: Single-Clock and Dual-Clock FIFO Core Parameters

Parameter	Legal Values	Description
<b>Bits per symbol</b>	1–32	These parameters determine the width of the FIFO.
<b>Symbols per beat</b>	1–32	FIFO width = <b>Bits per symbol</b> * <b>Symbols per beat</b> , where: <b>Bits per symbol</b> is the number of bits in a symbol, and <b>Symbols per beat</b> is the number of symbols transferred in a beat.
<b>Error width</b>	0–32	The width of the <code>error</code> signal.
<b>FIFO depth</b>	2 <sup>n</sup>	The FIFO depth. An output pipeline stage is added to the FIFO to increase performance, which increases the FIFO depth by one. <n> = n=1,2,3,4...
<b>Use packets</b>	—	Turn on this parameter to enable data packet support on the Avalon-ST data interfaces.
<b>Channel width</b>	1–32	The width of the <code>channel</code> signal.
<b>Avalon-ST Single Clock FIFO Only</b>		
<b>Use fill level</b>	—	Turn on this parameter to include the Avalon-MM control and status register interface (CSR). The CSR is enabled when <b>Use fill level</b> is set to 1.

Parameter	Legal Values	Description
<b>Use Store and Forward</b>		To turn on <b>Cut-through mode</b> , <b>Use store and forward</b> must be set to 0. Turning on <b>Use store and forward</b> prompts the user to turn on <b>Use fill level</b> , and then the CSR appears.
<b>Avalon-ST Dual Clock FIFO Only</b>		
<b>Use sink fill level</b>	—	Turn on this parameter to include the Avalon-MM control and status register interface in the input clock domain.
<b>Use source fill level</b>	—	Turn on this parameter to include the Avalon-MM control and status register interface in the output clock domain.
<b>Write pointer synchronizer length</b>	2–8	The length of the write pointer synchronizer chain. Setting this parameter to a higher value leads to better metastability while increasing the latency of the core.
<b>Read pointer synchronizer length</b>	2–8	The length of the read pointer synchronizer chain. Setting this parameter to a higher value leads to better metastability.
<b>Use Max Channel</b>	—	Turn on this parameter to specify the maximum channel number.
<b>Max Channel</b>	1–255	Maximum channel number.

**Note:** For more information on metastability in Altera devices, refer to *Understanding Metastability in FPGAs*. For more information on metastability analysis and synchronization register chains, refer to the *Managing Metastability*.

#### Related Information

[Understanding Metastability in FPGAs](#)

[Managing Metastability](#) on page 13-1

## Avalon-ST Single-Clock FIFO Registers

**Table 10-71: Avalon-ST Single-Clock FIFO Registers**

The CSR interface in the Avalon-ST Single Clock FIFO core provides access to registers.

32-Bit Word Offset	Name	Access	Reset	Description
0	fill_level	R	0	24-bit FIFO fill level. Bits 24 to 31 are not used.
1	Reserved	—	—	Reserved for future use.
2	almost_full_threshold	RW	<b>FIFO depth-1</b>	Set this register to a value that indicates the FIFO buffer is getting full.
3	almost_empty_threshold	RW	0	Set this register to a value that indicates the FIFO buffer is getting empty.
4	cut_through_threshold	RW	0	<p><b>0</b>—Enables store and forward mode.</p> <p><b>Greater than 0</b>—Enables cut-through mode and specifies the minimum of entries in the FIFO buffer before the valid signal on the Avalon-ST source interface is asserted. Once the FIFO core starts sending the data to the downstream component, it continues to do so until the end of the packet.</p> <p><b>Note:</b> To turn on <b>Cut-through mode</b>, <b>Use store and forward</b> must be set to 0. Turning on <b>Use store and forward mode</b> prompts the user to turn on <b>Use fill level</b>, and then the CSR appears.</p>
5	drop_on_error	RW	0	<p><b>0</b>—Disables drop-on error.</p> <p><b>1</b>—Enables drop-on error.</p> <p>This register applies only when the <b>Use packet</b> and <b>Use store and forward</b> parameters are turned on.</p>

**Table 10-72: Register Description for Avalon-ST Dual-Clock FIFO**

The in\_csr and out\_csr interfaces in the Avalon-ST Dual Clock FIFO core reports the FIFO fill level.

32-Bit Word Offset	Name	Access	Reset Value	Description
0	fill_level	R	0	24-bit FIFO fill level. Bits 24 to 31 are not used.

## Related Information

- [Avalon Interface Specifications](#)
- [Avalon Memory-Mapped Design Optimizations](#)

## Document Revision History

**Table 10-73: Document Revision History**

The table below indicates edits made to the *Qsys System Design Components* content since its creation.

Date	Version	Changes
December 2014	14.1.0	<ul style="list-style-type: none"> <li>• AXI Timeout Bridge.</li> <li>• Added notes to <i>Avalon-MM Clock Crossing Bridge</i> pertaining to:               <ul style="list-style-type: none"> <li>• SDC constraints for its internal asynchronous FIFOs.</li> <li>• FIFO-based clock crossing.</li> </ul> </li> </ul>
June 2014	14.0.0	<ul style="list-style-type: none"> <li>• AXI Bridge support.</li> <li>• Address Span Extender updates.</li> <li>• Avalon-MM Unaligned Burst Expansion Bridge support.</li> </ul>
November 2013	13.1.0	<ul style="list-style-type: none"> <li>• Address Span Extender</li> </ul>
May 2013	13.0.0	<ul style="list-style-type: none"> <li>• Added Streaming Pipeline Stage support.</li> <li>• Added AMBA APB support.</li> </ul>
November 2012	12.1.0	<ul style="list-style-type: none"> <li>• Moved relevant content from the <i>Embedded Peripherals IP User Guide</i>.</li> </ul>

## Related Information

[Quartus II Handbook Archive](#)

2014.12.15

QII5V1



Subscribe



Send Feedback

This chapter provides design recommendations for Altera® devices and describes the Quartus® II Design Assistant, which helps you check your design for violations of Altera's design recommendations.

Current FPGA applications have reached the complexity and performance requirements of ASICs. In the development of complex system designs, good design practices have an enormous impact on the timing performance, logic utilization, and system reliability of a device. Well-coded designs behave in a predictable and reliable manner even when retargeted to different families or speed grades. Good design practices also aid in successful design migration between FPGA and ASIC implementations for prototyping and production.

For optimal performance, reliability, and faster time-to-market when designing with Altera devices, you should adhere to the following guidelines:

- Understand the impact of synchronous design practices
- Follow recommended design techniques, including hierarchical design partitioning, and timing closure guidelines
- Take advantage of the architectural features in the targeted device

## Following Synchronous FPGA Design Practices

The first step in good design methodology is to understand the implications of your design practices and techniques. This section outlines the benefits of optimal synchronous design practices and the hazards involved in other techniques.

Good synchronous design practices can help you meet your design goals consistently. Problems with other design techniques can include reliance on propagation delays in a device, which can lead to race conditions, incomplete timing analysis, and possible glitches.

In a synchronous design, a clock signal triggers every event. As long as you ensure that all the timing requirements of the registers are met, a synchronous design behaves in a predictable and reliable manner for all process, voltage, and temperature (PVT) conditions. You can easily migrate synchronous designs to different device families or speed grades.

## Implementing Synchronous Designs

In a synchronous design, the clock signal controls the activities of all inputs and outputs.

On every active edge of the clock (usually the rising edge), the data inputs of registers are sampled and transferred to outputs. Following an active clock edge, the outputs of combinational logic feeding the data

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered





inputs of registers change values. This change triggers a period of instability due to propagation delays through the logic as the signals go through several transitions and finally settle to new values. Changes that occur on data inputs of registers do not affect the values of their outputs until after the next active clock edge.

Because the internal circuitry of registers isolates data outputs from inputs, instability in the combinational logic does not affect the operation of the design as long as you meet the following timing requirements:

- Before an active clock edge, you must ensure that the data input has been stable for at least the setup time of the register.
- After an active clock edge, you must ensure that the data input remains stable for at least the hold time of the register.

When you specify all of your clock frequencies and other timing requirements, the Quartus II TimeQuest Timing Analyzer reports actual hardware requirements for the setup times ( $t_{SU}$ ) and hold times ( $t_H$ ) for every pin in your design. By meeting these external pin requirements and following synchronous design techniques, you ensure that you satisfy the setup and hold times for all registers in your device.

**Tip:** To meet setup and hold time requirements on all input pins, any inputs to combinational logic that feed a register should have a synchronous relationship with the clock of the register. If signals are asynchronous, you can register the signals at the inputs of the device to help prevent a violation of the required setup and hold times.

When you violate the setup or hold time of a register, you might oscillate the output, or set the output to an intermediate voltage level between the high and low levels called a metastable state. In this unstable state, small perturbations such as noise in power rails can cause the register to assume either the high or low voltage level, resulting in an unpredictable valid state. Various undesirable effects can occur, including increased propagation delays and incorrect output states. In some cases, the output can even oscillate between the two valid states for a relatively long period of time.

#### Related Information

##### About TimeQuest Timing Analysis

For information about timing requirements and analysis in the Quartus II software, refer to About TimeQuest Timing Analysis in Quartus II Help.

## Asynchronous Design Hazards

Designers use asynchronous techniques such as ripple counters or pulse generators in programmable logic device (PLD) designs, enabling them to take “short cuts” to save device resources.

Asynchronous design techniques have inherent problems such as relying on propagation delays in a device, which can vary with temperature and voltage fluctuations, resulting in incomplete timing constraints and possible glitches and spikes.

Some asynchronous design structures rely on the relative propagation delays of signals to function correctly. In these cases, race conditions can arise where the order of signal changes can affect the output of the logic. PLD designs can have varying timing delays, depending on how the design is placed and routed in the device with each compilation. Therefore, it is almost impossible to determine the timing delay associated with a particular block of logic ahead of time. As devices become faster due to device process improvements, the delays in an asynchronous design may decrease, resulting in a design that does not function as expected. This chapter provides specific examples. Relying on a particular delay also makes asynchronous designs difficult to migrate to different architectures, devices, or speed grades.

The timing of asynchronous design structures is often difficult or impossible to model with timing assignments and constraints. If you do not have complete or accurate timing constraints, the timing-driven algorithms used by your synthesis and place-and-route tools may not be able to perform the best optimizations, and the reported results may not be complete.

Some asynchronous design structures can generate harmful glitches, which are pulses that are very short compared with clock periods. Most glitches are generated by combinational logic. When the inputs of combinational logic change, the outputs exhibit several glitches before they settle to their new values. These glitches can propagate through the combinational logic, leading to incorrect values on the outputs in asynchronous designs. In a synchronous design, glitches on the data inputs of registers are normal events that have no negative consequences because the data is not processed until the clock edge.

## HDL Design Guidelines

When designing with HDL code, you should understand how a synthesis tool interprets different HDL design techniques and what results to expect.

Your design techniques can affect logic utilization and timing performance, as well as the design's reliability. This section describes basic design techniques that ensure optimal synthesis results for designs targeted to Altera devices while avoiding several common causes of unreliability and instability. Altera recommends that you design your combinational logic carefully to avoid potential problems and pay attention to your clocking schemes so that you can maintain synchronous functionality and avoid timing problems.

### Optimizing Combinational Logic

Combinational logic structures consist of logic functions that depend only on the current state of the inputs. In Altera FPGAs, these functions are implemented in the look-up tables (LUTs) with either logic elements (LEs) or adaptive logic modules (ALMs).


For cases where combinational logic feeds registers, the register control signals can implement part of the logic function to save LUT resources. By following the recommendations in this section, you can improve the reliability of your combinational design.

### Avoid Combinational Loops

Combinational loops are among the most common causes of instability and unreliability in digital designs. Combinational loops generally violate synchronous design principles by establishing a direct feedback loop that contains no registers.

You should avoid combinational loops whenever possible. In a synchronous design, feedback loops should include registers. For example, a combinational loop occurs when the left-hand side of an arithmetic expression also appears on the right-hand side in HDL code. A combinational loop also occurs when you feed back the output of a register to an asynchronous pin of the same register through combinational logic.

**Figure 11-1: Combinational Loop Through Asynchronous Control Pin**



**Tip:** Use recovery and removal analysis to perform timing analysis on asynchronous ports, such as `clear` or `reset` in the Quartus II software.

Combinational loops are inherently high-risk design structures for the following reasons:

- Combinational loop behavior generally depends on relative propagation delays through the logic involved in the loop. As discussed, propagation delays can change, which means the behavior of the loop is unpredictable.
- Combinational loops can cause endless computation loops in many design tools. Most tools break open combinational loops to process the design. The various tools used in the design flow may open a given loop in a different manner, processing it in a way that is inconsistent with the original design intent.

#### Related Information

#### [Specifying Timing Constraints and Exceptions](#)

### Avoid Unintended Latch Inference

A latch is a small circuit with combinational feedback that holds a value until a new value is assigned. You can implement latches with the Quartus II Text Editor or Block Editor.

It is common for mistakes in HDL code to cause unintended latch inference; Quartus II Synthesis issues a warning message if this occurs. Unlike other technologies, a latch in FPGA architecture is not significantly smaller than a register. The architecture is not optimized for latch implementation and latches generally have slower timing performance compared to equivalent registered circuitry.

Latches have a transparent mode in which data flows continuously from input to output. A positive latch is in transparent mode when the enable signal is high (low for negative latch). In transparent mode, glitches on the input can pass through to the output because of the direct path created. This presents significant complexity for timing analysis. Typical latch schemes use multiple enable phases to prevent long transparent paths from occurring. However, timing analysis cannot identify these safe applications.

The TimeQuest analyzer analyzes latches as synchronous elements clocked on the falling edge of the positive latch signal by default, and allows you to treat latches as having nontransparent start and end points. Be aware that even an instantaneous transition through transparent mode can lead to glitch propagation. The TimeQuest analyzer cannot perform cycle-borrowing analysis.

Due to various timing complexities, latches have limited support in formal verification tools. Therefore, you should not rely on formal verification for a design that includes latches.

**Tip:** Avoid using latches to ensure that you can completely analyze the timing performance and reliability of your design.

### Avoid Delay Chains in Clock Paths

You require delay chains when you use two or more consecutive nodes with a single fan-in and a single fan-out to cause delay. Inverters are often chained together to add delay. Delay chains are sometimes used to resolve race conditions created by other asynchronous design practices.

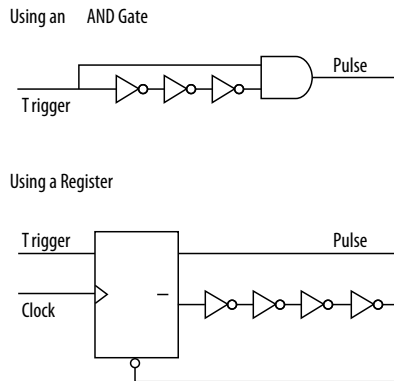
Delays in PLD designs can change with each placement and routing cycle. Effects such as rise and fall time differences and on-chip variation mean that delay chains, especially those placed on clock paths, can cause significant problems in your design. Avoid using delay chains to prevent these kinds of problems.

In some ASIC designs, delays are used for buffering signals as they are routed around the device. This functionality is not required in FPGA devices because the routing structure provides buffers throughout the device.

## Use Synchronous Pulse Generators

You can use delay chains to generate either one pulse (pulse generators) or a series of pulses (multivibrators). There are two common methods for pulse generation. These techniques are purely asynchronous and must be avoided.

**Figure 11-2: Asynchronous Pulse Generators**



A trigger signal feeds both inputs of a 2-input AND gate, but the design adds inverters to create a delay chain to one of the inputs. The width of the pulse depends on the time differences between path that feeds the gate directly, and the path that goes through the delay chain. This is the same mechanism responsible for the generation of glitches in combinational logic following a change of input values. This technique artificially increases the width of the glitch.


A register's output drives the same register's asynchronous reset signal through a delay chain. The register resets itself asynchronously after a certain delay.

The width of pulses generated in this way are difficult for synthesis and place-and-route to determine, set, or verify. The actual pulse width can only be determined after placement and routing, when routing and propagation delays are known. You cannot reliably create a specific pulse width when creating HDL code, and it cannot be set by EDA tools. The pulse may not be wide enough for the application under all PVT conditions. Also, the pulse width changes if you change to a different device. Additionally, verification is difficult because static timing analysis cannot verify the pulse width.

Multivibrators use a glitch generator to create pulses, together with a combinational loop that turns the circuit into an oscillator. This creates additional problems because of the number of pulses involved. Additionally, when the structures generate multiple pulses, they also create a new artificial clock in the design must be analyzed by design tools.

When you must use a pulse generator, use synchronous techniques.

**Figure 11-3: Recommended Pulse-Generation Technique**



The pulse width is always equal to the clock period. This pulse generator is predictable, can be verified with timing analysis, and is easily moved to other architectures, devices, or speed grades.

## Optimizing Clocking Schemes

Like combinational logic, clocking schemes have a large effect on the performance and reliability of a design.

Avoid using internally generated clocks (other than PLLs) wherever possible because they can cause functional and timing problems in the design. Clocks generated with combinational logic can introduce glitches that create functional problems, and the delay inherent in combinational logic can lead to timing problems.

**Tip:** Specify all clock relationships in the Quartus II software to allow for the best timing-driven optimizations during fitting and to allow correct timing analysis. Use clock setting assignments on any derived or internal clocks to specify their relationship to the base clock.

Use global device-wide, low-skew dedicated routing for all internally-generated clocks, instead of routing clocks on regular routing lines.

Avoid data transfers between different clocks wherever possible. If you require a data transfer between different clocks, use FIFO circuitry. You can use the clock uncertainty features in the Quartus II software to compensate for the variable delays between clock domains. Consider setting a clock setup uncertainty and clock hold uncertainty value of 10% to 15% of the clock delay.

The following sections provide specific examples and recommendations for avoiding clocking scheme problems.


### Register Combinational Logic Outputs

If you use the output from combinational logic as a clock signal or as an asynchronous reset signal, you can expect to see glitches in your design. In a synchronous design, glitches on data inputs of registers are normal events that have no consequences. However, a glitch or a spike on the clock input (or an asynchronous input) to a register can have significant consequences.

Narrow glitches can violate the register's minimum pulse width requirements. Setup and hold requirements might also be violated if the data input of the register changes when a glitch reaches the clock input. Even if the design does not violate timing requirements, the register output can change value unexpectedly and cause functional hazards elsewhere in the design.

To avoid these problems, you should always register the output of combinational logic before you use it as a clock signal.

**Figure 11-4: Recommended Clock-Generation Technique**



Registering the output of combinational logic ensures that glitches generated by the combinational logic are blocked at the data input of the register.

## Avoid Asynchronous Clock Division

Designs often require clocks that you create by dividing a master clock. Most Altera FPGAs provide dedicated phase-locked loop (PLL) circuitry for clock division. Using dedicated PLL circuitry can help you to avoid many of the problems that can be introduced by asynchronous clock division logic.

When you must use logic to divide a master clock, always use synchronous counters or state machines. Additionally, create your design so that registers always directly generate divided clock signals, and route the clock on global clock resources. To avoid glitches, do not decode the outputs of a counter or a state machine to generate clock signals.

## Avoid Ripple Counters

To simplify verification, avoid ripple counters in your design. In the past, FPGA designers implemented ripple counters to divide clocks by a power of two because the counters are easy to design and may use fewer gates than their synchronous counterparts.

Ripple counters use cascaded registers, in which the output pin of one register feeds the clock pin of the register in the next stage. This cascading can cause problems because the counter creates a ripple clock at each stage. These ripple clocks must be handled properly during timing analysis, which can be difficult and may require you to make complicated timing assignments in your synthesis and placement and routing tools.

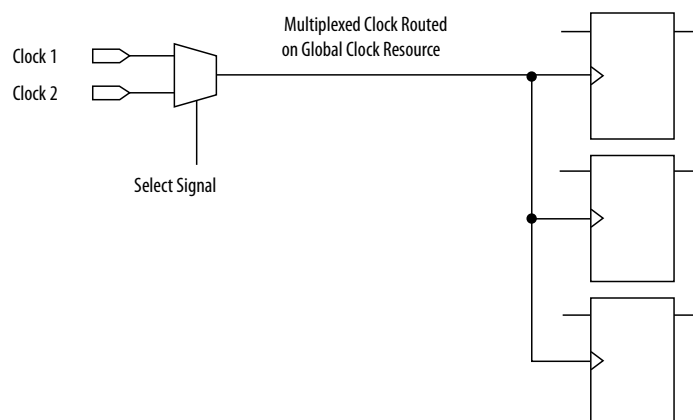
You can often use ripple clock structures to make ripple counters out of the smallest amount of logic possible. However, in all Altera devices supported by the Quartus II software, using a ripple clock structure to reduce the amount of logic used for a counter is unnecessary because the device allows you to construct a counter using one logic element per counter bit. You should avoid using ripple counters completely.

## Use Multiplexed Clocks

Use clock multiplexing to operate the same logic function with different clock sources. In these designs, multiplexing selects a clock source.

For example, telecommunications applications that deal with multiple frequency standards often use multiplexed clocks.

Figure 11-5: Multiplexing Logic and Clock Sources



Adding multiplexing logic to the clock signal can create the problems addressed in the previous sections, but requirements for multiplexed clocks vary widely, depending on the application. Clock multiplexing is acceptable when the clock signal uses global clock routing resources and if the following criteria are met:

- The clock multiplexing logic does not change after initial configuration
- The design uses multiplexing logic to select a clock for testing purposes
- Registers are always reset when the clock switches
- A temporarily incorrect response following clock switching has no negative consequences

If the design switches clocks in real time with no reset signal, and your design cannot tolerate a temporarily incorrect response, you must use a synchronous design so that there are no timing violations on the registers, no glitches on clock signals, and no race conditions or other logical problems. By default, the Quartus II software optimizes and analyzes all possible paths through the multiplexer and between both internal clocks that may come from the multiplexer. This may lead to more restrictive analysis than required if the multiplexer is always selecting one particular clock. If you do not require the more complete analysis, you can assign the output of the multiplexer as a base clock in the Quartus II software, so that all register-to-register paths are analyzed using that clock.


**Tip:** Use dedicated hardware to perform clock multiplexing when it is available, instead of using multiplexing logic. For example, you can use the clock-switchover feature or clock control block available in certain Altera devices. These dedicated hardware blocks ensure that you use global low-skew routing lines and avoid any possible hold time problems on the device due to logic delay on the clock line.

**Note:** For device-specific information about clocking structures, refer to the appropriate device data sheet or handbook on the Literature page of the Altera website.

## Use Gated Clocks

Gated clocks turn a clock signal on and off using an enable signal that controls gating circuitry. When a clock is turned off, the corresponding clock domain is shut down and becomes functionally inactive.

**Figure 11-6: Gated Clock**



You can use gated clocks to reduce power consumption in some device architectures by effectively shutting down portions of a digital circuit when they are not in use. When a clock is gated, both the clock network and the registers driven by it stop toggling, thereby eliminating their contributions to power consumption. However, gated clocks are not part of a synchronous scheme and therefore can significantly increase the effort required for design implementation and verification. Gated clocks contribute to clock skew and make device migration difficult. These clocks are also sensitive to glitches, which can cause design failure.

Use dedicated hardware to perform clock gating rather than an AND or OR gate. For example, you can use the clock control block in newer Altera devices to shut down an entire clock network. Dedicated hardware blocks ensure that you use global routing with low skew, and avoid any possible hold time problems on the device due to logic delay on the clock line.

From a functional point of view, you can shut down a clock domain in a purely synchronous manner using a synchronous clock enable signal. However, when using a synchronous clock enable scheme, the


clock network continues toggling. This practice does not reduce power consumption as much as gating the clock at the source does. In most cases, use a synchronous scheme.

### Use Synchronous Clock Enables

To turn off a clock domain in a synchronous manner, use a synchronous clock enable signal. FPGAs efficiently support clock enable signals because there is a dedicated clock enable signal available on all device registers.

This scheme does not reduce power consumption as much as gating the clock at the source because the clock network keeps toggling, and performs the same function as a gated clock by disabling a set of registers. Insert a multiplexer in front of the data input of every register to either load new data, or copy the output of the register.

Figure 11-7: Synchronous Clock Enable



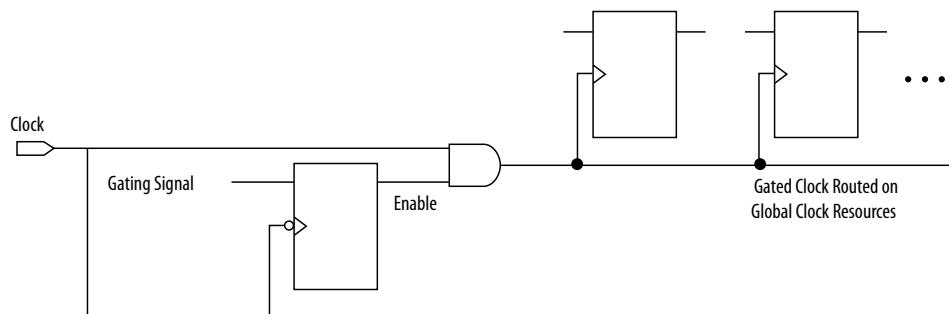
### Recommended Clock-Gating Methods

Use gated clocks only when your target application requires power reduction and when gated clocks are able to provide the required reduction in your device architecture.

If you must use clocks gated by logic, implement these clocks using the robust clock-gating technique and ensure that the gated clock signal uses dedicated global clock routing.

You can gate a clock signal at the source of the clock network, at each register, or somewhere in between. Because the clock network contributes to switching power consumption, gate the clock at the source whenever possible, so that you can shut down the entire clock network instead of gating it further along the clock network at the registers.

Figure 11-8: Recommended Clock-Gating Technique



A register generates the enable signal to ensure that the signal is free of glitches and spikes. The register that generates the enable signal is triggered on the inactive edge of the clock to be gated. Use the falling edge when gating a clock that is active on the rising edge. Using this technique, only one input of the gate that turns the clock on and off changes at a time. This prevents glitches or spikes on the output. Use an AND gate to gate a clock that is active on the rising edge. For a clock that is active on the falling edge, use an OR gate to gate the clock and register the enable command with a positive edge-triggered register.



When using this technique, pay close attention to the duty cycle of the clock and the delay through the logic that generates the enable signal because you must generate the enable command in one-half the clock cycle. This situation might cause problems if the logic that generates the enable command is particularly complex, or if the duty cycle of the clock is severely unbalanced. However, careful management of the duty cycle and logic delay may be an acceptable solution when compared with problems created by other methods of gating clocks.

Ensure that you apply a clock setting to the gated clock in the TimeQuest analyzer. Apply a clock setting to the output of the AND gate. Otherwise, the timing analyzer might analyze the circuit using the clock path through the register as the longest clock path and the path that skips the register as the shortest clock path, resulting in artificial clock skew.

In certain cases, converting the gated clocks to clock enables may help reduce glitch and clock skew, and eventually produce a more accurate timing analysis. You can set the Quartus II software to automatically convert gated clocks to clock enables by turning on the **Auto Gated Clock Conversion** option. The conversion applies to two types of gated clocking schemes: single-gated clock and cascaded-gated clock.

## Optimizing Physical Implementation and Timing Closure

This section provides design and timing closure techniques for high speed or complex core logic designs with challenging timing requirements. These techniques may also be helpful for low or medium speed designs.

### Planning Physical Implementation

When planning a design, consider the following elements of physical implementation:

- The number of unique clock domains and their relationships
- The amount of logic in each functional block
- The location and direction of data flow between blocks
- How data routes to the functional blocks between I/O interfaces

Interface-wide control or status signals may have competing or opposing constraints. For example, when a functional block's control or status signals interface with physical channels from both sides of the device. In such cases you must provide enough pipeline register stages to allow these signals to traverse the width of the device. In addition, you can structure the hierarchy of the design into separate logic modules for each side of the device. The side modules can generate and use registered control signals per side. This simplifies floorplanning, particularly in designs with transceivers, by placing per-side logic near the transceivers.

When adding register stages to pipeline control signals, turn off the **Auto Shift Register Replacement** option (**Assignments > Settings > Compiler Settings > Advanced Settings (Synthesis)**) for these registers. By default, chains of registers can be converted to a RAM-based implementation based on performance and resource estimates. Since pipelining helps meet timing requirements over long distance, this assignment ensures that control signals are not converted.

### Planning FPGA Resources

Your design requirements impact the use of FPGA resources. Plan functional blocks with appropriate global, regional, and dual-regional network signals in mind.

In general, after allocating the clocks in a design, use global networks for the highest fan-out control signals. When a global network signal distributes a high fan-out control signal, the global signal can drive logic anywhere in the device. Similarly, when using a regional network signal, the driven must be in one quadrant of the device, or half the device for a dual-regional network signal. Depending on data flow and physical locations of the data entry and exit between the I/Os and the device, restricting a functional block to a quadrant or half the device may not be practical for performance or resource requirements.

When floorplanning a design, consider the balance of different types of device resources, such as memory, logic, and DSP blocks in the main functional blocks. For example, if a design is memory intensive with a small amount of logic, it may be difficult to develop an effective floorplan. Logic that interfaces with the memory would have to spread across the chip to access the memory. In this case, it is important to use enough register stages in the data and control paths to allow signals to traverse the chip to access the physically disparate resources needed.

## Optimizing Timing Closure

You can make changes to your design and constraints that help you achieve timing closure.

Whenever you change the project settings, you must balance any performance improvement of the setting against any potential increase in compilation time associated with the setting. You can view the performance gain versus runtime cost by reviewing the Fitter messages after design processing.

You can use physical synthesis optimizations for combinational logic, register retiming, and register duplication techniques to optimize your design for timing closure.

Click **Assignments > Settings > Compiler Settings > Advanced Settings (Fitter)** to turn on physical synthesis options.

- Physical synthesis for combinational logic—When the **Perform physical synthesis for combinational logic** is turned on, the report panel identifies logic that physical synthesis can modify. You can use this information to modify the design so that the associated optimization can be turned off to save compile time.
- Register duplication—This technique is most useful where registers have high fan-out, or where the fan-out is in physically distant areas of the device. Review the netlist optimizations report and consider manually duplicating registers automatically added by physical synthesis. You can also locate the original and duplicate registers in the Chip Planner. Compare their locations, and if the fan-out is improved, modify the code and turn off register duplication to save compile time.
- Register retiming—This technique is particularly useful where some combinatorial paths between registers exceed the timing goal while other paths fall short. If a design is already heavily pipelined, register retiming is less likely to provide significant performance gains since there should not be significantly unbalanced levels of logic across pipeline stages.

The application of appropriate timing constraints is essential to timing closure. Use the following general guidelines in applying timing constraints:

- Apply multicycle constraints in your design wherever single-cycle timing analysis is not required.
- Apply False Path constraints to all asynchronous clock domain crossings or resets in the design. This technique prevents overconstraining and the Fitter focuses only on critical paths to reduce compile time. However, over constraining timing critical clock domains can sometimes provide better timing results and lower compile times than physical synthesis.
- Overconstrain rather than using physical synthesis when the slack improvement from physical synthesis is near zero. Overconstrain the frequency requirement on timing critical clock domains by using setup uncertainty.
- When evaluating the effect of constraint changes on performance and runtime, compile the design with at least three different seeds to determine the average performance and runtime effects. Different constraint combinations produce various results. Three samples or more establishes a performance trend. Modify your constraints based on performance improvement or decline.
- Leave settings at the default value whenever possible. Increasing performance constraints can increase the compile time significantly. While those increases may be necessary to close timing on a design, using the default settings whenever possible minimizes compile time.

## Optimizing Critical Timing Paths

To close timing in high speed designs, review paths with the largest timing failures. Correcting a single, large timing failure can result in a very significant timing improvement.

Review the register placement and routing paths by clicking **Tools > Chip Planner**. Large timing failures on high fan-out control signals can be caused by any of the following conditions:

- Sub-optimal use of global networks
- Signals that traverse the chip on local routing without pipelining
- Failure to correct high fan-out by register duplication

For high-speed and high-bandwidth designs, optimize speed by reducing bus width and wire usage. To reduce wire use, move the data as little as possible. For example, if a block of logic functions on a few bits of a word, store inactive bits in a fifo or memory. Memory is cheaper and denser than registers and reduces wire usage.

## Optimizing Power Consumption

The total FPGA power consumption is comprised of I/O power, core static power, and core dynamic power. Knowledge of the relationship between these components is fundamental in calculating the overall total power consumption.

You can use various optimization techniques and tools to minimize power consumption when applied during FPGA design implementation. The Quartus II software offers power-driven compilation features to fully optimize device power consumption. Power-driven compilation focuses on reducing your design's total power consumption using power-driven synthesis and power-driven placement and routing.

## Managing Design Metastability

Metastability in PLD designs can be caused by the synchronization of asynchronous signals. You can use the Quartus II software to analyze the mean time between failures (MTBF) due to metastability, thus optimizing the design to improve the metastability MTBF. A high metastability MTBF indicates a more robust design.

### Related Information

#### [Viewing Metastability Reports](#)

For more information about viewing metastability reports, refer to [Viewing Metastability Reports](#) in Quartus II Help.

## Checking Design Violations

To improve the reliability, timing performance, and logic utilization of your design, avoid design rule violations. The Quartus II software provides the Design Assistant tool that automatically checks for design rule violations and reports their location.

The Design Assistant is a design rule checking tool that allows you to check for design issues early in the design flow. The Design Assistant checks your design for adherence to Altera-recommended design guidelines. You can specify which rules you want the Design Assistant to apply to your design. This is useful if you know that your design violates particular rules that are not critical and you can allow these rule violations. The Design Assistant generates design violation reports with details about each violation based on the settings that you specified.

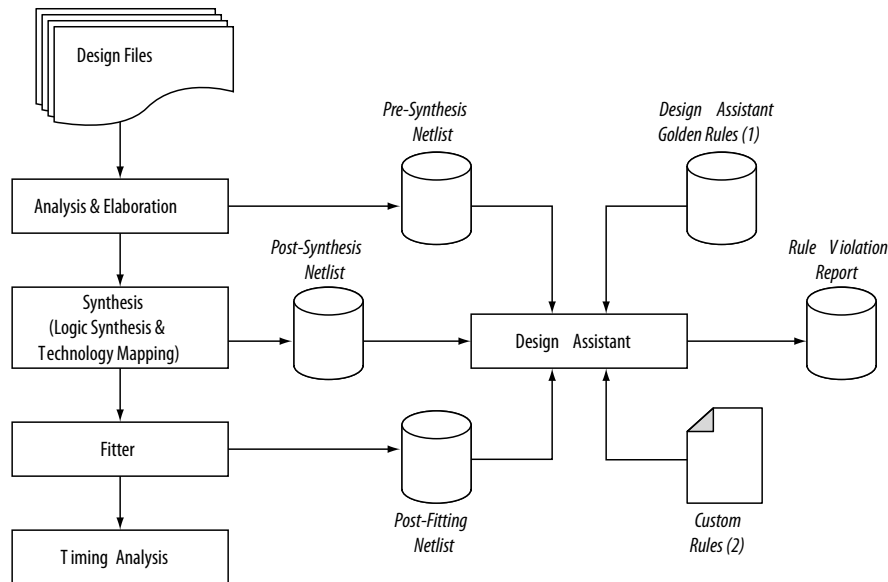
This section provides an introduction to the Quartus II design flow with the Design Assistant, message severity levels, and an explanation about how to set up the Design Assistant. The last parts of the section describe the design rules and the reports generated by the Design Assistant. The Design Assistant supports all Altera devices supported by the Quartus II software.

## Validating Against Design Rules

You can run the Design Assistant following design synthesis or compilation. The Design Assistant performs a post-fit netlist analysis of your design.

The default is to apply all of the rules to your project. If there are some rules that are unimportant to your design, you can turn off the rules that you do not want the Design Assistant to use.

Figure 11-9: Quartus II Design Flow with the Design Assistant



1. Database of the default rules for the Design Assistant.
2. A file that contains the **.xml** codes of the custom rules for the Design Assistant. For more details about how to create this file .

The Design Assistant analyzes your design netlist at different stages of the compilation flow and may yield different warnings or errors, even though the netlists are functionally the same. Your pre-synthesis, post-synthesis, and post-fitting netlists might be different due to optimizations performed by the Quartus II software. For example, a warning message in a pre-synthesis netlist may be removed after the netlist has been synthesized into a post-synthesis or post-fitting netlist.

The exact operation of the Design Assistant depends on when you run it:

- When you run the Design Assistant after running a full compilation or fitting, the Design Assistant performs a post-fitting analysis on the design.
- When you run the Design Assistant after performing Analysis and Synthesis, the Design Assistant performs post-synthesis analysis on the design.
- When you start the Design Assistant after performing Analysis and Elaboration, the Design Assistant performs a pre-synthesis analysis on the design. You can also perform pre-synthesis analysis with the Design Assistant using the command-line. You can use the `-rtl` option with the `quartus_drc` executable, as shown in the following example:

```
quartus_drc <project_name> --rtl=on
```

If your design violates a design rule, the Design Assistant generates warning messages and information messages about the violated rule. The Design Assistant displays these messages in the Messages window, in the Design Assistant Messages report, and in the Design Assistant report files. You can find the Design Assistant report files called `<project_name>.drc.rpt` in the `<project_name>` subdirectory of the project directory.

### Related Information

#### [About the Design Assistant](#)

## Creating Custom Design Rules

You can define and validate your design against your own custom set of design rules. You can save these rules in a text file (with any file extension) with the XML format.

You then specify the path to that file in the Design Assistant settings page and run the Design Assistant for violation checking.

Refer to the following location to locate the file that contains the default rules for the Design Assistant:

```
<Quartus II install path>\quartus\libraries\design-assistant\da_golden_rule.xml
```

## Custom Design Rule Examples

The following examples of custom rules show how to check node relationships and clock relationships in a design.

This example shows the XML codes for checking SR latch structures in a design.

### Example 11-1: Detecting SR Latches in a Design

```
<DA_RULE ID="EX01" SEVERITY="CRITICAL" NAME="Checking Design for SR Latch"
DEFAULT_RUN="YES">
<RULE_DEFINITION>
  <FORBID>
    <OR>
      <NODE NAME="NODE_1" TYPE="SRLATCH" />
      <HAS_NODE NODE_LIST="NODE_1" />
      <NODE NAME="NODE_1" TOTAL_FANIN="EQ2" />
      <NODE NAME="NODE_2" TOTAL_FANIN="EQ2" />
    <AND>
      <NODE_RELATIONSHIP FROM_NAME="NODE_1" FROM_TYPE="NAND"
TO_NAME="NODE_2" TO_TYPE="NAND" />
      <NODE_RELATIONSHIP FROM_NAME="NODE_2" FROM_TYPE="NAND"
TO_NAME="NODE_1" TO_TYPE="NAND" />
    </AND>
    <AND>
      <NODE_RELATIONSHIP FROM_NAME="NODE_1" FROM_TYPE="NOR"
TO_NAME="NODE_2" TO_TYPE="NOR" />
    </AND>
  </FORBID>
</RULE_DEFINITION>
```

```

        <NODE_RELATIONSHIP FROM_NAME="NODE_2" FROM_TYPE="NOR"
TO_NAME="NODE_1" TO_TYPE="NOR" />
    </AND>
    </OR>
</FORBID>
</RULE_DEFINITION>

<REPORTING_ROOT>
    <MESSAGE NAME="Rule %ARG1%: Found %ARG2% node(s) related to this rule.">
        <MESSAGE_ARGUMENT NAME="ARG1" TYPE="ATTRIBUTE" VALUE="ID" />
        <MESSAGE_ARGUMENT NAME="ARG2" TYPE="TOTAL_NODE" VALUE="NODE_1" />
    </MESSAGE>
</REPORTING_ROOT>
</DA_RULE>

```

The possible SR latch structures are specified in the rule definition section. Codes defined in the <AND></AND> block are tied together, meaning that each statement in the block must be true for the block to be fulfilled (AND gate similarity). In the <OR></OR> block, as long as one statement in the block is true, the block is fulfilled (OR gate similarity). If no <AND></AND> or <OR></OR> blocks are specified, the default is <AND></AND>.

The <FORBID></FORBID> section contains the undesirable condition for the design, which in this case is the SR latch structures. If the condition is fulfilled, the Design Assistant highlights a rule violation.


### Example 11-2: Detecting SR Latches in a Design

```

<AND>
    <NODE_RELATIONSHIP FROM_NAME="NODE_1" FROM_TYPE="NAND" TO_NAME="NODE_2"
    TO_TYPE="NAND" />
    <NODE_RELATIONSHIP FROM_NAME="NODE_2" FROM_TYPE="NAND" TO_NAME="NODE_1"
    TO_TYPE="NAND" />
</AND>

```

Figure 11-10: Undesired Condition 1




```

<AND>
    <NODE_RELATIONSHIP FROM_NAME="NODE_1" FROM_TYPE="NOR" TO_NAME="NODE_2"
TO_TYPE="NOR" />
    <NODE_RELATIONSHIP FROM_NAME="NODE_2" FROM_TYPE="NOR" TO_NAME="NODE_1"
TO_TYPE="NOR" />
</AND>

```

Figure 11-11: Undesired Condition 2



This example shows how to use the `CLOCK_RELATIONSHIP` attribute to relate nodes to clock domains. This example checks for correct synchronization in data transfer between asynchronous clock domains. Synchronization is done with cascaded registers, also called synchronizers, at the receiving clock domain. The code in This example checks for the synchronizer configuration based on the following guidelines:

- The cascading registers need to be triggered on the same clock edge
- There is no logic between the register output of the transmitting clock domain and the cascaded registers in the receiving asynchronous clock domain.

### Example 11-3: Detecting Incorrect Synchronizer Configuration

```
<DA_RULE ID="EX02" SEVERITY="HIGH" NAME="Data Transfer Not Synch Correctly"
DEFAULT_RUN="YES">

<RULE_DEFINITION>
<DECLARE>
  <NODE NAME="NODE_1" TYPE="REG" />
  <NODE NAME="NODE_2" TYPE="REG" />
  <NODE NAME="NODE_3" TYPE="REG" />
</DECLARE>
<FORBID>
  <NODE_RELATIONSHIP FROM_NAME="NODE_1" TO_NAME="NODE_2" TO_PORT="D_PORT"
CLOCK_RELATIONSHIP="ASYN" />
  <NODE_RELATIONSHIP FROM_NAME="NODE_2" TO_NAME="NODE_3" TO_PORT="D_PORT"
CLOCK_RELATIONSHIP="!ASYN" />
  <OR>
    <NODE_RELATIONSHIP FROM_NAME="NODE_1" TO_NAME="NODE_2"
TO_PORT="D_PORT" REQUIRED_THROUGH="YES" THROUGH_TYPE="COMB" CLOCK_RELATION-
SHIP="ASYN" />
    <CLOCK_RELATIONSHIP NAME="SEQ_EDGE|ASYN" NODE_LIST="NODE_2, NODE_3" />
  </OR>
</FORBID>
</RULE_DEFINITION>

<REPORTING_ROOT>
<MESSAGE NAME="Rule %ARG1%: Found %ARG2% node(s) related to this rule.">
  <MESSAGE_ARGUMENT NAME="ARG1" TYPE="ATTRIBUTE" VALUE="ID" />
  <MESSAGE_ARGUMENT NAME="ARG2" TYPE="TOTAL_NODE" VALUE="NODE_1" />
  <MESSAGE NAME="Source node(s): %ARG3%, Destination node(s): %ARG4%">
    <MESSAGE_ARGUMENT NAME="ARG3" TYPE="NODE" VALUE="NODE_1" />
    <MESSAGE_ARGUMENT NAME="ARG4" TYPE="NODE" VALUE="NODE_2" />
  </MESSAGE>
</MESSAGE>
```

```
</MESSAGE>  
</REPORTING_ROOT>  
</DA_RULE>
```

The codes differentiate the clock domains. `ASYN` means asynchronous, and `!ASYN` means non-asynchronous. This notation is useful for describing nodes that are in different clock domains. The following lines from the example state that `NODE_2` and `NODE_3` are in the same clock domain, but `NODE_1` is not.

```
<NODE_RELATIONSHIP FROM_NAME="NODE_1" TO_NAME="NODE_2" TO_PORT="D_PORT"  
CLOCK_RELATIONSHIP="ASYN" />  
  
<NODE_RELATIONSHIP FROM_NAME="NODE_2" TO_NAME="NODE_3" TO_PORT="D_PORT"  
CLOCK_RELATIONSHIP="!ASYN" />
```

The next line of code states that `NODE_2` and `NODE_3` have a clock relationship of either sequential edge or asynchronous.

```
<CLOCK_RELATIONSHIP NAME="SEQ_EDGE|ASYN" NODE_LIST="NODE_2, NODE_3" />
```

The `<FORBID></FORBID>` section contains the undesirable condition for the design, which in this case is the undesired configuration of the synchronizer. If the condition is fulfilled, the Design Assistant highlights a rule violation.

The possible SR latch structures are specified in the rule definition section. Codes defined in the `<AND></AND>` block are tied together, meaning that each statement in the block must be true for the block to be fulfilled (AND gate similarity). In the `<OR></OR>` block, as long as one statement in the block is true, the block is fulfilled (OR gate similarity). If no `<AND></AND>` or `<OR></OR>` blocks are specified, the default is `<AND></AND>`.

The `<FORBID></FORBID>` section contains the undesirable condition for the design, which in this case is the SR latch structures. If the condition is fulfilled, the Design Assistant highlights a rule violation.


The following examples show the undesired conditions from with their equivalent block diagrams:

### Example 11-4: Undesired Condition 3

```
<NODE_RELATIONSHIP FROM_NAME="NODE_1" TO_NAME="NODE_2" TO_PORT="D_PORT"  
CLOCK_RELATIONSHIP="ASYN" />  
  
<NODE_RELATIONSHIP FROM_NAME="NODE_2" TO_NAME="NODE_3" TO_PORT="D_PORT"  
CLOCK_RELATIONSHIP="!ASYN" />  
  
<NODE_RELATIONSHIP FROM_NAME="NODE_1" TO_NAME="NODE_2" TO_PORT="D_PORT"  
REQUIRED_THROUGH="YES"  
THROUGH_TYPE="COMB" CLOCK_RELATIONSHIP="ASYN" />
```



Figure 11-12: Undesired Condition 3




Example 11-5: Undesired Condition 4

```
<NODE_RELATIONSHIP FROM_NAME="NODE_1" TO_NAME="NODE_2" TO_PORT="D_PORT"
CLOCK_RELATIONSHIP="ASYN" />

<NODE_RELATIONSHIP FROM_NAME="NODE_2" TO_NAME="NODE_3" TO_PORT="D_PORT"
CLOCK_RELATIONSHIP="!ASYN" />

<CLOCK_RELATIONSHIP NAME="SEQ_EDGE|ASYN" NODE_LIST="NODE_2, NODE_3" />
```

Figure 11-13: Undesired Condition 4



## Use Clock and Register-Control Architectural Features

In addition to following general design guidelines, you must code your design with the device architecture in mind. FPGAs provide device-wide clocks and register control signals that can improve performance.

### Use Global Clock Network Resources

Altera FPGAs provide device-wide global clock routing resources and dedicated inputs. Use the FPGA's low-skew, high fan-out dedicated routing where available.

By assigning a clock input to one of these dedicated clock pins or with a Quartus II logic option to assign global routing, you can take advantage of the dedicated routing available for clock signals.

In an ASIC design, you should balance the clock delay as it is distributed across the device. Because Altera FPGAs provide device-wide global clock routing resources and dedicated inputs, there is no need to manually balance delays on the clock network.

You should limit the number of clocks in your design to the number of dedicated global clock resources available in your FPGA. Clocks feeding multiple locations that do not use global routing may exhibit clock skew across the device that could lead to timing problems. In addition, when you use combinational logic to generate an internal clock, it adds delays on the clock path. In some cases, delay on a clock line can result in a clock skew greater than the data path length between two registers. If the clock skew is greater than the data delay, you violate the timing parameters of the register (such as hold time requirements) and the design does not function correctly.

FPGAs offer a number of low-skew global routing resources to distribute high fan-out signals to help with the implementation of large designs with many clock domains. Many large FPGA devices provide dedicated global clock networks, regional clock networks, and dedicated fast regional clock networks. These clocks are organized into a hierarchical clock structure that allows many clocks in each device region with low skew and delay. There are typically several dedicated clock pins to drive either global or regional clock networks, and both PLL outputs and internal clocks can drive various clock networks.

To reduce clock skew in a given clock domain and ensure that hold times are met in that clock domain, assign each clock signal to one of the global high fan-out, low-skew clock networks in the FPGA device. The Quartus II software automatically uses global routing for high fan-out control signals, PLL outputs, and signals feeding the global clock pins on the device. You can make explicit Global Signal logic option settings by turning on the **Global Signal** option setting. Use this option when it is necessary to force the software to use the global routing for particular signals.

To take full advantage of these routing resources, the sources of clock signals in a design (input clock pins or internally-generated clocks) need to drive only the clock input ports of registers. In older Altera device families, if a clock signal feeds the data ports of a register, the signal may not be able to use dedicated routing, which can lead to decreased performance and clock skew problems. In general, allowing clock signals to drive the data ports of registers is not considered synchronous design and can complicate timing analysis.

## Use Global Reset Resources

ASIC designs may use local resets to avoid long routing delays. Take advantage of the device-wide asynchronous reset pin available on most FPGAs to eliminate these problems. This reset signal provides low-skew routing across the device.

The following are three types of resets used in synchronous circuits:

- Synchronous Reset
- Asynchronous Reset
- Synchronized Asynchronous Reset—preferred when designing an FPGA circuit

## Use Synchronous Resets

The synchronous reset ensures that the circuit is fully synchronous. You can easily time the circuit with the Quartus II TimeQuest analyzer.

Because clocks that are synchronous to each other launch and latch the reset signal, the data arrival and data required times are easily determined for proper slack analysis. The synchronous reset is easier to use with cycle-based simulators.

There are two methods by which a reset signal can reach a register; either by being gated in with the data input, or by using an LAB-wide control signal (`sync1r`). If you use the first method, you risk adding an additional gate delay to the circuit to accommodate the reset signal, which causes increased data arrival times and negatively impacts setup slack. The second method relies on dedicated routing in the LAB to each register, but this is slower than an asynchronous reset to the same register.

Figure 11-14: Synchronous Reset



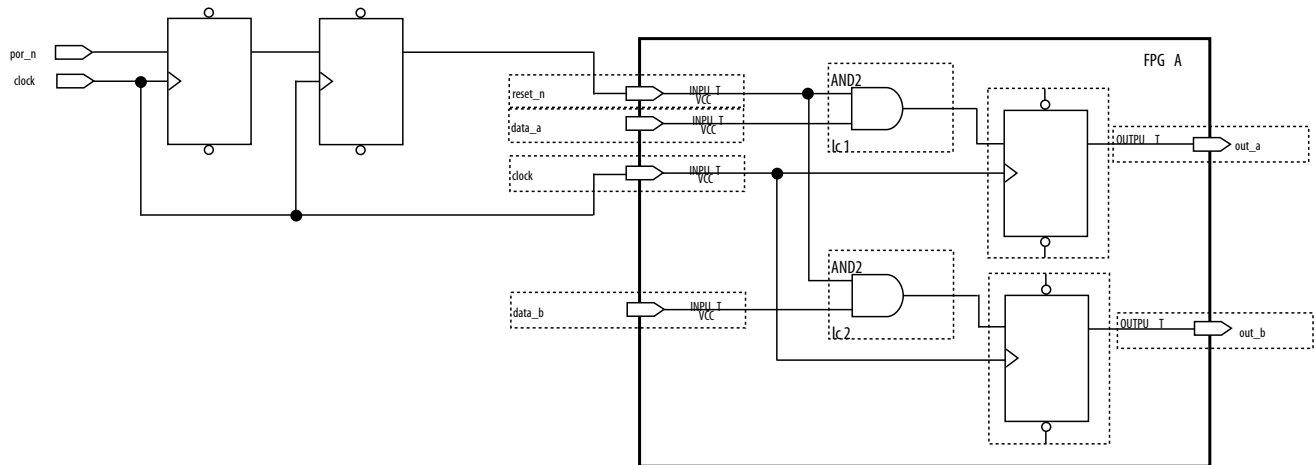


Figure 11-15: LAB-Wide Control Signals



Consider two types of synchronous resets when you examine the timing analysis of synchronous resets—externally synchronized resets and internally synchronized resets. Externally synchronized resets are synchronized to the clock domain outside the FPGA, and are not very common. A power-on asynchronous reset is dual-rank synchronized externally to the system clock and then brought into the FPGA. Inside the FPGA, gate this reset with the data input to the registers to implement a synchronous reset.

Figure 11-16: Externally Synchronized Reset



The following example shows the Verilog equivalent of the schematic. When you use synchronous resets, the reset signal is not put in the sensitivity list.

The following example shows the necessary modifications that you should make to the internally synchronized reset.

#### Example 11-6: Verilog Code for Externally Synchronized Reset

```

module sync_reset_ext (
    input  clock,
    input  reset_n,
    input  data_a,
    input  data_b,
    output out_a,
    output out_b
);
    reg  reg1, reg2;
    assign out_a = reg1;
    assign out_b = reg2;
    always @ (posedge clock)
    begin
        if (!reset_n)
            begin
                reg1    <= 1'b0;
                reg2    <= 1'b0;
            end
        else
            begin
                reg1    <= data_a;
                reg2    <= data_b;
            end
        end
    end
endmodule // sync_reset_ext

```

The following example shows the constraints for the externally synchronous reset. Because the external reset is synchronous, you only need to constrain the `reset_n` signal as a normal input signal with `set_input_delay` constraint for `-max` and `-min`.

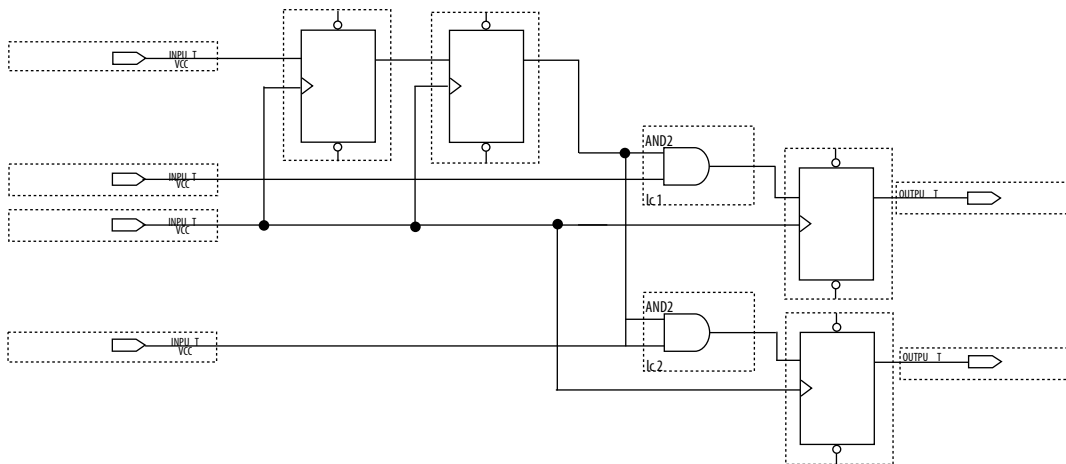
**Example 11-7: SDC Constraints for Externally Synchronized Reset**

```

# Input clock - 100 MHz
create_clock [get_ports {clock}] \
  -name {clock} \
  -period 10.0 \
  -waveform {0.0 5.0}
# Input constraints on low-active reset
# and data
set_input_delay 7.0 \
  -max \
  -clock [get_clocks {clock}] \
  [get_ports {reset_n data_a data_b}]
set_input_delay 1.0 \
  -min \
  -clock [get_clocks {clock}] \
  [get_ports {reset_n data_a data_b}]

```

More often, resets coming into the device are asynchronous, and must be synchronized internally before being sent to the registers.

**Figure 11-17: Internally Synchronized Reset**

The following example shows the Verilog equivalent of the schematic. Only the clock edge is in the sensitivity list for a synchronous reset.

**Example 11-8: Verilog Code for Internally Synchronized Reset**

```

module sync_reset_ext (
  input  clock,
  input  reset_n,
  input  data_a,
  input  data_b,
  output out_a,
  output out_b
);
  reg  reg1, reg2;
  assign out_a = reg1;

```

```
assign out_b = reg2;
always @ (posedge clock)
begin
    if (!reset_n)
    begin
        reg1    <= 1'b0;
        reg2    <= 1'b0;
    end
    else
    begin
        reg1    <= data_a;
        reg2    <= data_b;
    end
end
endmodule // sync_reset_ext
```


The SDC constraints are similar to the external synchronous reset, except that the input reset cannot be constrained because it is asynchronous and should be cut with a `set_false_path` statement to avoid these being considered as unconstrained paths.

### Example 11-9: SDC Constraints for Internally Synchronized Reset

```
# Input clock - 100 MHz
create_clock [get_ports {clock}] \
    -name {clock} \
    -period 10.0 \
    -waveform {0.0 5.0}
# Input constraints on data
set_input_delay 7.0 \
    -max \
    -clock [get_clocks {clock}] \
    [get_ports {data_a data_b}]
set_input_delay 1.0 \
    -min \
    -clock [get_clocks {clock}] \
    [get_ports {data_a data_b}]
# Cut the asynchronous reset input
set_false_path \
    -from [get_ports {reset_n}] \
    -to [all_registers]
```

An issue with synchronous resets is their behavior with respect to short pulses (less than a period) on the asynchronous input to the synchronizer flipflops. This can be a disadvantage because the asynchronous reset requires a pulse width of at least one period wide to guarantee that it is captured by the first flipflop. However, this can also be viewed as an advantage in that this circuit increases noise immunity. Spurious pulses on the asynchronous input have a lower chance of being captured by the first flipflop, so the pulses do not trigger a synchronous reset. In some cases, you might want to increase the noise immunity further and reject any asynchronous input reset that is less than  $n$  periods wide to debounce an asynchronous input reset.

Figure 11-18: Internally Synchronized Reset with Pulse Extender



1. Junction dots indicate the number of stages. You can have more flip flops to get a wider pulse that spans more clock cycles.

Many designs have more than one clock signal. In these cases, use a separate reset synchronization circuit for each clock domain in the design. When you create synchronizers for PLL output clocks, these clock domains are not reset until you lock the PLL and the PLL output clocks are stable. If you use the reset to the PLL, this reset does not have to be synchronous with the input clock of the PLL. You can use an asynchronous reset for this. Using a reset to the PLL further delays the assertion of a synchronous reset to the PLL output clock domains when using internally synchronized resets.


## Using Asynchronous Resets

Asynchronous resets are the most common form of reset in circuit designs, as well as the easiest to implement. Typically, you can insert the asynchronous reset into the device, turn on the global buffer, and connect to the asynchronous reset pin of every register in the device.

This method is only advantageous under certain circumstances—you do not need to always reset the register. Unlike the synchronous reset, the asynchronous reset is not inserted in the data path, and does not negatively impact the data arrival times between registers. Reset takes effect immediately, and as soon as the registers receive the reset pulse, the registers are reset. The asynchronous reset is not dependent on the clock.

However, when the reset is deasserted and does not pass the recovery ( $\mu\text{t}_{\text{SU}}$ ) or removal ( $\mu\text{t}_{\text{H}}$ ) time check (the TimeQuest analyzer recovery and removal analysis checks both times), the edge is said to have fallen into the metastability zone. Additional time is required to determine the correct state, and the delay can cause the setup time to fail to register downstream, leading to system failure. To avoid this, add a few follower registers after the register with the asynchronous reset and use the output of these registers in the design. Use the follower registers to synchronize the data to the clock to remove the metastability issues. You should place these registers close to each other in the device to keep the routing delays to a minimum, which decreases data arrival times and increases MTBF. Ensure that these follower registers themselves are not reset, but are initialized over a period of several clock cycles by “flushing out” their current or initial state.

**Figure 11-19: Asynchronous Reset with Follower Registers**



The following example shows the equivalent Verilog code. The active edge of the reset is now in the sensitivity list for the procedural block, which infers a clock enable on the follower registers with the inverse of the reset signal tied to the clock enable. The follower registers should be in a separate procedural block as shown using non-blocking assignments.

**Example 11-10: Verilog Code of Asynchronous Reset with Follower Registers**

```

module async_reset (
    input  clock,
    input  reset_n,
    input  data_a,
    output out_a,
);
    reg  reg1, reg2, reg3;
    assign out_a = reg3;
    always @ (posedge clock, negedge reset_n)
    begin
        if (!reset_n)
            reg1  <= 1'b0;
        else
            reg1  <= data_a;
    end
    always @ (posedge clock)
    begin
        reg2  <= reg1;
        reg3  <= reg2;
    end
endmodule // async_reset

```

You can easily constrain an asynchronous reset. By definition, asynchronous resets have a non-deterministic relationship to the clock domains of the registers they are resetting. Therefore, static timing analysis of these resets is not possible and you can use the `set_false_path` command to exclude the path from timing analysis. Because the relationship of the reset to the clock at the register is not known, you cannot run recovery and removal analysis in the TimeQuest analyzer for this path. Attempting to do so even without the false path statement results in no paths reported for recovery and removal.

**Example 11-11: SDC Constraints for Asynchronous Reset**

```

# Input clock - 100 MHz
create_clock [get_ports {clock}] \

```



```

        -name {clock} \
        -period 10.0 \
        -waveform {0.0 5.0}
# Input constraints on data
set_input_delay 7.0 \
    -max \
    -clock [get_clocks {clock}] \
    [get_ports {data_a}]
set_input_delay 1.0 \
    -min \
    -clock [get_clocks {clock}] \
    [get_ports {data_a}]
# Cut the asynchronous reset input
set_false_path \
    -from [get_ports {reset_n}] \
    -to [all_registers]

```

The asynchronous reset is susceptible to noise, and a noisy asynchronous reset can cause a spurious reset. You must ensure that the asynchronous reset is debounced and filtered. You can easily enter into a reset asynchronously, but releasing a reset asynchronously can lead to potential problems (also referred to as “reset removal”) with metastability, including the hazards of unwanted situations with synchronous circuits involving feedback.

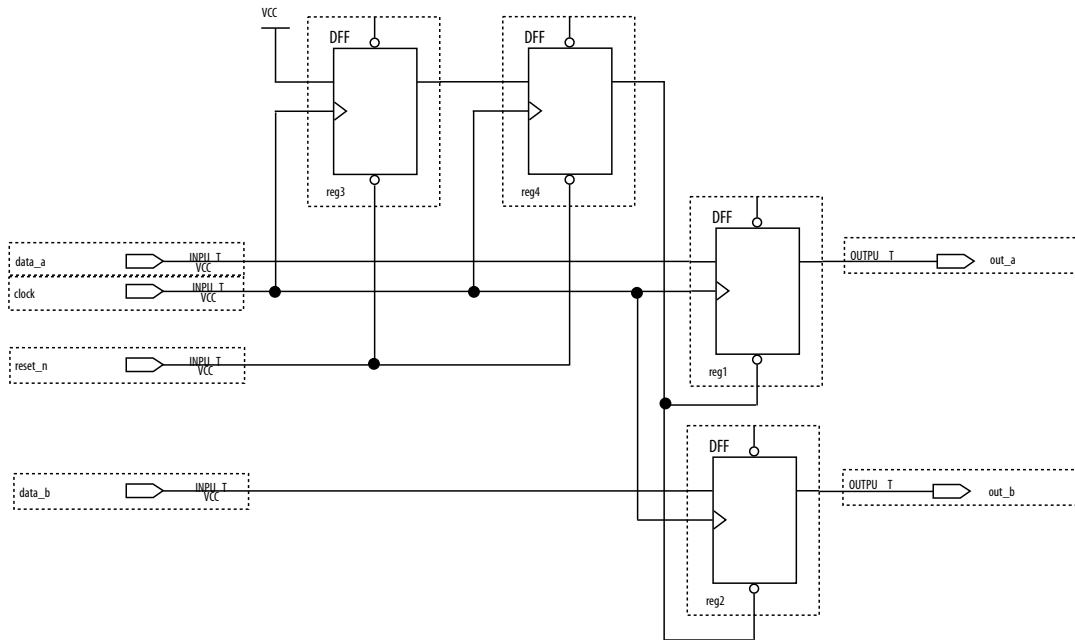
### Use Synchronized Asynchronous Reset

To avoid potential problems associated with purely synchronous resets and purely asynchronous resets, you can use synchronized asynchronous resets. Synchronized asynchronous resets combine the advantages of synchronous and asynchronous resets.

These resets are asynchronously asserted and synchronously deasserted. This takes effect almost instantaneously, and ensures that no data path for speed is involved, and that the circuit is synchronous for timing analysis and is resistant to noise.

The following example shows a method for implementing the synchronized asynchronous reset. You should use synchronizer registers in a similar manner as synchronous resets. However, the asynchronous reset input is gated directly to the `CLRn` pin of the synchronizer registers and immediately asserts the resulting reset. When the reset is deasserted, logic “1” is clocked through the synchronizers to synchronously deassert the resulting reset.

Figure 11-20: Schematic of Synchronized Asynchronous Reset



The following example shows the equivalent Verilog HDL code. Use the active edge of the reset in the sensitivity list for the blocks.

Example 11-12: Verilog Code for Synchronized Asynchronous Reset

```

module sync_async_reset (
    input    clock,
    input    reset_n,
    input    data_a,
    input    data_b,
    output   out_a,
    output   out_b
);
    reg    reg1, reg2;
    reg    reg3, reg4;
    assign out_a    = reg1;
    assign out_b    = reg2;
    assign rst_n    = reg4;
    always @ (posedge clock, negedge reset_n)
    begin
        if (!reset_n)
            begin
                reg3    <= 1'b0;
                reg4    <= 1'b0;
            end
        else
            begin
                reg3    <= 1'b1;
                reg4    <= reg3;
            end
        end
    end
    always @ (posedge clock, negedge rst_n)
    begin
        if (!rst_n)
    
```

```

begin
    reg1    <= 1'b0;
    reg2    <= 1'b0;
end
else
begin
    reg1    <= data_a;
    reg2    <= data_b;
end
end
endmodule // sync_async_reset

```

To minimize the metastability effect between the two synchronization registers, and to increase the MTBF, the registers should be located as close as possible in the device to minimize routing delay. If possible, locate the registers in the same logic array block (LAB). The input reset signal (`reset_n`) must be excluded with a `set_false_path` command:

```
set_false_path -from [get_ports {reset_n}] -to [all_registers]
```

The `set_false_path` command used with the specified constraint excludes unnecessary input timing reports that would otherwise result from specifying an input delay on the reset pin.

The instantaneous assertion of synchronized asynchronous resets is susceptible to noise and runt pulses. If possible, you should debounce the asynchronous reset and filter the reset before it enters the device. The circuit ensures that the synchronized asynchronous reset is at least one full clock period in length. To extend this time to  $n$  clock periods, you must increase the number of synchronizer registers to  $n + 1$ . You must connect the asynchronous input reset (`reset_n`) to the `CLRN` pin of all the synchronizer registers to maintain the asynchronous assertion of the synchronized asynchronous reset.

## Avoid Asynchronous Register Control Signals

Avoid using an asynchronous load signal if the design target device architecture does not include registers with dedicated circuitry for asynchronous loads. Also, avoid using both asynchronous clear and preset if the architecture provides only one of these control signals.

Some Altera devices directly support an asynchronous clear function, but not a preset or load function. When the target device does not directly support the signals, the synthesis or placement and routing software must use combinational logic to implement the same functionality. In addition, if you use signals in a priority other than the inherent priority in the device architecture, combinational logic may be required to implement the necessary control signals. Combinational logic is less efficient and can cause glitches and other problems; it is best to avoid these implementations.

## Implementing Embedded RAM

Altera's dedicated memory architecture offers many advanced features that you can enable with Altera-provided IP cores. Use synchronous memory blocks for your design, so that the blocks can be mapped directly into the device dedicated memory blocks.

You can use single-port, dual-port, or three-port RAM with a single- or dual-clocking method. You should not infer the asynchronous memory logic as a memory block or place the asynchronous memory logic in the dedicated memory block, but implement the asynchronous memory logic in regular logic cells.

Altera memory blocks have different read-during-write behaviors, depending on the targeted device family, memory mode, and block type. Read-during-write behavior refers to read and write from the same

memory address in the same clock cycle; for example, you read from the same address to which you write in the same clock cycle.

You should check how you specify the memory in your HDL code when you use read-during-write behavior. The HDL code that describes the read returns either the old data stored at the memory location, or the new data being written to the memory location.

In some cases, when the device architecture cannot implement the memory behavior described in your HDL code, the memory block is not mapped to the dedicated RAM blocks, or the memory block is implemented using extra logic in addition to the dedicated RAM block. Implement the read-during-write behavior using single-port RAM in Arria GX devices and the Cyclone and Stratix series of devices to avoid this extra logic implementation.

In many synthesis tools, you can specify that the read-during-write behavior is not important to your design; if, for example, you never read and write from the same address in the same clock cycle. For Quartus II integrated synthesis, add the synthesis attribute `ramstyle="no_rw_check"` to allow the software to choose the read-during-write behavior of a RAM, rather than using the read-during-write behavior specified in your HDL code. Using this type of attribute prevents the synthesis tool from using extra logic to implement the memory block and, in some cases, can allow memory inference when it would otherwise be impossible.

## Document Revision History

Table 11-1: Document Revision History

Date	Version	Changes
2014.12.15	14.1.0	Updated location of Fitter Settings, Analysis & Synthesis Settings, and Physical Optimization Settings to Compiler Settings.
June 2014	14.0.0	Removed references to obsolete MegaWizard Plug-In Manager.
November 2013	13.1.0	Removed HardCopy device information.
May 2013	13.0.0	Removed PrimeTime support.
June 2012	12.0.0	Removed survey link.
November 2011	11.0.1	Template update.
May 2011	11.0.0	Added information to Reset Resources .
December 2010	10.1.0	<ul style="list-style-type: none"> <li>Title changed from Design Recommendations for Altera Devices and the Quartus II Design Assistant.</li> <li>Updated to new template.</li> <li>Added references to Quartus II Help for “Metastability” on page 9–13 and “Incremental Compilation” on page 9–13.</li> <li>Removed duplicated content and added references to Quartus II Help for “Custom Rules” on page 9–15.</li> </ul>

Date	Version	Changes
July 2010	10.0.0	<ul style="list-style-type: none"> <li>Removed duplicated content and added references to Quartus II Help for Design Assistant settings, Design Assistant rules, Enabling and Disabling Design Assistant Rules, and Viewing Design Assistant reports.</li> <li>Removed information from “Combinational Logic Structures” on page 5–4</li> <li>Changed heading from “Design Techniques to Save Power” to “Power Optimization” on page 5–12</li> <li>Added new “Metastability” section</li> <li>Added new “Incremental Compilation” section</li> <li>Added information to “Reset Resources” on page 5–23</li> <li>Removed “Referenced Documents” section</li> </ul>
November 2009	9.1.0	<ul style="list-style-type: none"> <li>Removed documentation of obsolete rules.</li> </ul>
March 2009	9.0.0	<ul style="list-style-type: none"> <li>No change to content.</li> </ul>
November 2008	8.1.0	<ul style="list-style-type: none"> <li>Changed to 8-1/2 x 11 page size</li> <li>Added new section “Custom Rules Coding Examples” on page 5–18</li> <li>Added paragraph to “Recommended Clock-Gating Methods” on page 5–11</li> <li>Added new section: “Design Techniques to Save Power” on page 5–12</li> </ul>
May 2008	8.0.0	<ul style="list-style-type: none"> <li>Updated Figure 5–9 on page 5–13; added custom rules file to the flow</li> <li>Added notes to Figure 5–9 on page 5–13</li> <li>Added new section: “Custom Rules Report” on page 5–34</li> <li>Added new section: “Custom Rules” on page 5–34</li> <li>Added new section: “Targeting Embedded RAM Architectural Features” on page 5–38</li> <li>Minor editorial updates throughout the chapter</li> <li>Added hyperlinks to referenced documents throughout the chapter</li> </ul>

**Related Information**

[http://www.altera.com/literature/lit-qts\\_archive.jsp](http://www.altera.com/literature/lit-qts_archive.jsp)

# Recommended HDL Coding Styles 12

2014.12.15

QII5V1



Subscribe



Send Feedback

This chapter provides Hardware Description Language (HDL) coding style recommendations to ensure optimal synthesis results when targeting Altera devices.

HDL coding styles can have a significant effect on the quality of results that you achieve for program-mable logic designs. Synthesis tools optimize HDL code for both logic utilization and performance; however, synthesis tools have no information about the purpose or intent of the design. The best optimizations require your conscious interaction. The Altera website provides design examples for other types of functions and to target specific applications.

**Note:** For style recommendations, options, or HDL attributes specific to your synthesis tool (including Quartus II integrated synthesis and other EDA tools), refer to the tool vendor's documentation.

## Related Information

- [Recommended Design Practices](#) on page 11-1
- [Advanced Synthesis Cookbook](#)
- [Design Examples](#)
- [Reference Designs](#)
- [Quartus II Integrated Synthesis](#) on page 16-1

## Using Provided HDL Templates

You can use provided HDL templates to start your HDL designs.

Altera provides templates for Verilog HDL, SystemVerilog, and VHDL. Many of the HDL examples in this document correspond with the **Full Designs** examples in the **Quartus II Templates**. You can insert HDL code into your own design using the templates or examples.

## Inserting a HDL Code from the Template

Insert HDL code from a provided template, follow these steps:

1. On the **File** menu, click **New**.
2. In the **New** dialog box, select the type of design file corresponding to the type of HDL you want to use, SystemVerilog HDL File, VHDL File, or Verilog HDL File.
3. Right-click in the HDL file and then click **Insert Template**.
4. In the **Insert Template** dialog box, expand the section corresponding to the appropriate HDL, then expand the **Full Designs** section.


© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



5. Select a design. The HDL appears in the **Preview** pane.
6. Click **Insert** to paste the HDL design to the blank Verilog or VHDL file you created in step 2.
7. Click **Close** to close the **Insert Template** dialog box.

Figure 12-1: Inserting a RAM Template



**Note:** You can use any of the standard features of the Quartus II Text Editor to modify the HDL design or save the template as an HDL file to edit in your preferred text editor.

#### Related Information

[About the Quartus II Text Editor](#)

## Instantiating IP Cores in HDL

Altera provides parameterizable IP cores that are optimized for Altera device architectures. Using IP cores instead of coding your own logic saves valuable design time.

Additionally, the Altera-provided IP cores offer more efficient logic synthesis and device implementation. You can scale the IP core's size and specify various options by setting parameters. You can instantiate the IP core directly in your HDL file code by calling the IP core name and defining its parameters as you would any other module, component, or subdesign. Alternatively, you can use the IP Catalog (**Tools > IP**

**Catalog**) and parameter editor GUI to simplify customization of your IP core variation. You can infer or instantiate IP cores that optimize the following device architecture features:

- Transceivers
- LVDS drivers
- Memory and DSP blocks
- Phase-locked loops (PLLs)
- double-data rate input/output (DDIO) circuitry

For some types of logic functions, such as memories and DSP functions, you can infer device-specific dedicated architecture blocks instead of instantiating an IP core. Quartus II synthesis recognizes certain HDL code structures and automatically infers the appropriate IP core or map directly to device atoms.

#### Related Information

- [Inferring Multipliers and DSP Functions](#) on page 12-3
- [Inferring Memory Functions from HDL Code](#) on page 12-8
- [Altera IP Core Literature](#)

## Inferring Multipliers and DSP Functions

The following sections describe how to infer multiplier and DSP functions from generic HDL code, and, if applicable, how to target the dedicated DSP block architecture in Altera devices.

#### Related Information

[DSP Solutions Center](#)

### Inferring Multipliers

To infer multiplier functions, synthesis tools detect multiplier logic and implement this in Altera IP cores, or map the logic directly to device atoms.

For devices with DSP blocks, the software can implement the function in a DSP block instead of logic, depending on device utilization. The Quartus II Fitter can also place input and output registers in DSP blocks (that is, perform register packing) to improve performance and area utilization.

The Verilog HDL and VHDL code examples show, for unsigned and signed multipliers, that synthesis tools can infer as an IP core or DSP block atoms. Each example fits into one DSP block element. In addition, when register packing occurs, no extra logic cells for registers are required.

**Note:** The `signed` declaration in Verilog HDL is a feature of the Verilog 2001 Standard.

#### Example 12-1: Verilog HDL Unsigned Multiplier

```
module unsigned_mult (out, a, b);
    output [15:0] out;
    input [7:0] a;
    input [7:0] b;
    assign out = a * b;
endmodule
```



**Example 12-2: Verilog HDL Signed Multiplier with Input and Output Registers (Pipelining = 2)**

```

module signed_mult (out, clk, a, b);
    output [15:0] out;
    input clk;
    input signed [7:0] a;
    input signed [7:0] b;

    reg signed [7:0] a_reg;
    reg signed [7:0] b_reg;
    reg signed [15:0] out;
    wire signed [15:0] mult_out;

    assign mult_out = a_reg * b_reg;

    always @ (posedge clk)
    begin
        a_reg <= a;
        b_reg <= b;
        out <= mult_out;
    end
endmodule

```

**Example 12-3: VHDL Unsigned Multiplier with Input and Output Registers (Pipelining = 2)**

```

LIBRARY ieee;
USE ieee.std_logic_1164.all;
USE ieee.numeric_std.all;

ENTITY unsigned_mult IS
    PORT (
        a: IN UNSIGNED (7 DOWNTO 0);
        b: IN UNSIGNED (7 DOWNTO 0);
        clk: IN STD_LOGIC;
        aclr: IN STD_LOGIC;
        result: OUT UNSIGNED (15 DOWNTO 0)
    );
END unsigned_mult;

ARCHITECTURE rtl OF unsigned_mult IS
    SIGNAL a_reg, b_reg: UNSIGNED (7 DOWNTO 0);
BEGIN
    PROCESS (clk, aclr)
    BEGIN
        IF (aclr = '1') THEN
            a_reg <= (OTHERS => '0');
            b_reg <= (OTHERS => '0');
            result <= (OTHERS => '0');
        ELSIF (clk'event AND clk = '1') THEN
            a_reg <= a;
            b_reg <= b;
            result <= a_reg * b_reg;
        END IF;
    END PROCESS;
END rtl;

```

### Example 12-4: VHDL Signed Multiplier

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;
USE ieee.numeric_std.all;

ENTITY signed_mult IS
  PORT (
    a: IN SIGNED (7 DOWNT0 0);
    b: IN SIGNED (7 DOWNT0 0);
    result: OUT SIGNED (15 DOWNT0 0)
  );
END signed_mult;

ARCHITECTURE rtl OF signed_mult IS
BEGIN
  result <= a * b;
END rtl;
```

## Inferring Multiply-Accumulator and Multiply-Adder

Synthesis tools detect multiply-accumulate or multiply-add functions and implement them as Altera IP cores, respectively, or may map them directly to device atoms. The Quartus II software then places these functions in DSP blocks during placement and routing.

**Note:** Synthesis tools infer multiply-accumulator and multiply-adder functions only if the Altera device family has dedicated DSP blocks that support these functions.

A simple multiply-accumulator consists of a multiplier feeding an addition operator. The addition operator feeds a set of registers that then feeds the second input to the addition operator. A simple multiply-adder consists of two to four multipliers feeding one or two levels of addition, subtraction, or addition/subtraction operators. Addition is always the second-level operator, if it is used. In addition to the multiply-accumulator and multiply-adder, the Quartus II Fitter also places input and output registers into the DSP blocks to pack registers and improve performance and area utilization.

Some device families offer additional advanced multiply-add and accumulate functions, such as complex multiplication, input shift register, or larger multiplications.

The Verilog HDL and VHDL code samples infer multiply-accumulators and multiply-adders with input, output, and pipeline registers, as well as an optional asynchronous clear signal. Using the three sets of registers provides the best performance through the function, with a latency of three. You can remove the registers in your design to reduce the latency.

**Note:** To obtain high performance in DSP designs, use register pipelining and avoid unregistered DSP functions.

### Example 12-5: Verilog HDL Unsigned Multiply-Accumulator

```
module unsig_altmult_accum (dataout, dataa, datab, clk, aclr, clken);
  input [7:0] dataa, datab;
  input clk, aclr, clken;
  output reg[16:0] dataout;

  reg [7:0] dataa_reg, datab_reg;
  reg [15:0] multa_reg;
  wire [15:0] multa;
```

```

wire [16:0] adder_out;
assign multa = dataa_reg * datab_reg;
assign adder_out = multa_reg + dataout;

always @ (posedge clk or posedge aclr)
begin
  if (aclr)
  begin
    dataa_reg <= 8'b0;
    datab_reg <= 8'b0;
    multa_reg <= 16'b0;
    dataout <= 17'b0;
  end
  else if (clken)
  begin
    dataa_reg <= dataa;
    datab_reg <= datab;
    multa_reg <= multa;
    dataout <= adder_out;
  end
end
endmodule

```

### Example 12-6: Verilog HDL Signed Multiply-Adder

```

module sig_altmult_add (dataa, datab, datac, datad, clock, aclr, result);
  input signed [15:0] dataa, datab, datac, datad;
  input clock, aclr;
  output reg signed [32:0] result;

  reg signed [15:0] dataa_reg, datab_reg, datac_reg, datad_reg;
  reg signed [31:0] mult0_result, mult1_result;

  always @ (posedge clock or posedge aclr) begin
    if (aclr) begin
      dataa_reg <= 16'b0;
      datab_reg <= 16'b0;
      datac_reg <= 16'b0;
      datad_reg <= 16'b0;
      mult0_result <= 32'b0;
      mult1_result <= 32'b0;
      result <= 33'b0;
    end
    else begin
      dataa_reg <= dataa;
      datab_reg <= datab;
      datac_reg <= datac;
      datad_reg <= datad;
      mult0_result <= dataa_reg * datab_reg;
      mult1_result <= datac_reg * datad_reg;
      result <= mult0_result + mult1_result;
    end
  end
endmodule

```

### Example 12-7: VHDL Signed Multiply-Accumulator

```

LIBRARY ieee;

```

```

USE ieee.std_logic_1164.all;
USE ieee.numeric_std.all;

ENTITY sig_altmult_accum IS
  PORT (
    a: IN SIGNED(7 DOWNT0 0);
    b: IN SIGNED (7 DOWNT0 0);
    clk: IN STD_LOGIC;
    aclr: IN STD_LOGIC;
    accum_out: OUT SIGNED (15 DOWNT0 0)
  );
END sig_altmult_accum;

ARCHITECTURE rtl OF sig_altmult_accum IS
  SIGNAL a_reg, b_reg: SIGNED (7 DOWNT0 0);
  SIGNAL pdt_reg: SIGNED (15 DOWNT0 0);
  SIGNAL adder_out: SIGNED (15 DOWNT0 0);
BEGIN
  PROCESS (clk, aclr)
  BEGIN
    IF (aclr = '1') then
      a_reg <= (others => '0');
      b_reg <= (others => '0');
      pdt_reg <= (others => '0');
      adder_out <= (others => '0');
    ELSIF (clk'event and clk = '1') THEN
      a_reg <= (a);
      b_reg <= (b);
      pdt_reg <= a_reg * b_reg;
      adder_out <= adder_out + pdt_reg;
    END IF;
  END process;
  accum_out <= adder_out;
END rtl;

```

### Example 12-8: VHDL Unsigned Multiply-Adder

```

LIBRARY ieee;
USE ieee.std_logic_1164.all;
USE ieee.numeric_std.all;

ENTITY unsignedmult_add IS
  PORT (
    a: IN UNSIGNED (7 DOWNT0 0);
    b: IN UNSIGNED (7 DOWNT0 0);
    c: IN UNSIGNED (7 DOWNT0 0);
    d: IN UNSIGNED (7 DOWNT0 0);
    clk: IN STD_LOGIC;
    aclr: IN STD_LOGIC;
    result: OUT UNSIGNED (15 DOWNT0 0)
  );
END unsignedmult_add;

ARCHITECTURE rtl OF unsignedmult_add IS
  SIGNAL a_reg, b_reg, c_reg, d_reg: UNSIGNED (7 DOWNT0 0);
  SIGNAL pdt_reg, pdt2_reg: UNSIGNED (15 DOWNT0 0);
  SIGNAL result_reg: UNSIGNED (15 DOWNT0 0);
BEGIN
  PROCESS (clk, aclr)
  BEGIN
    IF (aclr = '1') THEN
      a_reg <= (OTHERS => '0');
      b_reg <= (OTHERS => '0');
    END IF;
  END process;
  result_reg <= c_reg * d_reg + pdt_reg;
END rtl;

```

```
c_reg <= (OTHERS => '0');
d_reg <= (OTHERS => '0');
pdt_reg <= (OTHERS => '0');
pdt2_reg <= (OTHERS => '0');

ELSIF (clk'event AND clk = '1') THEN
  a_reg <= a;
  b_reg <= b;
  c_reg <= c;
  d_reg <= d;
  pdt_reg <= a_reg * b_reg;
  pdt2_reg <= c_reg * d_reg;
  result_reg <= pdt_reg + pdt2_reg;
END IF;
END PROCESS;
result <= result_reg;
END rtl;
```

#### Related Information

- [DSP Design Examples](#)
- [AN639: Inferring Stratix V DSP Blocks for FIR Filtering](#)

## Inferring Memory Functions from HDL Code

The following sections describe how to infer memory functions and target dedicated memory architecture using HDL code.

Altera's dedicated memory architecture offers a number of advanced features that can be easily targeted by instantiating Altera various Altera memory IP Cores in HDL. The following coding recommendations provide portable examples of generic HDL code that infer the appropriate Altera memory IP core. However, if you want to use some of the advanced memory features in Altera devices, consider using the IP core directly so that you can customize the ports and parameters easily. You can also use the Quartus II templates provided in the Quartus II software as a starting point.

Most of these designs can also be found on the Design Examples page on the Altera website.

**Table 12-1: Altera Memory HDL Design Examples**

Language	Full Design Name
VHDL	Single-Port RAM
	Single-Port RAM with Initial Contents
	Simple Dual-Port RAM (single clock)Simple Dual-Port RAM (dual clock)
	True Dual-Port RAM (single clock)
	True Dual-Port RAM (dual clock)
	Mixed-Width RAM
	Mixed-Width True Dual-Port RAM
	Byte-Enabled Simple Dual-Port RAM
	Byte-Enabled True Dual-Port RAM
	Single-Port ROMDual-Port ROM
Verilog HDL	Single-Port RAM
	Single-Port RAM with Initial Contents
	Simple Dual-Port RAM (single clock)
	Simple Dual-Port RAM (dual clock)
	True Dual-Port RAM (single clock)
	True Dual-Port RAM (dual clock)
	Single-Port ROM
	Dual-Port ROM
System Verilog	Mixed-Width Port RAM
	Mixed-Width True Dual-Port RAM
	Mixed-Width True Dual-Port RAM (new data on same port read during write)
	Byte-Enabled Simple Dual Port RAM
	Byte-Enabled True Dual-Port RAM

**Related Information**

- [Instantiating Altera IP Cores in HDL Code](#)
- [Design Examples](#)

**Inferring RAM functions from HDL Code**

To infer RAM functions, synthesis tools detect sets of registers and logic that can be replaced with Altera IP cores for device families that have dedicated RAM blocks, or may map them directly to device memory atoms.

Synthesis tools typically consider all signals and variables that have a multi-dimensional array type and then create a RAM block, if applicable. This is based on the way the signals or variables are assigned or referenced in the HDL source description.

Standard synthesis tools recognize single-port and simple dual-port (one read port and one write port) RAM blocks. Some tools (such as the Quartus II software) also recognize true dual-port (two read ports and two write ports) RAM blocks that map to the memory blocks in certain Altera devices.

Some tools (such as the Quartus II software) also infer memory blocks for array variables and signals that are referenced (read/written) by two indices, to recognize mixed-width and byte-enabled RAMs for certain coding styles.

**Note:** If your design contains a RAM block that your synthesis tool does not recognize and infer, the design might require a large amount of system memory that can potentially cause compilation problems

When you use a formal verification flow, Altera recommends that you create RAM blocks in separate entities or modules that contain only the RAM logic. In certain formal verification flows, for example, when using Quartus II integrated synthesis, the entity or module containing the inferred RAM is put into a black box automatically because formal verification tools do not support RAM blocks. The Quartus II software issues a warning message when this situation occurs. If the entity or module contains any additional logic outside the RAM block, this logic cannot be verified because it also must be treated as a black box for formal verification.

## Use Synchronous Memory Blocks

Use synchronous memory blocks for Altera designs.

Because memory blocks in the newest devices from Altera are synchronous, RAM designs that are targeted towards architectures that contain these dedicated memory blocks must be synchronous to be mapped directly into the device architecture. For these devices, asynchronous memory logic is implemented in regular logic cells.

Synchronous memory offers several advantages over asynchronous memory, including higher frequencies and thus higher memory bandwidth, increased reliability, and less standby power. In many designs with asynchronous memory, the memory interfaces with synchronous logic so that the conversion to synchronous memory design is straightforward. To convert asynchronous memory you can move registers from the data path into the memory block.

Synchronous memories are supported in all Altera device families. A memory block is considered synchronous if it uses one of the following read behaviors:

- Memory read occurs in a Verilog `always` block with a clock signal or a VHDL clocked process. The recommended coding style for synchronous memories is to create your design with a registered read output.
- Memory read occurs outside a clocked block, but there is a synchronous read address (that is, the address used in the read statement is registered). This type of logic is not always inferred as a memory block, or may require external bypass logic, depending on the target device architecture.

**Note:** The synchronous memory structures in Altera devices can differ from the structures in other vendors' devices. For best results, match your design to the target device architecture.

Later sections provide coding recommendations for various memory types. All of these examples are synchronous to ensure that they can be directly mapped into the dedicated memory architecture available in Altera FPGAs.

## Avoid Unsupported Reset and Control Conditions

To ensure that your HDL code can be implemented in the target device architecture, avoid unsupported reset conditions or other control logic that does not exist in the device architecture.

The RAM contents of Altera memory blocks cannot be cleared with a reset signal during device operation. If your HDL code describes a RAM with a reset signal for the RAM contents, the logic is implemented in regular logic cells instead of a memory block. Altera recommends against putting RAM read or write operations in an `always` block or `process` block with a reset signal. If you want to specify memory contents, initialize the memory or write the data to the RAM during device operation.

In addition to reset signals, other control logic can prevent memory logic from being inferred as a memory block. For example, you cannot use a clock enable on the read address registers in some devices because this affects the output latch of the RAM, and therefore the synthesized result in the device RAM architecture would not match the HDL description. You can use the address stall feature as a read address clock enable to avoid this limitation. Check the documentation for your device architecture to ensure that your code matches the hardware available in the device.

### Example 12-9: Verilog RAM with Reset Signal that Clears RAM Contents: Not Supported in Device Architecture

```
module clear_ram
(
    input clock, reset, we,
    input [7:0] data_in,
    input [4:0] address,
    output reg [7:0] data_out
);

    reg [7:0] mem [0:31];
    integer i;

    always @ (posedge clock or posedge reset)
    begin
        if (reset == 1'b1)
            mem[address] <= 0;
        else if (we == 1'b1)
            mem[address] <= data_in;

        data_out <= mem[address];
    end
endmodule
```

### Example 12-10: Verilog RAM with Reset Signal that Affects RAM: Not Supported in Device Architecture

```
module bad_reset
(
    input clock,
    input reset,
    input we,
    input [7:0] data_in,
    input [4:0] address,
    output reg [7:0] data_out,
    input d,
    output reg q
);
```



```

reg [7:0] mem [0:31];
integer i;

always @ (posedge clock or posedge reset)
begin
    if (reset == 1'b1)
        q <= 0;
    else
        begin
            if (we == 1'b1)
                mem[address] <= data_in;

            data_out <= mem[address];
            q <= d;
        end
    end
end
endmodule

```

### Related Information

[Specifying Initial Memory Contents at Power-Up](#) on page 12-25

## Check Read-During-Write Behavior

It is important to check the read-during-write behavior of the memory block described in your HDL design as compared to the behavior in your target device architecture.

Your HDL source code specifies the memory behavior when you read and write from the same memory address in the same clock cycle. The code specifies that the read returns either the old data at the address, or the new data being written to the address. This behavior is referred to as the read-during-write behavior of the memory block. Altera memory blocks have different read-during-write behavior depending on the target device family, memory mode, and block type.

Synthesis tools map an HDL design into the target device architecture, with the goal of maintaining the functionality described in your source code. Therefore, if your source code specifies unsupported read-during-write behavior for the device RAM blocks, the software must implement the logic outside the RAM hardware in regular logic cells.

One common problem occurs when there is a continuous read in the HDL code, as in the following examples. You should avoid using these coding styles:

```

//Verilog HDL concurrent signal assignment
assign q = ram[raddr_reg];

-- VHDL concurrent signal assignment
q <= ram(raddr_reg);

```

When a write operation occurs, this type of HDL implies that the read should immediately reflect the new data at the address, independent of the read clock. However, that is not the behavior of synchronous memory blocks. In the device architecture, the new data is not available until the next edge of the read clock. Therefore, if the synthesis tool mapped the logic directly to a synchronous memory block, the device functionality and gate-level simulation results would not match the HDL description or functional simulation results. If the write clock and read clock are the same, the synthesis tool can infer memory blocks and add extra bypass logic so that the device behavior matches the HDL behavior. If the write and read clocks are different, the synthesis tool cannot reliably add bypass logic, so the logic is implemented in regular logic cells instead of dedicated RAM blocks. The examples in the following sections discuss some of these differences for read-during-write conditions.

In addition, the MLAB feature in certain device logic array blocks (LABs) does not easily support old data or new data behavior for a read-during-write in the dedicated device architecture. Implementing the extra logic to support this behavior significantly reduces timing performance through the memory.

**Note:** For best performance in MLAB memories, your design should not depend on the read data during a write operation.

In many synthesis tools, you can specify that the read-during-write behavior is not important to your design; for example, if you never read from the same address to which you write in the same clock cycle. For Quartus II integrated synthesis, add the synthesis attribute `ramstyle` set to `"no_rw_check"` to allow the software to choose the read-during-write behavior of a RAM, rather than use the behavior specified by your HDL code. In some cases, this attribute prevents the synthesis tool from using extra logic to implement the memory block, or can allow memory inference when it would otherwise be impossible.

Synchronous RAM blocks require a synchronous read, so Quartus II integrated synthesis packs either data output registers or read address registers into the RAM block. When the read address registers are packed into the RAM block, the read address signals connected to the RAM block contain the next value of the read address signals indexing the HDL variable, which impacts which clock cycle the read and the write occur, and changes the read-during-write conditions. Therefore, bypass logic may still be added to the design to preserve the read-during-write behavior, even if the `"no_rw_check"` attribute is set.

#### Related Information

- [Quartus II Integrated Synthesis](#) on page 16-1

## Controlling RAM Inference and Implementation

Synthesis tools usually do not infer small RAM blocks because small RAM blocks typically can be implemented more efficiently using the registers in regular logic.

If you are using Quartus II integrated synthesis, you can direct the software to infer RAM blocks for all sizes with the **Allow Any RAM Size for Recognition** option in the **Advanced Analysis & Synthesis Settings** dialog box.

Some synthesis tools provide options to control the implementation of inferred RAM blocks for Altera devices with synchronous memory blocks. For example, Quartus II integrated synthesis provides the `ramstyle` synthesis attribute to specify the type of memory block or to specify the use of regular logic instead of a dedicated memory block. Quartus II integrated synthesis does not map inferred memory into MLABs unless the HDL code specifies the appropriate `ramstyle` attribute, although the Fitter may map some memories to MLABs.

If you want to control the implementation after the RAM function is inferred during synthesis, you can set the `ram_block_type` parameter of the ALTSYNCRAM IP core. In the Assignment Editor, select **Parameters** in the **Categories** list. You can use the **Node Finder** or drag the appropriate instance from the Project Navigator window to enter the RAM hierarchical instance name. Type `ram_block_type` as the **Parameter Name** and type one of the following memory types supported by your target device family in the **Value** field: "M-RAM", "M512", "M4K", "M9K", "M10K", "M20K", "M144K", or "MLAB".

You can also specify the maximum depth of memory blocks used to infer RAM or ROM in your design. Apply the `max_depth` synthesis attribute to the declaration of a variable that represents a RAM or ROM in your design file. For example:

```
// Limit the depth of the memory blocks implement "ram" to 512
// This forces the software to use two M512 blocks instead of one M4K block to
implement this RAM
(* max_depth = 512 *) reg [7:0] ram[0:1023];
```

**Related Information**

- [Quartus II Integrated Synthesis](#) on page 16-1

**Single-Clock Synchronous RAM with Old Data Read-During-Write Behavior**

The code examples in this section show Verilog HDL and VHDL code that infers simple dual-port, single-clock synchronous RAM. Single-port RAM blocks use a similar coding style.

The read-during-write behavior in these examples is to read the old data at the memory address. Altera recommends that you use the Old Data Read-During-Write coding style for most RAM blocks as long as your design does not require the RAM location's new value when you perform a simultaneous read and write to that RAM location. For best performance in MLAB memories, use the appropriate attribute so that your design does not depend on the read data during a write operation. The simple dual-port RAM code samples map directly into Altera synchronous memory.

Single-port versions of memory blocks (that is, using the same read address and write address signals) can allow better RAM utilization than dual-port memory blocks, depending on the device family.

**Example 12-11: Verilog HDL Single-Clock Simple Dual-Port Synchronous RAM with Old Data Read-During-Write Behavior**

```

module single_clk_ram(
    output reg [7:0] q,
    input [7:0] d,
    input [6:0] write_address, read_address,
    input we, clk
);
    reg [7:0] mem [127:0];

    always @ (posedge clk) begin
        if (we)
            mem[write_address] <= d;
        q <= mem[read_address]; // q doesn't get d in this clock cycle
    end
endmodule

```

**Example 12-12: VHDL Single-Clock Simple Dual-Port Synchronous RAM with Old Data Read-During-Write Behavior**

```

LIBRARY ieee;
USE ieee.std_logic_1164.all;

ENTITY single_clock_ram IS
    PORT (
        clock: IN STD_LOGIC;
        data: IN STD_LOGIC_VECTOR (2 DOWNTO 0);
        write_address: IN INTEGER RANGE 0 to 31;
        read_address: IN INTEGER RANGE 0 to 31;
        we: IN STD_LOGIC;
        q: OUT STD_LOGIC_VECTOR (2 DOWNTO 0)
    );
END single_clock_ram;

ARCHITECTURE rtl OF single_clock_ram IS
    TYPE MEM IS ARRAY(0 TO 31) OF STD_LOGIC_VECTOR(2 DOWNTO 0);
    SIGNAL ram_block: MEM;

```

```
BEGIN
  PROCESS (clock)
  BEGIN
    IF (clock'event AND clock = '1') THEN
      IF (we = '1') THEN
        ram_block(write_address) <= data;
      END IF;
      q <= ram_block(read_address);
      -- VHDL semantics imply that q doesn't get data
      -- in this clock cycle
    END IF;
  END PROCESS;
END rtl;
```

### Related Information

- [Check Read-During-Write Behavior](#) on page 12-12
- [Single-Clock Synchronous RAM with New Data Read-During-Write Behavior](#) on page 12-15

## Single-Clock Synchronous RAM with New Data Read-During-Write Behavior

The examples in this section describe RAM blocks in which a simultaneous read and write to the same location reads the new value that is currently being written to that RAM location.

To implement this behavior in the target device, synthesis software adds bypass logic around the RAM block. This bypass logic increases the area utilization of the design and decreases the performance if the RAM block is part of the design's critical path.

Single-port versions of the Verilog memory block (that is, using the same read address and write address signals) do not require any logic cells to create bypass logic in the Arria, Stratix, and Cyclone series of devices, because the device memory supports new data read-during-write behavior when in single-port mode (same clock, same read address, and same write address).

For Quartus II integrated synthesis, if you do not require the read-through-write capability, add the synthesis attribute `ramstyle="no_rw_check"` to allow the software to choose the read-during-write behavior of a RAM, rather than using the behavior specified by your HDL code. This attribute may prevent generation of extra bypass logic, but it is not always possible to eliminate the requirement for bypass logic.

### Example 12-13: Verilog HDL Single-Clock Simple Dual-Port Synchronous RAM with New Data Read-During-Write Behavior

```
module single_clock_wr_ram(
  output reg [7:0] q,
  input [7:0] d,
  input [6:0] write_address, read_address,
  input we, clk
);
  reg [7:0] mem [127:0];

  always @ (posedge clk) begin
    if (we)
      mem[write_address] = d;
    q = mem[read_address]; // q does get d in this clock cycle
  if
    // we is high
  end
endmodule
```

It is possible to create a single-clock RAM using an assign statement to read the address of `mem` to create the output `q`. By itself, the code describes new data read-during-write behavior. However, if the RAM output feeds a register in another hierarchy, a read-during-write results in the old data. Synthesis tools may not infer a RAM block if the tool cannot determine which behavior is described, such as when the memory feeds a hard hierarchical partition boundary. Avoid this type of coding.

### Example 12-14: Avoid This Coding Style

```
reg [7:0] mem [127:0];
reg [6:0] read_address_reg;

always @ (posedge clk) begin
    if (we)
        mem[write_address] <= d;

    read_address_reg <= read_address;
end

assign q = mem[read_address_reg];
```

The following example uses a concurrent signal assignment to read from the RAM. By itself, this example describes new data read-during-write behavior. However, if the RAM output feeds a register in another hierarchy, a read-during-write results in the old data. Synthesis tools may not infer a RAM block if the tool cannot determine which behavior is described, such as when the memory feeds a hard hierarchical partition boundary.

### Example 12-15: VHDL Single-Clock Simple Dual-Port Synchronous RAM with New Data Read-During-Write Behavior

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;

ENTITY single_clock_rw_ram IS
    PORT (
        clock: IN STD_LOGIC;
        data: IN STD_LOGIC_VECTOR (2 DOWNTO 0);
        write_address: IN INTEGER RANGE 0 to 31;
        read_address: IN INTEGER RANGE 0 to 31;
        we: IN STD_LOGIC;
        q: OUT STD_LOGIC_VECTOR (2 DOWNTO 0)
    );
END single_clock_rw_ram;

ARCHITECTURE rtl OF single_clock_rw_ram IS
    TYPE MEM IS ARRAY(0 TO 31) OF STD_LOGIC_VECTOR(2 DOWNTO 0);
    SIGNAL ram_block: MEM;
    SIGNAL read_address_reg: INTEGER RANGE 0 to 31;
BEGIN
    PROCESS (clock)
    BEGIN
        IF (clock'event AND clock = '1') THEN
            IF (we = '1') THEN
                ram_block(write_address) <= data;
            END IF;
            read_address_reg <= read_address;
        END IF;
    END PROCESS;
END rtl;
```

```
END PROCESS;
  q <= ram_block(read_address_reg);
END rtl;
```

For Quartus II integrated synthesis, if you do not require the read-through-write capability, add the synthesis attribute `ramstyle="no_rw_check"` to allow the software to choose the read-during-write behavior of a RAM, rather than using the behavior specified by your HDL code. This attribute may prevent generation of extra bypass logic but it is not always possible to eliminate the requirement for bypass logic.

#### Related Information

- [Check Read-During-Write Behavior](#) on page 12-12
- [Check Read-During-Write Behavior](#) on page 12-12
- [Single-Clock Synchronous RAM with Old Data Read-During-Write Behavior](#) on page 12-14

### Simple Dual-Port, Dual-Clock Synchronous RAM

In dual clock designs, synthesis tools cannot accurately infer the read-during-write behavior because it depends on the timing of the two clocks within the target device.

Therefore, the read-during-write behavior of the synthesized design is undefined and may differ from your original HDL code. When Quartus II integrated synthesis infers this type of RAM, it issues a warning because of the undefined read-during-write behavior.

#### Example 12-16: Verilog HDL Simple Dual-Port, Dual-Clock Synchronous RAM

```
module dual_clock_ram(
  output reg [7:0] q,
  input [7:0] d,
  input [6:0] write_address, read_address,
  input we, clk1, clk2
);
  reg [6:0] read_address_reg;
  reg [7:0] mem [127:0];

  always @ (posedge clk1)
  begin
    if (we)
      mem[write_address] <= d;
  end

  always @ (posedge clk2) begin
    q <= mem[read_address_reg];
    read_address_reg <= read_address;
  end
endmodule
```

#### Example 12-17: VHDL Simple Dual-Port, Dual-Clock Synchronous RAM

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;
ENTITY dual_clock_ram IS
  PORT (
    clock1, clock2: IN STD_LOGIC;
```

```

        data: IN STD_LOGIC_VECTOR (3 DOWNT0 0);
        write_address: IN INTEGER RANGE 0 to 31;
        read_address: IN INTEGER RANGE 0 to 31;
        we: IN STD_LOGIC;
        q: OUT STD_LOGIC_VECTOR (3 DOWNT0 0)
    );
END dual_clock_ram;
ARCHITECTURE rtl OF dual_clock_ram IS
    TYPE MEM IS ARRAY(0 TO 31) OF STD_LOGIC_VECTOR(3 DOWNT0 0);
    SIGNAL ram_block: MEM;
    SIGNAL read_address_reg : INTEGER RANGE 0 to 31;
BEGIN
    PROCESS (clock1)
    BEGIN
        IF (clock1'event AND clock1 = '1') THEN
            IF (we = '1') THEN
                ram_block(write_address) <= data;
            END IF;
        END IF;
    END PROCESS;
    PROCESS (clock2)
    BEGIN
        IF (clock2'event AND clock2 = '1') THEN
            q <= ram_block(read_address_reg);
            read_address_reg <= read_address;
        END IF;
    END PROCESS;
END rtl;

```

### Related Information

[Check Read-During-Write Behavior](#) on page 12-12

## True Dual-Port Synchronous RAM

The code examples in this section show Verilog HDL and VHDL code that infers true dual-port synchronous RAM. Different synthesis tools may differ in their support for these types of memories.

Altera synchronous memory blocks have two independent address ports, allowing for operations on two unique addresses simultaneously. A read operation and a write operation can share the same port if they share the same address. The Quartus II software infers true dual-port RAMs in Verilog HDL and VHDL with any combination of independent read or write operations in the same clock cycle, with at most two unique port addresses, performing two reads and one write, two writes and one read, or two writes and two reads in one clock cycle with one or two unique addresses.

In the synchronous RAM block architecture, there is no priority between the two ports. Therefore, if you write to the same location on both ports at the same time, the result is indeterminate in the device architecture. You must ensure your HDL code does not imply priority for writes to the memory block, if you want the design to be implemented in a dedicated hardware memory block. For example, if both ports are defined in the same process block, the code is synthesized and simulated sequentially so that there is a priority between the two ports. If your code does imply a priority, the logic cannot be implemented in the device RAM blocks and is implemented in regular logic cells. You must also consider the read-during-write behavior of the RAM block to ensure that it can be mapped directly to the device RAM architecture.

When a read and write operation occurs on the same port for the same address, the read operation may behave as follows:

- **Read new data**—This mode matches the behavior of synchronous memory blocks.
- **Read old data**—This mode is supported only in device families that support M144K and M9K memory blocks.

When a read and write operation occurs on different ports for the same address (also known as mixed port), the read operation may behave as follows:

- **Read new data**—Quartus II integrated synthesis supports this mode by creating bypass logic around the synchronous memory block.
- **Read old data**—Synchronous memory blocks support this behavior.
- **Read don't care**—This behavior is supported on different ports in simple dual-port mode by synchronous memory blocks.

The Verilog HDL single-clock code sample maps directly into Altera synchronous memory. When a read and write operation occurs on the same port for the same address, the new data being written to the memory is read. When a read and write operation occurs on different ports for the same address, the old data in the memory is read. Simultaneous writes to the same location on both ports results in indeterminate behavior.

A dual-clock version of this design describes the same behavior, but the memory in the target device will have undefined mixed port read-during-write behavior because it depends on the relationship between the clocks.

### Example 12-18: Verilog HDL True Dual-Port RAM with Single Clock

```
module true_dual_port_ram_single_clock
(
    input [(DATA_WIDTH-1):0] data_a, data_b,
    input [(ADDR_WIDTH-1):0] addr_a, addr_b,
    input we_a, we_b, clk,
    output reg [(DATA_WIDTH-1):0] q_a, q_b
);

    parameter DATA_WIDTH = 8;
    parameter ADDR_WIDTH = 6;

    // Declare the RAM variable
    reg [DATA_WIDTH-1:0] ram[2**ADDR_WIDTH-1:0];

    always @ (posedge clk)
    begin // Port A
        if (we_a)
            begin
                ram[addr_a] <= data_a;
                q_a <= data_a;
            end
        else
            q_a <= ram[addr_a];
        end
    always @ (posedge clk)
    begin // Port b
        if (we_b)
            begin
                ram[addr_b] <= data_b;
                q_b <= data_b;
            end
        else
            q_b <= ram[addr_b];
        end
    end

endmodule
```

If you use the following Verilog HDL read statements instead of the `if-else` statements, the HDL code specifies that the read results in old data when a read operation and write operation occurs



at the same time for the same address on the same port or mixed ports. This mode is supported only in device families that support M144, M9k, and MLAB memory blocks.

### Example 12-19: VHDL Read Statement Example

```

always @ (posedge clk)
begin // Port A
  if (we_a)
    ram[addr_a] <= data_a;

  q_a <= ram[addr_a];
end

always @ (posedge clk)
begin // Port B
  if (we_b)
    ram[addr_b] <= data_b;

  q_b <= ram[addr_b];
end

```

The VHDL single-clock code sample maps directly into Altera synchronous memory. When a read and write operation occurs on the same port for the same address, the new data being written to the memory is read. When a read and write operation occurs on different ports for the same address, the old data in the memory is read. Simultaneous write operations to the same location on both ports results in indeterminate behavior.

A dual-clock version of this design describes the same behavior, but the memory in the target device will have undefined mixed port read-during-write behavior because it depends on the relationship between the clocks.

### Example 12-20: VHDL True Dual-Port RAM with Single Clock (part 1)

```

library ieee;
use ieee.std_logic_1164.all;

entity true_dual_port_ram_single_clock is
  generic (
    DATA_WIDTH : natural := 8;
    ADDR_WIDTH  : natural := 6;
  );
  port (
    clk      : in std_logic;
    addr_a   : in natural range 0 to 2**ADDR_WIDTH - 1;
    addr_b   : in natural range 0 to 2**ADDR_WIDTH - 1;
    data_a   : in std_logic_vector((DATA_WIDTH-1) downto 0);
    data_b   : in std_logic_vector((DATA_WIDTH-1) downto 0);
    we_a     : in std_logic := '1';
    we_b     : in std_logic := '1';
    q_a      : out std_logic_vector((DATA_WIDTH -1) downto 0);
    q_b      : out std_logic_vector((DATA_WIDTH -1) downto 0);
  );
end true_dual_port_ram_single_clock;

architecture rtl of true_dual_port_ram_single_clock is
  -- Build a 2-D array type for the RAM
  subtype word_t is std_logic_vector((DATA_WIDTH-1) downto 0);

```

```
type memory_t is array((2**ADDR_WIDTH - 1) downto 0) of word_t;
-- Declare the RAM signal.
shared variable ram : memory_t;
```

### Example 12-21: VHDL True Dual-Port RAM with Single Clock (part 2)

```
begin
process(clk)
begin
if(rising_edge(clk)) then -- Port A
  if(we_a = '1') then
    ram(addr_a) <= data_a;

    -- Read-during-write on the same port returns NEW data
    q_a <= data_a;
  else
    -- Read-during-write on the mixed port returns OLD data
    q_a <= ram(addr_a);
  end if;
end if;
end process;

process(clk)
begin
if(rising_edge(clk)) then -- Port B
  if(we_b = '1') then
    ram(addr_b) := data_b;
    -- Read-during-write on the same port returns NEW data
    q_b <= data_b;
  else
    -- Read-during-write on the mixed port returns OLD data
    q_b <= ram(addr_b);
  end if;
end if;
end process;

end rtl;
```

#### Related Information

[Check Read-During-Write Behavior](#) on page 12-12

### Mixed-Width Dual-Port RAM

The RAM code examples show SystemVerilog and VHDL code that infers RAM with data ports with different widths.

This type of logic is not supported in Verilog-1995 or Verilog-2001 because of the requirement for a multi-dimensional array to model the different read width, write width, or both. Different synthesis tools may differ in their support for these memories. This section describes the inference rules for Quartus II integrated synthesis.

The first dimension of the multi-dimensional packed array represents the ratio of the wider port to the narrower port, and the second dimension represents the narrower port width. The read and write port widths must specify a read or write ratio supported by the memory blocks in the target device, or the synthesis tool does not infer a RAM.

Refer to the Quartus II templates for parameterized examples that you can use for supported combinations of read and write widths, and true dual port RAM examples with two read ports and two write ports for mixed-width writes and reads.

### Example 12-22: SystemVerilog Mixed-Width RAM with Read Width Smaller than Write Width

```

module mixed_width_ram    // 256x32 write and 1024x8 read
(
    input [7:0] waddr,
    input [31:0] wdata,
    input we, clk,
    input [9:0] raddr,
    output [7:0] q
);
    logic [3:0][7:0] ram[0:255];
    always_ff@(posedge clk)
        begin
            if(we) ram[waddr] <= wdata;
            q <= ram[raddr / 4][raddr % 4];
        end
endmodule : mixed_width_ram

```

### Example 12-23: SystemVerilog Mixed-Width RAM with Read Width Larger than Write Width

```

module mixed_width_ram    // 1024x8 write and 256x32 read
(
    input [9:0] waddr,
    input [31:0] wdata,
    input we, clk,
    input [7:0] raddr,
    output [9:0] q
);
    logic [3:0][7:0] ram[0:255];
    always_ff@(posedge clk)
        begin
            if(we) ram[waddr / 4][waddr % 4] <= wdata;
            q <= ram[raddr];
        end
endmodule : mixed_width_ram

```

### Example 12-24: VHDL Mixed-Width RAM with Read Width Smaller than Write Width

```

library ieee;
use ieee.std_logic_1164.all;

package ram_types is
    type word_t is array (0 to 3) of std_logic_vector(7 downto 0);
    type ram_t is array (0 to 255) of word_t;
end ram_types;

library ieee;
use ieee.std_logic_1164.all;
library work;
use work.ram_types.all;

```

```

entity mixed_width_ram is
  port (
    we, clk : in  std_logic;
    waddr   : in  integer range 0 to 255;
    wdata   : in  word_t;
    raddr   : in  integer range 0 to 1023;
    q       : out std_logic_vector(7 downto 0));
end mixed_width_ram;

architecture rtl of mixed_width_ram is
  signal ram : ram_t;
begin -- rtl
  process(clk, we)
  begin
    if(rising_edge(clk)) then
      if(we = '1') then
        ram(waddr) <= wdata;
      end if;
      q <= ram(raddr / 4 )(raddr mod 4);
    end if;
  end process;
end rtl;

```

### Example 12-25: VHDL Mixed-Width RAM with Read Width Larger than Write Width

```

library ieee;
use ieee.std_logic_1164.all;

package ram_types is
  type word_t is array (0 to 3) of std_logic_vector(7 downto 0);
  type ram_t is array (0 to 255) of word_t;
end ram_types;

library ieee;
use ieee.std_logic_1164.all;
library work;
use work.ram_types.all;

entity mixed_width_ram is
  port (
    we, clk : in  std_logic;
    waddr   : in  integer range 0 to 1023;
    wdata   : in  std_logic_vector(7 downto 0);
    raddr   : in  integer range 0 to 255;
    q       : out word_t);
end mixed_width_ram;

architecture rtl of mixed_width_ram is
  signal ram : ram_t;
begin -- rtl
  process(clk, we)
  begin
    if(rising_edge(clk)) then
      if(we = '1') then
        ram(waddr / 4 )(waddr mod 4) <= wdata;
      end if;
      q <= ram(raddr);
    end if;
  end process;
end rtl;

```

## RAM with Byte-Enable Signals

The RAM code examples show SystemVerilog and VHDL code that infers RAM with controls for writing single bytes into the memory word, or byte-enable signals.

Byte enables are modeled by creating write expressions with two indices and writing part of a RAM "word." With these implementations, you can also write more than one byte at once by enabling the appropriate byte enables.

This type of logic is not supported in Verilog-1995 or Verilog-2001 because of the requirement for a multidimensional array. Different synthesis tools may differ in their support for these memories. This section describes the inference rules for Quartus II integrated synthesis.

Refer to the Quartus II templates for parameterized examples that you can use for different address widths, and true dual port RAM examples with two read ports and two write ports.

### Example 12-26: SystemVerilog Simple Dual-Port Synchronous RAM with Byte Enable

```

module byte_enabled_simple_dual_port_ram
(
    input we, clk,
    input [5:0] waddr, raddr, // address width = 6
    input [3:0] be,          // 4 bytes per word
    input [31:0] wdata,      // byte width = 8, 4 bytes per word
    output reg [31:0] q      // byte width = 8, 4 bytes per word
);
    // use a multi-dimensional packed array
    //to model individual bytes within the word
    logic [3:0][7:0] ram[0:63]; // # words = 1 << address width

    always_ff@(posedge clk)
    begin
        if(we) begin
            if(be[0]) ram[waddr][0] <= wdata[7:0];
            if(be[1]) ram[waddr][1] <= wdata[15:8];
            if(be[2]) ram[waddr][2] <= wdata[23:16];
            if(be[3]) ram[waddr][3] <= wdata[31:24];
        end
        q <= ram[raddr];
    end
endmodule

```

### Example 12-27: VHDL Simple Dual-Port Synchronous RAM with Byte Enable

```

library ieee;
use ieee.std_logic_1164.all;
library work;

entity byte_enabled_simple_dual_port_ram is
port (
    we, clk : in  std_logic;
    waddr, raddr : in  integer range 0 to 63 ; -- address width = 6
    be      : in  std_logic_vector (3 downto 0); -- 4 bytes per word
    wdata   : in  std_logic_vector(31 downto 0); -- byte width = 8
    q       : out std_logic_vector(31 downto 0) ); -- byte width = 8
end byte_enabled_simple_dual_port_ram;

architecture rtl of byte_enabled_simple_dual_port_ram is
-- build up 2D array to hold the memory

```

```
type word_t is array (0 to 3) of std_logic_vector(7 downto 0);
type ram_t is array (0 to 63) of word_t;

signal ram : ram_t;
signal q_local : word_t;

begin -- Re-organize the read data from the RAM to match the output
    unpack: for i in 0 to 3 generate
        q(8*(i+1) - 1 downto 8*i) <= q_local(i);
    end generate unpack;

process(clk)
begin
    if(rising_edge(clk)) then
        if(we = '1') then
            if(be(0) = '1') then
                ram(waddr)(0) <= wdata(7 downto 0);
            end if;
            if be(1) = '1' then
                ram(waddr)(1) <= wdata(15 downto 8);
            end if;
            if be(2) = '1' then
                ram(waddr)(2) <= wdata(23 downto 16);
            end if;
            if be(3) = '1' then
                ram(waddr)(3) <= wdata(31 downto 24);
            end if;
        end if;
        q_local <= ram(raddr);
    end if;
end process;
end rtl;
```

## Specifying Initial Memory Contents at Power-Up

Your synthesis tool may offer various ways to specify the initial contents of an inferred memory.

There are slight power-up and initialization differences between dedicated RAM blocks and the MLAB memory due to the continuous read of the MLAB. Altera dedicated RAM block outputs always power-up to zero and are set to the initial value on the first read. For example, if address 0 is pre-initialized to FF, the RAM block powers up with the output at 0. A subsequent read after power-up from address 0 outputs the pre-initialized value of FF. Therefore, if a RAM is powered up and an enable (read enable or clock enable) is held low, the power-up output of 0 is maintained until the first valid read cycle. The MLAB is implemented using registers that power-up to 0, but are initialized to their initial value immediately at power-up or reset. Therefore, the initial value is seen, regardless of the enable status. The Quartus II software maps inferred memory to MLABs when the HDL code specifies an appropriate `ramstyle` attribute.

In Verilog HDL, you can use an initial block to initialize the contents of an inferred memory. Quartus II integrated synthesis automatically converts the initial block into a `.mif` file for the inferred RAM.

### Example 12-28: Verilog HDL RAM with Initialized Contents

```
module ram_with_init(
    output reg [7:0] q,
    input [7:0] d,
    input [4:0] write_address, read_address,
    input we, clk
);
    reg [7:0] mem [0:31];
```

```

integer i;

initial begin
    for (i = 0; i < 32; i = i + 1)
        mem[i] = i[7:0];
end

always @ (posedge clk) begin
    if (we)
        mem[write_address] <= d;
    q <= mem[read_address];
end
endmodule

```

Quartus II integrated synthesis and other synthesis tools also support the `$readmemb` and `$readmemh` commands so that RAM initialization and ROM initialization work identically in synthesis and simulation.

### Example 12-29: Verilog HDL RAM Initialized with the `readmemb` Command

```

reg [7:0] ram[0:15];
initial
begin
    $readmemb("ram.txt", ram);
end

```

In VHDL, you can initialize the contents of an inferred memory by specifying a default value for the corresponding signal. Quartus II integrated synthesis automatically converts the default value into a `.mif` file for the inferred RAM.

### Example 12-30: VHDL RAM with Initialized Contents

```

LIBRARY ieee;
USE ieee.std_logic_1164.all;
use ieee.numeric_std.all;

ENTITY ram_with_init IS
    PORT(
        clock: IN STD_LOGIC;
        data: IN UNSIGNED (7 DOWNTO 0);
        write_address: IN integer RANGE 0 to 31;
        read_address: IN integer RANGE 0 to 31;
        we: IN std_logic;
        q: OUT UNSIGNED (7 DOWNTO 0));
END;

ARCHITECTURE rtl OF ram_with_init IS

    TYPE MEM IS ARRAY(31 DOWNTO 0) OF unsigned(7 DOWNTO 0);
    FUNCTION initialize_ram
        return MEM is
        variable result : MEM;
    BEGIN
        FOR i IN 31 DOWNTO 0 LOOP
            result(i) := to_unsigned(natural(i), natural'(8));
        END LOOP;
        RETURN result;
    END;

```

```
END initialize_ram;

SIGNAL ram_block : MEM := initialize_ram;
BEGIN
PROCESS (clock)
BEGIN
IF (clock'event AND clock = '1') THEN
IF (we = '1') THEN
ram_block(write_address) <= data;
END IF;
q <= ram_block(read_address);
END IF;
END PROCESS;
END rtl;
```

### Related Information

- [Quartus II Integrated Synthesis](#) on page 16-1

## Inferring ROM Functions from HDL Code

ROMs are inferred when a `CASE` statement exists in which a value is set to a constant for every choice in the `case` statement.

Because small ROMs typically achieve the best performance when they are implemented using the registers in regular logic, each ROM function must meet a minimum size requirement to be inferred and placed into memory.

**Note:** If you use Quartus II integrated synthesis, you can direct the software to infer ROM blocks for all sizes with the **Allow Any ROM Size for Recognition** option in the **Advanced Analysis & Synthesis Settings** dialog box.

Some synthesis tools provide options to control the implementation of inferred ROM blocks for Altera devices with synchronous memory blocks. For example, Quartus II integrated synthesis provides the `romstyle` synthesis attribute to specify the type of memory block or to specify the use of regular logic instead of a dedicated memory block.

**Note:** Because formal verification tools do not support ROM IP cores, Quartus II integrated synthesis does not infer ROM IP cores when a formal verification tool is selected. When you are using a formal verification flow, Altera recommends that you instantiate ROM IP core blocks in separate entities or modules that contain only the ROM logic, because you may need to treat the entity or module as a black box during formal verification. Depending on the device family's dedicated RAM architecture, the ROM logic may have to be synchronous; refer to the device family handbook for details.

For device architectures with synchronous RAM blocks, such as the Arria series, Cyclone series, or Stratix series devices and newer device families, either the address or the output must be registered for synthesis software to infer a ROM block. When your design uses output registers, the synthesis software implements registers from the input registers of the RAM block without affecting the functionality of the ROM. If you register the address, the power-up state of the inferred ROM can be different from the HDL design. In this scenario, the synthesis software issues a warning. The Quartus II Help explains the condition under which the functionality changes when you use Quartus II integrated synthesis.

The following ROM examples map directly to the Altera memory architecture.



**Example 12-31: Verilog HDL Synchronous ROM**

```

module sync_rom (clock, address, data_out);
    input clock;
    input [7:0] address;
    output [5:0] data_out;

    reg [5:0] data_out;

    always @ (posedge clock)
    begin
        case (address)
            8'b00000000: data_out = 6'b101111;
            8'b00000001: data_out = 6'b110110;
            ...
            8'b11111110: data_out = 6'b000001;
            8'b11111111: data_out = 6'b101010;
        endcase
    end
endmodule

```

**Example 12-32: VHDL Synchronous ROM**

```

LIBRARY ieee;
USE ieee.std_logic_1164.all;

ENTITY sync_rom IS
    PORT (
        clock: IN STD_LOGIC;
        address: IN STD_LOGIC_VECTOR(7 downto 0);
        data_out: OUT STD_LOGIC_VECTOR(5 downto 0)
    );
END sync_rom;

ARCHITECTURE rtl OF sync_rom IS
BEGIN
    PROCESS (clock)
    BEGIN
        IF rising_edge (clock) THEN
            CASE address IS
                WHEN "00000000" => data_out <= "101111";
                WHEN "00000001" => data_out <= "110110";
                ...
                WHEN "11111110" => data_out <= "000001";
                WHEN "11111111" => data_out <= "101010";
                WHEN OTHERS      => data_out <= "101111";
            END CASE;
        END IF;
    END PROCESS;
END rtl;

```

**Example 12-33: Verilog HDL Dual-Port Synchronous ROM Using readmemb**

```

module dual_port_rom (
    input [(addr_width-1):0] addr_a, addr_b,

```

```

input clk,
output reg [(data_width-1):0] q_a, q_b
);
parameter data_width = 8;
parameter addr_width = 8;

reg [data_width-1:0] rom[2**addr_width-1:0];

initial // Read the memory contents in the file
//dual_port_rom_init.txt.
begin
  $readmemb("dual_port_rom_init.txt", rom);
end

always @ (posedge clk)
begin
  q_a <= rom[addr_a];
  q_b <= rom[addr_b];
end
endmodule

```

### Example 12-34: VHDL Dual-Port Synchronous ROM Using Initialization Function

```

library ieee;
use ieee.std_logic_1164.all;
use ieee.numeric_std.all;

entity dual_port_rom is
  generic (
    DATA_WIDTH : natural := 8;
    ADDR_WIDTH   : natural := 8
  );
  port (
    clk      : in std_logic;
    addr_a   : in natural range 0 to 2**ADDR_WIDTH - 1;
    addr_b   : in natural range 0 to 2**ADDR_WIDTH - 1;
    q_a      : out std_logic_vector((DATA_WIDTH - 1) downto 0);
    q_b      : out std_logic_vector((DATA_WIDTH - 1) downto 0)
  );
end entity;

architecture rtl of dual_port_rom is
  -- Build a 2-D array type for the ROM
  subtype word_t is std_logic_vector((DATA_WIDTH-1) downto 0);
  type memory_t is array(addr_a'high downto 0) of word_t;

  function init_rom
    return memory_t is
    variable tmp : memory_t := (others => (others => '0'));
  begin
    for addr_pos in 0 to 2**ADDR_WIDTH - 1 loop
      -- Initialize each address with the address itself
      tmp(addr_pos) := std_logic_vector(to_unsigned(addr_pos,
DATA_WIDTH));
    end loop;
    return tmp;
  end init_rom;

  -- Declare the ROM signal and specify a default initialization value.
  signal rom : memory_t := init_rom;
begin
  process(clk)
  begin

```

```

if (rising_edge(clk)) then
    q_a <= rom(addr_a);
    q_b <= rom(addr_b);
end if;
end process;
end rtl;

```

### Related Information

- [Quartus II Integrated Synthesis](#) on page 16-1

## Inferring Shift Registers in HDL Code

To infer shift registers, synthesis tools detect a group of shift registers of the same length and convert them to an Altera shift register IP core.

To be detected, all the shift registers must have the following characteristics:

- Use the same clock and clock enable
- Do not have any other secondary signals
- Have equally spaced taps that are at least three registers apart

When you use a formal verification flow, Altera recommends that you create shift register blocks in separate entities or modules containing only the shift register logic, because you might have to treat the entity or module as a black box during formal verification.

**Note:** Because formal verification tools do not support shift register IP cores, Quartus II integrated synthesis does not infer the Altera shift register IP core when a formal verification tool is selected. You can select EDA tools for use with your design on the **EDA Tool Settings** page of the **Settings** dialog box in the Quartus II software.

Synthesis recognizes shift registers only for device families that have dedicated RAM blocks, and the software uses certain guidelines to determine the best implementation.

Quartus II integrated synthesis uses the following guidelines which are common in other EDA tools. The Quartus II software determines whether to infer the Altera shift register IP core based on the width of the registered bus ( $w$ ), the length between each tap ( $L$ ), and the number of taps ( $N$ ). If the **Auto Shift Register Recognition** setting is set to **Auto**, Quartus II integrated synthesis uses the **Optimization Technique** setting, logic and RAM utilization information about the design, and timing information from **Timing-Driven Synthesis** to determine which shift registers are implemented in RAM blocks for logic.

- If the registered bus width is one ( $w = 1$ ), the software infers shift register IP if the number of taps times the length between each tap is greater than or equal to 64 ( $N \times L > 64$ ).
- If the registered bus width is greater than one ( $w > 1$ ), the software infers Altera shift register IP core if the registered bus width times the number of taps times the length between each tap is greater than or equal to 32 ( $w \times N \times L > 32$ ).

If the length between each tap ( $L$ ) is not a power of two, the software uses more logic to decode the read and write counters. This situation occurs because for different sizes of shift registers, external decode logic that uses logic elements (LEs) or ALMs is required to implement the function. This decode logic eliminates the performance and utilization advantages of implementing shift registers in memory.

The registers that the software maps to the Altera shift register IP core and places in RAM are not available in a Verilog HDL or VHDL output file for simulation tools because their node names do not exist after synthesis.

**Note:** If your design uses a shift enable signal to infer a shift register, the shift register will not be implemented into MLAB memory, but can use only dedicated RAM blocks.

## Simple Shift Register

The code samples show a simple, single-bit wide, 64-bit long shift register.

The synthesis software implements the register ( $w = 1$  and  $m = 64$ ) in an ALTSHIFT\_TAPS IP core for supported devices and maps it to RAM in supported devices, which may be placed in dedicated RAM blocks or MLAB memory. If the length of the register is less than 64 bits, the software implements the shift register in logic.

### Example 12-35: Verilog HDL Single-Bit Wide, 64-Bit Long Shift Register

```
module shift_1x64 (clk, shift, sr_in, sr_out);
    input clk, shift;
    input sr_in;
    output sr_out;

    reg [63:0] sr;

    always @ (posedge clk)
    begin
        if (shift == 1'b1)
            begin
                sr[63:1] <= sr[62:0];
                sr[0] <= sr_in;
            end
        end
    assign sr_out = sr[63];
endmodule
```

### Example 12-36: VHDL Single-Bit Wide, 64-Bit Long Shift Register

```
LIBRARY IEEE;
USE IEEE.STD_LOGIC_1164.all;
ENTITY shift_1x64 IS
    PORT (
        clk: IN STD_LOGIC;
        shift: IN STD_LOGIC;
        sr_in: IN STD_LOGIC;
        sr_out: OUT STD_LOGIC
    );
END shift_1x64;

ARCHITECTURE arch OF shift_1x64 IS
    TYPE sr_length IS ARRAY (63 DOWNT0 0) OF STD_LOGIC;
    SIGNAL sr: sr_length;
BEGIN
    PROCESS (clk)
    BEGIN
        IF (clk'EVENT and clk = '1') THEN
            IF (shift = '1') THEN
                sr(63 DOWNT0 1) <= sr(62 DOWNT0 0);
                sr(0) <= sr_in;
            END IF;
        END IF;
    END PROCESS;
    sr_out <= sr(63);
END arch;
```

## Shift Register with Evenly Spaced Taps

The following examples show a Verilog HDL and VHDL 8-bit wide, 64-bit long shift register ( $w > 1$  and  $m = 64$ ) with evenly spaced taps at 15, 31, and 47.

The synthesis software implements this function in a single ALTSHIFT\_TAPS IP core and maps it to RAM in supported devices, which is allowed placement in dedicated RAM blocks or MLAB memory.

### Example 12-37: Verilog HDL 8-Bit Wide, 64-Bit Long Shift Register with Evenly Spaced Taps

```

module shift_8x64_taps (clk, shift, sr_in, sr_out, sr_tap_one, sr_tap_two,
sr_tap_three );
    input clk, shift;
    input [7:0] sr_in;
    output [7:0] sr_tap_one, sr_tap_two, sr_tap_three, sr_out;

    reg [7:0] sr [63:0];
    integer n;

    always @ (posedge clk)
    begin
        if (shift == 1'b1)
        begin
            for (n = 63; n>0; n = n-1)
            begin
                sr[n] <= sr[n-1];
            end
            sr[0] <= sr_in;
        end

        end
        assign sr_tap_one = sr[15];
        assign sr_tap_two = sr[31];
        assign sr_tap_three = sr[47];
        assign sr_out = sr[63];
    endmodule

```

### Example 12-38: VHDL 8-Bit Wide, 64-Bit Long Shift Register with Evenly Spaced Taps

```

LIBRARY IEEE;
USE IEEE.STD_LOGIC_1164.all;
ENTITY shift_8x64_taps IS
    PORT (
        clk: IN STD_LOGIC;
        shift: IN STD_LOGIC;
        sr_in: IN STD_LOGIC_VECTOR(7 DOWNTO 0);
        sr_tap_one: OUT STD_LOGIC_VECTOR(7 DOWNTO 0);
        sr_tap_two : OUT STD_LOGIC_VECTOR(7 DOWNTO 0);
        sr_tap_three: OUT STD_LOGIC_VECTOR(7 DOWNTO 0);
        sr_out: OUT STD_LOGIC_VECTOR(7 DOWNTO 0)
    );
END shift_8x64_taps;

ARCHITECTURE arch OF shift_8x64_taps IS
    SUBTYPE sr_width IS STD_LOGIC_VECTOR(7 DOWNTO 0);
    TYPE sr_length IS ARRAY (63 DOWNTO 0) OF sr_width;
    SIGNAL sr: sr_length;
BEGIN
    PROCESS (clk)
    BEGIN
        IF (clk'EVENT and clk = '1') THEN

```

```
        IF (shift = '1') THEN
            sr(63 DOWNT0 1) <= sr(62 DOWNT0 0);
            sr(0) <= sr_in;
        END IF;
    END IF;
END PROCESS;
sr_tap_one <= sr(15);
sr_tap_two <= sr(31);
sr_tap_three <= sr(47);
sr_out <= sr(63);
END arch;
```

## Register and Latch Coding Guidelines

This section provides device-specific coding recommendations for Altera registers and latches.

Understanding the architecture of the target Altera device helps ensure that your code produces the expected results and achieves the optimal quality of results.

### Register Power-Up Values in Altera Devices

Registers in the device core always power up to a low (0) logic level on all Altera devices.

If your design specifies a power-up level other than 0, synthesis tools can implement logic that causes registers to behave as if they were powering up to a high (1) logic level.

If your design uses a preset signal on a device that does not support presets in the register architecture, your synthesis tool may convert the preset signal to a clear signal, which requires synthesis to perform an optimization referred to as NOT gate push-back. NOT gate push-back adds an inverter to the input and the output of the register so that the reset and power-up conditions will appear to be high, and the device operates as expected. In this case, your synthesis tool may issue a message informing you about the power-up condition. The register itself powers up low, but the register output is inverted, so the signal that arrives at all destinations is high.

Due to these effects, if you specify a non-zero reset value, you may cause your synthesis tool to use the asynchronous clear (`aclr`) signals available on the registers to implement the high bits with NOT gate push-back. In that case, the registers look as though they power up to the specified reset value.

When an asynchronous load (`aload`) signal is available in the device registers, your synthesis tools can implement a reset of 1 or 0 value by using an asynchronous load of 1 or 0. When the synthesis tool uses a load signal, it is not performing NOT gate push-back, so the registers power up to a 0 logic level.

For additional details, refer to the appropriate device family handbook or the appropriate handbook on the Altera website.

Designers typically use an explicit reset signal for the design, which forces all registers into their appropriate values after reset. Altera recommends this practice to reset the device after power-up to restore the proper state.

You can make your design more stable and avoid potential glitches by synchronizing external or combinational logic of the device architecture before you drive the asynchronous control ports of registers.

#### Related Information

- [Design Recommendations for Altera Devices and the Quartus II Design Assistant](#) on page 11-1

## Specifying a Power-Up Value

If you want to force a particular power-up condition for your design, you can use the synthesis options available in your synthesis tool.

With Quartus II integrated synthesis, you can apply the **Power-Up Level** logic option. You can also apply the option with an `altera_attribute` assignment in your source code. Using this option forces synthesis to perform NOT gate push-back because synthesis tools cannot actually change the power-up states of core registers.

You can apply the Quartus II integrated synthesis **Power-Up Level** logic option to a specific register or to a design entity, module, or subdesign. If you do so, every register in that block receives the value. Registers power up to 0 by default; therefore, you can use this assignment to force all registers to power up to 1 using NOT gate push-back.

**Note:** Setting the **Power-Up Level** to a logic level of high for a large design entity could degrade the quality of results due to the number of inverters that are required. In some situations, issues are caused by enable signal inference or secondary control logic inference. It may also be more difficult to migrate such a design to an ASIC.

**Note:** You can simulate the power-up behavior in a functional simulation if you use initialization.

Some synthesis tools can also read the default or initial values for registered signals and implement this behavior in the device. For example, Quartus II integrated synthesis converts default values for registered signals into **Power-Up Level** settings. When the Quartus II software reads the default values, the synthesized behavior matches the power-up state of the HDL code during a functional simulation.

### Example 12-39: Verilog Register with High Power-Up Value

```
reg q = 1'b1; //q has a default value of '1'

always @ (posedge clk)
begin
    q <= d;
end
```

### Example 12-40: VHDL Register with High Power-Up Level

```
SIGNAL q : STD_LOGIC := '1'; -- q has a default value of '1'

PROCESS (clk, reset)
BEGIN
    IF (rising_edge(clk)) THEN
        q <= d;
    END IF;
END PROCESS;
```

There may also be undeclared default power-up conditions based on signal type. If you declare a VHDL register signal as an integer, Quartus II synthesis attempts to use the left end of the integer range as the power-up value. For the default signed integer type, the default power-up value is the highest magnitude negative integer (100...001). For an unsigned integer type, the default power-up value is 0.

**Note:** If the target device architecture does not support two asynchronous control signals, such as `aclr` and `aload`, you cannot set a different power-up state and

reset state. If the NOT gate push-back algorithm creates logic to set a register to 1, that register will power-up high. If you set a different power-up condition through a synthesis assignment or initial value, the power-up level is ignored during synthesis.

#### Related Information

- [Quartus II Integrated Synthesis](#) on page 16-1

## Secondary Register Control Signals Such as Clear and Clock Enable

The registers in Altera FPGAs provide a number of secondary control signals (such as clear and enable signals) that you can use to implement control logic for each register without using extra logic cells.

The registers in Altera FPGAs provide a number of secondary control signals (such as clear and enable signals) that you can use to implement control logic for each register without using extra logic cells. Device families vary in their support for secondary signals, so consult the device family data sheet to verify which signals are available in your target device.

To make the most efficient use of the signals in the device, your HDL code should match the device architecture as closely as possible. The control signals have a certain priority due to the nature of the architecture, so your HDL code should follow that priority where possible.

Your synthesis tool can emulate any control signals using regular logic, so achieving functionally correct results is always possible. However, if your design requirements are flexible in terms of which control signals are used and in what priority, match your design to the target device architecture to achieve the most efficient results. If the priority of the signals in your design is not the same as that of the target architecture, extra logic may be required to implement the control signals. This extra logic uses additional device resources and can cause additional delays for the control signals.

In addition, there are certain cases where using logic other than the dedicated control logic in the device architecture can have a larger impact. For example, the clock enable signal has priority over the synchronous reset or clear signal in the device architecture. The clock enable turns off the clock line in the LAB, and the clear signal is synchronous. Therefore, in the device architecture, the synchronous clear takes effect only when a clock edge occurs.

If you code a register with a synchronous clear signal that has priority over the clock enable signal, the software must emulate the clock enable functionality using data inputs to the registers. Because the signal does not use the clock enable port of a register, you cannot apply a Clock Enable Multicycle constraint. In this case, following the priority of signals available in the device is clearly the best choice for the priority of these control signals, and using a different priority causes unexpected results with an assignment to the clock enable signal.

**Note:** The priority order for secondary control signals in Altera devices differs from the order for other vendors' devices. If your design requirements are flexible regarding priority, verify that the secondary control signals meet design performance requirements when migrating designs between FPGA vendors and try to match your target device architecture to achieve the best results.



The signal order is the same for all Altera device families, although, as noted previously, not all device families provide every signal. The following priority order is observed:

1. Asynchronous Clear, `aclr`—highest priority
2. Asynchronous Load, `aload`
3. Enable, `ena`
4. Synchronous Clear, `sclr`
5. Synchronous Load, `sload`
6. Data In, `data`—lowest priority

The following examples provide Verilog HDL and VHDL code that creates a register with the `aclr`, `aload`, and `ena` control signals.

**Note:** The Verilog HDL example does not have `adata` on the sensitivity list, but the VHDL example does. This is a limitation of the Verilog HDL language—there is no way to describe an asynchronous load signal (in which `q` toggles if `adata` toggles while `aload` is high). All synthesis tools should infer an `aload` signal from this construct despite this limitation. When they perform such inference, you may see information or warning messages from the synthesis tool.

### Example 12-41: Verilog HDL D-Type Flipflop (Register) with `ena`, `aclr`, and `aload` Control Signals

```
module dff_control(clk, aclr, aload, ena, data, adata, q);
    input clk, aclr, aload, ena, data, adata;
    output q;

    reg q;

    always @ (posedge clk or posedge aclr or posedge aload)
    begin
        if (aclr)
            q <= 1'b0;
        else if (aload)
            q <= adata;
        else if (ena)
            q <= data;
    end
endmodule
```

### Example 12-42: VHDL D-Type Flipflop (Register) with `ena`, `aclr`, and `aload` Control Signals (part 1)

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;

ENTITY dff_control IS
    PORT (
        clk: IN STD_LOGIC;
        aclr: IN STD_LOGIC;
        aload: IN STD_LOGIC;
        adata: IN STD_LOGIC;
        ena: IN STD_LOGIC;
        data: IN STD_LOGIC;
        q: OUT STD_LOGIC
    );
END ENTITY;
```

```
);  
END dff_control;
```

### Example 12-43: VHDL D-Type Flipflop (Register) with ena, aclr, and aload Control Signals (part 2)

```
ARCHITECTURE rtl OF dff_control IS  
BEGIN  
    PROCESS (clk, aclr, aload, adata)  
    BEGIN  
        IF (aclr = '1') THEN  
            q <= '0';  
        ELSIF (aload = '1') THEN  
            q <= adata;  
        ELSE  
            IF (clk = '1' AND clk'event) THEN  
                IF (ena = '1') THEN  
                    q <= data;  
                END IF;  
            END IF;  
        END IF;  
    END PROCESS;  
END rtl;
```

## Latches

A latch is a small combinational loop that holds the value of a signal until a new value is assigned.

Latches can be inferred from HDL code when you did not intend to use a latch. If you do intend to infer a latch, it is important to infer it correctly to guarantee correct device operation.

**Note:** Altera recommends that you design without the use of latches whenever possible.

### Related Information

- [Recommended Design Practices](#) on page 11-1

## Avoid Unintentional Latch Generation

When you are designing combinational logic, certain coding styles can create an unintentional latch.

For example, when `CASE` or `IF` statements do not cover all possible input conditions, latches may be required to hold the output if a new output value is not assigned. Check your synthesis tool messages for references to inferred latches. If your code unintentionally creates a latch, make code changes to remove the latch.

A latch is required if a signal is assigned a value outside of a clock edge (for example, with an asynchronous reset), but is not assigned a value in an edge-triggered design block. An unintentional latch may be generated if your HDL code assigns a value to a signal in an edge-triggered design block, but that logic is removed during synthesis. For example, when a `CASE` or `IF` statement tests the value of a condition with a parameter or generic that evaluates to `FALSE`, any logic or signal assignment in that statement is not required and is optimized away during synthesis. This optimization may result in a latch being generated for the signal.

**Note:** Latches have limited support in formal verification tools. Therefore, ensure that you do not infer latches unintentionally.

The `full_case` attribute can be used in Verilog HDL designs to treat unspecified cases as don't care values (x). However, using the `full_case` attribute can cause simulation mismatches because this attribute is a synthesis-only attribute, so simulation tools still treat the unspecified cases as latches.

Omitting the final `else` or `when others` clause in an `if` or `case` statement can also generate a latch. Don't care (x) assignments on the default conditions are useful in preventing latch generation. For the best logic optimization, assign the default `case` or final `else` value to don't care (x) instead of a logic value.

Without the final `else` clause, the following code creates unintentional latches to cover the remaining combinations of the `sel` inputs. When you are targeting a Stratix device with this code, omitting the final `else` condition can cause the synthesis software to use up to six LEs, instead of the three it uses with the `else` statement. Additionally, assigning the final `else` clause to 1 instead of x can result in slightly more LEs, because the synthesis software cannot perform as much optimization when you specify a constant value compared to a don't care value.

### Example 12-44: VHDL Code Preventing Unintentional Latch Creation

```
LIBRARY ieee;
USE IEEE.std_logic_1164.all;

ENTITY nolatch IS
    PORT (a,b,c: IN STD_LOGIC;
          sel: IN STD_LOGIC_VECTOR (4 DOWNTO 0);
          oput: OUT STD_LOGIC);
END nolatch;

ARCHITECTURE rtl OF nolatch IS
BEGIN
    PROCESS (a,b,c,sel) BEGIN
        if sel = "00000" THEN
            oput <= a;
        ELSIF sel = "00001" THEN
            oput <= b;
        ELSIF sel = "00010" THEN
            oput <= c;
        ELSE
            oput <= 'X'; --/          --- Prevents latch inference
        END if;
    END PROCESS;
END rtl;
```

#### Related Information

- [Quartus II Integrated Synthesis](#) on page 16-1

### Inferring Latches Correctly

Synthesis tools can infer a latch that does not exhibit the glitch and timing hazard problems typically associated with combinational loops. When using Quartus II integrated synthesis, latches that are inferred by the software are reported in the **User-Specified and Inferred Latches** section of the Compilation Report. This report indicates whether the latch is considered safe and free of timing hazards.

**Note:** Timing analysis does not completely model latch timing in some cases. Do not use latches unless required by your design, and you fully understand the impact of using the latches.

If a latch or combinational loop in your design is not listed in the **User Specified and Inferred Latches** section, it means that it was not inferred as a safe latch by the software and is not considered glitch-free.

All combinational loops listed in the **Analysis & Synthesis Logic Cells Representing Combinational Loops** table in the Compilation Report are at risk of timing hazards. These entries indicate possible problems with your design that you should investigate. However, it is possible to have a correct design that includes combinational loops. For example, it is possible that the combinational loop cannot be sensitized. This can occur in cases where there is an electrical path in the hardware, but either the designer knows that the circuit never encounters data that causes that path to be activated, or the surrounding logic is set up in a mutually exclusive manner that prevents that path from ever being sensitized, independent of the data input.

For macrocell-based devices, all data (D-type) latches and set-reset (S-R) latches listed in the **Analysis & Synthesis User Specified and Inferred Latches** table have an implementation free of timing hazards, such as glitches. The implementation includes both a cover term to ensure there is no glitching and a single macrocell in the feedback loop.

For 4-input LUT-based devices, such as Stratix devices, the Cyclone series, and MAX II devices, all latches in the **User Specified and Inferred Latches** table with a single LUT in the feedback loop are free of timing hazards when a single input changes. Because of the hardware behavior of the LUT, the output does not glitch when a single input toggles between two values that are supposed to produce the same output value, such as a D-type input toggling when the enable input is inactive or a set input toggling when a reset input with higher priority is active. This hardware behavior of the LUT means that no cover term is required for a loop around a single LUT. The Quartus II software uses a single LUT in the feedback loop whenever possible. A latch that has data, enable, set, and reset inputs in addition to the output fed back to the input cannot be implemented in a single 4-input LUT. If the Quartus II software cannot implement the latch with a single-LUT loop because there are too many inputs, the **User Specified and Inferred Latches** table indicates that the latch is not free of timing hazards.

For 6-input LUT-based devices, the software can implement all latch inputs with a single adaptive look-up table (ALUT) in the combinational loop. Therefore, all latches in the **User-Specified and Inferred Latches** table are free of timing hazards when a single input changes.

If a latch is listed as a safe latch, other optimizations performed by the Quartus II software, such as physical synthesis netlist optimizations in the Fitter, maintain the hazard-free performance. To ensure hazard-free behavior, only one control input can change at a time. Changing two inputs simultaneously, such as deasserting set and reset at the same time, or changing data and enable at the same time, can produce incorrect behavior in any latch.

Quartus II integrated synthesis infers latches from `always` blocks in Verilog HDL and `process` statements in VHDL, but not from continuous assignments in Verilog HDL or concurrent signal assignments in VHDL. These rules are the same as for register inference. The software infers registers or flipflops only from `always` blocks and `process` statements.

### Example 12-45: Verilog HDL Set-Reset Latch

```
module simple_latch (
    input SetTerm,
    input ResetTerm,
    output reg LatchOut
);

    always @ (SetTerm or ResetTerm) begin
        if (SetTerm)
            LatchOut = 1'b1
        else if (ResetTerm)
            LatchOut = 1'b0
    end
endmodule
```

```

    end
endmodule

```

### Example 12-46: VHDL Data Type Latch

```

LIBRARY IEEE;
USE IEEE.std_logic_1164.all;

ENTITY simple_latch IS
    PORT (
        enable, data    : IN STD_LOGIC;
        q               : OUT STD_LOGIC
    );
END simple_latch;

ARCHITECTURE rtl OF simple_latch IS
BEGIN

    latch : PROCESS (enable, data)
        BEGIN
            IF (enable = '1') THEN
                q <= data;
            END IF;
        END PROCESS latch;
END rtl;

```

The following example shows a Verilog HDL continuous assignment that does not infer a latch in the Quartus II software:

### Example 12-47: VHDL Continuous Assignment Does Not Infer Latch

```

assign latch_out = (~en & latch_out) | (en & data);

```

The behavior of the assignment is similar to a latch, but it may not function correctly as a latch, and its timing is not analyzed as a latch. Quartus II integrated synthesis also creates safe latches when possible for instantiations of an Altera latch IP core. You can use an Altera latch IP core to define a latch with any combination of data, enable, set, and reset inputs. The same limitations apply for creating safe latches as for inferring latches from HDL code.

Inferring Altera latch IP core in another synthesis tool ensures that the implementation is also recognized as a latch in the Quartus II software. If a third-party synthesis tool implements a latch using the Altera latch IP core, the Quartus II integrated synthesis lists the latch in the **User-Specified and Inferred Latches** table in the same way as it lists latches created in HDL source code. The coding style necessary to produce an Altera latch IP core implementation may depend on your synthesis tool. Some third-party synthesis tools list the number of Altera latch IP cores that are inferred.

For LUT-based families, the Fitter uses global routing for control signals, including signals that Analysis and Synthesis identifies as latch enables. In some cases the global insertion delay may decrease the timing performance. If necessary, you can turn off the **Quartus II Global Signal** logic option to manually prevent the use of global signals. Global latch enables are listed in the **Global & Other Fast Signals** table in the Compilation Report.

## General Coding Guidelines

This section describes how coding styles impacts synthesis of HDL code into the target Altera device.

Following Altera recommended coding styles, and in some cases designing logic structures to match the appropriate device architecture, can provide significant improvements in your design's efficiency and performance.

### Tri-State Signals

When you target Altera devices, you should use tri-state signals only when they are attached to top-level bidirectional or output pins.

Avoid lower-level bidirectional pins, and avoid using the z logic value unless it is driving an output or bidirectional pin. Synthesis tools implement designs with internal tri-state signals correctly in Altera devices using multiplexer logic, but Altera does not recommend this coding practice.

**Note:** In hierarchical block-based or incremental design flows, a hierarchical boundary cannot contain any bidirectional ports, unless the lower-level bidirectional port is connected directly through the hierarchy to a top-level output pin without connecting to any other design logic. If you use boundary tri-states in a lower-level block, synthesis software must push the tri-states through the hierarchy to the top level to make use of the tri-state drivers on output pins of Altera devices. Because pushing tri-states requires optimizing through hierarchies, lower-level tri-states are restricted with block-based design methodologies.

### Clock Multiplexing


Clock multiplexing is sometimes used to operate the same logic function with different clock sources.

This type of logic can introduce glitches that create functional problems, and the delay inherent in the combinational logic can lead to timing problems. Clock multiplexers trigger warnings from a wide range of design rule check and timing analysis tools.

Altera recommends using dedicated hardware to perform clock multiplexing when it is available, instead of using multiplexing logic. For example, you can use the Clock Switchover feature or the Clock Control Block available in certain Altera devices. These dedicated hardware blocks avoid glitches, ensure that you use global low-skew routing lines, and avoid any possible hold time problems on the device due to logic delay on the clock line. Many Altera devices also support dynamic PLL reconfiguration, which is the safest and most robust method of changing clock rates during device operation.

If you implement a clock multiplexer in logic cells because the design has too many clocks to use the clock control block, or if dynamic reconfiguration is too complex for your design, it is important to consider simultaneous toggling inputs and ensure glitch-free transitions.

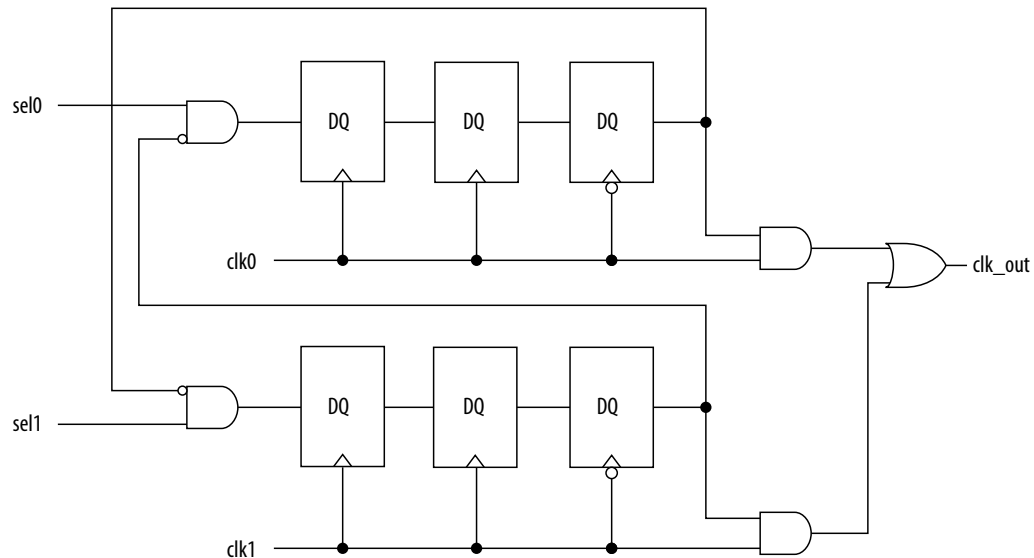
Figure 12-2: Simple Clock Multiplexer in a 6-Input LUT



The data sheet for your target device describes how LUT outputs may glitch during a simultaneous toggle of input signals, independent of the LUT function. Although, in practice, the 4:1 MUX function does not generate detectable glitches during simultaneous data input toggles, it is possible to construct cell implementations that do exhibit significant glitches, so this simple clock mux structure is not recommended. An additional problem with this implementation is that the output behaves erratically during a change in the `clk_select` signals. This behavior could create timing violations on all registers fed by the system clock and result in possible metastability.

A more sophisticated clock select structure can eliminate the simultaneous toggle and switching problems.

Figure 12-3: Glitch-Free Clock Multiplexer Structure



You can generalize this structure for any number of clock channels. The design ensures that no clock activates until all others are inactive for at least a few cycles, and that activation occurs while the clock is low. The design applies a `synthesis_keep` directive to the AND gates on the right side, which ensures there are no simultaneous toggles on the input of the `clk_out` OR gate.

**Note:** Switching from clock A to clock B requires that clock A continue to operate for at least a few cycles. If the old clock stops immediately, the design sticks. The select signals are implemented as a “one-hot” control in this example, but you can use other encoding if you prefer. The input side logic is asynchronous and is not critical. This design can tolerate extreme glitching during the switch process.

### Example 12-48: Verilog HDL Clock Multiplexing Design to Avoid Glitches

```
module clock_mux (clk,clk_select,clk_out);

    parameter num_clocks = 4;

    input [num_clocks-1:0] clk;
    input [num_clocks-1:0] clk_select; // one hot
    output clk_out;

    genvar i;

    reg [num_clocks-1:0] ena_r0;
    reg [num_clocks-1:0] ena_r1;
    reg [num_clocks-1:0] ena_r2;
    wire [num_clocks-1:0] qualified_sel;

    // A look-up-table (LUT) can glitch when multiple inputs
    // change simultaneously. Use the keep attribute to
    // insert a hard logic cell buffer and prevent
    // the unrelated clocks from appearing on the same LUT.

    wire [num_clocks-1:0] gated_clks /* synthesis keep */;

    initial begin
        ena_r0 = 0;
        ena_r1 = 0;
        ena_r2 = 0;
    end

    generate
        for (i=0; i<num_clocks; i=i+1)
            begin : lp0
                wire [num_clocks-1:0] tmp_mask;
                assign tmp_mask = {num_clocks{1'b1}} ^ (1 << i);

                assign qualified_sel[i] = clk_select[i] & (~|(ena_r2 &
tmp_mask));

                always @(posedge clk[i]) begin
                    ena_r0[i] <= qualified_sel[i];
                    ena_r1[i] <= ena_r0[i];
                end

                always @(negedge clk[i]) begin
                    ena_r2[i] <= ena_r1[i];
                end

                assign gated_clks[i] = clk[i] & ena_r2[i];
            end
    endgenerate

    // These will not exhibit simultaneous toggle by construction
    assign clk_out = |gated_clks;

endmodule
```



**Related Information**[Altera IP Core Literature](#)

## Adder Trees

Structuring adder trees appropriately to match your targeted Altera device architecture can provide significant improvements in your design's efficiency and performance.

A good example of an application using a large adder tree is a finite impulse response (FIR) correlator. Using a pipelined binary or ternary adder tree appropriately can greatly improve the quality of your results.

This section explains why coding recommendations are different for Altera 4-input LUT devices and 6-input LUT devices.

### Architectures with 4-Input LUTs in Logic Elements

Architectures such as Stratix devices and the Cyclone series of devices contain 4-input LUTs as the standard combinational structure in the LE.

If your design can tolerate pipelining, the fastest way to add three numbers  $A$ ,  $B$ , and  $C$  in devices that use 4-input lookup tables is to add  $A + B$ , register the output, and then add the registered output to  $C$ . Adding  $A + B$  takes one level of logic (one bit is added in one LE), so this runs at full clock speed. This can be extended to as many numbers as desired.

Adding five numbers in devices that use 4-input lookup tables requires four adders and three levels of registers for a total of 64 LEs (for 16-bit numbers).

### Architectures with 6-Input LUTs in Adaptive Logic Modules

High-performance Altera device families use a 6-input LUT in their basic logic structure. These devices benefit from a different coding style from the previous example presented for 4-input LUTs.

Specifically, in these devices, ALMs can simultaneously add three bits. Therefore, the tree must be two levels deep and contain just two add-by-three inputs instead of four add-by-two inputs.

Although the code in the previous example compiles successfully for 6-input LUT devices, the code is inefficient and does not take advantage of the 6-input adaptive ALUT. By restructuring the tree as a ternary tree, the design becomes much more efficient, significantly improving density utilization. Therefore, when you are targeting with ALUTs and ALMs, large pipelined binary adder trees designed for 4-input LUT architectures should be rewritten to take advantage of the advanced device architecture.

**Note:** You cannot pack a LAB full when using this type of coding style because of the number of LAB inputs. However, in a typical design, the Quartus II Fitter can pack other logic into each LAB to take advantage of the unused ALMs.

These examples show pipelined adders, but partitioning your addition operations can help you achieve better results in nonpipelined adders as well. If your design is not pipelined, a ternary tree provides much better performance than a binary tree. For example, depending on your synthesis tool, the HDL code `sum = (A + B + C) + (D + E)` is more likely to create the optimal implementation of a 3-input adder for  $A + B + C$  followed by a 3-input adder for  $sum1 + D + E$  than the code without the parentheses. If you do not add the parentheses, the synthesis tool may partition the addition in a way that is not optimal for the architecture.

#### Example 12-49: Verilog-2001 State Machine

```
module verilog_fsm (clk, reset, in_1, in_2, out);
```

```

input clk, reset;
input [3:0] in_1, in_2;
output [4:0] out;
parameter state_0 = 3'b000;
parameter state_1 = 3'b001;
parameter state_2 = 3'b010;
parameter state_3 = 3'b011;
parameter state_4 = 3'b100;

reg [4:0] tmp_out_0, tmp_out_1, tmp_out_2;
reg [2:0] state, next_state;

always @ (posedge clk or posedge reset)
begin
    if (reset)
        state <= state_0;
    else
        state <= next_state;
end
always @ (*)
begin
    tmp_out_0 = in_1 + in_2;
    tmp_out_1 = in_1 - in_2;
    case (state)
        state_0: begin
            tmp_out_2 = in_1 + 5'b00001;
            next_state = state_1;
        end
        state_1: begin
            if (in_1 < in_2) begin
                next_state = state_2;
                tmp_out_2 = tmp_out_0;
            end
            else begin
                next_state = state_3;
                tmp_out_2 = tmp_out_1;
            end
        end
        state_2: begin
            tmp_out_2 = tmp_out_0 - 5'b00001;
            next_state = state_3;
        end
        state_3: begin
            tmp_out_2 = tmp_out_1 + 5'b00001;
            next_state = state_0;
        end
        state_4: begin
            tmp_out_2 = in_2 + 5'b00001;
            next_state = state_0;
        end
        default: begin
            tmp_out_2 = 5'b00000;
            next_state = state_0;
        end
    endcase
end
assign out = tmp_out_2;
endmodule

```

An equivalent implementation of this state machine can be achieved by using ``define` instead of the parameter data type, as follows:

```

`define state_0 3'b000
`define state_1 3'b001
`define state_2 3'b010

```

```
`define state_3 3'b011
`define state_4 3'b100
```

In this case, the `state` and `next_state` assignments are assigned a `'state_x` instead of a `state_x`, for example:

```
next_state <= `state_3;
```

**Note:** Although the ``define` construct is supported, Altera strongly recommends the use of the `parameter` data type because doing so preserves the state names throughout synthesis.

## State Machine HDL Guidelines

Synthesis tools can recognize and encode Verilog HDL and VHDL state machines during synthesis. This section presents guidelines to ensure the best results when you use state machines.

Ensuring that your synthesis tool recognizes a piece of code as a state machine allows the tool to recode the state variables to improve the quality of results, and allows the tool to use the known properties of state machines to optimize other parts of the design. When synthesis recognizes a state machine, it is often able to improve the design area and performance.

To achieve the best results on average, synthesis tools often use one-hot encoding for FPGA devices and minimal-bit encoding for CPLD devices, although the choice of implementation can vary for different state machines and different devices. Refer to your synthesis tool documentation for specific ways to control the manner in which state machines are encoded.

To ensure proper recognition and inference of state machines and to improve the quality of results, Altera recommends that you observe the following guidelines, which apply to both Verilog HDL and VHDL:

- Assign default values to outputs derived from the state machine so that synthesis does not generate unwanted latches.
- Separate the state machine logic from all arithmetic functions and data paths, including assigning output values.
- If your design contains an operation that is used by more than one state, define the operation outside the state machine and cause the output logic of the state machine to use this value.
- Use a simple asynchronous or synchronous reset to ensure a defined power-up state. If your state machine design contains more elaborate reset logic, such as both an asynchronous reset and an asynchronous load, the Quartus II software generates regular logic rather than inferring a state machine.

If a state machine enters an illegal state due to a problem with the device, the design likely ceases to function correctly until the next reset of the state machine. Synthesis tools do not provide for this situation by default. The same issue applies to any other registers if there is some kind of fault in the system. A `default` or `when others` clause does not affect this operation, assuming that your design never deliberately enters this state. Synthesis tools remove any logic generated by a default state if it is not reachable by normal state machine operation.

Many synthesis tools (including Quartus II integrated synthesis) have an option to implement a safe state machine. The software inserts extra logic to detect an illegal state and force the state machine's transition to the reset state. It is commonly used when the state machine can enter an illegal state. The most

common cause of this situation is a state machine that has control inputs that come from another clock domain, such as the control logic for a dual-clock FIFO.

This option protects only state machines by forcing them into the reset state. All other registers in the design are not protected this way. If the design has asynchronous inputs, Altera recommends using a synchronization register chain instead of relying on the safe state machine option.

#### Related Information

- [Quartus II Integrated Synthesis](#) on page 16-1

## Verilog HDL State Machines

To ensure proper recognition and inference of Verilog HDL state machines, observe the following additional Verilog HDL guidelines.

Some of these guidelines may be specific to Quartus II integrated synthesis. Refer to your synthesis tool documentation for specific coding recommendations. If the state machine is not recognized and inferred by the synthesis software (such as Quartus II integrated synthesis), the state machine is implemented as regular logic gates and registers, and the state machine is not listed as a state machine in the **Analysis & Synthesis** section of the Quartus II Compilation Report. In this case, the software does not perform any of the optimizations that are specific to state machines.

- If you are using the SystemVerilog standard, use enumerated types to describe state machines.
- Represent the states in a state machine with the parameter data types in Verilog-1995 and Verilog-2001, and use the parameters to make state assignments. This parameter implementation makes the state machine easier to read and reduces the risk of errors during coding.
- Altera recommends against the direct use of integer values for state variables, such as `next_state <= 0`. However, using an integer does not prevent inference in the Quartus II software.
- No state machine is inferred in the Quartus II software if the state transition logic uses arithmetic similar to that in the following example:

```
case (state)
  0: begin
    if (ena) next_state <= state + 2;
    else next_state <= state + 1;
  end
  1: begin
    ...
  end
endcase
```

```
case (state)0: beginif (ena) next_state <= state + 2;else next_state <= state + 1;end1: begin...endcase
```

- No state machine is inferred in the Quartus II software if the state variable is an output.
- No state machine is inferred in the Quartus II software for signed variables.

#### Related Information

- [Verilog-2001 State Machine Coding Example](#) on page 12-47
- [SystemVerilog State Machine Coding Example](#) on page 12-49

### Verilog-2001 State Machine Coding Example

The following module `verilog_fsm` is an example of a typical Verilog HDL state machine implementation.

This state machine has five states. The asynchronous reset sets the variable `state` to `state_0`. The sum of `in_1` and `in_2` is an output of the state machine in `state_1` and `state_2`. The difference (`in_1 - in_2`) is

also used in `state_1` and `state_2`. The temporary variables `tmp_out_0` and `tmp_out_1` store the sum and the difference of `in_1` and `in_2`. Using these temporary variables in the various states of the state machine ensures proper resource sharing between the mutually exclusive states.

### Example 12-50: Verilog-2001 State Machine

```

module verilog_fsm (clk, reset, in_1, in_2, out);
    input clk, reset;
    input [3:0] in_1, in_2;
    output [4:0] out;
    parameter state_0 = 3'b000;
    parameter state_1 = 3'b001;
    parameter state_2 = 3'b010;
    parameter state_3 = 3'b011;
    parameter state_4 = 3'b100;

    reg [4:0] tmp_out_0, tmp_out_1, tmp_out_2;
    reg [2:0] state, next_state;

    always @ (posedge clk or posedge reset)
    begin
        if (reset)
            state <= state_0;
        else
            state <= next_state;
    end
    always @ (*)
    begin
        tmp_out_0 = in_1 + in_2;
        tmp_out_1 = in_1 - in_2;
        case (state)
            state_0: begin
                tmp_out_2 = in_1 + 5'b00001;
                next_state = state_1;
            end
            state_1: begin
                if (in_1 < in_2) begin
                    next_state = state_2;
                    tmp_out_2 = tmp_out_0;
                end
                else begin
                    next_state = state_3;
                    tmp_out_2 = tmp_out_1;
                end
            end
            state_2: begin
                tmp_out_2 = tmp_out_0 - 5'b00001;
                next_state = state_3;
            end
            state_3: begin
                tmp_out_2 = tmp_out_1 + 5'b00001;
                next_state = state_0;
            end
            state_4: begin
                tmp_out_2 = in_2 + 5'b00001;
                next_state = state_0;
            end
            default: begin
                tmp_out_2 = 5'b00000;
                next_state = state_0;
            end
        endcase
    end
end

```

```
    assign out = tmp_out_2;
endmodule
```

An equivalent implementation of this state machine can be achieved by using ``define` instead of the `parameter` data type, as follows:

```
`define state_0 3'b000
`define state_1 3'b001
`define state_2 3'b010
`define state_3 3'b011
`define state_4 3'b100
```

In this case, the `state` and `next_state` assignments are assigned a `'state_x` instead of a `state_x`, for example:

```
next_state <= `state_3;
```

**Note:** Although the ``define` construct is supported, Altera strongly recommends the use of the `parameter` data type because doing so preserves the state names throughout synthesis.

### SystemVerilog State Machine Coding Example

The module `enum_fsm` is an example of a SystemVerilog state machine implementation that uses enumerated types. Altera recommends using this coding style to describe state machines in SystemVerilog.

**Note:** In Quartus II integrated synthesis, the enumerated type that defines the states for the state machine must be of an unsigned integer type. If you do not specify the enumerated type as `int unsigned`, a signed `int` type is used by default. In this case, the Quartus II integrated synthesis synthesizes the design, but does not infer or optimize the logic as a state machine.

### Example 12-51: SystemVerilog State Machine Using Enumerated Types

```
module enum_fsm (input clk, reset, input int data[3:0], output int o);
    enum int unsigned { S0 = 0, S1 = 2, S2 = 4, S3 = 8 } state, next_state;

    always_comb begin : next_state_logic
        next_state = S0;
        case(state)
            S0: next_state = S1;
            S1: next_state = S2;
            S2: next_state = S3;
            S3: next_state = S3;
        endcase
    end

    always_comb begin
        case(state)
            S0: o = data[3];
            S1: o = data[2];
            S2: o = data[1];
            S3: o = data[0];
        endcase
    end
end
```

```

always_ff@(posedge clk or negedge reset) begin
    if(~reset)
        state <= S0;
    else
        state <= next_state;
    end
endmodule

```

## VHDL State Machines

To ensure proper recognition and inference of VHDL state machines, represent the states in a state machine with enumerated types and use the corresponding types to make state assignments.

This implementation makes the state machine easier to read and reduces the risk of errors during coding. If the state is not represented by an enumerated type, synthesis software (such as Quartus II integrated synthesis) does not recognize the state machine. Instead, the state machine is implemented as regular logic gates and registers and the state machine is not listed as a state machine in the **Analysis & Synthesis** section of the Quartus II Compilation Report. In this case, the software does not perform any of the optimizations that are specific to state machines.

### VHDL State Machine Coding Example

The following state machine has five states. The asynchronous reset sets the variable `state` to `state_0`.

The sum of `in1` and `in2` is an output of the state machine in `state_1` and `state_2`. The difference (`in1 - in2`) is also used in `state_1` and `state_2`. The temporary variables `tmp_out_0` and `tmp_out_1` store the sum and the difference of `in1` and `in2`. Using these temporary variables in the various states of the state machine ensures proper resource sharing between the mutually exclusive states.

### Example 12-52: VHDL State Machine

```

LIBRARY ieee;
USE ieee.std_logic_1164.all;
USE ieee.numeric_std.all;
ENTITY vhdl_fsm IS
    PORT(
        clk: IN STD_LOGIC;
        reset: IN STD_LOGIC;
        in1: IN UNSIGNED(4 downto 0);
        in2: IN UNSIGNED(4 downto 0);
        out_1: OUT UNSIGNED(4 downto 0)
    );
END vhdl_fsm;
ARCHITECTURE rtl OF vhdl_fsm IS
    TYPE Tstate IS (state_0, state_1, state_2, state_3, state_4);
    SIGNAL state: Tstate;
    SIGNAL next_state: Tstate;
BEGIN
    PROCESS(clk, reset)
    BEGIN
        IF reset = '1' THEN
            state <= state_0;
        ELSIF rising_edge(clk) THEN
            state <= next_state;
        END IF;
    END PROCESS;
    PROCESS (state, in1, in2)
        VARIABLE tmp_out_0: UNSIGNED (4 downto 0);
        VARIABLE tmp_out_1: UNSIGNED (4 downto 0);
    BEGIN

```

```
tmp_out_0 := in1 + in2;
tmp_out_1 := in1 - in2;
CASE state IS
  WHEN state_0 =>
    out_1 <= in1;
    next_state <= state_1;
  WHEN state_1 =>
    IF (in1 < in2) then
      next_state <= state_2;
      out_1 <= tmp_out_0;
    ELSE
      next_state <= state_3;
      out_1 <= tmp_out_1;
    END IF;
  WHEN state_2 =>
    IF (in1 < "0100") then
      out_1 <= tmp_out_0;
    ELSE
      out_1 <= tmp_out_1;
    END IF;
    next_state <= state_3;
  WHEN state_3 =>
    out_1 <= "11111";
    next_state <= state_4;
  WHEN state_4 =>
    out_1 <= in2;
    next_state <= state_0;
  WHEN OTHERS =>
    out_1 <= "00000";
    next_state <= state_0;
END CASE;
END PROCESS;
END rtl;
```

## Multiplexer HDL Guidelines

Multiplexers form a large portion of the logic utilization in many FPGA designs. By optimizing your multiplexer logic, you ensure the most efficient implementation in your Altera device.

This section addresses common problems and provides design guidelines to achieve optimal resource utilization for multiplexer designs. The section also describes various types of multiplexers, and how they are implemented.

For more information, refer to the *Advanced Synthesis Cookbook*.

### Related Information

[Advanced Synthesis Cookbook](#)

## Quartus II Software Option for Multiplexer Restructuring

Quartus II integrated synthesis provides the **Restructure Multiplexers** logic option that extracts and optimizes buses of multiplexers during synthesis.

The default setting **Auto** for this option uses the optimization when it is most likely to benefit the optimization targets for your design. You can turn the option on or off specifically to have more control over its use.

Even with this Quartus II-specific option turned on, it is beneficial to understand how your coding style can be interpreted by your synthesis tool, and avoid the situations that can cause problems in your design.



**Related Information**

- [Quartus II Integrated Synthesis](#) on page 16-1

**Multiplexer Types**

This section addresses how multiplexers are created from various types of HDL code. *CASE* statements, *IF* statements, and state machines are all common sources of multiplexer logic in designs.

These HDL structures create different types of multiplexers, including binary multiplexers, selector multiplexers, and priority multiplexers. Understanding how multiplexers are created from HDL code, and how they might be implemented during synthesis, is the first step toward optimizing multiplexer structures for best results.

**Binary Multiplexers**

Binary multiplexers select inputs based on binary-encoded selection bits.

Stratix series devices starting with the Stratix II device family feature 6-input look up tables (LUTs) which are perfectly suited for 4:1 multiplexer building blocks (4 data and 2 select inputs). The extended input mode facilitates implementing 8:1 blocks, and the fractured mode handles residual 2:1 multiplexer pairs. For device families using 4-input LUTs, such as the Cyclone series and Stratix devices, the 4:1 binary multiplexer is efficiently implemented by using two 4-input LUTs. Larger binary multiplexers are decomposed by the synthesis tool into 4:1 multiplexer blocks, possibly with a residual 2:1 multiplexer at the head.

**Example 12-53: Verilog HDL Binary-Encoded Multiplexers**

```
case (sel)
  2'b00: z = a;
  2'b01: z = b;
  2'b10: z = c;
  2'b11: z = d;
endcase
```

**Selector Multiplexers**

Selector multiplexers have a separate select line for each data input.

The select lines for the multiplexer are one-hot encoded. Selector multiplexers are commonly built as a tree of AND and OR gates. An N-input selector multiplexer of this structure is slightly less efficient in implementation than a binary multiplexer. However, in many cases the select signal is the output of a decoder, in which case Quartus II Synthesis will try to combine the selector and decoder into a binary multiplexer.

**Example 12-54: Verilog HDL One-Hot-Encoded Case Statement**

```
case (sel)
  4'b0001: z = a;
  4'b0010: z = b;
  4'b0100: z = c;
  4'b1000: z = d;
  default: z = 1'bx;
endcase
```

## Priority Multiplexers

In priority multiplexers, the select logic implies a priority. The options to select the correct item must be checked in a specific order based on signal priority.


These structures commonly are created from IF, ELSE, WHEN, SELECT, and ?: statements in VHDL or Verilog HDL.

### Example 12-55: VHDL IF Statement Implying Priority

```
IF cond1 THEN z <= a;
ELSIF cond2 THEN z <= b;
ELSIF cond3 THEN z <= c;
ELSE z <= d;
END IF;
```

The multiplexers form a chain, evaluating each condition or select bit sequentially.

Figure 12-4: Priority Multiplexer Implementation of an IF Statement



Depending on the number of multiplexers in the chain, the timing delay through this chain can become large, especially for device families with 4-input LUTs.

To improve the timing delay through the multiplexer, avoid priority multiplexers if priority is not required. If the order of the choices is not important to the design, use a CASE statement to implement a binary or selector multiplexer instead of a priority multiplexer. If delay through the structure is important in a multiplexed design requiring priority, consider recoding the design to reduce the number of logic levels to minimize delay, especially along your critical paths.

## Implicit Defaults in If Statements

The IF statements in Verilog HDL and VHDL can be a convenient way to specify conditions that do not easily lend themselves to a CASE-type approach.

However, using IF statements can result in complicated multiplexer trees that are not easy for synthesis tools to optimize. In particular, every IF statement has an implicit ELSE condition, even when it is not specified. These implicit defaults can cause additional complexity in a multiplexed design.

There are several ways you can simplify multiplexed logic and remove unneeded defaults. The optimal method may be to recode the design so the logic takes the structure of a 4:1 CASE statement. Alternatively, if priority is important, you can restructure the code to reduce default cases and flatten the multiplexer. Examine whether the default "ELSE IF" conditions are don't care cases. You may be able to create a

default `ELSE` statement to make the behavior explicit. Avoid unnecessary default conditions in your multiplexer logic to reduce the complexity and logic utilization required to implement your design.

## Default or Others Case Assignment

To fully specify the cases in a `CASE` statement, include a `default` (Verilog HDL) or `OTHERS` (VHDL) assignment.

This assignment is especially important in one-hot encoding schemes where many combinations of the select lines are unused. Specifying a case for the unused select line combinations gives the synthesis tool information about how to synthesize these cases, and is required by the Verilog HDL and VHDL language specifications.

Some designs do not require that the outcome in the unused cases be considered, often because designers assume these cases will not occur. For these types of designs, you can specify any value for the `default` or `OTHERS` assignment. However, be aware that the assignment value you choose can have a large effect on the logic utilization required to implement the design due to the different ways synthesis tools treat different values for the assignment, and how the synthesis tools use different speed and area optimizations.

To obtain best results, explicitly define invalid `CASE` selections with a separate `default` or `OTHERS` statement instead of combining the invalid cases with one of the defined cases.

If the value in the invalid cases is not important, specify those cases explicitly by assigning the `x` (don't care) logic value instead of choosing another value. This assignment allows your synthesis tool to perform the best area optimizations.

## Cyclic Redundancy Check Functions

CRC computations are used heavily by communications protocols and storage devices to detect any corruption of data.

These functions are highly effective; there is a very low probability that corrupted data can pass a 32-bit CRC check.

CRC functions typically use wide XOR gates to compare the data. The way synthesis tools flatten and factor these XOR gates to implement the logic in FPGA LUTs can greatly impact the area and performance results for the design. XOR gates have a cancellation property that creates an exceptionally large number of reasonable factoring combinations, so synthesis tools cannot always choose the best result by default.

The 6-input ALUT has a significant advantage over 4-input LUTs for these designs. When properly synthesized, CRC processing designs can run at high speeds in devices with 6-input ALUTs.

The following guidelines help you improve the quality of results for CRC designs in Altera devices.

### If Performance is Important, Optimize for Speed

Synthesis tools flatten XOR gates to minimize area and depth of levels of logic.

Synthesis tools such as Quartus II integrated synthesis target area optimization by default for these logic structures. Therefore, for more focus on depth reduction, set the synthesis optimization technique to speed.

Flattening for depth sometimes causes a significant increase in area.

## Use Separate CRC Blocks Instead of Cascaded Stages

Some designers optimize their CRC designs to use cascaded stages (for example, four stages of 8 bits). In such designs, intermediate calculations are used as required (such as the calculations after 8, 24, or 32 bits) depending on the data width.

This design is not optimal in FPGA devices. The XOR cancellations that can be performed in CRC designs mean that the function does not require all the intermediate calculations to determine the final result. Therefore, forcing the use of intermediate calculations increases the area required to implement the function, as well as increasing the logic depth because of the cascading. It is typically better to create full separate CRC blocks for each data width that you require in the design, and then multiplex them together to choose the appropriate mode at a given time.

## Use Separate CRC Blocks Instead of Allowing Blocks to Merge

Synthesis tools often attempt to optimize CRC designs by sharing resources and extracting duplicates in two different CRC blocks because of the factoring options in the XOR logic.

The CRC logic allows significant reductions, but this works best when each CRC function is optimized separately. Check for duplicate extraction behavior if you have different CRC functions that are driven by common data signals or that feed the same destination signals.

If you are having problems with the quality of results and you see that two CRC functions are sharing logic, ensure that the blocks are synthesized independently using one of the following methods:

- Define each CRC block as a separate design partition in an incremental compilation design flow.
- Synthesize each CRC block as a separate project in your third-party synthesis tool and then write a separate Verilog Quartus Mapping (**.vqm**) or EDIF netlist file for each.

### Related Information

- [Quartus II Incremental Compilation](#) on page 3-1

## Take Advantage of Latency if Available

If your design can use more than one cycle to implement the CRC functionality, adding registers and retiming the design can help reduce area, improve performance, and reduce power utilization.

If your synthesis tool offers a retiming feature (such as the Quartus II software **Perform gate-level register retiming** option), you can insert an extra bank of registers at the input and allow the retiming feature to move the registers for better results. You can also build the CRC unit half as wide and alternate between halves of the data in each clock cycle.

## Save Power by Disabling CRC Blocks When Not in Use

CRC designs are heavy consumers of dynamic power because the logic toggles whenever there is a change in the design.

To save power, use clock enables to disable the CRC function for every clock cycle that the logic is not required. Some designs don't check the CRC results for a few clock cycles while other logic is performed. It is valuable to disable the CRC function even for this short amount of time.

## Use the Device Synchronous Load (sload) Signal to Initialize

The data in many CRC designs must be initialized to 1's before operation. If your target device supports the use of the `sload` signal, you should use it to set all the registers in your design to 1's before operation.

To enable use of the `sload` signal, follow the coding guidelines presented in this chapter. You can check the register equations in the Chip Planner to ensure that the signal was used as expected.

If you must force a register implementation using an `sload` signal, you can use low-level device primitives as described in *Designing with Low-Level Primitives User Guide*.

#### Related Information

- [Secondary Register Control Signals Such as Clear and Clock Enable](#) on page 12-35
- [Designing with Low-Level Primitives User Guide](#)

## Comparator HDL Guidelines

Synthesis software, including Quartus II integrated synthesis, uses device and context-specific implementation rules for comparators (`<`, `>`, or `==`) and selects the best one for your design.

This section provides some information about the different types of implementations available and provides suggestions on how you can code your design to encourage a specific implementation.

The `==` comparator is implemented in general logic cells. The `<` comparison can be implemented using the carry chain or general logic cells. In devices with 6-input ALUTs, the carry chain is capable of comparing up to three bits per cell. In devices with 4-input LUTs, the capacity is one bit of comparison per cell, which is similar to an add/subtract chain. The carry chain implementation tends to be faster than the general logic on standalone benchmark test cases, but can result in lower performance when it is part of a larger design due to the increased restriction on the Fitter. The area requirement is similar for most input patterns. The synthesis software selects an appropriate implementation based on the input pattern.

If you are using Quartus II integrated synthesis, you can guide the synthesis by using specific coding styles. To select a carry chain implementation explicitly, rephrase your comparison in terms of addition. As a simple example, the following coding style allows the synthesis tool to select the implementation, which is most likely using general logic cells in modern device families:

```
wire [6:0] a,b;
wire alb = a<b;
```

In the following coding style, the synthesis tool uses a carry chain (except for a few cases, such as when the chain is very short or the signals `a` and `b` minimize to the same signal):


```
wire [6:0] a,b;
wire [7:0] tmp = a - b;
wire alb = tmp[7]
```

This second coding style uses the top bit of the `tmp` signal, which is 1 in twos complement logic if `a` is less than `b`, because the subtraction `a - b` results in a negative number.

If you have any information about the range of the input, you have “don’t care” values that you can use to optimize the design. Because this information is not available to the synthesis tool, you can often reduce the device area required to implement the comparator with specific hand implementation of the logic.

You can also check whether a bus value is within a constant range with a small amount of logic area by using the following logic structure. This type of logic occurs frequently in address decoders.

Figure 12-5: Example Logic Structure for Using Comparators to Check a Bus Value Range



## Counter HDL Guidelines

Implementing counters in HDL code is easy; they are implemented with an adder followed by registers.

Remember that the register control signals, such as enable (*ena*), synchronous clear (*sc1r*), and synchronous load (*sload*), are available. For the best area utilization, ensure that the up/down control or controls are expressed in terms of one addition instead of two separate addition operators.

If you use the following coding style, your synthesis tool may implement two separate carry chains for addition (if it doesn't detect the issue and optimize the logic):

```
out <= count_up ? out + 1 : out - 1;
```

The following coding style requires only one adder along with some other logic:

```
out <= out + (count_up ? 1 : -1);
```

In this case, the coding style better matches the device hardware because there is only one carry chain adder, and the  $-1$  constant logic is implemented in the LUT in front of the adder without adding extra area utilization.

## Designing with Low-Level Primitives

Low-level HDL design is the practice of using low-level primitives and assignments to dictate a particular hardware implementation for a piece of logic. Low-level primitives are small architectural building blocks that assist you in creating your design.

With the Quartus II software, you can use low-level HDL design techniques to force a specific hardware implementation that can help you achieve better resource utilization or faster timing results.

**Note:** Using low-level primitives is an advanced technique to help with specific design challenges, and is optional in the Altera design flow. For many designs, synthesizing generic HDL source code and Altera IP cores give you the best results.

Low-level primitives allow you to use the following types of coding techniques:

- Instantiate the logic cell or `LCELL` primitive to prevent Quartus II integrated synthesis from performing optimizations across a logic cell
- Create carry and cascade chains using `CARRY`, `CARRY_SUM`, and `CASCADE` primitives
- Instantiate registers with specific control signals using `DFE` primitives
- Specify the creation of LUT functions by identifying the LUT boundaries
- Use I/O buffers to specify I/O standards, current strengths, and other I/O assignments
- Use I/O buffers to specify differential pin names in your HDL code, instead of using the automatically-generated negative pin name for each pair

For details about and examples of using these types of assignments, refer to the *Designing with Low-Level Primitives User Guide*.

#### Related Information

[Designing with Low-Level Primitives User Guide](#)

## Document Revision History

The following revisions history applies to this chapter.

**Table 12-2: Document Revision History**

Date	Version	Changes
2014.12.15	14.1.0	Updated location of Fitter Settings, Analysis & Synthesis Settings, and Physical Optimization Settings to Compiler Settings.
2014.08.18	14.0.a10.0	<ul style="list-style-type: none"> <li>• Added recommendation to use register pipelining to obtain high performance in DSP designs.</li> </ul>
2014.06.30	14.0.0	Removed obsolete MegaWizard Plug-In Manager support.
November 2013	13.1.0	Removed HardCopy device support.
June 2012	12.0.0	<ul style="list-style-type: none"> <li>• Revised section on inserting Altera templates.</li> <li>• Code update for Example 11-51.</li> <li>• Minor corrections and updates.</li> </ul>
November 2011	11.1.0	<ul style="list-style-type: none"> <li>• Updated document template.</li> <li>• Minor updates and corrections.</li> </ul>
December 2010	10.1.0	<ul style="list-style-type: none"> <li>• Changed to new document template.</li> <li>• Updated Unintentional Latch Generation content.</li> <li>• Code update for Example 11-18.</li> </ul>
July 2010	10.0.0	<ul style="list-style-type: none"> <li>• Added support for mixed-width RAM</li> <li>• Updated support for <code>no_rw_check</code> for inferring RAM blocks</li> <li>• Added support for byte-enable</li> </ul>

Date	Version	Changes
November 2009	9.1.0	<ul style="list-style-type: none"> <li>Updated support for Controlling Inference and Implementation in Device RAM Blocks</li> <li>Updated support for Shift Registers</li> </ul>
March 2009	9.0.0	<ul style="list-style-type: none"> <li>Corrected and updated several examples</li> <li>Added support for Arria II GX devices</li> <li>Other minor changes to chapter</li> </ul>
November 2008	8.1.0	Changed to 8-1/2 x 11 page size. No change to content.
May 2008	8.0.0	Updates for the Quartus II software version 8.0 release, including: <ul style="list-style-type: none"> <li>Added information to “RAM Functions—Inferring ALTSYNCRAM and ALTDPRAM Megafunctions from HDL Code” on page 6–13</li> <li>Added information to “Avoid Unsupported Reset and Control Conditions” on page 6–14</li> <li>Added information to “Check Read-During-Write Behavior” on page 6–16</li> <li>Added two new examples to “ROM Functions—Inferring ALTSYNCRAM and LPM_ROM Megafunctions from HDL Code” on page 6–28: Example 6–24 and Example 6–25</li> <li>Added new section: “Clock Multiplexing” on page 6–46</li> <li>Added hyperlinks to references within the chapter</li> <li>Minor editorial updates</li> </ul>

**Related Information**

[Quartus II Handbook Archive](#)



# Managing Metastability with the Quartus II Software 13

2014.12.15

QI15V1



Subscribe



Send Feedback

You can use the Quartus® II software to analyze the average mean time between failures (MTBF) due to metastability caused by synchronization of asynchronous signals, and optimize the design to improve the metastability MTBF.

All registers in digital devices, such as FPGAs, have defined signal-timing requirements that allow each register to correctly capture data at its input ports and produce an output signal. To ensure reliable operation, the input to a register must be stable for a minimum amount of time before the clock edge (register setup time or  $t_{SU}$ ) and a minimum amount of time after the clock edge (register hold time or  $t_H$ ). The register output is available after a specified clock-to-output delay ( $t_{CO}$ ).

If the data violates the setup or hold time requirements, the output of the register might go into a metastable state. In a metastable state, the voltage at the register output hovers at a value between the high and low states, which means the output transition to a defined high or low state is delayed beyond the specified  $t_{CO}$ . Different destination registers might capture different values for the metastable signal, which can cause the system to fail.

In synchronous systems, the input signals must always meet the register timing requirements, so that metastability does not occur. Metastability problems commonly occur when a signal is transferred between circuitry in unrelated or asynchronous clock domains, because the signal can arrive at any time relative to the destination clock.

The MTBF due to metastability is an estimate of the average time between instances when metastability could cause a design failure. A high MTBF (such as hundreds or thousands of years between metastability failures) indicates a more robust design. You should determine an acceptable target MTBF in the context of your entire system and taking in account that MTBF calculations are statistical estimates.

The metastability MTBF for a specific signal transfer, or all the transfers in a design, can be calculated using information about the design and the device characteristics. Improving the metastability MTBF for your design reduces the chance that signal transfers could cause metastability problems in your device.

The Quartus II software provides analysis, optimization, and reporting features to help manage metastability in Altera designs. These metastability features are supported only for designs constrained with the Quartus II Timing Analyzer. Both typical and worst-case MBTF values are generated for select device families.

## Related Information

- **Understanding Metastability in FPGAs**

For more information about metastability due to signal synchronization, its effects in FPGAs, and how MTBF is calculated

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



- **Reliability Report**  
For information about Altera device reliability

## Metastability Analysis in the Quartus II Software

When a signal transfers between circuitry in unrelated or asynchronous clock domains, the first register in the new clock domain acts as a synchronization register.

To minimize the failures due to metastability in asynchronous signal transfers, circuit designers typically use a sequence of registers (a synchronization register chain or synchronizer) in the destination clock domain to resynchronize the signal to the new clock domain and allow additional time for a potentially metastable signal to resolve to a known value. Designers commonly use two registers to synchronize a new signal, but a standard of three registers provides better metastability protection.

The timing analyzer can analyze and report the MTBF for each identified synchronizer that meets its timing requirements, and can generate an estimate of the overall design MTBF. The software uses this information to optimize the design MTBF, and you can use this information to determine whether your design requires longer synchronizer chains.

### Related Information

- **Metastability and MTBF Reporting** on page 13-4
- **MTBF Optimization** on page 13-7

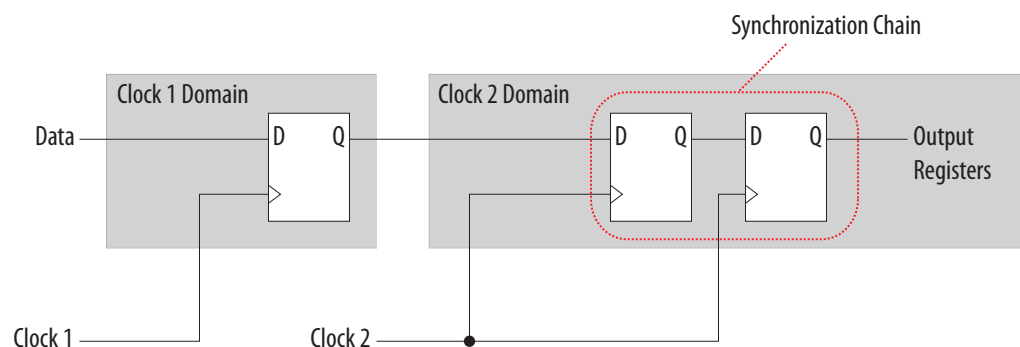
## Synchronization Register Chains

A synchronization register chain, or synchronizer, is defined as a sequence of registers that meets the following requirements:

- The registers in the chain are all clocked by the same clock or phase-related clocks.
- The first register in the chain is driven asynchronously or from an unrelated clock domain.
- Each register fans out to only one register, except the last register in the chain.

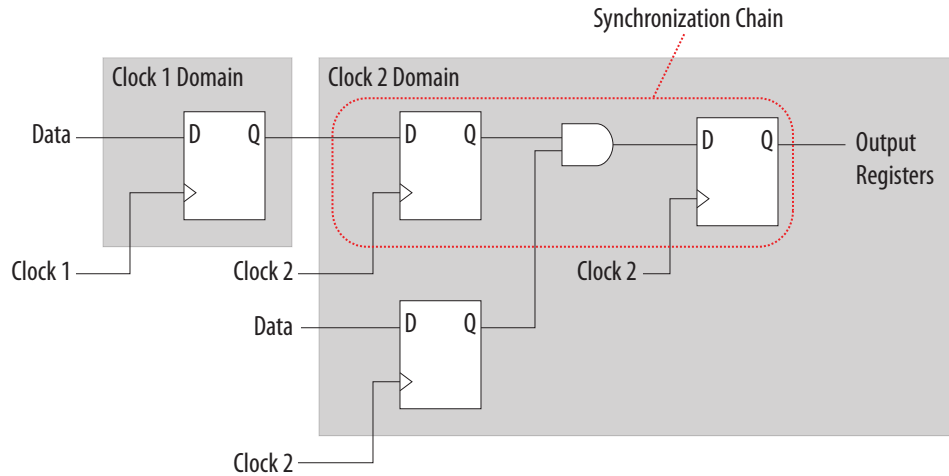
The length of the synchronization register chain is the number of registers in the synchronizing clock domain that meet the above requirements. The figure shows a sample two-register synchronization chain.

**Figure 13-1: Sample Synchronization Register Chain**



The path between synchronization registers can contain combinational logic as long as all registers of the synchronization register chain are in the same clock domain. The figure shows an example of a synchronization register chain that includes logic between the registers.

**Figure 13-2: Sample Synchronization Register Chain Containing Logic**



The timing slack available in the register-to-register paths of the synchronizer allows a metastable signal to settle, and is referred to as the available settling time. The available settling time in the MTBF calculation for a synchronizer is the sum of the output timing slacks for each register in the chain. Adding available settling time with additional synchronization registers improves the metastability MTBF.

**Related Information**

[How Timing Constraints Affect Synchronizer Identification and Metastability Analysis](#) on page 13-3

## Identifying Synchronizers for Metastability Analysis

The first step in enabling metastability MTBF analysis and optimization in the Quartus II software is to identify which registers are part of a synchronization register chain. You can apply synchronizer identification settings globally to automatically list possible synchronizers with the **Synchronizer identification** option on the **Timing Analyzer** page in the **Settings** dialog box.

Synchronization chains are already identified within most Altera intellectual property (IP) cores.

**Related Information**

[Identifying Synchronizers for Metastability](#)

For more information about how to enable metastability MTBF analysis and optimization in the Quartus II software, and more detailed descriptions of the synchronizer identification settings

## How Timing Constraints Affect Synchronizer Identification and Metastability Analysis

The timing analyzer can analyze metastability MTBF only if a synchronization chain meets its timing requirements. The metastability failure rate depends on the timing slack available in the synchronizer's register-to-register connections, because that slack is the available settling time for a potential metastable

signal. Therefore, you must ensure that your design is correctly constrained with the real application frequency requirements to get an accurate MTBF report.

In addition, the **Auto** and **Forced If Asynchronous** synchronizer identification options use timing constraints to automatically detect the synchronizer chains in the design. These options check for signal transfers between circuitry in unrelated or asynchronous clock domains, so clock domains must be related correctly with timing constraints.

The timing analyzer views input ports as asynchronous signals unless they are associated correctly with a clock domain. If an input port fans out to registers that are not acting as synchronization registers, apply a `set_input_delay` constraint to the input port; otherwise, the input register might be reported as a synchronization register. Constraining a synchronous input port with a `set_max_delay` constraint for a setup ( $t_{SU}$ ) requirement does not prevent synchronizer identification because the constraint does not associate the input port with a clock domain.

Instead, use the following command to specify an input setup requirement associated with a clock:

```
set_input_delay -max -clock <clock name> <latch - launch - tsu requirement> <input port name>
```

Registers that are at the end of false paths are also considered synchronization registers because false paths are not timing-analyzed. Because there are no timing requirements for these paths, the signal may change at any point, which may violate the  $t_{SU}$  and  $t_H$  of the register. Therefore, these registers are identified as synchronization registers. If these registers are not used for synchronization, you can turn off synchronizer identification and analysis. To do so, set **Synchronizer Identification** to **Off** for the first synchronization register in these register chains.

## Metastability and MTBF Reporting

The Quartus II software reports the metastability analysis results in the Compilation Report and Timing Analyzer reports.

The MTBF calculation uses timing and structural information about the design, silicon characteristics, and operating conditions, along with the data toggle rate.

If you change the **Synchronizer Identification** settings, you can generate new metastability reports by rerunning the timing analyzer. However, you should rerun the Fitter first so that the registers identified with the new setting can be optimized for metastability MTBF.

### Related Information

- [Metastability Reports](#) on page 13-4
- [MTBF Optimization](#) on page 13-7
- [Synchronizer Data Toggle Rate in MTBF Calculation](#) on page 13-6
- [Understanding Metastability in FPGAs](#)

For more information about how metastability MTBF is calculated

## Metastability Reports

Metastability reports provide summaries of the metastability analysis results. In addition to the MTBF Summary and Synchronizer Summary reports, the Timing Analyzer tool reports additional statistics in a report for each synchronizer chain.

**Note:** If the design uses only the **Auto Synchronizer Identification** setting, the reports list likely synchronizers but do not report MTBF. To obtain an MTBF for each register chain, force identification of synchronization registers.

**Note:** If the synchronizer chain does not meet its timing requirements, the reports list identified synchronizers but do not report MTBF. To obtain MTBF calculations, ensure that the design is properly constrained and that the synchronizer meets its timing requirements.

#### Related Information

- [Identifying Synchronizers for Metastability Analysis](#) on page 13-3
- [How Timing Constraints Affect Synchronizer Identification and Metastability Analysis](#) on page 13-3
- [Viewing Metastability Reports](#)  
For more information about how to access metastability reports in the Quartus II software

## MTBF Summary Report

The MTBF Summary reports an estimate of the overall robustness of cross-clock domain and asynchronous transfers in the design. This estimate uses the MTBF results of all synchronization chains in the design to calculate an MTBF for the entire design.

### Typical and Worst-Case MTBF of Design

The MTBF Summary Report shows the **Typical MTBF of Design** and the **Worst-Case MTBF of Design** for supported fully-characterized devices. The typical MTBF result assumes typical conditions, defined as nominal silicon characteristics for the selected device speed grade, as well as nominal operating conditions. The worst case MTBF result uses the worst case silicon characteristics for the selected device speed grade.

When you analyze multiple timing corners in the timing analyzer, the MTBF calculation may vary because of changes in the operating conditions, and the timing slack or available metastability settling time. Altera recommends running multi-corner timing analysis to ensure that you analyze the worst MTBF results, because the worst timing corner for MTBF does not necessarily match the worst corner for timing performance.

#### Related Information

[Timing Analyzer page](#)

### Synchronizer Chains

The MTBF Summary report also lists the **Number of Synchronizer Chains Found** and the length of the **Shortest Synchronizer Chain**, which can help you identify whether the report is based on accurate information.

If the number of synchronizer chains found is different from what you expect, or if the length of the shortest synchronizer chain is less than you expect, you might have to add or change **Synchronizer Identification** settings for the design. The report also provides the **Worst Case Available Settling Time**, defined as the available settling time for the synchronizer with the worst MTBF.

You can use the reported **Fraction of Chains for which MTBFs Could Not be Calculated** to determine whether a high proportion of chains are missing in the metastability analysis. A fraction of 1, for example, means that MTBF could not be calculated for any chains in the design. MTBF is not calculated if you have not identified the chain with the appropriate **Synchronizer identification** option, or if paths are not timing-analyzed and therefore have no valid slack for metastability analysis. You might have to correct your timing constraints to enable complete analysis of the applicable register chains.

## Increasing Available Settling Time

The MTBF Summary report specifies how an increase of 100ps in available settling time increases the MTBF values. If your MTBF is not satisfactory, this metric can help you determine how much extra slack would be required in your synchronizer chain to allow you to reach the desired design MTBF.

## Synchronizer Summary Report

The **Synchronizer Summary** lists the synchronization register chains detected in the design depending on the Synchronizer Identification setting.

The **Source Node** is the register or input port that is the source of the asynchronous transfer. The **Synchronization Node** is the first register of the synchronization chain. The **Source Clock** is the clock domain of the source node, and the **Synchronization Clock** is the clock domain of the synchronizer chain.

This summary reports the calculated **Worst-Case MTBF**, if available, and the **Typical MTBF**, for each appropriately identified synchronization register chain that meets its timing requirement.

### Related Information

[Synchronizer Chain Statistics Report in the Timing Analyzer](#) on page 13-6

## Synchronizer Chain Statistics Report in the Timing Analyzer

The timing analyzer provides an additional report for each synchronizer chain.

The **Chain Summary** tab matches the Synchronizer Summary information described in [Synchronizer Summary Report](#), while the **Statistics** tab adds more details, including whether the **Method of Synchronizer Identification** was **User Specified** (with the **Forced if Asynchronous** or **Forced** settings for the **Synchronizer Identification** setting), or **Automatic** (with the **Auto** setting). The **Number of Synchronization Registers in Chain** report provides information about the parameters that affect the MTBF calculation, including the **Available Settling Time** for the chain and the **Data Toggle Rate Used in MTBF Calculation**.

The following information is also included to help you locate the chain in your design:

- **Source Clock** and **Asynchronous Source** node of the signal.
- **Synchronization Clock** in the destination clock domain.
- Node names of the **Synchronization Registers** in the chain.

### Related Information

[Synchronizer Data Toggle Rate in MTBF Calculation](#) on page 13-6

## Synchronizer Data Toggle Rate in MTBF Calculation

The MTBF calculations assume the data being synchronized is switching at a toggle rate of 12.5% of the source clock frequency. That is, the arriving data is assumed to switch once every eight source clock cycles.

If multiple clocks apply, the highest frequency is used. If no source clocks can be determined, the data rate is taken as 12.5% of the synchronization clock frequency.

If you know an approximate rate at which the data changes, specify it with the **Synchronizer Toggle Rate** assignment in the Assignment Editor. You can also apply this assignment to an entity or the entire design. Set the data toggle rate, in number of transitions per second, on the first register of a synchronization chain. The timing analyzer takes the specified rate into account when computing the MTBF of that particular register chain. If a data signal never toggles and does not affect the reliability of the design, you

can set the **Synchronizer Toggle Rate** to 0 for the synchronization chain so the MTBF is not reported. To apply the assignment with Tcl, use the following command:

```
set_instance_assignment -name SYNCHRONIZER_TOGGLE_RATE <toggle rate in transitions/  
second> -to <register name>
```

In addition to **Synchronizer Toggle Rate**, there are two other assignments associated with toggle rates, which are not used for metastability MTBF calculations. The I/O Maximum Toggle Rate is only used for pins, and specifies the worst-case toggle rates used for signal integrity purposes. The Power Toggle Rate assignment is used to specify the expected time-averaged toggle rate, and is used by the PowerPlay Power Analyzer to estimate time-averaged power consumption.

## MTBF Optimization

In addition to reporting synchronization register chains and MTBF values found in the design, the Quartus II software can also protect these registers from optimizations that might negatively impact MTBF and can optimize the register placement and routing if the MTBF is too low.

Synchronization register chains must first be explicitly identified as synchronizers. Altera recommends that you set **Synchronizer Identification** to **Forced If Asynchronous** for all registers that are part of a synchronizer chain.

Optimization algorithms, such as register duplication and logic retiming in physical synthesis, are not performed on identified synchronization registers. The Fitter protects the number of synchronization registers specified by the **Synchronizer Register Chain Length** option.

In addition, the Fitter optimizes identified synchronizers for improved MTBF by placing and routing the registers to increase their output setup slack values. Adding slack in the synchronizer chain increases the available settling time for a potentially metastable signal, which improves the chance that the signal resolves to a known value, and exponentially increases the design MTBF. The Fitter optimizes the number of synchronization registers specified by the **Synchronizer Register Chain Length** option.

Metastability optimization is **on** by default. To view or change the **Optimize Design for Metastability** option, click **Assignments > Settings > Compiler Settings > Advanced Settings (Fitter)**. To turn the optimization on or off with Tcl, use the following command:

```
set_global_assignment -name OPTIMIZE_FOR_METASTABILITY <ON|OFF>
```

### Related Information

[Identifying Synchronizers for Metastability Analysis](#) on page 13-3

## Synchronization Register Chain Length

The **Synchronization Register Chain Length** option specifies how many registers should be protected from optimizations that might reduce MTBF for each register chain, and controls how many registers should be optimized to increase MTBF with the **Optimize Design for Metastability** option.

For example, if the **Synchronization Register Chain Length** option is set to 2, optimizations such as register duplication or logic retiming are prevented from being performed on the first two registers in all identified synchronization chains. The first two registers are also optimized to improve MTBF when the **Optimize Design for Metastability** option is turned on.

The default setting for the **Synchronization Register Chain Length** option is 2. The first register of a synchronization chain is always protected from operations that might reduce MTBF, but you should set the protection length to protect more of the synchronizer chain. Altera recommends that you set this option to the maximum length of synchronization chains you have in your design so that all synchronization registers are preserved and optimized for MTBF.

Click **Assignments > Settings > Compiler Settings > Advanced Settings (Synthesis)** to change the global **Synchronization Register Chain Length** option.

You can also set the **Synchronization Register Chain Length** on a node or an entity in the Assignment Editor. You can set this value on the first register in a synchronization chain to specify how many registers to protect and optimize in this chain. This individual setting is useful if you want to protect and optimize extra registers that you have created in a specific synchronization chain that has low MTBF, or optimize less registers for MTBF in a specific chain where the maximum frequency or timing performance is not being met. To make the global setting with Tcl, use the following command:

```
set_global_assignment -name SYNCHRONIZATION_REGISTER_CHAIN_LENGTH <number of registers>
```

To apply the assignment to a design instance or the first register in a specific chain with Tcl, use the following command:

```
set_instance_assignment -name SYNCHRONIZATION_REGISTER_CHAIN_LENGTH <number of registers> -to <register or instance name>
```

## Reducing Metastability Effects

You can check your design's metastability MTBF in the Metastability Summary report, and determine an acceptable target MTBF given the context of your entire system and the fact that MTBF calculations are statistical estimates. A high metastability MTBF (such as hundreds or thousands of years between metastability failures) indicates a more robust design.

This section provides guidelines to ensure complete and accurate metastability analysis, and some suggestions to follow if the Quartus II metastability reports calculate an unacceptable MTBF value. The Timing Optimization Advisor (available from the Tools menu) gives similar suggestions in the Metastability Optimization section.

### Related Information

[Metastability Reports](#) on page 13-4

## Apply Complete System-Centric Timing Constraints for the Timing Analyzer

To enable the Quartus II metastability features, make sure that the timing analyzer is used for timing analysis.

Ensure that the design is fully timing constrained and that it meets its timing requirements. If the synchronization chain does not meet its timing requirements, MTBF cannot be calculated. If the clock domain constraints are set up incorrectly, the signal transfers between circuitry in unrelated or asynchronous clock domains might be identified incorrectly.

Use industry-standard system-centric I/O timing constraints instead of using FPGA-centric timing constraints.



You should use `set_input_delay` constraints in place of `set_max_delay` constraints to associate each input port with a clock domain to help eliminate false positives during synchronization register identification.

#### Related Information

[How Timing Constraints Affect Synchronizer Identification and Metastability Analysis](#) on page 13-3

## Force the Identification of Synchronization Registers

Use the guidelines in [Identifying Synchronizers for Metastability Analysis](#) to ensure the software reports and optimizes the appropriate register chains.

You should identify synchronization registers with the **Synchronizer Identification** set to **Forced If Asynchronous** in the Assignment Editor. If there are any registers that the software detects as synchronous but you want to be analyzed for metastability, apply the **Forced** setting to the first synchronizing register. Set **Synchronizer Identification** to **Off** for registers that are not synchronizers for asynchronous signals or unrelated clock domains.

To help you find the synchronizers in your design, you can set the global **Synchronizer Identification** setting on the **Timing Analyzer** page of the **Settings** dialog box to **Auto** to generate a list of all the possible synchronization chains in your design.

#### Related Information

[Identifying Synchronizers for Metastability Analysis](#) on page 13-3

## Set the Synchronizer Data Toggle Rate

The MTBF calculations assume the data being synchronized is switching at a toggle rate of 12.5% of the source clock frequency.

To obtain a more accurate MTBF for a specific chain or all chains in your design, set the **Synchronizer Toggle Rate**.

#### Related Information

[Synchronizer Data Toggle Rate in MTBF Calculation](#) on page 13-6

## Optimize Metastability During Fitting

Ensure that the **Optimize Design for Metastability** setting is turned on.

#### Related Information

[MTBF Optimization](#) on page 13-7

## Increase the Length of Synchronizers to Protect and Optimize

Increase the Synchronizer Chain Length parameter to the maximum length of synchronization chains in your design. If you have synchronization chains longer than 2 identified in your design, you can protect the entire synchronization chain from operations that might reduce MTBF and allow metastability optimizations to improve the MTBF.

#### Related Information

[Synchronization Register Chain Length](#) on page 13-7

## Set Fitter Effort to Standard Fit instead of Auto Fit

If your design MTBF is too low after following the other guidelines, you can try increasing the Fitter effort to perform more metastability optimization. The default **Auto Fit** setting reduces the Fitter's effort after meeting the design's timing and routing requirements to reduce compilation time.

This effort reduction can result in less metastability optimization if the timing requirements are easy to meet. If **Auto Fit** reduces the Fitter's effort during your design compilation, setting the Fitter effort to **Standard Fit** might improve the design's MTBF results. To modify the **Fitter Effort**, click **Assignments > Settings > Compiler Settings > Advanced Settings (Fitter)**.

## Increase the Number of Stages Used in Synchronizers

Designers commonly use two registers in a synchronization chain to minimize the occurrence of metastable events, and a standard of three registers provides better metastability protection. However, synchronization chains with two or even three registers may not be enough to produce a high enough MTBF when the design runs at high clock and data frequencies.

If a synchronization chain is reported to have a low MTBF, consider adding an additional register stage to your synchronization chain. This additional stage increases the settling time of the synchronization chain, allowing more opportunity for the signal to resolve to a known state during a metastable event. Additional settling time increases the MTBF of the chain and improves the robustness of your design. However, adding a synchronization stage introduces an additional stage of latency on the signal.

If you use the Altera FIFO IP core with separate read and write clocks to cross clock domains, increase the metastability protection (and latency) for better MTBF. In the DCFIFO parameter editor, choose the **Best metastability protection, best fmax, unsynchronized clocks** option to add three or more synchronization stages. You can increase the number of stages to more than three using the **How many sync stages?** setting.

## Select a Faster Speed Grade Device

The design MTBF depends on process parameters of the device used. Faster devices are less susceptible to metastability issues. If the design MTBF falls significantly below the target MTBF, switching to a faster speed grade can improve the MTBF substantially.

## Scripting Support

You can run procedures and make settings described in this chapter in a Tcl script. You can also run some procedures at a command prompt. For detailed information about scripting command options, refer to the Quartus II Command-Line and Tcl API Help browser.

To run the Help browser, type the following command at the command prompt:

```
quartus_sh --qhelp r
```

### Related Information

- [Tcl Scripting](#)  
For more information about Tcl scripting
- [Quartus II Settings File Reference Manual](#)  
For more information about settings and constraints in the Quartus II software

- [Command-Line Scripting](#)  
For more information about command-line scripting
- [About Quartus II Scripting](#)  
For more information about command-line scripting

## Identifying Synchronizers for Metastability Analysis

To apply the global Synchronizer Identification assignment, use the following command:

```
set_global_assignment -name SYNCHRONIZER_IDENTIFICATION <OFF|AUTO|"FORCED IF ASYNCHRONOUS">
```

To apply the **Synchronizer Identification** assignment to a specific register or instance, use the following command:

```
set_instance_assignment -name SYNCHRONIZER_IDENTIFICATION <AUTO|"FORCED IF ASYNCHRONOUS"|FORCED|OFF> -to <register or instance name>
```

## Synchronizer Data Toggle Rate in MTBF Calculation

To specify a toggle rate for MTBF calculations as described on page [Synchronizer Data Toggle Rate in MTBF Calculation](#), use the following command:

```
set_instance_assignment -name SYNCHRONIZER_TOGGLE_RATE <toggle rate in transitions/second> -to <register name>
```

### Related Information

[Synchronizer Data Toggle Rate in MTBF Calculation](#) on page 13-6

## report\_metastability and Tcl Command

If you use a command-line or scripting flow, you can generate the metastability analysis reports described in [Metastability Reports](#) outside of the Quartus II and user interfaces.

The table describes the options for the `report_metastability` and Tcl command.

**Table 13-1: report\_metastability Command Options**

Option	Description
-append	If output is sent to a file, this option appends the result to that file. Otherwise, the file is overwritten.
-file <name>	Sends the results to an ASCII or HTML file. The extension specified in the file name determines the file type — either <b>*.txt</b> or <b>*.html</b> .
-panel_name <name>	Sends the results to the panel and specifies the name of the new panel.
-stdout	Indicates the report be sent to the standard output, via messages. This option is required only if you have selected another output format, such as a file, and would also like to receive messages.

**Related Information**[Metastability Reports](#) on page 13-4

## MTBF Optimization

To ensure that metastability optimization described on page [MTBF Optimization](#) is turned on (or to turn it off), use the following command:

```
set_global_assignment -name OPTIMIZE_FOR_METASTABILITY <ON|OFF>
```

**Related Information**[MTBF Optimization](#) on page 13-7

## Synchronization Register Chain Length

To globally set the number of registers in a synchronization chain to be protected and optimized as described on page [Synchronization Register Chain Length](#), use the following command:

```
set_global_assignment -name SYNCHRONIZATION_REGISTER_CHAIN_LENGTH <number of registers>
```

To apply the assignment to a design instance or the first register in a specific chain, use the following command:

```
set_instance_assignment -name SYNCHRONIZATION_REGISTER_CHAIN_LENGTH <number of registers> -to <register or instance name>
```

**Related Information**[Synchronization Register Chain Length](#) on page 13-7

## Managing Metastability

Altera's Quartus II software provides industry-leading analysis and optimization features to help you manage metastability in your FPGA designs. Set up your Quartus II project with the appropriate constraints and settings to enable the software to analyze, report, and optimize the design MTBF. Take advantage of these features in the Quartus II software to make your design more robust with respect to metastability.

## Document Revision History

**Table 13-2: Document Revision History**

Date	Version	Changes
2014.12.15	14.1.0	Updated location of Fitter Settings, Analysis & Synthesis Settings, and Physical Optimization Settings to Compiler Settings.
June 2014	14.0.0	Updated formatting.

Date	Version	Changes
June 2012	12.0.0	Removed survey link.
November 2011	10.0.2	Template update.
December 2010	10.0.1	Changed to new document template.
July 2010	10.0.0	Technical edit.
November 2009	9.1.0	Clarified description of synchronizer identification settings. Minor changes to text and figures throughout document.
March 2009	9.0.0	Initial release.

**Related Information**

[Quartus II Handbook Archive](#)

For previous versions of the Quartus II Handbook

# Best Practices for Incremental Compilation Partitions and Floorplan Assignments 14

2014.12.15

QI15V1



Subscribe



Send Feedback

## About Incremental Compilation and Floorplan Assignments

This manual provides guidelines to help you partition your design to take advantage of Quartus® II incremental compilation, and to help you create a design floorplan using LogicLock™ regions when they are recommended to support the compilation flow.

The Quartus II incremental compilation feature allows you to partition a design, compile partitions separately, and reuse results for unchanged partitions. Incremental compilation provides the following benefits:

- Reduces compilation times by an average of 75% for large design changes
- Preserves performance for unchanged design blocks
- Provides repeatable results and reduces the number of compilations
- Enables team-based design flows

### Related Information

[Quartus II Incremental Compilation for Hierarchical and Team-Based Design documentation](#) on page 3-1

[Incremental Compilation online help](#)

## Incremental Compilation Overview

Quartus II incremental compilation is an optional compilation flow that enhances the default Quartus II compilation. If you do not partition your design for incremental compilation, your design is compiled using the default “flat” compilation flow.

To prepare your design for incremental compilation, you first determine which logical hierarchy boundaries should be defined as separate partitions in your design, and ensure your design hierarchy and source code is set up to support this partitioning. You can then create design partition assignments in the Quartus II software to specify which hierarchy blocks are compiled independently as partitions (including empty partitions for missing or incomplete logic blocks).

During compilation, Quartus II Analysis & Synthesis and the Fitter create separate netlists for each partition. Netlists are internal post-synthesis and post-fit database representations of your design.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered

In subsequent compilations, you can select which netlist to preserve for each partition. You can either reuse the synthesis or fitting netlist, or instruct the Quartus II software to resynthesize the source files. You can also use compilation results exported from another Quartus II project.

When you make changes to your design, the Quartus II software recompiles only the designated partitions and merges the new compilation results with existing netlists for other partitions, according to the degree of results preservation you set with the netlist for each design partition.

In some cases, Altera recommends that you create a design floorplan with placement assignments to constrain parts of the design to specific regions of the device.

You must use the partial reconfiguration (PR) feature in conjunction with incremental compilation for Stratix® V device families. Partial reconfiguration allows you to reconfigure a portion of the FPGA dynamically, while the remainder of the device continues to operate as intended.

#### Related Information

- [Introduction to Design Floorplans](#) on page 14-36
- [Using the Incremental Compilation Design Flow online help](#)  
Step-by-step information about using incremental compilation to recompile only changed parts of your design
- [Design Planning for Partial Reconfiguration documentation](#) on page 4-1

The Partial Reconfiguration (PR) feature in the Quartus II software allows you to reconfigure a portion of the FPGA dynamically, while the remainder of the device continues to operate. The Quartus II software supports the PR feature for the Altera® Stratix® V device family.

## Recommendations for the Netlist Type

For subsequent compilations, you specify which post-compilation netlist you want to use with the netlist type for each partition.

Use the following general guidelines to set the netlist type for each partition:

- **Source File**—Use this setting to resynthesize the source code (with any new assignments, and replace any previous synthesis or Fitter results).
  - If you modify the design source, the software automatically resynthesizes the partitions with the appropriate netlist type, which makes the **Source File** setting optional in this case.
  - Most assignments do not trigger an automatic recompilation, so you must set the netlist type to **Source File** to compile the source files with new assignments or constraints that affect synthesis.
- **Post-Synthesis** (default)—Use this setting to re-fit the design (with any new Fitter assignments), but preserve the synthesis results when the source files have not changed. If it is difficult to meet the required timing performance, you can use this setting to allow the Fitter the most flexibility in placement and routing. This setting does not reduce compilation time as much as the **Post-Fit** setting or preserve timing performance from the previous compilation.
- **Post-Fit**—Use this setting to preserve Fitter and performance results when the source files have not changed. This setting reduces compilation time the most, and preserves timing performance from the previous compilation.
- **Post-Fit with Fitter Preservation Level set to Placement**—Use the **Advanced Fitter Preservation Level** setting on the **Advanced** tab in the **Design Partition Properties** dialog box to allow more flexibility and find the best routing for all partitions given their placement.

The Quartus II software Rapid Recompile feature instructs the Compiler to reuse the compatible compilation results if most of the design has not changed since the last compilation. This feature reduces compilation time and preserves performance when there are small and isolated design changes within a partition, and works with all netlist type settings. With this feature, you do not have control over which parts of the design are recompiled; the Compiler determines which parts of the design must be recompiled.

## Design Flows Using Incremental Compilation

The Quartus II incremental compilation feature supports various design flows. Your design flow affects design optimization and the amount of design planning required to obtain optimal results.

### Using Standard Flow

In the standard incremental compilation flow, the top-level design is divided into partitions, which can be compiled and optimized together in one Quartus II project. If another team member or IP provider is developing source code for the top-level design, they can functionally verify their partition independently, and then simply provide the partition's source code to the project lead for integration into the top-level design. If the project lead wants to compile the top-level design when source code is not yet complete for a partition, they can create an empty placeholder for the partition until the code is ready to be added to the top-level design.

Compiling all design partitions in a single Quartus II project ensures that all design logic is compiled with a consistent set of assignments, and allows the software to perform global placement and routing optimizations. Compiling all design logic together is beneficial for FPGA design flows because all parts of the design must use the same shared set of device resources. Therefore, it is often easier to ensure good quality of results when partitions are developed within a single top-level Quartus II project.

### Using Team-Based Flow

In the team-based incremental compilation flow, you can design and optimize partitions by accessing the top-level project from a shared source control system or creating copies of the top-level Quartus II project framework. As development continues, designers export their partition so that the post-synthesis netlist or post-fitting results can be integrated into the top-level design.

### Using Third-Party IP Delivery Flow

If required for third-party IP delivery, or in cases where designers cannot access a shared or copied top-level project framework, you can create and compile a design partition logic in isolation and export a partition that is included in the top-level project. If this type of design flow is necessary, planning and rigorous design guidelines might be required to ensure that designers have a consistent view of project assignments and resource allocations. Therefore, developing partitions in completely separate Quartus II projects can be more challenging than having all source code within one project or developing design partitions within the same top-level project framework.

### Combining Design Flows

You can also combine design flows and use exported partitions only when it is necessary to support your design environment. For example, if the top-level design includes one or more design blocks that will be optimized by remote designers or IP providers, you can integrate those blocks into the reserved partitions in the top-level design when the code is complete, but also have other partitions that will be developed within the top-level design.



If any partitions are developed independently, the project lead must ensure that top-level constraints (such as timing constraints, any relevant floorplan or pin assignments, and optimization settings) are consistent with those used by all designers.

## Project Management in Team-Based Design Flows

If possible, each team member should work within the same top-level project framework. Using the same project framework amongst team members ensures that designers have the settings and constraints needed for their partition and allows designers to analyze how their design block interacts with other partitions in the top-level design.

### Using a Source Control System

In a team-based environment where designers have access to the project through source control software, each designer can use project files as read-only and develop their partition within the source control system. As designers check in their completed partitions, other team members can see how their partitions interact.

### Using a Copy of the Top-Level Project

If designers do not have access to a source control system, the project lead can provide each designer with a copy of the top-level project framework to use as they develop their partitions. In both cases, each designer exports their completed design as a partition, and then the project lead integrates the partition into the top-level design. The project lead can choose to use only the post-synthesis netlist and rerun placement and routing, or to use the post-fitting results to preserve the placement and routing results from the other designer's projects. Using post-synthesis partitions gives the Fitter the most flexibility and is likely to achieve a good result for all partitions, but if one partition has difficulty meeting timing, the designer can choose to preserve their successful fitting results.

### Using a Separate Project

Alternatively, designers can use their own Quartus II project for their independent design block. You might use this design flow if a designer, such as a third-party IP provider, does not have access to the entire top-level project framework. In this case, each designer must create their own project with all the relevant assignments and constraints. This type of design flow requires more planning and rigorous design guidelines. If the project lead plans to incorporate the post-fitting compilation results for the partition, this design flow requires especially careful planning to avoid resource conflicts.

### Using Scripts

The project lead also has the option to generate design partition scripts to manage resource and timing budgets in the top-level design when partitions are developed outside the top-level project framework. Scripts make it easier for designers of independent Quartus II projects to follow instructions from the project lead. The Quartus II design partition scripts feature creates Tcl scripts or **.tcl** files and makefiles that an independent designer can run to set up an independent Quartus II project.

#### Related Information

[Generating Design Partition Scripts for Project Management online help](#)

### Using Constraints

If designers create Quartus II assignments or timing constraints for their partitions, they must ensure that the constraints are integrated into the top-level design. If partition designers use the same top-level project framework (and design hierarchy), the constraints or Synopsys Design Constraints File (**.sdc**) can

be easily copied or included in the top-level design. If partition designers use a separate Quartus II project with a different design hierarchy, they must ensure that constraints are applied to the appropriate level of hierarchy in the top-level design, and design the `.sdc` for easy delivery to the project lead.

#### Related Information

[Including SDC Constraints from Lower-Level Partitions for Third-Party IP Delivery](#) on page 14-32

[Quartus II Incremental Compilation for Hierarchical and Team-Based Design documentation](#) on page 3-1

Information about the different types of incremental design flows and example applications, as well as documented restrictions and limitations

## Why Plan Partitions and Floorplan Assignments?

Incremental design flows typically require more planning than flat compilations, and require you to be more rigorous about following good design practices. For example, you might need to structure your source code or design hierarchy to ensure that logic is grouped correctly for optimization. It is easier to implement the correct logic grouping early in the design cycle than to restructure the code later.

Planning involves setting up the design logic for partitioning and may also involve planning placement assignments to create a floorplan. Not all design flows require floorplan assignments. If you decide to add floorplan assignments later, when the design is close to completion, well-planned partitions make floorplan creation easier. Poor partition or floorplan assignments can worsen design area utilization and performance and make timing closure more difficult.

As FPGA devices get larger and more complex, following good design practices become more important for all design flows. Adhering to recommended synchronous design practices makes designs more robust and easier to debug. Using an incremental compilation flow adds additional steps and requirements to your project, but can provide significant benefits in design productivity by preserving the performance of critical blocks and reducing compilation time.

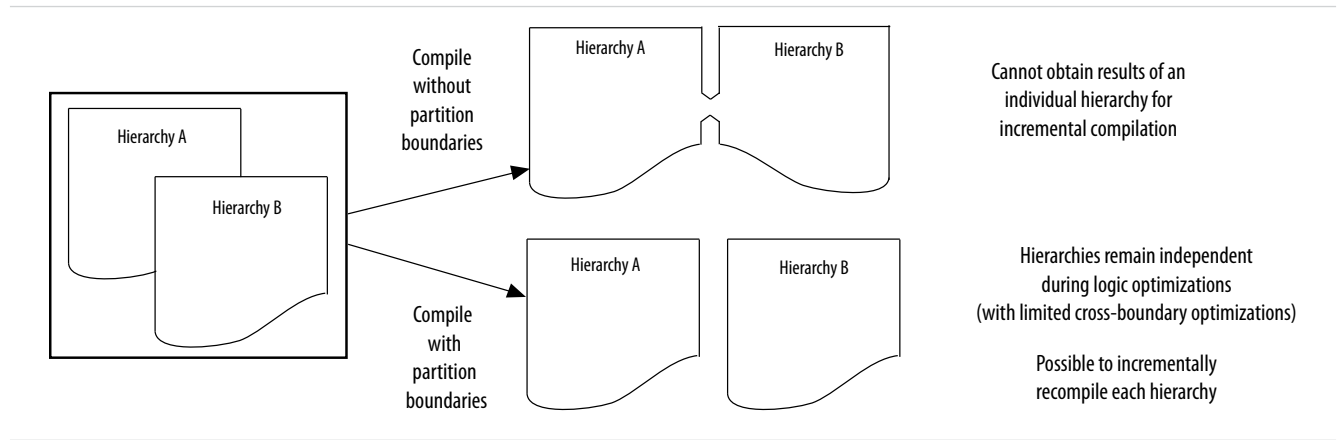
#### Related Information

[Introduction to Design Floorplans](#) on page 14-36

## Partition Boundaries and Optimization

The logical hierarchical boundaries between partitions are treated as hard boundaries for logic optimization (except for some limited cross-boundary optimizations) to allow the software to size and place each partition independently. The figure shows the effects of partition boundaries during logic optimization.

Figure 14-1: Effects of Partition Boundaries During Logic Optimization



## Merging Partitions

You can use the **Merge** command in the Design Partitions window to combine hierarchical partitions into a single partition, as long as they share the same immediate parent partition. Merging partitions allows additional optimizations for partition I/O ports that connect between or feed more than one of the merged hierarchical design blocks.

When partitions are placed together, the Fitter can perform placement optimizations on the design as a whole to optimize the placement of cross-boundary paths. However, the Fitter can never perform logic optimizations such as physical synthesis across the partition boundary. If partitions are fit separately in different projects, or if some partitions use previous post-fitting results, the Fitter does not place and route the entire cross-boundary path at the same time and cannot fully optimize placement across the partition boundaries. Good design partitions can be placed independently because cross-partition paths are not the critical timing paths in the design.

## Resource Utilization

There are possible timing performance utilization effects due to partitioning and creating a floorplan. Not all designs encounter these issues, but you should consider these effects if a flat version of your design is very close to meeting its timing requirements, or is close to using all the device resources, before adding partition or floorplan assignments:

- Partitions can increase resource utilization due to cross-boundary optimization limitations if the design does not follow partitioning guidelines. Floorplan assignments can also increase resource utilization because regions can lead to unused logic. If your device is full with the flat version of your design, you can focus on creating partitions and floorplan assignments for timing-critical or often-changing blocks to benefit most from incremental compilation.
- Partitions and floorplan assignments might increase routing utilization compared to a flat design. If long compilation times are due to routing congestion, you might not be able to use the incremental flow to reduce compilation time. Review the Fitter messages to check how much time is spent during routing optimizations to determine the percentage of routing utilization. When routing is difficult, you can use incremental compilation to lock the routing for routing-critical blocks only (with other partitions empty), and then compile the rest of the design after the critical blocks meet their requirements.
- Partitions can reduce timing performance in some cases because of the optimization and resource effects described above, causing longer logic delays. Floorplan assignments restrict logic placement, which can make it more difficult for the Fitter to meet timing requirements. Use the guidelines in this manual to reduce any effect on your design performance.

#### Related Information

- [Design Partition Guidelines](#) on page 14-9
- [Checking Floorplan Quality](#) on page 14-44

## Turning On Supported Cross-Boundary Optimizations

You can improve the optimizations performed between design partitions by turning on the cross-boundary optimizations feature. You can select the optimizations as individual assignments for each partition. This allows the cross-boundary optimization feature to give you more control over the optimizations that work best for your design.

You can turn on the cross-boundary optimizations for your design partitions on the **Advanced** tab of the **Design Partition Properties** dialog box. Once you change the optimization settings, the Quartus II software recompiles your partition from source automatically. Cross-boundary optimizations include the following: propagate constants, propagate inversions on partition inputs, merge inputs fed by a common source, merge electrically equivalent bidirectional pins, absorb internal paths, and remove logic connected to dangling outputs.

Cross-boundary optimizations are implemented top-down from the parent partition into the child partition, but not vice-versa. The cross-boundary optimization feature cannot be used with partitions with multiple personas (partial reconfiguration partitions).

Although more partitions allow for a greater reduction in compilation time, consider limiting the number of partitions to prevent degradation in the quality of results. Creating good design partitions and good floorplan location assignments helps to improve the design resource utilization and timing performance results for cross-partition paths.

#### Related Information

[Design Partition Properties Dialog Box online help](#)

# Guidelines for Incremental Compilation

## General Partitioning Guidelines

The first step in planning your design partitions is to organize your source code so that it supports good partition assignments. Although you can assign any hierarchical block of your design as a design partition or merge hierarchical blocks into the same partition, following the design guidelines presented below ensures better results.

## Plan Design Hierarchy and Design Files

You begin the partitioning process by planning the design hierarchy. When you assign a hierarchical instance as a design partition, the partition includes the assigned instance and entities instantiated below that are not defined as separate partitions. You can use the **Merge** command in the Design Partitions window to combine hierarchical partitions into a single partition, as long as they have the same immediate parent partition.

- When planning your design hierarchy, keep logic in the “leaves” of the hierarchy instead of having logic at the top-level of the design so that you can isolate partitions if required.
- Create entities that can form partitions of approximately equal size. For example, do not instantiate small entities at the same hierarchy level, because it is more difficult to group them to form reasonably-sized partitions.
- Create each entity in an independent file. The Quartus II software uses a file checksum to detect changes, and automatically recompiles a partition if its source file changes and its netlist type is set to either post-synthesis or post-fit. If the design entities for two partitions are defined in the same file, changes to the logic in one partition initiates recompilation for both partitions.
- Design dependencies also affect which partitions are compiled when a source file changes. If two partitions rely on the same lower-level entity definition, changes in that lower-level entity affect both partitions. Commands such as VHDL `use` and Verilog HDL `include` create dependencies between files, so that changes to one file can trigger recompilations in all dependent files. Avoid these types of file dependencies if possible. The Partition Dependent Files report for each partition in the Analysis & Synthesis section of the Compilation report lists which files contribute to each partition.

## Using Partitions with Third-Party Synthesis Tools

Incremental compilation works well with third-party synthesis tools in addition to Quartus II Integrated Synthesis. If you use a third-party synthesis tool, set up your tool to create a separate Verilog Quartus Mapping File (**.vqm**) or EDIF Input File (**.edf**) netlist for each hierarchical partition. In the Quartus II software, designate the top-level entity from each netlist as a design partition. The **.vqm** or **.edf** netlist file is treated as the source file for the partition in the Quartus II software.

### Related Information

- [Quartus II Incremental Compilation for Hierarchical and Team-Based Design documentation](#) on page 3-1

## Partition Design by Functionality and Block Size

Initially, you should partition your design along functional boundaries. In a top-level system block diagram, each block is often a natural design partition. Typically, each block of a system is relatively independent and has more signal interaction internally than interaction between blocks, which helps

reduce optimizations between partition boundaries. Keeping functional blocks together means that synthesis and fitting can optimize related logic as a whole, which can lead to improved optimization.

- Consider how many partitions you want to maintain in your design to determine the size of each partition. Your compilation time reduction goal is also a factor, because compiling small partitions is typically faster than compiling large partitions.
- There is no minimum size for partitions; however, having too many partitions can reduce the quality of results by limiting optimization. Ensure that the design partitions are not too small. As a general guideline, each partition should contain more than approximately 2,000 logic elements (LEs) or adaptive logic modules (ALMs). If your design is incomplete when you partition the design, use previous designs to help estimate the size of each block.

### Partition Design by Clock Domain and Timing Criticality

Consider which clock in your design feeds the logic in each partition. If possible, keep clock domains within one partition. When a clock signal is isolated to one partition, it reduces dependence on other partitions for timing optimization. Isolating a clock domain to one partition also allows better use of regional clock routing networks if the partition logic is constrained to one region of the design. Additionally, limiting the number of clocks within each partition simplifies the timing requirements for each partition during optimization. Use an appropriate subsystem to implement the required logic for any clock domain transfers (such as a synchronization circuit, dual-port RAM, or FIFO). You can include this logic inside the partition at one side of the transfer.

Try to isolate timing-critical logic from logic that you expect to easily meet timing requirements. Doing so allows you to preserve the satisfactory results for non-critical partitions and focus optimization iterations on only the timing-critical portions of the design to minimize compilation time.

#### Related Information

##### [Analyzing and Optimizing the Design Floorplan with the Chip Planner documentation](#)

Information about clock domains and their affect on partition design

### Consider What Is Changing

When assigning partitions, you should consider what is changing in the design. Is there intellectual property (IP) or reused logic for which the source code will not change during future design iterations? If so, define the logic in its own partition so that you can compile one time and immediately preserve the results and not have to compile that part of the design again. Is logic being tuned or optimized, or are specifications changing for part of the design? If so, define changing logic in its own partition so that you can recompile only the changing part while the rest of the design remains unchanged.

As a general rule, create partitions to isolate logic that will change from logic that will not change. Partitioning a design in this way maximizes the preservation of unchanged logic and minimizes compilation time.

### Design Partition Guidelines

Follow the design partition guidelines below when you create or modify the HDL code for each design block that you might want to assign as a design partition. You do not need to follow all the recommendations exactly to achieve a good quality of results with the incremental compilation flow, but adhering to as many as possible maximizes your chances for success.


The design partition guidelines include examples of the types of optimizations that are prevented by partition boundaries, and describes how you can structure or modify your partitions to avoid these limitations.

## Register Partition Inputs and Outputs

Use registers at partition input and output connections that are potentially timing-critical. Registers minimize the delays on inter-partition paths and prevent the need for cross-boundary optimizations.

If every partition boundary has a register as shown in the figure, every register-to-register timing path between partitions includes only routing delay. Therefore, the timing paths between partitions are likely not timing-critical, and the Fitter can generally place each partition independently from other partitions. This advantage makes it easier to create floorplan location assignments for each separate partition, and is especially important for flows in which partitions are placed independently in separate Quartus II projects. Additionally, the partition boundary does not affect combinational logic optimization because each register-to-register logic path is contained within a single partition.

**Figure 14-2: Registering Partition I/O**



If a design cannot include both input and output registers for each partition due to latency or resource utilization concerns, choose to register one end of each connection. If you register every partition output, for example, the combinational logic that occurs in each cross-partition path is included in one partition so that it can be optimized together.

It is a good synchronous design practice to include registers for every output of a design block. Registered outputs ensure that the input timing performance for each design block is controlled exclusively within the destination logic block.

### Related Information

- [Partition Statistics Report](#) on page 14-31
- [Incremental Compilation Advisor](#) on page 14-28


## Minimize Cross-Partition-Boundary I/O

Minimize the number of I/O paths that cross between partition boundaries to keep logic paths within a single partition for optimization. Doing so makes partitions more independent for both logic and placement optimization.

This guideline is most important for timing-critical and high-speed connections between partitions, especially in cases where the input and output of each partition is not registered. Slow connections that are not timing-critical are acceptable because they should not impact the overall timing performance of the design. If there are timing-critical paths between partitions, rework or merge the partitions to avoid these inter-partition paths.

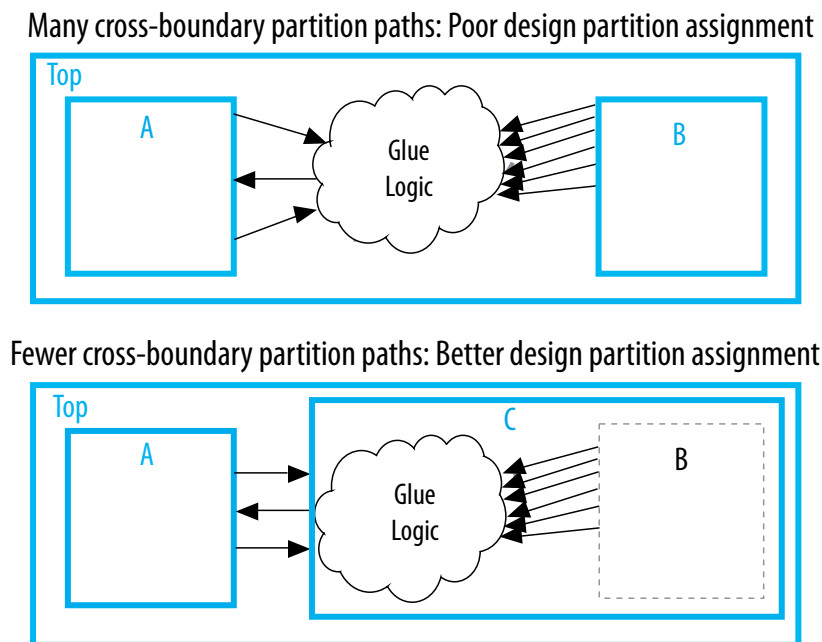
When dividing your design into partitions, consider the types of functions at the partition boundaries. The figure shows an expansive function with more outputs than inputs in the left diagram, which makes a poor partition boundary, and, on the right side, a better place to assign the partition boundary that minimizes cross-partition I/Os. Adding registers to one or both sides of the cross-partition path in this example would further improve partition quality.

**Figure 14-3: Minimizing I/O Between Partitions by Moving the Partition Boundary**



Another way to minimize connections between partitions is to avoid using combinational "glue logic" between partitions. You can often move the logic to the partition at one end of the connection to keep more logic paths within one partition. For example, the bottom diagram includes a new level of hierarchy C defined as a partition instead of block B. Clearly, there are fewer I/O connections between partitions A and C than between partitions A and B.

**Figure 14-4: Minimizing I/O between Partitions by Modifying Glue Logic**





**Related Information**

- [Partition Statistics Report](#) on page 14-31
- [Incremental Compilation Advisor](#) on page 14-28

**Examine the Need for Logic Optimization Across Partitions**

Partition boundaries prevent logic optimizations across partitions (except for some limited cross-boundary optimizations).

In some cases, especially if part of the design is complete or comes from another designer, the designer might not have followed these guidelines when the source code was created. These guidelines are not mandatory to implement an incremental compilation flow, but can improve the quality of results. If assigning a partition affects resource utilization or timing performance of a design block as compared to the flat design, it might be due to one of the issues described in the logic optimization across partitions guidelines below. Many of the examples suggest simple changes to your partition definitions or hierarchy to move the partition boundary to improve your results.

The following guidelines ensure that your design does not require logic optimization across partition boundaries:

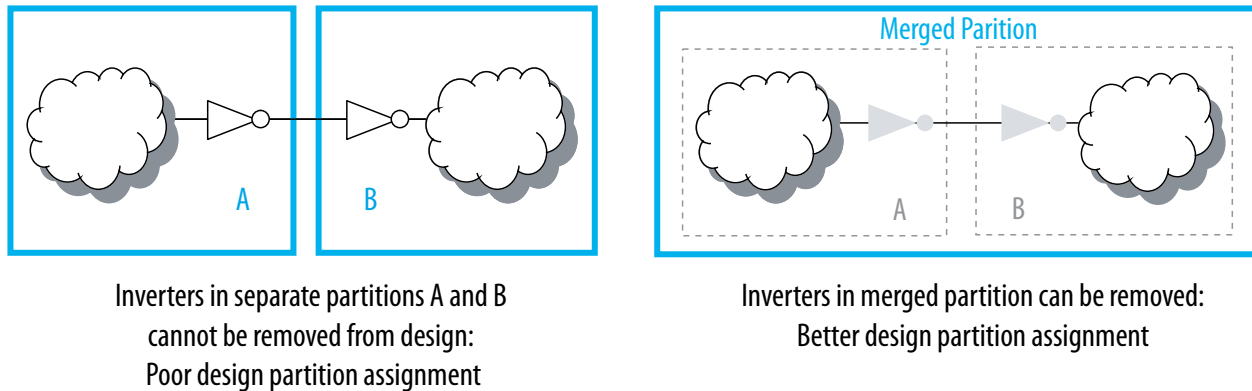
**Keep Logic in the Same Partition for Optimization and Merging**

If your design logic requires logic optimization or merging to obtain optimal results, ensure that all the logic is part of the same partition because only limited cross-boundary optimizations are permitted.

***Example—Combinational Logic Path***

If a combinational logic path is split across two partitions, the logic cannot be optimized or merged into one logic cell in the device. This effect can result in an extra logic cell in the path, increasing the logic delay. As a very simple example, consider two inverters on the same signal in two different partitions, A and B, as shown in the left diagram of the figure. To maintain correct incremental functionality, these two inverters cannot be removed from the design during optimization because they occur in different design partitions. The Quartus II software cannot use information about other partitions when it compiles each partition, because each partition is allowed to change independently from the other.

On the right side of the figure, partitions A and B are merged to group the logic in blocks A and B into one partition. If the two blocks A and B are not under the same immediate parent partition, you can create a wrapper file to define a new level of hierarchy that contains both blocks, and set this new hierarchy block as the partition. With the logic contained in one partition, the software can optimize the logic and remove the two inverters (shown in gray), which reduces the delay for that logic path. Removing two inverters is not a significant reduction in resource utilization because inversion logic is readily available in Altera device architecture. However, this example is a simple demonstration of the types of logic optimization that are prevented by partition boundaries.

**Figure 14-5: Keeping Logic in the Same Partition for Optimization****Example—Fitter Merging**

In a flat design, the Fitter can also merge logical instantiations into the same physical device resource. With incremental compilation, logic defined in different partitions cannot be merged to use the same physical device resource.

For example, the Fitter can merge two single-port RAMs from a design into one dedicated RAM block in the device. If the two RAMs are defined in different partitions, the Fitter cannot merge them into one dedicated device RAM block.

This limitation is a only a concern if merging is required to fit the design in the target device. Therefore, you are more likely to encounter this issue during troubleshooting rather than during planning, if your design uses more logic than is available in the device.

**Merging PLLs and Transceivers (GXB)**

Multiple instances of the ALTPLL IP core can use the same PLL resource on the device. Similarly, GXB transceiver instances can share high-speed serial interface (HSSI) resources in the same quad as other instances. The Fitter can merge multiple instantiations of these blocks into the same device resource, even if it requires optimization across partitions. Therefore, there are no restrictions for PLLs and high-speed transceiver blocks when setting up partitions.

**Keep Constants in the Same Partition as Logic**

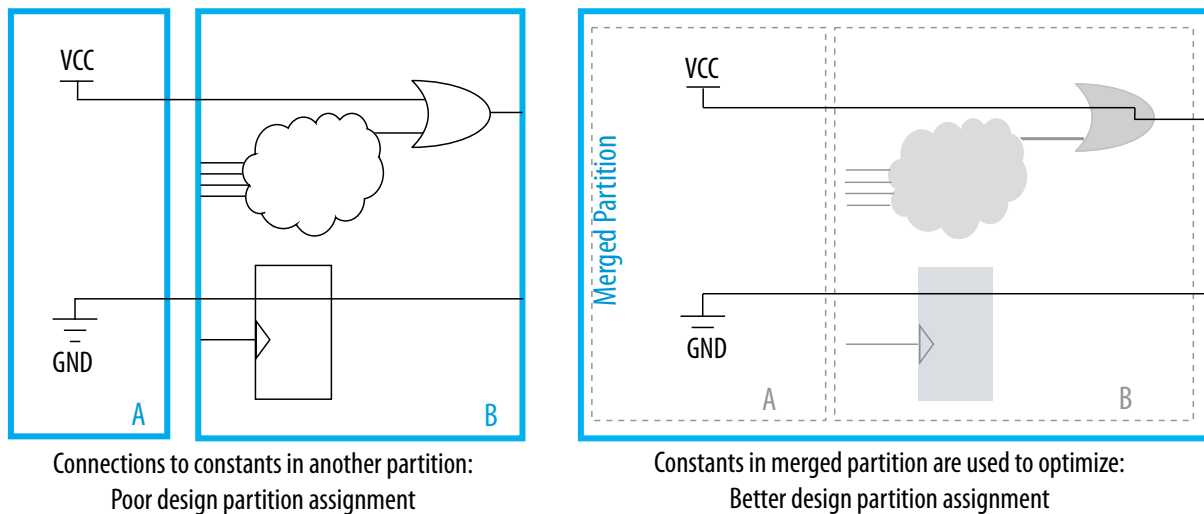
Because the Quartus II software cannot fully optimize across a partition boundary, constants are not propagated across partition boundaries, except from parent partition to child partition. A signal that is constant ( $1/V_{CC}$  or  $0/GND$ ) in one partition cannot affect another partition.

**Example—Constants in Merged Partitions**

For example, the left diagram of the figure shows part of a design in which partition A defines some signals as constants (and assumes that the other input connections come from elsewhere in the design and are not shown in the figure). Constants such as these could appear due to parameter or generic settings or configurations with parameters, setting a bus to a specific set of values, or could result from optimizations that occur within a group of logic. Because the blocks are independent, the software cannot optimize the logic in block B based on the information from block A. The right side of the figure shows a merged partition that groups the logic in blocks A and B. If the two blocks A and B are not under the same immediate parent partition, you can create a wrapper file to define a new level of hierarchy that contains both blocks, and set this new hierarchical block as the partition.

Within the single merged partition, the Quartus II software can use the constants to optimize and remove much of the logic in block B (shown in gray), as shown in the figure.

**Figure 14-6: Keeping Constants in the Same Partition as the Logic They Feed**



#### Related Information

- [Partition Statistics Report](#) on page 14-31
- [Incremental Compilation Advisor](#) on page 14-28

### Avoid Signals That Drive Multiple Partition I/O or Connect I/O Together

Do not use the same signal to drive multiple ports of a single partition or directly connect two ports of a partition. If the same signal drives multiple ports of a partition, or if two ports of a partition are directly connected, those ports are logically equivalent. However, the software has limited information about connections made in another partition (including the top-level partition), the compilation cannot take advantage of the equivalence. This restriction usually produces sub-optimal results.

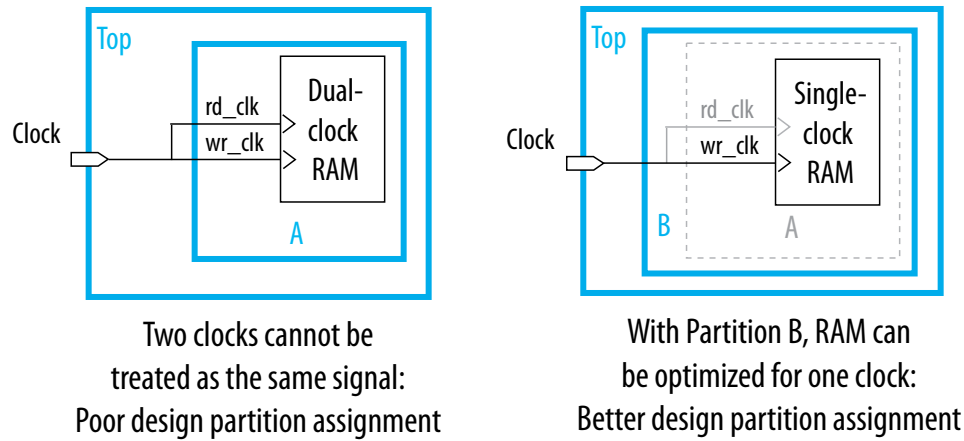
If your design has these types of connections, redefine the partition boundaries to remove the affected ports. If one signal from a higher-level partition feeds two input ports of the same partition, feed the one signal into the partition, and then make the two connections within the partition. If an output port drives an input port of the same partition, the connection can be made internally without going through any I/O ports. If an input port drives an output port directly, the connection can likely be implemented without the ports in the lower-level partition by connecting the signals in a higher-level partition.

#### Example—Single Signal Driving More Than One Port

The figure shows an example of one signal driving more than one port. The left diagram shows a design where a single clock signal is used to drive both the read and write clocks of a RAM block. Because the RAM block is compiled as a separate partition A, the RAM block is implemented as though there are two unique clocks. If you know that the port connectivity will not change (that is, the ports will always be driven by the same signal in the top-level partition), redefine the port interface so that there is only a single port that can drive both connections inside the partition. You can create a wrapper file to define a partition that has fewer ports, as shown in the diagram on the right side. With the single clock fed into the partition, the RAM can be optimized into a single-clock RAM instead of a dual-clock RAM. Single-clock RAM can provide better performance in the device architecture. Additionally, partition A might use two

global routing lines for the two copies of the clock signal. Partition B can use one global line that fans out to all destinations. Using just the single port connection prevents overuse of global routing resources.

**Figure 14-7: Preventing One Signal from Driving Multiple Partition Inputs**



**Related Information**

[Incremental Compilation Advisor](#) on page 14-28

**Invert Clocks in Destination Partitions**

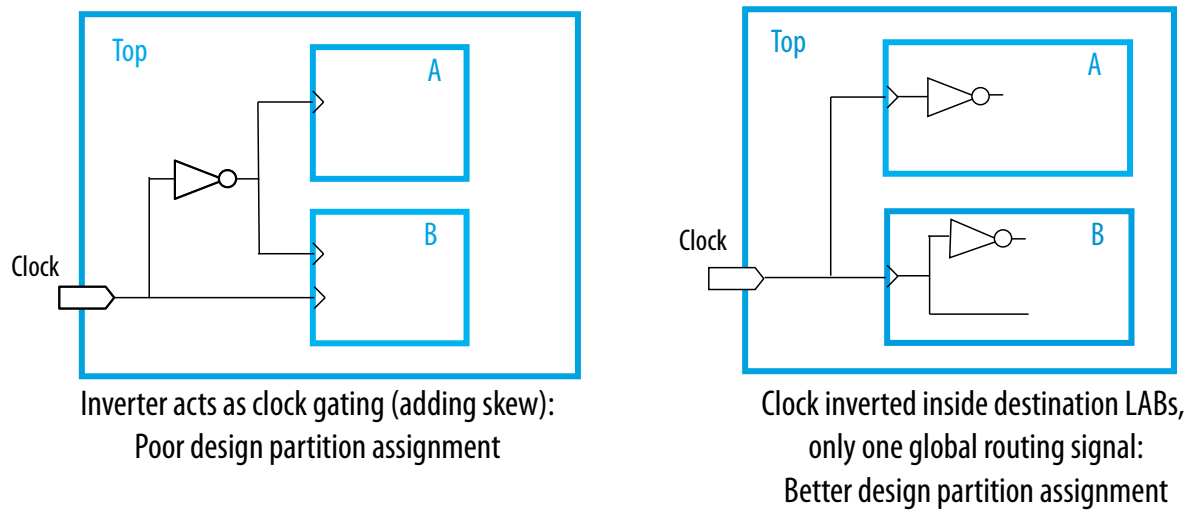
For best results, clock inversion should be performed in the destination logic array block (LAB) because each LAB contains clock inversion circuitry in the device architecture. In a flat compilation, the Quartus II software can optimize a clock inversion to propagate it to the destination LABs regardless of where the inversion takes place in the design hierarchy. However, clock inversion cannot propagate through a partition boundary (except from a parent partition to a child partition) to take advantage of the inversion architecture in the destination LABs.

**Example—Clock Signal Inversion**

With partition boundaries as shown in the left diagram of the figure, the Quartus II software uses logic to invert the signal in the partition that defines the inversion (the top-level partition in this example), and then routes the signal on a global clock resource to its destinations (in partitions A and B). The inverted clock acts as a gated clock with high skew. A better solution is to invert the clock signal in the destination partitions as shown on the right side of the diagram. In this case, the correct logic and routing resources can be used, and the signal does not behave like a gated clock.

The figure shows the clock signal inversion in the destination partitions.

Figure 14-8: Inverting Clock Signal in Destination Partitions



Notice that this diagram also shows another example of a single pin feeding two ports of a partition boundary. In the left diagram, partition B does not have the information that the clock and inverted clock come from the same source. In the right diagram, partition B has more information to help optimize the design because the clock is connected as one port of the partition.

### Connect I/O Pin Directly to I/O Register for Packing Across Partition Boundaries

The Quartus II software allows cross-partition register packing of I/O registers in certain cases where your input and output pins are defined in the top-level hierarchy (and the top-level partition), but the corresponding I/O registers are defined in other partitions.

Input pin cross-partition register packing requires the following specific circumstances:

- The input pin feeds exactly one register.
- The path between the input pin and register includes only input ports of partitions that have one fan-out each.

Output pin cross-partition register packing requires the following specific circumstances:


- The register feeds exactly one output pin.
- The output pin is fed by only one signal.
- The path between the register and output pin includes only output ports of partitions that have one fan-out each.

The following examples of I/O register packing illustrate this point using Block Design File (.bdf) schematics to describe the design logic.

#### Example 1—Output Register in Partition Feeding Multiple Output Pins

In this example, the subdesign contains a single register.


Figure 14-9: Subdesign with One Register, Designated as a Separate Partition



If the top-level design instantiates the subdesign with a single fan-out directly feeding an output pin, and designates the subdesign as a separate design partition, the Quartus II software can perform cross-partition register packing because the single partition port feeds the output pin directly.

In this example, the top-level design instantiates the subdesign as an output register with more than one fan-out signal.

Figure 14-10: Top-Level Design Instantiating the Subdesign with Two Output Pins



In this case, the Quartus II software does not perform output register packing. If there is a **Fast Output Register** assignment on pin `out`, the software issues a warning that the Fitter cannot pack the node to an I/O pin because the node and the I/O cell are connected across a design partition boundary.

This type of cross-partition register packing is not allowed because it requires modification to the interface of the subdesign partition. To perform incremental compilation, the Quartus II software must preserve the interface of design partitions.

To allow the Quartus II software to pack the register in the subdesign with the output pin `out` in the figure, restructure your HDL code so that output registers directly connect to output pins by making one of the following changes:

- Place the register in the same partition as the output pin. The simplest method is to move the register from the subdesign partition into the partition containing the output pin. Doing so guarantees that the Fitter can optimize the two nodes without violating partition boundaries.
- Duplicate the register in your subdesign HDL so that each register feeds only one pin, and then connect the extra output pin to the new port in the top-level design. Doing so converts the cross-partition register packing into the simplest case where each register has a single fan-out.

Figure 14-11: Modified Subdesign with Two Output Registers and Two Output Ports


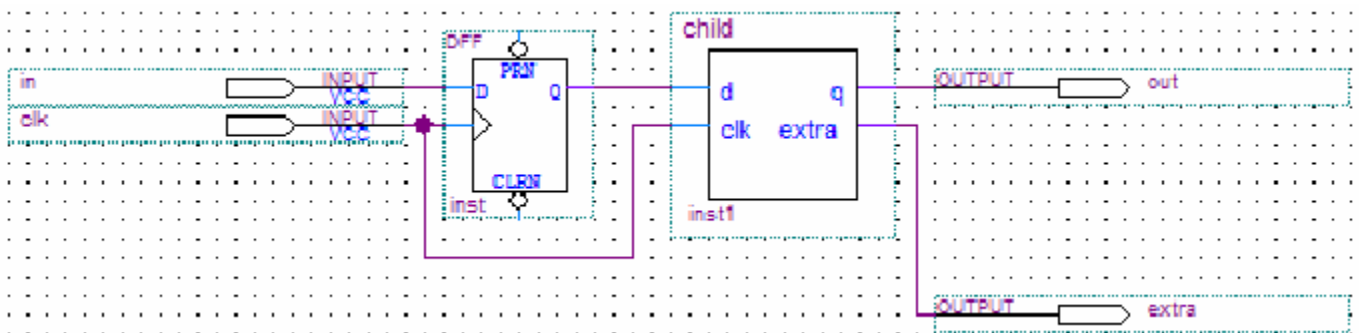


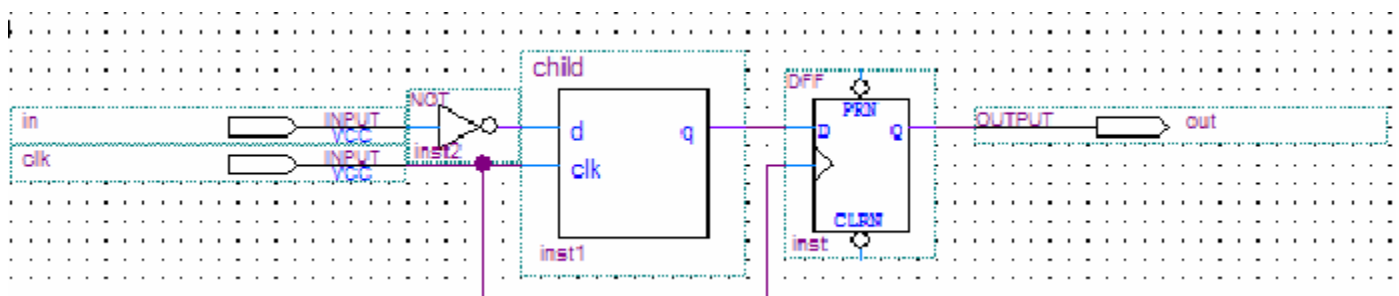
Figure 14-12: Modified Top-Level Design Connecting Two Output Ports to Output Pins




### Example 2—Input Register in Partition Fed by an Inverted Input Pin or Output Register in Partition Feeding an Inverted Output Pin

In this example, a subdesign designated as a separate partition contains a register. The top-level design in the figure instantiates the subdesign as an input register with the input pin inverted. The top-level design instantiates the subdesign as an output register with the signal inverted before feeding an output pin.

Figure 14-13: Top-Level Design Instantiating Subdesign as an Input Register with an Inverted Input Pin



**Figure 14-14: Top-Level Design Instantiating the Subdesign as an Output Register Feeding an Inverted Output Pin**



In these cases, the Quartus II software does not perform register packing. If there is a **Fast Input Register** assignment on pin `in`, as shown in the top figure, or a **Fast Output Register** assignment on pin `out`, as shown in the bottom figure, the Quartus II software issues a warning that the Fitter cannot pack the node to an I/O pin because the node and I/O cell are connected across a design partition boundary.

This type of register packing is not allowed because it requires moving logic across a design partition boundary to place into a single I/O device atom. To perform register packing, either the register must be moved out of the subdesign partition, or the inverter must be moved into the subdesign partition to be implemented in the register.

To allow the Quartus II software to pack the single register in the subdesign with the input pin `in`, as shown in top figure or the output pin `out`, as shown in the bottom figure, restructure your HDL code to place the register in the same partition as the inverter by making one of the following changes:

- Move the register from the subdesign partition into the top-level partition containing the pin. Doing so ensures that the Fitter can optimize the I/O register and inverter without violating partition boundaries.
- Move the inverter from the top-level block into the subdesign, and then connect the subdesign directly to a pin in the top-level design. Doing so allows the Fitter to optimize the inverter into the register implementation, so that the register is directly connected to a pin, which enables register packing.


### Do Not Use Internal Tri-States

Internal tri-state signals are not recommended for FPGAs because the device architecture does not include internal tri-state logic. If designs use internal tri-states in a flat design, the tri-state logic is usually converted to OR gates or multiplexing logic. If tri-state logic occurs on a hierarchical partition boundary, the Quartus II software cannot convert the logic to combinational gates because the partition could be connected to a top-level device I/O through another partition.

The figures below show a design with partitions that are not supported for incremental compilation due to the internal tri-state output logic on the partition boundaries. Instead of using internal tri-state logic for partition outputs, implement the correct logic to select between the two signals. Doing so is good practice even when there are no partitions, because such logic explicitly defines the behavior for the internal signals instead of relying on the Quartus II software to convert the tri-state signals into logic.

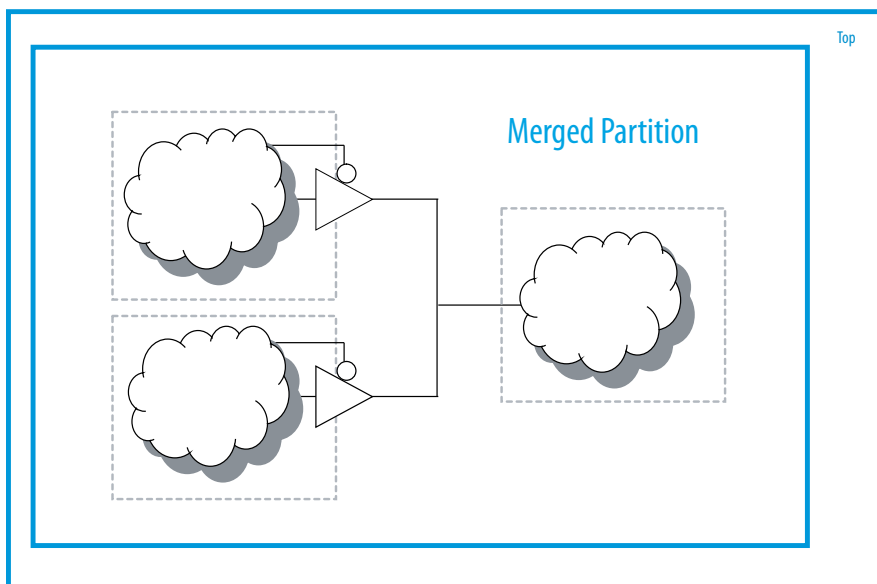


Figure 14-15: Unsupported Internal Tri-State Signals



Design results in Quartus II error message:  
The software cannot synthesize this  
design and maintain incremental functionality.

Figure 14-16: Merged Partition Allows Synthesis to Convert Internal Tri-State Logic to Combinational Logic



Merged partition allows synthesis to  
convert tri-state logic into  
combinational logic.

Do not use tri-state signals or bidirectional ports on hierarchical partition boundaries, unless the port is connected directly to a top-level I/O pin on the device. If you must use internal tri-state logic, ensure that all the control and destination logic is contained in the same partition, in which case the Quartus II software can convert the internal tri-state signals into combinational logic as in a flat design. In this example, you can also merge all three partitions into one partition, as shown in the bottom figure, to allow the Quartus II software to treat the logic as internal tri-state and perform the same type of optimization as a flat design. If possible, you should avoid using internal

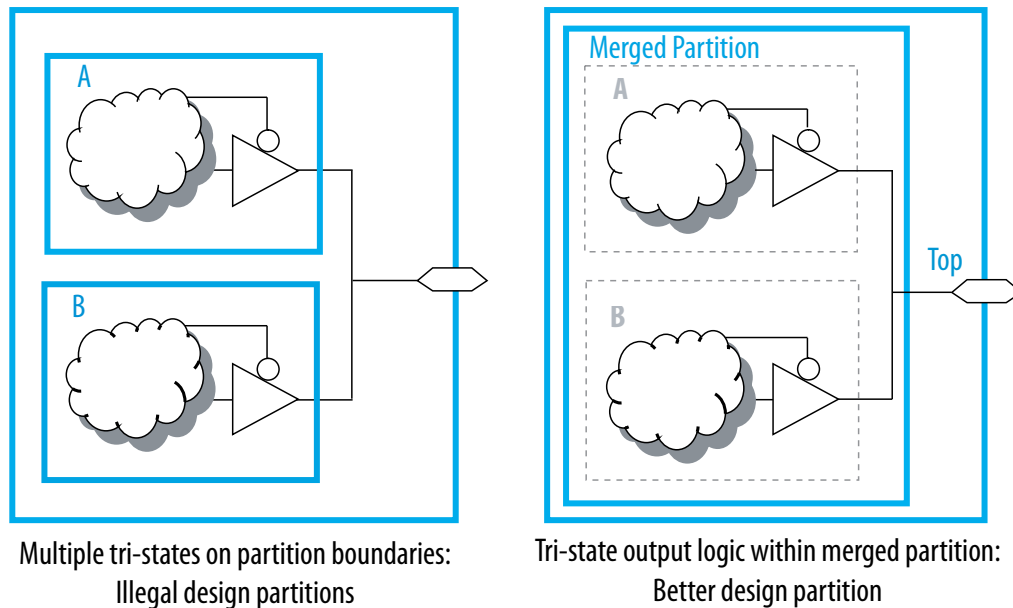
tri-state logic in any Altera FPGA design to ensure that you get the desired implementation when the design is compiled for the target device architecture.

### Include All Tri-State and Enable Logic in the Same Partition

When multiple output signals use tri-state logic to drive a device output pin, the Quartus II software merges the logic into one tri-state output pin. The Quartus II software cannot merge tri-state outputs into one output pin if any of the tri-state logic occurs on a partition boundary. Similarly, output pins with an output enable signal cannot be packed into the device I/O cell if the output enable logic is part of a different partition from the output register. To allow register packing for output pins with an output enable signal, structure your HDL code or design partition assignments so that the register and enable logic are defined in the same partition.

The figure shows a design with tri-state output signals that feed a device bidirectional I/O pin (assuming that the input connection feeds elsewhere in the design and is not shown in the figure). In the left diagram below, the tri-state output signals appear as the outputs of two separate partitions. In this case, the Quartus II software cannot implement the specified logic and maintain incremental functionality. In the right diagram, partitions A and B are merged to group the logic from the two blocks. With this single partition, the Quartus II software can merge the two tri-state output signals and implement them in the tri-state logic available in the device I/O element.

**Figure 14-17: Including All Tri-State Output Logic in the Same Partition**



### Summary of Guidelines Related to Logic Optimization Across Partitions

To ensure that your design does not require logic optimization across partitions, follow the guidelines below:

- Include logic in the same partition for optimization and merging
- Include constants in the same partition as logic
- Avoid signals that drive multiple partition I/O or connect I/O together
- Invert clocks in destination partitions

- Connect I/O directly to I/O register for packing across partition boundaries
- Do not use internal tri-states
- Include all tri-state and enable logic in the same partition

Remember that these guidelines are not mandatory when implementing an incremental compilation flow, but can improve the quality of results. When creating source design code, follow these guidelines and organize your HDL code to support good partition boundaries. For designs that are complete, assess whether assigning a partition affects the resource utilization or timing performance of a design block as compared to the flat design. Make the appropriate changes to your design or hierarchy, or merge partitions as required, to improve your results.

## Consider a Cascaded Reset Structure

Designs typically have a global asynchronous reset signal where a top-level signal feeds all partitions. To minimize skew for the high fan-out signal, the global reset signal is typically placed onto a global routing resource.

In some cases, having one global reset signal can lead to recovery and removal time problems. This issue is not specific to incremental flows; it could be applicable in any large high-speed design. In an incremental flow, the global reset signal creates a timing dependency between the top-level partition and lower-level partitions.

For incremental compilation, it is helpful to minimize the impact of global structures. To isolate each partition, consider adding reset synchronizers. Using cascaded reset structures, the intent is to reduce the inter-partition fan-out of the reset signal, thereby minimizing the effect of the global signal. Reducing the fan-out of the global reset signal also provides more flexibility in routing the cascaded signals, and might help recovery and removal times in some cases.

This recommendation can help in large designs, regardless of whether you are using incremental compilation. However, if one global signal can feed all the logic in its domain and meet recovery and removal times, this recommendation may not be applicable for your design. Minimizing global structures is more relevant for

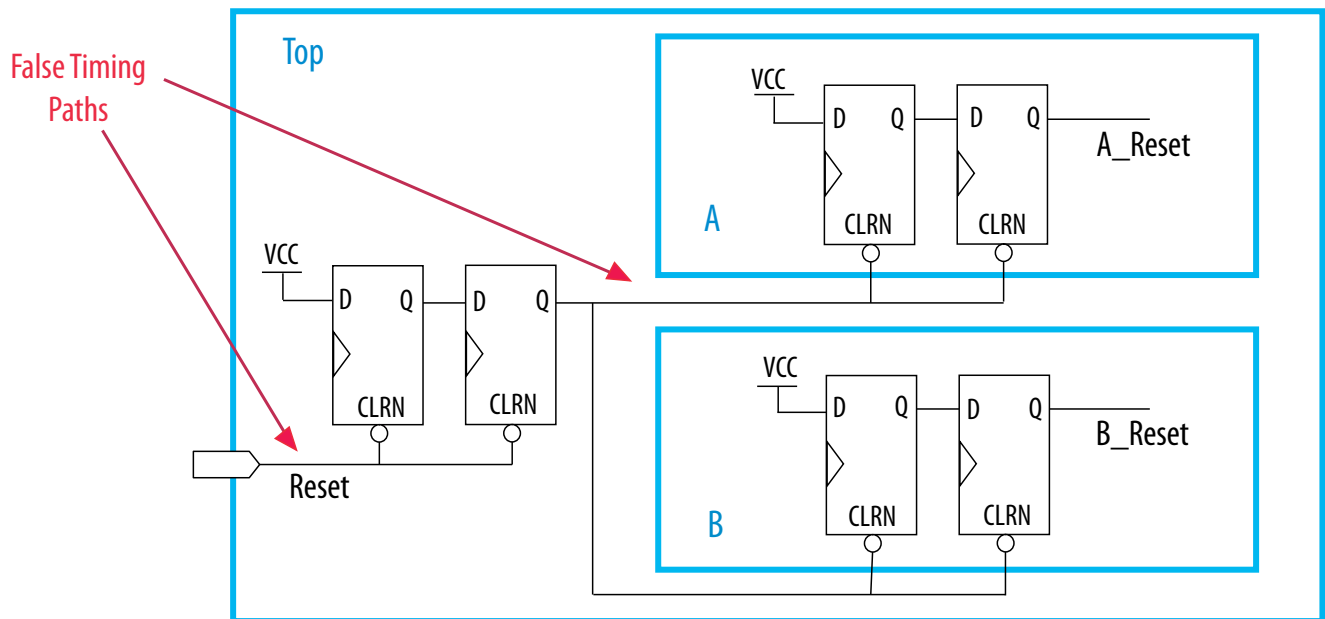
high-performance designs where meeting timing on the reset logic can be challenging. Isolating each partition and allowing more flexibility in global routing structures is an additional advantage in incremental flows.

If you add additional reset synchronizers to your design, latency is also added to the reset path, so ensure that this is acceptable in your design. Additionally, parts of the design may come out of the reset state in different clock cycles. You can balance the latency or add hand-shaking logic between partitions, if necessary, to accommodate these differences.

The signal is first synchronized on the chip following good synchronous design practices, meaning that the design asynchronously resets, but synchronously releases from reset to avoid any race conditions or metastability problems. Then, to minimize the impact of global structures, the circuit employs a divide-and-conquer approach for the reset structure. By implementing a cascaded reset structure, the reset paths for each partition are independent. This structure reduces the effect of inter-partition dependency because the inter-partition reset signals can now be treated as false paths for timing analysis. In some cases, the reset signal of the partition can be placed on local lines to reduce the delay added by routing to a global routing line. In other cases, the signal can be routed on a regional or quadrant clock signal.

The figure shows a cascaded reset structure.

Figure 14-18: Cascaded Reset Structure



This circuit design can help you achieve timing closure and partition independence for your global reset signal. Evaluate the circuit and consider how it works for your design.

**Related Information**

- [Recommended Design Practices documentation](#) on page 11-1  
Information and design recommendations for reset structures

## Design Partition Guidelines for Third-Party IP Delivery

There are additional design guidelines that can improve incremental compilation flows where exported partitions are developed independently. These guidelines are not always required, but are usually recommended if the design includes partitions compiled in a separate Quartus II project, such as when delivering intellectual property (IP). A unique challenge of IP delivery for FPGAs is the fact that the partitions developed independently must share a common set of resources. To minimize issues that might arise from sharing a common set of resources, you can design partitions within a single Quartus II project, or a copy of the top-level design. A common project ensures that designers have a consistent view of the top-level design framework.

Alternatively, an IP designer can export just the post-synthesis results to be integrated in the top-level design when the post-fitting results from the IP project are not required. Using a post-synthesis netlist provides more flexibility to the Quartus II Fitter, so that less resource allocation is required. If a common project is not possible, especially when the project lead plans to integrate the IP's post-fitting results, it is important that the project lead and IP designer clearly communicate their requirements.

**Related Information**

- [Project Management in Team-Based Design Flows](#) on page 14-4

## Allocate Logic Resources

In an incremental compilation design flow in which designers, such as third-party IP providers, optimize partitions and then export them to a top-level design, the Quartus II software places and routes each partition separately. In some cases, partitions can use conflicting resources when combined at the top level. Allocation of logic resources requires that you decide on a set of logic resources (including I/O, LAB logic blocks, RAM and DSP blocks) that the IP block will "own". This process can be interactive; the project lead and the IP designer might work together to determine what resources are required for the IP block and are available in the top-level design.

You can constrain logic utilization for the IP core using design floorplan location assignments. The design should specify I/O pin locations with pin assignments.

You can also specify limits for Quartus II synthesis to allocate and balance resources. This procedure can also help if device resources are overused in the individual partitions during synthesis.

In the standard synthesis flow, the Quartus II software can perform automated resource balancing for DSP blocks or RAM blocks and convert some of the logic into regular logic cells to prevent overuse.

You can use the Quartus II synthesis options to control inference of IP cores that use the DSP, or RAM blocks. You can also use the IP Catalog and Parameter Editor to customize your RAM or DSP IP cores to use regular logic instead of the dedicated hardware blocks.

### Related Information

- [Introduction to Design Floorplans](#) on page 14-36
- [Quartus II Integrated Synthesis documentation](#) on page 16-1  
Information about resource balancing DSP and RAM blocks when using Quartus II synthesis
- [Timing Closure and Optimization documentation](#)  
Tips about resource balancing and reducing resource utilization
- [More Analysis Synthesis Settings Dialog Box online help](#)  
Information about how to set global logic options for partitions

## Allocate Global Routing Signals and Clock Networks if Required

In most cases, you do not have to allocate global routing signals because the Quartus II software finds the best solution for the global signals. However, if your design is complex and has multiple clocks, especially for a partition developed by a third-party IP designer, you may have to allocate global routing resources between various partitions.

Global routing signals can cause conflicts when independent partitions are integrated into a top-level design. The Quartus II software automatically promotes high fan-out signals to use global routing resources available in the device. Third-party partitions can use the same global routing resources, thus causing conflicts in the top-level design. Additionally, LAB placement depends on whether the inputs to the logic cells within the LAB use a global clock signal. Problems can occur if a design does not use a global signal in a lower-level partition, but does use a global signal in the top-level design.

If the exported IP core is small, you can reduce the potential for problems by using constraints to promote clock and high fan-out signals to regional routing signals that cover only part of the device, instead of global routing signals. In this case, the Quartus II software is likely to find a routing solution in the top-level design because there are many regional routing signals available on most Altera devices, and designs do not typically overuse regional resources.

To ensure that an IP block can utilize a regional clock signal, view the resource coverage of regional clocks in the Chip Planner, and then align LogicLock regions that constrain partition placement with available global clock routing resources. For example, if the LogicLock region for a particular partition is limited to one device quadrant, that partition's clock can use a regional clock routing type that covers only one device quadrant. When all partition logic is available, the project lead can compile the entire design at the top level with floorplan assignments to allow the use of regional clocks that span only a part of the device.

If global resources are heavily used in the overall design, or the IP designer requires global clocks for their partition, you can set up constraints to avoid signal overuse at the top-level by assigning the appropriate type of global signals or setting a maximum number of clock signals for the partition.

You can use the **Global Signal** assignment to force or prevent the use of a global routing line, making the assignment to a clock source node or signal. You can also assign certain types of global clock resources in some device families, such as regional clocks. For example, if you have an IP core, such as a memory interface that specifies the use of a dual regional clock, you can constrain the IP to part of the device covered by a regional clock and change the **Global Signal** assignment to use a regional clock. This type of assignment can reduce clocking congestion and conflicts.

Alternatively, partition designers can specify the number of clocks allowed in the project using the maximum clocks allowed options in the **Advanced Settings (Fitter)** dialog box. Specify **Maximum number of clocks of any type allowed**, or use the **Maximum number of global clocks allowed**, **Maximum number of regional clocks allowed**, and **Maximum number of periphery clocks allowed** options to restrict the number of clock resources of a particular type in your design.

If you require more control when planning a design with integrated partitions, you can assign a specific signal to use a particular clock network in newer device families by assigning the clock control block instance called CLKCTRL. You can make a point-to-point assignment from a clock source node to a destination node, or a single-point assignment to a clock source node with the **Global Clock CLKCTRL Location** logic option. Set the assignment value to the name of the clock control block: `CLKCTRL_G<global network number>` for a global routing network, or `CLKCTRL_R<regional network number>` for a dedicated regional routing network in the device.

If you want to disable the automatic global promotion performed in the Fitter to prevent other signals from being placed on global (or regional) routing networks, turn off the **Auto Global Clock** and **Auto Global Register Control Signals** options in the **Advanced Settings (Fitter)** dialog box.

If you are using design partition scripts for independent partitions, the Quartus II software can automatically write the commands to pass global constraints and turn off automatic options.

Alternatively, to avoid problems when integrating partitions into the top-level design, you can direct the Fitter to discard the placement and routing of the partition netlist by using the post-synthesis netlist, which forces the Fitter to reassign all the global signals for the partition when compiling the top-level design.

#### Related Information

- [Advanced Settings \(Fitter\) Dialog Box online help](#)  
Information about how to disable automatic global promotion
- [Generating Design Partition Scripts for Project Management online help](#)
- [Analyzing and Optimizing the Design Floorplan with the Chip Planner documentation](#)  
Information about how clock networks affect partition design

## Assign Virtual Pins

Virtual pins map lower-level design I/Os to internal cells. If you are developing an IP block in an independent Quartus II project, use virtual pins when the number of I/Os on a partition exceeds the device I/O count, and to increase the timing accuracy of cross-partition paths.

You can create a virtual pin assignment in the Assignment Editor for partition I/Os that will become internal nodes in the top-level design. When you apply the Virtual Pin assignment to an input pin, the pin no longer appears as an FPGA pin, but is fixed to GND or VCC in the design. The assigned pin is not an open node. Leave the clock pins mapped to I/O pins to ensure proper routing.

You can specify locations for the virtual pins that correspond to the placement of other partitions, and also make timing assignments to the virtual pins to define a timing budget. Virtual pins are created automatically from the top-level design if you use design partition scripts. The scripts place the virtual pins to correspond with the placement of the other partitions in the top-level design.

**Note:** Tri-state outputs cannot be assigned as virtual pins because internal tri-state signals are not supported in Altera devices. Connect the signal in the design with regular logic, or allow the software to implement the signal as an external device I/O pin.

### Related Information

[Generating Design Partition Scripts for Project Management online help](#)

## Perform Timing Budgeting if Required

If you optimize partitions independently and integrate them to the top-level design, or compile with empty partitions, any unregistered paths that cross between partitions are not optimized as entire paths. In these cases, the Quartus II software has no information about the placement of the logic that connects to the I/O ports. If the logic in one partition is placed far away from logic in another partition, the routing delay between the logic can lead to problems in meeting timing requirements. You can reduce this effect by ensuring that input and output ports of the partitions are registered whenever possible. Additionally, using the same top-level project framework helps to avoid this problem by providing the software with full information about other design partitions in the top-level design.

To ensure that the software correctly optimizes the input and output logic in any independent partitions, you might be required to perform some manual timing budgeting. For each unregistered timing path that crosses between partitions, make timing assignments on the corresponding I/O path in each partition to constrain both ends of the path to the budgeted timing delay. Assigning a timing budget for each part of the connection ensures that the software optimizes the paths appropriately.

When performing manual timing budgeting in a partition for I/O ports that become internal partition connections in a top-level design, you can assign location and timing constraints to the virtual pin that represents each connection to further improve the quality of the timing budget.

**Note:** If you use design partition scripts, the Quartus II software can write I/O timing budget constraints automatically for virtual pins.

### Related Information

[Generating Design Partition Scripts for Project Management online help](#)

## Drive Clocks Directly

When partitions are exported from another Quartus II project, you should drive partition clock inputs directly with device clock input pins.

Connecting the clock signal directly avoids any timing analysis difficulties with gated clocks. Clock gating is never recommended for FPGA designs because of potential glitches and clock skew. Clock gating can be especially problematic with exported partitions because the partitions have no information about gating that takes place at the top-level design or in another partition. If a gated clock is required in a partition, perform the gating within that partition.

Direct connections to input clock pins also allows design partition scripts to send constraints from the top-level device pin to lower-level partitions.

**Related Information**

[Invert Clocks in Destination Partitions](#) on page 14-15

**Recreate PLLs for Lower-Level Partitions if Required**

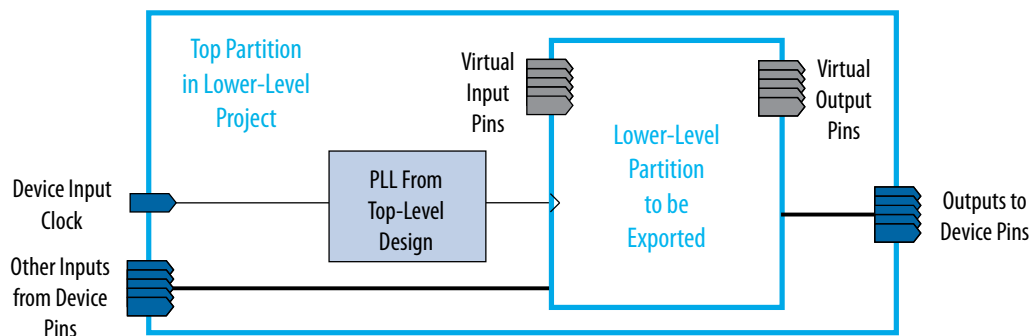
If you connect a PLL in your top-level design to partitions designed in separate Quartus II projects by third-party IP designers, the IP partitions do not have information about the multiplication, phase shift, or compensation delays for the PLL in the top-level design. To accommodate the PLL timing, you can make appropriate timing assignments in the projects created by IP designers to ensure that clocks are not left unconstrained or constrained with an incorrect frequency. Alternatively, you can duplicate the top-level PLL (or other derived clock logic) in the design file for the project created by the IP designer to ensure that you have the correct PLL parameters and clock delays for a complete and accurate timing analysis.

If the project lead creates a copy of the top-level project framework that includes all the settings and constraints needed for the design, this framework should include PLLs and other interface logic if this information is important to optimize partitions.

If you use a separate Quartus II project for an independent design block (such as when a designer or third-party IP provider does not have access to the entire design framework), include a copy of the top-level PLL in the lower-level partition as shown in figure.

In either case, the IP partition in the separate Quartus II project should contain just the partition logic that will be exported to the top-level design, while the full project includes more information about the top-level design. When the partition is complete, you can export just the partition without exporting the auxiliary PLL components to the top-level design. When you export a partition, the Quartus II software exports any hierarchy under the specified partition into the Quartus II Exported Partition File (.qxp), but does not include logic defined outside the partition (the PLL in this example).

**Figure 14-19: Recreating a Top-Level PLL in a Lower-Level Partition**





## Checking Partition Quality

There are several tools you can use to create and analyze partitions in the Quartus II software. Take advantage of these tools to assess your partition quality, and use the information to improve your design or assignments as required to achieve the best results.

### Incremental Compilation Advisor

You can use the Incremental Compilation Advisor to ensure that your design follows Altera's recommendations for creating design partitions and implementing the incremental compilation design flow methodology. Each recommendation in the Incremental Compilation Advisor provides an explanation, describes the effect of the recommendation, and provides the action required to make the suggested change.

#### Related Information

- [Incremental Compilation Advisor](#) on page 14-28
- [Incremental Compilation Advisor Command online help](#)
- [Example of Using the Incremental Compilation Advisor to Identify Non-Global Ports That Are Not Registered online help](#)  
For more information about the Incremental Compilation Advisor
- [Quartus II Incremental Compilation for Hierarchical and Team-Based Design documentation](#) on page 3-1

### Design Partition Planner

The Design Partition Planner allows you to view design connectivity and hierarchy, and can assist you in creating effective design partitions that follow the guidelines in this manual. You can also use the Design Partition Planner to optimize design performance by isolating and resolving failing paths on a partition-by-partition basis.

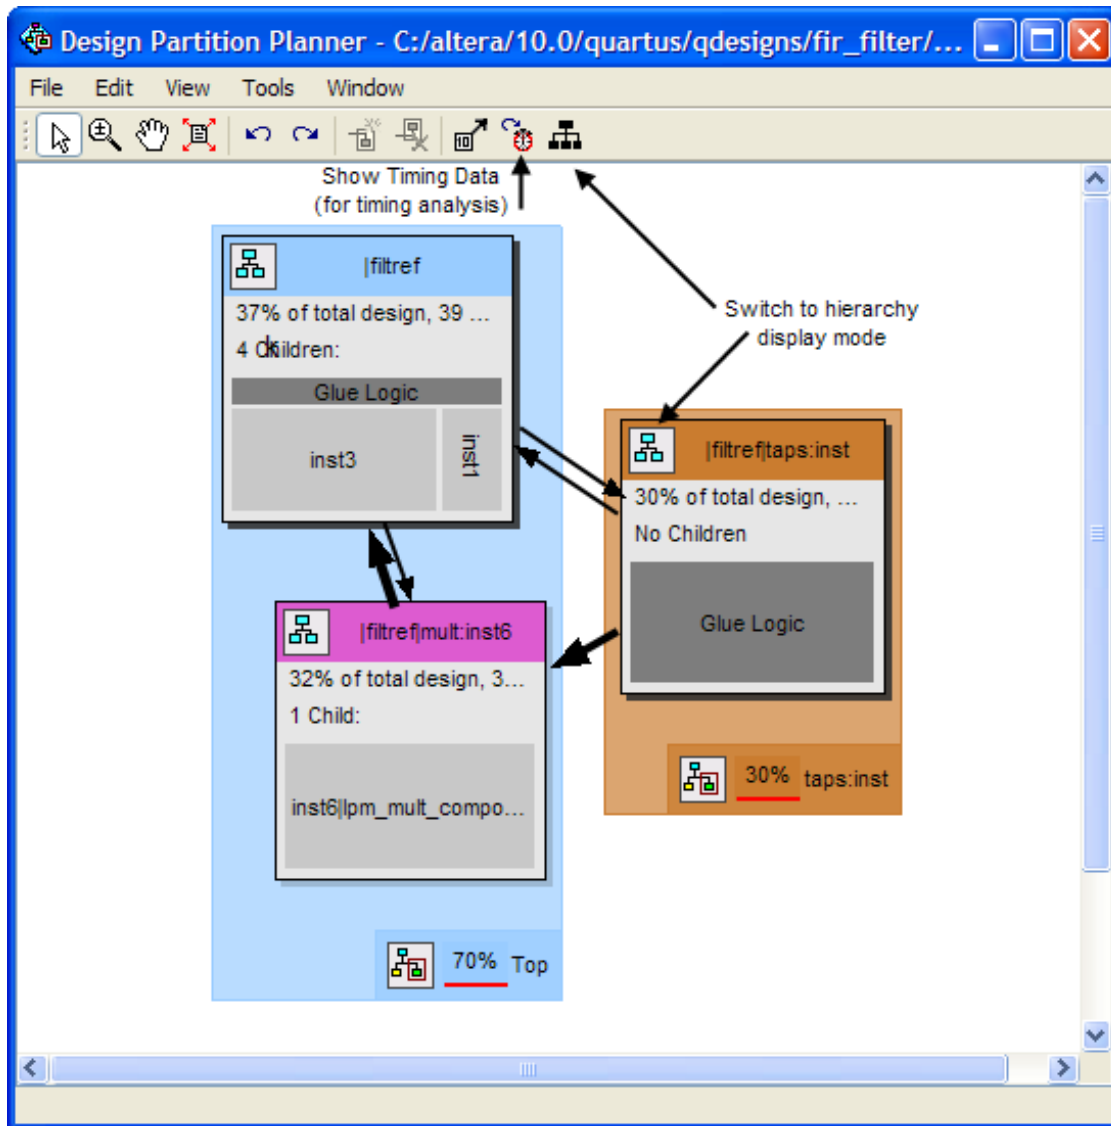
To view a design and create design partitions in the Design Partition Planner, you must first compile the design, or perform Analysis & Synthesis. In the Design Partition Planner, the design appears as a single top-level design block, with lower-level instances displayed as color-specific boxes.

In the Design Partition Planner, you can show connectivity between blocks and extract instances from the top-level design block. When you extract entities, connection bundles are drawn between entities, showing the number of connections existing between pairs of entities. When you have extracted a design block that you want to set as a design partition, right-click that design block, and then click **Create Design Partition**.

The Design Partition Planner also has an auto-partition feature that creates partitions based on the size and connectivity of the hierarchical design blocks. You can right-click the design block you want to partition (such as the top-level design hierarchy), and then click **Auto-Partition Children**. You can then analyze and adjust the partition assignments as required.

The figure shows the Design Partition Planner after making a design partition assignment to one instance and dragging another instance away from the top-level block within the same partition (two design blocks in the pale blue shaded box). The figure shows the connections between each partition and information about the size of each design instance.

Figure 14-20: Design Partition Planner



You can switch between connectivity display mode and hierarchical display mode, to examine the view-only hierarchy display. You can also remove the connection lines between partitions and I/O banks by turning off **Display connections to I/O banks**, or use the settings on the **Connection Counting** tab in the **Bundle Configuration** dialog box to adjust how the connections are counted in the bundles.

To optimize design performance, confine failing paths within individual design partitions so that there are no failing paths passing between partitions. In the top-level entity, child entities that contain failing paths are marked by a small red dot in the upper right corner of the entity box.

To view the critical timing paths from a timing analyzer report, first perform a timing analysis on your design, and then in the Design Partition Planner, click **Show Timing Data** on the View menu.

### Related Information

- [Design Partition Planner online help](#)
- [Using the Design Partition Planner online help](#)

## Viewing Design Partition Planner and Floorplan Side-by-Side

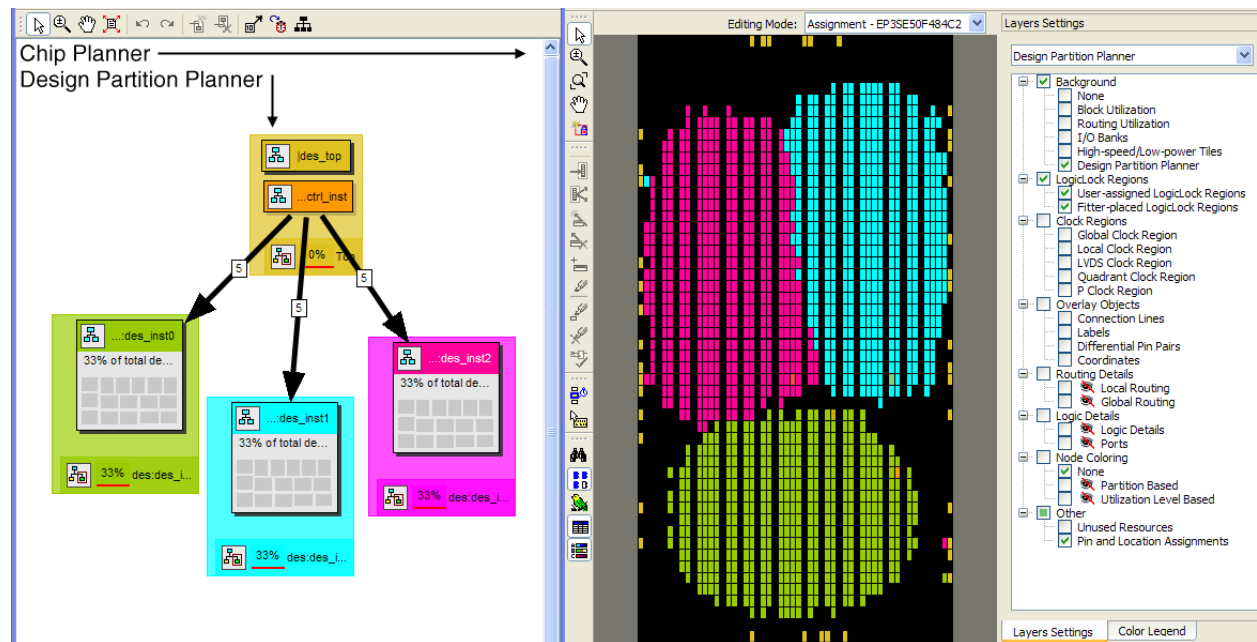
You can use the Design Partition Planner together with the Chip Planner to analyze natural placement groupings. This information can help you decide whether the design blocks should be grouped together in one partition, or whether they will make good partitions in the next compilation. It can also help determine whether the logic can easily be constrained by a LogicLock region. If logic naturally groups together when compiled without placement constraints, you can probably assign a reasonably sized LogicLock region to constrain the placement for subsequent compilations. You can experiment by extracting different design blocks in the Design Partition Planner and viewing the placement results of those design blocks from the previous compilation.

To view the Design Partition Planner and Chip Planner side-by-side, open the Design Partition Planner, and then open the Chip Planner and select the **Design Partition Planner** task. The **Design Partition Planner** task displays the physical locations of design entities with the same colors as in the Design Partition Planner.

In the Design Partition Planner, you can extract instances of interest from their parents by dragging and dropping, or with the **Extract from Parent** command. Evaluate the physical locations of instances in the Chip Planner and the connectivity between instances displayed in the Design Partition Planner. An entity is generally not suitable to be set as a separate design partition or constrained in a LogicLock region if the Chip Planner shows it physically dispersed over a noncontiguous area of the device after compilation. Use the Design Partition Planner to analyze the design connections. Child instances that are unsuitable to be set as separate design partitions or placed in LogicLock regions can be returned to their parent by dragging and dropping, or with the **Collapse to Parent** command.

The figure shows a design displayed in the Design Partition Planner and the Chip Planner with different colors for the top-level design and the three major design instances.

Figure 14-21: Design Partition Planner and Chip Planner



## Partition Statistics Report

You can view statistics about design partitions in the Partition Merge Partition Statistics report and the **Statistics** tab of the **Design Partitions Properties** dialog box. These reports are useful when optimizing your design partitions, or when compiling the completed top-level design in a team-based compilation flow to ensure that partitions meet the guidelines discussed in this manual.

The Partition Merge Partition Statistics report in the Partition Merge section of the Compilation report lists statistics about each partition. The statistics for each partition (each row in the table) include the number of logic cells, as well as the number of input and output pins and how many are registered. This report also lists how many ports are unconnected, or driven by a constant  $V_{CC}$  or GND. You can use this information to assess whether you have followed the guidelines for partition boundaries.

You can also view statistics about the resource and port connections for a particular partition on the **Statistics** tab of the **Design Partition Properties** dialog box. The **Show All Partitions** button allows you to view all the partitions in the same report. The Partition Merge Partition Statistics report also shows statistics for the **Internal Congestion: Total Connections and Registered Connections**. This information represents how many signals are connected within the partition. It then lists the inter-partition connections for each partition, which helps you to see how partitions are connected to each other.

### Related Information

[Partition Merge Reports online help](#)

## Report Partition Timing in the TimeQuest Timing Analyzer

The Report Partitions diagnostic report and the `report_partitions` SDC command in the TimeQuest analyzer produce a **Partition Timing Overview** and **Partition Timing Details** table, which lists the partitions, the number of failing paths, and the worst case timing slack within each partition.

You can use these reports to analyze the location of the critical timing paths in the design in relation to partitions. If a certain partition contains many failing paths, or failing inter-partition paths, you might be able to change your partitioning scheme and improve timing performance.

#### Related Information

#### [Quartus II TimeQuest Timing Analyzer documentation](#)

Information about the TimeQuest `report_timing` command and reports

## Check if Partition Assignments Impact the Quality of Results

You can ensure that you limit negative effect on the quality of results by following an iterative methodology during the partitioning process. In any incremental compilation flow where you can compile the source code for every partition during the partition planning phase, Altera recommends the following iterative flow:

1. Start with a complete design that is not partitioned and has no location or LogicLock region assignments.  
  
After Analysis & Synthesis and Partition Merge, perform a placement and timing analysis estimate with the **Start Early Timing Estimate** command. To run a full compilation instead, use the **Start Compilation** command.
2. Record the quality of results from the Compilation report (timing slack or  $f_{MAX}$ , area and any other relevant results).
3. Create design partitions following the guidelines described in this manual.
4. Perform another early timing estimate or a full compilation.
5. Record the quality of results from the Compilation report. If the quality of results is significantly worse than those obtained in the previous compilation, repeat step 3 through step 5 to change your partition assignments and use a different partitioning scheme.
6. Even if the quality of results is acceptable, you can repeat step 3 through step 5 by further dividing a large partition into several smaller partitions, which can improve compilation time in subsequent incremental compilations. You can repeat these steps until you achieve a good trade-off point (that is, all critical paths are localized within partitions, the quality of results is not negatively affected, and the size of each partition is reasonable).

You can also remove or disable partition assignments defined in the top-level design at any time during the design flow to compile the design as one flat compilation and get all possible design optimizations to assess the results. To disable the partitions without deleting the assignments, use the **Ignore partition assignments during compilation** option on the **Incremental Compilation** page of the **Settings** dialog box in the Quartus II software. This option disables all design partition assignments in your project and runs a full compilation, ignoring all partition boundaries and netlists. This option can be useful if you are using partitions to reduce compilation time as you develop various parts of the design, but can run a long compilation near the end of the design cycle to ensure the design meets its timing requirements.

## Including SDC Constraints from Lower-Level Partitions for Third-Party IP Delivery

When exported partitions are compiled in a separate Quartus II project, such as when a third-party designer is delivering IP, the project lead must transfer the top-level project framework information and constraints to the partitions, so that each designer has a consistent view of the constraints that apply to the

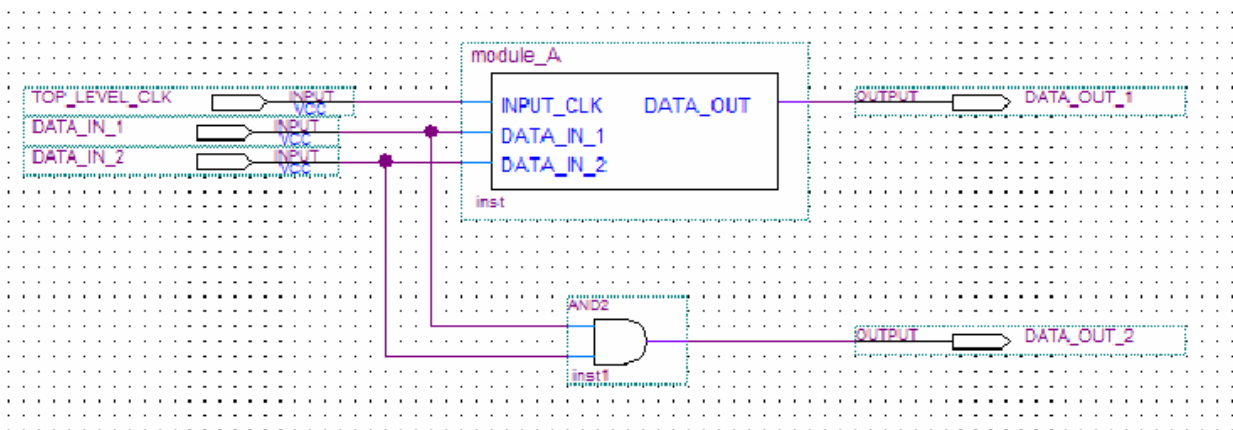
entire design. If the independent partition designers make any changes or add any constraints, they might have to transfer new constraints back to the project lead, so that these constraints are included in final timing sign-off of the entire design. Many assignments from the partition are carried with the partition into the top-level design; however, SDC format constraints for the TimeQuest analyzer are not copied into the top-level design automatically.

Passing additional timing constraints from a partition to the top-level design must be managed carefully. You can design within a single Quartus II project or a copy of the top-level design to simplify constraint management.

To ensure that there are no conflicts between the project lead's top-level constraints and those added by the third-party IP designer, use two .sdc files for each separate Quartus II project: an .sdc created by the project lead that includes project-wide constraints, and an .sdc created by the IP designer that includes partition-specific constraints.

The example design shown in the figure below is used to illustrate recommendations for managing the timing constraints in a third-party IP delivery flow. The top-level design instantiates a lower-level design block called `module_A` that is set as a design partition and developed by an IP designer in a separate Quartus II project.

Figure 14-22: Example Design to Illustrate SDC Constraints



In this top-level design, there is a single clock setting called `clk` associated with the FPGA input called `top_level_clk`. The top-level .sdc contains the following constraint for the clock:

```
create_clock -name {clk} -period 3.000 -waveform { 0.000 1.500 } \
[get_ports {TOP_LEVEL_CLK}]
```

## Creating an .sdc File with Project-Wide Constraints

The .sdc with project-wide constraints for the separate Quartus II project should contain all constraints that are not completely localized to the partition. The .sdc should be maintained by the project lead. The project lead must ensure that these timing constraints are delivered to the individual partition owners and that they are syntactically correct for each of the separate Quartus II projects. This communication can be challenging when the design is in flux and hierarchies change. The project lead can use design partition scripts to automatically pass some of these constraints to the separate Quartus II projects.

**Example Step 1—Project Lead Produces .sdc with Project-Wide Constraints for Lower-Level Partitions**

The **.sdc** with project-wide constraints is used in the partition, but is not exported back to the top-level design. The partition designer should not modify this file. If changes are necessary, they should be communicated to the project lead, who can then update the SDC constraints and distribute new files to all partition designers as required.

The **.sdc** should include clock creation and clock constraints for any clock used by more than one partition. These constraints are particularly important when working with complex clocking structures, such as the following:

- Cascaded clock multiplexers
- Cascaded PLLs
- Multiple independent clocks on the same clock pin
- Redundant clocking structures required for secure applications
- Virtual clocks and generated clocks that are consistently used for source synchronous interfaces
- Clock uncertainties

Additionally, the **.sdc** with project-wide constraints should contain all project-wide timing exception assignments, such as the following:

- Multicycle assignments, `set_multicycle_path`
- False path assignments, `set_false_path`
- Maximum delay assignments, `set_max_delay`
- Minimum delay assignments, `set_min_delay`

The project-wide **.sdc** can also contain any `set_input_delay` or `set_output_delay` constraints that are used for ports in separate Quartus II projects, because these represent delays external to a given partition. If the partition designer wants to set these constraints within the separate Quartus II projects, the team must ensure that the I/O port names are identical in all projects so that the assignments can be integrated successfully without changes.

Similarly, a constraint on a path that crosses a partition boundary should be in the project-wide **.sdc**, because it is not completely localized in a separate Quartus II project.

**Related Information**

[Generating Design Partition Scripts for Project Management online help](#)

**Example Step 1—Project Lead Produces .sdc with Project-Wide Constraints for Lower-Level Partitions**

The device input `top_level_clk` in [Figure 14-22](#) drives the `input_clk` port of `module_A`. To make sure the clock constraint is passed correctly to the partition, the project lead creates an **.sdc** with project-wide constraints for `module_A` that contains the following command:

```
create_clock -name {clk} -period 3.000 -waveform { 0.000 1.500 } [get_ports
{INPUT_CLK}]
```

The designer of `module_A` includes this **.sdc** as part of the separate Quartus II project.

**Creating an .sdc with Partition-Specific Constraints**

The **.sdc** with partition-specific constraints should contain all constraints that affect only the partition. For example, a `set_false_path` or `set_multicycle_path` constraint for a path entirely within the partition

should be in the partition-specific **.sdc**. These constraints are required for correct compilation of the partition, but do not need to be present in any other separate Quartus II projects.

The partition-specific **.sdc** should be maintained by the partition designer; they must add any constraints required to properly compile and analyze their partition.

The partition-specific **.sdc** is used in the separate Quartus II project and must be exported back to the project lead for the top-level design. The project lead must use the partition-specific constraints to properly constrain the placement, routing, or both, if the partition logic is fit at the top level, and to ensure that final timing sign-off is accurate. Use the following guidelines in the partition-specific **.sdc** to simplify these export and integration steps:

- Create a hierarchy variable for the partition (such as `module_A_hierarchy`) and set it to an empty string because the partition is the top-level instance in the separate Quartus II project. The project lead modifies this variable for the top-level hierarchy, reducing the effort of translating constraints on lower-level design hierarchies into constraints that apply in the top-level hierarchy. Use the following Tcl command first to check if the variable is already defined in the project, so that the top-level design does not use this empty hierarchy path: `if {[info exists module_A_hierarchy]}`.
- Use the hierarchy variable in the partition-specific **.sdc** as a prefix for assignments in the project. For example, instead of naming a particular instance of a register `reg:inst`, use `module_A_hierarchy:reg:inst`. Also, use the hierarchy variable as a prefix to any wildcard characters (such as `* *`).
- Pay attention to the location of the assignments to I/O ports of the partition. In most cases, these assignments should be specified in the **.sdc** with project-wide constraints, because the partition interface depends on the top-level design. If you want to set I/O constraints within the partition, the team must ensure that the I/O port names are identical in all projects so that the assignments can be integrated successfully without changes.
- Use caution with the `derive_clocks` and `derive_pll_clocks` commands. In most cases, the **.sdc** with project-wide constraints should call these commands. Because these commands impact the entire design, integrating them unexpectedly into the top-level design might cause problems.

If the design team follows these recommendations, the project lead should be able to include the **.sdc** with the partition-specific constraints provided by the partition designer directly in the top-level design.

### Example Step 2—Partition Designer Creates .sdc with Partition-Specific Constraints

The partition designer compiles the design with the **.sdc** with project-wide constraints and might want to add some additional constraints. In this example, the designer realizes that he or she must specify a false path between the register called `reg_in_1` and all destinations in this design block with the wildcard character (such as `* *`). This constraint applies entirely within the partition and must be exported to the top-level design, so it qualifies for inclusion in the **.sdc** with partition-specific constraints. The designer first defines the `module_A_hierarchy` variable and uses it when writing the constraint as follows:

```
if {[info exists module_A_hierarchy]} {  
    set module_A_hierarchy ""  
}  
set_false_path -from [get_registers ${module_A_hierarchy}reg_in_1] \  
-to [get_registers ${module_A_hierarchy}*]
```

### Consolidating the .sdc in the Top-Level Design

When the partition designers complete their designs, they export the results to the project lead. The project lead receives the exported **.qxp** files and a copy of the **.sdc** with partition-specific constraints.



To set up the top-level `.sdc` constraint file to accept the `.sdc` files from the separate Quartus II projects, the top-level `.sdc` should define the hierarchy variables specified in the partition `.sdc` files. List the variable for each partition and set it to the hierarchy path, up to and including the instantiation of the partition in the top-level design, including the final hierarchy character ”|”.

To ensure that the `.sdc` files are used in the correct order, the project lead can use the Tcl Source command to load each `.sdc`.

### Example Step 3—Project Lead Performs Final Timing Analysis and Sign-off

With these commands, the top-level `.sdc` file looks like the following example:

```
create_clock -name {clk} -period 3.000 -waveform { 0.000 1.500 } \
[get_ports {TOP_LEVEL_CLK}]
# Include the lower-level SDC file
set module_A_hierarchy "module_A:inst|" # Note the final '|' character
source <partition-specific constraint file such as ..\module_A
\module_A_constraints>.sdc
```

When the project lead performs top-level timing analysis, the false path assignment from the lower-level `module_A` project expands to the following:

```
set_false_path -from module_A:inst|reg_in_1 -to module_A:inst|*
```

Adding the hierarchy path as a prefix to the SDC command makes the constraint legal in the top-level design, and ensures that the wildcard does not affect any nodes outside the partition that it was intended to target.

## Introduction to Design Floorplans

A floorplan represents the layout of the physical resources on the device. Creating a design floorplan, or floorplanning, describes the process of mapping the logical design hierarchy onto physical regions in the device.

In the Quartus II software, LogicLock regions can be used to constrain blocks of a design to a particular region of the device. LogicLock regions represent an area on the device with a user-defined or Fitter-defined size and location in the device layout.

### Related Information

[Analyzing and Optimizing the Design Floorplan with the Chip Planner documentation](#)

## The Difference between Logical Partitions and Physical Regions

Design partitions are logical entities based on the design hierarchy. LogicLock regions are physical placement assignments that constrain logic to a particular region on the device.

A common misconception is that logic from a design partition is always grouped together on the device when you use incremental compilation. Actually, logic from a partition can be placed anywhere in the device if it is not constrained to a LogicLock region, although the Fitter can pack related logic together to improve timing performance. A logical design partition does not refer to any physical area on the device and does not directly control where instances are placed on the device.

If you want to control the placement of logic from a design partition and isolate it to a particular part of the device, you can assign the logical design partition to a physical region in the device floorplan with a

LogicLock region assignment. Altera recommends creating a design floorplan by assigning design partitions to LogicLock regions to improve the quality of results and avoid placement conflicts in some situations for incremental compilation.

Another misconception is that LogicLock assignments are used to preserve placement results for incremental compilation. Actually, LogicLock regions only constrain logic to a physical region on the device. Incremental compilation does not use LogicLock assignments or any location assignments to preserve the placement results; it simply reuses the results stored in the database netlist from a previous compilation.

## Why Create a Floorplan?

Creating a design floorplan is usually required if you want to preserve placement for partitions that will be exported, to avoid resource conflicts between partitions in the top-level design. Floorplan location planning can be important for a design that uses incremental compilation, for the following reasons:

- To avoid resource conflicts between partitions, predominantly when integrating partitions exported from another Quartus II project.
- To ensure good quality of results when recompiling individual timing-critical partitions.

Location assignments for each partition ensure that there are no placement conflicts between partitions. If there are no LogicLock region assignments, or if LogicLock regions are set to auto-size or floating location, no device resources are specifically allocated for the logic associated with the region. If you do not clearly define resource allocation, logic placement can conflict when you integrate the partitions in the top-level design if you reuse the placement information from the exported netlist.


Creating a floorplan is also recommended for timing-critical partitions that have little timing margin to maintain good quality of results when the design changes.

Floorplan assignments are not required for non-critical partitions compiled in the same Quartus II project. The logic for partitions that are not timing-critical can be placed anywhere in the device on each recompilation if that is best for your design.

Design floorplan assignments prevent the situation in which the Fitter must place a partition in an area of the device where most resources are used by other partitions. A LogicLock region provides a reasonable region to re-place logic after a change, so the Fitter does not have to scatter logic throughout the available space in the device.

The figure illustrates the problems that may be associated with refitting designs that do not have floorplan location assignments. The left floorplan shows the initial placement of a four-partition design (P1-P4) without any floorplan location assignments. The right floorplan shows the device if a change occurs to P3. After removing the logic for the changed partition, the Fitter must re-place and reroute the new logic for P3 in the scattered white space. The placement of the post-fit netlists for other partitions forces the Fitter to implement P3 with the device resources that have not been used.


Figure 14-23: Representation of Device Floorplan without Location Assignments



The Fitter has a more difficult task because of more difficult physical constraints, and as a result, compilation time often increases. The Fitter might not be able to find any legal placement for the logic in partition P3, even if it could in the initial compilation. Additionally, if the Fitter can find a legal placement, the quality of results often decreases in these cases, sometimes dramatically, because the new partition is now scattered throughout the device.

The figure below shows the initial placement of a four-partition design with floorplan location assignments. Each partition is assigned to a LogicLock region. The second part of the figure shows the device after partition P3 is removed. This placement presents a much more reasonable task to the Fitter and yields better results.

Figure 14-24: Representation of Device Floorplan with Location Assignments



Altera recommends that you create a LogicLock floorplan assignment for timing-critical blocks with little timing margin that will be recompiled as you make changes to the design.

## When to Create a Floorplan

It is important that you plan early to incorporate partitions into the design, and ensure that each partition follows partitioning guidelines. You can create floorplan assignments at different stages of the design flow, early or late in the flow. These guidelines help ensure better results as you begin creating floorplan location assignments.

## Early Floorplan

An early floorplan is created before the design stage. You can plan an early floorplan at the top level of a design to allocate each partition a portion of the device resources. Doing so allows the designer for each block to create the logic for their design partition without conflicting with other logic. Each partition can be optimized in a separate Quartus II project if required, and the design can still be easily integrated in the top-level design. Even within one Quartus II project, each partition can be locked down with a post-fit netlist, and you can be sure there is space in the device floorplan for other partitions.

When you have compiled your complete design, or after you have integrated the first versions of partitions developed in separate Quartus II projects, you can use the design information and Quartus II features to tune and improve the floorplan .

## Late Floorplan

A late floorplan is created or modified after the design is created, when the code is close to complete and the design structure is likely to remain stable. Creating a late floorplan is typically necessary only if you are starting to use incremental compilation late in the design flow, or need to reserve space for a logic block that becomes timing-critical but still has HDL changes to be integrated. When the design is complete, you can take advantage of the Quartus II analysis features to check the floorplan quality. To adjust the floorplan, you can perform iterative compilations as required and assess the results of different assignments.

**Note:** It may not be possible to create a good-quality late floorplan if you do not create partitions in the early stages of the design.

## Design Floorplan Placement Guidelines

The following guidelines are key to creating a good design floorplan:

- Capture correct resources in each region.
- Use good region placement to maintain design performance compared to flat compilation.

A common misconception is that creating a floorplan enhances timing performance, as compared to a flat compilation with no location assignments. The Fitter does not usually require guidance to get optimal results for a full design.

Floorplan assignments can help maintain good performance when designs change incrementally. However, poor placement assignments in an incremental compilation can often adversely affect performance results, as compared to a flat compilation, because the assignments limit the options for the Fitter. Investing time to find good region placement is required to match the performance of a full flat compilation.

## Flow for Creating a Floorplan

Use the following general procedure to create a floorplan:

1. Divide the design into partitions.
2. Assign the partitions to LogicLock regions.
3. Compile the design.
4. Analyze the results.
5. Modify the placement and size of regions, as required.

You might have to perform these steps several times to find the best combination of design partitions and LogicLock regions that meet the resource and timing goals of the design.

#### Related Information

- [Quartus II Incremental Compilation for Hierarchical and Team-Based Design documentation](#) on page 3-1

## Assigning Partitions to LogicLock Regions

Before compiling a design with new LogicLock assignments, ensure that the partition netlist type is set to **Post-Synthesis** or **Source File**, so that the Fitter does not reuse previous placement results.

In most cases, you should include logic from one partition in each LogicLock region. This organization helps to prevent resource conflicts when partitions are exported and can lead to better performance preservation when locking down parts of a design in a single project.

The Quartus II software is flexible and allows exceptions to this rule. For example, you can place more than one partition in the same LogicLock region if the partitions are tightly connected, but you do not want to merge the partitions into one larger partition. For best results, ensure that you recompile all partitions in the LogicLock region every time the logic in one partition changes. Additionally, if a partition contains multiple lower-level entities, you can place those entities in different areas of the device with multiple LogicLock regions, even if they are defined in the same partition.

You can use the **Reserved** LogicLock option to ensure that you avoid conflicts with other logic that is not locked into a LogicLock region. This option prevents other logic from being placed in the region, and is useful if you have empty partitions at any point during your design flow, so that you can reserve space in the floorplan. Do not make reserved regions too large to prevent unused area because no other logic can be placed in a region with the **Reserved** LogicLock option.

#### Related Information

[LogicLock Region Properties Dialog Box online help](#)

## How to Size and Place Regions

In an early floorplan, assign physical locations based on design specifications. Use information about the connections between partitions, the partition size, and the type of device resources required.

In a late floorplan, when the design is complete, you can use locations or regions chosen by the Fitter as a guideline. If you have compiled the full design, you can view the location of the partition logic in the Chip Planner. You can use the natural grouping of each unconstrained partition as a starting point for a LogicLock region constraint. View the placement for each partition that requires a floorplan constraint, and create a new LogicLock region by drawing a box around the area on the floorplan, and then assigning the partition to the region to constrain the partition placement.

Instead of creating regions based on the previous compilation results, you can start with the Fitter results for a default auto size and floating origin location for each new region when the design logic is complete. After compilation, lock the size and origin location. Instead of a full compilation, you can use the **Start Early Timing Estimate** command to perform a fast placement.

Alternatively, if the design logic is complete with auto-sized or floating location regions, you can specify the size based on the synthesis results and use the locations chosen by the Fitter with the **Set to Estimated Size** command. Like the previous option, start with floating origin location. After compilation, lock the origin location. Again, instead of a full compilation, you can use the **Start Early Timing Estimate**

command to perform a fast placement. You can also enable the **Fast Synthesis Effort** setting to reduce synthesis time.

After a compilation or early timing estimate, save the Fitter size and origin location of the Fitter with the **Set Size and Origin to Previous Fitter Results** command.

**Note:** It is important that you use the Fitter-chosen locations only as a starting point to give the regions a good fixed size and location. Ensure that all LogicLock regions in the design have a fixed size and have their origin locked to a specific location on the device. On average, regions with fixed size and location yield better timing performance than auto-sized regions.

#### Related Information

- [Checking Partition Quality](#) on page 14-28
- [Creating and Manipulating LogicLock Regions online help](#)

## Modifying Region Size and Origin

After saving the Fitter results from an initial compilation for a late floorplan, modify the regions using your knowledge of the design to set a specific size and location. If you have a good understanding of how the design fits together, you can often improve upon the regions placed in the initial compilation. In an early floorplan, when the design has not yet been created, you can use the guidelines in this section to set the size and origin, even though there is no initial Fitter placement.

The easiest way to move and resize regions is to drag the region location and borders in the Chip Planner. Make sure that you select the **User-Defined** region in the floorplan (as opposed to the **Fitter-Placed** region from the last compilation) so that you can change the region.

Generally, you can keep the Fitter-determined relative placement of the regions, but make adjustments if required to meet timing performance. If you find that the early timing estimate did not result in good relative placements, try performing a full compilation so that the Fitter can optimize for a full placement and routing.

If two LogicLock regions have several connections between them, ensure they are placed near each other to improve timing performance. By placing connected regions near each other, the Fitter has more opportunity to optimize inter-region paths when both partitions are recompiled. Reducing the criticality of inter-region paths also allows the Fitter more flexibility when placing other logic in each region.

If resource utilization is low in the overall device, enlarge the regions. Doing so usually improves the final results because it gives the Fitter more freedom to place additional or modified logic added to the partition during subsequent incremental compilations. It also allows room for optimizations such as pipelining and physical synthesis logic duplication.

Try to have each region evenly full, with the same "fullness" that the complete design would have without LogicLock regions; Altera recommends approximately 75% full.

Allow more area for regions that are densely populated, because overly congested regions can lead to poor results. Allow more empty space for timing-critical partitions to improve results. However, do not make regions too large for their logic. Regions that are too large can result in wasted resources and also lead to suboptimal results.

Ideally, almost the entire device should be covered by LogicLock regions if all partitions are assigned to regions.

Regions should not overlap in the device floorplan. If two partitions are allocated on an overlapping portion of the chip, each may independently claim common resources in this region. This leads to resource conflicts when integrating results into a top-level design. In a single project, overlapping regions give more difficult constraints to the Fitter and can lead to reduced quality of results.

You can create hierarchical LogicLock regions to ensure that the logic in a child partition is physically placed inside the LogicLock region for its parent partition. This can be useful when the parent partition does not contain registers at the boundary with the lower-level child partition and has a lot of signal connectivity. To create a hierarchical relationship between regions in the LogicLock Regions window, drag and drop the child region to the parent region.

## I/O Connections

Consider I/O timing when placing regions. Using I/O registers can minimize I/O timing problems, and using boundary registers on partitions can minimize problems connecting regions or partitions. However, I/O timing might still be a concern. It is most important for flows where each partition is compiled independently, because the Fitter can optimize the placement for paths between partitions if the partitions are compiled at the same time.

Place regions close to the appropriate I/O, if necessary. For example, DDR memory interfaces have very strict placement rules to meet timing requirements. Incorporate any specific placement requirements into your floorplan as required. You should create LogicLock regions for internal logic only, and provide pin location assignments for external device I/O pins (instead of including the I/O cells in a LogicLock region to control placement).

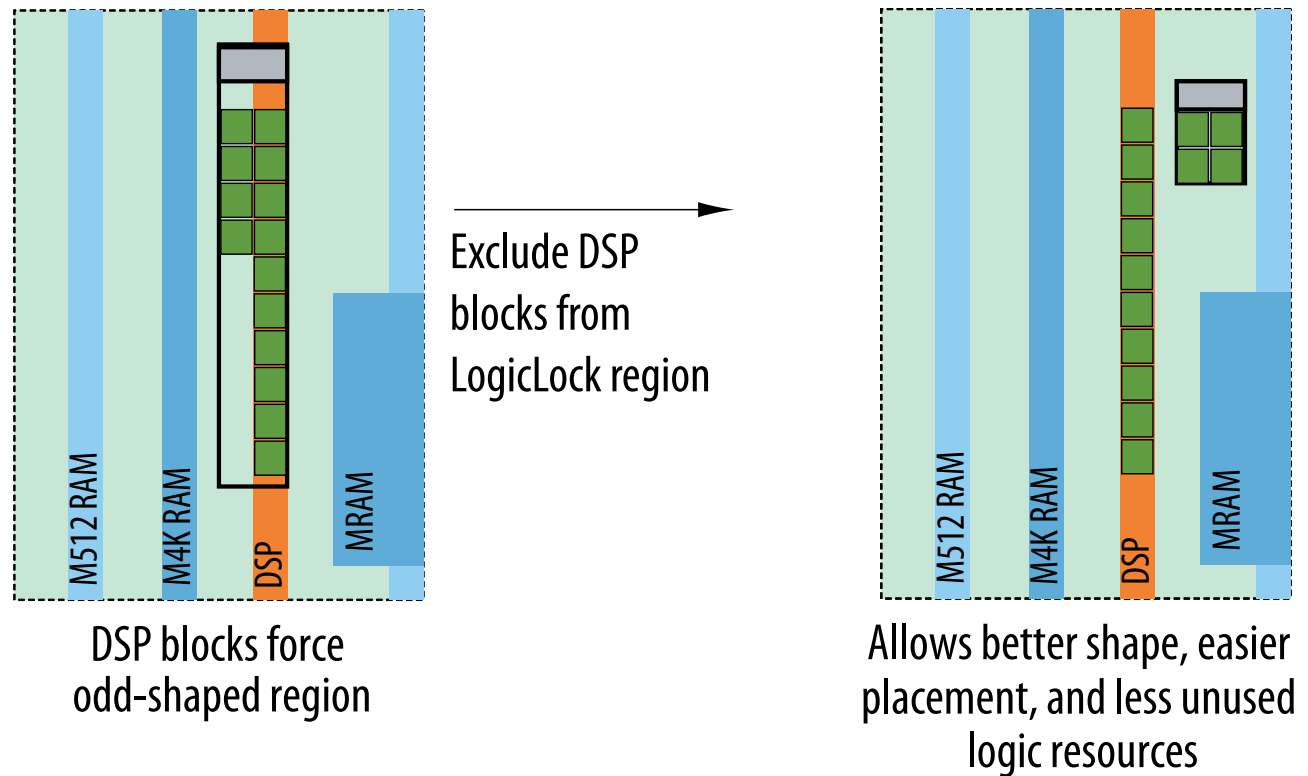
## LogicLock Resource Exclusions

You can exclude certain resource types from a LogicLock region to manage the ratio of logic to dedicated DSP and RAM resources in the region.

If your design contains memory or Digital Signal Processing (DSP) elements, you may want to exclude these elements from the LogicLock region. LogicLock resource exceptions prevent certain types of elements from being assigned to a region. Therefore, those elements are not required to be placed inside the region boundaries. The option does not prevent them from being placed inside the region boundaries unless the **Reserved** property of the region is turned on.

Resource exceptions are useful in cases where it is difficult to place rectangular regions for design blocks that contain memory and DSP elements, due to their placement in columns throughout the device floorplan. Exclude RAMs, DSPs, or logic cells to give the Fitter more flexibility with region sizing and placement. Excluding RAM or DSP elements can help to resolve no-fit errors that are caused by regions spanning too many resources, especially for designs that are memory-intensive, DSP-intensive, or both. The figure shows an example of a design with an odd-shaped region to accommodate DSP blocks for a region that does not contain very much logic. The right side of the figure shows the result after excluding DSP blocks from the region. The region can be placed more easily without wasting logic resources.

Figure 14-25: LogicLock Resource Exclusion Example



To view any resource exceptions, right-click in the LogicLock Regions window, and then click **LogicLock Regions Properties**. In the **LogicLock Regions Properties** dialog box, select the design element (module or entity) in the **Members** box, and then click **Edit**. In the **Edit Node** dialog box, to set up a resource exception, click the **Edit** button next to the **Excluded element types** box, and then turn on the design element types to be excluded from the region. You can choose to exclude combinational logic or registers from logic cells, or any of the sizes of TriMatrix memory blocks, or DSP blocks.

If the excluded logic is in its own lower-level design entity (even if it is within the same design partition), you can assign the entity to a separate LogicLock region to constrain its placement in the device.

You can also use this feature with the LogicLock **Reserved** property to reserve specific resources for logic that will be added to the design.

### Creating Floorplan Location Assignments With Tcl Commands—Excluding or Filtering Certain Device Elements (Such as RAM or DSP Blocks)

To assign a code block to a LogicLock region, with exclusions, use the following command:

```
set_logiclock_contents -region <LogicLock region name> \  
-to <block> -exceptions \"<keyword>:<keyword>\"
```



- *<LogicLock region name>*—The name of the LogicLock region to which the code block is assigned.
- *<block>*—A code block in a Quartus II project hierarchy, which can also be a design partition.
- *<keyword>*—The list of exceptions made during assignment. For example, if DSP was in the keyword list, the named block of code would be assigned to the LogicLock region, except for any DSP block within the code block. You can include the following exceptions in the `set_logiclock_contents` command:

Keyword variables:

- *REGISTER*—Any registers in the logic cells.
- *COMBINATIONAL*—Any combinational elements in the logic cells.
- *SMALL\_MEM*—Small TriMatrix memory blocks (M512 or MLAB).
- *MEDIUMMEM\_MEM*—Medium TriMatrix memory blocks (M4K or M9K).
- *LARGE\_MEM*—Large TriMatrix memory blocks (M-RAM or M144K).
- *DSP*—Any DSP blocks.
- *VIRTUAL\_PIN*—Any virtual pins.

**Note:** Resource filtering uses the optional Tcl argument `-exclude_resources` in the `set_logiclock_contents` function. If left unspecified, no resource filter is created. In the `.qsf`, resource filtering uses an extra LogicLock membership assignment called `LL_MEMBER_RESOURCE_EXCLUDE`. For example, the following line in the `.qsf` is used to specify a resource filter for the `alu:alu_unit` entity assigned to the ALU region.

```
set_instance_assignment -name LL_MEMBER_RESOURCE_EXCLUDE \
  "DSP:SMALL_MEM" -to "alu:alu_unit" -section_id ALU
```

## Creating Non-Rectangular Regions

To constrain placement to non-rectangular or non-contiguous areas of the device, you can connect multiple rectangular regions together using the **Merge** command.

For devices that do not support the **Merge** command (MAX<sup>TM</sup> II devices), you can limit entity placement to a sub-area of a LogicLock region to create non-rectangular constraints. In these devices, construct a LogicLock hierarchy by creating child regions inside of parent regions, and then use the **Reserved** option to control which logic can be placed inside these child regions. Setting the **Reserved** option for the region prevents the Fitter from placing nodes that are not assigned to the region inside the boundary of the region.

### Related Information

[Creating and Manipulating LogicLock Regions online help](#)

## Checking Floorplan Quality

The Quartus II software has several tools to help you create a floorplan. You can use these tools to assess your floorplan quality and use the information to improve your design or assignments as required to achieve the best results.

## Incremental Compilation Advisor

You can use the Incremental Compilation Advisor to check that your design follows the recommendations for creating floorplan location assignments that are presented in this manual.

## LogicLock Region Resource Estimates

You can view resource estimates for a LogicLock region to determine the region's resource coverage, and use this estimate before compilation to check region size. Using this estimate helps to ensure adequate resources when you are sizing or moving regions.

### Related Information

[LogicLock Region Properties Dialog Box online help](#)

## LogicLock Region Properties Statistics Report

LogicLock region statistics are similar to design partition properties, but also include resource usage details after compilation.

The statistics report the number of resources used and the total resources covered by the region, and also lists the number of I/O connections and how many I/Os are registered (good), as well as the number of internal connections and the number of inter-region connections (bad).

## Locate the Quartus II TimeQuest Timing Analyzer Path in the Chip Planner

In the TimeQuest analyzer user interface, you can locate a specific path in the Chip Planner to view its placement and perform a report timing operation (for example, report timing for all paths with less than 0 ns slack).

### Related Information

[Locate Dialog Box online help](#)

Information about how to locate paths between the TimeQuest analyzer and the Chip Planner

## Inter-Region Connection Bundles

The Chip Planner can display bundles of connections between LogicLock regions, with filtering options that allow you to choose the relevant data for display. These bundles can help you to visualize how many connections there are between each LogicLock region to improve floorplan assignments or to change partition assignments, if required.

### Related Information

[Inter-region Bundles Dialog Box online help](#)

Information about how to display bundles of connections between LogicLock regions

## Routing Utilization

The Chip Planner includes a feature to display a color map of routing congestion. This display helps identify areas of the chip that are too tightly packed.

In the Chip Planner, red LAB blocks indicate higher routing congestion. You can position the mouse pointer over a LAB to display a tooltip that reports the logic and routing utilization information.

### Related Information

[Chip Planner online help](#)

Information about how to view a color map of routing congestion in the Chip Planner

## Ensure Floorplan Assignments Do Not Significantly Impact Quality of Results

The end results of design partitioning and floorplan creation differ from design to design. However, it is important to evaluate your results to ensure that your scheme is successful. Compare your before and after results, and consider using another scheme if any of the following guidelines are not met:

- You should see only minor degradation in  $f_{MAX}$  after the design is partitioned and floorplan location assignments are created. There is some performance cost associated with setting up a design for incremental compilation; approximately 3% is typical.
- The area increase should be no more than 5% after the design is partitioned and floorplan location assignments are created.
- The time spent in the routing stage should not significantly increase.

The amount of compilation time spent in the routing stage is reported in the Messages window with an Info message that indicates the elapsed time for Fitter routing operations. If you notice a dramatic increase in routing time, the floorplan location assignments may be creating substantial routing congestion. In this case, decrease the number of LogicLock regions, which typically reduces the compilation time in subsequent incremental compilations and may also improve design performance.

## Recommended Design Flows and Application Examples

Listed below are application examples with design flows for partitioning and creating a design floorplan during common timing closure and team-based design scenarios. Each flow describes the situation in which it should be used, and provides a step-by-step description of the commands required to implement the flow.

### Create a Floorplan for Major Design Blocks

Use this incremental compilation flow for designs when you want to assign a floorplan location for each major block in your design. A full floorplan ensures that partitions do not interact as they are changed and recompiled— each partition has its own area of the device floorplan.

To create a floorplan for major design blocks, follow this general methodology:

1. In the Design Partitions window, ensure that all partitions have their netlist type set to **Source File** or **Post-Synthesis**. If the netlist type is set to **Post-Fit**, floorplan location assignments are not used when recompiling the design.
2. Create a LogicLock region for each partition (including the top-level entity, which is set as a partition by default).
3. Run a full compilation of your design to view the initial Fitter-chosen placement of the LogicLock regions as a guideline.
4. In the Chip Planner, view the placement results of each partition and LogicLock region on the device.
5. If required, modify the size and location of the LogicLock regions in the Chip Planner. For example, enlarge the regions to fill up the device and allow for future logic changes. You can also, if needed, create a new LogicLock region by drawing a box around an area on the floorplan.
6. Run an early timing estimate with the **Start Early Timing Estimate** command to estimate the timing performance of your design with the modified or new LogicLock regions.
7. Repeat steps 5 and 6 until you are satisfied with the quality of results for your design floorplan. Once you are satisfied with your results, run a full compilation of your design.

## Create a Floorplan Assignment for One Design Block with Difficult Timing

Use this flow when you have one timing-critical design block that requires more optimization than the rest of your design. You can take advantage of incremental compilation to reduce your compilation time without creating a full design floorplan.

In this scenario, you do not want to create floorplan assignments for the entire design. Instead, you can create a region to constrain the location of your critical design block, and allow the rest of the logic to be placed anywhere on the device. To create a region for critical design block, follow these steps:

1. Divide up your design into partitions. Ensure that you isolate the timing-critical logic in a separate partition.
2. Define a LogicLock region for the timing-critical partition. Ensure that you capture the correct amount of device resources in the region. Turn on the **Reserved** property to prevent any other logic from being placed in the region.
  - If the design block is not complete, reserve space in the design floorplan based on your knowledge of the design specifications, connectivity between design blocks, and estimates of the size of the partition based on any initial implementation numbers.
  - If the critical design block has initial source code ready, compile the design to place the LogicLock region. Save the Fitter-determined size and origin, and then enlarge the region to provide more flexibility and allow for future design changes.

As the rest of the design is completed, and the device fills up, the timing-critical region reserves an area of the floorplan. When you make changes to the design block, the logic will be re-placed in the same part of the device, which helps ensure good quality of results.

### Related Information

[Design Partition Guidelines](#) on page 14-9

## Create a Floorplan as the Project Lead in a Team-Based Flow

Use this approach when you have several designs that will be implemented in separate Quartus II projects by different designers, or third-party IP designers who want to optimize their designs independently and pass the results to the project lead.

As the project lead in this scenario, follow these steps to prepare the top-level design for a successful team-based design methodology with early floorplan planning:

1. Create a new Quartus II project that will ultimately contain the full implementation of the entire design.
2. Create a “skeleton” or framework of the design that defines the hierarchy for the subdesigns that will be implemented by separate designers. Consider the partitioning guidelines in this manual when determining the design hierarchy.
3. Make project-wide settings. Select the device, make global assignments for clocks and device I/O ports, and make any global signal constraints to specify which signals can use global routing resources.
4. Make design partition assignments for each major subdesign. Set the netlist type for each partition that will be implemented in a separate Quartus II project and later exported and integrated with the top-level design set to **Empty**.
5. Create LogicLock regions for each partition to create a design floorplan. This floorplan should consider the connectivity between partitions and estimates of the size of each partition based on any

initial implementation numbers and knowledge of the design specifications. Use the guidelines described in this chapter to choose a size and location for each LogicLock region.

6. Provide the constraints from the top-level design to partition designers using one of the following procedures:
  - a. Create a copy of the top-level Quartus II project framework by checking out the appropriate files from a source control system, using the **Copy Project** command, or creating a project archive. Provide each partition designer with the copy of the project.
  - b. Provide the constraints with documentation or scripts.

#### Related Information

[Generating Design Partition Scripts for Project Management online help](#)

## Document Revision History

Table 14-1: Document Revision History

Date	Version	Changes
2014.12.15	14.1.0	<ul style="list-style-type: none"> <li>• Updated location of Fitter Settings, Analysis &amp; Synthesis Settings, and Physical Optimization Settings to Compiler Settings.</li> <li>• Updated description of Virtual Pin assignment to clarify that assigned pins are no longer free as input pins.</li> </ul>
June 2014	14.0.0	<ul style="list-style-type: none"> <li>• Dita conversion.</li> <li>• Removed obsolete devices content for Arria GX, Cyclone, Cyclone II, Cyclone III, Stratix, Stratix GX, Stratix II, Stratix II GX,</li> <li>• Replace Megafunction content with IP Catalog and Parameter Editor content.</li> </ul>
November 2013	13.1.0	Removed HardCopy device information.
November 2012	12.1.0	Added Turning On Supported Cross-Boundary Optimizations.
June 2012	12.0.0	Removed survey link.
November 2011	11.0.1	Template update.
May 2011	11.0.0	Updated links.

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> <li>• Changed to new document template.</li> <li>• Moved "Creating Floorplan Location Assignments With Tcl Commands—Excluding or Filtering Certain Device Elements (Such as RAM or DSP Blocks)" from the Quartus II Incremental Compilation for Hierarchical and Team-Based Design chapter in volume 1 of the <i>Quartus II Handbook</i>.</li> <li>• Consolidated Design Partition Planner and Incremental Compilation Advisor information between the Quartus II Incremental Compilation for Hierarchical and Team-Based Design and Best Practices for Incremental Compilation Partitions and Floorplan Assignments handbook chapters.</li> </ul>
July 2010	10.0.0	<ul style="list-style-type: none"> <li>• Removed the explanation of the “bottom-up design flow” where designers work completely independently, and replaced with Altera’s recommendations for team-based environments where partitions are developed in the same top-level project framework, plus an explanation of the bottom-up process for including independent partitions from third-party IP designers.</li> <li>• Expanded the <b>Merge</b> command explanation to explain how it now accommodates cross-partition boundary optimizations.</li> <li>• Restructured Altera recommendations for when to use a floorplan.</li> </ul>
October 2009	9.1.0	<ul style="list-style-type: none"> <li>• Redefined the bottom-up design flow as team-based and reorganized previous design flow examples to include steps on how to pass top-level design information to lower-level projects.</li> <li>• Added "Including SDC Constraints from Lower-Level Partitions for Third-Party IP Delivery" from the Quartus II Incremental Compilation for Hierarchical and Team-Based Design chapter in volume 1 of the <i>Quartus II Handbook</i>.</li> <li>• Reorganized the "Recommended Design Flows and Application Examples" section.</li> <li>• Removed HardCopy APEX and HardCopy Stratix Devices section.</li> </ul>
March 2009	9.0.0	<ul style="list-style-type: none"> <li>• Added I/O register packing examples from <i>Incremental Compilation for Hierarchical and Team-Based Designs</i> chapter</li> <li>• Moved "Incremental Compilation Advisor" section</li> <li>• Added "Viewing Design Partition Planner and Floorplan Side-by-Side" section</li> <li>• Updated Figure 15-22</li> <li>• Chapter 8 was previously Chapter 7 in software release 8.1.</li> </ul>
November 2008	8.1.0	<ul style="list-style-type: none"> <li>• Changed to 8-1/2 x 11 page size. No change to content.</li> </ul>

Date	Version	Changes
May 2007	8.0.0	<ul style="list-style-type: none"><li data-bbox="461 239 662 273">• Initial release.</li></ul>

**Related Information**[Quartus II Handbook Archive](#)

For previous versions of the Quartus II Handbook

# Mitigating Single Event Upsets 15

2014.06.30

QI15V1



Subscribe



Send Feedback

The Quartus II software offers several features to enable the detection and correction of single event upsets (SEUs), or soft errors, as well as to characterize the effects of SEU on your designs.

## Understanding SEU

SEU can affect any semiconductor device.

SEUs are rare, unintended changes in the state of internal memory elements, caused by cosmic radiation effects. The change in state results in a soft error, so the affected device can be reset to its original value and there is no permanent damage to the device itself. Because of the unintended memory state, the device may operate erroneously until this upset is fixed.

The Soft Error Rate (SER) is expressed as Failure-in-Time (FIT) units, defined as one soft error occurrence every billion hours of operation. Often SEU mitigation is not required because of the low chance of occurrence. However, for highly complex systems, such as with multiple high-density components, error rate may be a significant system design factor. If your system includes multiple FPGAs and requires very high reliability and availability, you should consider the implications of soft errors, and use the available techniques for detecting and recovering from these types of errors. If your system is requiring high reliability and availability, consider the implications of soft errors, and use the techniques in this document to detect and recover from these types of errors.

FPGAs use memory both in user logic (bulk memory and registers) and in Configuration Random Access Memory (CRAM). CRAM configures the FPGA; this is the memory loaded with the contents of a .sof file by the Quartus II Programmer. The CRAM configures all logic and routing in the device. If an SEU strikes a CRAM bit, the effect can be harmless if the CRAM bit is not in use. However, the effect can be severe if it affects critical logic internal signal routing (such as a lookup table bit).

### Related Information

[Introduction to Single Event Upsets](#)

## Mitigating SEU Effects in Embedded User RAM

You can mitigate SEU effects for internal memories by using Error Correcting Codes (ECC) for internal SRAM blocks. This method is effective enough so that the FIT rate of ECC-protected memories is almost zero.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered





ECC use a Cyclic Redundancy Check (CRC) code for a given data word. These CRC bits provide redundancy on the data word which can detect the location of single-bit and double-bit flips in the data word. Since the location of the bit flips is known, the CRC can locate and correct the error. The CRC code can also account for bit flips in the CRC code itself. Altera FPGAs support both Single Error Correction Double Error Detection (SECEDED) and Double Error Correction Triple Error Detection (DECTED), depending on the device family.

Some Altera device memory blocks offer built-in CRC circuitry hardened in silicon. This is available in the M20K memory in Stratix V, and in the M144K block in Stratix IV, and Arria II devices. (Other device families can implement CRC functions using Altera IP cores). The CRC circuitry generates an EDCRC code at the data storage input of the RAM, and checks the CRC code at the output of the RAM. If an SEU affects any stored bits in the internal memory, the CRC automatically corrects the error when it is read from the memory. The ECC-enabled memory can report the occurrence of a single-bit flip, or adjacent double-bit and adjacent triple-bit flips, and will correct single- and double-bit flips. Adjacent triple-bit corruptions are detected and reported using a status bit, but not corrected.

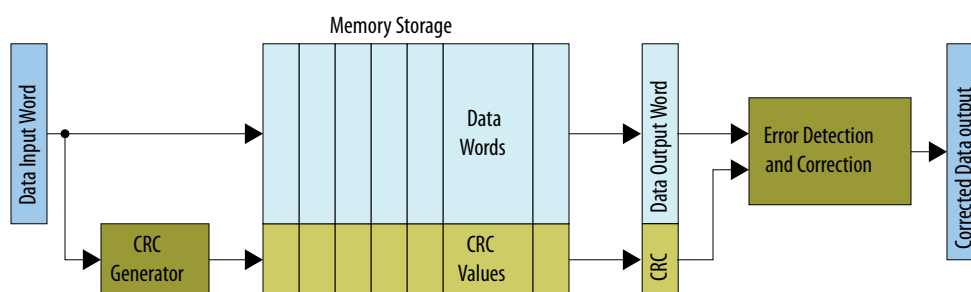
## Configuring the ECCRAM

You must configure the ECCRAM as a 2-port RAM (with independent read and write addresses). Use of these features does not reduce the amount of available logic.

While the CRC checking function results in some additional output delay, the hard ECC has a much higher  $f_{MAX}$  compared with an equivalent soft ECC implemented in general logic. Additionally, the hard IP can be pipelined in the M20K block by configuring the ECCRAM to use an output register at the corrected data output port. This increases performance while adding latency.

For devices without dedicated circuitry, you can implement the ECC by instantiating the ECC generation and checking functions as the IP core ALTECC.

Figure 15-1: Memory Storage

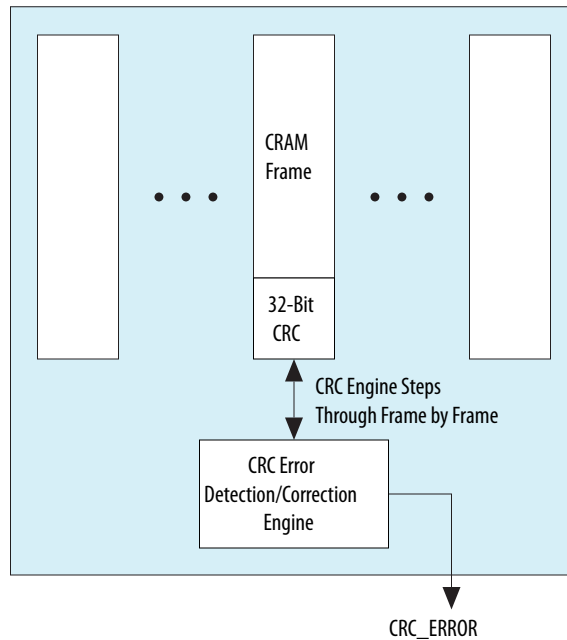


## Mitigating SEU Effects in Configuration RAM

Use EDCRC to detect and correct soft errors in CRAM. These EDCRC blocks are similar to those that protect internal user memory.

CRAM is organized into frames. The size of the frame and the number of frames is device specific. CRAM frames are continually checked for errors by loading each frame into a data register. The EDCRC block checks the frame for errors. Soft errors found trigger the assertion of a CRC\_ERROR pin on the device. Monitor this pin in your system. Take appropriate actions when this pin is asserted, indicating a soft error was detected in the configuration RAM.

Figure 15-2: CRAM Frame

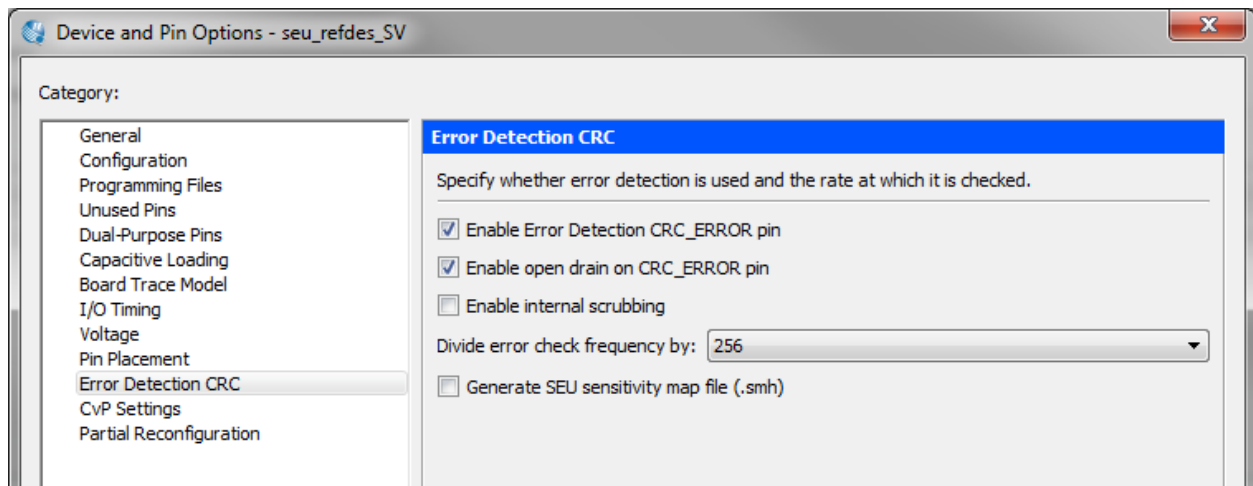


**Related Information**  
[Single Event Upsets](#)

## Scanning CRAM Frames

To enable the Quartus II software to scan CRAM frames, turn on **Enable Error Detection CRC\_ERROR** pin in the **Device and Pin Options** dialog box (**Assignments > Device > Device and Pin Options**).

Figure 15-3: Enable Error Detection CRC\_ERROR Pin



To enable the CRC\_ERROR pin as an open drain output, turn on **Enable open drain on CRC\_ERROR pin**.

To guarantee the availability of a clock, the EDCRC function operates on an independent clock generated internally on the FPGA itself. To enable EDCRC operation on a divided version of the clock select a value from the **Divide error check frequency by** value.

## Internal Scrubbing

Arria V, Cyclone V (including SoC devices), Stratix V, and later device families support automatic CRAM error correction, without resorting to the original CRAM contents from an external copy of the original SRAM Object File.

Automatic correction is possible because EDCRC calculates and stores redundancy fields along with the configuration bits. This automatic correction is known as scrubbing.

To enable internal scrubbing, turn on **Enable internal scrubbing** option in the **Device and Pin Options** dialog box.

If the Quartus II software finds a CRC error in a CRAM frame, the frame is reconstructed from the error correcting code calculated for that frame, and then the corrected frame is re-written into the CRAM.

**Note:** If you enable internal scrubbing, you must still plan a recovery sequence. Although scrubbing can restore the CRAM array to intended configuration, latency occurs between the soft error detection and correction. Because of the large number of configuration bits to be scanned, this latency may be up to 100 milliseconds for large devices. Therefore, the FPGA may operate with errors during that period.

### Related Information

[Error Detection CRC Page](#)

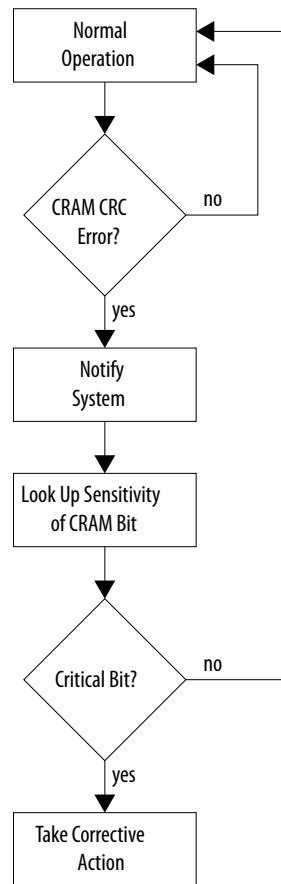
## Understanding SEU Sensitivity

Reconfiguring a running FPGA typically has a significant impact on the system using the FPGA. When planning for SEU recovery, account for the time required to bring the FPGA to a state consistent with the current state of the system. For example, if an internal state machine is in an illegal state, it may require reset. Also, the surrounding logic may need to account for this unexpected operation.

Often an SEU impacts CRAM bits not used by the implemented design. Many configuration bits are not used because they control logic and routing wires that are not used in a design. Depending on the implementation, 40% of all CRAM bits can be used even in the most heavily utilized devices. This means that only 40% of SEU events require intervention, and you can ignore 60% of SEU events.

You may determine that portions of the implemented design are not critical to the FPGA's function. Examples may include test circuitry implemented but not important to the operation of the device, or other non-critical functions that may be logged but do not need to be reprogrammed or reset.

Figure 15-4: Sensitivity Processing Flow



The ratio of SEU strikes versus functional interrupts is the Single Event Functional Interrupt (SEFI) ratio. Minimizing this ratio improves SEU mitigation.

**Related Information**

[Understanding Single Event Functional Interrupts in FPGA Designs](#)

## Designating the Sensitivity of your Design Hierarchy

The design hierarchy sensitivity processing depends on the contents of the Sensitivity Map Header File (.smf). This file determines the correct (least disruptive) recovery sequence for any CRAM bit flip. The .smf designates the sensitivity of each portion of the FPGA's logic design.

To generate the .smf, you must designate the sensitivity of the design from a functional logic view, using the hierarchy tagging procedure.

### Hierarchy Tagging

Hierarchy tagging is the process of classifying the sensitivity of the portions of your design.

The Quartus II software performs hierarchy tagging by creating a design partition, and then assigning the parameter ASD<sub>Region</sub> to that partition. The parameter can assume a value from 0 to 255, so there are 256

different classifications of system responses to the portions of your design. This sensitivity information is encoded into the `.smf` the running system uses to look-up the sensitivity of an SEU upset, and to perform the appropriate action to that CRAM location.

#### Related Information

[Altera Advanced SEU Detection IP Core User Guide](#)

## Altera Advanced SEU Detection IP Core

You must instantiate the Altera Advanced SEU Detection IP core to enable SEU detection and correction features.

When the EDCRC function detects an SEU, the Altera Advanced SEU Detection IP core determines the designer-designated sensitivity of that CRAM bit by looking up the sensitivity in the `.smf`.

When an EDCRC block detects an SEU, a sensitivity processor looks up the sensitivity of the affected CRAM bit in the `.smf`.

The user determines which version of the IP core to instantiate: on-chip or external. If the Altera Advanced SEU Detection IP core is configured for on-chip sensitivity processing, the IP core performs the lookup with the user-supplied memory interface. If the Altera Advanced SEU Detection IP core is configured for off-chip sensitivity processing, it notifies external logic (typically via a system CPU interrupt request), and provides cached event message register values to the off-chip sensitivity processor. The SMH information is stored in the external sensitivity processor's memory system.

#### Related Information

[Altera Advanced SEU Detection IP Core User Guide](#)

## On-Chip Sensitivity Processor

You can use the Advanced SEU Detection IP core to implement an on-chip sensitivity processor. The IP core interacts with user-supplied external memory access logic to read the Sensitivity Map Header file, stored on external memory.

Once it determines the sensitivity of the affected CRAM bit, the IP core can assert a Critical Error signal so the system provides an appropriate response. If the SEU is not critical, the Critical Error signal may be left un-asserted.

On-chip sensitivity processing is autonomous: the FPGA device determines whether it is affected by an SEU, without the need for external logic. However, this requires part of the FPGA's logic resources for the external memory interface.

#### Related Information

[Altera Advanced SEU Detection IP Core User Guide](#)

## External Sensitivity Processor

You can configure the Advanced SEU Detection IP core for use with an external sensitivity processor. In this case an external CPU, such as the ARM processor in Altera's SoC devices, receives an interrupt request when the FPGA detects an SEU. The CPU then reads the Error Message Register, and performs

the sensitivity lookup by referring to the Sensitivity Map Header file (**.smf**) stored in the CPU's memory space.

External sensitivity processing does not require on-board memory dedicated to the SMH storage function. Also, this technique relieves the FPGA of external memory interface requirements, along with the memory storage requirements for the sensitivity map itself. If a CPU is already present in the system, external sensitivity processing may be the more hardware-efficient way to implement sensitivity lookup.

#### Related Information

[Altera Advanced SEU Detection IP Core User Guide](#)

## Triple-Module Redundancy

If your system must suffer no downtime due to SEUs, consider Triple Module Redundancy as an SEU mitigation strategy.

Triple-Module-Redundancy (TMR) is an established technique for improving hardware fault tolerance. In TMR, three identical instances of hardware are supplied, along with voting hardware at the output of the hardware. If an SEU affects one of the instances, the voting logic notes the majority in a vote of the separate instances of the module to mask out any malfunctioning module.

The advantage of TMR is that there is no downtime in the case of a single SEU; if a module is found to be in faulty operation, that module can be scrubbed of its error by reprogramming it. The error detection and correction time is many orders of magnitude less than the MTBF due to SEU events. Therefore, you can repair a soft interrupt before another SEU affects another instance in the TMR triple.

The disadvantage of TMR is its extreme cost in hardware resources: it requires three times as much hardware, in addition to voting logic. This hardware cost can be minimized by judiciously implementing TMR only for the most critical part of the design.

There are several automated ways to generate TMR designs by automatically replicating designated functions and synthesizing the required voting logic. Synthesis vendors offering automated TMR synthesis include Synopsys and Mentor Graphics.

## Recovering from a Single-Event Upset

After correcting a bit flip in CRAM, the device is in its original configuration with respect to logic and routing. However, the internal state of the FPGA may be illegal.

The state of the device may be invalid because it may have been operating while SEUs corrupted its configuration. The errors from faulty operation may have propagated elsewhere within the FPGA or to the system outside the FPGA.

Forcing the FPGA into a known state is system dependent. Determining the possible outcomes from SEU, and designing a recovery response to SEU should be part of the FPGA and system design process.

## Evaluating Your System's Response to Functional Upsets

Because SEUs can randomly strike any memory element, system testing is especially important to ensure a comprehensive recovery response.

The Quartus II software includes the Fault Injection Debugger to aid in SEU recovery response. This feature is available for the Arria V, Cyclone V, and Stratix V device families.

The feature is available from the Quartus II GUI or at the command line. You must instantiate the Altera Fault Injection IP core into your FPGA design to use this feature. The IP core flips a CRAM bit by dynamically reconfiguring the frame containing that CRAM bit, flipping it to its opposite state.

The Fault Injection Debugger allows you to operate the FPGA in your system and inject random CRAM bit flips to test the ability of the FPGA and the system to detect and recover fully from an SEU. You should be able to observe your FPGA and your system recover from these simulated SEU strikes. You can then refine your FPGA and system recovery sequence by observing these strikes.

If you have recorded an SEU in the device's Error Message Register, the Fault Injection Debugger also allows you to specify a targeted fault to be injected (rather than inject the fault in a random location). This feature is available only from the command line.

### Related Information

[Debugging Single Event Upsets Using the Fault Injection Debugger](#)

## Document Revision History

Table 15-1: Document Revision History

Date	Version	Changes
June 2014	2014.06.30	<ul style="list-style-type: none"> <li>Updated formatting.</li> <li>Added "Mitigating SEU Effects in Embedded User RAM" section.</li> <li>Added "Altera Advanced SEU Detection IP Core" section.</li> </ul>
November 2012	2012.11.01	Preliminary release.

2014.12.15

QII5V1



Subscribe



Send Feedback

As programmable logic designs become more complex and require increased performance, advanced synthesis becomes an important part of a design flow. The Altera® Quartus® II software includes advanced Integrated Synthesis that fully supports VHDL, Verilog HDL, and Altera-specific design entry languages, and provides options to control the synthesis process. With this synthesis support, the Quartus II software provides a complete, easy-to-use solution.

### Related Information

[Recommended HDL Coding Styles](#) on page 12-1

For examples of Verilog HDL and VHDL code synthesized for specific logic functions

[Designing With Low-Level Primitives User Guide](#)

For more information about coding with primitives that describe specific low-level functions in Altera devices

## Design Flow

The Quartus II Analysis & Synthesis stage of the compilation flow runs Integrated Synthesis, which fully supports Verilog HDL, VHDL, and Altera-specific languages, and major features of the SystemVerilog language.

In the synthesis stage of the compilation flow, the Quartus II software performs logic synthesis to optimize design logic and performs technology mapping to implement the design logic in device resources such as logic elements (LEs) or adaptive logic modules (ALMs), and other dedicated logic blocks. The synthesis stage generates a single project database that integrates all your design files in a project (including any netlists from third-party synthesis tools).

You can use Analysis & Synthesis to perform the following compilation processes:

**Table 16-1: Compilation Process**

Compilation Process	Description
Analyze Current File	Parses your current design source file to check for syntax errors. This command does not report many semantic errors that require further design synthesis. To perform this analysis, on the Processing menu, click <b>Analyze Current File</b> .

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered





Compilation Process	Description
<b>Analysis &amp; Elaboration</b>	Checks your design for syntax and semantic errors and performs elaboration to identify your design hierarchy. To perform Analysis & Elaboration, on the Processing menu, point to <b>Start</b> , and then click <b>Start Analysis &amp; Elaboration</b> .
<b>Hierarchy Elaboration</b>	Parses HDL designs and generates a skeleton of hierarchies. Hierarchy Elaboration is similar to the Analysis & Elaboration flow, but without any elaborated logic, thus making it much faster to generate.
<b>Analysis &amp; Synthesis</b>	Performs complete Analysis & Synthesis on a design, including technology mapping. To perform Analysis & Synthesis, on the Processing menu, point to <b>Start</b> , and then click <b>Start Analysis &amp; Synthesis</b> .

**Related Information**

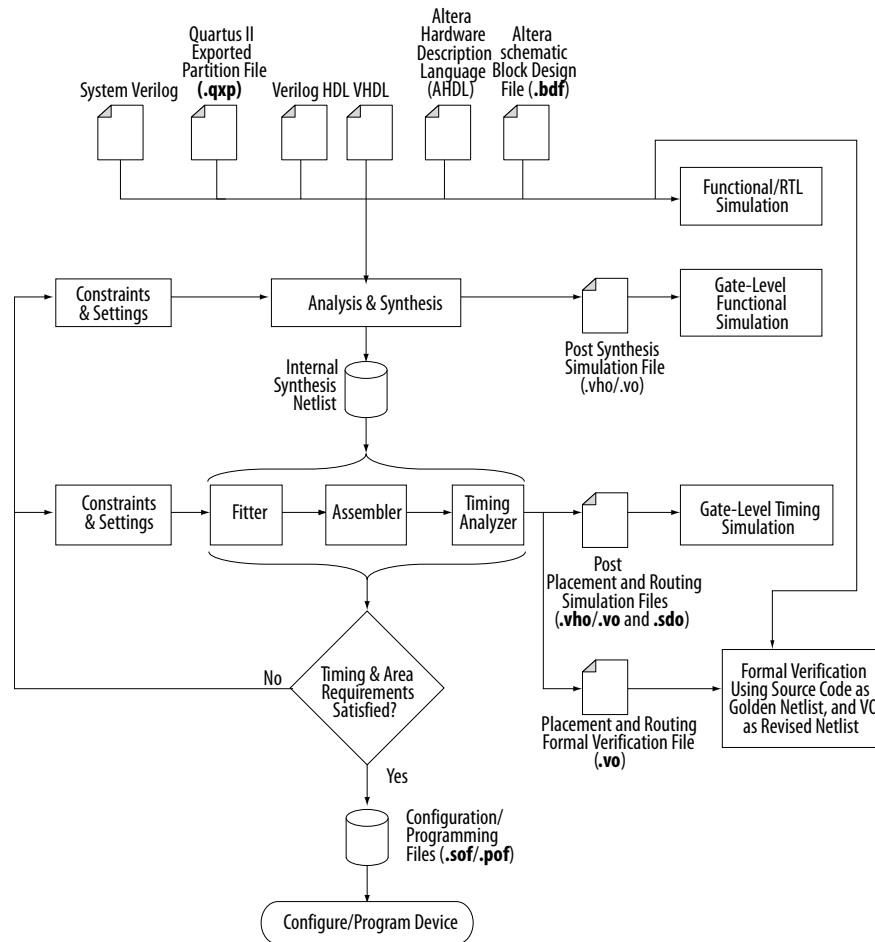
[Language Support](#) on page 16-4

**[Start Hierarchy Elaboration Command Processing Menu](#)**

For more information about the Hierarchy Elaboration flow

## Quartus II Integrated Synthesis Design and Compilation Flow

Figure 16-1: Basic Design Flow Using Quartus II Integrated Synthesis



The Quartus II Integrated Synthesis design and compilation flow consists of the following steps:

1. Create a project in the Quartus II software and specify the general project information, including the top-level design entity name.
2. Create design files in the Quartus II software or with a text editor.
3. On the Project menu, click **Add/Remove Files in Project** and add all design files to your Quartus II project using the **Files** page of the **Settings** dialog box.
4. Specify Compiler settings that control the compilation and optimization of your design during synthesis and fitting.
5. Add timing constraints to specify the timing requirements.
6. Compile your design. To synthesize your design, on the Processing menu, point to **Start**, and then click **Start Analysis & Synthesis**. To run a complete compilation flow including placement, routing, creation of a programming file, and timing analysis, click **Start Compilation** on the Processing menu.
7. After obtaining synthesis and placement and routing results that meet your requirements, program or configure your Altera device.

Integrated Synthesis generates netlists that enable you to perform functional simulation or gate-level timing simulation, timing analysis, and formal verification.

**Related Information**

- **Quartus II Synthesis Options** on page 16-21  
For more information about synthesis settings
- **Incremental Compilation** on page 16-20  
For more information about partitioning your design to reduce compilation time
- **Quartus II Exported Partition File as Source** on page 16-21  
For more information about using **.qxp** as a design source file
- **Introduction to the Quartus II Software**  
For an overall summary of features in the Quartus II software
- **Managing Files in a Project**  
For more information about Quartus II projects
- **About Compilation Flows**  
For more information about Quartus II the compilation flow

## Language Support

Quartus II Integrated Synthesis supports HDL. You can specify the Verilog HDL or VHDL language version in your design.

To ensure that the Quartus II software reads all associated project files, add each file to your Quartus II project by clicking **Add/Remove Files in Project** on the Project menu. You can add design files to your project. You can mix all supported languages and netlists generated by third-party synthesis tools in a single Quartus II project.

**Related Information**

- **Design Libraries** on page 16-12  
Describes how to compile and reference design units in custom libraries
- **Using Parameters/Generics** on page 16-15  
Describes how to use parameters or generics and pass them between languages
- **Insert Template Dialog Box**  
For more information about using the available templates in the Quartus II Text Editor for various Verilog and VHDL features

## Verilog HDL Support

The Quartus II Compiler's Analysis & Synthesis module supports the following Verilog HDL standards:

- Verilog-1995 (IEEE Standard 1364-1995)
- Verilog-2001 (IEEE Standard 1364-2001)
- SystemVerilog-2005 (IEEE Standard 1800-2005) (the Compiler does not support all constructs)

The Verilog HDL code samples provided in this document follow the Verilog-2001 standard unless otherwise specified. The Quartus II Compiler uses the Verilog-2001 standard by default for files that have the extension **.v**, and the SystemVerilog standard for files that have the extension **.sv**.

If you use scripts to add design files, you can use the `-HDL_VERSION` command to specify the HDL version for each design file.

The Quartus II software support for Verilog HDL is case sensitive in accordance with the Verilog HDL standard. The Quartus II software supports the compiler directive ``define`, in accordance with the Verilog HDL standard.

The Quartus II software supports the `include` compiler directive to include files with absolute paths (with either “/” or “\” as the separator), or relative paths. When searching for a relative path, the Quartus II software initially searches relative to the project directory. If the Quartus II software cannot find the file, the software then searches relative to all user libraries and then relative to the directory location of the current file.

#### Related Information

- [About Verilog HDL](#)  
For more information about Verilog HDL
- [Quartus II Verilog HDL Support](#)  
For more information about Quartus II Verilog HDL support
- [Specifying Verilog Input Settings](#)  
For more information about specifying a default Verilog HDL version for all files
- [verilog\\_input\\_version Synthesis Directive](#)  
For more information about controlling the Verilog HDL version that compiles your design in a design file with the `VERILOG_INPUT_VERSION` synthesis directive
- [Verilog HDL Synthesis Attributes and Directives](#)  
For more information about Verilog HDL synthesis attributes and directives
- [Adding an HDL File to a Project and Setting the HDL Version](#) on page 16-76
- [Synthesis Directives](#) on page 16-25  
For more information about specifying synthesis directives

## Verilog HDL Configuration

Verilog HDL configuration is a set of rules that specify the source code for particular instances.

Verilog HDL configuration allows you to perform the following tasks:

- Specify a library search order for resolving cell instances (as does a library mapping file)
- Specify overrides to the logical library search order for specified instances
- Specify overrides to the logical library search order for all instances of specified cells

For more information about these tasks, refer to [Table 16-2](#).

### Configuration Syntax

A Verilog HDL configuration contains the following statements:

```
config config_identifier;  
design [library_identifier.]cell_identifier;  
config_rule_statement;  
endconfig
```

Where:

- `config`—the keyword that begins the configuration.
- `config_identifier`—the name you enter for the configuration.
- `design`—the keyword that starts a design statement for specifying the top of the design.
- `[library_identifier.]cell_identifier`—specifies the top-level module (or top-level modules) in the design and the logical library for this module (modules).
- `config_rule_statement`—one or more of the following clauses: `default`, `instance`, or `cell`. For more information, refer to [Table 16-2](#).
- `endconfig`—the keyword that ends a configuration.

**Table 16-2: Type of Clauses for the `config_rule_statement` Keyword**

Clause Type	Description
default	<p>Specifies the logical libraries to search to resolve a default cell instance. A default cell instance is an instance in the design that is not specified in a subsequent instance or cell clause in the configuration.</p> <p>You specify these libraries with the <code>liblist</code> keyword. The following is an example of a default clause: <code>default liblist lib1 lib2;</code></p> <p>Also specifies resolving default instances in the logical libraries (<code>lib1</code> and <code>lib2</code>).</p> <p>Because libraries are inherited, some simulators (for example, VCS) also search the default (or current) library as well after the searching the logical libraries (<code>lib1</code> and <code>lib2</code>).</p>
instance	<p>Specifies a specific instance. The specified instance clause depends on the use of the following keywords:</p> <ul style="list-style-type: none"> <li>• <code>liblist</code>—specifies the logical libraries to search to resolve the instance.</li> <li>• <code>use</code>—specifies that the instance is an instance of the specified cell in the specified logical library.</li> </ul> <p>The following are examples of <code>instance</code> clauses:</p> <pre>instance top.dev1 liblist lib1 lib2;</pre> <p>This instance clause specifies to resolve <code>instance top.dev1</code> with the cells assigned to logical libraries <code>lib1</code> and <code>lib2</code>;</p> <pre>instance top.dev1.gm1 use lib2.gizmult;</pre> <p>This instance clause specifies that <code>top.dev1.gm1</code> is an instance of the cell named <code>gizmult</code> in logical library <code>lib2</code>.</p>
cell	<p>A <code>cell</code> clause is similar to an <code>instance</code> clause, except that the <code>cell</code> clause specifies all instances of a cell definition instead of specifying a particular instance. What it specifies depends on the use of the <code>liblist</code> or <code>use</code> keywords:</p> <ul style="list-style-type: none"> <li>• <code>liblist</code>—specifies the logical libraries to search to resolve all instances of the cell.</li> <li>• <code>use</code>—the specified cell's definition is in the specified library.</li> </ul>

## Hierarchical Configurations

A design can have more than one configuration. For example, you can define a configuration that specifies the source code you use in particular instances in a sub hierarchy, then define a configuration for a higher level of the design.

Suppose, for example, a sub hierarchy of a design is an eight-bit adder and the RTL Verilog code describes the adder in a logical library named `rtlLib` and the gate-level code describes the adder in a logical library named `gateLib`.

If you want to use the gate-level code for the 0 (zero) bit of the adder and the RTL level code for the other seven bits, the configuration might appear as shown in the following example:

```
config cfg1;
design aLib.eight_adder;
default liblist rtlLib;
instance adder.fulladd0 liblist gateLib;
endconfig
```

If you are instantiating this eight-bit adder eight times to create a 64-bit adder, use configuration `cfg1` for the first instance of the eight-bit adder, but not in any other instance. A configuration that would perform this function is shown in the following example:

```
config cfg2;
design bLib.64_adder;
default liblist bLib;
instance top.64add0 use work.cfg1:config;
endconfig
```

**Note:** The name of the unbound module may be different than the name of the cell that is bounded to the instance.

### Suffix :config

To distinguish between a module by the same name, use the optional extension `:config` to refer to configuration names. For example, you can always refer to a `cfg2` configuration as `cfg2:config` (even if the `cfg2` module does not exist).

## SystemVerilog Support

The Quartus II software supports the SystemVerilog constructs.

**Note:** Designs written to support the Verilog-2001 standard might not compile with the SystemVerilog setting because the SystemVerilog standard has several new reserved keywords.

### Related Information

- [Quartus II Support for SystemVerilog](#)  
For more information about the supported SystemVerilog constructs
- [Quartus II Support for Verilog 2001](#)  
For more information about the supported Verilog-2001 features

## Initial Constructs and Memory System Tasks

The Quartus II software infers power-up conditions from the Verilog HDL `initial` constructs. The Quartus II software also creates power-up settings for variables, including RAM blocks. If the Quartus II software encounters nonsynthesizable constructs in an `initial` block, it generates an error.

To avoid such errors, enclose nonsynthesizable constructs (such as those intended only for simulation) in `translate_off` and `translate_on` synthesis directives

Synthesis of initial constructs enables the power-up state of the synthesized design to match the power-up state of the original HDL code in simulation.

**Note:** Initial blocks do not infer power-up conditions in some third-party EDA synthesis tools. If you convert between synthesis tools, you must set your power-up conditions correctly.

Quartus II Integrated Synthesis supports the `$readmemb` and `$readmemh` system tasks to initialize memories.

This example shows an initial construct that initializes an inferred RAM with `$readmemb`.

### Example 16-1: Verilog HDL Code: Initializing RAM with the `readmemb` Command

```
reg [7:0] ram[0:15];
initial
begin
  $readmemb("ram.txt", ram);
end
```

When creating a text file to use for memory initialization, specify the address using the format `@<location>` on a new line, and then specify the memory word such as `110101` or `abcde` on the next line.

The following example shows a portion of a Memory Initialization File (.mif) for the RAM in [Example 16-1](#).

### Example 16-2: Text File Format: Initializing RAM with the `readmemb` Command

```
@0
00000000
@1
00000001
@2
00000010
...
@e
00001110
@f
00001111
```

#### Related Information

- [Translate Off and On / Synthesis Off and On](#) on page 16-58
- [Power-Up Level](#) on page 16-36

## Verilog HDL Macros

The Quartus II software fully supports Verilog HDL macros, which you can define with the `'define` compiler directive in your source code. You can also define macros in the Quartus II software or on the command line.

## Setting a Verilog HDL Macro Default Value in the Quartus II Software

To specify a macro in the Quartus II software, follow these steps:

1. Click **Assignments > Settings > Compiler Settings > Verilog HDL Input**
2. Under **Verilog HDL macro**, type the macro name in the **Name** box and the value in the **Setting** box.
3. Click **Add**.

## Setting a Verilog HDL Macro Default Value on the Command Line

To set a default value for a Verilog HDL macro on the command line, use the `--verilog_macro` option:

```
quartus_map <Design name> --verilog_macro= "<Macro name>=<Macro setting>"
```

The command in this example has the same effect as specifying ``define a 2` in the Verilog HDL source code:

```
quartus_map my_design --verilog_macro="a=2"
```

To specify multiple macros, you can repeat the option more than once.

```
quartus_map my_design --verilog_macro="a=2" --verilog_macro="b=3"
```

## VHDL Support

The Quartus II Compiler's Analysis & Synthesis module supports the following VHDL standards:

- VHDL 1987 (IEEE Standard 1076-1987)
- VHDL 1993 (IEEE Standard 1076-1993)
- VHDL 2008 (IEEE Standard 1076-2008)

The Quartus II Compiler uses the VHDL 1993 standard by default for files that have the extension `.vhdl` or `.vhd`.

**Note:** The VHDL code samples follow the VHDL 1993 standard.

To specify a default VHDL version for all files, follow these steps:

1. Click **Assignments > Settings > Compiler Settings > Verilog HDL Input**
2. On the **VHDL Input** page, under **VHDL version**, select the appropriate version, and then click **OK**.

To override the default VHDL version for each VHDL design file, follow these steps:

3. On the Project menu, click **Add/Remove Files in Project**.
4. On the **Files** page, select the appropriate file in the list, and then click **Properties**.
5. In the HDL version list, select **VHDL\_2008**, **VHDL\_1993**, or **VHDL\_1987**, and then click **OK**.

You can also specify the VHDL version that compiles your design for each design file with the `VHDL_INPUT_VERSION` synthesis directive. This directive overrides the default HDL version and any HDL version specified in the **File Properties** dialog box.

**Table 16-3: Controlling the VHDL Input Version with a Synthesis Directive**

HDL	Code
VHDL	<code>--synthesis VHDL_INPUT_VERSION &lt;language version&gt;</code>



HDL	Code
VHDL-2008	<pre>/* synthesis VHDL_INPUT_VERSION &lt;language version&gt; */</pre>

The variable <language version> requires one of the following values:

- VHDL\_1987
- VHDL\_1993
- VHDL\_2008

When the Quartus II software reads a VHDL\_INPUT\_VERSION synthesis directive, it changes the current language version as specified until after the file or until it reaches the next VHDL\_INPUT\_VERSION directive.

**Note:** You cannot change the language version in a VHDL design unit.

If you use scripts to add design files, you can use the -HDL\_VERSION command to specify the HDL version for each design file.

The Quartus II software reads default values for registered signals defined in the VHDL code and converts the default values into power-up level settings. This enables the power-up state of the synthesized design to match, as closely as possible, the power-up state of the original HDL code in simulation.

#### Related Information

- [Synthesis Directives](#) on page 16-25  
For more information about specifying synthesis directives
- [Adding an HDL File to a Project and Setting the HDL Version](#) on page 16-76  
For more information about using the -HDL\_VERSION command to specify the HDL version for each design file
- [Power-Up Level](#) on page 16-36  
For more information about power-up level

## VHDL-2008 Support

The Quartus II software contains support for VHDL 2008 with constructs defined in the IEEE Standard 1076-2008 version of the *IEEE Standard VHDL Language Reference Manual*.

#### Related Information

#### [Quartus II Support for VHDL 2008](#)

For more information about the Quartus II software support for VHDL-2008

## VHDL Standard Libraries and Packages

The Quartus II software includes the standard IEEE libraries and several vendor-specific VHDL libraries.

The IEEE library includes the standard VHDL packages `std_logic_1164`, `numeric_std`, `numeric_bit`, and `math_real`. The STD library is part of the VHDL language standard and includes the packages

standard (included in every project by default) and `textio`. For compatibility with older designs, the Quartus II software also supports the following vendor-specific packages and libraries:

- Synopsys packages such as `std_logic_arith` and `std_logic_unsigned` in the IEEE library
- Mentor Graphics® packages such as `std_logic_arith` in the ARITHMETIC library
- Altera primitive packages `altera_primitives_components` (for primitives such as `GLOBAL` and `DFFE`) and `maxplus2` (for legacy support of MAX+PLUS® II primitives) in the ALTERA library
- Altera IP core packages `altera_mf_components` and `stratixgx_mf_components` in the ALTERA\_MF library (for Altera-specific IP cores including LCELL), and `lpm_components` in the LPM library for library of parameterized modules (LPM) functions.

**Note:** Altera recommends that you import component declarations for Altera primitives such as `GLOBAL` and `DFFE` from the `altera_primitives_components` package and not the `altera_mf_components` package.

#### Related Information

- [Design Libraries](#) on page 16-12  
For information about organizing your own design units into custom libraries

## VHDL wait Constructs

The Quartus II software supports one VHDL `wait until` statement per process block. However, the Quartus II software does not support other VHDL wait constructs, such as `wait for` and `wait on` statements, or processes with multiple `wait` statements.

The following shows the `wait until` construct example for VHDL:

```
architecture dff_arch of ls_dff is
begin
  output: process begin
    wait until (CLK'event and CLK='1');
    Q <= D;
    Qbar <= not D;
  end process output;
end dff_arch;
```

## AHDL Support

The Quartus II Compiler's Analysis & Synthesis module fully supports the Altera Hardware Description Language (AHDL).

AHDL designs use Text Design Files (**.tdf**). You can import AHDL Include Files (**.inc**) into a **.tdf** with an AHDL `include` statement. Altera provides **.inc** files for all IP cores shipped with the Quartus II software.

**Note:** The AHDL language does not support the synthesis directives or attributes.

#### Related Information

##### [About AHDL](#)

For more information about AHDL

## Schematic Design Entry Support

The Quartus II Compiler's Analysis & Synthesis module fully supports **.bdf** for schematic design entry.

**Note:** Schematic entry methods do not support the synthesis directives or attributes.

#### Related Information

##### [About Schematic Design Entry](#)

For information about creating and editing schematic designs

## State Machine Editor

The Quartus II software supports graphical state machine entry. To create a new finite state machine (FSM) design, on the File menu, click **New**. In the **New** dialog box, expand the **Design Files** list, and then select **State Machine File**.

#### Related Information

##### [About the State Machine Editor](#)

For more information about the State Machine Editor

## Design Libraries

By default, the Quartus II software compiles all design files into the work library. If you do not specify a design library, if a file refers to a library that does not exist, or if the referenced library does not contain a referenced design unit, the Quartus II software searches the work library. This behavior allows the Quartus II software to compile most designs with minimal setup, but you have the option of creating separate custom design libraries.

To compile your design files into specific libraries (for example, when you have two or more functionally different design entities that share the same name), you can specify a destination library for each design file in various ways, as described in the following:

- [Specifying a Destination Library Name in the Settings Dialog Box](#) on page 16-13
- [Specifying a Destination Library Name in the Quartus II Settings File or with Tcl](#) on page 16-13

When the Quartus II Compiler analyzes the file, it stores the analyzed design units in the destination library of the file.

**Note:** A design can contain two or more entities with the same name if the Quartus II software compiles the entities into separate libraries.

When compiling a design instance, the Quartus II software initially searches for the entity in the library associated with the instance (which is the work library if you do not specify any library). If the Quartus II software could not locate the entity definition, the software searches for a unique entity definition in all design libraries. If the Quartus II software finds more than one entity with the same name, the software generates an error. If your design uses multiple entities with the same name, you must compile the entities into separate libraries.

In VHDL, you can associate an instance with an entity in several ways, as described in [Mapping a VHDL Instance to an Entity in a Specific Library](#) on page 16-14.

In Verilog HDL, BDF schematic entry, AHDL, VQM and EDIF netlists, you can use different libraries for each of the entities that have the same name, and compile the instantiation into the same library as the appropriate entity.

#### Related Information

[Mapping a VHDL Instance to an Entity in a Specific Library](#) on page 16-14

## Specifying a Destination Library Name in the Settings Dialog Box

To specify a library name for one of your design files, follow these steps:

1. On the Assignments menu, click **Settings**.
2. In the **Category** list, select **Files**.
3. Select the file in the **File Name** list.
4. Click **Properties**.
5. In the **File Properties** dialog box, select the type of design file from the **Type** list.
6. Type the library name in the **Library** field.
7. Click **OK**.

## Specifying a Destination Library Name in the Quartus II Settings File or with Tcl

You can specify the library name with the `-library` option to the `<language type>_FILE` assignment in the Quartus II Settings File (`.qsf`) or with Tcl commands.

For example, the following assignments specify that the Quartus II software analyzes the `my_file.vhd` and stores its contents (design units) in the VHDL library `my_lib`, and then analyzes the Verilog HDL file `my_header_file.h` and stores its contents in a library called `another_lib`.

```
set_global_assignment -name VHDL_FILE my_file.vhd -library my_lib
set_global_assignment -name VERILOG_FILE my_header_file.h -library another_lib
```

### Related Information

- [Scripting Support](#) on page 16-75  
For more information about Tcl scripting

## Specifying a Destination Library Name in a VHDL File

You can use the `library` synthesis directive to specify a library name in your VHDL source file. This directive takes the name of the destination library as a single string argument. Specify the `library` directive in a VHDL comment before the context clause for a primary design unit (that is, a package declaration, an entity declaration, or a configuration), with one of the supported keywords for synthesis directives, that is, `altera`, `synthesis`, `pragma`, `synopsys`, or `exemplar`.

The `library` directive overrides the default library destination `work`, the library setting specified for the current file in the **Settings** dialog box, any existing `.qsf` setting, any setting made through the Tcl interface, or any prior `library` directive in the current file. The directive remains effective until the end of the file or the next `library` synthesis directive.

The following example uses the `library` synthesis directive to create a library called `my_lib` containing the `my_entity` design unit:

```
-- synthesis library my_lib
library ieee;
use ieee.std_logic_1164.all;
entity my_entity(...)
end entity my_entity;
```

**Note:** You can specify a single destination library for all your design units in a given source file by specifying the library name in the **Settings** dialog box, editing the `.qsf`, or using the Tcl interface. To organize your design units in a single file

into different libraries rather than just a single library, you can use the `library` directive to change the destination VHDL library in a source file.

The Quartus II software generates an error if you use the `library` directive in a design unit.

### Related Information

- [Synthesis Directives](#) on page 16-25  
For more information about specifying synthesis directives

## Mapping a VHDL Instance to an Entity in a Specific Library

The VHDL language provides several ways to map or bind an instance to an entity in a specific library.

### Direct Entity Instantiation

In the direct entity instantiation method, the instantiation refers to an entity in a specific library.

The following shows the direct entity instantiation method for VHDL:

```
entity entity1 is
port(...);
end entity entity1;
architecture arch of entity1 is
begin
inst: entity lib1.foo
port map(...);
end architecture arch;
```

### Component Instantiation—Explicit Binding Instantiation

You can bind a component to an entity in several mechanisms. In an explicit binding indication, you bind a component instance to a specific entity.

The following shows the binding instantiation method for VHDL:

```
entity entity1 is
port(...);
end entity entity1;
package components is
component entity1 is
port map (...);
end component entity1;
end package components;
entity top_entity is
port(...);
end entity top_entity;
use lib1.components.all;
architecture arch of top_entity is
-- Explicitly bind instance I1 to entity1 from lib1
for I1: entity1 use entity lib1.entity1
port map(...);
end for;
begin
I1: entity1 port map(...);
end architecture arch;
```

## Component Instantiation—Default Binding

If you do not provide an explicit binding indication, the Quartus II software binds a component instance to the nearest visible entity with the same name. If no such entity is visible in the current scope, the Quartus II software binds the instance to the entity in the library in which you declare the component. For example, if you declare the component in a package in the `MY_LIB` library, an instance of the component binds to the entity in the `MY_LIB` library.

The code examples in the following examples show this instantiation method:

### Example 16-3: VHDL Code: Default Binding to the Entity in the Same Library as the Component Declaration

```
use mylib.pkg.foo; -- import component declaration from package "pkg" in

                                -- library "mylib"
architecture rtl of top
...
begin
-- This instance will be bound to entity "foo" in library "mylib"
inst: foo
port map(...);
end architecture rtl;
```

### Example 16-4: VHDL Code: Default Binding to the Directly Visible Entity

```
use mylib.foo; -- make entity "foo" in library "mylib" directly visible
architecture rtl of top
component foo is
generic (...)
port (...);
end component;
begin
-- This instance will be bound to entity "foo" in library "mylib"
inst: foo
port map(...);
end architecture rtl;
```

## Using Parameters/Generics

The Quartus II software supports parameters (known as generics in VHDL) and you can pass these parameters between design languages.

Click **Assignments > Settings > Compiler Settings > Default Parameters** to enter default parameter values for your design. In AHDL, the Quartus II software inherits parameters, so any default parameters apply to all AHDL instances in your design. You can also specify parameters for instantiated modules in a **.bdf**. To specify parameters in a **.bdf** instance, double-click the parameter value box for the instance symbol, or right-click the symbol and click **Properties**, and then click the **Parameters** tab.

You can specify parameters for instantiated modules in your design source files with the provided syntax for your chosen language. Some designs instantiate entities in a different language; for example, they might instantiate a VHDL entity from a Verilog HDL design file. You can pass parameters or generics between VHDL, Verilog HDL, AHDL, and BDF schematic entry, and from EDIF or VQM to any of these

languages. You do not require an additional procedure to pass parameters from one language to another. However, sometimes you must specify the type of parameter you are passing. In those cases, you must follow certain guidelines to ensure that the Quartus II software correctly interprets the parameter value.

#### Related Information

- [Setting Default Parameter Values and BDF Instance Parameter Values](#) on page 16-16  
For more information about the GUI-based entry methods, the interpretation of parameter values, and format recommendations
- [Passing Parameters Between Two Design Languages](#) on page 16-17  
For more information about parameter type rules

## Setting Default Parameter Values and BDF Instance Parameter Values

Default parameter values and BDF instance parameter values do not have an explicitly declared type. Usually, the Quartus II software can correctly infer the type from the value without ambiguity. For example, the Quartus II software interprets "ABC" as a string, 123 as an integer, and 15.4 as a floating-point value. In other cases, such as when the instantiated subdesign language is VHDL, the Quartus II software uses the type of the parameter, generic, or both in the instantiated entity to determine how to interpret the value, so that the Quartus II software interprets a value of 123 as a string if the VHDL parameter is of a type string. In addition, you can set the parameter value in a format that is legal in the language of the instantiated entity. For example, to pass an unsized bit literal value from **.bdf** to Verilog HDL, you can use `'1` as the parameter value, and to pass a 4-bit binary vector from **.bdf** to Verilog HDL, you can use `4'b1111` as the parameter value.

In a few cases, the Quartus II software cannot infer the correct type of parameter value. To avoid ambiguity, specify the parameter value in a type-encoded format in which the first or first and second characters of the parameter indicate the type of the parameter, and the rest of the string indicates the value in a quoted sub-string. For example, to pass a binary string 1001 from **.bdf** to Verilog HDL, you cannot use the value 1001, because the Quartus II software interprets it as a decimal value. You also cannot use the string "1001" because the Quartus II software interprets it as an ASCII string. You must use the type-encoded string `B"1001"` for the Quartus II software to correctly interpret the parameter value.

This table lists valid parameter strings and how the Quartus II software interprets the parameter strings. Use the type-encoded format only when necessary to resolve ambiguity.

**Table 16-4: Valid Parameter Strings and Interpretations**

Parameter String	Quartus II Parameter Type, Format, and Value
S"abc", s"abc"	String value abc
"abc123", "123abc"	String value abc123 or 123abc
F"12.3", f"12.3"	Floating point number 12.3
-5.4	Floating point number -5.4
D"123", d"123"	Decimal number 123
123, -123	Decimal number 123, -123
X"ff", H"ff"	Hexadecimal value FF

Parameter String	Quartus II Parameter Type, Format, and Value
Q"77", O"77"	Octal value 77
B"1010", b"1010"	Unsigned binary value 1010
SB"1010", sb"1010"	Signed binary value 1010
R"1", R"0", R"X", R"Z", r"1", r"0", r"X", r"Z"	Unsigned bit literal
E"apple", e"apple"	Enumeration type, value name is apple
P"1 unit"	Physical literal, the value is (1, unit)
A(...), a(...)	Array type or record type. The string (...) determines the array type or record type content

You can select the parameter type for global parameters or global constants with the pull-down list in the **Parameter** tab of the **Symbol Properties** dialog box. If you do not specify the parameter type, the Quartus II software interprets the parameter value and defines the parameter type. You must specify parameter type with the pull-down list to avoid ambiguity.

**Note:** If you open a **.bdf** in the Quartus II software, the software automatically updates the parameter types of old symbol blocks by interpreting the parameter value based on the language-independent format. If the Quartus II software does not recognize the parameter value type, the software sets the parameter type as **untyped**.

The Quartus II software supports the following parameter types:

- **Unsigned Integer**
- **Signed Integer**
- **Unsigned Binary**
- **Signed Binary**
- **Octal**
- **Hexadecimal**
- **Float**
- **Enum**
- **String**
- **Boolean**
- **Char**
- **Untyped/Auto**

## Passing Parameters Between Two Design Languages

When passing a parameter between two different languages, a design block that is higher in the design hierarchy instantiates a lower-level subdesign block and provides parameter information. The subdesign language (the design entity that you instantiate) must correctly interpret the parameter. Based on the information provided by the higher-level design and the value format, and sometimes by the parameter type of the subdesign entity, the Quartus II software interprets the type and value of the passed parameter.

When passing a parameter whose value is an enumerated type value or literal from a language that does not support enumerated types to one that does (for example, from Verilog HDL to VHDL), you must



ensure that the enumeration literal is in the correct spelling in the language of the higher-level design block (block that is higher in the hierarchy). The Quartus II software passes the parameter value as a string literal, and the language of the lower-level design correctly convert the string literal into the correct enumeration literal.

If the language of the lower-level entity is SystemVerilog, you must ensure that the `enum` value is in the correct case. In SystemVerilog, two enumeration literals differ in more than just case. For example, `enum {item, ITEM}` is not a good choice of item names because these names can create confusion and is more difficult to pass parameters from case-insensitive HDLs, such as VHDL.

Arrays have different support in different design languages. For details about the array parameter format, refer to the **Parameter** section in the Analysis & Synthesis Report of a design that contains array parameters or generics.

The following code shows examples of passing parameters from one design entry language to a subdesign written in another language.

**Table 16-5: VHDL Parameterized Subdesign Entity**

This table shows a VHDL subdesign that you instantiate in a top-level Verilog HDL design in [Table 16-6](#).

HDL	Code
VHDL	<pre> type fruit is (apple, orange, grape); entity vhd_sub is generic ( name : string := "default", width : integer := 8, number_string : string := "123", f : fruit := apple, binary_vector : std_logic_vector(3 downto 0) := "0101", signed_vector : signed (3 downto 0) := "1111"); </pre>

**Table 16-6: Verilog HDL Top-Level Design Instantiating and Passing Parameters to VHDL Entity**

This table shows a Verilog HDL Top-Level Design Instantiating and Passing Parameters to VHDL Entity from [Table 16-5](#).

HDL	Code
Verilog HDL	<pre> vhd_sub inst (...); defparam inst.name = "lower"; defparam inst.width = 3; defparam inst.num_string = "321"; defparam inst.f = "grape"; // Must exactly match enum value defparam inst.binary_vector = 4'b1010; defparam inst.signed_vector = 4'sb1010; </pre>

**Table 16-7: Verilog HDL Parameterized Subdesign Module**

This table shows a Verilog HDL subdesign that you instantiate in a top-level VHDL design in [Table 16-8](#).

HDL	Code
Verilog HDL	<pre> module veri_sub (...) parameter name = "default"; parameter width = 8; parameter number_string = "123"; parameter binary_vector = 4'b0101; parameter signed_vector = 4'sb1111; </pre>

**Table 16-8: VHDL Top-Level Design Instantiating and Passing Parameters to the Verilog HDL Module**

This table shows a VHDL Top-Level Design Instantiating and Passing Parameters to the Verilog HDL Module from [Table 16-7](#).


HDL	Code
VHDL	<pre> inst:veri_sub generic map ( name =&gt; "lower", width =&gt; 3, number_string =&gt; "321" binary_vector = "1010" signed_vector = "1010") </pre>

To use an HDL subdesign such as the one shown in [Table 16-7](#) in a top-level **.bdf** design, you must generate a symbol for the HDL file, as shown in [Figure 16-2](#). Open the HDL file in the Quartus II software, and then, on the File menu, point to **Create/Update**, and then click **Create Symbol Files for Current File**.

To specify parameters on a **.bdf** instance, double-click the parameter value box for the instance symbol, or right-click the symbol and click **Properties**, and then click the **Parameters** tab. Right-click the symbol and click **Update Design File from Selected Block** to pass the updated parameter to the HDL file.

**Figure 16-2: BDF Top-Level Design Instantiating and Passing Parameters to the Verilog HDL Module**

This figure shows BDF Top-Level Design Instantiating and Passing Parameters to the Verilog HDL Module from [Table 16-7](#)



## Incremental Compilation

Incremental compilation manages a design hierarchy for incremental design by allowing you to divide your design into multiple partitions. Incremental compilation ensures that the Quartus II software resynthesizes only the updated partitions of your design during compilation, to reduce the compilation time and the runtime memory usage. The feature maintains node names during synthesis for all registered and combinational nodes in unchanged partitions. You can perform incremental synthesis by setting the netlist type for all design partitions to **Post-Synthesis**.

You can also preserve the placement and routing information for unchanged partitions. This feature allows you to preserve performance of unchanged blocks in your design and reduces the time required for placement and routing, which significantly reduces your design compilation time.

### Related Information

- [About Incremental Compilation](#)  
For more information about incremental compilation
- [Best Practices for Incremental Compilation Partitions and Floorplan Assignments](#)  
For more information about incremental compilation best practices and floorplan assignments
- [Quartus II Incremental Compilation for Hierarchical and Team-Based Design](#) on page 3-1  
For more information about incremental compilation for hierarchical and team-based design

## Partitions for Preserving Hierarchical Boundaries

A design partition represents a portion of your design that you want to synthesize and fit incrementally.

If you want to preserve the **Optimization Technique** and **Restructure Multiplexers** logic options in any entity, you must create new partitions for the entity instead of using the **Preserve Hierarchical Boundary** logic option. If you have settings applied to specific existing design hierarchies, particularly those created in the Quartus II software versions before 9.0, you must create a design partition for the design hierarchy so that synthesis can optimize the design instance independently and preserve the hierarchical boundaries.

**Note:** The **Preserve Hierarchical Boundary** logic option is available only in Quartus II software versions 8.1 and earlier. Altera recommends using design partitions if you want to preserve hierarchical boundaries through the synthesis and fitting process, because incremental compilation maintains the hierarchical boundaries of design partitions.

## Parallel Synthesis

The **Parallel Synthesis** logic option reduces compilation time for synthesis. The option enables the Quartus II software to use multiple processors to synthesize multiple partitions in parallel.

This option is available when you perform the following tasks:

- Specify the maximum number of processors allowed under **Parallel Compilation** options in the **Compilation Process Settings** page of the **Settings** dialog box.
- Enable the incremental compilation feature.
- Use two or more partitions in your design.
- Turn on the **Parallel Synthesis** option.

By default, the Quartus II software enables the **Parallel Synthesis** option. To disable parallel synthesis, click **Assignments > Settings > Compiler Settings > Advanced Settings (Synthesis) > Parallel Synthesis**.

You can also set the **Parallel Synthesis** option with the following Tcl command:

```
set_global_assignment -name parallel_synthesis off
```

If you use the command line, you can differentiate among the interleaved messages by turning on the **Show partition that generated the message** option in the Messages page. This option shows the partition ID in parenthesis for each message.

You can view all the interleaved messages from different partitions in the Messages window. The **Partition** column in the Messages window displays the partition ID of the partition referred to in the message. After compilation, you can sort the messages by partition.

#### Related Information

##### [About the Messages Window](#)

For more information about displaying the Partition column

## Quartus II Exported Partition File as Source

You can use a **.qxp** as a source file in incremental compilation. The **.qxp** contains the precompiled design netlist exported as a partition from another Quartus II project, and fully defines the entity. Project team members or intellectual property (IP) providers can use a **.qxp** to send their design to the project lead, instead of sending the original HDL source code. The **.qxp** preserves the compilation results and instance-specific assignments. Not all global assignments can function in a different Quartus II project. You can override the assignments for the entity in the **.qxp** by applying assignments in the top-level design.

#### Related Information

##### [Quartus II Exported Partition File .qxp](#)

For more information about **.qxp**

##### [Quartus II Incremental Compilation for Hierarchical and Team-Based Design](#) on page 3-1

For more information about exporting design partitions and using **.qxp** files

## Quartus II Synthesis Options

The Quartus II software offers several options to help you control the synthesis process and achieve optimal results for your design.

**Note:** When you apply a Quartus II Synthesis option globally or to an entity, the option affects all lower-level entities in the hierarchy path, including entities instantiated with Altera and third-party IP.

#### Related Information

- [Setting Synthesis Options](#) on page 16-22

Describes the **Compiler Settings** page of the **Settings** dialog box, in which you can set the most common global settings and options, and defines the following types of synthesis options: Quartus II logic options, synthesis attributes, and synthesis directives.

## Setting Synthesis Options

You can set synthesis options in the **Settings** dialog box, or with logic options in the Quartus II software, or you can use synthesis attributes and directives in your HDL source code.

The **Compiler Settings** page of the **Settings** dialog box allows you to set global synthesis options that apply to the entire project. You can also use a corresponding Tcl command.

You can set some of the advanced synthesis settings in the **Advanced Settings** dialog box on the **Compiler Settings** page.

### Related Information

#### [Netlist Optimizations and Physical Synthesis](#)

For more information about Physical Synthesis options

## Quartus II Logic Options

The Quartus II logic options control many aspects of the synthesis and placement and routing process. To set logic options in the Quartus II software, on the Assignments menu, click **Assignment Editor**. You can also use a corresponding Tcl command to set global assignments. The Quartus II logic options enable you to set instance or node-specific assignments without editing the source HDL code.

### Related Information

#### [About the Assignment Editor](#)

For more information about using the Assignment Editor

## Synthesis Attributes

The Quartus II software supports synthesis attributes for Verilog HDL and VHDL, also commonly called pragmas. These attributes are not standard Verilog HDL or VHDL commands. Synthesis tools use attributes to control the synthesis process. The Quartus II software applies the attributes in the HDL source code, and attributes always apply to a specific design element. Some synthesis attributes are also available as Quartus II logic options via the Quartus II software or scripting. Each attribute description indicates a corresponding setting or a logic option that you can set in the Quartus II software. You can specify only some attributes with HDL synthesis attributes.

Attributes specified in your HDL code are not visible in the Assignment Editor or in the **.qsf**. Assignments or settings made with the Quartus II software, the **.qsf**, or the Tcl interface take precedence over assignments or settings made with synthesis attributes in your HDL code. The Quartus II software generates warning messages if the software finds invalid attributes, but does not generate an error or stop the compilation. This behavior is necessary because attributes are specific to various design tools, and attributes not recognized in the Quartus II software might be for a different EDA tool. The Quartus II software lists the attributes specified in your HDL code in the Source assignments table of the Analysis & Synthesis report.

The Verilog-2001, SystemVerilog, and VHDL language definitions provide specific syntax for specifying attributes, but in Verilog-1995, you must embed attribute assignments in comments. You can enter attributes in your code using the syntax in [Specifying Synthesis Attributes in Verilog-1995](#) on page 1-23 through [Synthesis Attributes in VHDL](#) on page 1-24, in which *<attribute>*, *<attribute type>*, *<value>*, *<object>*, and *<object type>* are variables, and the entry in brackets is optional. These examples demonstrate each syntax form.

**Note:** Verilog HDL is case sensitive; therefore, synthesis attributes in Verilog HDL files are also case sensitive.

In addition to the `synthesis` keyword shown above, the Quartus II software supports the `pragma`, `synopsys`, and `exemplar` keywords for compatibility with other synthesis tools. The software also supports the `altera` keyword, which allows you to add synthesis attributes that the Quartus II Integrated Synthesis feature recognizes and not by other tools that recognize the same synthesis attribute.

**Note:** Because formal verification tools do not recognize the `exemplar`, `pragma`, and `altera` keywords, avoid using these attribute keywords when using formal verification.

#### Related Information

- [Maximum Fan-Out](#) on page 16-42  
For more information about maximum fan-out attribute
- [Preserve Registers](#) on page 16-38  
For more information about `preserve` attribute

### Synthesis Attributes in Verilog-1995

You must use Verilog-1995 comment-embedded attributes as a suffix to the declaration of an item and must appear before a semicolon, when a semicolon is necessary.

**Note:** You cannot use the open one-line comment in Verilog HDL when a semicolon is necessary after the line, because it is not clear to which HDL element that the attribute applies. For example, you cannot make an attribute assignment such as `reg r; // synthesis <attribute>` because the Quartus II software could read the attribute as part of the next line.

### Specifying Synthesis Attributes in Verilog-1995

The following show an example of specifying synthesis attributes in Verilog-1995:

```
// synthesis <attribute> [ = <value> ]  
or  
/* synthesis <attribute> [ = <value> ] */
```

### Applying Multiple Attributes to the Same Instance in Verilog-1995

To apply multiple attributes to the same instance in Verilog-1995, separate the attributes with spaces.

```
//synthesis <attribute1> [ = <value> ] <attribute2> [ = <value> ]
```

For example, to set the `maxfan` attribute to 16 and set the `preserve` attribute on a register called `my_reg`, use the following syntax:

```
reg my_reg /* synthesis maxfan = 16 preserve */;
```

#### Related Information

- [Maximum Fan-Out](#) on page 16-42  
For more information about maximum fan-out attribute
- [Preserve Registers](#) on page 16-38  
For more information about `preserve` attribute

## Synthesis Attributes in Verilog-2001

You must use Verilog-2001 attributes as a prefix to a declaration, module item, statement, or port connection, and as a suffix to an operator or a Verilog HDL function name in an expression.

**Note:** Formal verification does not support the Verilog-2001 attribute syntax because the tools do not recognize the syntax.

## Specifying Synthesis Attributes in Verilog-2001 and SystemVerilog

```
(* <attribute> [ = <value> ] *)
```

## Applying Multiple Attributes

To apply multiple attributes to the same instance in Verilog-2001 or SystemVerilog, separate the attributes with commas.

```
(* <attribute1> [ = <value1> ], <attribute2> [ = <value2> ] *)
```

For example, to set the `maxfan` attribute to 16 and set the `preserve` attribute on a register called `my_reg`, use the following syntax:

```
(* maxfan = 16, preserve *) reg my_reg;
```

## Related Information

- [Maximum Fan-Out](#) on page 16-42  
For more information about maximum fan-out attribute
- [Preserve Registers](#) on page 16-38  
For more information about `preserve` attribute

## Synthesis Attributes in VHDL

VHDL attributes declare and apply the attribute type to the object you specify.

## Synthesis Attributes in VHDL

The following shows the synthesis attributes example in VHDL:

```
attribute <attribute> : <attribute type> ;
attribute <attribute> of <object> : <object type> is <value>;
```

## altera\_syn\_attributes

The Quartus II software defines and applies each attribute separately to a given node. For VHDL designs, the software declares all supported synthesis attributes in the `altera_syn_attributes` package in the Altera library. You can call this library from your VHDL code to declare the synthesis attributes:

```
LIBRARY altera;
USE altera.altera_syn_attributes.all;
```

## Synthesis Directives

The Quartus II software supports synthesis directives, also commonly called compiler directives or pragmas. You can include synthesis directives in Verilog HDL or VHDL code as comments. These directives are not standard Verilog HDL or VHDL commands. Synthesis tools use directives to control the synthesis process. Directives do not apply to a specific design node, but change the behavior of the synthesis tool from the point in which they occur in the HDL source code. Other tools, such as simulators, ignore these directives and treat them as comments.

**Table 16-9: Specifying Synthesis Directives**

You can enter synthesis directives in your code using the syntax in the following table, in which *<directive>* and *<value>* are variables, and the entry in brackets are optional. For synthesis directives, no equal sign before the value is necessary; this is different than the Verilog syntax for synthesis attributes. The examples demonstrate each syntax form.

Language	Syntax Example
Verilog HDL <sup>(10)</sup>	<pre>// synthesis &lt;directive&gt; [ &lt;value&gt; ] or /* synthesis &lt;directive&gt; [ &lt;value&gt; ] */</pre>
VHDL	<pre>-- synthesis &lt;directive&gt; [ &lt;value&gt; ]</pre>
VHDL-2008	<pre>/* synthesis &lt;directive&gt; [&lt;value&gt;] */</pre>

In addition to the `synthesis` keyword shown above, the software supports the `pragma`, `synopsys`, and `exemplar` keywords in Verilog HDL and VHDL for compatibility with other synthesis tools. The Quartus II software also supports the keyword `altera`, which allows you to add synthesis directives that only Quartus II Integrated Synthesis feature recognizes, and not by other tools that recognize the same synthesis directives.

**Note:** Because formal verification tools ignore the `exemplar`, `pragma`, and `altera` keywords, Altera recommends that you avoid using these directive keywords when you use formal verification to prevent mismatches with the Quartus II results.

## Optimization Technique

The **Optimization Technique** logic option specifies the goal for logic optimization during compilation; that is, whether to attempt to achieve maximum speed performance or minimum area usage, or a balance between the two.

### Related Information

#### [Optimization Technique logic option](#)

For more information about the Optimization Technique logic option

<sup>(10)</sup> Verilog HDL is case sensitive; therefore, all synthesis directives are also case sensitive.



## Auto Gated Clock Conversion

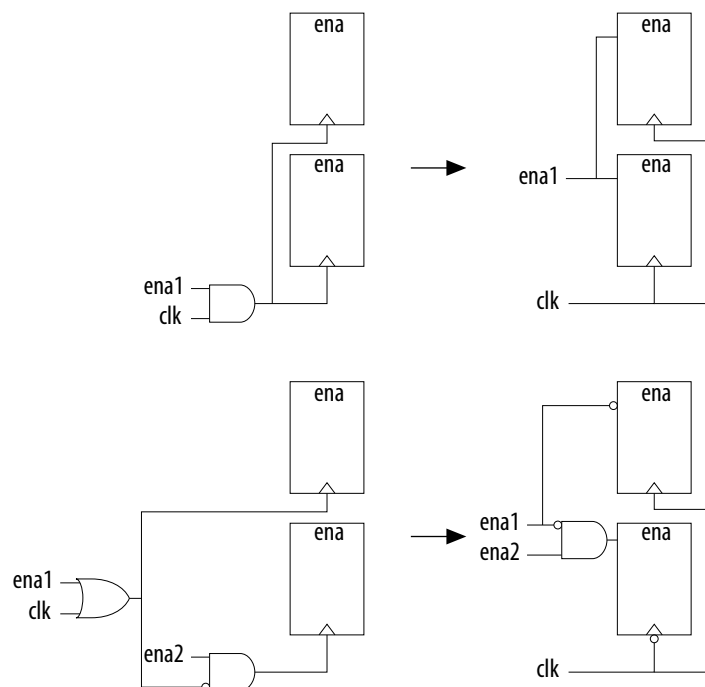
Clock gating is a common optimization technique in ASIC designs to minimize power consumption. You can use the **Auto Gated Clock Conversion** logic option to optimize your prototype ASIC designs by converting gated clocks into clock enables when you use FPGAs in your ASIC prototyping. The automatic conversion of gated clocks to clock enables is more efficient than manually modifying source code. The **Auto Gated Clock Conversion** logic option automatically converts qualified gated clocks (base clocks as defined in the Synopsys Design Constraints [SDC]) to clock enables. Click **AssignmentsSettingsCompiler SettingsAdvanced Settings (Synthesis)** to enable **Auto Gated Clock Conversion**.

The gated clock conversion occurs when all these conditions are met:

- Only one base clock drives a gated-clock
- For one set of gating input values, the value output of the gated clock remains constant and does not change as the base clock changes
- For one value of the base clock, changes in the gating inputs do not change the value output for the gated clock

The option supports combinational gates in clock gating network.

**Figure 16-3: Example Gated Clock Conversion**



**Note:** This option does not support registers in RAM, DSP blocks, or I/O related WYSIWYG primitives. Because the gated-clock conversion cannot trace the base clock from the gated clock, the gated clock conversion does not support multiple design partitions from incremental compilation in which the gated clock and base clock are not in the same hierarchical partition. A gated clock tree, instead of every gated clock, is the basis of each conversion. Therefore, if you cannot convert a

gated clock from a root gated clock of a multiple cascaded gated clock, the conversion of the entire gated clock tree fails.

The **Info** tab in the Messages window lists all the converted gated clocks. You can view a list of converted and nonconverted gated clocks from the Compilation Report under the **Optimization Results** of the Analysis & Synthesis Report. The **Gated Clock Conversion Details** table lists the reasons for nonconverted gated clocks.

#### Related Information

##### [Auto Gated Clock Conversion logic option](#)

For more information about Auto Gated Clock Conversion logic option and a list of supported devices

## Timing-Driven Synthesis

The **Timing-Driven Synthesis** logic option specifies whether Analysis & Synthesis should use the SDC timing constraints of your design to better optimize the circuit. When you turn on this option, Analysis & Synthesis runs timing analysis to obtain timing information about the netlist, and then considers the SDC timing constraints to focus on critical portions of your design when optimizing for performance, while optimizing noncritical portions for area. When you turn on this option, Analysis & Synthesis also protects SDC constraints by not merging duplicate registers that have incompatible timing constraints.

When you turn on the **Timing-Driven Synthesis** logic option, Analysis & Synthesis increases performance by improving logic depth on critical portions of your design, and improving area on noncritical portions of your design. The increased performance affects the amount of area used, specifically adaptive look-up tables (ALUTs) and registers in your design. Depending on how much of your design is timing critical, overall area can increase or decrease when you turn on the **Timing-Driven Synthesis** logic option. Runtime and peak memory use increases slightly if you turn on the **Timing-Driven Synthesis** logic option.

When you turn on the **Timing-Driven Synthesis** logic option, the **Optimization Technique** logic option has the following effect. With **Optimization Technique Speed**, Timing-Driven Synthesis optimizes timing-critical portions of your design for performance at the cost of increasing area (logic and register utilization). With an **Optimization Technique** of **Balanced**, Timing-Driven Synthesis also optimizes the timing-critical portions of your design for performance, but the option allows only limited area increase. With **Optimization Technique Area**, Timing-Driven Synthesis optimizes your design only for area. **Timing-Driven Synthesis** prevents registers with incompatible timing constraints from merging for any **Optimization Technique** setting. If your design contains multiple partitions, you can select **Timing-Driven Synthesis** unique options for each partition. If you use a **.qxp** as a source file, or if your design uses partitions developed in separate Quartus II projects, the software cannot properly compute timing of paths that cross the partition boundaries.

Even with the **Optimization Technique** logic option set to **Speed**, the **Timing-Driven Synthesis** option still considers the resource usage in your design when increasing area to improve timing. For example, the **Timing-Driven Synthesis** option checks if a device has enough registers before deciding to implement the shift registers in logic cells instead of RAM for better timing performance.

When using incremental compilation, Integrated Synthesis allows each partition to use up all the registers in a device. You can use the **Maximum Number of LABs** settings to specify the number of LABs that every partition can use. If your design has only one partition, you can also use the **Maximum Number of LABs** settings to limit the number of resources that your design can use. This limitation is useful when you add more logic to your design.

To turn on or turn off the **Timing-Driven Synthesis** logic option, follow these steps:

1. Click **Assignments > Settings > Compiler Settings > Advanced Settings (Synthesis)**.
2. Turn on or turn off **Timing-Driven Synthesis**.

**Note:** Altera recommends that you select a specific device for timing-driven synthesis to have the most accurate timing information. When you select auto device, timing-driven synthesis uses the smallest device for the selected family to obtain timing information.

#### Related Information

##### [Timing-Driven Synthesis logic option](#)

For more information about Timing-Driven Synthesis logic option and a list of supported devices

##### [SDC Constraint Protection](#) on page 16-28

For more information about SDC constraint protection

## SDC Constraint Protection

The **SDC Constraint Protection** option specifies whether Analysis & Synthesis should protect registers from merging when they have incompatible timing constraints. For example, when you turn on this option, the software does not merge two registers that are duplicates of each other but have different multicycle constraints on them. When you turn on the **Timing-Driven Synthesis** option, the software detects registers with incompatible constraints, and you do not need to turn on **SDC Constraint Protection**. Click **Assignments > Settings > Compiler Settings > Advanced Settings (Synthesis)** to enable the **SDC constraint protection** option.

## PowerPlay Power Optimization

The **PowerPlay Power Optimization** logic option controls the power-driven compilation setting of Analysis & Synthesis and determines how aggressively Analysis & Synthesis optimizes your design for power.

#### Related Information

- [PowerPlay Power Optimization logic option](#)  
For more information about the available settings for the PowerPlay power optimization logic option and a list of supported devices
- [Power Optimization](#)  
For more information about optimizing your design for power utilization
- [PowerPlay Power Analysis](#)  
For information about analyzing your power results

## Limiting Resource Usage in Partitions

Resource balancing is important when performing Analysis & Synthesis. During resource balancing, Quartus II Integrated Synthesis considers the amount of used and available DSP and RAM blocks in the device, and tries to balance these resources to prevent no-fit errors.

For DSP blocks, Resource balancing is important when performing Analysis & Synthesis. During resource balancing, Quartus II Integrated Synthesis considers the amount of used and available DSP and RAM blocks in the device, and tries to balance these resources to prevent no-fit errors. resource balancing converts the remaining DSP blocks to equivalent logic if there are more DSP blocks in your design that

the software can place in the device. For RAM blocks, resource balancing converts RAM blocks to different types of RAM blocks if there are not enough blocks of a certain type available in the device; however, Quartus II Integrated Synthesis does not convert RAM blocks to logic.

**Note:** The RAM balancing feature does not support Stratix V devices because Stratix V has only M20K memory blocks.

By default, Quartus II Integrated Synthesis considers the information in the targeted device to identify the number of available DSP or RAM blocks. However, in incremental compilation, each partition considers the information in the device independently and consequently assumes that the partition has all the DSP and RAM blocks in the device available for use, resulting in over allocation of DSP or RAM blocks in your design, which means that the total number of DSP or RAM blocks used by all the partitions is greater than the number of DSP or RAM blocks available in the device, leading to a no-fit error during the fitting process.

#### Related Information

- [Creating LogicLock Regions](#) on page 16-29  
For more information about preventing a no-fit error during the fitting process
- [Using Assignments to Limit the Number of RAM and DSP Blocks](#) on page 16-29  
For more information about preventing a no-fit error during the fitting process

## Creating LogicLock Regions

The floorplan-aware synthesis feature allows you to use LogicLock regions to define resource allocation for DSP blocks and RAM blocks. For example, if you assign a certain partition to a certain LogicLock region, resource balancing takes into account that all the DSP and RAM blocks in that partition need to fit in this LogicLock region. Resource balancing then balances the DSP and RAM blocks accordingly.

Because floorplan-aware balancing step considers only one partition at a time, it does not know that nodes from another partition may be using the same resources. When using this feature, Altera recommends that you do not manually assign nodes from different partitions to the same LogicLock region.

If you do not want the software to consider the LogicLock floorplan constraints when performing DSP and RAM balancing, you can turn off the floorplan-aware synthesis feature. Click **Assignments > Settings > Compiler Settings > Advanced Settings (Synthesis)** to disable **Use LogicLock Constraints During Resource Balancing** option.

#### Related Information

- [Quartus II Incremental Compilation for Hierarchical and Team-Based Design](#) on page 3-1  
For more information about using LogicLock regions to create a floorplan for incremental compilation

## Using Assignments to Limit the Number of RAM and DSP Blocks

For DSP and RAM block balancing, you can use assignments to limit the maximum number of blocks that the balancer allows. You can set these assignments globally or on individual partitions. For DSP block balancing, the **Maximum DSP Block Usage** logic option allows you to specify the maximum number of DSP blocks that the DSP block balancer assumes are available for the current partition. For RAM blocks, the floorplan-aware logic option allows you to specify maximum resources for different RAM types, such as **Maximum Number of M4K/M9K/M20K/M10K Memory Blocks**, **Maximum Number of M512 Memory Blocks**, **Maximum Number of M-RAM/M144K Memory Blocks**, or **Maximum Number of LABs**.

The partition-specific assignment overrides the global assignment, if any. However, each partition that does not have a partition-specific assignment uses the value set by the global assignment, or the value derived from the device size if no global assignment exists. This action can also lead to over allocation. Therefore, Altera recommends that you always set the assignment on each partition individually.

To select the **Maximum Number <block type> Memory Blocks** option or the **Maximum DSP Block Usage** option globally, click **Assignments > Settings > Compiler Settings > Advanced Settings (Synthesis)**. You can use the Assignment Editor to set this assignment on a partition by selecting the assignment, and setting it on the root entity of a partition. You can set any positive integer as the value of this assignment. If you set this assignment on a name other than a partition root, Analysis & Synthesis gives an error.

#### Related Information

- **Maximum DSP Block Usage logic option**  
For more information about the **Maximum DSP Block Usage** logic option, including a list of supported device families
- **Maximum Number of M4K/M9K/M20K/M10K Memory Blocks logic option**  
For more information about the **Maximum Number of M4K/M9K/M20K/M10K Memory Blocks** logic option, including a list of supported device families
- **Maximum Number of M512 Memory Blocks logic option**  
For more information about the **Maximum Number of M512 Memory Blocks** logic option, including a list of supported device families
- **Maximum Number of M-RAM/144K Memory Blocks logic option**  
For more information about **Maximum Number of M-RAM/144K Memory Blocks** logic option, including a list of supported device families
- **Maximum Number of LABs logic option**  
For more information about the **Maximum Number of LABs** logic option, including a list of supported device families

## Restructure Multiplexers

The **Restructure Multiplexers** logic option restructures multiplexers to create more efficient use of area, allowing you to implement multiplexers with a reduced number of LEs or ALMs.

When multiplexers from one part of your design feed multiplexers in another part of your design, trees of multiplexers form. Multiplexers may arise in different parts of your design through Verilog HDL or VHDL constructs such as the “if,” “case,” or “?:” statements. Multiplexer buses occur most often as a result of multiplexing together arrays in Verilog HDL, or `STD_LOGIC_VECTOR` signals in VHDL. The **Restructure Multiplexers** logic option identifies buses of multiplexer trees that have a similar structure. This logic option optimizes the structure of each multiplexer bus for the target device to reduce the overall amount of logic in your design.

Results of the multiplexer optimizations are design dependent, but area reductions as high as 20% are possible. The option can negatively affect your design's  $f_{MAX}$ .

#### Related Information

- **Recommended HDL Coding Styles** on page 12-1  
For more information about optimizing for multiplexers

- **Analysis Synthesis Optimization Results Reports**  
For more information about the Multiplexer Restructuring Statistics report table for each bus of multiplexers
- **Restructure Multiplexers logic option**  
For more information about the Restructure Multiplexers logic option, including the settings and a list of supported device families

## Synthesis Effort

The **Synthesis Effort** logic option specifies the overall synthesis effort level in the Quartus II software.

### Related Information

#### [Synthesis Effort logic option](#)

For more information about Synthesis Effort logic option, including a list of supported device families

## Fitter Initial Placement Seed

The **Fitter Initial Placement Seed** option specifies the seed that Synthesis uses to randomly run synthesis in a slightly different way. You can use this seed when your design is close to meeting requirements, to get a slightly different result. The seeds that produce the best result for a design might change if your design changes.

To set the **Synthesis Seed** option, click **Assignments > Settings > Compiler Settings > Advanced Settings (Fitter)**. The default value is **1**. You can specify a positive integer value.

## State Machine Processing

The **State Machine Processing** logic option specifies the processing style to synthesize a state machine.

The default state machine encoding, **Auto**, uses one-hot encoding for FPGA devices and minimal-bits encoding for CPLDs. These settings achieve the best results on average, but another encoding style might be more appropriate for your design, so this option allows you to control the state machine encoding.

For one-hot encoding, the Quartus II software does not guarantee that each state has one bit set to one and all other bits set to zero. Quartus II Integrated Synthesis creates one-hot register encoding with standard one-hot encoding and then inverts the first bit. This results in an initial state with all zero values, and the remaining states have two 1 values. Quartus II Integrated Synthesis encodes the initial state with all zeros for the state machine power-up because all device registers power up to a low value. This encoding has the same properties as true one-hot encoding: the software recognizes each state by the value of one bit. For example, in a one-hot-encoded state machine with five states, including an initial or reset state, the software uses the following register encoding:

State 0	0	0	0	0	0
State 1	0	0	0	1	1
State 2	0	0	1	0	1
State 3	0	1	0	0	1
State 4	1	0	0	0	1

If you set the **State Machine Processing** logic option to **User-Encoded** in a Verilog HDL design, the software starts with the original design values for the state constants. For example, a Verilog HDL design can contain the following declaration:

```
parameter S0 = 4'b1010, S1 = 4'b0101, ...
```

If the software infers the states `S0`, `S1`, ... the software uses the encoding `4'b1010`, `4'b0101`, ... . If necessary, the software inverts bits in a user-encoded state machine to ensure that all bits of the reset state of the state machine are zero.

**Note:** You can view the state machine encoding from the Compilation Report under the State Machines of the Analysis & Synthesis Report. The State Machine Viewer displays only a graphical representation of the state machines as interpreted from your design.

To assign your own state encoding with the **User-Encoded** setting of the **State Machine Processing** option in a VHDL design, you must apply specific binary encoding to the elements of an enumerated type because enumeration literals have no numeric values in VHDL. Use the `syn_encoding` synthesis attribute to apply your encoding values.

#### Related Information

- [Manually Specifying State Assignments Using the `syn\_encoding` Attribute](#) on page 16-32
- [Recommended HDL Coding Styles](#) on page 12-1  
For guidelines on how to correctly infer and encode your state machine
- [Analyzing Designs with Quartus II Netlist Viewers](#)  
For more information about the State Machine Viewer
- [State Machine Processing logic option](#)  
For information about the State Machine Processing logic option, including the settings and supported devices
- [Manually Specifying Enumerated Types Using the `enum\_encoding` Attribute](#) on page 16-33  
For more information about assigning your own state encoding with the **User-Encoded** setting of the **State Machine Processing** option in a VHDL design

## Manually Specifying State Assignments Using the `syn_encoding` Attribute

The Quartus II software infers state machines from enumerated types and automatically assigns state encoding based on [State Machine Processing](#) on page 16-31.

With this logic option, you can choose the value **User-Encoded** to use the encoding from your HDL code. However, in standard VHDL code, you cannot specify user encoding in the state machine description because enumeration literals have no numeric values in VHDL.

To assign your own state encoding for the **User-Encoded State Machine Processing** setting, use the `syn_encoding` synthesis attribute to apply specific binary encodings to the elements of an enumerated type or to specify an encoding style. The Quartus II software can implement Enumeration Types with different encoding styles, as listed in this table.

**Table 16-10: syn\_encoding Attribute Values**

Attribute Value	Enumeration Types
"default"	Use an encoding based on the number of enumeration literals in the Enumeration Type. If the number of literals is less than five, use the "sequential" encoding. If the number of literals is more than five, but fewer than 50, use a "one-hot" encoding. Otherwise, use a "gray" encoding.
"sequential"	Use a binary encoding in which the first enumeration literal in the Enumeration Type has encoding 0 and the second 1.
"gray"	Use an encoding in which the encodings for adjacent enumeration literals differ by exactly one bit. An N-bit gray code can represent $2^N$ values.
"johnson"	Use an encoding similar to a gray code. An N-bit Johnson code can represent at most $2^N$ states, but requires less logic than a gray encoding.
"one-hot"	The default encoding style requiring N bits, in which N is the number of enumeration literals in the Enumeration Type.
"compact"	Use an encoding with the fewest bits.
"user"	Encode each state using its value in the Verilog source. By changing the values of your state constants, you can change the encoding of your state machine.

The `syn_encoding` attribute must follow the enumeration type definition, but precede its use.

**Related Information**

[State Machine Processing](#) on page 16-31

**Manually Specifying Enumerated Types Using the enum\_encoding Attribute**

By default, the Quartus II software one-hot encodes all enumerated types you defined. With the `enum_encoding` attribute, you can specify the logic encoding for an enumerated type and override the default one-hot encoding to improve the logic efficiency.

**Note:** If an enumerated type represents the states of a state machine, using the `enum_encoding` attribute to specify a manual state encoding prevents the Compiler from recognizing state machines based on the enumerated type. Instead, the Compiler processes these state machines as regular logic with the encoding specified by the attribute, and the Report window for your project does not list these states machines as state machines. If you want to control the encoding for a recognized state machine, use the **State Machine Processing** logic option and the `syn_encoding` synthesis attribute.

To use the `enum_encoding` attribute in a VHDL design file, associate the attribute with the enumeration type whose encoding you want to control. The `enum_encoding` attribute must follow the enumeration type definition, but precede its use. In addition, the attribute value should be a string literal that specifies either an arbitrary user encoding or an encoding style of "default", "sequential", "gray", "johnson", or "one-hot".

An arbitrary user encoding consists of a space-delimited list of encodings. The list must contain as many encodings as the number of enumeration literals in your enumeration type. In addition,



the encodings should have the same length, and each encoding must consist solely of values from the `std_ulogic` type declared by the `std_logic_1164` package in the IEEE library.

In this example, the `enum_encoding` attribute specifies an arbitrary user encoding for the enumeration type `fruit`.

### Example 16-5: Specifying an Arbitrary User Encoding for Enumerated Type

```
type fruit is (apple, orange, pear, mango);
attribute enum_encoding : string;
attribute enum_encoding of fruit : type is "11 01 10 00";
```

This example shows the encoded enumeration literals:

### Example 16-6: Encoded Enumeration Literals

```
apple   = "11"
orange  = "01"
pear    = "10"
mango   = "00"
```

Altera recommends that you specify an encoding style, rather than a manual user encoding, especially when the enumeration type has a large number of enumeration literals. The Quartus II software can implement Enumeration Types with the different encoding styles, as shown in this table.

**Table 16-11: enum\_encoding Attribute Values**

Attribute Value	Enumeration Types
"default"	Use an encoding based on the number of enumeration literals in the enumeration type. If the number of literals are fewer than five, use the "sequential" encoding. If the number of literals are more than five, but fewer than 50 literals, use a "one-hot" encoding. Otherwise, use a "gray" encoding.
"sequential"	Use a binary encoding in which the first enumeration literal in the enumeration type has encoding 0 and the second 1.
"gray"	Use an encoding in which the encodings for adjacent enumeration literals differ by exactly one bit. An N-bit gray code can represent $2^N$ values.
"johnson"	Use an encoding similar to a gray code. An N-bit Johnson code can represent at most $2^N$ states, but requires less logic than a gray encoding.
"one-hot"	The default encoding style requiring N bits, in which N is the number of enumeration literals in the enumeration type.

In [Example 16-5](#), the `enum_encoding` attribute manually specified a gray encoding for the enumeration type `fruit`. You can also concisely write this example by specifying the "gray" encoding style instead of a manual encoding, as shown in the following example:

**Example 16-7: Specifying the “gray” Encoding Style or Enumeration Type**

```
type fruit is (apple, orange, pear, mango);
attribute enum_encoding : string;
attribute enum_encoding of fruit : type is "gray";
```

## Safe State Machine

The **Safe State Machine** logic option and corresponding `syn_encoding` attribute value `safe` specify that the software must insert extra logic to detect an illegal state, and force the transition of the state machine to the reset state.

A finite state machine can enter an illegal state—meaning the state registers contain a value that does not correspond to any defined state. By default, the behavior of the state machine that enters an illegal state is undefined. However, you can set the `syn_encoding` attribute to `safe` or use the **Safe State Machine** logic option if you want the state machine to recover deterministically from an illegal state. The software inserts extra logic to detect an illegal state, and forces the transition of the state machine to the reset state. You can use this logic option when the state machine enters an illegal state. The most common cause of an illegal state is a state machine that has control inputs that come from another clock domain, such as the control logic for a clock-crossing FIFO, because the state machine must have inputs from another clock domain. This option protects only state machines (and not other registers) by forcing them into the reset state. You can use this option if your design has asynchronous inputs. However, Altera recommends using a synchronization register chain instead of relying on the safe state machine option.

The `safe` state machine value does not use any user-defined default logic from your HDL code that corresponds to unreachable states. Verilog HDL and VHDL enable you to specify a behavior for all states in the state machine explicitly, including unreachable states. However, synthesis tools detect if state machine logic is unreachable and minimize or remove the logic. Synthesis tools also remove any flag signals or logic that indicate such an illegal state. If the software implements the state machine as `safe`, the recovery logic added by Quartus II Integrated Synthesis forces its transition from an illegal state to the reset state.

You can set the **Safe State Machine** logic option globally, or on individual state machines. To set this logic option, click **Assignments > Settings > Compiler Settings > Advanced Settings (Synthesis)**.

**Table 16-12: Setting the `syn_encoding safe` attribute on a State Machine in HDL**

HDL	Code
Verilog HDL	<pre>reg [2:0] my_fsm /* synthesis syn_encoding = "safe" */ ;</pre>
Verilog-2001 and SystemVerilog	<pre>(* syn_encoding = "safe" *) reg [2:0] my_fsm;</pre>
VHDL	<pre>ATTRIBUTE syn_encoding OF my_fsm : TYPE IS "safe";</pre>

If you specify an encoding style, separate the encoding style value in the quotation marks with the `safe` value with a comma, as follows: "safe, one-hot" or "safe, gray".

Safe state machine implementation can result in a noticeable area increase for your design. Therefore, Altera recommends that you set this option only on the critical state machines in your design in which the safe mode is necessary, such as a state machine that uses inputs from asynchronous clock domains. You may not need to use this option if you correctly synchronize inputs coming from other clock domains.

**Note:** If you create the `safe` state machine assignment on an instance that the software fails to recognize as a state machine, or an entity that contains a state machine, the software takes no action. You must restructure the code, so that the software recognizes and infers the instance as a state machine.

#### Related Information

- [Manually Specifying State Assignments Using the `syn\_encoding` Attribute](#) on page 16-32
- [Safe State Machine logic option](#)  
For more information about the Safe State Machine logic option
- [Recommended HDL Coding Styles](#) on page 12-1  
For guidelines to ensure that the software correctly infers your state machine

## Power-Up Level

This logic option causes a register (flipflop) to power up with the specified logic level, either high (1) or low (0). The registers in the core hardware power up to 0 in all Altera devices. For the register to power up with a logic level high, the Compiler performs an optimization referred to as NOT-gate push back on the register. NOT-gate push back adds an inverter to the input and the output of the register, so that the reset and power-up conditions appear to be high and the device operates as expected. The register itself still powers up to 0, but the register output inverts so the signal arriving at all destinations is 1.

The **Power-Up Level** option supports wildcard characters, and you can apply this option to any register, registered logic cell WYSIWYG primitive, or to a design entity containing registers, if you want to set the power level for all registers in your design entity. If you assign this option to a registered logic cell WYSIWYG primitive, such as an atom primitive from a third-party synthesis tool, you must turn on the **Perform WYSIWYG Primitive Resynthesis** logic option for the option to take effect. You can also apply the option to a pin with the logic configurations described in the following list:

- If you turn on this option for an input pin, the option transfers to the register that the pin drives, if all these conditions are present:
  - No logic, other than inversion, between the pin and the register.
  - The input pin drives the data input of the register.
  - The input pin does not fan-out to any other logic.
- If you turn on this option for an output or bidirectional pin, the option transfers to the register that feeds the pin, if all these conditions are present:
  - No logic, other than inversion, between the register and the pin.
  - The register does not fan out to any other logic.

### Related Information

#### [Power-Up Level logic option](#)

For more information about the Power-Up Level logic option, including information on the supported device families

## Inferred Power-Up Levels

Quartus II Integrated Synthesis reads default values for registered signals defined in Verilog HDL and VHDL code, and converts the default values into **Power-Up Level** settings. The software also synthesizes variables with assigned values in Verilog HDL initial blocks into power-up conditions. Synthesis of these default and initial constructs allows synthesized behavior of your design to match, as closely as possible, the power-up state of the HDL code during a functional simulation.

The following register declarations all set a power-up level of  $V_{CC}$  or a logic value “1”, as shown in this example:

```
signal q : std_logic = '1'; -- power-up to VCC

reg q = 1'b1; // power-up to VCC

reg q;
initial begin q = 1'b1; end // power-up to VCC
```

### Related Information

- [Recommended HDL Coding Styles](#) on page 12-1  
For more information about NOT-gate push back, the power-up states for Altera devices, and how set and reset control signals affect the power-up level

## Power-Up Don't Care

This logic option allows the Compiler to optimize registers in your design that do not have a defined power-up condition.

For example, your design might have a register with its D input tied to  $V_{CC}$ , and with no clear signal or other secondary signals. If you turn on this option, the Compiler can choose for the register to power up to  $V_{CC}$ . Therefore, the output of the register is always  $V_{CC}$ . The Compiler can remove the register and connect its output to  $V_{CC}$ . If you turn this option off or if you set a **Power-Up Level** assignment of **Low** for this register, the register transitions from GND to  $V_{CC}$  when your design starts up on the first clock signal. Thus, the register is at  $V_{CC}$  and you cannot remove the register. Similarly, if the register has a clear signal, the Compiler cannot remove the register because after asserting the clear signal, the register transitions again to GND and back to  $V_{CC}$ .

If the Compiler performs a **Power-Up Don't Care** optimization that allows it to remove a register, it issues a message to indicate that it is doing so.

This project-wide option does not apply to registers that have the **Power-Up Level** logic option set to either **High** or **Low**.

### Related Information

#### [Power-Up Don't Care logic option](#)

For more information about Power-Up Don't Care logic option and a list of supported devices

## Remove Duplicate Registers

The **Remove Duplicate Registers** logic option removes registers that are identical to other registers.

### Related Information

#### [Remove Duplicate Registers logic option](#)

For more information about Remove Duplicate Registers logic option and the supported devices

## Preserve Registers

This attribute and logic option directs the Compiler not to minimize or remove a specified register during synthesis optimizations or register netlist optimizations. Optimizations can eliminate redundant registers and registers with constant drivers; this option prevents the software from reducing a register to a constant or merging with a duplicate register. This option can preserve a register so you can observe the register during simulation or with the SignalTap<sup>®</sup> II Logic Analyzer. Additionally, this option can preserve registers if you create a preliminary version of your design in which you have not specified the secondary signals. You can also use the attribute to preserve a duplicate of an I/O register so that you can place one copy of the I/O register in an I/O cell and the second in the core.

**Note:** This option cannot preserve registers that have no fan-out.

The **Preserve Registers** logic option prevents the software from inferring a register as a state machine.

You can set the **Preserve Registers** logic option in the Quartus II software, or you can set the `preserve` attribute in your HDL code. In these examples, the Quartus II software preserves the `my_reg` register.

**Table 16-13: Setting the `syn_preserve` attribute in HDL Code**

HDL	Code <sup>(11)</sup>
Verilog HDL	<pre>reg my_reg /* synthesis syn_preserve = 1 */;</pre>
Verilog-2001	<pre>(* syn_preserve = 1 *) reg my_reg;</pre>

**Table 16-14: Setting the `preserve` attribute in HDL Code**

In addition to `preserve`, the Quartus II software supports the `syn_preserve` attribute name for compatibility with other synthesis tools.

HDL	Code
VHDL	<pre>signal my_reg : stdlogic; attribute preserve : boolean; attribute preserve of my_reg : signal is true;</pre>

### Related Information

#### [Preserve Registers logic option](#)

For more information about the Preserve Registers logic option and the supported devices

<sup>(11)</sup> The `= 1` after the `preserve` are optional, because the assignment uses a default value of 1 when you specify the assignment.

[Noprune Synthesis Attribute/Preserve Fan-out Free Register Node](#) on page 16-39

For more information about preventing the removal of registers with no fan-out

## Disable Register Merging/Don't Merge Register

This logic option and attribute prevents the specified register from merging with other registers and prevents other registers from merging with the specified register. When applied to a design entity, it applies to all registers in the entity.

You can set the **Disable Register Merging** logic option in the Quartus II software, or you can set the `dont_merge` attribute in your HDL code, as shown in these examples. In these examples, the logic option or the attribute prevents the `my_reg` register from merging.

**Table 16-15: Setting the `dont_merge` attribute in HDL code**

HDL	Code
Verilog HD	<pre>reg my_reg /* synthesis dont_merge */;</pre>
Verilog-2001 and SystemVerilog	<pre>(* dont_merge *) reg my_reg;</pre>
VHDL	<pre>signal my_reg : stdlogic; attribute dont_merge : boolean; attribute dont_merge of my_reg : signal is true;</pre>

### Related Information

#### [Disable Register Merging logic option](#)

For more information about the **Disable Register Merging** logic option and the supported devices

## Noprune Synthesis Attribute/Preserve Fan-out Free Register Node

This synthesis attribute and corresponding logic option direct the Compiler to preserve a fan-out-free register through the entire compilation flow. This option is different from the **Preserve Registers** option, which prevents the Quartus II software from reducing a register to a constant or merging with a duplicate register. Standard synthesis optimizations remove nodes that do not directly or indirectly feed a top-level output pin. This option can retain a register so you can observe the register in the Simulator or the SignalTap II Logic Analyzer. Additionally, this option can retain registers if you create a preliminary version of your design in which you have not specified the fan-out logic of the register.

You can set the **Preserve Fan-out Free Register Node** logic option in the Quartus II software, or you can set the `noprune` attribute in your HDL code, as shown in these examples. In these examples, the logic option or the attribute preserves the `my_reg` register.

**Note:** You must use the `noprune` attribute instead of the logic option if the register has no immediate fan-out in its module or entity. If you do not use the synthesis attribute, the software removes (or “prunes”) registers with no fan-out during Analysis & Elaboration before the logic synthesis stage applies any logic options. If the register has no fan-out in the full design, but has fan-out in its module or entity, you can use the logic option to retain the register through compilation.

The software supports the attribute name `syn_noprune` for compatibility with other synthesis tools.

Table 16-16: Setting the `noprune` attribute in HDL code

HDL	Code
Verilog HD	<pre>reg my_reg /* synthesis syn_noprune */;</pre>
Verilog-2001 and SystemVerilog	<pre>(* noprune *) reg my_reg;</pre>
VHDL	<pre>signal my_reg : stdlogic; attribute noprune: boolean; attribute noprune of my_reg : signal is true;</pre>

**Related Information****Preserve Fan-out Free Register logic option**

For more information about **Preserve Fan-out Free Register Node** logic option and a list of supported devices

**Keep Combinational Node/Implement as Output of Logic Cell**

This synthesis attribute and corresponding logic option direct the Compiler to keep a wire or combinational node through logic synthesis minimizations and netlist optimizations. A wire that has a `keep` attribute or a node that has the **Implement as Output of Logic Cell** logic option applied becomes the output of a logic cell in the final synthesis netlist, and the name of the logic cell remains the same as the name of the wire or node. You can use this directive to make combinational nodes visible to the SignalTap II Logic Analyzer.

**Note:** The option cannot keep nodes that have no fan-out. You cannot maintain node names for wires with tri-state drivers, or if the signal feeds a top-level pin of the same name (the software changes the node name to a name such as `<net name>-buf0`).

You can use the **Ignore LCELL Buffers** logic option to direct Analysis & Synthesis to ignore logic cell buffers that the **Implement as Output of Logic Cell** logic option or the `LCELL` primitive created. If you apply this logic option to an entity, it affects all lower-level entities in the hierarchy path.

**Note:** To avoid unintended design optimizations, ensure that any entity instantiated with Altera or third-party IP that relies on logic cell buffers for correct behavior does not inherit the **Ignore LCELL Buffers** logic option. For example, if an IP core uses logic cell buffers to manage high fan-out signals and inherits the **Ignore LCELL Buffers** logic option, the target device may no longer function properly.

You can turn off the **Ignore LCELL Buffers** logic option for a specific entity to override any assignments inherited from higher-level entities in the hierarchy path if logic cell buffers created by the **Implement as Output of Logic Cell** logic option or the `LCELL` primitive are required for correct behavior.

You can set the **Implement as Output of Logic Cell** logic option in the Quartus II software, or you can set the `keep` attribute in your HDL code, as shown in these tables. In these tables, the Compiler maintains the node name `my_wire`.

**Table 16-17: Setting the keep Attribute in HDL code**

HDL	Code
Verilog HD	<code>wire my_wire /* synthesis keep = 1 */;</code>
Verilog-2001	<code>(* keep = 1 *) wire my_wire;</code>

**Table 16-18: Setting the syn\_keep Attribute in HDL Code**

In addition to `keep`, the Quartus II software supports the `syn_keep` attribute name for compatibility with other synthesis tools.

HDL	Code
VHDL	<code>signal my_wire: bit; attribute syn_keep: boolean; attribute syn_keep of my_wire: signal is true;</code>

**Related Information**

**Implement as Output of Logic Cell logic option**

For more information about the **Implement as Output of Logic Cell** logic option and the supported devices

## Disabling Synthesis Netlist Optimizations with dont\_rettime Attribute

This attribute disables synthesis retiming optimizations on the register you specify. When applied to a design entity, it applies to all registers in the entity.

You can turn off retiming optimizations with this option and prevent node name changes, so that the Compiler can correctly use your timing constraints for the register.

You can set the **Netlist Optimizations** logic option to **Never Allow** in the Quartus II software to disable retiming along with other synthesis netlist optimizations, or you can set the `dont_rettime` attribute in your HDL code, as shown in the following table. In the following table, the code prevents `my_reg` register from being retimed.

**Table 16-19: Setting the dont\_rettime Attribute in HDL Code**

HDL	Code
Verilog HDL	<code>reg my_reg /* synthesis dont_rettime */;</code>
Verilog-2001 and SystemVerilo	<code>(* dont_rettime *) reg my_reg;</code>
VHD	<code>signal my_reg : std_logic; attribute dont_rettime : boolean; attribute dont_rettime of my_reg : signal is true;</code>



**Note:** For compatibility with third-party synthesis tools, Quartus II Integrated Synthesis also supports the attribute `syn_allow_retiming`. To disable retiming, set `syn_allow_retiming` to 0 (Verilog HDL) or `false` (VHDL). This attribute does not have any effect when you set the attribute to 1 or `true`.

## Disabling Synthesis Netlist Optimizations with dont\_replicate Attribute

This attribute disables synthesis replication optimizations on the register you specify. When applied to a design entity, it applies to all registers in the entity.

You can turn off register replication (or duplication) optimizations with this option, so that the Compiler uses your timing constraints for the register.

You can set the **Netlist Optimizations** logic option to **Never Allow** in the Quartus II software to disable replication along with other synthesis netlist optimizations, or you can set the `dont_replicate` attribute in your HDL code, as shown in these examples. In these examples, the code prevents the replication of the `my_reg` register.

**Table 16-20: Setting the `dont_replicate` attribute in HDL Code**

HDL	Code
Verilog HD	<pre>reg my_reg /* synthesis dont_replicate */;</pre>
Verilog-2001 and SystemVerilog	<pre>(* dont_replicate *) reg my_reg;</pre>
VHDL	<pre>signal my_reg : std_logic; attribute dont_replicate : boolean; attribute dont_replicate of my_reg : signal is true;</pre>

**Note:** For compatibility with third-party synthesis tools, Quartus II Integrated Synthesis also supports the attribute `syn_replicate`. To disable replication, set `syn_replicate` to 0 (Verilog HDL) or `false` (VHDL). This attribute does not have any effect when you set the attribute to 1 or `true`.

## Maximum Fan-Out

This **Maximum Fan-Out** attribute and logic option direct the Compiler to control the number of destinations that a node feeds. The Compiler duplicates a node and splits its fan-out until the individual fan-out of each copy falls below the maximum fan-out restriction. You can apply this option to a register or a logic cell buffer, or to a design entity that contains these elements. You can use this option to reduce the load of critical signals, which can improve performance. You can use the option to instruct the Compiler to duplicate a register that feeds nodes in different locations on the target device. Duplicating the register can enable the Fitter to place these new registers closer to their destination logic to minimize routing delay.

To turn off the option for a given node if you set the option at a higher level of the design hierarchy, in the **Netlist Optimizations** logic option, select **Never Allow**. If not disabled by the **Netlist Optimizations**

option, the Compiler acknowledges the maximum fan-out constraint as long as the following conditions are met:

- The node is not part of a cascade, carry, or register cascade chain.
- The node does not feed itself.
- The node feeds other logic cells, DSP blocks, RAM blocks, and pins through data, address, clock enable, and other ports, but not through any asynchronous control ports (such as asynchronous clear).

The Compiler does not create duplicate nodes in these cases, because there is no clear way to duplicate the node, or to avoid the small differences in timing which could produce functional differences in the implementation (in the third condition above in which asynchronous control signals are involved). If you cannot apply the constraint because you do not meet one of these conditions, the Compiler issues a message to indicate that the Compiler ignores the maximum fan-out assignment. To instruct the Compiler not to check node destinations for possible problems such as the third condition, you can set the **Netlist Optimizations** logic option to **Always Allow** for a given node.

**Note:** If you have enabled any of the Quartus II netlist optimizations that affect registers, add the `preserve` attribute to any registers to which you have set a `maxfan` attribute. The `preserve` attribute ensures that the netlist optimization algorithms, such as register retiming, do not affect the registers.

You can set the **Maximum Fan-Out** logic option in the Quartus II software. This option supports wildcard characters. You can also set the `maxfan` attribute in your HDL code, as shown in these examples. In these examples, the Compiler duplicates the `clk_gen` register, so its fan-out is not greater than 50.

**Table 16-21: Setting the `maxfan` attribute in HDL Code**

HDL	Code
Verilog HDL	<code>reg clk_gen /* synthesis syn_maxfan = 50 */;</code>
Verilog-2001	<code>(* maxfan = 50 *) reg clk_gen;</code>

**Table 16-22: Setting the `syn_maxfan` attribute in HDL Code**

The Quartus II software supports the `syn_maxfan` attribute for compatibility with other synthesis tools.

HDL	Code
VHDL	<code>signal clk_gen : stdlogic; attribute maxfan : signal ; attribute maxfan of clk_gen : signal is 50;</code>

**Related Information**

- [Netlist Optimizations and Physical Synthesis](#)  
For details about netlist optimizations
- [Maximum Fan-Out logic option](#)  
For more information about the Maximum Fan-Out logic option and the supported devices

## Controlling Clock Enable Signals with Auto Clock Enable Replacement and `direct_enable`

The **Auto Clock Enable Replacement** logic option allows the software to find logic that feeds a register and move the logic to the register's clock enable input port. To solve fitting or performance issues with designs that have many clock enables, you can turn off this option for individual registers or design entities. Turning the option off prevents the software from using the register's clock enable port. The software implements the clock enable functionality using multiplexers in logic cells.

If the software does not move the specific logic to a clock enable input with the **Auto Clock Enable Replacement** logic option, you can instruct the software to use a direct clock enable signal. The attribute ensures that the signal drives the clock enable port, and the software does not optimize or combine the signal with other logic.

These tables show how to set this attribute to ensure that the attribute preserves the signal and uses the signal as a clock enable.

**Table 16-23: Setting the `direct_enable` in HDL Code**

HDL	Code
Verilog HDL	<pre>wire my_enable /* synthesis direct_enable = 1 */ ;</pre>
VHDL	<pre>attribute direct_enable: boolean; attribute direct_enable of my_enable: signal is true;</pre>

**Table 16-24: Setting the `syn_direct_enable` in HDL Code**

The Quartus II software supports the `syn_direct_enable` attribute name for compatibility with other synthesis tools.

HDL	Code
Verilog-2001 and SystemVerilog	<pre>(* syn_direct_enable *) wire my_enable;</pre>

### Related Information

#### [Auto Clock Enable Replacement logic option](#)

For more information about the **Auto Clock Enable Replacement** logic option and the supported devices

## Inferring Multiplier, DSP, and Memory Functions from HDL Code

The Quartus II Compiler automatically recognizes multipliers, multiply-accumulators, multiply-adders, or memory functions described in HDL code, and either converts the HDL code into respective IP core or maps them directly to device atoms or memory atoms. If the software converts the HDL code into an IP core, the software uses the Altera IP core code when you compile your design, even when you do not specifically instantiate the IP core. The software infers IP cores to take advantage of logic that you optimize for Altera devices. The area and performance of such logic can be better than the results from inferring generic logic from the same HDL code.

Additionally, you must use IP cores to access certain architecture-specific features, such as RAM, DSP blocks, and shift registers that provide improved performance compared with basic logic cells.

The Quartus II software provides options to control the inference of certain types of IP cores.

#### Related Information

- [Recommended HDL Coding Styles](#) on page 12-1  
For details about coding style recommendations when targeting IP cores in Altera devices

## Multiply-Accumulators and Multiply-Adders

Use the **Auto DSP Block Replacement** logic option to control DSP block inference for multiply-accumulations and multiply-adders. To disable inference, turn off this option for the entire project on the **Advanced Analysis & Synthesis** dialog box of the **Compiler Settings** page.

#### Related Information

##### [Auto DSP Block Replacement logic option](#)

For more information about the Auto DSP Block Replacement logic option and the supported devices

## Shift Registers

Use the **Auto Shift Register Replacement** logic option to control shift register inference. This option has three settings: **Off**, **Auto** and **Always**. **Auto** is the default setting in which Quartus II Integrated Synthesis decides which shift registers to replace or leave in registers. Placing shift registers in memory saves logic area, but can have a negative effect on  $f_{max}$ . Quartus II Integrated Synthesis uses the optimization technique setting, logic and RAM utilization of your design, and timing information from **Timing-Driven Synthesis** to determine which shift registers are located in memory and which are located in registers. To disable inference, click **Assignments > Settings > Compiler Settings > Advanced Settings (Synthesis)**. You can also disable the option for a specific block with the Assignment Editor. Even if you set the logic option to **On** or **Auto**, the software might not infer small shift registers because small shift registers do not benefit from implementation in dedicated memory. However, you can use the **Allow Any Shift Register Size for Recognition** logic option to instruct synthesis to infer a shift register even when its size is too small.

You can use the **Allow Shift Register Merging across Hierarchies** option to prevent the Compiler from merging shift registers in different hierarchies into one larger shift register. The option has three settings: **On**, **Off**, and **Auto**. The **Auto** setting is the default setting, and the Compiler decides whether or not to merge shift registers across hierarchies. When you turn on this option, the Compiler allows all shift registers to merge across hierarchies, and when you turn off this option, the Compiler does not allow any shift registers to merge across hierarchies. You can set this option globally or on entities or individual nodes.

**Note:** The registers that the software maps to the RAM-based Shift Register IP core and places in RAM are not available in the Simulator because their node names do not exist after synthesis.

The Compiler turns off the **Auto Shift Register Replacement** logic option when you select a formal verification tool on the **EDA Tool Settings** page. If you do not select a formal verification tool, the Compiler issues a warning and the compilation report lists shift registers that the logic option might infer. To enable an IP core for the shift register in the formal verification flow, you can either instantiate a shift register explicitly with the IP catalog or make the shift register into a black box in a separate entity or module.

### Related Information

#### [Auto Shift Register Replacement logic option](#)

For more information about the Auto Shift Register Replacement logic option and the supported devices

#### [RAM-Based Shift Register \(ALTSHIFT\\_TAPS\) User Guide](#)

For more information about the RAM-based Shift Register IP core

## RAM and ROM

Use the **Auto RAM Replacement** and **Auto ROM Replacement** logic options to control RAM and ROM inference, respectively. To disable the inference, click **Assignments > Settings > Compiler Settings > Advanced Settings (Synthesis)**.

**Note:** Although the software implements inferred shift registers in RAM blocks, you cannot turn off the **Auto RAM Replacement** option to disable shift register replacement. Use the **Auto Shift Register Replacement** option.

The software might not infer very small RAM or ROM blocks because you can implement very small memory blocks with the registers in the logic. However, you can use the **Allow Any RAM Size for Recognition** and **Allow Any ROM Size for Recognition** logic options to instruct synthesis to infer a memory block even when its size is too small.

**Note:** The software turns off the **Auto ROM Replacement** logic option when you select a formal verification tool in the **EDA Tool Settings** page. If you do not select a formal verification tool, the software issues a warning and a report panel provides a list of ROMs that the logic option might infer. To enable an IP core for the shift register in the formal verification flow, you can either instantiate a ROM explicitly using the IP Catalog or create a black box for the ROM in a separate entity or in a separate module.

Although formal verification tools do not support inferred RAM blocks, due to the importance of inferring RAM in many designs, the software turns on the **Auto RAM Replacement** logic option when you select a formal verification tool in the **EDA Tool Settings** page. The software automatically performs black box instance for any module or entity that contains an inferred RAM block. The software issues a warning and lists the black box created in the compilation report. This black box allows formal verification tools to proceed; however, the formal verification tool cannot verify the entire module or entire entity that contains the RAM. Altera recommends that you explicitly instantiate RAM blocks in separate modules or in separate entities so that the formal verification tool can verify as much logic as possible.

### Related Information

- [Shift Registers](#) on page 16-45
- [Auto RAM Replacement logic option](#)  
For more information about the Auto RAM Replacement logic option and its supported devices
- [Auto ROM Replacement logic option](#)  
For more information about the Auto ROM Replacement logic option and its supported devices

## Resource Aware RAM, ROM, and Shift-Register Inference

The Quartus II Integrated Synthesis considers resource usage when inferring RAM, ROM, and shift registers. During RAM, ROM, and shift register inferencing, synthesis looks at the number of memories available in the current device and does not infer more memory than is available to avoid a no-fit error.

Synthesis tries to select the memories that are not inferred in a way that aims at the smallest increase in logic and registers.

Resource aware RAM, ROM and shift register inference is controlled by the **Resource Aware Inference for Block RAM** option. To disable this option for the entire project, click **Assignments > Settings > Compiler Settings > Advanced Settings (Synthesis)**.

When you select the **Auto** setting, resource aware RAM, ROM, and shift register inference use the resource counts from the largest device.

For designs with multiple partitions, Quartus II Integrated Synthesis considers one partition at a time. Therefore, for each partition, it assumes that all RAM blocks are available to that partition. If this causes a no-fit error, you can limit the number of RAM blocks available per partition with the **Maximum Number of M512 Memory Blocks**, **Maximum Number of M4K/M9K/M20K/M10K Memory Blocks**, **Maximum Number of M-RAM/M144K Memory Blocks** and **Maximum Number of LABs** settings in the Assignment Editor. The balancer also uses these options.

## Auto RAM to Logic Cell Conversion

The **Auto RAM to Logic Cell Conversion** logic option allows Quartus II Integrated Synthesis to convert small RAM blocks to logic cells if the logic cell implementation gives better quality of results. The software converts only single-port or simple-dual port RAMs with no initialization files to logic cells. You can set this option globally or apply it to individual RAM nodes. You can enable this option by turning on the appropriate option for the entire project in the **Advanced Analysis & Synthesis Settings** dialog box.

For Arria GX and Stratix family of devices, the software uses the following rules to determine the placement of a RAM, either in logic cells or a dedicated RAM block:

- If the number of words is less than 16, use a RAM block if the total number of bits is greater than or equal to 64.
- If the number of words is greater than or equal to 16, use a RAM block if the total number of bits is greater than or equal to 32.
- Otherwise, implement the RAM in logic cells.

For the Cyclone family of devices, the software uses the following rules:

- If the number of words is greater than or equal to 64, use a RAM block.
- If the number of words is greater than or equal to 16 and less than 64, use a RAM block if the total number of bits is greater than or equal to 128.
- Otherwise, implement the RAM in logic cells.

### Related Information

#### [Auto RAM to Logic Cell Conversion logic option](#)

For more information about the Auto RAM to Logic Cell Conversion logic options and the supported devices

## RAM Style and ROM Style—for Inferred Memory

These attributes specify the implementation for an inferred RAM or ROM block. You can specify the type of TriMatrix embedded memory block, or specify the use of standard logic cells (LEs or ALMs). The Quartus II software supports the attributes only for device families with TriMatrix embedded memory blocks.

The `ramstyle` and `romstyle` attributes take a single string value. The `M512`, `M4K`, `M-RAM`, `MLAB`, `M9K`, `M144K`, `M20K`, and `M10K` values (as applicable for the target device family) indicate the type of memory block to use for the inferred RAM or ROM. If you set the attribute to a block type that does not exist in the target device family, the software generates a warning and ignores the assignment. The `logic` value indicates that the Quartus II software implements the RAM or ROM in regular logic rather than dedicated memory blocks. You can set the attribute on a module or entity, in which case it specifies the default implementation style for all inferred memory blocks in the immediate hierarchy. You can also set the attribute on a specific signal (VHDL) or variable (Verilog HDL) declaration, in which case it specifies the preferred implementation style for that specific memory, overriding the default implementation style.

**Note:** If you specify a `logic` value, the memory appears as a RAM or ROM block in the RTL Viewer, but Integrated Synthesis converts the memory to regular logic during synthesis.

In addition to `ramstyle` and `romstyle`, the Quartus II software supports the `syn_ramstyle` attribute name for compatibility with other synthesis tools.

These tables specify that you must implement all memory in the module or the `my_memory_blocks` entity with a specific type of block.

**Table 16-25: Applying a `romstyle` Attribute to a Module Declaration**

HDL	Code
Verilog-1995	<pre>module my_memory_blocks (...) /* synthesis romstyle = "M4K" */ ;</pre>

**Table 16-26: Applying a `ramstyle` Attribute to a Module Declaration**

HDL	Code
Verilog-2001 and SystemVerilog	<pre>(* ramstyle = "M512" *) module my_memory_blocks (...);</pre>

**Table 16-27: Applying a `romstyle` Attribute to an Architecture**

HDL	Code
VHDL	<pre>architecture rtl of my_my_memory_blocks is attribute romstyle : string; attribute romstyle of rtl : architecture is "M-RAM"; begin</pre>

These tables specify that you must implement the inferred `my_ram` or `my_rom` memory with regular logic instead of a TriMatrix memory block.

**Table 16-28: Applying a `syn_ramstyle` Attribute to a Variable Declaration**

HDL	Code
Verilog-1995	<pre>reg [0:7] my_ram[0:63] /* synthesis syn_ramstyle = "logic" */ ;</pre>

**Table 16-29: Applying a `romstyle` Attribute to a Variable Declaration**

HDL	Code
Verilog-2001 and SystemVerilog	<pre>(* romstyle = "logic" *) reg [0:7] my_rom[0:63];</pre>

**Table 16-30: Applying a `ramstyle` Attribute to a Signal Declaration**

HDL	Code
VHDL	<pre>type memory_t is array (0 to 63) of std_logic_vector (0 to 7) ; signal my_ram : memory_t; attribute ramstyle : string; attribute ramstyle of my_ram : signal is "logic";</pre>

You can control the depth of an inferred memory block and optimize its usage with the `max_depth` attribute. You can also optimize the usage of the memory block with this attribute.

These tables specify the depth of the inferred memory `mem` using the `max_depth` synthesis attribute.

**Table 16-31: Applying a `max_depth` Attribute to a Variable Declaration**

HDL	Code
Verilog-1995	<pre>reg [7:0] mem [127:0] /* synthesis max_depth = 2048 */</pre>

**Table 16-32: Applying a `max_depth` Attribute to a Variable Declaration**

HDL	Code
Verilog-2001 and SystemVerilog	<pre>(* max_depth = 2048*) reg [7:0] mem [127:0];</pre>

**Table 16-33: Applying a `max_depth` Attribute to a Variable Declaration**

HDL	Code
VHDL	<pre>type ram_block is array (0 to 31) of std_logic_vector (2 downto 0); signal mem : ram_block; attribute max_depth : natural; attribute max_depth OF mem : signal is 2048;</pre>

The syntax for setting these attributes in HDL is the same as the syntax for other synthesis attributes, as shown in [Synthesis Attributes](#) on page 16-22.

**Related Information**

[Synthesis Attributes](#) on page 16-22



## RAM Style Attribute—For Shift Registers Inference

The RAM style attribute for shift register allows you to use the RAM style attribute for shift registers, just as you use them for RAM or ROMs. The Quartus II Synthesis uses the RAM style attribute during shift register inference. If synthesis infers the shift register to RAM, it will be sent to the requested RAM block type. Shift registers are merged only if the RAM style attributes are compatible. If the RAM style is set to logic, a shift register does not get inferred to RAM.

**Note:** You can also assign the RAM style attribute for shift registers globally, which will affect all shift registers.

**Table 16-34: Setting the RAM Style Attribute for Shift Registers**

HDL	Code
Verilog	<pre>(* ramstyle = "mlab" *)reg [N-1:0] sr;</pre>
VHDL	<pre>attribute ramstyle : string;attribute ramstyle of sr : signal is "M20K";</pre>

### Disabling Add Pass-Through Logic to Inferred RAMs `no_rw_check` Attribute

Use the `no_rw_check` value for the `ramstyle` attribute, or disable the **Add Pass-Through Logic to Inferred RAMs** logic option assignment to indicate that your design does not depend on the behavior of the inferred RAM, when there are reads and writes to the same address in the same clock cycle. If you specify the attribute or disable the logic option, the Quartus II software chooses a read-during-write behavior instead of the read-during-write behavior of your HDL source code.

You disable or edit the attributes of this option by modifying the `add_pass_through_logic_to_inferred_rams` option in the Quartus II Settings File (`.qsf`). There is no corresponding GUI setting for this option.

Sometimes, you must map an inferred RAM into regular logic cells because the inferred RAM has a read-during-write behavior that the TriMatrix memory blocks in your target device do not support. In other cases, the Quartus II software must insert extra logic to mimic read-during-write behavior of the HDL source to increase the area of your design and potentially reduce its performance. In some of these cases, you can use the attribute to specify that the software can implement the RAM directly in a TriMatrix memory block without using logic. You can also use the attribute to prevent a warning message for dual-clock RAMs in the case that the inferred behavior in the device does not exactly match the read-during-write conditions described in the HDL code.

To set the **Add Pass-Through Logic to Inferred RAMs** logic option with the Quartus II software, click **Assignments > Settings > Compiler Settings > Advanced Settings (Synthesis)**.

These examples use two addresses and normally require extra logic after the RAM to ensure that the read-during-write conditions in the device match the HDL code. If your design does not require a defined read-during-write condition, the extra logic is not necessary. With the `no_rw_check` attribute, Quartus II Integrated Synthesis does not generate the extra logic.

**Table 16-35: Inferred RAM Using no\_rw\_check Attribute**

HDL	Code
Verilog HDL	<pre> module ram_infer (q, wa, ra, d, we, clk);     output [7:0] q;     input [7:0] d;     input [6:0] wa;     input [6:0] ra;     input we, clk;     reg [6:0] read_add;     (* ramstyle = "no_rw_check" *) reg [7:0] mem [127:0];     always @ (posedge clk) begin         if (we)             mem[wa] &lt;= d;             read_add &lt;= ra;         end         assign q = mem[read_add];     endmodule </pre>
VHDL	<pre> LIBRARY ieee; USE ieee.std_logic_1164.ALL; ENTITY ram IS     PORT (         clock: IN STD_LOGIC;         data: IN STD_LOGIC_VECTOR (2 DOWNTO 0);         write_address: IN INTEGER RANGE 0 to 31;         read_address: IN INTEGER RANGE 0 to 31;         we: IN STD_LOGIC;         q: OUT STD_LOGIC_VECTOR (2 DOWNTO 0) ); END ram; ARCHITECTURE rtl OF ram IS     TYPE MEM IS ARRAY(0 TO 31) OF STD_LOGIC_VECTOR(2 DOWNTO 0);     SIGNAL ram_block: MEM;     ATTRIBUTE ramstyle : string;     ATTRIBUTE ramstyle of ram_block : signal is "no_rw_check";     SIGNAL read_address_reg: INTEGER RANGE 0 to 31; BEGIN     PROCESS (clock)     BEGIN         IF (clock'event AND clock = '1') THEN             IF (we = '1') THEN                 ram_block(write_address) &lt;= data;             END IF;             read_address_reg &lt;= read_address;         END IF;     END PROCESS;     q &lt;= ram_block(read_address_reg); END rtl; </pre>

You can use a ramstyle attribute with the MLAB value, so that the Quartus II software can infer a small RAM block and place it in an MLAB.

**Note:** You can use this attribute in cases in which some asynchronous RAM blocks might be coded with read-during-write behavior that does not match the Stratix IV and Stratix V architectures. Thus, the device behavior would not exactly match the behavior that the code describes. If the difference in behavior is acceptable in your design, use the ramstyle attribute with the no\_rw\_check value to specify that the software should not check the read-during-write behavior when inferring the RAM. When you set this attribute, Quartus II Integrated Synthesis allows the behavior of the output to differ when the asynchronous read occurs on an address that had a write on the most

recent clock edge. That is, the functional HDL simulation results do not match the hardware behavior if you write to an address that is being read. To include these attributes, set the value of the `ramstyle` attribute to `MLAB, no_rw_check`.

These examples show the method of setting two values to the `ramstyle` attribute with a small asynchronous RAM block, with the `ramstyle` synthesis attribute set, so that the software can implement the memory in the MLAB memory block and so that the read-during-write behavior is not important. Without the attribute, this design requires 512 registers and 240 ALUTs. With the attribute, the design requires eight memory ALUTs and only 15 registers.

**Table 16-36: Inferred RAM Using `no_rw_check` and `MLAB` Attributes**

HDL	Code
Verilog HDL	<pre> module async_ram (     input  [5:0] addr,     input  [7:0] data_in,     input      clk,     input      write,     output [7:0] data_out ); (* ramstyle = "MLAB, no_rw_check" *) reg [7:0] mem[0:63]; assign data_out = mem[addr]; always @ (posedge clk) begin     if (write)         mem[addr] = data_in;     end endmodule </pre>
VHDL	<pre> LIBRARY ieee; USE ieee.std_logic_1164.ALL; ENTITY ram IS     PORT (         clock: IN STD_LOGIC;         data: IN STD_LOGIC_VECTOR (2 DOWNTO 0);         write_address: IN INTEGER RANGE 0 to 31;         read_address: IN INTEGER RANGE 0 to 31;         we: IN STD_LOGIC;         q: OUT STD_LOGIC_VECTOR (2 DOWNTO 0)); END ram; ARCHITECTURE rtl OF ram IS     TYPE MEM IS ARRAY(0 TO 31) OF STD_LOGIC_VECTOR(2 DOWNTO 0);     SIGNAL ram_block: MEM;     ATTRIBUTE ramstyle : string;     ATTRIBUTE ramstyle of ram_block : signal is "MLAB , no_rw_ check";     SIGNAL read_address_reg: INTEGER RANGE 0 to 31; BEGIN     PROCESS (clock)     BEGIN         IF (clock'event AND clock = '1') THEN             IF (we = '1') THEN                 ram_block(write_address) &lt;= data;             END IF;             read_address_reg &lt;= read_address;         END IF;     END PROCESS;     q &lt;= ram_block(read_address_reg); END rtl; </pre>

**Related Information**

[Recommended HDL Coding Styles](#) on page 12-1

For more information about recommended styles for inferring RAM and some of the issues involved with different read-during-write conditions

[Add Pass-Through Logic to Inferred RAMs logic option](#)

For more information about the Add Pass-Through Logic to Inferred RAMs logic option and the supported devices

## RAM Initialization File—for Inferred Memory

The `ram_init_file` attribute specifies the initial contents of an inferred memory with a `.mif`. The attribute takes a string value containing the name of the RAM initialization file.

The `ram_init_file` attribute is supported for ROM too.

**Table 16-37: Applying a `ram_init_file` Attribute**

HDL	Code
Verilog-1995	<pre>reg [7:0] mem[0:255] /* synthesis ram_init_file = " my_init_file.mif" */;</pre>
Verilog-2001	<pre>(* ram_init_file = "my_init_file.mif" *) reg [7:0] mem[0:255];</pre>
VHDL <sup>(12)</sup>	<pre>type mem_t is array(0 to 255) of unsigned(7 downto 0); signal ram : mem_t; attribute ram_init_file : string; attribute ram_init_file of ram : signal is "my_init_file.mif";</pre>

**Related Information**

- [Recommended HDL Coding Styles](#) on page 12-1  
For more information about Inferring ROM Functions from HDL Code

## Multiplier Style—for Inferred Multipliers

The `multstyle` attribute specifies the implementation style for multiplication operations (\*) in your HDL source code. You can use this attribute to specify whether you prefer the Compiler to implement a multiplication operation in general logic or dedicated hardware, if available in the target device.

The `multstyle` attribute takes a string value of "logic" or "dsp", indicating a preferred implementation in logic or in dedicated hardware, respectively. In Verilog HDL, apply the attribute to a module declaration, a variable declaration, or a specific binary expression that contains the \* operator. In VHDL, apply the synthesis attribute to a signal, variable, entity, or architecture.

<sup>(12)</sup> You can also initialize the contents of an inferred memory by specifying a default value for the corresponding signal. In Verilog HDL, you can use an initial block to specify the memory contents. Quartus II Integrated Synthesis automatically converts the default value into a `.mif` for the inferred RAM.

**Note:** Specifying a `multstyle` of "dsp" does not guarantee that the Quartus II software can implement a multiplication in dedicated DSP hardware. The final implementation depends on several conditions, including the availability of dedicated hardware in the target device, the size of the operands, and whether or not one or both operands are constant.

In addition to `multstyle`, the Quartus II software supports the `syn_multstyle` attribute name for compatibility with other synthesis tools.

When applied to a Verilog HDL module declaration, the attribute specifies the default implementation style for all instances of the `*` operator in the module. For example, in the following code examples, the `multstyle` attribute directs the Quartus II software to implement all multiplications inside module `my_module` in the dedicated multiplication hardware.

**Table 16-38: Applying a `multstyle` Attribute to a Module Declaration**

HDL	Code
Verilog-1995	<pre>module my_module (...) /* synthesis multstyle = "dsp" */;</pre>
Verilog-2001	<pre>(* multstyle = "dsp" *) module my_module(...);</pre>

When applied to a Verilog HDL variable declaration, the attribute specifies the implementation style for a multiplication operator, which has a result directly assigned to the variable. The attribute overrides the `multstyle` attribute with the enclosing module, if present.

In these examples, the `multstyle` attribute applied to variable `result` directs the Quartus II software to implement `a * b` in logic rather than the dedicated hardware.

**Table 16-39: Applying a `multstyle` Attribute to a Variable Declaration**

HDL	Code
Verilog-2001	<pre>wire [8:0] a, b; (* multstyle = "logic" *) wire [17:0] result; assign result = a * b; //Multiplication must be                         //directly assigned to result</pre>
Verilog-1995	<pre>wire [8:0] a, b; wire [17:0] result /* synthesis multstyle = "logic" */; assign result = a * b; //Multiplication must be                         //directly assigned to result</pre>

When applied directly to a binary expression that contains the `*` operator, the attribute specifies the implementation style for that specific operator alone and overrides any `multstyle` attribute with the target variable or enclosing module.

In this example, the `multstyle` attribute indicates that you must implement `a * b` in the dedicated hardware.

**Table 16-40: Applying a `multstyle` Attribute to a Binary Expression**

HDL	Code
Verilog-2001	<pre>wire [8:0] a, b; wire [17:0] result; assign result = a * (* multstyle = "dsp" *) b;</pre>

**Note:** You cannot use Verilog-1995 attribute syntax to apply the `multstyle` attribute to a binary expression.

When applied to a VHDL entity or architecture, the attribute specifies the default implementation style for all instances of the `*` operator in the entity or architecture.

In this example, the `multstyle` attribute directs the Quartus II software to use dedicated hardware, if possible, for all multiplications inside architecture `rtl` of entity `my_entity`.

**Table 16-41: Applying a `multstyle` Attribute to an Architecture**

HDL	Code
VHDL	<pre>architecture rtl of my_entity is     attribute multstyle : string;     attribute multstyle of rtl : architecture is "dsp"; begin</pre>

When applied to a VHDL signal or variable, the attribute specifies the implementation style for all instances of the `*` operator, which has a result directly assigned to the signal or variable. The attribute overrides the `multstyle` attribute with the enclosing entity or architecture, if present.

In this example, the `multstyle` attribute associated with signal `result` directs the Quartus II software to implement `a * b` in logic rather than the dedicated hardware.

**Table 16-42: Applying a `multstyle` Attribute to a Signal or Variable**

HDL	Code
VHDL	<pre>signal a, b : unsigned(8 downto 0); signal result : unsigned(17 downto 0);  attribute multstyle : string; attribute multstyle of result : signal is "logic"; result &lt;= a * b;</pre>

## Full Case Attribute

A Verilog HDL case statement is full when its case items cover all possible binary values of the case expression or when a default case statement is present. A `full_case` attribute attached to a case statement header that is not full forces synthesis to treat the unspecified states as a don't care value. VHDL case statements must be full, so the attribute does not apply to VHDL.

Using this attribute on a case statement that is not full allows you to avoid the latch inference problems.

**Note:** Latches have limited support in formal verification tools. Do not infer latches unintentionally, for example, through an incomplete case statement when using formal verification.

Formal verification tools support the `full_case` synthesis attribute (with limited support for attribute syntax, as described in [Synthesis Attributes](#) on page 16-22).

Using the `full_case` attribute might cause a simulation mismatch between the Verilog HDL functional and the post-Quartus II simulation because unknown case statement cases can still function as latches during functional simulation. For example, a simulation mismatch can occur with the code in [Table 16-43](#) when `sel` is `2'b11` because a functional HDL simulation output behaves as a latch and the Quartus II simulation output behaves as a don't care value.

**Note:** Altera recommends making the case statement “full” in your regular HDL code, instead of using the `full_case` attribute.

**Table 16-43: A full\_case Attribute**

The case statement in this example is not full because you do not specify some `sel` binary values. Because you use the `full_case` attribute, synthesis treats the output as “don't care” when the `sel` input is `2'b11`.

HDL	Code
Verilog HDL	<pre> module full_case (a, sel, y);   input [3:0] a;   input [1:0] sel;   output y;   reg y;   always @ (a or sel)   case (sel)                                // synthesis full_case     2'b00: y=a[0];     2'b01: y=a[1];     2'b10: y=a[2];   endcase endmodule </pre>

Verilog-2001 syntax also accepts the statements in [Table 16-44](#) in the `case` header instead of the comment form as shown in [Table 16-43](#).

**Table 16-44: Syntax for the full\_case Attribute**

HDL	Syntax
Verilog-2001	<code>(* full_case *) case (sel)</code>

#### Related Information

[Synthesis Attributes](#) on page 16-22

[Recommended Design Practices](#) on page 11-1

For more information about avoiding latch inference problems

## Parallel Case

The `parallel_case` attribute indicates that you must consider a Verilog HDL case statement as parallel; that is, you can match only one case item at a time. Case items in Verilog HDL case statements might

overlap. To resolve multiple matching case items, the Verilog HDL language defines a priority among case items in which the case statement always executes the first case item that matches the case expression value. By default, the Quartus II software implements the extra logic necessary to satisfy this priority relationship.

Attaching a `parallel_case` attribute to a case statement header allows the Quartus II software to consider its case items as inherently parallel; that is, at most one case item matches the case expression value. Parallel case items simplify the generated logic.

In VHDL, the individual choices in a case statement might not overlap, so they are always parallel and this attribute does not apply.

Altera recommends that you use this attribute only when the `case` statement is truly parallel. If you use the attribute in any other situation, the generated logic does not match the functional simulation behavior of the Verilog HDL.

**Note:** Altera recommends that you avoid using the `parallel_case` attribute, because you may mismatch the Verilog HDL functional and the post-Quartus II simulation.

If you specify `SystemVerilog-2005` as the supported Verilog HDL version for your design, you can use the `SystemVerilog` keyword `unique` to achieve the same result as the `parallel_case` directive without causing simulation mismatches.

This example shows a `casez` statement with overlapping case items. In functional HDL simulation, the software prioritizes the three case items by the bits in `sel`. For example, `sel[2]` takes priority over `sel[1]`, which takes priority over `sel[0]`. However, the synthesized design can simulate differently because the `parallel_case` attribute eliminates this priority. If more than one bit of `sel` is high, more than one output (`a`, `b`, or `c`) is high as well, a situation that cannot occur in functional HDL simulation.

**Table 16-45: A `parallel_case` Attribute**

HDL	Code
Verilog HDL	<pre> module parallel_case (sel, a, b, c);   input [2:0] sel;   output a, b, c;   reg a, b, c;   always @ (sel)   begin     {a, b, c} = 3'b0;     casez (sel) // synthesis parallel_case       3'b1??: a = 1'b1;       3'b?1?: b = 1'b1;       3'b??1: c = 1'b1;     endcase   end endmodule </pre>



**Table 16-46: Verilog-2001 Syntax**

Verilog-2001 syntax also accepts the statements as shown in the following table in the `case` (or `casez`) header instead of the comment form, as shown in [Table 16-45](#).

HDL	Syntax
Verilog-2001	<code>(* parallel_case *) casez (sel)</code>

## Translate Off and On / Synthesis Off and On

The `translate_off` and `translate_on` synthesis directives indicate whether the Quartus II software or a third-party synthesis tool should compile a portion of HDL code that is not relevant for synthesis. The `translate_off` directive marks the beginning of code that the synthesis tool should ignore; the `translate_on` directive indicates that synthesis should resume. You can also use the `synthesis_on` and `synthesis_off` directives as a synonym for translate on and off.

You can use these directives to indicate a portion of code for simulation only. The synthesis tool reads synthesis-specific directives and processes them during synthesis; however, third-party simulation tools read the directives as comments and ignore them.

These examples show these directives.

**Table 16-47: Translate Off and On**

HDL	Code
Verilog HDL	<pre>// synthesis translate_off parameter tpd = 2; // Delay for simulation #tpd; // synthesis translate_on</pre>
VHDL	<pre>-- synthesis translate_off use std.textio.all; -- synthesis translate_on</pre>
VHDL 2008	<pre>/* synthesis translate_off */ use std.textio.all; /* synthesis translate_on */</pre>

If you want to ignore only a portion of code in Quartus II Integrated Synthesis, you can use the Altera-specific attribute keyword `altera`. For example, use the `// altera translate_off` and `// altera translate_on` directives to direct Quartus II Integrated Synthesis to ignore a portion of code that you intend only for other synthesis tools.

## Ignore `translate_off` and `synthesis_off` Directives

The **Ignore `translate_off` and `synthesis_off` Directives** logic option directs Quartus II Integrated Synthesis to ignore the `translate_off` and `synthesis_off` directives. Turning on this logic option allows you to compile code that you want the third-party synthesis tools to ignore; for example, IP core declarations that the other tools treat as black boxes but the Quartus II software can compile. To set the

Ignore `translate_off` and `synthesis_off` Directives logic option, click **Assignments > Settings > Compiler Settings > Advanced Settings (Synthesis)**.

**Related Information**

**Ignore `translate_off` and `synthesis_off` Directives logic option**

For more information about the **Ignore `translate_off` and `synthesis_off` Directives** logic option and the supported devices

## Read Comments as HDL

The `read_comments_as_HDL` synthesis directive indicates that the Quartus II software should compile a portion of HDL code that you commented out. This directive allows you to comment out portions of HDL source code that are not relevant for simulation, while instructing the Quartus II software to read and synthesize that same source code. Setting the `read_comments_as_HDL` directive to `on` indicates the beginning of commented code that the synthesis tool should read; setting the `read_comments_as_HDL` directive to `off` indicates the end of the code.

**Note:** You can use this directive with `translate_off` and `translate_on` to create one HDL source file that includes an IP core instantiation for synthesis and a behavioral description for simulation.

Formal verification tools do not support the `read_comments_as_HDL` directive because the tools do not recognize the directive.

In these examples, the Compiler synthesizes the commented code enclosed by `read_comments_as_HDL` because the directive is visible to the Quartus II Compiler. VHDL 2008 allows block comments, which comments are also supported for synthesis directives.

**Note:** Because synthesis directives are case sensitive in Verilog HDL, you must match the case of the directive, as shown in the following examples.

**Table 16-48: Read Comments as HDL**

HDL	Code
Verilog HDL	<pre>// synthesis read_comments_as_HDL on // my_rom lpm_rom      (.address (address), //                      .data    (data)); // synthesis read_comments_as_HDL off</pre>
VHDL	<pre>-- synthesis read_comments_as_HDL on -- my_rom : entity lpm_rom --   port map ( --     address =&gt; address, --     data    =&gt; data,    ); -- synthesis read_comments_as_HDL off</pre>
VHDL 2008	<pre>/* synthesis read_comments_as_HDL on */ /* my_rom : entity lpm_rom    port map (      address =&gt; address,      data =&gt; data, ); */ synthesis read_comments_as_HDL off */</pre>

## Use I/O Flipflops

The `useioff` attribute directs the Quartus II software to implement input, output, and output enable flipflops (or registers) in I/O cells that have fast, direct connections to an I/O pin, when possible. To improve I/O performance by minimizing setup, clock-to-output, and clock-to-output enable times, you can apply the `useioff` synthesis attribute. The **Fast Input Register**, **Fast Output Register**, and **Fast Output Enable Register** logic options support this synthesis attribute. You can also set this synthesis attribute in the Assignment Editor.

The `useioff` synthesis attribute takes a boolean value. You can apply the value only to the port declarations of a top-level Verilog HDL module or VHDL entity (it is ignored if applied elsewhere). Setting the value to 1 (Verilog HDL) or `TRUE` (VHDL) instructs the Quartus II software to pack registers into I/O cells. Setting the value to 0 (Verilog HDL) or `FALSE` (VHDL) prevents register packing into I/O cells.

In [Table 16-49](#) and [Table 16-50](#), the `useioff` synthesis attribute directs the Quartus II software to implement the `a_reg`, `b_reg`, and `o_reg` registers in the I/O cells corresponding to the `a`, `b`, and `o` ports, respectively.

**Table 16-49: Verilog HDL Code: The `useioff` Attribute**

HDL	Code
Verilog HDL	<pre> module top_level(clk, a, b, o);   input clk;   input [1:0] a, b /* synthesis useioff = 1 */;   output [2:0] o /* synthesis useioff = 1 */;   reg [1:0] a_reg, b_reg;   reg [2:0] o_reg;   always @ (posedge clk)   begin     a_reg &lt;= a;     b_reg &lt;= b;     o_reg &lt;= a_reg + b_reg;   end   assign o = o_reg; endmodule </pre>

[Table 16-50](#) and [Table 16-51](#) show that the Verilog-2001 syntax also accepts the type of statements instead of the comment form in [Table 16-49](#).

**Table 16-50: Verilog-2001 Code: the `useioff` Attribute**

HDL	Code
Verilog-2001	<pre> (* useioff = 1 *) input [1:0] a, b; (* useioff = 1 *) output [2:0] o; </pre>

**Table 16-51: VHDL Code: the useioff Attribute**

HDL	Code
VHDL	<pre> library ieee; use ieee.std_logic_1164.all; use ieee.numeric_std.all; entity useioff_example is   port (     clk : in std_logic;     a, b : in unsigned(1 downto 0);     o : out unsigned(1 downto 0));   attribute useioff : boolean;   attribute useioff of a : signal is true;   attribute useioff of b : signal is true;   attribute useioff of o : signal is true; end useioff_example; architecture rtl of useioff_example is   signal o_reg, a_reg, b_reg : unsigned(1 downto 0); begin   process(clk)   begin     if (clk = '1' AND clk'event) then       a_reg &lt;= a;       b_reg &lt;= b;       o_reg &lt;= a_reg + b_reg;     end if;   end process;   o &lt;= o_reg; end rtl; </pre>

## Specifying Pin Locations with chip\_pin

The `chip_pin` attribute allows you to assign pin locations in your HDL source. You can use the attribute only on the ports of the top-level entity or module in your design. You can assign pins only to single-bit or one-dimensional bus ports in your design.

For single-bit ports, the value of the `chip_pin` attribute is the name of the pin on the target device, as specified by the pin table of the device.

**Note:** In addition to the `chip_pin` attribute, the Quartus II software supports the `altera_chip_pin_lc` attribute name for compatibility with other synthesis tools. When using this attribute in other synthesis tools, some older device families require an “@” symbol in front of each pin assignment. In the Quartus II software, the “@” is optional.

**Table 16-52: Applying Chip Pin to a Single Pin**

These examples in this table show different ways of assigning `my_pin1` to Pin C1 and `my_pin2` to Pin 4 on a different target device.

HDL	Code
Verilog-1995	<pre> input my_pin1 /* synthesis chip_pin = "C1" */; input my_pin2 /* synthesis altera_chip_pin_lc = "@4" */; </pre>
Verilog-2001	<pre> (* chip_pin = "C1" *) input my_pin1; (* altera_chip_pin_lc = "@4" *) input my_pin2; </pre>

HDL	Code
VHDL	<pre>entity my_entity is port(my_pin1: in std_logic; my_pin2: in std_logic;...); end my_entity; attribute chip_pin : string; attribute altera_chip_pin_lc : string; attribute chip_pin of my_pin1 : signal is "C1"; attribute altera_chip_pin_lc of my_pin2 : signal is "@4";</pre>

For bus I/O ports, the value of the chip pin attribute is a comma-delimited list of pin assignments. The order in which you declare the range of the port determines the mapping of assignments to individual bits in the port. To leave a bit unassigned, leave its corresponding pin assignment blank.

**Table 16-53: Applying Chip Pin to a Bus of Pins**

The example in this table assigns `my_pin[2]` to `Pin_4`, `my_pin[1]` to `Pin_5`, and `my_pin[0]` to `Pin_6`.

HDL	Code
Verilog-1995	<pre>input [2:0] my_pin /* synthesis chip_pin = "4, 5, 6" */;</pre>

**Table 16-54: Applying Chip Pin to Part of a Bus**

The example in this table reverses the order of the signals in the bus, assigning `my_pin[0]` to `Pin_4` and `my_pin[2]` to `Pin_6` but leaves `my_pin[1]` unassigned.

HDL	Code
Verilog-1995	<pre>input [0:2] my_pin /* synthesis chip_pin = "4, ,6" */;</pre>

**Table 16-55: Applying Chip Pin to Part of a Bus of Pins**

The example in this table assigns `my_pin[2]` to `Pin 4` and `my_pin[0]` to `Pin 6`, but leaves `my_pin[1]` unassigned.

HDL	Code
VHDL	<pre>entity my_entity is port(my_pin: in std_logic_vector(2 downto 0);...); end my_entity; attribute chip_pin of my_pin: signal is "4, , 6";</pre>

**Table 16-56: VHDL and Verilog-2001 Examples: Assigning Pin Location and I/O Standard**

HDL	Code
VHDL	<pre>attribute altera_chip_pin_lc: string; attribute altera_attribute: string; attribute altera_chip_pin_lc of clk: signal is "B13"; attribute altera_attribute of clk:signal is "-name IO_STANDARD \"3.3-V LVCOS\"";</pre>

HDL	Code
Verilog-2001	<pre>(* altera_attribute = "-name IO_STANDARD \"3.3-V LVCMS\"")(* chip_pin = "L5" *)input clk; (* altera_attribute = "-name IO_STANDARD LVDS" *)(* chip_pin = "L4" *)input sel; output [3:0] data_o, input [3:0] data_i);</pre>

## Using altera\_attribute to Set Quartus II Logic Options

The `altera_attribute` attribute allows you to apply Quartus II logic options and assignments to an object in your HDL source code. You can set this attribute on an entity, architecture, instance, register, RAM block, or I/O pin. You cannot set it on an arbitrary combinational node such as a net. With `altera_attribute`, you can control synthesis options from your HDL source even when the options lack a specific HDL synthesis attribute. You can also use this attribute to pass entity-level settings and assignments to phases of the Compiler flow that follow Analysis & Synthesis, such as Fitting.

Assignments or settings made through the Quartus II software, the `.qsf`, or the Tcl interface take precedence over assignments or settings made with the `altera_attribute` synthesis attribute in your HDL code.

The attribute value is a single string containing a list of `.qsf` variable assignments separated by semicolons:

```
-name <variable_1> <value_1>;-name <variable_2> <value_2>[;...]
```

If the Quartus II option or assignment includes a target, source, and section tag, you must use the syntax in this example for each `.qsf` variable assignment:

```
-name <variable> <value>
-from <source> -to <target> -section_id <section>
```

This example shows the syntax for the full attribute value, including the optional target, source, and section tags for two different `.qsf` assignments:

```
" -name <variable_1> <value_1> [-from <source_1>] [-to <target_1>] [-section_id \
<section_1>]; -name <variable_2> <value_2> [-from <source_2>] [-to <target_2>] \
[-section_id <section_2>] "
```

### Table 16-57: Example Usage

If the assigned value of a variable is a string of text, you must use escaped quotes around the value in Verilog HDL or double-quotes in VHDL:

HDL	Code
Assigned Value of a Variable in Verilog HDL (With Nonexistent Variable and Value Terms)	"VARIABLE_NAME \"STRING_VALUE\" "
Assigned Value of a Variable in VHDL (With Nonexistent Variable and Value Terms)	"VARIABLE_NAME " "STRING_VALUE" " "

To find the `.qsf` variable name or value corresponding to a specific Quartus II option or assignment, you can set the option setting or assignment in the Quartus II software, and then make the changes in the `.qsf`.

## Applying altera\_attribute to an Instance

These examples use `altera_attribute` to set the power-up level of an inferred register.

**Table 16-58: Applying altera\_attribute to an Instance**

These examples use `altera_attribute` to set the power-up level of an inferred register.

HDL	Code
Verilog-1995	<pre>reg my_reg /* synthesis altera_attribute = "-name POWER_UP_LEVEL HIGH" */;</pre>
Verilog-2001	<pre>(* altera_attribute = "-name POWER_UP_LEVEL HIGH" *) reg my_reg;</pre>
VHDL	<pre>signal my_reg : std_logic; attribute altera_attribute : string; attribute altera_attribute of my_reg: signal is "-name POWER_UP_LEVEL HIGH";</pre>

**Note:** For inferred instances, you cannot apply the attribute to the instance directly. Therefore, you must apply the attribute to one of the output nets of the instance. The Quartus II software automatically moves the attribute to the inferred instance.

## Applying altera\_attribute to an Entity

These examples use the `altera_attribute` to disable the **Auto Shift Register Replacement** synthesis option for an entity. To apply the Altera Attribute to a VHDL entity, you must set the attribute on its architecture rather than on the entity itself.

**Table 16-59: Applying altera\_attribute to an Entity**

HDL	Code
Verilog-1995	<pre>module my_entity(...) /* synthesis altera_attribute = "-name AUTO_SHIFT_REGISTER_RECOGNITION OFF" */;</pre>
Verilog-2001	<pre>(* altera_attribute = "-name AUTO_SHIFT_REGISTER_RECOGNITION OFF" *) module my_entity(...) ;</pre>
VHDL	<pre>entity my_entity is -- Declare generics and ports end my_entity; architecture rtl of my_entity is attribute altera_attribute : string; -- Attribute set on architecture, not entity attribute altera_attribute of rtl: architecture is "-name AUTO_SHIFT_REGISTER_RECOGNITION OFF"; begin -- The architecture body end rtl;</pre>

## Applying altera\_attribute with the -to Option

You can also use `altera_attribute` for more complex assignments that have more than one instance. In [Table 16-60](#), the `altera_attribute` cuts all timing paths from `reg1` to `reg2`, equivalent to this Tcl or `.qsf` command, as shown in the example below:

```
set_instance_assignment -name CUT ON -from reg1 -to reg2
```

**Table 16-60: Applying altera\_attribute with the -to Option**

HDL	Code
Verilog-1995	<pre>reg reg2; reg reg1 /* synthesis altera_attribute = "-name CUT ON -to reg2" */;</pre>
Verilog-2001 and SystemVerilog	<pre>reg reg2; (* altera_attribute = "-name CUT ON -to reg2" *) reg reg1;</pre>
VHDL	<pre>signal reg1, reg2 : std_logic; attribute altera_attribute: string; attribute altera_attribute of reg1 : signal is "-name CUT ON -to reg2";</pre>

You can specify either the `-to` option or the `-from` option in a single `altera_attribute`; Integrated Synthesis automatically sets the remaining option to the target of the `altera_attribute`. You can also specify wildcards for either option. For example, if you specify "\*" for the `-to` option instead of `reg2` in these examples, the Quartus II software cuts all timing paths from `reg1` to every other register in this design entity.

You can use the `altera_attribute` only for entity-level settings, and the assignments (including wildcards) apply only to the current entity.

### Related Information

[Synthesis Attributes](#) on page 16-22

[Quartus II Settings File Manual](#)

Lists all variable names

## Analyzing Synthesis Results

After performing synthesis, you can check your synthesis results in the **Analysis & Synthesis** section of the Compilation Report and the Project Navigator.

### Analysis & Synthesis Section of the Compilation Report

The Compilation Report, which provides a summary of results for the project, appears after a successful compilation. After Analysis & Synthesis, the Summary section of the Compilation Report provides a summary of utilization based on synthesis data, before Fitter optimizations have occurred. The **Analysis & Synthesis** section lists synthesis-specific information.



Analysis & Synthesis includes various report sections, including a list of the source files read for the project, the resource utilization by entity after synthesis, and information about state machines, latches, optimization results, and parameter settings.

#### Related Information

##### [Analysis Synthesis Summary Reports](#)

For more information about each report section

## Project Navigator

The **Hierarchy** tab of the Project Navigator provides a view of the project hierarchy and a summary of resource and device information about the current project. After Analysis & Synthesis, before the Fitter begins, the Project Navigator provides a summary of utilization based on synthesis data, before Fitter optimizations have occurred.

If an entity in the Hierarchy tab contains parameter settings, a tooltip displays the settings when you hold the pointer over the entity.

## Upgrade IP Components Dialog Box

In the Quartus II software version 12.1 SP1 and later, the **Upgrade IP Components** dialog box allows you to upgrade all outdated IP in your project after you move to a newer version of the Quartus II software.

#### Related Information

##### [Upgrade IP Components dialog box](#)

For more information about the Upgrade IP Components dialog box

## Analyzing and Controlling Synthesis Messages

You can analyze the generated messages during synthesis and control which messages appear during compilation.

## Quartus II Messages

The messages that appear during Analysis & Synthesis describe many of the optimizations during the synthesis stage, and provide information about how the software interprets your design. Altera recommends checking the messages to analyze **Critical Warnings** and **Warnings**, because these messages can relate to important design problems. Read the **Info** messages to get more information about how the software processes your design.

The software groups the messages by following types: **Info**, **Warning**, **Critical Warning**, and **Error**.

You can specify the type of Analysis & Synthesis messages that you want to view by selecting the **Analysis & Synthesis Message Level** option. To specify the display level, click **Assignments > Settings > Compiler Settings > Advanced Settings (Synthesis)**

#### Related Information

- [About the Messages Window](#)  
For more information about the Messages window and message suppression
- [About Message Suppression](#)  
For more information about the Messages window and message suppression

- [Managing Quartus II Projects](#) on page 1-1  
For more information about the Messages

## VHDL and Verilog HDL Messages

The Quartus II software issues a variety of messages when it is analyzing and elaborating the Verilog HDL and VHDL files in your design. These HDL messages are a subset of all Quartus II messages that help you identify potential problems early in the design process.

HDL messages fall into the following categories:

- **Info message**—lists a property of your design.
- **Warning message**—indicates a potential problem in your design. Potential problems come from a variety of sources, including typos, inappropriate design practices, or the functional limitations of your target device. Though HDL warning messages do not always identify actual problems, Altera recommends investigating code that generates an HDL warning. Otherwise, the synthesized behavior of your design might not match your original intent or its simulated behavior.
- **Error message**—indicates an actual problem with your design. Your HDL code can be invalid due to a syntax or semantic error, or it might not be synthesizable as written.

In this example, the sensitivity list contains multiple copies of the variable `i`. While the Verilog HDL language does not prohibit duplicate entries in a sensitivity list, it is clear that this design has a typing error: Variable `j` should be listed on the sensitivity list to avoid a possible simulation or synthesis mismatch.

```
//dup.v
module dup(input i, input j, output reg o);
always @ (i or i)
    o = i & j;
endmodule
```

When processing the HDL code, the Quartus II software generates the following warning message.

```
Warning: (10276) Verilog HDL sensitivity list warning at dup.v(2): sensitivity list
contains multiple entries for "i".
```

In Verilog HDL, variable names are case sensitive, so the variables `my_reg` and `MY_REG` below are two different variables. However, declaring variables that have names in different cases is confusing, especially if you use VHDL, in which variables are not case sensitive.

```
// namecase.v
module namecase (input i, output o);
    reg my_reg;
    reg MY_REG;
    assign o = i;
endmodule
```

When processing the HDL code, the Quartus II software generates the following informational message:

```
Info: (10281) Verilog HDL information at namecase.v(3): variable name "MY_REG" and
variable name "my_reg" should not differ only in case.
```

In addition, the Quartus II software generates additional HDL info messages to inform you that this small design does not use neither `my_reg` nor `MY_REG`:

```
Info: (10035) Verilog HDL or VHDL information at namecase.v(3): object "my_reg"
declared but not used
Info: (10035) Verilog HDL or VHDL information at namecase.v(4): object "MY_REG"
declared but not used
```

The Quartus II software allows you to control how many HDL messages you can view during the Analysis & Elaboration of your design files. You can set the HDL Message Level to enable or disable groups of HDL messages, or you can enable or disable specific messages.

#### Related Information

- [Synthesis Directives](#) on page 16-25  
For more information about synthesis directives and their syntax

## Setting the HDL Message Level

The HDL Message Level specifies the types of messages that the Quartus II software displays when it is analyzing and elaborating your design files.

**Table 16-61: HDL Info Message Level**

Level	Purpose	Description
<b>Level1</b>	High-severity messages only	If you want to view only the HDL messages that identify likely problems with your design, select Level1. When you select Level1, the Quartus II software issues a message only if there is an actual problem with your design.
<b>Level2</b>	High-severity and medium-severity messages	If you want to view additional HDL messages that identify possible problems with your design, select Level2. Level2 is the default setting.
<b>Level3</b>	All messages, including low-severity messages	If you want to view all HDL info and warning messages, select Level3. This level includes extra “LINT” messages that suggest changes to improve the style of your HDL code.

You must address all issues reported at the **Level1** setting. The default HDL message level is **Level2**.

To set the HDL Message Level in the Quartus II software, follow these steps:

1. Click **Assignments > Settings > Compiler Settings > Advanced Settings (Synthesis)**
2. Set the necessary message level from the pull-down menu in the **HDL Message Level** list, and then click **OK**.

You can override this default setting in a source file with the `message_level synthesis` directive, which takes the values `level1`, `level2`, and `level3`, as shown in the following table.

**Table 16-62: HDL Examples of message\_level Directive**

HDL	Code
Verilog HDL	<pre>// altera message_level level1 or /* altera message_level level3 */</pre>
VHDL	<pre>-- altera message_level level2</pre>

A `message_level` synthesis directive remains effective until the end of a file or until the next `message_level` directive. In VHDL, you can use the `message_level` synthesis directive to set the HDL Message Level for entities and architectures, but not for other design units. An HDL Message Level for an entity applies to its architectures, unless overridden by another `message_level` directive. In Verilog HDL, you can use the `message_level` directive to set the HDL Message Level for a module.

### Enabling or Disabling Specific HDL Messages by Module/Entity

Message ID is in parentheses at the beginning of the message. Use the Message ID to enable or disable a specific HDL info or warning message. Enabling or disabling a specific message overrides its HDL Message Level. This method is different from the message suppression in the Messages window because you can disable messages for a specific module or a specific entity. This method applies only to the HDL messages, and if you disable a message with this method, the Quartus II software lists the message as a suppressed message.

To disable specific HDL messages in the Quartus II software, follow these steps:

1. Click **Assignments > Settings > Compiler Settings > Advanced Settings (Synthesis)**.
2. In the **Advanced Message Settings** dialog box, add the Message IDs you want to enable or disable.

To enable or disable specific HDL messages in your HDL, use the `message_on` and `message_off` synthesis directives. These directives require a space-separated list of Message IDs. You can enable or disable messages with these synthesis directives immediately before Verilog HDL modules, VHDL entities, or VHDL architectures. You cannot enable or disable a message during an HDL construct.

A message enabled or disabled via a `message_on` or `message_off` synthesis directive overrides its HDL Message Level or any `message_level` synthesis directive. The message remains disabled until the end of the source file or until you use another `message_on` or `message_off` directive to change the status of the message.

**Table 16-63: HDL message\_off Directive for Message with ID 10000**

HDL	Code
Verilog HDL	<pre>// altera message_off 10000 or /* altera message_off 10000 */</pre>
VHDL	<pre>-- altera message_off 10000</pre>

## Node-Naming Conventions in Quartus II Integrated Synthesis

Whenever possible, Quartus II Integrated Synthesis uses wire or signal names from your source code to name nodes such as LEs or ALMs. Some nodes, such as registers, have predictable names that do not change when a design is resynthesized, although certain optimizations can affect register names. The names of other nodes, particularly LEs or ALMs that contain only combinational logic, can change due to logic optimizations that the software performs.

### Hierarchical Node-Naming Conventions

To make each name in your design unique, the Quartus II software adds the hierarchy path to the beginning of each name. The “|” separator indicates a level of hierarchy. For each instance in the hierarchy, the software adds the entity name and the instance name of that entity, with the “:” separator between each entity name and its instance name. For example, if a design defines entity `A` with the name `my_A_inst`, the hierarchy path of that entity would be `A:my_A_inst`. You can obtain the full name of any node by starting with the hierarchical instance path, followed by a “|”, and ending with the node name inside that entity.

This example shows you the convention:

```
<entity 0>:<instance_name 0>|<entity 1>:<instance_name 1>|...|<instance_name n>|
<node_name>
```

For example, if entity `A` contains a register (DFF atom) called `my_dff`, its full hierarchy name would be `A:my_A_inst|my_dff`.

To instruct the Compiler to generate node names that do not contain entity names, on the **Compilation Process Settings** page of the **Settings** dialog box, click **More Settings**, and then turn off **Display entity name for node name**.

With this option turned off, the node names use the convention in shown in this example:

```
<instance_name 0>|<instance_name 1>|...|<instance_name n> |<node_name>
```

### Node-Naming Conventions for Registers (DFF or D Flipflop Atoms)

In Verilog HDL and VHDL, inferred registers use the names of the `reg` or `signal` connected to the output.

**Table 16-64: HDL Example of a Register that Creates `my_dff_out` DFF Primitive**

HDL Register	Code
Verilog HDL	<pre>wire dff_in, my_dff_out, clk; always @ (posedge clk) my_dff_out &lt;= dff_in;</pre>

HDL Register	Code
VHDL	<pre> signal dff_in, my_dff_out, clk; process (clk) begin if (rising_edge(clk)) then my_dff_out &lt;= dff_in; end if; end process; </pre>

AHDL designs explicitly declare DFF registers rather than infer, so the software uses the user-declared name for the register.

For schematic designs using a **.bdf**, your design names all elements when you instantiate the elements in your design, so the software uses the name you defined for the register or DFF.

In the special case that a wire or signal (such as `my_dff_out` in the preceding examples) is also an output pin of your top-level design, the Quartus II software cannot use that name for the register (for example, cannot use `my_dff_out`) because the software requires that all logic and I/O cells have unique names. Here, Quartus II Integrated Synthesis appends `~reg0` to the register name.

**Table 16-65: Verilog HDL Register Feeding Output Pin**

For example, the Verilog HDL code example in this table generates a register called `q~reg0`.

HDL	Code
Verilog HDL	<pre> module my_dff (input clk, input d, output q); always @ (posedge clk) q &lt;= d; endmodule </pre>

This situation occurs only for registers driving top-level pins. If a register drives a port of a lower level of the hierarchy, the software removes the port during hierarchy flattening and the register retains its original name, in this case, `q`.

## Register Changes During Synthesis

On some occasions, you might not find registers that you expect to view in the synthesis netlist. Logic optimization might remove registers and synthesis optimizations might change the names of the registers. Common optimizations include inference of a state machine, counter, adder-subtractor, or shift register from registers and surrounding logic. Other common register changes occur when the software packs these registers into dedicated hardware on the FPGA, such as a DSP block or a RAM block.

The following factors can affect register names:

- [Synthesis and Fitting Optimizations](#) on page 16-72
- [State Machines](#) on page 16-72
- [Inferred Adder-Subtractors, Shift Registers, Memory, and DSP Functions](#) on page 16-73
- [Packed Input and Output Registers of RAM and DSP Blocks](#) on page 16-73

## Synthesis and Fitting Optimizations

Logic optimization during synthesis might remove registers if you do not connect the registers to inputs or outputs in your design, or if you can simplify the logic due to constant signal values. Synthesis optimizations might change register names, such as when the software merges duplicate registers to reduce resource utilization.

NOT-gate push back optimizations can affect registers that use preset signals. This type of optimization can impact your timing assignments when the software uses registers as clock dividers. If this situation occurs in your design, change the clock settings to work on the new register name.

Synthesis netlist optimizations often change node names because the software can combine or duplicate registers to optimize your design.

The Quartus II Compilation Report provides a list of registers that synthesis optimizations remove, and a brief reason for the removal. To generate the Quartus II Compilation Report, follow these steps:

1. In the **Analysis & Synthesis** folder, open **Optimization Results**.
2. Open **Register Statistics**, and then click the **Registers Removed During Synthesis** report.
3. Click **Removed Registers Triggering Further Register Optimizations**.

The second report contains a list of registers that causes synthesis optimizations to remove other registers from your design. The report provides a brief reason for the removal, and a list of registers that synthesis optimizations remove due to the removal of the initial register.

Quartus II Integrated Synthesis creates synonyms for registers duplicated with the **Maximum Fan-Out** option (or `maxfan` attribute). Therefore, timing assignments applied to nodes that are duplicated with this option are applied to the new nodes as well.

The Quartus II Fitter can also change node names after synthesis (for example, when the Fitter uses register packing to pack a register into an I/O element, or when physical synthesis modifies logic). The Fitter creates synonyms for duplicated registers so timing analysis can use the existing node name when applying assignments.

You can instruct the Quartus II software to preserve certain nodes throughout compilation so you can use them for verification or making assignments.

### Related Information

#### [Netlist Optimizations and Physical Synthesis](#)

For more information about the type of optimizations performed by synthesis netlist optimizations

#### [Preserving Register Names](#) on page 16-73

For more information about preserving certain nodes throughout compilation

## State Machines

If your HDL code infers a state machine, the software maps the registers that represent the states into a new set of registers that implement the state machine. Most commonly, the software converts the state machine into a one-hot form in which one register represents each state. In this case, for Verilog HDL or VHDL designs, the registers take the name of the state register and the states.

For example, consider a Verilog HDL state machine in which the states are `parameter state0 = 1, state1 = 2, state2 = 3`, and in which the software declares the state machine register as `reg [1:0] my_fsm`. In this example, the three one-hot state registers are `my_fsm.state0`, `my_fsm.state1`, and `my_fsm.state2`.

An AHDL design explicitly specifies state machines with a state machine name. Your design names state machine registers with synthesized names based on the state machine name, but not the state names. For example, if a `my_fsm` state machine has four state bits, The software might synthesize these state bits with names such as `my_fsm~12`, `my_fsm~13`, `my_fsm~14`, and `my_fsm~15`.

## Inferred Adder-Subtractors, Shift Registers, Memory, and DSP Functions

The Quartus II software infers IP cores from Verilog HDL and VHDL code for logic that forms adder-subtractors, shift registers, RAM, ROM, and arithmetic functions that are placed in DSP blocks.

Because adder-subtractors are part of an IP core instead of generic logic, the combinational logic exists in the design with different names. For shift registers, memory, and DSP functions, the software implements the registers and logic inside the dedicated RAM or DSP blocks in the device. Thus, the registers are not visible as separate LEs or ALMs.

### Related Information

- [Recommended HDL Coding Styles](#) on page 12-1  
For information about inferring IP cores

## Packed Input and Output Registers of RAM and DSP Blocks

The software packs registers into the input registers and output registers of RAM and DSP blocks, so that they are not visible as separate registers in LEs or ALMs.

### Related Information

- [Recommended HDL Coding Styles](#) on page 12-1  
For information about packing registers into RAM and DSP IP cores

## Preserving Register Names

Altera recommends that you preserve certain register names for verification or debugging, or to ensure that you applied timing assignments correctly. Quartus II Integrated Synthesis preserves certain nodes automatically if the software uses the nodes in a timing constraint.

### Related Information

- [Preserve Registers](#) on page 16-38  
Use the `preserve` attribute to instruct the Compiler not to minimize or remove a specified register during synthesis optimizations or register netlist optimizations
- [Noprune Synthesis Attribute/Preserve Fan-out Free Register Node](#) on page 16-39  
Use the `noprune` attribute to preserve a fan-out-free register through the entire compilation flow
- [Disable Register Merging/Don't Merge Register](#) on page 16-39  
Use the synthesis attribute `syn_dont_merge` to ensure that the Compiler does not merge registers with other registers

## Node-Naming Conventions for Combinational Logic Cells

Whenever possible for Verilog HDL, VHDL, and AHDL code, the Quartus II software uses wire names that are the targets of assignments, but can change the node names due to synthesis optimizations.



For example, consider the Verilog HDL code in this example. Quartus II Integrated Synthesis uses the names `c`, `d`, `e`, and `f` for the generated combinational logic cells.

```

wire c;
reg d, e, f;
assign c = a | b;
always @ (a or b)
d = a & b;
always @ (a or b) begin : my_label
e = a ^ b;
end
always @ (a or b)
f = ~(a | b);

```

For schematic designs using a **.bdf**, your design names all elements when you instantiate the elements in your design and the software uses the name you defined when possible.

If logic cells are packed with registers in device architectures such as the Stratix and Cyclone device families, those names might not appear in the netlist after fitting. In other devices, such as newer families in the Stratix and Cyclone series device families, the register and combinational nodes are kept separate throughout the compilation, so these names are more often maintained through fitting.

When logic optimizations occur during synthesis, it is not always possible to retain the initial names as described. Sometimes, synthesized names are used, which are the wire names with a tilde (~) and a number appended. For example, if a complex expression is assigned to wire `w` and that expression generates several logic cells, those cells can have names such as `w`, `w~1`, and `w~2`. Sometimes the original wire name `w` is removed, and an arbitrary name such as `rt1~123` is created. Quartus II Integrated Synthesis attempts to retain user names whenever possible. Any node name ending with `~<number>` is a name created during synthesis, which can change if the design is changed and re-synthesized. Knowing these naming conventions helps you understand your post-synthesis results, helping you to debug your design or create assignments.

During synthesis, the software maintains combinational clock logic by not changing nodes that might be clocks. The software also maintains or protects multiplexers in clock trees, so that the TimeQuest analyzer has information about which paths are unate, to allow complete and correct analysis of combinational clocks. Multiplexers often occur in clock trees when the software selects between different clocks. To help with the analysis of clock trees, the software ensures that each multiplexer encountered in a clock tree is broken into 2:1 multiplexers, and each of those 2:1 multiplexers is mapped into one lookup table (independent of the device family). This optimization might result in a slight increase in area, and for some designs a decrease in timing performance. To disable the option, click **Assignments > Settings > Compiler Settings > Advanced Settings (Synthesis) > Clock MUX Protection**.

#### Related Information

##### [Clock MUX Protection logic option](#)

For more information about Clock MUX Protection logic option and a list of supported devices

## Preserving Combinational Logic Names

You can preserve certain combinational logic node names for verification or debugging, or to ensure that timing assignments are applied correctly.

Use the `keep` attribute to keep a wire name or combinational node name through logic synthesis minimizations and netlist optimizations.

For any internal node in your design clock network, use `keep` to protect the name so that you can apply correct clock settings. Also, set the attribute for combinational logic involved in `cut` and `-through` assignments.

**Note:** Setting the `keep` attribute for combinational logic can increase the area utilization and increase the delay of the final mapped logic because the attribute requires the insertion of extra combinational logic. Use the attribute only when necessary.

#### Related Information

- [Keep Combinational Node/Implement as Output of Logic Cell](#) on page 16-40

## Scripting Support

You can run procedures and make settings in a Tcl script. You can also run some procedures at a command prompt. For detailed information about scripting command options, refer to the Quartus II Command-Line and Tcl API Help browser.

To run the Help browser, type the command at the command prompt shown in this example:

```
quartus_sh --qhelp
```

You can specify many of the options either on an instance, at the global level, or both.

To make a global assignment, use the Tcl command shown in this example:

```
set_global_assignment -name <QSF Variable Name> <Value>
```

To make an instance assignment, use the Tcl command shown in this example:

```
set_instance_assignment -name <QSF Variable Name> <Value> \ -to <Instance Name>
```

To set the **Synthesis Effort** option at the command line, use the `--effort` option with the `quartus_map` executable shown in this example:

```
quartus_map <Design name> --effort= "auto | fast"
```

The early timing estimate feature gives you preliminary timing estimates before running a full compilation, which results in a quicker iteration time; therefore, you can save significant compilation time to get a good estimation of the final timing of your design.

If you want to run fast synthesis with the Fitter **Early Timing Estimate** option, use the command shown in this example. This command runs the full flow with timing analysis:

```
quartus_sh --flow early_timing_estimate_with_synthesis <Design name>
```

#### Related Information

- [Tcl Scripting](#)  
For more information about Tcl scripting
- [Quartus II Settings File Manual](#)  
For more information about all settings and constraints in the Quartus II software
- [Command-Line Scripting](#)  
For more information about command-line scripting

- [API Functions for Tcl](#)  
For more information about Tcl scripting
- [Synthesis Effort](#) on page 16-31

## Adding an HDL File to a Project and Setting the HDL Version

To add an HDL or schematic entry design file to your project, use the Tcl assignments shown in this example:

```
set_global_assignment -name VERILOG_FILE <file name>.<v|sv>
set_global_assignment -name SYSTEMVERILOG_FILE <file name>.sv
set_global_assignment -name VHDL_FILE <file name>.<vhd|vhdl>
set_global_assignment -name AHDL_FILE <file name>.tdf
set_global_assignment -name BDF_FILE <file name>.bdf
```

**Note:** You can use any file extension for design files, as long as you specify the correct language when adding the design file. For example, you can use **.h** for Verilog HDL header files.

To specify the Verilog HDL or VHDL version, use the option shown in this example, at the end of the VERILOG\_FILE or VHDL\_FILE command:

```
- HDL_VERSION <language version>
```

The variable <language version> takes one of the following values:

- VERILOG\_1995
- VERILOG\_2001
- SYSTEMVERILOG\_2005
- VHDL\_1987
- VHDL\_1993
- VHDL\_2008

For example, to add a Verilog HDL file called **my\_file.v** written in Verilog-1995, use the command shown in this example:

```
set_global_assignment -name VERILOG_FILE my_file.v -HDL_VERSION \ VERILOG_1995
```

In this example, the `syn_encoding` attribute associates a binary encoding with the states in the enumerated type `count_state`. In this example, the states are encoded with the following values: zero = "11", one = "01", two = "10", three = "00".

```
ARCHITECTURE rtl OF my_fsm IS
    TYPE count_state IS (zero, one, two, three);
    ATTRIBUTE syn_encoding : STRING;
    ATTRIBUTE syn_encoding OF count_state : TYPE IS "11 01 10 00";
    SIGNAL present_state, next_state : count_state;
BEGIN
```

You can also use the `syn_encoding` attribute in Verilog HDL to direct the synthesis tool to use the encoding from your HDL code, instead of using the **State Machine Processing** option.

The `syn_encoding` value "user" instructs the Quartus II software to encode each state with its corresponding value from the Verilog HDL source code. By changing the values of your state constants, you can change the encoding of your state machine.

In [Example 16-8](#), the states are encoded as follows:

```
init = "00"  
last = "11"  
next = "01"  
later = "10"
```

### Example 16-8: Verilog-2001 and SystemVerilog Code: Specifying User-Encoded States with the `syn_encoding` Attribute

```
(* syn_encoding = "user" *) reg [1:0] state;  
parameter init = 0, last = 3, next = 1, later = 2;  
always @ (state) begin  
  case (state)  
    init:  
      out = 2'b01;  
    next:  
      out = 2'b10;  
    later:  
      out = 2'b11;  
    last:  
      out = 2'b00;  
  endcase  
end
```

Without the `syn_encoding` attribute, the Quartus II software encodes the state machine based on the current value of the **State Machine Processing** logic option.

If you also specify a safe state machine (as described in [Safe State Machine](#) on page 16-35), separate the encoding style value in the quotation marks from the safe value with a comma, as follows: "safe, one-hot" or "safe, gray".

#### Related Information

- [Safe State Machine](#) on page 16-35
- [Manually Specifying State Assignments Using the `syn\_encoding` Attribute](#) on page 16-32

## Assigning a Pin

To assign a signal to a pin or device location, use the Tcl command shown in this example:

```
set_location_assignment -to <signal name> <location>
```

Valid locations are pin location names. Some device families also support edge and I/O bank locations. Edge locations are `EDGE_BOTTOM`, `EDGE_LEFT`, `EDGE_TOP`, and `EDGE_RIGHT`. I/O bank locations include `IOBANK_1` to `IOBANK_n`, where `n` is the number of I/O banks in a device.

## Creating Design Partitions for Incremental Compilation

To create a partition, use the command shown in this example:

```
set_instance_assignment -name PARTITION_HIERARCHY \  
<file name> -to <destination> -section_id <partition name>
```

The *<file name>* variable is the name used for internally generated netlist files during incremental compilation. If you create the partition in the Quartus II software, netlist files are named automatically by the Quartus II software based on the instance name. If you use Tcl to create your partitions, you must assign a custom file name that is unique across all partitions. For the top-level partition, the specified file name is ignored, and you can use any dummy value. To ensure the names are safe and platform independent, file names should be unique, regardless of case. For example, if a partition uses the file name `my_file`, no other partition can use the file name `MY_FILE`. To make file naming simple, Altera recommends that you base each file name on the corresponding instance name for the partition.

The *<destination>* is the short hierarchy path of the entity. A short hierarchy path is the full hierarchy path without the top-level name, for example: `"ram:ram_unit|altsyncram:altsyncram_component"` (with quotation marks). For the top-level partition, you can use the pipe (|) symbol to represent the top-level entity.

The *<partition name>* is the partition name you designate, which should be unique and less than 1024 characters long. The name may only consist of alphanumeric characters, as well as pipe (|), colon (:), and underscore (\_) characters. Altera recommends enclosing the name in double quotation marks ("").

#### Related Information

- [Node-Naming Conventions in Quartus II Integrated Synthesis](#) on page 16-70  
For more information about hierarchical naming conventions

## Quartus II Synthesis Options

#### Related Information

#### [Logic options](#)

For more information about the .qsf variable names and applicable values for the settings

## Document Revision History

Table 16-66: Document Revision History

Date	Version	Changes
2014.12.15	14.1.0	Updated location of Fitter Settings, Analysis & Synthesis Settings, and Physical Optimization Settings to Compiler Settings.
2014.06.30	14.0.0	Template update.
November 2013	13.1.0	<ul style="list-style-type: none"> <li>• Added a note regarding ROM inference using the <code>ram_init_file</code> in “RAM Initialization File—for Inferred Memory” on page 16-61.</li> </ul>
May 2013	13.0.0	<ul style="list-style-type: none"> <li>• Added “Verilog HDL Configuration” on page 16-6.</li> <li>• Added “RAM Style Attribute—For Shift Registers Inference” on page 16-57.</li> <li>• Added “Upgrade IP Components Dialog Box” on page 16-75.</li> </ul>
June 2012	12.0.0	<ul style="list-style-type: none"> <li>• Updated “Design Flow” on page 16-2.</li> </ul>

Date	Version	Changes
November 2011	11.1.0	<ul style="list-style-type: none"> <li>Updated “Language Support” on page 16–5, “Incremental Compilation” on page 16–22, “Quartus II Synthesis Options” on page 16–24.</li> </ul>
May 2011	11.0.0	<ul style="list-style-type: none"> <li>Updated “Specifying Pin Locations with chip_pin” on page 14–65, and “Shift Registers” on page 14–48.</li> <li>Added a link to Quartus II Help in “SystemVerilog Support” on page 14–5.</li> <li>Added Example 14–106 and Example 14–107 on page 14–67.</li> </ul>
December 2010	10.1.0	<ul style="list-style-type: none"> <li>Updated “Verilog HDL Support” on page 13–4 to include Verilog-2001 support.</li> <li>Updated “VHDL-2008 Support” on page 13–9 to include the condition operator (explicit and implicit) support.</li> <li>Rewrote “Limiting Resource Usage in Partitions” on page 13–32.</li> <li>Added “Creating LogicLock Regions” on page 13–32 and “Using Assignments to Limit the Number of RAM and DSP Blocks” on page 13–33.</li> <li>Updated “Turning Off the Add Pass-Through Logic to Inferred RAMs no_rw_check Attribute” on page 13–55.</li> <li>Updated “Auto Gated Clock Conversion” on page 13–28.</li> <li>Added links to Quartus II Help.</li> </ul>

Date	Version	Changes
July 2010	10.0.0	<ul style="list-style-type: none"> <li>• Removed Referenced Documents section.</li> <li>• Added “Synthesis Seed” on page 9–36 section.</li> <li>• Updated the following sections: <ul style="list-style-type: none"> <li>“SystemVerilog Support” on page 9–5</li> <li>“VHDL-2008 Support” on page 9–10</li> <li>“Using Parameters/Generics” on page 9–16</li> <li>“Parallel Synthesis” on page 9–21</li> <li>“Limiting Resource Usage in Partitions” on page 9–32</li> <li>“Synthesis Effort” on page 9–35</li> <li>“Synthesis Attributes” on page 9–25</li> <li>“Synthesis Directives” on page 9–27</li> <li>“Auto Gated Clock Conversion” on page 9–29</li> <li>“State Machine Processing” on page 9–36</li> <li>“Multiply-Accumulators and Multiply-Adders” on page 9–50</li> <li>“Resource Aware RAM, ROM, and Shift-Register Inference” on page 9–52</li> <li>“RAM Style and ROM Style—for Inferred Memory” on page 9–53</li> <li>“Turning Off the Add Pass-Through Logic to Inferred RAMs no_rw_check Attribute” on page 9–55</li> <li>“Using altera_attribute to Set Quartus II Logic Options” on page 9–68</li> <li>“Adding an HDL File to a Project and Setting the HDL Version” on page 9–83</li> <li>“Creating Design Partitions for Incremental Compilation” on page 9–85</li> <li>“Inferring Multiplier, DSP, and Memory Functions from HDL Code” on page 9–50</li> </ul> </li> <li>• Updated Table 9–9 on page 9–86.</li> </ul>
December 2009	9.1.1	<ul style="list-style-type: none"> <li>• Added information clarifying inheritance of Synthesis settings by lower-level entities, including Altera and third-party IP</li> <li>• Updated “Keep Combinational Node/Implement as Output of Logic Cell” on page 9–46</li> </ul>

Date	Version	Changes
November 2009	9.1.0	<ul style="list-style-type: none"> <li>• Updated the following sections:               <ul style="list-style-type: none"> <li>“Initial Constructs and Memory System Tasks” on page 9–7</li> <li>“VHDL Support” on page 9–9</li> <li>“Parallel Synthesis” on page 9–21</li> <li>“Synthesis Directives” on page 9–27</li> <li>“Timing-Driven Synthesis” on page 9–31</li> <li>“Safe State Machines” on page 9–40</li> <li>“RAM Style and ROM Style—for Inferred Memory” on page 9–53</li> <li>“Translate Off and On / Synthesis Off and On” on page 9–62</li> <li>“Read Comments as HDL” on page 9–63</li> <li>“Adding an HDL File to a Project and Setting the HDL Version” on page 9–81</li> </ul> </li> <li>• Removed “Remove Redundant Logic Cells” section</li> <li>• Added “Resource Aware RAM, ROM, and Shift-Register Inference” section</li> <li>• Updated Table 9–9 on page 9–83</li> </ul>
March 2009	9.0.0	<ul style="list-style-type: none"> <li>• Updated Table 9–9.</li> <li>• Updated the following sections:               <ul style="list-style-type: none"> <li>“Partitions for Preserving Hierarchical Boundaries” on page 9–20</li> <li>“Analysis &amp; Synthesis Settings Page of the Settings Dialog Box” on page 9–24</li> <li>“Timing-Driven Synthesis” on page 9–30</li> <li>“Turning Off Add Pass-Through Logic to Inferred RAMs/no_rw_check Attribute Setting” on page 9–54</li> </ul> </li> <li>• Added “Parallel Synthesis” on page 9–21</li> <li>• Chapter 9 was previously Chapter 8 in software version 8.1</li> </ul>

**Related Information**

- [Quartus II Handbook Archive](#)  
For previous versions of the Quartus II Handbook



2014.12.15

QII5V1



Subscribe



Send Feedback

## About Synplify Support

This manual delineates the support for the Synopsys Synplify software in the Quartus<sup>®</sup> II software, as well as key design flows, methodologies, and techniques for achieving optimal results in Altera<sup>®</sup> devices. The content in this manual applies to the Synplify, Synplify Pro, and Synplify Premier software unless otherwise specified. This manual assumes that you have set up, licensed, and are familiar with the Synplify software.

This manual includes the following information:

- General design flow with the Synplify and Quartus II software
- Exporting designs and constraints to the Quartus II software using NativeLink integration
- Synplify software optimization strategies, including timing-driven compilation settings, optimization options, and Altera-specific attributes
- Guidelines for Altera IP cores and library of parameterized module (LPM) functions, instantiating them with the IP Catalog, and tips for inferring them from hardware description language (HDL) code
- Incremental compilation and block-based design, including the MultiPoint flow in the Synplify Pro and Synplify Premier software

### Related Information

- [Synplify Synthesis Techniques with the Quartus II Software online training](#)
- [Synplify Pro Tips and Tricks online training](#)

## Design Flow

The following steps describe a basic Quartus II software design flow using the Synplify software:

1. Create Verilog HDL or VHDL design files.
2. Set up a project in the Synplify software and add the HDL design files for synthesis.
3. Select a target device and add timing constraints and compiler directives in the Synplify software to help optimize the design during synthesis.
4. Synthesize the project in the Synplify software.
5. Create a Quartus II project and import the following files generated by the Synplify software into the Quartus II software. Use the following files for placement and routing, and for performance evaluation:

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered

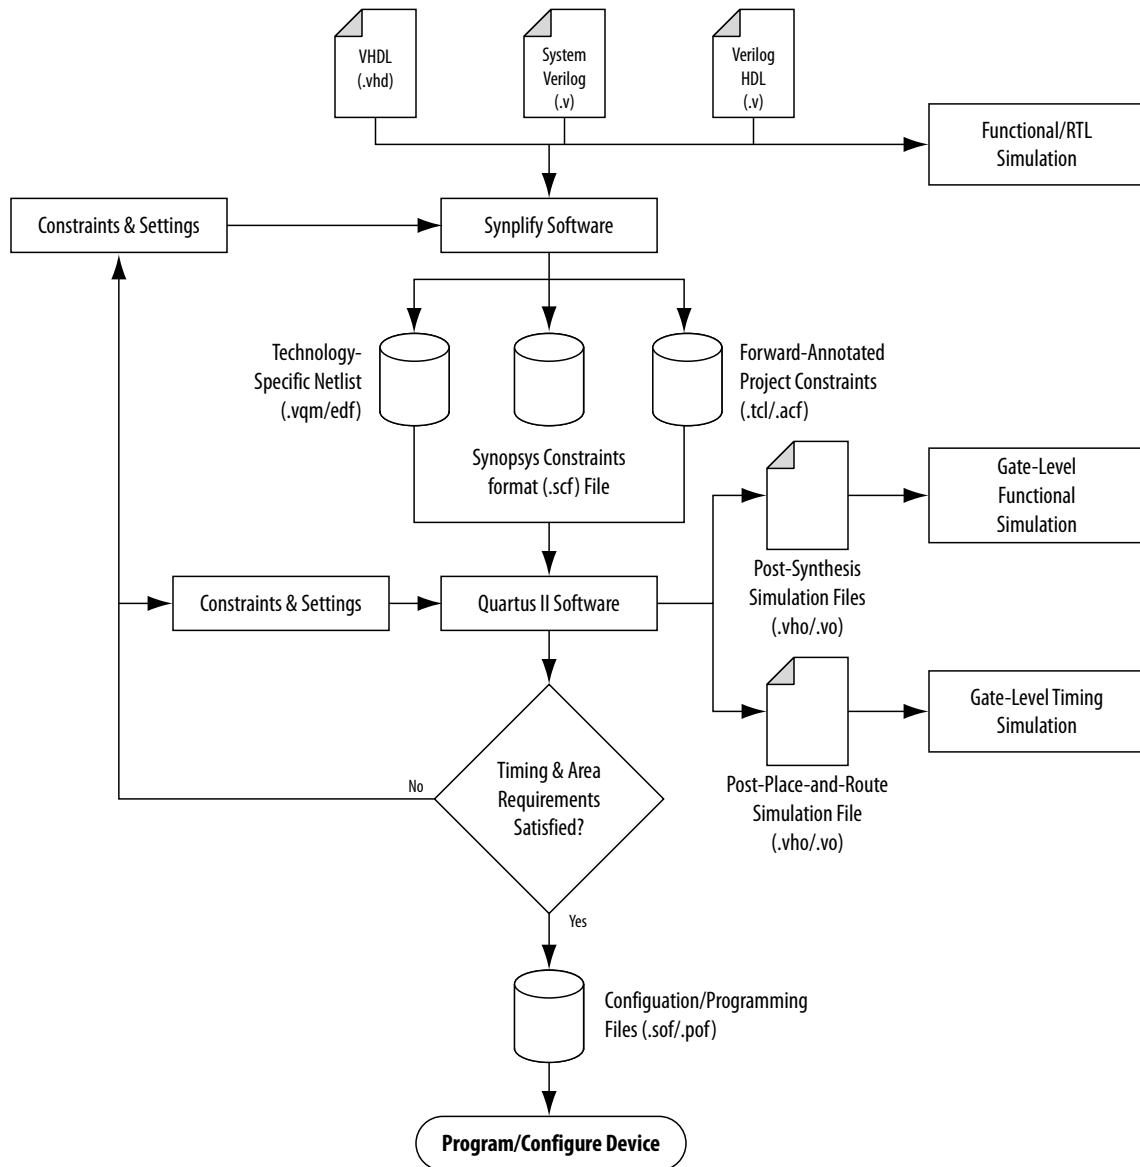


- Verilog Quartus Mapping File (**.vqm**) netlist.
- The Synopsys Constraints Format (**.scf**) file for TimeQuest Timing Analyzer constraints.
- The **.tcl** file to set up your Quartus II project and pass constraints.

**Note:** Alternatively, you can run the Quartus II software from within the Synplify software.

6. After obtaining place-and-route results that meet your requirements, configure or program the Altera device.

Figure 17-1: Recommended Design Flow



#### Related Information

- [Running the Quartus II Software from within the Synplify Software](#) on page 17-4
- [Synplify Software Generated Files](#) on page 17-5
- [Design Constraints Support](#) on page 17-6

## Hardware Description Language Support

The Synplify software supports VHDL, Verilog HDL, and SystemVerilog source files. However, only the Synplify Pro and Premier software support mixed synthesis, allowing a combination of VHDL and Verilog HDL or SystemVerilog format source files.

The HDL Analyst that is included in the Synplify software is a graphical tool for generating schematic views of the technology-independent RTL view netlist (.srs) and technology-view netlist (.srm) files. You can use the Synplify HDL Analyst to analyze and debug your design visually. The HDL Analyst supports cross-probing between the RTL and Technology views, the HDL source code, the Finite State Machine (FSM) viewer, and between the technology view and the timing report file in the Quartus II software. A separate license file is required to enable the HDL Analyst in the Synplify software. The Synplify Pro and Premier software include the HDL Analyst.

### Related Information

[Guidelines for Altera IP Cores and Architecture-Specific Features](#) on page 17-15

## Altera Device Family Support

Support for newly released device families may require an overlay. Contact Synopsys for more information.

### Related Information

[Synopsys Website](#)

## Tool Setup

### Specifying the Quartus II Software Version

You can specify your version of the Quartus II software in **Implementation Options** in the Synplify software. This option ensures that the netlist is compatible with the software version and supports the newest features. Altera recommends using the latest version of the Quartus II software whenever possible. If your Quartus II software version is newer than the versions available in the **Quartus Version** list, check if there is a newer version of the Synplify software available that supports the current Quartus II software version. Otherwise, select the latest version in the list for the best compatibility.

**Note:** The **Quartus Version** list is available only after selecting an Altera device.

#### Example 17-1: Specifying Quartus II Software Version at the Command Line

```
set_option -quartus_version <version number>
```

## Exporting Designs to the Quartus II Software Using NativeLink Integration

The NativeLink feature in the Quartus II software facilitates the seamless transfer of information between the Quartus II software and EDA tools, and allows you to run other EDA design entry or synthesis, simulation, and timing analysis tools automatically from within the Quartus II software. After a design is

synthesized in the Synplify software, a **.vqm** netlist file, an **.scf** file for TimeQuest Timing Analyzer timing constraints, and **.tcl** files are used to import the design into the Quartus II software for place-and-route. You can run the Quartus II software from within the Synplify software or as a stand-alone application. After you import the design into the Quartus II software, you can specify different options to further optimize the design.

**Note:** When you are using NativeLink integration, the path to your project must not contain empty spaces. The Synplify software uses Tcl scripts to communicate with the Quartus II software, and the Tcl language does not accept arguments with empty spaces in the path.

Use NativeLink integration to integrate the Synplify software and Quartus II software with a single GUI for both synthesis and place-and-route operations. NativeLink integration allows you to run the Quartus II software from within the Synplify software GUI, or to run the Synplify software from within the Quartus II software GUI.

## Running the Quartus II Software from within the Synplify Software

To run the Quartus II software from within the Synplify software, you must set the `QUARTUS_ROOTDIR` environment variable to the Quartus II software installation directory located in `<Quartus II system directory>\altera\ <version number>\quartus`. You must set this environment variable to use the Synplify and Quartus II software together. Synplify also uses this variable to open the Quartus II software in the background and obtain detailed information about the Altera IP cores used in the design.

For the Windows operating system, do the following:

1. Point to **Start**, and click **Control Panel**.
2. Click **System** > **Advanced system settings** > **Environment Variables**.
3. Create a `QUARTUS_ROOTDIR` system variable.

For the Linux operating system, do the following:

- Create an environment variable `QUARTUS_ROOTDIR` that points to the `<home directory>/altera <version number>` location.

You can create new place and route implementations with the **New P&R** button in the Synplify software GUI. Under each implementation, the Synplify Pro software creates a place-and-route implementation called `pr_<number> Altera Place and Route`. To run the Quartus II software in command-line mode after each synthesis run, use the text box to turn on the place-and-route implementation. The results of the place-and-route are written to a log file in the `pr_ <number>` directory under the current implementation directory.

You can also use the commands in the Quartus II menu to run the Quartus II software at any time following a successful completion of synthesis. In the Synplify software, on the Options menu, click **Quartus II** and then choose one of the following commands:

- **Launch Quartus** —Opens the Quartus II software GUI and creates a Quartus II project with the synthesized output file, forward-annotated timing constraints, and pin assignments. Use this command to configure options for the project and to execute any Quartus II commands.
- **Run Background Compile**—Runs the Quartus II software in command-line mode with the project settings from the synthesis run. The results of the place-and-route are written to a log file.

The `<project_name>_cons.tcl` file is used to set up the Quartus II project and directs the `<project_name>.tcl` file to pass constraints from the Synplify software to the Quartus II software. By default, the `<project_name>.tcl` file contains device, timing, and location assignments. The

`<project_name>.tcl` file contains the command to use the Synplify-generated `.scf` constraints file with the TimeQuest Timing Analyzer.

**Related Information**

[Design Flow](#) on page 17-1

## Using the Quartus II Software to Run the Synplify Software

You can set up the Quartus II software to run the Synplify software for synthesis with NativeLink integration. This feature allows you to use the Synplify software to quickly synthesize a design as part of a standard compilation in the Quartus II software. When you use this feature, the Synplify software does not use any timing constraints or assignments, such as incremental compilation partitions, that you have set in the Quartus II software.

**Note:** For best results, Synopsys recommends that you set constraints in the Synplify software and use a Tcl script to pass these constraints to the Quartus II software, instead of opening the Synplify software from within the Quartus II software.

To set up the Quartus II software to run the Synplify software, do the following:

1. On the Tools menu, click **Options**.
2. In the **Options** dialog box, click **EDA Tool Options** and specify the path of the Synplify or Synplify Pro software under **Location of Executable**.

Running the Synplify software with NativeLink integration is supported on both floating network and node-locked fixed PC licenses. Both types of licenses support batch mode compilation.

**Related Information**

[About Using the Synplify Software with the Quartus II Software Online Help](#)

## Synplify Software Generated Files

During synthesis, the Synplify software produces several intermediate and output files.

**Table 17-1: Synplify Intermediate and Output Files**

File Extensions	File Description
<code>.vqm</code>	Technology-specific netlist in <code>.vqm</code> file format. A <code>.vqm</code> file is created for all Altera device families supported by the Quartus II software.
<code>.scf</code> <sup>(13)</sup>	Synopsys Constraint Format file containing timing constraints for the TimeQuest Timing Analyzer.

<sup>(13)</sup> If your design uses the Classic Timing Analyzer for timing analysis in the Quartus II software versions 10.0 and earlier, the Synplify software generates timing constraints in the Tcl Constraints File (`.tcl`). If you are using the Quartus II software versions 10.1 and later, you must use the TimeQuest Timing Analyzer for timing analysis.

File Extensions	File Description
<b>.tcl</b>	Forward-annotated constraints file containing constraints and assignments. A <b>.tcl</b> file for the Quartus II software is created for all devices. The <b>.tcl</b> file contains the appropriate Tcl commands to create and set up a Quartus II project and pass placement constraints.
<b>.srs</b>	Technology-independent RTL netlist file that can be read only by the Synplify software.
<b>.srm</b>	Technology view netlist file.
<b>.acf</b>	Assignment and Configurations file for backward compatibility with the MAX+PLUS II software. For devices supported by the MAX+PLUS II software, the MAX+PLUS II assignments are imported from the MAX+PLUS II <b>.acf</b> file.
<b>.srr<sup>(14)</sup></b>	Synthesis Report file.

**Related Information**

[Design Flow](#) on page 17-1

## Design Constraints Support

You can specify timing constraints and attributes by using the SCOPE window of the Synplify software, by editing the **.sdc** file, or by defining the compiler directives in the HDL source file. The Synplify software forward-annotates many of these constraints to the Quartus II software.

After synthesis is complete, do the following steps:

1. Import the **.vqm** netlist to the Quartus II software for place-and-route.
2. Use the **.tcl** file generated by the Synplify software to forward-annotate your project constraints including device selection. The **.tcl** file calls the generated **.scf** to forward-annotate TimeQuest Timing Analyzer timing constraints.

**Related Information**

- [Design Flow](#) on page 17-1
- [Synplify Optimization Strategies](#) on page 17-8
- [Netlist Optimizations and Physical Synthesis Documentation](#)

<sup>(14)</sup> This report file includes performance estimates that are often based on pre-place-and-route information. Use the  $f_{MAX}$  reported by the Quartus II software after place-and-route—it is the only reliable source of timing information. This report file includes post-synthesis device resource utilization statistics that might inaccurately predict resource usage after place-and-route. The Synplify software does not account for black box functions nor for logic usage reduction achieved through register packing performed by the Quartus II software. Register packing combines a single register and look-up table (LUT) into a single logic cell, reducing logic cell utilization below the Synplify software estimate. Use the device utilization reported by the Quartus II software after place-and-route.

## Running the Quartus II Software Manually With the Synplify-Generated Tcl Script

You can run the Quartus II software with a Synplify-generated Tcl script.

To run the Tcl script to set up your project assignments, perform the following steps:

1. Ensure the **.vqm**, **.scf**, and **.tcl** files are located in the same directory.
2. In the Quartus II software, on the View menu, point to **Utility Windows** and click **Tcl Console**. The Quartus II Tcl Console opens.
3. At the Tcl Console command prompt, type the following:

```
source <path>/<project name>_cons.tcl
```

## Passing TimeQuest SDC Timing Constraints to the Quartus II Software

The TimeQuest Timing Analyzer is a powerful ASIC-style timing analysis tool that validates the timing performance of all logic in your design using an industry standard constraints format, Synopsys Design Constraints (SDC).

The Synplify-generated **.tcl** file contains constraints for the Quartus II software, such as the device specification and any location constraints. Timing constraints are forward-annotated in the Synopsys Constraints Format (**.scf**) file.

**Note:** Synopsys recommends that you modify constraints using the SCOPE constraint editor window, rather than using the generated **.sdc**, **.scf**, or **.tcl** file.

The following list of Synplify constraints are converted to the equivalent Quartus II SDC commands and are forward-annotated to the Quartus II software in the **.scf** file:

- `define_clock`
- `define_input_delay`
- `define_output_delay`
- `define_multicycle_path`
- `define_false_path`

All Synplify constraints described above are mapped to SDC commands for the TimeQuest Timing Analyzer.

For syntax and arguments for these commands, refer to the applicable topic in this manual or refer to Synplify Help. For a list of corresponding commands in the Quartus II software, refer to the Quartus II Help.

### Related Information

- [Timing-Driven Synthesis Settings](#) on page 17-9
- [Quartus II TimeQuest Timing Analyzer Documentation](#)
- [Quartus II Online Help](#)

## Individual Clocks and Frequencies

Specify clock frequencies for individual clocks in the Synplify software with the `define_clock` command. This command is passed to the Quartus II software with the `create_clock` command.

## Input and Output Delay

Specify input delay and output delay constraints in the Synplify software with the `define_input_delay` and `define_output_delay` commands, respectively. These commands are passed to the Quartus II software with the `set_input_delay` and `set_output_delay` commands.

## Multicycle Path

Specify a multicycle path constraint in the Synplify software with the `define_multicycle_path` command. This command is passed to the Quartus II software with the `set_multicycle_path` command.

## False Path

Specify a false path constraint in the Synplify software with the `define_false_path` command. This command is passed to the Quartus II software with the `set_false_path` command.

## Simulation and Formal Verification

You can perform simulation and formal verification at various stages in the design process. You can perform final timing analysis after placement and routing is complete.

If area and timing requirements are satisfied, use the files generated by the Quartus II software to program or configure the Altera device. If your area or timing requirements are not met, you can change the constraints in the Synplify software or the Quartus II software and rerun synthesis. Altera recommends that you provide timing constraints in the Synplify software and any placement constraints in the Quartus II software. Repeat the process until area and timing requirements are met.

You can also use other options and techniques in the Quartus II software to meet area and timing requirements, such as WYSIWYG Primitive Resynthesis, which can perform optimizations on your `.vqm` netlist within the Quartus II software.

**Note:** In some cases, you might be required to modify the source code if the area and timing requirements cannot be met using options in the Synplify and Quartus II software.

## Synplify Optimization Strategies

Combining Synplify software constraints with VHDL and Verilog HDL coding techniques and Quartus II software options can help you obtain the results that you require.

For more information about applying attributes, refer to the *Synopsys FPGA Synthesis Reference Manual*.

### Related Information

- [Design Constraints Support](#) on page 17-6
- [Recommended Design Practices Documentation](#) on page 11-1
- [Timing Closure and Optimization Documentation](#)

## Using Synplify Premier to Optimize Your Design

Compared to other Synplify products, the Synplify Premier software offers additional physical synthesis optimizations. After typical logic synthesis, the Synplify Premier software places and routes the design and attempts to restructure the netlist based on the physical location of the logic in the Altera device. The



Synplify Premier software forward-annotates the design netlist to the Quartus II software to perform the final placement and routing. In the default flow, the Synplify Premier software also forward-annotates placement information for the critical path(s) in the design, which can improve the compilation time in the Quartus II software.

The physical location annotation file is called `<design name>_plc.tcl`. If you open the Quartus II software from the Synplify Premier software user interface, the Quartus II software automatically uses this file for the placement information.

The Physical Analyst allows you to examine the placed netlist from the Synplify Premier software, which is similar to the HDL Analyst for a logical netlist. You can use this display to analyze and diagnose potential problems.

## Using Implementations in Synplify Pro or Premier

You can create different synthesis results without overwriting the existing results, in the Synplify Pro or Premier software, by creating a new implementation from the Project menu. For each implementation, specify the target device, synthesis options, and constraint files. Each implementation generates its own subdirectory that contains all the resulting files, including `.vqm`, `.scf`, and `.tcl` files, from a compilation of the particular implementation. You can then compare the results of the different implementations to find the optimal set of synthesis options and constraints for a design.

## Timing-Driven Synthesis Settings

The Synplify software supports timing-driven synthesis with user-assigned timing constraints to optimize the performance of the design.

The Quartus II NativeLink feature allows timing constraints that are applied in the Synplify software to be forward-annotated for the Quartus II software with an `.scf` file for timing-driven place and route.

The Synplify Synthesis Report File (`.srr`) contains timing reports of estimated place-and-route delays. The Quartus II software can perform further optimizations on a post-synthesis netlist from third-party synthesis tools. In addition, designs might contain black boxes or intellectual property (IP) functions that have not been optimized by the third-party synthesis software. Actual timing results are obtained only after the design has been fully placed and routed in the Quartus II software. For these reasons, the Quartus II post place-and-route timing reports provide a more accurate representation of the design. Use the statistics in these reports to evaluate design performance.

### Related Information

- [Passing TimeQuest SDC Timing Constraints to the Quartus II Software](#) on page 17-7
- [Exporting Designs to the Quartus II Software Using NativeLink Integration](#) on page 17-3

## Clock Frequencies

For single-clock designs, you can specify a global frequency when using the push-button flow. While this flow is simple and provides good results, it often does not meet the performance requirements for more advanced designs. You can use timing constraints, compiler directives, and other attributes to help optimize the performance of a design. You can enter these attributes and directives directly in the HDL code. Alternatively, you can enter attributes (not directives) into an `.sdc` file with the SCOPE window in the Synplify software.

Use the SCOPE window to set global frequency requirements for the entire design and individual clock settings. Use the **Clocks** tab in the SCOPE window to specify frequency (or period), rise times, fall times,

duty cycle, and other settings. Assigning individual clock settings, rather than over-constraining the global frequency, helps the Quartus II software and the Synplify software achieve the fastest clock frequency for the overall design. The `define_clock` attribute assigns clock constraints.

## Multiple Clock Domains

The Synplify software can perform timing analysis on unrelated clock domains. Each clock group is a different clock domain and is treated as unrelated to the clocks in all other clock groups. All clocks in a single clock group are assumed to be related, and the Synplify software automatically calculates the relationship between the clocks. You can assign clocks to a new clock group or put related clocks in the same clock group with the **Clocks** tab in the SCOPE window, or with the `define_clock` attribute.

## Input and Output Delays

Specify the input and output delays for the ports of a design in the **Input/Output** tab of the SCOPE window, or with the `define_input_delay` and `define_output_delay` attributes. The Synplify software does not allow you to assign the  $t_{CO}$  and  $t_{SU}$  values directly to inputs and outputs. However, a  $t_{CO}$  value can be inferred by setting an external output delay; a  $t_{SU}$  value can be inferred by setting an external input delay.

### Relationship Between $t_{CO}$ and the Output Delay

$$t_{CO} = \text{clock period} - \text{external output delay}$$

### Relationship Between $t_{SU}$ and the Input Delay

$$t_{SU} = \text{clock period} - \text{external input delay}$$

When the `syn_forward_io_constraints` attribute is set to 1, the Synplify software passes the external input and output delays to the Quartus II software using NativeLink integration. The Quartus II software then uses the external delays to calculate the maximum system frequency.

## Multicycle Paths

A multicycle path is a path that requires more than one clock cycle to propagate. Specify any multicycle paths in the design in the **Multi-Cycle Paths** tab of the SCOPE window, or with the `define_multicycle_path` attribute. You should specify which paths are multicycle to prevent the Quartus II and the Synplify compilers from working excessively on a non-critical path. Not specifying these paths can also result in an inaccurate critical path reported during timing analysis.

## False Paths

False paths are paths that should be ignored during timing analysis, or should be assigned low (or no) priority during optimization. Some examples of false paths include slow asynchronous resets, and test logic that has been added to the design. Set these paths in the **False Paths** tab of the SCOPE window, or use the `define_false_path` attribute.

## FSM Compiler

If the FSM Compiler is turned on, the compiler automatically detects state machines in a design, which are then extracted and optimized. The FSM Compiler analyzes state machines and implements sequential, gray, or one-hot encoding, based on the number of states. The compiler also performs unused-state analysis, optimization of unreachable states, and minimization of transition logic. Implementation is based on the number of states, regardless of the coding style in the HDL code.

If the FSM Compiler is turned off, the compiler does not optimize logic as state machines. The state machines are implemented as HDL code. Thus, if the coding style for a state machine is sequential, the implementation is also sequential.

Use the `syn_state_machine` compiler directive to specify or prevent a state machine from being extracted and optimized. To override the default encoding of the FSM Compiler, use the `syn_encoding` directive.

**Table 17-2: `syn_encoding` Directive Values**

Value	Description
Sequential	Generates state machines with the fewest possible flipflops. Sequential, also called binary, state machines are useful for area-critical designs when timing is not the primary concern.
Gray	Generates state machines where only one flipflop changes during each transition. Gray-encoded state machines tend to be glitches.
One-hot	Generates state machines containing one flipflop for each state. One-hot state machines typically provide the best performance and shortest clock-to-output delays. However, one-hot implementations are usually larger than sequential implementations.
Safe	Generates extra control logic to force the state machine to the reset state if an invalid state is reached. You can use the safe value in conjunction with any of the other three values, which results in the state machine being implemented with the requested encoding scheme and the generation of the reset logic.

### Example 17-2: Sample VHDL Code for Applying `syn_encoding` Directive

```
SIGNAL current_state : STD_LOGIC_VECTOR (7 DOWNTO 0);  
ATTRIBUTE syn_encoding : STRING;  
ATTRIBUTE syn_encoding OF current_state : SIGNAL IS "sequential";
```

By default, the state machine logic is optimized for speed and area, which may be potentially undesirable for critical systems. The safe value generates extra control logic to force the state machine to the reset state if an invalid state is reached.

## FSM Explorer in Synplify Pro and Premier

The Synplify Pro and Premier software use the FSM Explorer to explore different encoding styles for a state machine automatically, and then implement the best encoding based on the overall design constraints. The FSM Explorer uses the FSM Compiler to identify and extract state machines from a design. However, unlike the FSM Compiler, which chooses the encoding style based on the number of states, the FSM Explorer attempts several different encoding styles before choosing a specific one. The trade-off is that the compilation requires more time to analyze the state machine, but finds an optimal encoding scheme for the state machine.

## Optimization Attributes and Options

### Retiming in Synplify Pro and Premier

The Synplify Pro and Premier software can retime a design, which can improve the timing performance of sequential circuits by moving registers (register balancing) across combinational elements. Be aware that retimed registers incur name changes. You can retime your design from **Implementation Options** or you can use the `syn_allow_retiming` attribute.

### Maximum Fan-Out

When your design has critical path nets with high fan-out, use the `syn_maxfan` attribute to control the fan-out of the net. Setting this attribute for a specific net results in the replication of the driver of the net to reduce overall fan-out. The `syn_maxfan` attribute takes an integer value and applies it to inputs or registers. The `syn_maxfan` attribute cannot be used to duplicate control signals. The minimum allowed value of the attribute is 4. Using this attribute might result in increased logic resource utilization, thus straining routing resources, which can lead to long compilation times and difficult fitting.

If you must duplicate an output register or an output enable register, you can create a register for each output pin by using the `syn_useioff` attribute.

### Preserving Nets

During synthesis, the compiler maintains ports, registers, and instantiated components. However, some nets cannot be maintained to create an optimized circuit. Applying the `syn_keep` directive overrides the optimization of the compiler and preserves the net during synthesis. The `syn_keep` directive is a Boolean data type value and can be applied to wires (Verilog HDL) and signals (VHDL). Setting the value to **true** preserves the net through synthesis.

### Register Packing

Altera devices allow register packing into I/O cells. Altera recommends allowing the Quartus II software to make the I/O register assignments. However, you can control register packing with the `syn_useioff` attribute. The `syn_useioff` attribute is a Boolean data type value that can be applied to ports or entire modules. Setting the value to **1** instructs the compiler to pack the register into an I/O cell. Setting the value to **0** prevents register packing in both the Synplify and Quartus II software.

### Resource Sharing

The Synplify software uses resource sharing techniques during synthesis, by default, to reduce area. Turning off the **Resource Sharing** option on the **Options** tab of the **Implementation Options** dialog box improves performance results for some designs. You can also turn off the option for a specific module with the `syn_sharing` attribute. If you turn off this option, be sure to check the results to verify improvement in timing performance. If there is no improvement, turn on **Resource Sharing**.

### Preserving Hierarchy

The Synplify software performs cross-boundary optimization by default, which causes the design to flatten to allow optimization. You can use the `syn_hier` attribute to override the default compiler settings. The `syn_hier` attribute applies a string value to modules, architectures, or both. Setting the value to **hard** maintains the boundaries of a module, architecture, or both, but allows constant propagation. Setting the value to **locked** prevents all cross-boundary optimizations. Use the **locked** setting with the partition setting to create separate design blocks and multiple output netlists for incremental compilation.

By default, the Synplify software generates a hierarchical `.vqm` file. To flatten the file, set the `syn_netlist_hierarchy` attribute to `0`.

#### Related Information

[Using MultiPoint Synthesis with Incremental Compilation](#) on page 17-26

## Register Input and Output Delays

Two advanced options, `define_reg_input_delay` and `define_reg_output_delay`, can speed up paths feeding a register, or coming from a register, by a specific number of nanoseconds. The Synplify software attempts to meet the global clock frequency goals for a design as well as the individual clock frequency goals (set with the `define_clock` attribute). You can use these attributes to add a delay to paths feeding into or out of registers to further constrain critical paths. You can slow down a path that is too highly optimized by setting this attributes to a negative number.

The `define_reg_input_delay` and `define_reg_output_delay` options are useful to close timing if your design does not meet timing goals, because the routing delay after placement and routing exceeds the delay predicted by the Synplify software. Rerun synthesis using these options, specifying the actual routing delay (from place-and-route results) so that the tool can meet the required clock frequency. Synopsys recommends that for best results, do not make these assignments too aggressively. For example, you can increase the routing delay value, but do not also use the full routing delay from the last compilation.

In the SCOPE constraint window, the registers panel contains the following options:

- **Register**—Specifies the name of the register. If you have initialized a compiled design, select the name from the list.
- **Type**—Specifies whether the delay is an input or output delay.
- **Route**—Shrinks the effective period for the constrained registers by the specified value without affecting the clock period that is forward-annotated to the Quartus II software.

Use the following Tcl command syntax to specify an input or output register delay in nanoseconds.

### Example 17-3: Input and Output Register Delay

```
define_reg_input_delay {<register>} -route <delay in ns>  
define_reg_output_delay {<register>} -route <delay in ns>
```

## syn\_direct\_enable

This attribute controls the assignment of a clock-enable net to the dedicated enable pin of a register. With this attribute, you can direct the Synplify mapper to use a particular net as the only clock enable when the design has multiple clock enable candidates.

To use this attribute as a compiler directive to infer registers with clock enables, enter the `syn_direct_enable` directive in your source code, instead of the SCOPE spreadsheet.

The `syn_direct_enable` data type is Boolean. A value of `1` or `true` enables net assignment to the clock-enable pin. The following is the syntax for Verilog HDL:

```
object /* synthesis syn_direct_enable = 1 */ ;
```

## I/O Standard

For certain Altera devices, specify the I/O standard type for an I/O pad in the design with the **I/O Standard** panel in the Synplify SCOPE window.

The Synplify SDC syntax for the `define_io_standard` constraint, in which the `delay_type` must be either `input_delay` or `output_delay`.

### Example 17-4: define\_io\_standard Constraint

```
define_io_standard [-disable|-enable] {<objectName>} -delay_type \  
[input_delay|output_delay] <columnTclName>{<value>} \  
[<columnTclName>{<value>}...]
```

For details about supported I/O standards, refer to the *Synopsys FPGA Synthesis Reference Manual*.

## Altera-Specific Attributes

You can use the `altera_chip_pin_lc`, `altera_io_powerup`, and `altera_io_opendrain` attributes with specific Altera device features, which are forward-annotated to the Quartus II project, and are used during place-and-route.

### altera\_chip\_pin\_lc

Use the `altera_chip_pin_lc` attribute to make pin assignments. This attribute applies a string value to inputs and outputs. Use the attribute only on the ports of the top-level entity in the design. Do not use this attribute to assign pin locations from entities at lower levels of the design hierarchy.

**Note:** The `altera_chip_pin_lc` attribute is not supported for any MAX series device.

In the SCOPE window, set the value of the `altera_chip_pin_lc` attribute to a pin number or a list of pin numbers.

You can use VHDL code for making location assignments for supported Altera devices. Pin location assignments for these devices are written to the output `.tcl` file.

**Note:** The `data_out` signal is a 4-bit signal; `data_out[3]` is assigned to pin 14 and `data_out[0]` is assigned to pin 15.

### Example 17-5: Making Location Assignments in VHDL

```
ENTITY sample (data_in : IN STD_LOGIC_VECTOR (3 DOWNTO 0);  
              data_out: OUT STD_LOGIC_VECTOR (3 DOWNTO 0));  
  ATTRIBUTE altera_chip_pin_lc : STRING;  
  ATTRIBUTE altera_chip_pin_lc OF data_out : SIGNAL IS "14, 5, 16, 15";
```

### altera\_io\_powerup

Use the `altera_io_powerup` attribute to define the power-up value of an I/O register that has no set or reset. This attribute applies a string value (**high**/**low**) to ports with I/O registers. By default, the power-up value of the I/O register is set to **low**.

## altera\_io\_opendrain

Use the `altera_io_opendrain` attribute to specify open-drain mode I/O ports. This attribute applies a boolean data type value to outputs or bidirectional ports for devices that support open-drain mode.

## Guidelines for Altera IP Cores and Architecture-Specific Features

Altera provides parameterizable IP cores, including LPMs, device-specific Altera IP cores, and IP available through the Altera Megafunction Partners Program (AMPP<sup>SM</sup>). You can use IP cores by instantiating them in your HDL code, or by inferring certain IP cores from generic HDL code.

You can instantiate an IP core in your HDL code with the IP Catalog and configure the IP core with the Parameter Editor, or instantiate the IP core using the port and parameter definition. The IP Catalog and Parameter Editor provide a graphical interface within the Quartus II software to customize any available Altera IP core for the design.

The Synplify software also automatically recognizes certain types of HDL code, and infers the appropriate Altera IP core when an IP core provides optimal results. The Synplify software provides options to control inference of certain types of IP cores.

### Related Information

- [Hardware Description Language Support](#) on page 17-3
- [Recommended HDL Coding Styles Documentation](#) on page 12-1
- [About the IP Catalog Online Help](#)

## Instantiating Altera IP Cores with the IP Catalog

When you use the IP Catalog and Parameter Editor to set up and configure an IP core, the IP Catalog creates a VHDL or Verilog HDL wrapper file `<output file>.v|vhd` that instantiates the IP core.

The Synplify software uses the Quartus II timing and resource estimation netlist feature to report more accurate resource utilization and timing performance estimates, and leverages timing-driven optimization, instead of treating the IP core as a “black box.” Including the generated IP core variation wrapper file in your Synplify project, gives the Synplify software complete information about the IP core.

**Note:** There is an option in the Parameter Editor to generate a netlist for resource and timing estimation. This option is not recommended for the Synplify software because the software automatically generates this information in the background without a separate netlist. If you do create a separate netlist `<output file>_syn.v` and use that file in your synthesis project, you must also include the `<output file>.v|vhd` file in your Quartus II project.

Verify that the correct Quartus II version is specified in the Synplify software before compiling the generated file to ensure that the software uses the correct library definitions for the IP core. The **Quartus Version** setting must match the version of the Quartus II software used to generate the customized IP core.

In addition, ensure that the `QUARTUS_ROOTDIR` environment variable specifies the installation directory location of the correct Quartus II version. The Synplify software uses this information to launch the Quartus II software in the background. The environment variable setting must match the version of the Quartus II software used to generate the customized IP core.

### Related Information

- [Specifying the Quartus II Software Version](#) on page 17-3
- [Using the Quartus II Software to Run the Synplify Software](#) on page 17-5

## Instantiating Altera IP Cores with IP Catalog Generated Verilog HDL Files

If you turn on the `<output file>_inst.v` option on the Parameter Editor, the IP Catalog generates a Verilog HDL instantiation template file for use in your Synplify design. The instantiation template file, `<output file>_inst.v`, helps to instantiate the IP core variation wrapper file, `<output file>.v`, in your top-level design. Include the IP core variation wrapper file `<output file>.v` in your Synplify project. The Synplify software includes the IP core information in the output `.vqm` netlist file. You do not need to include the generated IP core variation wrapper file in your Quartus II project.

## Instantiating Altera IP Cores with IP Catalog Generated VHDL Files

If you turn on the `<output file>.cmp` and `<output file>_inst.vhd` options on the Parameter Editor, the IP catalog generates a VHDL component declaration file and a VHDL instantiation template file for use in your Synplify design. These files can help you instantiate the IP core variation wrapper file, `<output file>.vhd`, in your top-level design. Include the `<output file>.vhd` in your Synplify project. The Synplify software includes the IP core information in the output `.vqm` netlist file. You do not need to include the generated IP core variation wrapper file in your Quartus II project.

## Changing Synplify's Default Behavior for Instantiated Altera IP Cores

By default, the Synplify software automatically opens the Quartus II software in the background to generate a resource and timing estimation netlist for IP cores.

You might want to change this behavior to reduce run times in the Synplify software, because generating the netlist files can take several minutes for large designs, or if the Synplify software cannot access your Quartus II software installation to generate the files. Changing this behavior might speed up the compilation time in the Synplify software, but the Quality of Results (QoR) might be reduced.

The Synplify software directs the Quartus II software to generate information in two ways:

- Some IP cores provide a “clear box” model—the Synplify software fully synthesizes this model and includes the device architecture-specific primitives in the output `.vqm` netlist file.
- Other IP cores provide a “grey box” model—the Synplify software reads the resource information, but the netlist does not contain all the logic functionality.

**Note:** You need to turn on **Generate netlist** when using the grey box model. For more information, see the Quartus II online help.

For these IP cores, the Synplify software uses the logic information for resource and timing estimation and optimization, and then instantiates the IP core in the output `.vqm` netlist file so the Quartus II software can implement the appropriate device primitives. By default, the Synplify software uses the clear box model when available, and otherwise uses the grey box model.

### Related Information

- [Including Files for Quartus II Placement and Routing Only](#) on page 17-19
- [Synplify Synthesis Techniques with the Quartus II Software online training](#)  
Includes more information about design flows using clear box model and grey box model.
- [Generating a Netlist for 3rd Party Synthesis Tools online help](#)



## Instantiating Intellectual Property with the IP Catalog and Parameter Editor

Many Altera IP cores include a resource and timing estimation netlist that the Synplify software uses to report more accurate resource utilization and timing performance estimates, and leverage timing-driven optimization rather than a black box function.

To create this netlist file, perform the following steps:

1. Select the IP core in the IP Catalog.
2. Click **Next** to open the Parameter Editor.
3. Click **Set Up Simulation**, which sets up all the EDA options.
4. Turn on the **Generate netlist** option to generate a netlist for resource and timing estimation and click **OK**.
5. Click **Generate** to generate the netlist file.

The Quartus II software generates a file `<output file>_syn.v`. This netlist contains the grey box information for resource and timing estimation, but does not contain the actual implementation. Include this netlist file in your Synplify project. Next, include the IP core variation wrapper file `<output file>.v|vhd` in the Quartus II project along with your Synplify `.vqm` output netlist.

If your IP core does not include a resource and timing estimation netlist, the Synplify software must treat the IP core as a black box.

### Related Information

[Including Files for Quartus II Placement and Routing Only](#) on page 17-19

## Instantiating Black Box IP Cores with Generated Verilog HDL Files

Use the `syn_black_box` compiler directive to declare a module as a black box. The top-level design files must contain the IP port-mapping and a hollow-body module declaration. Apply the `syn_black_box` directive to the module declaration in the top-level file or a separate file included in the project so that the Synplify software recognizes the module is a black box. The software compiles successfully without this directive, but reports an additional warning message. Using this directive allows you to add other directives.

The example shows a top-level file that instantiates `my_verilogIP.v`, which is a simple customized variation generated by the IP Catalog.

### Example 17-6: Sample Top-Level Verilog HDL Code with Black Box Instantiation of IP

```
module top (clk, count);
    input clk;
    output [7:0] count;
    my_verilogIP verilogIP_inst (.clock (clk), .q (count));
endmodule
// Module declaration
// The following attribute is added to create a
// black box for this module.
module my_verilogIP (clock, q) /* synthesis syn_black_box */;
    input clock;
    output [7:0] q;
endmodule
```

## Instantiating Black Box IP Cores with Generated VHDL Files

Use the `syn_black_box` compiler directive to declare a component as a black box. The top-level design files must contain the IP core variation component declaration and port-mapping. Apply the

`syn_black_box` directive to the component declaration in the top-level file. The software compiles successfully without this directive, but reports an additional warning message. Using this directive allows you to add other directives.

The example shows a top-level file that instantiates `my_vhdlIP.vhd`, which is a simplified customized variation generated by the IP Catalog.

### Example 17-7: Sample Top-Level VHDL Code with Black Box Instantiation of IP

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;

ENTITY top IS
  PORT (
    clk: IN STD_LOGIC ;
    count: OUT STD_LOGIC_VECTOR (7 DOWNTO 0)
  );
END top;

ARCHITECTURE rtl OF top IS
  COMPONENT my_vhdlIP
    PORT (
      clock: IN STD_LOGIC ;
      q: OUT STD_LOGIC_VECTOR (7 DOWNTO 0)
    );
  end COMPONENT;
  attribute syn_black_box : boolean;
  attribute syn_black_box of my_vhdlIP: component is true;
BEGIN
  vhdlIP_inst : my_vhdlIP PORT MAP (
    clock => clk,
    q => count
  );
END rtl;
```

### Other Synplify Software Attributes for Creating Black Boxes

Instantiating IP as a black box does not provide visibility into the IP for the synthesis tool. Thus, it does not take full advantage of the synthesis tool's timing-driven optimization. For better timing optimization, especially if the black box does not have registered inputs and outputs, add timing models to black boxes by adding the `syn_tpd`, `syn_tsu`, and `syn_tco` attributes.

### Example 17-8: Adding Timing Models to Black Boxes in Verilog HDL

```
module ram32x4(z,d,addr,we,clk);
  /* synthesis syn_black_box syn_tco1="clk->z[3:0]=4.0"
     syn_tpd1="addr[3:0]->[3:0]=8.0"
     syn_tsu1="addr[3:0]->clk=2.0"
     syn_tsu2="we->clk=3.0" */
  output [3:0]z;
  input [3:0]d;
  input [3:0]addr;
  input we;
  input clk;
endmodule
```

The following additional attributes are supported by the Synplify software to communicate details about the characteristics of the black box module within the HDL code:

- `syn_resources`—Specifies the resources used in a particular black box.
- `black_box_pad_pin`—Prevents mapping to I/O cells.
- `black_box_tri_pin`—Indicates a tri-stated signal.

For more information about applying these attributes, refer to the *Synopsys FPGA Synthesis Reference Manual*.

## Including Files for Quartus II Placement and Routing Only

In the Synplify software, you can add files to your project that are used only during placement and routing in the Quartus II software. This can be useful if you have grey or black boxes for Synplify synthesis that require the full design files to be compiled in the Quartus II software.

You can also set the option in a script using the `-job_owner par` option.

The example shows how to define files for a Synplify project that includes a top-level design file, a grey box netlist file, an IP wrapper file, and an encrypted IP file. With these files, the Synplify software writes an empty instantiation of “core” in the `.vqm` file and uses the grey box netlist for resource and timing estimation. The files `core.v` and `core_enc8b10b.v` are not compiled by the Synplify software, but are copied into the place-and-route directory. The Quartus II software compiles these files to implement the “core” IP block.

### Example 17-9: Commands to Define Files for a Synplify Project

```
add_file -verilog -job_owner par "core_enc8b10b.v"  
add_file -verilog -job_owner par "core.v"  
add_file -verilog "core_gb.v"  
add_file -verilog "top.v"
```

## Inferring Altera IP Cores from HDL Code

The Synplify software uses Behavior Extraction Synthesis Technology (BEST) algorithms to infer high-level structures such as RAMs, ROMs, operators, FSMs, and DSP multiplication operations. Then, the Synplify software keeps the structures abstract for as long as possible in the synthesis process. This allows the use of technology-specific resources to implement these structures by inferring the appropriate Altera IP core when an IP core provides optimal results.


### Related Information

- [Recommended HDL Coding Styles Documentation](#) on page 12-1

## Inferring Multipliers

The figure shows the HDL Analyst view of an unsigned  $8 \times 8$  multiplier with two pipeline stages after synthesis in the Synplify software. This multiplier is converted into an `ALTMULT_ADD` or `ALTMULT_ACCUM` IP core. For devices with DSP blocks, the software might implement the function in a DSP block instead of regular logic, depending on device utilization. For some devices, the software maps directly to DSP block device primitives instead of instantiating an IP core in the `.vqm` file.

Figure 17-2: HDL Analyst View of LPM\_MULT IP Core (Unsigned 8x8 Multiplier with Pipeline=2)



### Resource Balancing

While mapping multipliers to DSP blocks, the Synplify software performs resource balancing for optimum performance.

Altera devices have a fixed number of DSP blocks, which includes a fixed number of embedded multipliers. If the design uses more multipliers than are available, the Synplify software automatically maps the extra multipliers to logic elements (LEs), or adaptive logic modules (ALMs).

If a design uses more multipliers than are available in the DSP blocks, the Synplify software maps the multipliers in the critical paths to DSP blocks. Next, any wide multipliers, which might or might not be in the critical paths, are mapped to DSP blocks. Smaller multipliers and multipliers that are not in the critical paths might then be implemented in the logic (LEs or ALMs). This ensures that the design fits successfully in the device.

### Controlling the DSP Block Inference

You can implement multipliers in DSP blocks or in logic in Altera devices that contain DSP blocks. You can control this implementation through attribute settings in the Synplify software.

### Signal Level Attribute

You can control the implementation of individual multipliers by using the `syn_multstyle` attribute as shown in the following Verilog HDL code (where `<signal_name>` is the name of the signal):

```
<signal_name> /* synthesis syn_multstyle = "logic" */;
```

The `syn_multstyle` attribute applies to wires only; it cannot be applied to registers.

**Table 17-3: DSP Block Attribute Setting in the Synplify Software**

Attribute Name	Value	Description
syn_multstyle	lpm_mult	LPM function inferred and multipliers implemented in DSP blocks.
	logic	LPM function not inferred and multipliers implemented as LEs by the Synplify software.
	block_mult	DSP IP core is inferred and multipliers are mapped directly to DSP block device primitives (for supported devices).

**Example 17-10: Signal Attributes for Controlling DSP Block Inference in Verilog HDL Code**

```

module mult(a,b,c,r,en);
  input [7:0] a,b;
  output [15:0] r;
  input [15:0] c;
  input en;
  wire [15:0] temp /* synthesis syn_multstyle="logic" */;

  assign temp = a*b;
  assign r = en ? temp : c;
endmodule

```

**Example 17-11: Signal Attributes for Controlling DSP Block Inference in VHDL Code**

```

library ieee;
use ieee.std_logic_1164.all;
use ieee.std_logic_unsigned.all;

entity onereg is port (
  r : out std_logic_vector (15 downto 0);
  en : in std_logic;
  a : in std_logic_vector (7 downto 0);
  b : in std_logic_vector (7 downto 0);
  c : in std_logic_vector (15 downto 0);
);
end onereg;

architecture beh of onereg is
  signal temp : std_logic_vector (15 downto 0);
  attribute syn_multstyle : string;
  attribute syn_multstyle of temp : signal is "logic";

begin
  temp <= a * b;
  r <= temp when en='1' else c;
end beh;

```

## Inferring RAM

When a RAM block is inferred from an HDL design, the Synplify software uses an Altera IP core to target the device memory architecture. For some devices, the Synplify software maps directly to memory block device primitives instead of instantiating an IP core in the `.vqm` file.

Follow these guidelines for the Synplify software to successfully infer RAM in a design:

- The address line must be at least two bits wide.
- Resets on the memory are not supported. Refer to the device family documentation for information about whether read and write ports must be synchronous.
- Some Verilog HDL statements with blocking assignments might not be mapped to RAM blocks, so avoid blocking statements when modeling RAMs in Verilog HDL.

For some device families, the `syn_ramstyle` attribute specifies the implementation to use for an inferred RAM. You can apply the `syn_ramstyle` attribute globally to a module or a RAM instance, to specify `registers` or `block_ram` values. To turn off RAM inference, set the attribute value to `registers`.

When inferring RAM for some Altera device families, the Synplify software generates additional bypass logic. This logic is generated to resolve a half-cycle read/write behavior difference between the RTL and post-synthesis simulations. The RTL simulation shows the memory being updated on the positive edge of the clock; the post-synthesis simulation shows the memory being updated on the negative edge of the clock. To eliminate bypass logic, the output of the RAM must be registered. By adding this register, the output of the RAM is seen after a full clock cycle, by which time the update has occurred, thus eliminating the need for bypass logic.

For devices with TriMatrix memory blocks, disable the creation of glue logic by setting the `syn_ramstyle` value to `no_rw_check`. Set `syn_ramstyle` to `no_rw_check` to disable the creation of glue logic in dual-port mode.

### Example 17-12: VHDL Code for Inferred Dual-Port RAM

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;
USE ieee.std_logic_signed.all;

ENTITY dualport_ram IS
PORT ( data_out: OUT STD_LOGIC_VECTOR (7 DOWNTO 0);
      data_in: IN STD_LOGIC_VECTOR (7 DOWNTO 0)
      wr_addr, rd_addr: IN STD_LOGIC_VECTOR (6 DOWNTO 0);
      we: IN STD_LOGIC);
      clk: IN STD_LOGIC);
END dualport_ram;

ARCHITECTURE ram_infer OF dualport_ram IS
TYPE Mem_Type IS ARRAY (127 DOWNTO 0) OF STD_LOGIC_VECTOR (7 DOWNTO 0);
SIGNAL mem: Mem_Type;
SIGNAL addr_reg: STD_LOGIC_VECTOR (6 DOWNTO 0);

BEGIN
  data_out <= mem (CONV_INTEGER(rd_addr));
  PROCESS (clk, we, data_in) BEGIN
    IF (clk='1' AND clk'EVENT) THEN
      IF (we='1') THEN
        mem(CONV_INTEGER(wr_addr)) <= data_in;
      END IF;
    END IF;
  END PROCESS;
END ram_infer;
```

```

    END PROCESS;
END ram_infer;

```

### Example 17-13: VHDL Code for Inferred Dual-Port RAM Preventing Bypass Logic

```

LIBRARY ieee;
USE ieee.std_logic_1164.all;
USE ieee.std_logic_signed.all;

ENTITY dualport_ram IS
PORT ( data_out: OUT STD_LOGIC_VECTOR (7 DOWNTO 0);
      data_in : IN  STD_LOGIC_VECTOR (7 DOWNTO 0);
      wr_addr, rd_addr : IN STD_LOGIC_VECTOR (6 DOWNTO 0);
      we : IN STD_LOGIC;
      clk : IN STD_LOGIC);
END dualport_ram;

ARCHITECTURE ram_infer OF dualport_ram IS
TYPE Mem_Type IS ARRAY (127 DOWNTO 0) OF STD_LOGIC_VECTOR (7 DOWNTO 0);
SIGNAL mem : Mem_Type;
SIGNAL addr_reg : STD_LOGIC_VECTOR (6 DOWNTO 0);
SIGNAL tmp_out : STD_LOGIC_VECTOR (7 DOWNTO 0); --output register

BEGIN
    tmp_out <= mem (CONV_INTEGER (rd_addr));
    PROCESS (clk, we, data_in) BEGIN
        IF (clk='1' AND clk'EVENT) THEN
            IF (we='1') THEN
                mem(CONV_INTEGER(wr_addr)) <= data_in;
            END IF;
            data_out <= tmp_out; --registers output preventing
                                -- bypass logic generation
        END IF;
    END PROCESS;
END ram_infer;

```

## RAM Initialization

Use the Verilog HDL `$readmemb` or `$readmemh` system tasks in your HDL code to initialize RAM memories. The Synplify compiler forward-annotates the initialization values in the `.srs` (technology-independent RTL netlist) file and the mapper generates the corresponding hexadecimal memory initialization (`.hex`) file. One `.hex` file is created for each of the `altsyncram` IP cores that are inferred in the design. The `.hex` file is associated with the `altsyncram` instance in the `.vqm` file using the `init_file` attribute.

The examples show how RAM can be initialized through HDL code, and how the corresponding `.hex` file is generated using Verilog HDL.

### Example 17-14: Using `$readmemb` System Task to Initialize an Inferred RAM in Verilog HDL Code

```

initial
begin
    $readmemb("mem.ini", mem);
end
always @(posedge clk)
begin
    raddr_reg <= raddr;
    if(we)

```

```

        mem[waddr] <= data;
    end

```

### Example 17-15: Sample of .vqm Instance Containing Memory Initialization File

```

altsyncram mem_hex( .wren_a(we), .wren_b(GND), ...);

defparam mem_hex.lpm_type = "altsyncram";
defparam mem_hex.operation_mode = "Dual_Port";
...
defparam mem_hex.init_file = "mem_hex.hex";

```

## Inferring ROM

When a ROM block is inferred from an HDL design, the Synplify software uses an Altera IP core to target the device memory architecture. For some devices, the Synplify software maps directly to memory block device atoms instead of instantiating an IP core in the **.vqm** file.

Follow these guidelines for the Synplify software to successfully infer ROM in a design:

- The address line must be at least two bits wide.
- The ROM must be at least half full.
- A CASE or IF statement must make 16 or more assignments using constant values of the same width.

## Inferring Shift Registers

The Synplify software infers shift registers for sequential shift components so that they can be placed in dedicated memory blocks in supported device architectures using the ALTSHIFT\_TAPS IP core.

If necessary, set the implementation style with the `syn_srlstyle` attribute. If you do not want the components automatically mapped to shift registers, set the value to `registers`. You can set the value globally, or on individual modules or registers.

For some designs, turning off shift register inference improves the design performance.

## Incremental Compilation and Block-Based Design

As designs become more complex and designers work in teams, a block-based incremental design flow is often an effective design approach. In an incremental compilation flow, you can make changes to part of the design while maintaining the placement and performance of unchanged parts of the design. Design iterations are made dramatically faster by focusing new compilations on particular design partitions and merging results with previous compilation results of other partitions. You can perform optimization on individual subblocks and then preserve the results before you integrate the blocks into a final design and optimize it at the top-level.

MultiPoint synthesis, which is available for certain device technologies in the Synplify Pro and Premier software, provides an automated block-based incremental synthesis flow. The MultiPoint feature manages a design hierarchy to let you design incrementally and synthesize designs that take too long for synthesis of the entire project. MultiPoint synthesis allows different netlist files to be created for different sections of a design hierarchy and supports the Quartus II incremental compilation methodology. This feature also ensures that only those sections of a design that have been updated are resynthesized when the design is



compiled, reducing synthesis run time and preserving the results for the unchanged blocks. You can change and resynthesize one section of a design without affecting other sections.

You can also partition your design and create different netlist files manually with the Synplify software by creating a separate project for the logic in each partition of the design. Creating different netlist files for each partition of the design also means that each partition can be independent of the others.

Hierarchical design methodologies can improve the efficiency of your design process, providing better design reuse opportunities and fewer integration problems when working in a team environment. When you use these incremental synthesis methodologies, you can take advantage of incremental compilation in the Quartus II software. You can perform placement and routing on only the changed partitions of the design, which reduces place-and-route time and preserves your fitting results.

#### Related Information

- [Quartus II Incremental Compilation for Hierarchical and Team-Based Design Documentation](#) on page 3-1

## Design Flow for Incremental Compilation

The following steps describe the general incremental compilation flow when using these features of the Quartus II software:

1. Create Verilog HDL or VHDL design files.
2. Determine which hierarchical blocks you want to treat as separate partitions in your design.
3. Set up your design using the MultiPoint synthesis feature or separate projects so that a separate netlist file is created for each design partition.
4. If using separate projects, disable I/O pad insertion in the implementations for lower-level partitions.
5. Compile and map each partition in the Synplify software, making constraints as you would in a non-incremental design flow.
6. Import the **.vqm** netlist and **.tcl** file for each partition into the Quartus II software and set up the Quartus II project(s) for incremental compilation.
7. Compile your design in the Quartus II software and preserve the compilation results with the post-fit netlist in incremental compilation.
8. When you make design or synthesis optimization changes to part of your design, resynthesize only the partition you modified to generate a new netlist and **.tcl** file. Do not regenerate netlist files for the unmodified partitions.
9. Import the new netlist and **.tcl** file into the Quartus II software and recompile the design in the Quartus II software with incremental compilation.

## Creating a Design with Separate Netlist Files for Incremental Compilation

The first stage of a hierarchical or incremental design flow is to ensure that different parts of your design do not affect each other. Ensure that you have separate netlists for each partition in your design so you can take advantage of incremental compilation in the Quartus II software. If the entire design is in one netlist file, changes in one partition might affect other partitions because of possible node name changes when you resynthesize the design.

To ensure proper functionality of the synthesis flow, create separate netlist files only for modules and entities. In addition, each module or entity requires its own design file. If two different modules are in the same design file, but are defined as being part of different partitions, incremental compilation cannot be maintained since both partitions must be recompiled when one module is changed.

Altera recommends that you register all inputs and outputs of each partition. This makes logic synchronous, and avoids any delay penalty on signals that cross partition boundaries.

If you use boundary tri-states in a lower-level block, the Synplify software pushes, or bubbles, the tri-states through the hierarchy to the top-level to use the tri-state drivers on output pins of Altera devices. Because bubbling tri-states requires optimizing through hierarchies, lower-level tri-states are not supported with a block-based compilation methodology. Use tri-state drivers only at the external output pins of the device and in the top-level block in the hierarchy.

You can generate multiple **.vqm** netlist files with the MultiPoint synthesis flow in the Synplify Pro and Premier software, or by manually creating separate Synplify projects and creating a black box for each block that you want to designate as a separate design partition.

In the MultiPoint synthesis flow in the Synplify Pro and Premier software, you create multiple **.vqm** netlist files from one easy-to-manage, top-level synthesis project. By using the manual black box method, you have multiple synthesis projects, which might be required for certain team-based or bottom-up designs where a single top-level project is not desired.

After you have created multiple **.vqm** files using one of these two methods, you must create the appropriate Quartus II projects to place-and-route the design.

#### Related Information

- [Best Practices for Incremental Compilation Partitions and Floorplan Assignments Documentation](#) on page 14-1

## Using MultiPoint Synthesis with Incremental Compilation

This topic describes how to generate multiple **.vqm** files using the Synplify Pro and Premier software MultiPoint synthesis flow. You must first set up your constraint file and Synplify options, then apply the appropriate Compile Point settings to write multiple **.vqm** files and create design partition assignments for incremental compilation.

#### Related Information

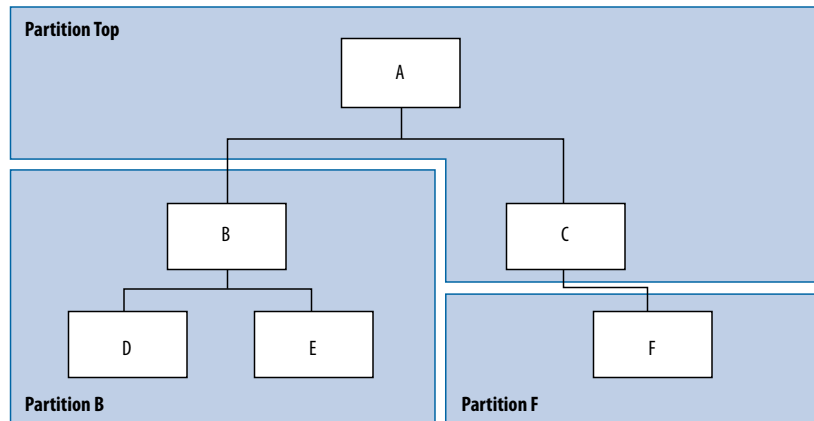
[Preserving Hierarchy](#) on page 17-12

### Set Compile Points and Create Constraint Files

The MultiPoint flow lets you segment a design into smaller synthesis units, called Compile Points. The synthesis software treats each Compile Point as a partition for incremental mapping, which allows you to isolate and work on each Compile Point module as independent segments of the larger design without impacting other design modules. A design can have any number of Compile Points, and Compile Points can be nested. The top-level module is always treated as a Compile Point.

Compile Points are optimized in isolation from their parent, which can be another Compile Point or a top-level design. Each block created with a Compile Point is unaffected by critical paths or constraints on its parent or other blocks. A Compile Point is independent, with its own individual constraints. During synthesis, any Compile Points that have not yet been synthesized are synthesized before the top level. Nested Compile Points are synthesized before the parent Compile Points in which they are contained. When you apply the appropriate setting for the Compile Point, a separate netlist is created for that Compile Point, isolating that logic from any other logic in the design.

The figure shows an example of a design hierarchy that is split into multiple partitions. The top-level block of each partition can be synthesized as a separate Compile Point.

**Figure 17-3: Partitions in a Hierarchical Design**

In this case, modules A, B, and F are Compile Points. The top-level Compile Point consists of the top-level block in the design (that is, block A in this example), including the logic that is not defined under another Compile Point. In this example, the design for top-level Compile Point A also includes the logic in one of its subblocks, C. Because block F is defined as its own Compile Point, it is not treated as part of the top-level Compile Point A. Another separate Compile Point B contains the logic in blocks B, D, and E. One netlist is created for the top-level module A and submodule C, another netlist is created for B and its submodules D and E, while a third netlist is created for F.

Apply Compile Points to the module, or to the architecture in the Synplify Pro SCOPE spreadsheet, or to the .sdc file. You cannot set a Compile Point in the Verilog HDL or VHDL source code. You can set the constraints manually using Tcl, by editing the .sdc file, or you can use the GUI.

### Defining Compile Points With .tcl or .sdc Files

To set Compile Points with a .tcl or .sdc file, use the `define_compile_point` command.

#### Example 17-16: The `define_compile_point` Command

```
define_compile_point [-disable] {<objname>} -type {locked, partition}
```

`<objname>` represents any module in the design. The Compile Point type `{locked, partition}` indicates that the Compile Point represents a partition for the Quartus II incremental compilation flow.

Each Compile Point has a set of constraint files that begin with the `define_current_design` command to set up the SCOPE environment, as follows:

```
define_current_design {<my_module>}
```

### Additional Considerations for Compile Points

To ensure that changes to a Compile Point do not affect the top-level parent module, turn off the **Update Compile Point Timing Data** option in the **Implementation Options** dialog box. If this option is turned on, updates to a child module can impact the top-level module.

You can apply the `syn_allowed_resources` attribute to any Compile Point view to restrict the number of resources for a particular module.

When using Compile Points with incremental compilation, be aware of the following restrictions:

- To use Compile Points effectively, you must provide timing constraints (timing budgeting) for each Compile Point; the more accurate the constraints, the better your results are. Constraints are not automatically budgeted, so manual time budgeting is essential. Altera recommends that you register all inputs and outputs of each partition. This avoids any logic delay penalty on signals that cross-partition boundaries.
- When using the Synplify attribute `syn_useioff` to pack registers in the I/O Elements (IOEs) of Altera devices, these registers must be in the top-level module. Otherwise, you must direct the Quartus II software to perform I/O register packing instead of the `syn_useioff` attribute. You can use the **Fast Input Register** or **Fast Output Register** options, or set I/O timing constraints and turn on **Optimize I/O cell register placement for timing** on the **Advanced Settings (Fitter)** dialog box in the Quartus II software.
- There is no incremental synthesis support for top-level logic; any logic in the top-level is resynthesized during every compilation in the Synplify software.

For more information about using Compile Points and setting Synplify attributes and constraints for both top-level and lower-level Compile Points, refer to the *Synopsys FPGA Synthesis User Guide* and the *Synopsys FPGA Synthesis Reference Manual*.

## Creating a Quartus II Project for Compile Points and Multiple .vqm Files

During compilation, the Synplify Pro and Premier software creates a `<top-level project>.tcl` file that provides the Quartus II software with the appropriate constraints and design partition assignments, creating a partition for each `.vqm` file along with the information to set up a Quartus II project.

Depending on your design methodology, you can create one Quartus II project for all netlists or a separate Quartus II project for each netlist. In the standard incremental compilation design flow, you create design partition assignments and optional LogicLock™ floorplan location assignments for each partition in the design within a single Quartus II project. This methodology allows for the best quality of results and performance preservation during incremental changes to your design.

You might require a bottom-up design flow if each partition must be optimized separately, such as for third-party IP delivery. If you use this flow, Altera recommends you create a design floorplan to avoid placement conflicts between each partition. To follow this design flow in the Quartus II software, create separate Quartus II projects, export each design partition and incorporate them into a top-level design using the incremental compilation features to maintain placement results.

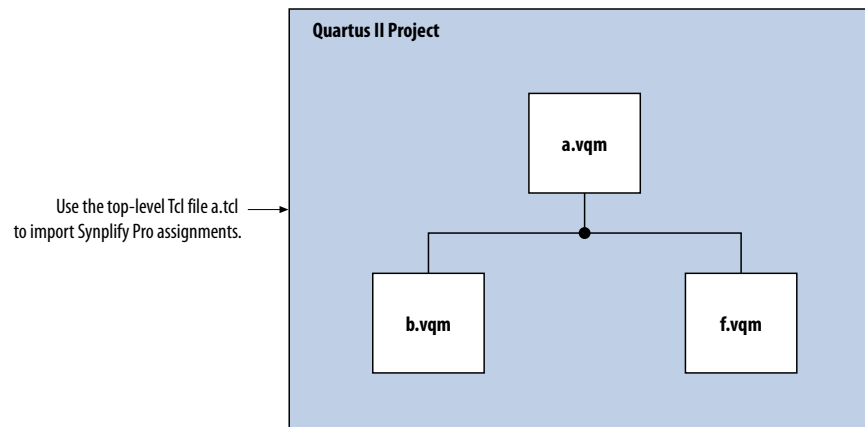
### Related Information

[Running the Quartus II Software Manually With the Synplify-Generated Tcl Script](#) on page 17-7

## Creating a Single Quartus II Project for a Standard Incremental Compilation Flow

Use the `<top-level project>.tcl` file that contains the Synplify assignments for all partitions within the project. This method allows you to import all the partitions into one Quartus II project and optimize all modules within the project at once, while taking advantage of the performance preservation and compilation-time reduction that incremental compilation offers.


Figure 17-4: Design Flow Using Multiple .vqm Files with One Quartus II Project



### Creating Multiple Quartus II Projects for a Bottom-Up Incremental Compilation Flow

Use the <lower-level compile point>.tcl files that contain the Synplify assignments for each Compile Point. Generate multiple Quartus II projects, one for each partition and netlist in the design. The designers in the project can optimize their own partitions separately within the Quartus II software and export the results for their own partitions. You can export the optimized subdesigns and then import them into one top-level Quartus II project using incremental compilation to complete the design.

Figure 17-5: Design Flow Using Multiple .vqm Files with Multiple Quartus II Projects




### Creating Multiple .vqm Files for a Incremental Compilation Flow With Separate Synplify Projects

You can manually generate multiple .vqm files for a incremental compilation flow with black boxes and separate Synplify projects for each design partition. This manual flow is supported by versions of the Synplify software without the MultiPoint Synthesis feature.

## Manually Creating Multiple .vqm Files With Black Boxes

To create multiple .vqm files manually in the Synplify software, create a separate project for each lower-level module and top-level design that you want to maintain as a separate .vqm file for an incremental compilation partition. Implement black box instantiations of lower-level partitions in your top-level project.

Figure 17-6: Partitions in a Hierarchical Design



The partition top contains the top-level block in the design (block A) and the logic that is not defined as part of another partition. In this example, the partition for top-level block A also includes the logic in one of its sub-blocks, block C. Because block F is contained in its own partition, it is not treated as part of the top-level partition A. Another separate partition, partition B, contains the logic in blocks B, D, and E. In a team-based design, engineers can work independently on the logic in different partitions. One netlist is created for the top-level module A and its submodule C, another netlist is created for module B and its submodules D and E, while a third netlist is created for module F.

### Creating Multiple .vqm Files for this Design

To create multiple .vqm files for this design, follow these steps:

1. Generate a .vqm file for module B. Use **B.v.vhd**, **D.v.vhd**, and **E.v.vhd** as the source files.
2. Generate a .vqm file for module F. Use **F.v.vhd** as the source files.
3. Generate a top-level .vqm file for module A. Use **A.v.vhd** and **C.v.vhd** as the source files. Ensure that you use black box modules B and F, which were optimized separately in the previous steps.

### Creating Black Boxes in Verilog HDL

Any design block that is not defined in the project, or included in the list of files to be read for a project, is treated as a black box by the software. Use the `syn_black_box` attribute to indicate that you intend to create a black box for the module. In Verilog HDL, you must provide an empty module declaration for a module that is treated as a black box.

The example shows the **A.v** top-level file. Follow the same procedure for lower-level files that also contain a black box for any module beneath the current level hierarchy.

### Example 17-17: Verilog HDL Black Box for Top-Level File A.v

```

module A (data_in, clk, e, ld, data_out);
    input data_in, clk, e, ld;
  
```

```

output [15:0] data_out;

wire [15:0] cnt_out;

B U1 (.data_in (data_in), .clk(clk), .ld (ld), .data_out(cnt_out));
F U2 (.d(cnt_out), .clk(clk), .e(e), .q(data_out));

// Any other code in A.v goes here.
endmodule

// Empty Module Declarations of Sub-Blocks B and F follow here.
// These module declarations (including ports) are required for black boxes.

module B (data_in, clk, ld, data_out) /* synthesis syn_black_box */ ;
    input data_in, clk, ld;
    output [15:0] data_out;
endmodule

module F (d, clk, e, q) /* synthesis syn_black_box */ ;
    input [15:0] d;
    input clk, e;
    output [15:0] q;
endmodule

```

## Creating Black Boxes in VHDL

Any design that is not defined in the project, or included in the list of files to be read for a project, is treated as a black box by the software. Use the `syn_black_box` attribute to indicate that you intend to treat the component as a black box. In VHDL, you must have a component declaration for the black box.

Although VHDL is not case-sensitive, a **.vqm** (a subset of Verilog HDL) file is case-sensitive. Entity names and their port declarations are forwarded to the **.vqm** file. Black box names and port declarations are also passed to the **.vqm** file. To prevent case-based mismatches, use the same capitalization for black box and entity declarations in VHDL designs.

The example shows the **A.vhd** top-level file. Follow this same procedure for any lower-level files that contain a black box for any block beneath the current level of hierarchy.

### Example 17-18: VHDL Black Box for Top-Level File A.vhd

```

LIBRARY ieee;
USE ieee.std_logic_1164.all;
LIBRARY synplify;
USE synplify.attributes.all;

ENTITY A IS
PORT (data_in : IN INTEGER RANGE 0 TO 15;
      clk, e, ld : IN STD_LOGIC;
      data_out : OUT INTEGER RANGE 0 TO 15 );
END A;

ARCHITECTURE a_arch OF A IS

COMPONENT B PORT(
    data_in : IN INTEGER RANGE 0 TO 15;
    clk, ld : IN STD_LOGIC;
    d_out : OUT INTEGER RANGE 0 TO 15);
END COMPONENT;

COMPONENT F PORT(
    d : IN INTEGER RANGE 0 TO 15;
    clk, e: IN STD_LOGIC;
    q : OUT INTEGER RANGE 0 TO 15);

```

```

END COMPONENT;

attribute syn_black_box of B: component is true;
attribute syn_black_box of F: component is true;

-- Other component declarations in A.vhd go here
signal cnt_out : INTEGER RANGE 0 TO 15;

BEGIN

U1 : B
PORT MAP (
    data_in => data_in,
    clk => clk,
    ld => ld,
    d_out => cnt_out );

U2 : F
PORT MAP (
    d => cnt_out,
    clk => clk,
    e => e,
    q => data_out );

-- Any other code in A.vhd goes here

END a_arch;

```

After you complete the steps above, you have a netlist for each partition of the design. These files are ready for use with the incremental compilation flow in the Quartus II software.

## Creating a Quartus II Project for Multiple .vqm Files

The Synplify software creates a **.tcl** file for each **.vqm** file that provides the Quartus II software with the appropriate constraints and information to set up a project.

Depending on your design methodology, you can create one Quartus II project for all netlists or a separate Quartus II project for each netlist. In the standard incremental compilation design flow, you create design partition assignments and optional LogicLock floorplan location assignments for each partition in the design within a single Quartus II project. This methodology allows for the best quality of results and performance preservation during incremental changes to your design. You might require a bottom-up design flow where each partition must be optimized separately, such as for third-party IP delivery.

To perform this design flow in the Quartus II software, create separate Quartus II projects, export each design partition and incorporate it into a top-level design using the incremental compilation features to maintain the results.

### Related Information

[Running the Quartus II Software Manually With the Synplify-Generated Tcl Script](#) on page 17-7


## Creating a Single Quartus II Project for a Standard Incremental Compilation Flow

Use the *<top-level project>.tcl* file that contains the Synplify assignments for the top-level design. This method allows you to import all of the partitions into one Quartus II project and optimize all modules within the project at once, taking advantage of the performance preservation and compilation time reduction offered by incremental compilation.

All of the constraints from the top-level project are passed to the Quartus II software in the top-level **.tcl** file, but constraints made in the lower-level projects within the Synplify software are not forward-annotated. Enter these constraints manually in your Quartus II project.



Figure 17-7: Design Flow Using Multiple .vqm Files with One Quartus II Project

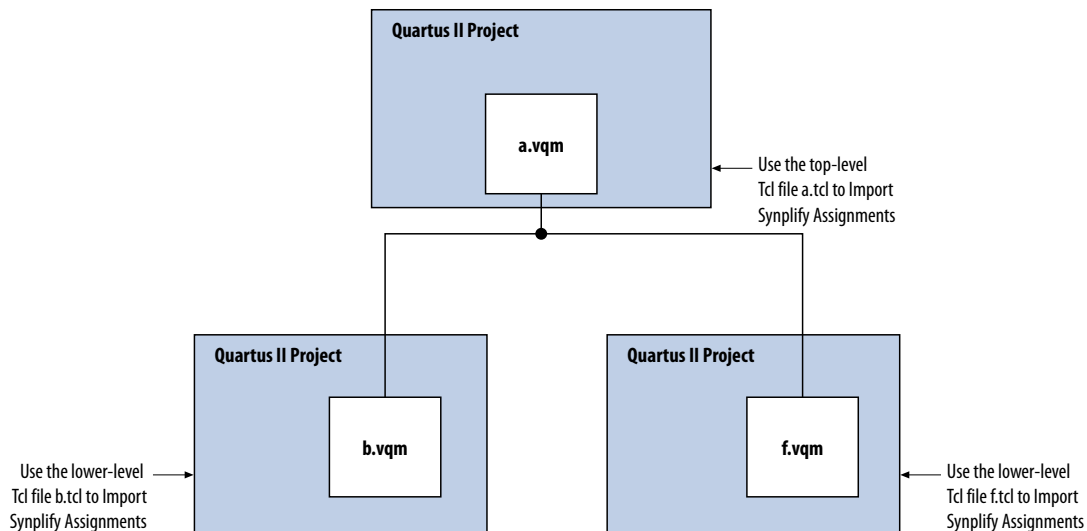


### Creating Multiple Quartus II Projects for a Bottom-Up Incremental Compilation Flow

Use the .tcl file that is created for each .vqm file by the Synplify software for each Synplify project. This method generates multiple Quartus II projects, one for each block in the design. The designers in the project can optimize their own blocks separately within the Quartus II software and export the placement of their own blocks.

Designers should create a LogicLock region to create a design floorplan for each block to avoid conflicts between partitions. The top-level designer then imports all the blocks and assignments into the top-level project. This method allows each block in the design to be optimized separately and then imported into one top-level project.

Figure 17-8: Design Flow Using Multiple Synplify Projects and Multiple Quartus II Projects



### Performing Incremental Compilation in the Quartus II Software

In a standard design flow using Multipoint Synthesis, the Synplify software uses the Quartus II top-level .tcl file to ensure that the two tools databases stay synchronized. The Tcl file creates, changes, or

deletes partition assignments in the Quartus II software for Compile Points that you create, change, or delete in the Synplify software. However, if you create, change, or delete a partition in the Quartus II software, the Synplify software does not change your Compile Point settings. Make any corresponding change in your Synplify project to ensure that you create the correct **.vqm** files.

**Note:** If you use the NativeLink integration feature, the Synplify software does not use any information about design partition assignments that you have set in the Quartus II software.

If you create netlist files with multiple Synplify projects, or if you do not use the Synplify Pro or Premier-generated **.tcl** files to update constraints in your Quartus II project, you must ensure that your Synplify **.vqm** netlists align with your Quartus II partition settings.

After you have set up your Quartus II project with **.vqm** netlist files as separate design partitions, set the appropriate Quartus II options to preserve your compilation results. On the Assignments menu, click **Design Partitions Window**. Change the **Netlist Type** to **Post-Fit** to preserve the previous compilation's post-fit placement results. If you do not make these settings, the Quartus II software does not reuse the placement or routing results from the previous compilation.

You can take advantage of incremental compilation with your Synplify design to reduce compilation time in the Quartus II software and preserve the results for unchanged design blocks.

#### Related Information

[Using the Quartus II Software to Run the Synplify Software](#) on page 17-5

[Quartus II Incremental Compilation for Hierarchical and Team-Based Design Documentation](#) on page 3-1

## Document Revision History

Table 17-4: Document Revision History

Date	Version	Changes
2014.12.15	14.1.0	Updated location of Fitter Settings, Analysis & Synthesis Settings, and Physical Optimization Settings to Compiler Settings.
November 2013	13.1.0	Dita conversion. Restructured content.
June 2012	12.0.0	Removed survey link.
November 2011	10.1.1	Template update.

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> <li>• Changed to new document template.</li> <li>• Removed Classic Timing Analyzer support.</li> <li>• Removed the “altera_ implement_in_esb or altera_ implement_in_eab” section.</li> <li>• Edited the “Creating a Quartus II Project for Compile Points and Multiple .vqm Files” on page 14–33 section for changes with the incremental compilation flow.</li> <li>• Edited the “Creating a Quartus II Project for Multiple .vqm Files” on page 14–39 section for changes with the incremental compilation flow.</li> <li>• Editorial changes.</li> </ul>
July 2010	10.0.0	<ul style="list-style-type: none"> <li>• Minor updates for the Quartus II software version 10.0 release.</li> </ul>
November 2009	9.1.0	<ul style="list-style-type: none"> <li>• Minor updates for the Quartus II software version 9.1 release.</li> </ul>
March 2009	9.0.0	<ul style="list-style-type: none"> <li>• Added new section “Exporting Designs to the Quartus II Software Using NativeLink Integration” on page 14–14.</li> <li>• Minor updates for the Quartus II software version 9.0 release.</li> <li>• Chapter 10 was previously Chapter 9 in software version 8.1.</li> </ul>

Date	Version	Changes
November 2008	8.1.0	<ul style="list-style-type: none"><li>• Changed to 8-1/2 x 11 page size</li><li>• Changed the chapter title from “Synplicity Synplify &amp; Synplify Pro Support” to “Synopsys Synplify Support”</li><li>• Replaced references to Synplicity with references to Synopsys</li><li>• Added information about Synplify Premier</li><li>• Updated supported device list</li><li>• Added SystemVerilog information to Figure 14-1</li></ul>

Date	Version	Changes
May 2008	8.0.0	<ul style="list-style-type: none"> <li>• Updated supported device list</li> <li>• Updated constraint annotation information for the TimeQuest Timing Analyzer</li> <li>• Updated RAM and MAC constraint limitations</li> <li>• Revised Table 9-1</li> <li>• Added new section “Changing Synplify’s Default Behavior for Instantiated Altera Megafunctions”</li> <li>• Added new section “Instantiating Intellectual Property Using the MegaWizard Plug-In Manager and IP Toolbench”</li> <li>• Added new section “Including Files for Quartus II Placement and Routing Only”</li> <li>• Added new section “Additional Considerations for Compile Points”</li> <li>• Removed section “Apply the LogicLock Attributes”</li> <li>• Modified Figure 9-4, 9-43, 9-47, and 9-48</li> <li>• Added new section “Performing Incremental Compilation in the Quartus II Software”</li> <li>• Numerous text changes and additions throughout the chapter</li> <li>• Renamed several sections</li> <li>• Updated “Referenced Documents” section</li> </ul>

**Related Information**

[Quartus II Handbook Archive Website](#)

2014.06.30

QII5V1



Subscribe



Send Feedback

## About Precision RTL Synthesis Support

This manual delineates the support for the Mentor Graphics® Precision RTL Synthesis and Precision RTL Plus Synthesis software in the Quartus® II software, as well as key design flows, methodologies and techniques for improving your results for Altera® devices. This manual assumes that you have set up, licensed, and installed the Precision Synthesis software and the Quartus II software. You must set up, license, and install the Precision RTL Plus Synthesis software if you want to use the incremental synthesis feature for incremental compilation and block-based design.

To obtain and license the Precision Synthesis software, refer to the Mentor Graphics website. To install and run the Precision Synthesis software and to set up your work environment, refer to the *Precision Synthesis Installation Guide* in the Precision Manuals Bookcase. To access the Manuals Bookcase in the Precision Synthesis software, click **Help** and select **Open Manuals Bookcase**.

### Related Information

[Mentor Graphics website](#)

## Design Flow

The following steps describe a basic Quartus II design flow using the Precision Synthesis software:

1. Create Verilog HDL or VHDL design files.
2. Create a project in the Precision Synthesis software that contains the HDL files for your design, select your target device, and set global constraints.
3. Compile the project in the Precision Synthesis software.
4. Add specific timing constraints, optimization attributes, and compiler directives to optimize the design during synthesis. With the design analysis and cross-probing capabilities of the Precision Synthesis software, you can identify and improve circuit area and performance issues using prelayout timing estimates.

**Note:** For best results, Mentor Graphics recommends specifying constraints that are as close as possible to actual operating requirements. Properly setting clock and I/O constraints, assigning clock domains, and indicating false and multicycle paths guide the synthesis algorithms more accurately toward a suitable solution in the shortest synthesis time.

5. Synthesize the project in the Precision Synthesis software.
6. Create a Quartus II project and import the following files generated by the Precision Synthesis software into the Quartus II project:

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered

- The Verilog Quartus Mapping File ( **.vqm** ) netlist
- Synopsys Design Constraints File ( **.sdc** ) for TimeQuest Timing Analyzer constraints
- Tcl Script Files ( **.tcl** ) to set up your Quartus II project and pass constraints

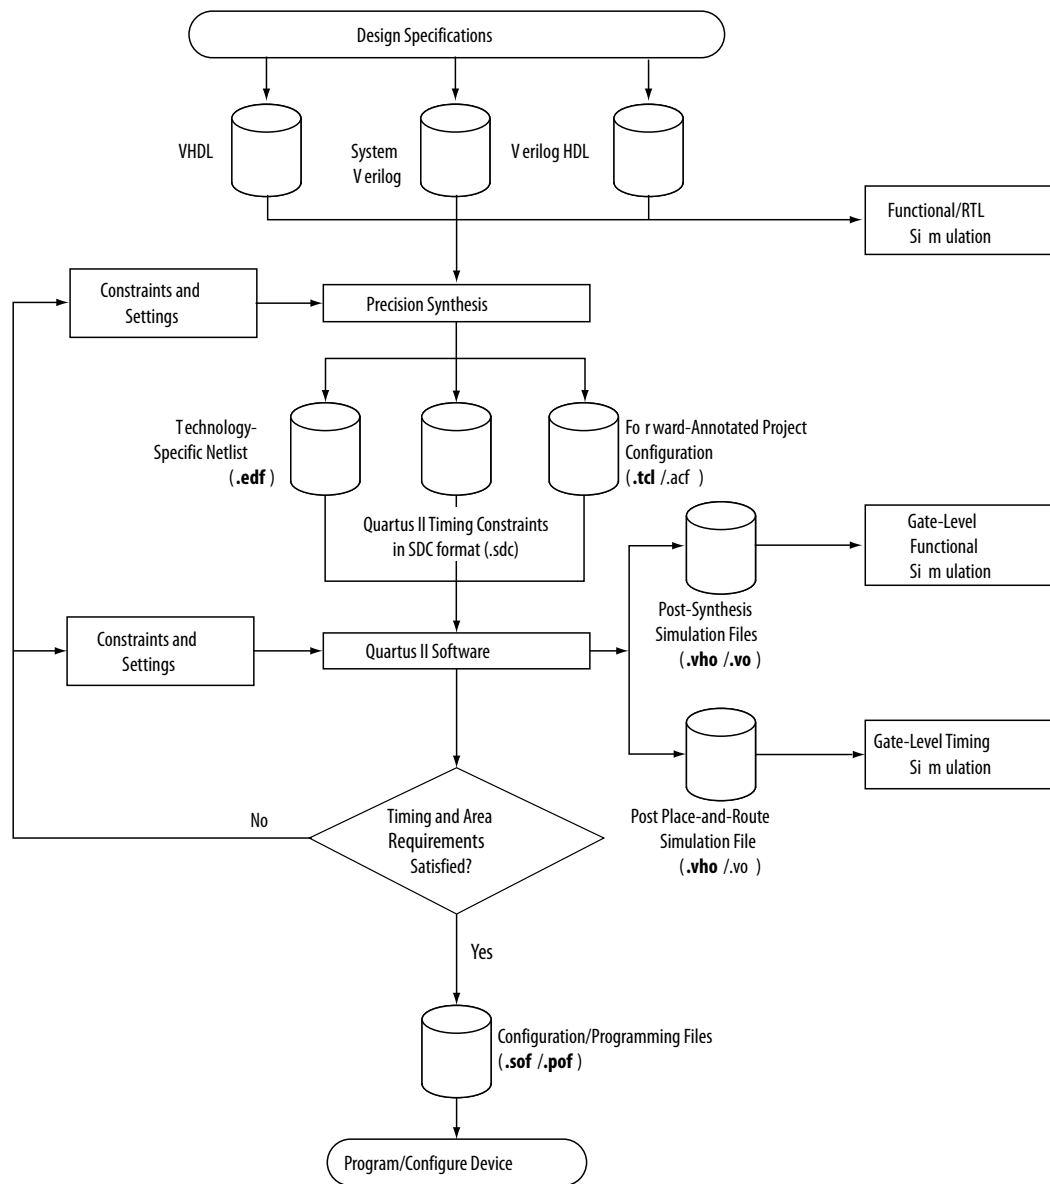
**Note:** If your design uses the Classic Timing Analyzer for timing analysis in the Quartus II software versions 10.0 and earlier, the Precision Synthesis software generates timing constraints in the Tcl Constraints File ( **.tcl** ). If you are using the Quartus II software versions 10.1 and later, you must use the TimeQuest Timing Analyzer for timing analysis.

7. After obtaining place-and-route results that meet your requirements, configure or program the Altera device.

You can run the Quartus II software from within the Precision Synthesis software, or run the Precision Synthesis software using the Quartus II software.



Figure 18-1: Design Flow Using the Precision Synthesis Software and Quartus II Software



**Related Information**

- [Running the Quartus II Software from within the Precision Synthesis Software](#) on page 18-9
- [Using the Quartus II Software to Run the Precision Synthesis Software](#) on page 18-10

**Timing Optimization**

If your area or timing requirements are not met, you can change the constraints and resynthesize the design in the Precision Synthesis software, or you can change the constraints to optimize the design during place-and-route in the Quartus II software. Repeat the process until the area and timing requirements are met.



You can use other options and techniques in the Quartus II software to meet area and timing requirements. For example, the **WYSIWYG Primitive Resynthesis** option can perform optimizations on your EDIF netlist in the Quartus II software.

While simulation and analysis can be performed at various points in the design process, final timing analysis should be performed after placement and routing is complete.

#### Related Information

- [Netlist Optimizations and Physical Synthesis documentation](#)
- [Timing Closure and Optimization documentation](#)

## Altera Device Family Support

The Precision Synthesis software supports active devices available in the current version of the Quartus II software. Support for newly released device families may require an overlay. Contact Mentor Graphics for more information.

## Precision Synthesis Generated Files

During synthesis, the Precision Synthesis software produces several intermediate and output files.

**Table 18-1: Precision Synthesis Software Intermediate and Output Files**

File Extension	File Description
<b>.psp</b>	Precision Synthesis Project File.
<b>.xdb</b>	Mentor Graphics Design Database File.
<b>.rep<sup>(15)</sup></b>	Synthesis Area and Timing Report File.
<b>.vqm<sup>(16)</sup></b>	Technology-specific netlist in <b>.vqm</b> file format.  By default, the Precision Synthesis software creates <b>.vqm</b> files for Arria series, Cyclone series, and Stratix series devices. The Precision Synthesis software defaults to creating <b>.vqm</b> files when the device is supported.

<sup>(15)</sup> The timing report file includes performance estimates that are based on pre-place-and-route information. Use the  $f_{MAX}$  reported by the Quartus II software after place-and-route for accurate post-place-and-route timing information. The area report file includes post-synthesis device resource utilization statistics that can differ from the resource usage after place-and-route due to black boxes or further optimizations performed during placement and routing. Use the device utilization reported by the Quartus II software after place-and-route for final resource utilization results.

<sup>(16)</sup> The Precision Synthesis software-generated VQM file is supported by the Quartus II software version 10.1 and later.

File Extension	File Description
<b>.tcl</b>	Forward-annotated Tcl assignments and constraints file. The <i>&lt;project name&gt;.tcl</i> file is generated for all devices. The <b>.tcl</b> file acts as the Quartus II Project Configuration file and is used to make basic project and placement assignments, and to create and compile a Quartus II project.
<b>.acf</b>	Assignment and Configurations file for backward compatibility with the MAX+PLUS II software. For devices supported by the MAX+PLUS II software, the MAX+PLUS II assignments are imported from the MAX+PLUS II <b>.acf</b> file.
<b>.sdc</b>	Quartus II timing constraints file in Synopsys Design Constraints format. This file is generated automatically if the device uses the TimeQuest Timing Analyzer by default in the Quartus II software, and has the naming convention <i>&lt;project name&gt;_pnr_constraints.sdc</i> .

#### Related Information

- [Exporting Designs to the Quartus II Software Using NativeLink Integration](#) on page 18-9
- [Synthesizing the Design and Evaluating the Results](#) on page 18-8

## Creating and Compiling a Project in the Precision Synthesis Software

After creating your design files, create a project in the Precision Synthesis software that contains the basic settings for compiling the design.

## Mapping the Precision Synthesis Design

In the next steps, you set constraints and map the design to technology-specific cells. The Precision Synthesis software maps the design by default to the fastest possible implementation that meets your timing constraints. To accomplish this, you must specify timing requirements for the automatically determined clock sources. With this information, the Precision Synthesis software performs static timing analysis to determine the location of the critical timing paths. The Precision Synthesis software achieves the best results for your design when you set as many realistic constraints as possible. Be sure to set constraints for timing, mapping, false paths, multicycle paths, and other factors that control the structure of the implemented design.

Mentor Graphics recommends creating an **.sdc** file and adding this file to the **Constraint Files** section of the **Project Files** list. You can create this file with a text editor, by issuing command-line constraint parameters, or by directing the Precision Synthesis software to generate the file automatically the first time you synthesize your design. By default, the Precision Synthesis software saves all timing constraints and attributes in two files: **precision\_rtl.sdc** and **precision\_tech.sdc**. The **precision\_rtl.sdc** file contains constraints set on the RTL-level database (post-compilation) and the **precision\_tech.sdc** file contains constraints set on the gate-level database (post-synthesis) located in the current implementation directory.

You can also enter constraints at the command line. After adding constraints at the command line, update the `.sdc` file with the `update_constraint_file` command. You can add constraints that change infrequently directly to the HDL source files with HDL attributes or pragmas.

**Note:** The Precision `.sdc` file contains all the constraints for the Precision Synthesis project. For the Quartus II software, placement constraints are written in a `.tcl` file and timing constraints for the TimeQuest Timing Analyzer are written in the Quartus II `.sdc` file.

For details about the syntax of Synopsys Design Constraint commands, refer to the *Precision RTL Synthesis User's Manual* and the *Precision Synthesis Reference Manual*. For more details and examples of attributes, refer to the *Attributes* chapter in the *Precision Synthesis Reference Manual*.

## Setting Timing Constraints

The Precision Synthesis software uses timing constraints, based on the industry-standard `.sdc` file format, to deliver optimal results. Missing timing constraints can result in incomplete timing analysis and might prevent timing errors from being detected. The Precision Synthesis software provides constraint analysis prior to synthesis to ensure that designs are fully and accurately constrained. The `<project name>_pnr_constraints.sdc` file, which contains timing constraints in SDC format, is generated in the Quartus II software.

**Note:** Because the `.sdc` file format requires that timing constraints be set relative to defined clocks, you must specify your clock constraints before applying any other timing constraints.

You also can use multicycle path and false path assignments to relax requirements or exclude nodes from timing requirements, which can improve area utilization and allow the software optimizations to focus on the most critical parts of the design.

For details about the syntax of Synopsys Design Constraint commands, refer to the *Precision RTL Synthesis User's Manual* and the *Precision Synthesis Reference Manual*.

## Setting Mapping Constraints

Mapping constraints affect how your design is mapped into the target Altera device. You can set mapping constraints in the user interface, in HDL code, or with the `set_attribute` command in the constraint file.

## Assigning Pin Numbers and I/O Settings

The Precision Synthesis software supports assigning device pin numbers, I/O standards, drive strengths, and slew-rate settings to top-level ports of the design. You can set these timing constraints with the `set_attribute` command, the GUI, or by specifying synthesis attributes in your HDL code. These constraints are forward-annotated in the `<project name>.tcl` file that is read by the Quartus II software during place-and-route and do not affect synthesis.

You can use the `set_attribute` command in the Precision Synthesis software `.sdc` file to specify pin number constraints, I/O standards, drive strengths, and slow slew-rate settings. The table below describes the format to use for entries in the Precision Synthesis software constraint file.

**Table 18-2: Constraint File Settings**

Constraint	Entry Format for Precision Constraint File
Pin number	<code>set_attribute -name PIN_NUMBER -value "&lt;pin number&gt;" -port &lt;port name&gt;</code>
I/O standard	<code>set_attribute -name IOSTANDARD -value "&lt;I/O Standard&gt;" -port &lt;port name&gt;</code>
Drive strength	<code>set_attribute -name DRIVE -value "&lt;drive strength in mA&gt;" -port &lt;port name&gt;</code>
Slew rate	<code>set_attribute -name SLEW -value "TRUE   FALSE" -port &lt;port name&gt;</code>

You also can use synthesis attributes or pragmas in your HDL code to make these assignments.

### Example 18-1: Verilog HDL Pin Assignment

```
//pragma attribute clk pin_number P10;
```

### Example 18-2: VHDL Pin Assignment

```
attribute pin_number : string
attribute pin_number of clk : signal is "P10";
```

You can use the same syntax to assign the I/O standard using the `IOSTANDARD` attribute, drive strength using the attribute `DRIVE`, and slew rate using the `SLEW` attribute.

For more details about attributes and how to set these attributes in your HDL code, refer to the *Precision Synthesis Reference Manual*.

## Assigning I/O Registers

The Precision Synthesis software performs timing-driven I/O register mapping by default. You can force a register to the device IO element (IOE) using the Complex I/O constraint. This option does not apply if you turn off **I/O pad insertion**.

**Note:** You also can make the assignment by right-clicking on the pin in the Schematic Viewer.

For the Stratix series, Cyclone series, and the MAX II device families, the Precision Synthesis software can move an internal register to an I/O register without any restrictions on design hierarchy.

For more mature devices, the Precision Synthesis software can move an internal register to an I/O register only when the register exists in the top-level of the hierarchy. If the register is buried in the hierarchy, you must flatten the hierarchy so that the buried registers are moved to the top-level of the design.

## Disabling I/O Pad Insertion

The Precision Synthesis software assigns I/O pad atoms (device primitives used to represent the I/O pins and I/O registers) to all ports in the top-level of a design by default. In certain situations, you might not

want the software to add I/O pads to all I/O pins in the design. The Quartus II software can compile a design without I/O pads; however, including I/O pads provides the Precision Synthesis software with more information about the top-level pins in the design.

### Preventing the Precision Synthesis Software from Adding I/O Pads

If you are compiling a subdesign as a separate project, I/O pins cannot be primary inputs or outputs of the device; therefore, the I/O pins should not have an I/O pad associated with them.

To prevent the Precision Synthesis software from adding I/O pads:

- You can use the Precision Synthesis GUI or add the following command to the project file:

```
setup_design -addio=false
```

### Preventing the Precision Synthesis Software from Adding an I/O Pad on an Individual Pin

To prevent I/O pad insertion on an individual pin when you are using a black box, such as DDR or a phase-locked loop (PLL), at the external ports of the design, perform the following steps:

1. Compile your design.
2. Use the Precision Synthesis GUI to select the individual pin and turn off I/O pad insertion.

**Note:** You also can make this assignment by attaching the `nopad` attribute to the port in the HDL source code.

### Controlling Fan-Out on Data Nets

Fan-out is defined as the number of nodes driven by an instance or top-level port. High fan-out nets can cause significant delays that result in an unroutable net. On a critical path, high fan-out nets can cause longer delays in a single net segment that result in the timing constraints not being met. To prevent this behavior, each device family has a global fan-out value set in the Precision Synthesis software library. In addition, the Quartus II software automatically routes high fan-out signals on global routing lines in the Altera device whenever possible.

To eliminate routability and timing issues associated with high fan-out nets, the Precision Synthesis software also allows you to override the library default value on a global or individual net basis. You can override the library value by setting a `max_fanout` attribute on the net.

## Synthesizing the Design and Evaluating the Results

During synthesis, the Precision Synthesis software optimizes the compiled design, and then writes out netlists and reports to the implementation subdirectory of your working directory after the implementation is saved, using the following naming convention:

```
<project name>_impl_<number>
```

After synthesis is complete, you can evaluate the results for area and timing. The *Precision RTL Synthesis User's Manual* describes different results that can be evaluated in the software.

There are several schematic viewers available in the Precision Synthesis software: RTL schematic, Technology-mapped schematic, and Critical Path schematic. These analysis tools allow you to quickly and easily isolate the source of timing or area issues, and to make additional constraint or code changes to optimize the design.

## Obtaining Accurate Logic Utilization and Timing Analysis Reports

Historically, designers have relied on post-synthesis logic utilization and timing reports to determine the amount of logic their design requires, the size of the device required, and how fast the design runs. However, today's FPGA devices provide a wide variety of advanced features in addition to basic registers and look-up tables (LUTs). The Quartus II software has advanced algorithms to take advantage of these features, as well as optimization techniques to increase performance and reduce the amount of logic required for a given design. In addition, designs can contain black boxes and functions that take advantage of specific device features. Because of these advances, synthesis tool reports provide post-synthesis area and timing estimates, but you should use the place-and-route software to obtain final logic utilization and timing reports.

## Exporting Designs to the Quartus II Software Using NativeLink Integration

The NativeLink feature in the Quartus II software facilitates the seamless transfer of information between the Quartus II software and EDA tools, which allows you to run other EDA design entry/synthesis, simulation, and timing analysis tools automatically from within the Quartus II software.

After a design is synthesized in the Precision Synthesis software, the technology-mapped design is written to the current implementation directory as an EDIF netlist file, along with a Quartus II Project Configuration File and a place-and-route constraints file. You can use the Project Configuration script, `<project name>.tcl`, to create and compile a Quartus II project for your EDIF or VQM netlist. This script makes basic project assignments, such as assigning the target device specified in the Precision Synthesis software. If you select a newer Altera device, the constraints are written in SDC format to the `<project name>_pnr_constraints.sdc` file by default, which is used by the Fitter and the TimeQuest Timing Analyzer in the Quartus II software.

Use the following Precision Synthesis software command before compilation to generate the `<project name>_pnr_constraints.sdc`:

```
setup_design -timequest_sdc
```

With this command, the file is generated after synthesis.

## Running the Quartus II Software from within the Precision Synthesis Software

The Precision Synthesis software also has a built-in place-and-route environment that allows you to run the Quartus II Fitter and view the results in the Precision Synthesis GUI. This feature is useful when performing an initial compilation of your design to view post-place-and-route timing and device utilization results. Not all the advanced Quartus II options that control the compilation process are available when you use this feature.

Two primary Precision Synthesis software commands control the place-and-route process. Use the `setup_place_and_route` command to set the place-and-route options. Start the process with the `place_and_route` command.

Precision Synthesis software uses individual Quartus II executables, such as analysis and synthesis (`quartus_map`), Fitter (`quartus_fit`), and the TimeQuest Timing Analyzer (`quartus_sta`) for improved runtime and memory utilization during place and route. This flow is referred to as the **Quartus II Modular** flow option in the Precision Synthesis software. By default, the Precision Synthesis software generates a Quartus II Project Configuration File (`.tcl` file) for current device families. Timing constraints

**Tcl Script**

that you set during synthesis are exported to the Quartus II place-and-route constraints file `<project name>_pnr_constraints.sdc`.

After you compile the design in the Quartus II software from within the Precision Synthesis software, you can invoke the Quartus II GUI manually and then open the project using the generated Quartus II project file. You can view reports, run analysis tools, specify options, and run the various processing flows available in the Quartus II software.

For more information about running the Quartus II software from within the Precision Synthesis software, refer to the *Altera Quartus II Integration* chapter in the *Precision Synthesis Reference Manual*.

## Running the Quartus II Software Manually Using the Precision Synthesis-Generated Tcl Script

You can run the Quartus II software using a Tcl script generated by the Precision Synthesis software. To run the Tcl script generated by the Precision Synthesis software to set up your project and start a full compilation, perform the following steps:

1. Ensure the `.vqm` file, `.tcl` files, and `.sdc` file are located in the same directory. The files should be located in the implementation directory by default.
2. In the Quartus II software, on the View menu, point to **Utility Windows** and click **Tcl Console**.
3. At the Tcl Console command prompt, type the command:

```
source <path>/<project name>.tcl
```

4. On the File menu, click **Open Project**. Browse to the project name and click **Open**.
5. Compile the project in the Quartus II software.

## Using the Quartus II Software to Run the Precision Synthesis Software

With NativeLink integration, you can set up the Quartus II software to run the Precision Synthesis software. This feature allows you to use the Precision Synthesis software to synthesize a design as part of a standard compilation. When you use this feature, the Precision Synthesis software does not use any timing constraints or assignments, such as incremental compilation partitions, that you have set in the Quartus II software.

### Related Information

- [Exporting Designs to the Quartus II Software Using NativeLink Integration](#) on page 18-9
- [Using the NativeLink Feature with Other EDA Tools online help](#)

## Passing Constraints to the Quartus II Software

The place-and-route constraints script forward-annotates timing constraints that you made in the Precision Synthesis software. This integration allows you to enter these constraints once in the Precision Synthesis software, and then pass them automatically to the Quartus II software.

The following constraints are translated by the Precision Synthesis software and are applicable to the TimeQuest Timing Analyzer:

- `create_clock`
- `set_input_delay`
- `set_output_delay`
- `set_max_delay`

- set\_min\_delay
- set\_false\_path
- set\_multicycle\_path

## create\_clock

You can specify a clock in the Precision Synthesis software.

### Example 18-3: Specifying a Clock Using create\_clock

```
create_clock -name <clock_name> -period <period in ns> \  
-waveform {<edge_list>} -domain <ClockDomain> <pin>
```

The period is specified in units of nanoseconds (ns). If no clock domain is specified, the clock belongs to a default clock domain `main`. All clocks in the same clock domain are treated as synchronous (related) clocks. If no `<clock_name>` is provided, the default name `virtual_default` is used. The `<edge_list>` sets the rise and fall edges of the clock signal over an entire clock period. The first value in the list is a rising transition, typically the first rising transition after time zero. The waveform can contain any even number of alternating edges, and the edges listed should alternate between rising and falling. The position of any edge can be equal to or greater than zero but must be equal to or less than the clock period.

If `-waveform <edge_list>` is not specified and `-period <period in ns>` is specified, the default waveform has a rising edge of 0.0 and a falling edge of `<period_value>/2`.

The Precision Synthesis software maps the clock constraint to the TimeQuest `create_clock` setting in the Quartus II software.

The Quartus II software supports only clock waveforms with two edges in a clock cycle. If the Precision Synthesis software finds a multi-edge clock, it issues an error message when you synthesize your design in the Precision Synthesis software.

## set\_input\_delay

This port-specific input delay constraint is specified in the Precision Synthesis software.

### Example 18-4: Specifying set\_input\_delay

```
set_input_delay {<delay_value> <port_pin_list>} \  
-clock <clock_name> -rise -fall -add_delay
```

This constraint is mapped to the `set_input_delay` setting in the Quartus II software.

When the reference clock `<clock_name>` is not specified, all clocks are assumed to be the reference clocks for this assignment. The input pin name for the assignment can be an input pin name of a time group. The software can use the `clock_fall` option to specify delay relative to the falling edge of the clock.

**Note:** Although the Precision Synthesis software allows you to set input delays on pins inside the design, these constraints are not sent to the Quartus II software, and a message is displayed.



## set\_output\_delay

This port-specific output delay constraint is specified in the Precision Synthesis software.

### Example 18-5: Using the set\_output\_delay Constraint

```
set_output_delay {<delay_value> <port_pin_list>} \  
-clock <clock_name> -rise -fall -add_delay
```

This constraint is mapped to the `set_output_delay` setting in the Quartus II software.

When the reference clock `<clock_name>` is not specified, all clocks are assumed to be the reference clocks for this assignment. The output pin name for the assignment can be an output pin name of a time group.

**Note:** Although the Precision Synthesis software allows you to set output delays on pins inside the design, these constraints are not sent to the Quartus II software.

## set\_max\_delay and set\_min\_delay

The maximum delay and minimum delay for a point-to-point timing path constraint is specified in the Precision Synthesis software.

### Example 18-6: Using the set\_max\_delay Constraint

```
set_max_delay -from {<from_node_list>} -to {<to_node_list>} <delay_value>
```

### Example 18-7: Using the set\_min\_delay Constraint

```
set_min_delay -from {<from_node_list>} -to {<to_node_list>} <delay_value>
```

The `set_max_delay` and `set_min_delay` commands specify that the maximum and minimum respectively, required delay for any start point in `<from_node_list>` to any endpoint in `<to_node_list>` must be less than or greater than `<delay_value>`. Typically, you use these commands to override the default setup constraint for any path with a specific maximum or minimum time value for the path.

The node lists can contain a collection of clocks, registers, ports, pins, or cells. The `-from` and `-to` parameters specify the source (start point) and the destination (endpoint) of the timing path, respectively. The source list (`<from_node_list>`) cannot include output ports, and the destination list (`<to_node_list>`) cannot include input ports. If you include more than one node on a list, you must enclose the nodes in quotes or in braces (`{ }`).

If you specify a clock in the source list, you must specify a clock in the destination list. Applying `set_max_delay` or `set_min_delay` setting between clocks applies the exception from all registers or ports driven by the source clock to all registers or ports driven by the destination clock. Applying exceptions between clocks is more efficient than applying them for specific node-to-node, or node-to-clock paths. If you want to specify pin names in the list, the source must be a clock pin and the destination must be any non-clock input pin to a register. Assignments from clock pins, or to and from cells, apply to all registers in the cell or for those driven by the clock pin.

## set\_false\_path

The false path constraint is specified in the Precision Synthesis software.

### Example 18-8: Using the set\_false\_path Constraint

```
set_false_path -to <to_node_list> -from <from_node_list> -reset_path
```

The node lists can be a list of clocks, ports, instances, and pins. Multiple elements in the list can be represented using wildcards such as \* and ?.

In a place-and-route Tcl constraints file, this false path setting in the Precision Synthesis software is mapped to a `set_false_path` setting. The Quartus II software supports `setup`, `hold`, `rise`, or `fall` options for this assignment.

The node lists for this assignment represents top-level ports and/or nets connected to instances (end points of timing assignments).

Any false path setting in the Precision Synthesis software can be mapped to a setting in the Quartus II software with a `through` path specification.

## set\_multicycle\_path

The multicycle path constraint is specified in the Precision Synthesis software.

### Example 18-9: Using the set\_multicycle\_path Constraint

```
set_multicycle_path <multiplier_value> [-start] [-end] \  
-to <to_node_list> -from <from_node_list> -reset_path
```

The node list can contain clocks, ports, instances, and pins. Multiple elements in the list can be represented using wildcards such as \* and ?. Paths without multicycle path definitions are identical to paths with multipliers of 1. To add one additional cycle to the datapath, use a multiplier value of 2. The option `start` indicates that source clock cycles should be considered for the multiplier. The option `end` indicates that destination clock cycles should be considered for the multiplier. The default is to reference the end clock.

In the place-and-route Tcl constraints file, the multicycle path setting in the Precision Synthesis software is mapped to a `set_multicycle_path` setting. The Quartus II software supports the `rise` or `fall` options on this assignment.

The node lists represent top-level ports and/or nets connected to instances (end points of timing assignments). The node lists can contain wildcards (such as \*); the Quartus II software automatically expands all wildcards.

Any multicycle path setting in Precision Synthesis software can be mapped to a setting in the Quartus II software with a `through` specification.

## Guidelines for Altera IP Cores and Architecture-Specific Features

Altera provides parameterizable IP cores, including the LPMs, device-specific Altera IP cores, and IP available through the Altera Megafunction Partners Program (AMPP<sup>SM</sup>). You can use IP cores by instantiating them in your HDL code or by inferring certain functions from generic HDL code.

If you want to instantiate an IP core such as a PLL in your HDL code, you can instantiate and parameterize the function using the port and parameter definitions, or you can customize a function with the Parameter Editor. Altera recommends using the IP Catalog and Parameter Editor, which provides a graphical interface within the Quartus II software for customizing and parameterizing any available IP core for the design.

The Precision Synthesis software automatically recognizes certain types of HDL code and infers the appropriate IP core.

### Related Information

- [Inferring Altera IP Cores from HDL Code](#) on page 18-16
- [Recommended HDL Coding Styles documentation](#) on page 12-1
- [Introduction to Altera IP Cores documentation](#)

## Instantiating IP Cores With IP Catalog-Generated Verilog HDL Files

The IP Catalog generates a Verilog HDL instantiation template file `<output file>_inst.v` and a hollow-body black box module declaration `<output file>_bb.v` for use in your Precision Synthesis design. Incorporate the instantiation template file, `<output file>_inst.v`, into your top-level design to instantiate the IP core wrapper file, `<output file>.v`.

Include the hollow-body black box module declaration `<output file>_bb.v` in your Precision Synthesis project to describe the port connections of the black box. Adding the IP core wrapper file `<output file>.v` in your Precision Synthesis project is optional, but you must add it to your Quartus II project along with the Precision Synthesis-generated EDIF or VQM netlist.

Alternatively, you can include the IP core wrapper file `<output file>.v` in your Precision Synthesis project and turn on the **Exclude file from Compile Phase** option in the Precision Synthesis software to exclude the file from compilation and to copy the file to the appropriate directory for use by the Quartus II software during place-and-route.

## Instantiating IP Cores With IP Catalog-Generated VHDL Files

The IP Catalog generates a VHDL component declaration file `<output file>.cmp` and a VHDL instantiation template file `<output file>_inst.vhd` for use in your Precision Synthesis design. Incorporate the component declaration and instantiation template into your top-level design to instantiate the IP core wrapper file, `<output file>.vhd`.

Adding the IP core wrapper file `<output file>.vhd` in your Precision Synthesis project is optional, but you must add the file to your Quartus II project along with the Precision Synthesis-generated EDIF or VQM netlist.

Alternatively, you can include the IP core wrapper file `<output file>.v` in your Precision Synthesis project and turn on the **Exclude file from Compile Phase** option in the Precision Synthesis software to exclude the file from compilation and to copy the file to the appropriate directory for use by the Quartus II software during place-and-route.

## Instantiating Intellectual Property With the IP Catalog and Parameter Editor

Many Altera IP functions include a resource and timing estimation netlist that the Precision Synthesis software can use to synthesize and optimize logic around the IP efficiently. As a result, the Precision Synthesis software provides better timing correlation, area estimates, and Quality of Results (QoR) than a black box approach.

To create this netlist file, perform the following steps:

1. Select the IP function in the IP Catalog.
2. Click **Next** to open the Parameter Editor.
3. Click **Set Up Simulation**, which sets up all the EDA options.
4. Turn on the **Generate netlist** option to generate a netlist for resource and timing estimation and click **OK**.
5. Click **Generate** to generate the netlist file.

The Quartus II software generates a file `<output file>_syn.v`. This netlist contains the “grey box” information for resource and timing estimation, but does not contain the actual implementation. Include this netlist file into your Precision Synthesis project as an input file. Then include the IP core wrapper file `<output file>.v|vhd` in the Quartus II project along with your EDIF or VQM output netlist.

The generated “grey box” netlist file, `<output file>_syn.v`, is always in Verilog HDL format, even if you select VHDL as the output file format.

**Note:** For information about creating a grey box netlist file from the command line, search Altera's Knowledge Database.

### Related Information

[Altera Knowledge Center website](#)

## Instantiating Black Box IP Functions With Generated Verilog HDL Files

You can use the `syn_black_box` or `black_box` compiler directives to declare a module as a black box. The top-level design files must contain the IP port mapping and a hollow-body module declaration. You can apply the directive to the module declaration in the top-level file or a separate file included in the project so that the Precision Synthesis software recognizes the module is a black box.

**Note:** The `syn_black_box` and `black_box` directives are supported only on module or entity definitions.

The example below shows a sample top-level file that instantiates `my_verilogIP.v`, which is a simplified customized variation generated by the IP Catalog and Parameter Editor.

### Example 18-10: Top-Level Verilog HDL Code with Black Box Instantiation of IP

```
module top (clk, count);
    input clk;
    output[7:0] count;

    my_verilogIP verilogIP_inst (.clock (clk), .q (count));
endmodule

// Module declaration
// The following attribute is added to create a
// black box for this module.
module my_verilogIP (clock, q) /* synthesis syn_black_box */;
    input clock;
```

```
output[7:0] q;  
endmodule
```

## Instantiating Black Box IP Functions With Generated VHDL Files

You can use the `syn_black_box` or `black_box` compiler directives to declare a component as a black box. The top-level design files must contain the IP core variation component declaration and port mapping. Apply the directive to the component declaration in the top-level file.

**Note:** The `syn_black_box` and `black_box` directives are supported only on module or entity definitions.

The example below shows a sample top-level file that instantiates `my_vhdlIP.vhd`, which is a simplified customized variation generated by the IP Catalog and Parameter Editor.

### Example 18-11: Top-Level VHDL Code with Black Box Instantiation of IP

```
LIBRARY ieee;  
USE ieee.std_logic_1164.all;  
ENTITY top IS  
  PORT (  
    clk: IN STD_LOGIC ;  
    count: OUT STD_LOGIC_VECTOR (7 DOWNTO 0)  
  );  
END top;  
  
ARCHITECTURE rtl OF top IS  
  COMPONENT my_vhdlIP  
  PORT (  
    clock: IN STD_LOGIC ;  
    q: OUT STD_LOGIC_VECTOR (7 DOWNTO 0)  
  );  
  end COMPONENT;  
  attribute syn_black_box : boolean;  
  attribute syn_black_box of my_vhdlIP: component is true;  
  BEGIN  
    vhdlIP_inst : my_vhdlIP PORT MAP (  
      clock => clk,  
      q => count  
    );  
  END rtl;
```

## Inferring Altera IP Cores from HDL Code

The Precision Synthesis software automatically recognizes certain types of HDL code and maps arithmetical operators, relational operators, and memory (RAM and ROM), to technology-specific implementations. This functionality allows technology-specific resources to implement these structures by inferring the appropriate Altera function to provide optimal results. In some cases, the Precision Synthesis software has options that you can use to disable or control inference.

For coding style recommendations and examples for inferring technology-specific architecture in Altera devices, refer to the *Precision Synthesis Style Guide*.

### Related Information

- [Recommended HDL Coding Styles documentation](#) on page 12-1

## Multipliers

The Precision Synthesis software detects multipliers in HDL code and maps them directly to device atoms to implement the multiplier in the appropriate type of logic. The Precision Synthesis software also allows you to control the device resources that are used to implement individual multipliers.

### Controlling DSP Block Inference for Multipliers

By default, the Precision Synthesis software uses DSP blocks available in Stratix series devices to implement multipliers. The default setting is **AUTO**, which allows the Precision Synthesis software to map to logic look-up tables (LUTs) or DSP blocks, depending on the size of the multiplier. You can use the Precision Synthesis GUI or HDL attributes for direct mapping to only logic elements or to only DSP blocks.

**Table 18-3: Options for dedicated\_mult Parameter to Control Multiplier Implementation in Precision Synthesis**

Value	Description
<b>ON</b>	Use only DSP blocks to implement multipliers, regardless of the size of the multiplier.
<b>OFF</b>	Use only logic (LUTs) to implement multipliers, regardless of the size of the multiplier.
<b>AUTO</b>	Use logic (LUTs) or DSP blocks to implement multipliers, depending on the size of the multipliers.

### Setting the Use Dedicated Multiplier Option

To set the `Use Dedicated Multiplier` option in the Precision Synthesis GUI, compile the design, and then in the Design Hierarchy browser, right-click the operator for the desired multiplier and click **Use Dedicated Multiplier**.

### Setting the dedicated\_mult Attribute

To control the implementation of a multiplier in your HDL code, use the `dedicated_mult` attribute with the appropriate value as shown in the examples below.

#### Example 18-12: Setting the dedicated\_mult Attribute in Verilog HDL

```
//synthesis attribute <signal name> dedicated_mult <value>
```

#### Example 18-13: Setting the dedicated\_mult Attribute in VHDL

```
ATTRIBUTE dedicated_mult: STRING;
ATTRIBUTE dedicated_mult OF <signal name>: SIGNAL IS <value>;
```

The `dedicated_mult` attribute can be applied to signals and wires; it does not work when applied to a register. This attribute can be applied only to simple multiplier code, such as `a = b * c`.

Some signals for which the `dedicated_mult` attribute is set can be removed during synthesis by the Precision Synthesis software for design optimization. In such cases, if you want to force the

implementation, you should preserve the signal by setting the `preserve_signal` attribute to `TRUE`.

#### Example 18-14: Setting the `preserve_signal` Attribute in Verilog HDL

```
//synthesis attribute <signal name> preserve_signal TRUE
```

#### Example 18-15: Setting the `preserve_signal` Attribute in VHDL

```
ATTRIBUTE preserve_signal: BOOLEAN;
ATTRIBUTE preserve_signal OF <signal name>: SIGNAL IS TRUE;
```

#### Example 18-16: Verilog HDL Multiplier Implemented in Logic

```
module unsigned_mult (result, a, b);
    output [15:0] result;
    input [7:0] a;
    input [7:0] b;
    assign result = a * b;
    //synthesis attribute result dedicated_mult OFF
endmodule
```

#### Example 18-17: VHDL Multiplier Implemented in Logic

```
LIBRARY ieee;
USE ieee.std_logic_1164.ALL;
USE ieee.std_logic_arith.ALL;
USE ieee.std_logic_unsigned.ALL;

ENTITY unsigned_mult IS
    PORT(
        a: IN std_logic_vector (7 DOWNTO 0);
        b: IN std_logic_vector (7 DOWNTO 0);
        result: OUT std_logic_vector (15 DOWNTO 0));
    ATTRIBUTE dedicated_mult: STRING;
END unsigned_mult;

ARCHITECTURE rtl OF unsigned_mult IS
    SIGNAL a_int, b_int: UNSIGNED (7 downto 0);
    SIGNAL pdt_int: UNSIGNED (15 downto 0);
    ATTRIBUTE dedicated_mult OF pdt_int: SIGNAL IS "OFF";
BEGIN
    a_int <= UNSIGNED (a);
    b_int <= UNSIGNED (b);
    pdt_int <= a_int * b_int;
    result <= std_logic_vector(pdt_int);
END rtl;
```

### Multiplier-Accumulators and Multiplier-Adders

The Precision Synthesis software also allows you to control the device resources used to implement multiply-accumulators or multiply-adders in your project or in a particular module.

The Precision Synthesis software detects multiply-accumulators or multiply-adders in HDL code and infers an ALTMULT\_ACCUM or ALTMULT\_ADD IP cores so that the logic can be placed in DSP blocks, or the software maps these functions directly to device atoms to implement the multiplier in the appropriate type of logic.

**Note:** The Precision Synthesis software supports inference for these functions only if the target device family has dedicated DSP blocks.

For more information about DSP blocks in Altera devices, refer to the appropriate Altera device family handbook and device-specific documentation. For details about which functions a given DSP block can implement, refer to the DSP Solutions Center on the Altera website.

For more information about inferring multiply-accumulator and multiply-adder IP cores in HDL code, refer to the Altera *Recommended HDL Coding Styles* and the Mentor Graphics *Precision Synthesis Style Guide*.

**Related Information**

[Altera DSP Solutions website](#)

[Recommended HDL Coding Styles documentation](#) on page 12-1

**Controlling DSP Block Inference**

By default, the Precision Synthesis software infers the ALTMULT\_ADD or ALTMULT\_ACCUM IP cores appropriately in your design. These IP cores allow the Quartus II software to select either logic or DSP blocks, depending on the device utilization and the size of the function.

You can use the `extract_mac` attribute to prevent inference of an ALTMULT\_ADD or ALTMULT\_ACCUM IP cores in a certain module or entity.

**Table 18-4: Options for `extract_mac` Attribute Controlling DSP Implementation**

Value	Description
TRUE	The ALTMULT_ADD or ALTMULT_ACCUM IP core is inferred.
FALSE	The ALTMULT_ADD or ALTMULT_ACCUM IP core is not inferred.

To control inference, use the `extract_mac` attribute with the appropriate value from the examples below in your HDL code.

**Example 18-18: Setting the `extract_mac` Attribute in Verilog HDL**

```
//synthesis attribute <module name> extract_mac <value>
```

**Example 18-19: Setting the `extract_mac` Attribute in VHDL**

```
ATTRIBUTE extract_mac: BOOLEAN;
ATTRIBUTE extract_mac OF <entity name>: ENTITY IS <value>;
```



To control the implementation of the multiplier portion of a multiply-accumulator or multiply-adder, you must use the `dedicated_mult` attribute.

You can use the `extract_mac`, `dedicated_mult`, and `preserve_signal` attributes (in Verilog HDL and VHDL) to implement the given DSP function in logic in the Quartus II software.

### Example 18-20: Using `extract_mac`, `dedicated_mult`, and `preserve_signal` in Verilog HDL

```
module unsig_altmult_accuml (dataout, dataa, datab, clk, aclr, clken);
  input [7:0] dataa, datab;
  input clk, aclr, clken;
  output [31:0] dataout;

  reg    [31:0] dataout;
  wire   [15:0] multa;
  wire   [31:0] adder_out;

  assign multa = dataa * datab;

  //synthesis attribute multa preserve_signal TRUE
  //synthesis attribute multa dedicated_mult OFF
  assign adder_out = multa + dataout;

  always @ (posedge clk or posedge aclr)
  begin
    if (aclr)
      dataout <= 0;
    else if (clken)
      dataout <= adder_out;
  end

  //synthesis attribute unsig_altmult_accuml extract_mac FALSE
endmodule
```

### Example 18-21: Using `extract_mac`, `dedicated_mult`, and `preserve_signal` in VHDL

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;
USE ieee.std_logic_arith.all;
USE ieee.std_logic_signed.all;
ENTITY signedmult_add IS
  PORT(
    a, b, c, d: IN STD_LOGIC_VECTOR (7 DOWNTO 0);
    result: OUT STD_LOGIC_VECTOR (15 DOWNTO 0));
  ATTRIBUTE preserve_signal: BOOLEAN;
  ATTRIBUTE dedicated_mult: STRING;
  ATTRIBUTE extract_mac: BOOLEAN;
  ATTRIBUTE extract_mac OF signedmult_add: ENTITY IS FALSE;
END signedmult_add;
ARCHITECTURE rtl OF signedmult_add IS
  SIGNAL a_int, b_int, c_int, d_int : signed (7 DOWNTO 0);
  SIGNAL pdt_int, pdt2_int : signed (15 DOWNTO 0);
  SIGNAL result_int: signed (15 DOWNTO 0);
  ATTRIBUTE preserve_signal OF pdt_int: SIGNAL IS TRUE;
  ATTRIBUTE dedicated_mult OF pdt_int: SIGNAL IS "OFF";
  ATTRIBUTE preserve_signal OF pdt2_int: SIGNAL IS TRUE;
  ATTRIBUTE dedicated_mult OF pdt2_int: SIGNAL IS "OFF";
```

```
BEGIN
  a_int <= signed (a);
  b_int <= signed (b);
  c_int <= signed (c);
  d_int <= signed (d);
  pdt_int <= a_int * b_int;
  pdt2_int <= c_int * d_int;
  result_int <= pdt_int + pdt2_int;
  result <= STD_LOGIC_VECTOR(result_int);
END rtl;
```

## RAM and ROM

The Precision Synthesis software detects memory structures in HDL code and converts them to an operator that infers an ALTSYNCRAM or LPM\_RAM\_DP IP cores, depending on the device family. The software then places these functions in memory blocks.

The software supports inference for these functions only if the target device family has dedicated memory blocks.

For more information about inferring RAM and ROM IP cores in HDL code, refer to the *Precision Synthesis Style Guide*.

### Related Information

- [Recommended HDL Coding Styles documentation](#) on page 12-1

## Incremental Compilation and Block-Based Design

As designs become more complex and designers work in teams, a block-based incremental design flow is often an effective design approach. In an incremental compilation flow, you can make changes to one part of the design while maintaining the placement and performance of unchanged parts of the design. Design iterations can be made dramatically faster by focusing new compilations on particular design partitions and merging results with the results of previous compilations of other partitions. You can perform optimization on individual blocks and then integrate them into a final design and optimize the design at the top-level.

The first step in an incremental design flow is to make sure that different parts of your design do not affect each other. You must ensure that you have separate netlists for each partition in your design. If the whole design is in one netlist file, changes in one partition affect other partitions because of possible node name changes when you resynthesize the design.

You can create different implementations for each partition in your Precision Synthesis project, which allows you to switch between partitions without leaving the current project file. You can also create a separate project for each partition if you require separate projects for a team-based design flow. Alternatively, you can use the incremental synthesis capability in the Precision RTL Plus software.

### Related Information

- [Quartus II Incremental Compilation for Hierarchical and Team-Based Design documentation](#) on page 3-1

## Creating a Design with Precision RTL Plus Incremental Synthesis

The Precision RTL Plus incremental synthesis flow for Quartus II incremental compilation uses a partition-based approach to achieve faster design cycle time.

Using the incremental synthesis feature, you can create different netlist files for different partitions of a design hierarchy within one partition implementation, which makes each partition independent of the others in an incremental compilation flow. Only the portions of a design that have been updated must be recompiled during design iterations. You can make changes and resynthesize one partition in a design to create a new netlist without affecting the synthesis results or fitting of other partitions.

The following steps show a general flow for partition-based incremental synthesis with Quartus II incremental compilation:

1. Create Verilog HDL or VHDL design files.
2. Determine which hierarchical blocks you want to treat as separate partitions in your design, and designate the partitions with the `incr_partition` attribute.
3. Create a project in the Precision RTL Plus Synthesis software and add the HDL design files to the project.
4. Enable incremental synthesis in the Precision RTL Plus Synthesis software using one of these methods:
  - Use the Precision RTL Plus Synthesis GUI to turn on **Enable Incremental Synthesis**.
  - Run the following command in the Transcript Window:

```
setup_design -enable_incr_synth
```

5. Run the basic Precision Synthesis flow of compilation, synthesis, and place-and-route on your design. In subsequent runs, the Precision RTL Plus Synthesis software processes only the parts of the design that have changed, resulting in a shorter iteration than the initial run. The performance of the unchanged partitions is preserved.

The Precision RTL Plus Synthesis software sets the netlist types of the unchanged partitions to **Post Fit** and the changed partitions to **Post Synthesis**. You can change the netlist type during timing closure in the Quartus II software to obtain the best QoR.

6. Import the EDIF or VQM netlist for each partition and the top-level `.tcl` file into the Quartus II software, and set up the Quartus II project to use incremental compilation.
7. Compile your Quartus II project.
8. If you want, you can change the Quartus II incremental compilation netlist type for a partition with the **Design Partitions Window**. You can change the **Netlist Type** to one of the following options:
  - To preserve the previous post-fit placement results, change the **Netlist Type** of the partition to **Post-Fit**.
  - To preserve the previous routing results, set the **Fitter Preservation Level** of the partition to **Placement and Routing**.

### Creating Partitions with the `incr_partition` Attribute

Partitions are set using the HDL `incr_partition` attribute. The Precision Synthesis software creates or deletes partitions by reading this attribute during compilation iterations. The attribute can be attached to either the design unit definition or an instance.

To delete partitions, you can remove the attribute or set the attribute value to false.

**Note:** The Precision Synthesis software ignores partitions set in a black box.

### Example 18-22: Using incr\_partition Attribute to Create a Partition in Verilog HDL

```
Design unit partition:

module my_block(
    input clk;
    output reg [31:0] data_out) /* synthesis incr_partition */ ;

Instance partition:

my_block my_block_inst(.clk(clk), .data_out(data_out));
// synthesis attribute my_block_inst incr_partition true
```

### Example 18-23: Using incr\_partition Attribute to a Create Partition in VHDL

```
Design unit partition:

entity my_block is
    port(
        clk : in std_logic;
        data_out : out std_logic_vector(31 downto 0)
    );
    attribute incr_partition : boolean;
    attribute incr_partition of my_block : entity is true;
end entity my_block;

Instance partition:

component my_block is
    port(
        clk : in std_logic;
        data_out : out std_logic_vector(31 downto 0)
    );
end component;

attribute incr_partition : boolean;
attribute incr_partition of my_block_inst : label is true;

my_block_inst my_block
    port map(clk, data_out);
```

## Creating Multiple Mapped Netlist Files With Separate Precision Projects or Implementations

You can manually generate multiple netlist files, which can be VQM or EDIF files, for incremental compilation using black boxes and separate Precision projects or implementations for each design partition. This manual flow is supported in versions of the Precision software that do not include the incremental synthesis feature. You might also use this feature if you perform synthesis in a team-based environment without a top-level synthesis project that includes all of the lower-level design blocks.

In the Precision Synthesis software, create a separate implementation, or a separate project, for each lower-level module and for the top-level design that you want to maintain as a separate netlist file. Implement black box instantiations of lower-level modules in your top-level implementation or project.

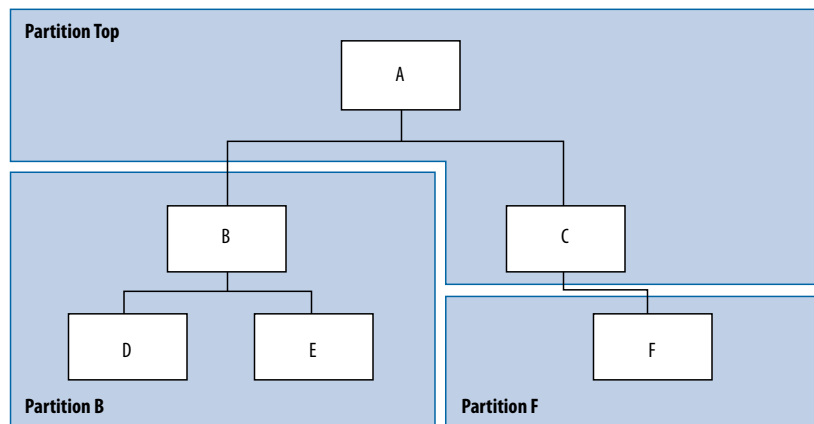
For more information about managing implementations and projects, refer to the *Precision RTL Synthesis User's Manual*.

**Note:** In a standard Quartus II incremental compilation flow, Precision Synthesis software constraints made on lower-level modules are not passed to the Quartus II software. Ensure that appropriate constraints are made in the top-level Precision Synthesis project, or in the Quartus II project.

## Creating Black Boxes to Create Netlists

In the figure below, the top-level partition contains the top-level block in the design (block A) and the logic that is not defined as part of another partition. In this example, the partition for top-level block A also includes the logic in the sub-block C. Because block F is contained in its own partition, it is not treated as part of the top-level partition A. Another separate partition, B, contains the logic in blocks B, D, and E. In a team-based design, different engineers may work on the logic in different partitions. One netlist is created for the top-level module A and its submodule C, another netlist is created for module B and its submodules D and E, while a third netlist is created for module F.

**Figure 18-2: Partitions in a Hierarchical Design**



To create multiple EDIF netlist files for this design, follow these steps:

1. Generate a netlist file for module B. Use **B.v/.vhd**, **D.v/.vhd**, and **E.v/.vhd** as the source files.
2. Generate a netlist file for module F. Use **F.v/.vhd** as the source file.
3. Generate a top-level netlist file for module A. Use **A.v/.vhd** and **C.v/.vhd** as the source files. Ensure that you create black boxes for modules B and F, which were optimized separately in the previous steps.

The goal is to individually synthesize and generate a netlist file for each lower-level module and then instantiate these modules as black boxes in the top-level file. You can then synthesize the top-level file to generate the netlist file for the top-level design. Finally, both the lower-level and top-level netlist files are provided to your Quartus II project.

**Note:** When you make design or synthesis optimization changes to part of your design, resynthesize only the changed partition to generate the new netlist file. Do not resynthesize the implementations or projects for the unchanged partitions.

## Creating Black Boxes in Verilog HDL

Any design block that is not defined in the project or included in the list of files to be read for a project is treated as a black box by the software. In Verilog HDL, you must provide an empty module declaration for any module that is treated as a black box.

A black box for the top-level file **A.v** is shown in the following example. Provide an empty module declaration for any lower-level files, which also contain a black box for any module beneath the current level of hierarchy.

#### Example 18-24: Verilog HDL Black Box for Top-Level File A.v

```
module A (data_in, clk, e, ld, data_out);
    input data_in, clk, e, ld;
    output [15:0] data_out;
    wire [15:0] cnt_out;
    B U1 (.data_in (data_in), .clk(clk), .ld (ld), .data_out(cnt_out));
    F U2 (.d(cnt_out), .clk(clk), .e(e), .q(data_out));
    // Any other code in A.v goes here.
endmodule
//Empty Module Declarations of Sub-Blocks B and F follow here.
// These module declarations (including ports) are required for black
boxes.
module B (data_in, clk, ld, data_out);
    input data_in, clk, ld;
    output [15:0] data_out;
endmodule
module F (d, clk, e, q);
    input [15:0] d;
    input clk, e;
    output [15:0] q;
endmodule
```

### Creating Black Boxes in VHDL

Any design block that is not defined in the project or included in the list of files to be read for a project is treated as a black box by the software. In VHDL, you must provide a component declaration for the black box.

A black box for the top-level file **A.vhd** is shown in the example below. Provide a component declaration for any lower-level files that also contain a black box or for any block beneath the current level of hierarchy.

#### Example 18-25: VHDL Black Box for Top-Level File A.vhd

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;
ENTITY A IS
    PORT ( data_in : IN INTEGER RANGE 0 TO 15;
          clk, e, ld : IN STD_LOGIC;
          data_out : OUT INTEGER RANGE 0 TO 15);
END A;
ARCHITECTURE a_arch OF A IS
    COMPONENT B PORT(
        data_in : IN INTEGER RANGE 0 TO 15;
        clk, ld : IN STD_LOGIC;
        d_out : OUT INTEGER RANGE 0 TO 15);
    END COMPONENT;
    COMPONENT F PORT(
        d : IN INTEGER RANGE 0 TO 15;
        clk, e: IN STD_LOGIC;
        q : OUT INTEGER RANGE 0 TO 15);
    END COMPONENT;
```

```
-- Other component declarations in A.vhd go here
signal cnt_out : INTEGER RANGE 0 TO 15;
BEGIN
  U1 : B
  PORT MAP (
    data_in => data_in,
    clk => clk,
    ld => ld,
    d_out => cnt_out);
  U2 : F
  PORT MAP (
    d => cnt_out,
    clk => clk,
    e => e,
    q => data_out);
  -- Any other code in A.vhd goes here
END a_arch;
```

After you complete the steps outlined above, you have different netlist files for each partition of the design. These files are ready for use with incremental compilation in the Quartus II software.

## Creating Quartus II Projects for Multiple Netlist Files

The Precision Synthesis software creates a **.tcl** file for each implementation, and provides the Quartus II software with the appropriate constraints and information to set up a project. When using incremental synthesis, the Precision RTL Plus Synthesis software creates only a single **.tcl** file, *<project name>\_incr\_partitions.tcl*, to pass the partition information to the Quartus II software.

Depending on your design methodology, you can create one Quartus II project for all netlists, or a separate Quartus II project for each netlist. In the standard incremental compilation design flow, you create design partition assignments for each partition in the design within a single Quartus II project. This methodology provides the best QoR and performance preservation during incremental changes to your design. You might require a bottom-up design flow if each partition must be optimized separately, such as for third-party IP delivery.

To follow this design flow in the Quartus II software, create separate Quartus II projects and export each design partition and incorporate it into a top-level design using the incremental compilation features to maintain placement results.

### Related Information

[Running the Quartus II Software Manually Using the Precision Synthesis-Generated Tcl Script](#) on page 18-10

## Creating a Single Quartus II Project for a Standard Incremental Compilation Flow

Use the *<top-level project>.tcl* file generated for the top-level partition to create your Quartus II project and import all the netlists into this one Quartus II project for an incremental compilation flow. You can optimize all partitions within the single Quartus II project and take advantage of the performance preservation and compilation time reduction that incremental compilation provides.

All the constraints from the top-level implementation are passed to the Quartus II software in the top-level **.tcl** file, but any constraints made only in the lower-level implementations within the Precision Synthesis software are not forward-annotated. Enter these constraints manually in your Quartus II project.

## Creating Multiple Quartus II Projects for a Bottom-Up Flow

Use the .tcl files generated by the Precision Synthesis software for each Precision Synthesis software implementation or project to generate multiple Quartus II projects, one for each partition in the design. Each designer in the project can optimize their block separately in the Quartus II software and export the placement of their blocks using incremental compilation. Designers should create a LogicLock region to provide a floorplan location assignment for each block; the top-level designer should then import all the blocks and assignments into the top-level project.

## Hierarchy and Design Considerations

To ensure the proper functioning of the synthesis flow, you can create separate partitions only for modules, entities, or existing netlist files. In addition, each module or entity must have its own design file. If two different modules are in the same design file, but are defined as being part of different partitions, incremental synthesis cannot be maintained because both regions must be recompiled when you change one of the modules.

Altera recommends that you register all inputs and outputs of each partition. This makes logic synchronous and avoids any delay penalty on signals that cross partition boundaries.

If you use boundary tri-states in a lower-level block, the Precision Synthesis software pushes the tri-states through the hierarchy to the top-level to make use of the tri-state drivers on output pins of Altera devices. Because pushing tri-states requires optimizing through hierarchies, lower-level tri-states are not supported with a block-based compilation methodology. You should use tri-state drivers only at the external output pins of the device and in the top-level block in the hierarchy.

### Related Information

- [Best Practices for Incremental Compilation Partitions and Floorplan Assignments documentation](#) on page 14-1

## Document Revision History

Table 18-5: Document Revision History

Date	Version	Changes
June 2014	14.0.0	<ul style="list-style-type: none"> <li>• Dita conversion.</li> <li>• Removed obsolete devices.</li> <li>• Replaced Megafunction, MegaWizard, and IP Toolbench content with IP Catalog and Parameter Editor content.</li> </ul>
June 2012	12.0.0	<ul style="list-style-type: none"> <li>• Removed survey link.</li> </ul>
November 2011	10.1.1	<ul style="list-style-type: none"> <li>• Template update.</li> <li>• Minor editorial changes.</li> </ul>



Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> <li>• Changed to new document template.</li> <li>• Removed Classic Timing Analyzer support.</li> <li>• Added support for .vqm netlist files.</li> <li>• Edited the “Creating Quartus II Projects for Multiple EDIF Files” on page 15–30 section for changes with the incremental compilation flow.</li> <li>• Editorial changes.</li> </ul>
July 2010	10.0.0	<ul style="list-style-type: none"> <li>• Minor updates for the Quartus II software version 10.0 release</li> </ul>
November 2009	9.1.0	<ul style="list-style-type: none"> <li>• Minor updates for the Quartus II software version 9.1 release</li> </ul>
March 2009	9.0.0	<ul style="list-style-type: none"> <li>• Updated list of supported devices for the Quartus II software version 9.0 release</li> <li>• Chapter 11 was previously Chapter 10 in software version 8.1</li> </ul>

Date	Version	Changes
November 2008	8.1.0	<ul style="list-style-type: none"> <li>• Changed to 8-1/2 x 11 page size</li> <li>• Title changed to <i>Mentor Graphics Precision Synthesis Support</i></li> <li>• Updated list of supported devices</li> <li>• Added information about the Precision RTL Plus incremental synthesis flow</li> <li>• Updated Figure 10-1 to include SystemVerilog</li> <li>• Updated “Guidelines for Altera Megafunctions and Architecture-Specific Features” on page 10–19</li> <li>• Updated “Incremental Compilation and Block-Based Design” on page 10–28</li> <li>• Added section “Creating Partitions with the incr_partition Attribute” on page 10–29</li> </ul>
May 2008	8.0.0	<ul style="list-style-type: none"> <li>• Removed Mercury from the list of supported devices</li> <li>• Changed Precision version to 2007a update 3</li> <li>• Added note for Stratix IV support</li> <li>• Renamed “Creating a Project and Compiling the Design” section to “Creating and Compiling a Project in the Precision RTL Synthesis Software”</li> <li>• Added information about constraints in the Tcl file</li> <li>• Updated document based on the Quartus II software version 8.0</li> </ul>

For previous versions of the *Quartus II Handbook*, refer to the Quartus II Handbook Archive.

**Related Information**

[Quartus II Handbook Archive](#)

# Analyzing Designs with Quartus II Netlist Viewers 19

2014.06.30

QII5V1



Subscribe



Send Feedback

This chapter describes how you can use the Quartus® II Netlist Viewers to analyze and debug your designs.

As FPGA designs grow in size and complexity, the ability to analyze, debug, optimize, and constrain your design is critical. With today's advanced designs, several design engineers are involved in coding and synthesizing different design blocks, making it difficult to analyze and debug the design. The Quartus II RTL Viewer, State Machine Viewer, and Technology Map Viewer provide powerful ways to view your initial and fully mapped synthesis results during the debugging, optimization, and constraint entry processes.

This chapter contains the following sections:

## Related Information

- [When to Use the Netlist Viewers: Analyzing Design Problems](#) on page 19-1
- [Introduction to the User Interface](#) on page 19-5
- [Quartus II Design Flow with the Netlist Viewers](#) on page 19-2
- [State Machine Viewer Overview](#) on page 19-4
- [RTL Viewer Overview](#) on page 19-3
- [Technology Map Viewer Overview](#) on page 19-5
- [Filtering in the Schematic View](#) on page 19-16
- [Probing to a Source Design File and Other Quartus II Windows](#) on page 19-22
- [Probing to the Netlist Viewers from Other Quartus II Windows](#) on page 19-22
- [Viewing a Timing Path](#) on page 19-23

## When to Use the Netlist Viewers: Analyzing Design Problems

You can use the Netlist Viewers to analyze and debug your design. This section provides simple examples of how to use the RTL Viewer, State Machine Viewer, and Technology Map Viewer to analyze problems encountered in the design process.

Using the RTL Viewer is a good way to view your initial synthesis results to determine whether you have created the necessary logic, and that the logic and connections have been interpreted correctly by the software. You can use the RTL Viewer and State Machine Viewer to check your design visually before simulation or other verification processes. Catching design errors at this early stage of the design process can save you valuable time.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



If you see unexpected behavior during verification, use the RTL Viewer to trace through the netlist and ensure that the connections and logic in your design are as expected. You can also view state machine transitions and transition equations with the State Machine Viewer. Viewing your design helps you find and analyze the source of design problems. If your design looks correct in the RTL Viewer, you know to focus your analysis on later stages of the design process and investigate potential timing violations or issues in the verification flow itself.

You can use the Technology Map Viewer to look at the results at the end of Analysis and Synthesis. If you have compiled your design through the Fitter stage, you can view your post-mapping netlist in the Technology Map Viewer (Post-Mapping) and your post-fitting netlist in the Technology Map Viewer. If you perform only Analysis and Synthesis, both the Netlist Viewers display the same post-mapping netlist.

In addition, you can use the RTL Viewer or Technology Map Viewer to locate the source of a particular signal, which can help you debug your design. Use the navigation techniques described in this chapter to search easily through your design. You can trace back from a point of interest to find the source of the signal and ensure the connections are as expected.

The Technology Map Viewer can help you locate post-synthesis nodes in your netlist and make assignments when optimizing your design. This functionality is useful when making a multicycle clock timing assignment between two registers in your design. Start at an I/O port and trace forward or backward through the design and through levels of hierarchy to find nodes of interest, or locate a specific register by visually inspecting the schematic.

You can use the RTL Viewer, State Machine Viewer, and Technology Map Viewer in many other ways throughout the design, debug, and optimization stages. This chapter shows you how to use the various features of the Netlist Viewers to increase your productivity when analyzing a design.

#### Related Information

- [Quartus II Design Flow with the Netlist Viewers](#) on page 19-2
- [State Machine Viewer Overview](#) on page 19-4
- [RTL Viewer Overview](#) on page 19-3
- [Technology Map Viewer Overview](#) on page 19-5

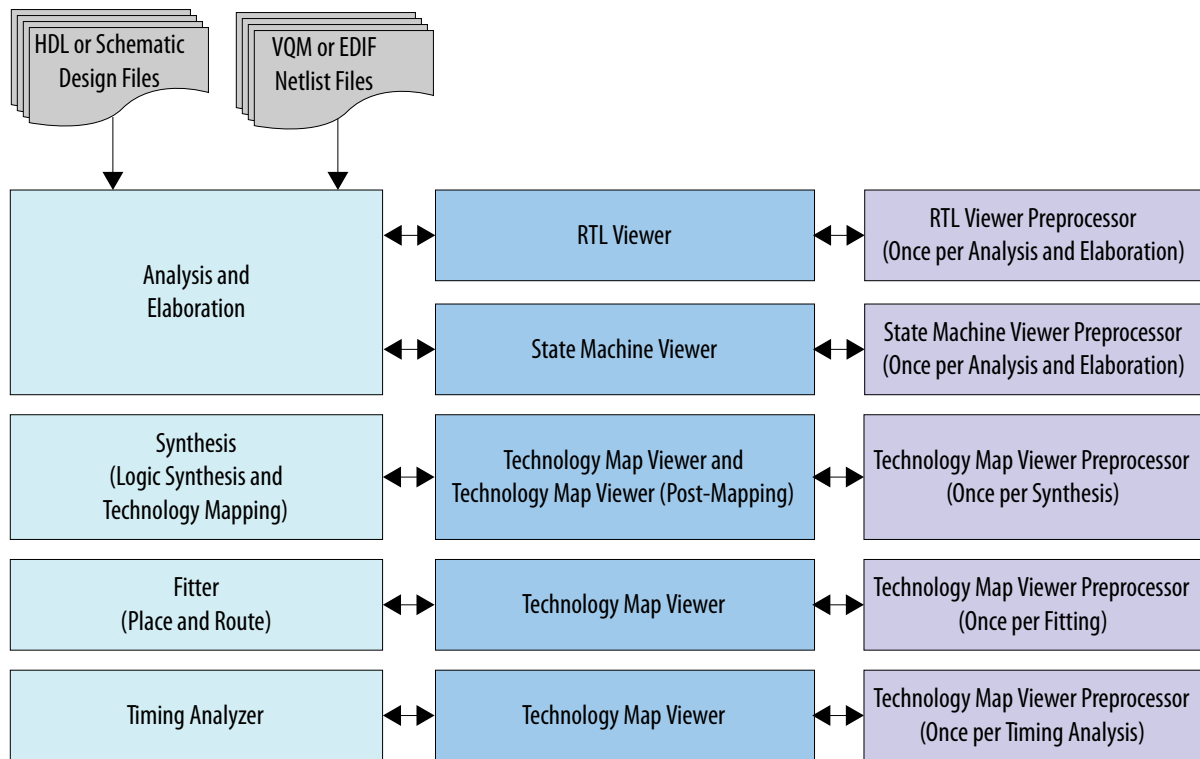
## Quartus II Design Flow with the Netlist Viewers

When you first open one of the Netlist Viewers after compiling the design, a preprocessor stage runs automatically before the Netlist Viewer opens.

The preprocessor process box contains a link to the **Settings > Compilation Process Settings** page where you can turn on the **Run Netlist Viewers preprocessing during compilation** option. When this option is turned on, the preprocessing becomes part of the full project compilation flow and the Netlist Viewer opens immediately without displaying the preprocessing dialog.

**Figure 19-1: Quartus II Design Flow Including the RTL Viewer and Technology Map Viewer**

This figure shows how Netlist Viewers fit into the basic Quartus II design flow.



Before the Netlist Viewer can run the preprocessor stage, you must compile your design:

- To open the RTL Viewer or State Machine Viewer, first perform Analysis and Elaboration.
- To open the Technology Map Viewer (Post-Fitting) or the Technology Map Viewer (Post-Mapping), first perform Analysis and Synthesis.

The Netlist Viewers display the results of the last successful compilation. Therefore, if you make a design change that causes an error during Analysis and Elaboration, you cannot view the netlist for the new design files, but you can still see the results from the last successfully compiled version of the design files. If you receive an error during compilation and you have not yet successfully run the appropriate compilation stage for your project, the Netlist Viewer cannot be displayed; in this case, the Quartus II software issues an error message when you try to open the Netlist Viewer.

**Note:** If the Netlist Viewer is open when you start a new compilation, the Netlist Viewer closes automatically. You must open the Netlist Viewer again to view the new design netlist after compilation completes successfully.

## RTL Viewer Overview

The Quartus II RTL Viewer allows you to view a register transfer level (RTL) graphical representation of your Quartus II integrated synthesis results or your third-party netlist file in the Quartus II software.

You can view results after Analysis and Elaboration when your design uses any supported Quartus II design entry method, including Verilog HDL Design Files (**.v**), SystemVerilog Design Files (**.sv**), VHDL Design Files (**.vhd**), AHDL Text Design Files (**.tdf**), or schematic Block Design Files (**.bdf**). You can also view the hierarchy of atom primitives (such as device logic cells and I/O ports) when your design uses a synthesis tool to generate a Verilog Quartus Mapping File (**.vqm**) or Electronic Design Interchange Format (**.edf**) file.

The Quartus II RTL Viewer displays a schematic view of the design netlist after Analysis and Elaboration or netlist extraction is performed by the Quartus II software, but before technology mapping and any synthesis or fitter optimizations. This view is not the final design structure because optimizations have not yet occurred. This view most closely represents your original source design. If you synthesized your design with the Quartus II integrated synthesis, this view shows how the Quartus II software interpreted your design files. If you use a third-party synthesis tool, this view shows the netlist written by your synthesis tool.

When displaying your design, the RTL Viewer optimizes the netlist to maximize readability in the following ways:

- Logic with no fan-out (its outputs are unconnected) and logic with no fan-in (its inputs are unconnected) are removed from the display.
- Default connections such as  $V_{CC}$  and GND are not shown.
- Pins, nets, wires, module ports, and certain logic are grouped into buses where appropriate.
- Constant bus connections are grouped.
- Values are displayed in hexadecimal format.
- NOT gates are converted to bubble inversion symbols in the schematic.
- Chains of equivalent combinational gates are merged into a single gate. For example, a 2-input AND gate feeding a 2-input AND gate is converted to a single 3-input AND gate.
- State machine logic is converted into a state diagram, state transition table, and state encoding table, which are displayed in the State Machine Viewer.

To run the RTL Viewer for a Quartus II project, first analyze the design to generate an RTL netlist. To analyze the design and generate an RTL netlist, on the Processing menu, point to **Start** and click **Start Analysis & Elaboration**. You can also perform a full compilation on any process that includes the initial Analysis and Elaboration stage of the Quartus II compilation flow.

To run the RTL Viewer, on the Tools menu, point to **Netlist Viewers** and click **RTL Viewer**.

## State Machine Viewer Overview

The State Machine Viewer presents a high-level view of finite state machines in your design. The State Machine Viewer provides a graphical representation of the states and their related transitions, as well as a state transition table that displays the condition equation for each of the state transitions, and encoding information for each state.

To run the State Machine Viewer, on the Tools menu, point to **Netlist Viewers** and click **State Machine Viewer**. To open the State Machine Viewer for a particular state machine, double-click the state machine instance in the RTL Viewer.

### Related Information

[State Machine Viewer](#) on page 19-20

## Technology Map Viewer Overview

The Quartus II Technology Map Viewer provides a technology-specific, graphical representation of your design after Analysis and Synthesis or after the Fitter has mapped your design into the target device.

The Technology Map Viewer shows the hierarchy of atom primitives (such as device logic cells and I/O ports) in your design. For supported families, you can also view internal registers and look-up tables (LUTs) inside logic cells (LCELLs) and registers in I/O atom primitives.

Where possible, the port names of each hierarchy are maintained throughout synthesis; however, port names might change or be removed from the design. For example, if a port is unconnected or driven by GND or  $V_{CC}$ , it is removed during synthesis. When a port name changes, the port is assigned a related user logic name in the design or a generic port name such as `IN1` or `OUT1`.

You can view your Quartus II technology-mapped results after synthesis, fitting, or timing analysis. To run the Technology Map Viewer for a Quartus II project, on the Processing menu, point to **Start** and click **Start Analysis & Synthesis** to synthesize and map the design to the target technology. At this stage, the Technology Map Viewer shows the same post-mapping netlist as the Technology Map Viewer (Post-Mapping). You can also perform a full compilation, or any process that includes the synthesis stage in the compilation flow.

If you have completed the Fitter stage, the Technology Map Viewer shows the changes made to your netlist by the Fitter, such as physical synthesis optimizations, while the Technology Map Viewer (Post-Mapping) shows the post-mapping netlist. If you have completed the Timing Analysis stage, you can locate timing paths from the Timing Analyzer report in the Technology Map Viewer.

To open the Technology Map Viewer, on the Tools menu, point to **Netlist Viewers** and click **Technology Map Viewer (Post-Fitting)** or **Technology Map Viewer (Post Mapping)**.

### Related Information

- [View Contents of Nodes in the Schematic View](#) on page 19-16
- [Viewing a Timing Path](#) on page 19-23

## Introduction to the User Interface

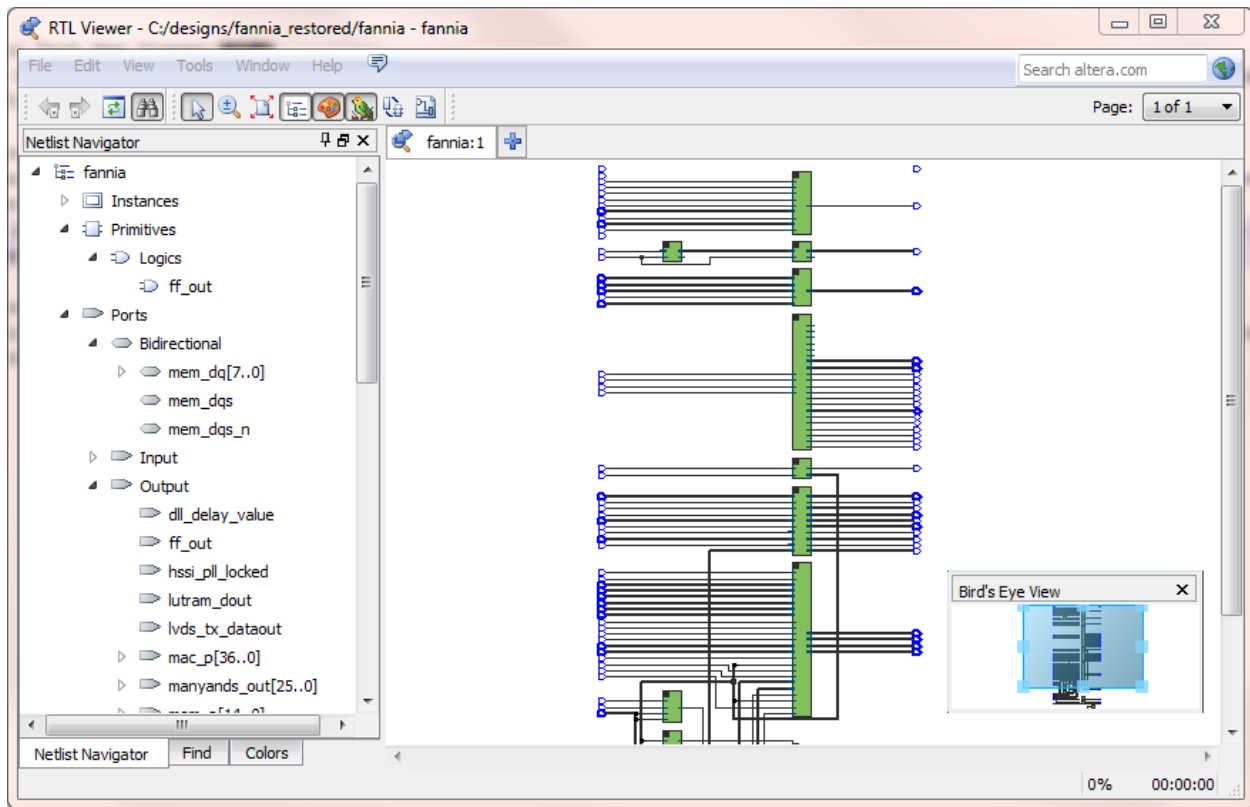
The Netlist Viewer is a graphical user-interface for viewing and manipulating nodes and nets in the netlist.

The RTL Viewer and Technology Map Viewer each consist of these main parts:

- The **Netlist Navigator** pane—displays a representation of the project hierarchy.
- The **Find** pane—allows you to find and locate specific design elements in the schematic view.
- The **Properties** pane displays the properties of the selected block when you select **Properties** from the shortcut menu.
- The schematic view—displays a graphical representation of the internal structure of your design.

Figure 19-2: RTL Viewer

This figure shows the schematic view and the **Netlist Navigator** pane of the RTL Viewer.



Netlist Viewers also contain a toolbar that provides tools to use in the schematic view.

- Use the **Back** and **Forward** buttons to switch between schematic views. You can go forward only if you have not made any changes to the view since going back. These commands do not undo an action, such as selecting a node. The Netlist Viewer caches up to ten actions including filtering, hierarchy navigation, netlist navigation, and zooming.
- The **Refresh** button restores the schematic view and optimizes the layout. **Refresh** does not reload the database if you change your design and recompile.
- The **Find** button opens and closes the **Find** pane.
- The **Selection tool** and **Zoom tool** buttons toggle between the selection mode and zoom mode.
- The **Fit in Page** button resets the schematic view to encompass the entire design.
- The **Netlist Navigator** button opens or closes the **Netlist Navigator** pane.
- The **Color Settings** button opens the **Colors** pane where you can customize the color scheme used in the Netlist Viewer.



- The **Bird's Eye View** button opens the **Bird's Eye View** window which displays a miniature version of your design and allows you to navigate within the design and adjust the magnification in the schematic view quickly.
- The **Show/Hide Instance Pins** button can toggle the display of instance pins not displayed by functions such as cross-probing between a Netlist Viewer and TimeQuest. You can also use it to hide unconnected instance pins when filtering a node results in large numbers of unconnected or unused pins.
- The **Show Netlist on One Page** button displays the netlist on a single page if the Netlist Viewer has split the design across several pages. This can make netlist tracing easier.

You can have only one RTL Viewer, one Technology Map Viewer (Post-Fitting), one Technology Map Viewer (Post-Mapping), and one State Machine Viewer window open at the same time, although each window can show multiple pages, each with multiple tabs. For example, you cannot have two RTL Viewer windows open at the same time.

**Related Information**

- [Netlist Navigator Pane](#) on page 19-7
- [Netlist Viewers Find Pane](#) on page 19-10
- [Properties Pane](#) on page 19-8

## Netlist Navigator Pane

The **Netlist Navigator** pane displays the entire netlist in a tree format based on the hierarchical levels of the design. In each level, similar elements are grouped into subcategories.

You can use the **Netlist Navigator** pane to traverse through the design hierarchy to view the logic schematic for each level. You can also select an element in the **Netlist Navigator** to highlight in the schematic view.

**Note:** Nodes inside atom primitives are not listed in the **Netlist Navigator** pane.

For each module in the design hierarchy, the **Netlist Navigator** pane displays the applicable elements listed in the following table. Click the “+” icon to expand an element.

**Table 19-1: Netlist Navigator Pane Elements**

Elements	Description
Instances	Modules or instances in the design that can be expanded to lower hierarchy levels.
State Machines	State machine instances in the design that can be viewed in the State Machine Viewer.

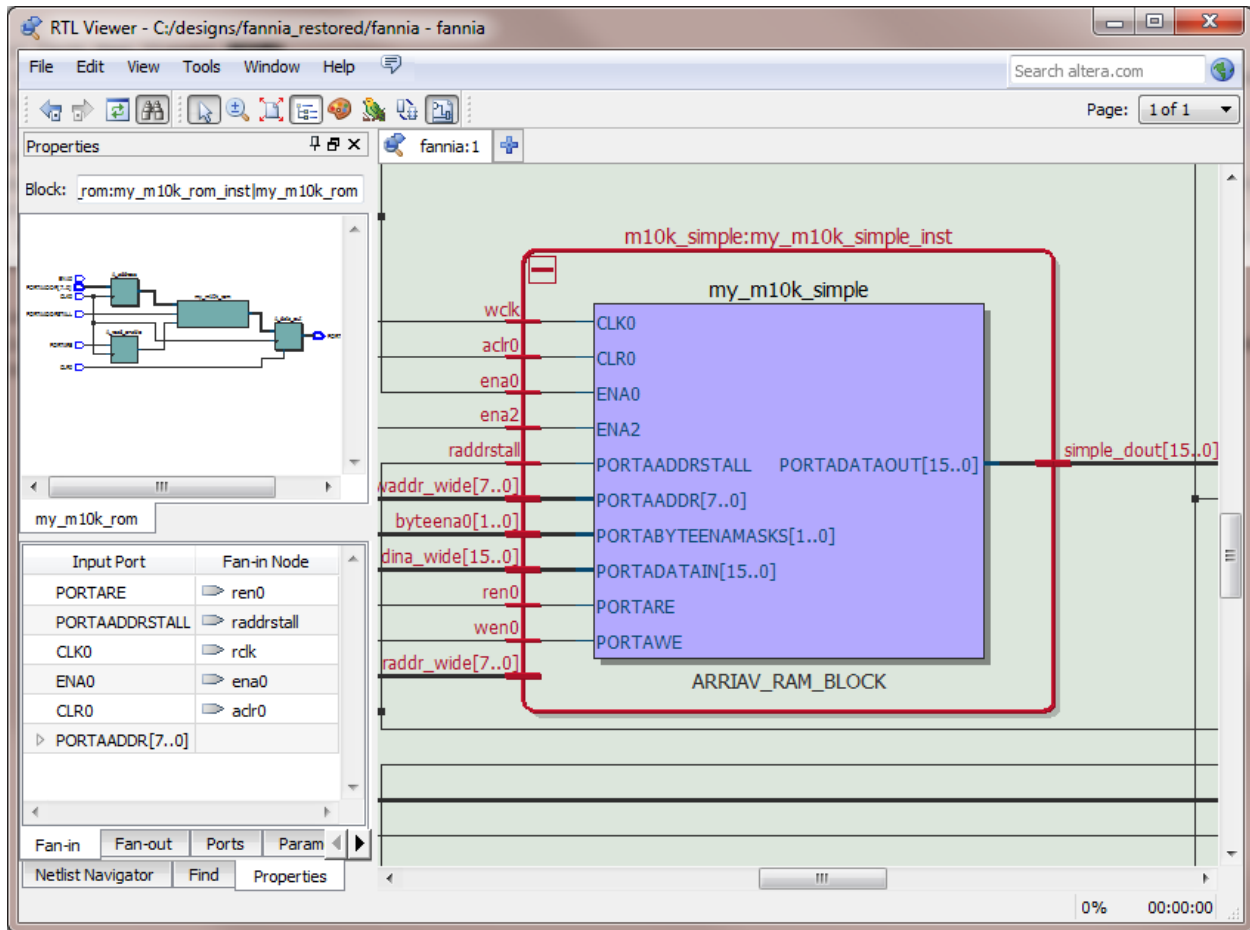
Elements	Description
Primitives	<p>Low-level nodes that cannot be expanded to any lower hierarchy level. These primitives include:</p> <ul style="list-style-type: none"> <li>• Registers and gates that you can view in the RTL Viewer when using Quartus II integrated synthesis</li> <li>• Logic cell atoms in the Technology Map Viewer or in the RTL Viewer when using a VQM or EDIF from third-party synthesis software</li> </ul> <p>In the Technology Map Viewer, you can view the internal implementation of certain atom primitives, but you cannot traverse into a lower-level of hierarchy.</p>
Ports	<p>The I/O ports in the current level of hierarchy.</p> <ul style="list-style-type: none"> <li>• Pins are device I/O pins when viewing the top hierarchy level and are I/O ports of the design when viewing the lower-levels.</li> <li>• When a pin represents a bus or an array of pins, expand the pin entry in the list view to see individual pin names.</li> </ul>

## Properties Pane

You can view the properties of an instance or primitive using the **Properties** pane.

**Figure 19-3: Properties Pane**

To view the properties of an instance or primitive in the RTL Viewer or Technology Map Viewer, right-click the node and click **Properties**.



The **Properties** pane contains tabs with the following information about the selected node:

- The **Fan-in** tab displays the **Input port** and **Fan-in Node**.
- The **Fan-out** tab displays the **Output port** and **Fan-out Node**.
- The **Parameters** tab displays the **Parameter Name** and **Values** of an instance.
- The **Ports** tab displays the **Port Name** and **Constant** value (for example,  $V_{CC}$  or GND). The possible value of a port are listed below.

**Table 19-2: Possible Port Values**

Value	Description
$V_{CC}$	The port is not connected and has $V_{CC}$ value (tied to $V_{CC}$ )
GND	The port is not connected and has GND value (tied to GND)

Value	Description
--	The port is connected and has value (other than V <sub>CC</sub> or GND)
Unconnected	The port is not connected and has no value (hanging)

If the selected node is an atom primitive, the **Properties** pane displays a schematic of the internal logic.

## Netlist Viewers Find Pane

You can narrow the range of the search process by setting the following options in the **Find** pane:

- Click **Browse** in the **Find** pane to specify the hierarchy level of the search. In the **Select Hierarchy Level** dialog box, select the particular instance you want to search.
- Turn on the **Include subentities** option to include child hierarchies of the parent instance during the search.
- Click **Options** to open the **Find Options** dialog box. Turn on **Instances**, **Nodes**, **Ports**, or any combination of the three to further refine the parameters of the search.

When you click the **List** button, a progress bar appears below the **Find** box.

All results that match the criteria you set are listed in a table. When you double-click an item in the table, the related node is highlighted in red in the schematic view.

## Schematic View

The schematic view is shown on the right side of the RTL Viewer and Technology Map Viewer. The schematic view contains a schematic representing the design logic in the netlist. This view is the main screen for viewing your gate-level netlist in the RTL Viewer and your technology-mapped netlist in the Technology Map Viewer.

The RTL Viewer and Technology Map Viewer attempt to display schematic in a single page view by default. If the schematic crosses over to several pages, you can highlight a net and use connectors to trace the signal in a single page.

## Display Schematics in Multiple Tabbed View

The RTL Viewer and Technology Map Viewer support multiple tabbed views.

With multiple tabbed view, schematics can be displayed in different tabs. Selection is independent between tabbed views, but selection in the tab in focus is synchronous with the Netlist Navigator pane.

To create a new blank tab, click the **New Tab** button at the end of the tab row. You can now drag a node from the **Netlist Navigator** pane into the schematic view.

You can right-click in a tab to see a shortcut menu where you can:

- Create a blank view with **New Tab**
- Create a **Duplicate Tab** of the tab in focus
- Choose to **Cascade Tabs**
- Choose to **Tile Tabs**
- Choose **Close Tab** to close the tab in focus
- Choose **Close Other Tabs** to close all tabs except the tab in focus




## Schematic Symbols







The symbols for nodes in the schematic represent elements of your design netlist. These elements include input and output ports, registers, logic gates, Altera® primitives, high-level operators, and hierarchical instances





The following table lists and describes the primitives and basic symbols that you can display in the schematic view of the RTL Viewer and Technology Map Viewer.

**Note:** The logic gates and operator primitives appear only in the RTL Viewer. Logic in the Technology Map Viewer is represented by atom primitives, such as registers and LCELLs.

**Table 19-3: Symbols in the Schematic View**


Symbol	Description
<p>I/O Ports</p> 	<p>An input, output, or bidirectional port in the current level of hierarchy. A device input, output, or bidirectional pin when viewing the top-level hierarchy. The symbol can also represent a bus. Only one wire is shown connected to the bidirectional symbol, representing the input and output paths.</p> <p>Input symbols appear on the left-most side of the schematic. Output and bidirectional symbols appear on the right-most side of the schematic.</p>
<p>I/O Connectors</p> 	<p>An input or output connector, representing a net that comes from another page of the same hierarchy. To go to the page that contains the source or the destination, double-click on the connector to jump to the appropriate page.</p>
<p>OR, AND, XOR Gates</p> 	<p>An OR, AND, or XOR gate primitive (the number of ports can vary). A small circle (bubble symbol) on an input or output port indicates the port is inverted.</p>

Symbol	Description
<p>MULTIPLEXER</p> 	<p>A multiplexer primitive with a selector port that selects between port 0 and port 1. A multiplexer with more than two inputs is displayed as an operator.</p>
<p>BUFFER</p> 	<p>A buffer primitive. The figure shows the tri-state buffer, with an inverted output enable port. Other buffers without an enable port include LCELL, SOFT, CARRY, and GLOBAL. The NOT gate and EXP expander buffers use this symbol without an enable port and with an inverted output port.</p>
<p>LATCH</p> 	<p>A latch/DFF (data flipflop) primitive. A DFF has the same ports as a latch and a clock trigger. The other flipflop primitives are similar:</p> <ul style="list-style-type: none"> <li>• DFFEA (data flipflop with enable and asynchronous load) primitive with additional <code>ALOAD</code> asynchronous load and <code>ADATA</code> data signals</li> <li>• DFFEAS (data flipflop with enable and synchronous and asynchronous load), which has <code>ASDATA</code> as the secondary data port</li> </ul>
<p>Atom Primitive</p> 	<p>An atom primitive. The symbol displays the atom name, the port names, and the atom type. The blue shading indicates an atom primitive for which you can view the internal details.</p>
<p>Other Primitive</p> 	<p>Any primitive that does not fall into the previous categories. Primitives are low-level nodes that cannot be expanded to any lower hierarchy. The symbol displays the port names, the primitive or operator type, and its name.</p>
<p>Instance</p> 	<p>An instance in the design that does not correspond to a primitive or operator (a user-defined hierarchy block). The symbol displays the port name and the instance name.</p>

Symbol	Description
<p>Encrypted Instance</p> 	<p>A user-defined encrypted instance in the design. The symbol displays the instance name. You cannot open the schematic for the lower-level hierarchy, because the source design is encrypted.</p>
<p>State Machine Instance</p> 	<p>A finite state machine instance in the design.</p>
<p>RAM</p> 	<p>A synchronous memory instance with registered inputs and optionally registered outputs. The symbol shows the device family and the type of memory block. This figure shows a true dual-port memory block in a Stratix M-RAM block.</p>
<p>Constant</p> 	<p>A constant signal value that is highlighted in gray and displayed in hexadecimal format by default throughout the schematic.</p>









The following table lists and describes the symbol open only in the State Machine Viewer.

Table 19-4: Symbol Available Only in the State Machine Viewer




Symbol	Description
State Node 	The node representing a state in a finite state machine. State transitions are indicated with arcs between state nodes. The double circle border indicates the state connects to logic outside the state machine, and a single circle border indicates the state node does not feed outside logic.

The following lists and describes the additional higher level operator symbols in the RTL Viewer schematic view.

Table 19-5: Operator Symbols in the RTL Viewer Schematic View

Symbol	Description
	An adder operator: $OUT = A + B$
	A multiplier operator: $OUT = A \times B$
	A divider operator: $OUT = A / B$
	Equals
	A left shift operator: $OUT = (A \ll COUNT)$
	A right shift operator: $OUT = (A \gg COUNT)$
	A modulo operator: $OUT = (A \% B)$
	A less than comparator: $OUT = (A < B : A > B)$



Symbol	Description
	<p>A multiplexer:</p> $\text{OUT} = \text{DATA} [\text{SEL}]$ <p>The data range size is <math>2^{\text{sel range size}}</math></p>
	<p>A selector:</p> <p>A multiplexer with one-hot select input and more than two input signals</p>
	<p>A binary number decoder:</p> $\text{OUT} = (\text{binary\_number} (\text{IN}) == x)$ <p>for <math>x = 0</math> to <math>x = 2^{(n+1)} - 1</math></p>

**Related Information**

- [Partition the Schematic into Pages](#) on page 19-19
- [Follow Nets Across Schematic Pages](#) on page 19-19
- [State Machine Viewer](#) on page 19-20

## Select Items in the Schematic View

To select an item in the schematic view, ensure that the **Selection Tool** is enabled in the Netlist Viewer toolbar (this tool is enabled by default). Click an item in the schematic view to highlight it in red.

Select multiple items by pressing the Shift key while selecting with your mouse.

Items selected in the schematic view are automatically selected in the **Netlist Navigator** pane. The folder then expands automatically if it is required to show the selected entry; however, the folder does not collapse automatically when you are not using or you have deselected the entries.

When you select a hierarchy box, node, or port in the schematic view, the item is highlighted in red but none of the connecting nets are highlighted. When you select a net (wire or bus) in the schematic view, all connected nets are highlighted in red.

Once you have selected an item, you can perform different actions on it based on the contents of the shortcut menu which appears when you right-click on your selection.

**Related Information**

[Netlist Navigator Pane](#) on page 19-7

## Shortcut Menu Commands in the Schematic View

When you right-click on an instance or primitive selected in the schematic view, the Netlist Viewer displays a shortcut menu.

If the selected item is a node, you see the following options:

- Click **Expand to Upper Hierarchy** to displays the parent hierarchy of the node in focus.
- Click **Copy ToolTip** to copy the selected item name to the clipboard. This command does not work on nets.
- Click **Hide Selection** to remove the selected item from the schematic view. This does not delete the item from the design, merely masks it in the current view.
- Click **Filtering** to display a sub-menu with options for filtering your selection.

When the selected item is a net, the shortcut menu displays the option to **Unbundle Net**. When you unbundle a net, all connected bus pins and ports are ungrouped and displayed. You can use this to trace bundled connections more easily.

## Filtering in the Schematic View

Filtering allows you to filter out nodes and nets in your netlist to view only the logic elements of interest to you.

You can filter your netlist by selecting hierarchy boxes, nodes, ports of a node, or states in a state machine that are part of the path you want to see. The following filter commands are available:

- **Sources**—Displays the sources of the selection.
- **Destinations**—Displays the destinations of the selection.
- **Sources & Destinations**—displays the sources and destinations of the selection.
- **Selected Nodes**—Displays only the selected nodes.
- **Between Selected Nodes**—Displays nodes and connections in the path between the selected nodes .
- **Bus Index**—Displays the sources or destinations for one or more indices of an output or input bus port .
- **Filtering Options**—Displays the **Filtering Options** dialog box:
  - **Stop filtering at register**—Turning this option on directs the Netlist Viewer to filter out to the nearest register boundary.
  - **Filter across hierarchies**—Turning this option on directs the Netlist Viewer to filter across hierarchies.
  - **Maximum number of hierarchy levels**—Sets the maximum number of hierarchy levels displayed in the schematic view.

To filter your netlist, select a hierarchy box, node, port, net, or state node, right-click in the window, point to **Filter** and click the appropriate filter command. The Netlist Viewer generates a new page showing the netlist that remains after filtering.

When filtering in a state diagram in the State Machine Viewer, sources and destinations refer to the previous and next transition states or paths between transition states in the state diagram. The transition table and encoding table also reflect the filtering.


## View Contents of Nodes in the Schematic View

In the RTL Viewer and the Technology Map Viewer, you can view the contents of nodes to see their underlying implementation details.

You can view LUTs, registers, and logic gates. You can also view the implementation of RAM and DSP blocks in certain devices in the RTL Viewer or Technology Map Viewer. In the Technology Map Viewer, you can view the contents of primitives to see their underlying implementation details.

**Figure 19-4: Wrapping and Unwrapping Objects**

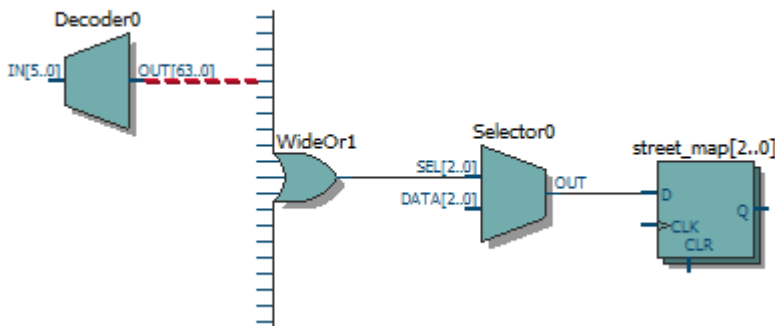
If you can unwrap the contents of an instance, a plus symbol appears in the upper right corner of the object in the schematic view. To wrap the contents (and revert to the compact format), click the minus symbol in the upper right corner of the unwrapped instance.



**Note:** In the schematic view, the internal details in an atom instance cannot be selected as individual nodes. Any mouse action on any of the internal details is treated as a mouse action on the atom instance.

**Figure 19-5: Nodes with Connections Outside the Hierarchy**

In some cases, the selected instance connects to something outside the visible level of the hierarchy in the schematic view. In this case, the net appears as a dotted line. Double-click on the dotted line to expand the view to display the destination of the connection.

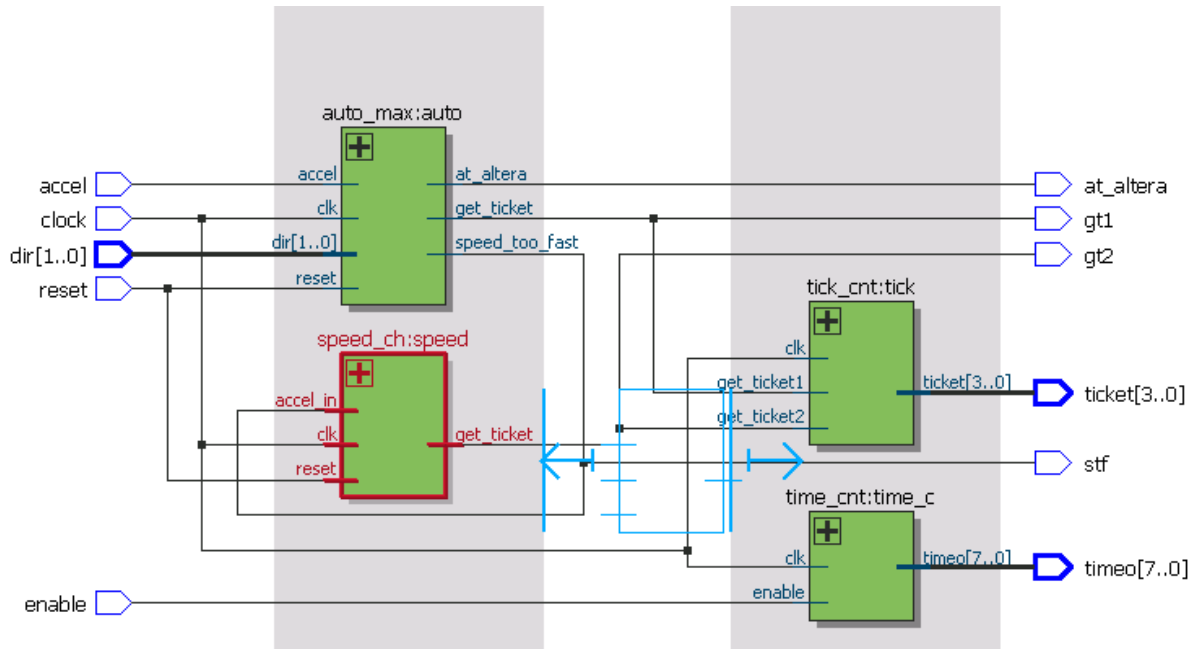


## Moving Nodes in the Schematic View

You can drag and drop items in the schematic view to rearrange them.

### Figure 19-6: Drag and Drop Movement of Nodes

To move a node from one area of the netlist to another, select the node and hold down the Shift key. Legal placements appear as shaded areas within the hierarchy. Click to drop the selected node.



Right-click and click on **Refresh** to restore the schematic view to its default arrangement.

## View LUT Representations in the Technology Map Viewer

You can view different representations of a LUT by right-clicking the selected LUT and clicking **Properties**.

You can view the LUT representations in the following three tabs in the **Properties** dialog box:

- The **Schematic** tab—the equivalent gate representations of the LUT.
- The **Truth Table** tab—the truth table representations.

### Related Information

[Properties Pane](#) on page 19-8

## Zoom Controls

You can control the magnification of your schematic on the View menu, with the Zoom Tool in the toolbar, or with mouse gestures.

By default, the Netlist Viewer displays most pages sized to fit in the window. If the schematic page is very large, the schematic is displayed at the minimum zoom level, and the view is centered on the first node. Click **Zoom In** to view the image at a larger size, and click **Zoom Out** to view the image (when the entire image is not displayed) at a smaller size. The **Zoom** command allows you to specify a magnification percentage (100% is considered the normal size for the schematic symbols).

Within the schematic view, you can also use the following mouse gestures to zoom in on a specific section:

- **zoom in**—Dragging a box around an area starting in the upper-left and dragging to the lower right zooms in on that area.
- **zoom -0.5**—Dragging a line from lower-left to upper-right zooms out 0.5 levels of magnification.
- **zoom 0.5**—Dragging a line from lower-right to upper-left zooms in 0.5 levels of magnification.
- **zoom fit**—Dragging a line from upper-right to lower-left fits the schematic view in the page.

You can also use the Zoom Tool on the Netlist Viewer toolbar to control magnification in the schematic view. When you select the Zoom Tool in the toolbar, clicking in the schematic zooms in and centers the view on the location you clicked. Right-click in the schematic to zoom out and center the view on the location you clicked. When you select the Zoom Tool, you can also zoom into a certain portion of the schematic by selecting a rectangular box area with your mouse cursor. The schematic is enlarged to show the selected area.

#### Related Information

[Filtering in the Schematic View](#) on page 19-16

## Navigating with the Bird's Eye View

To open the Bird's Eye View, on the View menu, click **Bird's Eye View**, or click on the **Bird's Eye View** icon in the toolbar.

Viewing the entire schematic can be useful when debugging and tracing through a large netlist. The Quartus II software allows you to quickly navigate to a specific section of the schematic using the Bird's Eye View feature, which is available in the RTL Viewer and Technology Map Viewer.

The Bird's Eye View shows the current area of interest.

- Select an area by clicking and dragging the indicator or right-clicking to form a rectangular box around an area.
- Click and drag the rectangular box to move around the schematic.
- Resize the rectangular box to zoom-in or zoom-out in the schematic view.

## Partition the Schematic into Pages

For large design hierarchies, the RTL Viewer and Technology Map Viewer partition your netlist into multiple pages in the schematic view.

When a hierarchy level is partitioned into multiple pages, the title bar for the schematic window indicates which page is displayed and how many total pages exist for this level of hierarchy. The schematic view displays this as **Page** <current page number> **of** <total number of pages>.

#### Related Information

[Introduction to the User Interface](#) on page 19-5

## Follow Nets Across Schematic Pages

Input and output connector symbols indicate nodes that connect across pages of the same hierarchy. Double-click a connector to trace the net to the next page of the hierarchy.

**Note:** After you double-click to follow a connector port, the Netlist Viewer opens a new page, which centers the view on the particular source or destination net using the same zoom factor as the previous page. To trace a specific net to the new page of the hierarchy, Altera recommends that you first select the necessary net, which highlights it in red, before you double-click to traverse pages.

## Related Information

[Schematic Symbols](#) on page 19-11

## State Machine Viewer

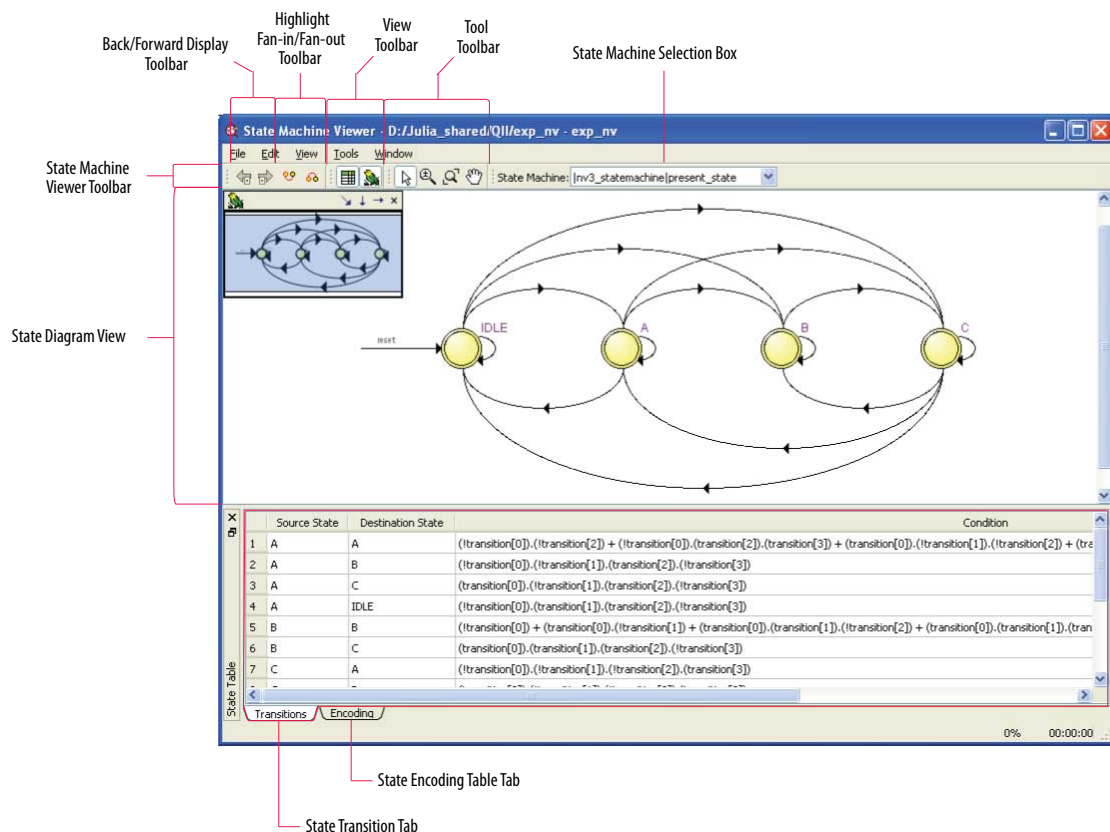
The State Machine Viewer displays a graphical representation of the state machines in your design.

You can open the State Machine Viewer in any of the following ways:

- On the Tools menu, point to **Netlist Viewers** and click **State Machine Viewer**.
- Double-click a state machine instance in the RTL Viewer

**Figure 19-7: The State Machine Viewer**

The following figure shows an example of the State Machine Viewer for a simple state machine and lists the components of the viewer.



### State Diagram View

The state diagram view appears at the top of the State Machine Viewer. It contains a diagram of the states and state transitions.

The nodes that represent each state are arranged horizontally in the state diagram view with the initial state (the state node that receives the reset signal) in the left-most position. Nodes that connect to logic outside of the state machine instance are represented by a double circle. The state transition is represented by an arc with an arrow pointing in the direction of the transition.

When you select a node in the state diagram view, and turn on the **Highlight Fan-in** or **Highlight Fan-out** command from the View menu or the State Machine Viewer toolbar, the respective fan-in or fan-out transitions from the node are highlighted in red.

**Note:** An encrypted block with a state machine displays encoding information in the state encoding table, but does not display a state transition diagram or table.

## State Transition Table

The state transition table on the **Transitions** tab at the bottom of the State Machine Viewer displays the condition equation for each state transition.

Each row in the table represents a transition (each arc in the state diagram view). The table has the following columns:

- **Source State**—the name of the source state for the transition
- **Destination State**—the name of the destination state for the transition
- **Condition**—the condition equation that causes the transition from source state to destination state

To see all of the transitions to and from each state name, click the appropriate column heading to sort on that column.

The text in each column is left-aligned by default; to change the alignment and to make it easier to see the relevant part of the text, right-click the column and click **Align Right**. To revert to left alignment, click **Align Left**.

Click in any cell in the table to select it. To select all cells, right-click in the cell and click **Select All**; or, on the Edit menu, click **Select All**. To copy selected cells to the clipboard, right-click the cells and click **Copy Table**; or, on the Edit menu, point to **Copy** and click **Copy Table**. You can paste the table into any text editor as tab-separated columns.

## State Encoding Table

The state encoding table on the **Encoding** tab at the bottom of the State Machine Viewer displays encoding information for each state transition.

To view state encoding information in the State Machine Viewer, you must synthesize your design with the **Start Analysis & Synthesis** command. If you have only elaborated your design with the **Start Analysis & Elaboration** command, the encoding information is not displayed.

## Select Items in the State Machine Viewer

You can select and highlight each state node and transition in the State Machine Viewer. To select a state transition, click the arc that represents the transition.

When you select a node or transition arc in the state diagram view, the matching state node or equation conditions in the state transition table are highlighted; conversely, when you select a state node or equation condition in the state transition table, the corresponding state node or transition arc is highlighted in the state diagram view.

## Switch Between State Machines

A design may contain multiple state machines. To choose which state machine to view, use the **State Machine** selection box located at the top of the State Machine Viewer. Click in the drop-down box and select the necessary state machine.

## Probing to a Source Design File and Other Quartus II Windows

The RTL Viewer, Technology Map Viewer, and State Machine Viewer allow you to cross-probe to the source design file and to various other windows in the Quartus II software.

You can select one or more hierarchy boxes, nodes, state nodes, or state transition arcs that interest you in the Netlist Viewer and locate the corresponding items in another applicable Quartus II software window. You can then view and make changes or assignments in the appropriate editor or floorplan.

To locate an item from the Netlist Viewer in another window, right-click the items of interest in the schematic or state diagram, point to **Locate**, and click the appropriate command. The following commands are available:

- **Locate in Assignment Editor**
- **Locate in Pin Planner**
- **Locate in Chip Planner**
- **Locate in Resource Property Editor**
- **Locate in Technology Map Viewer**
- **Locate in RTL Viewer**
- **Locate in Design File**

The options available for locating an item depend on the type of node and whether it exists after placement and routing. If a command is enabled in the menu, it is available for the selected node. You can use the **Locate in Assignment Editor** command for all nodes, but assignments might be ignored during placement and routing if they are applied to nodes that do not exist after synthesis.

The Netlist Viewer automatically opens another window for the appropriate editor or floorplan and highlights the selected node or net in the newly opened window. You can switch back to the Netlist Viewer by selecting it in the Window menu or by closing, minimizing, or moving the new window.

## Probing to the Netlist Viewers from Other Quartus II Windows

You can cross-probe to the RTL Viewer and Technology Map Viewer from other windows in the Quartus II software. You can select one or more nodes or nets in another window and locate them in one of the Netlist Viewers.

You can locate nodes between the RTL Viewer, State Machine Viewer, and Technology Map Viewer, and you can locate nodes in the RTL Viewer and Technology Map Viewer from the following Quartus II software windows:

- Project Navigator
- Timing Closure Floorplan
- Chip Planner
- Resource Property Editor
- Node Finder



- Assignment Editor
- Messages Window
- Compilation Report
- TimeQuest Timing Analyzer (supports the Technology Map Viewer only)

To locate elements in the Netlist Viewer from another Quartus II window, select the node or nodes in the appropriate window; for example, select an entity in the **Entity** list on the **Hierarchy** tab in the Project Navigator, or select nodes in the Timing Closure Floorplan, or select node names in the **From** or **To** column in the Assignment Editor. Next, right-click the selected object, point to **Locate**, and click **Locate in RTL Viewer** or **Locate in Technology Map Viewer**. After you click this command, the Netlist Viewer opens, or is brought to the foreground if the Netlist Viewer is open.

When cross-probing from the Time

**Note:** The first time the window opens after a compilation, the preprocessor stage runs before the Netlist Viewer opens.

The Netlist Viewer shows the selected nodes and, if applicable, the connections between the nodes. The display is similar to what you see if you right-click the object, point to **Filter**, and click **Selected Nodes** using **Filter across hierarchy**. If the nodes cannot be found in the Netlist Viewer, a message box displays the message: **Can't find requested location**.

## Viewing a Timing Path

You can cross-probe from a report panel in the TimeQuest Timing Analyzer to see a visual representation of a timing path.

To take advantage of this feature, you must complete a full compilation of your design, including the timing analyzer stage. To see the timing results for your design, on the Processing menu, click **Compilation Report**. On the left side of the Compilation Report, select **TimeQuest Timing Analyzer**. When you select a detailed report, the timing information is listed in a table format on the right side of the Compilation Report; each row of the table represents a timing path in the design. You can also view timing paths in TimeQuest analyzer report panels. To view a particular timing path in the Technology Map Viewer or RTL Viewer, right-click the appropriate row in the table, point to **Locate**, and click **Locate in Technology Map Viewer** or **Locate in RTL Viewer**.

- To locate a path, on the **Tasks** pane, in the **Custom Reports** folder, double-click **Report Timing**.
- In the **Report Timing** dialog box, make necessary settings, and then click the **Report Timing** button.
- After the TimeQuest analyzer generates the report, right-click on the node in the table and select **Locate Path**. In the Technology Map Viewer, the schematic page displays the nodes along the timing path with a summary of the total delay.

When you locate the timing path from the TimeQuest analyzer to the Technology Map Viewer, the interconnect and cell delay associated with each node is displayed on top of the schematic symbols. The total slack of the selected timing path is displayed in the Page Title section of the schematic.

In the RTL Viewer, the schematic page displays the nodes in the paths between the source and destination registers with a summary of the total delay.

The RTL Viewer netlist is based on an initial stage of synthesis, so the post-fitting nodes might not exist in the RTL Viewer netlist. Therefore, the internal delay numbers are not displayed in the RTL Viewer as they are in the Technology Map Viewer, and the timing path might not be displayed exactly as it appears in the timing analysis report. If multiple paths exist between the source and destination registers, the RTL Viewer might display more than just the timing path. There are also some cases in which the path cannot be displayed, such as paths through state machines, encrypted intellectual property (IP), or registers that

are created during the fitting process. In cases where the timing path displayed in the RTL Viewer might not be the correct path, the compiler issues messages.

## Document Revision History

Date	Version	Changes
2014.06.30	14.0.0	Added Show Netlist on One Page and show/Hide Instance Pins commands.
November 2013	13.1.0	Removed HardCopy device information. Reorganized and migrated to new template. Added support for new Netlist viewer.
November 2012	12.1.0	Added sections to support Global Net Routing feature.
June 2012	12.0.0	Removed survey link.
November 2011	10.0.2	Template update.
December 2010	10.0.1	Changed to new document template.
July 2010	10.0.0	<ul style="list-style-type: none"> <li>• Updated screenshots</li> <li>• Updated chapter for the Quartus II software version 10.0, including major user interface changes</li> </ul>
November 2009	9.1.0	<ul style="list-style-type: none"> <li>• Updated devices</li> <li>• Minor text edits</li> </ul>

Date	Version	Changes
March 2009	9.0.0	<ul style="list-style-type: none"><li>• Chapter 13 was formerly Chapter 12 in version 8.1.0</li><li>• Updated Figure 13–2, Figure 13–3, Figure 13–4, Figure 13–14, and Figure 13–30</li><li>• Added “Enable or Disable the Auto Hierarchy List” on page 13–15</li><li>• Updated “Find Command” on page 13–44</li></ul>
November 2008	8.1.0	Changed page size to 8.5” × 11”

Date	Version	Changes
May 2008	8.0.0	<ul style="list-style-type: none"> <li>• Added Arria GX support</li> <li>• Updated operator symbols</li> <li>• Updated information about the radial menu feature</li> <li>• Updated zooming feature</li> <li>• Updated information about probing from schematic to SignalTap II Analyzer</li> <li>• Updated constant signal information</li> <li>• Added .png and .gif to the list of supported image file formats</li> <li>• Updated several figures and tables</li> <li>• Added new sections “Enabling and Disabling the Radial Menu”, “Changing the Time Interval”, “Changing the Constant Signal Value Formatting”, “Logic Clouds in the RTL Viewer”, “Logic Clouds in the Technology Map Viewer”, “Manually Group and Ungroup Logic Clouds”, “Customizing the Shortcut Commands”</li> <li>• Renamed several sections</li> <li>• Removed section “Customizing the Radial Menu”</li> <li>• Moved section “Grouping Combinational Logic into Logic Clouds”</li> <li>• Updated document content based on the Quartus II software version 8.0</li> </ul>

**Related Information****[Quartus II Handbook Archive](#)**

For previous versions of the *Quartus II Handbook*, refer to the Quartus II Handbook Archive.

# Quartus II Handbook Volume 2: Design Implementation and Optimization



Subscribe



Send Feedback

**QII5V2**  
2014.12.15

101 Innovation Drive  
San Jose, CA 95134  
[www.altera.com](http://www.altera.com)



2014.06.30

QI15V2



Subscribe



Send Feedback

Constraints, sometimes known as assignments or logic options, control the way the Quartus II software implements a design for an FPGA. Constraints are also central in the way that the TimeQuest Timing Analyzer and the PowerPlay Power Analyzer inform synthesis, placement, and routing.

There are several types of constraints:

- Global design constraints and software settings, such as device family selection, package type, and pin count.
- Entity-level constraints, such as logic options and placement assignments.
- Instance-level constraints.
- Pin assignments and I/O constraints.

User-created constraints are contained in one of two files: the Quartus II Settings File (**.qsf**) or, in the case of timing constraints, the Synopsys Design Constraints file (**.sdc**). Constraints and assignments made with the **Device** dialog box, **Settings** dialog box, Assignment Editor, Chip Planner, and Pin Planner are contained in the Quartus II Settings File. The **.qsf** file contains project-wide and instance-level assignments for the current revision of the project in Tcl syntax. You can create separate revisions of your project with different settings, and there is a separate **.qsf** file for each revision.

The TimeQuest Timing Analyzer uses industry-standard Synopsys Design Constraints, also using Tcl syntax, that are contained in Synopsys Design Constraints (**.sdc**) files. The TimeQuest Timing Analyzer GUI is a tool for making timing constraints and viewing the results of subsequent analysis.

There are several ways to constrain a design, each potentially more appropriate than the others, depending on your tool chain and design flow. You can constrain designs for compilation and analysis in the Quartus II software using the GUI, as well as using Tcl syntax and scripting. By combining the Tcl syntax of the **.qsf** files and the **.sdc** files with procedural Tcl, you can automate iteration over several different settings, changing constraints and recompiling.

## Constraining Designs with the Quartus II GUI

In the Quartus II GUI, the New Project Wizard, **Device** dialog box, and **Settings** dialog box allow you to make global constraints and software settings. The Assignment Editor and Pin Planner are spreadsheet-style interfaces for constraining your design at the instance or entity level.

The Assignment Editor and Pin Planner make constraint types and values available based on global design characteristics such as the targeted device. These tools help you verify that your constraints are valid before compilation by allowing you to pick only from valid values for each constraint.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



The TimeQuest Timing Analyzer GUI allows you to make timing constraints in SDC format and view the effects of those constraints on the timing in your design. Before running the TimeQuest timing analyzer, you must specify initial timing constraints that describe the clock characteristics, timing exceptions, and external signal arrival and required times. The Quartus II Fitter optimizes the placement of logic in the device to meet your specified constraints.

#### Related Information

#### [About TimeQuest Timing Analysis](#)

For more information about timing constraints and the TimeQuest Timing Analyzer

## Global Constraints

Global constraints affect the entire Quartus II project and all of the applicable logic in the design. Many of these constraints are simply project settings, such as the targeted device selected for the design.

Synthesis optimizations and global timing and power analysis settings can also be applied with globally. Global constraints are often made when running the New Project Wizard, or in the **Device** dialog box or the **Settings** dialog box, early project development.

### Common Types of Global Constraints

The following are the most common types of global constraints:

- Target device specification
- Top-level entity of your design, and the names of the design files included in the project
- Operating temperature limits and conditions
- Physical synthesis optimizations
- Analysis and synthesis options and optimization techniques
- Verilog HDL and VHDL language versions used in your project
- Fitter effort and timing driven compilation settings
- **.sdc** files for the TimeQuest timing analyzer to use during analysis as part of a full compilation flow

### Settings That Direct Compilation and Analysis Flows

Settings that direct compilation and analysis flows in the Quartus II software are also stored in the Quartus II Settings File for your project, including the following global software settings:

- Early Timing Estimate mode
- Settings for EDA tool integration such as third-party synthesis tools, simulation tools, timing analysis tools, and formal verification tools.
- Settings and settings file specifications for the Quartus II Assembler, SignalTap II Logic Analyzer, PowerPlay power analyzer, and SSN Analyzer.

### Global Constraints and Software Settings

Global constraints and software settings stored in the Quartus II settings file are specific to each revision of your design, allowing you to control the operation of the software differently for different revisions. For example, different revisions can specify different operating temperatures and different devices, so that you can compare results.

Only the valid assignments made in the Assignment Editor are saved in the Quartus II Settings File, which is located in the project directory. When you make a design constraint, the new assignment is placed on a new line at the end of the file.

When you create or update a constraint in the GUI, the Quartus II software displays the equivalent Tcl command in the **System** tab of the Messages window. You can use the displayed messages as references when making assignments using Tcl commands.

#### Related Information

- [Setting Up and Running a Compilation](#)  
For more information about specifying initial global constraints and software settings
- [Managing Quartus II Projects](#)  
For more information about how the Quartus II software uses Quartus II Settings Files

## Node, Entity, and Instance-Level Constraints

Node, entity, and instance-level constraints constrain a particular segment of the design hierarchy, as opposed to the entire design. In the Quartus II software GUI, most instance-level constraints are made with the Assignment Editor, Pin Planner, and Chip Planner.

Both the Assignment Editor and Pin Planner aid you in correctly constraining your design, both passively, through device-and-assignment-determined pick lists, and actively, through live I/O checking.

You can assign logic functions to physical resources on the device, using location assignments with the Assignment Editor or the Chip Planner. Node, entity, and instance-level constraints take precedence over any global constraints that affect the same sections of the design hierarchy. You can edit and view all node and entity-level constraints you created in the Assignment Editor, or you can filter the assignments by choosing to view assignments only for specific locations, such as DSP blocks.

#### Related Information

- [Assigning Device I/O Pins in Pin Planner](#)
- [About the Chip Planner](#)
- [About the Assignment Editor](#)

## Constraining Designs with the Pin Planner

The Pin Planner helps you visualize, plan, and assign device I/O pins to ensure compatibility with your PCB layout. The Pin Planner provides a graphical view of the I/O resources in the target device package. You can quickly locate various I/O pins and assign them design elements or other properties.

The Quartus II software uses these assignments to place and route your design during device programming. The Pin Planner also helps with early pin planning by allowing you to plan and assign IP interface or user nodes not yet defined in the design.

The Pin Planner Task window provides one-click access to common pin planning tasks. After clicking a pin planning task, you view and highlight the results in the Report window by selecting or deselecting I/O types. You can quickly identify I/O banks, VREF groups, edges, and differential pin pairings to assist you in the pin planning process. You can verify the legality of new and existing pin assignments with the live I/O check feature and view the results in the Live I/O Check Status window.

## Constraining Designs with the Chip Planner

The Chip Planner allows you to view the device from a variety of different perspectives, and you can make precise assignments to specific floorplan locations.

With the Chip Planner, you can adjust existing assignments to device resources, such as pins, logic cells, and LABs using drag and drop features and a graphical interface. You can also view equations and routing



information, and demote assignments by dragging and dropping assignments to various regions in the Regions window.

## Probing Between Components of the Quartus II GUI

The Assignment Editor, Chip Planner, and Pin Planner let you locate nodes and instances in the source files for your design in other Quartus II viewers.

You can select a cell in the Assignment Editor spreadsheet and locate the corresponding item in another applicable Quartus II software window, such as the Chip Planner. To locate an item from the Assignment Editor in another window, right-click the item of interest in the spreadsheet, point to **Locate**, and click the appropriate command.

You can also locate nodes in the Assignment Editor and other constraint tools from other windows within the Quartus II software. First, select the node or nodes in the appropriate window. For example, select an entity in the **Entity** list in the **Hierarchy** tab in the Project Navigator, or select nodes in the Chip Planner. Next, right-click the selected object, point to **Locate**, and click **Locate in Assignment Editor**. The Assignment Editor opens, or it is brought to the foreground if it is already open.

### Related Information

- [Assigning Device I/O Pins in Pin Planner](#)
- [About the Chip Planner](#)
- [About the Assignment Editor](#)

## SDC and the TimeQuest Timing Analyzer

You can make individual timing constraints for individual entities, nodes, and pins with the Constraints menu of the TimeQuest Timing Analyzer. The TimeQuest Timing Analyzer GUI provides easy access to timing constraints, and reporting, without requiring knowledge of SDC syntax.

As you specify commands and options in the GUI, the corresponding SDC or Tcl command appears in the Console. This lets you know exactly what constraint you have added to your Synopsys Design Constraints file, and also enables you to learn SDC syntax for use in scripted flows. The GUI also provides enhanced graphical reporting features.

Individual timing assignments override project-wide requirements. You can also assign timing exceptions to nodes and paths to avoid reporting of incorrect or irrelevant timing violations. The TimeQuest timing analyzer supports point-to-point timing constraints, wildcards to identify specific nodes when making constraints, and assignment groups to make individual constraints to groups of nodes.

### Related Information

#### [About TimeQuest Timing Analysis](#)

For more information about timing constraints and the TimeQuest Timing Analyzer

## Constraining Designs with Tcl

Because **.sdc** files and **.qsf** files are both in Tcl syntax, you can modify these files to be part of a scripted constraint and compilation flow.

With Quartus II Tcl packages, Tcl scripts can open projects, make the assignments procedurally that would otherwise be specified in a **.qsf** file, compile a design, and compare compilation results against



known goals and benchmarks for the design. Such a script can further automate the iterative process by modifying design constraints and recompiling the design.

### Related Information

#### About Quartus II Tcl Scripting

## Quartus II Settings Files and Tcl

QSF files use Tcl syntax, but, unmodified, are not executable scripts. However, you can embed QSF constraints in a scripted iterative compilation flow, where the script that automates compilation and custom results reporting also contains the design constraints.

```
set_global_assignment -name FAMILY "Cyclone II"
set_global_assignment -name DEVICE EP2C35F672C6
set_global_assignment -name TOP_LEVEL_ENTITY chiptrip
set_global_assignment -name ORIGINAL_QUARTUS_VERSION 10.0
set_global_assignment -name PROJECT_CREATION_TIME_DATE "11:45:02 JUNE 08, 2010"
set_global_assignment -name LAST_QUARTUS_VERSION 10.0
set_global_assignment -name MIN_CORE_JUNCTION_TEMP 0
set_global_assignment -name MAX_CORE_JUNCTION_TEMP 85
set_instance_assignment -name PARTITION_HIERARCHY root_partition -to | -section_id
Top
set_global_assignment -name PARTITION_NETLIST_TYPE SOURCE -section_id Top
set_global_assignment -name PARTITION_FITTER_PRESERVATION_LEVEL
PLACEMENT_AND_ROUTING \ -section_id Top
set_global_assignment -name PARTITION_COLOR 16764057 -section_id Top
set_global_assignment -name LL_ROOT_REGION ON -section_id "Root Region"
set_global_assignment -name LL_MEMBER_STATE LOCKED -section_id "Root Region"
set_global_assignment -name STRATIX_DEVICE_IO_STANDARD "3.3-V LVTTL"
set_location_assignment PIN_P2 -to clk2
set_location_assignment PIN_AE4 -to ticket[0]
set_location_assignment PIN_J23 -to ticket[2]
set_location_assignment PIN_Y12 -to timeo[1]
set_location_assignment PIN_N2 -to reset
set_location_assignment PIN_R2 -to timeo[7]
set_location_assignment PIN_P1 -to clk1
set_location_assignment PIN_M3 -to ticket[1]
set_location_assignment PIN_AE24 -to ~LVDS150p/nCEO~
set_location_assignment PIN_C2 -to accel
set_location_assignment PIN_K4 -to ticket[3]
set_location_assignment PIN_B3 -to stf
set_location_assignment PIN_T9 -to timeo[0]
set_location_assignment PIN_M5 -to timeo[6]
set_location_assignment PIN_J8 -to dir[1]
set_location_assignment PIN_C5 -to timeo[5]
set_location_assignment PIN_F6 -to gt1
set_location_assignment PIN_P24 -to timeo[2]
set_location_assignment PIN_B2 -to at_altera
set_location_assignment PIN_P3 -to timeo[4]
set_location_assignment PIN_M4 -to enable
set_location_assignment PIN_E3 -to ~ASDO~
set_location_assignment PIN_E5 -to dir[0]
set_location_assignment PIN_R25 -to timeo[3]
set_location_assignment PIN_D3 -to ~nCSO~
set_location_assignment PIN_G4 -to gt2
set_global_assignment -name MISC_FILE "D:/altera/chiptrip/chiptrip.dpf"
set_global_assignment -name USE_TIMEQUEST_TIMING_ANALYZER ON
set_global_assignment -name POWER_PRESET_COOLING_SOLUTION \
"23 MM HEAT SINK WITH 200 LFPM AIRFLOW"
set_global_assignment -name POWER_BOARD_THERMAL_MODEL "NONE (CONSERVATIVE)"
set_global_assignment -name SDC_FILE chiptrip.sdc
```

The example shows the way that the `set_global_assignment` Quartus II Tcl command makes all global constraints and software settings, with `set_location_assignment` constraining each I/O node in the design to a physical pin on the device.

However, after you initially create the Quartus II Settings File for your design, you can export the contents to a procedural, executable Tcl (**.tcl**) file. You can then use that generated script to restore certain settings after experimenting with other constraints. You can also use the generated Tcl script to archive your assignments instead of archiving the Quartus II Settings file itself.

To export your constraints as an executable Tcl script, on the Project menu, click **Generate Tcl File for Project**.

```
# Quartus II: Generate Tcl File for Project
# File: chiptrip.tcl
# Generated on: Tue Jun 08 13:08:48 2010
# Load Quartus II Tcl Project package
package require ::quartus::project
set need_to_close_project 0
set make_assignments 1
# Check that the right project is open
if {[is_project_open]} {
    if {[string compare $quartus(project) "chiptrip"]} {
        puts "Project chiptrip is not open"
        set make_assignments 0
    }
} else {
    # Only open if not already open
    if {[project_exists chiptrip]} {
        project_open -revision chiptrip chiptrip
    } else {
        project_new -revision chiptrip chiptrip
    }
    set need_to_close_project 1
}
# Make assignments
if {$make_assignments} {
    set_global_assignment -name FAMILY "Cyclone II"
    set_global_assignment -name DEVICE EP2C35F672C6
    set_global_assignment -name TOP_LEVEL_ENTITY chiptrip
    set_global_assignment -name ORIGINAL_QUARTUS_VERSION 10.0
    set_global_assignment -name PROJECT_CREATION_TIME_DATE "11:45:02 JUNE 08, 2010"
    set_global_assignment -name LAST_QUARTUS_VERSION 10.0
    set_global_assignment -name MIN_CORE_JUNCTION_TEMP 0
    set_global_assignment -name MAX_CORE_JUNCTION_TEMP 85
    set_instance_assignment -name PARTITION_HIERARCHY root_partition -to | -section_id
    Top
    set_global_assignment -name PARTITION_NETLIST_TYPE SOURCE -section_id Top
    set_global_assignment -name PARTITION_FITTER_PRESERVATION_LEVEL
    PLACEMENT_AND_ROUTING \ -section_id Top
    set_global_assignment -name PARTITION_COLOR 16764057 -section_id Top
    set_global_assignment -name LL_ROOT_REGION ON -section_id "Root Region"
    set_global_assignment -name LL_MEMBER_STATE LOCKED -section_id "Root Region"
    set_global_assignment -name STRATIX_DEVICE_IO_STANDARD "3.3-V LVTTL"
    set_location_assignment PIN_P2 -to clk2
    set_location_assignment PIN_AE4 -to ticket[0]
    set_location_assignment PIN_J23 -to ticket[2]
    set_location_assignment PIN_Y12 -to timeo[1]
    set_location_assignment PIN_N2 -to reset
    set_location_assignment PIN_R2 -to timeo[7]
    set_location_assignment PIN_P1 -to clk1
    set_location_assignment PIN_M3 -to ticket[1]
    set_location_assignment PIN_AE24 -to ~LVDS150p/nCEO~
    set_location_assignment PIN_C2 -to accel
    set_location_assignment PIN_K4 -to ticket[3]
```

```

set_location_assignment PIN_B3 -to stf
set_location_assignment PIN_T9 -to timeo[0]
set_location_assignment PIN_M5 -to timeo[6]
set_location_assignment PIN_J8 -to dir[1]
set_location_assignment PIN_C5 -to timeo[5]
set_location_assignment PIN_F6 -to gt1
set_location_assignment PIN_P24 -to timeo[2]
set_location_assignment PIN_B2 -to at_altera
set_location_assignment PIN_P3 -to timeo[4]
set_location_assignment PIN_M4 -to enable
set_location_assignment PIN_E3 -to ~ASDO~
set_location_assignment PIN_E5 -to dir[0]
set_location_assignment PIN_R25 -to timeo[3]
set_location_assignment PIN_D3 -to ~nCSO~
set_location_assignment PIN_G4 -to gt2
set_global_assignment -name MISC_FILE "D:/altera/chiptrip/chiptrip.dpf"
set_global_assignment -name USE_TIMEQUEST_TIMING_ANALYZER ON
set_global_assignment -name POWER_PRESET_COOLING_SOLUTION \
"23 MM HEAT SINK WITH 200 LFPM AIRFLOW"
set_global_assignment -name POWER_BOARD_THERMAL_MODEL "NONE (CONSERVATIVE)"
set_global_assignment -name SDC_FILE chiptrip.sdc
# Commit assignments
export_assignments
# Close project
if {$need_to_close_project} {
    project_close
}
}

```

After setting initial values for variables to control constraint creation and whether or not the project needs to be closed at the end of the script, the generated script checks to see if a project is open. If a project is open but it is not the correct project, in this case, **chiptrip**, the script prints `Project chiptrip is not open` to the console and does nothing else.

If no project is open, the script determines if **chiptrip** exists in the current directory. If the project exists, the script opens the project. If the project does not exist, the script creates a new project and opens the project.

The script then creates the constraints. After creating the constraints, the script writes the constraints to the Quartus II Settings File and then closes the project.

## Timing Analysis with Synopsys Design Constraints and Tcl

Timing constraints used in analysis by the Quartus II TimeQuest Timing Analyzer are stored in **.sdc** files. Because they use Tcl syntax, the constraints in **.sdc** files can be incorporated into other scripts for iterative timing analysis.

```

# -----
set_time_unit ns
set_decimal_places 3
# -----
#
create_clock -period 10.0 -waveform { 0 5.0 } clk2 -name clk2
create_clock -period 4.0 -waveform { 0 2.0 } clk1 -name clk1
# clk1 -> dir* : INPUT_MAX_DELAY = 1 ns
set_input_delay -max 1ns -clock clk1 [get_ports dir*]
# clk2 -> time* : OUTPUT_MAX_DELAY = -2 ns
set_output_delay -max -2ns -clock clk2 [get_ports time*]

```

Similar to the constraints in the Quartus II Settings File, you can make the SDC constraints part of an executable timing analysis script.

```
project_open chiptrip
create_timing_netlist
#
# Create Constraints
#
create_clock -period 10.0 -waveform { 0 5.0 } clk2 -name clk2
create_clock -period 4.0 -waveform { 0 2.0 } clk1 -name clk1
# clk1 -> dir* : INPUT_MAX_DELAY = 1 ns
set_input_delay -max 1ns -clock clk1 [get_ports dir*]
# clk2 -> time* : OUTPUT_MAX_DELAY = -2 ns
set_output_delay -max -2ns -clock clk2 [get_ports time*]
#
# Perform timing analysis for several different sets of operating conditions
#
foreach_in_collection oc [get_available_operating_conditions] {
    set_operating_conditions $oc
    update_timing_netlist
    report_timing -setup -npaths 1
    report_timing -hold -npaths 1
    report_timing -recovery -npaths 1
    report_timing -removal -npaths 1
    report_min_pulse_width -nworst 1
}
delete_timing_netlist
project_close
```

The script opens the project, creates a timing netlist, then constrains the two clocks in the design and applies input and output delay constraints. The clock settings and delay constraints are identical to those in the `.sdc` file shown in the first example. The next section of the script updates the timing netlist for the constraints and performs multi-corner timing analysis on the design.

## A Fully Iterative Scripted Flow

You can use the `::quartus::flow` Tcl package and other packages in the Quartus II Tcl API to add flow control to modify constraints and recompile your design in an automated flow. You can combine your timing constraints with the other constraints for your design, and embed them in an executable Tcl script that also iteratively compiles your design as different constraints are applied.

Each time such a modified generated script is run, it can modify the `.qsf` file and `.sdc` file for your project based on the results of iterative compilations, effectively replacing these files for the purposes of archiving and version control using industry-standard source control methods and practices.

This type of scripted flow can include automated compilation of a design, modification of design constraints, and recompilation of the design, based on how you foresee results and pre-determine next-step constraint changes in response to those results.

### Related Information

- [API Functions for Tcl](#)
- [About Quartus II Tcl Scripting](#)

# Document Revision History

Table 1-1: Document Revision History

Date	Version	Changes
June 2014	14.0.0	Formatting updates.
November 2012	12.1.0	Update Pin Planner description for task and report windows.
June 2012	12.0.0	Removed survey link.
November 2011	10.0.2	Template update.
December 2010	10.0.1	Template update.
July 2010	10.0.0	Rewrote chapter to more broadly cover all design constraint methods. Removed procedural steps and user interface details, and replaced with links to Quartus II Help.
November 2009	9.1.0	<ul style="list-style-type: none"> <li>Added two notes.</li> <li>Minor text edits.</li> </ul>
March 2009	9.0.0	<ul style="list-style-type: none"> <li>Revised and reorganized the entire chapter.</li> <li>Added section “Probing to Source Design Files and Other Quartus II Windows” on page1–2.</li> <li>Added description of node type icons (Table1–3).</li> <li>Added explanation of wildcard characters.</li> </ul>
November 2008	8.1.0	Changed to 8½” × 11” page size. No change to content.
May 2008	8.0.0	Updated Quartus II software 8.0 revision and date.

**Related Information**

[Quartus II Handbook Archive](#)

For previous versions of the Quartus II Handbook

2014.12.15

QII5V2



Subscribe



Send Feedback

FPGA design software that easily integrates into your design flow saves time and improves productivity. The Altera® Quartus® II software provides you with a command-line executable for each step of the FPGA design flow to make the design process customizable and flexible.

The benefits provided by command-line executables include:

- Command-line control over each step of the design flow
- Easy integration with scripted design flows including makefiles
- Reduced memory requirements
- Improved performance

The command-line executables are also completely interchangeable with the Quartus II GUI, allowing you to use the exact combination of tools that you prefer.

## Benefits of Command-Line Executables

The Quartus II command-line executables provide control over each step of the design flow. Each executable includes options to control commonly used software settings. Each executable also provides detailed, built-in help describing its function, available options, and settings.

Command-line executables allow for easy integration with scripted design flows. You can easily create scripts with a series of commands. These scripts can be batch-processed, allowing for integration with distributed computing in server farms. You can also integrate the Quartus II command-line executables in makefile-based design flows. These features enhance the ease of integration between the Quartus II software and other EDA synthesis, simulation, and verification software.

Command-line executables add flexibility without sacrificing the ease-of-use of the Quartus II GUI. You can use the Quartus II GUI and command-line executables at different stages in the design flow. For example, you might use the Quartus II GUI to edit the floorplan for the design, use the command-line executables to perform place-and-route, and return to the Quartus II GUI to perform debugging with the Chip Editor.

Command-line executables reduce the amount of memory required during each step in the design flow. Because each executable targets only one step in the design flow, the executables themselves are relatively compact, both in file size and the amount of memory used during processing. This memory usage reduction improves performance, and is particularly beneficial in design environments where heavy usage of computing resources results in reduced memory availability.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



**Related Information**[Using the Quartus II Executables in Shell Scripts](#)

## Introductory Example

The following introduction to command-line executables demonstrates how to create a project, fit the design, and generate programming files.

The tutorial design included with the Quartus II software is used to demonstrate this functionality. If installed, the tutorial design is found in the <Quartus II directory>/**qdesigns/fir\_filter** directory.

Before making changes, copy the tutorial directory and type the four commands shown in the introductory example below at a command prompt in the new project directory.

The <Quartus II directory>/**quartus/bin** directory must be in your `PATH` environment variable.

```
quartus_map filtref --source=filtref.bdf --family="Cyclone V"  
quartus_fit filtref --part=EP3C10F256C8 --pack_register=minimize_area  
quartus_asm filtref  
quartus_sta filtref
```

The `quartus_map filtref --source=filtref.bdf --family="Cyclone V"` command creates a new Quartus II project called **filtref** with **filtref.bdf** as the top-level file. It targets the Cyclone® V device family and performs logic synthesis and technology mapping on the design files.

The `quartus_fit filtref --part=EP3C10F256C8 --pack_register=minimize_area` command performs fitting on the **filtref** project. This command specifies an EP3C10F256C8 device, and the `--pack_register=minimize_area` option causes the Fitter to pack sequential and combinational functions into single logic cells to reduce device resource usage.

The `quartus_asm filtref` command creates programming files for the **filtref** project.

The `quartus_sta filtref` command performs basic timing analysis on the **filtref** project using the Quartus II TimeQuest Timing Analyzer, reporting worst-case setup slack, worst-case hold slack, and other measurements.

You can put the four commands from the introductory example into a batch file or script file, and run them. For example, you can create a simple UNIX shell script called **compile.sh**, which includes the code shown in the UNIX shell script example below.

```
#!/bin/sh  
PROJECT=filtref  
TOP_LEVEL_FILE=filtref.bdf  
FAMILY="Cyclone V"  
PART=EP3C10F256C8  
PACKING_OPTION=minimize_area  
quartus_map $PROJECT --source=$TOP_LEVEL_FILE --family=$FAMILY  
quartus_fit $PROJECT --part=$PART --pack_register=$PACKING_OPTION  
quartus_asm $PROJECT  
quartus_sta $PROJECT
```

Edit the script as necessary and compile your project.



**Related Information**

**[TimeQuest Timing Analyzer Quick Start Tutorial](#)**

For more information about using all of the features of the quartus\_sta executable. The TimeQuest Timing Analyzer employs Synopsys Design Constraints to fully analyze the timing of your design.

**Command-Line Scripting Help**

Help for command-line executables is available through different methods. You can access help built in to the executables with command-line options. You can use the Quartus II Command-Line and Tcl API Help browser for an easy graphical view of the help information.

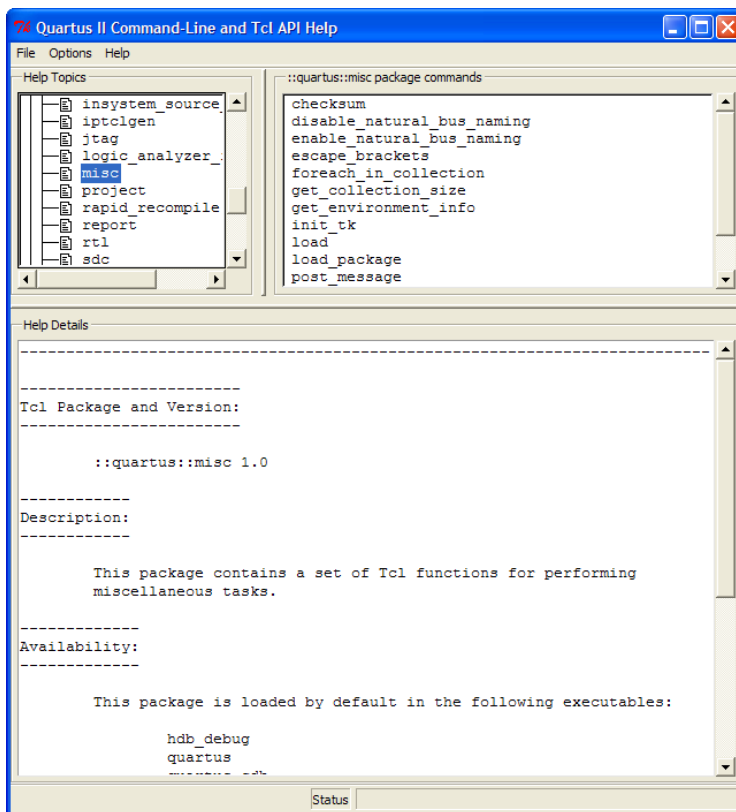
To use the Quartus II Command-Line and Tcl API Help browser, type the following command:

```
quartus_sh --qhhelp
```

This command starts the Quartus II Command-Line and Tcl API Help browser, a viewer for information about the Quartus II Command-Line executables and Tcl API.

Use the -h option with any of the Quartus II Command-Line executables to get a description and list of supported options. Use the --help=<option name> option for detailed information about each option.

**Figure 2-1: Quartus II Command-Line and Tcl API Help Browser**



## Project Settings with Command-Line Options

Command-line options are provided for many common global project settings and for performing common tasks.

You can use either of two methods to make assignments to an individual entity. If the project exists, open the project in the Quartus II GUI, change the assignment, and close the project. The changed assignment is updated in the **.qsf**. Any command-line executables that are run after this update use the updated assignment. You can also make assignments using the Quartus II Tcl scripting API. If you want to completely script the creation of a Quartus II project, choose this method.

### Related Information

- [Option Precedence](#) on page 2-4
- [Tcl Scripting](#) on page 3-1
- [QSF Reference Manual](#)

## Option Precedence

If you use command-line executables, you must be aware of the precedence of various project assignments and how to control the precedence. Assignments for a particular project exist in the Quartus II Settings File (**.qsf**) for the project. Before the **.qsf** is updated after assignment changes, the updated assignments are reflected in compiler database files that hold intermediate compilation results.

All command-line options override any conflicting assignments found in the **.qsf** or the compiler database files. There are two command-line options to specify whether the **.qsf** or compiler database files take precedence for any assignments not specified as command-line options.

Any assignment not specified as a command-line option or found in the **.qsf** or compiler database file is set to its default value.

The file precedence command-line options are `--read_settings_files` and `--write_settings_files`.

By default, the `--read_settings_files` and `--write_settings_files` options are turned on. Turning on the `--read_settings_files` option causes a command-line executable to read assignments from the **.qsf** instead of from the compiler database files. Turning on the `--write_settings_files` option causes a command-line executable to update the **.qsf** to reflect any specified options, as happens when you close a project in the Quartus II GUI.

If you use command-line executables, be aware of the precedence of various project assignments and how to control the precedence. Assignments for a particular project can exist in three places:

- The **.qsf** for the project
- The result of the last compilation, in the **/db** directory, which reflects the assignments that existed when the project was compiled
- Command-line options

The precedence for reading assignments depends on the value of the `--read_settings_files` option.

**Table 2-1: Precedence for Reading Assignments**

Option Specified	Precedence for Reading Assignments
--read_settings_files = on (Default)	Command-line options The <b>.qsf</b> for the project Project database (db directory, if it exists) Quartus II software defaults
--read_settings_files = off	Command-line options Project database (db directory, if it exists) Quartus II software defaults

The table lists the locations to which assignments are written, depending on the value of the `--write_settings_files` command-line option.

**Table 2-2: Location for Writing Assignments**

Option Specified	Location for Writing Assignments
--write_settings_files = on (Default)	<b>.qsf</b> and compiler database
--write_settings_files = off	Compiler database

The example assumes that a project named **fir\_filter** exists, and that the analysis and synthesis step has been performed (using the `quartus_map` executable).

```
quartus_fit fir_filter --pack_register=off
quartus_sta fir_filter
mv fir_filter_sta.rpt fir_filter_1_sta.rpt
quartus_fit fir_filter --pack_register=minimize_area --write_settings_files=off
quartus_sta fir_filter
mv fir_filter_sta.rpt fir_filter_2_sta.rpt
```

The first command, `quartus_fit fir_filter --pack_register=off`, runs the `quartus_fit` executable with no aggressive attempts to reduce device resource usage.

The second command, `quartus_sta fir_filter`, performs basic timing analysis for the results of the previous fit.

The third command uses the UNIX `mv` command to copy the report file output from **quartus\_sta** to a file with a new name, so that the results are not overwritten by subsequent timing analysis.

The fourth command runs **quartus\_fit** a second time, and directs it to attempt to pack logic into registers to reduce device resource usage. With the `--write_settings_files=off` option, the command-line executable does not update the **.qsf** to reflect the changed register packing setting. Instead, only the compiler database files reflect the changed setting. If the `--write_settings_files=off` option is not specified, the command-line executable updates the **.qsf** to reflect the register packing setting.

The fifth command reruns timing analysis, and the sixth command renames the report file, so that it is not overwritten by subsequent timing analysis.

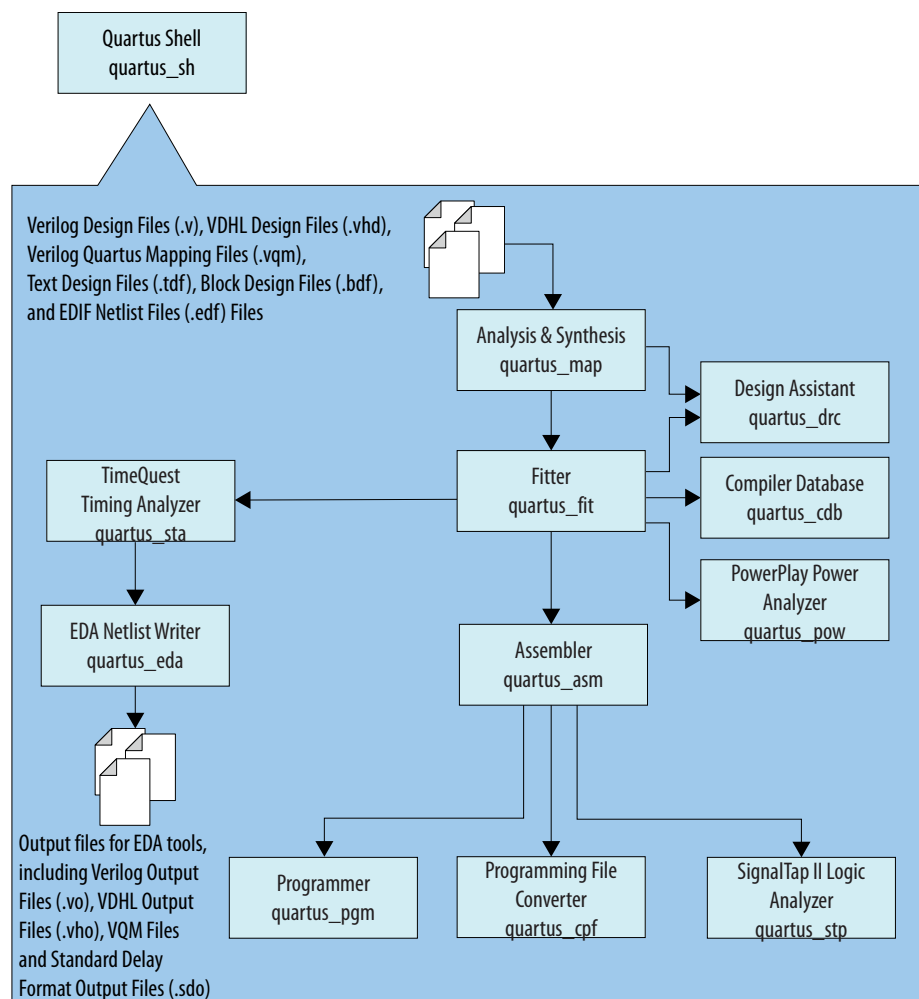
Use the options `--read_settings_files=off` and `--write_settings_files=off` (where appropriate) to optimize the way that the Quartus II software reads and updates settings files. In the following example, the `quartus_asm` executable does not read or write settings files because doing so would not change any settings for the project.

```
quartus_map filtref --source=filtref --part=EP3C10F256C8
quartus_fit filtref --pack_register=off --read_settings_files=off
quartus_asm filtref --read_settings_files=off --write_settings_files=off
```

## Compilation with `quartus_sh --flow`

The figure shows a typical Quartus II FPGA design flow using command-line executables.

Figure 2-2: Typical Design Flow



Use the `quartus_sh` executable with the `--flow` option to perform a complete compilation flow with a single command. The `--flow` option supports the smart recompile feature and efficiently sets command-line arguments for each executable in the flow.

The following example runs compilation, timing analysis, and programming file generation with a single command:

```
quartus_sh --flow compile filtref
```

**Tip:** For information about specialized flows, type `quartus_sh --help=flow` at a command prompt.

## Text-Based Report Files

Each command-line executable creates a text report file when it is run. These files report success or failure, and contain information about the processing performed by the executable.

Report file names contain the revision name and the short-form name of the executable that generated the report file, in the format `<revision>.<executable>.rpt`. For example, using the `quartus_fit` executable to place and route a project with the revision name `design_top` generates a report file named `design_top.fit.rpt`. Similarly, using the `quartus_sta` executable to perform timing analysis on a project with the revision name `fir_filter` generates a report file named `fir_filter.sta.rpt`.

As an alternative to parsing text-based report files, you can use the `::quartus::report` Tcl package.

### Related Information

[Tcl Scripting](#) on page 3-1

[About Quartus II Tcl Scripting](#)

## Using Command-Line Executables In Scripts

You can use command-line executables in scripts that control other software in addition to the Quartus II software. For example, if your design flow uses third-party synthesis or simulation software, and if you can run the other software at a command prompt, you can include those commands with Quartus II executables in a single script.

The Quartus II command-line executables include options for common global project settings and operations, but you must use a Tcl script or the Quartus II GUI to set up a new project and apply individual constraints, such as pin location assignments and timing requirements. Command-line executables are very useful for working with existing projects, for making common global settings, and for performing common operations. For more flexibility in a flow, use a Tcl script, which makes it easier to pass data between different stages of the design flow and have more control during the flow.

For example, a UNIX shell script could run other synthesis software, then place-and-route the design in the Quartus II software, then generate output netlists for other simulation software. This script shows a script that synthesizes a design with the Synopsys Synplify software, simulates the design using the Mentor Graphics ModelSim<sup>®</sup> software, and then compiles the design targeting a Cyclone V device.

```
#!/bin/sh
# Run synthesis first.
# This example assumes you use Synplify software
synplify -batch synthesize.tcl
```

```
# If your Quartus II project exists already, you can just
# recompile the design.
# You can also use the script described in a later example to
# create a new project from scratch
quartus_sh --flow compile myproject
# Use the quartus_sta executable to do fast and slow-model
# timing analysis
quartus_sta myproject --model=slow
quartus_sta myproject --model=fast
# Use the quartus_eda executable to write out a gate-level
# Verilog simulation netlist for ModelSim
quartus_eda my_project --simulation --tool=modelsim --format=verilog
# Perform the simulation with the ModelSim software
vlib cycloneV_ver
vlog -work cycloneV_ver /opt/quartusii/eda/sim_lib/cycloneV_atoms.v
vlib work
vlog -work work my_project.vo
vsim -L cycloneV_ver -t lps work.my_project
```

### Related Information

[Tcl Scripting](#) on page 3-1

[About Quartus II Tcl Scripting](#)

## Makefile Implementation

You can use the Quartus II command-line executables in conjunction with the `make` utility to automatically update files when other files they depend on change. The file dependencies and commands used to update files are specified in a text file called a makefile.

To facilitate easier development of efficient makefiles, the following “smart action” scripting command is provided with the Quartus II software:

```
quartus_sh --determine_smart_action
```

Because assignments for a Quartus II project are stored in the `.qsf`, including it in every rule results in unnecessary processing steps. For example, updating a setting related to programming file generation, which requires re-running only `quartus_asm`, modifies the `.qsf`, requiring a complete recompilation if the `.qsf` is included in every rule.

The smart action command determines the earliest command-line executable in the compilation flow that must be run based on the current `.qsf`, and generates a change file corresponding to that executable. For example, if `quartus_map` must be re-run, the smart action command creates or updates a file named `map.chg`. Thus, rather than including the `.qsf` in each makefile rule, include only the appropriate change file.

The sample makefile uses change files and the smart action command. You can copy and modify it for your own use. A copy of this example is included in the help for the makefile option, which is available by typing:

```
quartus_sh --help=makefiles

#####
# Project Configuration:
#
# Specify the name of the design (project), the Quartus II Settings
# File (.qsf), and the list of source files used.
#####
```

```
PROJECT = chiptrip
SOURCE_FILES = auto_max.v chiptrip.v speed_ch.v tick_cnt.v time_cnt.v
ASSIGNMENT_FILES = chiptrip.qpf chiptrip.qsf
#####
# Main Targets
#
# all: build everything
# clean: remove output files and database
#####
all: smart.log $(PROJECT).asm.rpt $(PROJECT).sta.rpt
clean:
    rm -rf *.rpt *.chg smart.log *.htm *.eqn *.pin *.sof *.pof db
map: smart.log $(PROJECT).map.rpt
fit: smart.log $(PROJECT).fit.rpt
asm: smart.log $(PROJECT).asm.rpt
sta: smart.log $(PROJECT).sta.rpt
smart: smart.log
#####
# Executable Configuration
#####
MAP_ARGS = --family=Stratix
FIT_ARGS = --part=EP1S20F484C6
ASM_ARGS =
STA_ARGS =
#####
# Target implementations
#####
STAMP = echo done >
$(PROJECT).map.rpt: map.chg $(SOURCE_FILES)
    quartus_map $(MAP_ARGS) $(PROJECT)
    $(STAMP) fit.chg
#####
# Project initialization
#####
$(ASSIGNMENT_FILES):
    quartus_sh --prepare $(PROJECT)
map.chg:
    $(STAMP) map.chg
fit.chg:
    $(STAMP) fit.chg
sta.chg:
    $(STAMP) sta.chg
asm.chg:
    $(STAMP) asm.chg
```

A Tcl script is provided with the Quartus II software to create or modify files that are specified as dependencies in the make rules, assisting you in makefile development. Complete information about this Tcl script and how to integrate it with makefiles is available by running the following command:

```
quartus_sh --help=determine_smart_action
```

## Common Scripting Examples

You can create scripts including command line executable to control common Quartus II processes.

### Create a Project and Apply Constraints

The command-line executables include options for common global project settings and commands. To apply constraints such as pin locations and timing assignments, run a Tcl script with the constraints in it. You can write a Tcl constraint file yourself, or generate one for an existing project.

From the Project menu, click **Generate Tcl File for Project**.

The example creates a project with a Tcl script and applies project constraints using the tutorial design files in the `<Quartus II installation directory>/qdesigns/fir_filter/` directory.

```
project_new filtref -overwrite
# Assign family, device, and top-level file
set_global_assignment -name FAMILY Cyclone
set_global_assignment -name DEVICE EP1C12F256C6
set_global_assignment -name BDF_FILE filtref.bdf
# Assign pins
set_location_assignment -to clk Pin_28
set_location_assignment -to clkx2 Pin_29
set_location_assignment -to d[0] Pin_139
set_location_assignment -to d[1] Pin_140
# Other assignments could follow
project_close
```

Save the script in a file called **setup\_proj.tcl** and type the commands illustrated in the example at a command prompt to create the design, apply constraints, compile the design, and perform fast-corner and slow-corner timing analysis. Timing analysis results are saved in two files, **filtref\_sta\_1.rpt** and **filtref\_sta\_2.rpt**.

```
quartus_sh -t setup_proj.tcl
quartus_map filtref
quartus_fit filtref
quartus_asm filtref
quartus_sta filtref --model=fast --export_settings=off
mv filtref_sta.rpt filtref_sta_1.rpt
quartus_sta filtref --export_settings=off
mv filtref_sta.rpt filtref_sta_2.rpt
```

Type the following commands to create the design, apply constraints, and compile the design, without performing timing analysis:

```
quartus_sh -t setup_proj.tcl
quartus_sh --flow compile filtref
```

The `quartus_sh --flow compile` command performs a full compilation, and is equivalent to clicking the **Start Compilation** button in the toolbar.

## Check Design File Syntax

The UNIX shell script example shown in the example assumes that the Quartus II software **fir\_filter** tutorial project exists in the current directory. You can find the **fir\_filter** project in the `<Quartus II directory>/qdesigns/fir_filter` directory unless the Quartus II software tutorial files are not installed.

The `--analyze_file` option causes the **quartus\_map** executable to perform a syntax check on each file. The script checks the exit code of the **quartus\_map** executable to determine whether there is an error during the syntax check. Files with syntax errors are added to the `FILES_WITH_ERRORS` variable, and when all files are checked, the script prints a message indicating syntax errors.

When options are not specified, the executable uses the project database values. If not specified in the project database, the executable uses the Quartus II software default values. For example, the **fir\_filter** project is set to target the Cyclone device family, so it is not necessary to specify the `--family` option.

```
#!/bin/sh
FILES_WITH_ERRORS=""
# Iterate over each file with a .bdf or .v extension
for filename in `ls *.bdf *.v`
do
```



```
# Perform a syntax check on the specified file
quartus_map fir_filter --analyze_file=$filename
# If the exit code is non-zero, the file has a syntax error
if [ $? -ne 0 ]
then
    FILES_WITH_ERRORS="$FILES_WITH_ERRORS $filename"
fi
done
if [ -z "$FILES_WITH_ERRORS" ]
then
    echo "All files passed the syntax check"
    exit 0
else
    echo "There were syntax errors in the following file(s)"
    echo $FILES_WITH_ERRORS
    exit 1
fi
```

## Create a Project and Synthesize a Netlist Using Netlist Optimizations

This example creates a new Quartus II project with a file **top.edf** as the top-level entity. The `--enable_register_retiming=on` and `--enable_wysiwyg_resynthesis=on` options cause **quartus\_map** to optimize the design using gate-level register retiming and technology remapping.

The `--part` option causes **quartus\_map** to target an EP3C10F256C8 device. To create the project and synthesize it using the netlist optimizations described above, type the command shown in this example at a command prompt.

```
quartus_map top --source=top.edf --enable_register_retiming=on
--enable_wysiwyg_resynthesis=on --part=EP3C10F256C8
```

## Archive and Restore Projects

You can archive or restore a Quartus II Archive File (**.qar**) with a single command. This makes it easy to take snapshots of projects when you use batch files or shell scripts for compilation and project management.

Use the `--archive` or `--restore` options for `quartus_sh` as appropriate. Type the command shown in the example at a command prompt to archive your project.

```
quartus_sh --archive <project name>
```

The archive file is automatically named `<project name>.qar`. If you want to use a different name, type the command with the `-output` option as shown in example the example.

```
quartus_sh --archive <project name> -output <filename>
```

To restore a project archive, type the command shown in the example at a command prompt.

```
quartus_sh --restore <archive name>
```

The command restores the project archive to the current directory and overwrites existing files.

### Related Information

#### [Managing Quartus II Projects](#)

## Perform I/O Assignment Analysis

You can perform I/O assignment analysis with a single command. I/O assignment analysis checks pin assignments to ensure they do not violate board layout guidelines. I/O assignment analysis does not require a complete place and route, so it can quickly verify that your pin assignments are correct.

```
quartus_fit --check_ios <project name> --rev=<revision name>
```

## Update Memory Contents Without Recompiling

You can use two commands to update the contents of memory blocks in your design without recompiling. Use the `quartus_cdb` executable with the `--update_mif` option to update memory contents from `.mif` or `.hexout` files. Then, rerun the assembler with the `quartus_asm` executable to regenerate the `.sof`, `.pof`, and any other programming files.

```
quartus_cdb --update_mif <project name> [--rev=<revision name>]
quartus_asm <project name> [--rev=<revision name>]
```

The example shows the commands for a DOS batch file for this example. With a DOS batch file, you can specify the project name and the revision name once for both commands. To create the DOS batch file, paste the following lines into a file called **update\_memory.bat**.

```
quartus_cdb --update_mif %1 --rev=%2
quartus_asm %1 --rev=%2
```

To run the batch file, type the following command at a command prompt:

```
update_memory.bat <project name> <revision name>
```

## Create a Compressed Configuration File

You can create a compressed configuration file in two ways. The first way is to run `quartus_cpf` with an option file that turns on compression.

To create an option file that turns on compression, type the following command at a command prompt:

```
quartus_cpf -w <filename>.opt
```

This interactive command guides you through some questions, then creates an option file based on your answers. Use `--option` to cause **quartus\_cpf** to use the option file. For example, the following command creates a compressed `.pof` that targets an EPCS64 device:

```
quartus_cpf --convert --option=<filename>.opt --device=EPCS64 <file>.sof <file>.pof
```

Alternatively, you can use the Convert Programming Files utility in the Quartus II software GUI to create a Conversion Setup File (`.cof`). Configure any options you want, including compression, then save the conversion setup. Use the following command to run the conversion setup you specified.

```
quartus_cpf --convert <file>.cof
```

## Fit a Design as Quickly as Possible

This example assumes that a project called **top** exists in the current directory, and that the name of the top-level entity is **top**. The `--effort=fast` option causes the **quartus\_fit** to use the fast fit algorithm to

increase compilation speed, possibly at the expense of reduced  $f_{MAX}$  performance. The `--one_fit_attempt=on` option restricts the Fitter to only one fitting attempt for the design.

To attempt to fit the project called **top** as quickly as possible, type the command shown at a command prompt.

```
quartus_fit top --effort=fast --one_fit_attempt=on
```

## Fit a Design Using Multiple Seeds

This shell script example assumes that the Quartus II software tutorial project called **fir\_filter** exists in the current directory (defined in the file **fir\_filter.qpf**). If the tutorial files are installed on your system, this project exists in the <Quartus II directory>/qdesigns<quartus\_version\_number>/**fir\_filter** directory.

Because the top-level entity in the project does not have the same name as the project, you must specify the revision name for the top-level entity with the `--rev` option. The `--seed` option specifies the seeds to use for fitting.

A seed is a parameter that affects the random initial placement of the Quartus II Fitter. Varying the seed can result in better performance for some designs.

After each fitting attempt, the script creates new directories for the results of each fitting attempt and copies the complete project to the new directory so that the results are available for viewing and debugging after the script has completed.

```
#!/bin/sh
ERROR_SEEDS=""
quartus_map fir_filter --rev=filtref
# Iterate over a number of seeds
for seed in 1 2 3 4 5
do
echo "Starting fit with seed=$seed"
# Perform a fitting attempt with the specified seed
quartus_fit fir_filter --seed=$seed --rev=filtref
# If the exit-code is non-zero, the fitting attempt was
# successful, so copy the project to a new directory
if [ $? -eq 0 ]
then
    mkdir ../fir_filter-seed_$seed
    mkdir ../fir_filter-seed_$seed/db
    cp * ../fir_filter-seed_$seed
    cp db/* ../fir_filter-seed_$seed/db
else
    ERROR_SEEDS="$ERROR_SEEDS $seed"
fi
done
if [ -z "$ERROR_SEEDS" ]
then
echo "Seed sweeping was successful"
exit 0
else
echo "There were errors with the following seed(s)"
echo $ERROR_SEEDS
exit 1
fi
```

**Tip:** Use Design Space Explorer II (DSE) included with the Quartus II software script (by typing `quartus_dse` at a command prompt) to improve design performance by performing automated seed sweeping.

**Related Information****Design Space Explorer II**

For more information about DSE II, type `quartus_dse --help` at a command prompt.

**The QFlow Script**

A Tcl/Tk-based graphical interface called QFlow is included with the command-line executables. You can use the QFlow interface to open projects, launch some of the command-line executables, view report files, and make some global project assignments.

The QFlow interface can run the following command-line executables:

- `quartus_map` (Analysis and Synthesis)
- `quartus_fit` (Fitter)
- `quartus_sta` (TimeQuest timing analyzer)
- `quartus_asm` (Assembler)
- `quartus_eda` (EDA Netlist Writer)

To view floorplans or perform other GUI-intensive tasks, launch the Quartus II software.

Start QFlow by typing the following command at a command prompt:

```
quartus_sh -g
```

**Tip:** The QFlow script is located in the `<Quartus II directory>/common/tcl/apps/qflow/` directory.

**Document Revision History****Table 2-3: Document Revision History**

Date	Version	Changes
December 2014	14.1.0	Updated DSE II commands.
June 2014	14.0.0	Updated formatting.
November 2013	13.1.0	Removed information about <code>-silnet qmegawiz</code> command
June 2012	12.0.0	Removed survey link.
November 2011	11.0.1	Template update.
May 2011	11.0.0	Corrected <code>quartus_qpf</code> example usage. Updated examples.

Date	Version	Changes
December 2010	10.1.0	<p>Template update.</p> <p>Added section on using a script to regenerate megafunction variations.</p> <p>Removed references to the Classic Timing Analyzer (quartus_tan).</p> <p>Removed Qflow illustration.</p>
July 2010	10.0.0	<p>Updated script examples to use quartus_sta instead of quartus_tan, and other minor updates throughout document.</p>
November 2009	9.1.0	<p>Updated Table 2–1 to add quartus_jli and quartus_jbcc executables and descriptions, and other minor updates throughout document.</p>
March 2009	9.0.0	<p>No change to content.</p>
November 2008	8.1.0	<p>Added the following sections:</p> <ul style="list-style-type: none"> <li>• “The MegaWizard Plug-In Manager” on page 2–11</li> <li>• “Command-Line Support” on page 2–12</li> <li>• “Module and Wizard Names” on page 2–13</li> <li>• “Ports and Parameters” on page 2–14</li> <li>• “Invalid Configurations” on page 2–15</li> <li>• “Strategies to Determine Port and Parameter Values” on page 2–15</li> <li>• “Optional Files” on page 2–15</li> <li>• “Parameter File” on page 2–16</li> <li>• “Working Directory” on page 2–17</li> <li>• “Variation File Name” on page 2–17</li> <li>• “Create a Compressed Configuration File” on page 2–21</li> <li>• Updated “Option Precedence” on page 2–5 to clarify how to control precedence</li> <li>• Corrected Example 2–5 on page 2–8</li> <li>• Changed Example 2–1, Example 2–2, Example 2–4, and Example 2–7 to use the EP1C12F256C6 device</li> <li>• Minor editorial updates</li> <li>• Updated entire chapter using 8½” × 11” chapter template</li> </ul>
May 2008	8.0.0	<ul style="list-style-type: none"> <li>• Updated “Referenced Documents” on page 2–20.</li> <li>• Updated references in document.</li> </ul>

**Related Information**

[Quartus II Handbook Archive](#)

For previous versions of the Quartus II Handbook

2014.06.30

QI15V2



Subscribe



Send Feedback

## Tcl Scripting

Developing and running Tcl scripts to control the Altera® Quartus® II software allows you to perform a wide range of functions, such as compiling a design or writing procedures to automate common tasks.

You can use Tcl scripts to manage a Quartus II project, make assignments, define design constraints, make device assignments, compile your design, perform timing analysis, and access reports. Tcl scripts also facilitate project or assignment migration. For example, when designing in different projects with the same prototype or development board, you can automate reassignment of pin locations in each new project. The Quartus II software can also generate a Tcl script based on all the current assignments in the project, which aids in switching assignments to another project.

The Quartus II software Tcl commands follow the EDA industry Tcl application programming interface (API) standards for command-line options. This simplifies learning and using Tcl commands. If you encounter an error with a command argument, the Tcl interpreter includes help information showing correct usage.

This chapter includes sample Tcl scripts for automating the Quartus II software. You can modify these example scripts for use with your own designs. You can find more Tcl scripts in the Design Examples section of the Support area on the Altera website.

## Tool Command Language

Tcl (pronounced “tickle”) stands for Tool Command Language, a popular scripting language that is similar to many shell scripting and high-level programming languages. It provides support for control structures, variables, network socket access, and APIs.

Tcl is the EDA industry-standard scripting language used by Synopsys, Mentor Graphics®, and Altera software. It allows you to create custom commands and works seamlessly across most development platforms.

You can create your own procedures by writing scripts containing basic Tcl commands and Quartus II API functions. You can then automate your design flow, run the Quartus II software in batch mode, or execute the individual Tcl commands interactively in the Quartus II Tcl interactive shell.

The Quartus II software supports Tcl/Tk version 8.5, supplied by the Tcl DeveloperXchange.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



**Related Information**

- **External References** on page 3-22  
For a list of recommended literature on Tcl.
- **Tcl Scripting Basics** on page 3-17  
For more information on Tcl scripting, or if you are are a Tcl beginner.
- [tcl.activestate.com/](http://tcl.activestate.com/)

## Quartus II Tcl Packages

The Quartus II Tcl commands are grouped in packages by function.

**Table 3-1: Quartus Tcl Packages**

Package Name	Package Description
<b>backannotate</b>	Back annotate assignments
<b>chip_planner</b>	Identify and modify resource usage and routing with the Chip Editor
<b>database_manager</b>	Manage version-compatible database files
<b>device</b>	Get device and family information from the device database
<b>flow</b>	Compile a project, run command-line executables and other common flows
<b>incremental_compilation</b>	Manipulate design partitions and LogicLock regions, and settings related to incremental compilation
<b>insystem_memory_edit</b>	Read and edit memory contents in Altera devices
<b>insystem_source_probe</b>	interact with the In-System Sources and Probes tool in an Altera device
<b>jtag</b>	Control the JTAG chain
<b>logic_analyzer_interface</b>	Query and modify the logic analyzer interface output pin state
<b>misc</b>	Perform miscellaneous tasks such as enabling natural bus naming, package loading, and message posting
<b>project</b>	Create and manage projects and revisions, make any project assignments including timing assignments
<b>report</b>	Get information from report tables, create custom reports
<b>rtl</b>	Traversing and querying the RTL netlist of your design
<b>sdc</b>	Specifies constraints and exceptions to the TimeQuest Timing Analyzer
<b>sdc_ext</b>	Altera-specific SDC commands
<b>simulator</b>	Configure and perform simulations
<b>sta</b>	Contains the set of Tcl functions for obtaining advanced information from the Quartus II TimeQuest Timing Analyzer

Package Name	Package Description
<b>stp</b>	Run the SignalTap® II Logic Analyzer

By default, only the minimum number of packages is loaded automatically with each Quartus II executable. This keeps the memory requirement for each executable as low as possible. Because the minimum number of packages is automatically loaded, you must load other packages before you can run commands in those packages.

Because different packages are available in different executables, you must run your scripts with executables that include the packages you use in the scripts. For example, if you use commands in the **sd\_c\_ext** package, you must use the **quartus\_sta** executable to run the script because the **quartus\_sta** executable is the only one with support for the **sd\_c\_ext** package.

The following command prints lists of the packages loaded or available to load for an executable, to the console:

```
<executable name> --tcl_eval help
```

For example, type the following command to list the packages loaded or available to load by the **quartus\_fit** executable:

```
quartus_fit --tcl_eval help
```

## Loading Packages

To load a Quartus II Tcl package, use the `load_package` command as follows:

```
load_package [-version <version number>] <package name>
```

This command is similar to the `package require` Tcl command, but you can easily alternate between different versions of a Quartus II Tcl package with the `load_package` command because of the `-version` option.

### Related Information

- [Command-Line Scripting](#) on page 2-1  
For additional information about these and other Quartus II command-line executables.

## Quartus II Tcl API Help

Access the Quartus II Tcl API Help reference by typing the following command at a system command prompt:

```
quartus_sh --qhhelp
```

This command runs the Quartus II Command-Line and Tcl API help browser, which documents all commands and options in the Quartus II Tcl API.

Quartus II Tcl help allows easy access to information about the Quartus II Tcl commands. To access the help information, type `help` at a Tcl prompt.



## Tcl Help Output

```
tcl> help
-----
-----
Available Quartus II Tcl Packages:
-----
Loaded                               Not Loaded
-----
::quartus::misc                      ::quartus::device
::quartus::old_api                   ::quartus::backannotate
::quartus::project                   ::quartus::flow
::quartus::timing_assignment          ::quartus::logiclock
::quartus::timing_report              ::quartus::report
* Type "help -tcl"
to get an overview on Quartus II Tcl usages.
```

Table 3-2: Help Options Available in the Quartus II Tcl Environment

Help Command	Description
help	To view a list of available Quartus II Tcl packages, loaded and not loaded.
help -tcl	To view a list of commands used to load Tcl packages and access command-line help.
help -pkg <package_name> [-version <version number>]	To view help for a specified Quartus II package that includes the list of available Tcl commands. For convenience, you can omit the ::quartus:: package prefix, and type help -pkg <package name>. <p>If you do not specify the -version option, help for the currently loaded package is displayed by default. If the package for which you want help is not loaded, help for the latest version of the package is displayed by default.</p> <p>Examples:</p> <pre>help -pkg ::quartus::project help -pkg project help -pkg project -version 1.0</pre>
<command_name> -h or <command_name> -help	To view short help for a Quartus II Tcl command for which the package is loaded. <p>Examples:</p> <pre>project_open -h project_open -help</pre>

Help Command	Description
<pre>package require ::quartus::&lt;pack age name&gt; [&lt;version&gt;]</pre>	<p>To load a Quartus II Tcl package with the specified version. If <code>&lt;version&gt;</code> is not specified, the latest version of the package is loaded by default.</p> <p>Example:</p> <pre>package require ::quartus::project 1.0</pre> <p>This command is similar to the <code>load_package</code> command.</p> <p>The advantage of the <code>load_package</code> command is that you can alternate freely between different versions of the same package.</p> <p>Type <code>load_package &lt;package name&gt; [-version &lt;version number&gt;]</code> to load a Quartus II Tcl package with the specified version. If the <code>-version</code> option is not specified, the latest version of the package is loaded by default.</p> <p>Example:</p> <pre>load_package ::quartus::project -version 1.0</pre>
<pre>help -cmd &lt;command_name&gt; [-version &lt;version&gt;]</pre> <p>or</p> <pre>&lt;command_name&gt; -long_ help</pre>	<p>To view complete help text for a Quartus II Tcl command.</p> <p>If you do not specify the <code>-version</code> option, help for the command in the currently loaded package version is displayed by default.</p> <p>If the package version for which you want help is not loaded, help for the latest version of the package is displayed by default.</p> <p>Examples:</p> <pre>project_open -long_help help -cmd project_open help -cmd project_open -version 1.0</pre>
<pre>help -examples</pre>	<p>To view examples of Quartus II Tcl usage.</p>
<pre>help -quartus</pre>	<p>To view help on the predefined global Tcl array that contains project information and information about the Quartus II executable that is currently running.</p>
<pre>quartus_sh --qhelp</pre>	<p>To launch the Tk viewer for Quartus II command-line help and display help for the command-line executables and Tcl API packages.</p>

The Tcl API help is also available in Quartus II online help. Search for the command or package name to find details about that command or package.

#### Related Information

- [Command-Line Scripting](#) on page 2-1  
For more information about the Tk viewer for Quartus II command-line help.

## Command-Line Options: -t, -s, and --tcl\_eval

There are three command-line options you can use with executables that support Tcl.

Table 3-3: Command-Line Options Supporting Tcl Scripting

Command-Line Option	Description
<code>--script=&lt;script file&gt; [&lt;script args&gt;]</code>	Run the specified Tcl script with optional arguments.
<code>-t &lt;script file&gt; [&lt;script args&gt;]</code>	Run the specified Tcl script with optional arguments. The <code>-t</code> option is the short form of the <code>--script</code> option.
<code>--shell</code>	Open the executable in the interactive Tcl shell mode.
<code>-s</code>	Open the executable in the interactive Tcl shell mode. The <code>-s</code> option is the short form of the <code>--shell</code> option.
<code>--tcl_eval &lt;tcl command&gt;</code>	Evaluate the remaining command-line arguments as Tcl commands. For example, the following command displays help for the project package: <code>quartus_sh --tcl_eval help -pkg project</code>

### Run a Tcl Script

Running an executable with the `-t` option runs the specified Tcl script. You can also specify arguments to the script. Access the arguments through the `argv` variable, or use a package such as **cmdline**, which supports arguments of the following form:

```
-<argument name> <argument value>
```

The **cmdline** package is included in the `<Quartus II directory>/common/tcl/packages` directory.

For example, to run a script called **myscript.tcl** with one argument, `Stratix`, type the following command at a system command prompt:

```
quartus_sh -t myscript.tcl Stratix
```

#### Related Information

[Accessing Command-Line Arguments](#) on page 3-14

### Interactive Shell Mode

Running an executable with the `-s` option starts an interactive Tcl shell. For example, to open the Quartus II TimeQuest Timing Analyzer executable in interactive shell mode, type the following command:

```
quartus_sta -s
```

Commands you type in the Tcl shell are interpreted when you click **Enter**. You can run a Tcl script in the interactive shell with the following command:

```
source <script name>
```

If a command is not recognized by the shell, it is assumed to be an external command and executed with the `exec` command.

### Evaluate as Tcl

Running an executable with the `--tcl_eval` option causes the executable to immediately evaluate the remaining command-line arguments as Tcl commands. This can be useful if you want to run simple Tcl commands from other scripting languages.

For example, the following command runs the Tcl command that prints out the commands available in the project package.

```
quartus_sh --tcl_eval help -pkg project
```

## The Quartus II Tcl Console Window

You can run Tcl commands directly in the Quartus II Tcl Console window. On the View menu, click **Utility Windows**. By default, the Tcl Console window is docked in the bottom-right corner of the Quartus II GUI. All Tcl commands typed in the Tcl Console are interpreted by the Quartus II Tcl shell.

**Note:** Some shell commands such as `cd`, `ls`, and others can be run in the Tcl Console window, with the Tcl `exec` command. However, for best results, run shell commands and Quartus II executables from a system command prompt outside of the Quartus II software GUI.

Tcl messages appear in the **System** tab (Messages window). Errors and messages written to `stdout` and `stderr` also are shown in the Quartus II Tcl Console window.

## End-to-End Design Flows

You can use Tcl scripts to control all aspects of the design flow, including controlling other software, when the other software also includes a scripting interface.

Typically, EDA tools include their own script interpreters that extend core language functionality with tool-specific commands. For example, the Quartus II Tcl interpreter supports all core Tcl commands, and adds numerous commands specific to the Quartus II software. You can include commands in one Tcl script to run another script, which allows you to combine or chain together scripts to control different tools. Because scripts for different tools must be executed with different Tcl interpreters, it is difficult to pass information between the scripts unless one script writes information into a file and another script reads it.

Within the Quartus II software, you can perform many different operations in a design flow (such as synthesis, fitting, and timing analysis) from a single script, making it easy to maintain global state information and pass data between the operations. However, there are some limitations on the operations you can perform in a single script due to the various packages supported by each executable.

There are no limitations on running flows from any executable. Flows include operations found in the Start section of the Processing menu in the Quartus II GUI, and are also documented as options for the `execute_flow` Tcl command. If you can make settings in the Quartus II software and run a flow to get your desired result, you can make the same settings and run the same flow in a Tcl script.

## Creating Projects and Making Assignments

You can easily create a script that makes all the assignments for an existing project, and then use the script at any time to restore your project settings to a known state. From the Project menu, click **Generate Tcl File for Project** to automatically generate a `.tcl` file with all of your assignments. You can source this file to recreate your project, and you can edit the file to add other commands, such as compiling the design. The file is a good starting point to learn about project management commands and assignment commands.

The following example shows how to create a project, make assignments, and compile the project. It uses the `fir_filter` tutorial design files in the `qdesigns` installation directory. Run this script in the `fir_filter` directory, with the `quartus_sh` executable.

## Create and Compile a Project

```
load_package flow
# Create the project and overwrite any settings
# files that exist
project_new fir_filter -revision filtref -overwrite
# Set the device, the name of the top-level BDF,
# and the name of the top level entity
set_global_assignment -name FAMILY Cyclone
set_global_assignment -name DEVICE EP1C6F256C6
set_global_assignment -name BDF_FILE filtref.bdf
set_global_assignment -name TOP_LEVEL_ENTITY filtref
# Add other pin assignments here
set_location_assignment -to clk Pin_G1
# compile the project
execute_flow -compile
project_close
```

**Note:** The assignments created or modified while a project is open are not committed to the Quartus II Settings File (.qsf) unless you explicitly call `export_assignments` or `project_close` (unless `-dont_export_assignments` is specified). In some cases, such as when running `execute_flow`, the Quartus II software automatically commits the changes.

### Related Information

- [Interactive Shell Mode](#) on page 3-6
- [Constraining Designs](#) on page 1-1  
For more information on making assignments.
- [QSF Reference Manual](#)  
For more information on scripting for all Quartus II project settings and assignments.

## Compiling Designs

You can run the Quartus II command-line executables from Tcl scripts. Use the included **flow** package to run various Quartus II compilation flows, or run each executable directly.

### The flow Package

The **flow** package includes two commands for running Quartus II command-line executables, either individually or together in standard compilation sequence. The `execute_module` command allows you to run an individual Quartus II command-line executable. The `execute_flow` command allows you to run some or all of the executables in commonly-used combinations. Use the **flow** package instead of system calls to run Quartus II executables from scripts or from the Quartus II Tcl Console.

### Compile All Revisions

You can use a simple Tcl script to compile all revisions in your project. Save the following script in a file called **compile\_revisions.tcl** and type the following to run it:

```
quartus_sh -t compile_revisions.tcl <project name>
```

### Compile All Revisions

```
load_package flow
project_open [lindex $quartus(args) 0]
```



```
set original_revision [get_current_revision]
foreach revision [get_project_revisions] {
    set_current_revision $revision
    execute flow -compile
}
set_current_revision $original_revision
project_close
```

## Reporting

It is sometimes necessary to extract information from the Compilation Report to evaluate results. The Quartus II Tcl API provides easy access to report data so you do not have to write scripts to parse the text report files.

If you know the exact cell or cells you want to access, use the `get_report_panel_data` command and specify the row and column names (or x and y coordinates) and the name of the appropriate report panel. You can often search for data in a report panel. To do this, use a loop that reads the report one row at a time with the `get_report_panel_row` command.

Column headings in report panels are in row 0. If you use a loop that reads the report one row at a time, you can start with row 1 to skip the row with column headings. The `get_number_of_rows` command returns the number of rows in the report panel, including the column heading row. Because the number of rows includes the column heading row, continue your loop as long as the loop index is less than the number of rows.

Report panels are hierarchically arranged and each level of hierarchy is denoted by the string “||” in the panel name. For example, the name of the Fitter Settings report panel is `Fitter||Fitter Settings` because it is in the Fitter folder. Panels at the highest hierarchy level do not use the “||” string. For example, the Flow Settings report panel is named `Flow Settings`.

The following Tcl code prints a list of all report panel names in your project. You can run this code with any executable that includes support for the report package.

### Print All Report Panel Names

```
load_package report
project_open myproject
load_report
set panel_names [get_report_panel_names]
foreach panel_name $panel_names {
    post_message "$panel_name"
}
```

## Viewing Report Data in Excel

The Microsoft Excel software is sometimes used to view or manipulate timing analysis results. You can create a Comma Separated Value (`.csv`) file from any Quartus II report to open with Excel. The following Tcl code shows a simple way to create a `.csv` file with data from the Fitter panel in a report. You could modify the script to use command-line arguments to pass in the name of the project, report panel, and output file to use. You can run this script example with any executable that supports the report package.

### Create .csv Files from Reports

```
load_package report
project_open my-project
load_report
```

```
# This is the name of the report panel to save as a CSV file
set panel_name "Fitter|Fitter Settings"
set csv_file "output.csv"
set fh [open $csv_file w]
set num_rows [get_number_of_rows -name $panel_name]
# Go through all the rows in the report file, including the
# row with headings, and write out the comma-separated data
for { set i 0 } { $i < $num_rows } { incr i } {
    set row_data [get_report_panel_row -name $panel_name \
        -row $i]
    puts $fh [join $row_data ","]
}
close $fh
unload_report
```

## Timing Analysis

The Quartus II TimeQuest Timing Analyzer includes support for industry-standard SDC commands in the **sdc** package. The Quartus II software also includes comprehensive Tcl APIs and SDC extensions for the TimeQuest Timing Analyzer in the **sta**, and **sdc\_ext** packages.

### Related Information

#### [Quartus II TimeQuest Timing Analyzer](#)

For information about how to perform timing analysis with the Quartus II TimeQuest Timing Analyzer.

## Automating Script Execution

You can configure scripts to run automatically at various points during compilation. Use this capability to automatically run scripts that perform custom reporting, make specific assignments, and perform many other tasks.

The following three global assignments control when a script is run automatically:

- `PRE_FLOW_SCRIPT_FILE` —before a flow starts
- `POST_MODULE_SCRIPT_FILE` —after a module finishes
- `POST_FLOW_SCRIPT_FILE` —after a flow finishes

A module is another term for a Quartus II executable that performs one step in a flow. For example, two modules are Analysis and Synthesis (**quartus\_map**), and timing analysis (**quartus\_sta**).

A flow is a series of modules that the Quartus II software runs with predefined options. For example, compiling a design is a flow that typically consists of the following steps (performed by the indicated module):

1. Analysis and synthesis (**quartus\_map**)
2. Fitter (**quartus\_fit**)
3. Assembler (**quartus\_asm**)
4. Timing Analyzer (**quartus\_sta**)

Other flows are described in the help for the `execute_flow` Tcl command. In addition, many commands in the Processing menu of the Quartus II GUI correspond to this design flow.

To make an assignment automatically run a script, add an assignment with the following form to the **.qsf** for your project:

```
set_global_assignment -name <assignment name> <executable>:<script name>
```

The Quartus II software runs the scripts.

```
<executable> -t <script name> <flow or module name> <project name> <revision name>
```

The first argument passed in the `argv` variable (or `quartus(args)` variable) is the name of the flow or module being executed, depending on the assignment you use. The second argument is the name of the project and the third argument is the name of the revision.

When you use the `POST_MODULE_SCRIPT_FILE` assignment, the specified script is automatically run after every executable in a flow. You can use a string comparison with the module name (the first argument passed in to the script) to isolate script processing to certain modules.

## Execution Example

To illustrate how automatic script execution works in a complete flow, assume you have a project called **top** with a current revision called **rev\_1**, and you have the following assignments in the **.qsf** for your project.

```
set_global_assignment -name PRE_FLOW_SCRIPT_FILE quartus_sh:first.tcl
set_global_assignment -name POST_MODULE_SCRIPT_FILE quartus_sh:next.tcl
set_global_assignment -name POST_FLOW_SCRIPT_FILE quartus_sh:last.tcl
```

When you compile your project, the `PRE_FLOW_SCRIPT_FILE` assignment causes the following command to be run before compilation begins:

```
quartus_sh -t first.tcl compile top rev_1
```

Next, the Quartus II software starts compilation with analysis and synthesis, performed by the **quartus\_map** executable. After the analysis and synthesis finishes, the `POST_MODULE_SCRIPT_FILE` assignment causes the following command to run:

```
quartus_sh -t next.tcl quartus_map top rev_1
```

Then, the Quartus II software continues compilation with the Fitter, performed by the **quartus\_fit** executable. After the Fitter finishes, the `POST_MODULE_SCRIPT_FILE` assignment runs the following command:

```
quartus_sh -t next.tcl quartus_fit top rev_1
```

Corresponding commands are run after the other stages of the compilation. When the compilation is over, the `POST_FLOW_SCRIPT_FILE` assignment runs the following command:

```
quartus_sh -t last.tcl compile top rev_1
```

## Controlling Processing

The `POST_MODULE_SCRIPT_FILE` assignment causes a script to run after every module. Because the same script is run after every module, you might have to include some conditional statements that restrict processing in your script to certain modules.

For example, if you want a script to run only after timing analysis, use a conditional test like the following example. It checks the flow or module name passed as the first argument to the script and executes code when the module is **quartus\_sta**.



## Restrict Processing to a Single Module

```
set module [lindex $quartus(args) 0]
if [string match "quartus_sta" $module] {
    # Include commands here that are run
    # after timing analysis
    # Use the post-message command to display
    # messages
    post_message "Running after timing analysis"
}
```

## Displaying Messages

Because of the way the Quartus II software runs the scripts automatically, you must use the `post_message` command to display messages, instead of the `puts` command. This requirement applies only to scripts that are run by the three assignments listed in “Automating Script Execution”.

### Related Information

- [The `post\_message` Command](#) on page 3-14  
For more information about this command.
- [Automating Script Execution](#) on page 3-10  
For more information on the three scripts capable of scripting-message automation.

## Other Scripting Features

The Quartus II Tcl API includes other general-purpose commands and features described in this section.

### Natural Bus Naming

The Quartus II software supports natural bus naming. Natural bus naming allows you to use square brackets to specify bus indexes in HDL without including escape characters to prevent Tcl from interpreting the square brackets as containing commands. For example, one signal in a bus named `address` can be identified as `address[0]` instead of `address\[0\]`. You can take advantage of natural bus naming when making assignments.

```
set_location_assignment -to address[10] Pin_M20
```

The Quartus II software defaults to natural bus naming. You can turn off natural bus naming with the `disable_natural_bus_naming` command. For more information about natural bus naming, type the following at a Quartus II Tcl prompt:

```
enable_natural_bus_naming -h
```

### Short Option Names

You can use short versions of command options, as long as they are unambiguous. For example, the `project_open` command supports two options: `-current_revision` and `-revision`.

You can use any of the following abbreviations of the `-revision` option:

- -r
- -re
- -rev
- -revi
- -revis
- -revisio

You can use an option as short as `-r` because in the case of the `project_open` command no other option starts with the letter `r`. However, the `report_timing` command includes the options `-recovery` and `-removal`. You cannot use `-r` or `-re` to shorten either of those options, because the abbreviation would not be unique to only one option.

## Collection Commands

Some Quartus II Tcl functions return very large sets of data that would be inefficient as Tcl lists. These data structures are referred to as collections. The Quartus II Tcl API uses a collection ID to access the collection.

There are two Quartus II Tcl commands for working with collections, `foreach_in_collection` and `get_collection_size`. Use the `set` command to assign a collection ID to a variable.

### Related Information

#### [foreach\\_in\\_collection](#)

For information about which Quartus II Tcl commands return collection IDs

## The `foreach_in_collection` Command

The `foreach_in_collection` command is similar to the `foreach` Tcl command. Use it to iterate through all elements in a collection. The following example prints all instance assignments in an open project.

### `foreach_in_collection` Example

```
set all_instance_assignments [get_all_instance_assignments -name *]
foreach_in_collection asgn $all_instance_assignments {
    # Information about each assignment is
    # returned in a list. For information
    # about the list elements, refer to Help
    # for the get-all-instance-assignments command.
    set to [lindex $asgn 2]
    set name [lindex $asgn 3]
    set value [lindex $asgn 4]
    puts "Assignment to $to: $name = $value"
}
```

## The `get_collection_size` Command

Use the `get_collection_size` command to get the number of elements in a collection. The following example prints the number of global assignments in an open project.

### `get_collection_size` Example

```
set all_global_assignments [get_all_global_assignments -name *]
set num_global_assignments [get_collection_size $all_global_assignments]
puts "There are $num_global_assignments global assignments in your project"
```

## The `post_message` Command

To print messages that are formatted like Quartus II software messages, use the `post_message` command. Messages printed by the `post_message` command appear in the **System** tab of the Messages window in the Quartus II GUI, and are written to standard at when scripts are run. Arguments for the `post_message` command include an optional message type and a required message string.

The message type can be one of the following:

- `info` (default)
- `extra_info`
- `warning`
- `critical_warning`
- `error`

If you do not specify a type, Quartus II software defaults to `info`.

With the Quartus II software in Windows, you can color code messages displayed at the system command prompt with the `post_message` command. Add the following line to your `quartus2.ini` file:

```
DISPLAY_COMMAND_LINE_MESSAGES_IN_COLOR = on
```

The following example shows how to use the `post_message` command.

```
post_message -type warning "Design has gated clocks"
```

## Accessing Command-Line Arguments

Many Tcl scripts are designed to accept command-line arguments, such as the name of a project or revision. The global variable `quartus(args)` is a list of the arguments typed on the command-line following the name of the Tcl script. The following Tcl example prints all of the arguments in the `quartus(args)` variable.

### Simple Command-Line Argument Access

```
set i 0
foreach arg $quartus(args) {
    puts "The value at index $i is $arg"
    incr i
}
```

If you copy the script in the previous example to a file named `print_args.tcl`, it displays the following output when you type the following at a command prompt.

### Passing Command-Line Arguments to Scripts

```
quartus_sh -t print_args.tcl my_project 100MHz
The value at index 0 is my_project
The value at index 1 is 100MHz
```

## The `cmdline` Package

You can use the `cmdline` package included with the Quartus II software for more robust and self-documenting command-line argument passing. The `cmdline` package supports command-line arguments with the form `-<option><value>>`.

## cmdline Package

```
package require cmdline
variable ::argv0 $::quartus(args)
set options {
  { "project.arg" "" "Project name" }
  { "frequency.arg" "" "Frequency" }
}
set usage "You need to specify options and values"
array set optshash [::cmdline::getoptions ::argv $options $usage]
puts "The project name is $optshash(project)"
puts "The frequency is $optshash(frequency)"
```

If you save those commands in a Tcl script called **print\_cmd\_args.tcl** you see the following output when you type the following command at a command prompt.

### Passing Command-Line Arguments for Scripts

```
quartus_sh -t print_cmd_args.tcl -project my_project -frequency 100MHz
The project name is my_project
The frequency is 100MHz
```

Virtually all Quartus II Tcl scripts must open a project. You can open a project, and you can optionally specify a revision name with code like the following example. The example checks whether the specified project exists. If it does, the example opens the current revision, or the revision you specify.

### Full-Featured Method to Open Projects

```
package require cmdline
variable ::argv0 $::quartus(args)
set options { \
  { "project.arg" "" "Project Name" } \
  { "revision.arg" "" "Revision Name" } \
}
array set optshash [::cmdline::getoptions ::argv0 $options]
# Ensure the project exists before trying to open it
if {[project_exists $optshash(project)]} {
  if {[string equal "" $optshash(revision)]} {
    # There is no revision name specified, so default
    # to the current revision
    project_open $optshash(project) -current_revision
  } else {
    # There is a revision name specified, so open the
    # project with that revision
    project_open $optshash(project) -revision \
      $optshash(revision)
  }
} else {
  puts "Project $optshash(project) does not exist"
  exit 1
}
# The rest of your script goes here
```

If you do not require this flexibility or error checking, you can use just the `project_open` command.

### Simple Method to Open Projects

```
set proj_name [lindex $argv 0]
project_open $proj_name
```

## The quartus() Array

The scripts in the preceding examples parsed command line arguments found in `quartus(args)`. The global `quartus()` Tcl array includes other information about your project and the current Quartus II executable that might be useful to your scripts. For information on the other elements of the `quartus()` array, type the following command at a Tcl prompt:

```
help -quartus
```

## The Quartus II Tcl Shell in Interactive Mode

This section presents how to make project assignments and then compile the finite impulse response (FIR) filter tutorial project with the `quartus_sh` interactive shell. This example assumes that you already have the `fir_filter` tutorial design files in a project directory.

To begin, type the following at the system command prompt to run the interactive Tcl shell:

```
quartus_sh -s
```

Create a new project called `fir_filter`, with a revision called `filtref` by typing the following command at a Tcl prompt:

```
project_new -revision filtref fir_filter
```

**Note:** If the project file and project name are the same, the Quartus II software gives the revision the same name as the project.

Because the revision named `filtref` matches the top-level file, all design files are automatically picked up from the hierarchy tree.

Next, set a global assignment for the device with the following command:

```
set_global_assignment -name family Cyclone
```

To learn more about assignment names that you can use with the `-name` option, refer to Quartus II Help.

**Note:** For assignment values that contain spaces, enclose the value in quotation marks.

To quickly compile a design, use the `::quartus::flow` package, which properly exports the new project assignments and compiles the design with the proper sequence of the command-line executables. First, load the package:

```
load_package flow
```

It returns the following:

```
1.0
```

To perform a full compilation of the FIR filter design, use the `execute_flow` command with the `-compile` option:

```
execute_flow -compile
```

This command compiles the FIR filter tutorial project, exporting the project assignments and running `quartus_map`, `quartus_fit`, `quartus_asm`, and `quartus_sta`. This sequence of events is the same as selecting **Start Compilation** from the Processing menu in the Quartus II GUI.

When you are finished with a project, close it with the `project_close` command.

To exit the interactive Tcl shell, type `exit` at a Tcl prompt.

## The tclsh Shell

On the UNIX and Linux operating systems, the tclsh shell included with the Quartus II software is initialized with a minimal `PATH` environment variable. As a result, system commands might not be available within the tclsh shell because certain directories are not in the `PATH` environment variable.

To include other directories in the path searched by the tclsh shell, set the `QUARTUS_INIT_PATH` environment variable before running the tclsh shell. Directories in the `QUARTUS_INIT_PATH` environment variable are searched by the tclsh shell when you execute a system command.

## Tcl Scripting Basics

The core Tcl commands support variables, control structures, and procedures. Additionally, there are commands for accessing the file system and network sockets, and running other programs. You can create platform-independent graphical interfaces with the Tk widget set.

Tcl commands are executed immediately as they are typed in an interactive Tcl shell. You can also create scripts (including the examples in this chapter) in files and run them with the Quartus II executables or with the tclsh shell.

## Hello World Example

The following shows the basic “Hello world” example in Tcl:

```
puts "Hello world"
```

Use double quotation marks to group the words `hello` and `world` as one argument. Double quotation marks allow substitutions to occur in the group. Substitutions can be simple variable substitutions, or the result of running a nested command. Use curly braces `{ }` for grouping when you want to prevent substitutions.

## Variables

Assign a value to a variable with the `set` command. You do not have to declare a variable before using it. Tcl variable names are case-sensitive.

```
set a 1
```

To access the contents of a variable, use a dollar sign (“\$”) before the variable name. The following example prints “Hello world” in a different way.

```
set a Hello  
set b world  
puts "$a $b"
```

## Substitutions

Tcl performs three types of substitution:

- Variable value substitution
- Nested command substitution
- Backslash substitution

## Variable Value Substitution

Variable value substitution, refers to accessing the value stored in a variable with a dollar sign (“\$”) before the variable name.

## Nested Command Substitution

Nested command substitution refers to how the Tcl interpreter evaluates Tcl code in square brackets. The Tcl interpreter evaluates nested commands, starting with the innermost nested command, and commands nested at the same level from left to right. Each nested command result is substituted in the outer command.

```
set a [string length foo]
```

## Backslash Substitution

Backslash substitution allows you to quote reserved characters in Tcl, such as dollar signs (“\$”) and braces (“{ }”). You can also specify other special ASCII characters like tabs and new lines with backslash substitutions. The backslash character is the Tcl line continuation character, used when a Tcl command wraps to more than one line.

```
set this_is_a_long_variable_name [string length "Hello \  
world."]
```

## Arithmetic

Use the `expr` command to perform arithmetic calculations. Use curly braces (“{ }”) to group the arguments of this command for greater efficiency and numeric precision.

```
set a 5  
set b [expr { $a + sqrt(2) }]
```

Tcl also supports boolean operators such as `&&` (AND), `||` (OR), `!` (NOT), and comparison operators such as `<` (less than), `>` (greater than), and `==` (equal to).

## Lists

A Tcl list is a series of values. Supported list operations include creating lists, appending lists, extracting list elements, computing the length of a list, sorting a list, and more.

```
set a { 1 2 3 }
```

You can use the `lindex` command to extract information at a specific index in a list. Indexes are zero-based. You can use the `index end` to specify the last element in the list, or the `index end-<n>` to count from the end of the list. For example to print the second element (at index 1) in the list stored in `a` use the following code.

```
puts [lindex $a 1]
```

The `llength` command returns the length of a list.

```
puts [llength $a]
```

The `lappend` command appends elements to a list. If a list does not already exist, the list you specify is created. The list variable name is not specified with a dollar sign (“\$”).

```
lappend a 4 5 6
```

## Arrays

Arrays are similar to lists except that they use a string-based index. Tcl arrays are implemented as hash tables. You can create arrays by setting each element individually or with the `array set` command.

To set an element with an index of `Mon` to a value of `Monday` in an array called `days`, use the following command:

```
set days(Mon) Monday
```

The `array set` command requires a list of index/value pairs. This example sets the array called `days`:

```
array set days { Sun Sunday Mon Monday Tue Tuesday \
  Wed Wednesday Thu Thursday Fri Friday Sat Saturday }
```

```
set day_abbreviation Mon
puts $days($day_abbreviation)
```

Use the `array names` command to get a list of all the indexes in a particular array. The index values are not returned in any specified order. The following example is one way to iterate over all the values in an array.

```
foreach day [array names days] {
  puts "The abbreviation $day corresponds to the day \
name $days($day)"
}
```

Arrays are a very flexible way of storing information in a Tcl script and are a good way to build complex data structures.

## Control Structures

Tcl supports common control structures, including if-then-else conditions and `for`, `foreach`, and `while` loops. The position of the curly braces as shown in the following examples ensures the control structure commands are executed efficiently and correctly. The following example prints whether the value of variable `a` is positive, negative, or zero.

### If-Then-Else Structure

```
if { $a > 0 } {
  puts "The value is positive"
} elseif { $a < 0 } {
  puts "The value is negative"
} else {
  puts "The value is zero"
}
```



The following example uses a `for` loop to print each element in a list.

### For Loop

```
set a { 1 2 3 }
for { set i 0 } { $i < [llength $a] } { incr i } {
    puts "The list element at index $i is [lindex $a $i]"
}
```

The following example uses a `foreach` loop to print each element in a list.

### foreach Loop

```
set a { 1 2 3 }
foreach element $a {
    puts "The list element is $element"
}
```

The following example uses a `while` loop to print each element in a list.

### while Loop

```
set a { 1 2 3 }
set i 0
while { $i < [llength $a] } {
    puts "The list element at index $i is [lindex $a $i]"
    incr i
}
```

You do not have to use the `expr` command in boolean expressions in control structure commands because they invoke the `expr` command automatically.

## Procedures

Use the `proc` command to define a Tcl procedure (known as a subroutine or function in other scripting and programming languages). The scope of variables in a procedure is local to the procedure. If the procedure returns a value, use the `return` command to return the value from the procedure. The following example defines a procedure that multiplies two numbers and returns the result.

### Simple Procedure

```
proc multiply { x y } {
    set product [expr { $x * $y }]
    return $product
}
```

The following example shows how to use the `multiply` procedure in your code. You must define a procedure before your script calls it.

### Using a Procedure

```
proc multiply { x y } {
    set product [expr { $x * $y }]
    return $product
}
set a 1
```

```
set b 2
puts [multiply $a $b]
```

Define procedures near the beginning of a script. If you want to access global variables in a procedure, use the `global` command in each procedure that uses a global variable.

### Accessing Global Variables

```
proc print_global_list_element { i } {
    global my_data
    puts "The list element at index $i is [lindex $my_data $i]"
}
set my_data { 1 2 3}
print_global_list_element 0
```

## File I/O

Tcl includes commands to read from and write to files. You must open a file before you can read from or write to it, and close it when the read and write operations are done. To open a file, use the `open` command; to close a file, use the `close` command. When you open a file, specify its name and the mode in which to open it. If you do not specify a mode, Tcl defaults to read mode. To write to a file, specify `w` for write mode.

### Open a File for Writing

```
set output [open myfile.txt w]
```

Tcl supports other modes, including appending to existing files and reading from and writing to the same file.

The `open` command returns a file handle to use for read or write access. You can use the `puts` command to write to a file by specifying a filehandle.

### Write to a File

```
set output [open myfile.txt w]
puts $output "This text is written to the file."
close $output
```

You can read a file one line at a time with the `gets` command. The following example uses the `gets` command to read each line of the file and then prints it out with its line number.

### Read from a File

```
set input [open myfile.txt]
set line_num 1
while { [gets $input line] >= 0 } {
    # Process the line of text here
    puts "$line_num: $line"
    incr line_num
}
close $input
```

## Syntax and Comments

Arguments to Tcl commands are separated by white space, and Tcl commands are terminated by a newline character or a semicolon. You must use backslashes when a Tcl command extends more than one line.

Tcl uses the hash or pound character (#) to begin comments. The # character must begin a comment. If you prefer to include comments on the same line as a command, be sure to terminate the command with a semicolon before the # character. The following example is a valid line of code that includes a `set` command and a comment.

```
set a 1;# Initializes a
```

Without the semicolon, it would be an invalid command because the `set` command would not terminate until the new line after the comment.

The Tcl interpreter counts curly braces inside comments, which can lead to errors that are difficult to track down. The following example causes an error because of unbalanced curly braces.

```
# if { $x > 0 } {
if { $y > 0 } {
    # code here
}
```

## External References

For more information about Tcl, refer to the following sources:

- Practical Programming in Tcl and Tk, Brent B. Welch
- Tcl and the TK Toolkit, John Ousterhout
- Effective Tcl/TK Programming, Michael McLennan and Mark Harrison

### Related Information

- [Quartus II Tcl Examples](#)  
For Quartus II Tcl example scripts
- [tcl.activestate.com](http://tcl.activestate.com)  
Tcl Developer Xchange

## Document Revision History

Table 3-4: Document Revision History

Date	Version	Changes
June 2014	14.0.0	Updated the format.
June 2012	12.0.0	<ul style="list-style-type: none"> <li>• Removed survey link.</li> </ul>

Date	Version	Changes
November 2011	11.0.1	<ul style="list-style-type: none"> <li>• Template update</li> <li>• Updated supported version of Tcl in the section “Tool Command Language.”</li> <li>• minor editorial changes</li> </ul>
May 2011	11.0.0	Minor updates throughout document.
December 2010	10.1.0	Template update Updated to remove tcl packages used by the Classic Timing Analyzer
July 2010	10.0.0	Minor updates throughout document.
November 2009	9.1.0	<ul style="list-style-type: none"> <li>• Removed LogicLock example.</li> <li>• Added the incremental_compilation, insystem_source_probe, and rtl packages to Table 3-1 and Table 3-2.</li> <li>• Added quartus_map to table 3-2.</li> </ul>
March 2009	9.0.0	<ul style="list-style-type: none"> <li>• Removed the “EDA Tool Assignments” section</li> <li>• Added the section “Compile All Revisions” on page 3–9</li> <li>• Added the section “Using the tclsh Shell” on page 3–20</li> </ul>
November 2008	8.1.0	Changed to 8½” × 11” page size. No change to content.
May 2008	8.0.0	Updated references.

**Related Information**

[Quartus II Handbook Archive](#)

For previous versions of the Quartus II Handbook.

2014.12.15

QI15V2



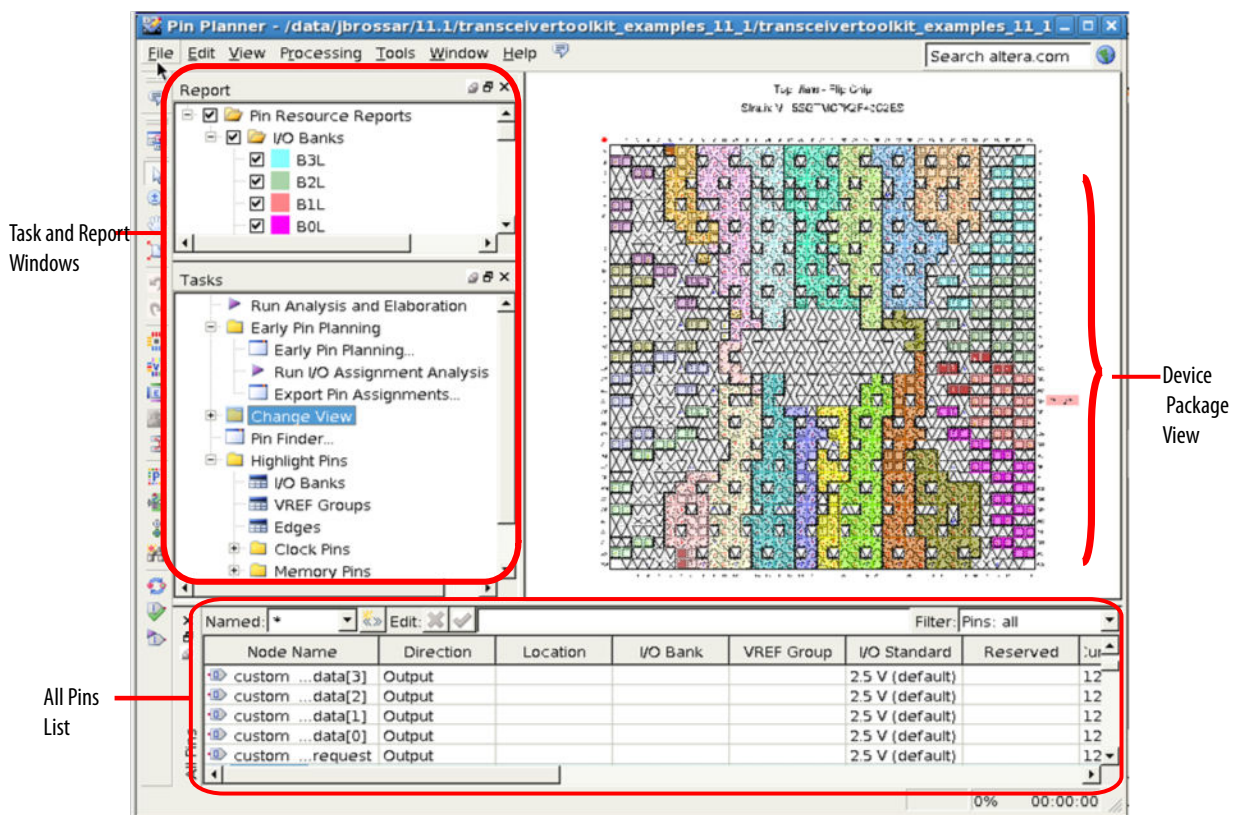
Subscribe



Send Feedback

This document describes efficient planning and assignment of the I/O pins in your target device. You should consider I/O standards, pin placement rules, and your PCB characteristics early in the design phase.

**Figure 4-1: Pin Planner GUI**



**Table 4-1: Quartus II I/O Pin Planning Tools**

I/O Planning Task	Click to Access
Edit, validate, or export pin assignments	<b>Assignments &gt; Pin Planner</b>

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



I/O Planning Task	Click to Access
View tailored pin planning advice	<b>Tools &gt; Advisors &gt; Pin Advisor</b>
Validate pin assignments against design rules	<b>Processing &gt; Start I/O Assignment Analysis</b>

For more information about special pin assignment features for the Arria 10 SoC devices, refer to *Instantiating the HPS Component* in the *Arria 10 Hard Processor System Technical Reference Manual*.

#### Related Information

- [Instantiating the HPS Component](#)

## I/O Planning Overview

You should plan and assign I/O pins in your design for compatibility with your target device and PCB characteristics. Plan I/O pins early to reduce design iterations and develop an accurate PCB layout sooner. You can assign expected nodes not yet defined in design files, including interface IP core signals, and then generate a top-level file. Specify interfaces for memory, high-speed I/O, device configuration, and debugging tools in your top-level file. The top-level file instantiates the next level of design hierarchy and includes interface port information.

Use the Pin Planner to view, assign, and validate device I/O pin logic and properties. Alternatively, you can enter I/O assignments in a Tcl script, or directly in HDL code. The Pin Planner Task window provides one-click access to I/O planning steps. You can filter and search the nodes in the design. You can define custom groups of pins for assignment. Instantly locate and highlight specific pin types for assignment or evaluation, such as I/O banks, VREF groups, edges, DQ/DQS pins, hard memory interface pins, PCIe hard IP interface pins, hard processor system pins, and clock region input pins. Assign design elements, I/O standards, interface IP, and other properties to the device I/O pins by name or by drag and drop. You can then generate a top-level design file for I/O validation.

Use live I/O check to verify the legality of pin assignments in real time. Use I/O assignment analysis to run a full I/O analysis against VCCIO, VREF, electromigration (current density), Simultaneous Switching Output (SSO), drive strength, I/O standard, PCI\_IO clamp diode, and I/O pin direction compatibility rules.

## Basic I/O Planning Flow

The following steps describe the basic flow for assigning and verifying I/O pin assignments:

1. Click **Assignments > Device** and select a target device that meets your logic, performance, and I/O requirements. Consider and specify I/O standards, voltage and power supply requirements, and available I/O pins.
2. Click **Assignments > Pin Planner**.
3. To setup a top-level HDL wrapper file that defines early port and interface information for your design, click **Early Pin Planning** in the Tasks pane.

- a. Click **Import IP Core** to import any defined IP core, and then assign signals to the interface IP nodes.
- b. Click **Set Up Top-Level File** and assign user nodes to device pins. User nodes become virtual pins in the top-level file and are not assigned to device pins.
- c. Click **Generate Top-Level File**. Use this file to validate I/O assignments.
4. Click **Run I/O Assignment Analysis** in the Tasks pane to validate any early assignments and generate a synthesized design netlist.
5. Assign I/O properties to match your device and PCB characteristics, including assigning logic, I/O standards, output loading, slew rate, and current strength.
6. Click **Run I/O Assignment Analysis** in the Tasks pane to validate assignments and generate a synthesized design netlist. Correct any problems reported.
7. Click **Processing > Start Compilation**. During compilation, the Quartus II® software runs I/O assignment analysis and advanced I/O timing analysis.

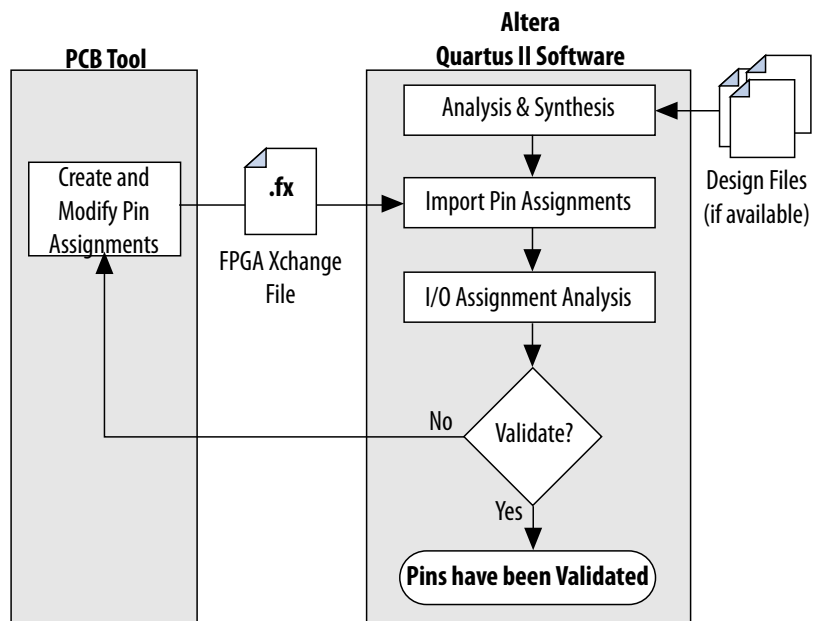
## Integrating PCB Design Tools

You can integrate PCB design tools into your work flow to help correctly map pin assignments to the symbols your system circuit schematics and board layout. The Quartus II software integrates with board layout tools by allowing import and export of pin assignment information in Quartus II Settings Files (.qsf), Pin-Out File (.pin), and FPGA Xchange-Format File (.fx) files. You can integrate PCB tools in the the following ways:

**Table 4-2: Integrating PCB Design Tools**

PCB Tool Integration	Supported PCB Tool
Define and validate I/O assignments in the Pin Planner, and then export the assignments to the PCB tool for validation	Mentor Graphics® I/O Designer Cadence Allegro
Define I/O assignments in your PCB tool, and then import the assignments into the Pin Planner for validation	Mentor Graphics® I/O Designer Cadence Allegro

Figure 4-2: PCB Tool Integration



For more information about incorporating PCB design tools, refer to the *Cadence PCB Design Tools Support* and *Mentor Graphics PCB Design Tools Support* chapters in volume 2 of the *Quartus II Handbook*.






#### Related Information

- [Mentor Graphics PCB Design Tools Support](#) on page 7-1



## Altera Device Terms

The following terms describe Altera device and I/O structures:

Terms	Description	Diagram
Device Package (BGA example)	Ceramic or plastic heat sink surface mounted with FPGA die and I/O pins or solder balls. In a wire bond BGA example, copper wires connect the bond pads to the solder balls of the package. Click <b>View &gt; Show &gt; Package Top</b> or <b>View &gt; Show &gt; Package Bottom</b> in Pin Planner	
I/O Bank	I/O pins are grouped in I/O banks for assignment of I/O standards. Each numbered bank has its own voltage source pins, called VCCIO pins, for high I/O performance. The specified VCCIO pin voltage is between 1.5 V and 3.3 V. Each bank supports multiple pins with different I/O standards. All pins in a bank must use the same VCCIO signal. Click <b>View &gt; Show &gt; I/O Banks</b> in Pin Planner.	
I/O Pin	A wire lead or small solder ball on the package bottom or periphery. Each pin has an alphanumeric row and column number. I, O, Q, S, X, and Z are never used. The alphabet is repeated and prefixed with the letter A when exceeded. All I/O pins display by default.	
Pad	I/O pins are connected to pads located on the perimeter of the top metal layer of the silicon die. Each pad is numbered with an ID starting at 0, and increments by one in a counterclockwise direction around the device. Click <b>View &gt; Pad View</b> in Pin Planner.	
VREF Pin Group	A group of pins including one dedicated VREF pin required by voltage-referenced I/O standards. A VREF group contains a smaller number of pins than an I/O bank. This maintains the signal integrity of the VREF pin. One or more VREF groups exist in an I/O bank. The pins in a VREF group share the same VCCIO and VREF voltages. Click <b>View &gt; Show &gt; Show VREF Groups</b> in Pin Planner..	

## Assigning I/O Pins

Use the Pin Planner to visualize, modify, and validate I/O assignments in a graphical representation of the target device. To assign I/O pins, locate the device I/O pin(s) for assignment, enter properties for the pin(s), and validate the legality of the assignment. You can increase the accuracy of I/O assignment analysis by reserving specific device pins to accommodate undefined but expected I/O.

To assign I/O pins in the Pin Planner, follow these steps:

1. Open a Quartus II project, and then click **Assignments > Pin Planner**.
2. (Optional) To validate I/O pin assignments in real time, click **Processing > Enable Live I/O Check**.
3. Click **Processing > Start Analysis & Elaboration** to elaborate the design and display **All Pins** in the device view.
4. To locate or highlight pins for assignment, click **Pin Finder** or a pin type under **Highlight Pins** in the Tasks pane.
5. (Optional) To define a custom group of nodes for assignment, select one or more nodes in the **Groups** or **All Pins** list, and then click **Create Group**.
6. Enter assignments of logic, I/O standards, interface IP, and properties for device I/O pins in the **All Pins** spreadsheet, or by drag and drop into the package view.
7. To assign properties to differential pin pairs, click **Show Differential Pin Pair Connections**. A red connection line appears between positive (p) and negative (n) differential pins.
8. (Optional) To create board trace model assignments, right-click an output or bidirectional pin, and then click **Board Trace Model**. For differential I/O standards, the board trace model uses a differential pin pair with two symmetrical board trace models. Specify board trace parameters on the positive end of the differential pin pair. The assignment applies to the corresponding value on the negative end of the differential pin pair.
9. To run a full I/O assignment analysis, click **Run I/O Assignment Analysis**. The Fitter reports analysis results. Only reserved pins are analyzed prior to design synthesis.

## Assigning to Exclusive Pin Groups

You can designate groups of pins for exclusive assignment. When you assign pins to an **Exclusive I/O Group**, the Fitter does not place the signals in the same I/O bank with any other exclusive I/O group. For example, if you have a set of signals assigned exclusively to `group_a`, and another set of signals assigned to `group_b`, the Fitter ensures placement of each group in different I/O banks.

## Assigning Slew Rate and Drive Strength

You can designate the device pin slew rate and drive strength. These properties affect the pin's outgoing signal integrity. Use either the **Slew Rate** or **Slow Slew Rate** assignment to adjust the drive strength of a pin with the **Current Strength** assignment. The slew rate and drive strength apply during live I/O check and I/O assignment analysis.


## Assigning Differential Pins

When you use the Pin Planner to assign a differential I/O standard to a single-ended top-level pin in your design, it automatically recognizes the negative pin as part of the differential pin pair assignment and creates the negative pin for you. The Quartus II software writes the location assignment for the negative pin to the `.qsf`; however, the I/O standard assignment is not added to the `.qsf` for the negative pin of the differential pair.

The following example shows a design with `lvds_in` top-level pin, to which you assign a differential I/O standard. The Pin Planner automatically creates the differential pin, `lvds_in(n)` to complete the differential pin pair.


**Note:** If you have a single-ended clock that feeds a PLL, assign the pin only to the positive clock pin of a differential pair in the target device. Single-ended pins that feed a PLL and are assigned to the negative clock pin device cause the design to not fit.

**Figure 4-3: Creating a Differential Pin Pair in the Pin Planner**



If your design contains a large bus that exceeds the pins available in a particular I/O bank, you can use edge location assignments to place the bus. Edge location assignments improve the circuit board routing ability of large buses, because they are close together near an edge. The following shows Altera device package edges.

**Figure 4-4: Die View and Package View of the Four Edges on an Altera Device**



## Overriding I/O Placement Rules on Differential Pins

Each device family has predefined I/O placement rules. The I/O placement rules ensure that noisy signals do not corrupt neighboring signals. For example, I/O placement rules define the allowed placement of single-ended I/O with respect to differential pins, or how many output and bidirectional pins can be placed within a VREF group when using voltage referenced input standards. You can use the `IO_MAXIMUM_TOGGLE_RATE` assignment to override I/O placement rules on pins, such as for system reset pins that do not switch during normal design activity. Setting a value of 0 MHz for this assignment causes the Fitter to recognize the pin at a DC state throughout device operation. The Fitter excludes the assigned pin from placement rule analysis. Do not assign an `IO_MAXIMUM_TOGGLE_RATE` of 0 MHz to any actively switching pin or your design may not function as intended.

## Entering Pin Assignments with Tcl Commands

You can use Tcl scripts to apply pin assignments rather than using the GUI. Enter individual Tcl commands in the Tcl Console, or type the following to apply the assignments contained in a Tcl script:

### Example 4-1: Applying Tcl Script Assignments

```
quartus_sh -t <my_tcl_script>.tcl
```

The following example shows use of the `set_location_assignment` and `set_instance_assignment` Tcl commands to assign a pin to a specific location, I/O standard, and drive strength.

### Example 4-2: Scripted Pin Assignment

```
set_location_assignment PIN M20 -to address[10]  
set_instance_assignment -name IO_STANDARD "2.5 V" -to address[10]  
set_instance_assignment -name  
CURRENT_STRENGTH_NEW "MAXIMUM CURRENT" -to address[10]
```

#### Related Information

[Tcl Scripting](#) on page 3-1

#### API Functions

## Entering Pin Assignments in HDL Code

You can use synthesis attributes or low-level I/O primitives to embed I/O pin assignments directly in your HDL code. When you analyze and synthesize the HDL code, the information is converted into the appropriate I/O pin assignments. You can use either of the following methods to specify pin-related assignments with HDL code:

- Assigning synthesis attributes for signal names that are top-level pins
- Using low-level I/O primitives, such as `ALT_BUF_IN`, to specify input, output, and differential buffers, and for setting parameters or attributes



## Using Synthesis Attributes

The Quartus II software translates synthesis attributes into standard assignments during compilation. The assignments appear in the Pin Planner. If you modify or delete these assignments in the Pin Planner and then recompile your design, the Pin Planner changes override the synthesis attributes. Quartus II synthesis supports the `chip_pin`, `useioff`, and `altera_attribute` synthesis attributes.

Use the `chip_pin` and `useioff` synthesis attributes to create pin location assignments and to assign **Fast Input Register**, **Fast Output Register**, and **Fast Output Enable Register** logic options. The following examples use the `chip_pin` and `useioff` attributes to embed location and **Fast Input Register** logic option assignments in Verilog HDL and VHDL design files.

### Example 4-3: Verilog HDL Synthesis Attribute

```
input my_pin1 /* synthesis altera_attribute = "-name FAST_INPUT_REGISTER ON;
-name IO_STANDARD \"2.5 V\" " */ ;
```

### Example 4-4: VHDL Synthesis Attribute

```
VHDL Example
entity my_entity is
  port(
    my_pin1: in std_logic
  );
end my_entity;

architecture rtl of my_entity is
  attribute useioff : boolean;
  attribute useioff of my_pin1 : signal is true;
  attribute chip_pin : string;
  attribute chip_pin of my_pin1 : signal is "C1";
begin -- The architecture body
end rtl;
```

Use the `altera_attribute` synthesis attribute to create other pin-related assignments in your HDL code. The `altera_attribute` attribute is understood only by Quartus II integrated synthesis and supports all types of instance assignments. The following examples use the `altera_attribute` attribute to embed **Fast Input Register** logic option assignments and I/O standard assignments in both a Verilog HDL and a VHDL design file.

### Example 4-5: Verilog HDL Synthesis Attribute

```
input my_pin1 /* synthesis chip_pin = "C1" useioff = 1 */;
```

### Example 4-6: VHDL Synthesis Attribute

```
entity my_entity is
  port(
    my_pin1: in std_logic
  );
end my_entity;
```

```

architecture rtl of my_entity is
begin

attribute altera_attribute : string;
attribute altera_attribute of my_pin1: signal is "-name FAST_INPUT_REGISTER
ON;
-- The architecture body
end rtl;

```

## Using Low-Level I/O Primitives

You can alternatively enter I/O pin assignments using low-level I/O primitives. You can assign pin locations, I/O standards, drive strengths, slew rates, and on-chip termination (OCT) value assignments. You can also use low-level differential I/O primitives to define both positive and negative pins of a differential pair in the HDL code for your design.

Primitive-based assignments do not appear in the Pin Planner until after you perform a full compilation and back-annotate pin assignments (**Assignments > Back Annotate Assignments**).

### Related Information

[Designing with Low Level Primitives User Guide](#)

## Importing and Exporting I/O Pin Assignments

The Quartus II software supports transfer of I/O pin assignments across projects, or for analysis in third-party PCB tools. You can import or export I/O pin assignments in the following ways:

**Table 4-3: Importing and Exporting I/O Pin Assignments**

	Import Assignments	Export Assignments
<b>Scenario</b>	<ul style="list-style-type: none"> <li>From your PCB design tool or spreadsheet into Pin Planner during early pin planning or after optimization in PCB tool</li> <li>From another Quartus II project with common constraints</li> </ul>	<ul style="list-style-type: none"> <li>From Quartus II project for optimization in a PCB design tool</li> <li>From Quartus II project for spreadsheet analysis or use in scripting assignments</li> <li>From Quartus II project for import into another Quartus II project with similar constraints</li> </ul>
<b>Command</b>	<b>Assignments &gt; Import Assignments</b>	<b>Assignments &gt; Export Assignments</b>
<b>File formats</b>	<b>.qsf, .esf, .acf, .csv, .txt, .sdc</b>	<b>.pin, .fx, .csv, .tcl, .qsf</b>
<b>Notes</b>	N/A	Exported <b>.csv</b> files retain column and row order and format. Do not modify the row of column headings if importing the <b>.csv</b> file

## Importing and Exporting for PCB Tools

The Pin Planner supports import and export of assignments with PCB tools. You can export valid assignments as a **.pin** file for analysis in other supported PCB tools. You can also import optimized assignment from supported PCB tools. The **.pin** file contains pin name, number, and detailed properties.

Mentor Graphics I/O Designer requires you to generate and import both an **.fx** and a **.pin** file to transfer assignments. However, the Quartus II software requires only the **.fx** to import pin assignments from I/O Designer.

**Table 4-4: Contents of .pin File**

File Column Name	Description
Pin Name/Usage	The name of the design pin, or whether the pin is GND or V <sub>CC</sub> pin
Location	The pin number of the location on the device package
Dir	The direction of the pin
I/O Standard	The name of the I/O standard to which the pin is configured
Voltage	The voltage level that is required to be connected to the pin
I/O Bank	The I/O bank to which the pin belongs
User Assignment	Y or N indicating if the location assignment for the design pin was user assigned (Y) or assigned by the Fitter (N)

### Related Information

[Pin-Out Files for Altera Devices](#)

[Mentor Graphics PCB Tools Support](#) on page 7-1

## Migrating Assignments to Another Target Device

You can migrate compatible pin assignments from one target device to another. You can migrate to a different density and the same device package. You can also migrate between device packages with different densities and pin counts. Click **View > Pin Migration Window** to verify whether your pin assignments are compatible with migration to a different Altera device.

The Quartus II software ignores invalid assignments and generates an error message during compilation. After evaluating migration compatibility, modify any incompatible assignments, and then click **Export** to export the assignments to another project.

Figure 4-5: Device Migration Compatibility (AC24 does not exist in migration device)

Pin Migration View

Current Device: EP2530F672C4

	Pin Number	Migration Result				Migration Devices											
		Pin Function	I/O Bank	VREF Group	Clock Pin	EP2530F672C4				EP2515F672C4				EP2560F672C4			
						Pin Function	I/O Bank	VREF Group	Clock Pin	Pin Function	I/O Bank	VREF Group	Clock Pin	Pin Function	I/O Bank	VREF Group	Clock Pin
87	PIN_AC11	VREFB7N0	7	B7_NO		VREFB7N0	7	B7_NO		VREFB7N0	7	B7_NO		VREFB7N0	7	B7_NO	
88	PIN_AC12	Column I/O	10	B7_NO	Yes	Column I/O	10	B7_NO	Yes	Column I/O	10	B7_NO	Yes	Column I/O	10	B7_NO	Yes
89	PIN_AC13	Column I/O	7	B7_NO	Yes	Column I/O	7	B7_NO	Yes	Column I/O	7	B7_NO	Yes	Column I/O	7	B7_NO	Yes
90	PIN_AC14	Column I/O	8	B8_N1	Yes	Column I/O	8	B8_N1	Yes	Column I/O	8	B8_N1	Yes	Column I/O	8	B8_N2	Yes
91	PIN_AC15	NC				Column I/O	8	B8_N1		NC				Column I/O	12	B8_N2	Yes
92	PIN_AC16	VREFB8N1	8	B8_N1		VREFB8N1	8	B8_N1		VREFB8N1	8	B8_N1		VREFB8N2	8	B8_N2	
93	PIN_AC17	Column I/O	8	B8_N1		Column I/O	8	B8_N1		Column I/O	8	B8_N1		Column I/O	8	B8_N1	
94	PIN_AC18	Column I/O	8	B8_NO		Column I/O	8	B8_NO		Column I/O	8	B8_N1		Column I/O	8	B8_NO	
95	PIN_AC19	Column I/O	8	B8_NO		Column I/O	8	B8_NO		Column I/O	8	B8_NO		Column I/O	8	B8_NO	
96	PIN_AC20	Column I/O	8	B8_NO		Column I/O	8	B8_NO		Column I/O	8	B8_NO		Column I/O	8	B8_NO	
97	PIN_AC21	Column I/O	8	B8_NO		Column I/O	8	B8_NO		Column I/O	8	B8_NO		Column I/O	8	B8_NO	
98	PIN_AC22	VREFB8N0	8	B8_NO		VREFB8N0	8	B8_NO		VREFB8N0	8	B8_NO		VREFB8N0	8	B8_NO	
99	PIN_AC23	VREFB1N2	1	B1_N2		Column I/O	8	B8_NO		NC				VREFB1N2	1	B1_N2	
100	PIN_AC24	NC				Row I/O	1	B1_N1		NC				Row I/O	1	B1_N1	
101	PIN_AC25	NC				Row I/O	1	B1_N1		NC				Row I/O	1	B1_N1	
102	PIN_AC26	VCCIO1	1			VCCIO1	1			VCCIO1	1			VCCIO1	1		
103	PIN_AD1	NC				Row I/O	6	B6_NO		NC				Row I/O	6	B6_N1	
104	PIN_AD2	NC				Row I/O	6	B6_NO		NC				Row I/O	6	B6_N1	
105	PIN_AD3	Column I/O	7	B7_N1		Column I/O	7	B7_N1		Column I/O	7	B7_N1		Column I/O	7	B7_N2	
106	PIN_AD4	Column I/O	7	B7_N1		Column I/O	7	B7_N1		Column I/O	7	B7_N1		Column I/O	7	B7_N2	
107	PIN_AD5	Column I/O	7	B7_N1		Column I/O	7	B7_N1		Column I/O	7	B7_N1		Column I/O	7	B7_N2	
108	PIN_AD6	Column I/O	7	B7_N1		Column I/O	7	B7_N1		Column I/O	7	B7_N1		Column I/O	7	B7_N2	
109	PIN_AD7	Column I/O	7	B7_N1		Column I/O	7	B7_N1		Column I/O	7	B7_N1		Column I/O	7	B7_N1	
110	PIN_AD8	Column I/O	7	B7_NO		Column I/O	7	B7_NO		Column I/O	7	B7_NO		Column I/O	7	B7_N1	
111	PIN_AD9	Column I/O	7	B7_NO		Column I/O	7	B7_NO		Column I/O	7	B7_NO		Column I/O	7	B7_N1	
112	PIN_AD10	Column I/O	7	B7_NO		Column I/O	7	B7_NO		Column I/O	7	B7_NO		Column I/O	7	B7_NO	
113	PIN_AD11	Column I/O	7	B7_NO		Column I/O	7	B7_NO		Column I/O	7	B7_NO		Column I/O	7	B7_NO	
114	PIN_AD12	Column I/O	10	B7_NO	Yes	Column I/O	10	B7_NO	Yes	Column I/O	10	B7_NO	Yes	Column I/O	10	B7_NO	Yes
115	PIN_AD13	Column I/O	10	B7_NO	Yes	Column I/O	10	B7_NO	Yes	Column I/O	10	B7_NO	Yes	Column I/O	10	B7_NO	Yes
116	PIN_AD14	Column I/O	7	B7_NO	Yes	Column I/O	7	B7_NO	Yes	Column I/O	7	B7_NO	Yes	Column I/O	7	B7_NO	Yes

Device... Pin Finder...  Show only highlighted pins  Show migration differences Export...

The migration result for the pin function of highlighted PIN\_AC23 is not an NC but a voltage reference VREFB1N2 even though the pin is an NC in the migration device. VREF standards have a higher priority than an NC, thus the migration result display the voltage reference. Even if you do not use that pin for a port connection in your design, you must use the VREF standard for I/O standards that require it on the actual board for the migration device.

If one of the migration devices has pins intended for connection to  $V_{CC}$  or GND and these same pins are I/O pins on a different device in the migration path, the Quartus II software ensures these pins are not used for I/O. Ensure that these pins are connected to the correct PCB plane.

When migrating between two devices in the same package, pins that are not connected to the smaller die may be intended to connect to  $V_{CC}$  or GND on the larger die. To facilitate migration, you can connect these pins to  $V_{CC}$  or GND in your original design because the pins are not physically connected to the smaller die.

#### Related Information

[AN90: SameFrame PinOut Design for FineLine BGA Packages](#)



## Validating Pin Assignments

The Quartus II software validates I/O pin assignments against predefined I/O rules for your target device. You can use the following tools to validate your I/O pin assignments throughout the pin planning process:

**Table 4-5: I/O Validation Tools**

I/O Validation Tool	Description	Click to Run
Live I/O Check	Verifies preliminary, basic I/O legality as you enter assignments	<b>Processing &gt; Enable Live I/O Check</b>
I/O Assignment Analysis	Verifies I/O assignment legality of synthesized design against full set of I/O rules for the target device	<b>Processing &gt; Start I/O Assignment Analysis</b>
Advanced I/O Timing	Fully validates I/O assignments against all I/O and timing checks during compilation	<b>Processing &gt; Start Compilation</b>

## I/O Assignment Validation Rules

I/O Assignment Analysis validates your assignments against the following rules:

**Table 4-6: Examples of I/O Rule Checks**

Rule	Description	HDL Required?
I/O bank capacity	Checks the number of pins assigned to an I/O bank against the number of pins allowed in the I/O bank.	No
I/O bank VCCIO voltage compatibility	Checks that no more than one VCCIO is required for the pins assigned to the I/O bank.	No
I/O bank VREF voltage compatibility	Checks that no more than one VREF is required for the pins assigned to the I/O bank.	No
I/O standard and location conflicts	Checks whether the pin location supports the assigned I/O standard.	No
I/O standard and signal direction conflicts	Checks whether the pin location supports the assigned I/O standard and direction. For example, certain I/O standards on a particular pin location can only support output pins.	No
Differential I/O standards cannot have open drain turned on	Checks that open drain is turned off for all pins with a differential I/O standard.	No

Rule	Description	HDL Required?
I/O standard and drive strength conflicts	Checks whether the drive strength assignments are within the specifications of the I/O standard.	No
Drive strength and location conflicts	Checks whether the pin location supports the assigned drive strength.	No
BUSHOLD and location conflicts	Checks whether the pin location supports BUSHOLD. For example, dedicated clock pins do not support BUSHOLD.	No
WEAK_PULLUP and location conflicts	Checks whether the pin location supports WEAK_PULLUP (for example, dedicated clock pins do not support WEAK_PULLUP).	No
Electromigration check	Checks whether combined drive strength of consecutive pads exceeds a certain limit. For example, the total current drive for 10 consecutive pads on a Stratix II device cannot exceed 200 mA.	No
PCI_IO clamp diode, location, and I/O standard conflicts	Checks whether the pin location along with the I/O standard assigned supports PCI_IO clamp diode.	No
SERDES and I/O pin location compatibility check	Checks that all pins connected to a SERDES in your design are assigned to dedicated SERDES pin locations.	Yes
PLL and I/O pin location compatibility check	Checks whether pins connected to a PLL are assigned to the dedicated PLL pin locations.	Yes

Table 4-7: Signal Switching Noise Rules

Rule	Description	HDL Required?
I/O bank can not have single-ended I/O when DPA exists	Checks that no single-ended I/O pin exists in the same I/O bank as a DPA.	No
A PLL I/O bank does not support both a single-ended I/O and a differential signal simultaneously	Checks that there are no single-ended I/O pins present in the PLL I/O Bank when a differential signal exists.	No
Single-ended output is required to be a certain distance away from a differential I/O pin	Checks whether single-ended output pins are a certain distance away from a differential I/O pin.	No

Rule	Description	HDL Required?
Single-ended output has to be a certain distance away from a VREF pad	Checks whether single-ended output pins are a certain distance away from a VREF pad.	No
Single-ended input is required to be a certain distance away from a differential I/O pin	Checks whether single-ended input pins are a certain distance away from a differential I/O pin.	No
Too many outputs or bidirectional pins in a VREFGROUP when a VREF is used	Checks that there are no more than a certain number of outputs or bidirectional pins in a VREFGROUP when a VREF is used.	No
Too many outputs in a VREFGROUP	Checks whether too many outputs are in a VREFGROUP.	No

## Checking I/O Pin Assignments In Real-Time

Live I/O check validates I/O assignments against basic I/O buffer rules in real time. The Pin Planner immediately reports warnings or errors about assignments as you enter them. The Live I/O Check Status window displays the total number of errors and warnings. Use this analysis to quickly correct basic errors before proceeding. Run full I/O assignment analysis when you are ready to validate pin assignments against the complete set of I/O system rules.

**Note:** Live I/O check is supported only for Arria II, Cyclone IV, MAX II, and Stratix IV device families.

Live I/O check validates against the following basic I/O buffer rules:

- $V_{CCIO}$  and  $V_{REF}$  voltage compatibility rules
- Electromigration (current density) rules
- Simultaneous Switching Output (SSO) rules
- I/O property compatibility rules, such as drive strength compatibility, I/O standard compatibility,  $PCI\_IO$  clamp diode compatibility, and I/O direction compatibility
- Illegal location assignments:
  - An I/O bank or VREF group with no available pins
  - The negative pin of a differential pair if the positive pin of the differential pair is assigned with a node name with a differential I/O standard
  - Pin locations that do not support the I/O standard assigned to the selected node name
  - For HSTL- and SSTL-type I/O standards, VREF groups of a different  $V_{REF}$  voltage than the selected node name.

### Related Information

[Assigning Device I/O Pins in Pin Planner](#)

## Running I/O Assignment Analysis

I/O assignment analysis validates I/O assignments against the complete set of I/O system and board layout rules. Full I/O assignment analysis validates blocks that directly feed or are fed by resources such as

a PLL, LVDS, or gigabit transceiver blocks. In addition, the checker validates the legality of proper  $V_{REF}$  pin use, pin locations, and acceptable mixed I/O standards

Run I/O assignment analysis during early pin planning to validate initial reserved pin assignments before compilation. Once you define design files, run I/O assignment analysis to perform more thorough legality checks with respect to the synthesized netlist. Run I/O assignment analysis whenever you modify I/O assignments.

The Fitter assigns pins to accommodate your constraints. For example, if you assign an edge location to a group of LVDS pins, the Fitter assigns pin locations for each LVDS pin in the specified edge location and then performs legality checks. To display the Fitter-placed pins, click **Show Fitter Placements** in the Pin Planner. To accept these suggested pin locations, you must back-annotate your pin assignments.

View the I/O Assignment Warnings report to view and resolve all assignment warnings. For example, a warning that some design pins have undefined drive strength or slew rate. The Fitter recognizes undefined, single-ended output and bidirectional pins as non-calibrated OCT. To resolve the warning, assign the **Current Strength**, **Slew Rate** or **Slow Slew Rate** for the reported pin. Alternatively, you could assign the **Termination** to the pin. You cannot assign drive strength or slew rate settings when a pin has an OCT assignment.

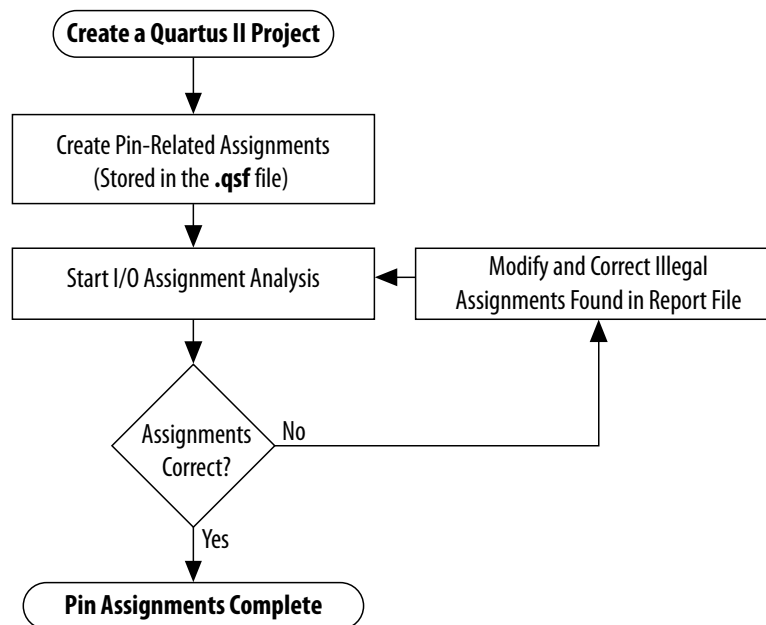
#### Related Information

#### [Back-Annotating Assignments for A Project](#)

### Running Early I/O Assignment Analysis (without Design Files)

You can perform basic I/O legality checks before defining HDL design files. This technique produces a preliminary board layout. For example, you can specify a target device and enter pin assignments that correspond to PCB characteristics. You can reserve and assign an I/O standards to each pin, and then run I/O assignment analysis to ensure that there are no I/O standard conflicts in each I/O bank.

**Figure 4-6: Assigning and Analyzing Pin-Outs without Design Files**

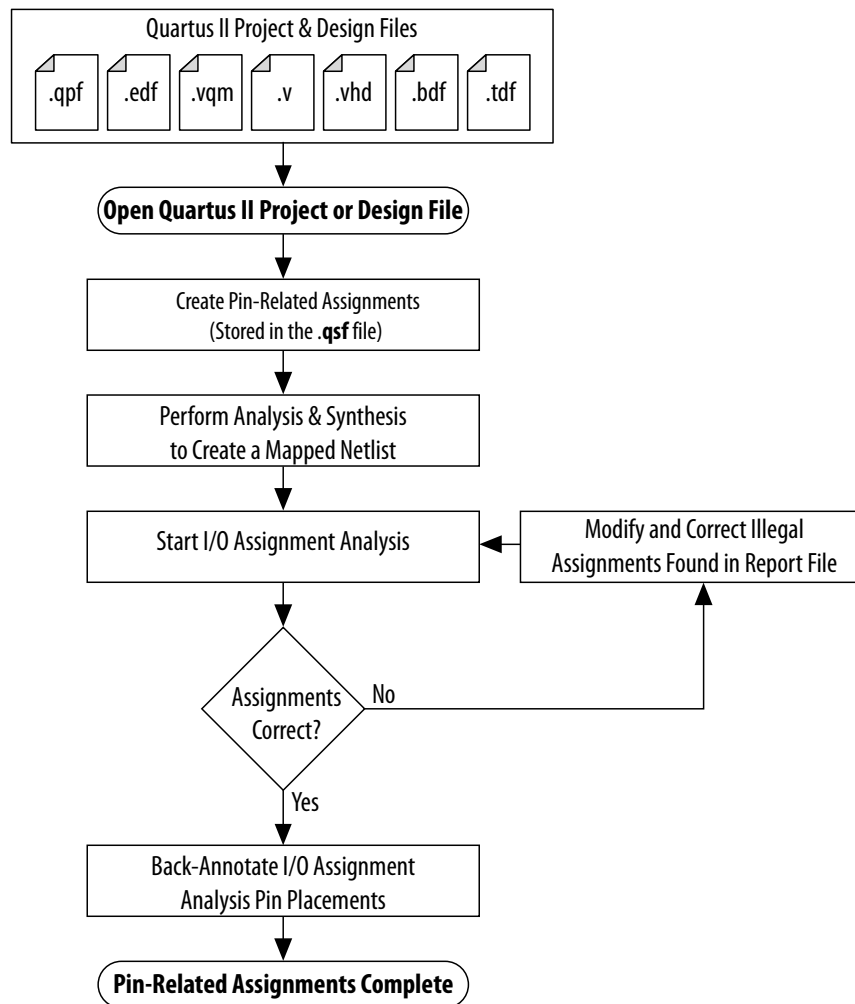


You must reserve all pins you intend to use as I/O pins, so that the Fitter can determine each pin type. After performing I/O assignment analysis, correct any errors reported by the Fitter and rerun I/O assignment analysis until all errors are corrected. A complete I/O assignment analysis requires all design files.

### Running I/O Assignment Analysis (with Design Files)

Use I/O assignment analysis to perform full I/O legality checks after fully defining HDL design files. When you run I/O assignment analysis on a complete design, the tool verifies all I/O pin assignments against all I/O rules. When you run I/O assignment analysis on a partial designs, the tool checks legality only for defined portions of the design. The following figure shows the work flow for analyzing pin-outs with design files.

Figure 4-7: I/O Assignment Analysis Flow



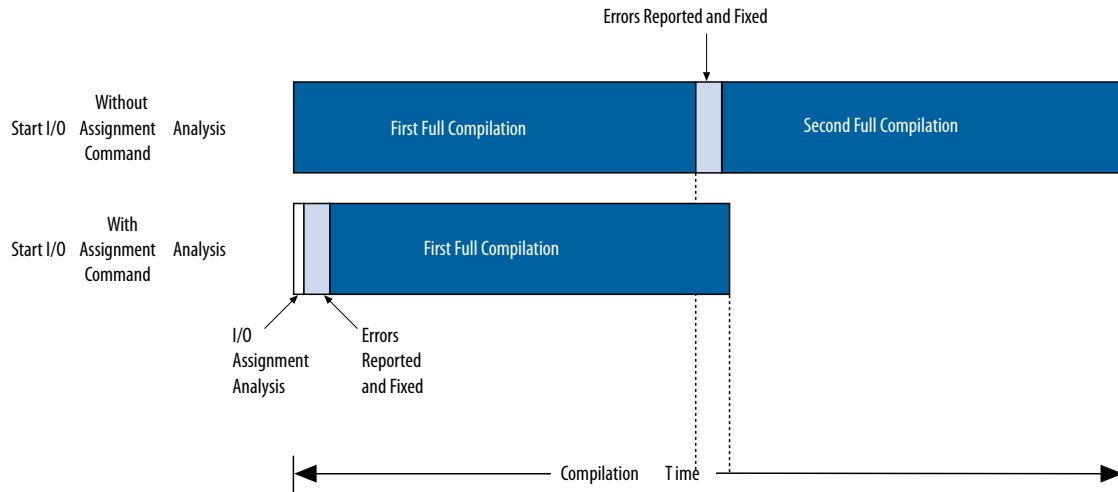
Even if I/O assignment analysis passes on incomplete design files, you may still encounter errors during full compilation. For example, you can assign a clock to a user I/O pin instead of assigning it to a dedicated clock pin, or design the clock to drive a PLL that you have not yet instantiated in the design.

This occurs because I/O assignment analysis does not account for the logic that the pin drives, and does not verify that only dedicated clock inputs can drive the a PLL clock port.

To obtain better coverage, analyze as much of the design as possible over time, especially logic that connects to pins. For example, if your design includes PLLs or LVDS blocks, define these files prior to full analysis. After performing I/O assignment analysis, correct any errors reported by the Fitter and rerun I/O assignment analysis until all errors are corrected.

The following figure shows the compilation time benefit of performing I/O assignment analysis before running a full compilation.

**Figure 4-8: I/O Assignment Analysis Reduces Compilation Time**



## Overriding Default I/O Pin Analysis

You can override the default I/O analysis of various pins to accommodate I/O rule exceptions, such as for analyzing VREF or inactive pins.

Each device contains a number of VREF pins, each supporting a number of I/O pins. A VREF pin and its I/O pins comprise a VREF bank. The VREF pins are typically assigned inputs with VREF I/O standards, such as HSTL- and SSTL-type I/O standards. Conversely, VREF outputs do not require the VREF pin. When a voltage-referenced input is present in a VREF bank, only a certain number of outputs can be present in that VREF bank. I/O assignment analysis treats bidirectional signals controlled by different output enables as independent output enables.

To assign the **Output Enable Group** option to bidirectional signals to analyze the signals as a single output enable group, follow these steps:

1. To access this assignment in the Pin Planner, right-click the **All pins** list and click **Customize Columns**.
2. Under **Available columns**, add **Output Enable Group** to **Show these columns in this order**. The column appears in the **All Pins** list.
3. Enter the same integer value for the **Output Enable Group** assignment for all sets of signals that are driving in the same direction.

This assignment is especially important for external memory interfaces. For example, consider a DDR2 interface in a Stratix II device. The device allows 30 pins in a VREF group. Each byte lane for a  $\times 8$  DDR2 interface includes one DQS pin and eight DQ pins, for a total of nine pins per byte lane. The DDR2 interface uses the **SSTL 18 Class I VREF I/O** standard. In typical interfaces, each byte lane has its own output enable. In this example, the DDR2 interface has four byte lanes. Using 30 I/O pins in a VREF group, there are three byte lanes and an extra byte lane that supports the three remaining pins. Without the **Output Enable Group** assignment, the Fitter analyzes each byte lane as an independent group driven by a unique output enable. In this worst-case scenario the three pins are inputs, and the other 27 pins are outputs violating the 20 output pin limit.

Because DDR2 DQS and DQ pins are always driven in the same direction, the analysis reports an error that is not applicable to your design. The **Output Enable Group** assignment designates the DQS and DQ pins as a single group driven by a common output enable for I/O assignment analysis. When you use the **Output Enable Group** logic option assignment, the DQS and DQ pins are checked as all input pins or all output pins and are not in violation of the I/O rules.

You can also use the **Output Enable Group** assignment to designate pins that are driven only at certain times. For example, the data mask signal in DDR2 interfaces is an output signal, but it is driven only when the DDR2 is writing (bidirectional signals are outputs). To avoid I/O assignment analysis errors, use the **Output Enable Group** logic option assignment to assign the data mask to the same value as the DQ and DQS signals.

You can also use the **Output Enable Group** to designate VREF input pins that are inactive during the time the outputs are driving. This assignment removes the VREF input pins from the VREF analysis. For example, the QVLD signal for an RLDRAM II interface is active only during a read. During a write, the QVLD pin is not active and does not count as an active VREF input pin in the VREF group. Place the QVLD pins in the same output enable group as the RLDRAM II data pins.

#### Related Information

[The TimeQuest Timing Analyzer](#)

## Understanding I/O Analysis Reports

The detailed I/O assignment analysis reports include the affected pin name and a problem description. The Fitter section of the Compilation report contains information generated during I/O assignment analysis, including the following reports:

- I/O Assignment Warnings—lists warnings generated for each pin
- Resource Section—quantifies use of various pin types and I/O banks
- I/O Rules Section—lists summary, details, and matrix information about the I/O rules tested

The **Status** column indicates whether rules passed, failed, or could not be checked. A severity rating indicates the rule's importance for effective analysis. "Inapplicable" rules do not apply to the target device family.

Figure 4-9: I/O Rules Matrix

Pin/Rules	ID_000001	ID_000002	ID_000003	ID_000004	ID_000005	ID_000006	ID_000007	ID_000008	ID_000009	ID_000010	ID_000011
1 Total Pass	21	0	21	0	0	21	21	0	21	21	20
2 Total Unchecked	1	0	1	0	0	1	1	0	1	1	1
3 Total Inapplicable	0	22	0	22	22	0	0	22	0	0	0
4 Total Fail	0	0	0	0	0	0	0	0	0	0	1
5 yvalid	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
6 iclrow	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Fail
7 yrn_out[7]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
8 yrn_out[6]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
9 yrn_out[5]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
10 yrn_out[4]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
11 yrn_out[3]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
12 yrn_out[2]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
13 yrn_out[1]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
14 yrn_out[0]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
15 clk	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
16 reset	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
17 clkx2	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
18 newt	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
19 d[7]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
20 d[6]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
21 d[5]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
22 d[4]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
23 d[3]	Unchecked	Inapplicable	Unchecked	Inapplicable	Inapplicable	Unchecked	Unchecked	Inapplicable	Unchecked	Unchecked	Unchecked
24 d[2]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
25 d[1]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
26 d[0]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass

## Verifying I/O Timing

You must verify board-level signal integrity and I/O timing when assigning I/O pins. High-speed interface operation requires a quality signal and low propagation delay at the far end of the board route. Click **Tools > TimeQuest Timing Analyzer** to confirm timing after making I/O pin assignments. For example, if you change the slew rates or drive strengths of some I/O pins with ECOs, you can verify timing without recompiling the design. You must understand I/O timing and what factors affect I/O timing paths in your design. The accuracy of the output load specification of the output and bidirectional pins affects the I/O timing results.

The Quartus II software supports three different methods of I/O timing analysis:

Table 4-8: I/O Timing Analysis Methods

I/O Timing Analysis	Description
Advanced I/O timing analysis	Analyze I/O timing with your board trace model to report accurate, “board-aware” simulation models. Configures a complete board trace model for each I/O standard or pin. TimeQuest applies simulation results of the I/O buffer, package, and board trace model to generate accurate I/O delays and system level signal information. Use this information to improve timing and signal integrity.
I/O timing analysis	Analyze I/O timing with default or specified capacitive load without signal integrity analysis. TimeQuest reports tCO to an I/O pin using a default or user-specified value for a capacitive load.



I/O Timing Analysis	Description
Full board routing simulation	Use Altera-provided or Quartus II software-generated IBIS or HSPICE I/O models for simulation in Mentor Graphics HyperLynx and Synopsys HSPICE.

**Note:** Advanced I/O timing analysis is supported only for .28nm and larger device families. For devices that support advanced I/O timing, it is the default method of I/O timing analysis. For all other devices, you must use a default or user-specified capacitive load assignment to determine  $t_{CO}$  and power measurements.

For more information about advanced I/O timing support, refer to the appropriate device handbook for your target device. For more information about board-level signal integrity and tips on how to improve signal integrity in your high-speed designs, refer to the Altera Signal Integrity Center page of the Altera website.

For information about creating IBIS and HSPICE models with the Quartus II software and integrating those models into HyperLynx and HSPICE simulations, refer to the *Signal Integrity Analysis with Third Party Tools* chapter in volume 2 of the *Quartus II Handbook*.

#### Related Information

- [Literature and Technical Documentation](#)
- [Altera Signal Integrity Center](#)
- [Signal Integrity Analysis with Third-Party Tools](#) on page 6-1

## Running Advanced I/O Timing

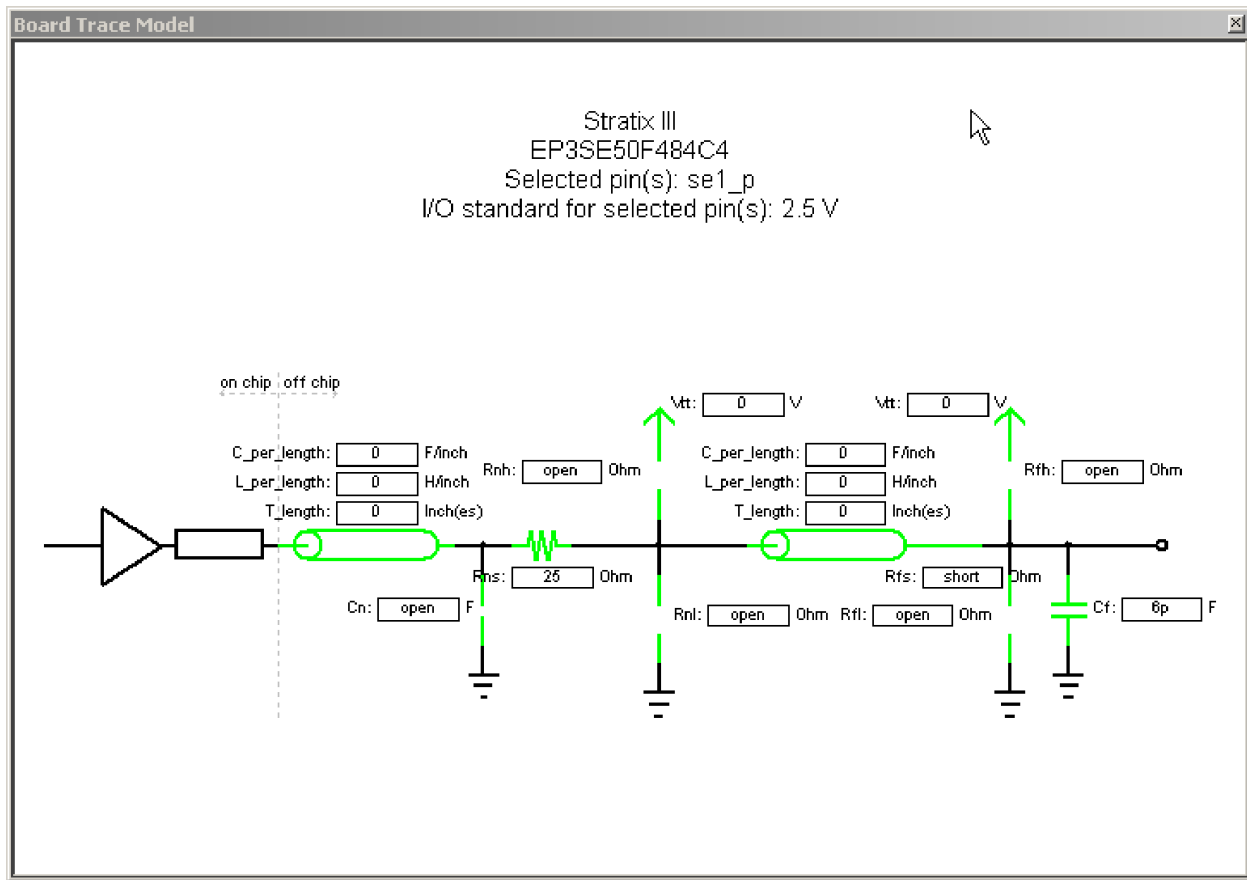
Advanced I/O timing analysis uses your board trace model and termination network specification to report accurate output buffer-to-pin timing estimates, FPGA pin and board trace signal integrity and delay values. Advanced I/O timing runs automatically for supported devices during compilation.

### Understanding the Board Trace Models

The Quartus II software provides board trace model templates for various I/O standards. The following figure shows the template for a 2.5 V I/O standard. This model consists of near-end and far-end board component parameters.

Near-end board trace modeling includes the elements which are close to the device. Far-end modeling includes the elements which are at the receiver end of the link, closer to the receiving device. Board trace model topology is conceptual and does not necessarily match the actual board trace for every component. For example, near-end model parameters can represent device-end discrete termination and breakout traces. Far-end modeling can represent the bulk of the board trace to discrete external memory components, and the far end termination network. You can analyze the same circuit with near-end modeling of the entire board, including memory component termination, and far-end modeling of the actual memory component.

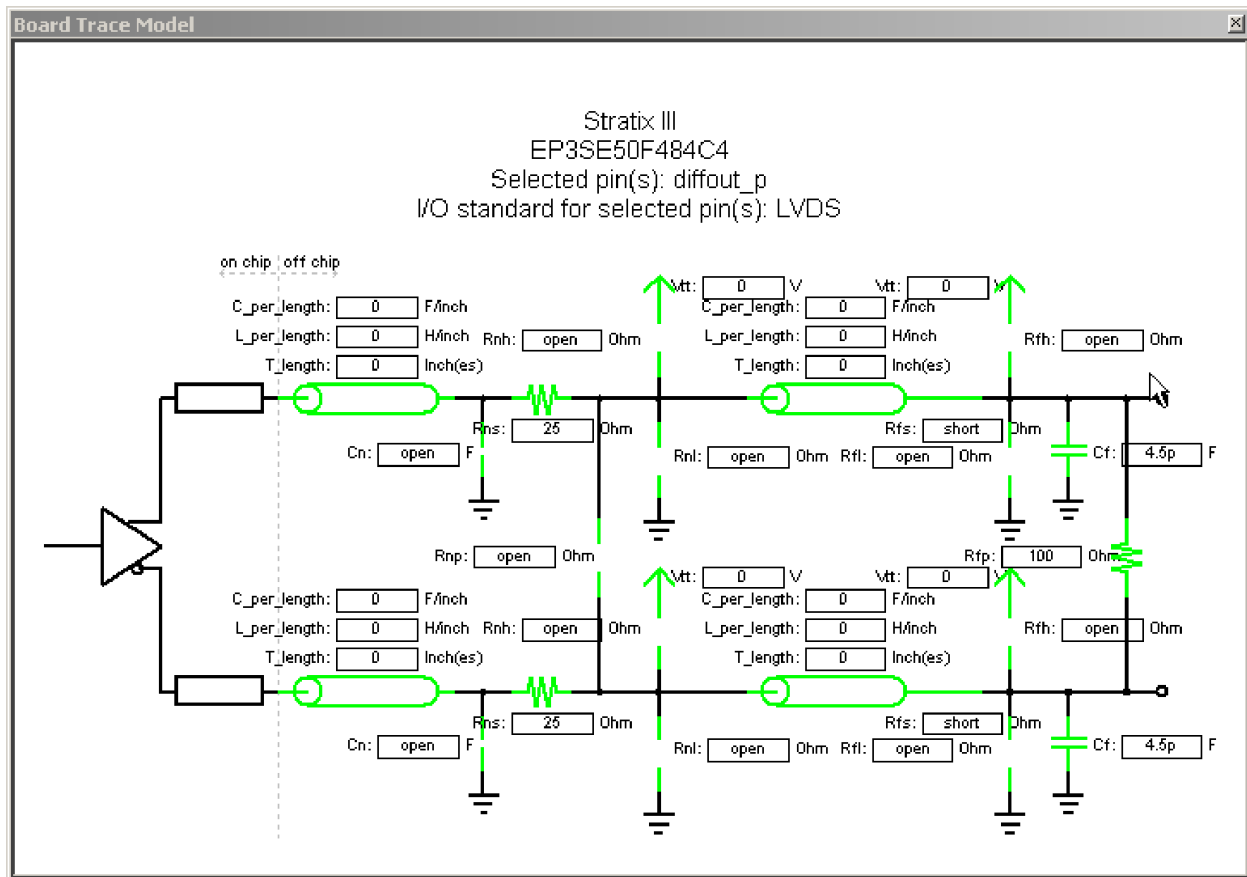
Figure 4-10: 2.5-V I/O Standard Board Trace Model



The following figure shows the template for the LVDS I/O standard. The far-end capacitance ( $C_f$ ) represents the external-device or multiple-device capacitive load. If you have multiple devices on the far-end, you must find the equivalent capacitance at the far-end, taking into account all receiver capacitances. The far-end capacitance can be the sum of all the receiver capacitances.

The Quartus II software models lossless transmission lines, and does not require a transmission-line resistance value. Only distributed inductance ( $L$ ) and capacitance ( $C$ ) values are needed. The distributed  $L$  and  $C$  values of transmission lines must be entered on a per-inch basis, and can be obtained from the PCB vendor or manufacturer, the CAD Design tool, or a signal integrity tool, such as the Mentor Graphics Hyperlynx software.

Figure 4-11: LVDS Differential Board Trace Model



## Defining the Board Trace Model

The board trace model describes a board trace and termination network as a set of capacitive, resistive, and inductive parameters. Advanced I/O Timing and the SSN Analyzer use the model to simulate the output signal from the output buffer to the far end of the board trace. You can define the capacitive load, any termination components, and trace impedances in the board routing for any output pin or bidirectional pin in output mode. You can configure an overall board trace model for each I/O standard or for specific pins. Define an overall board trace model for each I/O standard in your design. Use that model for all pins that use the I/O standard. You can customize the model for specific pins using the Board Trace Model window in the Pin Planner.

1. Click **Assignments > Device** and then click **Device and Pin Options**.
2. Click **Board Trace Model** and define board trace model values for each I/O standard.
3. Click **I/O Timing** and define default I/O timing options at board trace near and far ends.
4. Click **Assignments > Pin Planner** and assign board trace model values to individual pins.

## Example 4-7: Specifying Board Trace Model

```
## setting the near end series resistance model of sel_p output pin to 25 ohms
```

```
set_instance_assignment -name BOARD_MODEL_NEAR_SERIES_R 25 -to sel_p
## Setting the far end capacitance model for sel_p output signal to 6
picofarads
set_instance_assignment -name BOARD_MODEL_FAR_C 6P -to sel_p
```

#### Related Information

- [Using Advanced I/O Timing](#)
- [Board Trace Models](#)

## Modifying the Board Trace Model

To modify the board trace model, click **View > Board Trace Model** in the Pin Planner. You can modify any of the board trace model parameters within a graphical representation of the board trace model.

The Board Trace Model window displays the routing and components for positive and negative signals in a differential signal pair. Only modify the positive signal of the pair, as the setting automatically applies to the negative signal. Use standard unit prefixes such as *p*, *n*, and *k* to represent pico, nano, and kilo, respectively. Use the **short** or **open** value to designate a short or open circuit for a parallel component.

## Specifying Near End vs Far End I/O Timing Analysis

You can select a nearend or far end point for I/O timing analysis. Near end timing analysis extends to the device pin. You can apply the `set_output_delay` constraint during near end analysis to account for the delay across the board.

Far end I/O timing analysis, then advanced I/O timing analysis extends to the external device input, at the far end of the board trace. Whether you choose a near end or far end timing endpoint, the board trace models are taken into account during timing analysis.

## Understanding Advanced I/O Timing Analysis Reports

View I/O timing analysis information in the following reports:

**Table 4-9: Advanced I/O Timing Reports**

I/O Timing Report	Description
TimeQuest Report	Reports signal integrity and board delay data.
Board Trace Model Assignments report	Summarizes the board trace model component settings for each output and bidirectional signal.
Signal Integrity Metrics report	Contains all the signal integrity metrics calculated during advanced I/O timing analysis based on the board trace model settings for each output or bidirectional pin. Includes measurements at both the FPGA pin and at the far-end load of board delay, steady state voltages, and rise and fall times.

**Note:** By default, the TimeQuest analyzer generates the Slow-Corner Signal Integrity Metrics report. To generate a Fast-Corner Signal Integrity Metrics report you must change the delay model by clicking **Tools > TimeQuest Timing Analyzer**.

**Related Information**[The TimeQuest Timing Analyzer](#)

## Adjusting I/O Timing and Power with Capacitive Loading

When calculating  $t_{CO}$  and power for output and bidirectional pins, the TimeQuest analyzer and the PowerPlay Power Analyzer use a bulk capacitive load. You can adjust the value of the capacitive load per I/O standard to obtain more precise  $t_{CO}$  and power measurements, reflecting the behavior of the output or bidirectional net on your PCB. The Quartus II software ignores capacitive load settings on input pins. You can adjust the capacitive load settings per I/O standard, in picofarads (pF), for your entire design. During compilation, the Compiler measures power and  $t_{CO}$  measurements based on your settings. You can also adjust the capacitive load on an individual pin with the **Output Pin Load** logic option.

## Viewing Routing and Timing Delays

Right-click any node and click **Locate > Locate in Chip Planner** to visualize and adjust I/O timing delays and routing between user I/O pads and  $V_{CC}$ , GND, and  $V_{REF}$  pads. The Chip Planner graphically displays logic placement, LogicLock regions, relative resource usage, detailed routing information, fan-in and fan-out, register paths, and high-speed transceiver channels. You can view physical timing estimates, routing congestion, and clock regions. Use the Chip Planner to change connections between resources and make post-compilation changes to logic cell and I/O atom placement. When you select items in the Pin Planner, the corresponding item is highlighted in Chip Planner.

## Analyzing Simultaneous Switching Noise

Click **Processing > Start > Start SSN Analyzer** to estimate the voltage noise for each pin in the design. The simultaneous switching noise (SSN) analysis accounts for the pin placement, I/O standard, board trace, output enable group, timing constraint, and PCB characteristics that you specify. The analysis produces a voltage noise estimate for each pin in the design. View the SSN results in the Pin Planner and adjust your I/O assignments to optimize signal integrity.

**Related Information**[Simultaneous Switching Noise \(SSN\) Analysis and Optimization](#)

## Scripting API

You can alternatively use Tcl commands to access I/O management functions, rather than using the GUI. For detailed information about specific scripting command options and Tcl API packages, type the following command at a system command prompt to view the Tcl API Help browser:

```
quartus_sh --qhelp
```

**Related Information**

- [Tcl Scripting](#) on page 3-1
- [Command Line Scripting](#) on page 2-1

## Run I/O Assignment Analysis

Enter the following in the Tcl console or a Tcl script:

```
execute_flow -check_ios
```

Type the following at a system command prompt:

```
quartus_fit <project name> --check_ios
```

## Generate Mapped Netlist

Enter the following in the Tcl console or in a Tcl script:

```
execute_module -tool map
```

The `execute_module` command is in the flow package.

Type the following at a system command prompt:

```
quartus_map <project name>
```

## Reserve Pins

Use the following Tcl command to reserve a pin:

```
set_instance_assignment -name RESERVE_PIN <value> -to <signal name>
```

Use one of the following valid reserved pin values:

- "AS BIDIRECTIONAL"
- "AS INPUT TRI STATED"
- "AS OUTPUT DRIVING AN UNSPECIFIED SIGNAL"
- "AS OUTPUT DRIVING GROUND"
- "AS SIGNALPROBE OUTPUT"

**Note:** You must include the quotation marks when specifying the reserved pin value.

## Set Location

Use the following Tcl command to assign a signal to a pin or device location:

```
set_location_assignment <location> -to <signal name>
```

Valid locations are pin locations, I/O bank locations, or edge locations. Pin locations include pin names, such as `PIN_A3`. I/O bank locations include `IOBANK_1` up to `IOBANK_n`, where *n* is the number of I/O banks in the device.

Use one of the following valid edge location values:

- `EDGE_BOTTOM`
- `EDGE_LEFT`
- `EDGE_TOP`
- `EDGE_RIGHT`

## Exclusive I/O Group

Use the following Tcl command to create an exclusive I/O group assignments:

```
set_instance_assignment -name "EXCLUSIVE_IO_GROUP" -to pin
```

## Slew Rate and Current Strength

Use the following Tcl commands to create an slew rate and drive strength assignments:

```
set_instance_assignment -name CURRENT_STRENGTH_NEW 8MA -to e[0]
set_instance_assignment -name SLEW_RATE 2 -to e[0]
```

### Related Information

[Altera Device Package Information Data Sheet](#)

## Document Revision History

The following table shows the revision history for this chapter.

**Table 4-10: Document Revision History**

Date	Version	Changes
2014.12.15	14.1.0	<ul style="list-style-type: none"> <li>Updated Live I/O check device support to include only limited device families.</li> </ul>
2014.08.30	14.0a10.0	<ul style="list-style-type: none"> <li>Added link to information about special pin assignment features for Arria 10 SoC devices.</li> </ul>
2014.06.30	14.0.0	<ul style="list-style-type: none"> <li>Replaced MegaWizard Plug-In Manager information with IP Catalog.</li> </ul>
November 2013	13.1.0	<ul style="list-style-type: none"> <li>Reorganization and conversion to DITA.</li> </ul>
May 2013	13.0.0	<ul style="list-style-type: none"> <li>Added information about overriding I/O placement rules.</li> </ul>
November 2012	12.1.0	<ul style="list-style-type: none"> <li>Updated Pin Planner description for new task and report windows.</li> </ul>
June 2012	12.0.0	<ul style="list-style-type: none"> <li>Removed survey link.</li> </ul>
November 2011	11.1.0	<ul style="list-style-type: none"> <li>Minor updates and corrections.</li> <li>Updated the document template.</li> </ul>
December 2010	10.0.1	Template update

Date	Version	Changes
July 2010	10.0.0	<ul style="list-style-type: none"> <li>• Reorganized and edited the chapter</li> <li>• Added links to Help for procedural information previously included in the chapter</li> <li>• Added information on rules marked Inapplicable in the I/O Rules Matrix Report</li> <li>• Added information on assigning slew rate and drive strength settings to pins to fix I/O assignment warnings</li> </ul>
November 2009	9.1.0	<ul style="list-style-type: none"> <li>• Reorganized entire chapter to include links to Help for procedural information previously included in the chapter</li> <li>• Added documentation on near-end and far-end advanced I/O timing</li> </ul>
March 2009	9.0.0	<ul style="list-style-type: none"> <li>• Updated “Pad View Window” on page 5–20</li> <li>• Added new figures: <ul style="list-style-type: none"> <li>• Figure 5–15</li> <li>• Figure 5–16</li> </ul> </li> <li>• Added new section “Viewing Simultaneous Switching Noise (SSN) Results” on page 5–17</li> <li>• Added new section “Creating Exclusive I/O Group Assignments” on page 5–18</li> </ul>

**Related Information**

[Quartus II Handbook Archive](#)



2014.12.15

QII5V2



Subscribe



Send Feedback

## Simultaneous Switching Noise (SSN) Analysis and Optimizations

FPGA design has evolved from small programmable circuits to designs that compete with multimillion-gate ASICs. At the same time, the I/O counts on FPGAs and logic density requirements of designs have increased exponentially.

The higher-speed interfaces in FPGAs, including high-speed serial interfaces and memory interfaces, require careful interface design on the PCB. Designers must address the timing and signal integrity requirements of these interfaces early in the design cycle. Simultaneous switching noise (SSN) often leads to the degradation of signal integrity by causing signal distortion, thereby reducing the noise margin of a system.

Today's complex FPGA system design is incomplete without addressing the integrity of signals coming in to and out of the FPGA. Altera recommends that you perform SSN analysis early in your FPGA design and prior to the layout of your PCB with complete SSN analysis of your FPGA in the Quartus® II software. This chapter describes the Quartus II SSN Analyzer tool and covers the following topics:

### Definitions

The terminology used in this chapter includes the following terms:

- **Aggressor:** An output or bidirectional signal that contributes to the noise for a victim I/O pin
- **PDN:** Power distribution network
- **QH:** Quiet high signal level on a pin
- **QHN:** Quiet high noise on a pin, measured in volts
- **QL:** Quiet low signal level on a pin
- **QLN:** Quiet low noise on a pin, measured in volts
- **SI:** Signal integrity (a superset of SSN, covering all noise sources)
- **SSN:** Simultaneous switching noise
- **SSO:** Simultaneous switching output (which are either the output or bidirectional pins)
- **Victim:** An input, output, or bidirectional pin that is analyzed during SSN analysis. During SSN analysis, each pin is analyzed as a victim. If a pin is an output or bidirectional pin, the same pin acts as an aggressor signal for other pins.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered


## Understanding SSN

SSN is defined as a noise voltage induced onto a single victim I/O pin on a device due to the switching behavior of other aggressor I/O pins on the device. SSN can be divided into two types of noise: voltage noise and timing noise.

In a sample system with three pins, two of the pins (A and C) are switching, while one pin (B) is quiet. If the pins are driven in isolation, the voltage waveforms at the output of the buffers appear without noise interference, as shown by the solid curves at the left of the figure. However, when pins A and C are switching simultaneously, the noise generated by the switching is injected onto other pins. This noise manifests itself as a voltage noise on pin B and timing noise on pins A and C.

**Figure 5-1: System with Three Pins**

In this figure, the dotted curves show the voltage noise on pin B and timing noise on pins A and C.



Voltage noise is measured as the change in voltage of a signal due to SSN. When a signal is QH, it is measured as the change in voltage toward 0 V. When a signal is QL, it is measured as the change in voltage toward  $V_{CC}$ .

In the Quartus<sup>®</sup> II software, only voltage noise is analyzed. Voltage noise can be caused by SSOs under two worst-case conditions:

- The victim pin is high and the aggressor pins (SSOs) are switching from low to high
- The victim pin is low and the aggressor pins (SSOs) are switching from high to low

**Figure 5-2: Quiet High Output Noise Estimation on Pin B**

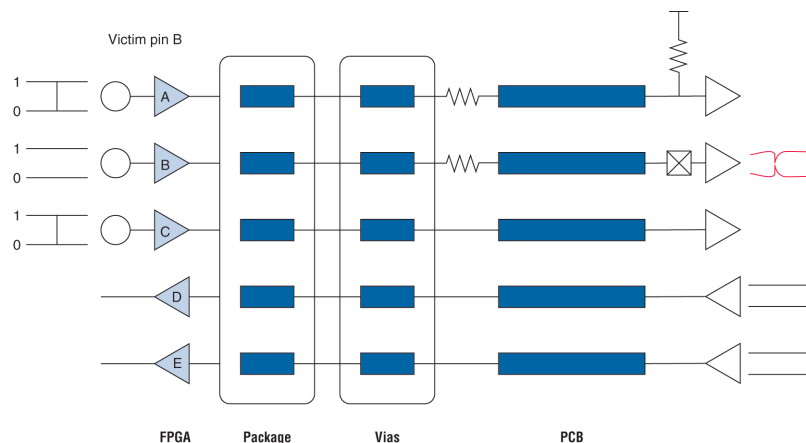
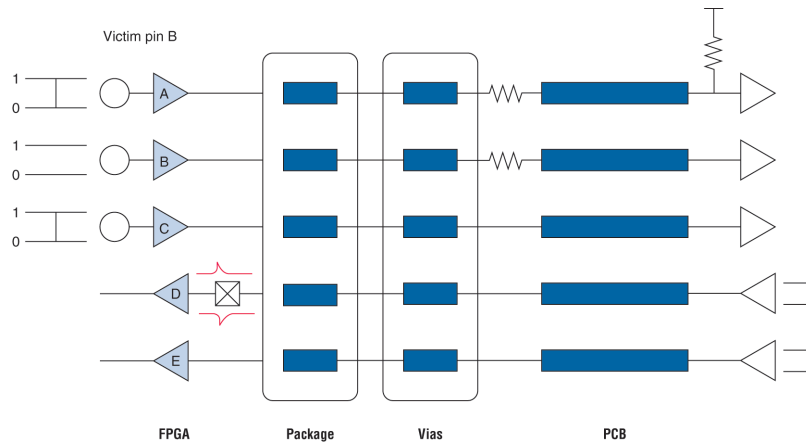



Figure 5-3: Quiet Low Input Noise Estimation for Pin D



SSN can occur in any system, but the induced noise does not always result in failures. Voltage functional errors are caused by SSN on quiet victim pins only when the voltage values on the quiet pins change by a large voltage that the logic listening to that signal reads a change in the logic value. For QH signals, a voltage functional error occurs when noise events cause the voltage to fall below  $V_{IH}$ . Similarly, for QL signals, a voltage functional error occurs when noise events cause the voltage to rise above  $V_{IL}$ . Because  $V_{IH}$  and  $V_{IL}$  of the Altera device are different for different I/O standards, and because signals have different quiet voltage values, the absolute amount of SSN, measured in volts, cannot be used to determine if a voltage failure occurs. Instead, to assess the level of impact by SSN in the SSN analysis, the Quartus II software quantifies the SSN in terms of the percentage of signal margin in Altera devices.


Figure 5-4: Reporting Noise Margins



The figure shows four noise events, two on QH signals and two on QL signals. The two noise events on the right-side of the figure consume 50 percent of the signal margin and do not cause voltage functional errors. However, the two noise events on the left side of the figure consume 100 percent of the signal margin, which can cause a voltage functional error.

Noise caused by aggressor signals are synchronously related to the victim pin outside of the sampling window of a receiver. This noise affects the switching time of a victim pin, but are not considered an input threshold violation failure.

Figure 5-5: Synchronous Voltage Noise with No Functional Error

**Related Information**

[SSN Analysis Overview](#) on page 5-5

## SSN Estimation Tools

Addressing SSN early in your FPGA design and PCB layout can help you avoid costly board respins and lost time, both of which can impact your time-to-market.

Altera provides many tools for SSN analysis and estimation, including the following tools:

- SSN characterization reports
- An early SSN estimation (ESE) tool
- The SSN Analyzer in the Quartus II software

The ESE tool is useful for preliminary SSN analysis of your FPGA design; for more accurate results, however, you must use the SSN Analyzer in the Quartus II software.

**Table 5-1: Comparison of ESE Tool and SSN Analyzer Tool**

ESE Tool	SSN Analyzer
Is not integrated with the Quartus II software.	Integrated with the Quartus II software, allowing you to perform preliminary SSN analysis while making I/O assignment changes in the Quartus II software.
QL and QH levels are computed assuming a worst-case pattern of I/O placements.	QL and QH levels are computed based on the I/O placements in your design.
No support for entering board information.	Supports board trace models and board layer information, resulting in a more accurate SSN analysis.
No graphical representation.	Integrated with the Quartus II Pin Planner, in which an SSN map shows the QL and QH levels on victim pins.
Good for doing an early SSN estimate. Does not require you to use the Quartus II software.	Requires you to create a Quartus II software project and provide the top-level port information.

### Related Information

- [Signal Integrity Center](#)  
For more information on the Altera website about the SSN characterization reports and the ESE tool, including device support information.
- [About the SSN Analyzer](#)  
For more information about the devices for which you can run the SSN Analyzer, refer to Quartus II Help.

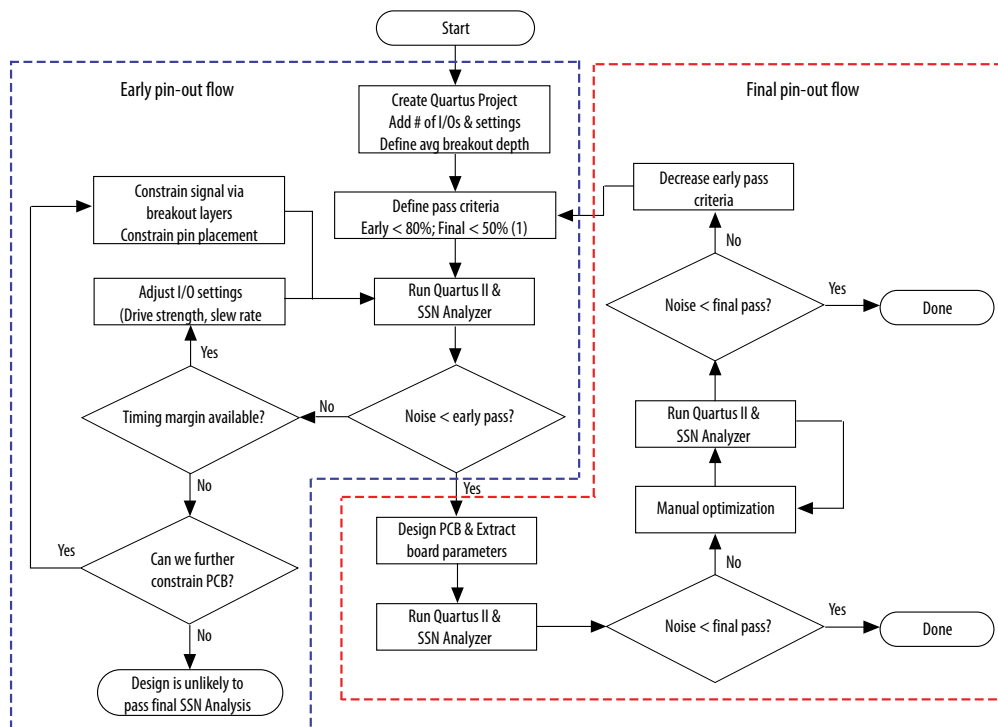
## SSN Analysis Overview

You can run the SSN Analyzer at different stages in your design cycle to obtain SSN results. The accuracy of the results depends on the completeness of your design information.

Altera recommends that you start SSN analysis early in the design cycle to obtain preliminary results and make adjustments to your I/O assignments, and iterate through the design cycle to finally perform a fully constrained SSN analysis with complete information about your board.

The early pin-out flow assumes conservative design rules initially, and then lets you analyze the design and iteratively apply tighter design rules until SSN analysis indicates your design meets SSN constraints. You must define pass criteria for SSN analysis as a percentage of signal margin in both the early pin-out flow and the final pin-out flow. The pass criteria you define is specific to your design requirements. For example, a pass criterion you might define is a condition that verifies you have sufficient SSN margins in your design. You may require that the acceptable voltage noise on a pin must be below 70% of the voltage level for that pin. The pass criteria for the early-pin out flow may be higher than the final pin-out flow criteria, so that you do not spend too much time optimizing the on-FPGA portions of your design when the SSN metrics for the design may improve after the design is fully specified.

Figure 5-6: Early Flow and Final Pin-Out SSN Analysis



Note :

1. Pass criteria determined by customer requirements.

## Performing Early Pin-Out SSN Analysis

In the early stages of your design cycle, before you create pin location for your design, use the early pin-out flow to obtain preliminary SSN analysis results.

In order to obtain useful SSN results, you must define the top-level ports of your design, but your design files do not have to be complete.

### Performing Early Pin-Out SSN Analysis with the ESE Tool

If you know the I/O standards and signaling standards for your design, you can use the ESE tool to perform an initial SSN evaluation.

#### Related Information

##### [Signal Integrity Center](#)

For more information on the Altera website about the ESE tool.

### Performing Early Pin-Out SSN Analysis with the SSN Analyzer

In the early stages of your design cycle, you may not have complete board information, such as board trace parameters, layer information, and the signal breakout layers. If you run the SSN Analyzer without this specific information, it uses default board trace models and board layer information for SSN analysis, and as a result the SSN Analyzer confidence level is low. If the noise amounts are larger than the pass criteria for early pin-out SSN analysis, verify whether the SSN noise violations are true failures or false failures.

For example, sometimes the SSN Analyzer can determine whether pins are switching synchronously and use that information to filter false positives; however, it may not be able to determine all the synchronous groups. You can improve the SSN analysis results by adjusting your I/O assignments and other design settings. After you optimize your design such that it meets the pass criteria for the early pin-out flow, you can then begin to design your PCB.

If you have complete information for the top-level ports of your design, you can use the SSN Analyzer to perform an initial SSN evaluation. Use the following steps to perform early pin-out SSN analysis:

1. Create a project in the Quartus II software.
2. Specify your top-level design information either in schematic form or in HDL code.
3. Perform Analysis and Synthesis.
4. Create I/O assignments, such as I/O standard assignments, for the top-level ports in your design.

**Note:** Do not create pin location assignments. The Fitter automatically creates optimized pin location assignments.

5. If you do not have completed design files and timing constraints, run I/O assignment analysis.

**Note:** During I/O assignment analysis, the Fitter places all the unplaced pins on the device, and checks all the I/O placement rules.

6. Run the SSN Analyzer.

#### Related Information

- [Optimizing Your Design for SSN Analysis](#) on page 5-8
- [Managing Quartus II Projects](#)  
For more information about creating and managing projects.
- [I/O Management](#) on page 4-1  
For more about generating a top-level design file in the Quartus II software and I/O assignment analysis.

## Performing Final Pin-Out SSN Analysis

You perform final pin-out SSN analysis after you place all the pins in your design, or the Fitter places them for you, and you have complete information about the board trace models and PCB layers.

Even if your design achieves sufficient SSN results during early pin-out SSN analysis, you should run SSN analysis with the complete PCB information to ensure that SSN does not cause failures in your final design. You must specify the board parameters in the Quartus II software, including the PCB layer thicknesses, the signal breakout layers, and the board trace models, before you can run SSN analysis on your final assignments.

If the SSN analysis results meet the pass criteria for final pin-out SSN analysis, SSN analysis is complete. If the SSN analysis results do not meet the pass criteria, you must further optimize your design by changing the board and design parameters and then rerun the SSN Analyzer. If the design still does not meet the pass criteria, reduce the pass criteria for early pin-out SSN analysis, and restart the process. By reducing the pass criteria for early pin-out SSN analysis, you place a greater emphasis on reducing SSN through I/O settings and I/O placement. Changing the drive strength and slew rate of output and bidirectional pins, as well as adjusting the placement of different SSOs, can affect SSN results. Adjusting I/O settings and placement allows the design to meet the pass criteria for final pin-out SSN analysis after you specify the actual PCB board parameters.

#### Related Information

- [Optimizing Your Design for SSN Analysis](#) on page 5-8

## Design Factors Affecting SSN Results

There are many factors that affect the SSN levels in your design. The two main factors are the drive strength and slew rate settings of the output and bidirectional pins in your design.

### Related Information

#### [Signal Integrity Center](#)

For more information on the Altera website about the factors that contribute to SSN voltage noise in your FPGA design and managing SSN in your system.

## Optimizing Your Design for SSN Analysis

The SSN Analyzer gives you flexibility to precisely define your system to obtain accurate SSN results.

The SSN Analyzer produces a voltage noise estimate for each input, output, and bidirectional pin in the design. It allows you to estimate the SSN levels, comprised of QLN and QHN levels, for your FPGA pins. Performing SSN analysis helps you optimize your design for SSN during compilation.

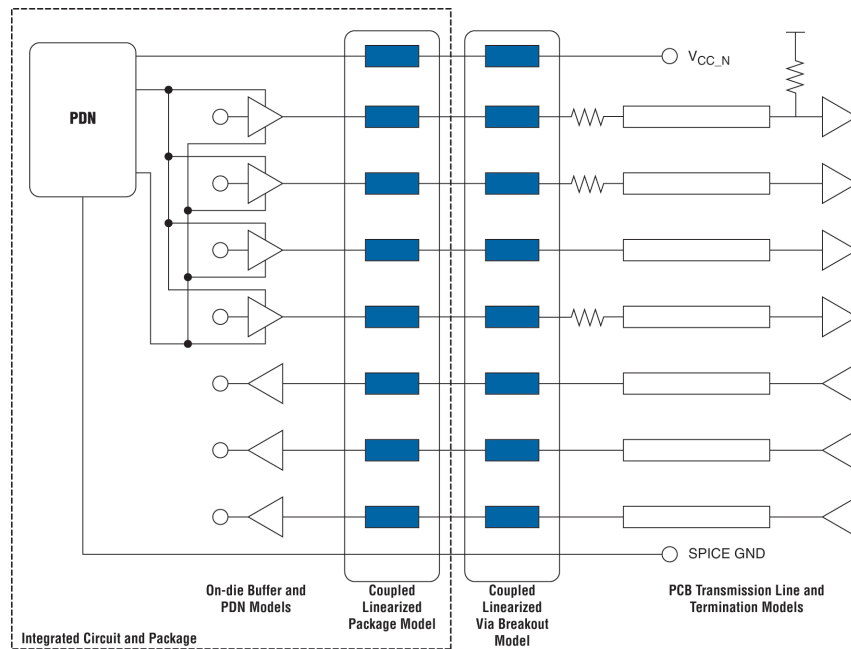
Because the SSN Analyzer is integrated into the Quartus II software, it can automatically set up a system topology that matches your design. The SSN Analyzer accounts for different I/O standards and slew rate settings for each buffer in the design and models different board traces for each signal. Also, it correctly models the state of the unused pins in the design. The SSN Analyzer leverages any custom board trace assignments you set up for use by the advanced I/O timing feature.

The SSN Analyzer also models the package and vias in the design. Models for the different packages that Altera devices support are integrated into the Quartus II software. In the Quartus II software, you can specify different layers on which signals break out, each with its own thickness, and then specify which signal breaks out on which layer.

After constructing the circuit topology, the SSN Analyzer uses a simulation-based methodology to determine the SSN for each victim pin in the design.



Figure 5-7: Circuit Topology for SSN Analysis



## Optimizing Pin Placements for Signal Integrity


You can take advantage of a built-in SSN optimization feature in the Quartus II software with the **SSN Optimization** logic option.

The I/O placements in your design may be affected when you use this option. Setting this option to **Normal compilation** does not affect the  $f_{MAX}$  of your design during compilation, however setting this option to **Extra effort** level may impact your design  $f_{MAX}$ .

**Note:** In order to use the **SSN Optimization** logic option, Altera recommends that you do not create location assignments for your pins; instead, let the Fitter place the pins during compilation so that it places the pins to meet the timing performance of your design. To display the Fitter-placed pins use the Show Fitter Placements feature in the Pin Planner. To accept these suggested pin locations, you must back-annotate your pin assignments.

The image on the left shows the placement of the pins without the **SSN Optimization** logic option, and the image on the right shows the adjustments the Fitter made to pin placements to reduce the amount of SSN in the design when the **SSN Optimization** logic option is turned on.

Figure 5-8: SSN Analysis Results Before and After Using the SSN Optimization Logic Option



**Related Information**

- **Setting Up and Running the Fitter**  
For more information about creating project-wide logic option assignments, refer to Quartus II Help.
- **Show Commands**  
For more information about the Show Fitter Placements feature, refer to Quartus II Help.
- **Back-Annotating Assignments for A Project**  
For more information about back-annotating assignments, refer to Quartus II Help.
- **Area, Timing, and Compilation Time Optimization**  
For more information about design optimization features, refer to the *Quartus II Handbook*.

## Specifying Board Trace Model Settings

The SSN Analyzer uses circuit models to determine voltage noise during SSN analysis. The circuit topology is incomplete without board trace information and PCB layer information.

You must describe the board trace and PCB layer parameters in your design to accurately compute the SSN in your FPGA device. However, if you do not specify some or all of the board trace parameters and PCB layer information, the SSN Analyzer uses default parameters during SSN analysis. When you use the default parameters, the SSN confidence level is low.

The board trace models required for the SSN Analyzer include the board trace termination resistors, pin loads (capacitance), and transmission line parameters. You can define the board circuit models, which are also known as board trace models, in the Quartus II software. The board trace model settings are shared with the models used during advanced I/O timing.

You can define an overall board trace model for each I/O standard in your design; this overall board trace model is the default model for all pins that use a particular I/O standard. After configuring the overall board trace model, you can customize the model for specific pins. The parameters you specify for the board trace model are also used in during advanced I/O timing analysis with the TimeQuest Timing Analyzer. If you already specified the board trace models as part of your advanced I/O timing assignments, the same parameters are used during SSN analysis.

All the assignments for board trace models you specify are saved to the **.qsf**. You can also use Tcl commands to create board trace model assignments.

### Tcl Commands for Specifying Board Trace Models

```
set_instance_assignment -name BOARD_MODEL_TLINE_L_PER_LENGTH "3.041E-7" -to e[0]
set_instance_assignment -name BOARD_MODEL_TLINE_LENGTH 0.1391 -to e[0]
set_instance_assignment -name BOARD_MODEL_TLINE_C_PER_LENGTH "1.463E-10" -to e[0]
```

The best way to calculate transmission line parameters is to use a two-dimensional solver to estimate the inductance per inch and capacitance per inch for the transmission line. The termination resistor topology information can be obtained from the PCB schematics. The near-end and far-end pin load (capacitance) values can be obtained from the PCB schematic and other device data sheets. For example, if you know that an FPGA pin is driving a DIMM, you can obtain the far-end loading information in the data sheet for your target device.

**Related Information**

- **Understanding the SSN Reports** on page 5-15  
For more information about the default parameters used by the SSN Analyzer and SSN confidence levels reported in the Confidence Metric Details Report.

- **I/O Management** on page 4-1  
For more information about defining board trace models and advanced I/O timing, refer to the *Quartus II Handbook*.
- **Using Advanced I/O Timing**  
For more information about defining a board trace model for your entire design, refer to Quartus II Help.
- **Board Trace Model**  
For more information about configuring component values for a board trace model, including a complete list of the supported unit prefixes and setting the values with Tcl scripts, refer to Quartus II Help.
- **Literature and Technical Documentation**  
For more information, refer to the *Device Family Data Sheet* in the appropriate device handbook, available on the Altera website.

## Defining PCB Layers and PCB Layer Thickness

Every PCB is fabricated using a number of layers. To remove some of the pessimism from your SSN results, Altera recommends that you create assignments describing your PCB layers in the Quartus II software.

You can specify the number of layers on your PCB, and their thickness. The PCB layer information is used only during SSN analysis and is not used in other processes run by the Quartus II software. If a custom PCB breakout region is not described you can select the default thickness, which directs the SSN Analyzer to use a single-layer PCB breakout region during SSN analysis.

All the assignments you create for the PCB layers are saved to the **.qsf**. You can also use Tcl commands to create PCB layer assignments. You can create any number of PCB layers, however, the layers must be consecutive.

### Tcl Commands for Specifying PCB Layer Assignments

```
set_global_assignment -name PCB_LAYER_THICKNESS 0.00099822M -section_id 1
set_global_assignment -name PCB_LAYER_THICKNESS 0.00034036M -section_id 2
set_global_assignment -name PCB_LAYER_THICKNESS 0.00034036M -section_id 3
```

The cross-section shows the stackup information of a PCB, which tells you the number of layers used in your PCB. The PCB shown in this example consists of various signal and circuit layers on which FPGA pins are routed, as well as the power and ground layers.

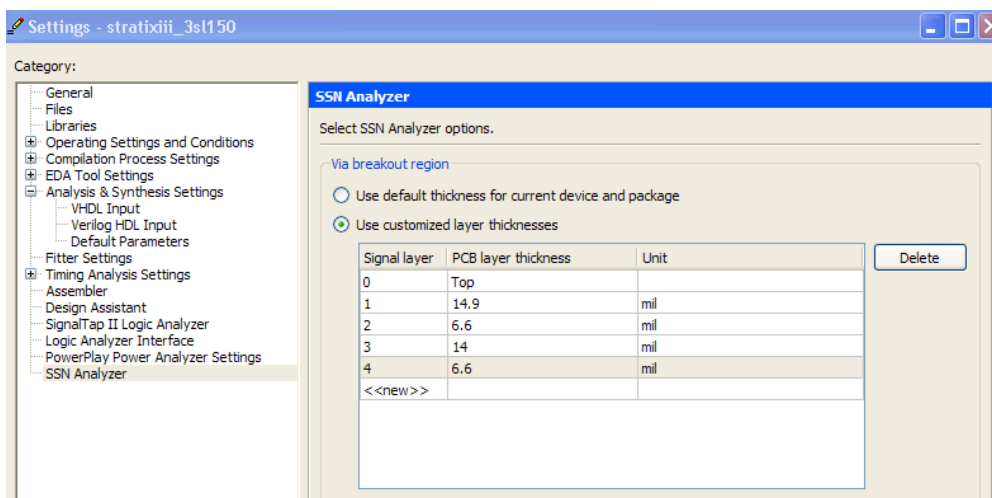
Figure 5-9: Snapshot of Stackup of a PCB Shown in the Allegro Board Design Environment

	Subclass Name	Type	Material	Thickness (MIL)
1		SURFACE	AIR	
2	TOP	CONDUCTOR	PLATED_COPPER_FOIL	1.9
3		DIELECTRIC	FR-4	3.6
4	L2-PWR	PLANE	COPPER	1.2
5		DIELECTRIC	FR-4	3
6	L3-GND	PLANE	COPPER	1.2
7		DIELECTRIC	FR-4	4
8	L4-SIGNAL	CONDUCTOR	COPPER	0.6
9		DIELECTRIC	FR-4	6
10	L5-SIGNAL	CONDUCTOR	COPPER	0.6
11		DIELECTRIC	FR-4	4
12	L6 GND	PLANE	COPPER	1.2
13		DIELECTRIC	FR-4	3
14	L7-PWR	PLANE	COPPER	1.2
15		DIELECTRIC	FR-4	4
16	L8-SIGNAL	CONDUCTOR	COPPER	0.6
17		DIELECTRIC	FR-4	6
18	L9-SIGNAL	CONDUCTOR	COPPER	0.6
19		DIELECTRIC	FR-4	4
20	L10-GND	PLANE	COPPER	1.2
21		DIELECTRIC	FR-4	3

In this example, each of the four signal layers are a different thickness, with the depths shown in the **Thickness (MIL)** column. The layer thickness for each signal layer is computed as follows:

- Signal Layer 1 is the L4-SIGNAL, at thickness  $(1.9+3.6+1.2+3+1.2+4=)$  14.9 mils
- Signal Layer 2 is the L5-SIGNAL, at thickness  $(0.6+6=)$  6.6 mils
- Signal Layer 3 is the L8-SIGNAL, at thickness  $(0.6+4+1.2+3+1.2+4=)$  14 mils
- Signal Layer 4 is the L9-SIGNAL, at thickness  $(0.6+6=)$  6.6 mils

Figure 5-10: PCB Layers and Thickness Assignments Specified in the Quartus II Software



### Related Information

#### [Running the SSN Analyzer](#)

For more information about specifying PCB layer information, refer to Quartus II Help.

## Specifying Signal Breakout Layers

Each user I/O pin in your FPGA device can break out at different layers on your PCB. In the Pin Planner, you can specify on which layers the I/O pins in your design break out.

The breakout layer information is used only during SSN analysis and is not used in other processes run by the Quartus II software. To assign a pin to PCB layer, follow these steps:

1. On the Assignments menu, click **Pin Planner**.
2. If necessary, perform Analysis & Elaboration, Analysis & Synthesis, or fully compile the design to populate the Pin Planner with the node names in the design.
3. Right-click anywhere in the **All Pins** or **Groups** list, and then click **Customize Columns**.
4. Select the **PCB layer** column and move it from the **Available columns** list to the **Show these columns in this order** list.
5. Click **OK**.
6. In the **PCB layer** column, specify the PCB layer to which you want to connect the signal.
7. On the File menu, click **Save Project** to save the changes.

**Note:** When you create PCB breakout layer assignments in the Pin Planner, you can assign the pin to any layer, even if you did not yet define the PCB layer.

## Creating I/O Assignments

I/O assignments are required in FPGA design and are also used during SSN analysis to estimate voltage noise.

Each input, output, or bidirectional signal in your design is assigned a physical pin location on the device using pin location assignments. Each signal has a physical I/O buffer that has a specific I/O standard, pin location, drive strength, and slew rate. The SSN Analyzer supports most I/O standards in a device family, such as the **LVTTL** and **LVC MOS** I/O standards.

**Note:** The SSN Analyzer does not support differential I/O standards, such as the **LVDS** I/O standard and its variations, because differential I/O standards contribute a small amount of SSN.

### Related Information

#### [Literature and Technical Documentation](#)

For more information on the Altera website about supported I/O standards.

[I/O Management](#) on page 4-1

For more information about creating and managing I/O assignments, refer to the *Quartus II Handbook*.

## Decreasing Pessimism in SSN Analysis

In the absence of specific timing information, the SSN Analyzer analyzes your design under worst-case conditions.

Worst-case conditions include all pins acting as aggressor signals on all possible victim pins and all aggressor pins switching with the worst possible timing relationship. The results of SSN analysis under worst-case conditions are very pessimistic. You can improve the results of SSN Analysis by creating group

assignments for specific types of pins. Use the following group assignments to decrease the pessimism in SSN analysis results:

- Assign pins to an output enable group—All pins in an output enable group must be either all input pins or all output pins. If all the pins in a group are always either all inputs or all outputs, it is impossible for an output pin in the group to cause SSN noise on an input pin in the group. You can assign pins to an output enable group with the **Output Enable Group** logic option.
- Assign pins to a synchronous group—I/O pins that are part of a synchronous group (signals that switch at the same time) may cause SSN, but do not result in any failures because the noise glitch occurs during the switching period of the signal. The noise, therefore, does not occur in the sampling window of that signal. You can assign pins to an output enable group with the **Synchronous Group** logic option. For example, in your design you have a bus with 32 pins that all belong to the same group. In a real operation, the bus switches at the same time, so any voltage noise induced by a pin on its groupmates does not matter, because it does not fall in the sampling window. If you do not assign the bus to a synchronous group, the other 31 pins can act as aggressors for the first pin in that group, leading to higher QL and QH noise levels during SSN analysis.

In some cases, the SSN Analyzer can detect the grouping for bidirectional pins by looking at the output enable signal of the bidirectional pins. However, Altera recommends that you explicitly specify the bidirectional groups and output groups in your design.

#### Related Information

##### [Assigning Device I/O Pins](#)

For more information about creating logic option assignments, refer to Quartus II Help.

## Excluding Pins as Aggressor Signals

The SSN Analyzer uses the following conditions to exclude pins as aggressor signals for a specific victim pin:

- A pin that is a complement of the victim pin. For example, any pin that is assigned a differential I/O standard cannot be an aggressor pin.
- A programming pin or JTAG pin because these pins are not active in user mode.
- Pins that have the same output enable signal as a bidirectional victim pin that the SSN Analyzer analyzes as an input pin. Pins with the same output enable signal also act as input pins and therefore cannot be aggressor pins at the same time.
- Pins in the same synchronous group as a victim output pin.
- A pin assigned the **I/O Maximum Toggle Rate** logic option with a frequency setting of zero. The SSN Analyzer does not consider pins with this setting as aggressor pins.

#### Related Information

- [Performing SSN Analysis and Viewing Results](#) on page 5-14

For information about grouping output pins, or about grouping bidirectional pins.

- [Assigning Device I/O Pins](#)

For more information about creating pin assignments with the Pin Planner, refer to Quartus II Help.

## Performing SSN Analysis and Viewing Results

You can perform SSN analysis either on your entire design, or you can limit the analysis to specific I/O banks.

If you know the problem area for SSN is within one I/O bank and you are changing pin assignments only in that bank, you can run SSN analysis for just that one I/O bank to reduce analysis time.

## Related Information

- [Running the SSN Analyzer](#)

For more information on performing SSN analysis, refer to Quartus II Help.

- [Literature and Technical Documentation](#)

For more information about I/O bank numbering refer to the appropriate device handbook available on the Altera website.

## Understanding the SSN Reports

When SSN analysis is complete, you can view detailed analysis reports. The detailed messages in the reports help you understand and resolve SSN problems.

The SSN Analyzer section of the Compilation report contains information generated during SSN analysis, including the following reports:

- Summary
- Output Pins
- Input Pins
- Unanalyzed Pins
- Confidence Metric Details

### Summary Report

The Summary report summarizes the SSN Analyzer status and rates the SSN Analyzer confidence level as low, medium, or high.

The confidence level depends on the completeness of your board trace model assignments. The more assignments you complete, the higher the confidence level. However, the confidence level does not always contribute to the accuracy of the QL and QH noise levels on a victim pin. The accuracy of QH and QL noise levels depends the accuracy of your board trace model assignments.

### Output Pins and Input Pins Reports

The Output Pins report lists all of the output pins and bidirectional pins that are treated as output pins during SSN analysis.

The Input Pins report lists all of the input pins and bidirectional pins that are treated as inputs during SSN analysis. Both reports list the location assignments for the pins treated as SSN outputs or inputs during SSN analysis, the QL and QH noise in volts, and what percentage the QL and QH margins are for the I/O standard used for that signal. The QH and QL noise margins that fall in the critical range (> 90%) are shown in red. The QH and QL noise margins that fall in the range of 70% to 90% are shown in gray.

### Unanalyzed Pins Report

Not all pins are analyzed for SSN analysis. The following pins are not analyzed and are reported in the Unanalyzed Pins report:

- Pins assigned the LVDS I/O standard or any LVDS variations, such as the **mini-LVDS** I/O standard.
- Pins created in the migration flow that cover power and supply pins in other packages.
- The negative terminals of pseudo-differential I/O standards; the noise on differential standards is reported as the differential noise and is reported on the positive terminal.

## Confidence Metric Details

The Confidence Metric Details Report lists the values used during SSN Analysis for unspecified I/O, board, and PCB assignments.

## Viewing SSN Analysis Results in the Pin Planner

After SSN analysis completes, you can analyze the results in the Pin Planner. In the Pin Planner you can identify the SSN hotspots in your device, as well as the QL and QH noise levels.

The QL and QH results for each pin are displayed with a different color that represents whether the pin is below the warning threshold, below the critical threshold, or above the critical threshold. This color representation is also referred to as the SSN map of your FPGA device.

When you view the SSN map, you can customize which details to display, including input pins, output pins, QH signals, QL signals, and noise levels. You can also adjust the threshold levels for QH and QL noise voltages. Adjusting the threshold levels in the Pin Planner does not change the threshold levels reported during SSN analysis and does not change the data in any of the SSN reports.

You can also change I/O assignments and board trace information and rerun the SSN Analyzer to view the SSN analysis results based on those modified settings.

### Related Information

- [Show SSN Analyzer Results](#)
- [Running the SSN Analyzer](#)

For more information, refer to Quartus II Help.

## Decreasing Processing Time for SSN Analysis

FPGA designs are getting larger in density, logic, and I/O count. The time it takes to complete SSN analysis and other Quartus II software processes affects your development time.

Faster processing times can reduce your design cycle time. Use the following guidelines to reduce processing time:

- Direct the Quartus II software to use more than one processor for parallel executables, including the SSN Analyzer
- Perform SSN analysis after I/O assignment analysis if your design files and constraints are complete, and you are interested in generating the SSN results early in the design process and want to adjust I/O placements to see if you can obtain better results
- Perform SSN analysis after fitting if you want to view preliminary SSN results that do not take into account complete I/O assignment and I/O timing results
- Perform engineering change orders (ECOs) on your design, rather than recompiling the entire design, if you want to rerun SSN analysis after changing I/O assignments

### Related Information

- [Setting Up and Running Analysis and Synthesis](#)  
For more information about using parallel processors, refer to Quartus II Help.
- [Compilation Process Settings Page](#)  
For more information about using parallel processors, refer to Quartus II Help.
- [Assigning Device I/O Pins](#)  
For more information about performing I/O assignment analysis, refer to Quartus II Help.



- **Setting Up and Running the Fitter**  
For more information about running the Fitter, refer to Quartus II Help.
- **Engineering Change Management with the Chip Planner** on page 17-1  
For more information about performing ECOs on your design, refer to the *Quartus II Handbook*.

## Scripting Support

A Tcl script allows you to run procedures and determine settings. You can also run some of these procedures at a command prompt.

The Quartus II software provides several packages to compile your design and create I/O assignments for analysis and fitting. You can create a custom Tcl script that maps the design and runs SSN analysis on your design.

For detailed information about specific scripting command options and Tcl API packages, type the following command at a system command prompt to run the Quartus II Command-Line and Tcl API Help browser:

```
quartus_sh --qhelp
```

### Related Information

[Tcl Scripting](#) on page 3-1

[Command-Line Scripting](#) on page 2-1

For more information about Quartus II scripting support, including examples, refer to the *Quartus II Handbook*.

### API Functions for Tcl

For more information about Quartus II scripting support, including examples, refer to Quartus II Help.

## Optimizing Pin Placements for Signal Integrity

You can create an assignment that directs the Fitter to optimize pin placements for signal integrity with a Tcl command.

The following Tcl command directs the Fitter to optimize pin placement for signal integrity without affecting design  $f_{MAX}$ :

```
set_global_assignment -name OPTIMIZE_SIGNAL_INTEGRITY "Normal Compilation"
```

### Related Information

[Optimizing Pin Placements for Signal Integrity](#) on page 5-9

## Defining PCB Layers and PCB Layer Thickness

You can create PCB layer and thickness assignments with a Tcl command.

### Tcl Commands for Specifying PCB Layer Assignments

```
set_global_assignment -name PCB_LAYER_THICKNESS 0.00099822M -section_id 1
set_global_assignment -name PCB_LAYER_THICKNESS 0.00034036M -section_id 2
set_global_assignment -name PCB_LAYER_THICKNESS 0.00034036M -section_id 3
set_global_assignment -name PCB_LAYER_THICKNESS 0.00055372M -section_id 4
set_global_assignment -name PCB_LAYER_THICKNESS 0.00034036M -section_id 5
```

```
set_global_assignment -name PCB_LAYER_THICKNESS 0.00034036M -section_id 6
set_global_assignment -name PCB_LAYER_THICKNESS 0.00082042M -section_id 7
```

These Tcl commands specify that there are seven PCB layers in the design, each with a different thickness. In each assignment, the letter *M* indicates the unit of measurement is millimeters. When you specify PCB layer assignments with Tcl commands, you must list the layers in consecutive order. For example, you would receive an error during SSN Analysis if your Tcl commands created the following assignments:

```
set_global_assignment -name PCB_LAYER_THICKNESS 0.00099822M -section_id 1
set_global_assignment -name PCB_LAYER_THICKNESS 0.00082042M -section_id 7
```

To create assignments with the unit of measurement in mils, refer to the syntax in the following Tcl commands.

```
set_global_assignment -name PCB_LAYER_THICKNESS 14.9MIL -section_id 1
set_global_assignment -name PCB_LAYER_THICKNESS 6.6MIL -section_id 2
set_global_assignment -name PCB_LAYER_THICKNESS 14MIL -section_id 3
set_global_assignment -name PCB_LAYER_THICKNESS 6.6MIL -section_id 4
```

### Related Information

[Defining PCB Layers and PCB Layer Thickness](#) on page 5-11

## Specifying Signal Breakout Layers

You can create signal breakout layer assignments with a Tcl command.:

```
set_instance_assignment -name PCB_LAYER 10 -to e[2] set_instance_assignment -name
PCB_LAYER 3 -to e[3]
```

When you create PCB breakout layer assignments with Tcl commands, if you do not specify a PCB layer, or if you specify a PCB layer that does not exist, the SSN Analyzer breaks out the signal at the bottommost PCB layer.

**Note:** If you create a PCB layer breakout assignment to a layer that does not exist, the SSN Analyzer will generate a warning message.

## Decreasing Pessimism in SSN Analysis

You can create output enable group and synchronous group assignments to help decrease pessimism during SSN Analysis with a Tcl command.

The following Tcl command assigns the bidirectional bus `DATAINOUT` to an output enable group:

```
set_instance_assignment -name OUTPUT_ENABLE_GROUP 1 -to DATAINOUT
```

The following Tcl command assigns the bus `PCI_ADD_io` to a synchronous group:

```
set_instance_assignment -name SYNCHRONOUS_GROUP 1 -to PCI_AD_io
```

### Related Information

[Decreasing Pessimism in SSN Analysis](#) on page 5-13

## Performing SSN Analysis

You can perform SSN analysis with a command-line command. Use the `quartus_si` package that is provided with the Quartus II software.

Type the following command at a system command prompt to start the SSN Analyzer:

```
quartus_si <project name>
```

To analyze just one I/O bank, type the following command at a system command prompt:

```
quartus_si <project revision> --bank = bank id>
```

For example, to run analyze the I/O bank 2A type the following command:

```
quartus_si counter --bank=2A
```

For more information about the `quartus_si` package, type `quartus_si -h` at a system command prompt.

**Related Information**

[Performing SSN Analysis and Viewing Results](#) on page 5-14

## Document Revision History

Date	Version	Changes
December 2014	14.1.0]	<ul style="list-style-type: none"> <li>Minimal text edits for clarity in the topic about understanding SSN.</li> </ul>
June 2014	14.0.0	Updated format.
June 2012	12.0.0	Removed survey link.
November 2011	10.0.2	Template update
December 2010	10.0.1	Template update
July 2010	10.0.0	<ul style="list-style-type: none"> <li>Reorganized and edited the chapter</li> <li>Added links to Quartus II Help for procedural information previously included in the chapter</li> </ul>
November 2009	9.1.0	<ul style="list-style-type: none"> <li>Added “Figure 6–9 shows the layout cross-section of a PCB in the Cadence Allegro PCB tool. The cross-section shows the stackup information of a PCB, which tells you the number of layers used in your PCB. The PCB shown in this example consists of various signal and circuit layers on which FPGA pins are routed, as well as the power and ground layers.” on page 6–12</li> <li>Updated for the Quartus II software 9.1 release</li> </ul>
March 2009	9.0.0	Initial release

**Related Information**

[Quartus II Handbook Archive](#)

For previous versions of the Quartus II Handbook on the Altera website.

2014.06.30

QI15V2



Subscribe



Send Feedback

## Signal Integrity Analysis with Third-Party Tools

With the ever-increasing operating speed of interfaces in traditional FPGA design, the timing and signal integrity margins between the FPGA and other devices on the board must be within specification and tolerance before a single PCB is built.

If the board trace is designed poorly or the route is too heavily loaded, noise in the signal can cause data corruption, while overshoot and undershoot can potentially damage input buffers over time.

As FPGA devices are used in high-speed applications, signal integrity and timing margin between the FPGA and other devices on the printed circuit board (PCB) are important aspects to consider to ensure proper system operation. To avoid time-consuming redesigns and expensive board respins, the topology and routing of critical signals must be simulated. The high-speed interfaces available on current FPGA devices must be modeled accurately and integrated into timing models and board-level signal integrity simulations. The tools used in the design of an FPGA and its integration into a PCB must be “board-aware”—able to take into account properties of the board routing and the connected devices on the board.

The Quartus<sup>®</sup> II software provides methodologies, resources, and tools to ensure good signal integrity and timing margin between Altera<sup>®</sup> FPGA devices and other components on the board. Three types of analysis are possible with the Quartus II software:

- I/O timing with a default or user-specified capacitive load and no signal integrity analysis (default)
- The Quartus II **Enable Advanced I/O Timing** option utilizing a user-defined board trace model to produce enhanced timing reports from accurate “board-aware” simulation models
- Full board routing simulation in third-party tools using Altera-provided or generated Input/Output Buffer Information Specification (IBIS) or HSPICE I/O models

I/O timing using a specified capacitive test load requires no special configuration other than setting the size of the load. I/O timing reports from the Quartus II TimeQuest or the Quartus II Classic Timing Analyzer are generated based only on point-to-point delays within the I/O buffer and assume the presence of the capacitive test load with no other details about the board specified. The default size of the load is based on the I/O standard selected for the pin. Timing is measured to the FPGA pin with no signal integrity analysis details.

The **Enable Advanced I/O Timing** option expands the details in I/O timing reports by taking board topology and termination components into account. A complete point-to-point board trace model is defined and accounted for in the timing analysis. This ability to define a board trace model is an example of how the Quartus II software is “board-aware.”

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



In this case, timing and signal integrity metrics between the I/O buffer and the defined far end load are analyzed and reported in enhanced reports generated by the Quartus II TimeQuest Timing Analyzer.

#### Related Information

- [I/O Management](#) on page 4-1  
For more information about defining capacitive test loads or how to use the **Enable Advanced I/O Timing** option to configure a board trace model.

## Signal Integrity Simulations with HSPICE and IBIS Models

The Quartus II software can export accurate HSPICE models with the built-in HSPICE Writer. You can run signal integrity simulations with these complete HSPICE models in Synopsys HSPICE. IBIS models of the FPGA I/O buffers are also created easily with the Quartus II IBIS Writer.

You can run signal integrity simulations with these complete HSPICE models in Synopsys HSPICE.

You can integrate IBIS models into any third-party simulation tool that supports them, such as the Mentor Graphics® Hyperlynx software. With the ability to create industry-standard model definition files quickly, you can build accurate simulations that can provide data to help improve board-level signal integrity.

The I/O's IBIS and HSPICE model creation available in the Quartus II software can help prevent problems before a costly board respin is required. In general, creating and running accurate simulations is difficult and time consuming. The tools in the Quartus II software automate the I/O model setup and creation process by configuring the models specifically for your design. With these tools, you can set up and run accurate simulations quickly and acquire data that helps guide your FPGA and board design.

The information about signal integrity in this chapter refers to board-level signal integrity based on I/O buffer configuration and board parameters, not simultaneous switching noise (SSN), also known as ground bounce or  $V_{CC}$  sag. SSN is a product of multiple output drivers switching at the same time, causing an overall drop in the voltage of the chip's power supply. This can cause temporary glitches in the specified level of ground or  $V_{CC}$  for the device.

This chapter is intended for FPGA and board designers and includes details about the concepts and steps involved in getting designs simulated and how to adjust designs to improve board-level timing and signal integrity. Also included is information about how to create accurate models from the Quartus II software and how to use those models in simulation software.

The information in this chapter is meant for those who are familiar with the Quartus II software and basic concepts of signal integrity and the design techniques and components in good PCB design. Finally, you should know how to set up simulations and use your selected third-party simulation tool.

#### Related Information

- [AN 315: Guidelines for Designing High-Speed FPGA PCBs](#)  
For a more information about SSN and ways to prevent it.
- [Altera Signal Integrity Center](#)  
For information about basic signal integrity concepts and signal integrity details pertaining to Altera FPGA devices.

## I/O Model Selection: IBIS or HSPICE

The Quartus II software can export two different types of I/O models that are useful for different simulation situations, IBIS models and HSPICE models.

IBIS models define the behavior of input or output buffers through the use of voltage-current (V-I) and voltage-time (V-t) data tables. HSPICE models, often referred to as HSPICE decks, include complete physical descriptions of the transistors and parasitic capacitances that make up an I/O buffer along with all the parameter settings required to run a simulation. The HSPICE decks generated by the Quartus II software are preconfigured with the I/O standard, voltage, and pin loading settings for each pin in your design.

The choice of I/O model type is based on many factors.

**Table 6-1: IBIS and HSPICE Model Comparison**

Feature	IBIS Model	HSPICE Model
I/O Buffer Description	<b>Behavioral</b> —I/O buffers are described by voltage-current and voltage-time tables in typical, minimum, and maximum supply voltage cases.	<b>Physical</b> —I/O buffers and all components in a circuit are described by their physical properties, such as transistor characteristics and parasitic capacitances, as well as their connections to one another.
Model Customization	<b>Simple and limited</b> —The model completely describes the I/O buffer and does not usually have to be customized.	<b>Fully customizable</b> —Unless connected to an arbitrary board description, the description of the board trace model must be customized in the model file. All parameters of the simulation are also adjustable.
Simulation Set Up and Run Time	<b>Fast</b> —Simulations run quickly after set up correctly.	<b>Slow</b> —Simulations take time to set up and take longer to run and complete.
Simulation Accuracy	<b>Good</b> —For most simulations, accuracy is sufficient to make useful adjustments to the FPGA and/or board design to improve signal integrity.	<b>Excellent</b> —Simulations are highly accurate, making HSPICE simulation almost a requirement for any high-speed design where signal integrity and timing margins are tight.
Third-Party Tool Support	<b>Excellent</b> —Almost all third-party board simulation tools support IBIS.	<b>Good</b> —Most third-party tools that support SPICE support HSPICE. However, Synopsys HSPICE is required for simulations of Altera’s encrypted HSPICE models.

### Related Information

#### [AN 283: Simulating Altera Devices with IBIS Models](#)

For more information about IBIS files created by the Quartus II IBIS Writer and IBIS files in general, as well as links to websites with detailed information.

## FPGA to Board Signal Integrity Analysis Flow

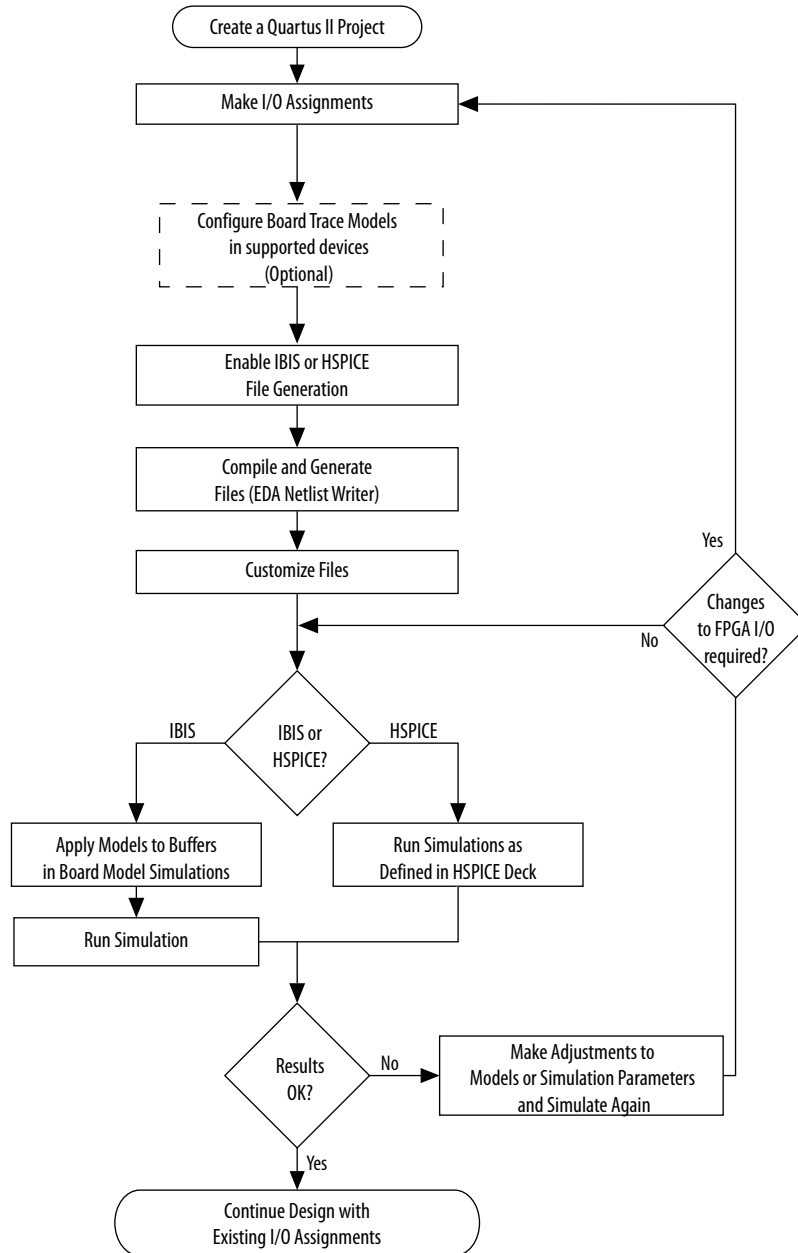
Board signal integrity analysis can take place at any point in the FPGA design process and is often performed before and after board layout. If it is performed early in the process as part of a pre-PCB layout analysis, the models used for simulations can be more generic.

These models can be changed as much as required to see how adjustments improve timing or signal integrity and help with the design and routing of the PCB. Simulations and the resulting changes made at this stage allow you to analyze “what if” scenarios to plan and implement your design better. To assist with early board signal integrity analysis, you can download generic IBIS model files for each device family and obtain HSPICE buffer simulation kits from the “Board Level Tools” section of the EDA Tool Support Resource Center.

Typically, if board signal integrity analysis is performed late in the design, it is used for a post-layout verification. The inputs and outputs of the FPGA are defined, and required board routing topologies and constraints are known. Simulations can help you find problems that might still exist in the FPGA or board design before fabrication and assembly. In either case, a simple process flow illustrates how to create accurate IBIS and HSPICE models from a design in the Quartus II software and transfer them to third-party simulation tools.

Your design depends on the type of model, IBIS or HSPICE, that you use for your simulations. When you understand the steps in the analysis flow, refer to the section of this chapter that corresponds to the model type you are using.

Figure 6-1: Third-Party Board Signal Integrity Analysis Flow



**Related Information**

**[EDA Tool Support Resource Center](#)**

For more information, generic IBIS model files for each device family, and to obtain HSPICE buffer simulation kits.

**Create I/O and Board Trace Model Assignments**

You can configure a board trace model for output signals or for bidirectional signals in output mode. You can then automatically transfer its description to HSPICE decks generated by the HSPICE Writer. This helps improve simulation accuracy.



To configure a board trace model, in the **Settings** dialog box, in the **TimeQuest Timing Analyzer** page, turn on the **Enable Advanced I/O Timing** option and configure the board trace model assignment settings for each I/O standard used in your design. You can add series or parallel termination, specify the transmission line length, and set the value of the far-end capacitive load. You can configure these parameters either in the Board Trace Model view of the Pin Planner, or click **SettingsDeviceDevice and Pin Options**.

The Quartus II software can generate IBIS models and HSPICE decks without having to configure a board trace model with the **Enable Advanced I/O Timing** option. In fact, IBIS models ignore any board trace model settings other than the far-end capacitive load. If any load value is set other than the default, the delay given by IBIS models generated by the IBIS Writer cannot be used to account correctly for the double counting problem. The load value mismatch between the IBIS delay and the  $t_{CO}$  measurement of the Quartus II software prevents the delays from being safely added together. Warning messages displayed when the EDA Netlist Writer runs indicate when this mismatch occurs.

#### Related Information

- [I/O Management](#) on page 4-1  
For information about how to use the **Enable Advanced I/O Timing** option and configure board trace models for the I/O standards used in your design.

## Output File Generation

IBIS and HSPICE model files are not generated by the Quartus II software by default. To generate or update the files automatically during each project compilation, select the type of file to generate and a location where to save the file in the project settings.

The IBIS and HSPICE Writers in the Quartus II software are run as part of the EDA Netlist Writer during normal project compilation. If either writer is turned on in the project settings, IBIS or HSPICE files are created and stored in the specified location. For IBIS, a single file is generated containing information about all assigned pins. HSPICE file generation creates separate files for each assigned pin. You can run the EDA Netlist Writer separately from a full compilation in the Quartus II software or at the command line. However, you must fully compile the project or perform I/O Assignment Analysis at least once for the IBIS and HSPICE Writers to have information about the I/O assignments and settings in the design.

**Note:** These settings can also be specified with commands in a Tcl script.

## Customize the Output Files

The files generated by either the IBIS or HSPICE Writer are text files that you can edit and customize easily for design or experimentation purposes.

IBIS files downloaded from the Altera website must be customized with the correct RLC values for the specific device package you have selected for your design. IBIS files generated by the IBIS Writer do not require this customization because they are configured automatically with the RLC values for your selected device. HSPICE decks require modification to include a detailed description of your board. With **Enable Advanced I/O Timing** turned on and a board trace model defined in the Quartus II software, generated HSPICE decks automatically include that model's parameters. However, Altera recommends that you replace that model with a more detailed model that describes your board design more accurately. A default simulation included in the generated HSPICE decks measures delay between the FPGA and the far-end device. You can make additions or adjustments to the default simulation in the generated files to change the parameters of the default simulation or to perform additional measurements.

## Set Up and Run Simulations in Third-Party Tools

When you have generated the files, you can use them to perform simulations in your selected simulation tool.

With IBIS models, you can apply them to input, output, or bidirectional buffer entities and quickly set up and run simulations. For HSPICE decks, the simulation parameters are included in the files. Open the files in Synopsys HSPICE and run simulations for each pin as required.

With HSPICE decks generated from the HSPICE Writer, the double counting problem is accounted for, which ensures that your simulations are accurate. Simulations that involve IBIS models created with anything other than the default loading settings in the Quartus II software must take the change in the size of the load between the IBIS delay and the Quartus II  $t_{CO}$  measurement into account. Warning messages during compilation alert you to this change.

## Interpret Simulation Results

If you encounter timing or signal integrity issues with your high-speed signals after running simulations, you can make adjustments to I/O assignment settings in the Quartus II software.

These could include such things as drive strength or I/O standard, or making changes to your board routing or topology. After regenerating models in the Quartus II software based on the changes you have made, rerun the simulations to check whether your changes corrected the problem.

## Simulation with IBIS Models

IBIS models provide a way to run accurate signal integrity simulations quickly. IBIS models describe the behavior of I/O buffers with voltage-current and voltage-time data curves.


Because of their behavioral nature, IBIS models do not have to include any information about the internal circuit design of the I/O buffer. Most component manufacturers, including Altera, provide IBIS models for free download and use in signal integrity analysis simulation tools. You can download generic device family IBIS models from the Altera website for early design simulation or use the IBIS Writer to create custom IBIS models for your existing design.

## Elements of an IBIS Model

An IBIS model file (**.ibs**) is a text file that describes the behavior of an I/O buffer across minimum, typical, and maximum temperature and voltage ranges with a specified test load.

The tables and values specified in the IBIS file describe five basic elements of the I/O buffer.

Figure 6-2: Five Basic Elements of an I/O Buffer in IBIS Models



The following elements correspond to each numbered block.

- 1. Pulldown**—A voltage-current table describes the current when the buffer is driven low based on a pull-down voltage range of  $-V_{CC}$  to  $2 V_{CC}$ .
- 2. Pullup**—A voltage-current table describes the current when the buffer is driven high based on a pull-up voltage range of  $-V_{CC}$  to  $V_{CC}$ .
- 3. Ground and Power Clamps**—Voltage-current tables describe the current when clamping diodes for electrostatic discharge (ESD) are present. The ground clamp voltage range is  $-V_{CC}$  to  $V_{CC}$ , and the power clamp voltage range is  $-V_{CC}$  to ground.
- 4. Ramp and Rising/Falling Waveform**—A voltage-time (dv/dt) ratio describes the rise and fall time of the buffer during a logic transition. Optional rising and falling waveform tables can be added to more accurately describe the characteristics of the rising and falling transitions.
- 5. Total Output Capacitance and Package RLC**—The total output capacitance includes the parasitic capacitances of the output pad, clamp diodes (if present), and input transistors. The package RLC is device package-specific and defines the resistance, inductance, and capacitance of the bond wire and pin of the I/O.

#### Related Information

##### [AN 283: Simulating Altera Devices with IBIS Models](#)

For more information about IBIS models and Altera-specific features, including links to the official IBIS specification.

## Creating Accurate IBIS Models

There are two methods to obtain Altera device IBIS files for your board-level signal integrity simulations. You can download generic IBIS models from the Altera website. You can also use the IBIS writer in the Quartus II software to create design-specific models.

The IBIS file generated by the Quartus II software contains models of both input and output termination, and is supported for IBIS model versions of 4.2 and later. Arria V, Cyclone V, and Stratix V device families allow the use of bidirectional I/O with dynamic on-chip termination (OCT).

Dynamic OCT is used where a signal uses a series on-chip termination during output operation and a parallel on-chip termination during input operation. Typically this is used in Altera External Memory Interface IP.

The Quartus II IBIS dynamic OCT IBIS model names end in **g50c\_r50c**. For example: **sst115i\_ctnio\_g50c\_r50c**.

In the simulation tool, the IBIS model is attached to a buffer.

- When the buffer is assigned as an output, use the series termination r50c.
- When the buffer is assigned as an input, use the parallel termination g50c.

## Download IBIS Models

Altera provides IBIS models for almost all FPGA and FPGA configuration devices. You can use the IBIS models from the website to perform early simulations of the I/O buffers you expect to use in your design as part of a pre-layout analysis.

Downloaded IBIS models have the RLC package values set to one particular device in each device family.

The **.ibs** file can be customized for your device package and can be used for any simulation. IBIS models downloaded and used for simulations in this manner are generic. They describe only a certain set of models listed for each device on the Altera IBIS Models page of the Altera website. To create customized models for your design, use the IBIS Writer as described in the next section.

To simulate your design with the model accurately, you must adjust the RLC values in the IBIS model file to match the values for your particular device package by performing the following steps:

1. Download and expand the ZIP file (**.zip**) of the IBIS model for the device family you are using for your design. The **.zip** file contains the **.ibs** file along with an IBIS model user guide and a model data correlation report.
2. Download the Package RLC Values spreadsheet for the same device family.
3. Open the spreadsheet and locate the row that describes the device package used in your design.
4. From the package's **I/O** row, copy the minimum, maximum, and typical values of resistance, inductance, and capacitance for your device package.
5. Open the **.ibs** file in a text editor and locate the [Package] section of the file.
6. Overwrite the listed values copied with the values from the spreadsheet and save the file.

### Related Information

#### [Altera IBIS Models](#)

For information about whether models for your selected device are available.

## Generate Custom IBIS Models with the IBIS Writer

If you have started your FPGA design and have created custom I/O assignments, you can use the Quartus II IBIS Writer to create custom IBIS models to accurately reflect your assignments.

Examples of custom assignments include drive strength settings or the enabling of clamping diodes for ESD protection. IBIS models created with the IBIS Writer take I/O assignment settings into account.

If the **Enable Advanced I/O Timing** option is turned off, the generated **.ibs** files are based on the load value setting for each I/O standard on the **Capacitive Loading** page of the **Device and Pin Options** dialog box in the **Device** dialog box. With the **Enable Advanced I/O Timing** option turned on, IBIS models use an effective capacitive load based on settings found in the board trace model on the **Board Trace Model** page in the **Device and Pin Options** dialog box or the **Board Trace Model** view in the Pin Planner. The effective capacitive load is based on the sum of the **Near capacitance**, **Transmission line distributed capacitance**, and the **Far capacitance** settings in the board trace model. Resistance values and transmission line inductance values are ignored.

**Note:** If you made any changes from the default load settings, the delay in the generated IBIS model cannot safely be added to the Quartus II  $t_{CO}$  measurement to account for the double counting problem. This is because the load values between the two delay measurements do not match. When this happens, the Quartus II software displays warning messages when the EDA Netlist Writer runs to remind you about the load value mismatch.

**Related Information**

- **Generating IBIS Output Files with the Quartus II Software**  
For step-by-step instructions on how to generate IBIS models with the Quartus II software, refer to Quartus II Help.
- **AN 283: Simulating Altera Devices with IBIS Models**  
For more information about IBIS model generation.

## Design Simulation Using the Mentor Graphics HyperLynx® Software

You must integrate IBIS models downloaded from the Altera website or created with the Quartus II IBIS Writer into board design simulations to accurately model timing and signal integrity.

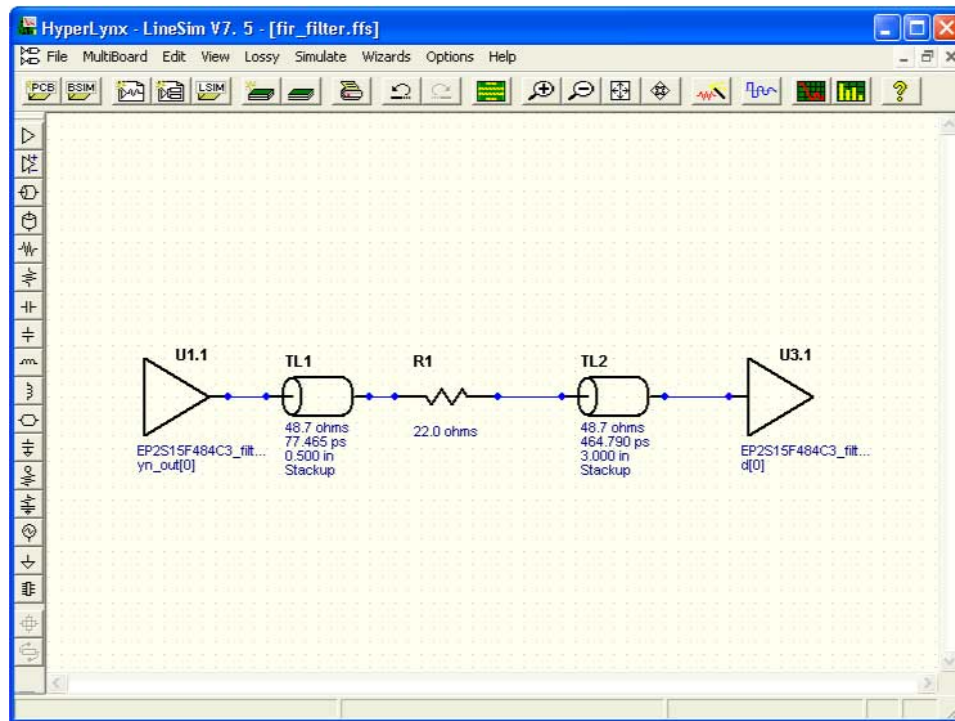
The HyperLynx software from Mentor Graphics is one of the most popular tools for design simulation. The HyperLynx software makes it easy to integrate IBIS models into simulations.

The HyperLynx software is a PCB analysis and simulation tool for high-speed designs, consisting of two products, LineSim and BoardSim. LineSim is an early simulation tool. Before any board routing takes place, LineSim is used to simulate “what if” scenarios to assist in creating routing rules and defining board parameters. BoardSim is a post-layout tool used to analyze existing board routing. Specific nets are selected from a board layout file and simulated in a manner similar to LineSim. With board and routing parameters, and surrounding signal routing known, highly accurate simulations of the final fabricated PCB are possible. This section focuses on LineSim. Because the process of creating and running simulations is very similar for both LineSim and BoardSim, the details of IBIS model use in LineSim applies to simulations in BoardSim.

Simulations in LineSim are configured using a schematic GUI to create connections and topologies between I/O buffers, route trace segments, and termination components. LineSim provides two methods for creating routing schematics: cell-based and free-form. Cell-based schematics are based on fixed cells consisting of typical placements of buffers, trace impedances, and components. Parts of the grid-based cells are filled with the desired objects to create the topology. A topology in a cell-based schematic is limited by the available connections within and between the cells.

A more robust and expandable way to create a circuit schematic for simulation is to use the free-form schematic format in LineSim. The free-form schematic format makes it easy to place parts into any configuration and edit them as required. This section describes the use of IBIS models with free-form schematics, but the process is nearly identical for cell-based schematics.

Figure 6-3: HyperLynx LineSim Free-Form Schematic Editor



When you use HyperLynx software to perform simulations, you typically perform the following steps:

1. Create a new LineSim free-form schematic document and set up the board stackup for your PCB using the Stackup Editor. In this editor, specify board layer properties including layer thickness, dielectric constant, and trace width.
2. Create a circuit schematic for the net you want to simulate. The schematic represents all the parts of the routed net including source and destination I/O buffers, termination components, transmission line segments, and representations of impedance discontinuities such as vias or connectors.
3. Assign IBIS models to the source and destination I/O buffers to represent their behavior during operation.
4. Attach probes from the digital oscilloscope that is built in to LineSim to points in the circuit that you want to monitor during simulation. Typically, at least one probe is attached to the pin of a destination I/O buffer. For differential signals, you can attach a differential probe to both the positive and negative pins at the destination.
5. Configure and run the simulation. You can simulate a rising or falling edge and test the circuit under different drive strength conditions.
6. Interpret the results and make adjustments. Based on the waveforms captured in the digital oscilloscope, you can adjust anything in the circuit schematic to correct any signal integrity issues, such as overshoot or ringing. If necessary, you can make I/O assignment changes in the Quartus II software, regenerate the IBIS file with the IBIS Writer, and apply the updated IBIS model to the buffers in your HyperLynx software schematic.
7. Repeat the simulations and circuit adjustments until you are satisfied with the results. When the operation of the net meets your design requirements, implement changes to your I/O assignments in

the Quartus II software and/or adjust your board routing constraints, component values, and placement to match the simulation.

#### Related Information

[www.mentor.com](http://www.mentor.com)

For more information about HyperLynx software, including schematic creation, simulation setup, model usage, product support, licensing, and training.

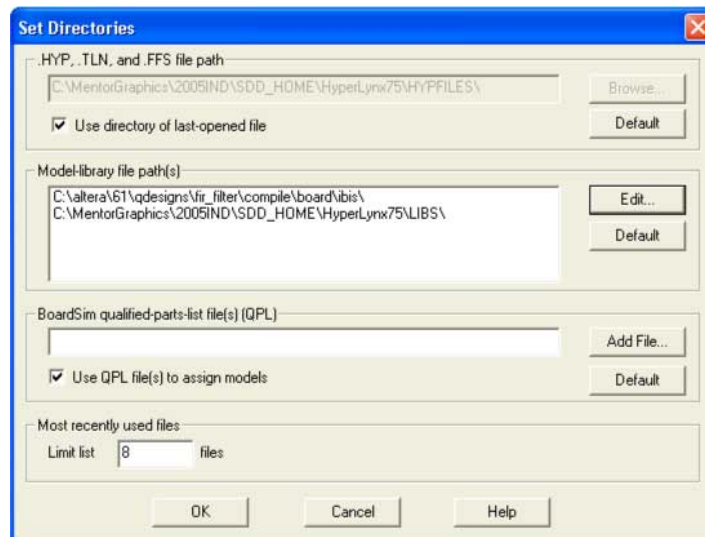
## Configuring LineSim to Use Altera IBIS Models

You must configure LineSim to find and use the downloaded or generated IBIS models for your design. To do this, add the location of your **.ibs** file or files to the LineSim Model Library search path. Then you apply a selected model to a buffer in your schematic.

To add the Quartus II software's default IBIS model location, *<project directory>/board/ibis*, to the HyperLynx LineSim model library search path, perform the following steps in LineSim:

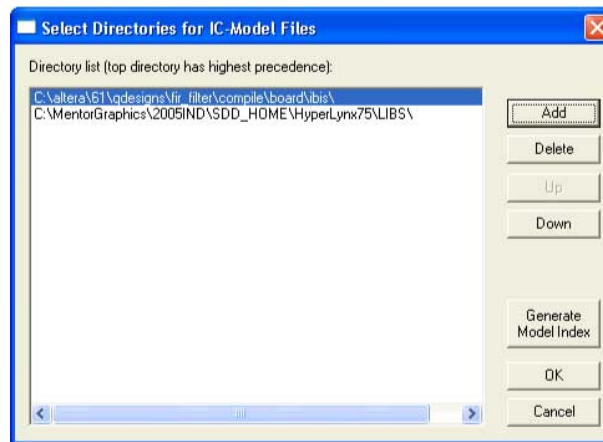
1. From the Options menu, click **Directories**. The **Set Directories** dialog box appears. The **Model-library file path(s)** list displays the order in which LineSim searches file directories for model files.

Figure 6-4: LineSim Set Directories Dialog Box



2. Click **Edit**. A dialog box appears where you can add directories and adjust the order in which LineSim searches them.

Figure 6-5: LineSim Select Directories Dialog Box



3. Click **Add**
4. Browse to the default IBIS model location, *<project directory>/board/ibis*. Click **OK**.
5. Click **Up** to move the IBIS model directory to the top of the list. Click **Generate Model Index** to update LineSim's model database with the models found in the added directory.
6. Click **OK**. The IBIS model directory for your project is added to the top of the Model-library file path(s) list.
7. To close the **Set Directories** dialog box, click **OK**.

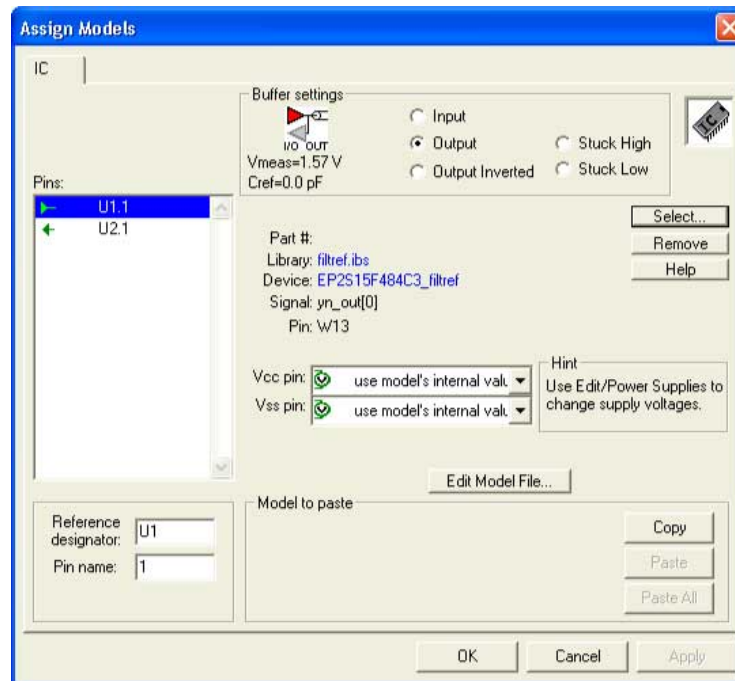
## Integrating Altera IBIS Models into LineSim Simulations

When the location for IBIS files has been set, you can assign the downloaded or generated IBIS models to the buffers in your schematic. To do this, perform the following steps:

1. Double-click a buffer symbol in your schematic to open the **Assign Models** dialog box. You can also click **Assign Models** from the buffer symbol's right-click menu.




Figure 6-6: LineSim Assign Model Dialog Box



2. The pin of the buffer symbol you selected should be highlighted in the **Pins** list. If you want to assign a model to a different symbol or pin, select it from the list.
3. Click **Select**. The **Select IC Model** dialog box appears.

Figure 6-7: LineSim Select IC Model Dialog Box



4. To filter the list of available libraries to display only IBIS models, select **.IBS**. Scroll through the **Libraries** list, and click the name of the library for your design. By default, this is *<project name>.ibs*.


5. The device for your design should be selected as the only item in the **Devices** list. If not, select your device from the list.
6. From the **Signal** list, select the name of the signal you want to simulate. You can also choose to select by device pin number.
7. Click **OK**. The **Assign Models** dialog box displays the selected **.ibs** file and signal.
8. If applicable to the signal you chose, adjust the buffer settings as required for the simulation.
9. Select and configure other buffer pins from the **Pins** list in the same manner.
10. Click **OK** when all I/O models are assigned.

## Running and Interpreting LineSim Simulations

You can now run any desired simulations and make adjustments to the I/O assignments or simulation parameters as required.


For example, if you see too much overshoot in the simulated signal at the destination buffer after running a simulation, you could adjust the drive strength I/O assignment setting to a lower value. Regenerate the **.ibs** file, and run the simulation again to verify whether the change fixed the problem.

**Figure 6-8: Example of Overshoot in HyperLynx with IBIS Models**



If you see a discontinuity or other anomalies at the destination, such as slow rise and fall times, adjust the termination scheme or termination component values. After making these changes, rerun the simulation to check whether your adjustments solved the problem. In this case, it is not necessary to regenerate the **.ibs** file.

Figure 6-9: Example of Signal Integrity Anomaly in HyperLynx with IBIS Models



#### Related Information

##### [Altera Signal Integrity Center](#)

For more information about board-level signal integrity and to learn about ways to improve it with simple changes to your design.

## Simulation with HSPICE Models

HSPICE decks are used to perform highly accurate simulations by describing the physical properties of all aspects of a circuit precisely. HSPICE decks describe I/O buffers, board components, and all of the connections between them, as well as defining the parameters of the simulation to be run.

By their nature, HSPICE decks are highly customizable and require a detailed description of the circuit under simulation. For devices that support advanced I/O timing, when **Enable Advanced I/O Timing** is turned on, the HSPICE decks generated by the Quartus II HSPICE Writer automatically include board components and topology defined in the Board Trace Model. Configure the board components and topology in the Pin Planner or in the **Board Trace Model** tab of the **Device and Pin Options** dialog box. All HSPICE decks generated by the Quartus II software include compensation for the double count problem. You can simulate with the default simulation parameters built in to the generated HSPICE decks or make adjustments to customize your simulation.

#### Related Information

[The Double Counting Problem in HSPICE Simulations](#) on page 6-17

## Supported Devices and Signaling

The HSPICE Writer in the Quartus II software supports Arria, Cyclone, and Stratix devices for the creation of a board trace model in the Quartus II software for automatic inclusion in an HSPICE deck.

The HSPICE files include the board trace description you create in the Board Trace Model view in the Pin Planner or the **Board Trace Model** tab in the **Device and Pin Options** dialog box.

**Note:** Note that for Arria 10 devices, you may need to download the Encrypted HSPICE model from the Altera website.

**Related Information**

[I/O Management](#) on page 4-1

For more information about the **Enable Advanced I/O Timing** option and configuring board trace models for the I/O standards in your design.

[SPICE Models for Altera Devices](#)

For more information about the Encrypted HSPICE model.

## Accessing HSPICE Simulation Kits

You can access the available HSPICE models with the Quartus II software's HSPICE Writer tool and also at the Spice Models for Altera Devices web page.

The Quartus II software HSPICE Writer tool removes many common sources of user error from the I/O simulation process. The HSPICE Writer tool automatically creates preconfigured I/O simulation spice decks that only require the addition of a user board model. All the difficult tasks required to configure the I/O modes and interpret the timing results are handled automatically by the HSPICE Writer tool.

**Related Information**

[Spice Models for Altera Devices](#)

For more information about downloadable HSPICE models.

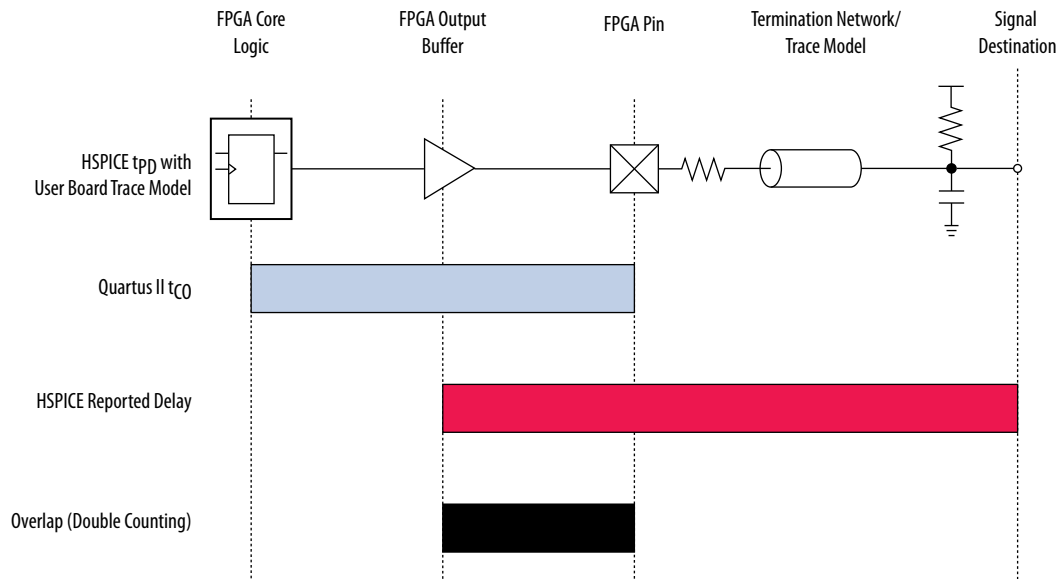
## The Double Counting Problem in HSPICE Simulations

Simulating I/Os using accurate models is extremely helpful for finding and fixing FPGA I/O timing and board signal integrity issues before any boards are built. However, the usefulness of such simulations is directly related to the accuracy of the models used and whether the simulations are set up and performed correctly. To ensure accuracy in models and simulations created for FPGA output signals, the timing hand-off between  $t_{CO}$  timing in the Quartus II software and simulation-based board delay must be taken into account. If this hand-off is not handled correctly, the calculated delay could either count some of the delay twice or even miss counting some of the delay entirely.

### Defining the Double Counting Problem

The double counting problem is inherent to the method output timing is analyzed versus the method used for HSPICE models. The timing analyzer tools in the Quartus II software measure delay timing for an output signal from the core logic of the FPGA design through the output buffer ending at the FPGA pin with a default capacitive load or a specified value for the selected I/O standard. This measurement is the  $t_{CO}$  timing variable.

Figure 6-10: Double Counting Problem



HSPICE models for board simulation measure  $t_{PD}$  (propagation delay) from an arbitrary reference point in the output buffer, through the device pin, out along the board routing, and ending at the signal destination.

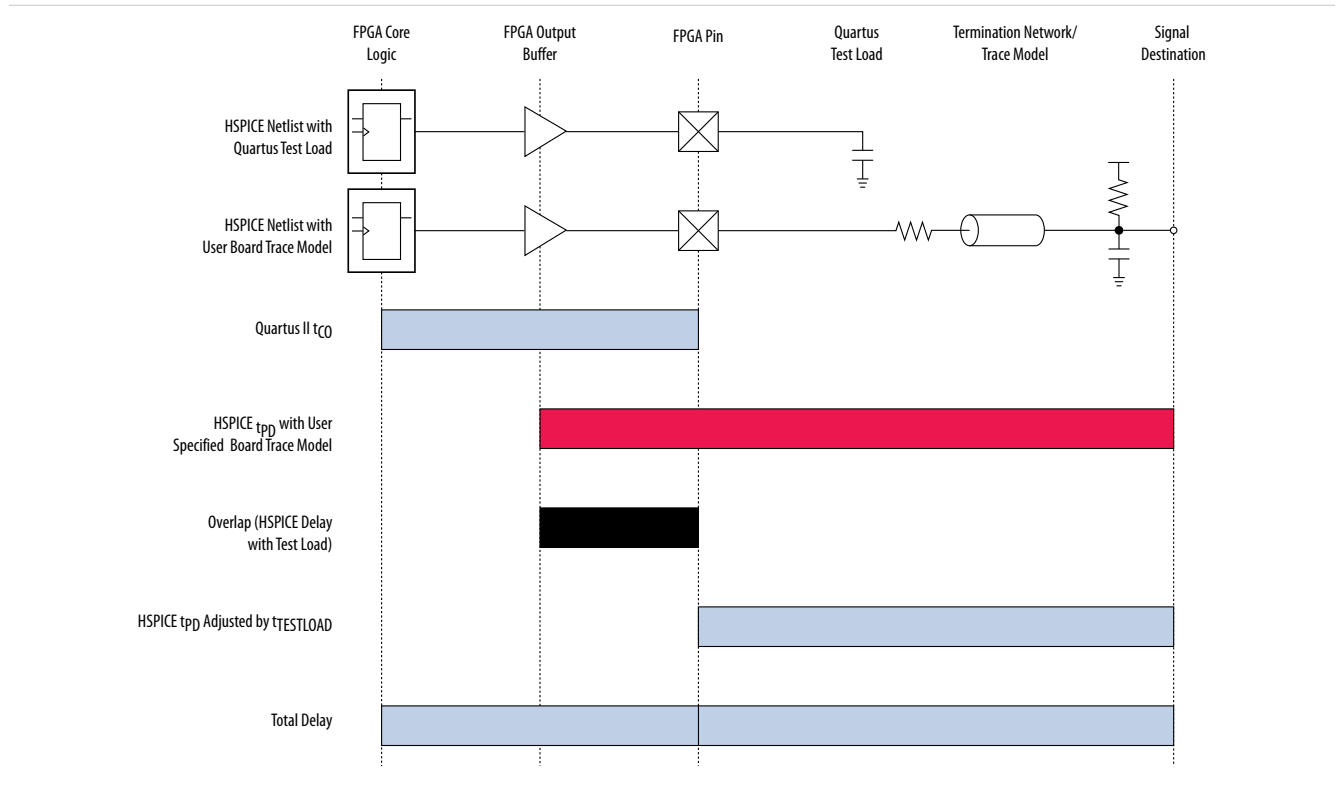
It is apparent immediately that if these two delays were simply added together, the delay between the output buffer and the device pin would be counted twice in the calculation. A model or simulation that does not account for this double count would create overly pessimistic simulation results, because the double-counted delay can limit I/O performance artificially. To fix the problem, it might seem that simply subtracting the overlap between  $t_{CO}$  and  $t_{PD}$  would account for the double count. However, this adjustment would not be accurate because each measurement is based on a different load.

**Note:** Input signals do not exhibit this problem because the HSPICE models for inputs stop at the FPGA pin instead of at the input buffer. In this case, simply adding the delays together produces an accurate measurement of delay timing.

### The Solution to Double Counting

To adjust the measurements to account for the double-counting, the delay between the arbitrary point in the output buffer selected by the HSPICE model and the FPGA pin must be subtracted from either  $t_{CO}$  or  $t_{PD}$  before adding the results together. The subtracted delay must also be based on a common load between the two measurements. This is done by repeating the HSPICE model measurement, but with the same load used by the Quartus II software for the  $t_{CO}$  measurement.

Figure 6-11: Common Test Loads Used for Output Timing



With  $t_{TESTLOAD}$  known, the total delay is calculated for the output signal from the FPGA logic to the signal destination on the board, accounting for the double count.

$$t_{\text{delay}} = t_{CO} + (t_{PD} - t_{TESTLOAD})$$

The preconfigured simulation files generated by the HSPICE Writer in the Quartus II software are designed to account for the double-counting problem based on this calculation automatically. Performing accurate timing simulations is easy without having to make adjustments for double counting manually.

## HSPICE Writer Tool Flow

This section includes information to help you get started using the Quartus II software HSPICE Writer tool. The information in this section assumes you have a basic knowledge of the standard Quartus II software design flow, such as project and assignment creation, compilation, and timing analysis.

### Related Information

#### [Quartus II Handbook](#)

For additional information about standard design flows.

## Applying I/O Assignments

The first step in the HSPICE Writer tool flow is to configure the I/O standards and modes for each of the pins in your design properly. In the Quartus II software, these settings are represented by assignments that map I/O settings, such as pin selection, and I/O standard and drive strength, to corresponding signals in your design.

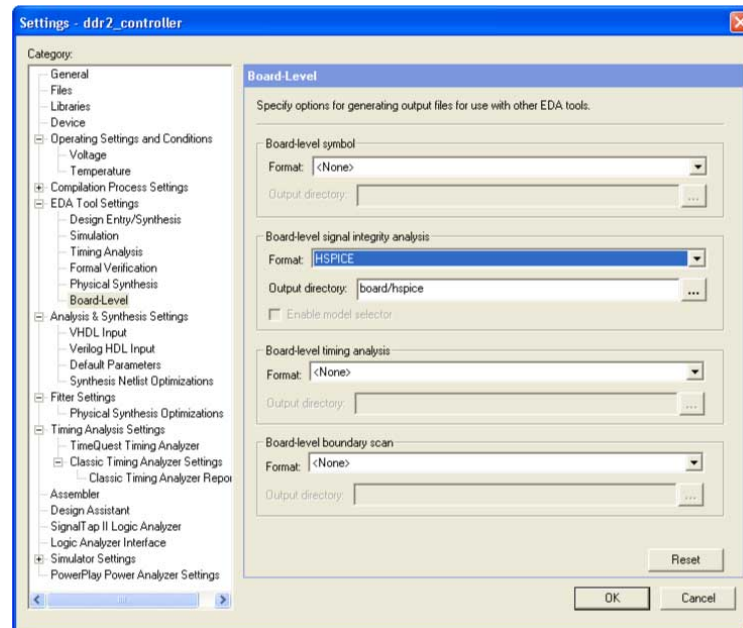
The Quartus II software provides multiple methods for creating these assignments:

- Using the Pin Planner
- Using the assignment editor
- Manually editing the .qsf file
- By making assignments in a scripted Quartus II flow using Tcl

## Enabling HSPICE Writer

You must enable the HSPICE Writer in the **Settings** dialog box of the Quartus II software to generate the HSPICE decks from the Quartus II software.

**Figure 6-12: EDA Tool Settings: Board Level Options Dialog Box**



## Enabling HSPICE Writer Using Assignments

You can also use HSPICE Writer in conjunction with a scripted Tcl flow. To enable HSPICE Writer during a full compile, include the following lines in your Tcl script.

### Enable HSPICE Writer

```
set_global_assignment -name EDA_BOARD_DESIGN_SIGNAL_INTEGRITY_TOOL \
  "HSPICE (Signal Integrity)"
set_global_assignment -name EDA_OUTPUT_DATA_FORMAT HSPICE \
  -section_id eda_board_design_signal_integrity
set_global_assignment -name EDA_NETLIST_WRITER_OUTPUT_DIR <output_directory> \
  -section_id eda_board_design_signal_integrity
```

As with command-line invocation, specifying the output directory is optional. If not specified, the output directory defaults to **board/hspice**.

## Naming Conventions for HSPICE Files

HSPICE Writer automatically generates simulation files and names them using the following naming convention: `<device>_<pin #>_<pin_name>_<in/out>.sp`.

For bidirectional pins, two spice decks are produced; one with the I/O buffer configured as an input, and the other with the I/O buffer configured as an output.

The Quartus II software supports alphanumeric pin names that contain the underscore (\_) and dash (-) characters. Any illegal characters used in file names are converted automatically to underscores.

#### Related Information

- [Sample Output for I/O HSPICE Simulation Deck](#) on page 6-31
- [Sample Input for I/O HSPICE Simulation Deck](#) on page 6-27

### Invoking HSPICE Writer

After HSPICE Writer is enabled, the HSPICE simulation files are generated automatically each time the project is completely compiled. The Quartus II software also provides an option to generate a new set of simulation files without having to recompile manually. In the Processing menu, click **Start EDA Netlist Writer** to generate new simulation files automatically.

**Note:** You must perform both Analysis & Synthesis and Fitting on a design before invoking the HSPICE Writer tool.

### Invoking HSPICE Writer from the Command Line

If you use a script-based flow to compile your project, you can create HSPICE model files by including the following commands in your Tcl script (**.tcl** file).

#### Create HSPICE Model Files

```
set_global_assignment -name EDA_BOARD_DESIGN_SIGNAL_INTEGRITY_TOOL \  
"HSPICE (Signal Integrity)"  
set_global_assignment -name EDA_OUTPUT_DATA_FORMAT HSPICE \  
-section_ideda_board_design_signal_integrity  
set_global_assignment -name EDA_NETLIST_WRITER_OUTPUT_DIR <output_directory> \  
-section_id eda_board_design_signal_integrity
```

The *<output\_directory>* option specifies the location where HSPICE model files are saved. By default, the *<project\_directory>/board/hspice* directory is used.

#### Invoke HSPICE Writer

To invoke the HSPICE Writer tool through the command line, type:

```
quartus_eda.exe <project_name> --board_signal_integrity=on --format=HSPICE \  
--output_directory=<output_directory>
```

*<output\_directory>* specifies the location where the generated spice decks will be written (relative to the design directory). This is an optional parameter and defaults to **board/hspice**.

### Customizing Automatically Generated HSPICE Decks

HSPICE models generated by the HSPICE Writer can be used for simulation as generated.

A default board description is included, and a default simulation is set up to measure rise and fall delays for both input and output simulations, which compensates for the double counting problem. However, Altera recommends that you customize the board description to more accurately represent your routing and termination scheme.



The sample board trace loading in the generated HSPICE model files must be replaced by your actual trace model before you can run a correct simulation. To do this, open the generated HSPICE model files for all pins you want to simulate and locate the following section.

### Sample Board Trace Section

- \* I/O Board Trace and Termination Description
- \* - Replace this with your board trace and termination description

You must replace the example load with a load that matches the design of your PCB board. This includes a trace model, termination resistors, and, for output simulations, a receiver model. The spice circuit node that represents the pin of the FPGA package is called **pin**. The node that represents the far pin of the external device is called **load-in** (for output SPICE decks) and **source-in** (for input SPICE decks).

For an input simulation, you must also modify the stimulus portion of the spice file. The section of the file that must be modified is indicated in the following comment block.

### Sample Source Stimulus Section


- \* Sample source stimulus placeholder
- \* - Replace this with your I/O driver model

Replace the sample stimulus model with a model for the device that will drive the FPGA.

## Running an HSPICE Simulation

Because simulation parameters are configured directly in the HSPICE model files, running a simulation requires only that you open an HSPICE file in the HSPICE user interface and start the simulation.

Figure 6-13: HSPICE User Interface Window



Click **Open** and browse to the location of the HSPICE model files generated by the Quartus II HSPICE Writer. The default location for HSPICE model files is *<project directory>/board/hspice*. Select the **.sp** file generated by the HSPICE Writer for the signal you want to simulate. Click **OK**.

To run the simulation, click **Simulate**. The status of the simulation is displayed in the window and saved in an **.lis** file with the same name as the **.sp** file when the simulation is complete. Check the **.lis** file if an error occurs during the simulation requiring a change in the **.sp** file to fix.

## Interpreting the Results of an Output Simulation

By default, the automatically generated output simulation spice decks are set up to measure three delays for both rising and falling transitions. Two of the measurements, `tpd_rise` and `tpd_fall`, measure the double-counting corrected delay from the FPGA pin to the load pin. To determine the complete clock-edge to load-pin delay, add these numbers to the Quartus II software reported default loading  $t_{CO}$  delay.

The remaining four measurements, `tpd_uncomp_rise`, `tpd_uncomp_fall`, `t_dblcnt_rise`, and `t_dblcnt_fall`, are required for the double-counting compensation process and are not required for further timing usage.

### Related Information

[Simulation Analysis](#) on page 6-31

## Interpreting the Results of an Input Simulation

By default, the automatically generated input simulation SPICE decks are set up to measure delays from the source's driver pin to the FPGA's input pin for both rising and falling transitions.

The propagation delay is reported by HSPICE measure statements as `tpd_rise` and `tpd_fall`. To determine the complete source driver pin-to-FPGA register delay, add these numbers to the Quartus II software reported  $T_H$  and  $T_{SU}$  input timing numbers.

## Viewing and Interpreting Tabular Simulation Results

The `.lis` file stores the collected simulation data in tabular form. The default simulation configured by the HSPICE Writer produces delay measurements for rising and falling transitions on both input and output simulations.


These measurements are found in the `.lis` file and named `tpd_rise` and `tpd_fall`. For output simulations, these values are already adjusted for the double count. To determine the complete delay from the FPGA logic to the load pin, add either of these measurements to the Quartus II  $t_{CO}$  delay. For input simulations, add either of these measurements to the Quartus II  $t_{SU}$  and  $t_H$  delay values to calculate the complete delay from the far end stimulus to the FPGA logic. Other values found in the `.lis` file, such as `tpd_uncomp_rise`, `tpd_uncomp_fall`, `t_dblcnt_rise`, and `t_dblcnt_fall`, are parts of the double count compensation calculation. These values are not necessary for further analysis.

## Viewing Graphical Simulation Results

You can view the results of the simulation quickly as a graphical waveform display using the AvanWaves viewer included with HSPICE. With the default simulation configured by the HSPICE Writer, you can view the simulated waveforms at both the source and destination in input and output simulations.

To see the waveforms for the simulation, in the HSPICE user interface window, click **AvanWaves**. The AvanWaves viewer opens and displays the **Results Browser**.

Figure 6-14: HSPICE AvanWaves Results Browser




The **Results Browser** lets you select which waveform to view quickly in the main viewing window. If multiple simulations are run on the same signal, the list at the top of the **Results Browser** displays the results of each simulation. Click the simulation description to select which simulation to view. By default, the descriptions are derived from the first line of the HSPICE file, so the description might appear as a line of asterisks.

Select the type of waveform to view, by performing the following steps:

1. To see the source and destination waveforms with the default simulation, from the **Types** list, select **Voltages**.
2. On the **Curves** list, double-click the waveform you want to view. The waveform appears in the main viewing window.

You can zoom in and out and adjust the view as desired.


Figure 6-15: AvanWaves Waveform Viewer



## Making Design Adjustments Based on HSPICE Simulations


Based on the results of your simulations, you can make adjustments to the I/O assignments or simulation parameters if required. For example, after you run a simulation and see overshoot or ringing in the simulated signal at the destination buffer, you can adjust the drive strength I/O assignment setting to a lower value. Regenerate the HSPICE deck, and run the simulation again to verify that the change fixed the problem.

Figure 6-16: Example of Overshoot in the AvanWaves Waveform Viewer



If there is a discontinuity or any other anomalies at the destination, adjust the board description in the Quartus II Board Trace Model, or in the generated HSPICE model files to change the termination scheme or adjust termination component values. After making these changes, regenerate the HSPICE files if necessary, and rerun the simulation to verify whether your adjustments solved the problem.

Figure 6-17: Example of Signal Integrity Anomaly in the AvanWaves Waveform Viewer



#### Related Information

#### [Altera Signal Integrity Center](#)

For more information about board-level signal integrity and to learn about ways to improve it with simple changes to your FPGA design.

## Sample Input for I/O HSPICE Simulation Deck

The following sections examine a typical HSPICE simulation spice deck for an I/O of type input. Each section presents the simulation file one block at a time.

### Header Comment

The first block of an input simulation spice deck is the header comment. The purpose of this block is to provide an easily readable summary of how the simulation file has been automatically configured by the Quartus II software.

This block has two main components: The first component summarizes the I/O configuration relevant information such as device, speed grade, and so on. The second component specifies the exact test condition that the Quartus II software assumes for the given I/O standard.

### Sample Header Comment Block

```
* Quartus II HSPICE Writer I/O Simulation Deck*
*
* This spice simulation deck was automatically generated by
* Quartus for the following IO settings:
*
```

```

* Device:          EP2S60F1020C3
* Speed Grade:    C3
* Pin:            AA4 (out96)
* Bank:           IO Bank 6 (Row I/O)
* I/O Standard:   LVTTL, 12mA
* OCT:            Off
*
* Quartus II's default I/O timing delays assume the following slow
* corner simulation conditions.
*
* Specified Test Conditions For Quartus II Tco
* Temperature:     85C (Slowest Temperature Corner)
* Transistor Model: TT (Typical Transistor Corner)
* Vccn:            3.135V (Vccn_min = Nominal - 5%)
* Vccpd:           2.97V (Vccpd_min = Nominal - 10%)
* Load:           No Load
* Vtt:            1.5675V (Voltage reference is Vccn/2)
*
* Note: The I/O transistors are specified to operate at least as
* fast as the TT transistor corner, actual production
* devices can be as fast as the FF corner. Any simulations
* for hold times should be conducted using the fast process
* corner with the following simulation conditions.
*   Temperature:   0C (Fastest Commercial Temperature Corner **)
*   Transistor Model: FF (Fastest Transistor Corner)
*   Vccn:          1.98V (Vccn_hold = Nominal + 10%)
*   Vccpd:         3.63V (Vccpd_hold = Nominal + 10%)
*   Vtt:          0.95V (Vtt_hold = Vccn/2 - 40mV)
*   Vcc:          1.25V (Vcc_hold = Maximum Recommended)
*   Package Model: Short-circuit from pad to pin (no parasitics)
*
* Warnings:

```

## Simulation Conditions

The simulation conditions block loads the appropriate process corner models for the transistors. This condition is automatically set up for the slow timing corner and is modified only if other simulation corners are desired.

### Simulation Conditions Block

```

* Process Settings

.options brief
.inc 'sii_tt.inc' * TT process corner

```

## Simulation Options

The simulation options block configures the simulation temperature and configures HSPICE with typical simulation options.

### Simulation Options Block

```

* Simulation Options

.options brief=0
.options badchr co=132 scale=1e-6 acct ingold=2 nomod dv=1.0
+ dcstep=1 absv=1e-3 absi=1e-8 probe csdf=2 accurate=1
+ converge=1
.temp 85

```

**Note:** For a detailed description of these options, consult your *HSPICE* manual.

## Constant Definition

The constant definition block of the simulation file instantiates the voltage sources that controls the configuration modes of the I/O buffer.

### Constant Definition Block

```
* Constant Definition

voeb      oeb      0      vc * Set to 0 to enable buffer output
vopdrain  opdrain  0      0  * Set to vc to enable open drain
vrambh    rambh    0      0  * Set to vc to enable bus hold
vrpullup  rpullup  0      0  * Set to vc to enable weak pullup
vpcdp5    rpcdp5   0      rp5 * Set the IO standard
vpcdp4    rpcdp4   0      rp4
vpcdp3    rpcdp3   0      rp3
vpcdp2    rpcdp2   0      rp2
vpcdp1    rpcdp1   0      rp1
vpcdp0    rpcdp0   0      rp0
vpcdn4    rpcdn4   0      rn4
vpcdn3    rpcdn3   0      rn3
vpcdn2    rpcdn2   0      rn2
vpcdn1    rpcdn1   0      rn1
vpcdn0    rpcdn0   0      rn0
vdin din   0      0
```

Where:

- Voltage source `voeb` controls the output enable of the buffer and is set to disabled for inputs.
- `vopdrain` controls the open drain mode for the I/O.
- `vrambh` controls the bus hold circuitry in the I/O.
- `vrpullup` controls the weak pullup.
- The next 11 voltages sources control the I/O standard of the buffer and are configured through a later library call.
- `vdin` is not used on input pins because it is the data pin for the output buffer.

### Buffer Netlist

The buffer netlist block of the simulation spice deck loads all the load models required for the corresponding input pin.

#### Buffer Netlist Block

```
* IO Buffer Netlist

.include `vio_buffer.inc`
```

### Drive Strength

The drive strength block of the simulation SPICE deck loads the configuration bits necessary to configure the I/O into the proper I/O standard and drive strengths.

Although these settings are not relevant to an input buffer, they are provided to allow the SPICE deck to be modifiable to support bidirectional simulations.



## Drive Strength Block

```
* Drive Strength Settings
.lib 'drive_select_hio.lib' 3p3ttl_12ma
```

## I/O Buffer Instantiation

The I/O buffer instantiation block of the simulation SPICE deck instantiates the necessary power supplies and I/O model components that are necessary to simulate the given I/O.

### I/O Buffer Instantiation

```
I/O Buffer Instantiation

* Supply Voltages Settings
.param vcn=3.135
.param vpd=2.97
.param vc=1.15

* Instantiate Power Supplies|
vvcc      vcc      0      vc      * FPGA core voltage
vvss      vss      0      0      * FPGA core ground
vvccn     vccn     0      vcn     * IO supply voltage
vvssn     vssn     0      0      * IO ground
vvccpd    vccpd    0      vpd     * Pre-drive supply voltage

* Instantiate I/O Buffer
xvio_buf  din  oeb  opdrain  die  rambh
+ rpcdn4  rpcdn3  rpcdn2  rpcdn1  rpcdn0
+ rpcdp5  rpcdp4  rpcdp3  rpcdp2  rpcdp1  rpcdp0
+ rpullup vccn  vccpd  vcpad0  vio_buf

* Internal Loading on Pad
* - No loading on this pad due to differential buffer/support
*   circuitry

* I/O Buffer Package Model
* - Single-ended I/O standard on a Row I/O
.lib 'lib/package.lib' hio
xpkg die pin hio_pkg
```

## Board Trace and Termination

The board trace and termination block of the simulation SPICE deck is provided only as an example. Replace this block with your own board trace and termination models.

### Board Trace and Termination Block

```
* I/O Board Trace and Termination Description
* - Replace this with your board trace and termination description

wtline pin vssn load vssn N=1 L=1 RLGCMODEL=tlinemodel
.MODEL tlinemodel W MODELTYPE=RLGC N=1 Lo=7.13n Co=2.85p
Rterm2 load vssn 1x
```

## Stimulus Model

The stimulus model block of the simulation spice deck is provided only as a place holder example. Replace this block with your own stimulus model. Options for this include an IBIS or HSPICE model, among others.

### Stimulus Model Block

```
* Sample source stimulus placeholder
* - Replace this with your I/O driver model

Vsource source 0 pulse(0 vcn 0s 0.4ns 0.4ns 8.5ns 17.4ns)
```

## Simulation Analysis

The simulation analysis block of the simulation file is configured to measure the propagation delay from the source to the FPGA pin. Both the source and end point of the delay are referenced against the 50%  $V_{CCN}$  crossing point of the waveform.

### Simulation Analysis Block

```
* Simulation Analysis Setup

* Print out the voltage waveform at both the source and the pin
.print tran v(source) v(pin)
.tran 0.020ns 17ns

* Measure the propagation delay from the source pin to the pin
* referenced against the 50% voltage threshold crossing point

.measure TRAN tpd_rise TRIG v(source) val='vcn*0.5' rise=1
+ TARG v(pin) val = 'vcn*0.5' rise=1
.measure TRAN tpd_fall TRIG v(source) val='vcn*0.5' fall=1
+ TARG v(pin) val = 'vcn*0.5' fall=1
```

## Sample Output for I/O HSPICE Simulation Deck

A typical HSPICE simulation SPICE deck for an I/O-type output has several sections. Each section presents the simulation file one block at a time.

### Header Comment

The first block of an output simulation SPICE deck is the header comment. The purpose of this block is to provide a readable summary of how the simulation file has been automatically configured by the Quartus II software.

This block has two main components:

- The first component summarizes the I/O configuration relevant information such as device, speed grade, and so on.
- The second component specifies the exact test condition that the Quartus II software assumes when generating  $t_{CO}$  delay numbers. This information is used as part of the double-counting correction circuitry contained in the simulation file.

The SPICE decks are preconfigured to calculate the slow process corner delay but can also be used to simulate the fast process corner as well. The fast corner conditions are listed in the header under the notes section.

The final section of the header comment lists any warning messages that you must consider when you use the SPICE decks.

### Header Comment Block

```
* Quartus II HSPICE Writer I/O Simulation Deck
*
* This spice simulation deck was automatically generated by
* Quartus II for the following IO settings:
*
* Device:          EP2S60F1020C3
* Speed Grade:    C3
* Pin:            AA4 (out96)
* Bank:           IO Bank 6 (Row I/O)
* I/O Standard:  LVTTL, 12mA
* OCT:            Off
*
* Quartus' default I/O timing delays assume the following slow
* corner simulation conditions.
* Specified Test Conditions For Quartus II Tco
* Temperature:    85C (Slowest Temperature Corner)
* Transistor Model: TT (Typical Transistor Corner)
* Vccn:           3.135V (Vccn_min = Nominal - 5%)
* Vccpd:          2.97V (Vccpd_min = Nominal - 10%)
* Load:          No Load
* Vtt:            1.5675V (Voltage reference is Vccn/2)
* For C3 devices, the TT transistor corner provides an
* approximation for worst case timing. However, for functionality
* simulations, it is recommended that the SS corner be simulated
* as well.
*
* Note: The I/O transistors are specified to operate at least as
* fast as the TT transistor corner, actual production
* devices can be as fast as the FF corner. Any simulations
* for hold times should be conducted using the fast process
* corner with the following simulation conditions.
* Temperature:    0C (Fastest Commercial Temperature Corner **)
* Transistor Model: FF (Fastest Transistor Corner)
* Vccn:           1.98V (Vccn_hold = Nominal + 10%)
* Vccpd:          3.63V (Vccpd_hold = Nominal + 10%)
* Vtt:            0.95V (Vtt_hold = Vccn/2 - 40mV)
* Vcc:            1.25V (Vcc_hold = Maximum Recommended)
* Package Model:  Short-circuit from pad to pin
* Warnings:
```

### Simulation Conditions

The simulation conditions block loads the appropriate process corner models for the transistors. This condition is automatically set up for the slow timing corner and must be modified only if other simulation corners are desired.

### Simulation Conditions Block

```
* Process Settings

.options brief
.inc 'sii_tt.inc' * typical-typical process corner
```

**Note:** Two separate corners cannot be simulated at the same time. Instead, simulate the base case using the Quartus corner as one simulation and then perform a second simulation using the desired customer corner. The results of the two simulations can be manually added together.

## Simulation Options

The simulation options block configures the simulation temperature and configures HSPICE with typical simulation options.

### Simulation Options Block

```
* Simulation Options
.options brief=0
.options badchr co=132 scale=1e-6 acct ingold=2 nomod dv=1.0
+         dcstep=1 absv=1e-3 absi=1e-8 probe csdf=2 accurate=1
+         converge=1
.temp 85
```

**Note:** For a detailed description of these options, consult your *HSPICE* manual.

## Constant Definition

The constant definition block of the output simulation SPICE deck instantiates the voltage sources that controls the configuration modes of the I/O buffer.

### Constant Definition Block

```
* Constant Definition

voeb      oeb      0      0 * Set to 0 to enable buffer output
vopdrain  opdrain  0      0 * Set to vc to enable open drain
vrambh    rambh    0      0 * Set to vc to enable bus hold
vrpullup  rpullup  0      0 * Set to vc to enable weak pullup
vpci      rpci     0      0 * Set to vc to enable pci mode
vpcdp4    rpcdp4   0      rp4 * These control bits set the IO standard
vpcdp3    rpcdp3   0      rp3
vpcdp2    rpcdp2   0      rp2
vpcdp1    rpcdp1   0      rp1
vpcdp0    rpcdp0   0      rp0
vpcdn4    rpcdn4   0      rn4
vpcdn3    rpcdn3   0      rn3
vpcdn2    rpcdn2   0      rn2
vpcdn1    rpcdn1   0      rn1
vpcdn0    rpcdn0   0      rn0
vdin      din      0      pulse(0 vc 0s 0.2ns 0.2ns 8.5ns 17.4ns)
```

Where:

- Voltage source `voeb` controls the output enable of the buffer.
- `vopdrain` controls the open drain mode for the I/O.
- `vrambh` controls the bus hold circuitry in the I/O.
- `vrpullup` controls the weak pullup.
- `vpci` controls the PCI clamp.
- The next ten voltage sources control the I/O standard of the buffer and are configured through a later library call.
- `vdin` is connected to the data input of the I/O buffer.
- The edge rate of the input stimulus is automatically set to the correct value by the Quartus II software.

## I/O Buffer Netlist

The I/O buffer netlist block loads all of the models required for the corresponding pin. These include a model for the I/O output buffer, as well as any loads that might be present on the pin.

## I/O Buffer Netlist Block

```
*IO Buffer Netlist

.include 'hio_buffer.inc'
.include 'lvds_input_load.inc'
.include 'lvds_oct_load.inc'
```

## Drive Strength

The drive strength block of the simulation spice deck loads the configuration bits for configuring the I/O to the proper I/O standard and drive strength. These options are set by the HSPIICE Writer tool and are not changed for expected use.

## Drive Strength Block

```
* Drive Strength Settings

.lib 'drive_select_hio.lib' 3p3ttl_12ma
```

## Slew Rate and Delay Chain

Stratix and Cyclone devices have sections for configuring the slew rate and delay chain settings.

## Slew Rate and Delay Chain Settings

```
* Programmable Output Delay Control Settings

.lib 'lib/output_delay_control.lib' no_delay

* Programmable Slew Rate Control Settings

.lib 'lib/slew_rate_control.lib' slow_slow
```

## I/O Buffer Instantiation

The I/O buffer instantiation block of the output simulation spice deck instantiates the necessary power supplies and I/O model components that are necessary to simulate the given I/O.

## I/O Buffer Instantiation Block

```
* I/O Buffer Instantiation

* Supply Voltages Settings
.param vcn=3.135
.param vpd=2.97
.param vc=1.15

* Instantiate Power Supplies
vvcc      vcc      0      vc      * FPGA core voltage
vvss      vss      0      0      * FPGA core ground
vvccn     vccn     0      vcn     * IO supply voltage
vvssn     vssn     0      0      * IO ground
vvccpd    vccpd    0      vpd     * Pre-drive supply voltage

* Instantiate I/O Buffer
xhio_buf  din  oeb  opdrain  die  rambh
+ rpcdn4  rpcdn3  rpcdn2  rpcdn1  rpcdn0
+ rpcdp4  rpcdp3  rpcdp2  rpcdp1  rpcdp0
+ rpullup vccn  vccpd  vcpad0  hio_buf
```

```
* Internal Loading on Pad
* - This pad has an LVDS input buffer connected to it, along
*   with differential OCT circuitry. Both are disabled but
*   introduce loading on the pad that is modeled below.
xlvs_input_load die vss vccn lvds_input_load
xlvs_oct_load die vss vccpd vccn vcpad0 vccn lvds_oct_load

* I/O Buffer Package Model
* - Single-ended I/O standard on a Row I/O
.lib 'lib/package.lib' hio
xpkg die pin hio_pkg
```

## Board and Trace Termination

The board trace and termination block of the simulation SPICE deck is provided only as an example. Replace this block with your specific board loading models.

### Board Trace and Termination Block

```
* I/O Board Trace And Termination Description
* - Replace this with your board trace and termination description
wtline pin vssn load vssn N=1 L=1 RLGCMODEL=tlinemodel
.MODEL tlinemodel W MODELTYPE=RLGC N=1 Lo=7.13n Co=2.85p
Rterm2 load vssn 1x
```

## Double-Counting Compensation Circuitry

The double-counting compensation circuitry block of the simulation SPICE deck instantiates a second I/O buffer that is used to measure double-counting. The buffer is configured identically to the user I/O buffer but is connected to the Quartus II software test load. The simulated delay of this second buffer can be interpreted as the amount of double-counting between the Quartus II software and HSPICE Writer simulated results.

As the amount of double-counting is constant for a given I/O standard on a given pin, consider separating the double-counting circuitry from the simulation file. In doing so, you can perform any number of I/O simulations while referencing the delay only once.

### (Part of )Double-Counting Compensation Circuitry Block

```
* Double Counting Compensation Circuitry
*
* The following circuit is designed to calculate the amount of
* double counting between Quartus II and the HSPICE models. If
* you have not changed the default simulation temperature or
* transistor corner the double counting will be automatically
* compensated by this spice deck. In the event you wish to
* simulate an IO at a different temperature or transistor corner
* you will need to remove this section of code and manually
* account for double counting. A description of Altera's
* recommended procedure for this process can be found in the
* Quartus II HSPICE Writer AppNote.
*
* Supply Voltages Settings
.param vcn_tl=3.135
.param vpd_tl=2.97

* Test Load Constant Definition
vopdrain_tl   opdrain_tl   0   0
vrambh_tl    rambh_tl     0   0
```

```

vrpullup_t1  rpullup_t1  0    0

* Instantiate Power Supplies
vvccn_t1     vccn_t1      0    vcn_t1
vvssn_t1     vssn_t1      0    0
vvccpd_t1    vccpd_t1    0    vpd_t1

* Instantiate I/O Buffer
xhio_testload din oeb opdrain_t1 die_t1 rambh_t1
+ rpcdn4 rpcdn3 rpcdn2 rpcdn1 rpcdn0
+ rpcdp4 rpcdp3 rpcdp2 rpcdp1 rpcdp0
+ rpullup_t1 vcn_t1 vccpd_t1 vcpad0_t1 hio_buf

* Internal Loading on Pad
xlvs_input_testload die_t1 vss vccn_t1 lvds_input_load
xlvs_oct_testload die_t1 vss vccpd_t1 vccn_t1 vcpad0_t1 vccn_t1
lvds_oct_load

* I/O Buffer Package Model
* - Single-ended I/O standard on a Row I/O
.lib 'lib/package.lib' hio
xpkg die pin hio_pkg

* Default Altera Test Load
* - 3.3V LVTTTL default test condition is an open load

```

### Related Information

[The Double Counting Problem in HSPICE Simulations](#) on page 6-17

## Simulation Analysis

The simulation analysis block is set up to measure double-counting corrected delays. This is accomplished by measuring the uncompensated delay of the I/O buffer when connected to the user load, and when subtracting the simulated amount of double-counting from the test load I/O buffer.

### Simulation Analysis Block

```

*Simulation Analysis Setup

* Print out the voltage waveform at both the pin and far end load
.print tran v(pin) v(load)
.tran 0.020ns 17ns

* Measure the propagation delay to the load pin. This value will
* include some double counting with Quartus II's Tco
.measure TRAN tpd_uncomp_rise TRIG v(din) val='vc*0.5' rise=1
+ TARG v(load) val='vcn*0.5' rise=1
.measure TRAN tpd_uncomp_fall TRIG v(din) val='vc*0.5' fall=1
+ TARG v(load) val='vcn*0.5' fall=1

* The test load buffer can calculate the amount of double counting
.measure TRAN t_dbcnt_rise TRIG v(din) val='vc*0.5' rise=1
+ TARG v(pin_t1) val='vcn_t1*0.5' rise=1
.measure TRAN t_dbcnt_fall TRIG v(din) val='vc*0.5' fall=1
+ TARG v(pin_t1) val='vcn_t1*0.5' fall=1

* Calculate the true propagation delay by subtraction
.measure TRAN tpd_rise PARAM='tpd_uncomp_rise-t_dbcnt_rise'
.measure TRAN tpd_fall PARAM='tpd_uncomp_fall-t_dbcnt_fall'

```

## Advanced Topics

The information in this section describes some of the more advanced topics and methods employed when setting up and running HSPICE simulation files.

### PVT Simulations

The automatically generated HSPICE simulation files are set up to simulate the slow process corner using low voltage, high temperature, and slow transistors. To ensure a fully robust link, Altera recommends that you run simulations over all process corners.

To perform process, voltage, and temperature (PVT) simulations, manually modify the spice decks in a two step process:

1. Remove the double-counting compensation circuitry from the simulation file. This is required as the amount of double-counting is dependant upon how the Quartus II software calculates delays and is not based on which PVT corner is being simulated. By default, the Quartus II software provides timing numbers using the slow process corner.
2. Select the proper corner for the PVT simulation by setting the correct HSPICE temperature, changing the supply voltage sources, and loading the correct transistor models.

A more detailed description of HSPICE process corners can be found in the family-specific HSPICE model documentation.

#### Related Information

[Accessing HSPICE Simulation Kits](#) on page 6-17

### Hold Time Analysis

Altera recommends performing worst-case hold time analysis using the fast corner models, which use fast transistors, high voltage, and low temperature. This involves modifying the SPICE decks to select the correct temperature option, change the supply voltage sources, and load the correct fast transistor models. The values of these parameters are located in the header comment section of the corresponding simulation deck files.

For a truly worst-case analysis, combine the HSPICE Writer hold time analysis results with the Quartus II software fast timing model. This requires that you change the double-counting compensation circuitry in the simulations files to also simulate the fast process corners, as this is what the Quartus II software uses for the fast timing model.

**Note:** This method of hold time analysis is recommended only for globally synchronous buses. Do not apply this method of hold-time analysis to source synchronous buses. This is because the source synchronous clocking scheme is designed to cancel out some of the PVT timing effects. If this is not taken into account, the timing results will not be accurate. Proper source synchronous timing analysis is beyond the scope of this document.

### I/O Voltage Variations

Use each of the FPGA family datasheets to verify the recommended operating conditions for supply voltages. For current FPGA families, the maximum recommended voltage corresponds to the fast corner, while the minimum recommended voltage corresponds to the slow corner. These voltage recommendations are specified at the power pins of the FPGA and are not necessarily the same voltage that are seen by the I/O buffers due to package IR drops.

The automatically generated HSPICE simulation files model this IR effect pessimistically by including a 50-mV IR drop on the  $V_{CCPD}$  supply when a high drive strength standard is being used.



## Correlation Report

Correlation reports for the HSPICE I/O models are located in the family-specific HSPICE I/O buffer simulation kits.

### Related Information

[Accessing HSPICE Simulation Kits](#) on page 6-17

## Document Revision History

Table 6-2: Document Revision History

Date	Version	Changes
June 2014	14.0.0	Updated format.
December 2010	10.0.1	Template update.
July 2010	10.0.0	Updated device support.
November 2009	9.1.0	No change to content.
March 2009	9.0.0	<ul style="list-style-type: none"> <li>Was volume 3, chapter 12 in the 8.1.0 release.</li> <li>No change to content.</li> </ul>
November 2008	8.1.0	<ul style="list-style-type: none"> <li>Changed to 8-1/2 x 11 page size.</li> <li>Added information for Stratix III devices.</li> <li>Input signals for Cyclone III devices are supported.</li> </ul>
May 2008	8.0.0	<ul style="list-style-type: none"> <li>Updated “Introduction” on page 12–1.</li> <li>Updated Figure 12–1.</li> <li>Updated Figure 12–3.</li> <li>Updated Figure 12–13.</li> <li>Updated “Output File Generation” on page 12–6.</li> <li>Updated “Simulation with HSPICE Models” on page 12–17.</li> <li>Updated “Invoking HSPICE Writer from the Command Line” on page 12–22.</li> <li>Added “Sample Input for I/O HSPICE Simulation Deck” on page 12–29.</li> <li>Added “Sample Output for I/O HSPICE Simulation Deck” on page 12–33.</li> <li>Updated “Correlation Report” on page 12–41.</li> <li>Added hyperlinks to referenced documents and websites throughout the chapter.</li> <li>Made minor editorial updates.</li> </ul>

**Related Information**

[Quartus II Handbook Archive](#)

For previous versions of the Quartus II Handbook.

2014.06.30

QII5V2



Subscribe



Send Feedback

You can integrate the Mentor Graphics® I/O Designer or DxDesigner PCB design tools into the Quartus® II design flow. This combination provides a complete FPGA-to-board design workflow.

With today's large, high-pin-count and high-speed FPGA devices, good and correct PCB design practices are essential to ensure correct system operation. The PCB design takes place concurrently with the design and programming of the FPGA. The FPGA or ASIC designer initially creates signal and pin assignments, and the board designer must correctly transfer these assignments to the symbols in their system circuit schematics and board layout. As the board design progresses, Altera recommends reassigning pins to optimize the PCB layout. Ensure that you inform the FPGA designer of the pin reassignments so that the new assignments are included in an updated placement and routing of the design.

The Mentor Graphics I/O Designer software allows you to take advantage of the full FPGA symbol design, creation, editing, and back-annotation flow supported by the Mentor Graphics tools.

This chapter covers the following topics:

- Mentor Graphics and Altera software integration flow
- Generating supporting files
- Adding Quartus II I/O assignments to I/O Designer
- Updating assignment changes between the I/O Designer the Quartus II software
- Generating I/O Designer symbols
- Creating DxDesigner symbols from the Quartus II output files

This chapter is intended for board design and layout engineers who want to start the FPGA board integration while the FPGA is still in the design phase. Alternatively, the board designer can plan the FPGA pin-out and routing requirements in the Mentor Graphics tools and pass the information back to the Quartus II software for placement and routing. Part librarians can also benefit from this chapter by learning how to use output from the Quartus II software to create new library parts and symbols.

The procedures in this chapter require the following software:

- The Quartus II software version 5.1 or later
- DxDesigner software version 2004 or later
- Mentor Graphics I/O Designer software (optional)

**Note:** To obtain and license the Mentor Graphics tools and for product information, support, and training, refer to the Mentor Graphics website.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

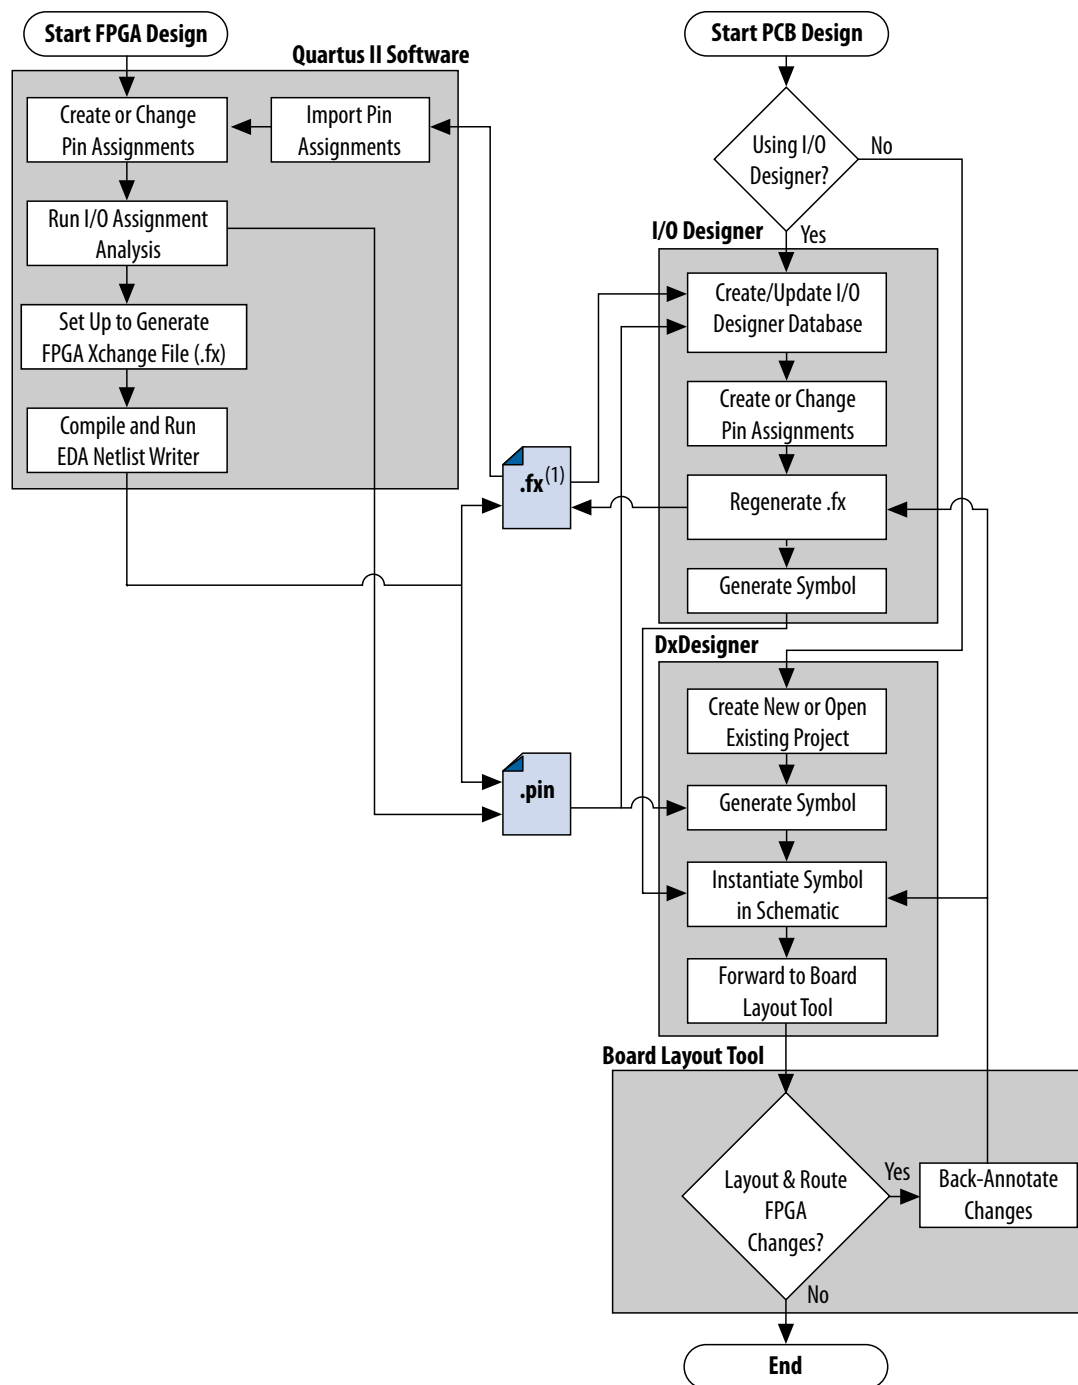
ISO  
9001:2008  
Registered



## FPGA-to-PCB Design Flow

You can create a design flow integrating an Altera® FPGA design from the Quartus II software, and a circuit schematic in the DxDesigner software.

Figure 7-1: Design Flow with and Without the I/O Designer Software



**Note:** The Quartus II software generates the **.fx** in the output directory you specify in the **Board-Level** page of the **Settings** dialog box. However, the Quartus II software and the I/O Designer software can import pin assignments from an **.fx** located in any directory. Use a backup **.fx** to prevent overwriting existing assignments or importing invalid assignments.

To integrate the I/O Designer into your design flow, follow these steps:

1. In the Quartus II software, click **Assignments > Settings > EDA Tool Settings > Board-Level** to specify settings for **.fx** symbol file generation.
2. Compile your design to generate the **.fx** and Pin-Out File (**.pin**) in the Quartus II project directory.
3. Create a board design with the DxDesigner software and the I/O Designer software by performing the following steps:
  - a. Create a new I/O Designer database based on the **.fx** and the **.pin** files.
  - b. In the I/O Designer software, make adjustments to signal and pin assignments.
  - c. Regenerate the **.fx** in the I/O Designer software to export the I/O Designer software changes to the Quartus II software.
  - d. Generate a single or fractured symbol for use in the DxDesigner software.
  - e. Add the symbol to the **sym** directory of a DxDesigner project, or specify a new DxDesigner project with the new symbol.
  - f. Instantiate the symbol in your DxDesigner schematic and export the design to the board layout tool.
  - g. Back-annotate pin changes created in the board layout tool to the DxDesigner software and back to the I/O Designer software and the Quartus II software.
4. Create a board design with the DxDesigner software without the I/O Designer software by performing the following steps:
  - a. Create a new DxBoardLink symbol with the **Symbol** wizard and reference the **.pin** from the Quartus II software in an existing DxDesigner project.
  - b. Instantiate the symbol in your DxDesigner schematic and export the design to a board layout tool.


**Note:** You can update these symbols with design changes with or without the I/O Designer software. If you use the Mentor Graphics I/O Designer software and you change symbols with the DxDesigner software, you must reimport the symbols into I/O Designer to avoid overwriting your symbol changes.

## Integrating with I/O Designer

You can integrate the Mentor Graphics I/O Designer software into the Quartus II design flow. Pin and signal assignment changes can be made anywhere in the design flow with either the Quartus II Pin Planner or the I/O Designer software. The I/O Designer software facilitates moving these changes, as well as synthesis, placement, and routing changes, between the Quartus II software, an external synthesis tool (if used), and a schematic capture tool such as the DxDesigner software.

This section describes how to use the I/O Designer software to transfer pin and signal assignment information to and from the Quartus II software with an **.fx**, and how to create symbols for the DxDesigner software.

Figure 7-2: I/O Designer Design Flow



**Note:** (2) DxDesigner software-specific steps in the design flow are not part of the I/O Designer flow.

## Generating Pin Assignment Files

You transfer I/O pin assignments from the Quartus II software to the Mentor Graphics PCB tools by generating optional **.pin** and **.fx** files during Quartus II compilation. These files contain pin assignment information for use in other tools. Click **Assignments > Settings > Board-Level** to specify settings for optional PCB tool file generation. Click **Processing > Start Compilation** to compile the design to generate the file(s) in the project directory.

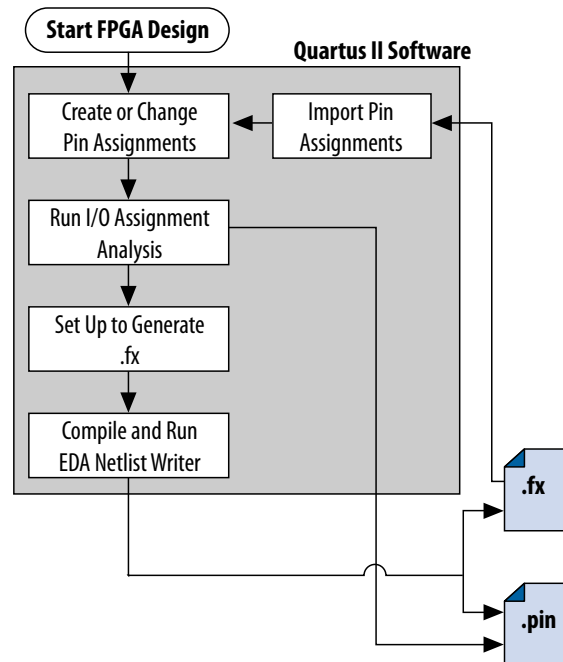
The Quartus II-generated **.pin** contains the I/O pin name, number, location, direction, and I/O standard for all used and unused pins in the design. Click **Assignments > Pin Planner** to modify I/O pin

assignments. You cannot import pin assignment changes from a Mentor Graphics **.pin** into the Quartus II software.

The **.fx** is an input or output of either the Quartus II or I/O Designer software. You can generate an **.fx** in the Quartus II software for symbol generation in the Mentor Graphics I/O Designer software. A Quartus II **.fx** contains the pin name, number, location, direction, I/O standard, drive strength, termination, slew rate, IOB delay, and differential pins. An I/O Designer **.fx** additionally includes information about unused pins and pin set groups.

The I/O Designer software can also read from or update a Quartus II Settings File (**.qsf**). You can use the **.qsf** in the same way as use of the **.fx**, but pin swap group information does not transfer between I/O Designer and the Quartus II software. Use the **.fx** rather than the **.qsf** for transferring I/O assignment information.

Figure 7-3: Generating **.pin** and **.fx** files



## I/O Designer Settings

You can directly export I/O Designer symbols to the DxDesigner software. To set options for integrating I/O Designer with Dx Designer, follow these steps:

1. Start the I/O Designer software.
2. Click **Tools > Preferences**.
3. Click **Paths**, and then double-click the **DxDesigner executable file path** field to select the location of the DxDesigner application.
4. Click **Apply**.
5. Click **Symbol Editor**, and then click **Export**. In the Export type menu, under **General**, select **DxDesigner/PADS-Designer**.
6. Click **Apply**, and then click **OK**.

7. Click **File > Properties**.
8. Click the **PCB Flow** tab, and then click **Path to a DxDesigner project directory**.
9. Click **OK**.

If you do not have a new DxDesigner project in the Database wizard and a DxDesigner project, you must create a new database with the DxDesigner software, and then specify the project location in I/O Designer.

## Transferring I/O Assignments

You can transfer Quartus II signal and pin assignments contained in **.pin** and **.fx** files into an I/O Designer database. Use the I/O Designer Database Wizard to create a new database incorporating the **.fx** and **.pin** files. You can also create a new, empty database and manually add the assignment information. If there is no available signal or pin assignment information, you can create an empty database containing only a selection of the target device. This technique is useful if you know the signals in your design and the pins you want to assign. You can subsequently transfer this information to the Quartus II software for placement and routing.

### Before you begin

You may create a very simple I/O Designer database that includes only the **.pin** or **.fx** file information. However, when using only a **.pin**, you cannot import I/O assignment changes from I/O Designer back into the Quartus II software without also generating an **.fx**. If your I/O Designer database includes only **.fx** file information, the database may not contain all the available I/O assignment information. The Quartus II **.fx** file only lists assigned pins. The **.pin** lists all assigned and unassigned device pins. Use both the **.pin** and the **.fx** to produce the most complete set of I/O Designer database information.

To create a new I/O Designer database using the Database wizard, follow these steps:

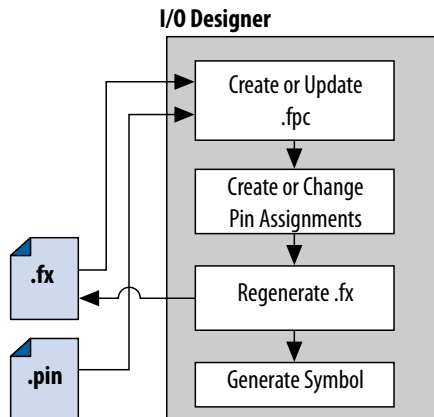
1. Start the I/O Designer software. The **Welcome to I/O Designer** dialog box appears. Select **Wizard to create new database** and click **OK**.  
If the **Welcome to I/O Designer** dialog box does not appear, you can access the wizard through the menu. To access the wizard, click **File > Database Wizard**.
2. Click **Next**. The **Define HDL source file** page appears.  
If no HDL files are available, or if the **.fx** contains your signal and pin assignments, you can skip Step 3 and proceed to Step 4.
3. If your design includes a Verilog HDL or VHDL file, you can add a top-level Verilog HDL or VHDL file in the I/O Designer software. Adding a file allows you to create functional blocks or get signal names from your design. You must create all physical pin assignments in I/O Designer if you are not using an **.fx** or a **.pin**. Click **Next**. The **Database Name** page appears.
4. In the **Database Name** page, type your database file name. Click **Next**. The Database Location window appears.
5. Add a path to the new or an existing database in the **Location** field, or browse to a database location. Click **Next**. The **FPGA flow** page appears.
6. In the Vendor menu, click **Altera**.
7. In the Tool/Library menu, click **Quartus II <version>** to select your version of the Quartus II software.  
**Note:** The Quartus II software version listed may not match your actual software version. If your version is not listed, select the latest version. If your target device is not available, the device may not yet be supported by the I/O Designer software.
8. Select the appropriate device family, device, package, and speed (if applicable), from the corresponding menus. Click **Next**. The **Place and route** page appears.



9. In the **FPGAX file name** field, type or browse to the backup copy of the **.fx** generated by the Quartus II software.
10. In the **Pin report file name** field, type or browse to the **.pin** generated by the Quartus II software. Click **Next**.  
You can also select a **.qsf** for update. The I/O Designer software can update the pin assignment information in the **.qsf** without affecting any other information in the file.  
  
**Note:** You can import a **.pin** without importing an **.fx**. The I/O Designer software does not generate a **.pin**. To transfer assignment information to the Quartus II software, select an additional file and file type. Altera recommends selecting an **.fx** in addition to a **.pin** for transferring all the assignment information in the **.fx** and **.pin** files. In some versions of the I/O Designer software, the standard file picker may incorrectly look for a **.pin** instead of an **.fx**. In this case, select **All Files (\*.\*)** from the **Save as type** list and select the file from the list.
11. On the **Synthesis** page, specify an external synthesis tool and a synthesis constraints file for use with the tool. If you do not use an external synthesis tool, click **Next**.
12. On the **PCB Flow** page, you can select an existing schematic project or create a new project as a symbol information destination.
  - To select an existing project, select **Choose existing project** and click **Browse** after the **Project Path** field. The **Select project** dialog box appears. Select the project.
  - To create a new project, in the **Select project** dialog box, select **Create new empty project**. Type the project file name in the **Name** field and browse to the location where you want to save the file. Click **OK**.
13. If you have not specified a design tool to which you can send symbol information in the I/O Designer software, click **Advanced** in the **PCB Flow** page and select your design tool. If you select the DxDesigner software, you have the option to specify a Hierarchical Occurrence Attributes (**.oat**) file to import into the I/O Designer software. Click **Next** and then click **Finish** to create the database. Updating

## Updating I/O Designer with Quartus II Pin Assignments

As you fine tune your design in the Quartus II software, changes to design logic and pin assignments are likely. You must transfer any pin assignment changes made during design iterations for correct analysis in your circuit schematic and board layout tools. You transfer Quartus II pin assignment changes to I/O Designer by updating the **.fx** and the **.pin** files in the Quartus II software. When you update the **.fx** or the **.pin**, the I/O Designer database imports the changes automatically when configured according to the following instructions.

**Before you begin****Figure 7-4: Updating Quartus II Pin Assignments in I/O Designer**

To update the **.fx** in your selected output directory and the **.pin** in your project directory after making changes to the design, perform the following tasks:

1. In the I/O Designer software, click **File > Properties**.
2. Under **FPGA Xchange**, specify the **.fx** file name and location.
3. Under **Place and Route**, specify the **.pin** file name and location.  
After you have set up these file locations, the I/O Designer software monitors these files for changes. If the specified **.fx** or **.pin** is modified during design processing, three indicators flash red in the lower right corner of the I/O Designer GUI. You can click the indicators to open the **I/O Designer Update Wizard** dialog box. The **I/O Designer Update Wizard** dialog box lists the updated files in the database.
4. Make logic or pin assignment changes in your design.
5. Click **Processing > Start > Start I/O Assignment Analysis** to validate your latest assignment changes.
6. To preserve your changes and update the corresponding the **.fx** and **.pin** files, click **Processing > Start > Start EDA Netlist Writer** or **Processing > Start Compilation**.

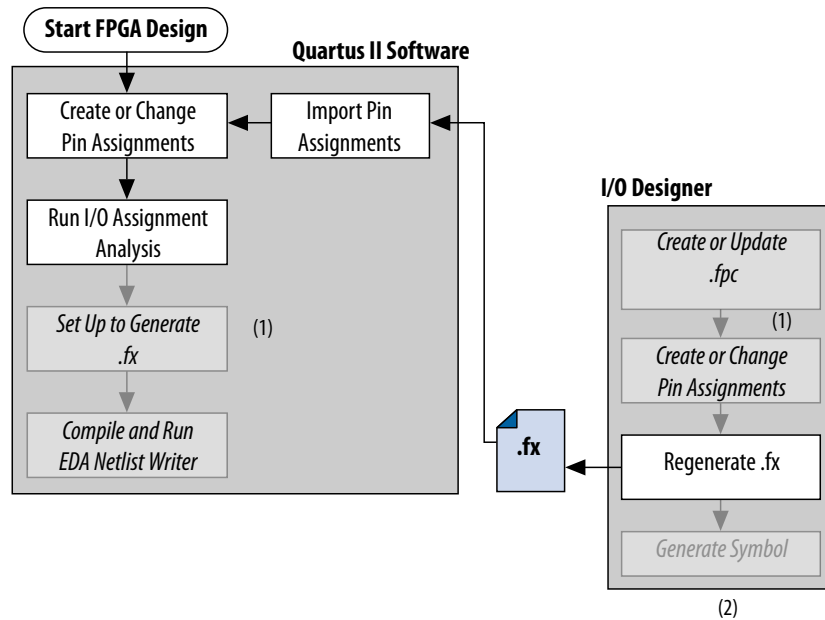
**Note:** Your I/O Designer database should use a backup copy of the **.fx** generated by the Quartus II software. Otherwise, updating the file in the Quartus II software overwrites any changes made to the file by the I/O Designer software. If there are I/O Designer assignments in the **.fx** that you want to preserve, create a backup copy of the file before updating it in the Quartus II software, and verify that your I/O Designer database points to the backup copy.

## Updating Quartus II with I/O Designer Pin Assignments

As you fine tune your board design in I/O Designer, changes to signal routing and layout are likely. You must import any routing and layout changes into the Quartus II software for accurate place and route to match the new pin-out. The I/O Designer tool supports this flow.

## Before you begin

Figure 7-5: Importing I/O Designer Pin Assignments



To import I/O Designer pin assignments, follow these steps:

1. Make pin assignment changes directly in the I/O Designer software, or the software can automatically update changes made in a board layout tool that are back-annotated to a schematic entry program such as the DxDesigner software.
2. To update the .fx with the changes, click **Generate > FPGA Xchange File**.
3. Open your Quartus II project.
4. Click **Assignments > Import Assignments**.
5. (Optional) To preserve original assignments before import, turn on **Copy existing assignments into <project name>.qsf.bak** before importing before importing the .fx.
6. Select the .fx and click **Open**.
7. Click **OK**.

## Generating Schematic Symbols in I/O Designer

Circuit board schematic creation is one of the first tasks required in the design of a new PCB. You can use the I/O Designer software to generate schematic symbols for your Quartus II FPGA design for use in the DXDesigner schematic entry tools. The I/O Designer software can generate symbols for use in various Mentor Graphics schematic entry tools, and can import changes back-annotated by board layout tools to update the database and update the Quartus II software with the .fx

Most FPGA devices contain hundreds of pins, requiring large schematic symbols that may not fit on a single schematic page. Symbol designs in the I/O Designer software can be split or fractured into various functional blocks, allowing multiple part fractures on the same schematic page or across multiple pages. In the DxDesigner software, these part fractures join together with the use of the HETERO attribute.

You can use the I/O Designer **Symbol** wizard to quickly create symbols that you can subsequently refine. Alternatively, you can import symbols from another DXDesigner project, and then assign an FPGA to the symbol. To import symbols in the I/O Designer software, **File > Import Symbol**.

I/O Designer symbols are either functional, physical (PCB), or both. Signals imported into the database, usually from Verilog HDL or VHDL files, are the basis of a functional symbol. No physical device pins must be associated with the signals to generate a functional symbol. This section focuses on board-level PCB symbols with signals directly mapped to physical device pins through assignments in either the Quartus II Pin Planner or in the I/O Designer database.

## Generating Schematic Symbols

To create a symbol based on a selected Altera FPGA device, follow these steps:

1. Start the I/O Designer software.
2. Click **Symbol > Symbol Wizard**.
3. In the **Symbol name** field, type the symbol name. The **DEVICE** and **PKG\_TYPE** fields display the device and package information.

**Note:** If **DEVICE** and **PKG\_TYPE** are blank or incorrect, close the Symbol wizard and specify the correct device information (**File > Properties > FPGA Flow**).

4. Under **Symbol type**, click **PCB**. Under **Use signals**, click **All**, then click **Next**.
5. Select fracturing options for your symbol. If you are using the Symbol wizard to edit a previously created fractured symbol, you must turn on **Reuse existing fractures** to preserve your current fractures. Select other options on this page as appropriate for your symbol. Click **Next**.
6. Select additional fracturing options for your symbol. Click **Next**.
7. Select the options for the appearance of the symbols. Click **Next**.
8. Define the information you want to label for the entire symbol and for individual pins. Click **Next**.
9. Add any additional signals and pins to the symbol. Click **Finish**.

You can view your symbol and any fractures you created with the Symbol Editor. You can edit parts of the symbol, delete fractures, or rerun the Symbol wizard. When you modify pin assignments in I/O Designer database, I/O Designer symbols automatically reflect these changes. Modify assignments in the I/O Designer software by supplying and updated **.fx** from the Quartus II software, or by back-annotating changes in your board layout tool.

## Exporting Schematic Symbols to DxDesigner

You can export your I/O Designer schematic symbols for to DxDesigner for further design entry work. To generate all fractures of a symbol, click **Generate > All Symbols**. To generate only the currently displayed symbol, click **Generate > Current Symbol Only**. The DxDesigner project **/sym** directory preserves each symbol in the database as a separate file. You can instantiate the symbols in your DxDesigner schematics.


## Integrating with DxDesigner

You can integrate the Mentor Graphics DxDesigner schematic capture tool into the Quartus II design flow. Use DxDesigner to create flat circuit schematics or to create hierarchical schematics that facilitate design reuse and a team-based design for all PCB types. Use DxDesigner in conjunction with I/O Designer software for a complete FPGA I/O and PCB design flow.

If you use DxDesigner without the I/O Designer software, the design flow is one-way, using only the **.pin** generated by the Quartus II software. You can only make signal and pin assignment changes in the

Quartus II software. You cannot back-annotate changes made in a board layout tool or in a DxDesigner symbol to the Quartus II software.

Figure 7-6: DxDesigner-only Flow (without I/O Designer)



## DxDesigner Project Settings

DxDesigner new projects automatically create FPGA symbols by default. However, if you are using the I/O Designer with DxDesigner, you must enable DxBoardLink Flow options for integration with the I/O Designer software. To enable the DxBoardLink flow design configuration when creating a new DxDesigner project, follow these steps:

1. Start the DxDesigner software.
2. Click **File > New**, and then click the **Project** tab.
3. Click **More**. Turn on **DxBoardLink**. To enable the DxBoardLink Flow design configuration for an existing project, click **Design Configurations** in the Design Configuration toolbar and turn on **DxBoardLink**.

## Creating Schematic Symbols in DxDesigner

You can create schematic symbols in the DxDesigner software manually or with the Symbol wizard. The DxDesigner Symbol wizard is similar to the I/O Designer Symbol wizard, but with fewer fracturing options. The DxDesigner Symbol wizard creates, fractures, and edits FPGA symbols based on the specified Altera device. To create a symbol with the Symbol wizard, follow these steps;

1. Start the DxDesigner software.
2. Click **Symbol Wizard** in the toolbar.
3. Type the new symbol name in the name field and click **OK**.
4. Specify creation of a new symbol or modification of an existing symbol. To modify an existing symbol, specify the library path or alias, and select the existing symbol. To create a new symbol, select DxBoardLink for the symbol source. The DxDesigner block type defaults to Module because the FPGA design does not have an underlying DxDesigner schematic. Choose whether or not to fracture the symbol. Click **Next**.
5. Type a name for the symbol, an overall part name for all the symbol fractures, and a library name for the new library created for this symbol. By default, the part and library names are the same as the symbol name. Click **Next**.

6. Specify the appearance of the generated symbol and how itthe grid you have set in your DxDesigner project schematic. After making your selections. Click **Next**.
7. In the **FPGA vendor** list, select **Altera Quartus**. In the **Pin-Out file to import** field, select the **.pin** from your Quartus II project directory. You can also specify Fracturing Scheme, Bus pin, and Power pin options. Click **Next**.
8. Select to create or modify symbol attributes for use in the DxDesigner software. Click **Next**.
9. On the **Pin Settings** page, make any final adjustments to pin and label location and information. Each tabbed spreadsheet represents a fracture of your symbol. Click **Save Symbol**.  
After creating the symbol, you can examine and place any fracture of the symbol in your schematic. You can locate separate files of all the fractures you created in the library you specified or created in the **/sym** directory in your DxDesigner project. You can add the symbols to your schematics or you can manually edit the symbols or with the Symbol wizard.

## Analyzing FPGA Simultaneous Switching Noise (SSN)

With the Quartus II software, you can extract pin assignment data and perform SSN analysis of your design. Perform SSN analysis early in the board layout stage as part of your overall pin planning process. Use the Quartus II SSN Analyzer to optimize the pin assignments for better SSN performance.

## Scripting API

The I/O Designer software includes a command line Tcl interpreter. All commands input through the I/O Designer GUI translate into Tcl commands run by the tool. You can run individual Tcl commands or scripts in the I/O Designer Console window, rather than using the GUI.

You can use the following Tcl commands to control I/O Designer.

- `set_fpga_xchange_file <file name>`—specifies the **.fx** from which the I/O Designer software updates assignments.
- `update_from_fpga_xchange_file`—updates the I/O Designer database with assignment updates from the currently specified **.fx**.
- `generate_fpga_xchange_file`—updates the **.fx** with I/O Designer software changes for transfer back into the Quartus II software.
- `set_pin_report_file -quartus_pin <file name>`—imports assignment data from a Quartus II software **.pin** file.
- `symbolwizard`—runs the I/O Designer Symbol wizard.
- `set_dx_designer_project -path <path>`

## Document Revision History

**Table 7-1: Document Revision History**

Date	Version	Changes
2014.06.30	14.0.0	<ul style="list-style-type: none"> <li>• Replaced MegaWizard Plug-In Manager information with IP Catalog.</li> <li>• Added standard information about upgrading IP cores.</li> <li>• Added standard installation and licensing information.</li> <li>• Removed outdated device support level information. IP core device support is now available in IP Catalog and parameter editor.</li> </ul>
June 2012	12.0.0	<ul style="list-style-type: none"> <li>• Removed survey link.</li> </ul>
December 2010	10.1.0	<ul style="list-style-type: none"> <li>• Changed to new document template.</li> </ul>

**Related Information**

[Quartus II Handbook Archive](#)

2014.06.30

QI15V2



Subscribe



Send Feedback

## Cadence PCB Design Tools Support

The Quartus<sup>®</sup> II software interacts with the following software to provide a complete FPGA-to-board integration design workflow: the Cadence Allegro Design Entry HDL software and the Cadence Allegro Design Entry CIS (Component Information System) software (also known as OrCAD Capture CIS). The information is useful for board design and layout engineers who want to begin the FPGA board integration process while the FPGA is still in the design phase. Part librarians can also benefit by learning the method to use output from the Quartus II software to create new library parts and symbols.

With today's large, high-pin-count and high-speed FPGA devices, good PCB design practices are important to ensure the correct operation of your system. The PCB design takes place concurrently with the design and programming of the FPGA. An FPGA or ASIC designer initially creates the signal and pin assignments and the board designer must transfer these assignments to the symbols used in their system circuit schematics and board layout correctly. As the board design progresses, you must perform pin reassignments to optimize the layout. You must communicate pin reassignments to the FPGA designer to ensure the new assignments are processed through the FPGA with updated placement and routing.

You require the following software:

- The Quartus II software version 5.1 or later
- The Cadence Allegro Design Entry HDL software or the Cadence Allegro Design Entry CIS software version 15.2 or later
- The OrCAD Capture software with the optional CIS option version 10.3 or later (optional)

**Note:** These programs are very similar because the Cadence Allegro Design Entry CIS software is based on the OrCAD Capture software. Any procedural information can also apply to the OrCAD Capture software unless otherwise noted.

### Related Information

- [www.cadence.com](http://www.cadence.com)  
For more information about obtaining and licensing the Cadence tools and for product information, support, and training
- [www.cadence.com](http://www.cadence.com)  
For more information about the OrCAD Capture software and the CIS option
- [www.ema-eda.com](http://www.ema-eda.com)  
For more information about Cadence and OrCAD support and training

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



## Product Comparison

Table 8-1: Cadence and OrCAD Product Comparison

Description	Cadence Allegro Design Entry HDL	Cadence Allegro Design Entry CIS	OrCAD Capture CIS
Former Name	Concept HDL Expert	Capture CIS Studio	—
History	More commonly known by its former name, Cadence renamed all board design tools in 2004 under the Allegro name.	Based directly on OrCAD Capture CIS, the Cadence Allegro Design Entry CIS software is still developed by OrCAD but sold and marketed by Cadence. EMA provides support and training.	The basis for Design Entry CIS is still developed by OrCAD for continued use by existing OrCAD customers. EMA provides support and training for all OrCAD products.
Vendor Design Flow	Cadence Allegro 600 series, formerly known as the Expert Series, for high-end, high-speed design.	Cadence Allegro 200 series, formerly known as the Studio Series, for small- to medium-level design.	—

### Related Information

- [www.cadence.com](http://www.cadence.com)
- [www.ema-eda.com](http://www.ema-eda.com)

## FPGA-to-PCB Design Flow

You can create a design flow integrating an Altera FPGA design from the Quartus II software through a circuit schematic in the Cadence Allegro Design Entry HDL software or the Cadence Allegro Design Entry CIS software.

Figure 8-1: Design Flow with the Cadence Allegro Design Entry HDL Software

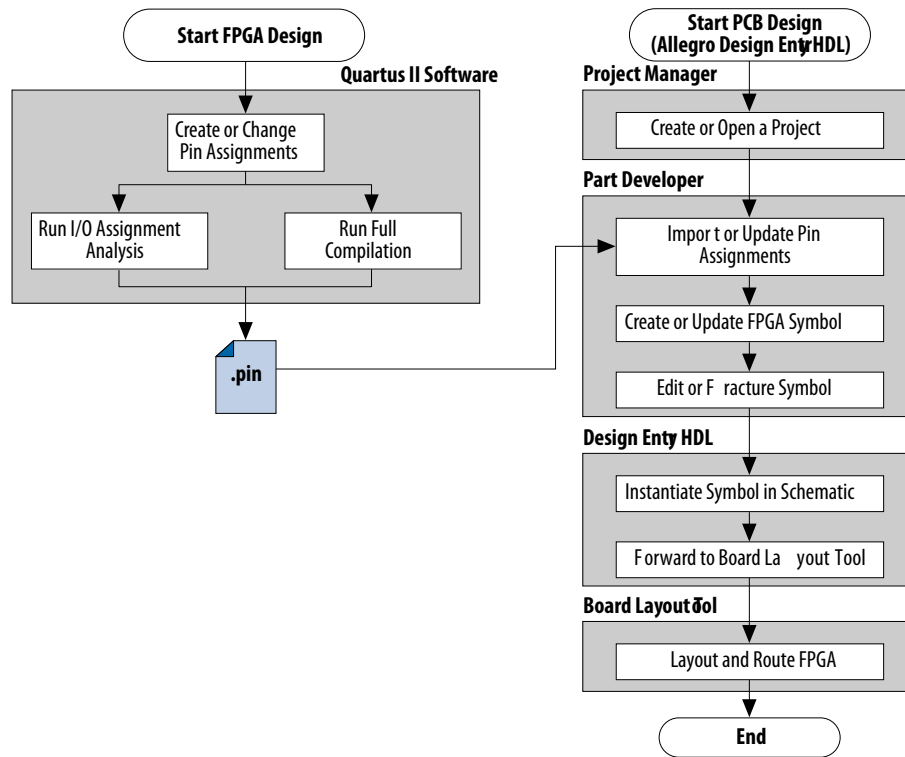
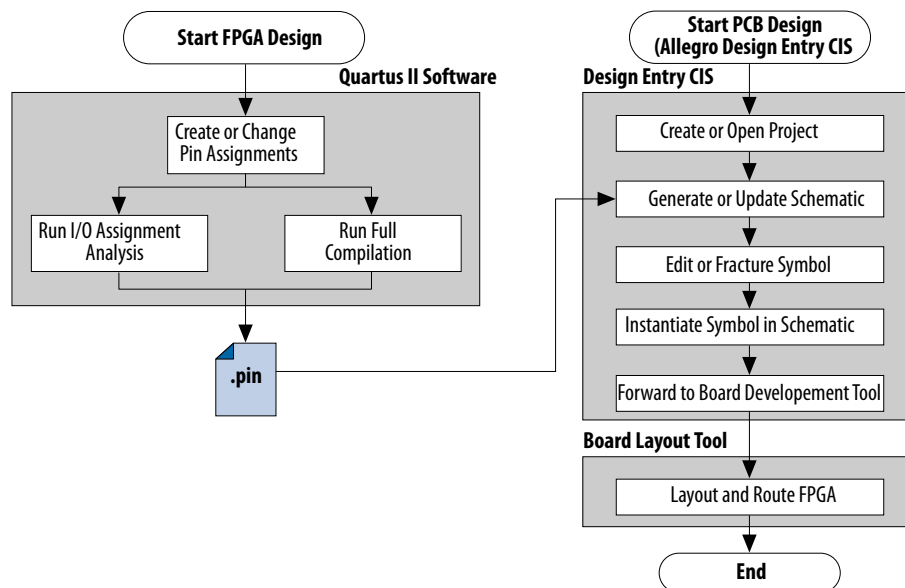


Figure 8-2: Design Flow with the Cadence Allegro Design Entry CIS Software



To create FPGA symbols using the Cadence Allegro PCB Librarian Part Developer tool, you must obtain the Cadence PCB Librarian Expert license. You can update symbols with changes made to the FPGA design using any of these tools.

## Integrating Altera FPGA Design

To integrate an Altera FPGA design starting in the Quartus II software through to a circuit schematic in the Cadence Allegro Design Entry HDL software or the Cadence Allegro Design Entry CIS software, follow these steps:

1. In the Quartus II software, compile your design to generate a Pin-Out File (**.pin**) to transfer the assignments to the Cadence software.
2. If you are using the Cadence Allegro Design Entry HDL software for your schematic design, follow these steps:
  - a. Open an existing project or create a new project in the Cadence Allegro Project Manager tool.
  - b. Construct a new symbol or update an existing symbol using the Cadence Allegro PCB Librarian Part Developer tool.
  - c. With the Cadence Allegro PCB Librarian Part Developer tool, edit your symbol or fracture it into smaller parts (optional).
  - d. Instantiate the symbol in your Cadence Allegro Design Entry HDL software schematic and transfer the design to your board layout tool.

or

If you are using the Cadence Allegro Design Entry CIS software for your schematic design, follow these steps:

- e. Generate a new part in a new or existing Cadence Allegro Design Entry CIS project, referencing the **.pin** output file from the Quartus II software. You can also update an existing symbol with a new **.pin**.
- f. Split the symbol into smaller parts as necessary.
- g. Instantiate the symbol in your Cadence Allegro Design Entry CIS schematic and transfer the design to your board layout tool.

## Performing Simultaneous Switching Noise (SSN) Analysis of Your FPGA

With the Quartus II software, you can extract pin assignment data and perform SSN analysis of your FPGA design for designs targeting the Stratix III device family. You can analyze SSN in your device early in the board layout stage as part of your overall pin planning process; however, you do not have to perform SSN analysis to generate pin assignment data from the Quartus II software. You can use the SSN Analyzer tool to optimize the pin assignments for better SSN performance of your device.

### Related Information

[About the SSN Analyzer](#)

[Simultaneous Switching Noise \(SSN\) Analysis and Optimizations](#) on page 5-1

## Setting Up the Quartus II Software

You can transfer pin and signal assignments from the Quartus II software to the Cadence design tools by generating the Quartus II project **.pin**. The **.pin** is an output file generated by the Quartus II Fitter containing pin assignment information. You can use the Quartus II Pin Planner to set and change the assignments in the **.pin** and then transfer the assignments to the Cadence design tools. You cannot, however, import pin assignment changes from the Cadence design tools into the Quartus II software with the **.pin**.

The **.pin** lists all used and unused pins on your selected Altera device. The **.pin** also provides the following basic information fields for each assigned pin on the device:

- Pin signal name and usage
- Pin number
- Signal direction
- I/O standard
- Voltage
- I/O bank
- User or Fitter-assigned

#### Related Information

- [I/O Management](#) on page 4-1  
For more information about using the Quartus II Pin Planner to create or change pin assignment details.

## Generating a .pin File

To generate a **.pin**, follow these steps:

1. Compile or perform I/O assignment analysis on your design.
2. Locate the **.pin** in your Quartus II project directory with the name <project name>.**.pin**.

#### Related Information

- [I/O Management](#) on page 4-1  
For more information about pin and signal assignment transfer and the files that the Quartus II software can import and export.

## FPGA-to-Board Integration with the Cadence Allegro Design Entry HDL Software

The Cadence Allegro Design Entry HDL software is a schematic capture tool and is part of the Cadence 600 series design flow. Use the Cadence Allegro Design Entry HDL software to create flat circuit schematics for all types of PCB design. The Cadence Allegro Design Entry HDL software can also create hierarchical schematics to facilitate design reuse and team-based design. With the Cadence Allegro Design Entry HDL software, the design flow from FPGA-to-board is one-way, using only the **.pin** generated by the Quartus II software. You can only make signal and pin assignment changes in the Quartus II software and these changes reflect as updated symbols in a Cadence Allegro Design Entry HDL project.

For more information about the design flow with the Cadence Allegro Design Entry HDL software, refer to [Figure 8-1](#).

**Note:** Routing or pin assignment changes made in a board layout tool or a Cadence Allegro Design Entry HDL software symbol cannot be back-annotated to the Quartus II software.

**Related Information**[www.cadence.com](http://www.cadence.com)

Provides information about the Cadence Allegro Design Entry HDL software and the Cadence Allegro PCB Librarian Part Developer tool, including licensing, support, usage, training, and product updates.

## Creating Symbols

In addition to circuit simulation, circuit board schematic creation is one of the first tasks required when designing a new PCB. Schematics must understand how the PCB works, and to generate a netlist for a board layout tool for board design and routing. The Cadence Allegro PCB Librarian Part Developer tool allows you to create schematic symbols based on FPGA designs exported from the Quartus II software.

You can create symbols for the Cadence Allegro Design Entry HDL project with the Cadence Allegro PCB Librarian Part Developer tool, which is available in the Cadence Allegro Project Manager tool. Altera recommends using the Cadence Allegro PCB Librarian Part Developer tool to import FPGA designs into the Cadence Allegro Design Entry HDL software.

You must obtain a PCB Librarian Expert license from Cadence to run the Cadence Allegro PCB Librarian Part Developer tool. The Cadence Allegro PCB Librarian Part Developer tool provides a GUI with many options for creating, editing, fracturing, and updating symbols. If you do not use the Cadence Allegro PCB Librarian Part Developer tool, you must create and edit symbols manually in the Symbol Schematic View in the Cadence Allegro Design Entry HDL software.

**Note:** If you do not have a PCB Librarian Expert license, you can automatically create FPGA symbols using the programmable IC (PIC) design flow found in the Cadence Allegro Project Manager tool.

Before creating a symbol from an FPGA design, you must open a Cadence Allegro Design Entry HDL project with the Cadence Allegro Project Manager tool. If you do not have an existing Cadence Allegro Design Entry HDL project, you can create one with the Cadence Allegro Design Entry HDL software. The Cadence Allegro Design Entry HDL project directory with the name <project name>.cpm contains your Cadence Allegro Design Entry HDL projects.

While the Cadence Allegro PCB Librarian Part Developer tool refers to symbol fractures as slots, the other tools use different names to refer to symbol fractures.

**Table 8-2: Symbol Fracture Naming Conventions**

	Cadence Allegro PCB Librarian Part Developer Tool	Cadence Allegro Design Entry HDL Software	Cadence Allegro Design Entry CIS Software
During symbol generation	Slots	—	Sections
During symbol schematic instantiation	—	Versions	Parts

**Related Information**[www.cadence.com](http://www.cadence.com)

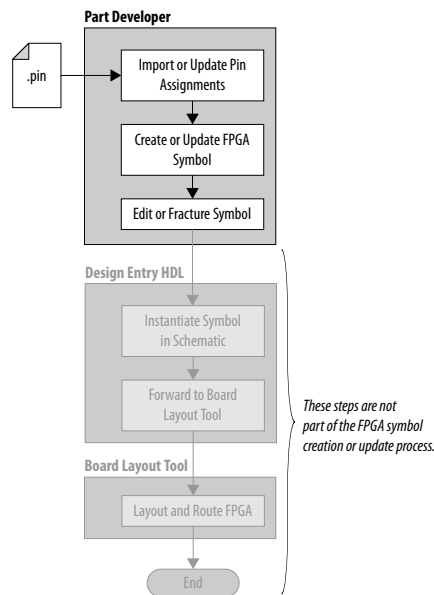
Provides information about using the PIC design flow.

## Cadence Allegro PCB Librarian Part Developer Tool

You can create, fracture, and edit schematic symbols for your designs using the Cadence Allegro PCB Librarian Part Developer tool. Symbols designed in the Cadence Allegro PCB Librarian Part Developer

tool can be split or fractured into several functional blocks called slots, allowing multiple smaller part fractures to exist on the same schematic page or across multiple pages.

## Cadence Allegro PCB Librarian Part Developer Tool in the Design Flow



To run the Cadence Allegro PCB Librarian Part Developer tool, you must open a Cadence Allegro Design Entry HDL project in the Cadence Allegro Project Manager tool. To open the Cadence Allegro PCB Librarian Part Developer tool, on the Flows menu, click **Library Management**, and then click **Part Developer**.

### Related Information

- [FPGA-to-PCB Design Flow](#) on page 8-2

### Import and Export Wizard

After starting the Cadence Allegro PCB Librarian Part Developer tool, use the **Import and Export** wizard to import your pin assignments from the Quartus II software.

**Note:** Altera recommends using your PCB Librarian Expert license file. To point to your PCB Librarian Expert license file, on the File menu, click **Change Product** and then select the correct product license.

To access the Import and Export wizard, follow these steps:

1. On the File menu, click **Import and Export**.
2. Select **Import ECO-FPGA**, and then click **Next**.
3. In the **Select Source** page of the **Import and Export** wizard, specify the following settings:

- a. In the **Vendor** list, select **Altera**.
  - b. In the **PnR Tool** list, select **quartusII**.
  - c. In the **PR File** box, browse to select the **.pin** in your Quartus II project directory.
  - d. Click **Simulation Options** to select simulation input files.
  - e. Click **Next**.
4. In the **Select Destination** dialog box, specify the following settings:
- a. Under **Select Component**, click **Generate Custom Component** to create a new component in a library,  
  
or  
  
Click **Use standard component** to base your symbol on an existing component.
- Note:** Altera recommends creating a new component if you previously created a generic component for an FPGA device. Generic components can cause some problems with your design. When you create a new component, you can place your pin and signal assignments from the Quartus II software on this component and reuse the component as a base when you have a new FPGA design.
- b. In the **Library** list, select an existing library. You can select from the cells in the selected library. Each cell represents all the symbol versions and part fractures for a particular part. In the **Cell** list, select the existing cell to use as a base for your part.
  - c. In the **Destination Library** list, select a destination library for the component. Click **Next**.
  - d. Review and edit the assignments you import into the Cadence Allegro PCB Librarian Part Developer tool based on the data in the **.pin** and then click **Finish**. The location of each pin is not included in the **Preview of Import Data** page of the **Import and Export** wizard, but input pins are on the left side of the created symbol, output pins on the right, power pins on the top, and ground pins on the bottom.


## Editing and Fracturing Symbol

After creating your new symbol in the Cadence Allegro PCB Librarian Part Developer tool, you can edit the symbol graphics, fracture the symbol into multiple slots, and add or change package or symbol properties.

The Part Developer Symbol Editor contains many graphical tools to edit the graphics of a particular symbol. To edit the symbol graphics, select the symbol in the cell hierarchy. The **Symbol Pins** tab appears. You can edit the preview graphic of the symbol in the **Symbol Pins** tab.

Fracturing a Cadence Allegro PCB Librarian Part Developer package into separate symbol slots is useful for FPGA designs. A single symbol for most FPGA packages might be too large for a single schematic page. Splitting the part into separate slots allows you to organize parts of the symbol by function, creating cleaner circuit schematics. For example, you can create one slot for an I/O symbol, a second slot for a JTAG symbol, and a third slot for a power/ground symbol.

**Figure 8-3: Splitting a Symbol into Multiple Slots**



- This diagram represents a Cyclone device with JTAG or passive serial (PS) mode configuration option settings. Symbols created for other devices or other configuration modes may have different sets of configuration pins, but can be fractured in a similar manner.  
- The power/ground slot shows only a representation of power and ground pins because the device contains a large number of power and ground pins.

To fracture a part into separate slots, or to modify the slot locations of pins on parts fractured in the Cadence Allegro PCB Librarian Part Developer tool, follow these steps:

1. Start the Cadence Allegro Design Project Manager.
2. On the Flows menu, click **Library Management**.
3. Click **Part Developer**.
4. Click the name of the package you want to change in the cell hierarchy.
5. Click **Functions/Slots**. If you are not creating new slots but want to change the slot location of some pins, proceed to Step 6. If you are creating new slots, click **Add**. A dialog box appears, allowing you to add extra symbol slots. Set the number of extra slots you want to add to the existing symbol, not the total number of desired slots for the part. Click **OK**.
6. Click **Distribute Pins**. Specify the slot location for each pin. Use the checkboxes in each column to move pins from one slot to another. Click **OK**.
7. After distributing the pins, click the **Package Pin** tab and click **Generate Symbol(s)**.
8. Select whether to create a new symbol or modify an existing symbol in each slot. Click **OK**.

The newly generated or modified slot symbols appear as separate symbols in the cell hierarchy. Each of these symbols can be edited individually.

**Caution:** The Cadence Allegro PCB Librarian Part Developer tool allows you to remap pin assignments in the **Package Pin** tab of the main Cadence Allegro PCB Librarian Part Developer window. If signals remap to different pins in the Cadence Allegro PCB Librarian Part Developer tool, the changes reflect only in regenerated symbols for use in your schematics. You cannot transfer pin assignment changes to the Quartus II software from the Cadence Allegro PCB Librarian Part Developer tool, which creates a potential mismatch of the schematic symbols and assignments in the FPGA design. If pin assignment changes are necessary, make the changes in the Quartus II Pin Planner instead of the Cadence Allegro PCB Librarian Part Developer tool, and update the symbol as described in the following sections.

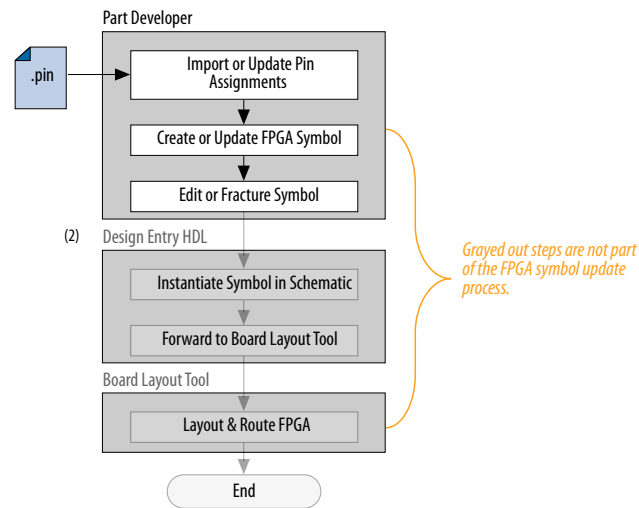
For more information about creating, editing, and organizing component symbols with the Cadence Allegro PCB Librarian Part Developer tool, refer to the Part Developer Help.



## Updating FPGA Symbols

As the design process continues, you must make logic changes in the Quartus II software, placing signals on different pins after recompiling the design, or use the Quartus II Pin Planner to make changes manually. The board designer can request such changes to improve the board routing and layout. To ensure signals connect to the correct pins on the FPGA, you must carry forward these types of changes to the circuit schematic and board layout tools. Updating the **.pin** in the Quartus II software facilitates this flow.

**Figure 8-4: Updating the FPGA Symbol in the Design Flow**



To update the symbol using the Cadence Allegro PCB Librarian Part Developer tool after updating the **.pin**, follow these steps:

1. On the File menu, click **Import and Export**. The Import and Export wizard appears.
2. In the list of actions to perform, select **Import ECO - FPGA**. Click **Next**. The **Select Source** dialog box appears.
3. Select the updated source of the FPGA assignment information. In the **Vendor** list, select **Altera**. In the **PnR Tool** list, select **quartusII**. In the **PR File** field, click **browse** to specify the updated **.pin** in your Quartus II project directory. Click **Next**. The Select Destination window appears.
4. Select the source component and a destination cell for the updated symbol. To create a new component based on the updated pin assignment data, select **Generate Custom Component**. Selecting **Generate Custom Component** replaces the cell listed under the **Specify Library and Cell** name header with a new, nonfractured cell. You can preserve these edits by selecting **Use standard component and select the existing library and cell**. Select the destination library for the component and click **Next**. The **Preview of Import Data** dialog box appears.
5. Make any additional changes to your symbol. Click **Next**. A list of ECO messages appears summarizing the changes made to the cell. To accept the changes and update the cell, click **Finish**.
6. The main Cadence Allegro PCB Librarian Part Developer window appears. You can edit, fracture, and generate the updated symbols as usual from the main Cadence Allegro PCB Librarian Part Developer window.

**Note:** If the Cadence Allegro PCB Librarian Part Developer tool is not set up to point to your PCB Librarian Expert license file, an error message appears in red at the bottom of the message text

window of the Part Developer when you select the **Import and Export** command. To point to your PCB Librarian Expert license, on the File menu, click **Change Product**, and select the correct product license.

#### Related Information

- [FPGA-to-PCB Design Flow](#) on page 8-2

## Instantiating the Symbol in the Cadence Allegro Design Entry HDL Software

To instantiate the symbol in your Cadence Allegro Design Entry HDL schematic after saving the new symbol in the Cadence Allegro PCB Librarian Part Developer tool, follow these steps:

1. In the Cadence Allegro Project Manager tool, switch to the board design flow.
2. On the Flows menu, click **Board Design**.
3. To start the Cadence Allegro Design Entry HDL software, click **Design Entry**.
4. To add the newly created symbol to your schematic, on the Component menu, click **Add**. The **Add Component** dialog box appears.
5. Select the new symbol library location, and select the name of the cell you created from the list of cells.

The symbol attaches to your cursor for placement in the schematic. To fracture the symbol into slots, right-click the symbol and choose **Version** to select one of the slots for placement in the schematic.

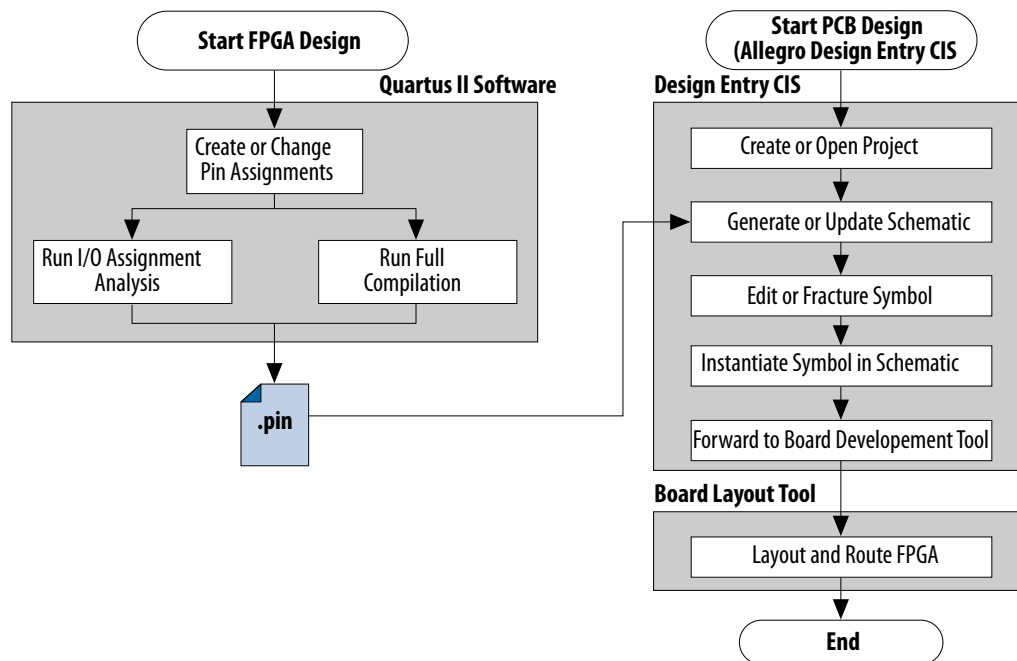
#### Related Information

- Provides more information about the Cadence Allegro Design Entry HDL software, including licensing, support, usage, training, and product updates.

## FPGA-to-Board Integration with Cadence Allegro Design Entry CIS Software

The Cadence Allegro Design Entry CIS software is a schematic capture tool (part of the Cadence 200 series design flow based on OrCAD Capture CIS). Use the Cadence Allegro Design Entry CIS software to create flat circuit schematics for all types of PCB design. You can also create hierarchical schematics to facilitate design reuse and team-based design using the Cadence Allegro Design Entry CIS software. With the Cadence Allegro Design Entry CIS software, the design flow from FPGA-to-board is unidirectional using only the **.pin** generated by the Quartus II software. You can only make signal and pin assignment changes in the Quartus II software. These changes reflect as updated symbols in a Cadence Allegro Design Entry CIS schematic project.

Figure 8-5: Design Flow with the Cadence Allegro Design Entry CIS Software



**Note:** Routing or pin assignment changes made in a board layout tool or a Cadence Allegro Design Entry CIS symbol cannot be back-annotated to the Quartus II software.

#### Related Information

- [www.cadence.com](http://www.cadence.com)  
For more information about the Cadence Allegro Design Entry CIS software, including licensing, support, usage, training, and product updates.
- [www.ema-eda.com](http://www.ema-eda.com)  
For more information about the Cadence Allegro Design Entry CIS software, including licensing, support, usage, training, and product updates.

## Creating a Cadence Allegro Design Entry CIS Project

The Cadence Allegro Design Entry CIS software has built-in support for creating schematic symbols using pin assignment information imported from the Quartus II software.

To create a new project in the Cadence Allegro Design Entry CIS software, follow these steps:

1. On the File menu, point to **New** and click **Project**. The New Project wizard starts.

When you create a new project, you can select the PC Board wizard, the Programmable Logic wizard, or a blank schematic.

2. Select the PC Board wizard to create a project where you can select which part libraries to use, or select a blank schematic.

The Programmable Logic wizard only builds an FPGA logic design in the Cadence Allegro Design Entry CIS software.

Your new project is in the specified location and consists of the following files:

- OrCAD Capture Project File (.opj)
- Schematic Design File (.dsn)

## Generating a Part

After you create a new project or open an existing project in the Cadence Allegro Design Entry CIS software, you can generate a new schematic symbol based on your Quartus II FPGA design. You can also update an existing symbol. The Cadence Allegro Design Entry CIS software stores component symbols in OrCAD Library File (.olb). When you place a symbol in a library attached to a project, it is immediately available for instantiation in the project schematic.

You can add symbols to an existing library or you can create a new library specifically for the symbols generated from your FPGA designs. To create a new library, follow these steps:

1. On the File menu, point to **New** and click **Library** in the Cadence Allegro Design Entry CIS software to create a default library named **library1.olb**. This library appears in the **Library** folder in the Project Manager window of the Cadence Allegro Design Entry CIS software.
2. To specify a desired name and location for the library, right-click the new library and select **Save As**. Saving the new library creates the library file.

## Generating Schematic Symbol

You can now create a new symbol to represent your FPGA design in your schematic.

To generate a schematic symbol, follow these steps:

1. Start the Cadence Allegro Design Entry CIS software.
2. On the Tools menu, click **Generate Part**. The **Generate Part** dialog box appears.
3. To specify the **.pin** from your Quartus II design, in the **Netlist/source file type** field, click **Browse**.
4. In the **Netlist/source file type** list, select **Altera Pin File**
5. Type the new part name.
6. Specify the **Destination part library** for the symbol. Failing to select an existing library for the part creates a new library with a default name that matches the name of your Cadence Allegro Design Entry CIS project.
7. To create a new symbol for this design, select **Create new part**. If you updated your **.pin** in the Quartus II software and want to transfer any assignment changes to an existing symbol, select **Update pins on existing part in library**.
8. Select any other desired options and set **Implementation type** to **<none>**. The symbol is for a primitive library part based only on the **.pin** and does not require special implementation. Click **OK**.
9. Review the Undo warning and click **Yes** to complete the symbol generation.

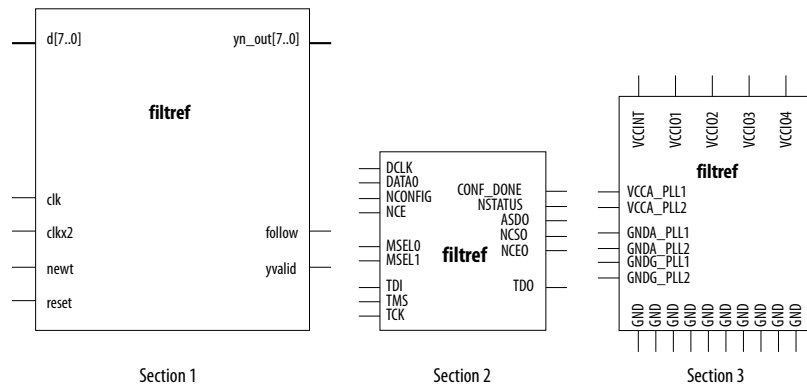
You can locate the generated symbol in the selected library or in a new library found in the **Outputs** folder of the design in the Project Manager window. Double-click the name of the new symbol to see its graphical representation and edit it manually using the tools available in the Cadence Allegro Design Entry CIS software.

**Note:** For more information about creating and editing symbols in the Cadence Allegro Design Entry CIS software, refer to the Help in the software.

## Splitting a Part

After saving a new symbol in a project library, you can fracture the symbol into multiple parts called sections. Fracturing a part into separate sections is useful for FPGA designs. A single symbol for most FPGA packages might be too large for a single schematic page. Splitting the part into separate sections allows you to organize parts of the symbol by function, creating cleaner circuit schematics. For example, you can create one slot for an I/O symbol, a second slot for a JTAG symbol, and a third slot for a power/ground symbol.

**Figure 8-6: Splitting a Symbol into Multiple Sections**



- This diagram represents a Cyclone device with JTAG or passive serial (PS) mode configuration option settings. Symbols created for other devices or other configuration modes might have different sets of configuration pins, but can be fractured in a similar manner.  
- The power/ground section shows only a representation of power and ground pins because the device contains a high number of power and ground pins.

**Note:** Although symbol generation in the Design Entry CIS software refers to symbol fractures as sections, other tools use different names to refer to symbol fractures.

To split a part into sections, select the part in its library in the Project Manager window of the Cadence Allegro Design Entry CIS software. On the Tools menu, click **Split Part** or right-click the part and choose **Split Part**. The **Split Part Section Input Spreadsheet** appears.

Figure 8-7: Split Part Section Input Spreadsheet

	Number	Name	Type	Order	Length	User Assign	I/O Bank	Voltage	I/O Standard	Location	Section
1	H1	clk	Input	0	Line		1			Left	1
2	G1	clkx2	Input	1	Line		1			Left	1
3	K13	CONF_DONE	Passive	2	Line		3			Left	2
4	C7	d[0]	Input	3	Line		2			Left	1
5	A6	d[1]	Input	4	Line		2			Left	1
6	D7	d[2]	Input	5	Line		2			Left	1
7	B7	d[3]	Input	6	Line		2			Left	1
8	B8	d[4]	Input	7	Line		2			Left	1
9	M7	d[5]	Input	8	Line		4			Left	1
10	A8	d[6]	Input	9	Line		2			Left	1
11	B6	d[7]	Input	10	Line		2			Left	1
12	H2	DATA0	Input	11	Line		1			Left	2
13	K4	DCLK	Bidirectional	12	Line		1			Left	2
14	C6	follow	Output	13	Line		2			Right	1
15	J3	MSEL0	Passive	14	Line		1			Left	2
16	J2	MSEL1	Passive	15	Line		1			Left	2
17	J4	nCE	Passive	16	Line		1			Left	2
18	H4	nCEO	Passive	17	Line		1			Left	2
19	H3	nCONFIG	Passive	18	Line		1			Left	2
20	H5	newt	Input	19	Line		1			Left	1
21	J13	nSTATUS	Passive	20	Line		3			Left	2
22	G16	reset	Input	21	Line		3			Left	1

Each row in the spreadsheet represents a pin in the symbol. The **Section** column indicates the section of the symbol to which each pin is assigned. You can locate all pins in a new symbol in section 1. You can change the values in the **Section** column to assign pins to various sections of the symbol. You can also specify the side of a section on the location of the pin by changing the values in the **Location** column. When you are ready, click **Split**. A new symbol appears in the same library as the original with the name <original part name>\_Split1.

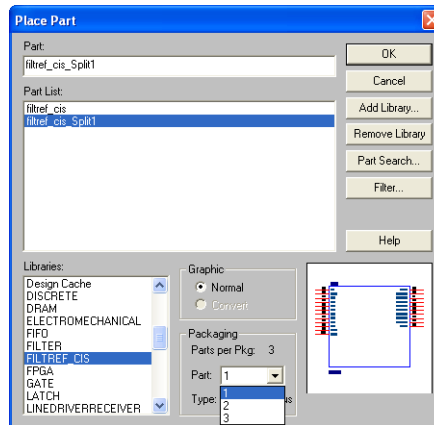
View and edit each section individually. To view the new sections of the part, double-click the part. The Part Symbol Editor window appears and the first section of the part displays for editing. On the View menu, click **Package** to view thumbnails of all the part sections. To edit the section of the symbol, double-click the thumbnail.

For more information about splitting parts into sections and editing symbol sections in the Cadence Allegro Design Entry CIS software, refer to the Help in the software.

## Instantiating a Symbol in a Design Entry CIS Schematic

After saving a new symbol in a library in your Cadence Allegro Design Entry CIS project, you can instantiate the new symbol on a page in your schematic. Open a schematic page in the Project Manager window of the Cadence Allegro Design Entry CIS software. To add the newly created symbol to your schematic on the schematic page, on the Place menu, click **Part**. The **Place Part** dialog box appears.

Figure 8-8: Place Part Dialog Box



Select the new symbol library location and the newly created part name. If you select a part that is split into sections, you can select the section to place from the **Part** pop-up menu. Click **OK**. The symbol attaches to your cursor for placement in the schematic. To place the symbol, click on the schematic page.

For more information about using the Cadence Allegro Design Entry CIS software, refer to the Help in the software.

## Altera Libraries for the Cadence Allegro Design Entry CIS Software

Altera provides downloadable **.olb** for many of its device packages. You can add these libraries to your Cadence Allegro Design Entry CIS project and update the symbols with the pin assignments contained in the **.pin** generated by the Quartus II software. You can use the downloaded library symbols as a base for creating custom schematic symbols with your pin assignments that you can edit or fracture. This method increases productivity by reducing the amount of time it takes to create and edit a new symbol.

### Using the Altera-provided Libraries with your Cadence Allegro Design Entry CIS Project

To use the Altera-provided libraries with your Cadence Allegro Design Entry CIS project, follow these steps:

1. Download the library of your target device from the Download Center page found through the Support page on the Altera website.
2. Create a copy of the appropriate **.olb** to maintain the original symbols. Place the copy in a convenient location, such as your Cadence Allegro Design Entry CIS project directory.
3. In the Project Manager window of the Cadence Allegro Design Entry CIS software, click once on the **Library** folder to select it. On the Edit menu, click **Project** or right-click the **Library** folder and choose **Add File** to select the copy of the downloaded **.olb** and add it to your project. You can locate the new library in the list of part libraries for your project.
4. On the Tools menu, click **Generate Part**. The **Generate Part** dialog box appears.
5. In the **Netlist/source file** field, click **Browse** to specify the **.pin** in your Quartus II design.
6. From the **Netlist/source file type** list, select **Altera Pin File**.
7. For **Part name**, type the name of the target device the same as it appears in the downloaded library file. For example, if you are using a device from the **CYCLONE06.OLB** library, type the part name to match one of the devices in this library such as **ep1c6f256**. You can rename the symbol in the Project Manager window after updating the part.

8. Set the **Destination part library** to the copy of the downloaded library you added to the project.
9. Select **Update pins on existing part in library**. Click **OK**.
10. Click **Yes**.

The symbol is updated with your pin assignments. Double-click the symbol in the Project Manager window to view and edit the symbol. On the View menu, click **Package** if you want to view and edit other sections of the symbol. If the symbol in the downloaded library is fractured into sections, you can edit each section but you cannot further fracture the part. You can generate a new part without using the downloaded part library if you require additional sections.

For more information about creating, editing, and fracturing symbols in the Cadence Allegro Design Entry CIS software, refer to the Help in the software.

## Document Revision History

Table 8-3: Document Revision History

Date	Version	Changes
June 2014	14.0.0	Converted to DITA format.
June 2012	12.0.0	Removed survey link.
November 2011	10.0.2	Template update.
December 2010	10.0.1	Template update.
July 2010	10.0.0	<ul style="list-style-type: none"> <li>• General style editing.</li> <li>• Removed Referenced Document Section.</li> <li>• Added a link to Help in “Performing Simultaneous Switching Noise (SSN) Analysis of Your FPGA” on page 9–5.</li> </ul>
November 2009	9.1.0	<ul style="list-style-type: none"> <li>• Added “Performing Simultaneous Switching Noise (SSN) Analysis of Your FPGA” on page 9–5.</li> <li>• General style editing.</li> <li>• Edited Figure 9–4 on page 9–10 and Figure 9–8 on page 9–16.</li> </ul>
March 2009	9.0.0	<ul style="list-style-type: none"> <li>• Chapter 9 was previously Chapter 7 in the 8.1 software release.</li> <li>• No change to content.</li> </ul>
November 2008	8.1.0	Changed to 8-1/2 x 11 page size.
May 2008	8.0.0	Updated references.

### Related Information

#### [Quartus II Handbook Archive](#)

For previous versions of the Quartus II Handbook



# Reviewing Printed Circuit Board Schematics with the Quartus II Software

# 9

2014.06.30

QIISV2



Subscribe



Send Feedback

Altera FPGAs and CPLDs offer a multitude of configurable options to allow you to implement a custom application-specific circuit on your PCB.

Your Quartus II project provides important information specific to your programmable logic design, which you can use in conjunction with the device literature available on Altera's website to ensure that you implement the correct board-level connections in your schematic.

Refer to the **Settings** dialog box options, the Fitter report, and Messages window when creating and reviewing your PCB schematic. The Quartus II software also provides the Pin Planner and the SSN Analyzer to assist you during your PCB schematic review process.

## Related Information

- [Schematic Review Worksheets](#)
- [Pin Connection Guidelines](#)

## Reviewing Quartus II Software Settings

Review these settings in the Quartus II software to help you review your PCB schematic.

The **Device** dialog box in the Quartus II software allows you to specify device-specific assignments and settings. You can use the **Device** dialog box to specify general project-wide options, including specific device and pin options, which help you to implement correct board-level connections in your PCB schematic.

The **Device** dialog box provides project-specific device information, including the target device and any migration devices you specify. Using migration devices can impact the number of available user I/O pins and internal resources, as well as require connection of some user I/O pins to power/ground pins to support migration.

If you want to use vertical migration, which allows you to use different devices with the same package, you can specify your list of migration devices in the **Migration Devices** dialog box. The Fitter places the pins in your design based on your targeted migration devices, and allows you to use only I/O pins that are common to all of the migration devices.

If a migration device has pins that are power or ground, but the pins are also user I/O pins on a different device in the migration path, the Fitter ensures that these pins are not used as user I/O pins. You must ensure that these pins are connected to the appropriate plane on the PCB.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered

If you are migrating from a smaller device with NC (no-connect) pins to a larger device with power or ground pins in the same package, you can safely connect the NC pins to power or ground pins to facilitate successful migration.

#### Related Information

#### [Migration Devices Dialog Box](#)

For more information about the **Migration Devices** dialog box in the Quartus II software

## Device and Pins Options Dialog Box Settings

You can set device and pin options and verify important design-specific data in the **Device and Pin Options** dialog box, including options found on the **General**, **Configuration**, **Unused Pin**, **Dual-Purpose Pins**, and **Voltage** pages.

### Configuration Settings

The **Configuration** page of the **Device and Pin Options** dialog box specifies the configuration scheme and configuration device for the target device. Use the **Configuration** page settings to verify the configuration scheme with the MSEL pin settings used on your PCB schematic and the I/O voltage of the configuration scheme.

Your specific configuration settings may impact the availability of some dual-purpose I/O pins in user mode.

Refer to [Dual-Purpose Pins Settings](#) on page 9-2 for more information.

#### Related Information

[Dual-Purpose Pins Settings](#) on page 9-2

### Unused Pin Settings

The **Unused Pin** page specifies the behavior of all unused pins in your design. Use the **Unused Pin** page to ensure that unused pin settings are compatible with your PCB. For example, if you reserve all unused pins as outputs driving ground, you must ensure that you do not connect unused I/O pins to VCC pins on your PCB. Connecting unused I/O pins to VCC pins may result in contention that could lead to higher than expected current draw and possible device overstress.

The **Reserve all unused pins** list shows available unused pin state options for the target device. The default state for each pin is the recommended setting for each device family.

When you reserve a pin as output driving ground, the Fitter connects a ground signal to the output pin internally. You should connect the output pin to the ground plane on your PCB, although you are not required to do so. Connecting the output driving ground to the ground plane is known as creating a virtual ground pin, which helps to minimize simultaneous switching noise (SSN) and ground bounce effects.

### Dual-Purpose Pins Settings

The **Dual-Purpose Pins** page specifies how configuration pins should be used after device configuration completes. You can set the function of the dual-purpose pins by selecting a value for a specific pin in the **Dual-purpose pins** list. Pin functions should match your PCB schematic. The available options on the **Dual-Purpose Pins** page may differ depending on the selected configuration mode.

## Voltage Settings

The **Voltage** page specifies the default VCCIO I/O bank voltage and the default I/O bank voltage for the pins on the target device. VCCIO I/O bank voltage settings made in the **Voltage** page are overridden by I/O standard assignments made on I/O pins in their respective banks.

Refer to the [Reviewing Device Pin-Out Information in the Fitter Report](#) on page 9-3 for more details about the I/O bank voltages for your design.

### Related Information

[Reviewing Device Pin-Out Information in the Fitter Report](#) on page 9-3

## Error Detection CRC Settings

The **Error Detection CRC** page specifies error detection cyclic redundancy check (CRC) use for the target device. When **Enable error detection CRC** is turned on, the device checks the validity of the programming data in the devices. Any changes made in the data while the device is in operation generates an error.

Turning on the **Enable open drain on CRC error pin** option allows the CRC ERROR pin to be set as an open-drain pin in some devices, which decouples the voltage level of the CRC ERROR pin from VCCIO voltage. You must connect a pull-up resistor to the CRC ERROR pin on your PCB if you turn on this option.

In addition to settings in the **Device** dialog box, you should verify settings in the **Voltage** page of the **Settings** dialog box.

### Related Information

#### [Device and Pin Options Dialog Box](#)

For more information about the **Device and Pins Options** dialog box in the Quartus II software

## Voltage Settings

The **Voltage** page, under **Operating Settings and Conditions** in the **Settings** dialog box, allows you to specify voltage operating conditions for timing and power analyses. Ensure that the settings in the **Voltage** page match the settings in your PCB schematic, especially if the target device includes transceivers.

The **Voltage** page settings requirements differ depending on the settings of the transceiver instances in the design. Refer to the Fitter report for the required settings, and verify that the voltage settings are correctly set up for your PCB schematic.

After verifying your settings in the **Device** and **Settings** dialog boxes, you can verify your device pin-out with the Fitter report.

### Related Information

#### [Pin Connection Guidelines](#)

For more information about voltage settings

## Reviewing Device Pin-Out Information in the Fitter Report


After you compile your design, you can use the reports in the Resource section of the Fitter report to check your device pin-out in detail.

The Input Pins, Output Pins, and Bidirectional Pins reports identify all the user I/O pins in your design and the features enabled for each I/O pin. For example, you can find use of weak internal pull-ups, PCI clamp diodes, and on-chip termination (OCT) pin assignments in these sections of the Fitter report. You can check the pin assignments reported in the Input Pins, Output Pins, and Bidirectional Pins reports against your PCB schematic to determine whether your PCB requires external components.

These reports also identify whether you made pin assignments or if the Fitter automatically placed the pins. If the Fitter changed your pin assignments, you should make these changes user assignments because the location of pin assignments made by the Fitter may change with subsequent compilations.

**Figure 9-1: Resource Section Report**

This figure shows the pins the Fitter chose for the OCT external calibration resistor connections (RUP/RDN) and the name of the associated termination block in the Input Pins report. You should make these types of assignments user assignments.



Compilation Report - Input Pins					
Resource Section					
Input Pins					
	Name	Pin #	I/O Bank	X coordin...	Y coordin
1	clock_source	AB39	2C	0	59
2	global_reset_n	AB41	2C	0	60
3	termination_blk0~_rdn_pad	C40	1A	0	113
4	termination_blk0~_rup_pad	D40	1A	0	113

The I/O Bank Usage report provides a high-level overview of the VCCIO and VREF requirements for your design, based on your I/O assignments. Verify that the requirements in this report match the settings in your PCB schematic. All unused I/O banks, and all banks with I/O pins with undefined I/O standards, default the VCCIO voltage to the voltage defined in the **Voltage** page of the **Device and Pin Options** dialog box.

The All Package Pins report lists all the pins on your device, including unused pins, dedicated pins and power/ground pins. You can use this report to verify pin characteristics, such as the location, name, usage, direction, I/O standard and voltage for each pin with the pin information in your PCB schematic. In particular, you should verify the recommended voltage levels at which you connect unused dedicated inputs and I/O and power pins, especially if you selected a migration device. Use the All Package Pins report to verify that you connected all the device voltage rails to the voltages reported.

Errors commonly reported include connecting the incorrect voltage to the predriver supply (VCCPD) pin in a specific bank, or leaving dedicated clock input pins floating. Unused input pins that should be connected to ground are designated as **GND+** in the **Pin Name/Usage** column in the All Package Pins report.

You can also use the All Package Pins report to check transceiver-specific pin connections and verify that they match the PCB schematic. Unused transceiver pins have the following requirements, based on the pin designation in the Fitter report:

- **GXB\_GND**—Unused GXB receiver or dedicated reference clock pin. This pin must be connected to GXB\_GND through a 10k Ohm resistor.
- **GXB\_NC**—Unused GXB transmitter or dedicated clock output pin. This pin must be disconnected.

Some transceiver power supply rails have dual voltage capabilities, such as VCCA\_L/R and VCCH\_L/R, that depend on the settings you created for the ALTGX parameter editor. Because these user-defined settings overwrite the default settings, you should use the All Package Pins report to verify that these power pins on the device symbol in the PCB schematics are connected to the voltage required by the transceiver. An incorrect connection may cause the transceiver to function not as expected.

If your design includes a memory interface, the DQS Summary report provides an overview of each DQ pin group. You can use this report to quickly confirm that the correct DQ/DQS pins are grouped together.

Finally, the Fitter Device Options report summarizes some of the settings made in the **Device and Pin Options** dialog box. Verify that these settings match your PCB schematics.

## Reviewing Compilation Error and Warning Messages

If your project does not compile without error or warning messages, you should resolve the issues identified by the Compiler before signing off on your pin-out or PCB schematic. Error messages often indicate illegal or unsupported use of the device resources and IP.

Additionally, you should cross-reference fitting and timing analysis warnings with the design implementation. Timing may be constrained due to nonideal pin placement. You should investigate if you can reassign pins to different locations to prevent fitting and timing analysis warnings. Ensure that you review each warning and consider its potential impact on the design.

## Using Additional Quartus II Software Features

You can generate IBIS files, which contain models specific to your design and selected I/O standards and options, with the Quartus II software.

Because board-level simulation is important to verify, you should check for potential signal integrity issues. You can turn on the **Board-Level Signal Integrity** feature in the **EDA Tool Settings** page of the **Settings** dialog box.

Additionally, using advanced I/O timing allows you to enter physical PCB information to accurately model the load seen by an output pin. This feature facilitates accurate I/O timing analysis.

### Related Information

- [Signal Integrity Analysis with Third-Party Tools](#) on page 6-1  
For more information about signal integrity analysis in the Quartus II software

- [I/O Management](#) on page 4-1  
For more information about advanced I/O timing

## Using Additional Quartus II Software Tools

Use the Pin Planner and the SSN Analyzer to assist you with reviewing your PCB schematics.

### Pin Planner

The Quartus II Pin Planner helps you visualize, plan, and assign device I/O pins in a graphical view of the target device package. You can quickly locate various I/O pins and assign them design elements or other properties to ensure compatibility with your PCB layout.

You can use the Pin Planner to verify the location of clock inputs, and whether they have been placed on dedicated clock input pins, which is recommended when your design uses PLLs.

You can also use the Pin Planner to verify the placement of dedicated SERDES pins. SERDES receiver inputs can be placed only on DIFFIO\_RX pins, while SERDES transmitter outputs can be placed only on DIFFIO\_TX pins.

The Pin Planner gives a visual indication of signal-to-signal proximity in the Pad View window, and also provides information about differential pin pair placement, such as the placement of pseudo-differential signals.

#### Related Information

- [I/O Management](#) on page 4-1  
For more information about the Pin Planner

### SSN Analyzer

The SSN Analyzer supports pin planning by estimating the voltage noise caused by the simultaneous switching of output pins on the device. Because of the importance of the potential SSN performance for a specific I/O placement, you can use the SSN Analyzer to analyze the effects of aggressor I/O signals on a victim I/O pin.

#### Related Information

- [Simultaneous Switching Noise \(SSN\) Analysis and Optimizations](#) on page 5-1  
For more information about the SSN Analyzer

## Document Revision History

Table 9-1: Document Revision History

Date	Version	Changes
June 2014	14.0.0	Template update.
November 2012	12.1.0	Minor update of Pin Planner description for task and report windows.

Date	Version	Changes
June 2012	12.0.0	Removed survey link.
November 2011	10.0.2	Template update.
December 2010	10.0.1	Changed to new document template. No change to content.
July 2010	10.0.0	Initial release.

**Related Information**

[Quartus II Handbook Archive](#)

For previous versions of the Quartus II Handbook

2014.12.15

QI15V2



Subscribe



Send Feedback

## Design Optimization Overview

This chapter introduces features in Altera's Quartus® II software that you can use to achieve the highest design performance when you design for programmable logic devices (PLDs), especially high density FPGAs.

Physical implementation can be an intimidating and challenging phase of the design process. The Quartus II software provides a comprehensive environment for FPGA designs, delivering unmatched performance, efficiency, and ease-of-use.

In a typical design flow, you must synthesize your design with Quartus II integrated synthesis or a third-party tool, place and route your design with the Fitter, and use the TimeQuest timing analyzer to ensure your design meets the timing requirements. With the PowerPlay Power Analyzer, you ensure the design's power consumption is within limits.

## Initial Compilation: Required Settings

There are basic assignments and settings Altera recommends for your initial compilation. Check the following settings before compiling your design in the Quartus II software. Significantly varied compilation results can occur depending on the assignments that you set.

## Device Settings

Device assignments determine the timing model that the Quartus II software uses during compilation.

Choose the correct speed grade to obtain accurate results and the best optimization. The device size and the package determine the device pin-out and the available resources in the device.

## Device Migration Settings

If you anticipate a change to the target device later in the design cycle, either because of changes in your design or other considerations, plan for the change at the beginning of your design cycle.

Whenever you select a target device, you can also list any other compatible devices you can migrate by clicking on the **Migration Devices** button in the **Device** dialog box.

Selecting the migration device and companion device early in the design cycle helps to minimize changes to your design at a later stage.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered





## I/O Assignments

The I/O standards and drive strengths specified for a design affect I/O timing. Specify I/O assignments so that the Quartus II software uses accurate I/O timing delays in timing analysis and Fitter optimizations.

If there is no PCB layout requirement, then you do not need to specify pin locations. If your pin locations are not fixed due to PCB layout requirements, then leave the pin locations unconstrained. If your pin locations are already fixed, then make pin assignments to constrain the compilation appropriately.

Use the Assignment Editor and Pin Planner to assign I/O standards and pin locations.

### Related Information

[Timing Closure and Optimization](#) on page 12-1

For more information about recommendations for making pin assignments that can have a large effect on your results in smaller macrocell-based architectures.

[I/O Management](#) on page 4-1

For more information about I/O standards and pin constraints, refer to the appropriate device handbook. For more information about planning and checking I/O assignments.

### About the Assignment Editor

For information about using the Assignment Editor, refer to Quartus II Help.

## Timing Requirement Settings

Use your real requirements to get the best results. If you apply more demanding timing requirements than you need, then increased resource usage, higher power utilization, increased compilation time, or all of these may result.

You must use comprehensive timing requirement settings to achieve the best results for the following reasons:

- Correct timing assignments enable the software to work hardest to optimize the performance of the timing-critical parts of your design and make trade-offs for performance. This optimization can also save area or power utilization in non-critical parts of your design.
- If enabled, the Quartus II software performs physical synthesis optimizations based on timing requirements.
- Depending on the **Fitter Effort** setting, the Fitter can reduce runtime if your design meets the timing requirements.

The Quartus II TimeQuest Timing Analyzer determines if the design implementation meets the timing requirement. The Compilation Report shows whether your design meets the timing requirements, while the timing analysis reporting commands provide detailed information about the timing paths.

To create timing constraints for the TimeQuest analyzer, create a Synopsys Design Constraints File (**.sdc**). You can also enter constraints in the TimeQuest GUI. Use the `write_sdc` command, or the Constraints menu in the TimeQuest analyzer. Click **Write SDC File** to write your constraints to an **.sdc**. You can add an **.sdc** to your project on the **Quartus II Settings** page under **Timing Analysis Settings**.

**Note:** If you already have an **.sdc** in your project, using the `write_sdc` command from the command line or using the **Write SDC File** option from the TimeQuest GUI allows you to create a new **.sdc** that combines the constraints from your current **.sdc** and any new constraints added through the GUI or command window, or overwrites the existing **.sdc** with your newly applied constraints.

Ensure that every clock signal has an accurate clock setting constraint. If clocks arrive from a common oscillator, then they are related. Ensure that you set up all related or derived clocks in the constraints correctly. You must constrain all I/O pins that require I/O timing optimization. Specify both minimum

and maximum timing constraints as applicable. If your design contains more than one clock or contains pins with different I/O requirements, make multiple clock settings and individual I/O assignments instead of using a global constraint.

Make any complex timing assignments required in your design, including false path and multicycle path assignments. Common situations for these types of assignments include reset or static control signals (when the time required for a signal to reach a destination is not important) or paths that have more than one clock cycle available for operation in a design. These assignments enable the Quartus II software to make appropriate trade-offs between timing paths and can enable the Compiler to improve timing performance in other parts of your design.

**Note:** To ensure that you apply constraints or assignments to all design nodes, you can report all unconstrained paths in your design with the **Report Unconstrained Paths** command in the **Task** pane of the Quartus II TimeQuest Timing Analyzer or the `report_ucp` Tcl command.

#### Related Information

- [Timing Closure and Optimization](#) on page 12-1  
For more information about optimization with physical synthesis.
- [Advanced Settings \(Fitter\)](#)  
For more information about reducing runtime by changing Fitter effort.
- [The Quartus II TimeQuest Timing Analyzer](#)  
For more information about timing assignments and timing analysis.
- [Quartus II TimeQuest Timing Analyzer Cookbook](#)  
For more information about timing assignments and timing analysis.

## Partitions and Floorplan Assignments for Incremental Compilation

The Quartus II incremental compilation feature enables hierarchical and team-based design flows in which you can compile parts of your design while other parts of your design remain unchanged. You can also Import parts of your design from separate Quartus II projects.

Using incremental compilation for your design with good design partitioning methodology helps to achieve timing closure. Creating design partitions on some of the major blocks in your design and assigning them to LogicLock™ regions, reduces Fitter time and improves the quality and repeatability of the results. LogicLock regions are flexible, reusable floorplan location constraints that help you place logic on the target device. When you assign entity instances or nodes to a LogicLock region, you direct the Fitter to place those entity instances or nodes inside the region during fitting.

Using incremental compilation helps you achieve timing closure block by block and preserve the timing performance between iterations, which aid in achieving timing closure for the entire design. Incremental compilation may also help reduce compilation times.

**Note:** If you plan to use incremental compilation, you must create a floorplan for your design. If you are not using incremental compilation, creating a floorplan is optional.

#### Related Information

- [About LogicLock Regions](#)  
For more information about LogicLock regions.
- [Reducing Compilation Time](#) on page 11-1  
For more information about using incremental compilation to reduce compilation time.
- [Best Practices for Incremental Compilation Partitions and Floorplan Assignments](#)  
For more information about guidelines to create partition and floorplan assignments for your design.

## Physical Implementation

Most optimization issues involve preserving previous results, reducing area, reducing critical path delay, reducing power consumption, and reducing runtime.

The Quartus II software includes advisors to address each of these issues and helps you optimize your design. Run these advisors during physical implementation for advice about your specific design.

You can reduce the time spent on design iterations by following the recommended design practices for designing with Altera® devices. Design planning is critical for successful design timing implementation and closure.

### Related Information

[Design Planning with the Quartus II Software](#)

## Trade-Offs and Limitations

Many optimization goals can conflict with one another, so you might need to resolve conflicting goals. For example, one major trade-off during physical implementation is between resource usage and critical path timing, because certain techniques (such as logic duplication) can improve timing performance at the cost of increased area. Similarly, a change in power requirements can result in area and timing trade-offs, such as if you reduce the number of available high-speed tiles, or if you attempt to shorten high-power nets at the expense of critical path nets.

In addition, system cost and time-to-market considerations can affect the choice of device. For example, a device with a higher speed grade or more clock networks can facilitate timing closure at the expense of higher power consumption and system cost.

Finally, not all designs can be realized in a hardware circuit with limited resources and given constraints. If you encounter resource limitations, timing constraints, or power constraints that cannot be resolved by the Fitter, consider rewriting parts of the HDL code.

### Related Information

- [Timing Closure and Optimization](#) on page 12-1

## Preserving Results and Enabling Teamwork

For some Quartus II Fitter algorithms, small changes to the design can have a large impact on the final result. For example, a critical path delay can change by 10% or more because of seemingly insignificant changes. If you are close to meeting your timing objectives, you can use the Fitter algorithm to your advantage by changing the fitter seed, which changes the pseudo-random result of the Fitter.

Conversely, if you cannot meet timing on a portion of your design, you can partition that portion and prevent it from recompiling if an unrelated part of the design is changed. This feature, known as incremental compilation, can reduce the Fitter runtimes by up to 70% if the design is partitioned, such that only small portions require recompilation at any one time.

When you use incremental compilation, you can apply design optimization options to individual design partitions and preserve performance in other partitions by leaving them untouched. Many optimization techniques often result in longer compilation times, but by applying them only on specific partitions, you can reduce this impact and complete iterations more quickly.

In addition, by physically floorplanning your partitions with LogicLock regions, you can enable team-based flows and allow multiple people to work on different portions of the design.

### Related Information

- [Quartus II Incremental Compilation for Hierarchical and Team-Based Designs](#)
- [About Incremental Compilation](#)

## Reducing Area

By default, the Quartus II Fitter might physically spread a design over the entire device to meet the set timing constraints. If you prefer to optimize your design to use the smallest area, you can change this behavior. If you require reduced area, you can enable certain physical synthesis options to modify your netlist to create a more area-efficient implementation, but at the cost of increased runtime and decreased performance.

### Related Information

[Netlist Optimizations and Physical Synthesis](#) on page 16-1

[Timing Closure and Optimization](#) on page 12-1

[Recommended HDL Coding Styles](#)

## Reducing Critical Path Delay

To meet complex timing requirements involving multiple clocks, routing resources, and area constraints, the Quartus II software offers a close interaction between synthesis, timing analysis, floorplan editing, and place-and-route processes.

By default, the Quartus II Fitter tries to meet the specified timing requirements and stops trying when the requirements are met. Therefore, using realistic constraints is important to successfully close timing. If you under-constrain your design, you may get sub-optimal results. By contrast, if you over-constrain your design, the Fitter might over-optimize non-critical paths at the expense of true critical paths. In addition, you might incur an increased area penalty. Compilation time may also increase because of excessively tight constraints.

If your resource usage is very high, the Quartus II Fitter might have trouble finding a legal placement. In such circumstances, the Fitter automatically modifies some of its settings to try to trade off performance for area.

The Quartus II Fitter offers a number of advanced options that can help you improve the performance of your design when you properly set constraints. Use the Timing Optimization Advisor to determine which options are best suited for your design.

If you use incremental compilation, you can help resolve inter-partition timing requirements by locking down results, one partition at a time, or by guiding the placement of the partitions with LogicLock regions. You might be able to improve the timing on such paths by placing the partitions optimally to reduce the length of critical paths. Once your inter-partition timing requirements are met, use incremental compilation to preserve the results and work on partitions that have not met timing requirements.

In high-density FPGAs, routing accounts for a major part of critical path timing. Because of this, duplicating or retiming logic can allow the Fitter to reduce delay on critical paths. The Quartus II software offers push-button netlist optimizations and physical synthesis options that can improve design performance at the expense of considerable increases of compilation time and area. Turn on only those options that help you keep reasonable compilation times and resource usage. Alternately, you can modify your HDL to manually duplicate or adjust the timing logic.

## Reducing Power Consumption

The Quartus II software has features that help reduce design power consumption. The PowerPlay power optimization options control the power-driven compilation settings for Synthesis and the Fitter.

### Related Information

- [Power Optimization](#) on page 13-1

## Reducing Runtime

Many Fitter settings influence compilation time. Most of the default settings in the Quartus II software are set for reduced compilation time. You can modify these settings based on your project requirements.

The Quartus II software supports parallel compilation in computers with multiple processors. This can reduce compilation times by up to 15% while giving the identical result as serial compilation.

You can also reduce compilation time with your iterations by using incremental compilation. Use incremental compilation when you want to change parts of your design, while keeping most of the remaining logic unchanged.

## Using Quartus II Tools

The following sections describe several Quartus II tools that you can use to help optimize your design.

### Design Analysis

The Quartus II software provides tools that help with a visual representation of your design. You can use the RTL Viewer to see a schematic representation of your design before synthesis and place-and-route. The Technology Map Viewer provides a schematic representation of the design implementation in the selected device architecture after synthesis and place-and-route. It can also include timing information.

With incremental compilation, the Design Partition Planner and the Chip Planner allow you to partition and layout your design at a higher level. In addition, you can perform many different tasks with the Chip Planner, including: making floorplan assignments, implementing engineering change orders (ECOs), and performing power analysis. Also, you can analyze your design and achieve a faster timing closure with the Chip Planner. The Chip Planner provides physical timing estimates, critical path display, and a routing congestion view to help guide placement for optimal performance.

### Related Information

- [Quartus II Incremental Compilation for Hierarchical and Team-Based Designs](#)
- [Best Practices for Incremental Compilation Partitions and Floorplan Assignments](#)
- [Engineering Change Management with the Chip Planner](#) on page 17-1

For more information on design analysis for incremental compilation.

### Advisors

The Quartus II software includes several advisors to help you optimize your design and reduce compilation time.

You can complete your design faster by following the recommendations in the Compilation Time Advisor, Incremental Compilation Advisor, Timing Optimization Advisor, Area Optimization Advisor, Resource Optimization Advisor, and Power Optimization Advisor. These advisors give recommendations based on your project settings and your design constraints.

**Related Information**

[Running Advisors in the Quartus II Software](#)

For more information about advisors, refer to Quartus II Help.

## Design Space Explorer II

Use Design Space Explorer II (DSE) to find optimal settings in the Quartus II software.

DSE II automatically tries different combinations of netlist optimizations and advanced Quartus II software compiler settings, and reports the best settings for your design, based on your chosen primary optimization goal. You can try different seeds with DSE II if you are fairly close to meeting your timing or area requirements and find one seed that meets timing or area requirements. Finally, DSE II can run compilations on a remote compute farm, which shortens the timing closure process.

**Related Information**

[About Design Space Explorer II](#)

## Document Revision History

**Table 10-1: Document Revision History**

Date	Version	Changes
2014.12.15	14.1.0	<ul style="list-style-type: none"> <li>Updated location of Fitter Settings, Analysis &amp; Synthesis Settings, and Physical Synthesis Optimizations to Compiler Settings.</li> <li>Updated DSE II content.</li> </ul>
June 2014	14.0.0	Updated format.
November 2013	13.1.0	Minor changes for HardCopy.
May 2013	13.0.0	Added the information about initial compilation requirements. This section was moved from the Area Optimization chapter of the Quartus II Handbook. Minor updates to delineate division of Timing and Area optimization chapters.
June 2012	12.0.0	Removed survey link.
November 2011	10.0.3	Template update.
December 2010	10.0.2	Changed to new document template. No change to content.
August 2010	10.0.1	Corrected link
July 2010	10.0.0	Initial release. Chapter based on topics and text in Section III of volume 2.

**Related Information**[Quartus II Handbook Archive](#)

For previous versions of the Quartus II Handbook.

2014.12.15

QI15V2



Subscribe



Send Feedback

## Reducing Compilation Time

The Analysis and Synthesis and Fitter modules consume the majority of time in a compilation. The Quartus® II software offers several features and techniques to help reduce compilation time.

The Analysis and Synthesis module includes physical synthesis optimizations performed during synthesis, if you have turned on physical synthesis optimizations. The Fitter includes two steps, placement and routing, and also includes physical synthesis if you turned on the physical synthesis option with **Normal** or **Extra** effort levels. The **Flow Elapsed Time** section of the Compilation Report shows the duration of the Analysis and Synthesis and Fitter modules. The Fitter Messages report in the **Fitter** section of the Compilation Report displays the elapsed time for placement and routing processes.

Placement is the process of finding optimum locations for the logic in your design. Placement includes Quartus II pre-Fitter operations, which place dedicated logic such as clocks, PLLs, and transceiver blocks. Routing is the process of connecting the nets between the logic in your design. Finding better placement for the logic in your design requires more compilation time. Good logic placement allows you to more easily meet your timing requirements and makes your design easier to route.

```
Info: Fitter placement operations ending: elapsed time =  
<days:hours:minutes:seconds>  
Info: Fitter routing operations ending: elapsed time = <days:hours:minutes:seconds>
```

The Quartus II software displays info messages while the Fitter is running (including Placement and Routing). The Message window displays this message every hour to indicate Fitter operations are progressing normally.

```
Info: Placement optimizations have been running for 4 hour(s)
```

## Compilation Time Advisor

A Compilation Time Advisor is available to help you to reduce compilation time. Run the Compilation Time Advisor on the Tools menu by pointing to **Advisors** and clicking **Compilation Time Advisor**. You can find all the compilation time optimizing techniques described in this section in the Compilation Time Advisor as well.



## Strategies to Reduce the Overall Compilation Time

You can use the following strategies to reduce the overall time required to compile your design: , incremental compilation, and use of the Rapid Recompile and Smart Compilation features.

- Parallel compilation (for systems with multiple processor cores)
- Incremental compilation reduces compilation time by only recompiling design partitions that have not met design requirements.
- Rapid Recompile and Smart Compilation reuse results from a previous compilation to reduce overall compilation time

### Using Rapid Recompile

Rapid Recompile automatically reuses previous synthesis, placement, and routing results to reduce subsequent recompilation time and timing variations after making small design changes. Rapid Recompile is supported only for Arria V, Cyclone V, and Stratix V devices. You can use Rapid Recompile to implement HDL-based functional ECO changes that affect a small subset of a large or complex design (less than 5% of total design logic), without full recompilation. Rapid Recompile can achieve up to 4x reduction in compilation time for impacted portions of the design. Rapid Recompile works in conjunction with incremental compilation to isolate and compile design changes within a design partition, rather than recompiling the entire design partition for a small change.

To start Rapid Recompile following initial compilation, simply click **Processing > Start > Start Rapid Recompile**. Alternatively, you can type the following command to start rapid recompile at the command line: `quartus_sh -flow recompile <project name>`. Rapid Recompile implements the following type of design changes without full recompilation:

- Changes to nodes tapped by the SignalTap II Logic Analyzer
- Changes to combinational logic functions
- Changes to state machine logic (for example, new states, state transition changes)
- Changes to signal or bus latency or addition of pipeline registers
- Changes to coefficients of an adder or multiplier
- Changes register packing behavior of DSP, RAM, or I/O
- Removal of unnecessary logic
- Changes to synthesis directives

The Partition Merge Rapid Recompile Summary report provides detailed information about the recompilation of each design partition. The Incremental Compilation Preservation Summary report provides details about placement and routing implementation.

### Using Parallel Compilation with Multiple Processors

The Quartus II software can detect the number of processors available on a computer and use multiple processors to reduce compilation time.

You can control the number of processors used during a compilation on a per user basis. The Quartus II software can use up to 16 processors to run algorithms in parallel and reduce compilation time. The Quartus II software turns on parallel compilation by default to enable the software to detect available multiple processors. You can specify the maximum number of processors that the software can use if you want to reserve some of the available processors for other tasks.

**Note:** Do not consider processors with Intel Hyper-Threading as more than one processor. If you have a single processor with Intel Hyper-Threading enabled, you should set the number of processors to

one. Altera recommends that you do not use the Intel Hyper-Threading feature for Quartus II compilations, because it can increase runtimes.

The software does not necessarily use all the processors that you specify during a given compilation. Additionally, the software never uses more than the specified number of processors, enabling you to work on other tasks on your computer without it becoming slow or less responsive.

If you have partitioned your design and enabled parallel compilation, the Quartus II software can use different processors to compile those partitions simultaneously during Analysis and Synthesis. This can cause higher peak memory usage during Analysis and Synthesis.

You can reduce the compilation time by up to 10% on systems with two processing cores and by up to 20% on systems with four cores. With certain design flows in which timing analysis runs alone, multiple processors can reduce the time required for timing analysis by an average of 10% when using two processors. This reduction can reach an average of 15% when using four processors.

The actual reduction in compilation time when using incremental compilation partitions depends on your design and on the specific compilation settings. For example, compilations with multi-corner optimization turned on benefit more from using multiple processors than do compilations without multi-corner optimization. The runtime requirement is not reduced for some other compilation goals, such as Analysis and Synthesis. The Fitter (`quartus_fit`) and the Quartus II TimeQuest Timing Analyzer (`quartus_sta`) stages in the compilation can, in certain cases, benefit from the use of multiple processors. The **Flow Elapsed Time** panel of the Compilation Report shows the average number of processors for these stages. The Parallel Compilation panel of the appropriate report, such as the Fitter report, shows a more detailed breakdown of processor usage. This panel is displayed only if parallel compilation is enabled.

Parallel compilation is available for Arria<sup>®</sup> series, Cyclone<sup>®</sup>, MAX<sup>®</sup> II, MAX V (limited support), and Stratix<sup>®</sup> series devices.

You can also set the number of processors available for Quartus II compilation using the following Tcl command in your script:

```
set_global_assignment -name NUM_PARALLEL_PROCESSORS <value>
```

In this case, `<value>` is an integer from 1 to 16.

If you want the Quartus II software to detect the number of processors and use all the processors for the compilation, include the following Tcl command in your script:

```
set_global_assignment -name NUM_PARALLEL_PROCESSORS ALL
```

The use of multiple processors does not affect the quality of the fit. For a given Fitter seed on a specific design, the fit is exactly the same, regardless of whether the Quartus II software uses one processor or multiple processors. The only difference between compilations using a different number of processors is the compilation time.

#### Related Information

- [Processing Page \(Options Dialog Box\)](#)
- [Compilation Process Settings Page \(Settings Dialog Box\)](#)

For more information about how to control the number of processors used during compilation for a specific project, refer to Quartus II Help.

## Using Incremental Compilation

The incremental compilation feature can accelerate design iteration time by up to 70% for small design changes, and helps you reach design timing closure more efficiently.

You can speed up design iterations by recompiling only a particular design partition and merging results with previous compilation results from other partitions. You can also use physical synthesis optimization techniques for specific design partitions while leaving other parts of your design untouched to preserve performance.

If you are using a third-party synthesis tool, you can create separate atom netlist files for the parts of your design that you already have synthesized and optimized so that you update only the parts of your design that change.

In the standard incremental compilation design flow, you can divide the top-level design into partitions, which the software can compile and optimize in the top-level Quartus II project. You can preserve fitting results and performance for completed partitions while other parts of your design are changing. Incremental compilation reduces the compilation time for each design iteration because the software does not recompile the unchanged partitions in your design.

The incremental compilation feature facilitates team-based design flows by enabling designers to create and optimize design blocks independently, when necessary, and supports third-party IP integration.

#### Related Information

- [Quartus II Incremental Compilation for Hierarchical and Team-Based Design](#)  
For more information about the full incremental compilation flow in the Quartus II software.
- [Section IV. Synthesis](#)  
For more information about creating multiple netlist files in third-party tools for use with incremental compilation refer to the appropriate chapter in volume 1 of the *Quartus II Handbook*
- [About Incremental Compilation](#)  
For more information about incremental compilation, refer to Quartus II Help.

## Reducing Synthesis Time and Synthesis Netlist Optimization Time

You can reduce synthesis time without affecting the Fitter time by reducing your use of netlist optimizations and by using incremental compilation (with **Netlist Type** set to **Post-Synthesis**). For tips on reducing synthesis time when using third-party EDA synthesis tools, refer to your synthesis software's documentation.

### Settings to Reduce Synthesis Time and Synthesis Netlist Optimization Time

You can use Quartus II integrated synthesis to synthesize and optimize HDL designs, and you can use synthesis netlist optimizations to optimize netlists that were synthesized by third-party EDA software. When using Quartus II Integrated Synthesis, you can also enable Physical Synthesis Optimization before performing Analysis and Synthesis. Netlist optimizations can cause the Analysis and Synthesis module to take much longer to run. Read the Analysis and Synthesis messages to determine how much time these optimizations take. The compilation time spent in Analysis and Synthesis is usually short compared to the compilation time spent in the Fitter.

If your design meets your performance requirements without synthesis netlist optimizations, turn off the optimizations to save time. If you require synthesis netlist optimizations to meet performance, you can optimize parts of your design hierarchy separately to reduce the overall time spent in Analysis and Synthesis.

Turn off settings that are not useful. In general, if you carry over compilation settings from a previous project, evaluate all settings and keep only those that you need.

## Use Appropriate Coding Style to Reduce Synthesis Time

Your HDL coding style can also affect the synthesis time. For example, if you want to infer RAM blocks from your code, you must follow the guidelines for inferring RAMs. If RAM blocks are not inferred properly, the software implements those blocks as registers.

If you are trying to infer a large memory block, the software consumes more resources in the FPGA. This can cause routing congestion and increasing compilation time significantly. If you see high routing utilizations in certain blocks, it is a good idea to review the code for such blocks.

### Related Information

#### [Recommended HDL Coding Styles](#)

For more information about coding guidelines.

## Reducing Placement Time

The time required to place a design depends on two factors: the number of ways the logic in your design can be placed in the device and the settings that control the amount of effort required to find a good placement. You can reduce the placement time in two ways:

- Change the settings for the placement algorithm.
- Use incremental compilation to preserve the placement for the unchanged parts of your design.

Sometimes there is a trade-off between placement time and routing time. Routing time can increase if the placer does not run long enough to find a good placement. When you reduce placement time, ensure that it does not increase routing time and negate the overall time reduction.

## Fitter Effort Setting

The highest Fitter effort setting, **Standard Fit**, requires the most runtime, but does not always yield a better result than using the default **Auto Fit**.

For designs with very tight timing requirements, both **Auto Fit** and **Standard Fit** use the maximum effort during optimization. Altera recommends using **Auto Fit** for reducing compilation time. If you are certain that your design has only easy-to-meet timing constraints, you can select **Fast Fit** for an even greater runtime savings.

## Placement Effort Multiplier Settings

You can control the amount of time the Fitter spends in placement by reducing with the **Placement Effort Multiplier** option.

Click **Assignments > Settings > Compiler Settings > Advanced Settings (Fitter)** and specify a value for Placement Effort Multiplier. The default is 1.0. Legal values must be greater than 0 and can be non-integer values. Numbers between 0 and 1 can reduce fitting time, but also can reduce placement quality and design performance.

## Physical Synthesis Effort Settings

Physical synthesis options enable you to optimize your post-synthesis netlist and improve your timing performance. These options, which affect placement, can significantly increase compilation time.

If your design meets your performance requirements without physical synthesis options, turn them off to reduce compilation time. For example, if some or all of the physical synthesis algorithm information messages display an improvement of 0 ps, turning off physical synthesis can reduce compilation time.

You also can use the **Physical synthesis effort** setting on the **Advanced Fitter Settings** dialog box to reduce the amount of extra compilation time used by these optimizations.

The **Fast** setting directs the Quartus II software to use a lower level of physical synthesis optimization. Compared to the **Normal** physical synthesis effort level, using the **Fast** setting can cause a smaller increase in compilation time. However, the lower level of optimization can result in a smaller increase in design performance.

## Preserving Placement with Incremental Compilation

Preserving information about previous placements can make future placements faster. The incremental compilation feature provides an easy-to-use method for preserving placement results.

### Related Information

[Using Incremental Compilation](#) on page 11-3

## Reducing Routing Time

The time required to route a design depends on three factors: the device architecture, the placement of your design in the device, and the connectivity between different parts of your design.

The routing time is usually not a significant amount of the compilation time. If your design requires a long time to route, perform one or more of the following actions:

- Check for routing congestion.
- Turn off **Fitter Aggressive Routability Optimization**.
- Use incremental compilation to preserve routing information for parts of your design.

## Identifying Routing Congestion in the Chip Planner

To identify areas of routing congestion in your design, open the Chip Planner from the Tools menu.

To view the routing congestion in the Chip Planner, double-click the **Report Routing Utilization** command in the **Tasks** list. Click **Preview** in the **Report Routing Utilization** dialog box to preview the default congestion display. Change the **Routing utilization type** to display congestion for specific resources. The default display uses dark blue for 0% congestion and red for 100%. Adjust the slider for **Threshold percentage** to change the congestion threshold level.

Even if average congestion is not very high, your design may have areas where congestion is very high in a specific type of routing. You can use the Chip Planner to identify areas of high congestion for specific interconnect types. You can change the connections in your design to reduce routing congestion. If the area with routing congestion is in a LogicLock region or between LogicLock regions, change or remove the LogicLock regions and recompile your design. If the routing time remains the same, the time is a characteristic of your design and the placement. If the routing time decreases, consider changing the size, location, or contents of LogicLock regions to reduce congestion and decrease routing time.

Sometimes, routing congestion may be a result of the HDL coding style used in your design. After you identify congested areas using the Chip Planner, review the HDL code for the blocks placed in those areas to determine whether you can reduce interconnect usage by code changes.

The Quartus II compilation messages contain information about average and peak interconnect usage. Peak interconnect usage over 75%, or average interconnect usage over 60%, could be an indication that it might be difficult to fit your design. Similarly, peak interconnect usage over 90%, or average interconnect usage over 75%, are likely to have increased chances of not getting a valid fit.

## Preserving Routing with Incremental Compilation

Preserving the previous routing results for part of your design can reduce future routing time. Incremental compilation provides an easy-to-use methodology that preserves placement and routing results.

**Related Information**

[Using Incremental Compilation](#) on page 11-3

[Analyzing and Optimizing the Design Floorplan](#) on page 15-1

For more information about identifying areas of congested routing using the Chip Planner.

## Reducing Static Timing Analysis Time

If you are performing timing-driven synthesis, the Quartus II software runs the TimeQuest analyzer during Analysis and Synthesis.

The Quartus II Fitter also runs the TimeQuest analyzer during placement and routing. If there are incorrect constraints in the Synopsys Design Constraints File (.sdc), the Quartus II software may spend unnecessary time processing constraints several times.

- If you do not specify false paths and multicycle paths in your design, the TimeQuest analyzer may analyze paths that are not relevant to your design.
- If you redefine constraints in the .sdc files, the TimeQuest analyzer may spend additional time processing them. To avoid this situation, look for indications that Synopsys design constraints are being redefined in the compilation messages, and update the .sdc file.
- Ensure that you provide the correct timing constraints to your design, because the software cannot assume design intent, such as which paths to consider as false paths or multicycle paths. When you specify these assignments correctly, the TimeQuest analyzer skips analysis for those paths, and the Fitter does not spend additional time optimizing those paths.

## Setting Process Priority

It might be necessary to reduce the computing resources allocated to the compilation at the expense of increased compilation time. It can be convenient to reduce the resource allocation to the compilation with single processor machines if you must run other tasks at the same time.

**Related Information**

[Processing Page \(Options Dialog Box\)](#)

For more information about setting process priority, refer to Quartus II Help.

## Document Revision History

**Table 11-1: Document Revision History**

Date	Version	Changes
2014.12.15	14.1.0	<ul style="list-style-type: none"> <li>• Updated location of Fitter Settings, Analysis &amp; Synthesis Settings, and Physical Synthesis Optimizations to Compiler Settings.</li> <li>• Added information about Rapid Recompile feature.</li> </ul>
2014.08.18	14.0a10.0	Added restriction about smart compilation in Arria 10 devices.
June 2014	14.0.0	Updated format.

Date	Version	Changes
May 2013	13.0.0	<p>Removed the “Limit to One Fitting Attempt”, “Using Early Timing Estimation”, “Final Placement Optimizations”, and “Using Rapid Recompile” sections.</p> <p>Updated “Placement Effort Multiplier Settings” section.</p> <p>Updated “Identifying Routing Congestion in the Chip Planner” section.</p> <p>General editorial changes throughout the chapter.</p>
June 2012	12.0.0	Removed survey link.
November 2011	11.0.1	Template update.
May 2011	11.0.0	<ul style="list-style-type: none"> <li>• Updated “Using Parallel Compilation with Multiple Processors”.</li> <li>• Updated “Identifying Routing Congestion in the Chip Planner”.</li> <li>• General editorial changes throughout the chapter.</li> </ul>
December 2010	10.1.0	<ul style="list-style-type: none"> <li>• Template update.</li> <li>• Added details about peak and average interconnect usage.</li> <li>• Added new section “Reducing Static Timing Analysis Time”.</li> <li>• Minor changes throughout chapter.</li> </ul>
July 2010	10.0.0	Initial release.

**Related Information****[Quartus II Handbook Archive](#)**

For previous versions of the Quartus II Handbook.

2014.12.15

QI15V2



Subscribe



Send Feedback

## About Timing Closure and Optimization

This manual describes techniques to improve timing performance when designing for Altera® devices.

The application techniques vary between designs. Applying each technique does not always improve results. Settings and options in the Quartus® II software have default values that provide the best trade-off between compilation time, resource utilization, and timing performance. You can adjust these settings to determine whether other settings provide better results for your design.

## Initial Compilation: Optional Fitter Settings

The Fitter offers many optional settings; however, this section focuses on the optional timing-optimization related Fitter settings only, which are the **Optimize Hold Timing**, **Optimize Multi-Corner Timing**, and **Fitter Aggressive Routability Optimization** settings.

**Caution:** The settings required to optimize different designs could be different. The group of settings that work best for one design may not produce the best result for another design.

### Related Information

[Advanced Fitter Setting Dialog Box online help](#)

For scripting and device family support information of the **Optimize Hold Timing** and **Optimize Multi-Corner Timing** settings

## Optimize Hold Timing

The **Optimize Hold Timing** option directs the Quartus II software to optimize minimum delay timing constraints. By default, the Quartus II software optimizes hold timing for all paths for designs for supported devices. By default, the Quartus II software optimizes hold timing only for I/O paths and minimum  $t_{PD}$  paths for older devices.

When you turn on **Optimize Hold Timing** in the **Advanced Fitter Settings** dialog box, the Quartus II software adds delay to paths to ensure that your design meets the minimum delay requirements. If you select **I/O Paths and Minimum TPD Paths**, the Fitter works to meet the following criteria:

- Hold times ( $t_H$ ) from the device input pins to the registers
- Minimum delays from I/O pins to I/O registers or from I/O registers to I/O pins
- Minimum clock-to-out time ( $t_{CO}$ ) from registers to output pins

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

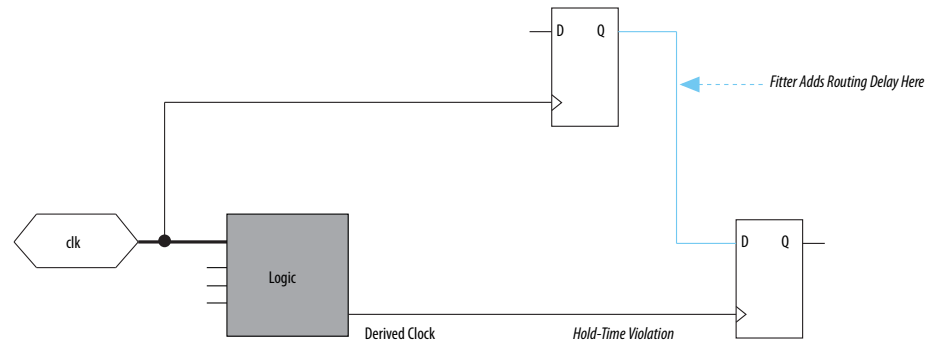
ISO  
9001:2008  
Registered





If you select **All Paths**, the Fitter also works to meet hold requirements from registers to registers, as highlighted in blue in the figure, in which a derived clock generated with logic causes a hold time problem on another register.

**Figure 12-1: Optimize Hold Timing Option Fixing an Internal Hold Time Violation**



However, if your design still has internal hold time violations between registers, correct the violations by manually adding some delays by instantiating LCELL primitives, or by making changes to your design, such as using a clock enable signal instead of a derived or gated clock.

#### Related Information

#### [Recommended Design Practices documentation](#)

For design practices that help to eliminate internal hold time violations

## Optimize Multi-Corner Timing

Due to process variations and changes in operating conditions, delays on some paths can be significantly smaller than those in the slow corner timing model. This can result in hold time violations on those paths, and in rare cases, additional setup time violations.

Also, because of the small process geometries of newer device families, the slowest circuit performance of designs targeting these devices does not necessarily occur at the highest operating temperature. The temperature at which the circuit is slowest depends on the selected device, the design, and the compilation results. Therefore, the Quartus II software provides newer device families with three different timing corners—Slow 85°C corner, Slow 0°C corner, and Fast 0°C corner. For other device families, two timing corners are available—Fast 0°C and Slow 85°C corner.

The **Optimize multi-corner timing** option directs the Fitter to consider all corner timing delays, including both fast-corner timing and slow-corner timing, during optimization to meet timing requirements at all process corners and operating conditions. By default, this option is on, and the Fitter optimizes designs considering multi-corner delays in addition to slow-corner delays, for example, from the fast-corner timing model, which is based on the fastest manufactured device, operating under high-voltage conditions

The **Optimize multi-corner timing** option helps to create a design implementation that is more robust across process, temperature, and voltage variations. Turning on this option increases compilation time by approximately 10%.

When this option is off, the Fitter optimizes designs considering only slow-corner delays from the slow-corner timing model (slowest manufactured device for a given speed grade, operating in low-voltage conditions).

## Fitter Aggressive Routability Optimization

The **Fitter Aggressive Routability Optimizations** logic option allows you to specify whether the Fitter aggressively optimizes for routability. Performing aggressive routability optimizations may decrease design speed, but may also reduce routing wire usage and routing time.

This option is useful if routing resources are resulting in no-fit errors, and you want to reduce routing wire use.

The table lists the settings for the **Fitter Aggressive Routability Optimizations** logic option.

**Table 12-1: Fitter Aggressive Routability Optimizations Logic Option Settings**

Settings	Description
Always	The Fitter always performs aggressive routability optimizations. If you set the Fitter Aggressive Routability Optimizations logic option to Always, reducing wire utilization may affect the performance of your design.
Never	The Fitter never performs aggressive routability optimizations. If improving timing is more important than reducing wire usage, then set this option to Automatically or Never.
Automatically	The Fitter performs aggressive routability optimizations automatically, based on the routability and timing requirements of the design. If improving timing is more important than reducing wire usage, then set this option to Automatically or Never.

## Design Analysis

The initial compilation establishes whether the design achieves a successful fit and meets the specified timing requirements. This section describes how to analyze your design results in the Quartus II software.

### Ignored Timing Constraints

The Quartus II software ignores illegal, obsolete, and conflicting constraints.

You can view a list of ignored constraints by clicking **Report Ignored Constraints** in the Reports menu in the TimeQuest GUI or by typing the following command to generate a list of ignored timing constraints:

```
report_sdc -ignored -panel_name "Ignored Constraints"
```

You should analyze any constraints that the Quartus II software ignores. If necessary, correct the constraints and recompile your design before proceeding with design optimization.

You can view a list of ignored assignment in the **Ignored Assignment Report** generated by the Fitter.

### Related Information

- [Quartus II TimeQuest Timing Analyzer documentation](#)  
For more information about the `report_sdc` command and its options
- [Fitter Summary Reports online help](#)

## I/O Timing (Including $t_{PD}$ )

TimeQuest analyzer supports the Synopsys Design Constraints (SDC) format for constraining your design. When using the TimeQuest analyzer for timing analysis, use the `set_input_delay` constraint to specify the data arrival time at an input port with respect to a given clock. For output ports, use the `set_output_delay` command to specify the data arrival time at an output port's receiver with respect to a given clock. You can use the `report_timing` Tcl command to generate the I/O timing reports.

The I/O paths that do not meet the required timing performance are reported as having negative slack and are highlighted in red in the TimeQuest analyzer **Report** pane. In cases where you do not apply an explicit I/O timing constraint to an I/O pin, the Quartus II timing analysis software still reports the **Actual** number, which is the timing number that must be met for that timing parameter when the device runs in your system.

### Related Information

#### [Quartus II TimeQuest Timing Analyzer documentation](#)

Information about how timing numbers are calculated

## Register-to-Register Timing

### Timing Analysis with the TimeQuest Timing Analyzer

Analyze all valid register-to-register paths by using the appropriate constraints in the TimeQuest analyzer. To view all timing summaries, run the **Report All Summaries** command by double-clicking **Report All Summaries** in the Tasks pane in the TimeQuest analyzer.

If any clock domains have failing paths (highlighted in red in the Report panel), right-click the Clock Name listed in the Clocks Summary panel and go to **Report Timing** to get more details. Your design meets timing requirements when you do not have negative slack on any register-to-register path on any of the clock domains.

When timing requirements are not met, a report on the failed paths (highlighted in red) can uncover more detail.

When you select a path listed in the TimeQuest **Report Timing** pane, the tabs in the corresponding path detail pane show a path summary of source and destination registers and their timing, statistics about the path delay, detailed information about the complete data path with all nodes in the path, and the waveforms of the relevant signals. The **Extra Fitter Information** tab will show a Graphical Data Path of where the offending path lies on the physical device. This can reveal whether the timing failure may be distance related, due to the source and destination node being too close or too far. The Chip Planner can also be used to investigate the physical layout of a failing path in more detail. To locate a selected path in the Chip Planner, right-click a node, point to **Locate**, and select **Locate in Chip Planner**. The Chip Planner appears with the path highlighted. Use this to show fanout, fanin, routing congestion, and region assignments information, and to determine whether those factors might be contributing to the timing critical path. Additionally, if you know that a path is not a valid path, you can set it to be a false path using the shortcut menu.

The **Data Path** tab can also be useful for determining contributions to timing critical paths. The **Data Path** tab shows details of the paths that the clock and data took to get from source to destination nodes, and the time it took on an incremental and cumulative basis. It also provides information about the routing types and elements used, and their locations.

To view the path details of any selected path, click the **Data Path** tab in the path details pane. The **Data Path** tab displays the details of the Data Arrival Path, as well as the Data Required Path.

The **Waveform** tab will show the slack relationship between arrival data and required data. This could be useful for determining how close or far off the path is from meeting timing.

To aid in timing debug, the RTL Viewer or Technology Map Viewer allow you to see schematic representations of your design. These viewers allow you to view a gate-level or technology-mapped representation of your design netlist. By providing a view of the path from source and destination nodes, the viewers can help identify areas in a design that may benefit from reducing the number of logic levels between the nodes. To locate a timing path in one of the viewers, right-click a path in the report, point to **Locate**, and click **Locate in RTL Viewer** or **Locate in Technology Map Viewer**.

#### Related Information

- [Quartus II TimeQuest Timing Analyzer documentation](#)  
Information about how timing analysis results are calculated
- [Analyzing Designs with Quartus II Netlist Viewers documentation](#)

### Tips for Analyzing Failing Paths

When you are analyzing failing paths, examine the reports and waveforms to determine if the correct constraints are being applied, and add timing exceptions as appropriate. A multicycle constraint relaxes setup or hold relationships by the specified number of clock cycles. A false path constraint specifies paths that can be ignored during timing analysis. Both constraints allow the Fitter to work harder on affected paths.

Focus on improving the paths that show the worst slack. The Fitter works hardest on paths with the worst slack. If you fix these paths, the Fitter might be able to improve the other failing timing paths in the design.

Check for particular nodes that appear in many failing paths. These nodes will appear in a timing report panel at the top of the list, along with their minimum slacks. Look for paths that have common source registers, destination registers, or common intermediate combinational nodes. In some cases, the registers might not be identical, but are part of the same bus.

In the timing analysis report panels, clicking on the **From** or **To** column headings can help to sort the paths by the source or destination registers. Clicking first on **From**, then on **To**, uses the registers in the **To** column as the primary sort and the registers in the **From** column as the secondary sort. If you see common nodes, these nodes indicate areas of your design that might be improved through source code changes or Quartus II optimization settings. Constraining the placement for just one of the paths might decrease the timing performance for other paths by moving the common node further away in the device.

#### Related Information

[Design Evaluation for Timing Closure](#) on page 12-23

## Tips for Analyzing Failing Clock Paths that Cross Clock Domains

When analyzing clock path failures, check whether these paths cross two clock domains. This is the case if the **From Clock** and **To Clock** in the timing analysis report are different.

**Figure 12-2: Different Value in From Clock and To Clock Field**

Setup Transfers						
	From Clock	To Clock	RR Paths	FR Paths	RF Paths	FF Paths
1	clkin	clkin	21	0	0	0
2	clkin	clkout	false path	0	0	0
3	clkout	clkout	31	0	0	0

There can also be paths that involve a different clock in the middle of the path, even if the source and destination register clock are the same.

When you run Report Timing on your design, the report shows the launch clock and latch clock for each failing path. Check whether these failing paths between these clock domains should be analyzed synchronously. If the failing paths are not to be analyzed synchronously, they must be set as false paths. Also check the relationship between the launch clock and latch clock to make sure it is realistic and what you expect from your knowledge of the design. For example, the path can start at a rising edge and end at a falling edge, which reduces the setup relationship by one half clock cycle.

Review the clock skew reported in the Timing Report. A large skew may indicate a problem in your design, such as a gated clock or a problem in the physical layout (for example, a clock using local routing instead of dedicated clock routing). When you have made sure the paths are analyzed synchronously and that there is no large skew on the path, and that the constraints are correct, you can analyze the data path. These steps help you fine tune your constraints for paths across clock domains to ensure you get an accurate timing report.

Check if the PLL phase shift is reducing the setup requirement. You might be able to adjust this using PLL parameters and settings.

Paths that cross clock domains are generally protected with synchronization logic (for example, FIFOs or double-data synchronization registers) to allow asynchronous interaction between the two clock domains. In such cases, you can ignore the timing paths between registers in the two clock domains while running timing analysis, even if the clocks are related.

The Fitter attempts to optimize all failing timing paths. If there are paths that can be ignored for optimization and timing analysis, but the paths do not have constraints that instruct the Fitter to ignore them, the Fitter tries to optimize those paths as well. In some cases, optimizing unnecessary paths can prevent the Fitter from meeting the timing requirements on timing paths that are critical to the design. It is beneficial to specify all paths that can be ignored by setting false path constraints on them, so that the Fitter can put more effort into the paths that must meet their timing requirements instead of optimizing paths that can be ignored.

### Related Information

#### [Quartus II TimeQuest Timing Analyzer](#)

Details about how to ignore timing paths that cross clock domains

## Tips for Analyzing Paths from/to the Source and Destination of Critical Path

When analyzing the failing paths in a design, it is often helpful to get a fuller picture of the many interactions the fitter may be working on around the paths. To understand what may be pulling on a critical path, the following `report_timing` command can be useful.

In the project directory, run the `report_timing` command, shown in the example below, in a `.tcl` file to analyze the nodes in a critical path.

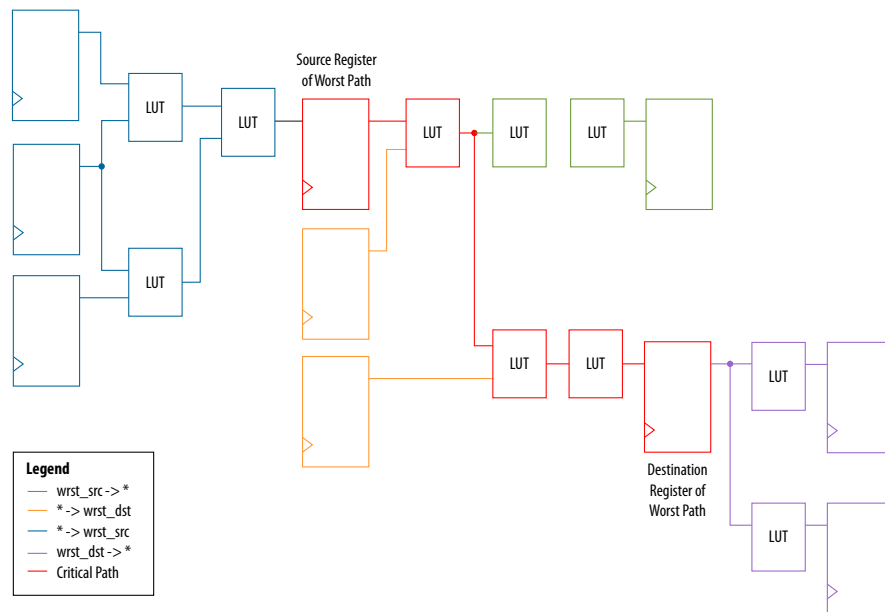
```
set wrst_src <insert_source_of_worst_path_here>
set wrst_dst <insert_destination_of_worst_path_here>
report_timing -setup -npaths 50 -detail path_only -from $wrst_src \
-panel_name "Worst Path||wrst_src -> *"
report_timing -setup -npaths 50 -detail path_only -to $wrst_dst \
-panel_name "Worst Path||* -> wrst_dst"
report_timing -setup -npaths 50 -detail path_only -to $wrst_src \
-panel_name "Worst Path||* -> wrst_src"
report_timing -setup -npaths 50 -detail path_only -from $wrst_dst \
-panel_name "Worst Path||wrst_dst -> *"
```

Copy the node names from the **From Node** and **To Node** columns of the worst path into the first two variables, and then in the TimeQuest timing analyzer, in the Script menu, source the `.tcl` script.

In the resulting timing panel, timing failed paths (highlighted in red) can be located in the Chip Planner, where information such as distance between the nodes and large fanouts can be viewed.

The figure shows a simplified example of what these reports analyzed.

**Figure 12-3: Timing Report**



The critical path of the design is in red. The script analyzes the path between the worst source and destination registers. The first `report_timing` command analyzes other path that the source is driving, as shown in green. The second `report_timing` command analyzes the critical path and other path going to the destination, shown in yellow. These commands report everything inside these two endpoints that are pulling them in different directions. The last two `report_timing` commands show everything outside of

the endpoints pulling them in other directions. If any of these reports have slacks near the critical path, then the Fitter is balancing these paths with the critical path, trying to achieve the best slack. The figure is quite simple compared to the critical path in most designs, but it is easy to see how this can get very complicated quickly.

### Tips for Locating Multiple Paths to the Chip Planner

The Chip Planner can be used as a visual aid in locating timing critical paths. To view these paths from timing reports, do the following:

1. Run `report_timing` to show multiple paths.
2. Select multiple rows of the timing report.
3. Right-click, select **Locate Path**, and then click **Chip Planner**.
4. The **Locate History** window in the Chip Planner displays the selected paths and the worst path.
5. Double-click **Locate Paths** to show all paths at once, or select individual paths to view the path in the Chip Planner.

This will show whether timing failures may be due to large distances between the nodes or large fanouts.

### Tips for Creating a .tcl Script to Monitor Critical Paths Across Compiles

Many designs have the same critical paths show up after each compile, but some suffer from having critical paths bounce around between different hierarchies, changing with each compile.

This could happen in high speed designs where many register to register paths have very little slack. Different placements can then result in timing failures in the marginal paths. In designs like this, create a **TQ\_critical\_paths.tcl** script in the project directory. For a given compile, view the critical paths and then write a generic `report_timing` command to capture those paths. For example, if several paths fail in a low-level hierarchy, you can add the following command:

```
report_timing -setup -npaths 50 -detail path_only \  
-to "main_system: main_system_inst|app_cpu:cpu|*" \  
-panel_name "Critical Paths||s: * -> app_cpu"
```

If there is a specific path, such as a bit of a state-machine going to other `*count_sync*` registers, you can add a command as shown by the following:

```
report_timing -setup -npaths 50 -detail path_only \  
-from "main_system: main_system_inst|egress_count_sm:egress_inst|update" \  
-to "*count_sync*" -panel_name "Critical Paths||s: egress_sm|update -> count_sync"
```

This file can be sourced in the TimeQuest timing analyzer after every compilation, and new `report_timing` commands can be added as new critical paths appear. This helps you monitor paths that consistently fail and paths that are only marginal, so you can prioritize effectively.

### Global Routing Resources

Global routing resources are designed to distribute high fan-out, low-skew signals (such as clocks) without consuming regular routing resources. Depending on the device, these resources can span the entire chip, or some smaller portion, such as a quadrant. The Quartus II software attempts to assign signals to global routing resources automatically, but you might be able to make more suitable assignments manually.

For details about the number and types of global routing resources available, refer to the relevant device handbook.

Check the global signal utilization in your design to ensure that the appropriate signals have been placed on the global routing resources. In the Compilation Report, open the Fitter report and click **Resource Section**. Analyze the Global & Other Fast Signals and Non-Global High Fan-out Signals reports to determine whether any changes are required.

You might be able to reduce skew for high fan-out signals by placing them on global routing resources. Conversely, you can reduce the insertion delay of low fan-out signals by removing them from global routing resources. Doing so can improve clock enable timing and control signal recovery/removal timing, but increases clock skew. Use the **Global Signal** setting in the Assignment Editor to control global routing resources.

## Optimizing Timing (LUT-Based Devices)

You can use the following guidelines if your design does not meet its timing requirements:

### Debugging Timing Failures in the TimeQuest Analyzer

A **Report Timing Closure Recommendations** task is available in the Custom Reports section of the **Tasks** pane of the TimeQuest analyzer. Use this report to get more information and help on the failing paths in your design.

When you run the **Report Timing Closure Recommendations** task, you get specific recommendations about failing paths in your design and changes that you can make to potentially fix the failing paths.

Selecting the **Report Timing Closure Recommendations** task opens the **Report Timing Closure Recommendations** dialog box.

From the **Report Timing Closure Recommendations** dialog box, you can select paths based on the clock domain, filter by nodes on path, and choose the number of paths to analyze.

After running the **Report Timing Closure Recommendations** task in the TimeQuest analyzer, examine the reports in the **Report Timing Closure Recommendations** folder in the **Report** pane of the TimeQuest analyzer GUI. Each recommendation has star symbols (\*) associated with it. Recommendations with more stars are more likely to help you close timing on your design.

The reports give you the most probable causes of failure for each path being analyzed. The reports are organized into sections, depending on the type of issues found in the design, such as large clock skew, restricted optimizations, unbalanced logic, skipped optimizations, coding style that has too many levels of logic between registers, or region or partition constraints specific to your project.

You will see recommendations that may help you fix the failing paths. For detailed analysis of the critical paths, run the `report_timing` command on specified paths. In the **Extra Fitter Information** tab of the Path report panel, you will also see detailed Fitter-related information that may help you visualize the issue and take the appropriate action if your constraints cause a specific placement.

#### Related Information

[Report Timing Closure Recommendations Dialog Box online help](#)



## Timing Optimization Advisor


While the TimeQuest **Report Timing Closure Recommendations** task gives specific recommendations to fix failing paths, the **Timing Optimization Advisor** gives more general recommendations to improve timing performance for a design.

The **Timing Optimization Advisor** guides you in making settings that optimize your design to meet your timing requirements. To run the Timing Optimization Advisor, on the Tools menu, point to **Advisors** and click **Timing Optimization Advisor**. This advisor describes many of the suggestions made in this section.

When you open the Timing Optimization Advisor after compilation, you can find recommendations to improve the timing performance of your design. Some of the recommendations in these advisors can contradict each other. Altera recommends evaluating these options and choosing the settings that best suit the given requirements.

The example shows the Timing Optimization Advisor after compiling a design that meets its frequency requirements, but requires setting changes to improve the timing.

Figure 12-4: Timing Optimization Advisor



When you expand one of the categories in the Timing Optimization Advisor, such as **Maximum Frequency (fmax)** or **I/O Timing (tsu, tco, tpd)**, the recommendations are divided into stages. The stages show the order in which to apply the recommended settings. The first stage contains the options that are easiest to change, make the least drastic changes to your design optimization, and have the least effect on compilation time. Icons indicate whether each recommended setting has been made in the current project. In the figure, the checkmark icons in the list of recommendations for Stage 1 indicate recommendations that are already implemented. The warning icons indicate recommendations that are not followed for this compilation. The information icons indicate general suggestions. For these entries, the advisor

does not report whether these recommendations were followed, but instead explains how you can achieve better performance. For a legend that provides more information for each icon, refer to the “How to use” page in the Timing Optimization Advisor.

There is a link from each recommendation to the appropriate location in the Quartus II GUI where you can change the settings. For example, consider the **Synthesis Netlist Optimizations** page of the **Settings** dialog box or the **Global Signals category** in the Assignment Editor. This approach provides the most control over which settings are made and helps you learn about the settings in the software. In some cases, you can also use the **Correct the Settings** button to automatically make the suggested change to global settings.

For some entries in the Timing Optimization Advisor, a button appears that allows you to further analyze your design and gives you more information. The advisor provides a table with the clocks in the design and indicates whether they have been assigned a timing constraint.

## I/O Timing Optimization

This stage of design optimization focuses on I/O timing. Ensure that you have made the appropriate assignments described in the “Initial Compilation: Required Settings” section in the **Design Optimization Overview** chapter of the *Quartus II Handbook*. You must also ensure that resource utilization is satisfactory before proceeding with I/O timing optimization. The suggestions provided in this section are applicable to all Altera FPGA families and to the MAX II family of CPLDs.

Because changes to the I/O paths affect the internal register-to-register timing, complete this stage before proceeding to the register-to-register timing optimization stage as described in **Register-to-Register Timing Optimization Techniques (LUT-Based Devices)**.

The options presented in this section address how to improve I/O timing, including the setup delay ( $t_{SU}$ ), hold time ( $t_H$ ), and clock-to-output ( $t_{CO}$ ) parameters.

### Improving Setup and Clock-to-Output Times Summary

The table lists the recommended order in which to use techniques to reduce  $t_{SU}$  and  $t_{CO}$  times. “Yes” indicates which timing parameters are affected by each technique. Reducing  $t_{SU}$  times increases hold ( $t_H$ ) times.

**Table 12-2: Improving Setup and Clock-to-Output Times**

Technique	Affects $t_{SU}$	Affects $t_{CO}$
Ensure that the appropriate constraints are set for the failing I/Os (refer to the “Initial Compilation: Required Settings” section in the <b>Design Optimization Overview</b> chapter of the <i>Quartus II Handbook</i> .)	Yes	Yes
Use timing-driven compilation for I/O ( <b>Fast Input, Output, and Output Enable Registers</b> )	Yes	Yes
Use fast input register ( <b>Programmable Delays</b> )	Yes	N/A
Use fast output register, fast output enable register, and fast OCT register ( <b>Programmable Delays</b> )	N/A	Yes
Decrease the value of <b>Input Delay from Pin to Input Register</b> or set <b>Decrease Input Delay to Input Register = ON</b>	Yes	N/A

Technique	Affects $t_{SU}$	Affects $t_{CO}$
Decrease the value of <b>Input Delay from Pin to Internal Cells</b> or set <b>Decrease Input Delay to Internal Cells = ON</b>	Yes	N/A
Decrease the value of <b>Delay from Output Register to Output Pin</b> or set <b>Increase Delay to Output Pin = OFF</b> ( <b>Fast Input, Output, and Output Enable Registers</b> )	N/A	Yes
Increase the value of <b>Input Delay from Dual-Purpose Clock Pin to Fan-Out Destinations</b> ( <b>Fast Input, Output, and Output Enable Registers</b> )	Yes	N/A
Use PLLs to shift clock edges ( <b>Use PLLs to Shift Clock Edges</b> )	Yes	Yes
Use the <b>Fast Regional Clock</b> ( <b>Change How Hold Times are Optimized for MAX II Devices</b> )	N/A	Yes
For MAX II or MAX V family devices, set <b>Guarantee I/O Paths Have Zero Hold Time at Fast Corner</b> to OFF, or <b>When <math>T_{SU}</math> and <math>T_{PD}</math> Constraints Permit</b> ( <b>Change How Hold Times are Optimized for MAX II Devices</b> )	Yes	N/A
Increase the value of <b>Delay to output enable pin</b> or set <b>Increase delay to output enable pin</b> ( <b>Use PLLs to Shift Clock Edges</b> )	N/A	Yes

Note to table :

1. These options may not apply to all device families.

## Timing-Driven Compilation

This option moves registers into I/O elements if required to meet  $t_{SU}$  or  $t_{CO}$  assignments, duplicating the register if necessary (as in the case in which a register fans out to multiple output locations). This option is turned on by default and is a global setting. The option does not apply to MAX II series devices because they do not contain I/O registers.

The **Optimize IOC Register Placement for Timing** option affects only pins that have a  $t_{SU}$  or  $t_{CO}$  requirement. Using the I/O register is possible only if the register directly feeds a pin or is fed directly by a pin. This setting does not affect registers with any of the following characteristics:

- Have combinational logic between the register and the pin
- Are part of a carry or cascade chain
- Have an overriding location assignment
- Use the asynchronous load port and the value is not 1 (in device families where the port is available)

Registers with the characteristics listed are optimized using the regular Quartus II Fitter optimizations.

### Related Information

[Optimize IOC Register Placement for Timing Logic Option online help](#)

## Fast Input, Output, and Output Enable Registers

Normally, with correct timing assignments, the Fitter already places the I/O registers in the correct I/O cell or in the core, to meet the performance requirement. However, you can place individual registers in I/O cells manually by making fast I/O assignments with the Assignment Editor.

For more information about the **Fast Input Register** option, **Fast Output Register** option, **Fast Output Enable Register** option, and **Fast OCT (on-chip termination) Register** option, refer to Quartus II Help.

In MAX II series devices, which have no I/O registers, these assignments lock the register into the LAB adjacent to the I/O pin if there is a pin location assignment for that I/O pin.

If the fast I/O setting is on, the register is always placed in the I/O element. If the fast I/O setting is off, the register is never placed in the I/O element. This is true even if the **Optimize IOC Register Placement for Timing** option is turned on. If there is no fast I/O assignment, the Quartus II software determines whether to place registers in I/O elements if the **Optimize IOC Register Placement for Timing** option is turned on.

You can also use the four fast I/O options (**Fast Input Register**, **Fast Output Register**, **Fast Output Enable Register**, and **Fast OCT Register**) to override the location of a register that is in a LogicLock region and force it into an I/O cell. If you apply this assignment to a register that feeds multiple pins, the register is duplicated and placed in all relevant I/O elements. In MAX II series devices, the register is duplicated and placed in each distinct LAB location that is next to an I/O pin with a pin location assignment.

## Programmable Delays

You can use various programmable delay options to minimize the  $t_{SU}$  and  $t_{CO}$  times. For Arria, Cyclone, MAX II, MAX V, and Stratix series devices, the Quartus II software automatically adjusts the applicable programmable delays to help meet timing requirements. Programmable delays are advanced options to use only after you compile a project, check the I/O timing, and determine that the timing is unsatisfactory. For detailed information about the effect of these options, refer to the device family handbook or data sheet.

After you have made a programmable delay assignment and compiled the design, you can view the implemented delay values for every delay chain for every I/O pin in the **Delay Chain Summary** section of the Compilation Report.

You can assign programmable delay options to supported nodes with the Assignment Editor. You can also view and modify the delay chain setting for the target device with the Chip Planner and Resource Property Editor. When you use the Resource Property Editor to make changes after performing a full compilation, recompiling the entire design is not necessary; you can save changes directly to the netlist. Because these changes are made directly to the netlist, the changes are not made again automatically when you recompile the design. The change management features allow you to reapply the changes on subsequent compilations.

Although the programmable delays in newer devices are user-controllable, Altera recommends their use for advanced users only. However, the Quartus II software might use the programmable delays internally during the Fitter phase.

For details about the programmable delay logic options available for Altera devices, refer to the following Quartus II Help topics:

[Input Delay from Pin to Input Register logic option](#)

[Input Delay from Pin to Internal Cells logic option](#)

[Output Enable Pin Delay logic option](#)

[Delay from Output Register to Output Pin logic option](#)

[Input Delay from Dual-Purpose Clock Pin to Fan-Out Destinations logic option](#)


**Related Information**

- [Engineering Change Management with the Chip Planner documentation](#) on page 17-1  
Information about using the Chip Planner and Resource Property Editor

**Use PLLs to Shift Clock Edges**

Using a PLL typically improves I/O timing automatically. If the timing requirements are still not met, most devices allow the PLL output to be phase shifted to change the I/O timing. Shifting the clock backwards gives a better  $t_H$  at the expense of  $t_{SU}$ , while shifting it forward gives a better  $t_{SU}$  at the expense of  $t_H$ . You can use this technique only in devices that offer PLLs with the phase shift option.

**Figure 12-5: Shift Clock Edges Forward to Improve  $t_{SU}$  at the Expense of  $t_H$**



You can achieve the same type of effect in certain devices by using the programmable delay called **Input Delay from Dual Purpose Clock Pin to Fan-Out Destinations**.

**Related Information**

[Input Delay from Dual-Purpose Clock Pin to Fan-Out Destinations Logic Option online help](#)

**Use Fast Regional Clock Networks and Regional Clocks Networks**

Altera devices have a variety of hierarchical clock structures. These include dedicated global clock networks, regional clock networks, fast regional clock networks, and periphery clock networks. The available resources differ between the various Altera device families.

For the number of clocking resources available in your target device, refer to the appropriate device handbook.

In general, fast regional clocks have less delay to I/O elements than regional and global clocks, and are used for high fan-out control signals. Regional clocks provide the lowest clock delay and skew for logic contained in a single quadrant. Placing clocks on these low-skew and low-delay clock nets provides better  $t_{CO}$  performance.

**Spine Clock Limitations**

Global clock networks, regional clock networks, and periphery clock networks have an additional level of clock hierarchy known as spine clocks. Spine clocks drive the final row and column clocks to their registers; thus, the clock to every register in the chip is reached through spine clocks. Spine clocks are not directly user controllable.

If your project has high clock routing demands, due to limitations in the Quartus II software, you may see spine clock errors. These errors are often seen with designs using multiple memory interfaces and high-speed serial interface (HSSI) channels (especially PMA Direct mode).

To reduce these spine clock errors, you can constrain your design to better use your regional clock resources using the following techniques:

- If your design does not use LogicLock regions, or if the LogicLock regions are not aligned to your clock region boundaries, create additional LogicLock regions and further constrain your logic.

**Note:** Register packing, a Fitter optimization option, may ignore LogicLock regions. If this occurs, disable register packing for specific instances through the Quartus II Assignment Editor.

- Some periphery features may ignore LogicLock region assignments. When this happens, the global promotion process may not function properly. To ensure that the global promotion process uses the correct locations, assign specific pins to the I/Os using these periphery features.
- By default, some IP MegaCore functions apply a global signal assignment with a value of dual-regional clock. If you constrain your logic to a regional clock region and set the global signal assignment to **Regional** instead of **Dual-Regional**, you can reduce clock resource contention.

## Change How Hold Times are Optimized for MAX II Devices

For MAX II devices, you can use the **Guarantee I/O Paths Have Zero Hold Time at Fast Corner** option to control how hold time is optimized by the Quartus II software.

### Related Information

[Guarantee I/O Paths Have Zero Hold Time at Fast Corner Logic Option online help](#)

## Register-to-Register Timing Optimization Techniques (LUT-Based Devices)

The next stage of design optimization is to improve register-to-register ( $f_{MAX}$ ) timing. The following sections provide available options if the performance requirements are not achieved after compilation.

Coding style affects the performance of your design to a greater extent than other changes in settings. Always evaluate your code and make sure to use synchronous design practices.

**Note:** When using the TimeQuest analyzer, register-to-register timing optimization is the same as maximizing the slack on the clock domains in your design. You can use the techniques described in this section to improve the slack on different timing paths in your design.

Before optimizing your design, understand the structure of your design as well as the type of logic affected by each optimization. An optimization can decrease performance if the optimization does not benefit your logic structure.

### Related Information

[Recommended Design Practices documentation](#)

Details about synchronous design practices and coding styles

## Optimize Source Code

In many cases, optimizing the design's source code can have a very significant effect on your design performance. In fact, optimizing your source code is typically the most effective technique for improving the quality of your results and is often a better choice than using LogicLock or location assignments.

Be aware of the number of logic levels needed to implement your logic while you are coding. Too many levels of logic between registers could result in critical paths failing timing. Try restructuring the design to use pipelining or more efficient coding techniques. Also, try limiting high fan-out signals in the source code. When possible, duplicate and pipeline control signals. Make sure the duplicate registers are protected by a preserve attribute, to avoid merging during synthesis.

If the critical path in your design involves memory or DSP functions, check whether you have code blocks in your design that describe memory or functions that are not being inferred and placed in dedicated logic. You might be able to modify your source code to cause these functions to be placed into high-performance dedicated memory or resources in the target device. When using RAM/DSP blocks, enable the optional input and output registers.

Ensure that your state machines are recognized as state machine logic and optimized appropriately in your synthesis tool. State machines that are recognized are generally optimized better than if the synthesis tool treats them as generic logic. In the Quartus II software, you can check the State Machine report under **Analysis & Synthesis** in the Compilation Report. This report provides details, including state encoding for each state machine that was recognized during compilation. If your state machine is not recognized, you might have to change your source code to enable it to be recognized.

#### Related Information

- [Recommended HDL Coding Styles documentation](#)  
Coding style guidelines including examples of HDL code for inferring memory, functions, guidelines, and sample HDL code for state machines
- [AN 584: Timing Closure Methodology for Advanced FPGA Designs application note.](#)

### Improving Register-to-Register Timing Summary

The choice of options and settings to improve the timing margin (slack) or to improve register-to-register timing depends on the failing paths in the design. To achieve the results that best approximate your performance requirements, apply the following techniques and compile the design after each step:

1. Ensure that your timing assignments are complete and correct. For details, refer to the “Initial Compilation: Required Settings” section in the [Design Optimization Overview](#) chapter of the *Quartus II Handbook*.
2. Ensure that you have reviewed all warning messages from your initial compilation and check for ignored timing assignments.
3. Apply netlist synthesis optimization options.
4. To optimize for speed, apply the following synthesis options:
  - [Optimize Synthesis for Speed, Not Area](#)
  - [Flatten the Hierarchy During Synthesis](#)
  - [Set the Synthesis Effort to High](#)
  - [Change State Machine Encoding](#)
  - [Prevent Shift Register Inference](#)
  - [Use Other Synthesis Options Available in Your Synthesis Tool](#)
5. To optimize for performance using physical synthesis, apply the following options:
  - Perform physical synthesis for combinational logic
  - Perform automatic asynchronous signal pipelining
  - Perform register duplication
  - Perform register retiming
  - Perform logic to memory mapping
6. Try different Fitter seeds. If there are very few paths that are failing by small negative slack, then you can try with a different seed to see if there is a fit that meets constraints in the Fitter seed noise.  
**Note:** Omit this step if a large number of critical paths are failing or if the paths are failing badly.
7. To control placement, make LogicLock assignments.
8. Make design source code modifications to fix areas of the design that are still failing timing requirements by significant amounts.
9. Make location assignments, or as a last resort, perform manual placement by back-annotating the design.

You can use Design Space Explorer II (DSE) to automate the process of running several different compilations with different settings.

If these techniques do not achieve performance requirements, additional design source code modifications might be required.

#### Related Information

[Design Space Explorer II online help](#)

## Physical Synthesis Optimizations

The Quartus II software offers physical synthesis optimizations that can help improve the performance of many designs regardless of the synthesis tool used. Physical synthesis optimizations can be applied both during synthesis and during fitting.

Physical synthesis optimizations that occur during the synthesis stage of the Quartus II compilation operate either on the output from another EDA synthesis tool or as an intermediate step in Quartus II integrated synthesis. These optimizations make changes to the synthesis netlist to improve either area or speed, depending on your selected optimization technique and effort level.

To view and modify the synthesis netlist optimization options, click **Assignments > Settings > Compiler Settings > Advanced Settings (Fitter)**.

If you use a third-party EDA synthesis tool and want to determine if the Quartus II software can remap the circuit to improve performance, you can use the **Perform WYSIWYG Primitive Resynthesis** option. This option directs the Quartus II software to unmap the LEs in an atom netlist to logic gates and then map the gates back to Altera-specific primitives. Using Altera-specific primitives enables the Fitter to remap the circuits using architecture-specific techniques.

The Quartus II technology mapper optimizes the design to achieve maximum speed performance, minimum area usage, or balances high performance and minimal logic usage, according to the setting of the **Optimization Technique** option. Set this option to **Speed** or **Balanced**.

The physical synthesis optimizations occur during the Fitter stage of the Quartus II compilation. Physical synthesis optimizations make placement-specific changes to the netlist that improve speed performance results for a specific Altera device.

The following physical synthesis optimizations are available during the Fitter stage for improving performance:

- Physical synthesis for combinational logic
- Automatic asynchronous signal pipelining
- Physical synthesis for registers
  - Register duplication
  - Register retiming

**Note:** If you want the performance gain from physical synthesis only on parts of your design, you can apply the physical synthesis options on specific instances.



To apply physical synthesis assignments for fitting on a per-instance basis, use the Quartus II Assignment Editor. The following assignments are available as instance assignments:

- Perform physical synthesis for combinational logic
- Perform register duplication for performance
- Perform register retiming for performance
- Perform automatic asynchronous signal pipelining

#### Related Information

- [Perform WYSIWYG Primitive Resynthesis Logic Option online help](#)
- [Optimization Technique Logic Option online help](#)
- [Working With Assignments in the Assignment Editor online help](#)

## Turn Off Extra-Effort Power Optimization Settings

If PowerPlay power optimization settings are set to **Extra Effort**, your design performance can be affected. If improving timing performance is more important than reducing power use, set the PowerPlay power optimization setting to **Normal**.

#### Related Information

[PowerPlay Power Optimization Logic Option online help](#)

[Power Optimization documentation](#) on page 13-1

## Optimize Synthesis for Speed, Not Area

The manner in which the design is synthesized has a large impact on design performance. Design performance varies depending on the way the design is coded, the synthesis tool used, and the options specified when synthesizing. Change your synthesis options if a large number of paths are failing or if specific paths are failing badly and have many levels of logic.

Set your device and timing constraints in your synthesis tool. Synthesis tools are timing-driven and optimized to meet specified timing requirements. If you do not specify a target frequency, some synthesis tools optimize for area.

Some synthesis tools offer an easy way to instruct the tool to focus on speed instead of area.

You can also specify this logic option for specific modules in your design with the Assignment Editor while leaving the default **Optimization Technique** setting at **Balanced** (for the best trade-off between area and speed for certain device families) or **Area** (if area is an important concern). You can also use the **Speed Optimization Technique for Clock Domains** option in the Assignment Editor to specify that all combinational logic in or between the specified clock domain(s) is optimized for speed.

To achieve best performance with push-button compilation, follow the recommendations in the following sections for other synthesis settings. You can use DSE II to experiment with different Quartus II synthesis options to optimize your design for the best performance.

#### Related Information

- [Optimization Technique Logic Option online help](#)
- [Synthesis documentation](#)  
Information about setting timing requirements and synthesis options in Quartus II integrated synthesis and third-party synthesis tools
- [Design Space Explorer II online help](#)

## Flatten the Hierarchy During Synthesis

Synthesis tools typically let you preserve hierarchical boundaries, which can be useful for verification or other purposes. However, the best optimization results generally occur when the synthesis tool optimizes across hierarchical boundaries, because doing so often allows the synthesis tool to perform the most logic minimization, which can improve performance. Whenever possible, flatten your design hierarchy to achieve the best results. If you are using Quartus II incremental compilation, you cannot flatten your design across design partitions. Incremental compilation always preserves the hierarchical boundaries between design partitions. Follow Altera's recommendations for design partitioning, such as registering partition boundaries to reduce the effect of cross-boundary optimizations.

### Related Information

[Quartus II Incremental Compilation for Hierarchical and Team-Based Design documentation](#)

## Set the Synthesis Effort to High

Some synthesis tools offer varying synthesis effort levels to trade off compilation time with synthesis results. Set the synthesis effort to **high** to achieve best results when applicable.

## Change State Machine Encoding

State machines can be encoded using various techniques. One-hot encoding, which uses one register for every state bit, usually provides the best performance. If your design contains state machines, changing the state machine encoding to one-hot can improve performance at the cost of area.

### Related Information

[State Machine Processing Logic Option online help](#)

## Duplicate Logic for Fan-Out Control

Duplicating logic or registers can help improve timing in cases where moving a register in a failing timing path to reduce routing delay creates other failing paths or where there are timing problems due to the fan-out of the registers. Most often, timing failures occur not because of the high fan-out registers, but because of the location of those registers. Duplicating registers, where source and destination registers are physically close, can help improve slack on critical paths.

Many synthesis tools support options or attributes that specify the maximum fan-out of a register. When using Quartus II integrated synthesis, you can set the **Maximum Fan-Out** logic option in the Assignment Editor to control the number of destinations for a node so that the fan-out count does not exceed a specified value. You can also use the `maxfan` attribute in your HDL code. The software duplicates the node as required to achieve the specified maximum fan-out.

Logic duplication using **Maximum Fan-Out** assignments normally increases resource utilization and can potentially increase compilation time, depending on the placement and the total resource usage within the selected device. The improvement in timing performance that results because of **Maximum Fan-Out** assignments is very design-specific. This is because when you use the **Maximum Fan-Out** assignment, although the Fitter duplicates the source logic to limit the fan-out, it may not be able to control the destinations that each of the duplicated sources drive. Since the **Maximum Fan-Out** destination does not specify which of the destinations the duplicated source should drive, it is possible that it might still be driving logic located all around the device. To avoid this situation, you could use the **Manual Logic Duplication** logic option.

If you are using **Maximum Fan-Out** assignments, Altera recommends benchmarking your design with and without these assignments to evaluate whether they give the expected improvement in timing performance. Use the assignments only when you get improved results.

You can manually duplicate registers in the Quartus II software regardless of the synthesis tool used. To duplicate a register, apply the **Manual Logic Duplication** logic option to the register with the Assignment Editor.

**Note:** Various Fitter optimizations may cause a small violation to the **Maximum Fan-Out** assignments to improve timing.

#### Related Information

[Manual Logic Duplication Logic Option online help](#)

## Prevent Shift Register Inference

In some cases, turning off the inference of shift registers increases performance. Doing so forces the software to use logic cells to implement the shift register instead of implementing the registers in memory blocks using the ALTSHIFT\_TAPS IP core. If you implement shift registers in logic cells instead of memory, logic utilization is increased.

## Use Other Synthesis Options Available in Your Synthesis Tool

With your synthesis tool, experiment with the following options if they are available:

- Turn on register balancing or retiming
- Turn on register pipelining
- Turn off resource sharing

These options can increase performance, but typically increase the resource utilization of your design.

## Fitter Seed

The Fitter seed affects the initial placement configuration of the design. Changing the seed value changes the Fitter results because the fitting results change whenever there is a change in the initial conditions. Each seed value results in a somewhat different fit, and you can experiment with several different seeds to attempt to obtain better fitting results and timing performance.

When there are changes in your design, there is some random variation in performance between compilations. This variation is inherent in placement and routing algorithms—there are too many possibilities to try them all and get the absolute best result, so the initial conditions change the compilation result.

**Note:** Any design change that directly or indirectly affects the Fitter has the same type of random effect as changing the seed value. This includes any change in source files, **Compiler Settings** or **Timing Analyzer Settings**. The same effect can appear if you use a different computer processor type or different operating system, because different systems can change the way floating point numbers are calculated in the Fitter.

If a change in optimization settings slightly affects the register-to-register timing or number of failing paths, you cannot always be certain that your change caused the improvement or degradation, or whether it could be due to random effects in the Fitter. If your design is still changing, running a seed sweep (compiling your design with multiple seeds) determines whether the average result has improved after an optimization change and whether a setting that increases compilation time has benefits worth the increased time (such as setting the **Physical Synthesis Effort** to **Extra**). The sweep also shows the amount of random variation to expect for your design.

If your design is finalized, you can compile your design with different seeds to obtain one optimal result. However, if you subsequently make any changes to your design, you might need to perform seed sweep again.

On the Assignments menu, select **Compiler Settings** to control the initial placement with the seed. You can use the DSE II to perform a seed sweep easily.

You can use the following Tcl command from a script to specify a Fitter seed:

```
set_global_assignment -name SEED <value>
```

#### Related Information

##### [Design Space Explorer II online help](#)

Information about compiling your design with different seeds using Design Space Explorer II

## Set Maximum Router Timing Optimization Level

To improve routability in designs where the router did not pick up the optimal routing lines, set the **Router Timing Optimization Level** to **Maximum**. This setting determines how aggressively the router tries to meet the timing requirements. Setting this option to **Maximum** can increase design speed slightly at the cost of increased compilation time. Setting this option to **Minimum** can reduce compilation time at the cost of slightly reduced design speed. The default value is **Normal**.

#### Related Information

##### [Router Timing Optimization Level Logic Option online help](#)

## LogicLock Assignments

Using LogicLock assignments to improve timing performance is only recommended for older Altera devices, such as the MAX II family. For other device families, especially for larger devices such as Arria and Stratix series devices, Altera does not recommend using LogicLock assignments to improve timing performance. For these devices, use the LogicLock feature for performance preservation and to floorplan your design.

LogicLock assignments do not always improve the performance of the design. In many cases, you cannot improve upon results from the Fitter by making location assignments. If there are existing LogicLock assignments in your design, remove the assignments if your design methodology permits it. Recompile the design, and then check if the assignments are making the performance worse.

When making LogicLock assignments, it is important to consider how much flexibility to give the Fitter. LogicLock assignments provide more flexibility than hard location assignments. Assignments that are more flexible require higher Fitter effort, but reduce the chance of design overconstraint. The following types of LogicLock assignments are available, listed in the order of decreasing flexibility:

- Auto size, floating location regions
- Fixed size, floating location regions
- Fixed size, locked location regions

If you are unsure of how big or where a LogicLock region should go, the **Auto/Floating** options are useful for your first pass. After you determine where a LogicLock region must go, modify the Fixed/Locked regions, as Auto/Floating LogicLock regions can hurt your overall performance. To determine what to put into a LogicLock region, refer to the timing analysis results and analyze the critical paths in the Chip

Planner. The register-to-register timing paths in the Timing Analyzer section of the Compilation Report help you recognize patterns.

#### Related Information

- [Analyzing and Optimizing the Design Floorplan with the Chip Planner documentation](#) on page 15-1

## Hierarchy Assignments

For a design with the hierarchy shown in the figure, which has failing paths in the timing analysis results similar to those shown in the table, `mod_A` is probably a problem module. In this case, a good strategy to fix the failing paths is to place the `mod_A` hierarchy block in a LogicLock region so that all the nodes are closer together in the floorplan.

Figure 12-6: Design Hierarchy




Table 12-3: Failing Paths in a Module Listed in Timing Analysis

From	To
mod_A reg1	mod_A reg9
mod_A reg3	mod_A reg5
mod_A reg4	mod_A reg6
mod_A reg7	mod_A reg10
mod_A reg0	mod_A reg2

Hierarchical LogicLock regions are also important if you are using an incremental compilation flow. Place each design partition for incremental compilation in a separate LogicLock region to reduce conflicts and ensure good results as the design develops. You can use the auto size and floating location regions to find a good design floorplan, but fix the size and placement to achieve the best results in future compilations.

#### Related Information

- [Quartus II Incremental Compilation for Hierarchical and Team-Based Design documentation](#)
- [Best Practices for Incremental Compilation and Floorplan Assignments documentation](#)
- [Analyzing and Optimizing the Design Floorplan with the Chip Planner documentation](#) on page 15-1

## Location Assignments

If a small number of paths are failing to meet their timing requirements, you can use hard location assignments to optimize placement. Location assignments are less flexible for the Quartus II Fitter than LogicLock assignments. In some cases, when you are familiar with your design, you can enter location constraints in a way that produces better results.

**Note:** Improving fitting results, especially for larger devices, such as Arria and Stratix series devices, can be difficult. Location assignments do not always improve the performance of the design. In many cases, you cannot improve upon the results from the Fitter by making location assignments.

## Metastability Analysis and Optimization Techniques

Metastability problems can occur when a signal is transferred between circuitry in unrelated or asynchronous clock domains, because the designer cannot guarantee that the signal will meet its setup and hold time requirements. The mean time between failures (MTBF) is an estimate of the average time between instances when metastability could cause a design failure.

You can use the Quartus II software to analyze the average MTBF due to metastability when a design synchronizes asynchronous signals and to optimize the design to improve the MTBF. These metastability features are supported only for designs constrained with the TimeQuest analyzer, and for select device families.

If the MTBF of your design is low, refer to the Metastability Optimization section in the Timing Optimization Advisor, which suggests various settings that can help optimize your design in terms of metastability.

This chapter describes how to enable metastability analysis and identify the register synchronization chains in your design, provides details about metastability reports, and provides additional guidelines for managing metastability.

### Related Information

- [Understanding Metastability in FPGAs white paper](#)
- [Managing Metastability with the Quartus II Software documentation](#)

## Design Evaluation for Timing Closure

Follow the guidelines in this section when you encounter timing failures in a design. The guidelines show you how to evaluate compilation results of a design and how to address some of the problems. While the guideline does not cover specific examples of restructuring RTL to improve design speed, the analysis techniques help you to evaluate changes that may have to be made to RTL to close timing.

## Review Compilation Results

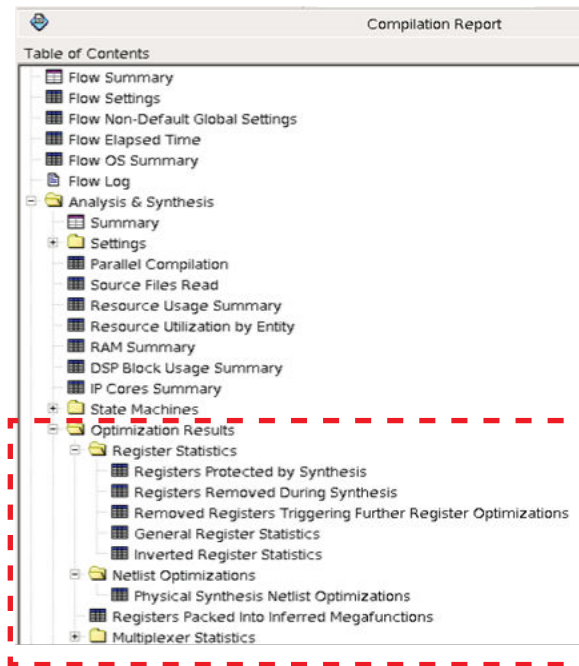
### Review Messages

After compiling your design, review the messages in each section of the compilation report. Most designs that fail timing start out with other problems that are reported as warning messages during compilation. Determine what causes a warning message, and whether the warning should be fixed or ignored. After reviewing the warning messages, review the informational messages. Take note of anything unexpected, for example, unconnected ports, ignored constraints, missing files, and assumptions or optimizations that the software made.

## Evaluate Physical Synthesis Results

If physical synthesis is enabled, the software can duplicate and retiming registers, and modify combinatorial logic during synthesis. After compilation, review the Optimization Results reports in the Analysis & Synthesis section. The reports list the optimizations performed by the physical synthesis optimizations, such as register duplication, retiming, and removal. These reports can be found in the Compilation Report panel.

Figure 12-7: Optimization Results Reports



When physical synthesis is enabled, compilation messages include a summary of the physical synthesis algorithms that were run, the performance improvement each algorithm achieved, and the elapsed time. The reported improvement is the sum of the largest improvement estimated to be achievable in each timing-critical clock domain. The values for the slack improvements can vary between compiles because of the random starting point of the compilation algorithms, but the values should be similar. The figure shows an example of the messages.

Figure 12-8: Compilation Messages

```
Starting physical synthesis algorithm register retiming
Physical synthesis algorithm register retiming complete: estimated slack improvement of 657 ps
Starting physical synthesis algorithm combinational resynthesis using boolean division
Physical synthesis algorithm combinational resynthesis using boolean division complete: estimated slack improvement of 471 ps
Starting physical synthesis algorithm register retiming
Physical synthesis algorithm register retiming complete: estimated slack improvement of 0 ps
Starting physical synthesis algorithm combinational resynthesis using boolean division
Physical synthesis algorithm combinational resynthesis using boolean division complete: estimated slack improvement of 0 ps
Physical synthesis optimizations for speed complete: elapsed CPU time is 00:31:42
```

## Evaluate Fitter Netlist Optimizations

The Fitter can also perform netlist optimizations to the design netlist. Major changes include register packing, duplicating or deleting logic cells, retiming registers, inverting signals, or modifying nodes in a

general way such as moving an input from one logic cell to another. These reports can be found in the Netlist Optimizations results of the Fitter section, and they should also be reviewed.

### Evaluate Optimization Results

After checking what optimizations were done and how they improved performance, evaluate the runtime it took to get the extra performance. To reduce compilation time, review the physical synthesis and netlist optimizations over a couple of compilations, and edit the RTL to reflect the changes that physical synthesis performed. If a particular set of registers consistently get retimed, edit the RTL to retime the registers the same way. If the changes are made to match what the physical synthesis algorithms did, the physical synthesis options can be turned off to save compile time while getting the same type of performance improvement.

### Evaluate Resource Usage

Evaluate a variety of resources used in the design, including global and non-global signal usage, routing utilization, and clustering difficulty.

#### Global and Non-global Usage

If your design contains a lot of clocks, evaluate global and non-global signals. Determine whether global resources are being used effectively, and if not, consider making changes. These reports can be found in the Resource Section under Fitter in the Compilation Report panel. The figure shows an example of inefficient use of a global clock. The highlighted line has a single fan-out from a global clock. Assigning it to a Regional Clock would make the Global Clock available for another signal. You can ignore signals with an empty value in the **Global Line Name** column as the signal uses dedicated routing, and not a clock buffer.

Figure 12-9: Inefficient Use of a Global Clock

Global & Other Fast Signals			
Location	Fan-Out	Global Resource Used	Global Line Name
FRACTIONALPLL_X98_Y2_N0	1	Global Clock	--
PLLOUTPUTCOUNTER_X98_Y2_N1	29044	Global Clock	GCLK7
PLLOUTPUTCOUNTER_X98_Y13_N1	253103	Global Clock	GCLK6
FF_X185_Y66_N13	280349	Global Clock	GCLK8
PIN_AE17	4887	Global Clock	GCLK4
FRACTIONALPLL_X98_Y11_N0	1	Global Clock	--
PLLOUTPUTCOUNTER_X98_Y3_N1	1	Global Clock	GCLK5
PLLOUTPUTCOUNTER_X98_Y1_N1	1691	Regional Clock	RCLK29
PLLOUTPUTCOUNTER_X98_Y8_N1	302	Regional Clock	RCLK23
PLLOUTPUTCOUNTER_X98_Y11_N1	141	Regional Clock	RCLK25
PLLOUTPUTCOUNTER_X98_Y10_N1	17	Regional Clock	RCLK22

The Non-Global High Fan-Out Signals report lists the highest fan-out nodes that are not routed on global signals. Reset and enable signals are at the top of the list. If there is routing congestion in the design, and there are high fan-out non-global nodes in the congested area, consider using global or regional signals to fan-out the nodes, or duplicate the high fan-out registers so that each of the duplicates can have fewer fan-outs. Use the Chip Planner to locate high fan-out nodes, to report routing congestion, and to determine whether the alternatives are viable.

#### Routing Usage

Review routing usage reported in the **Fitter Resource Usage Summary** report. The figure shows an example of the report.




Figure 12-10: Fitter Resource Usage Summary Report

	Resource	Usage
43	10G TX PCSs	12 / 36 ( 33 % )
44	HSSI FMA TX Serializers	12 / 36 ( 33 % )
45	CHANNEL PLLs	12 / 36 ( 33 % )
46	Impedance control blocks	1 / 4 ( 25 % )
47	Average interconnect usage (total/H/W)	55% / 55% / 55%
48	Peak interconnect usage (total/H/W)	88% / 88% / 90%
49		

The average interconnect usage reports the average amount of interconnect that is used, out of what is available on the device. The peak interconnect usage reports the largest amount of interconnect used in the most congested areas. Designs with an average value below 50% typically do not have any problems with routing. Designs with an average between 50-65% may have difficulty routing. Designs with an average over 65% typically have difficulty meeting timing unless the RTL is well designed to tolerate a highly utilized chip. Peak values at or above 90% are likely to have problems with timing closure; a 100% peak value indicates that all routing in an area of the device has been used, so there is a high possibility of degradation in timing performance. The figure shows the **Report Routing Utilization** report.

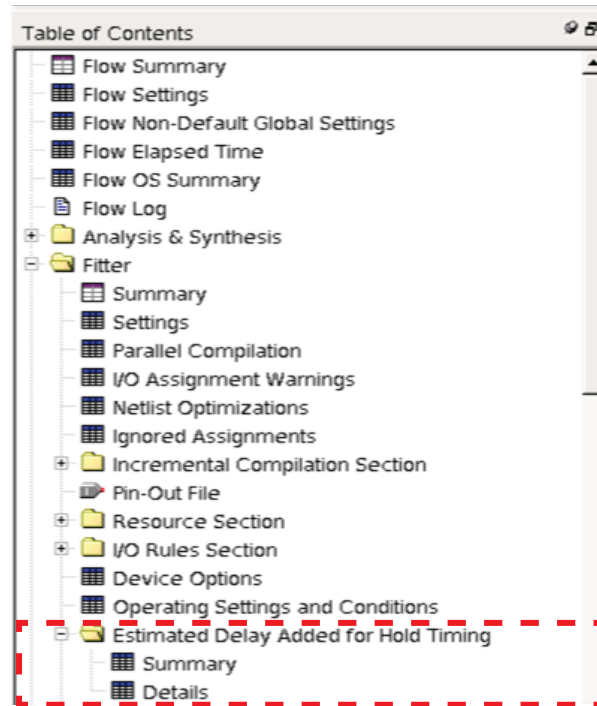
Figure 12-11: Report Routing Utilization Report



### Wires Added for Hold

As part of the fitting process, the router can add wire between register paths to increase delay to meet hold time requirements. During the routing process, the router reports how much extra wire was used to meet hold time requirements. Excessive amounts of added wire can indicate problems with the constraint. Typically it would be caused by incorrect multicycle transfers, particularly between different rate clocks, and between different clock networks. The Fitter reports how much routing delay was added in the **Estimated Delay Added for Hold Timing** report. Specific register paths can be reviewed to view whether a delay was added to meet hold requirements.

Figure 12-12: Estimated Delay Added for Hold Timing Report




An example of an incorrect constraint which can cause the router to add wire for hold requirements is when there is data transfer from 1x to 2x clocks. Assume the design intent is to allow two cycles per transfer. Data can arrive any time in the two destination clock cycles by adding a multicycle setup constraint as shown in the example:

```
set_multicycle_path -from 1x -to 2x -setup -end 2
```

The timing requirement is relaxed by one 2x clock cycle, as shown in the black line in the waveform in the figure.

Figure 12-13: Timing Requirement Relaxed Waveform



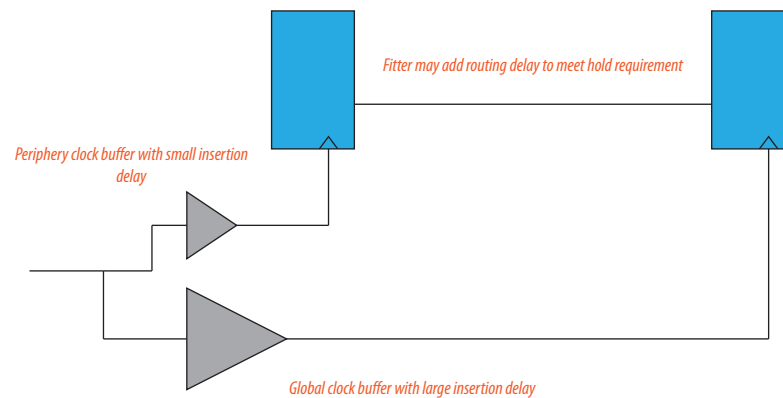
However, the default hold requirement, shown with the dashed blue line, may cause the router to add wire to guarantee that data is delayed by one cycle. To correct the hold requirement, add a multicycle constraint with a hold option.

```
set_multicycle_path -from 1x -to 2x -setup -end 2
set_multicycle_path -from 1x -to 2x -hold -end 1
```

The orange dashed line in the figure above represents the hold relationship, and no extra wire is required to delay the data.

The router can also add wire for hold timing requirements when data is transferred in the same clock domain, but between clock branches that use different buffering. Transferring between clock network types happens more often between the periphery and the core. The figure below shows a case where data is coming into a device, and uses a periphery clock to drive the source register, and a global clock to drive the destination register. A global clock buffer has larger insertion delay than a periphery clock buffer. The clock delay to the destination register is much larger than to the source register, hence extra delay is necessary on the data path to ensure that it meets its hold requirement.

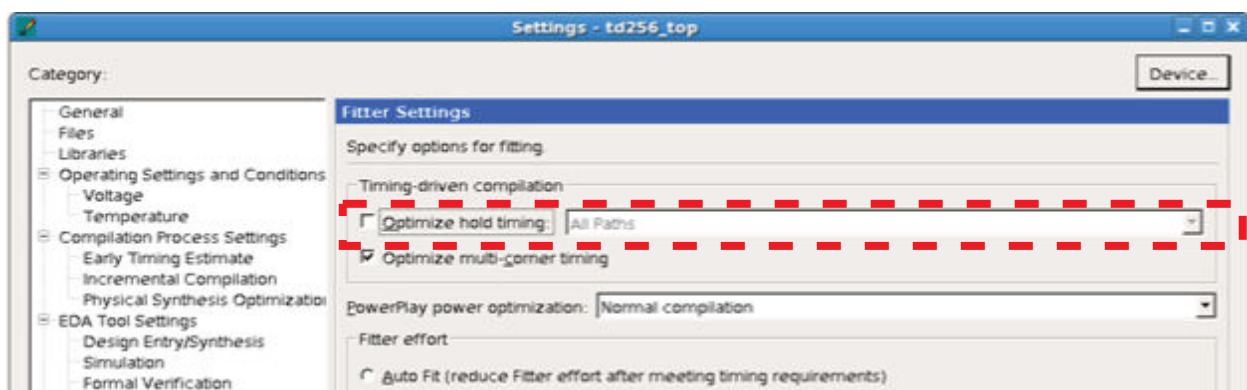
Figure 12-14: Clock Delay



To identify cases where a path has different clock network types, review the path in the TimeQuest timing analyzer, and check nodes along the source and destination clock paths. Also, check the source and destination clock frequencies to see whether they are the same, or multiples, and whether there are multicycle exceptions on the paths. In some cases, cross-domain paths may also be false by intent, so make sure there are false path exceptions on those.

If you suspect that routing is added to fix real hold problems, then disable the **Optimize hold timing** option. Recompile the design and rerun timing analysis to uncover paths that fail hold time.

Figure 12-15: Optimize Hold Timing Option



Disabling the **Optimize hold timing** option is a debug step, and should be left enabled (default state) during normal compiles. Wire added for hold is a normal part of timing optimization during routing and is not always a problem.

## Evaluate Other Reports and Adjust Settings Accordingly

### Difficulty Packing Design

In the Fitter Resource Section, under the **Resource Usage Summary**, review the **Difficulty Packing Design** report. The **Difficulty Packing Design** report details the effort level (low, medium, or high) of the Fitter to fit the design into the device, partition, and LogicLock region. As the effort level of **Difficulty Packing Design** increases, timing closure gets harder. Going from medium to high can result in significant drop in performance or increase in compile time. Consider reducing logic to reduce packing difficulty.

### Review Ignored Assignments


The **Compilation Report** includes details of any assignments ignored by the Fitter. Assignments typically get ignored if design names change, but assignments are not updated. Make sure any intended assignments are not being ignored.

### Review Non-Default Settings

The reports from Synthesis and Fitter show non-default settings used in a compilation. Review the non-default settings to ensure the design benefits from the change.

### Review Floorplan

Use the Chip Planner for reviewing placement. The Chip Planner can be used to locate hierarchical entities, and colors each located entity in the floorplan. Look for logic that seems out of place, based on where you would expect it to be. For example, logic that interfaces with I/Os should be close to the I/Os, and logic that interfaces with an IP or memory must be close to the IP or memory. The figure shows an example of a floorplan with color-coded entities. In the floorplan, the green block is spread apart. Check to see if those paths are failing timing, and if so, what connects to that module that could affect placement. The blue and aqua blocks are spread out and mixed together. Check and see if there are many connections between the two modules that may contribute to this. The pink logic at the bottom should interface with I/Os at the bottom edge.

**Figure 12-16: Floorplan with Color-Coded Entities**

Check fan-in and fan-out of a highlighted module by using the buttons on the task bar shown in the figure below.

**Figure 12-17: Fan-in and Fan-Out Buttons**

Look for signals that go a long way across the chip and see if they are contributing to timing failures.

Check global signal usage for signals that may affect logic placement. Logic feeding a global buffer may be pulled close to the buffer, away from related logic. High fan-out on non-global resource may pull logic together.

Check for routing congestion. Highly congested areas may cause logic to be spread out, and the design may be difficult to route.

### Evaluate Placement and Routing

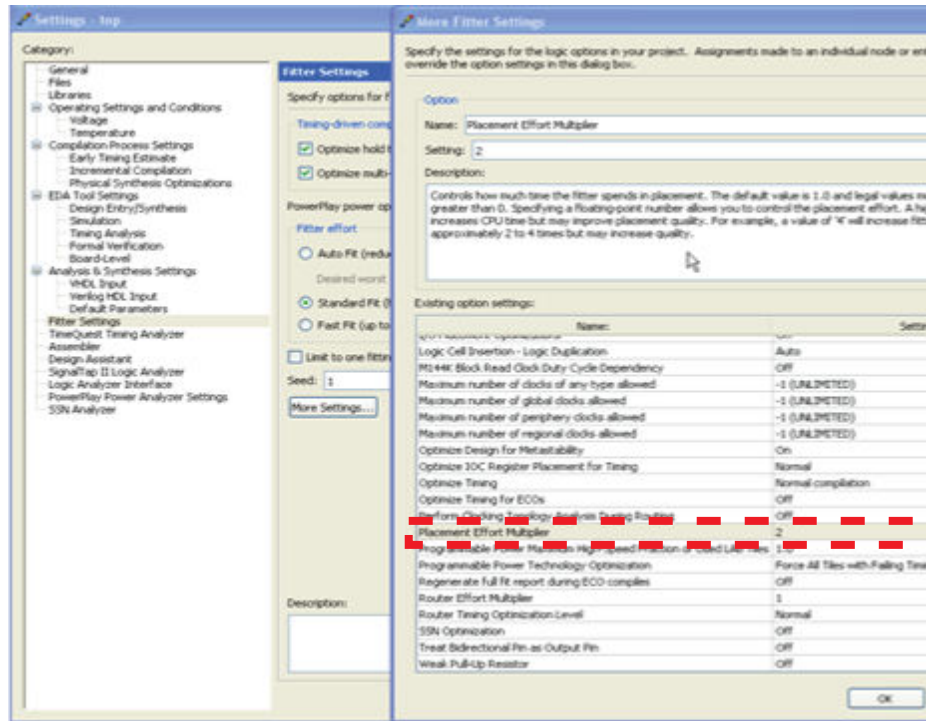
Review duration of parts of compile time in Fitter messages. If routing takes much more time than placement, then meeting timing may be more difficult than the placer predicted.

### Adjust Placement Effort

Increasing the **Placement Effort Multiplier** to improve placement quality may be a good tradeoff at the cost of higher compile time, but the benefit is design dependent. The value should be adjusted after

reviewing and optimizing other settings and RTL. Try an increased value, up to 4, and reset to default if performance or compile time does not improve.

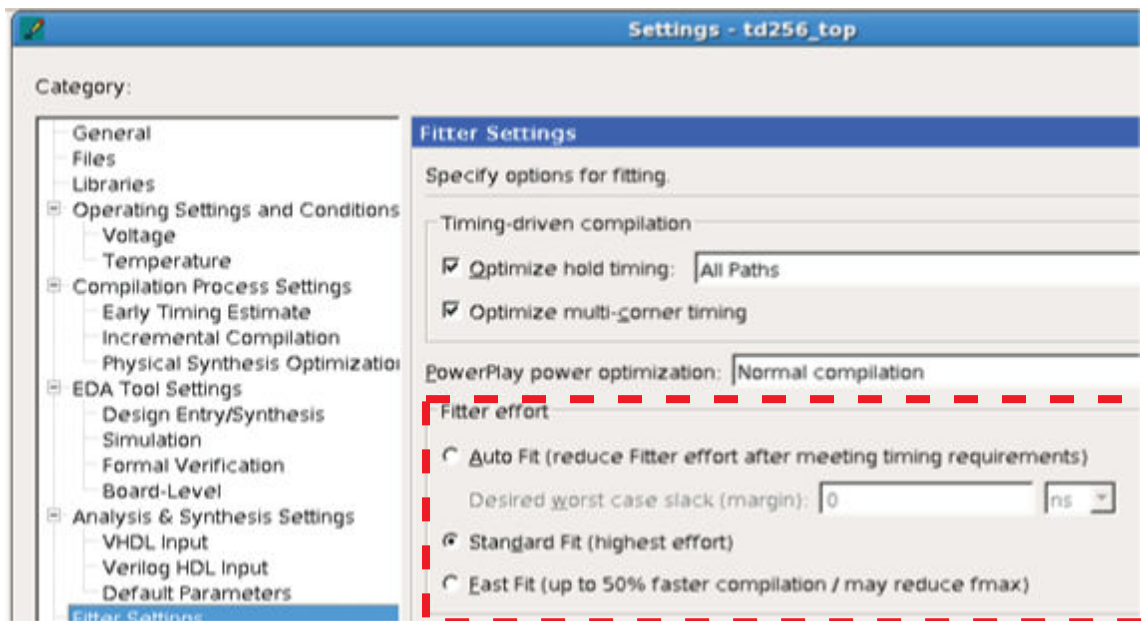
Figure 12-18: Placement Effort Multiplier



### Adjust Fitter Effort

To increase effort, enable the **Standard Fit (highest effort)** option. The default **Auto Fit** option reduces Fitter effort when it estimates timing requirements are met.

Figure 12-19: Fitter Effort



### Review Timing Constraints

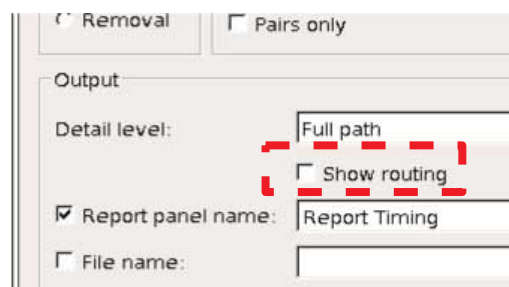
Ensure that clocks are constrained with the correct frequency requirements. Using the `derive_pll_clocks` assignment keeps generated clock settings updated. TimeQuest can be useful in reviewing SDC constraints. For example, under **Diagnostic** in the Task panel, the **Report Ignored Constraints** report shows any incorrect names in the design, most commonly caused by changes in the design hierarchy. Use the **Report Unconstrained Paths** report to locate unconstrained paths. Add constraints as necessary so that the design can be optimized.

## Review Details of Timing Paths

### Show Timing Path Routing

Showing routing for a path can help uncover unusual routing delays. In the TimeQuest Tasks panel, enable the **Report panel name** option, and then select **Report Timing**. Then, turn on the **Show routing** option to show routing wires in the path.

Figure 12-20: Show Routing



The **Extra Fitter Information** tab shows a miniature floorplan with the path highlighted. The path can also be located in the Chip Planner for viewing routing congestion, and to view whether nodes in a path are placed close together or far apart.

## Global Network Buffers

A routing path can be used to identify global network buffers that fail timing. Buffer locations are named according to the network they drive.

- CLK\_CTRL\_Gn—for Global driver
- CLK\_CTRL\_Rn—for Regional driver

Buffers to access the global networks are located in the center of each side of the device. The buffering to route a core logic signal on a global signal network will cause insertion delay. Some trade offs to consider for global and non-global routing are source location, insertion delay, fan-out, distance a signal travels, and possible congestion if the signal is demoted to local routing.


### Source Location

If the register feeding the global buffer cannot be moved closer, then consider changing either the design logic or the routing type.


### Insertion Delay

If a global signal is required, consider adding half a cycle to timing by using a negative-edge triggered register to generate the signal (top figure) and use a multicycle setup constraint (bottom figure).

**Figure 12-21: Negative-Edge Triggered Register**



**Figure 12-22: Multicycle Setup Constraint**



`set_multicycle_path -from <generating register> -setup -end 2`

### Fan-Out

Nodes with very high fan-out that use local routing tend to pull logic that they drive close to the source node. This can make other paths fail timing. Duplicating registers can help reduce the impact of high fan-out paths. Consider manually duplicating and preserving these registers. Using a MAX\_FANOUT assignment



may make arbitrary groups of fan-out nodes, whereas a designer can make more intelligent fan-out groups.

### Global Networks

If a signal should use a different type of global signal than it has automatically been assigned, use the Global Signal assignment to control the global signal usage on a per-signal basis. For example, if local routing is desired, set the Global Signal assignment to **OFF**.

**Figure 12-23: Global Signal Assignment**

To	Assignment Name /	Value	Enabled
reg_clk	Global Signal	Off	Yes

### Resets and Global Networks

Reset signals are often routed on global networks. Sometimes, the use of a global network causes recovery failures. Consider reviewing the placement of the register that generates the reset and the routing path of the signal.

### Suspicious Setup

Suspicious setup failures include paths with very small or very large requirements. One typical cause is math precision error. For example,  $10\text{Mhz}/3 = 33.33$  ns per period. In three cycles, the time would be 99.999 ns vs 100.000 ns. Setting a maximum delay could provide an appropriate setup relationship.

Another cause of failure would be paths that should be false by design intent, such as:

- asynchronous paths that are handled through FIFOs, or
- slow asynchronous paths that rely on handshaking for data that remain available for multiple clock cycles.

To prevent the Fitter from having to meet unnecessarily restrictive timing requirements, consider adding false or multicycle path statements.

### Logic Depth

The **Statistics** tab in the TimeQuest path report shows the levels of logic in a path. If the path fails timing and the number of logic levels is high, consider adding pipelining in that part of the design.

### Auto Shift Register Replacement

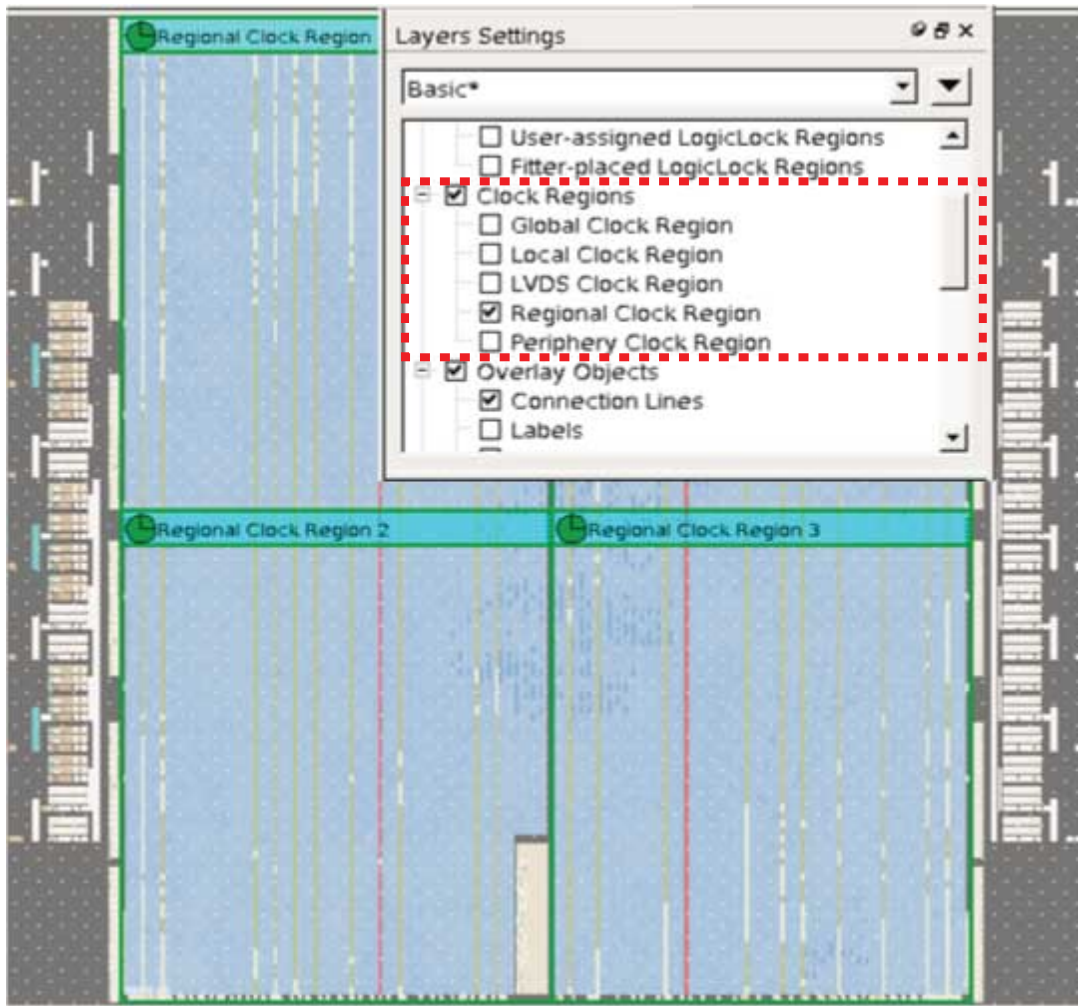
Shift registers or register chains can be converted to RAM during synthesis to save area. However, conversion to RAM often reduces speed. The names of the converted registers will include "altshift\_taps".

If paths that fail timing begin or end in shift registers, consider disabling the **Auto Shift Register Replacement** option. Registers that are intended for pipelining should not be converted. For shift registers that are converted to a chain, evaluate area/speed trade off of implementing in RAM or logic cells. If a design is close to full, shift register conversion to RAM may benefit non-critical clock domains by saving area. The settings can be changed globally or on a register or hierarchy basis from the default of **AUTO** to **OFF**.

### Clocking Architecture

Review the clock region boundaries in the Chip Planner. You must place registers driven by a regional clock in one quadrant of the chip.

Figure 12-24: Clock Regions



Timing failure can occur when the I/O interface at the top of the device connects to logic driven by a regional clock which is in one quadrant of the device, and placement restrictions force long paths to and from some of the I/Os to logic across quadrants.

Use a different type of clock source to drive the logic - global, which covers the whole device, or dual-regional which covers half the device. Alternatively, you can reduce the frequency of the I/O interface to accommodate the long path delays. You can also redesign the pinout of the device to place all the specified I/Os adjacent to the regional clock quadrant. This issue can happen when register locations are restricted, such as with LogicLock regions, clocking resources, or hard blocks (memories, DSPs, IPs). The Extra Fitter Information tab in the TimeQuest report informs you when placement is restricted for nodes in a path.

### Timing Closure Recommendations

The Report Timing Closure Recommendations task in the TimeQuest analyzer analyzes paths and provides specific recommendations based on path characteristics.

## Making Adjustments and Recompiling

Look for obvious problems that you can fix with minimal effort. To identify where the Compiler had trouble meeting timing, perform seed sweeping with about five compiles. Doing so shows consistently failing paths. Consider recoding or redesigning that part of the design.

To reach timing closure, a well written RTL can be more effective than changing your compilation settings. Seed sweeping can also be useful if the timing failure is very small, and the design has already been optimized for performance improvements and is close to final release. Additionally, seed sweeping can be used for evaluating changes to compilation settings. Compilation results vary due to the random nature of fitter algorithms. If a compilation setting change produces lower average performance, undo the change.

Sometimes, settings or constraints can cause more problems than they fix. When significant changes to the RTL or design architecture have been made, compile periodically with default settings and without LogicLock regions, and re-evaluate paths that fail timing.

Partitioning often does not help timing closure, and should be done at the beginning of the design process. Adding partitions can increase logic utilization if it prevents cross-boundary optimizations, making timing closure harder and increasing compile times.

Adding LogicLock regions can be an effective part of timing closure, but must be done at the beginning of a design. Adding new LogicLock regions at the end of the design cycle can restrict placement, hence lowering the performance.

## Scripting Support

You can run procedures and make settings described in this manual in a Tcl script. You can also run some procedures at a command prompt. For detailed information about scripting command options, refer to the Quartus II command-line and Tcl API Help browser. To run the Help browser, type the following command at the command prompt:

```
quartus_sh --qhhelp
```

You can specify many of the options described in this section either in an instance, or at a global level, or both.

Use the following Tcl command to make a global assignment:

```
set_global_assignment -name <.qsf variable name> <value>
```

Use the following Tcl command to make an instance assignment:

```
set_instance_assignment -name <.qsf variable name> <value> -to <instance name>
```

**Note:** If the <value> field includes spaces (for example, 'Standard Fit'), you must enclose the value in straight double quotation marks.

### Related Information

[Tcl Scripting documentation](#) on page 3-1

[Quartus II Settings Reference File Manual](#)

[Command-Line Scripting documentation](#) on page 2-1

## Initial Compilation Settings

Use the Quartus II Settings File (.qsf) variable name in the Tcl assignment to make the setting along with the appropriate value. The **Type** column indicates whether the setting is supported as a global setting, an instance setting, or both.

The top table lists the .qsf variable name and applicable values for the settings described in the “Initial Compilation: Required Settings” section in the [Design Optimization Overview](#) chapter in the *Quartus II Handbook*. The bottom table lists the advanced compilation settings.

**Table 12-4: Initial Compilation Settings**

Setting Name	.qsf File Variable Name	Values	Type
Optimize IOC Register Placement For Timing	OPTIMIZE_IOC_REGISTER_PLACEMENT_FOR_TIMING	ON, OFF	Global
Optimize Hold Timing	OPTIMIZE_HOLD_TIMING	OFF, IO PATHS AND MINIMUM TPD PATHS, ALL PATHS	Global

**Table 12-5: Advanced Compilation Settings**

Setting Name	.qsf File Variable Name	Values	Type
Router Timing Optimization level	ROUTER_TIMING_OPTIMIZATION_LEVEL	NORMAL, MINIMUM, MAXIMUM	Global

## Resource Utilization Optimization Techniques (LUT-Based Devices)

The table lists the .qsf file variable name and applicable values for the settings described in [Optimizing Timing \(LUT-Based Devices\)](#).

**Table 12-6: Resource Utilization Optimization Settings**

Setting Name	.qsf File Variable Name	Values	Type
Auto Packed Registers (1)	AUTO_PACKED_REGISTERS_<device family name>	OFF, NORMAL, MINIMIZE AREA, MINIMIZE AREA WITH CHAINS, AUTO	Global, Instance
Perform WYSIWYG Primitive Resynthesis	ADV_NETLIST_OPT_SYNTH_WYSIWYG_REMAP	ON, OFF	Global, Instance

Setting Name	.qsf File Variable Name	Values	Type
Physical Synthesis for Combinational Logic for Reducing Area	PHYSICAL_SYNTHESIS_COMBO_LOGIC_FOR_AREA	ON, OFF	Global, Instance
Physical Synthesis for Mapping Logic to Memory	PHYSICAL_SYNTHESIS_MAP_LOGIC_TO_MEMORY_FOR_AREA	ON, OFF	Global, Instance
Optimization Technique	<device family name>_OPTIMIZATION_TECHNIQUE	AREA, SPEED, BALANCED	Global, Instance
Speed Optimization Technique for Clock Domains	SYNTH_CRITICAL_CLOCK	ON, OFF	Instance
State Machine Encoding	STATE_MACHINE_PROCESSING	AUTO, ONE-HOT, GRAY, JOHNSON, MINIMAL BITS, SEQUENTIAL, USER-ENCODE	Global, Instance
Auto RAM Replacement	AUTO_RAM_RECOGNITION	ON, OFF	Global, Instance
Auto ROM Replacement	AUTO_ROM_RECOGNITION	ON, OFF	Global, Instance
Auto Shift Register Replacement	AUTO_SHIFT_REGISTER_RECOGNITION	ON, OFF	Global, Instance
Auto Block Replacement	AUTO_DSP_RECOGNITION	ON, OFF	Global, Instance
Number of Processors for Parallel Compilation	NUM_PARALLEL_PROCESSORS	Integer between 1 and 16 inclusive, or ALL	Global

Note to table :

1. Allowed values for this setting depend on the device family that you select.

## I/O Timing Optimization Techniques (LUT-Based Devices)

The table lists the .qsf file variable name and applicable values for the I/O timing optimization settings.

**Table 12-7: I/O Timing Optimization Settings**

Setting Name	.qsf File Variable Name	Values	Type
Optimize IOC Register Placement For Timing	OPTIMIZE_IOC_REGISTER_PLACEMENT_FOR_TIMING	ON, OFF	Global
Fast Input Register	FAST_INPUT_REGISTER	ON, OFF	Instance
Fast Output Register	FAST_OUTPUT_REGISTER	ON, OFF	Instance
Fast Output Enable Register	FAST_OUTPUT_ENABLE_REGISTER	ON, OFF	Instance
Fast OCT Register	FAST_OCT_REGISTER	ON, OFF	Instance

## Register-to-Register Timing Optimization Techniques (LUT-Based Devices)

The table lists the .qsf file variable name and applicable values for the settings described in [Register-to-Register Timing Optimization Techniques \(LUT-Based Devices\)](#).

**Table 12-8: Register-to-Register Timing Optimization Settings**

Setting Name	.qsf File Variable Name	Values	Type
Perform WYSIWYG Primitive Resynthesis	ADV_NETLIST_OPT_SYNTH_WYSIWYG_REMAP	ON, OFF	Global, Instance
Perform Physical Synthesis for Combinational Logic	PHYSICAL_SYNTHESIS_COMBO_LOGIC	ON, OFF	Global, Instance
Perform Register Duplication	PHYSICAL_SYNTHESIS_REGISTER_DUPLICATION	ON, OFF	Global, Instance
Perform Register Retiming	PHYSICAL_SYNTHESIS_REGISTER_RETIMING	ON, OFF	Global, Instance
Perform Automatic Asynchronous Signal Pipelining	PHYSICAL_SYNTHESIS_ASYNCHRONOUS_SIGNAL_PIPELINING	ON, OFF	Global, Instance
Physical Synthesis Effort	PHYSICAL_SYNTHESIS_EFFORT	NORMAL, EXTRA, FAST	Global
Fitter Seed	SEED	<integer>	Global
Maximum Fan-Out	MAX_FANOUT	<integer>	Instance
Manual Logic Duplication	DUPLICATE_ATOM	<node name>	Instance
Optimize Power during Synthesis	OPTIMIZE_POWER_DURING_SYNTHESIS	NORMAL, OFF EXTRA_EFFORT	Global

Setting Name	.qsf File Variable Name	Values	Type
Optimize Power during Fitting	OPTIMIZE_POWER_DURING_FITTING	NORMAL, OFF EXTRA_EFFORT	Global

## Document Revision History

Table 12-9: Document Revision History

Date	Version	Changes
2014.12.15	14.1.0	<ul style="list-style-type: none"> <li>Updated location of Fitter Settings, Analysis &amp; Synthesis Settings, and Physical Synthesis Optimizations to Compiler Settings.</li> <li>Updated DSE II content.</li> </ul>
June 2014	14.0.0	<ul style="list-style-type: none"> <li>Dita conversion.</li> <li>Removed content about obsolete devices that are no longer supported in QII software v14.0: Arria GX, Arria II, Cyclone III, Stratix II, Stratix III.</li> <li>Replaced Megafunction content with IP core content.</li> </ul>
November 2013	13.1.0	<ul style="list-style-type: none"> <li>Added Design Evaluation for Timing Closure section.</li> <li>Removed Optimizing Timing (Macrocell-Based CPLDs) section.</li> <li>Updated Optimize Multi-Corner Timing and Fitter Aggressive Routability Optimization.</li> <li>Updated Timing Analysis with the TimeQuest Timing Analyzer to show how to access the <b>Report All Summaries</b> command.</li> <li>Updated Ignored Timing Constraints to include a help link to <i>Fitter Summary Reports</i> with the <b>Ignored Assignment Report</b> information.</li> </ul>
May 2013	13.0.0	<ul style="list-style-type: none"> <li>Renamed chapter title from Area and Timing Optimization to Timing Closure and Optimization.</li> <li>Removed design and area/resources optimization information.</li> <li>Added the following sections: <ul style="list-style-type: none"> <li>Fitter Aggressive Routability Optimization.</li> <li>Tips for Analyzing Paths from/to the Source and Destination of Critical Path.</li> <li>Tips for Locating Multiple Paths to the Chip Planner.</li> <li>Tips for Creating a .tcl Script to Monitor Critical Paths Across Compiles.</li> </ul> </li> </ul>

Date	Version	Changes
November 2012	12.1.0	<ul style="list-style-type: none"> <li>Updated “Initial Compilation: Optional Fitter Settings” on page 13–2, “I/O Assignments” on page 13–2, “Initial Compilation: Optional Fitter Settings” on page 13–2, “Resource Utilization” on page 13–9, “Routing” on page 13–21, and “Resolving Resource Utilization Problems” on page 13–43.</li> </ul>
June 2012	12.0.0	<ul style="list-style-type: none"> <li>Updated “Optimize Multi-Corner Timing” on page 13–6, “Resource Utilization” on page 13–10, “Timing Analysis with the TimeQuest Timing Analyzer” on page 13–12, “Using the Resource Optimization Advisor” on page 13–15, “Increase Placement Effort Multiplier” on page 13–22, “Increase Router Effort Multiplier” on page 13–22 and “Debugging Timing Failures in the TimeQuest Analyzer” on page 13–24.</li> <li>Minor text edits throughout the chapter.</li> </ul>
November 2011	11.1.0	<ul style="list-style-type: none"> <li>Updated the “Timing Requirement Settings”, “Standard Fit”, “Fast Fit”, “Optimize Multi-Corner Timing”, “Timing Analysis with the TimeQuest Timing Analyzer”, “Debugging Timing Failures in the TimeQuest Analyzer”, “LogicLock Assignments”, “Tips for Analyzing Failing Clock Paths that Cross Clock Domains”, “Flatten the Hierarchy During Synthesis”, “Fast Input, Output, and Output Enable Registers”, and “Hierarchy Assignments” sections</li> <li>Updated Table 13–6</li> <li>Added the “Spine Clock Limitations” section</li> <li>Removed the Change State Machine Encoding section from page 19</li> <li>Removed Figure 13-5</li> <li>Minor text edits throughout the chapter</li> </ul>
May 2011	11.0.0	<ul style="list-style-type: none"> <li>Reorganized sections in “Initial Compilation: Optional Fitter Settings” section</li> <li>Added new information to “Resource Utilization” section</li> <li>Added new information to “Duplicate Logic for Fan-Out Control” section</li> <li>Added links to Help</li> <li>Additional edits and updates throughout chapter</li> </ul>
December 2010	10.1.0	<ul style="list-style-type: none"> <li>Added links to Help</li> <li>Updated device support</li> <li>Added “Debugging Timing Failures in the TimeQuest Analyzer” section</li> <li>Removed Classic Timing Analyzer references</li> <li>Other updates throughout chapter</li> </ul>
August 2010	10.0.1	Corrected link



Date	Version	Changes
July 2010	10.0.0	<ul style="list-style-type: none"><li>• Moved Compilation Time Optimization Techniques section to new <i>Reducing Compilation Time</i> chapter</li><li>• Removed references to Timing Closure Floorplan</li><li>• Moved Smart Compilation Setting and Early Timing Estimation sections to new <i>Reducing Compilation Time</i> chapter</li><li>• Added Other Optimization Resources section</li><li>• Removed outdated information</li><li>• Changed references to DSE chapter to Help links</li><li>• Linked to Help where appropriate</li><li>• Removed Referenced Documents section</li></ul>

**Related Information**[Quartus II Handbook Archive](#)

For previous versions of the Quartus II Handbook

2014.12.15

QI15V2



Subscribe



Send Feedback

## Power Optimization

The Quartus® II software offers power-driven compilation to fully optimize device power consumption. Power-driven compilation focuses on reducing your design's total power consumption using power-driven synthesis and power-driven place-and-route.

This chapter describes the power-driven compilation feature and flow in detail, as well as low power design techniques that can further reduce power consumption in your design. The techniques primarily target Arria®, Stratix®, and Cyclone® series of devices. These devices utilize a low-k dielectric material that dramatically reduces dynamic power and improves performance. Arria series, Stratix IV, and Stratix V device families include efficient logic structures called adaptive logic modules (ALMs) that obtain maximum performance while minimizing power consumption. Cyclone device families offer the optimal blend of high performance and low power in a low-cost FPGA.

Altera provides the Quartus II PowerPlay Power Analyzer to aid you during the design process by delivering fast and accurate estimations of power consumption. You can minimize power consumption, while taking advantage of the industry's leading FPGA performance, by using the tools and techniques described in this chapter.

Total FPGA power consumption is comprised of I/O power, core static power, and core dynamic power. This chapter focuses on design optimization options and techniques that help reduce core dynamic power and I/O power. In addition to these techniques, there are additional power optimization techniques available for specific devices. These techniques include:

- Programmable Power Technology
- Device Speed Grade Selection

### Related Information

- **Literature and Technical Documentation**

For more information about a device-specific architecture, refer to the device handbook on the Altera website.

- **PowerPlay Power Analysis**

For more information about the PowerPlay Power Analyzer, refer to volume 3 of the *Quartus II Handbook*.

- **AN 514: Power Optimization in Stratix IV FPGAs**

For more information about power optimization techniques available for Stratix IV devices.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered




## Power Dissipation

You can refine techniques that reduce power consumption in your design by understanding the sources of power dissipation.

The following figure shows the power dissipation of Stratix and Cyclone devices in different designs. All designs were analyzed at a fixed clock rate of 100 MHz and exhibited varied logic resource utilization across available resources.

**Figure 13-1: Average Core Dynamic Power Dissipation**



### Notes:

1. 103 different designs were used to obtain these results.
2. 96 different designs were used to obtain these results.
3. In designs using DSP blocks, DSPs consumed 5% of core dynamic power.

In Stratix and Cyclone device families, a series of column and row interconnect wires of varying lengths provide signal interconnections between logic array blocks (LABs), memory block structures, and digital signal processing (DSP) blocks or multiplier blocks. These interconnects dissipate the largest component of device power.

FPGA combinational logic is another source of power consumption. The basic building block of logic in the latest Stratix series devices is the ALM, and in Cyclone IV GX devices, it is the logic element (LE).

For more information about ALMs and LEs in Cyclone or Stratix devices, refer to the respective device handbook.

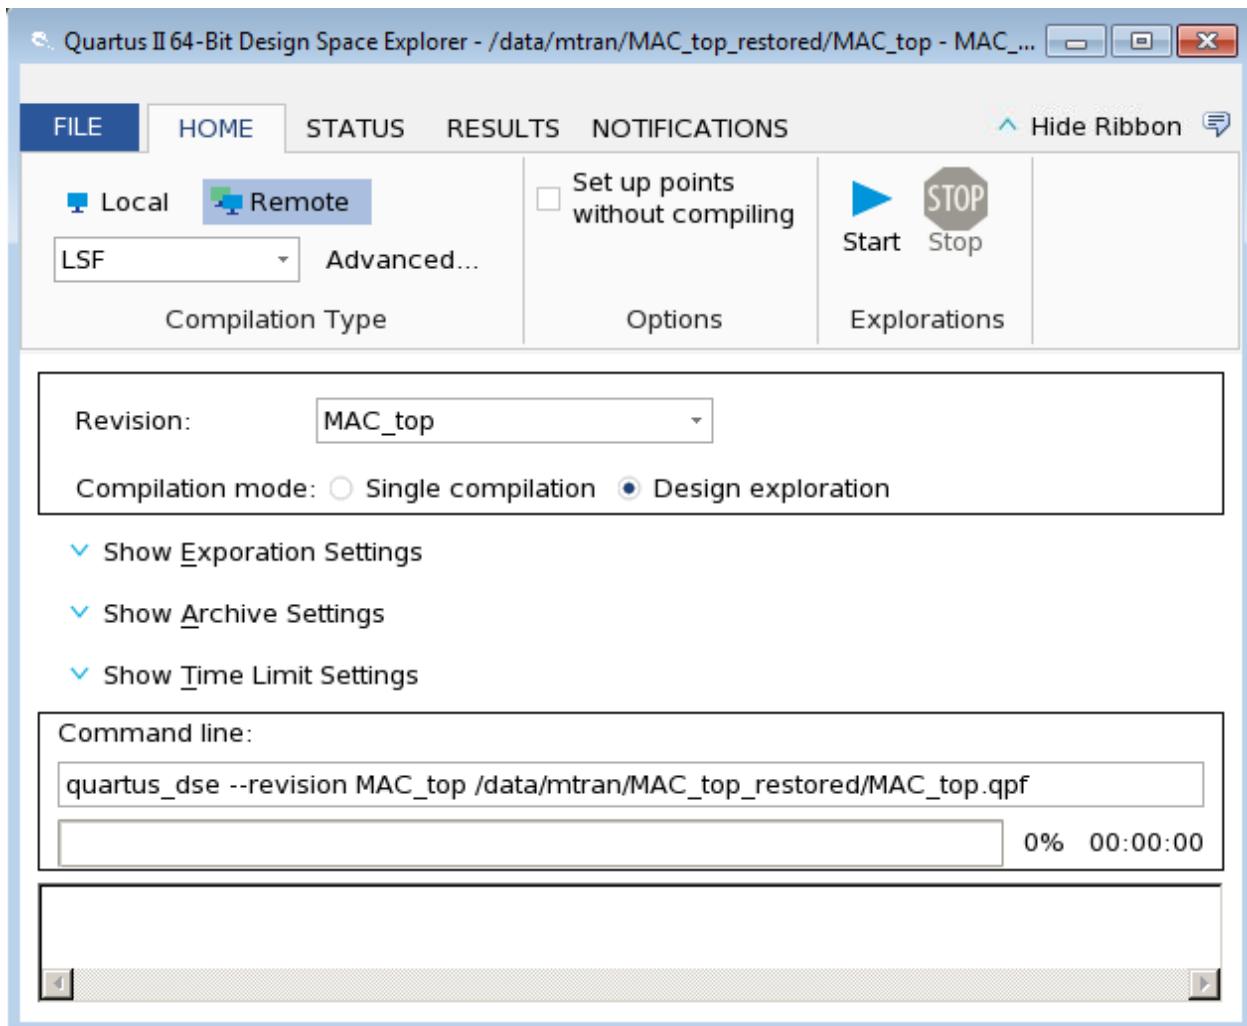
Memory and clock resources are other major consumers of power in FPGAs. Stratix devices feature the TriMatrix memory architecture. TriMatrix memory includes 512-bit M512 blocks, 4-Kbit M4K blocks, and 512-Kbit M-RAM blocks, which are configurable to support many features. Stratix IV TriMatrix on-chip memory is an enhancement based upon the Stratix II FPGA TriMatrix memory and includes three sizes of memory blocks: MLAB blocks, M9K blocks, and M144K blocks. Stratix IV and Stratix V devices feature Programmable Power Technology, an advanced architecture that enables a smooth trade-off between speed and power. The core of each Stratix IV and Stratix V device is divided into tiles, each of which may be put into a high-speed or low-power mode. The primary benefit of Programmable Power

Technology is to reduce static power, with a secondary benefit being a small reduction in dynamic power. Cyclone IV GX devices have 9-Kbit M9K memory blocks.

## Design Space Explorer II

Design Space Explorer II (DSE) is a simple, easy-to-use, design optimization utility that is included in the Quartus II software. DSE II explores and reports optimal Quartus II software options for your design, targeting either power optimization, design performance, or area utilization improvements. You can use DSE II to implement the techniques described in this chapter.

Figure 13-2: Design Space Explorer II User Interface



The power optimizations, under **Explore for**, target overall design power improvements. These settings focus on applying different options that specifically reduce total design thermal power.

By default, the Quartus II PowerPlay Power Analyzer is run for every exploration performed by DSE II when power optimizations are selected. This helps you debug your design and determine trade-offs between power requirements and performance optimization.

**Related Information****About Design Space Explorer II**

For more information about the DSE II, refer to Quartus II Help.

## Power-Driven Compilation

The standard Quartus II compilation flow consists of Analysis and Synthesis, placement and routing, Assembly, and Timing Analysis. Power-driven compilation takes place at the Analysis and Synthesis and Place-and-Route stages.

Quartus II software settings that control power-driven compilation are located in the **PowerPlay power optimization during synthesis** list in the **Advanced Settings (Synthesis)** dialog box, and the **PowerPlay power optimization during fitting** list on the **Advanced Fitter Settings** dialog box. The following sections describes these power optimization options at the Analysis and Synthesis and Fitter levels.

## Power-Driven Synthesis

Synthesis netlist optimization occurs during the synthesis stage of the compilation flow. The optimization technique makes changes to the synthesis netlist to optimize your design according to the selection of area, speed, or power optimization. This section describes power optimization techniques at the synthesis level.

To access the PowerPlay Power Optimization During Synthesis option, click **Assignments > Settings > Compiler Settings > Advanced Settings (Synthesis)**.

You can apply these settings on a project or entity level.

**Table 13-1: Optimize Power During Synthesis Options**

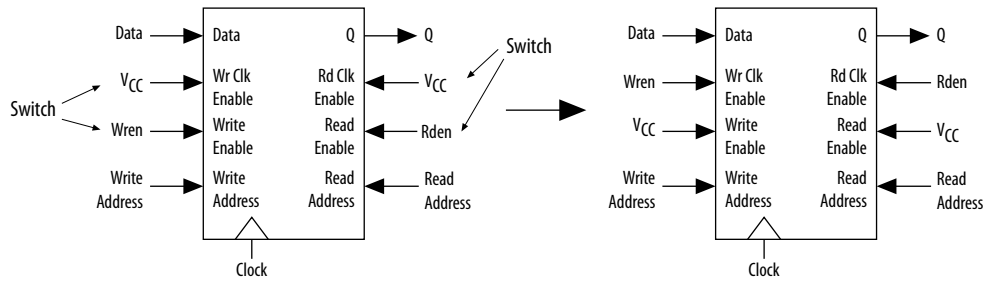
Settings	Description
Off	No netlist, placement, or routing optimizations are performed to minimize power.
Normal compilation (Default)	Low compute effort algorithms are applied to minimize power through netlist optimizations as long as they are not expected to reduce design performance.
Extra effort	High compute effort algorithms are applied to minimize power through netlist optimizations. Max performance might be impacted.

The **Normal compilation** setting is turned on by default. This setting performs memory optimization and power-aware logic mapping during synthesis.

Memory blocks can represent a large fraction of total design dynamic power. Minimizing the number of memory blocks accessed during each clock cycle can significantly reduce memory power. Memory optimization involves effective movement of user-defined read/write enable signals to associated read-and-write clock enable signals for all memory types.

A default implementation of a simple dual-port memory block in which write-clock enable signals and read-clock enable signals are connected to  $V_{CC}$ , making both read and write memory ports active during each clock cycle.

Figure 13-3: Memory Transformation



Memory transformation effectively moves the read-enable and write-enable signals to the respective read-clock enable and write-clock enable signals. By using this technique, memory ports are shut down when they are not accessed. This significantly reduces your design’s memory power consumption. For Stratix IV and Stratix V devices, the memory transformation takes place at the Fitter level by selecting the **Normal compilation** settings for the power optimization option.

In Cyclone IV GX and Stratix IV devices, the specified read-during-write behavior can significantly impact the power of single-port and bidirectional dual-port RAMs. It is best to set the read-during-write parameter to “Don’t care” (at the HDL level), as it allows an optimization whereby the `read-enable` signal can be set to the inversion of the existing `write-enable` signal (if one exists). This allows the core of the RAM to shut down (that is, not toggle), which saves a significant amount of power.

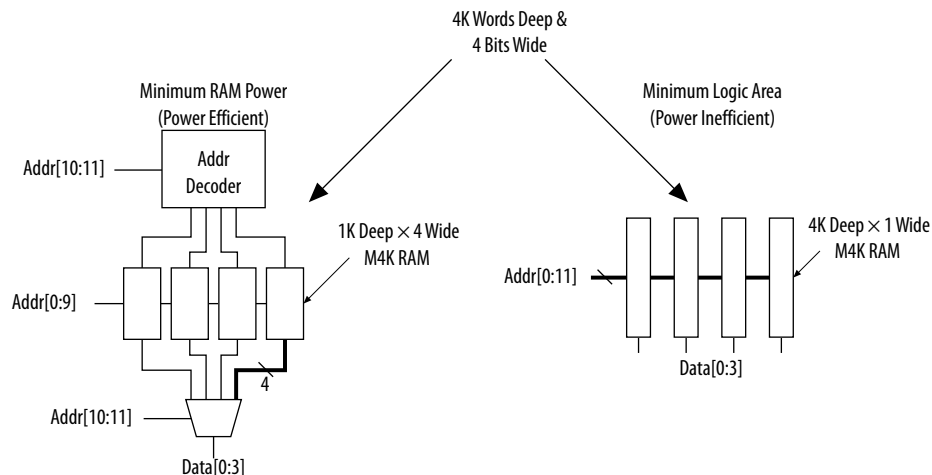
The other type of power optimization that takes place with the **Normal compilation** setting is power-aware logic mapping. The power-aware logic mapping reduces power by rearranging the logic during synthesis to eliminate nets with high toggle rates.

The **Extra effort** setting performs the functions of the **Normal compilation** setting and other memory optimizations to further reduce memory power by shutting down memory blocks that are not accessed. This level of memory optimization can require extra logic, which can reduce design performance.

The **Extra effort** setting also performs power-aware memory balancing. Power-aware memory balancing automatically chooses the best memory configuration for your memory implementation and provides optimal power saving by determining the number of memory blocks, decoder, and multiplexer circuits required. If you have not previously specified target-embedded memory blocks for your design’s memory functions, the power-aware balancer automatically selects them during memory implementation.

The following figure is an example of a  $4k \times 4$  (4k deep and 4 bits wide) memory implementation in two different configurations using M4K memory blocks available in some Stratix devices.

Figure 13-4: 4K × 4 Memory Implementation Using Multiple M4K Blocks



The minimum logic area implementation uses M4K blocks configured as  $4k \times 1$ . This implementation is the default in the Quartus II software because it has the minimum logic area (0 logic cells) and the highest speed. However, all four M4K blocks are active on each memory access in this implementation, which increases RAM power. The minimum RAM power implementation is created by selecting **Extra effort** in the **PowerPlay power optimization** list. This implementation automatically uses four M4K blocks configured as  $1k \times 4$  for optimal power saving. An address decoder is implemented by the RAM megafunction to select which of the four M4K blocks should be activated on a given cycle, based on the state of the top two user address bits. The RAM megafunction automatically implements a multiplexer to feed the downstream logic by choosing the appropriate M4K output. This implementation reduces RAM power because only one M4K block is active on any cycle, but it requires extra logic cells, costing logic area and potentially impacting design performance.

There is a trade-off between power saved by accessing fewer memories and power consumed by the extra decoder and multiplexor logic. The Quartus II software automatically balances the power savings against the costs to choose the lowest power configuration for each logical RAM. The benchmark data shows that the power-driven synthesis can reduce memory power consumption by as much as 60% in Stratix devices.

Memory optimization options can also be controlled by the `Low_Power_Mode` parameter in the **Default Parameters** page of the **Settings** dialog box. The settings for this parameter are **None**, **Auto**, and **ALL**. **None** corresponds to the **Off** setting in the **PowerPlay power optimization** list. **Auto** corresponds to the **Normal compilation** setting and **ALL** corresponds to the **Extra effort** setting, respectively. You can apply PowerPlay power optimization either on a compiler basis or on individual entities. The `Low_Power_Mode` parameter always takes precedence over the **Optimize Power for Synthesis** option for power optimization on memory.

You can also set the `MAXIMUM_DEPTH` parameter manually to configure the memory for low power optimization. This technique is the same as the power-aware memory balancer, but it is manual rather than automatic like the **Extra effort** setting in the **PowerPlay power optimization** list. You can set the `MAXIMUM_DEPTH` parameter for memory modules manually in the megafunction instantiation or in the IP Catalog for power optimization. The `MAXIMUM_DEPTH` parameter always takes precedence over the **Optimize Power for Synthesis** options for power optimization on memory optimization.

#### Related Information

- [Reducing Memory Power Consumption](#) on page 13-11  
For more information about clock enable signals.

- **Running a Power-Optimized Compilation**

For step-by-step instructions on how to perform power-driven synthesis, refer to Quartus II Help.

## Power-Driven Fitter

The **Compiler Settings** page provides access to **PowerPlay power optimization** settings.

You can apply these settings only on a project-wide basis. The **Extra effort** setting for the Fitter requires extensive effort to optimize the design for power and can increase the compilation time.

**Table 13-2: Power-Driven Fitter Option**

Settings	Description
Off	No netlist, placement, or routing optimizations are performed to minimize power.
Normal compilation (Default)	Low compute effort algorithms are applied to minimize power through placement and routing optimizations as long as they are not expected to reduce design performance.
Extra effort	High compute effort algorithms are applied to minimize power through placement and routing optimizations. Max performance might be impacted.

The **Normal compilation** setting is selected by default and performs DSP optimization by creating power-efficient DSP block configurations for your DSP functions. For Stratix IV and Stratix V devices, this setting, which is based on timing constraints entered for the design, enables the Programmable Power Technology to configure tiles as high-speed mode or low-power mode. Programmable Power Technology is always turned **ON** even when the **OFF** setting is selected for the **PowerPlay power optimization** option. Tiles are the combination of LAB and MLAB pairs (including the adjacent routing associated with LAB and MLAB), which can be configured to operate in high-speed or low-power mode. This level of power optimization does not have any affect on the fitting, timing results, or compile time.

The **Extra effort** setting performs the functions of the **Normal compilation** setting and other place-and-route optimizations during fitting to fully optimize the design for power. The Fitter applies an extra effort to minimize power even after timing requirements have been met by effectively moving the logic closer during placement to localize high-toggling nets, and using routes with low capacitance. However, this effort can increase the compilation time.

The **Extra effort** setting uses a Value Change Dump File (.vcd) that guides the Fitter to fully optimize the design for power, based on the signal activity of the design. The best power optimization during fitting results from using the most accurate signal activity information. Signal activities from full post-fit netlist (timing) simulation provide the highest accuracy because all node activities reflect the actual design behavior, provided that supplied input vectors are representative of typical design operation. If you do not have a .vcd file, the Quartus II software uses assignments, clock assignments, and vectorless estimation values (PowerPlay Power Analyzer Tool settings) to estimate the signal activities. This information is used to optimize your design for power during fitting. The benchmark data shows that the power-driven Fitter technique can reduce power consumption by as much as 19% in Stratix devices. On average, you can reduce core dynamic power by 16% with the Extra effort synthesis and Extra effort fitting settings, as compared to the **Off** settings in both synthesis and Fitter options for power-driven compilation.

**Note:** Only the **Extra effort** setting in the **PowerPlay power optimization** list for the Fitter option uses the signal activities (from .vcd files) during fitting. The settings made in the **PowerPlay Power Analyzer Settings** page in the **Settings** dialog box are used to calculate the signal activity of your design.



### Related Information

- **AN 514: Power Optimization in Stratix IV FPGAs**  
For more information about Stratix IV power optimization.
- **PowerPlay Power Analysis**  
For more information about `.vcd` files and how to create them, refer to the *Quartus II Handbook*.
- **Running a Power-Optimized Compilation**  
For step-by-step instructions on how to perform power-driven fitting, refer to Quartus II Help.

## Area-Driven Synthesis

Using area optimization rather than timing or delay optimization during synthesis saves power because you use fewer logic blocks. Using less logic usually means less switching activity. The Quartus II integrated synthesis tool provides **Speed**, **Balanced**, or **Area** for the **Optimization Technique** option. You can also specify this logic option for specific modules in your design with the Assignment Editor in cases where you want to reduce area using the **Area** setting (potentially at the expense of register-to-register timing performance) while leaving the default **Optimization Technique** setting at **Balanced** (for the best trade-off between area and speed for certain device families). The **Speed Optimization Technique** can increase the resource usage of your design if the constraints are too aggressive, and can also result in increased power consumption.


The benchmark data shows that the area-driven technique can reduce power consumption by as much as 31% in Stratix devices and as much as 15% in Cyclone devices.

## Gate-Level Register Retiming

You can also use gate-level register retiming to reduce circuit switching activity. Retiming shuffles registers across combinational blocks without changing design functionality. The **Perform gate-level register retiming** option in the Quartus II software enables the movement of registers across combinational logic to balance timing, allowing the software to trade off the delay between timing critical and noncritical timing paths.

Retiming uses fewer registers than pipelining. In this example of gate-level register retiming, the 10 ns critical delay is reduced by moving the register relative to the combinational logic, resulting in the reduction of data depth and switching activity.

Figure 13-5: Gate-Level Register Retiming



**Note:** Gate-level register retiming makes changes at the gate level. If you are using an atom netlist from a third-party synthesis tool, you must also select the **Perform WYSIWYG primitive resynthesis** option to undo the atom primitives to gates mapping (so that register retiming can be performed), and then to remap gates to Altera primitives. When using Quartus II integrated synthesis, retiming occurs during synthesis before the design is mapped to Altera primitives. The benchmark data shows that the combination of WYSIWYG remapping and gate-level register retiming techniques can reduce power consumption by as much as 6% in Stratix devices and as much as 21% in Cyclone devices.

#### Related Information

- [Netlist Optimizations and Physical Synthesis](#) on page 16-1  
For more information about register retiming, refer to the *Quartus II Handbook*.

## Design Guidelines

Several low-power design techniques can reduce power consumption when applied during FPGA design implementation. This section provides detailed design techniques for Cyclone IV GX devices that affect overall design power. The results of these techniques might be different from design to design.


### Clock Power Management

Clocks represent a significant portion of dynamic power consumption due to their high switching activity and long paths. Actual clock-related power consumption is higher than this because the power consumed by local clock distribution within logic, memory, and DSP or multiplier blocks is included in the power consumption for the respective blocks.

Clock routing power is automatically optimized by the Quartus II software, which enables only those portions of the clock network that are required to feed downstream registers. Power can be further reduced by gating clocks when they are not required. It is possible to build clock-gating logic, but this approach is not recommended because it is difficult to generate a glitch free clock in FPGAs using ALMs or LEs.

Cyclone IV, Stratix IV, and Stratix V devices use clock control blocks that include an enable signal. A clock control block is a clock buffer that lets you dynamically enable or disable the clock network and dynamically switch between multiple sources to drive the clock network. You can use the Quartus II IP Catalog to create this clock control block with the ALTCLKCTRL megafunction. Cyclone IV, Stratix IV, and Stratix V devices provide clock control blocks for global clock networks. In addition, Stratix IV, and Stratix V devices have clock control blocks for regional clock networks. The dynamic clock enable feature lets internal logic control the clock network. When a clock network is powered down, all the logic fed by that clock network does not toggle, thereby reducing the overall power consumption of the device. For example, the following shows a 4-input clock control block diagram.


**Figure 13-6: Clock Control Block Diagram**



The enable signal is applied to the clock signal before being distributed to global routing. Therefore, the enable signal can either have a significant timing slack (at least as large as the global routing delay) or it can reduce the  $f_{MAX}$  of the clock signal.

Another contributor to clock power consumption is the LAB clock that distributes a clock to the registers within a LAB. LAB clock power can be the dominant contributor to overall clock power. For example, in Cyclone devices, each LAB can use two clocks and two clock enable signals, as shown in the following figure. Each LAB's clock signal and clock enable signal are linked. For example, an LE in a particular LAB using the `labclk1` signal also uses the `labclkena1` signal.

Figure 13-7: LAB-Wide Control Signals



To reduce LAB-wide clock power consumption without disabling the entire clock tree, use the LAB-wide clock enable to gate the LAB-wide clock. The Quartus II software automatically promotes register-level clock enable signals to the LAB-level. All registers within an LAB that share a common clock and clock enable are controlled by a shared gated clock. To take advantage of these clock enables, use a clock enable construct in the relevant HDL code for the registered logic.

#### Related Information

#### [Clock Control Block Megafunction User Guide \(ALTCLKCTRL\)](#)

For more information about using clock control blocks.

### LAB-Wide Clock Enable Example

This VHDL code makes use of a LAB-wide clock enable. This clock-gating logic is automatically turned into an LAB-level clock enable signal.

```
IF clk'event AND clock = '1' THEN
  IF logic_is_enabled = '1' THEN
    reg <= value;
  ELSE
    reg <= reg;
  END IF;
END IF;
```

## Reducing Memory Power Consumption

The memory blocks in FPGA devices can represent a large fraction of typical core dynamic power. Memory consumes approximately 20% of the core dynamic power in typical some device designs. Memory blocks are unlike most other blocks in the device because most of their power is tied to the clock rate, and is insensitive to the toggle rate on the data and address lines.

### Reducing Memory Power Consumption


The memory blocks in FPGA devices can represent a large fraction of typical core dynamic power. Memory consumes approximately 20% of the core dynamic power in typical some device designs. Memory blocks are unlike most other blocks in the device because most of their power is tied to the clock rate, and is insensitive to the toggle rate on the data and address lines.

When a memory block is clocked, there is a sequence of timed events that occur within the block to execute a read or write. The circuitry controlled by the clock consumes the same amount of power regardless of whether or not the address or data has changed from one cycle to the next. Thus, the toggle rate of input data and the address bus have no impact on memory power consumption.

The key to reducing memory power consumption is to reduce the number of memory clocking events. You can achieve this through clock network-wide gating, or on a per-memory basis through use of the clock enable signals on the memory ports.

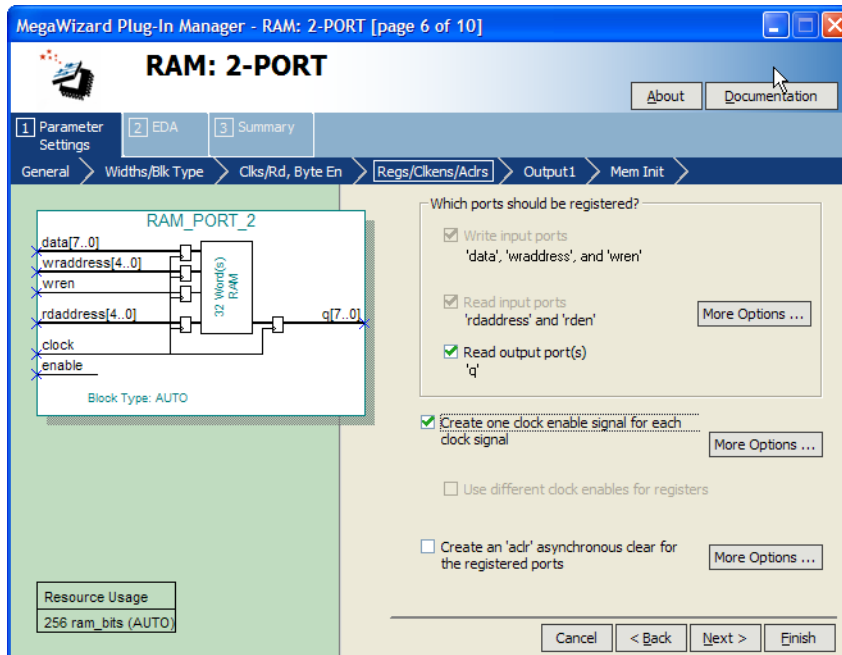
The logical view of the internal clock of the memory block. Use the appropriate enable signals on the memory to make use of the clock enable signal instead of gating the clock.

**Figure 13-8: Memory Clock Enable Signal**



Using the clock enable signal enables the memory only when necessary and shuts it down for the rest of the time, reducing the overall memory power consumption. You can create these enable signals by selecting the **Clock enable signal** option for the appropriate port when generating the memory block function.

Figure 13-9: RAM 2-Port Clock Enable Signal Selectable Option




For example, consider a design that contains a 32-bit-wide M4K memory block in ROM mode that is running at 200 MHz. Assuming that the output of this block is only required approximately every four cycles, this memory block will consume 8.45 mW of dynamic power according to the demands of the downstream logic. By adding a small amount of control logic to generate a read clock enable signal for the memory block only on the relevant cycles, the power can be cut 75% to 2.15 mW.

You can also use the `MAXIMUM_DEPTH` parameter in your memory megafunction to save power in Cyclone IV GX, Stratix IV, and Stratix V devices; however, this approach might increase the number of LEs required to implement the memory and affect design performance.

You can set the `MAXIMUM_DEPTH` parameter for memory modules manually in the megafunction instantiation. The Quartus II software automatically chooses the best design memory configuration for optimal power.

Figure 13-10: RAM 2-Port Maximum Depth Selectable Option



**Related Information**

- [Power-Driven Compilation](#) on page 13-4
- [Clock Power Management](#) on page 13-9  
For more information on clock network-wide gating.

**Memory Power Reduction Example**


Power usage measurements for a 4K × 36 simple dual-port memory implemented using multiple M4K blocks in a Stratix device. For each implementation, the M4K blocks are configured with a different memory depth.

**Table 13-3: 4K × 36 Simple Dual-Port Memory Implemented Using Multiple M4K Blocks**

M4K Configuration	Number of M4K Blocks	ALUTs
4K × 1 (Default setting)	36	0
2K × 2	36	40
1K × 4	36	62
512 × 9	32	143
256 × 18	32	302
128 × 36	32	633

Using the `MAXIMUM_DEPTH` parameter can save power. For all implementations, a user-provided read enable signal is present to indicate when read data is required. Using this power-saving technique can reduce power consumption by as much as 60%.

Figure 13-11: Power Savings Using the MAXIMUM\_DEPTH Parameter



As the memory depth becomes more shallow, memory dynamic power decreases because unaddressed M4K blocks can be shut off using a decoded combination of address bits and the read enable signal. For a 128-deep memory block, power used by the extra LEs starts to outweigh the power gain achieved by using a more shallow memory block depth. The power consumption of the memory blocks and associated LEs depends on the memory configuration.


**Note:** The SOPC Builder and Qsys system do not offer specific power savings control for on-chip memory block. There is no read enable, write enable, or clock enable that you can enable in the on-chip RAM megafunction to shut down the RAM block in the SOPC Builder and Qsys system.

## Pipelining and Retiming

Designs with many glitches consume more power because of faster switching activity. Glitches cause unnecessary and unpredictable temporary logic switches at the output of combinational logic. A glitch usually occurs when there is a mismatch in input signal timing leading to unequal propagation delay.

For example, consider an input change on one input of a 2-input XOR gate from 1 to 0, followed a few moments later by an input change from 0 to 1 on the other input. For a moment, both inputs become 1 (high) during the state transition, resulting in 0 (low) at the output of the XOR gate. Subsequently, when the second input transition takes place, the XOR gate output becomes 1 (high). During signal transition, a glitch is produced before the output becomes stable.

Figure 13-12: XOR Gate Showing Glitch at the Output




This glitch can propagate to subsequent logic and create unnecessary switching activity, increasing power consumption. Circuits with many XOR functions, such as arithmetic circuits or cyclic redundancy check (CRC) circuits, tend to have many glitches if there are several levels of combinational logic between registers.

Pipelining can reduce design glitches by inserting flipflops into long combinational paths. Flipflops do not allow glitches to propagate through combinational paths. Therefore, a pipelined circuit tends to have less glitching. Pipelining has the additional benefit of generally allowing higher clock speed operations, although it does increase the latency of a circuit (in terms of the number of clock cycles to a first result).

An example where pipelining is applied to break up a long combinational path.

Figure 13-13: Pipelining Example



Pipelining is very effective for glitch-prone arithmetic systems because it reduces switching activity, resulting in reduced power dissipation in combinational logic. Additionally, pipelining allows higher-speed operation by reducing logic-level numbers between registers. The disadvantage of this technique is that if there are not many glitches in your design, pipelining can increase power consumption by adding unnecessary registers. Pipelining can also increase resource utilization. The benchmark data shows that pipelining can reduce dynamic power consumption by as much as 30% in Cyclone and Stratix devices.

## Architectural Optimization

You can use design-level architectural optimization by taking advantage of specific device architecture features. These features include dedicated memory and DSP or multiplier blocks available in FPGA devices to perform memory or arithmetic-related functions. You can use these blocks in place of LUTs to reduce power consumption. For example, you can build large shift registers from RAM-based FIFO buffers instead of building the shift registers from the LE registers.

The Stratix device family allows you to efficiently target small, medium, and large memories with the TriMatrix memory architecture. Each TriMatrix memory block is optimized for a specific function. M512 memory blocks are more power-efficient than the distributed memory structures in some competing FPGAs. The M4K memory blocks are used to implement buffers for a wide variety of applications, including processor code storage, large look-up table implementation, and large memory applications. The M-RAM blocks are useful in applications where a large volume of data must be stored on-chip. Effective utilization of these memory blocks can have a significant impact on power reduction in your design.

The latest Stratix and Cyclone device families have configurable M9K memory blocks that provide various memory functions such as RAM, FIFO buffers, and ROM.

### Related Information

- [Area and Timing Optimization](#) on page 12-1  
For more information about using DSP and memory blocks efficiently, refer to the *Quartus II Handbook*.



## I/O Power Guidelines

Nonterminated I/O standards such as LVTTTL and LVCMOS have a rail-to-rail output swing. The voltage difference between logic-high and logic-low signals at the output pin is equal to the  $V_{CCIO}$  supply voltage. If the capacitive loading at the output pin is known, the dynamic power consumed in the I/O buffer can be calculated.

$$P = 0.5 \times F \times C \times V^2$$

In this equation,  $F$  is the output transition frequency and  $C$  is the total load capacitance being switched.  $V$  is equal to  $V_{CCIO}$  supply voltage. Because of the quadratic dependence on  $V_{CCIO}$ , lower voltage standards consume significantly less dynamic power.

Transistor-to-transistor logic (TTL) I/O buffers consume very little static power. As a result, the total power consumed by a LVTTTL or LVCMOS output is highly dependent on load and switching frequency.

When using resistively terminated I/O standards like SSTL and HSTL, the output load voltage swings by a small amount around some bias point. The same dynamic power equation is used, where  $V$  is the actual load voltage swing. Because this is much smaller than  $V_{CCIO}$ , dynamic power is lower than for nonterminated I/O under similar conditions. These resistively terminated I/O standards dissipate significant static (frequency-independent) power, because the I/O buffer is constantly driving current into the resistive termination network. However, the lower dynamic power of these I/O standards means they often have lower total power than LVCMOS or LVTTTL for high-frequency applications. Use the lowest drive strength I/O setting that meets your speed and waveform requirements to minimize I/O power when using resistively terminated standards.

You can save a small amount of static power by connecting unused I/O banks to the lowest possible  $V_{CCIO}$  voltage of 1.2 V.

For more information about I/O standards, refer to the *Stratix IV Device Handbook*, or the *Cyclone IV GX Handbook* on the Altera website.

When calculating I/O power, the PowerPlay Power Analyzer uses the default capacitive load set for the I/O standard in the **Capacitive Loading** page of the **Device and Pin Options** dialog box. Any other components defined in the board trace model are not taken into account for the power measurement.

For Cyclone IV GX, Stratix IV, and Stratix V, devices, Advanced I/O Timing is always used, which uses the full board trace model.

### Related Information

- [I/O Management](#) on page 4-1  
For information about using Advanced I/O Timing and configuring a board trace model, refer to the *Quartus II Handbook*.

## Dynamically Controlled On-Chip Terminations

Stratix IV and Stratix V FPGAs offer dynamic on-chip termination (OCT). Dynamic OCT enables series termination (RS) and parallel termination (RT) to dynamically turn on/off during the data transfer. This feature is especially useful when Stratix IV and Stratix V FPGAs are used with external memory interfaces, such as interfacing with DDR memories.

Compared to conventional termination, dynamic OCT reduces power consumption significantly as it eliminates the constant DC power consumed by parallel termination when transmitting data. Parallel termination is extremely useful for applications that interface with external memories where I/O standards, such as HSTL and SSTL, are used. Parallel termination supports dynamic OCT, which is useful for bidirectional interfaces.

The following is an example of power saving for a DDR3 interface using on-chip parallel termination.

The static current consumed by parallel OCT is equal to the  $V_{CCIO}$  voltage divided by  $100\ \Omega$ . For DDR3 interfaces that use SSTL-15, the static current is  $1.5\ \text{V}/100\ \Omega = 15\ \text{mA}$  per pin. Therefore, the static power is  $1.5\ \text{V} \times 15\ \text{mA} = 22.5\ \text{mW}$ . For an interface with 72 DQ and 18 DQS pins, the static power is  $90\ \text{pins} \times 22.5\ \text{mW} = 2.025\ \text{W}$ . Dynamic parallel OCT disables parallel termination during write operations, so if writing occurs 50% of the time, the power saved by dynamic parallel OCT is  $50\% \times 2.025\ \text{W} = 1.0125\ \text{W}$ .

**Related Information**

**Stratix IV Device I/O Features**

For more information about dynamic OCT in Stratix IV devices, refer to the chapter in the *Stratix IV Device Handbook*.

**Power Optimization Advisor**

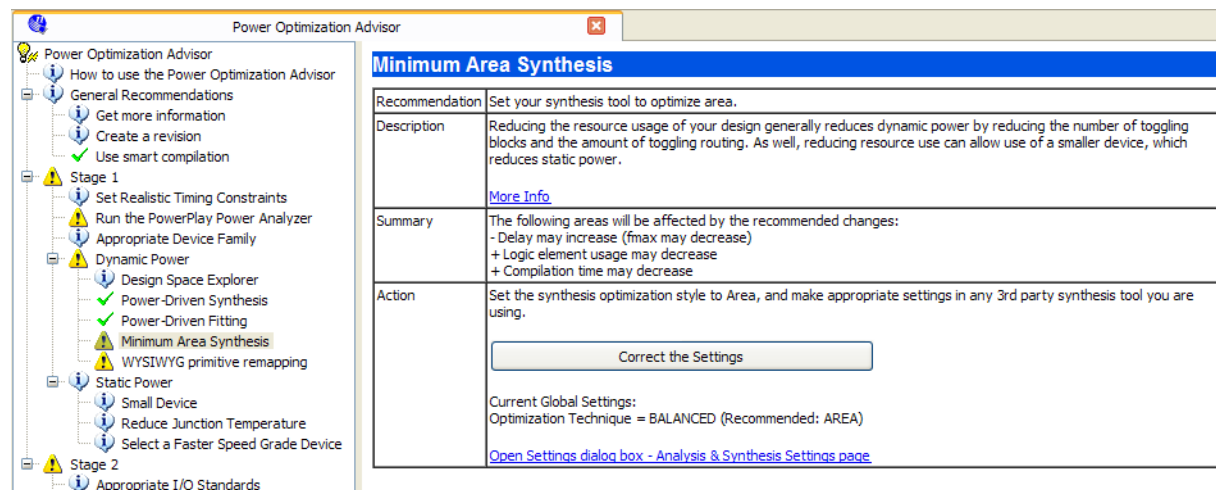
The Quartus II software includes the Power Optimization Advisor, which provides specific power optimization advice and recommendations based on the current design project settings and assignments. The advisor covers many of the suggestions listed in this chapter. The following example shows how to reduce your design power with the Power Optimization Advisor.

**Power Optimization Advisor Example**

After compiling your design, run the PowerPlay Power Analyzer to determine your design power and to see where power is dissipated in your design. Based on this information, you can run the Power Optimization Advisor to implement recommendations that can reduce design power.

The Power Optimization Advisor after compiling a design that is not fully optimized for power.

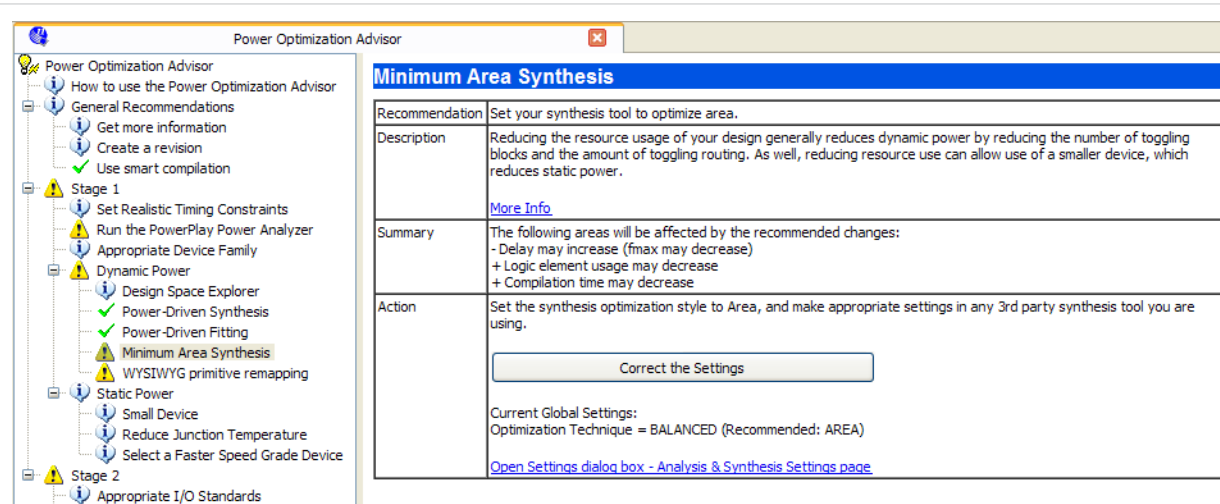
**Figure 13-14: Power Optimization Advisor**



The Power Optimization Advisor shows the recommendations that can reduce power in your design. The recommendations are split into stages to show the order in which you should apply the recommended settings. The first stage shows mostly CAD setting options that are easy to implement and highly effective in reducing design power. An icon indicates whether each recommended setting is made in the current project. The checkmark icons for Stage 1 shows the recommendations that are already implemented. The warning icons indicate recommendations that are not followed for this compilation. The information icon shows the general suggestions. Each recommendation includes the description, summary of the effect of the recommendation, and the action required to make the appropriate setting.

There is a link from each recommendation to the appropriate location in the Quartus II user interface where you can change the setting. After making the recommended changes, recompile your design. The Power Optimization Advisor indicates with green check marks that the recommendations were implemented successfully. You can use the PowerPlay Power Analyzer to verify your design power results.

**Figure 13-15: Implementation of Power Optimization Advisor Recommendations**



The recommendations listed in Stage 2 generally involve design changes, rather than CAD settings changes as in Stage 1. You can use these recommendations to further reduce your design power consumption. Altera recommends that you implement Stage 1 recommendations first, then the Stage 2 recommendations.

## Document Revision History

**Table 13-4: Document Revision History**

Date	Version	Changes
2014.12.15	14.1.0	<ul style="list-style-type: none"> <li>Updated location of Fitter Settings, Analysis &amp; Synthesis Settings, and Physical Synthesis Optimizations to Compiler Settings.</li> <li>Updated DSE II GUI and optimization settings.</li> </ul>
2014.06.30	14.0.0	Updated the format.
May 2013	13.0.0	Added a note to “Memory Power Reduction Example” on Qsys and SOPC Builder power savings limitation for on-chip memory block.
June 2012	12.0.0	Removed survey link.
November 2011	10.0.2	Template update.
December 2010	10.0.1	Template update.

Date	Version	Changes
July 2010	10.0.0	<ul style="list-style-type: none"> <li>• Was chapter 11 in the 9.1.0 release</li> <li>• Updated Figures 14-2, 14-3, 14-6, 14-18, 14-19, and 14-20</li> <li>• Updated device support</li> <li>• Minor editorial updates</li> </ul>
November 2009	9.1.0	<ul style="list-style-type: none"> <li>• Updated Figure 11-1 and associated references</li> <li>• Updated device support</li> <li>• Minor editorial update</li> </ul>
March 2009	9.0.0	<ul style="list-style-type: none"> <li>• Was chapter 9 in the 8.1.0 release</li> <li>• Updated for the Quartus II software release</li> <li>• Added benchmark results</li> <li>• Removed several sections</li> <li>• Updated Figure 13-1, Figure 13-17, and Figure 13-18</li> </ul>
November 2008	8.1.0	<ul style="list-style-type: none"> <li>• Changed to 8½" × 11" page size</li> <li>• Changed references to altsyncram to RAM</li> <li>• Minor editorial updates</li> </ul>
May 2008	8.0.0	<ul style="list-style-type: none"> <li>• Added support for Stratix IV devices</li> <li>• Updated Table 9-1 and 9-9</li> <li>• Updated "Architectural Optimization" on page 9-22</li> <li>• Added "Dynamically-Controlled On-Chip Terminations" on page 9-26</li> <li>• Updated "Referenced Documents" on page 9-29</li> <li>• Updated references</li> </ul>

**Related Information**

[Quartus II Handbook Archive](#)

For previous versions of the Quartus II Handbook on the Altera website.

2014.12.15

QI15V2



Subscribe



Send Feedback

This chapter describes techniques to reduce resource usage when designing for Altera® devices.

## Resource Utilization

Determining device utilization is important regardless of whether your design achieved a successful fit. If your compilation results in a no-fit error, resource utilization information is important for analyzing the fitting problems in your design. If your fitting is successful, review the resource utilization information to determine whether the future addition of extra logic or other design changes might introduce fitting difficulties. Also, review the resource utilization information to determine if it is impacting timing performance.

To determine resource usage, refer to the **Flow Summary** section of the Compilation Report. This section reports resource utilization, including pins, memory bits, digital signal processing (DSP) blocks, and phase-locked loops (PLLs). **Flow Summary** indicates whether your design exceeds the available device resources. More detailed information is available by viewing the reports under **Resource Section** in the **Fitter** section of the Compilation Report.

**Flow Summary** shows the overall logic utilization. The Fitter can spread logic throughout the device, which may lead to higher overall utilization.

As the device fills up, the Fitter automatically searches for logic functions with common inputs to place in one ALM. The number of packed registers also increases. Therefore, a design that has high overall utilization might still have space for extra logic if the logic and registers can be packed together more tightly.

The reports under the **Resource Section** in the **Fitter** section of the Compilation Report provide more detailed resource information. The Fitter Resource Usage Summary report breaks down the logic utilization information and provides other resource information, including the number of bits in each type of memory block. This panel also contains a summary of the usage of global clocks, PLLs, DSP blocks, and other device-specific resources.

You can also view reports describing some of the optimizations that occurred during compilation. For example, if you use Quartus® II integrated synthesis, the reports in the Optimization Results folder in the **Analysis & Synthesis** section include information about registers that integrated synthesis removed during synthesis. Use this report to estimate device resource utilization for a partial design to ensure that registers were not removed due to missing connections with other parts of the design.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered

If a specific resource usage is reported as less than 100% and a successful fit cannot be achieved, either there are not enough routing resources or some assignments are illegal. In either case, a message appears in the **Processing** tab of the Messages window describing the problem.

If the Fitter finishes unsuccessfully and runs much faster than on similar designs, a resource might be over-utilized or there might be an illegal assignment. If the Quartus II software seems to run for an excessively long time compared to runs on similar designs, a legal placement or route probably cannot be found. In the Compilation Report, look for errors and warnings that indicate these types of problems.

You can use the Chip Planner to find areas of the device that have routing congestion on specific types of routing resources. If you find areas with very high congestion, analyze the cause of the congestion. Issues such as high fan-out nets not using global resources, an improperly chosen optimization goal (speed versus area), very restrictive floorplan assignments, or the coding style can cause routing congestion. After you identify the cause, modify the source or settings to reduce routing congestion.

#### Related Information

- [Fitter Resources Report](#)  
For more information about Fitter Resources Report
- [Displaying Resources and Information](#)  
For information about how to view routing congestion
- [About the Chip Planner](#)  
For information about using the Chip Planner tool
- [Analyzing and Optimizing the Design Floorplan with the Chip Planner](#) on page 15-1  
For details about using the Chip Planner tool

## Optimizing Resource Utilization (LUT-Based Devices)

The following lists the stages after design analysis:

- Optimize resource utilization—Ensure that you have already set the basic constraints
- I/O timing optimization—Optimize I/O timing after you optimize resource utilization and your design fits in the desired target device
- Register-to-register timing optimization

#### Related Information

- [Design Optimization Overview](#) on page 10-1  
Provides information about setting basic constraints
- [Timing Closure and Optimization](#) on page 12-1  
Provides information about optimizing I/O timing. These tips are valid for all FPGA families and the MAX II family of CPLDs.

## Using the Resource Optimization Advisor

The Resource Optimization Advisor provides guidance in determining settings that optimize resource usage. To run the Resource Optimization Advisor, on the Tools menu, point to **Advisors**, and click **Resource Optimization Advisor**.

The Resource Optimization Advisor provides step-by-step advice about how to optimize resource usage (logic element, memory block, DSP block, I/O, and routing) of your design. Some of the recommenda-

tions in these categories might conflict with each other. Altera recommends evaluating the options and choosing the settings that best suit your requirements.

**Related Information**

**[Resource Optimization Advisor Command Tools Menu](#)**

For more information about the Resource Optimization Advisor

## Resolving Resource Utilization Issues Summary

Resource utilization issues can be divided into the following three categories:

**Table 14-1: Resource Utilization Issues**

Types of Resource Utilization Issues	Refer to
Issues relating to I/O pin utilization or placement, including dedicated I/O blocks such as PLLs or LVDS transceivers	<a href="#">I/O Pin Utilization or Placement</a> on page 14-3
Issues relating to logic utilization or placement, including logic cells containing registers and LUTs as well as dedicated logic, such as memory blocks and DSP blocks	<a href="#">Logic Utilization or Placement</a> on page 14-4
Issues relating to routing	<a href="#">Routing</a> on page 14-8

**Related Information**

- [I/O Pin Utilization or Placement](#) on page 14-3
- [Logic Utilization or Placement](#) on page 14-4
- [Routing](#) on page 14-8

### I/O Pin Utilization or Placement

Resolve I/O resource problems with these guidelines.

**Guideline: Use I/O Assignment Analysis**

To help with pin placement, on the Processing menu, point to **Start** and click **Start I/O Assignment Analysis**. The **Start I/O Assignment Analysis** command allows you to check your I/O assignments early in the design process. You can use this command to check the legality of pin assignments before, during, or after compilation of your design. If design files are available, you can use this command to accomplish more thorough legality checks on your design’s I/O pins and surrounding logic. These checks include proper reference voltage pin usage, valid pin location assignments, and acceptable mixed I/O standards.

Common issues with I/O placement relate to the fact that differential standards have specific pin pairings and certain I/O standards might be supported only on certain I/O banks.

If your compilation or I/O assignment analysis results in specific errors relating to I/O pins, follow the recommendations in the error message. Right-click the message in the Messages window and click **Help** to open the Quartus II Help topic for this message.

## Guideline: Modify Pin Assignments or Choose a Larger Package

If a design that has pin assignments fails to fit, compile the design without the pin assignments to determine whether a fit is possible for the design in the specified device and package. You can use this approach if a Quartus II error message indicates fitting problems due to pin assignments.

If the design fits when all pin assignments are ignored or when several pin assignments are ignored or moved, you might have to modify the pin assignments for the design or select a larger package.

If the design fails to fit because insufficient I/Os pins are available, a successful fit can often be obtained by using a larger device package (which can be the same device density) that has more available user I/O pins.

### Related Information

- [I/O Management](#) on page 4-1  
For more information about I/O assignment analysis

## Logic Utilization or Placement

Resolve logic resource problems, including logic cells containing registers and LUTs, as well as dedicated logic such as memory blocks and DSP blocks, with these guidelines.

### Guideline: Optimize Source Code

If your design does not fit because of logic utilization, then evaluate and modify the design at the source. You can often improve logic significantly by making design-specific changes to your source code. This is typically the most effective technique for improving the quality of your results.

If your design does not fit into available logic elements (LEs) or ALMs, but you have unused memory or DSP blocks, check if you have code blocks in your design that describe memory or DSP functions that are not being inferred and placed in dedicated logic. You might be able to modify your source code to allow these functions to be placed into dedicated memory or DSP resources in the target device.

Ensure that your state machines are recognized as state machine logic and optimized appropriately in your synthesis tool. State machines that are recognized are generally optimized better than if the synthesis tool treats them as generic logic. In the Quartus II software, you can check for the State Machine report under **Analysis & Synthesis** in the Compilation Report. This report provides details, including the state encoding for each state machine that was recognized during compilation. If your state machine is not being recognized, you might have to change your source code to enable it to be recognized.

### Related Information

- [Recommended HDL Coding Styles](#)  
For coding style guidelines, including examples of HDL code for inferring memory and DSP functions
- [Recommended HDL Coding Styles](#)  
For guidelines and sample HDL code for state machines
- [AN 584: Timing Closure Methodology for Advanced FPGA Designs.](#)  
For additional HDL coding examples

### Guideline: Optimize Synthesis for Area, Not Speed

If your design fails to fit because it uses too much logic, resynthesize the design to improve the area utilization. First, ensure that you have set your device and timing constraints correctly in your synthesis tool. Particularly when area utilization of the design is a concern, ensure that you do not over-constrain the timing requirements for the design. Synthesis tools generally try to meet the specified requirements, which can result in higher device resource usage if the constraints are too aggressive.



If resource utilization is an important concern, some synthesis tools offer an easy way to optimize for area instead of speed. If you are using Quartus II integrated synthesis, select **Balanced** or **Area** for the **Optimization Technique**. You can also specify an **Optimization Technique** logic option for specific modules in your design with the Assignment Editor in cases where you want to reduce area using the **Area** setting (potentially at the expense of register-to-register timing performance) while leaving the default **Optimization Technique** setting at **Balanced** (for the best trade-off between area and speed for certain device families) or **Speed**. You can also use the **Speed Optimization Technique for Clock Domains** logic option to specify that all combinational logic in or between the specified clock domain(s) is optimized for speed.

In some synthesis tools, not specifying an  $f_{MAX}$  requirement can result in less resource utilization.

**Note:** In the Quartus II software, the **Balanced** setting typically produces utilization results that are very similar to those produced by the **Area** setting, with better performance results. The **Area** setting can give better results in some cases.

The Quartus II software provides additional attributes and options that can help improve the quality of your synthesis results.

#### Related Information

##### [Synthesis](#)

For information about setting the timing requirements and synthesis options in Quartus II integrated synthesis and other synthesis tools

### Guideline: Restructure Multiplexers

Multiplexers form a large portion of the logic utilization in many FPGA designs. By optimizing your multiplexed logic, you can achieve a more efficient implementation in your Altera device.

#### Related Information

- [Restructure Multiplexers logic option](#)  
For more information about the Restructure Multiplexers option
- [Recommended HDL Coding Styles](#)  
For design guidelines to achieve optimal resource utilization for multiplexer designs

### Guideline: Perform WYSIWYG Primitive Resynthesis with Balanced or Area Setting

The **Perform WYSIWYG Primitive Resynthesis** logic option specifies whether to perform WYSIWYG primitive resynthesis during synthesis. This option uses the setting specified in the **Optimization Technique** logic option. The **Perform WYSIWYG Primitive Resynthesis** logic option is useful for resynthesizing some or all of the WYSIWYG primitives in your design for better area or performance. However, WYSIWYG primitive resynthesis can be done only when you use third-party synthesis tools.

**Note:** The **Balanced** setting typically produces utilization results that are very similar to the **Area** setting with better performance results. The **Area** setting can give better results in some cases. Performing WYSIWYG resynthesis for area in this way typically reduces register-to-register timing performance.

#### Related Information

##### [Perform WYSIWYG Primitive Resynthesis logic option](#)

For information about this logic option

### Guideline: Use Register Packing

The **Auto Packed Registers** option implements the functions of two cells into one logic cell by combining the register of one cell in which only the register is used with the LUT of another cell in which only the LUT is used.

#### Related Information

##### [Auto Packed Registers logic option](#)

For more information about the Auto Packed Registers logic option

### Guideline: Remove Fitter Constraints

A design with conflicting constraints or constraints that are difficult to meet may not fit in the targeted device. For example, a design might fail to fit if the location or LogicLock assignments are too strict and not enough routing resources are available on the device.

To resolve routing congestion caused by restrictive location constraints or LogicLock region assignments, use the **Routing Congestion** task in the Chip Planner to locate routing problems in the floorplan, then remove any internal location or LogicLock region assignments in that area. If your design still does not fit, the design is over-constrained. To correct the problem, remove all location and LogicLock assignments and run successive compilations, incrementally constraining the design before each compilation. You can delete specific location assignments in the Assignment Editor or the Chip Planner. To remove LogicLock assignments in the Chip Planner, in the LogicLock Regions Window, or on the Assignments menu, click **Remove Assignments**. Turn on the assignment categories you want to remove from the design in the **Available assignment categories** list.

#### Related Information

- [Analyzing and Optimizing the Design Floorplan with the Chip Planner](#) on page 15-1  
For more information about the **Routing Congestion** task in the Chip Planner

### Guideline: Flatten the Hierarchy During Synthesis

Synthesis tools typically provide the option of preserving hierarchical boundaries, which can be useful for verification or other purposes. However, the Quartus II software optimizes across hierarchical boundaries so as to perform the most logic minimization, which can reduce area in a design with no design partitions.

If you are using Quartus II incremental compilation, you cannot flatten your design across design partitions. Incremental compilation always preserves the hierarchical boundaries between design partitions, and the synthesis does not flatten it across partitions. Follow Altera's recommendations for design partitioning, such as registering partition boundaries to reduce the effect of cross-boundary optimizations.

#### Related Information

- [Quartus II Incremental Compilation for Hierarchical and Team-Based Design](#)  
For more information about using incremental compilation and recommendations for design partitioning
- [Best Practices for Incremental Compilation Partitions and Floorplan Assignments](#)  
For more information about using incremental compilation and recommendations for design partitioning

### Guideline: Retarget Memory Blocks

If your design fails to fit because it runs out of device memory resources, your design may require a certain type of memory that the device does not have. For example, a design that requires two M-RAM

blocks cannot be targeted to a device with only one M-RAM block. You might be able to obtain a fit by building one of the memories with a different size memory block, such as an M4K memory block.

If the memory block was created with a parameter editor, open the parameter editor and edit the RAM block type so it targets a new memory block size.

ROM and RAM memory blocks can also be inferred from your HDL code, and your synthesis software can place large shift registers into memory blocks by inferring the Shift register (RAM-based) IP core. This inference can be turned off in your synthesis tool to cause the memory or shift registers to be placed in logic instead of in memory blocks. Also, for improved timing performance, you can turn this inference off to prevent registers from being moved into RAM.

Depending on your synthesis tool, you can also set the RAM block type for inferred memory blocks. In Quartus II integrated synthesis, set the **ramstyle** attribute to the desired memory type for the inferred RAM blocks, or set the option to **logic**, to implement the memory block in standard logic instead of a memory block.

Consider the Resource Utilization by Entity report in the report file and determine whether there is an unusually high register count in any of the modules. Some coding styles can prevent the Quartus II software from inferring RAM blocks from the source code because of the blocks' architectural implementation, and force the software to implement the logic in flipflops. As an example, a function such as an asynchronous reset on a register bank might make the resistor bank incompatible with the RAM blocks in the device architecture, so that the register bank is implemented in flipflops. It is often possible to move a large register bank into RAM by slight modification of associated logic.

#### Related Information

- [Auto RAM Replacement logic option](#)
- [Auto ROM Replacement logic option](#)
- [Auto Shift Register Replacement logic option](#)
- [Synthesis](#)  
For more information about memory inference control in other synthesis tools
- [Recommended HDL Coding Styles](#)  
For more information about coding styles and HDL examples that ensure memory inference

### Guideline: Use Physical Synthesis Options to Reduce Area

The physical synthesis options for fitting help you decrease resource usage. When you enable these options, the Quartus II software makes placement-specific changes to the netlist that reduce resource utilization for a specific Altera device.

**Note:** The compilation time might increase considerably when you use physical synthesis options.

With the Quartus II software, you can apply physical synthesis options to specific instances, which can reduce the impact on compilation time. Physical synthesis instance assignments allow you to enable physical synthesis algorithms for specific portions of your design.

The following physical synthesis optimizations for fitting are available:

- Physical synthesis for combinational logic
- Map logic into memory

#### Related Information

[Physical Synthesis Optimizations Page Settings Dialog Box](#)

## Guideline: Retarget or Balance DSP Blocks

A design might not fit because it requires too many DSP blocks. You can implement all DSP block functions with logic cells, so you can retarget some of the DSP blocks to logic to obtain a fit.

If the DSP function was created with the parameter editor, open the parameter editor and edit the function so it targets logic cells instead of DSP blocks. The Quartus II software uses the `DEDICATED_MULTIPLIER_CIRCUITRY` IP core parameter to control the implementation.

DSP blocks also can be inferred from your HDL code for multipliers, multiply-adders, and multiply-accumulators. You can turn off this inference in your synthesis tool. When you are using Quartus II integrated synthesis, you can disable inference by turning off the **Auto DSP Block Replacement** logic option for your entire project. Click **Assignments > Settings > Compiler Settings > Advanced Settings (Synthesis)**. Turn off **Auto DSP Block Replacement**. Alternatively, you can disable the option for a specific block with the Assignment Editor.

The Quartus II software also offers the **DSP Block Balancing** logic option, which implements DSP block elements in logic cells or in different DSP block modes. The default **Auto** setting allows DSP block balancing to convert the DSP block slices automatically as appropriate to minimize the area and maximize the speed of the design. You can use other settings for a specific node or entity, or on a project-wide basis, to control how the Quartus II software converts DSP functions into logic cells and DSP blocks. Using any value other than **Auto** or **Off** overrides the `DEDICATED_MULTIPLIER_CIRCUITRY` parameter used in IP core variations.

### Related Information

- [Synthesis](#)  
For more information about disabling DSP block inference in other synthesis tools
- [Auto DSP Block Replacement logic option](#)
- [DSP Block Balancing logic option](#)

## Guideline: Use a Larger Device

If a successful fit cannot be achieved because of a shortage of routing resources, you might require a larger device.

## Routing

Resolve routing resource problems with these guidelines.

### Guideline: Set Auto Packed Registers to Sparse or Sparse Auto

The **Auto Packed Registers** option reduces LE or ALM count in a design. You can set this option by clicking **Assignment > Settings > Compiler Settings > Advanced Settings (Fitter)**.

### Related Information

[Auto Packed Registers logic option](#)

### Guideline: Set Fitter Aggressive Routability Optimizations to Always

The **Fitter Aggressive Routability Optimization** option is useful if your design does not fit due to excessive routing wire utilization.

If there is a significant imbalance between placement and routing time (during the first fitting attempt), it might be because of high wire utilization. Turning on the **Fitter Aggressive Routability Optimizations** option can reduce your compilation time.

On average, this option can save up to 6% wire utilization, but can also reduce performance by up to 4%, depending on the device.

#### Related Information

[Fitter Aggressive Routability Optimizations logic option](#)

### Guideline: Increase Router Effort Multiplier

The Router Effort Multiplier controls how quickly the router tries to find a valid solution. The default value is 1.0 and legal values must be greater than 0. Numbers higher than 1 help designs that are difficult to route by increasing the routing effort. Numbers closer to 0 (for example, 0.1) can reduce router runtime, but usually reduce routing quality slightly. Experimental evidence shows that a multiplier of 3.0 reduces overall wire usage by approximately 2%. Using a Router Effort Multiplier higher than the default value could be beneficial for designs with complex datapaths with more than five levels of logic. However, congestion in a design is primarily due to placement, and increasing the Router Effort Multiplier does not necessarily reduce congestion.

**Note:** Any Router Effort Multiplier value greater than 4 only increases by 10% for every additional 1. For example, a value of 10 is actually 4.6.

#### Related Information

[Router Effort Multiplier logic option](#)

### Guideline: Remove Fitter Constraints

A design with conflicting constraints or constraints that are difficult to achieve may not fit the targeted device. Conflicting or difficult-to-achieve constraints can occur when location or LogicLock assignments are too strict and there are not enough routing resources.

In this case, use the **Routing Congestion** task in the Chip Planner to locate routing problems in the floorplan, then remove all location and LogicLock region assignments from that area. If the local constraints are removed, and the design still does not fit, the design is over-constrained. To correct the problem, remove all location and LogicLock assignments and run successive compilations, incrementally constraining the design before each compilation. You can delete specific location assignments in the Assignment Editor or the Chip Planner. To remove LogicLock assignments in the Chip Planner, in the LogicLock Regions Window, or on the Assignments menu, click **Remove Assignments**. Turn on the assignment categories you want to remove from the design in the **Available assignment categories** list.

#### Related Information

- [About the Chip Planner](#)
- [Analyzing and Optimizing the Design Floorplan with the Chip Planner](#) on page 15-1  
For more information about the Routing Congestion task in the Chip Planner

### Guideline: Optimize Synthesis for Area, Not Speed

In some cases, resynthesizing the design to improve the area utilization can also improve the routability of the design. First, ensure that you have set your device and timing constraints correctly in your synthesis tool. Ensure that you do not overconstrain the timing requirements for the design, particularly when the area utilization of the design is a concern. Synthesis tools generally try to meet the specified requirements, which can result in higher device resource usage if the constraints are too aggressive.

If resource utilization is important to improve the routing results in your design, some synthesis tools offer an easy way to optimize for area instead of speed. If you are using Quartus II integrated synthesis, click **Assignments > Settings > Compiler Settings > Advanced Settings (Synthesis)**. For **Optimization Technique**, select **Balanced** or **Area**.

You can also specify this logic option for specific modules in your design with the Assignment Editor in cases where you want to reduce area using the **Area** setting (potentially at the expense of register-to-register timing performance). You can apply the setting to specific modules while leaving the default **Optimization Technique** setting at **Balanced** (for the best trade-off between area and speed for certain device families) or **Speed**. You can also use the **Speed Optimization Technique for Clock Domains** logic option to specify that all combinational logic in or between the specified clock domain(s) is optimized for speed.

**Note:** In the Quartus II software, the **Balanced** setting typically produces utilization results that are very similar to those obtained with the **Area** setting, with better performance results. The **Area** setting can yield better results in some unusual cases.

In some synthesis tools, not specifying an  $f_{MAX}$  requirement can result in less resource utilization, which can improve routability.

#### Related Information

##### Synthesis

For information about setting the timing requirements and synthesis options in Quartus II integrated synthesis and other synthesis tools

#### Guideline: Optimize Source Code

If your design does not fit because of routing problems and the methods described in the preceding sections do not sufficiently improve the routability of the design, modify the design at the source to achieve the desired results. You can often improve results significantly by making design-specific changes to your source code, such as duplicating logic or changing the connections between blocks that require significant routing resources.

#### Guideline: Use a Larger Device

If a successful fit cannot be achieved because of a shortage of routing resources, you might require a larger device.

## Scripting Support

You can run procedures and make settings described in this chapter in a Tcl script. You can also run some procedures at a command prompt. For detailed information about scripting command options, refer to the Quartus II command-line and Tcl API Help browser. To run the Help browser, type the following command at the command prompt:

```
quartus_sh --qhelp
```

You can specify many of the options described in this section either in an instance, or at a global level, or both.

Use the following Tcl command to make a global assignment:

```
set_global_assignment -name <.qsf variable name> <value>
```

Use the following Tcl command to make an instance assignment:

```
set_instance_assignment -name <.qsf variable name> <value> \  
-to <instance name>
```

**Note:** If the *<value>* field includes spaces (for example, 'Standard Fit'), you must enclose the value in straight double quotation marks.

**Related Information**

[Tcl Scripting](#) on page 3-1

For more information about Tcl scripting

[Quartus II Settings File Manual](#)

For more information about all settings and constraints in the Quartus II software

[Command-Line Scripting](#) on page 2-1

For more information about command-line scripting

## Initial Compilation Settings

Use the Quartus II Settings File (.qsf) variable name in the Tcl assignment to make the setting along with the appropriate value. The **Type** column indicates whether the setting is supported as a global setting, an instance setting, or both.

**Table 14-2: Advanced Compilation Settings**

Setting Name	.qsf File Variable Name	Values	Type
Placement Effort Multiplier	PLACEMENT_EFFORT_MULTIPLIER	Any positive, non-zero value	Global
Router Effort Multiplier	ROUTER_EFFORT_MULTIPLIER	Any positive, non-zero value	Global
Router Timing Optimization level	ROUTER_TIMING_OPTIMIZATION_LEVEL	NORMAL, MINIMUM, MAXIMUM	Global
Final Placement Optimization	FINAL_PLACEMENT_OPTIMIZATION	ALWAYS, AUTOMATICALLY, NEVER	Global

## Resource Utilization Optimization Techniques (LUT-Based Devices)

This table lists the .qsf file variable name and applicable values for the settings described in [Optimizing Resource Utilization \(LUT-Based Devices\)](#) on page 14-2.

**Table 14-3: Resource Utilization Optimization Settings**

Setting Name	.qsf File Variable Name	Values	Type
Auto Packed Registers <sup>(1)</sup>	AUTO_PACKED_REGISTERS_ <device family name>	OFF, NORMAL, MINIMIZE AREA, MINIMIZE AREA WITH CHAINS, AUTO	Global, Instance
Perform WYSIWYG Primitive Resynthesis	ADV_NETLIST_OPT_SYNTH_WYSIWYG_REMAP	ON, OFF	Global, Instance

<sup>(1)</sup> Allowed values for this setting depend on the device family that you select.

Setting Name	.qsf File Variable Name	Values	Type
Physical Synthesis for Combinational Logic for Reducing Area	PHYSICAL_SYNTHESIS_COMBO_LOGIC_FOR_AREA	ON, OFF	Global, Instance
Physical Synthesis for Mapping Logic to Memory	PHYSICAL_SYNTHESIS_MAP_LOGIC_TO_MEMORY_FOR_AREA	ON, OFF	Global, Instance
Optimization Technique	<device family name>_OPTIMIZATION_TECHNIQUE	AREA, SPEED, BALANCED	Global, Instance
Speed Optimization Technique for Clock Domains	SYNTH_CRITICAL_CLOCK	ON, OFF	Instance
State Machine Encoding	STATE_MACHINE_PROCESSING	AUTO, ONE-HOT, GRAY, JOHNSON, MINIMAL BITS, ONE-HOT, SEQUENTIAL, USER-ENCODE	Global, Instance
Auto RAM Replacement	AUTO_RAM_RECOGNITION	ON, OFF	Global, Instance
Auto ROM Replacement	AUTO_ROM_RECOGNITION	ON, OFF	Global, Instance
Auto Shift Register Replacement	AUTO_SHIFT_REGISTER_RECOGNITION	ON, OFF	Global, Instance
Auto Block Replacement	AUTO_DSP_RECOGNITION	ON, OFF	Global, Instance
Number of Processors for Parallel Compilation	NUM_PARALLEL_PROCESSORS	Integer between 1 and 16 inclusive, or ALL	Global

## Document Revision History

Table 14-4: Document Revision History

Date	Version	Changes
2014.12.15	14.1.0	Updated location of Fitter Settings, Analysis & Synthesis Settings, and Physical Synthesis Optimizations to Compiler Settings.



Date	Version	Changes
June 2014	14.0.0	<ul style="list-style-type: none"><li>• Removed Cyclone III and Stratix III devices references.</li><li>• Removed Macrocell-Based CPLDs related information.</li><li>• Updated template.</li></ul>
May 2013	13.0.0	Initial release.

# Analyzing and Optimizing the Design Floorplan with the Chip Planner 15

2014.12.15

QII5V2



Subscribe



Send Feedback

## Analyzing and Optimizing the Design Floorplan with the Chip Planner

As FPGA designs grow larger in density, the ability to analyze the design for performance, routing congestion, and logic placement to meet the design requirements becomes critical. This chapter discusses how to analyze the design floorplan with the Chip Planner.

Design floorplan analysis is a valuable method for achieving timing closure and optimal performance in highly complex designs. With analysis capability, the Quartus II<sup>®</sup> Chip Planner helps you close timing quickly on your designs. Using the Chip Planner together with LogicLock and Incremental Compilation enables you to compile your designs hierarchically, preserving the timing results from individual compilation runs. You can use LogicLock regions as part of an incremental compilation methodology to improve your productivity.

You can perform design analysis, as well as creating and optimizing the design floorplan with the Chip Planner. To make I/O assignments, use the Pin Planner.

### Related Information

- **I/O Management** on page 4-1  
For information about the Pin Planner.
- **Quartus II Incremental Compilation for Hierarchical and Team-Based Design**
- **Best Practices for Incremental Compilation Partitions and Floorplan Assignments**  
You can use the Design Partition Planner with the Chip Planner to customize the floorplan of your design.
- **About the Chip Planner**  
For a list of devices supported by the Chip Planner.
- **Altera Training**  
For more information about the Chip Planner.

## Chip Planner Overview

The Chip Planner provides a visual display of chip resources. The Chip Planner can show logic placement, LogicLock regions, relative resource usage, detailed routing information, fan-in and fan-out connections between nodes, timing paths between registers, delay estimates for paths, and routing congestion information.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



You can also make assignment changes with the Chip Planner, such as creating and deleting resource assignments, and you can perform post-compilation changes such as creating, moving, and deleting logic cells and I/O atoms. With the Chip Planner, you can view and create assignments for a design floorplan, perform power and design analyses, and implement ECOs. With the Chip Planner and Resource Property Editor, you can change connections between resources and make post-compilation changes to the properties of logic cells, I/O elements, PLLs, and RAM and digital signal processing (DSP) blocks.

#### Related Information

- [Engineering Change Management with the Chip Planner](#) on page 17-1  
For details about how to implement ECOs in your design using the Chip Planner in the Quartus II software.

## Starting the Chip Planner

To start the Chip Planner, on the Tools menu, click **Chip Planner (Floorplan & Chip Editor)**. You can also start the Chip Planner by the following methods:

- Click the **Chip Planner** icon on the Quartus II software toolbar
- On the Shortcut menu in the following tools, click **Locate > Locate in Chip Planner (Floorplan and Chip Editor)** to locate :
  - Design Partition Planner
  - Compilation Report
  - LogicLock Regions window
  - Technology Map Viewer
  - Project Navigator window
  - RTL source code
  - Node Finder
  - Simulation Report
  - RTL Viewer
  - Report Timing panel of the TimeQuest Timing Analyzer

## Chip Planner Toolbar

The Chip Planner provides powerful tools for design analysis with a GUI. You can access Chip Planner commands from the View menu and the Shortcut menu, or by clicking the icons on the toolbar.

## Chip Planner Presets, Layers, and Editing Modes

The Chip Planner models types of resource objects as unique display layers, and uses presets— which are predefined sets of layer settings—to control the display of resources.

The Chip Planner provides a set of default presets, and you can create custom presets to customize the display for your particular needs. The Basic, Detailed, and Floorplan Editing presets provided with the Chip Planner are useful for general ECO and assignment-related activities, while the Design Partition Planner preset is optimized for specific activities.

The Chip Planner has two editing modes, which determine the types of operations that you can perform. The Assignment editing mode allows you to make assignment changes that are applied by the Fitter during the next place and route operation. The ECO editing mode allows you to make post-compilation changes, commonly referred to as engineering change orders (ECOs).

You should choose the editing mode appropriate for the work that you want to perform, and a preset that displays the resources that you want to view, in a level of detail appropriate for your design.

## Locate History Window

As you optimize your design floorplan, you might have to locate a path or node in the Chip Planner many times.

The Locate History window lists all the nodes and paths you have displayed using a **Locate in Chip Planner (Floorplan and Chip Editor)** command, providing easy access to the nodes and paths of interest to you. If you locate a required path from the **TimeQuest Timing Analyzer Report Timing** pane, the **Locate History** window displays the required clock path. If you locate an arrival path from the **TimeQuest Timing Analyzer Report Timing** pane, the **Locate History** window displays the path from the arrival clock to the arrival data. Double-clicking a node or path in the **Locate History** window displays the selected node or path in the Chip Planner.

### Related Information

- [Layers Settings Dialog Box](#)
- [About the Chip Planner](#)  
For more information about the Chip Planner, refer to Quartus II Help.
- [Engineering Change Management with the Chip Planner](#) on page 17-1  
For more information about the ECO editing mode

## LogicLock Regions

LogicLock regions are floorplan location constraints that help you place logic on the target device. When you assign entity instances or nodes to a LogicLock region, you direct the Fitter to place those entity instances or nodes within the region during fitting. Your floorplan can contain several LogicLock regions.

A LogicLock region is defined by its height, width, and location; you can specify the size or location of a region, or both, or the Quartus II software can generate these properties automatically. The Quartus II software bases the size and location of a region on the contents of the region and the timing requirements of the module.

**Table 15-1: Types of LogicLock Regions**

Property	Value	Behavior
State	Floating   Locked	Floating allows the Quartus II software to determine the location of the region on the device. Floating regions are shown with a dashed boundary in the floorplan. Locked allows you to specify the location of the region. Locked regions are shown with a solid boundary in the floorplan. A locked region must have a fixed size.
Size	Auto   Fixed	Auto allows the Quartus II software to determine the appropriate size of a region given its contents. Fixed regions have a shape and size that you define.
Reserved	Off   On	Allows you to define whether the Fitter can use the resources within a region for entities that are not assigned to the region. If the reserved property is turned on, only items assigned to the region can be placed within its boundaries.
Origin	Any Floorplan Location	Specifies the location of the LogicLock region on the floorplan. For Arria series, Stratix series, Cyclone series, MAX II, and MAX V devices, the origin is located at the lower left corner of the LogicLock region. For other Altera® device families, the origin is located at the upper left corner of the LogicLock region.

Property	Value	Behavior
----------	-------	----------

**Note:** The Quartus II software cannot automatically define the size of a region if the location is locked. Therefore, if you want to specify the exact location of the region, you must also specify the size.

You can use the Design Partition Planner in conjunction with LogicLock regions to create a floorplan for your design.

#### Related Information

- [Quartus II Incremental Compilation for Hierarchical and Team-Based Designs](#)
- [Quartus II Incremental Compilation for Hierarchical and Team-Based Designs](#)

For more information about using the Design Partition Planner.

## Creating LogicLock Regions

You can create LogicLock Regions with the Project Navigator, the LogicLock Regions window, the Design Partition Planner, the Chip Planner, and with Tcl commands.

### Creating LogicLock Regions with the Project Navigator

After you perform either a full compilation or analysis and elaboration on the design, the Quartus II software displays the hierarchy of the design. On the View menu, **Utility Windows > Project Navigator**. With the hierarchy of the design fully expanded, right-click on any design entity in the design, and click **Create New LogicLock Region** to create a LogicLock region and assign the entity to the new region.

### Creating LogicLock Regions with the LogicLock Regions window

To create a LogicLock region with the LogicLock Regions window, on the Assignments menu, click **LogicLock Regions Window**. In the LogicLock Regions window, click <<new>>.

### Creating LogicLock Regions with the Design Partition Planner

To create a LogicLock region and assign a partition to it with the Design Partition Planner, right-click the partition and then click **Create LogicLock Region**.

### Creating LogicLock Regions with the Chip Planner

To create a LogicLock region in the Chip Planner, click **LogicLock Regions > Create LogicLock Region** on the View menu, then click and drag on the Chip Planner floorplan to create a region of your preferred location and size.

#### Related Information

[Creating or Modifying LogicLock Regions](#) on page 15-20

## Creating Non-rectangular LogicLock Regions

When you create a floorplan for your design, you may want to create nonrectangular LogicLock regions to exclude certain resources from the LogicLock region.

You might also create a nonrectangular LogicLock region to place certain parts of your design around specific device resources to improve performance.

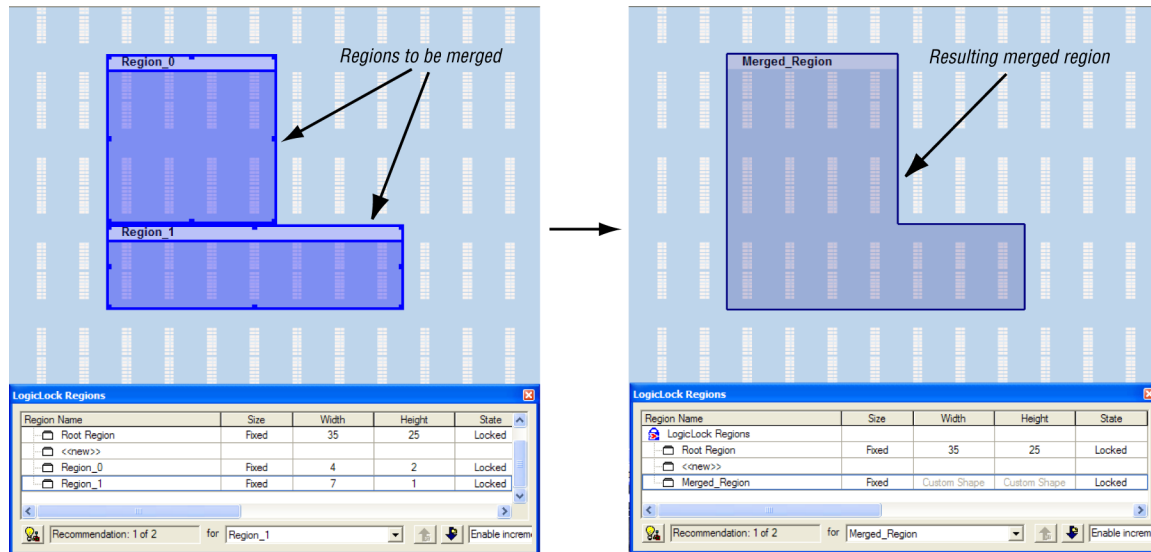
To create a nonrectangular region with the **Merge LogicLock Region** command, follow these steps:

1. In the Chip Planner, create two or more contiguous or non-contiguous rectangular regions.
2. Arrange the regions that you have created into the locations where you want the nonrectangular region to be.
3. Select all the individual regions that you want to merge by clicking each of them while pressing the Shift key.
4. Right-click the title bar of any of the LogicLock regions that you want to merge, point to **LogicLock Regions**, and then click **Merge LogicLock Region**. The individual regions that you select merge to create a single new region.

By default, the new LogicLock region has the same name as the component region containing the greatest number of resources; however, you can rename the new region. In the **LogicLock Regions Window**, the new region is shown as having a **Custom Shape**.

You can use the **Merge LogicLock Region** command to form a nonrectangular LogicLock region by merging two rectangular LogicLock regions.

**Figure 15-1: Using the Merge LogicLock Region command to create a nonrectangular region**



### Related Information

[Creating LogicLock Regions](#) on page 15-4

## Hierarchical (Parent and Child) LogicLock Regions

To further constrain module locations, you can define a hierarchy for a group of regions by declaring parent and child regions.

The Quartus II software places a child region completely within the boundaries of its parent region; a child region must be placed entirely within the boundary of its parent. Additionally, parent and child regions allow you to further improve the performance of a module by constraining nodes in the critical path of a module.

To make one LogicLock region a child of another LogicLock region, in the LogicLock Regions window, select the new child region and drag and drop the new child region into its new parent region.

**Note:** The LogicLock region hierarchy does not have to be the same as the design hierarchy.

You can create both auto-sized and fixed-sized LogicLock regions within a parent LogicLock region; however, the parent of a fixed-sized child region must also be fixed-sized. The location of a locked parent region is locked relative to the device; the location of a locked child region is locked relative to its parent region. If you change the parent's location, the locked child's origin changes, but maintains the same placement relative to the origin of its parent. The location of a floating child region can float within its parent. Complex region hierarchies might result in some LABs not being used, effectively increasing the resource utilization in the device. Do not create more levels of hierarchy than you need.

## Placing LogicLock Regions

A fixed region must contain all resources required by the design block assigned to the region. Although the Quartus II software can automatically place and size LogicLock regions to meet resource and timing requirements, you can manually place and size regions to meet your design requirements.

You should consider the following if you manually place or size a LogicLock region:

- LogicLock regions with pin assignments must be placed on the periphery of the device, adjacent to the pins. For the Arria series, Cyclone series, Stratix series, MAX II, and MAX V devices, you must also include the I/O block within the LogicLock Region.
- Floating LogicLock regions can overlap with their ancestors or descendants, but not with other floating LogicLock regions.

## Placing Device Resources into LogicLock Regions

A LogicLock region includes all device resources within its boundaries, including memory and pins.

The Quartus II software does not include pins automatically when you assign an entity to a region—you can manually assign pins to LogicLock regions; however, this placement puts location constraints on the region. The software only obeys pin assignments to locked regions that border the periphery of the device. For the Arria series, Cyclone series, Stratix series, MAX II, and MAX V devices, the locked regions must include the I/O pins as resources.

**Note:** Pin assignments to LogicLock regions are effective only in fixed and locked regions. Pin assignments to floating regions do not influence the placement of the region.

Only one LogicLock region can claim a device resource. If a LogicLock region boundary includes part of a device resource, the Quartus II software allocates the entire resource to that LogicLock region. When the Quartus II software places a floating auto-sized region, it places the region in an area that meets the requirements of the contents of the LogicLock region.

**Note:** If you want to import multiple instances of a module into a top-level design, you must ensure that the device has two or more locations with exactly the same device resources. (You can determine this from the applicable device handbook.) If the device does not have another area with exactly the same resources, the Quartus II software generates a fitting error during compilation of the top-level design.


## LogicLock Regions Window

You can use the LogicLock Regions window to create LogicLock regions, assign nodes and entities to them, and modify the properties of a LogicLock region such as size, state, width, height, origin, and whether the region is a reserved region.

The LogicLock Regions window also has a recommendations toolbar; select a LogicLock region from the drop-down list in the recommendations toolbar to display the relevant suggestions to optimize that LogicLock region. You can customize the LogicLock Regions window by dragging and dropping the

columns to change their order; you can also show and hide optional columns by right-clicking any column heading and then selecting the appropriate columns in the shortcut menu.

**Figure 15-2: LogicLock Regions Window**



Region Name	Size	Width	Height	State	Origin	Reserved	Enabled	Members
LogicLock Regions								
Root Region	Fixed	54	52	Locked	X0_Y0	Off	Enabled	None
<<new>>								
chiptrip	Auto	4	2	Floating	X32_Y4	Off	Enabled	chiptrip
auto_max:auto	Auto	1	1	Floating	X33_Y4	Off	Enabled	auto_max:auto
speed_ch:speed	Auto	1	1	Floating	X32_Y4	Off	Enabled	speed_ch:speed
tick_cnt:tick	Auto	1	1	Floating	X32_Y5	Off	Enabled	tick_cnt:tick
time_cnt:time_c	Auto	1	1	Floating	X10_Y34	Off	Enabled	time_cnt:time_c

Recommendation: 2 of 2 for chiptrip  
Tighten region "chiptrip" to improve utilization level. Utilization is < 70%. Details...

The **LogicLock Region Properties** dialog box provides a summary of all LogicLock regions in your design. Use the **LogicLock Region Properties** dialog box to obtain detailed information about your LogicLock region, such as which entities and nodes are assigned to your region and which resources are required. The **LogicLock Region Properties** dialog box shows the properties of the current selected regions and allows you to modify them. To open the **LogicLock Region Properties** dialog box, double-click any region in the **LogicLock Regions** window, or right-click the region and click **Properties**.

**Note:** For designs that target Arria series, Cyclone series, Stratix series, MAX II, and MAX V devices, the Quartus II software automatically creates a LogicLock region that encompasses the entire device. This default region is labelled `Root_Region`, and is locked and fixed.

**Note:** For Arria series, Cyclone series, Stratix series, MAX II, and MAX V devices, the origin of the LogicLock region is located at the lower-left corner of the region. For all other supported devices, the origin is located at the upper-left corner of the region.

## Reserved LogicLock Region

The Quartus II software honors all entity and node assignments to LogicLock regions. Occasionally, entities and nodes do not occupy an entire region, which leaves some of the region's resources unoccupied.

To increase the region's resource utilization and performance, the Quartus II software's default behavior fills the unoccupied resources with other nodes and entities that have not been assigned to another region. You can prevent this behavior by turning on **Reserved** on the **LogicLock Region Properties > General** tab. When you turn on this option, your LogicLock region contains only the entities and nodes that you specifically assigned to your LogicLock region.

## Excluded Resources

The Excluded Resources feature allows you to easily exclude specific device resources such as DSP blocks or M4K memory blocks from a LogicLock region.

For example, you can assign a specific entity to a LogicLock region but allow the DSP blocks of that entity to be placed anywhere on the device. Use the Excluded Resources feature on a per-LogicLock region member basis.

To exclude certain device resources from an entity, in the **LogicLock Region Properties** dialog box, highlight the entity in the **Design Element** column, and click **Edit**. In the **Edit Node** dialog box, under



**Excluded Element Types**, click the **Browse** button. In the **Excluded Resources Element Types** dialog box, you can select the device resources you want to exclude from the entity. When you have selected the resources to exclude, the **Excluded Resources** column is updated in the **LogicLock Region Properties** dialog box to reflect the excluded resources.

**Note:** The Excluded Resources feature prevents certain resource types from being included in a region, but it does not prevent the resources from being placed inside the region unless you set the region's **Reserved** property to **On**. To indicate to the Fitter that certain resources are not required inside a LogicLock region, define a resource filter.

#### Related Information

#### [Best Practices for Incremental Compilation Partitions and Floorplan Assignments](#)

For more information about resource filters.

## Additional Quartus II LogicLock Design Features

To complement the **LogicLock Regions** window, the Quartus II software has additional features to help you design with LogicLock regions.

### Analysis and Synthesis Resource Utilization by Entity

The Compilation Report contains an **Analysis and Synthesis Resource Utilization by Entity** section, which reports resource usage statistics, including entity-level information. You can use this feature to verify that any LogicLock region you manually create contains enough resources to accommodate all the entities you assign to it.

### Quartus II Revisions Feature

When you evaluate different LogicLock regions in your design, you might want to experiment with different configurations to achieve your desired results. The Quartus II Revisions feature allows you to organize the same project with different settings until you find an optimum configuration.

To use the Revisions feature, on the Project menu, click **Revisions**. In the **Revisions** dialog box, you can create and specify revisions. You can create a revision from the current design or any previously created revisions. Each revision can have an associated description. You can use revisions to organize the placement constraints created for your LogicLock regions.

### LogicLock Assignment Precedence

You can encounter conflicts during the assignment of entities and nodes to LogicLock regions. For example, an entire top-level entity might be assigned to one region and a node within this top-level entity assigned to another region. To resolve conflicting assignments, the Quartus II software maintains an order of precedence for LogicLock assignments. The following order of precedence, from highest to lowest, applies:

1. Exact node-level assignments
2. Path-based and wildcard assignments
3. Hierarchical assignments

**Note:** To open the **Priority** dialog box, select **LogicLock Regions Properties > General > Priority**. You can change the priority of path-based and wildcard assignments with the **Up** and **Down** buttons in the **Priority** dialog box. To prioritize assignments between regions, you must select multiple LogicLock regions and then open the **Priority** dialog box from the **LogicLock Regions Properties** dialog box.

## Related Information

### Understanding Assignment Priority

For more information about LogicLock assignment precedence.

## Virtual Pins

A virtual pin is an I/O element that is temporarily mapped to a logic element and not to a pin during compilation, and is then implemented as a LUT.

When you apply the Virtual Pin assignment to an input pin, the pin no longer appears as an FPGA pin, but is fixed to GND or VCC in the design. The assigned pin is not an open node.

Virtual pins should be used only for I/O elements in lower-level design entities that become nodes when imported to the top-level design. You can create virtual pins by assigning the Virtual Pin logic option to an I/O element.

You might use virtual pin assignments when you compile a partial design, because not all the I/Os from a partial design drive chip pins at the top level.

The virtual pin assignment identifies the I/O ports of a design module that are internal nodes in the top-level design. These assignments prevent the number of I/O ports in the lower-level modules from exceeding the total number of available device pins. Every I/O port that you designate as a virtual pin becomes mapped to either a logic cell or an adaptive logic module (ALM), depending on the target device.

**Note:** The Virtual Pin logic option must be assigned to an input or output pin. If you assign this option to a bidirectional pin, tri-state pin, or registered I/O element, Analysis and Synthesis ignores the assignment. If you assign this option to a tri-state pin, the Fitter inserts an I/O buffer to account for the tri-state logic; therefore, the pin cannot be a virtual pin. You can use multiplexer logic instead of a tri-state pin if you want to continue to use the assigned pin as a virtual pin. Do not use tri-state logic except for signals that connect directly to device I/O pins.

In the top-level design, you connect these virtual pins to an internal node of another module. By making assignments to virtual pins, you can place those pins in the same location or region on the device as that of the corresponding internal nodes in the top-level module. You can use the **Virtual Pin** option when compiling a LogicLock module with more pins than the target device allows. The **Virtual Pin** option can enable timing analysis of a design module that more closely matches the performance of the module after you integrate it into the top-level design.

**Note:** In the Node Finder, you can set **Filter Type** to **Pins: Virtual** to display all assigned virtual pins in the design. Alternatively, to access the Node Finder from the Assignment Editor, double-click the **To** field; when the arrow appears on the right side of the field, click the arrow and select **Node Finder**.

## Using LogicLock Regions in the Chip Planner

You can easily create LogicLock regions in the Chip Planner and assign resources to them.

## Viewing Connections Between LogicLock Regions in the Chip Planner

You can view and edit LogicLock regions using the Chip Planner. To view and edit LogicLock regions, select the **Floorplan Editing > Layers Settings**, or any Layers setting mode that has the **User-assigned LogicLock regions** setting enabled.

The Chip Planner shows the connections between LogicLock regions. By default, you can view each connection as an individual line. You can choose to display connections between two LogicLock regions

as a single bundled connection rather than as individual connection lines. To use this option, open the Chip Planner and on the View menu, click **Inter-region Bundles**.

#### Related Information

##### [Inter-region Bundles Dialog Box](#)

For more information about the Inter-region Bundles dialog box, refer to Quartus II Help.

## Using LogicLock Regions with the Design Partition Planner

You can optimize timing in a design by placing entities that share significant logical connectivity close to each other on the device.

By default, the Fitter usually places closely connected entities in the same area of the device; however, you can use LogicLock regions, together with the Design Partition Planner and the Chip Planner, to help ensure that logically connected entities retain optimal placement from one compilation to the next.

You can view the logical connectivity between entities with the Design Partition Planner, and the physical placement of those entities with the Chip Planner. In the Design Partition Planner, you can identify entities that are highly interconnected, and place those entities in a partition. In the Chip Planner, you can create LogicLock regions and assign each partition to a LogicLock region, thereby preserving the placement of the entities.

#### Related Information

- [Best Practices for Incremental Compilation Partition and Floorplan Assignments](#)
- [Quartus II Incremental Compilation for Hierarchical and Team-Based Design](#)
- [About the Design Partition Planner](#)

For more information about using LogicLock regions with design partitions

For more information about using the Design Partition Planner with the Chip Planner, refer to Quartus II Help.

## Design Floorplan Analysis Using the Chip Planner

The Chip Planner helps you visually analyze the floorplan of your design at any stage of your design cycle. With the Chip Planner, you can view post-compilation placement, connections, and routing paths.

You can also create LogicLock regions and location assignments. The Chip Planner allows you to create new logic cells and I/O atoms and to move existing logic cells and I/O atoms in your design. You can also see global and regional clock regions within the device, and the connections between I/O atoms, PLLs and the different clock regions.

From the Chip Planner, you can launch the Resource Property Editor, which you can use to change the properties and parameters of device resources, and modify connectivity between certain types of device resources. The Change Manager records any changes that you make to your design floorplan so that you can selectively undo changes if necessary.

The following sections present Chip Planner floorplan views and design analysis procedures which you can use with any Chip Planner preset, unless a procedure requires a specific preset or editing mode.

#### Related Information

- [About the Change Manager](#) on page 17-1

- [About the Resource Property Editor](#) on page 17-1  
For more information about the Resource Property Editor and the Change Manager, refer to Quartus II Help.
- [Engineering Change Management with the Chip Planner](#) on page 17-1  
For more information about the Resource Property Editor and the Change Manager.

## Chip Planner Floorplan Views

The Chip Planner uses a hierarchical zoom viewer that shows various abstraction levels of the targeted Altera device. As you zoom in, the level of abstraction decreases, revealing more details about your design.

### Bird's Eye View

The Bird's Eye View displays a high-level picture of resource usage for the entire chip and provides a fast and efficient way to navigate between areas of interest in the Chip Planner.

The Bird's Eye View is particularly useful when the parts of your design that you want to view are at opposite ends of the chip and you want to quickly navigate between resource elements without losing your frame of reference.

### Properties Window

The Properties Window displays detailed properties of the objects (such as atoms, paths, LogicLock regions, or routing elements) currently selected in the Chip Planner. To display the Properties Window, click **Properties** on the **View** menu in the Chip Planner.

### Related Information

- [Engineering Change Management with the Chip Planner](#) on page 17-1  
For more information about Chip Planner floorplan views.
- [Displaying Resources and Information](#)
- [Bird's Eye View](#)  
For more information about displaying information in the Bird's Eye View.

## Viewing Architecture-Specific Design Information

By adjusting the **Layers Settings** in the Chip Planner, you can view the following architecture-specific information related to your design:

- **Device routing resources used by your design**—View how blocks are connected, as well as the signal routing that connects the blocks.
- **LE configuration**—View logic element (LE) configuration in your design. For example, you can view which LE inputs are used; if the LE utilizes the register, the look-up table (LUT), or both; as well as the signal flow through the LE.
- **ALM configuration**—View ALM configuration in your design. For example, you can view which ALM inputs are used, if the ALM utilizes the registers, the upper LUT, the lower LUT, or all of them. You can also view the signal flow through the ALM.
- **I/O configuration**—View device I/O resource usage. For example, you can view which components of the I/O resources are used, if the delay chain settings are enabled, which I/O standards are set, and the signal flow through the I/O.
- **PLL configuration**—View phase-locked loop (PLL) configuration in your design. For example, you can view which control signals of the PLL are used with the settings for your PLL.
- **Timing**—View the delay between the inputs and outputs of FPGA elements. For example, you can analyze the timing of the `DATAB` input to the `COMBOUT` output.

In addition, you can modify the following device properties with the Chip Planner:

- LEs and ALMs
- I/O cells
- PLLs
- Registers in RAM and DSP blocks
- Connections between elements
- Placement of elements

For more information about LEs, ALMs, and other resources of an FPGA device, refer to the relevant device handbook.

## Viewing Available Clock Networks in the Device

When you select a task with clock region layer preset enabled, you can display the areas of the chip that are driven by global and regional clock networks. This global clock display feature is available for Arria V, Cyclone V, Stratix IV, and Stratix V device families.

Depending on the clock layers activated in the selected preset, the Chip Planner displays regional and global clock regions in the device, and the connectivity between clock regions, pins, and PLLs. Clock regions appear as rectangular overlay boxes with labels indicating the clock type and index. You can select each clock network region by clicking on the clock region. The clock-shaped icon at the top-left corner indicates that the region represents a clock network region. You can change the color in which the Chip Planner displays clock regions on the **Options** dialog box of the Tools menu.

The **Layers Settings** dialog box lists layers for different clock region types; when the selected device does not contain a given clock region, the option for that category is unavailable in the dialog box. You can customize the Chip Planner's display of clock regions by creating a custom preset with selected clock layers enabled in the Layers Settings dialog box.

### Related Information

#### [Displaying Resources and Information](#)

For more information about displaying clock regions.

## Viewing Critical Paths

Critical paths are timing paths in your design that have a negative slack. These timing paths can span from device I/Os to internal registers, registers to registers, or from registers to device I/Os.

The slack of a path determines its criticality; slack appears in the timing analysis report. Design analysis for timing closure is a fundamental requirement for optimal performance in highly complex designs. The analytical capability of the Chip Planner helps you close timing on complex designs.

Viewing critical paths in the Chip Planner helps you understand why a specific path is failing. You can see if any modification in the placement can reduce the negative slack. You can display details of a path (to expand/collapse the path to/from the connections in the path) by clicking **Expand Connections** in the toolbar, or by clicking on the "+/-" on the label.

You can locate failing paths from the timing report in the TimeQuest Timing Analyzer. To locate the critical paths, run the Report Timing task from the Custom Reports group in the Tasks pane of the TimeQuest Timing Analyzer. From the View pane, which lists the failing paths, right-click on any failing path or node, and select **Locate Path**. From the **Locate** dialog box, select **Chip Planner** to see the failing path in the Chip Planner.

**Note:** To display paths in the floorplan, you must first make timing settings and perform a timing analysis.

### Related Information

#### [The Quartus II TimeQuest Timing Analyzer](#)

For more information about performing static timing analysis with the Quartus II TimeQuest Timing Analyzer.

## Viewing Routing Congestion

The **Report Routing Utilization** task allows you to determine the percentage of routing resources in use following a compilation. This feature can identify where there is a lack of routing resources, helping you to make design changes to meet routing congestion design requirements.

Open the Chip Planner from the Tools menu. To view the routing congestion in the Chip Planner, double-click the **Report Routing Utilization** command in the **Tasks** list. Click **Preview** in the **Report Routing Utilization** dialog box to preview the default congestion display. Change the **Routing Utilization Type** to display congestion for specific resources. The default display uses dark blue for 0% congestion (blue indicates zero utilization) and red for 100%. You can adjust the slider for **Threshold percentage** to change the congestion threshold level.

The routing congestion map uses the color and shading of logic resources to indicate relative resource utilization; darker shading represents a greater utilization of routing resources. Areas where routing utilization exceeds the threshold value specified in the **Report Routing Utilization** dialog box appear in red. The congestion map can help you determine whether you can modify the floorplan, or make changes to the RTL to reduce routing congestion.

To identify a lack of routing resources, it is necessary to investigate each routing interconnect type separately by selecting each interconnect type in turn in the **Routing Utilization Settings** dialog box.

The Quartus II compilation messages contain information about average and peak interconnect usage. Peak interconnect usage over 75%, or average interconnect usage over 60%, could be an indication that it might be difficult to fit your design. Similarly, peak interconnect usage over 90%, or average interconnect usage over 75%, are likely to have increased chances of not getting a valid fit.

### Related Information

#### [Displaying Resources and Information](#)

For more information about displaying routing congestion


## Viewing I/O Banks

The Chip Planner can show all of the I/O banks of the device. To see the I/O bank map of the device, select **Report All I/O Banks** in the **Tasks** pane.

## Viewing High-Speed Serial Interfaces (HSSI)

For the Stratix V device family, the Chip Planner displays a detailed block view of the receiver and transmitter channels of the high-speed serial interfaces. To display the HSSI block view, select **Report HSSI Block Connectivity**.

Figure 15-3: Stratix V HSSI Receiver Channel Blocks



## Generating Fan-In and Fan-Out Connections

The ability to display fan-in and fan-out connections enables you to view the atoms that fan-in to or fan-out from the selected atom. To remove the connections displayed, use the **Clear Unselected Connections** icon in the Chip Planner toolbar.


## Generating Immediate Fan-In and Fan-Out Connections

The ability to display immediate fan-in and fan-out connections enables you to view the resource that is the immediate fan-in or fan-out connection for the selected atom. For example, if you select a logic resource and choose to view the immediate fan-in for that resource, you can see the routing resource that drives the logic resource. You can generate immediate fan-ins and fan-outs for all logic resources and routing resources. To remove the displayed connections from the screen, click the **Clear Unselected Connections** icon in the toolbar.

## Highlight Routing

The **Show Physical Routing** command in the **Locate History** pane enables you to highlight the routing resources used by a selected path or connection.

Figure 15-4: Highlight Routing



#### Related Information

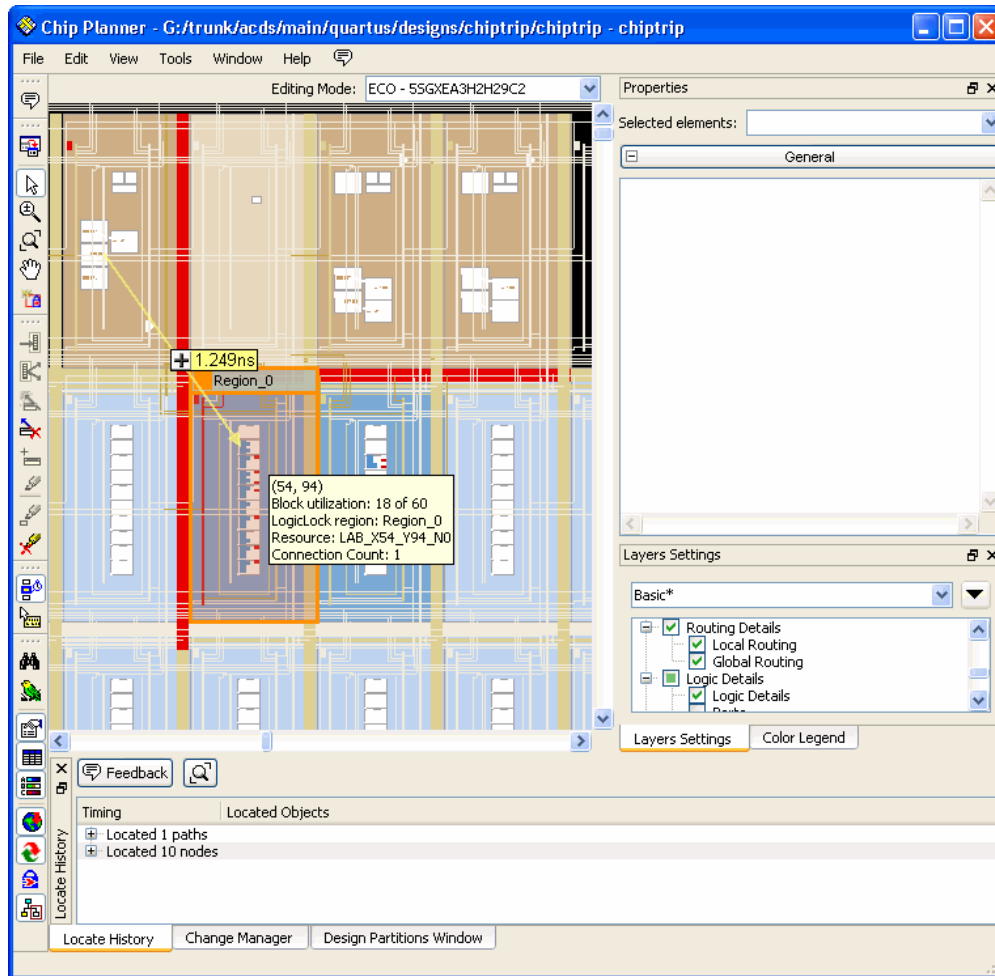
- [Engineering Change Management with the Chip Planner](#) on page 17-1  
For more information on how you can view and edit resources in the FPGA using the Resource Property Editor.

### Show Delays

With the Show Delays command, you can view timing delays for paths located from TimeQuest Timing Analyzer reports. For example, you can view the delay between two logic resources or between a logic resource and a routing resource.



Figure 15-5: Show Delays Associated in a TimeQuest Timing Analyzer Path



## Exploring Paths in the Chip Planner

You can use the Chip Planner to explore paths between logic elements. The following example uses the Chip Planner to traverse paths from the Timing Analysis report.

### Locate Path from the Timing Analysis Report to the Chip Planner


To locate a path from the Timing Analysis report to the Chip Planner, perform the following steps:

1. Select the path you want to locate.
2. Right-click the path in the Timing Analysis report, point to **Locate Path**, and click **Locate in Chip Planner**. The path is displayed with its timing data in the Chip Planner main window and is listed in the Locate History window.
3. To view the routing resources taken for a path you have located in the Chip Planner, select the path and then click the **Highlight Routing** icon in the Chip Planner toolbar, or from the View menu, click **Highlight Routing**.

## Analyzing Connections for a Path

To determine the connections between items in the Chip Planner, click the **Expand Connections** icon on the toolbar. To add the timing delays for paths located from the TimeQuest Timing Analyzer, click the **Show Delays** icon on the toolbar. To see the constituent delays on the selected path, click on the “+” sign next to the path delay displayed in the Chip Planner.

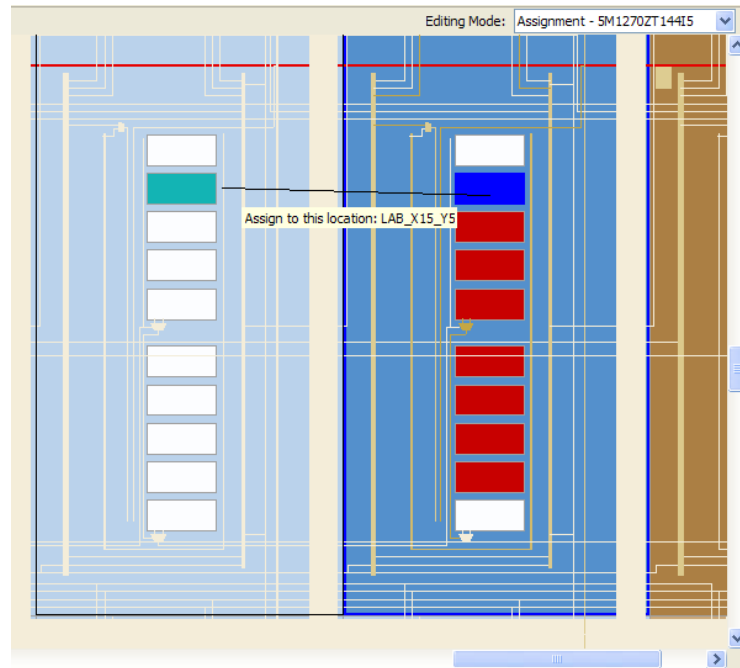
Figure 15-6: Path Analysis in the Chip Planner of a Path Located from TimeQuest



## Viewing Assignments in the Chip Planner

You can view location assignments by selecting the appropriate layer set in the Chip Planner. To view location assignments, select the **Floorplan Editing** preset or any custom preset that displays block utilization, and the Assignment editing mode. The Chip Planner shows location assignments graphically, by displaying assigned resources in a particular color (gray, by default). You can create or move an assignment by dragging the selected resource to a new location.

Figure 15-7: Viewing Assignments in the Chip Planner



You can make node and pin location assignments to LogicLock regions and custom regions using the drag-and-drop method in the Chip Planner. The Fitter applies the assignments that you create during the next place-and-route operation.

#### Related Information

##### [Working With Assignments in the Chip Planner](#)

For more information about managing assignments in the Chip Planner.

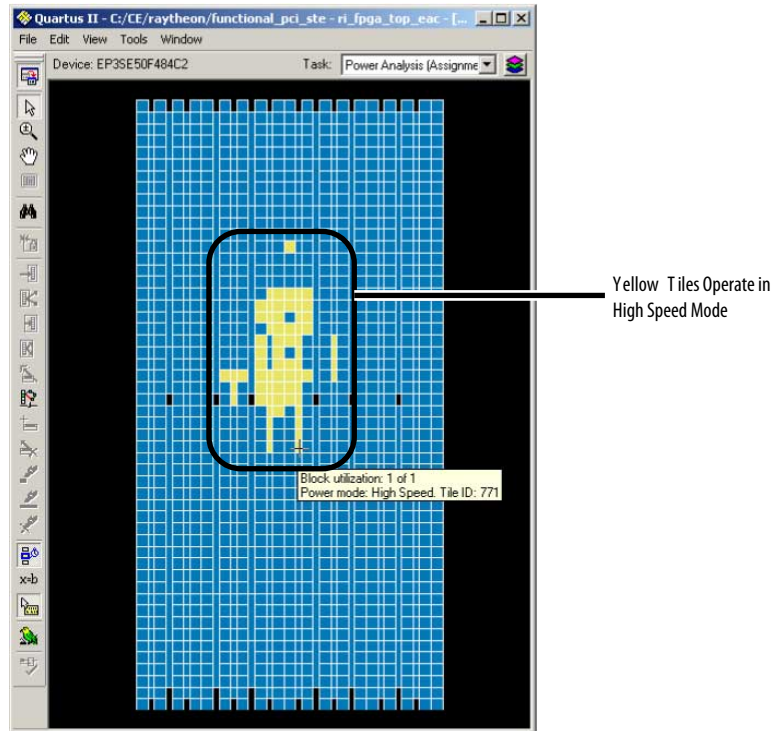
## Viewing High-Speed and Low-Power Tiles in the Chip Planner

To view a power map of designs that specify Stratix IV or Stratix V devices, select **Tasks > Report High-Speed/Low-Power Tiles** after running the Fitter.

Stratix IV, or Stratix V devices have ALMs that can operate in either high-speed mode or low-power mode. The power mode is set during the fitting process in the Quartus II software. These ALMs are grouped together to form larger blocks, called “tiles.”

When you select the **Report High-Speed/Low-Power Tiles** command for Stratix IV or Stratix V devices, the Chip Planner displays low-power and high-speed tiles in contrasting colors; yellow tiles operate in a high-speed mode, while blue tiles operate in a low-power mode. When you select the **Power** task, you can perform all floorplanner-related functions for this task; however, you cannot edit tiles to change the power mode.

Figure 15-8: Viewing High-Speed and Low Power Tiles in a Stratix Device



**Related Information**

[AN 514: Power Optimization in Stratix IV FPGAs](#)

To learn more about power analyses and optimizations in Stratix IV devices.

## Scripting Support

You can run procedures and specify the settings described in this chapter in a Tcl script. You can also run some procedures at a command prompt.

For detailed information about scripting command options, refer to the Quartus II command-line and Tcl API Help browser. To run the Help browser, type the following command at the command prompt:

```
quartus_sh --qhelp
```

**Related Information**

[API Functions for Tcl](#) on page 3-1

For more information about Tcl scripting in Quartus II Help.

[Tcl Scripting](#) on page 3-1

[API Functions for Tcl](#) For more information about Tcl scripting .

[Command-Line Scripting](#) on page 2-1

For more information about command-line scripting.

[Quartus II Settings File Manual](#)

For information about all settings and constraints in the Quartus II software.

## Initializing and Uninitializing a LogicLock Region

You must initialize the LogicLock data structures before creating or modifying any LogicLock regions and before executing any of the Tcl commands listed below.

Use the following Tcl command to initialize the LogicLock data structures:

```
initialize_logiclock
```

Use the following Tcl command to uninitialize the LogicLock data structures before closing your project:

```
uninitialize_logiclock
```

## Creating or Modifying LogicLock Regions

Use the following Tcl command to create or modify a LogicLock region:

```
set_logiclock -auto_size true -floating true -region <my_region-name>
```

**Note:** The command in the above example sets the size of the region to auto and the state to floating.

If you specify a region name that does not exist in the design, the command creates the region with the specified properties. If you specify the name of an existing region, the command changes all properties you specify and leaves unspecified properties unchanged.

### Related Information

[Creating LogicLock Regions](#) on page 15-4

## Obtaining LogicLock Region Properties

Use the following Tcl command to obtain LogicLock region properties. This example returns the height of the region named `my_region`:

```
get_logiclock -region my_region -height
```

## Assigning LogicLock Region Content

Use the following Tcl commands to assign or change nodes and entities in a LogicLock region. This example assigns all nodes with names matching `fifo*` to the region named `my_region`.

```
set_logiclock_contents -region my_region -to fifo*
```

You can also make path-based assignments with the following Tcl command:

```
set_logiclock_contents -region my_region -from fifo -to ram*
```

## Save a Node-Level Netlist for the Entire Design into a Persistent Source File

Make the following assignments to cause the Quartus II Fitter to save a node-level netlist for the entire design into a `.vqm` file:

```
set_global_assignment -name LOGICLOCK_INCREMENTAL_COMPILE_ASSIGNMENT ON  
set_global_assignment -name LOGICLOCK_INCREMENTAL_COMPILE_FILE <file name>
```

Any path specified in the file name is relative to the project directory. For example, specifying **atom\_netlists/top.vqm** places **top.vqm** in the **atom\_netlists** subdirectory of your project directory.

A **.vqm** file is saved in the directory specified at the completion of a full compilation.

**Note:** The saving of a node-level netlist to a persistent source file is not supported for designs targeting newer devices such as MAX V, Stratix IV, or Stratix V.

## Setting LogicLock Assignment Priority

Use the following Tcl code to set the priority for a LogicLock region's members. This example reverses the priorities of the LogicLock region in your design.

```
set reverse [list]
for each member [get_logiclock_member_priority] {
    set reverse [insert $reverse 0 $member]
}
set_logiclock_member_priority $reverse
```

## Assigning Virtual Pins

Use the following Tcl command to turn on the virtual pin setting for a pin called `my_pin`:

```
set_instance_assignment -name VIRTUAL_PIN ON -to my_pin
```

### Related Information

[Virtual Pins](#) on page 15-9

[Tcl Scripting](#) on page 3-1

For more information about Tcl scripting.

## Document Revision History

Table 15-2: Document Revision History

Date	Version	Changes
2014.12.15	14.1.0	Updated description of Virtual Pins assignment to clarify that assigned input is not available.
June 2014	14.0.0	Updated format
November 2013	13.1.0	Removed HardCopy device information.
May 2013	13.0.0	Updated "Viewing Routing Congestion" section Updated references to Quartus UI controls for the Chip Planner
June 2012	12.0.0	Removed survey link.
November 2011	11.0.1	Template update.

Date	Version	Changes
May 2011	11.0.0	<ul style="list-style-type: none"> <li>• Updated for the 11.0 release.</li> <li>Edited “LogicLock Regions”</li> <li>Updated “Viewing Routing Congestion”</li> <li>Updated “Locate History”</li> <li>Updated Figures 15-4, 15-9, 15-10, and 15-13</li> <li>Added Figure 15-6</li> </ul>
December 2010	10.1.0	<ul style="list-style-type: none"> <li>• Updated for the 10.1 release.</li> </ul>
July 2010	10.0.0	<ul style="list-style-type: none"> <li>• Updated device support information</li> <li>• Removed references to Timing Closure Floorplan; removed “Design Analysis Using the Timing Closure Floorplan” section</li> <li>• Added links to online Help topics</li> <li>• Added “Using LogicLock Regions with the Design Partition Planner” section</li> <li>• Updated “Viewing Critical Paths” section</li> <li>• Updated several graphics</li> <li>• Updated format of Document revision History table</li> </ul>
November 2009	9.1.0	<ul style="list-style-type: none"> <li>• Updated supported device information throughout</li> <li>• Removed deprecated sections related to the Timing Closure Floorplan for older device families. (For information on using the Timing Closure Floorplan with older device families, refer to previous versions of the Quartus II Handbook, available in the Quartus II Handbook Archive.)</li> <li>• Updated “Creating Nonrectangular LogicLock Regions” section</li> <li>• Added “Selected Elements Window” section</li> <li>• Updated table 12-1</li> </ul>

Date	Version	Changes
May 2008	8.0.0	<ul style="list-style-type: none"> <li>• Updated the following sections:               <ul style="list-style-type: none"> <li>“Chip Planner Tasks and Layers”</li> <li>“LogicLock Regions”</li> <li>“Back-Annotating LogicLock Regions”</li> <li>“LogicLock Regions in the Timing Closure Floorplan”</li> </ul> </li> <li>• Added the following sections:               <ul style="list-style-type: none"> <li>“Reserve LogicLock Region”</li> <li>“Creating Nonrectangular LogicLock Regions”</li> <li>“Viewing Available Clock Networks in the Device”</li> </ul> </li> <li>• Updated Table 10-1</li> <li>• Removed the following sections:               <ul style="list-style-type: none"> <li>Reserve LogicLock Region Design Analysis Using the Timing Closure Floorplan</li> </ul> </li> </ul>

**Related Information**

[Quartus II Handbook Archive](#)

For previous versions of the Quartus II Handbook.



# Netlist Optimizations and Physical Synthesis 16

2014.12.15

QI15V2



Subscribe



Send Feedback

## Netlist Optimizations and Physical Synthesis

The Quartus<sup>®</sup> II software offers physical synthesis optimizations to improve your design beyond the optimization performed in the normal course of the Quartus II compilation flow.

Physical synthesis optimizations can help improve the performance of your design regardless of the synthesis tool used, although the effect of physical synthesis optimizations depends on the structure of your design.

Netlist optimization options work with the atom netlist of your design, which describes a design in terms of Altera<sup>®</sup>-specific primitives. An atom netlist file can be an Electronic Design Interchange Format (**.edf**) file or a Verilog Quartus Mapping (**.vqm**) file generated by a third-party synthesis tool, or a netlist used internally by the Quartus II software. Physical synthesis optimizations are applied at different stages of the Quartus II compilation flow, either during synthesis, fitting, or both.

This chapter explains how the physical synthesis optimizations in the Quartus II software can modify your design's netlist to improve the quality of results. This chapter also provides information about preserving compilation results through back-annotation and writing out a new netlist, and provides guidelines for applying the various options.

**Note:** Because the node names for primitives in the design can change when you use physical synthesis optimizations, you should evaluate whether your design flow requires fixed node names. If you use a verification flow that might require fixed node names, such as the SignalTap<sup>®</sup> II Logic Analyzer, formal verification, or the LogicLock based optimization flow (for legacy devices), you must turn off physical synthesis options.

## WYSIWYG Primitive Resynthesis

If you use a third-party tool to synthesize your design, use the **Perform WYSIWYG primitive resynthesis** option to apply optimizations to the synthesized netlist.

The **Perform WYSIWYG primitive resynthesis** option directs the Quartus II software to un-map the logic elements (LEs) in an atom netlist to logic gates, and then re-map the gates back to Altera-specific primitives. Third-party synthesis tools generate either an **.edf** or **.vqm** atom netlist file using Altera-specific primitives. When you turn on the **Perform WYSIWYG primitive resynthesis** option, the Quartus II software can work on different techniques specific to the device architecture during the re-mapping

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



process. This feature re-maps the design using the **Optimization Technique** specified for your project (**Speed**, **Area**, or **Balanced**).

The **Perform WYSIWYG primitive resynthesis** option unmaps and remaps only logic cells, also referred to as LCELL or LE primitives, and regular I/O primitives (which may contain registers). Double data rate (DDR) I/O primitives, memory primitives, digital signal processing (DSP) primitives, and logic cells in carry/cascade chains are not remapped. Logic specified in an encrypted **.vqm** file or an **.edf** file, such as third-party intellectual property (IP), is not touched.

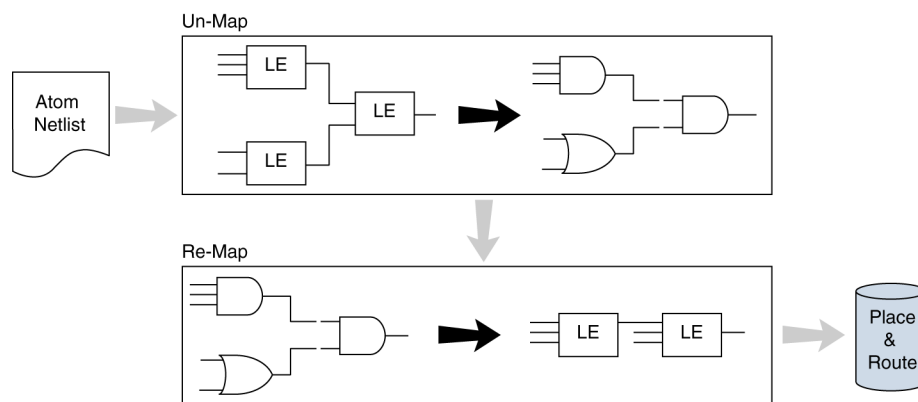
The **Perform WYSIWYG primitive resynthesis** option can change node names in the **.vqm** file or **.edf** file from your third-party synthesis tool, because the primitives in the atom netlist are broken apart and then remapped by the Quartus II software. The remapping process removes duplicate registers, but registers that are not removed retain the same name after remapping.

Any nodes or entities that have the **Netlist Optimizations** logic option set to **Never Allow** are not affected during WYSIWYG primitive resynthesis. You can use the Assignment Editor to apply the **Netlist Optimizations** logic option. This option disables WYSIWYG resynthesis for parts of your design.

**Note:** Primitive node names are specified during synthesis. When netlist optimizations are applied, node names might change because primitives are created and removed. HDL attributes applied to preserve logic in third-party synthesis tools cannot be maintained because those attributes are not written into the atom netlist read by the Quartus II software.

If you use the Quartus II software to synthesize, you can use the **Preserve Register (preserve)** and **Keep Combinational Logic (keep)** attributes to maintain certain nodes in the design.

Figure 16-1: Quartus II Flow for WYSIWYG Primitive Resynthesis



## Perform WYSIWYG Primitive Resynthesis

**Note:** The **Perform WYSIWYG primitive resynthesis** option has no effect if you are using Quartus II integrated synthesis to synthesize your design.

To turn on the **Perform WYSIWYG primitive resynthesis** option, click **Assignments > Settings > Compiler Settings > Advanced Settings (Synthesis)**.

### Related Information

#### [Quartus II Integrated Synthesis](#)

For more information on optimization techniques.

## Performing Physical Synthesis Optimizations

The Quartus II design flow involves separate steps of synthesis and fitting.

The synthesis step optimizes the logical structure of a circuit for area, speed, or both. The Fitter then places and routes the logic cells to ensure critical portions of logic are close together and use the fastest possible routing resources. While you are using this push-button flow, the synthesis stage is unable to anticipate the routing delays seen in the Fitter. Because routing delays are a significant part of the typical critical path delay, the physical synthesis optimizations available in the Quartus II software take those routing delays into consideration and focus timing-driven optimizations at those parts of the design. This tight integration of the fitting and synthesis processes is known as physical synthesis.

The following sections describe the physical synthesis optimizations available in the Quartus II software, and how they can help improve your performance results. Physical synthesis optimization options can be used with Arria series, Cyclone, and Stratix series device families.

**Note:** You cannot target optimizations to both device architectures individually because doing so results in a different post-fitting netlist for each device.

To choose physical synthesis optimization options, click **Assignments > Settings > Compiler Settings > Advanced Settings (Fitter)**. The settings include optimizations for improving performance and fitting in the selected device.

You can also set the effort level for physical synthesis optimizations. Normally, physical synthesis optimizations increase the compilation time; however, you can select the **Fast** effort level if you want to limit the increase in compilation time. When you select the **Fast** effort level, the Quartus II software performs limited register retiming operations during fitting. The **Extra** effort level runs additional algorithms to get the best circuit performance, but results in increased compilation time.

To optimize performance, the following options are available:

- **Perform physical synthesis for combinational logic for fitting**
- **Perform register retiming for performance**
- **Perform asynchronous signal pipelining**
- **Perform register duplication for performance**

To optimize for better fitting, you can choose from the following options:

- **Perform physical synthesis for combinational logic for fitting**

Nodes or entities that have the **Netlist Optimizations** logic option set to **Never Allow** are not affected by the physical synthesis algorithms. You can use the Assignment Editor to apply the **Netlist Optimizations** logic option. Use this option to disable physical synthesis optimizations for parts of your design.

### Related Information

#### [Compiler Settings Page \(Settings Dialog Box\)](#)

For more information about physical synthesis optimizations, refer to Quartus II Help.

## View and Modify Physical Synthesis Optimization Options

Some physical synthesis options affect only registered logic and some options affect only combinational logic. Select options based on whether you want to keep the registers intact or not. For example, if your

verification flow involves formal verification, you might have to keep the registers intact. To view and modify physical synthesis optimization options, perform the following steps:

1. On the Assignments menu, click **Settings**. The **Settings** dialog box appears.
2. Click **Compiler Settings > Advanced Settings (Fitter)**.
3. Specify the options for performing physical synthesis optimizations.

## Viewing Netlist Optimization Reports

All Physical Synthesis optimizations write results to the **Netlist Optimizations** report, which provides a list of atom netlist files that were modified, created, and deleted during physical synthesis.

To access the **Netlist Optimizations** report, perform the following steps:

1. On the Processing menu, click **Compilation Report**.
2. In the **Compilation Report** list, select **Netlist Optimizations** under **Fitter**.

## Viewing Synthesis Reports

Physical synthesis optimizations performed during synthesis write results to the synthesis report. To access this report, perform the following steps:

1. On the Processing menu, click **Compilation Report**.
2. In the **Compilation Report** list, select **Analysis & Synthesis**.
3. In the **Optimization Results** folder, select **Netlist Optimizations**. The **Physical Synthesis Netlist Optimizations** table appears, listing the physical synthesis netlist optimizations performed during synthesis.

## Automatic Asynchronous Signal Pipelining

The Quartus II Fitter to perform automatic insertion of pipeline stages for asynchronous clear and asynchronous load signals during fitting.

To enable **Perform asynchronous signal pipelining**, click **Assignments > Settings > Compiler Settings > Advanced Settings (Fitter)**. The Quartus II Fitter performs automatic insertion of pipeline stages for asynchronous clear and asynchronous load signals during fitting when these signals negatively affect performance. You can use this option if asynchronous control signal recovery and removal times are not achieving their requirements.

The **Perform automatic asynchronous signal pipelining** option improves performance for designs in which asynchronous signals in very fast clock domains cannot be distributed across the chip fast enough due to long global network delays. This optimization performs automatic pipelining of these signals, while attempting to minimize the total number of registers inserted.

**Note:** The **Perform automatic asynchronous signal pipelining** option adds registers to nets driving the asynchronous clear or asynchronous load ports of registers. These additional registers add register delays (adds latency) to the reset, adding the same number of register delays for each destination using the reset. The additional register delays can change the behavior of the signal in the design; therefore, you should use this option only if additional latency on the reset signals does not violate any design requirements. This option also prevents the promotion of signals to global routing resources.

The Quartus II software performs automatic asynchronous signal pipelining only if **Enable Recovery/Removal analysis** is turned on. If you use the TimeQuest Timing Analyzer, **Enable Recovery/Removal**

**analysis** is turned on by default. Pipelining is allowed only on asynchronous signals that have the following properties:

- The asynchronous signal is synchronized to a clock (a synchronization register drives the signal)
- The asynchronous signal fans-out only to asynchronous control ports of registers

The Quartus II software does not perform automatic asynchronous signal pipelining on asynchronous signals that have the **Netlist Optimization** logic option set to **Never Allow**.

## Physical Synthesis for Combinational Logic

To optimize the design and reduce delay along critical paths, you can turn on the **Perform physical synthesis for combinational logic** option.

This option swaps the look-up table (LUT) ports within LEs so that the critical path has fewer layers through which to travel. The **Perform physical synthesis for combinational logic** option also allows the duplication of LUTs to enable further optimizations on the critical path.

The **Perform physical synthesis for combinational logic** option affects only combinational logic in the form of LUTs. These transformations might occur during the synthesis stage or the Fitter stage during compilation. The registers contained in the affected logic cells are not modified. Inputs into memory blocks, DSP blocks, and I/O elements (IOEs) are not swapped.

The Quartus II software does not perform combinational optimization on logic cells that have the following properties:

- Are part of a chain
- Drive global signals
- Are constrained to a single logic array block (LAB) location
- Have the **Netlist Optimizations** option set to **Never Allow**

If you want to consider logic cells with any of these conditions for physical synthesis, you can override these rules by setting the **Netlist Optimizations** logic option to **Always Allow** on a given set of nodes.

### Related Information

#### [Setting Up and Running the Fitter](#)

For more information about using the Perform physical synthesis for combinational logic option, refer to Quartus II Help.

## Physical Synthesis for Registers—Register Duplication

The **Perform register duplication** option duplicates registers based on Fitter placement information. You can also duplicate combinational logic when this option is enabled. A logic cell that fans out to multiple locations can be duplicated to reduce the delay of one path without degrading the delay of another. The new logic cell can be placed closer to critical logic without affecting the other fan-out paths of the original logic cell.

The Quartus II software does not perform register duplication on logic cells that have the following properties:

- Are part of a chain
- Contain registers that drive asynchronous control signals on another register
- Contain registers that drive the clock of another register
- Contain registers that drive global signals
- Contain registers that are constrained to a single LAB location
- Contain registers that are driven by input pins without a  $t_{SU}$  constraint

- Contain registers that are driven by a register in another clock domain
- Are considered virtual I/O pins
- Have the **Netlist Optimizations** option set to **Never Allow**

#### Related Information

#### Setting Up and Running the Fitter


[Analyzing and Optimizing the Design Floorplan](#) on page 15-1

For more information about virtual I/O pins.

### Physical Synthesis for Registers—Register Retiming


The **Perform Register Retiming** option enables the movement of registers across combinational logic, allowing the Quartus II software to trade off the delay between timing-critical paths and non-critical paths. Register retiming can be done during Quartus II integrated synthesis or during the Fitter stages of design compilation.

**Figure 16-2: Reducing Critical Delay by Moving the Register Relative to Combinational Logic**



Retiming can create multiple registers at the input of a combinational block from a register at the output of a combinational block. In this case, the new registers have the same clock and clock enable. The asynchronous control signals and power-up level are derived from previous registers to provide equivalent functionality. Retiming can also combine multiple registers at the input of a combinational block to a single register.

**Figure 16-3: Combining Registers with Register Retiming**



To move registers across combinational logic to balance timing, click **Assignments > Settings > Compiler Settings > Advanced Settings (Fitter)**. Specify your preferred option under **Optimize for performance (physical synthesis)** and **Effort level**.

## Preventing Register Movement During Retiming

If you want to prevent register movement during register retiming, you can set the **Netlist Optimizations** logic option to **Never Allow**. You can apply this option to either individual registers or entities in the design using the Assignment Editor.

In digital circuits, synchronization registers are instantiated on cross clock domain paths to reduce the possibility of metastability. The Quartus II software detects such synchronization registers and does not move them, even if register retiming is turned on.

The following sets of registers are not moved during register retiming:

- Both registers in a direct connection from input pin-to-register-to-register if both registers have the same clock and the first register does not fan-out to anywhere else. These registers are considered synchronization registers.
- Both registers in a direct connection from register-to-register if both registers have the same clock, the first register does not fan out to anywhere else, and the first register is fed by another register in a different clock domain (directly or through combinational logic). These registers are considered synchronization registers.

The Quartus II software does not perform register retiming on logic cells that have the following properties:

- Are part of a cascade chain
- Contain registers that drive asynchronous control signals on another register
- Contain registers that drive the clock of another register
- Contain registers that drive a register in another clock domain
- Contain registers that are driven by a register in another clock domain

**Note:** The Quartus II software does not usually retime registers across different clock domains; however, if you use the Classic Timing Analyzer and specify a global  $f_{MAX}$  requirement, the Quartus II software interprets all clocks as related. Consequently, the Quartus II software might try to retime register-to-register paths associated with different clocks.

To avoid this circumstance, provide individual  $f_{MAX}$  requirements to each clock when using Classic Timing Analysis. When you constrain each clock individually, the Quartus II software assumes no relationship between different clock domains and considers each clock domain to be asynchronous to other clock domains; hence no register-to-register paths crossing clock domains are retimed.

When you use the TimeQuest Timing Analyzer, register-to-register paths across clock domains are never retimed, because the TimeQuest Timing Analyzer treats all clock domains as asynchronous to each other unless they are intentionally grouped.

- Contain registers that are constrained to a single LAB location
- Contain registers that are connected to SERDES
- Are considered virtual I/O pins
- Registers that have the **Netlist Optimizations** logic option set to **Never Allow**

The Quartus II software assumes that a synchronization register chain consists of two registers. If your design has synchronization register chains with more than two registers, you must indicate the number of

registers in your synchronization chains so that they are not affected by register retiming. To do this, perform the following steps:

1. Click **Assignments > Settings > Compiler Settings > Advanced Settings (Synthesis)**.
2. Modify the **Synchronization Register Chain Length** setting to match the synchronization register length used in your design. If you set a value of 1 for the **Synchronization Register Chain Length**, it means that any registers connected to the first register in a register-to-register connection can be moved during retiming. A value of  $n > 1$  means that any registers in a sequence of length 1, 2, ...  $n$  are not moved during register retiming.

If you want to consider logic cells that meet any of these conditions for physical synthesis, you can override these rules by setting the **Netlist Optimizations** logic option to **Always Allow** on a given set of registers.

#### Related Information

- [Analyzing and Optimizing the Design Floorplan](#) on page 15-1  
For more information about virtual I/O pins.

## Preserving Your Physical Synthesis Results

The Quartus II software generates the same results on every compilation for the same source code and settings on a given system, hence you do not need to preserve your results from compilation to compilation.

When you make changes to the source code or to the settings, you usually get the best results by allowing the software to compile without using previous compilation results or location assignments. In some cases, if you avoid performing analysis and synthesis or `quartus_map`, and run the Fitter or another desired Quartus II executable instead, you can skip the synthesis stage of the compilation.

When you use the Quartus II incremental compilation flow, you can preserve synthesis results for a particular partition of your design by choosing a netlist type of post-synthesis. If you want to preserve fitting results between compilation runs, choose a netlist type of post-fit during incremental compilation.

#### Related Information

- [Quartus II Incremental Compilation for Hierarchical and Team-Based Design](#)
- [About Incremental Compilation](#)  
For information about the incremental compilation design methodology.

## Node Preservation for Older Device Families

You can preserve the resulting nodes from physical synthesis in older devices that do not support incremental compilation. You might need to preserve nodes if you use the LogicLock flow to back-annotate placement, import one design into another, or both. For all device families that support incremental compilation, use that feature to preserve results.

The **Save a node-level netlist of the entire design into a persistent source file** option saves your final results as an atom-based netlist in `.vqm` file format. By default, the Quartus II software places the `.vqm` in the `atom_netlists` directory under the current project directory. To create a different `.vqm` file using different Quartus II settings, in the **Compilation Process Settings** page, change the **File name** setting.

If you use the physical synthesis optimizations and want to lock down the location of all LEs and other device resources in the design with the **Back-Annotate Assignments** command, a `.vqm` file netlist is required. The `.vqm` file preserves the changes that you made to your original netlist. Because the physical synthesis optimizations depend on the placement of the nodes in the design, back-annotating the



placement changes the results from physical synthesis. Changing the results means that node names are different, and your back-annotated locations are no longer valid.

You should not use a Quartus II-generated **.vqm** file or back-annotated location assignments with physical synthesis optimizations unless you have finalized the design. Making any changes to the design invalidates your physical synthesis results and back-annotated location assignments. If you require changes later, use the new source HDL code as your input files, and remove the back-annotated assignments corresponding to the Quartus II-generated **.vqm** file.

To back-annotate logic locations for a design that was compiled with physical synthesis optimizations, first create a **.vqm** file. When recompiling the design with the hard logic location assignments, use the new **.vqm** file as the input source file and turn off the physical synthesis optimizations for the new compilation.

If you are importing a **.vqm** file and back-annotated locations into another project that has any **Netlist Optimizations** turned on, you must apply the **Never Allow** constraint to make sure node names don't change; otherwise, the back-annotated location or LogicLock assignments are invalid.

To preserve the nodes from Quartus II physical synthesis optimization options for older devices that do not support incremental compilation (such as Max II devices), perform the following steps:

1. On the Assignments menu, click **Settings**. The **Settings** dialog box appears.
2. In the **Category** list, select **Compilation Process Settings**. The **Compilation Process Settings** page appears.
3. Turn on **Save a node-level netlist of the entire design into a persistent source file**. This setting is not available for Cyclone III, Stratix III, and newer devices.

Click **OK**.

**Note:** For newer devices use incremental compilation to preserve compilation results instead of using logic back-annotation.

## Physical Synthesis Options for Fitting

The Quartus II software provides physical synthesis optimization options for improving fitting results. To access these options, To choose physical synthesis optimization options, click **Assignments > Settings > Compiler Settings > Advanced Settings (Fitter)**.

Table 16-1: Physical Synthesis Optimizations Options

Option	Function
<b>Physical Synthesis for Combinational Logic for Fitting</b>	When you select this option, the Fitter detects duplicate combinational logic and optimizes combinational logic to improve the fit.
<b>Perform Logic to Memory Mapping</b>	When you select this option, the Fitter can remap registers and combinational logic in your design into unused memory blocks and achieves a fit.

## Applying Netlist Optimization Options

The improvement in performance when using netlist optimizations is design dependent. If you have restructured your design to balance critical path delays, netlist optimizations might yield minimal improvement in performance.

You may have to experiment with available options to see which combination of settings works best for a particular design. Refer to the messages in the compilation report to see the magnitude of improvement with each option, and to help you decide whether you should turn on a given option or specific effort level.

Turning on more netlist optimization options can result in more changes to the node names in the design; bear this in mind if you are using a verification flow, such as the SignalTap II Logic Analyzer or formal verification that requires fixed or known node names.

Applying all of the physical synthesis options at the **Extra** effort level generally produces the best results for those options, but adds significantly to the compilation time. You can also use the **Physical synthesis effort level** options to decrease the compilation time. The WYSIWYG primitive resynthesis option does not add much compilation time relative to the overall design compilation time.

To find the best results, you can use the Quartus II Design Space Explorer II (DSE) to apply various sets of netlist optimization options.

#### Related Information

##### [About Design Space Explorer II](#)

For more information about DSE II, refer to Quartus II Help.

## Scripting Support

You can run procedures and make settings described in this chapter in a Tcl script. You can also run some procedures at a command prompt. For detailed information about scripting command options, refer to the Quartus II Command-Line and Tcl API Help browser. To run the Help browser, type the following command at the command prompt:

```
quartus_sh --qhelp
```

You can specify many of the options described in this section on either an instance or global level, or both. Use the following Tcl command to make a global assignment:

```
set_global_assignment -name <QSF variable name> <value>
```

Use the following Tcl command to make an instance assignment:

```
set_instance_assignment -name <QSF variable name> <value> \  
-to <instance name>
```

#### Related Information

- [Tcl Scripting](#) on page 3-1
- [API Functions for Tcl](#)  
For more information about Tcl scripting.
- [Command-Line Scripting](#) on page 2-1  
For more information about command-line scripting, refer to Quartus II Help.
- [Quartus II Settings File Manual](#)  
For information about all settings and constraints in the Quartus II software.

## Synthesis Netlist Optimizations

The **.qsf** file variable name is used in the Tcl assignment to make the setting along with the appropriate value. The **Type** column indicates whether the setting is supported as a global setting, an instance setting, or both.

**Table 16-2: Synthesis Netlist Optimizations and Associated Settings**

Setting Name	Quartus II Settings File Variable Name	Values	Type
Perform WYSIWYG Primitive Resynthesis	ADV_NETLIST_OPT_SYNTH_WYSIWYG_REMAP	ON, OFF	Global, Instance
Optimization Technique	<Device Family Name>_OPTIMIZATION_TECHNIQUE	AREA, SPEED, BALANCED	Global, Instance
Power-Up Don't Care	ALLOW_POWER_UP_DONT_CARE	ON, OFF	Global
Save a node-level netlist into a persistent source file	LOGICLOCK_INCREMENTAL_COMPILE_ASSIGNMENT	ON, OFF	Global
	LOGICLOCK_INCREMENTAL_COMPILE_FILE	<file name>	
Allow Netlist Optimizations	ADV_NETLIST_OPT_ALLOWED	"ALWAYS ALLOW", DEFAULT, "NEVER ALLOW"	Instance

### Related Information

[WYSIWYG Primitive Resynthesis](#) on page 16-1

For more information on the Quartus II Settings File (**.qsf**), variable names and applicable values for the settings.

## Physical Synthesis Optimizations

The **.qsf** file variable name is used in the Tcl assignment to make the setting, along with the appropriate value. The **Type** column indicates whether the setting is supported as a global setting, an instance setting, or both.

**Table 16-3: Physical Synthesis Optimizations and Associated Settings**

Setting Name	Quartus II Settings File Variable Name	Values	Type
Physical Synthesis for Combinational Logic	PHYSICAL_SYNTHESIS_COMBO_LOGIC	ON, OFF	Global

Setting Name	Quartus II Settings File Variable Name	Values	Type
Automatic Asynchronous Signal Pipelining	PHYSICAL_SYNTHESIS_ASYNCHRONOUS_SIGNAL_PIPELINING	ON, OFF	Global
Perform Register Duplication	PHYSICAL_SYNTHESIS_REGISTER_DUPLICATION	ON, OFF	Global
Perform Register Retiming	PHYSICAL_SYNTHESIS_REGISTER_RETIMING	ON, OFF	Global
Power-Up Don't Care	ALLOW_POWER_UP_DONT_CARE	ON, OFF	Global, Instance
Power-Up Level	POWER_UP_LEVEL	HIGH, LOW	Instance
Allow Netlist Optimizations	ADV_NETLIST_OPT_ALLOWED	"ALWAYS ALLOW", DEFAULT, "NEVER ALLOW"	Instance
Save a node-level netlist into a persistent source file	LOGICLOCK_INCREMENTAL_COMPILE_ASSIGNMENT	ON, OFF	Global
	LOGICLOCK_INCREMENTAL_COMPILE_FILE	<file name>	

#### Related Information

- [Performing Physical Synthesis Optimizations](#) on page 16-3  
For more information on the Quartus II Settings File (.qsf), variable names and applicable values for the settings.
- [Quartus II Incremental Compilation for Hierarchical and Team-Based Design](#)  
For information about scripting and command line usage for incremental compilation.

## Back-Annotating Assignments

You can use the `logiclock_back_annotate` Tcl command to back-annotate resources in your design. This command can back-annotate resources in LogicLock regions, and resources in designs without LogicLock regions.

The following Tcl command back-annotates all registers in your design:

```
logiclock_back_annotate -resource_filter "REGISTER"
```

The `logiclock_back_annotate` command is in the `backannotate` package.

**Related Information**

[Preserving Your Physical Synthesis Results](#) on page 16-8  
For more information about back-annotating assignments.

## Document Revision History

**Table 16-4: Document Revision History**

Date	Version	Changes
2014.12.15	14.1.0	<ul style="list-style-type: none"> <li>Updated location of Fitter Settings, Analysis &amp; Synthesis Settings, and Physical Synthesis Optimizations Settings to Compiler Settings.</li> <li>Updated DSE II content.</li> </ul>
June 2014	14.0.0	Updated format.
November 2013	13.1.0	Removed HardCopy device information.
June 2012	12.0.0	Removed survey link.
November 2011	10.0.2	Template update.
December 2010	10.0.1	Template update.
July 2010	10.0.0	<ul style="list-style-type: none"> <li>Added links to Quartus II Help in several sections.</li> <li>Removed Referenced Documents section.</li> <li>Reformatted Document Revision History</li> </ul>
November 2009	9.1.0	<ul style="list-style-type: none"> <li>Added information to “Physical Synthesis for Registers—Register Retiming”</li> <li>Added information to “Applying Netlist Optimization Options”</li> <li>Made minor editorial updates</li> </ul>
March 2009	9.0.0	<ul style="list-style-type: none"> <li>Was chapter 11 in the 8.1.0 release.</li> <li>Updated the “Physical Synthesis for Registers—Register Retiming” and “Physical Synthesis Options for Fitting”</li> <li>Updated “Performing Physical Synthesis Optimizations”</li> <li>Deleted Gate-Level Register Retiming section.</li> <li>Updated the referenced documents</li> </ul>
November 2008	8.1.0	Changed to 8½” × 11” page size. No change to content.
May 2008	8.0.0	<ul style="list-style-type: none"> <li>Updated “Physical Synthesis Optimizations for Performance on page 11-9</li> <li>Added Physical Synthesis Options for Fitting on page 11-16</li> </ul>

**Related Information****[Quartus II Handbook Archive](#)**

For previous versions of the Quartus II Handbook.

# Engineering Change Orders with the Chip Planner 17

2014.06.30

QII5V2



Subscribe



Send Feedback

Programmable logic can accommodate changes to a system specification late in the design cycle. In a typical engineering project development cycle, the specification of the programmable logic portion is likely to change after engineering begins or while integrating all system elements. Last-minute design changes, commonly referred to as engineering change orders (ECOs), are small targeted changes to the functionality of a design after the design has been fully compiled.

The Chip Planner supports ECOs by allowing quick and efficient changes to your logic late in the design cycle. The Chip Planner provides a visual display of your post-place-and-route design mapped to the device architecture of your chosen FPGA and allows you to create, move, and delete logic cells and I/O atoms.

**Note:** In addition to making ECOs, the Chip Planner allows you to perform detailed analysis on routing congestion, relative resource usage, logic placement, LogicLock™ regions, fan-ins and fan-outs, paths between registers, and delay estimates for paths.

ECOs directly apply to atoms in the target device. As such, performing an ECO relies on your understanding of the device architecture of the target device.

## Related Information

- [About the Chip Planner](#)  
For a list of supported devices
- [Analyzing and Optimizing the Design Floorplan](#) on page 15-1  
For more information about using the Chip Planner for design analysis
- [Literature](#)  
For more information about the architecture of your device

## Engineering Change Orders

In the context of an FPGA design, you can apply an ECO directly to a physical resource on the device to modify its behavior. ECOs are typically made during the verification stage of a design cycle. When a small change is required on a design (such as modifying a PLL for a different clock frequency or routing a signal out to a pin for analysis) recompilation of the entire design can be time consuming, especially for larger designs.

Because several iterations of small design changes can occur during the verification cycle, recompilation times can quickly add up. Furthermore, a full recompilation due to a small design change can result in the

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



loss of previous design optimizations. Making ECOs, instead of performing a full recompilation on your design, limits the change only to the affected portions of logic.

## Performance Preservation

You can preserve the results of previous design optimizations when you make changes to an existing design with one of the following methods:

- Incremental compilation
- Rapid recompile
- ECOs

Choose the method to modify your design based on the scope of the change. The methods above are arranged from the larger scale change to the smallest targeted change to a compiled design.

The incremental compilation feature allows you to preserve compilation results at an RTL component or module level. After the initial compilation of your design, you can assign modules in your design hierarchy to partitions. Upon subsequent compilations, incremental compilation recompiles changed partitions based on the chosen preservation levels.

The rapid recompilation feature leverages results from the latest post-fit netlist to determine the changes required to honor modifications you have made to the source code. If you run a rapid recompilation, the Compiler refits only changed portion of the netlist.

ECOs provide a finer granularity of control compared to the incremental compilation and the rapid recompilation feature. All modifications are performed directly on the architectural elements of the device. You should use ECOs for targeted changes to the post-fit netlist.

**Note:** In the Quartus II software versions 10.0 and later, the software does not preserve ECO modifications to the netlist when you recompile a design with the incremental compilation feature turned on. You can reapply ECO changes made during a previous compilation with the Change Manager.

### Related Information

[Quartus II Incremental Compilation for Hierarchical and Team-Based Design](#)

## Compilation Time

In the traditional programmable logic design flow, a small change in the design requires a complete recompilation of the design. A complete recompilation of the design consists of synthesis and place-and-route. Making small changes to the design to reach the final implementation on a board can be a long process. Because the Chip Planner works only on the post-place-and-route database, you can implement your design changes in minutes without performing a full compilation.

## Verification

After you make a design change, you can verify the impact on your design. To verify that your changes do not violate timing requirements, perform static timing analysis with the Quartus II TimeQuest Timing Analyzer after you check and save your netlist changes in the Chip Planner.

Additionally, you can perform a gate-level or timing simulation of the ECO-modified design with the post-place-and-route netlist generated by the Quartus II software.

### Related Information

[Quartus II TimeQuest Timing Analyzer](#)



## Change Modification Record

All ECOs made with the Chip Planner are logged in the Change Manager to track all changes. With the Change Manager, you can easily revert to the original post-fit netlist or you can pick and choose which ECOs to apply.

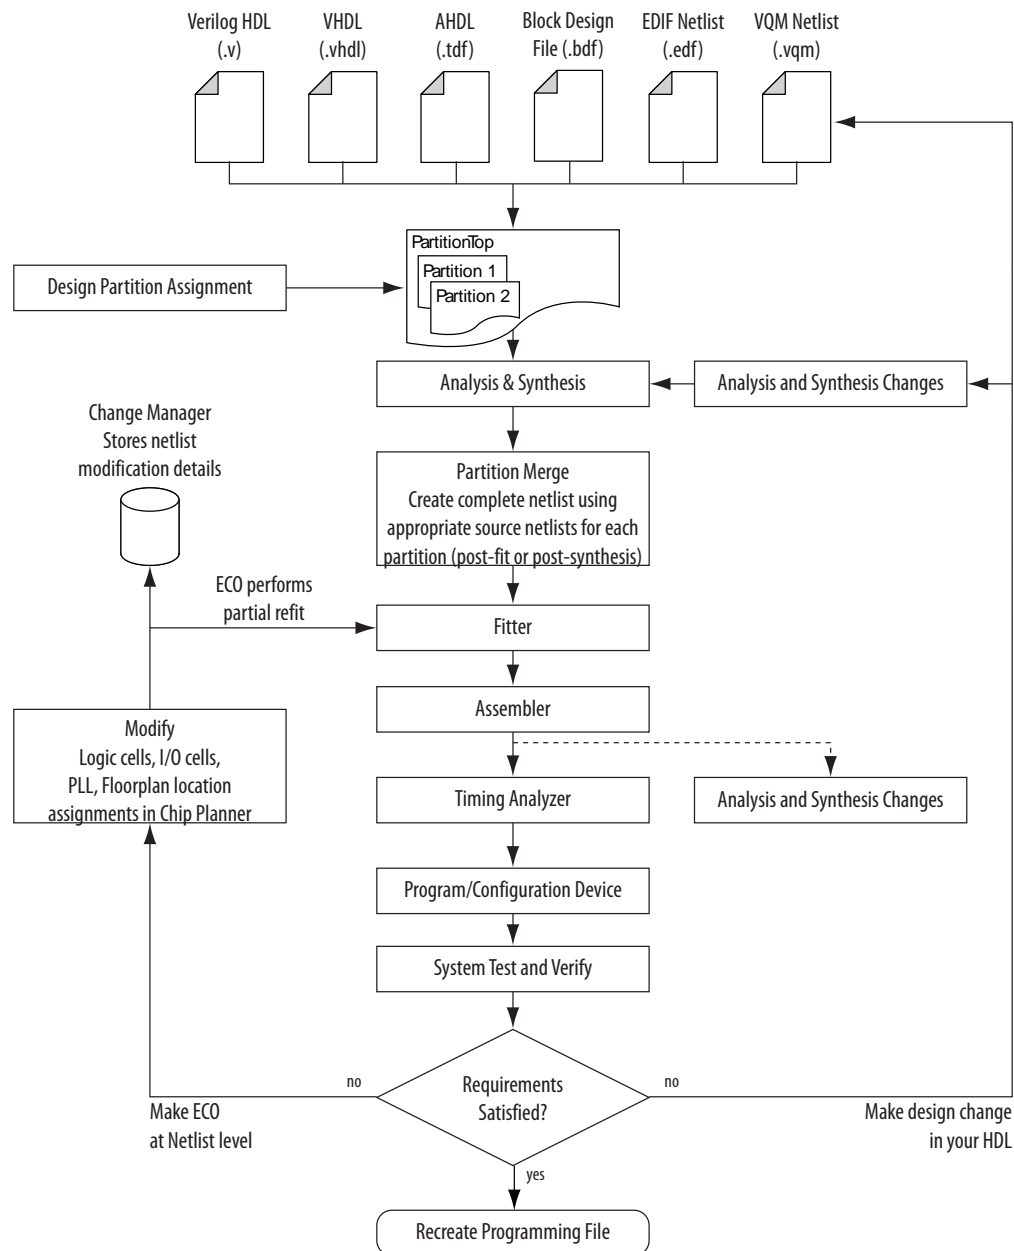
Additionally, the Quartus II software provides support for multiple compilation revisions of the same project. You can use ECOs made with the Chip Planner in conjunction with revision support to compare several different ECO changes and revert back to previous project revisions when required.

## ECO Design Flow

For iterative verification cycles, implementing small design changes at the netlist level can be faster than making an RTL code change. As such, making ECO changes are especially helpful when you debug the design on silicon and require a fast turnaround time to generate a programming file for debugging the design.

The figure shows the design flow for making ECOs.

Figure 17-1: Design Flow to Support ECOs



A typical ECO application occurs when you uncover a problem on the board and isolate the problem to the appropriate nodes or I/O cells on the device. You must be able to correct the functionality quickly and generate a new programming file. By making small changes with the Chip Planner, you can modify the post-place-and-route netlist directly without having to perform synthesis and logic mapping, thus decreasing the turnaround time for programming file generation during the verification cycle. If the

change corrects the problem, no modification of the HDL source code is necessary. You can use the Chip Planner to perform the following ECO-related changes to your design:

- Document the changes made with the Change Manager
- Easily recreate the steps taken to produce design changes
- Generate EDA simulation netlists for design verification

**Note:** For more complex changes that require HDL source code modifications, the incremental compilation feature can help reduce recompilation time.

## The Chip Planner Overview

The Chip Planner provides a visual display of device resources. It shows the arrangement and usage of the resource atoms in the device architecture that you are targeting. Resource atoms are the building blocks for your device, such as ALMs, LEs, PLLs, DSP blocks, memory blocks, or I/O elements.

The Chip Planner also provides an integrated platform for design analysis and for making ECOs to your design after place-and-route. The toolset consists of the Chip Planner (providing a device floorplan view of your mapped design) and two integrated subtools—the Resource Property Editor and the Change Manager.

For analysis, the Chip Planner can show logic placement, LogicLock regions, relative resource usage, detailed routing information, routing congestion, fan-ins and fan-outs, paths between registers, and delay estimates for paths. Additionally, the Chip Planner allows you to create location constraints or resource assignment changes, such as moving or deleting logic cells or I/O atoms with the device floorplan. For ECO changes, the Chip Planner enables you to create, move, or delete logic cells in the post-place-and-route netlist for fast programming file generation. Additionally, you can open the Resource Property Editor from the Chip Planner to edit the properties of resource atoms or to edit the connections between resource atoms. All changes to resource atoms and connections are logged automatically with the Change Manager.

## Opening the Chip Planner

To open the Chip Planner, on the Tools menu, click **Chip Planner**. Alternatively, click the **Chip Planner** icon on the Quartus II software toolbar.

Optionally, you can open the Chip Planner by cross-probing from the shortcut menu in the following tools:

- Design Partition Planner
- Compilation Report
- LogicLock Regions window
- Technology Map Viewer
- Project Navigator window
- RTL source code
- Node Finder
- Simulation Report
- RTL Viewer
- Report Timing panel of the TimeQuest Timing Analyzer

## The Chip Planner Tasks and Layers

The Chip Planner allows you to set up tasks to quickly implement ECO changes or manipulate assignments for the floorplan of the device. Each task consists of an editing mode and a set of customized layer settings.

### Related Information

- [Performing ECOs in the Resource Property Editor](#) on page 17-6
- [About the Chip Planner](#)
- [Analyzing and Optimizing the Design Floorplan](#) on page 15-1

## Performing ECOs with the Chip Planner (Floorplan View)

You can manipulate resource atoms in the Chip Planner when you select the ECO editing mode.

The following ECO changes can be made with the Chip Planner Floorplan view:

- Create atoms
- Delete atoms
- Move existing atoms

**Note:** To configure the properties of atoms, such as managing the connections between different LEs/ALMs, use the Resource Property Editor.

To select the ECO editing mode in the Chip Planner, in the **Editing Mode** list at the top of the Chip Planner, select the ECO editing mode.

### Related Information

[Performing ECOs in the Resource Property Editor](#) on page 17-6

## Creating, Deleting, and Moving Atoms

You can use the Chip Planner to create, delete, and move atoms in the post-compilation design.

### Related Information

[Creating, Deleting, and Moving Atoms](#)

## Check and Save Netlist Changes

After making all the ECOs, you can run the Fitter to incorporate the changes by clicking the **Check and Save Netlist Changes** icon in the Chip Planner toolbar. The Fitter compiles the ECO changes, performs design rule checks on the design, and generates a programming file.

## Performing ECOs in the Resource Property Editor

You can view and edit the following resources with the Resource Property Editor.

### Logic Elements

An Altera<sup>®</sup> LE contains a four-input LUT, which is a function generator that can implement any function of four variables. In addition, each LE contains a register fed by the output of the LUT or by an independent function generated in another LE.

You can use the Resource Property Editor to view and edit any LE in the FPGA. To open the Resource Property Editor for an LE, on the Project menu, point to **Locate**, and then click **Locate in Resource Property Editor** in one of the following views:

- RTL Viewer
- Technology Map Viewer
- Node Finder
- Chip Planner

For more information about LE architecture for a particular device family, refer to the device family handbook or data sheet.


You can use the Resource Property Editor to change the following LE properties:

- Data input to the LUT
- LUT mask or LUT

## Logic Element Properties

To view logic element properties, on the View menu, click **View Properties**.

Figure 17-2: LE Properties in the Resource Property Editor



## Modes of Operation

LUTs in an LE can operate in either normal or arithmetic mode.

When an LE is configured in normal mode, the LUT in the LE can implement a function of four inputs.

When the LE is configured in arithmetic mode, the LUT in the LE is divided into two 3-input LUTs. The first LUT generates the signal that drives the output of the LUT, while the second LUT generates the carry-out signal. The carry-out signal can drive only a carry-in signal of another LE.

For more information about LE modes of operation, refer to volume 1 of the appropriate device handbook.

## Sum and Carry Equations

You can change the logic function implemented by the LUT by changing the sum and carry equations.

When the LE is configured in normal mode, you can change only the sum equation. When the LE is configured in arithmetic mode, you can change both the sum and the carry equations.

The LUT mask is the hexadecimal representation of the LUT equation output. When you change the LUT equation, the Quartus II software automatically changes the LUT mask. Conversely, when you change the LUT mask, the Quartus II software automatically computes the LUT equation.

## sload and sclr Signals

Each LE register contains a synchronous load (`sload`) signal and a synchronous clear (`sclr`) signal. You can invert either the `sload` or `sclr` signal feeding into the LE.

If the design uses the `sload` signal in an LE, the signal and its inversion state must be the same for all other LEs in the same LAB. For example, if two LEs in a LAB have the `sload` signal connected, both LEs must have the `sload` signal set to the same value. This is also true for the `sclr` signal.

## Register Cascade Mode

When register cascade mode is enabled, the cascade-in port feeds the input to the register. The register cascade mode is used most often when the design implements shift registers.

You can change the register cascade mode by connecting (or disconnecting) the cascade in the port. However, if you create this port, you must ensure that the source port LE is directly above the destination LE.

## Cell Delay Table

The cell delay table describes the propagation delay from all inputs to all outputs for the selected LE.

## Logic Element Connections

To view the connections that feed in and out of an LE, on the View menu, click **View Port Connections**.

Figure 17-3: View LE Connections in the Connectivity Window

Input Port Name	Signal Name	Inverted	Output Port Name	Signal Name
DATAA	<Disconnected>	False	COLUT	<Disconnected>
DATAB	<Disconnected>	False	COMBOUT	<Disconnected>
SDATA	filter state_m:inst1 filter.tap4	False	REGOUT	filter inst4
DATAD	<Disconnected>	False		
CIN	<Disconnected>	False		
INVERTA	<Disconnected>	False		
REGCASCIN	<Disconnected>	False		
SLOAD	VCC	False		
SCLR	<Disconnected>	False		
IACLR	VCC	False		
ALOAD	<Disconnected>	False		
CLK	filter clk	False		
ENA	<Disconnected>	False		

## Deleting a Logic Element

To delete an LE, follow these steps:

1. Right-click the desired LE in the Chip Planner, point to **Locate**, and click **Locate in Resource Property Editor**.
2. You must remove all fan-out connections from an LE prior to deletion. To delete fan-out connections, right-click each connected output signal, point to **Remove**, and click **Fanouts**. Select all of the fan-out signals in the **Remove Fan-outs** dialog box and click **OK**.
3. To delete an atom after all fan-out connections are removed, right-click the atom in the Chip Planner and click **Delete Atom**.

## Adaptive Logic Modules

Each ALM contains LUT-based resources that can be divided between two adaptive LUTs (ALUTs).

With up to eight inputs to the two ALUTs, each ALM can implement various combinations of two functions. This adaptability allows the ALM to be completely backward-compatible with four-input LUT architectures. One ALM can implement any function with up to six inputs and certain seven-input functions. In addition to the ALUT-based resources, each ALM contains two programmable registers, two dedicated full adders, a carry chain, a shared arithmetic chain, and a register chain. The ALM can efficiently implement various arithmetic functions and shift registers with these dedicated resources.

You can implement the following types of functions in a single ALM:

- Two independent 4-input functions
- An independent 5-input function and an independent 3-input function
- A 5-input function and a 4-input function, if they share one input
- Two 5-input functions, if they share two inputs
- An independent 6-input function
- Two 6-input functions, if they share four inputs and share the same functions
- Certain 7-input functions

You can use the Resource Property Editor to change the following ALM properties:

- Data input to the LUT
- LUT mask or LUT equation


### Adaptive Logic Module Schematic

You can view and edit any ALM atom with the Resource Property Editor by right-clicking the ALM in the RTL Viewer, the Node Finder, or the Chip Planner, and clicking **Locate in Resource Property Editor**.

For a detailed description of the ALM, refer to the device handbooks of devices based on an ALM architecture.

By default, the Quartus II software displays the used resources in blue and the unused in gray. For the figure, the used resources are in blue and the unused resources are in gray.

Figure 17-4: Adaptive Logic Module



### Adaptive Logic Module Properties

The properties that you can display for the ALM include an equations table that shows the name and location of each of the two combinational nodes and two register nodes in the ALM, the individual LUT equations for each of the combinational nodes, and the `combout`, `sumout`, `carryout`, and `shareout` equations for each combinational node.

### Adaptive Logic Module Connections

Click **View > View Connectivity** to view the input and output connections for the ALM.

### FPGA I/O Elements

Altera FPGAs that have high-performance I/O elements, including up to six registers, are equipped with support for a number of I/O standards that allow you to run your design at peak speeds. Use the Resource Property Editor to view, change connectivity, and edit the properties of the I/O elements. Use the Chip Planner (Floorplan view) to change placement, delete, and create new I/O elements.

For a detailed description of the device I/O elements, refer to the applicable device handbook.

You can change the following I/O properties:

- Delay chain
- Bus hold
- Weak pull up
- Slow slew rate
- I/O standard
- Current strength



- Extend OE disable
- PCI I/O
- Register reset mode
- Register synchronous reset mode
- Register power up
- Register mode

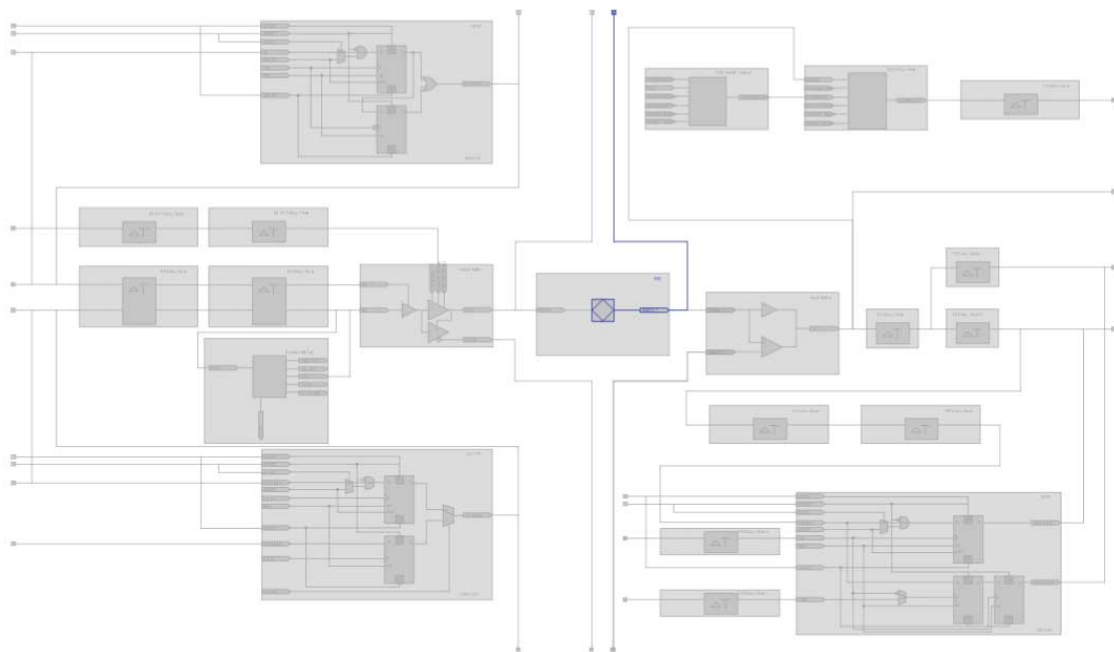
## Stratix V I/O Elements

The I/O elements in Stratix® V devices contain a bidirectional I/O buffer and I/O registers to support a complete embedded bidirectional single data rate (SDR) or double data rate (DDR) transfer.

I/O registers are composed of the input path for handling data from the pin to the core, the output path for handling data from the core to the pin, and the output enable path for handling the output enable signal to the output buffer. These registers allow faster source-synchronous register-to-register transfers and resynchronization. The input path consists of the DDR input registers, alignment and synchronization registers, and half data rate blocks; you can bypass each block in the input path. The input path uses the deskew delay to adjust the input register clock delay across process, voltage, and temperature (PVT) variations.

By default, the Quartus II software displays the used resources in blue and the unused resources in gray.

**Figure 17-5: Stratix V Device I/O Element Structure**



### Related Information

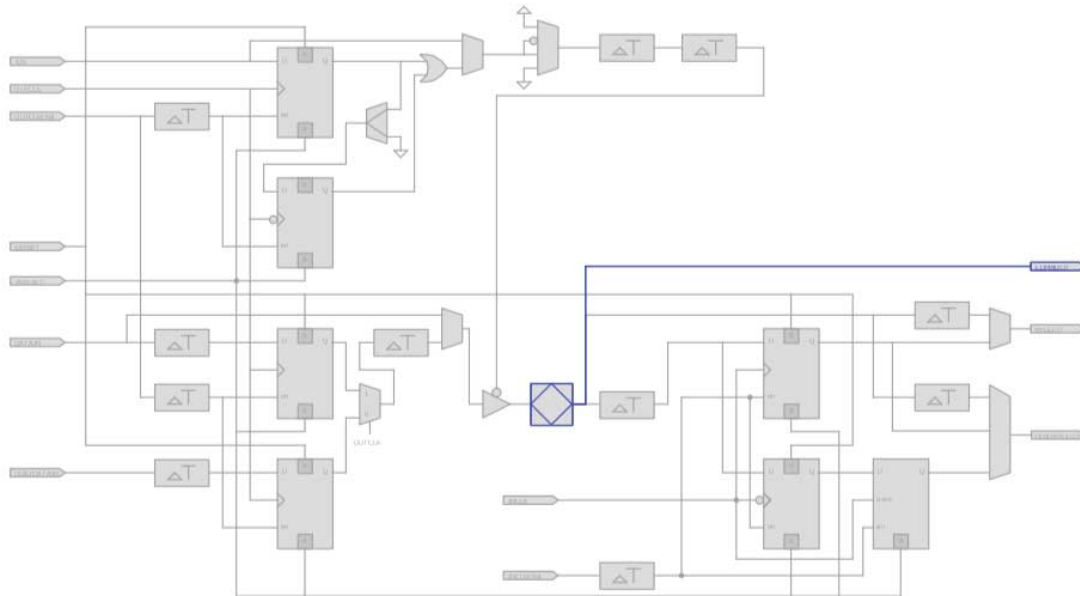
[Stratix V Device Handbook](#)

## Stratix IV I/O Elements

The I/O elements in Stratix IV devices contain a bidirectional I/O buffer and I/O registers to support a complete embedded bidirectional SDR or DDR transfer.

The I/O registers are composed of the input path for handling data from the pin to the core, the output path for handling data from the core to the pin, and the output enable path for handling the output enable signal for the output buffer. Each path consists of a set of delay elements that allow you to fine-tune the timing characteristics of each path for skew management. By default, the Quartus II software displays the used resources in blue and the unused resources in gray.

**Figure 17-6: Stratix IV I/O Element and Structure**



#### Related Information

##### Literature


For more information about I/O elements in Stratix IV devices

#### Arria V I/O Elements

The I/O elements in Arria<sup>®</sup> V devices contain a bidirectional I/O buffer and I/O registers to support a complete embedded bidirectional SDR or DDR transfer.

The I/O registers are composed of the input path for handling data from the pin to the core, the output path for handling data from the core to the pin, and the output enable path for handling the output enable signal for the output buffer. Each path consists of a set of delay elements that allow you to fine-tune the timing characteristics of each path for skew management. By default, the Quartus II software displays the used resources in blue and the unused resources in gray.

Figure 17-7: Arria V Device I/O Element and Structure

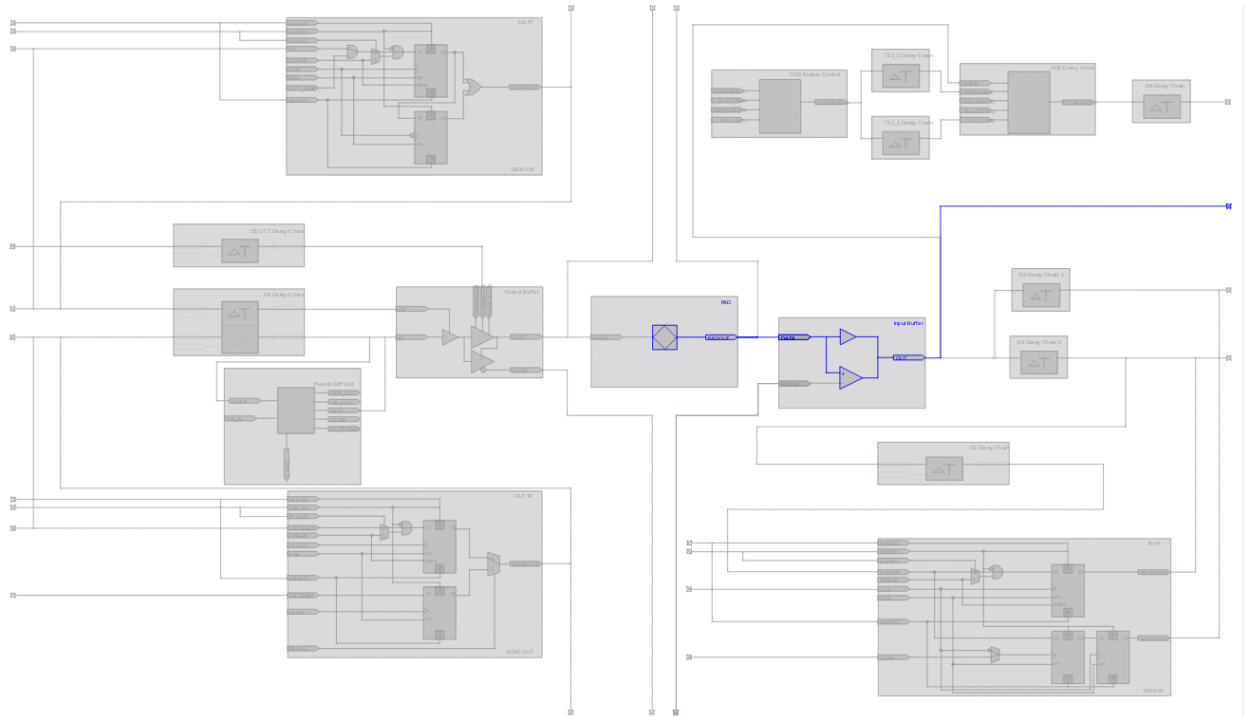


## Cyclone V I/O Elements

The I/O elements in Cyclone V devices contain a bidirectional I/O buffer and registers for complete embedded bidirectional single data rate transfer. The I/O element contains three input registers, two output registers, and two output-enable registers. The two output registers and two output-enable registers are utilized for double-data rate (DDR) applications.

You can use the input registers for fast setup times and the output registers for fast clock-to-output times. Additionally, you can use the output-enable (OE) registers for fast clock-to-output enable timing. You can use I/O elements for input, output, or bidirectional data paths. By default, the Quartus II software displays the used resources in blue and the unused resources in gray.


Figure 17-8: Cyclone V Device I/O Elements and Structure



## MAX V I/O Elements

The I/O elements in MAX<sup>®</sup> V devices contain a bidirectional I/O buffer. You can drive registers from adjacent LABs to or from the bidirectional I/O buffer of the I/O element. By default, the Quartus II software displays the used resources in blue and the unused resources in gray.

Figure 17-9: MAX V Device I/O Elements and Structure

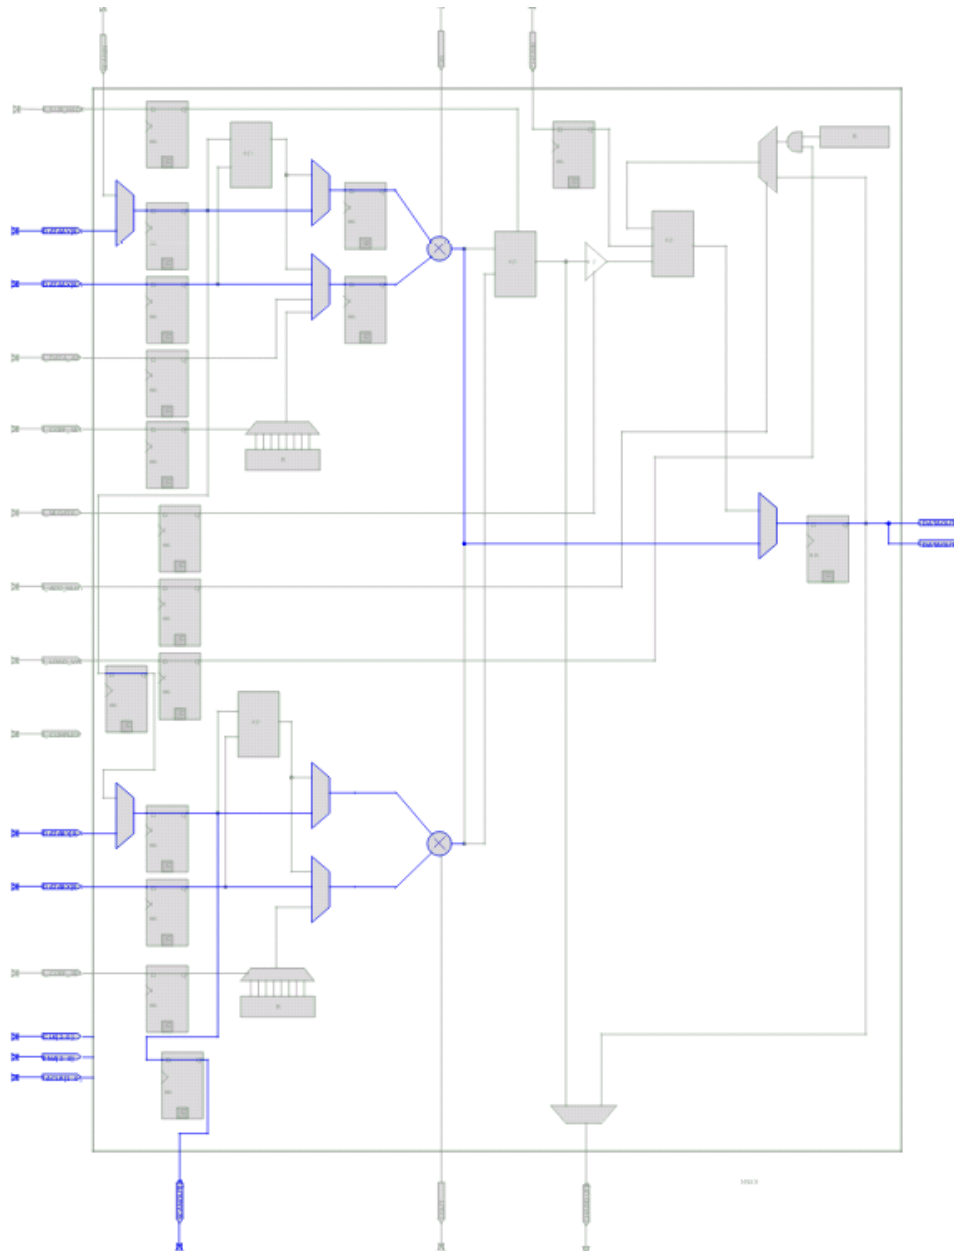


## FPGA RAM Blocks

With the Resource Property Editor, you can view the architecture of different RAM blocks in the device, modify the input and output registers to and from the RAM blocks, and modify the connectivity of the

input and output ports. By default, the Quartus II software displays the used resources in blue and the unused resources in gray.

**Figure 17-10: M9K RAM View in a Stratix V Device**



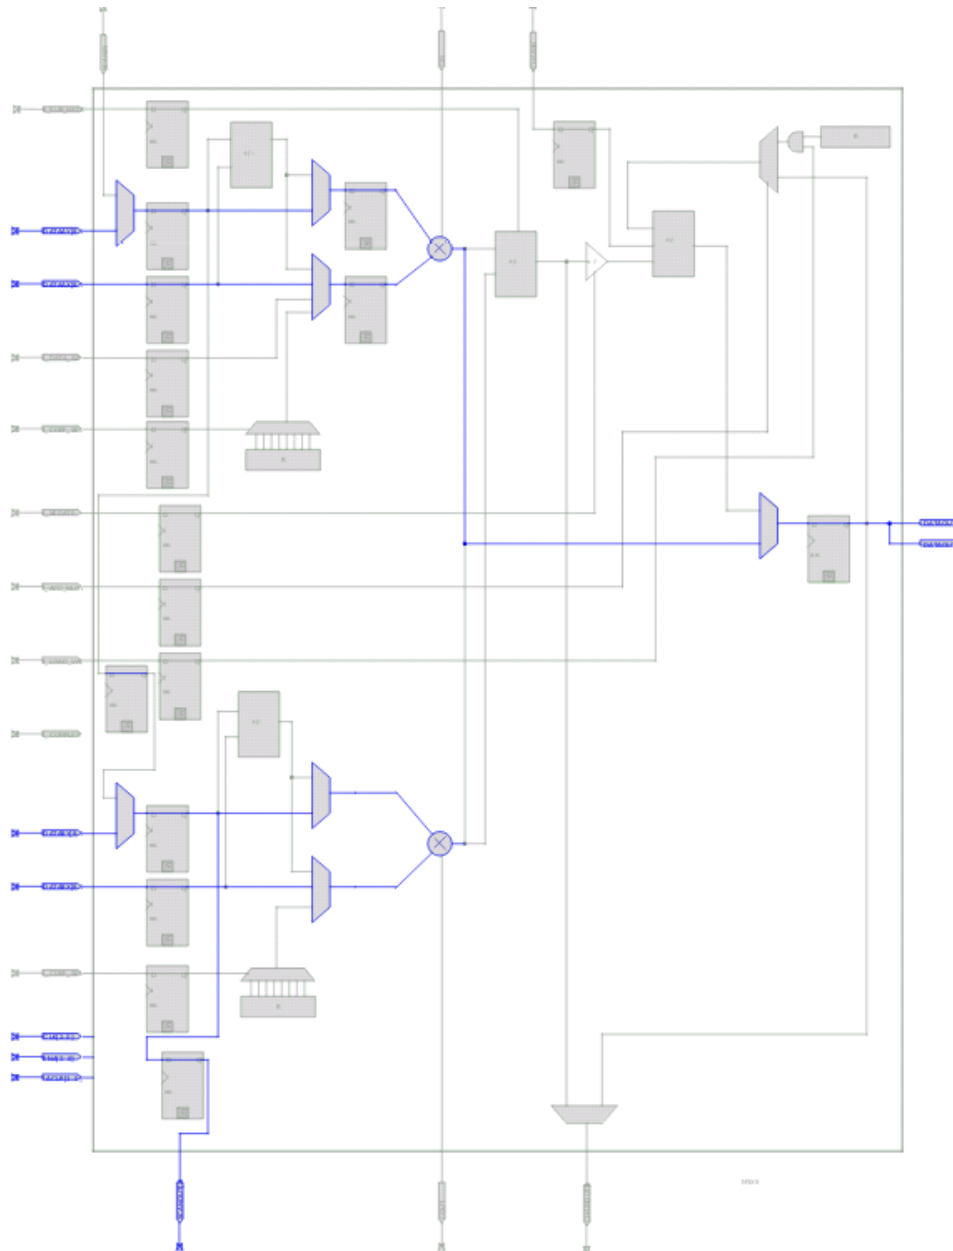
## FPGA DSP Blocks

Dedicated hardware DSP circuit blocks in Altera devices provide performance benefits for the critical DSP functions in your design.

The Resource Property Editor allows you to view the architecture of DSP blocks in the Resource Property Editor for the Cyclone and Stratix series of devices. The Resource Property Editor also allows you to modify the signal connections to and from the DSP blocks and modify the input and output registers to

and from the DSP blocks. By default, the Quartus II software displays the used resources in blue and the unused resources in gray.

**Figure 17-11: DSP Block View in a Stratix V Device**



## Change Manager

The Change Manager maintains a record of every change you perform with the Chip Planner, the Resource Property Editor, the SignalProbe feature, or a Tcl script. Each row of data in the Change Manager represents one ECO.

The Change Manager allows you to apply changes, roll back changes, delete changes, and export change records to a Text File (.txt), a Comma-Separated Value File (.csv), or a Tcl Script File (.tcl). The Change Manager tracks dependencies between changes, so that when you apply, roll back, or delete a change, any prerequisite or dependent changes are also applied, rolled back, or deleted.

#### Related Information

[About the Change Manager](#)

## Complex Changes in the Change Manager

Certain changes in the Change Manager (including creating or deleting atoms and changing connectivity) can appear to be self-contained, but are actually composed of multiple actions. The Change Manager marks such complex changes with a plus icon in the **Index** column.

You can click the plus icon to expand the change record and show all the component actions performed as part of that complex change.

#### Related Information

[Example of Managing Changes With the Change Manager](#)

## Managing SignalProbe Signals

The SignalProbe pins that you create from the **SignalProbe Pins** dialog box are recorded in the Change Manager. After you have made a SignalProbe assignment, you can use the Change Manager to quickly disable SignalProbe assignments by selecting **Revert to Last Saved Netlist** on the shortcut menu in the Change Manager.

#### Related Information

[Quick Design Debugging Using SignalProbe](#)

## Exporting Changes

You can export changes to a .txt, a .csv, or a .tcl. Tcl scripts allow you to reapply changes that were deleted during compilation.

#### Related Information

- [Managing Changes With the Change Manager](#)
- [Quartus II Incremental Compilation for Hierarchical and Team-Based Design](#)

## Scripting Support

You can run procedures and make settings described in this chapter in a Tcl script. You can also run some procedures at a command prompt. The Tcl commands for controlling the Chip Planner are located in the `chip_planner` package of the `quartus_cdb` executable.

#### Related Information

- [About Quartus II Scripting](#)
- [Tcl Scripting](#) on page 3-1
- [Quartus II Settings File Manual](#)
- [Command Line Scripting](#) on page 2-1

## Common ECO Applications

You can use an ECO to make a post-compilation change to your design.

To help build your system quickly, you can use Chip Planner functions to perform the following activities:

- Adjust the drive strength of an I/O with the Chip Planner
- Modify the PLL properties with the Resource Property Editor, see [Modify the PLL Properties With the Chip Planner](#)
- Modify the connectivity between new resource atoms with the Chip Planner and Resource Property Editor

### Related Information

[Modify the PLL Properties With the Chip Planner](#) on page 17-19

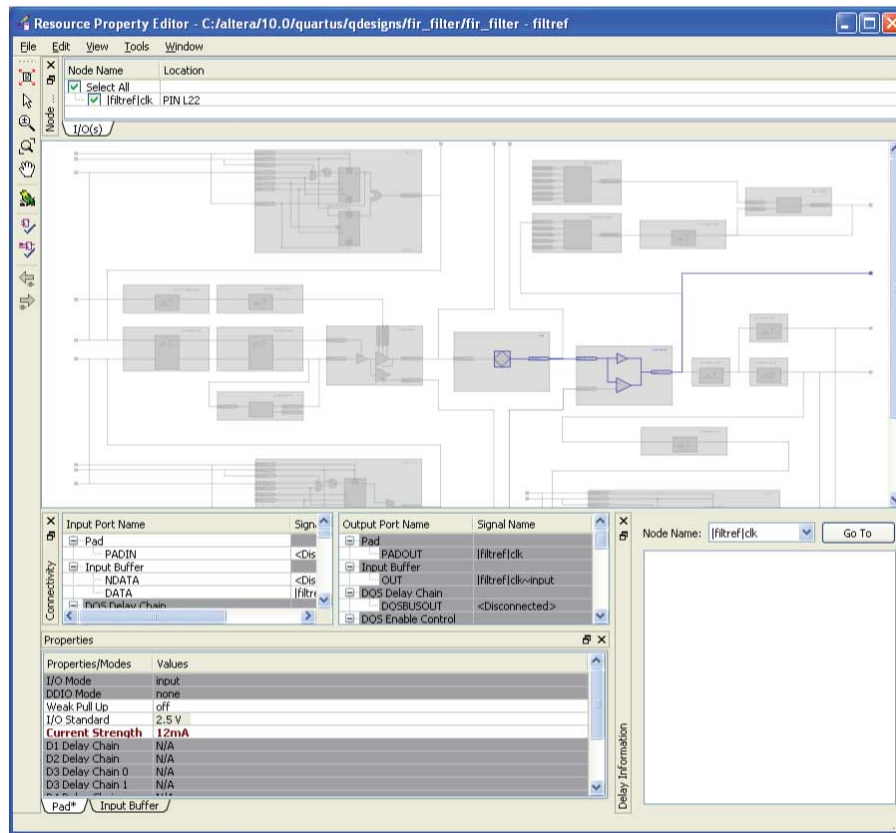
## Adjust the Drive Strength of an I/O with the Chip Planner

To adjust the drive strength of an I/O, follow these steps to incorporate the ECO changes into the netlist of the design.

1. In the **Editing Mode** list at the top of the Chip Planner, select the ECO editing mode.
2. Locate the I/O in the **Resource Property Editor**.
3. In the **Resource Property Editor**, point to the **Current Strength** option in the **Properties** pane and double-click the value to enable the drop-down list.
4. Change the value for the **Current Strength** option.
5. Right-click the ECO change in the Change Manager and click **Check & Save All Netlist Changes** to apply the ECO change.



Figure 17-12: I/O in the Resource Property Editor



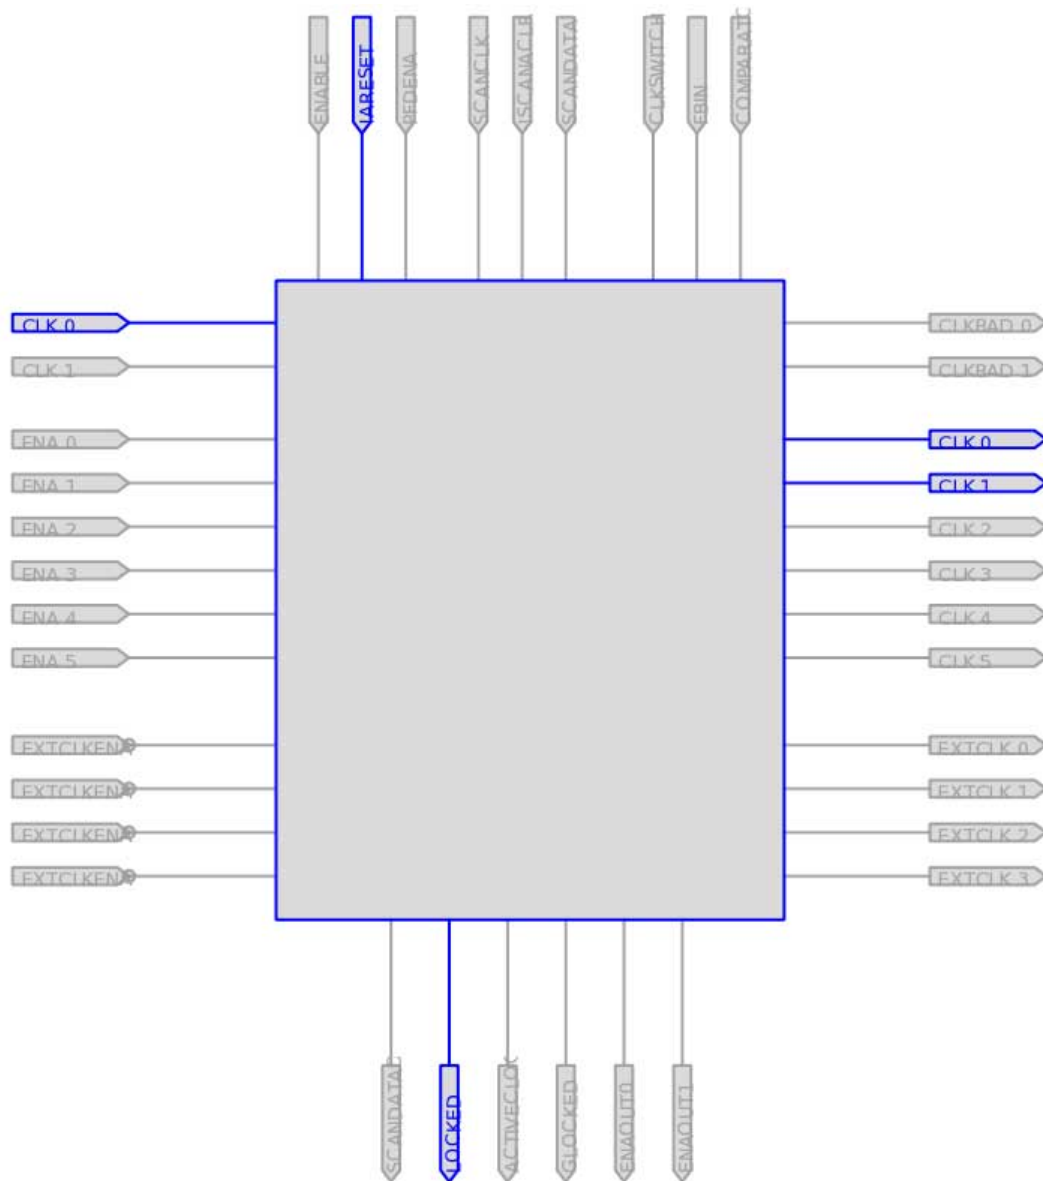
**Note:** You can change the pin locations of input or output ports with the ECO flow. You can drag and move the signal from an existing pin location to a new location while in the Post Compilation Editing (ECO) task in the Chip Planner. You can then click **Check & Save All Netlist Changes** to compile the ECO.

## Modify the PLL Properties With the Chip Planner

You use PLLs to modify and generate clock signals to meet design requirements. Additionally, you can use PLLs to distribute clock signals to different devices in a design, reducing clock skew between devices, improving I/O timing, and generating internal clock signals.

The Resource Property Editor allows you to view and modify PLL properties to meet your design requirements.

Figure 17-13: PLL View in the Resource Property Editor of a Stratix Device



## PLL Properties

The Resource Property Editor allows you to modify PLL options, such as phase shift, output clock frequency, and duty cycle.

You can also change the following PLL properties with the Resource Property Editor:

- Input frequency
- $M V_{CO}$  tap
- $M$  initial
- $M$  value
- $N$  value

- M counter delay
- N counter delay
- M2 value
- N2 value
- SS counter
- Charge pump current
- Loop filter resistance
- Loop filter capacitance
- Counter delay
- Counter high
- Counter low
- Counter mode
- Counter initial
- V<sub>CO</sub> tap

You can also view post-compilation PLL properties in the Compilation Report. To do so, in the Compilation Report, select **Fitter** and then select **Resource Section**.

### Adjusting the Duty Cycle

Use the equation to adjust the duty cycle of individual output clocks.

$$\text{High \%} = \frac{\text{Counter High}}{(\text{Counter High} + \text{Counter Low})}$$

### Adjusting the Phase Shift

Use the equation to adjust the phase shift of an output clock of a PLL.

$$\text{Phase Shift} = (\text{Period } V_{CO} \times 0.125 \times \text{Tap } V_{CO}) + (\text{Initial } V_{CO} \times \text{Period } V_{CO})$$

For normal mode, Tap V<sub>CO</sub>, Initial V<sub>CO</sub>, and Period V<sub>CO</sub> are governed by the following settings:

$$\text{Tap } V_{CO} = \text{Counter Delay} - M \text{ Tap } V_{CO}$$

$$\text{Initial } V_{CO} = \text{Counter Delay} - M \text{ Initial}$$

$$\text{Period } V_{CO} = \text{In Clock Period} \times N \div M$$

For external feedback mode, Tap V<sub>CO</sub>, Initial V<sub>CO</sub>, and Period V<sub>CO</sub> are governed by the following settings:

$$\text{Tap } V_{CO} = \text{Counter Delay} - M \text{ Tap } V_{CO}$$

$$\text{Initial } V_{CO} = \text{Counter Delay} - M \text{ Initial}$$

$$\text{Period } V_{CO} = \frac{\text{In Clock Period} \times N}{M + \text{Counter High} + \text{Counter Low}}$$

$$(M + \text{Counter High} + \text{Counter Low})$$

**Figure 17-14:**

#### Related Information

[Stratix Device Handbook](#)

## Adjusting the Output Clock Frequency

Use the equation to adjust the PLL output clock in normal mode.

$$\text{Output Clock Frequency} = \text{Input Frequency} \cdot \frac{\text{M Value}}{\text{N Value} + \text{Counter High} + \text{Counter Low}}$$

Use the equation to adjust the PLL output clock in external feedback mode.

$$\text{OUTCLK} = \frac{\text{M Value} + \text{External Feedback Counter High} + \text{External Feedback Counter Low}}{\text{N Value} + \text{Counter High} + \text{Counter Low}}$$

## Adjusting the Spread Spectrum

Use the equation to adjust the spread spectrum for your PLL.

$$\% \text{ Spread} = \frac{M_2 N_1}{M_1 N_2}$$

## Modify the Connectivity between Resource Atoms

The Chip Planner and Resource Property Editor allow you to create new resource atoms and manipulate the existing connection between resource atoms in the post-fit netlist. These features are useful for small changes when you are debugging a design, such as manually inserting pipeline registers into a combinational path that fails timing, or routing a signal to a spare I/O pin for analysis.

Use the following procedure to create a new register in a Cyclone V device and route register output to a spare I/O pin. This example illustrates how to create a new resource atom and modify the connections between resource atoms.

To create new resource atoms and manipulate the existing connection between resource atoms in the post-fit netlist, follow these steps:

1. Create a new register in the Chip Planner.
2. Locate the atom in the Resource Property Editor.
3. To assign a clock signal to the register, right-click the clock input port for the register, point to **Edit connection**, and click **Other**. Use the Node Finder to assign a clock signal from your design.
4. To tie the SLOAD input port to V<sub>CC</sub>, right-click the clock input port for the register, point to **Edit connection**, and click **VCC**.
5. Assign a data signal from your design to the SDATA port.
6. In the Connectivity window, under the output port names, copy the port name of the register.
7. In the Chip Planner, locate a free I/O resource and create an output buffer.
8. Locate the new I/O atom in the Resource Property Editor.
9. On the input port to the output buffer, right-click, point to **Edit connection**, and click **Other**.
10. In the **Edit Connection** dialog box, type the output port name of the register you have created.
11. Run the ECO Fitter to apply the changes by clicking **Check and Save Netlist Changes**.

**Note:** A successful ECO connection is subject to the available routing resources. You can view the relative routing utilization by selecting **Routing Utilization** as the Background Color Map in the **Layers Settings** dialog box of the Chip Planner. Also, you can view individual routing channel utilization from local, row, and column interconnects with the tooltips created when

you position your mouse pointer over the appropriate resource. Refer to the device data sheet for more information about the architecture of the routing interconnects of your device.

## Post ECO Steps

After you make an ECO change with the Chip Planner, you must perform static timing analysis of your design with the TimeQuest analyzer to ensure that your changes did not adversely affect the timing performance of your design.

For example, when you turn on one of the delay chain settings for a specific pin, you change the I/O timing. Therefore, to ensure that the design still meets all timing requirements, you should perform static timing analysis.

### Related Information

#### [Quartus II TimeQuest Timing Analyzer](#)

For more information about performing a static timing analysis of your design

## Document Revision History

Table 17-1: Document Revision History

Date	Version	Changes
June 2014	14.0.0	<ul style="list-style-type: none"> <li>Updated formatting.</li> <li>Removed references to Stratix, Stratix II, Stratix III, Arria GX, Arria II GX, Cyclone, Cyclone II, Cyclone III, and MAX II devices.</li> <li>Added MAX V, Cyclone V, Arria V I/O elements</li> </ul>
June 2012	12.0.0	Removed survey link.
November 2011	10.1.1	Template update.
December 2010	10.1.0	<ul style="list-style-type: none"> <li>Updated chapter to new template</li> <li>Removed “The Chip Planner FloorPlan Views” section</li> <li>Combined “Creating Atoms”, “Deleting Atoms”, and “Moving Atoms” sections, and linked to Help.</li> <li>Added Stratix V I/O elements in “FPGA I/O Elements”.</li> </ul>

Date	Version	Changes
July 2010	10.0.0	<ul style="list-style-type: none"> <li>• Added information to page 17–1.</li> <li>• Added information to “Engineering Change Orders” on page 17–2.</li> <li>• Changed heading from “Performance” to “Performance Preservation” on page 7–2.</li> <li>• Updated information in “Performance Preservation” on page 17–2.</li> <li>• Changed heading from “Documentation” to “Change Modification Record” on page 17–3.</li> <li>• Changed heading from “Resource Property Editor” to “Performing ECOs in the Resource Property Editor” on page 17–15.</li> <li>• Removed “Using Incremental Compilation in the ECO Flow” section. Preservation support for ECOs with the incremental compilation flow has been removed in the Quartus II software version 10.0.</li> <li>• Removed “Referenced Documents” section.</li> </ul>
November 2009	9.1.0	<ul style="list-style-type: none"> <li>• Updated device support list</li> <li>• Made minor editorial updates</li> </ul>
March 2009	9.0.0	<ul style="list-style-type: none"> <li>• Updated Figure 17–17.</li> <li>• Made minor editorial updates.</li> <li>• Chapter 15 was previously Chapter 13 in the 8.1.0 release.</li> </ul>
November 2008	8.1.0	<ul style="list-style-type: none"> <li>• Corrected preservation attributes for ECOs in the section “Using Incremental Compilation in the ECO Flow” on page 15–32.</li> <li>• Minor editorial updates.</li> <li>• Changed to 8½” x 11” page size.</li> </ul>
May 2008	8.0.0	<ul style="list-style-type: none"> <li>• Updated device support list</li> <li>• Modified description for ECO support for block RAMs and DSP blocks</li> <li>• Corrected Stratix PLL ECO example</li> <li>• Added an application example to show modifying the connectivity between resource atoms</li> </ul>

**Related Information****[Quartus II Handbook Archive](#)**

For previous versions of the Quartus II Handbook

# Quartus II Handbook Volume 3: Verification



Subscribe



Send Feedback

**QII5V3**  
2014.12.15

101 Innovation Drive  
San Jose, CA 95134  
[www.altera.com](http://www.altera.com)



2014.08.18

QI15V3



Subscribe



Send Feedback

This document describes simulating designs that target Altera devices. Simulation verifies design behavior before device programming. The Quartus II software supports RTL- and gate-level design simulation in supported EDA simulators. Simulation involves setting up your simulator working environment, compiling simulation model libraries, and running your simulation.

## Simulator Support

The Quartus II software supports specific EDA simulator versions for RTL and gate-level simulation.

**Table 1-1: Supported Simulators**

Vendor	Simulator	Version	Platform
Aldec	Active-HDL	9.3sp1	Windows
Aldec	Riviera-PRO	2014.06	Windows, Linux
Cadence	Incisive Enterprise	13.2	Linux
Mentor Graphics	ModelSim-Altera (provided)	10.3c	Windows, Linux
Mentor Graphics	ModelSim PE	10.3c	Windows
Mentor Graphics	ModelSim SE	10.3c	Windows, Linux
Mentor Graphics	QuestaSim	10.3c	Windows, Linux
Synopsys	VCS/VCS MX	2014.03-sp1	Linux

## Simulation Levels

The Quartus II software supports various levels of simulation in supported EDA simulators.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered





Table 1-2: Supported Simulation Levels

Simulation Level	Description	Simulation Input
RTL	Cycle-accurate simulation using Verilog HDL, SystemVerilog, and VHDL design source code with simulation models provided by Altera and other IP providers.	<ul style="list-style-type: none"> <li>• Design source/testbench</li> <li>• Altera simulation libraries</li> <li>• Altera IP plain text or IEEE encrypted RTL models</li> <li>• IP simulation models</li> <li>• Altera IP functional simulation models</li> <li>• Altera IP bus functional models</li> <li>• Qsys-generated models</li> <li>• Verification IP</li> </ul>
Gate-level functional	Simulation using a post-synthesis or post-fit functional netlist testing the post-synthesis functional netlist, or post-fit functional netlist.	<ul style="list-style-type: none"> <li>• Testbench</li> <li>• Altera simulation libraries</li> <li>• Post-synthesis or post-fit functional netlist</li> <li>• Altera IP bus functional models</li> </ul>
Gate-level timing	Simulation using a post-fit timing netlist, testing functional and timing. Not supported for Arria V, Cyclone V, or Stratix V devices.	<ul style="list-style-type: none"> <li>• Testbench</li> <li>• Altera simulation libraries</li> <li>• Post-fit timing netlist</li> <li>• Post-fit Standard Delay Output File (.sdo)</li> </ul>

**Note:** Gate-level timing simulation of an entire design can be slow and should be avoided. Gate-level timing simulation is not supported for Arria<sup>®</sup> V, Cyclone<sup>®</sup> V, or Stratix<sup>®</sup> V devices. Use TimeQuest static timing analysis rather than gate-level timing simulation.

## HDL Support

The Quartus II software provides the following HDL support for EDA simulators.

**Table 1-3: HDL Support**

Language	Description
VHDL	<ul style="list-style-type: none"> <li>For VHDL RTL simulation, compile design files directly in your simulator. To use NativeLink automation, analyze and elaborate your design in the Quartus II software, and then use the NativeLink simulator scripts to compile the design files in your simulator. You must also compile simulation models from the Altera simulation libraries and simulation models for the IP cores in your design. Use the Simulation Library Compiler or NativeLink to compile simulation models.</li> <li>For gate-level simulation, the EDA Netlist Writer generates a synthesized design netlist VHDL Output File (.vho). Compile the .vho in your simulator. You may also need to compile models from the Altera simulation libraries.</li> <li>IEEE 1364-2005 encrypted Verilog HDL simulation models are encrypted separately for each Altera-supported simulation vendor. If you want to simulate the model in a VHDL design, you need either a simulator that is capable of VHDL/Verilog HDL co-simulation, or any Mentor Graphics single language VHDL simulator.</li> </ul>
Verilog HDL SystemVerilog	<ul style="list-style-type: none"> <li>For RTL simulation in Verilog HDL or SystemVerilog, compile your design files in your simulator. To use NativeLink automation, analyze and elaborate your design in the Quartus II software, and then use the NativeLink simulator scripts to compile your design files in your simulator. You must also compile simulation models from the Altera simulation libraries and simulation models for the IP cores in your design. Use the Simulation Library Compiler or NativeLink to compile simulation models.</li> <li>For gate-level simulation, the EDA Netlist Writer generates a synthesized design netlist Verilog Output File (.vo). Compile the .vo in your simulator.</li> </ul>
Mixed HDL	<ul style="list-style-type: none"> <li>If your design is a mix of VHDL, Verilog HDL, and SystemVerilog files, you must use a mixed language simulator. Since Altera supports both languages, choose the most convenient language for any Altera IP core in your design.</li> <li>Altera provides Arria V, Cyclone V, Stratix V, and newer simulation model libraries and IP simulation models in Verilog HDL and IEEE encrypted Verilog. Your simulator's co-simulation capabilities support VHDL simulation of these models using VHDL “wrapper” files. Altera provides the wrapper for Verilog models to instantiate these models directly from your VHDL design.</li> </ul>
Schematic	<p>You must convert schematics to HDL format before simulation. You can use the converted VHDL or Verilog HDL files for RTL simulation.</p>

## Simulation Flows

The Quartus II software supports integration with EDA simulators.

**Table 1-4: Simulation Flows**

Simulation Flow	Description
NativeLink flow	<p>The NativeLink automated flow supports a variety of design flows. Do not use NativeLink if you require direct control over every aspect of simulation.</p> <ul style="list-style-type: none"> <li>• Use NativeLink to generate simulation scripts to compile your design and simulation libraries, and to automatically launch your simulator.</li> <li>• Specify your own compilation, elaboration, and simulation scripts for testbench and simulation model files that have not been analyzed by the Quartus II software.</li> <li>• Use NativeLink to supplement your scripts by automatically compiling design files, IP simulation model files, and Altera simulation library models.</li> </ul>
Custom flows	<p>Custom flows support manual control of all aspects of simulation, including the following:</p> <ul style="list-style-type: none"> <li>• Manually compile and simulate testbench, design, IP, and simulation model libraries, or write scripts to automate compilation and simulation in your simulator.</li> <li>• Use the Simulation Library Compiler to compile simulation libraries for all Altera devices and supported third-party simulators and languages.</li> </ul> <p>Use the custom flow if you require any of the following:</p> <ul style="list-style-type: none"> <li>• Custom compilation commands for design, IP, or simulation library model files (for example, macros, debugging or optimization options, or other simulator-specific options).</li> <li>• Multi-pass simulation flows.</li> <li>• Flows that use dynamically generated simulation scripts.</li> </ul>
Specialized flows	<p>Altera supports specialized flows for various design variations, including the following:</p> <ul style="list-style-type: none"> <li>• For simulation of Altera example designs, refer to the documentation for the example design or to the IP core user guide.</li> <li>• For simulation of Qsys designs, refer to <i>Creating a System with Qsys</i>.</li> <li>• For simulation of designs that include the Nios II embedded processor, refer to <i>Simulating a Nios II Embedded Processor</i>.</li> </ul>

### Related Information

- [IP User Guide Documentation](#)
- [Creating a System with Qsys](#)
- [Simulating a Nios II Embedded Processor](#)

## Preparing for Simulation

Preparing for RTL or gate-level simulation involves compiling the RTL or gate-level representation of your design and testbench. You must also compile IP simulation models, models from the Altera simulation libraries, and any other model libraries required for your design.

### Compiling Simulation Models

The Quartus II software includes simulation models for Altera IP cores.

These models include IP functional simulation models, and device family-specific models in the *<Quartus II installation path>/eda/sim\_lib* directory. These models include IEEE encrypted Verilog HDL models for both Verilog HDL and VHDL simulation. Before running simulation, you must compile the appropriate simulation models from the Altera simulation libraries.

Use any of the following methods to compile Altera simulation models:

- Use the NativeLink feature to automatically compile your design, Altera IP, simulation model libraries, and testbench.
- Run the Simulation Library Compiler to compile all RTL and gate-level simulation model libraries for your device, simulator, and design language.
- Compile Altera simulation models manually with your simulator.

After you compile the simulation model libraries, you can reuse these libraries in subsequent simulations to avoid having to compile them again.

**Note:** The specified timescale precision must be within 1ps when using Altera simulation models.

#### Related Information

[Altera Simulation Models](#)

### Generating IP Simulation Files for RTL Simulation

The Quartus II software supports both Verilog HDL and VHDL simulation of encrypted and unencrypted Altera IP cores. If your design includes Altera IP cores, you must compile any corresponding IP simulation models in your simulator with the rest of your design and testbench. The Quartus II software generates and copies the simulation models for IP cores to your project directory.

You can use the following files to simulate your Altera IP variation.

Table 1-5: Altera IP Simulation Files

File Type	Description	File Name
Simulator setup script	Simulator-specific script to compile, elaborate, and simulate Altera IP models and simulation model library files. Copy the commands into your simulation script, or edit these files to compile, elaborate, and simulate your design and testbench.	Cadence <ul style="list-style-type: none"> <li>• <b>cds.lib</b></li> <li>• <b>ncsim_setup.sh</b></li> <li>• <b>hdl.var</b></li> </ul> Mentor Graphics <ul style="list-style-type: none"> <li>• <b>msim_setup.tcl</b></li> </ul> Synopsys <ul style="list-style-type: none"> <li>• <b>synopsys_sim.setup</b></li> <li>• <b>vcs_setup.sh</b></li> <li>• <b>vcsmx_setup.sh</b></li> </ul> Aldec <ul style="list-style-type: none"> <li>• <b>rivierapro_setup.tcl</b></li> </ul>
Quartus II Simulation IP File (.sip)	Contains IP core simulation library mapping information. The .sip files enable NativeLink simulation and the Quartus II Archiver for IP cores.	<design name>.sip
IP functional simulation models	IP functional simulation models are cycle-accurate VHDL or Verilog HDL models generated by the Quartus II software for some Altera IP cores. IP functional simulation models support fast functional simulation of IP using industry-standard VHDL and Verilog HDL simulators.	<my_ip>.vho  <my_ip>.vo
IEEE encrypted models	Arria V, Cyclone V, Stratix V, and newer simulation model libraries and IP simulation models are provided in Verilog HDL and IEEE encrypted Verilog HDL. VHDL simulation of these models is supported using your simulator's co-simulation capabilities. IEEE encrypted Verilog HDL models are significantly faster than IP functional simulation models.	<my_ip>.v

### Generating IP Functional Simulation Models for RTL Simulation

Altera provides IP functional simulation models for some Altera IP cores. To generate IP functional simulation models, follow these steps:

- Turn on the **Generate Simulation Model** option when parameterizing the IP core.
- When you simulate your design, compile only the .vo or .vho for these IP cores in your simulator. In this case you should not compile the corresponding HDL file. The encrypted HDL file supports synthesis by only the Quartus II software.

**Note:** Altera IP cores that do not require IP functional simulation models for simulation, do not provide the **Generate Simulation Model** option in the IP core parameter editor.

**Note:** Many recently released Altera IP cores support RTL simulation using IEEE Verilog HDL encryption. IEEE encrypted models are significantly faster than IP functional simulation models. You can simulate the models in both Verilog HDL and VHDL designs.

**Related Information**

[AN 343: OpenCore Evaluation of AMPP Megafunctions](#)

## Running a Simulation (NativeLink Flow)

The NativeLink feature integrates your EDA simulator with the Quartus II software and automates the following simulation steps:

- Set and reuse simulation settings
- Generate simulator-specific files and simulation scripts
- Compile Altera simulation libraries
- Launch your simulator automatically following Quartus II Analysis & Elaboration, Analysis & Synthesis, or after a full compilation.

## Setting Up Simulation (NativeLink Flow)

Before running simulation using the NativeLink flow, you must specify settings for your simulator in the Quartus II software. To specify simulation settings in the Quartus II software, follow these steps:

1. Open a Quartus II project.
2. Click **Tools > Options** and specify the location of your simulator executable file .

**Table 1-6: Execution Paths for EDA Simulators**

Simulator	Path
Mentor Graphics ModelSim-Altera	<drive letter>:\<simulator install path>\ <b>win32aloem</b> (Windows)  /<simulator install path>/ <b>bin</b> (Linux)
Mentor Graphics ModelSim Mentor Graphics QuestaSim	<drive letter>:\<simulator install path>\ <b>win32</b> (Windows)  <simulator install path>/ <b>bin</b> (Linux)
Synopsys VCS/VCS MX	<simulator install path>/ <b>bin</b> (Linux)
Cadence Incisive Enterprise	<simulator install path>/ <b>tools/bin</b> (Linux)

Simulator	Path
Aldec Active-HDL Aldec Riviera-PRO	<drive letter>:\<simulator install path>\bin (Windows) <simulator install path>/bin (Linux)

3. Click **Assignments > Settings** and specify options on the **Simulation** page and **More NativeLink Settings** dialog box. Specify default options for simulation library compilation, netlist and tool command script generation, and for launching RTL or gate-level simulation automatically following Quartus II processing.
4. If your design includes a testbench, turn on **Compile test bench** and then click **Test Benches** to specify options for each testbench. Alternatively, turn on **Use script to compile testbench** and specify the script file.
5. If you want to use a script to setup simulation, turn on **Use script to setup simulation**.

## Running RTL Simulation (NativeLink Flow)

To run RTL simulation using the NativeLink flow, follow these steps:

1. Set up the simulation environment.
2. Click **Processing > Start > Analysis and Elaboration**.
3. Click **Tools > Run Simulation Tool > RTL Simulation**.

NativeLink compiles simulation libraries and launches and runs your RTL simulator automatically according to the NativeLink settings.

4. Review and analyze the simulation results in your simulator. Correct any functional errors in your design. If necessary, re-simulate the design to verify correct behavior.

## Running Gate-Level Simulation (NativeLink Flow)

To run gate-level simulation with the NativeLink flow, follow these steps:

1. Prepare for simulation.
2. Set up the simulation environment. To generate only a functional (rather than timing) gate-level netlist, click **More EDA Netlist Writer Settings**, and turn on **Generate netlist for functional simulation only**.
3. To synthesize the design, follow one of these steps:
  - To generate a post-fit functional or post-fit timing netlist and then automatically simulate your design according to your NativeLink settings, Click **Processing > Start Compilation**. Skip to step 6.
  - To synthesize the design for post-synthesis functional simulation only, click **Processing > Start > Start Analysis and Synthesis**.
4. To generate the simulation netlist, click **Start EDA Netlist Writer**.
5. Click **Tools > Run Simulation Tool > Gate Level Simulation**.
6. Review and analyze the simulation results in your simulator. Correct any unexpected or incorrect conditions found in your design. Simulate the design again until you verify correct behavior.

## Running a Simulation (Custom Flow)

Use a custom simulation flow to support any of the following more complex simulation scenarios:

- Custom compilation, elaboration, or run commands for your design, IP, or simulation library model files (for example, macros, debugging/optimization options, simulator-specific elaboration or run-time options)
- Multi-pass simulation flows
- Flows that use dynamically generated simulation scripts

Use these to compile libraries and generate simulation scripts for custom simulation flows:

- NativeLink-generated scripts—use NativeLink only to generate simulation script templates to develop your own custom scripts.
- Simulation Library Compiler—compile Altera simulation libraries for your device, HDL, and simulator. Generate scripts to compile simulation libraries as part of your custom simulation flow. This tool does not compile your design, IP, or testbench files.
- IP and Qsys simulation scripts—use the scripts generated for Altera IP cores and Qsys systems as templates to create simulation scripts. If your design includes multiple IP cores or Qsys systems, you can combine the simulation scripts into a single script, manually or by using the `ip-make-simscript` utility.

Use the following steps in a custom simulation flow:

1. Compile the design and testbench files in your simulator.
2. Run the simulation in your simulator.

Post-synthesis and post-fit gate-level simulations run significantly slower than RTL simulation. Altera recommends that you verify your design using RTL simulation for functionality and use the TimeQuest timing analyzer for timing. Timing simulation is not supported for Arria V, Cyclone V, Stratix V, and newer families.

### Related Information

#### [Running EDA Simulators](#)

## Generating Simulation Scripts

You can automatically generate simulation scripts to set up supported simulators. These scripts compile the required device libraries and system design files in the correct order, and then elaborate or load the top-level design for simulation. You can also use scripts to modify the top-level simulation environment, independent of IP simulation files that are replaced during regeneration. You can modify the scripts to set up supported simulators.

Use the NativeLink feature to generate simulation scripts to automate simulation steps. You can reuse these generated files and simulation scripts in a custom simulation flow. NativeLink optionally generates scripts for your simulator in the project subdirectory.

1. Click **Assignments > Settings**.
2. Under **EDA Tool Settings**, click **Simulation**.
3. Select the **Tool name** of your simulator.
4. Click **More NativeLink Settings**.
5. Turn on **Generate third-party EDA tool command scripts without running the EDA tool**.



Table 1-7: NativeLink Generated Scripts for RTL Simulation

Simulator(s)	Simulation File	Use
Mentor Graphics ModelSim QuestaSim	<code>/simulation/modelsim/&lt;my_ip&gt;.do</code>	Source directly with your simulator.
Aldec Riviera Pro	<code>/simulation/modelsim/&lt;my_ip&gt;.do</code>	Source directly with your simulator.
Synopsys VCS	<code>/simulation/modelsim/&lt;revision name&gt;_&lt;rtl or gate&gt;.vcs</code>	Add your testbench file name to this options file to pass the file to VCS using the <code>-file</code> option. If you specify a testbench file to NativeLink, NativeLink generates an <code>.sh</code> script that runs VCS.
Synopsys VCS MX	<code>/simulation/scsim/&lt;revision name&gt;_vcsmx_&lt;rtl or gate&gt;_&lt;verilog or vhd&gt;.tcl</code>	Run this script at the command line using the command: <code>quartus_sh -t &lt;script&gt;</code> Any testbench you specify with NativeLink is included in this script.
Cadence Incisive (NC SIM)	<code>/simulation/ncsim/&lt;revision name&gt;_ncsim_&lt;rtl or gate&gt;_&lt;verilog or vhd&gt;.tcl</code>	Run this script at the command line using the command: <code>quartus_sh -t &lt;script&gt;</code> . Any testbench you specify with NativeLink is included in this script.

You can use the following script variables:

- `TOP_LEVEL_NAME`—The top-level entity of your simulation is often a testbench that instantiates your design, and then your design instantiates IP cores and/or Qsys systems. Set the value of `TOP_LEVEL_NAME` to the top-level entity.
- `QSYS_SIMDIR`—Specifies the top-level directory containing the simulation files.
- Other variables control the compilation, elaboration, and simulation process.

## Generating Custom Simulation Scripts with ip-make-simscript

Use the `ip-make-simscript` utility to generate simulation command scripts for multiple IP cores or Qsys systems. Specify all Simulation Package Descriptor files (`.spd`), each of which lists the required simulation files for the corresponding IP core or Qsys system. The IP parameter editor generates the `.spd` files.

`ip-make-simscript` compiles IP simulation models into various simulation libraries. Use the `compile-to-work` option to compile all simulation files into a single work library. Use this option only if you require a simplified library structure.

When you specify multiple `.spd` files, the `ip-make-simscript` utility generates a single simulation script containing all required simulation information. The default value of `TOP_LEVEL_NAME` is the `TOP_LEVEL_NAME` defined in the IP core or Qsys `.spd` file.

Set appropriate variables in the script, or edit the variable assignment directly in the script. If the simulation script is a Tcl file that is sourced in the simulator, set the variables before sourcing the script. If the simulation script is a shell script, pass in the variables as command-line arguments to the shell script.

- To run `ip-make-simscript`, type the following at the command prompt:

```
<Quartus installation path>\quartus\sopc_builder\bin\ip-make-simscript
```

**Table 1-8: ip-make-simscript Examples**

Option	Description	Status
<code>--spd=&lt;file&gt;</code>	Describes the list of compiled files and memory model hierarchy. If your design includes multiple IP cores or Qsys systems that include <code>.spd</code> files, use this option for each file. For example:  <code>ip-make-simscript --spd=ip1.spd --spd=ip2.spd</code>	Required
<code>--output-directory=&lt;directory&gt;</code>	Specifies the location of output files. If unspecified, the default setting is the directory from which <code>ip-make-simscript</code> is run.	Optional
<code>--compile-to-work</code>	Compiles all design files to the default work library. Use this option only if you encounter problems managing your simulation with multiple libraries.	Optional
<code>--use-relative-paths</code>	Uses relative paths whenever possible.	Optional

## Document Revision History

Date	Version	Changes
2014.06.30	14.0.0	<ul style="list-style-type: none"> <li>Replaced MegaWizard Plug-In Manager information with IP Catalog.</li> </ul>
May 2013	13.0.0	<ul style="list-style-type: none"> <li>Updated introductory section and system and IP file locations.</li> </ul>
November 2012	12.1.0	<ul style="list-style-type: none"> <li>Revised chapter to reflect latest changes to other simulation documentation.</li> </ul>
June 2012	12.0.0	<ul style="list-style-type: none"> <li>Reorganization of chapter to reflect various simulation flows.</li> <li>Added NativeLink support for newer IP cores.</li> </ul>
November 2011	11.1.0	<ul style="list-style-type: none"> <li>Added information about encrypted Altera simulation model files.</li> <li>Added information about IP simulation and NativeLink.</li> </ul>

**Related Information**

[Quartus II Handbook Archive](#)



2014.06.30

QIISV3



Subscribe



Send Feedback

You can integrate your supported EDA simulator into the Quartus II design flow. This document provides guidelines for simulation of Quartus® II designs with Mentor Graphics® ModelSim-Altera®, ModelSim, or QuestaSim software. Altera provides the entry-level ModelSim-Altera software, along with precompiled Altera simulation libraries, to simplify simulation of Altera designs.

## Related Information

[Simulating Altera Designs](#) on page 1-1

[Managing Quartus II Projects](#)

## Quick Start Example (ModelSim with Verilog)

You can adapt the following RTL simulation example to get started quickly with ModelSim:

1. Type the following to specify your EDA simulator and executable path in the Quartus II software:

```
set_user_option -name EDA_TOOL_PATH_MODELSIM <modelsim executable path>
set_global_assignment -name EDA_SIMULATION_TOOL "MODELSIM (verilog)"
```

2. Compile simulation model libraries using one of the following methods:

- Run NativeLink RTL simulation to compile required design files, simulation models, and run your simulator. Verify results in your simulator. If you complete this step you can ignore the remaining steps.
- Use Quartus II Simulation Library Compiler to automatically compile all required simulation models for your design.
- Type the following commands to create and map Altera simulation libraries manually, and then compile the models manually:

```
vlib <lib1>_ver
vmap <lib1>_ver <lib1>_ver
vlog -work <lib1> <lib1>
```

3. Compile your design and testbench files:

```
vlog -work work <design or testbench name>.v
```

4. Load the design:

```
sim -L work -L <lib1>_ver -L <lib2>_ver work.<testbench name>
```

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered

## ModelSim, ModelSim-Altera, and QuestaSim Guidelines

The following guidelines apply to simulation of Altera designs in the ModelSim, ModelSim-Altera, or QuestaSim software.

### Using ModelSim-Altera Precompiled Libraries

Precompiled libraries for both functional and gate-level simulations are provided for the ModelSim-Altera software. You should not compile these library files before running a simulation. No precompiled libraries are provided for ModelSim or QuestaSim. You must compile the necessary libraries to perform functional or gate-level simulation with these tools.

The precompiled libraries provided in `<ModelSim-Altera path>/altera/` must be compatible with the version of the Quartus II software that creates the simulation netlist. To verify compatibility of precompiled libraries with your version of the Quartus II software, refer to the `<ModelSim-Altera path>/altera/version.txt` file. This file indicates the Quartus II software version and build of the precompiled libraries.

**Note:** Encrypted Altera simulation model files shipped with the Quartus II software version 10.1 and later can only be read by ModelSim-Altera Edition Software version 6.6c and later. These encrypted simulation model files are located at the `<Quartus II System directory>/quartus/eda/sim_lib/mentor` directory.

#### Related Information

[ModelSim-Altera Precompiled Libraries](#)

[Altera Simulation Models](#)

### Disabling Timing Violation on Registers

In certain situations, you may want to ignore timing violations on registers and disable the “X” propagation that occurs. For example, this technique may be helpful to eliminate timing violations in internal synchronization registers in asynchronous clock-domain crossing.

#### Before you begin

By default, the `x_on_violation_option` logic option is enabled for all design registers, resulting in an output of “X” at timing violation. To disable “X” propagation at timing violations on a specific register, disable the `x_on_violation_option` logic option for the specific register, as shown in the following example from the Quartus II Settings File (`.qsf`).

```
set_instance_assignment -name X_ON_VIOLATION_OPTION OFF -to \ <register_name>
```

### Passing Parameter Information from Verilog HDL to VHDL

You must use in-line parameters to pass values from Verilog HDL to VHDL.

#### Before you begin

By default, the `x_on_violation_option` logic option is enabled for all design registers, resulting in an output of “X” at timing violation. To disable “X” propagation at timing violations on a specific register, disable the `x_on_violation_option` logic option for the specific register, as shown in the following example from the Quartus II Settings File (`.qsf`).

```
set_instance_assignment -name X_ON_VIOLATION_OPTION OFF -to \ <register_name>
```

### Example 2-1: In-line Parameter Passing Example

```
lpm_add_sub#(.lpm_width(12), .lpm_direction("Add"),
.lpm_type("LPM_ADD_SUB"),
.lpm_hint("ONE_INPUT_IS_CONSTANT=NO,CIN_USED=NO" ))

lpm_add_sub_component (
    .dataa (dataa),
    .datab (datab),
    .result (sub_wire0)
);
```

**Note:** The sequence of the parameters depends on the sequence of the GENERIC in the VHDL component declaration.

## Increasing Simulation Speed

By default, the ModelSim and QuestaSim software runs in a debug-optimized mode.

### Before you begin

To run the ModelSim and QuestaSim software in speed-optimized mode, add the following two vlog command-line switches. In this mode, module boundaries are flattened and loops are optimized, which eliminates levels of debugging hierarchy and may result in faster simulation. This switch is not supported in the ModelSim-Altera simulator.

```
vlog -fast -05
```

## Simulating Transport Delays

By default, the ModelSim and QuestaSim software filter out all pulses that are shorter than the propagation delay between primitives.

Turning on the **transport delay** options in the ModelSim and QuestaSim software prevents the simulator from filtering out these pulses.

**Table 2-1: Transport Delay Simulation Options (ModelSim and QuestaSim)**

Option	Description
+transport_path_delays	Use when simulation pulses are shorter than the delay in a gate-level primitive. You must include the +pulse_e/number and +pulse_r/number options.
+transport_int_delays	Use when simulation pulses are shorter than the interconnect delay between gate-level primitives. You must include the +pulse_int_e/number and +pulse_int_r/number options.

**Note:** The +transport\_path\_delays and +transport\_path\_delays options apply automatically during NativeLink gate-level timing simulation. For more information about either of these options, refer

to the ModelSim-Altera Command Reference installed with the ModelSim and QuestaSim software.

The following ModelSim and QuestaSim software command shows the command line syntax to perform a gate-level timing simulation with the device family library:

```
vsim -t lps -L stratixii -sdftyp /il=filtref_vhd.sdo work.filtref_vhd_vec_tst \  
+transport_int_delays +transport_path_delays
```

## Viewing Error Messages

ModelSim and QuestaSim error and warning messages are tagged with a vsim or vcom code. To determine the cause and resolution for a vsim or vcom error or warning, use the verror command.

For example, ModelSim may return the following error:

```
# ** Error: C:/altera_trn/DUALPORT_TRY/simulation/modelsim/DUALPORT_TRY.vho(31):  
(vcom-1136) Unknown identifier "stratixiv"
```

In this case, type the following command:

```
verror 1136
```

The following description appears:

```
# vcom Message # 1136:  
# The specified name was referenced but was not found. This indicates  
# that either the name specified does not exist or is not visible at  
# this point in the code.
```

## Generating Power Analysis Files

To generate a timing Value Change Dump File (.vcd) for power analysis, you must first generate a `<filename>_dump_all_vcd_nodes.tcl` script file in the Quartus II software. You can then run the script from the ModelSim, QuestaSim, or ModelSim-Altera software to generate a timing `<filename>.vcd`. for use in the Quartus II PowerPlay power analyzer.

### Before you begin

To generate and use a .vcd for power analysis, follow these steps:

1. In the Quartus II software, click **Assignments > Settings**.
2. Under **EDA Tool Settings**, click **Simulation**.
3. Turn on **Generate Value Change Dump file script**, specify the type of output signals to include, and specify the top-level design instance name in your testbench.
4. Click **Processing > Start Compilation**.
5. Click **Tools > Run EDA Simulation > EDA Gate Level Simulation**. The Compiler creates the `<filename>_dump_all_vcd_nodes.tcl` file, the ModelSim simulation `<filename>_run_msim_gate_vhdl/verilog.do` file (including the .vcd and .tcl execution lines), and all other files for simulation. ModelSim then automatically runs the generated .do to start the simulation.
6. Stop the simulation if your testbench does not have a break point. ModelSim generates the .vcd only after simulation ends with the End Simulation function.

## Viewing Simulation Waveforms

ModelSim-Altera, ModelSim, and QuestaSim automatically generate a Wave Log Format File (.wlf) following simulation. You can use the .wlf to generate a waveform view.

### Before you begin

To view a waveform from a .wlf through ModelSim-Altera, ModelSim, or QuestaSim, perform the following steps:

1. Type `vsim` at the command line. The **ModelSim/QuestaSim** or **ModelSim-Altera** dialog box appears.
2. Click **File > Datasets**. The **Datasets Browser** dialog box appears.
3. Click **Open** and select your .wlf.
4. Click **Done**.
5. In the Object browser, select the signals that you want to observe.
6. Click **Add > Wave**, and then click **Selected Signals**.  
You must first convert the .vcd to a .wlf before you can view a waveform in ModelSim-Altera, ModelSim, or QuestaSim.
7. To convert the the .vcd to a .wlf, type the following at the command-line:

```
vcd2wlf <example>.vcd <example>.wlf
```

8. After conversion, view the .wlf waveform in ModelSim or QuestaSim.  
You can convert your .wlf to a .vcd by using the `wlf2vcd` command

## Simulating with ModelSim-Altera Waveform Editor

You can use the ModelSim-Altera Waveform Editor as a simple method to create stimulus vectors for simulation. You can create this design stimulus via interactive manipulation of waveforms from the wave window in ModelSim-Altera. With the ModelSim-Altera waveform editor, you can create and edit waveforms, drive simulation directly from created waveforms, and save created waveforms into a stimulus file.

### Related Information

[ModelSim Web Page](#)

## ModelSim Simulation Setup Script Example

The Quartus II software can generate a `msim_setup.tcl` simulation setup script for IP cores in your design. The script compiles the required device library models, compiles the design files, and elaborates the design with or without simulator optimization. To run the script, type `source msim_setup.tcl` in the simulator Transcript window.

Alternatively, if you are using the simulator at the command line, you can type the following command:

```
vsim -c -do msim_setup.tcl
```

In this example the **top-level-simulate.do** custom top-level simulation script sets the hierarchy variable `TOP_LEVEL_NAME` to `top_testbench` for the design, and sets the variable `QSYS_SIMDIR` to the location of the generated simulation files.

```
# Set hierarchy variables used in the IP-generated files
```



```

set TOP_LEVEL_NAME "top_testbench"
set QSYS_SIMDIR "./ip_top_sim"
# Source generated simulation script which defines aliases used below
source $QSYS_SIMDIR/mentor/msim_setup.tcl
# dev_com alias compiles simulation libraries for device library files
dev_com
# com alias compiles IP simulation or Qsys model files and/or Qsys model files in
the correct order
com
# Compile top level testbench that instantiates your IP
vlog -sv ./top_testbench.sv
# elab alias elaborates the top-level design and testbench
elab
# Run the full simulation
run - all

```

In this example, the top-level simulation files are stored in the same directory as the original IP core, so this variable is set to the IP-generated directory structure. The `QSYS_SIMDIR` variable provides the relative hierarchy path for the generated IP simulation files. The script calls the generated `msim_setup.tcl` script and uses the alias commands from the script to compile and elaborate the IP files required for simulation along with the top-level simulation testbench. You can specify additional simulator elaboration command options when you run the `elab` command, for example, `elab +nowarnTFMPC`. The last command run in the example starts the simulation.

## Unsupported Features

The Quartus II software does not support the following ModelSim simulation features:

- Altera does not support companion licensing for ModelSim AE.
- The USB software guard is not supported by versions earlier than Mentor Graphics ModelSim software version 5.8d.
- For ModelSim-Altera software versions prior to 5.5b, use the **PCLS** utility included with the software to set up the license.
- Some versions of ModelSim and QuestaSim support SystemVerilog, PSL assertions, SystemC, and more. For more information about specific feature support, refer to Mentor Graphics literature

### Related Information

- [ModelSim-Altera Software Web Page](#)

## Document Revision History

Table 2-2: Document Revision History

Date	Version	Changes
2014.06.30	14.0.0	<ul style="list-style-type: none"> <li>• Replaced MegaWizard Plug-In Manager information with IP Catalog.</li> </ul>
November 2012	12.1.0	<ul style="list-style-type: none"> <li>• Relocated general simulation information to Simulating Altera Designs.</li> </ul>

Date	Version	Changes
June 2012	12.0.0	<ul style="list-style-type: none"><li>Removed survey link.</li></ul>
November 2011	11.0.1	<ul style="list-style-type: none"><li>Changed to new document template.</li></ul>

**Related Information**

[Quartus II Handbook Archive](#)

2014.06.30

QI15V3



Subscribe



Send Feedback

You can integrate your supported EDA simulator into the Quartus II design flow. This document provides guidelines for simulation of Quartus® II designs with the Synopsys VCS or VCS MX software.

## Quick Start Example (VCS with Verilog)

You can adapt the following RTL simulation example to get started quickly with VCS:

1. Type the following to specify your EDA simulator and executable path in the Quartus II software:

```
set_user_option -name EDA_TOOL_PATH_VCS <VCS executable path>
set_global_assignment -name EDA_SIMULATION_TOOL "VCS"
```

2. Compile simulation model libraries using one of the following methods:

- Run NativeLink RTL simulation to compile required design files, simulation models, and run your simulator. Verify results in your simulator. If you complete this step you can ignore the remaining steps.
- Use Quartus II Simulation Library Compiler to automatically compile all required simulation models for your design.

3. Modify the **simlib\_comp.vcs** file to specify your design and testbench files.
4. Type the following to run the VCS simulator:

```
vcs -R -file simlib_comp.vcs
```

## VCS and QuestaSim Guidelines

The following guidelines apply to simulation of Altera designs in the VCS or VCS MX software:

- Do not specify the **-v** option for **altera\_insim.sv** because it defines a systemverilog package.
- Add **-verilog** and **+verilog2001ext+.v** options to make sure all **.v** files are compiled as verilog 2001 files, and all other files are compiled as systemverilog files.
- Add the **-lca** option for Stratix V and later families because they include IEEE-encrypted simulation files for VCS and VCS MX.
- Add **-timescale=1ps/1ps** to ensure picosecond resolution.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered

## Simulating Transport Delays

By default, the VCS and VCS MX software filter out all pulses that are shorter than the propagation delay between primitives. Turning on the **transport delay** options in the VCS and VCS MX software prevents the simulator from filtering out these pulses.

**Table 3-1: Transport Delay Simulation Options (VCS and VCS MX)**

Option	Description
+transport_path_delays	Use when simulation pulses are shorter than the delay in a gate-level primitive. You must include the +pulse_e/number and +pulse_r/number options.
+transport_int_delays	Use when simulation pulses are shorter than the interconnect delay between gate-level primitives. You must include the +pulse_int_e/number and +pulse_int_r/number options.

**Note:** The +transport\_path\_delays and +transport\_int\_delays options apply automatically during NativeLink gate-level timing simulation.

The following VCS and VCS MX software command runs a post-synthesis simulation:

```
vcs -R <testbench>.v <gate-level netlist>.v -v <Altera device family \
library>.v +transport_int_delays +pulse_int_e/0 +pulse_int_r/0 \
+transport_path_delays +pulse_e/0 +pulse_r/0
```

## Disabling Timing Violation on Registers

In certain situations, you may want to ignore timing violations on registers and disable the “X” propagation that occurs. For example, this technique may be helpful to eliminate timing violations in internal synchronization registers in asynchronous clock-domain crossing.

### Before you begin

By default, the **x\_on\_violation\_option** logic option is enabled for all design registers, resulting in an output of “X” at timing violation. To disable “X” propagation at timing violations on a specific register, disable the **x\_on\_violation\_option** logic option for the specific register, as shown in the following example from the Quartus II Settings File (.qsf).

```
set_instance_assignment -name X_ON_VIOLATION_OPTION OFF -to \ <register_name>
```

## Generating Power Analysis Files

You can generate a Verilog Value Change Dump File (.vcd) for power analysis in the Quartus II software, and then run the .vcd from the VCS software. Use this .vcd for power analysis in the Quartus II PowerPlay power analyzer.

## Before you begin

To generate and use a `.vcd` for power analysis, follow these steps:

1. In the Quartus II software, click **Assignments > Settings**.
2. Under **EDA Tool Settings**, click **Simulation**.
3. Turn on **Generate Value Change Dump file script**, specify the type of output signals to include, and specify the top-level design instance name in your testbench.
4. Click **Processing > Start Compilation**.
5. Use the following command to include the script in your testbench where the design under test (DUT) is instantiated:

```
include <revision_name>_dump_all_vcd_nodes.v
```

**Note:** Include the script within the testbench module block. If you include the script outside of the testbench module block, syntax errors occur during compilation.

6. Run the simulation with the VCS command. Exit the VCS software when the simulation is finished and the `<revision_name>.vcd` file is generated in the simulation directory.

## VCS Simulation Setup Script Example

The Quartus II software can generate a simulation setup script for IP cores in your design. The scripts contain shell commands that compile the required simulation models in the correct order, elaborate the top-level design, and run the simulation for 100 time units by default. You can run these scripts from a Linux command shell.

The scripts for VCS and VCS MX are `vcs_setup.sh` (for Verilog HDL or SystemVerilog) and `vcsmx_setup.sh` (combined Verilog HDL and SystemVerilog with VHDL). Read the generated `.sh` script to see the variables that are available for override when sourcing the script or redefining directly if you edit the script. To set up the simulation for a design, use the command-line to pass variable values to the shell script.

### Example 3-1: Using Command-line to Pass Simulation Variables

```
sh vcsmx_setup.sh\  
USER_DEFINED_ELAB_OPTIONS=+rad\  
USER_DEFINED_SIM_OPTIONS=+vcs+lic+wait
```

### Example 3-2: Example Top-Level Simulation Shell Script for VCS-MX

```
# Run generated script to compile libraries and IP simulation files  
# Skip elaboration and simulation of the IP variation  
sh ./ip_top_sim/synopsys/vcsmx/vcsmx_setup.sh SKIP_ELAB=1 SKIP_SIM=1  
QSYS_SIMDIR="./ip_top_sim"  
#Compile top-level testbench that instantiates IP  
vlogan -sverilog ./top_testbench.sv  
#Elaborate and simulate the top-level design  
vcs -lca -t ps <elaboration control options> top_testbench  
simv <simulation control options>
```

**Example 3-3: Example Top-Level Simulation Shell Script for VCS**

```
# Run script to compile libraries and IP simulation files
sh ./ip_top_sim/synopsys/vcs/vcs_setup.sh TOP_LEVEL_NAME="top_testbench"\
# Pass VCS elaboration options to compile files and elaborate top-level
  passed to the script as the TOP_LEVEL_NAME
USER_DEFINED_ELAB_OPTIONS="top_testbench.sv"\
# Pass in simulation options and run the simulation for specified amount of
time.
USER_DEFINED_SIM_OPTIONS="<simulation control options>
```

## Document Revision History

**Table 3-2: Document Revision History**

Date	Version	Changes
2014.06.30	14.0.0	<ul style="list-style-type: none"> <li>Replaced MegaWizard Plug-In Manager information with IP Catalog.</li> </ul>
November 2012	12.1.0	<ul style="list-style-type: none"> <li>Relocated general simulation information to Simulating Altera Designs.</li> </ul>
June 2012	12.0.0	<ul style="list-style-type: none"> <li>Removed survey link.</li> </ul>
November 2011	11.0.1	<ul style="list-style-type: none"> <li>Changed to new document template.</li> </ul>

**Related Information**

[Quartus II Handbook Archive](#)

2014.08.18

QI15V3



Subscribe



Send Feedback

You can integrate your supported EDA simulator into the Quartus II design flow. This chapter provides specific guidelines for simulation of Quartus® II designs with the Cadence Incisive Enterprise (IES) software.

## Quick Start Example (NC-Verilog)

You can adapt the following RTL simulation example to get started quickly with IES:

1. Type the following to specify your EDA simulator and executable path in the Quartus II software:

```
set_user_option -name EDA_TOOL_PATH_NCSIM <ncsim executable path>
set_global_assignment -name EDA_SIMULATION_TOOL "NC-Verilog (Verilog)"
```

2. Compile simulation model libraries using one of the following methods:

- Run NativeLink RTL simulation to compile required design files, simulation models, and run your simulator. Verify results in your simulator. If you complete this step you can ignore the remaining steps.
- Use Quartus II Simulation Library Compiler to automatically compile all required simulation models for your design.
- Map Altera simulation libraries by adding the following commands to a cds.lib file:

```
include ${CDS_INST_DIR}/tools/inca/files/cds.lib
DEFINE <lib1>_ver <lib1_ver>
```

Then, compile Altera simulation models manually:

```
vlog -work <lib1_ver>
```

3. Elaborate your design and testbench with IES:

```
ncelab <work library>.<top-level entity name>
```

4. Run the simulation:

```
ncsim <work library>.<top-level entity name>
```

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



## Cadence Incisive Enterprise (IES) Guidelines

The following guidelines apply to simulation of Altera designs in the IES software:

- Do not specify the `-v` option for `altera_Insim.sv` because it defines a systemverilog package.
- Add `-verilog` and `+verilog2001ext+.v` options to make sure all `.v` files are compiled as verilog 2001 files, and all other files are compiled as systemverilog files.
- Add the `-lca` option for Stratix V and later families because they include IEEE-encrypted simulation files for IES.
- Add `-timescale=1ps/1ps` to ensure picosecond resolution.

### Using GUI or Command-Line Interfaces

Altera supports both the IES GUI and command-line simulator interfaces.

To start the IES GUI, type `nc1aunch` at a command prompt.

**Table 4-1: Simulation Executables**

Program	Function
<code>ncvlog</code> <code>ncvhdl</code>	<code>ncvlog</code> compiles your Verilog HDL code and performs syntax and static semantics checks.  <code>ncvhdl</code> compiles your VHDL code and performs syntax and static semantics checks.
<code>ncelab</code>	Elaborates the design hierarchy and determines signal connectivity.
<code>ncsdfc</code>	Performs back-annotation for simulation with VHDL simulators.
<code>ncsim</code>	Runs mixed-language simulation. This program is the simulation kernel that performs event scheduling and executes the simulation code.

### Elaborating Your Design

The simulator automatically reads the `.sdo` file during elaboration of the Quartus II-generated Verilog HDL or SystemVerilog HDL netlist file. The `ncelab` command recognizes the embedded system task `$$sdf_annotate` and automatically compiles and annotates the `.sdo` file by running `ncsdfc` automatically.

VHDL netlist files do not contain system task calls to locate your `.sdf` file; therefore, you must compile the standard `.sdo` file manually. Locate the `.sdo` file in the same directory where you run elaboration or simulation. Otherwise, the `$$sdf_annotate` task cannot reference the `.sdo` file correctly. If you are starting an elaboration or simulation from a different directory, you can either comment out the `$$sdf_annotate` and annotate the `.sdo` file with the GUI, or add the full path of the `.sdo` file.

**Note:** If you use NC-Sim for post-fit VHDL functional simulation of a Stratix V design that includes RAM, an elaboration error might occur if the component declaration parameters are not in the same order as the architecture parameters. Use the `-namemap_mixgen` option with the `ncelab` command to match the component declaration parameter and architecture parameter names.



## Back-Annotating Simulation Timing Data (VHDL Only)

You can back annotate timing information in a Standard Delay Output File (.sdo) for VHDL simulators. To back annotate the .sdo timing data at the command line, follow these steps:

1. To compile the .sdo with the `ncsdfc` program, type the following command at the command prompt. The `ncsdfc` program generates an `<output name>.sdf.X` compiled .sdo file

```
ncsdfc <project name>_vhd.sdo -output <output name>
```

**Note:** If you do not specify an output name, `ncsdfc` uses `<project name>.sdo.X`

2. Specify the compiled .sdf file for the project by adding the following command to an ASCII SDF command file for the project:

```
COMPILED_SDF_FILE = "<project name>.sdf.X" SCOPE = <instance path>
```

3. After compiling the .sdf file, type the following command to elaborate the design:

```
ncelab worklib.<project name>:entity -SDF_CMD_FILE <SDF Command File>
```

### Example 4-1: Example SDF Command File

```
// SDF command file sdf_file  
COMPILED_SDF_FILE = "lpm_ram_dp_test_vhd.sdo.X",  
SCOPE = :tb,  
MTM_CONTROL = "TYPICAL",  
SCALE_FACTORS = "1.0:1.0:1.0",  
SCALE_TYPE = "FROM_MTM";
```

## Disabling Timing Violation on Registers

In certain situations, you may want to ignore timing violations on registers and disable the “X” propagation that occurs. For example, this technique may be helpful to eliminate timing violations in internal synchronization registers in asynchronous clock-domain crossing.

### Before you begin

By default, the `x_on_violation_option` logic option is enabled for all design registers, resulting in an output of “X” at timing violation. To disable “X” propagation at timing violations on a specific register, disable the `x_on_violation_option` logic option for the specific register, as shown in the following example from the Quartus II Settings File (.qsf).

```
set_instance_assignment -name X_ON_VIOLATION_OPTION OFF -to \ <register_name>
```

## Simulating Pulse Reject Delays

By default, the IES software filters out all pulses that are shorter than the propagation delay between primitives.

Setting the pulse reject delays options in the IES software prevents the simulation tool from filtering out these pulses. Use the following options to ensure that all signal pulses are seen in the simulation results.

Table 4-2: Pulse Reject Delay Options

Program	Function
-PULSE_R	Use when simulation pulses are shorter than the delay in a gate-level primitive. The argument is the percentage of delay for pulse reject limit for the path
-PULSE_INT_R	Use when simulation pulses are shorter than the interconnect delay between gate-level primitives. The argument is the percentage of delay for pulse reject limit for the path

## Viewing Simulation Waveforms

IES generates a **.trn** file automatically following simulation. You can use the **.trn** for generating the SimVision waveform view.

### Before you begin

To view a waveform from a **.trn** file through SimVision, follow these steps:

1. Type `simvision` at the command line. The **Design Browser** dialog box appears.
2. Click **File > Open Database** and click the **.trn** file.
3. In the **Design Browser** dialog box, select the signals that you want to observe from the Hierarchy.
4. Right-click the selected signals and click **Send to Waveform Window**.

You cannot view a waveform from a **.vcd** file in SimVision, and the **.vcd** file cannot be converted to a **.trn** file.

## IES Simulation Setup Script Example

The Quartus II software can generate a **ncsim\_setup.sh** simulation setup script for IP cores in your design. The script contains shell commands that compile the required device libraries, IP, or Qsys simulation models in the correct order. The script then elaborates the top-level design and runs the simulation for 100 time units by default. You can run these scripts from a Linux command shell. To set up the simulation script for a design, you can use the command-line to pass variable values to the shell script.

Read the generated **.sh** script to see the variables that are available for you to override when you source the script or that you can redefine directly in the generated **.sh** script. For example, you can specify additional elaboration and simulation options with the variables `USER_DEFINED_ELAB_OPTIONS` and `USER_DEFINED_SIM_OPTIONS`.

### Example 4-2: Example Top-Level Simulation Shell Script for Incisive (NCSIM)

```
# Run script to compile libraries and IP simulation files
# Skip elaboration and simulation of the IP variation
sh ./ip_top_sim/cadence/ncsim_setup.sh SKIP_ELAB=1 SKIP_SIM=1 QSYS_SIMDIR="./ip_top_sim"

#Compile the top-level testbench that instantiates your IP
ncvlog -sv ./top_testbench.sv
#Elaborate and simulate the top-level design
```

```
ncelab <elaboration control options> top_testbench  
ncsim <simulation control options> top_testbench
```

## Document Revision History

Table 4-3: Document Revision History

Date	Version	Changes
2014.08.18	14.0.a10.0	<ul style="list-style-type: none"><li>Corrected incorrect references to VCS and VCS MX.</li></ul>
2014.06.30	14.0.0	<ul style="list-style-type: none"><li>Replaced MegaWizard Plug-In Manager information with IP Catalog.</li></ul>
November 2012	12.1.0	<ul style="list-style-type: none"><li>Relocated general simulation information to Simulating Altera Designs.</li></ul>
June 2012	12.0.0	<ul style="list-style-type: none"><li>Removed survey link.</li></ul>
November 2011	11.0.1	<ul style="list-style-type: none"><li>Changed to new document template.</li></ul>

### Related Information

[Quartus II Handbook Archive](#)

2014.06.30

QI15V3



Subscribe



Send Feedback

You can integrate your supported EDA simulator into the Quartus II design flow. This chapter provides specific guidelines for simulation of Quartus® II designs with the Aldec Active-HDL or Riviera-PRO software.

## Quick Start Example (Active-HDL VHDL)

You can adapt the following RTL simulation example to get started quickly with Active-HDL:

1. Type the following to specify your EDA simulator and executable path in the Quartus II software:

```
set_user_option -name EDA_TOOL_PATH_ACTIVEHDL <Active HDL executable path>
set_global_assignment -name EDA_SIMULATION_TOOL "Active-HDL (VHDL)"
```

2. Compile simulation model libraries using one of the following methods:

- Run NativeLink RTL simulation to compile required design files, simulation models, and run your simulator. Verify results in your simulator. If you complete this step you can ignore the remaining steps.
- Use Quartus II Simulation Library Compiler to automatically compile all required simulation models for your design.
- Compile Altera simulation models manually:

```
vlib <library1> <altera_library1>
vcom -strict93 -dbg -work <library1> <lib1_component/pack.vhd> <lib1.vhd>
```

3. Create and open the workspace:

```
createdesign <workspace name> <workspace path>
opendesign -a <workspace name>.adf
```

4. Create the work library and compile the netlist and testbench files:

```
vlib work
vcom -strict93 -dbg -work work <output netlist> <testbench file>
```

5. Load the design:

```
vsim +access+r -t lps +transport_int_delays +transport_path_delays \
-L work -L <lib1> -L <lib2> work.<testbench module name>
```

6. Run the simulation in the Active-HDL simulator.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



## Aldec Active-HDL and Riviera-PRO Guidelines

The following guidelines apply to simulating Altera designs in the Active-HDL or Riviera-PRO software.

### Compiling SystemVerilog Files

If your design includes multiple SystemVerilog files, you must compile the System Verilog files together with a single `alog` command. If you have Verilog files and SystemVerilog files in your design, you must first compile the Verilog files, and then compile only the SystemVerilog files in the single `alog` command.

### Simulating Transport Delays

By default, the Active-HDL or Riviera-PRO software filters out all pulses that are shorter than the propagation delay between primitives. Turning on the **transport delay** options in the in the Active-HDL or Riviera-PRO software prevents the simulator from filtering out these pulses.

**Table 5-1: Transport Delay Simulation Options**

Option	Description
<code>+transport_path_delays</code>	Use when simulation pulses are shorter than the delay in a gate-level primitive. You must include the <code>+pulse_e/number</code> and <code>+pulse_r/number</code> options.
<code>+transport_int_delays</code>	Use when simulation pulses are shorter than the interconnect delay between gate-level primitives. You must include the <code>+pulse_int_e/number</code> and <code>+pulse_int_r/number</code> options.

**Note:** The `+transport_path_delays` and `+transport_int_delays` options apply automatically during NativeLink gate-level timing simulation.

To perform a gate-level timing simulation with the device family library, type the Active-HDL command:

```
vsim -t lps -L stratixii -sdftyp /il=filtref_vhd.sdo \
work.filtref_vhd_vec_tst +transport_int_delays +transport_path_delays
```

### Disabling Timing Violation on Registers

In certain situations, you may want to ignore timing violations on registers and disable the “X” propagation that occurs. For example, this technique may be helpful to eliminate timing violations in internal synchronization registers in asynchronous clock-domain crossing.

#### Before you begin

By default, the `x_on_violation_option` logic option is enabled for all design registers, resulting in an output of “X” at timing violation. To disable “X” propagation at timing violations on a specific register, disable the `x_on_violation_option` logic option for the specific register, as shown in the following example from the Quartus II Settings File (`.qsf`).

```
set_instance_assignment -name X_ON_VIOLATION_OPTION OFF -to \ <register_name>
```

## Using Simulation Setup Scripts

The Quartus II software can generate a `rivierapro_setup.tcl` simulation setup script for IP cores in your design. The use and content of the script file is similar to the `msim_setup.tcl` file used by the ModelSim simulator.

## Document Revision History

Table 5-2: Document Revision History

Date	Version	Changes
2014.06.30	14.0.0	<ul style="list-style-type: none"> <li>Replaced MegaWizard Plug-In Manager information with IP Catalog.</li> </ul>
November 2012	12.1.0	<ul style="list-style-type: none"> <li>Relocated general simulation information to Simulating Altera Designs.</li> </ul>
June 2012	12.0.0	<ul style="list-style-type: none"> <li>Removed survey link.</li> </ul>
November 2011	11.0.1	<ul style="list-style-type: none"> <li>Changed to new document template.</li> </ul>

### Related Information

[Quartus II Handbook Archive](#)

2014.06.30

QI15V3



Subscribe



Send Feedback

## Timing Analysis Overview

Comprehensive static timing analysis involves analysis of register-to-register, I/O, and asynchronous reset paths. Timing analysis with the TimeQuest Timing Analyzer uses data required times, data arrival times, and clock arrival times to verify circuit performance and detect possible timing violations.

The TimeQuest analyzer determines the timing relationships that must be met for the design to correctly function, and checks arrival times against required times to verify timing. This chapter is an overview of the concepts you need to know to analyze your designs with the TimeQuest analyzer.

### Related Information

- [The Quartus II TimeQuest Timing Analyzer](#) on page 7-1  
For more information about the TimeQuest analyzer flow and TimeQuest examples.

## TimeQuest Terminology and Concepts

Table 6-1: TimeQuest Analyzer Terminology

Term	Definition
nodes	Most basic timing netlist unit. Used to represent ports, pins, and registers.
cells	Look-up tables (LUT), registers, digital signal processing (DSP) blocks, memory blocks, input/output elements, and so on. <sup>(1)</sup>
pins	Inputs or outputs of cells.
nets	Connections between pins.
ports	Top-level module inputs or outputs; for example, device pins.
clocks	Abstract objects representing clock domains inside or outside of your design.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



Term	Definition
------	------------

Notes:

1. For Stratix® devices, the LUTs and registers are contained in logic elements (LE) and modeled as cells.


## Timing Netlists and Timing Paths

The TimeQuest analyzer requires a timing netlist to perform timing analysis on any design. After you generate a timing netlist, the TimeQuest analyzer uses the data to help determine the different design elements in your design and how to analyze timing.

### The Timing Netlist

A sample design for which the TimeQuest analyzer generates a timing netlist equivalent.


Figure 6-1: Sample Design



The timing netlist for the sample design shows how different design elements are divided into cells, pins, nets, and ports.



Figure 6-2: The TimeQuest Analyzer Timing Netlist




### Timing Paths

Timing paths connect two design nodes, such as the output of a register to the input of another register.

Understanding the types of timing paths is important to timing closure and optimization. The TimeQuest analyzer uses the following commonly analyzed paths:

- **Edge paths**—connections from ports-to-pins, from pins-to-pins, and from pins-to-ports.
- **Clock paths**—connections from device ports or internally generated clock pins to the clock pin of a register.
- **Data paths**—connections from a port or the data output pin of a sequential element to a port or the data input pin of another sequential element.
- **Asynchronous paths**—connections from a port or asynchronous pins of another sequential element such as an asynchronous reset or asynchronous clear.

Figure 6-3: Path Types Commonly Analyzed by the TimeQuest Analyzer



In addition to identifying various paths in a design, the TimeQuest analyzer analyzes clock characteristics to compute the worst-case requirement between any two registers in a single register-to-register path. You must constrain all clocks in your design before analyzing clock characteristics.


## Data and Clock Arrival Times

After the TimeQuest analyzer identifies the path type, it can report data and clock arrival times at register pins.

The TimeQuest analyzer calculates data arrival time by adding the launch edge time to the delay from the clock source to the clock pin of the source register, the micro clock-to-output delay ( $\mu t_{CO}$ ) of the source register, and the delay from the source register's data output (Q) to the destination register's data input (D).

The TimeQuest analyzer calculates data required time by adding the latch edge time to the sum of all delays between the clock port and the clock pin of the destination register, including any clock port buffer delays, and subtracts the micro setup time ( $\mu t_{SU}$ ) of the destination register, where the  $\mu t_{SU}$  is the intrinsic setup time of an internal register in the FPGA.

**Figure 6-4: Data Arrival and Data Required Times**



The basic calculations for data arrival and data required times including the launch and latch edges.

**Figure 6-5: Data Arrival and Data Required Time Equations**


Data Arrival Time	=	Launch Edge + Source Clock Delay + $\mu t_{CO}$ + Register-to-Register Delay
Data Required Time	=	Latch Edge + Destination Clock Delay - $\mu t_{SU}$

## Launch and Latch Edges

All timing relies on one or more clocks. In addition to analyzing paths, the TimeQuest analyzer determines clock relationships for all register-to-register transfers in your design.

The following figure shows the launch edge, which is the clock edge that sends data out of a register or other sequential element, and acts as a source for the data transfer. A latch edge is the active clock edge that captures data at the data port of a register or other sequential element, acting as a destination for the data transfer. In this example, the launch edge sends the data from register *reg1* at 0 ns, and the register *reg2* captures the data when triggered by the latch edge at 10 ns. The data arrives at the destination register before the next latch edge.

Figure 6-6: Setup and Hold Relationship for Launch and Latch Edges 10ns Apart



In timing analysis, and with the TimeQuest analyzer specifically, you create clock constraints and assign those constraints to nodes in your design. These clock constraints provide the structure required for repeatable data relationships. The primary relationships between clocks, in the same or different domains, are the setup relationship and the hold relationship.


**Note:** If you do not constrain the clocks in your design, the Quartus II software analyzes in terms of a 1 GHz clock to maximize timing based Fitter effort. To ensure realistic slack values, you must constrain all clocks in your design with real values.

## Clock Setup Check

To perform a clock setup check, the TimeQuest analyzer determines a setup relationship by analyzing each launch and latch edge for each register-to-register path.

For each latch edge at the destination register, the TimeQuest analyzer uses the closest previous clock edge at the source register as the launch edge. The following figure shows two setup relationships, setup A and setup B. For the latch edge at 10 ns, the closest clock that acts as a launch edge is at 3 ns and is labeled setup A. For the latch edge at 20 ns, the closest clock that acts as a launch edge is 19 ns and is labeled setup B. TimeQuest analyzes the most restrictive setup relationship, in this case setup B; if that relationship meets the design requirement, then setup A meets it by default.

Figure 6-7: Setup Check



The TimeQuest analyzer reports the result of clock setup checks as slack values. Slack is the margin by which a timing requirement is met or not met. Positive slack indicates the margin by which a requirement is met; negative slack indicates the margin by which a requirement is not met.

Figure 6-8: Clock Setup Slack for Internal Register-to-Register Paths

Clock Setup Slack	=	Data Required Time – Data Arrival Time
Data Arrival Time	=	Launch Edge + Clock Network Delay to Source Register + $\mu t_{co}$ + Register-to-Register Delay
Data Required Time	=	Latch Edge + Clock Network Delay to Destination Register – $\mu t_{su}$ – Setup Uncertainty

The TimeQuest analyzer performs setup checks using the maximum delay when calculating data arrival time, and minimum delay when calculating data required time.

**Figure 6-9: Clock Setup Slack from Input Port to Internal Register**

Clock Setup Slack	= Data Required Time – Data Arrival Time
Data Arrival Time	= Launch Edge + Clock Network Delay + Input Maximum Delay + Port-to-Register Delay
Data Required Time	= Latch Edge + Clock Network Delay to Destination Register – $\mu_{t_{su}}$ – Setup Uncertainty

**Figure 6-10: Clock Setup Slack from Internal Register to Output Port**


Clock Setup Slack	= Data Required Time – Data Arrival Time
Data Required Time	= Latch Edge + Clock Network Delay to Output Port – Output Maximum Delay
Data Arrival Time	= Launch Edge + Clock Network Delay to Source Register + $\mu_{t_{co}}$ + Register-to-Port Delay

## Clock Hold Check

To perform a clock hold check, the TimeQuest analyzer determines a hold relationship for each possible setup relationship that exists for all source and destination register pairs. The TimeQuest analyzer checks all adjacent clock edges from all setup relationships to determine the hold relationships.

The TimeQuest analyzer performs two hold checks for each setup relationship. The first hold check determines that the data launched by the current launch edge is not captured by the previous latch edge. The second hold check determines that the data launched by the next launch edge is not captured by the current latch edge. From the possible hold relationships, the TimeQuest analyzer selects the hold relationship that is the most restrictive. The most restrictive hold relationship is the hold relationship with the smallest difference between the latch and launch edges and determines the minimum allowable delay for the register-to-register path. In the following example, the TimeQuest analyzer selects hold check A2 as the most restrictive hold relationship of two setup relationships, setup A and setup B, and their respective hold checks.

**Figure 6-11: Setup and Hold Check Relationships**



**Figure 6-12: Clock Hold Slack for Internal Register-to-Register Paths**

Clock Hold Slack	= Data Arrival Time – Data Required Time
Data Arrival Time	= Launch Edge + Clock Network Delay to Source Register + $\mu_{t_{co}}$ + Register-to-Register Delay
Data Required Time	= Latch Edge + Clock Network Delay to Destination Register + $\mu_{t_{h}}$ + Hold Uncertainty

The TimeQuest analyzer performs hold checks using the minimum delay when calculating data arrival time, and maximum delay when calculating data required time.

**Figure 6-13: Clock Hold Slack Calculation from Input Port to Internal Register**

$$\begin{aligned} \text{Clock Hold Slack} &= \text{Data Arrival Time} - \text{Data Required Time} \\ \text{Data Arrival Time} &= \text{Launch Edge} + \text{Clock Network Delay} + \text{Input Minimum Delay} + \text{Pin-to-Register Delay} \\ \text{Data Required Time} &= \text{Latch Edge} + \text{Clock Network Delay to Destination Register} + \mu t_{\text{h}} \end{aligned}$$

**Figure 6-14: Clock Hold Slack Calculation from Internal Register to Output Port**

$$\begin{aligned} \text{Clock Hold Slack} &= \text{Data Arrival Time} - \text{Data Required Time} \\ \text{Data Arrival Time} &= \text{Latch Edge} + \text{Clock Network Delay to Source Register} + \mu t_{\text{co}} + \text{Register-to-Pin Delay} \\ \text{Data Required Time} &= \text{Latch Edge} + \text{Clock Network Delay} - \text{Output Minimum Delay} \end{aligned}$$

## Recovery and Removal Time

Recovery time is the minimum length of time for the deassertion of an asynchronous control signal relative to the next clock edge.

For example, signals such as `clear` and `preset` must be stable before the next active clock edge. The recovery slack calculation is similar to the clock setup slack calculation, but it applies to asynchronous control signals.

**Figure 6-15: Recovery Slack Calculation if the Asynchronous Control Signal is Registered**

$$\begin{aligned} \text{Recovery Slack Time} &= \text{Data Required Time} - \text{Data Arrival Time} \\ \text{Data Required Time} &= \text{Latch Edge} + \text{Clock Network Delay to Destination Register} - \mu t_{\text{su}} \\ \text{Data Arrival Time} &= \text{Launch Edge} + \text{Clock Network Delay to Source Register} + \mu t_{\text{co}} + \text{Register-to-Register Delay} \end{aligned}$$

**Figure 6-16: Recovery Slack Calculation if the Asynchronous Control Signal is not Registered**

$$\begin{aligned} \text{Recovery Slack Time} &= \text{Data Required Time} - \text{Data Arrival Time} \\ \text{Data Required Time} &= \text{Latch Edge} + \text{Clock Network Delay to Destination Register} - \mu t_{\text{su}} \\ \text{Data Arrival Time} &= \text{Launch Edge} + \text{Clock Network Delay} + \text{Input Maximum Delay} + \text{Port-to-Register Delay} \end{aligned}$$

**Note:** If the asynchronous reset signal is from a device I/O port, you must create an input delay constraint for the asynchronous reset port for the TimeQuest analyzer to perform recovery analysis on the path.

Removal time is the minimum length of time the deassertion of an asynchronous control signal must be stable after the active clock edge. The TimeQuest analyzer removal slack calculation is similar to the clock hold slack calculation, but it applies asynchronous control signals.

**Figure 6-17: Removal Slack Calculation if the Asynchronous Control Signal is Registered**

Removal Slack Time	= Data Arrival Time – Data Required Time
Data Arrival Time	= Launch Edge + Clock Network Delay to Source Register + $\mu t_{co}$ of Source Register + Register-to-Register Delay
Data Required Time	= Latch Edge + Clock Network Delay to Destination Register + $\mu t_h$

**Figure 6-18: Removal Slack Calculation if the Asynchronous Control Signal is not Registered**


Removal Slack Time	= Data Arrival Time – Data Required Time
Data Arrival Time	= Launch Edge + Clock Network Delay + Input Minimum Delay of Pin + Minimum Pin-to-Register Delay
Data Required Time	= Latch Edge + Clock Network Delay to Destination Register + $\mu t_h$

If the asynchronous reset signal is from a device pin, you must assign the **Input Minimum Delay** timing assignment to the asynchronous reset pin for the TimeQuest analyzer to perform removal analysis on the path.

## Multicycle Paths


Multicycle paths are data paths that require a non-default setup and/or hold relationship for proper analysis.

For example, a register may be required to capture data on every second or third rising clock edge. An example of a multicycle path between the input registers of a multiplier and an output register where the destination latches data on every other clock edge.

**Figure 6-19: Multicycle Path**

A register-to-register path used for the default setup and hold relationship, the respective timing diagrams for the source and destination clocks, and the default setup and hold relationships, when the source clock, `src_clk`, has a period of 10 ns and the destination clock, `dst_clk`, has a period of 5 ns. The default setup relationship is 5 ns; the default hold relationship is 0 ns.


Figure 6-20: Register-to-Register Path and Default Setup and Hold Timing Diagram



To accommodate the system requirements you can modify the default setup and hold relationships with a multicycle timing exception.

The actual setup relationship after you apply a multicycle timing exception. The exception has a multicycle setup assignment of two to use the second occurring latch edge; in this example, to 10 ns from the default value of 5 ns.

Figure 6-21: Modified Setup Diagram



**Related Information**

- [The Quartus II TimeQuest Timing Analyzer](#) on page 7-1  
For more information about creating exceptions with multicycle paths.

**Metastability**

Metastability problems can occur when a signal is transferred between circuitry in unrelated or asynchronous clock domains because the designer cannot guarantee that the signal will meet setup and hold time requirements.

To minimize the failures due to metastability, circuit designers typically use a sequence of registers, also known as a synchronization register chain, or synchronizer, in the destination clock domain to resynchronize the data signals to the new clock domain.

The mean time between failures (MTBF) is an estimate of the average time between instances of failure due to metastability.

The TimeQuest analyzer analyzes the potential for metastability in your design and can calculate the MTBF for synchronization register chains. The MTBF of the entire design is then estimated based on the synchronization chains it contains.

In addition to reporting synchronization register chains found in the design, the Quartus II software also protects these registers from optimizations that might negatively impact MTBF, such as register duplication and logic retiming. The Quartus II software can also optimize the MTBF of your design if the MTBF is too low.

#### Related Information


- **Understanding Metastability in FPGAs**  
For more information about metastability, its effects in FPGAs, and how MTBF is calculated.
- **Managing Metastability with the Quartus II Software**  
For more information about metastability analysis, reporting, and optimization features in the Quartus II software.

## Common Clock Path Pessimism Removal

Common clock path pessimism removal accounts for the minimum and maximum delay variation associated with common clock paths during static timing analysis by adding the difference between the maximum and minimum delay value of the common clock path to the appropriate slack equation.

Minimum and maximum delay variation can occur when two different delay values are used for the same clock path. For example, in a simple setup analysis, the maximum clock path delay to the source register is used to determine the data arrival time. The minimum clock path delay to the destination register is used to determine the data required time. However, if the clock path to the source register and to the destination register share a common clock path, both the maximum delay and the minimum delay are used to model the common clock path during timing analysis. The use of both the minimum delay and maximum delay results in an overly pessimistic analysis since two different delay values, the maximum and minimum delays, cannot be used to model the same clock path.

**Figure 6-22: Typical Register to Register Path**




Segment A is the common clock path between `reg1` and `reg2`. The minimum delay is 5.0 ns; the maximum delay is 5.5 ns. The difference between the maximum and minimum delay value equals the common clock path pessimism removal value; in this case, the common clock path pessimism is 0.5 ns. The TimeQuest analyzer adds the common clock path pessimism removal value to the appropriate slack equation to determine overall slack. Therefore, if the setup slack for the register-to-register path in the example equals 0.7 ns without common clock path pessimism removal, the slack would be 1.2 ns with common clock path pessimism removal.

You can also use common clock path pessimism removal to determine the minimum pulse width of a register. A clock signal must meet a register's minimum pulse width requirement to be recognized by the register. A minimum high time defines the minimum pulse width for a positive-edge triggered register. A minimum low time defines the minimum pulse width for a negative-edge triggered register.



Clock pulses that violate the minimum pulse width of a register prevent data from being latched at the data pin of the register. To calculate the slack of the minimum pulse width, the TimeQuest analyzer subtracts the required minimum pulse width time from the actual minimum pulse width time. The TimeQuest analyzer determines the actual minimum pulse width time by the clock requirement you specified for the clock that feeds the clock port of the register. The TimeQuest analyzer determines the required minimum pulse width time by the maximum rise, minimum rise, maximum fall, and minimum fall times.

**Figure 6-23: Required Minimum Pulse Width time for the High and Low Pulse**



With common clock path pessimism, the minimum pulse width slack can be increased by the smallest value of either the maximum rise time minus the minimum rise time, or the maximum fall time minus the minimum fall time. In the example, the slack value can be increased by 0.2 ns, which is the smallest value between 0.3 ns (0.8 ns – 0.5 ns) and 0.2 ns (0.9 ns – 0.7 ns).

**Related Information**

[TimeQuest Timing Analyzer Page \(Settings Dialog Box\)](#)

For more information, refer to the Quartus II Help.

## Clock-As-Data Analysis


The majority of FPGA designs contain simple connections between any two nodes known as either a data path or a clock path.

A data path is a connection between the output of a synchronous element to the input of another synchronous element.

A clock is a connection to the clock pin of a synchronous element. However, for more complex FPGA designs, such as designs that use source-synchronous interfaces, this simplified view is no longer sufficient. Clock-as-data analysis is performed in circuits with elements such as clock dividers and DDR source-synchronous outputs.

The connection between the input clock port and output clock port can be treated either as a clock path or a data path. A design where the path from port `clk_in` to port `clk_out` is both a clock and a data path. The clock path is from the port `clk_in` to the register `reg_data` clock pin. The data path is from port `clk_in` to the port `clk_out`.


Figure 6-24: Simplified Source Synchronous Output



With clock-as-data analysis, the TimeQuest analyzer provides a more accurate analysis of the path based on user constraints. For the clock path analysis, any phase shift associated with the phase-locked loop (PLL) is taken into consideration. For the data path analysis, any phase shift associated with the PLL is taken into consideration rather than ignored.

The clock-as-data analysis also applies to internally generated clock dividers. An internally generated clock divider. In this figure, waveforms are for the inverter feedback path, analyzed during timing analysis. The output of the divider register is used to determine the launch time and the clock port of the register is used to determine the latch time.

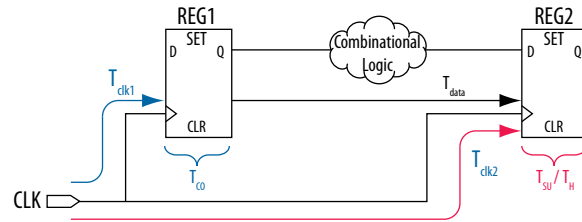
Figure 6-25: Clock Divider



## Multicycle Clock Setup Check and Hold Check Analysis

You can modify the setup and hold relationship when you apply a multicycle exception to a register-to-register path.

Figure 6-26: Register-to-Register Path




## Multicycle Clock Setup

The setup relationship is defined as the number of clock periods between the latch edge and the launch edge. By default, the TimeQuest analyzer performs a single-cycle path analysis, which results in the setup relationship being equal to one clock period (latch edge – launch edge). Applying a multicycle setup assignment, adjusts the setup relationship by the multicycle setup value. The adjustment value may be negative.


An end multicycle setup assignment modifies the latch edge of the destination clock by moving the latch edge the specified number of clock periods to the right of the determined default latch edge. The following figure shows various values of the end multicycle setup (EMS) assignment and the resulting latch edge.

Figure 6-27: End Multicycle Setup Values



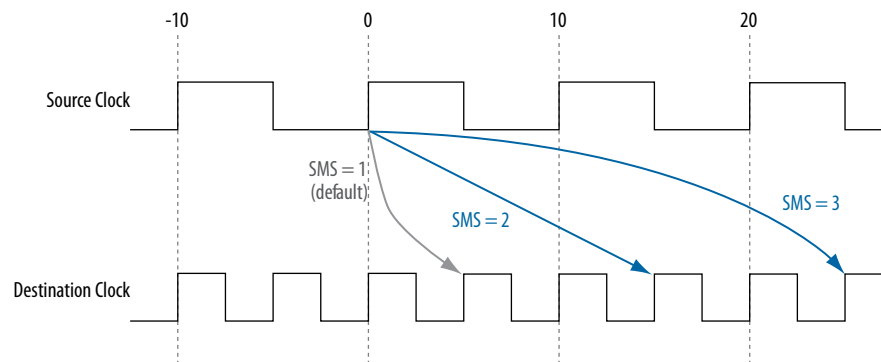
A start multicycle setup assignment modifies the launch edge of the source clock by moving the launch edge the specified number of clock periods to the left of the determined default launch edge. A start multicycle setup (SMS) assignment with various values can result in a specific launch edge.

Figure 6-28: Start Multicycle Setup Values



The setup relationship reported by the TimeQuest analyzer for the negative setup relationship.

Figure 6-29: Start Multicycle Setup Values Reported by the TimeQuest Analyzer



## Multicycle Clock Hold

The setup relationship is defined as the number of clock periods between the launch edge and the latch edge.

By default, the TimeQuest analyzer performs a single-cycle path analysis, which results in the hold relationship being equal to one clock period (launch edge – latch edge). When analyzing a path, the TimeQuest analyzer performs two hold checks. The first hold check determines that the data launched by the current launch edge is not captured by the previous latch edge. The second hold check determines that the data launched by the next launch edge is not captured by the current latch edge. The TimeQuest analyzer reports only the most restrictive hold check. The TimeQuest analyzer calculates the hold check by comparing launch and latch edges.

The calculation the TimeQuest analyzer performs to determine the hold check.

Figure 6-30: Hold Check


hold check 1 = current launch edge – previous latch edge

hold check 2 = next launch edge – current latch edge

**Tip:** If a hold check overlaps a setup check, the hold check is ignored.


A start multicycle hold assignment modifies the launch edge of the destination clock by moving the latch edge the specified number of clock periods to the right of the determined default launch edge. The following figure shows various values of the start multicycle hold (SMH) assignment and the resulting launch edge.

**Figure 6-31: Start Multicycle Hold Values**




An end multicycle hold assignment modifies the latch edge of the destination clock by moving the latch edge the specific ed number of clock periods to the left of the determined default latch edge. The following figure shows various values of the end multicycle hold (EMH) assignment and the resulting latch edge.

**Figure 6-32: End Multicycle Hold Values**



The hold relationship reported by the TimeQuest analyzer for the negative hold relationship shown in the figure above would look like this:

**Figure 6-33: End Multicycle Hold Values Reported by the TimeQuest Analyzer**



## Multicorner Analysis

The TimeQuest analyzer performs multicorner timing analysis to verify your design under a variety of operating conditions—such as voltage, process, and temperature—while performing static timing analysis.

To change the operating conditions or speed grade of the device used for timing analysis, use the `set_operating_conditions` command.

If you specify an operating condition Tcl object, the `-model`, `speed`, `-temperature`, and `-voltage` options are optional. If you do not specify an operating condition Tcl object, the `-model` option is required; the `-speed`, `-temperature`, and `-voltage` options are optional.

**Tip:** To obtain a list of available operating conditions for the target device, use the `get_available_operating_conditions -all` command.

To ensure that no violations occur under various conditions during the device operation, perform static timing analysis under all available operating conditions.

**Table 6-2: Operating Conditions for Slow and Fast Models**

Model	Speed Grade	Voltage	Temperature
Slow	Slowest speed grade in device density	$V_{cc}$ minimum supply <sup>(1)</sup>	Maximum $T_J$ <sup>(1)</sup>
Fast	Fastest speed grade in device density	$V_{cc}$ maximum supply <sup>(1)</sup>	Minimum $T_J$ <sup>(1)</sup>

Note :

1. Refer to the DC & Switching Characteristics chapter of the applicable device Handbook for  $V_{cc}$  and  $T_J$  values

In your design, you can set the operating conditions for to the slow timing model, with a voltage of 1100 mV, and temperature of 85° C with the following code:

```
set_operating_conditions -model slow -temperature 85 -voltage 1100
```

You can set the same operating conditions with a Tcl object:

```
set_operating_conditions 3_slow_1100mv_85c
```

The following block of code shows how to use the `set_operating_conditions` command to generate different reports for various operating conditions.

### Example 6-1: Script Excerpt for Analysis of Various Operating Conditions

```
#Specify initial operating conditions
set_operating_conditions -model slow -speed 3 -grade c -temperature 85 -
voltage 1100
#Update the timing netlist with the initial conditions
update_timing_netlist
#Perform reporting
#Change initial operating conditions. Use a temperature of 0C
set_operating_conditions -model slow -speed 3 -grade c -temperature 0 -
voltage 1100
#Update the timing netlist with the new operating condition
update_timing_netlist
#Perform reporting
#Change initial operating conditions. Use a temperature of 0C and a model of
fast
set_operating_conditions -model fast -speed 3 -grade c -temperature 0 -
voltage 1100
#Update the timing netlist with the new operating condition
update_timing_netlist
#Perform reporting
```

#### Related Information

[set\\_operating\\_conditions](#)

[get\\_available\\_operating\\_conditions](#)

For more information about the `get_available_operating_conditions` command

## Document Revision History

Table 6-3: Document Revision History

Date	Version	Changes
2014.12.15	14.1.0	Moved Multicycle Clock Setup Check and Hold Check Analysis section from the TimeQuest Timing Analyzer chapter.
June 2014	14.0.0	Updated format
June 2012	12.0.0	Added social networking icons, minor text updates
November 2011	11.1.0	Initial release.

#### Related Information

[Quartus II Handbook Archive](#)

For previous versions of the *Quartus II Handbook*.

2014.12.15

QI15V3



Subscribe



Send Feedback

The Quartus<sup>®</sup> II TimeQuest Timing Analyzer is a powerful ASIC-style timing analysis tool that validates the timing performance of all logic in your design using an industry-standard constraint, analysis, and reporting methodology. Use the TimeQuest analyzer GUI or command-line interface to constrain, analyze, and report results for all timing paths in your design.

This document is organized to allow you to refer to specific subjects relating to the TimeQuest analyzer and timing analysis. The sections cover the following topics:

**Recommended Flow** on page 7-2

The Quartus II TimeQuest analyzer performs constraint validation to timing verification as part of the compilation flow.

**Timing Constraints** on page 7-6

Timing analysis in the Quartus II software with the TimeQuest Timing Analyzer relies on constraining your design to make it meet your timing requirements.

**Running the TimeQuest Analyzer** on page 7-50

When you compile a design, the TimeQuest timing analyzer automatically performs multi-corner signoff timing analysis after the Fitter has finished.

**Understanding Results** on page 7-53

Knowing how your constraints are displayed when analyzing a path is one of the most important skills of timing analysis.

**Constraining and Analyzing with Tcl Commands** on page 7-60

You can use Tcl commands from the Quartus II software Tcl Application Programming Interface (API) to constrain, analyze, and collect information for your design.

**Generating Timing Reports** on page 7-64

The TimeQuest analyzer provides real-time static timing analysis result reports.

**Document Revision History** on page 7-66

## Related Information

- **Timing Analysis Overview** on page 6-1  
For more information about basic timing analysis concepts and how they pertain to the TimeQuest analyzer.
- **TimeQuest Timing Analyzer Resource Center**  
For more information about Altera resources available for the TimeQuest analyzer.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



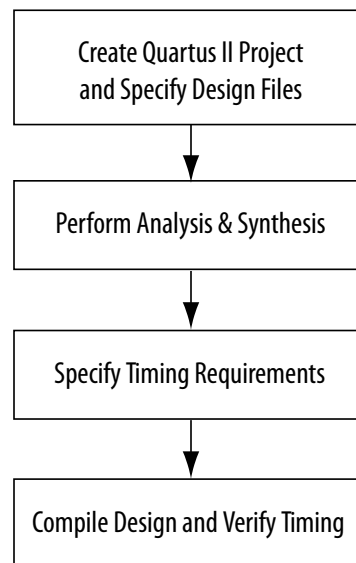


- [Altera Training](#)  
For more information about the TimeQuest analyzer.

## Recommended Flow for First Time Users

The Quartus II TimeQuest analyzer performs constraint validation to timing verification as part of the compilation flow. Both the TimeQuest analyzer and the Fitter use of constraints contained in a Synopsis Design Constraints (**.sdc**) file. The following flow is recommended if you have not created a project and do not have a SDC file with timing constraints for your design.

Figure 7-1: Design Flow with the TimeQuest Timing Analyzer



## Creating and Setting Up your Design

You must first create your project in the Quartus II software. Include all the necessary design files, including any existing Synopsis Design Constraints (**.sdc**) files, also referred to as SDC files, that contain timing constraints for your design. Some reference designs, or Altera or partner IP cores may already include one or more SDC files.

All SDC files must be added to your project so that your constraints are processed when the Quartus II software performs Fitting and Timing Analysis. Typically you must create an SDC file to constrain your design.

### Related Information

- [SDC File Precedence](#) on page 7-52
- [Managing Files in a Project](#)  
For more information on project file management, refer to Quartus II Help.

## Specifying Timing Requirements

Before running timing analysis with the TimeQuest analyzer, you must specify timing constraints, describe the clock frequency requirements and other characteristics, timing exceptions, and I/O timing

requirements of your design. You can use the TimeQuest Timing Analyzer Wizard to enter initial constraints for your design, and then refine timing constraints with the TimeQuest analyzer GUI.

Both the TimeQuest analyzer and the Fitter use of constraints contained in a Synopsis Design Constraints (.sdc) file.

The constraints in the SDC file are read in sequence. You must first make a constraint before making any references to that constraint. For example, if a generated clock references a base clock, the base clock constraint must be made before the generated clock constraint.

If you are new to timing analysis with the TimeQuest analyzer, you can use template files included with the Quartus II software and the interactive dialog boxes to create your initial SDC file. To use this method, refer to Performing an Initial Analysis and Synthesis.

If you are familiar with timing analysis, you can also create an SDC file in your preferred text editor. Don't forget to include the SDC file in the project when you are finished.

#### Related Information

- [Creating a Constraint File from Quartus II Templates with the Quartus II Text Editor](#) on page 7-4  
For more information on using the Quartus II Text Editor templates for SDC constraints.
- [Identifying the Quartus II Software Executable from the SDC File](#) on page 7-64
- [Specifying Timing Constraints and Exceptions](#)  
For more information, refer to Quartus II Help.

## Performing an Initial Analysis and Synthesis

Perform Analysis and Synthesis on your design so that you can find design entry names in the **Node Finder** to simplify creating constraints.

The Quartus II software populates an internal database with design element names. You must synthesize your design in order for the Quartus II software to assign names to your design elements, for example, pins, nodes, hierarchies, and timing paths.

If you have already compiled your design, you do not need need to perform the synthesis step again, because compiling the design automatically performs synthesis. You can either perform Analysis and Synthesis to create a post-map database, or perform a full compilation to create a post-fit database. Creating a post-map database is faster than a post-fit database, and is sufficient for creating initial timing constraints.

**Note:** If you are using incremental compilation, you must merge your design partitions after performing Analysis and Synthesis to create a post-map database.

**Note:** When compiling for the Arria® 10 device family, the following commands are required to perform initial synthesis and enable you to use the **Node Finder** to find names in your design:

```
quartus_map <design>
quartus_fit <design> --floorplan
quartus_sta <design> --post_map
```

When compiling for other devices, you can exclude the `quartus_fit <design> --floorplan` step:

```
quartus_map <design>
quartus_sta <design> --post_map
```

**Related Information**

- [Using the Node Finder](#)
- [Setting up and Running Analysis and Synthesis](#)
- [Setting up and Running a Compilation](#)

For more information, refer to Quartus II Help.

**Creating a Constraint File from Quartus II Templates with the Quartus II Text Editor**

You can create an SDC file from constraint templates in the Quartus II software with the Quartus II Text Editor, or with your preferred text editor.

1. On the **File** menu, click **New**.
2. In the **New** dialog box, select the **Synopsys Design Constraints File** type from the **Other Files** group. Click **OK**.
3. Right-click in the blank SDC file in the Quartus II Text Editor, then click **Insert Constraint**. Choose **Clock Constraint** followed by **Set Clock Groups** since they are the most widely used constraints. The Quartus II Text Editor displays a dialog box with interactive fields for creating constraints. For example, the **Create Clock** dialog box shows you the waveform for your `create_clock` constraint while you adjust the **Period** and **Rising** and **Falling** waveform edge settings. The actual constraint is displayed in the **SDC command** field. Click **Insert** to use the constraint in your SDC.

*or*

4. Click the **Insert Template** button on the text editor menu, or, right-click in the blank SDC file in the Quartus II Text Editor, then click **Insert TemplateTimeQuest**.
  - a. In the **Insert Template** dialog box, expand the **TimeQuest** section, then expand the **SDC Commands** section.
  - b. Expand a command category, for example, **Clocks**.
  - c. Select a command. The SDC constraint appears in the **Preview** pane.
  - d. Click **Insert** to paste the SDC constraint into the blank SDC file you created in step 2. This creates a generic constraint for you to edit manually, replacing variables such as clock names, period, rising and falling edges, ports, etc.
5. Repeat as needed with other constraints, or click **Close** to close the **Insert Template** dialog box.

You can now use any of the standard features of the Quartus II Text Editor to modify the SDC file, or save the SDC file to edit in a text editor. Your SDC can be saved with the same name as the project, and generally should be stored in the project directory.

**Related Information**

- [About the Quartus II Text Editor](#)  
For more information on inserting a template with the Quartus II Text Editor, refer to Quartus II Help.
- [Specifying Timing Constraints and Exceptions \(TimeQuest Timing Analyzer\)](#)  
For more information on using the TimeQuest analyzer GUI to modify timing constraints.
- [Create Clocks Dialog Box](#)
- [Set Clock Groups Dialog Box](#)  
For more information on Create Clocks and Set Clock Groups, refer to the Quartus II Help.

## Performing a Full Compilation

After creating initial timing constraints, compile your design.

During a full compilation, the Fitter uses the TimeQuest analyzer repeatedly to perform timing analysis with your timing constraints. By default, the Fitter can stop early if it meets your timing requirements, instead of attempting to achieve the maximum performance. You can modify this by changing the Fitter effort settings in the Quartus II software.

### Related Information

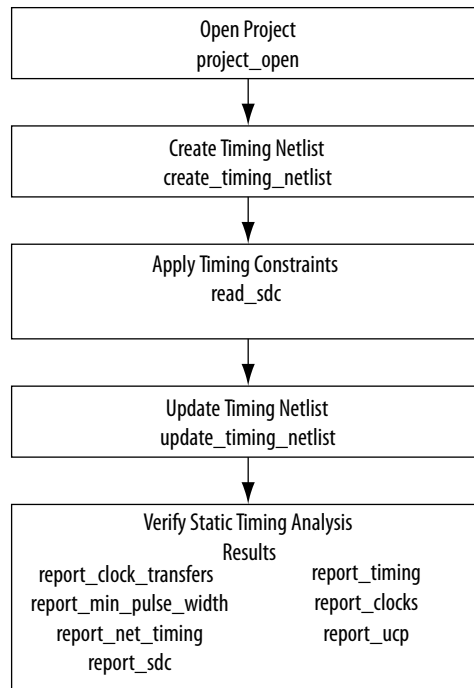
- [Analyzing Timing in Designs Compiled in Previous Versions](#) on page 7-6  
For more information about importing databases compiled in previous versions of the software.
- [Fitter Settings Page \(Settings Dialog Box\)](#)  
For more information about changing Fitter effort, refer to the Quartus II Help.

## Verifying Timing

The TimeQuest analyzer examines the timing paths in the design, calculates the propagation delay along each path, checks for timing constraint violations, and reports timing results as positive slack or negative slack. Negative slack indicates a timing violation. If you encounter violations along timing paths, use the timing reports to analyze your design and determine how best to optimize your design. If you modify, remove, or add constraints, you should perform a full compilation again. This iterative process helps resolve timing violations in your design.

There is a recommended flow for constraining and analyzing your design within the TimeQuest analyzer, and each part of the flow has a corresponding Tcl command.

Figure 7-2: The TimeQuest Timing Analyzer Flow



**Related Information****[Viewing Timing Analysis Results](#)**

For more information, refer to Quartus II Help.

## Analyzing Timing in Designs Compiled in Previous Versions

Performing a full compilation can be a lengthy process, however, once your design meets timing you can export the design database for later use. This can include operations such as verification of subsequent timing models, or use in a later version of the Quartus II software.

When you re-open the project, the Quartus II software opens the exported database from the export directory. You can then run TimeQuest on the design without having to recompile the project.

To export the database in the previous version of the Quartus II software, click **Project > Export Database** and select the export directory to contain the exported database.

To import a database in a later version of the Quartus II software, click **File > Open** and select the Quartus II Project file (**.qpf**) for the project.

Once you import the database, you can perform any TimeQuest analyzer functions on the design without recompiling.

## Timing Constraints

Timing analysis in the Quartus II software with the TimeQuest Timing Analyzer relies on constraining your design to make it meet your timing requirements. When discussing these constraints, they can be referred to as timing constraints, SDC constraints, or SDC commands interchangeably.

## Recommended Starting SDC Constraints

Almost every beginning SDC file should contain the following four commands:

[create\\_clock](#) on page 7-6

[derive\\_pll\\_clocks](#) on page 7-7

[derive\\_clock\\_uncertainty](#) on page 7-8

[SDC Constraint Creation Summary](#) on page 7-8

Those are the first three steps, which can usually be done very quickly. For a sample design with two clocks coming into it, your SDC file might look like this example:

[set\\_clock\\_groups](#) on page 7-8

**Related Information**

[Creating a Constraint File from Quartus II Templates with the Quartus II Text Editor](#) on page 7-4

### create\_clock

The first statements in a SDC file should be constraints for clocks, for example, constrain the external clocks coming into the FPGA with `create_clock`. An example of the basic syntax is:

```
create_clock -name sys_clk -period 8.0 [get_ports fpga_clk]
```

This command creates a clock called `sys_clk` with an 8ns period and applies it to the port called `fpga_clk`.

**Note:** Both Tcl files and SDC files are case-sensitive, so make sure references to pins, ports, or nodes, such as `fpga_clk` match the case used in your design.

By default, the clock has a rising edge at time 0ns, and a 50% duty cycle, hence a falling edge at time 4ns. If you require a different duty cycle or to represent an offset, use the `-waveform` option, however, this is seldom necessary.

It is common to create a clock with the same name as the port it is applied to. In the example above, this would be accomplished by:

```
create_clock -name fpga_clk -period 8.0 [get_ports fpga_clk]
```

There are now two unique objects called `fpga_clk`, a port in your design and a clock applied to that port.

**Note:** In Tcl syntax, square brackets execute the command inside them, so `[get_ports fpga_clk]` executes a command that finds all ports in the design that match `fpga_clk` and returns a collection of them. You can enter the command without using the `get_ports` collection command, as shown in the following example. There are benefits to using collection commands, which are described in "Collection Commands".

```
create_clock -name sys_clk -period 8.0 fpga_clk
```

Repeat this process, using one `create_clock` command for each known clock coming into your design. Later on you can use **Report Unconstrained Paths** to identify any unconstrained clocks.

**Note:** Rather than typing constraints, users can enter constraints through the GUI. After launching TimeQuest, open the SDC file from TimeQuest or Quartus II, place the cursor where the new constraint will go, and go to **Edit > Insert Constraint**, and choose the constraint.

**Warning:** Using the **Constraints** menu option in the TimeQuest GUI applies constraints directly to the timing database, but makes no entry in the SDC file. An advanced user may find reasons to do this, but if you are new to TimeQuest, Altera recommends entering your constraints directly into your SDC with the **Edit > Insert Constraint** command.

#### Related Information

[Creating Base Clocks](#) on page 7-11

## derive\_pll\_clocks

After the `create_clock` commands add the following command into your SDC file:

```
derive_pll_clocks
```

This command automatically creates a generated clock constraint on each output of the PLLs in your design..

When PLLs are created, you define how each PLL output is configured. Because of this, the TimeQuest analyzer can automatically constrain them, with the `derive_pll_clocks` command.

This command also creates other constraints. It constrains transceiver clocks. It adds multicycles between LVDS SERDES and user logic.

The **derive\_pll\_clocks** command prints an Info message to show each generated clock it creates.

If you are new to the TimeQuest analyzer, you may decide not to use `derive_pll_clocks`, and instead cut-and-paste each `create_generated_clock` assignment into the SDC file. There is nothing wrong with this, since the two are identical. The problem is that when you modify a PLL setting, you must remember to change its generated clock in the SDC file. Examples of this type of change include modifying an existing output clock, adding a new PLL output, or making a change to the PLL's hierarchy. Too many

designers forget to modify the SDC file and spend time debugging something that `derive_pll_clocks` would have updated automatically.

#### Related Information

- [Creating Base Clocks](#) on page 7-11
- [Deriving PLL Clocks](#) on page 7-17

## derive\_clock\_uncertainty

Add the following command to your SDC file:

```
derive_clock_uncertainty
```

This command calculates clock-to-clock uncertainties within the FPGA, due to characteristics like PLL jitter, clock tree jitter, etc. This should be in all SDC files and the TimeQuest analyzer generates a warning if this command is not found in your SDC files.

#### Related Information

[Accounting for Clock Effect Characteristics](#) on page 7-21

## SDC Constraint Creation Summary

Those are the first three steps, which can usually be done very quickly. For a sample design with two clocks coming into it, your SDC file might look like this example:

```
create_clock -period 20.00 -name adc_clk [get_ports adc_clk]
create_clock -period 8.00 -name sys_clk [get_ports sys_clk]

derive_pll_clocks

derive_clock_uncertainty
```

## set\_clock\_groups

With the constraints discussed previously, most, if not all, of the clocks in the design are now constrained. In the TimeQuest analyzer, all clocks are related by default, and you must indicate which clocks are not related. For example, if there are paths between an 8ns clock and 10ns clock, even if the clocks are completely asynchronous, the TimeQuest analyzer attempts to meet a 2ns setup relationship between these clocks unless you indicate that they are not related. The TimeQuest analyzer analyzes everything known, rather than assuming that all clocks are unrelated and requiring that you relate them. The SDC language has a powerful constraint for setting unrelated clocks called `set_clock_groups`. A template for the typical use of the `set_clock_groups` command is:

```
set_clock_groups -asynchronous -group {<clock1>...<clockn>} ... -group
{<clocka>...<clockn>}
```

The `set_clock_groups` command does not actually group clocks. Since the TimeQuest analyzer assumes all clocks are related by default, all clocks are effectively in one big group. Instead, the `set_clock_groups` command cuts timing between clocks in different groups.

There is no limit to the number of times you can specify a group option with `-group {<group of clocks>}`. When entering constraints through the GUI with **Edit > Insert Constraint**, the **Set Clock Groups** dialog box only permits two clock groups, but this is only a limitation of that dialog box. You can always manually add more into the SDC file.

Any clock not listed in the assignment is related to all clocks. If you forget a clock, the TimeQuest analyzer acts conservatively and analyzes that clock in context with all other domains to which it connects.

The `set_clock_groups` command requires either the `-asynchronous` or `-exclusive` option. The `-asynchronous` flag means the clocks are both toggling, but not in a way that can synchronously pass data. The `-exclusive` flag means the clocks do not toggle at the same time, and hence are mutually exclusive. An example of this might be a clock multiplexer that has two generated clock assignments on its output. Since only one can toggle at a time, these clocks are `-exclusive`. TimeQuest does not currently analyze crosstalk explicitly. Instead, the timing models use extra guard bands to account for any potential crosstalk-induced delays. TimeQuest treats the `-asynchronous` and `-exclusive` options the same.

A clock cannot be within multiple `-group` groupings in a single assignment, however, you can have multiple `set_clock_groups` assignments.

Another way to cut timing between clocks is to use `set_false_path`. To cut timing between `sys_clk` and `dsp_clk`, a user might enter:

```
set_false_path -from [get_clocks sys_clk] -to [get_clocks dsp_clk]

set_false_path -from [get_clocks dsp_clk] -to [sys_clk]
```

This works fine when there are only a few clocks, but quickly grows to a huge number of assignments that are completely unreadable. In a simple design with three PLLs that have multiple outputs, the `set_clock_groups` command can clearly show which clocks are related in less than ten lines, while `set_false_path` may be over 50 lines and be very non-intuitive on what is being cut.

#### Related Information

- [Creating Generated Clocks](#) on page 7-15
- [Relaxing Setup with `set\_multicycle\_path`](#) on page 7-27
- [Accounting for a Phase Shift](#) on page 7-28

#### Tips for Writing a `set_clock_groups` Constraint

Since `derive_pll_clocks` creates many of the clock names, you may not know all of the clock names to use in the clock groups.

A quick way to make this constraint is to use the SDC file you have created so far, with the three basic constraints described in previous topics. Make sure you have added it to your project, then open the TimeQuest timing analyzer GUI.

In the **Task** panel of the **TimeQuest** analyzer, double-click on **Report Clocks**. This reads your existing SDC and applies it to your design, then reports all the clocks. From that report, highlight all of the clock names in the first column, and copy the names.

You have just copied all the clock names in your design in the exact format the TimeQuest analyzer recognizes. Paste them into your SDC file to make a list of all clock names, one per line..

Format that list into the `set_clock_groups` command by cutting and pasting clock names into appropriate groups. Then enter the following empty template in your SDC file::

```
set_clock_groups -asynchronous -group { \
} \
-group { \
} \
-group { \
} \
-group { \
}
```

Cut and paste clocks into groups to define how they're related, adding or removing groups as necessary. Format to make the code readable.



**Note:** This command can be difficult to read on a single line. Instead, you should make use of the Tcl line continuation character "\". By putting a space after your last character and then "\", the end-of-line character is escaped. (And be careful not to have any whitespace after the escape character, or else it will escape the whitespace, not the end-of-line character).

```
set_clock_groups -asynchronous \
  -group {adc_clk \
    the_adc_pll|altpll_component_autogenerated|pll|clk[0] \
    the_adc_pll|altpll_component_autogenerated|pll|clk[1] \
    the_adc_pll|altpll_component_autogenerated|pll|clk[2] \
  } \
  -group {sys_clk \
    the_system_pll|altpll_component_autogenerated|pll|clk[0] \
    the_system_pll|altpll_component_autogenerated|pll|clk[1] \
  } \
  -group {the_system_pll|altpll_component_autogenerated|pll|clk[2] \
} \
```

**Note:** The last group has a PLL output `system_pll|..|clk[2]` while the input clock and other PLL outputs are in different groups. If PLLs are used, and the input clock frequency is not related to the frequency of the PLL's outputs, they must be treated asynchronously. Usually most outputs of a PLL are related and hence in the same group, but this is not a requirement, and depends on the requirements of your design.

For designs with complex clocking, writing this constraint can be an iterative process. For example, a design with two DDR3 cores and high-speed transceivers could easily have thirty or more clocks. In those cases, you can just add the clocks you've created. Since clocks not in the command are still related to every clock, you are conservatively grouping what is known. If there are still failing paths in the design between unrelated clock domains, you can start adding in the new clock domains as necessary. In this case, a large number of the clocks won't actually be in the `set_clock_groups` command, since they are either cut in the SDC file for the IP core (such as the SDC files generated by the DDR3 cores), or they only connect to clock domains to which they are related.

For many designs, that is all that's necessary to constrain the core. Some common core constraints that will not be covered in this quick start section that user's do are:

- Add multicycles between registers which can be analyzed at a slower rate than the default analysis, in other words, increasing the time when data can be read, or 'opening the window'. For example, a 10ns clock period will have a 10ns setup relationship. If the data changes at a slower rate, or perhaps the registers toggle at a slower rate due to a clock enable, then you should apply a multicycle that relaxes the setup relationship (opens the the window so that valid data can pass). This is a multiple of the clock period, making the setup relationship 20ns, 40ns, etc., while keeping the hold relationship at 0ns. These types of multicycles are generally applied to paths.
- The second common form of multicycle is when the user wants to advance the cycle in which data is read, or 'shift the window'. This generally occurs when your design performs a small phase-shift on a clock. For example, if your design has two 10ns clocks exiting a PLL, but the second clock has a 0.5ns phase-shift, the default setup relationship from the main clock to the phase-shifted clock is 0.5ns and the hold relationship is -9.5ns. It is almost impossible to meet a 0.5ns setup relationship, and most likely you intend the data to transfer in the next window. By adding a multicycle from the main clock to the phase-shifted clock, the setup relationship becomes 10.5ns and the hold relationship becomes 0.5ns. This multicycle is generally applied between clocks and is something the user should think about as soon as they do a small phase-shift on a clock. This type of multicycle is called shifting the window.
- Add a `create_generated_clock` to ripple clocks. When a register's output drives the `clk` port of another register, that is a ripple clock. Clocks do not propagate through registers, so all ripple clocks must have a `create_generated_clock` constraint applied to them for correct analysis. Unconstrained ripple clocks appear in the **Report Unconstrained Paths** report, so they are easily recognized. In general, ripple clocks should be avoided for many reasons, and if possible, a clock enable should be used instead.
- Add a `create_generated_clock` to clock mux outputs. Without this, all clocks propagate through the mux and are related. TimeQuest analyze paths downstream from the mux where one clock input feeds the source register and the other clock input feeds the destination, and vice-versa. Although it could be valid, this is usually not preferred behavior. By putting `create_generated_clock` constraints on the mux output, which relates them to the clocks coming into the mux, you can correctly group these clocks with other clocks.

## Creating Clocks and Clock Constraints

Clocks specify timing requirements for synchronous transfers and guide the Fitter optimization algorithms to achieve the best possible placement for your design. You must define all clocks and any associated clock characteristics, such as uncertainty or latency. The TimeQuest analyzer supports SDC commands that accommodate various clocking schemes such as:

- Base clocks
- Virtual clocks
- Multifrequency clocks
- Generated clocks

### Creating Base Clocks

Base clocks are the primary input clocks to the device. Unlike clocks that are generated in the device (such as an on-chip PLL), base clocks are generated by off-chip oscillators or forwarded from an external device. Define base clocks at the top of your SDC file, because generated clocks and other constraints often reference base clocks. The TimeQuest timing analyzer ignores any constraints that reference a clock that has not been defined.

Use the `create_clock` command to create a base clock. Use other constraints, such as those described in *Accounting for Clock Effect Characteristics*, to specify clock characteristics such as uncertainty and latency.

The following examples show the most common uses of the `create_clock` constraint:

### create\_clock Command

To specify a 100 MHz requirement on a `clk_sys` input clock port you would enter the following in your SDC file:

```
create_clock -period 10 -name clk_sys [get_ports clk_sys]
```

### 100 MHz Shifted by 90 Degrees Clock Creation

This example creates a 10 ns clock with a 50% duty cycle that is phase shifted by 90 degrees applied to port `clk_sys`. This type of clock definition is most commonly used when the FPGA receives source synchronous, double-rate data that is center-aligned with respect to the clock.

```
create_clock -period 10 -waveform { 2.5 7.5 } [get_ports clk_sys]
```

### Two Oscillators Driving the Same Clock Port

You can apply multiple clocks to the same target with the `-add` option. For example, to specify that the same clock input can be driven at two different frequencies, enter the following commands in your SDC file:

```
create_clock -period 10 -name clk_100 [get_ports clk_sys]
create_clock -period 5 -name clk_200 [get_ports clk_sys] -add
```

Although it is not common to have more than two base clocks defined on a port, you can define as many as are appropriate for your design, making sure you specify `-add` for all clocks after the first.

### Creating Multifrequency Clocks

You must create a multifrequency clock if your design has more than one clock source feeding a single clock node in your design. The additional clock may act as a low-power clock, with a lower frequency than the primary clock. If your design uses multifrequency clocks, use the `set_clock_groups` command to define clocks that are exclusive.

To create multifrequency clocks, use the `create_clock` command with the `-add` option to create multiple clocks on a clock node. You can create a 10 ns clock applied to clock port `clk`, and then add an additional 15 ns clock to the same clock port. The TimeQuest analyzer uses both clocks when it performs timing analysis.

```
create_clock -period 10 -name clock_primary -waveform { 0 5 } \
  [get_ports clk]
create_clock -period 15 -name clock_secondary -waveform { 0 7.5 } \
  [get_ports clk] -add
```

### Related Information

- [Accounting for Clock Effect Characteristics](#) on page 7-21
- [create\\_clock](#)
- [get\\_ports](#)

For more information about these commands, refer to Quartus II Help.

## Automatically Detecting Clocks and Creating Default Clock Constraints

To automatically create base clocks in your design, use the `derive_clocks` command. The `derive_clocks` command is equivalent to using the `create_clock` command for each register or port feeding the clock pin of a register. The `derive_clocks` command creates clock constraints on ports or registers to ensure every register in your design has a clock constraints, and it applies one period to all base clocks in your design.

You can have the TimeQuest analyzer create a base clock with a 100 Mhz requirement for unconstrained base clock nodes.

```
derive_clocks -period 10
```

**Warning:** Do not use the `derive_clocks` command for final timing sign-off; instead, you should create clocks for all clock sources with the `create_clock` and `create_generated_clock` commands. If your design has more than a single clock, the `derive_clocks` command constrains all the clocks to the same specified frequency. To achieve a thorough and realistic analysis of your design's timing requirements, you should make individual clock constraints for all clocks in your design.

If you want to have some base clocks created automatically, you can use the `-create_base_clocks` option to `derive_pll_clocks`. With this option, the `derive_pll_clocks` command automatically creates base clocks for each PLL, based on the input frequency information specified when the PLL was instantiated. The base clocks are named matching the port names. This feature works for simple port-to-PLL connections. Base clocks are not automatically generated for complex PLL connectivity, such as cascaded PLLs. You can also use the command `derive_pll_clocks -create_base_clocks` to create the input clocks for all PLL inputs automatically.

### Related Information

#### [derive\\_clocks](#)

For more information about this command, refer to Quartus II Help.

## Creating Virtual Clocks

A virtual clock is a clock that does not have a real source in the design or that does not interact directly with the design.

To create virtual clocks, use the `create_clock` command with no value specified for the `<targets>` option.

This example defines a 100Mhz virtual clock because no target is specified.

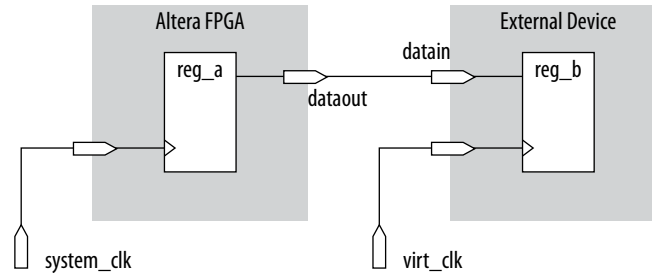
```
create_clock -period 10 -name my_virt_clk
```

## I/O Constraints with Virtual Clocks

Virtual clocks are most commonly used in I/O constraints; they represent the clock at the external device connected to the FPGA.

For the output circuit shown in the following figure, you should use a base clock to constrain the circuit in the FPGA, and a virtual clock to represent the clock driving the external device. Examples of the base clock, virtual clock, and output delay constraints for such a circuit are shown below.

Figure 7-3: Virtual Clock Board Topology



You can create a 10 ns virtual clock named `virt_clk` with a 50% duty cycle where the first rising edge occurs at 0 ns by adding the following code to your SDC file. The virtual clock is then used as the clock source for an output delay constraint.

### Example 7-1: Virtual Clock

```
#create base clock for the design
create_clock -period 5 [get_ports system_clk]
#create the virtual clock for the external register
create_clock -period 10 -name virt_clk
#set the output delay referencing the virtual clock
set_output_delay -clock virt_clk -max 1.5 [get_ports dataout]
set_output_delay -clock virt_clk -min 0.0 [get_ports dataout]
```

#### Related Information

- [set\\_input\\_delay](#)
- [set\\_output\\_delay](#)

For more information about these commands, refer to Quartus II Help.


### Example of Specifying an I/O Interface Clock

To specify I/O interface uncertainty, you must create a virtual clock and constrain the input and output ports with the `set_input_delay` and `set_output_delay` commands that reference the virtual clock.

When the `set_input_delay` or `set_output_delay` commands reference a clock port or PLL output, the virtual clock allows the `derive_clock_uncertainty` command to apply separate clock uncertainties for internal clock transfers and I/O interface clock transfers

Create the virtual clock with the same properties as the original clock that is driving the I/O port.

Figure 7-4: I/O Interface Clock Specifications



### Example 7-2: SDC Commands to Constrain the I/O Interface

```
# Create the base clock for the clock port
create_clock -period 10 -name clk_in [get_ports clk_in]
# Create a virtual clock with the same properties of the base clock
# driving the source register
create_clock -period 10 -name virt_clk_in
# Create the input delay referencing the virtual clock and not the base
# clock
# DO NOT use set_input_delay -clock clk_in <delay value>
# [get_ports data_in]
set_input_delay -clock virt_clk_in <delay value> [get_ports data_in]
```

### I/O Interface Uncertainty

Virtual clocks are recommended for I/O constraints because they most accurately represent the clocking topology of the design. An additional benefit is that you can specify different uncertainty values on clocks that interface with external I/O ports and clocks that feed register-to-register paths inside the FPGA.

#### Related Information

[Clock Uncertainty](#) on page 7-21

For more information about clock uncertainty and clock transfers.

### Creating Generated Clocks

Define generated clocks on any nodes in your design which modify the properties of a clock signal, including phase, frequency, offset, and duty cycle. Generated clocks are most commonly used on the outputs of PLLs, on register clock dividers, clock muxes, and clocks forwarded to other devices from an FPGA output port, such as source synchronous and memory interfaces. In the SDC file, create generated clocks after the base clocks have been defined. Generated clocks automatically account for all clock delays and clock latency to the generated clock target.

Use the `create_generated_clock` command to constrain generated clocks in your design.

The `-source` option specifies the name of a node in the clock path which is used as reference for your generated clock. The source of the generated clock must be a node in your design netlist and not the name of a previously defined clock. You can use any node name on the clock path between the input clock pin of the target of the generated clock and the target node of its reference clock as the source node. A good practice is to specify the input clock pin of the target node as the source of your new generated clock. That

way, the source of the generated clock is decoupled from the naming and hierarchy of its clock source. If you change its clock source, you don't have to edit the generated clock constraint.

If you have multiple base clocks feeding a node that is the source for a generated clock, you must define multiple generated clocks. Each generated clock is associated to one base clock using the `-master_clock` option in each generated clock statement. In some cases, generated clocks are generated with combinational logic. Depending on how your clock-modifying logic is synthesized, the name can change from compile to compile. If the name changes after you write the generated clock constraint, the generated clock is ignored because its target name no longer exists in the design. To avoid this problem, use a synthesis attribute or synthesis assignment to keep the final combinational node of the clock-modifying logic. Then use the kept name in your generated clock constraint. For details on keeping combinational nodes or wires, refer to the *Implement as Output of Logic Cell logic option* topic in Quartus II Help.

The TimeQuest analyzer provides the `derive_pll_clocks` command to automatically generate clocks for all PLL clock outputs. The properties of the generated clocks on the PLL outputs match the properties defined for the PLL.

### Related Information

- [Deriving PLL Clocks](#) on page 7-17  
For more information about deriving PLL clock outputs.
- [Implement as Output of Logic Cell logic option](#)  
For more information on keeping combinational nodes or wires, refer to Quartus II Help.
- [create\\_generate\\_clock](#)
- [derive\\_pll\\_clocks](#)
- [create\\_generated\\_clocks](#)  
For information about these commands, refer to Quartus II Help.
- [Specifying Timing Constraints and Exceptions](#)  
For more information about creating generated clocks, refer to Quartus II Help.

### Clock Divider Example

A common form of generated clock is a divide-by-two register clock divider. The following constraint creates a half-rate clock on the divide-by-two register.

Figure 7-5: Clock Divider





Figure 7-6: Clock Divider Waveform



```
create_clock -period 10 -name clk_sys [get_ports clk_sys]
create_generated_clock -name clk_div_2 -divide_by 2 -source \
  [get_ports clk_sys] [get_pins reg|q]
```


Or in order to have the clock source be the clock pin of the register you can use:

```
create_clock -period 10 -name clk_sys [get_ports clk_sys]
create_generated_clock -name clk_div_2 -divide_by 2 -source \
  [get_pins reg|clk] [get_pins reg|q]
```

### Clock Multiplexor Example

Another common form of generated clock is on the output of a clock mux. One generated clock on the output is required for each input clock. The SDC example also includes the `set_clock_groups` command to indicate that the two generated clocks can never be active simultaneously in the design, so the TimeQuest analyzer does not analyze cross-domain paths between the generated clocks on the output of the clock mux.

Figure 7-7: Clock Mux



```
create_clock -name clock_a -period 10 [get_ports clk_a]
create_clock -name clock_b -period 10 [get_ports clk_b]
create_generated_clock -name clock_a_mux -source [get_ports clk_a] [get_pins
  clk_mux|mux_out]
create_generated_clock -name clock_b_mux -source [get_ports clk_b] [get_pins
  clk_mux|mux_out] -add
set_clock_groups -exclusive -group clock_a_mux -group clock_b_mux
```

### Deriving PLL Clocks

Use the `derive_pll_clocks` command to direct the TimeQuest analyzer to automatically search the timing netlist for all unconstrained PLL output clocks. The `derive_pll_clocks` command detects your current PLL settings and automatically creates generated clocks on the outputs of every PLL by calling the `create_generated_clock` command.



## Create Base Clock for PLL input Clock Ports


```
create_clock -period 10.0 -name fpga_sys_clk [get_ports fpga_sys_clk]
derive_pll_clocks
```

If your design contains transceivers, LVDS transmitters, or LVDS receivers, you must use the `derive_pll_clocks` command. The command automatically constrains this logic in your design and creates timing exceptions for those blocks.

Include the `derive_pll_clocks` command in your SDC file after any `create_clock` command. Each time the TimeQuest analyzer reads your SDC file, the appropriate generate clock is created for each PLL output clock pin. If a clock exists on a PLL output before running `derive_pll_clocks`, the pre-existing clock has precedence, and an auto-generated clock is not created for that PLL output.

A simple PLL design with a register-to-register path.

**Figure 7-8: Simple PLL Design**



The TimeQuest analyzer generates messages when you use the `derive_pll_clocks` command to automatically constrain the PLL for a design similar to the previous image.

### Example 7-3: `derive_pll_clocks` Command Messages

```
Info:
Info: Deriving PLL Clocks:
Info: create_generated_clock -source pll_inst|altpll_component|pll|inclk[0] -
divide_by 2 -name
pll_inst|altpll_component|pll|clk[0] pll_inst|altpll_component|pll|clk[0]
Info:
```

The input clock pin of the PLL is the node `pll_inst|altpll_component|pll|inclk[0]` which is used for the `-source` option. The name of the output clock of the PLL is the PLL output clock node, `pll_inst|altpll_component|pll|clk[0]`.

If the PLL is in clock switchover mode, multiple clocks are created for the output clock of the PLL; one for the primary input clock (for example, `inclk[0]`), and one for the secondary input clock (for example, `inclk[1]`). You should create exclusive clock groups for the primary and secondary output clocks since they are not active simultaneously.

#### Related Information

- [Creating Clock Groups](#) on page 7-19  
For more information about creating exclusive clock groups.
- [derive\\_pll\\_clocks](#)
- [Derive PLL Clocks](#)  
For more information about the `derive_pll_clocks` command.

## Creating Clock Groups

The TimeQuest analyzer assumes all clocks are related unless constrained otherwise.

To specify clocks in your design that are exclusive or asynchronous, use the `set_clock_groups` command. The `set_clock_groups` command cuts timing between clocks in different groups, and performs the same analysis regardless of whether you specify `-exclusive` or `-asynchronous`. A group is defined with the `-group` option. The TimeQuest analyzer excludes the timing paths between clocks for each of the separate groups.

The following tables show examples of various group options for the `set_clock_groups` command.

**Table 7-1: set\_clock\_groups -group A**

Dest\Source	A	B	C	D
A	Analyzed	Cut	Cut	Cut
B	Cut	Analyzed	Analyzed	Analyzed
C	Cut	Analyzed	Analyzed	Analyzed
D	Cut	Analyzed	Analyzed	Analyzed

**Table 7-2: set\_clock\_groups -group {A B}**

Dest\Source	A	B	C	D
A	Analyzed	Analyzed	Cut	Cut
B	Analyzed	Analyzed	Cut	Cut
C	Cut	Cut	Analyzed	Analyzed
D	Cut	Cut	Analyzed	Analyzed

**Table 7-3: set\_clock\_groups -group A -group B**

Dest\Source	A	B	C	D
A	Analyzed	Cut	Cut	Cut
B	Cut	Analyzed	Cut	Cut
C	Cut	Cut	Analyzed	Analyzed
D	Cut	Cut	Analyzed	Analyzed

**Table 7-4: set\_clock\_groups -group {A C} -group {B D}**

Dest\Source	A	B	C	D
A	Analyzed	Cut	Analyzed	Cut
B	Cut	Analyzed	Cut	Analyzed
C	Analyzed	Cut	Analyzed	Cut
D	Cut	Analyzed	Cut	Analyzed

Table 7-5: set\_clock\_groups -group {A C D}

Dest\Source	A	B	C	D
A	Analyzed	Cut	Analyzed	Analyzed
B	Cut	Analyzed	Cut	Cut
C	Analyzed	Cut	Analyzed	Analyzed
D	Analyzed	Cut	Analyzed	Analyzed

**Related Information**[set\\_clock\\_groups](#)


For more information about this command, refer to Quartus II Help.

**Exclusive Clock Groups**

Use the `-exclusive` option to declare that two clocks are mutually exclusive. You may want to declare clocks as mutually exclusive when multiple clocks are created on the same node. This case occurs for multiplexed clocks.

For example, an input port may be clocked by either a 25-MHz or a 50-MHz clock. To constrain this port, create two clocks on the port, and then create clock groups to declare that they do not coexist in the design at the same time. Declaring the clocks as mutually exclusive eliminates clock transfers that are derived between the 25-MHz clock and the 50-MHz clock.

Figure 7-9: Clock Mux with Synchronous Path Across the Mux



```
create_clock -period 40 -name clk_a [get_ports {port_a}]
create_clock -add -period 20 -name clk_b [get_ports {port_a}]
set_clock_groups -exclusive -group {clk_a} -group {clk_b}
```

**Asynchronous Clock Groups**

Use the `-asynchronous` option to create asynchronous clock groups. Asynchronous clock groups are commonly used to break the timing relationship where data is transferred through a FIFO between clocks running at different rates.

**Related Information**[set\\_clock\\_groups](#)

For more information about this command, refer to Quartus II Help.

## Accounting for Clock Effect Characteristics

The clocks you create with the TimeQuest analyzer are ideal clocks that do not account for any board effects. You can account for clock effect characteristics with clock latency and clock uncertainty.

### Clock Latency

There are two forms of clock latency, clock source latency and clock network latency. Source latency is the propagation delay from the origin of the clock to the clock definition point (for example, a clock port). Network latency is the propagation delay from a clock definition point to a register's clock pin. The total latency at a register's clock pin is the sum of the source and network latencies in the clock path.

To specify source latency to any clock ports in your design, use the `set_clock_latency` command.

**Note:** The TimeQuest analyzer automatically computes network latencies; therefore, you only can characterize source latency with the `set_clock_latency` command. You must use the `-source` option.

#### Related Information

##### [set\\_clock\\_latency](#)

For more information about this command, refer to Quartus II Help.

### Clock Uncertainty

When clocks are created, they are ideal and have perfect edges. It is important to add uncertainty to those perfect edges, to mimic clock-level effects like jitter. You should include the `derive_clock_uncertainty` command in your SDC file so that appropriate setup and hold uncertainties are automatically calculated and applied to all clock transfers in your design. If you don't include the command, the TimeQuest analyzer performs it anyway; it is a critical part of constraining your design correctly.

The TimeQuest analyzer subtracts setup uncertainty from the data required time for each applicable path and adds the hold uncertainty to the data required time for each applicable path. This slightly reduces the setup and hold slack on each path.

The TimeQuest analyzer accounts for uncertainty clock effects for three types of clock-to-clock transfers; intraclock transfers, interclock transfers, and I/O interface clock transfers.

- Intraclock transfers occur when the register-to-register transfer takes place in the device and the source and destination clocks come from the same PLL output pin or clock port.
- Interclock transfers occur when a register-to-register transfer takes place in the core of the device and the source and destination clocks come from a different PLL output pin or clock port.
- I/O interface clock transfers occur when data transfers from an I/O port to the core of the device or from the core of the device to the I/O port.

To manually specify clock uncertainty, use the `set_clock_uncertainty` command. You can specify the uncertainty separately for setup and hold. You can also specify separate values for rising and falling clock transitions, although this is not commonly used. You can override the value that was automatically applied by the `derive_clock_uncertainty` command, or you can add to it.

The `derive_clock_uncertainty` command accounts for PLL clock jitter if the clock jitter on the input to a PLL is within the input jitter specification for PLL's in the specified device. If the input clock jitter for the PLL exceeds the specification, you should add additional uncertainty to your PLL output clocks to account for excess jitter with the `set_clock_uncertainty -add` command. Refer to the device handbook for your device for jitter specifications.

Another example is to use `set_clock_uncertainty -add` to add uncertainty to account for peak-to-peak jitter from a board when the jitter exceeds the jitter specification for that device. In this case you would add uncertainty to both setup and hold equal to 1/2 the jitter value:

```
set_clock_uncertainty -setup -to <clock name> \
  -setup -add <p2p jitter/2>

set_clock_uncertainty -hold -enable_same_physical_edge -to <clock name> \
  -add <p2p jitter/2>
```

There is a complex set of precedence rules for how the TimeQuest analyzer applies values from `derive_clock_uncertainty` and `set_clock_uncertainty`, which depend on the order the commands appear in your SDC files, and various options used with the commands. The Help topics referred to below contain complete descriptions of these rules. These precedence rules are much simpler to understand and implement if you follow these recommendations:

- If you want to assign your own clock uncertainty values to any clock transfers, the best practice is to put your `set_clock_uncertainty` exceptions after the `derive_clock_uncertainty` command in your SDC file.
- When you use the `-add` option for `set_clock_uncertainty`, the value you specify is added to the value from `derive_clock_uncertainty`. If you don't specify `-add`, the value you specify replaces the value from `derive_clock_uncertainty`.

#### Related Information

- [set\\_clock\\_uncertainty](#)
- [derive\\_clock\\_uncertainty](#)
- [remove\\_clock\\_uncertainty](#)

For more information about these commands, refer to Quartus II Help.

## Creating I/O Requirements

The TimeQuest analyzer reviews setup and hold relationships for designs in which an external source interacts with a register internal to the design. The TimeQuest analyzer supports input and output external delay modeling with the `set_input_delay` and `set_output_delay` commands. You can specify the clock and minimum and maximum arrival times relative to the clock.


You must specify timing requirements, including internal and external timing requirements, before you fully analyze a design. With external timing requirements specified, the TimeQuest analyzer verifies the I/O interface, or periphery of the device, against any system specification.

### Input Constraints

Input constraints allow you to specify all the external delays feeding into the device. Specify input requirements for all input ports in your design.

You can use the `set_input_delay` command to specify external input delay requirements. Use the `-clock` option to reference a virtual clock. Using a virtual clock allows the TimeQuest analyzer to correctly derive clock uncertainties for interclock and intraclock transfers. The virtual clock defines the launching clock for the input port. The TimeQuest analyzer automatically determines the latching clock inside the device that captures the input data, because all clocks in the device are defined.

Figure 7-10: Input Delay



The calculation the TimeQuest analyzer performs to determine the typical input delay.

Figure 7-11: Input Delay Calculation


$$\begin{aligned} \text{input delay}_{\text{MAX}} &= (\text{cd\_ext}_{\text{MAX}} - \text{cd\_altr}_{\text{MIN}}) + \text{tco\_ext}_{\text{MAX}} + \text{dd}_{\text{MAX}} \\ \text{input delay}_{\text{MIN}} &= (\text{cd\_ext}_{\text{MIN}} - \text{cd\_altr}_{\text{MAX}}) + \text{tco\_ext}_{\text{MIN}} + \text{dd}_{\text{MIN}} \end{aligned}$$

## Output Constraints

Output constraints allow you to specify all external delays from the device for all output ports in your design.

You can use the `set_output_delay` command to specify external output delay requirements. Use the `-clock` option to reference a virtual clock. The virtual clock defines the latching clock for the output port. The TimeQuest analyzer automatically determines the launching clock inside the device that launches the output data, because all clocks in the device are defined. The following figure is an example of an output delay referencing a virtual clock.

Figure 7-12: Output Delay



The calculation the TimeQuest analyzer performs to determine the typical out delay.

Figure 7-13: Output Delay Calculation

$$\begin{aligned} \text{output delay}_{\text{MAX}} &= \text{dd}_{\text{MAX}} + \text{tsu\_ext} + (\text{cd\_altr}_{\text{MAX}} - \text{cd\_ext}_{\text{MIN}}) \\ \text{output delay}_{\text{MIN}} &= (\text{dd}_{\text{MIN}} - \text{th\_ext} + (\text{cd\_altr}_{\text{MIN}} - \text{cd\_ext}_{\text{MAX}})) \end{aligned}$$

**Related Information**

- [set\\_input\\_delay](#)
- [set\\_output\\_delay](#)

For more information about these commands, refer to Quartus II Help.

## Creating Delay and Skew Constraints

The TimeQuest analyzer supports the Synopsys Design Constraint format for constraining timing for the ports in your design. These constraints allow the TimeQuest analyzer to perform a system static timing analysis that includes not only the device internal timing, but also any external device timing and board timing parameters.

### Advanced I/O Timing and Board Trace Model Delay

The TimeQuest analyzer can use advanced I/O timing and board trace model assignments to model I/O buffer delays in your design.

If you change any advanced I/O timing settings or board trace model assignments, recompile your design before you analyze timing, or use the `-force_dat` option to force delay annotation when you create a timing netlist.

#### Example 7-4: Forcing Delay Annotation

```
create_timing_netlist -force_dat
```

**Related Information**

- [Using Advanced I/O Timing](#)
- [I/O Management](#)

For more information about advanced I/O timing.

### Maximum Skew

To specify the maximum path-based skew requirements for registers and ports in the design and report the results of maximum skew analysis, use the `set_max_skew` command in conjunction with the `report_max_skew` command.

By default, the `set_max_skew` command excludes any input or output delay constraints.

**Related Information**

- [set\\_max\\_skew](#)
- [report\\_max\\_skew](#)

For more information about these commands, refer to Quartus II Help.

## Creating Timing Exceptions

Timing exceptions in the TimeQuest analyzer provide a way to modify the default timing analysis behavior to match the analysis required by your design. Specify timing exceptions after clocks and input and output delay constraints because timing exceptions can modify the default analysis.

### Precedence

If a conflict of node names occurs between timing exceptions, the following order of precedence applies:

1. False path
2. Minimum delays and maximum delays
3. Multicycle path

The false path timing exception has the highest precedence. Within each category, assignments to individual nodes have precedence over assignments to clocks. Finally, the remaining precedence for additional conflicts is order-dependent, such that the assignments most recently created overwrite, or partially overwrite, earlier assignments.

## False Paths

Specifying a false path in your design removes the path from timing analysis.

Use the `set_false_path` command to specify false paths in your design. You can specify either a point-to-point or clock-to-clock path as a false path. For example, a path you should specify as false path is a static configuration register that is written once during power-up initialization, but does not change state again. Although signals from static configuration registers often cross clock domains, you may not want to make false path exceptions to a clock-to-clock path, because some data may transfer across clock domains. However, you can selectively make false path exceptions from the static configuration register to all endpoints.

To make false path exceptions from all registers beginning with A to all registers beginning with B, use the following code in your SDC file.

```
set_false_path -from [get_pins A*] -to [get_pins B*]
```

The TimeQuest analyzer assumes all clocks are related unless you specify otherwise. Clock groups are a more efficient way to make false path exceptions between clocks, compared to writing multiple `set_false_path` exceptions between every clock transfer you want to eliminate.

### Related Information

- [Creating Clock Groups](#) on page 7-19  
For more information about creating exclusive clock groups.
- [set\\_false\\_path](#)  
For more information about this command, refer to Quartus II Help.

## Minimum and Maximum Delays

To specify an absolute minimum or maximum delay for a path, use the `set_min_delay` command or the `set_max_delay` commands, respectively. Specifying minimum and maximum delay directly overwrites existing setup and hold relationships with the minimum and maximum values.

Use the `set_max_delay` and `set_min_delay` commands to create constraints for asynchronous signals that do not have a specific clock relationship in your design, but require a minimum and maximum path delay. You can create minimum and maximum delay exceptions for port-to-port paths through the device without a register stage in the path. If you use minimum and maximum delay exceptions to constrain the path delay, specify both the minimum and maximum delay of the path; do not constrain only the minimum or maximum value.

If the source or destination node is clocked, the TimeQuest analyzer takes into account the clock paths, allowing more or less delay on the data path. If the source or destination node has an input or output delay, that delay is also included in the minimum or maximum delay check.

If you specify a minimum or maximum delay between timing nodes, the delay applies only to the path between the two nodes. If you specify a minimum or maximum delay for a clock, the delay applies to all paths where the source node or destination node is clocked by the clock.



You can create a minimum or maximum delay exception for an output port that does not have an output delay constraint. You cannot report timing for the paths associated with the output port; however, the TimeQuest analyzer reports any slack for the path in the setup summary and hold summary reports. Because there is no clock associated with the output port, no clock is reported for timing paths associated with the output port.

**Note:** To report timing with clock filters for output paths with minimum and maximum delay constraints, you can set the output delay for the output port with a value of zero. You can use an existing clock from the design or a virtual clock as the clock reference.

#### Related Information

- [set\\_max\\_delay](#)
- [set\\_min\\_delay](#)

For more information about these commands, refer to Quartus II Help.

## Delay Annotation

To modify the default delay values used during timing analysis, use the `set_annotated_delay` and `set_timing_derate` commands. You must update the timing netlist to see the results of these commands.

To specify different operating conditions in a single SDC file, rather than having multiple SDC files that specify different operating conditions, use the `set_annotated_delay -operating_conditions` command.

#### Related Information

- [set\\_timing\\_derate](#)
- [set\\_annotated\\_delay](#)

For more information about these commands, refer to the Quartus II Help.

## Multicycle Paths

By default, the TimeQuest analyzer performs a single-cycle analysis, which is the most restrictive type of analysis. When analyzing a path, the setup launch and latch edge times are determined by finding the closest two active edges in the respective waveforms.

For a hold analysis, the timing analyzer analyzes the path against two timing conditions for every possible setup relationship, not just the worst-case setup relationship. Therefore, the hold launch and latch times may be completely unrelated to the setup launch and latch edges. The TimeQuest analyzer does not report negative setup or hold relationships. When either a negative setup or a negative hold relationship is calculated, the TimeQuest analyzer moves both the launch and latch edges such that the setup and hold relationship becomes positive.

A multicycle constraint adjusts setup or hold relationships by the specified number of clock cycles based on the source (`-start`) or destination (`-end`) clock. An end setup multicycle constraint of 2 extends the worst-case setup latch edge by one destination clock period. If `-start` and `-end` values are not specified, the default constraint is `-end`.

Hold multicycle constraints are based on the default hold position (the default value is 0). An end hold multicycle constraint of 1 effectively subtracts one destination clock period from the default hold latch edge.

When the objects are timing nodes, the multicycle constraint only applies to the path between the two nodes. When an object is a clock, the multicycle constraint applies to all paths where the source node (`-from`) or destination node (`-to`) is clocked by the clock. When you adjust a setup relationship with a multicycle constraint, the hold relationship is adjusted automatically.

You can use TimeQuest analyzer commands to modify either the launch or latch edge times that the uses to determine a setup relationship or hold relationship.

**Table 7-6: Commands to Modify Edge Times**

Command	Description of Modification
<code>set_multicycle_path -setup -end &lt;value&gt;</code>	Latch edge time of the setup relationship
<code>set_multicycle_path -setup -start &lt;value&gt;</code>	Launch edge time of the setup relationship
<code>set_multicycle_path -hold -end &lt;value&gt;</code>	Latch edge time of the hold relationship
<code>set_multicycle_path -hold -start &lt;value&gt;</code>	Launch edge time of the hold relationship

### Common Multicycle Variations

Multicycle exceptions adjust the timing requirements for a register-to-register path, allowing the Fitter to optimally place and route a design in a device. Multicycle exceptions also can reduce compilation time and improve the quality of results, and can be used to change timing requirements. Two common multicycle variations are relaxing setup to allow a slower data transfer rate, and altering the setup to account for a phase shift.

#### Relaxing Setup with `set_multicycle_path`

A common type of multicycle exception occurs when the data transfer rate is slower than the clock cycle. Relaxing the setup relationship opens the window when data is accepted as valid.


In this example, the source clock has a period of 10 ns, but a group of registers are enabled by a toggling clock, so they only toggle every other cycle. Since they are fed by a 10 ns clock, the TimeQuest analyzer reports a set up of 10 ns and a hold of 0 ns. However, since the data is transferring every other cycle, the relationships should be analyzed as if the clock were operating at 20 ns, which would result in a setup of 20 ns, while the hold remains 0 ns, in essence, extending the window of time when the data can be recognized.

The following pair of multicycle assignments relax the setup relationship by specifying the `-setup` value of N and the `-hold` value as N-1. You must specify the hold relationship with a `-hold` assignment to prevent a positive hold requirement.

#### Relaxing Setup while Maintaining Hold

```
set_multicycle_path -setup -from src_reg* -to dst_reg* 2
set_multicycle_path -hold -from src_reg* -to dst_reg* 1
```

Figure 7-14: Relaxing Setup by Multiple Cycles



This pattern can be extended to create larger setup relationships in order to ease timing closure requirements. A common use for this exception is when writing to asynchronous RAM across an I/O interface. The delay between address, data, and a write enable may be several cycles. A multicycle exception to I/O ports can allow extra time for the address and data to resolve before the enable occurs.

You can relax the setup by three cycles with the following code in your SDC file.

### Three Cycle I/O Interface Exception

```
set_multicycle_path -setup -to [get_ports {SRAM_ADD[*] SRAM_DATA[*]}] 3
set_multicycle_path -hold -to [get_ports {SRAM_ADD[*] SRAM_DATA[*]}] 2
```

### Accounting for a Phase Shift

In this example, the design contains a PLL that performs a phase-shift on a clock whose domain exchanges data with domains that do not experience the phase shift. A common example is when the destination clock is phase-shifted forward and the source clock is not, the default setup relationship becomes that phase-shift, thus shifting the window when data is accepted as valid.


For example, the following code is a circumstance where a PLL phase-shifts one output forward by a small amount, in this case 0.2 ns.

### Cross Domain Phase-Shift

```
create_generated_clock -source pll|inclk[0] -name pll|clk[0] pll|clk[0]
create_generated_clock -source pll|inclk[0] -name pll|clk[1] -phase 30 pll|clk[1]
```

The default setup relationship for this phase-shift is 0.2 ns, shown in Figure A, creating a scenario where the hold relationship is negative, which makes achieving timing closure nearly impossible.

Figure 7-15: Phase-Shifted Setup and Hold



Adding the following constraint in your SDC allows the data to transfer to the following edge.

```
set_multicycle_path -setup -from [get_clocks clk_a] -to [get_clocks clk_b] 2
```

The hold relationship is derived from the setup relationship, making a multicycle hold constraint unnecessary.

#### Related Information

- [Same Frequency Clocks with Destination Clock Offset](#) on page 7-37  
Refer to this topic for a more complete example.
- [Same Frequency Clocks with Destination Clock Offset](#) on page 7-37  
Refer to this topic for a more complete example.
- [set\\_multicycle\\_path](#)  
For more information about this command, refer to the Quartus II Help.

### Examples of Basic Multicycle Exceptions


Each example explains how the multicycle exceptions affect the default setup and hold analysis in the TimeQuest analyzer. The multicycle exceptions are applied to a simple register-to-register circuit. Both the source and destination clocks are set to 10 ns.

#### Default Settings

By default, the TimeQuest analyzer performs a single-cycle analysis to determine the setup and hold checks. Also, by default, the TimeQuest analyzer sets the end multicycle setup assignment value to one and the end multicycle hold assignment value to zero.

The source and the destination timing waveform for the source register and destination register, respectively where HC1 and HC2 are hold checks one and two and SC is the setup check.

Figure 7-16: Default Timing Diagram



The calculation that the TimeQuest analyzer performs to determine the setup check.

Figure 7-17: Setup Check

$$\begin{aligned}
 \text{setup check} &= \text{current latch edge} - \text{closest previous launch edge} \\
 &= 10 \text{ ns} - 0 \text{ ns} \\
 &= 10 \text{ ns}
 \end{aligned}$$

The most restrictive setup relationship with the default single-cycle analysis, that is, a setup relationship with an end multicycle setup assignment of one, is 10 ns.

The setup report for the default setup in the TimeQuest analyzer with the launch and latch edges highlighted.

Figure 7-18: Setup Report

Path #1: Setup slack is 9.077						
Path Summary   Statistics   Data Path   Waveform						
Data Arrival Path						
	Total	Incr	RF	Type	Fanout	Element
1	0.000	0.000				launch edge time
2	2.522	2.522	R			clock network delay
3	2.606	0.084		uTco	1	src
4	2.606	0.000	RR	CELL	1	src1q
5	2.864	0.258	RR	IC	1	dst~feeder1dataf
6	2.960	0.096	RR	CELL	1	dst~feeder1combout
7	2.960	0.000	RR	IC	1	dst1d
8	3.065	0.105	RR	CELL	1	dst

Path #1: Setup slack is 9.077						
Path Summary   Statistics   Data Path   Waveform						
Data Required Path						
	Total	Incr	RF	Type	Fanout	Element
1	10.000	10.000				latch edge time
2	12.248	2.248	R			clock network delay
3	12.142	-0.106		uTsu	1	dst

Path #1: Setup slack is 9.077		
Path Summary   Statistics   Data Path   Waveform		
Property	Value	
1 From Node	src	
2 To Node	dst	
3 Launch Clock	clk_src	
4 Latch Clock	clk_dst	
5 Data Arrival Time	3.065	
6 Data Required Time	12.142	
7 Slack	9.077	

The calculation that the TimeQuest analyzer performs to determine the hold check. Both hold checks are equivalent.

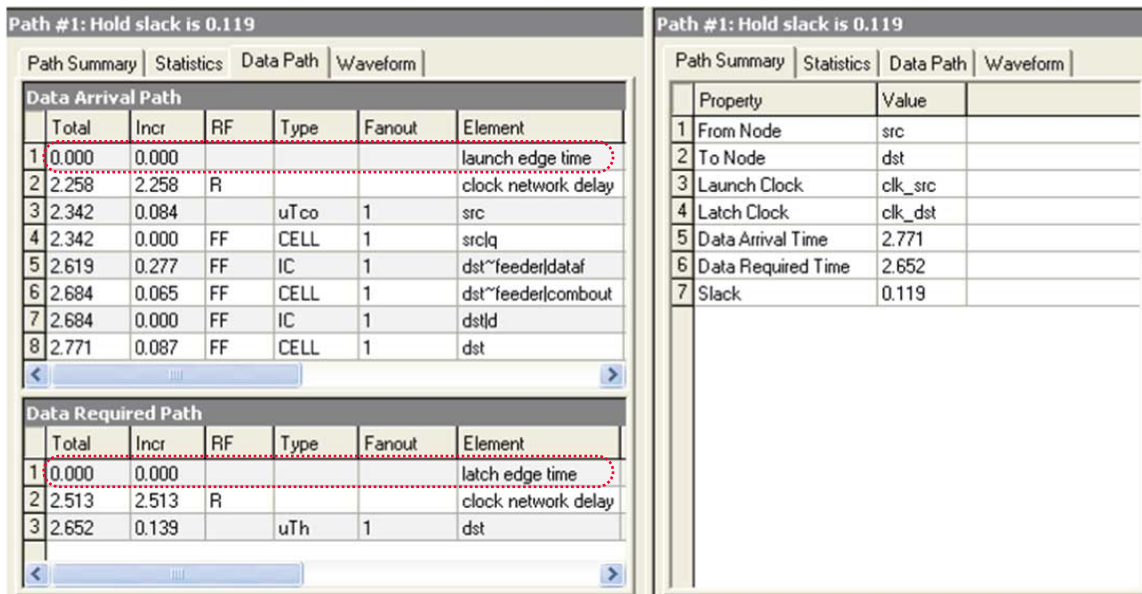
**Figure 7-19: Hold Check**

$$\begin{aligned}
 \text{hold check 1} &= \text{current launch edge} - \text{previous latch edge} \\
 &= 0 \text{ ns} - 0 \text{ ns} \\
 &= 0 \text{ ns} \\
 \\
 \text{hold check 2} &= \text{next launch edge} - \text{current latch edge} \\
 &= 10 \text{ ns} - 10 \text{ ns} \\
 &= 0 \text{ ns}
 \end{aligned}$$

The most restrictive hold relationship with the default single-cycle analysis, that a hold relationship with an end multicycle hold assignment of zero, is 0 ns.

The hold report for the default setup in the TimeQuest analyzer with the launch and latch edges highlighted.

**Figure 7-20: Hold Report**



**End Multicycle Setup = 2 and End Multicycle Hold = 0**

In this example, the end multicycle setup assignment value is two, and the end multicycle hold assignment value is zero.

**Multicycle Exceptions**

```


set_multicycle_path -from [get_clocks clk_src] -to [get_clocks clk_dst] \
    -setup -end 2
    
```

**Note:** An end multicycle hold value is not required because the default end multicycle hold value is zero.

In this example, the setup relationship is relaxed by a full clock period by moving the latch edge to the next latch edge. The hold analysis is unchanged from the default settings.

The setup timing diagram for the analysis that the TimeQuest analyzer performs. The latch edge is a clock cycle later than in the default single-cycle analysis.

**Figure 7-21: Setup Timing Diagram**



The calculation that the TimeQuest analyzer performs to determine the setup check.

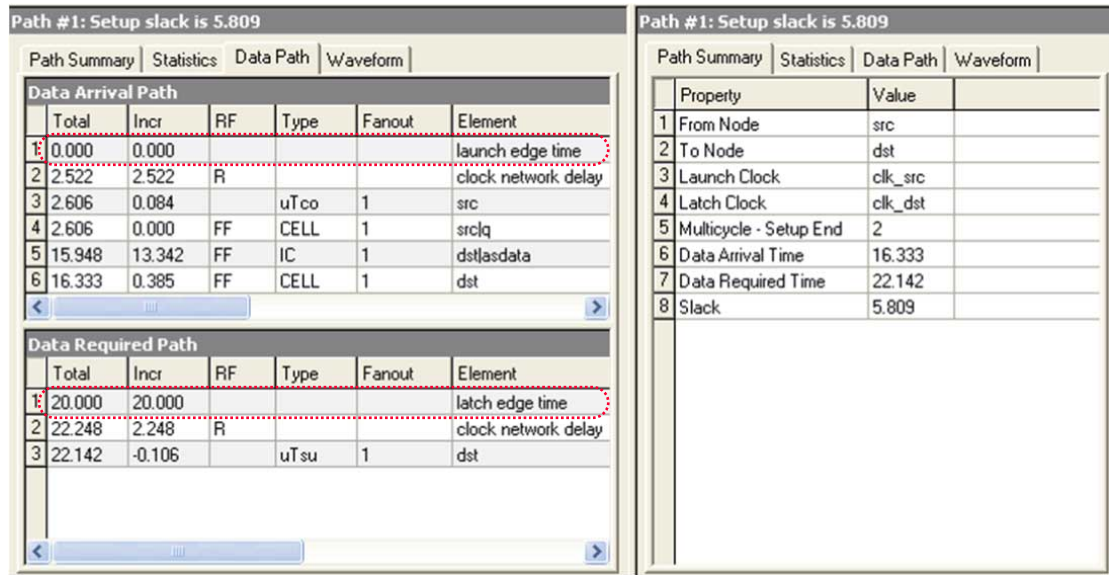
**Figure 7-22: Setup Check**

$$\begin{aligned}
 \text{setup check} &= \text{current latch edge} - \text{closest previous launch edge} \\
 &= 20 \text{ ns} - 0 \text{ ns} \\
 &= 20 \text{ ns}
 \end{aligned}$$

The most restrictive setup relationship with an end multicycle setup assignment of two is 20 ns.


The setup report in the TimeQuest analyzer with the launch and latch edges highlighted.

Figure 7-23: Setup Report



Because the multicycle hold latch and launch edges are the same as the results of hold analysis with the default settings, the multicycle hold analysis in this example is equivalent to the single-cycle hold analysis. The hold checks are relative to the setup check. Usually, the TimeQuest analyzer performs hold checks on every possible setup check, not only on the most restrictive setup check edges.

Figure 7-24: Hold Timing Diagram



The calculation that the TimeQuest analyzer performs to determine the hold check. Both hold checks are equivalent.



Figure 7-25:

hold check 1 = current launch edge – previous latch edge  
= 0 ns – 10 ns  
= –10 ns

hold check 2 = next launch edge – current latch edge  
= 10 ns – 20 ns  
= –10 ns

The most restrictive hold relationship with an end multicycle setup assignment value of two and an end multicycle hold assignment value of zero is 10 ns.

The hold report for this example in the TimeQuest analyzer with the launch and latch edges highlighted.

Figure 7-26: Hold Report

Path #1: Hold slack is 3.196

Data Arrival Path						
	Total	Incr	RF	Type	Fanout	Element
1	0.000	0.000				launch edge time
2	2.258	2.258	R			clock network delay
3	2.342	0.084		uTco	1	src
4	2.342	0.000	RR	CELL	1	src1q
5	15.606	13.264	RR	IC	1	dst1asdata
6	15.848	0.242	RR	CELL	1	dst

Data Required Path						
	Total	Incr	RF	Type	Fanout	Element
1	10.000	10.000				latch edge time
2	12.513	2.513	R			clock network delay
3	12.652	0.139		uTh	1	dst

Path #1: Hold slack is 3.196		
Property	Value	
1 From Node	src	
2 To Node	dst	
3 Launch Clock	clk_src	
4 Latch Clock	clk_dst	
5 Multicycle - Setup End	2	
6 Data Arrival Time	15.848	
7 Data Required Time	12.652	
8 Slack	3.196	

### End Multicycle Setup = 2 and End Multicycle Hold = 1

In this example, the end multicycle setup assignment value is two, and the end multicycle hold assignment value is one.


### Multicycle Exceptions

```
set_multicycle_path -from [get_clocks clk_src] -to [get_clocks clk_dst] -setup -end
2
set_multicycle_path -from [get_clocks clk_src] -to [get_clocks clk_dst] -hold -end 1
```

In this example, the setup relationship is relaxed by two clock periods by moving the latch edge to the left two clock periods. The hold relationship is relaxed by a full period by moving the latch edge to the previous latch edge.

The setup timing diagram for the analysis that the TimeQuest analyzer performs.

**Figure 7-27: Setup Timing Diagram**



The calculation that the TimeQuest analyzer performs to determine the setup check.

**Figure 7-28: Setup Check**

$$\begin{aligned}
 \text{setup check} &= \text{current latch edge} - \text{closest previous launch edge} \\
 &= 20 \text{ ns} - 0 \text{ ns} \\
 &= 20 \text{ ns}
 \end{aligned}$$

The most restrictive hold relationship with an end multicycle setup assignment value of two is 20 ns.

The setup report for this example in the TimeQuest analyzer with the launch and latch edges highlighted.

Figure 7-29: Setup Report

Path #1: Setup slack is 19.077						
Path Summary   Statistics   Data Path   Waveform						
Data Arrival Path						
	Total	Incr	RF	Type	Fanout	Element
1	0.000	0.000				launch edge time
2	2.522	2.522	R			clock network delay
3	2.606	0.084		uTco	1	src
4	2.606	0.000	RR	CELL	1	srcdq
5	2.864	0.258	RR	IC	1	dst~feeder dataf
6	2.960	0.096	RR	CELL	1	dst~feeder combout
7	2.960	0.000	RR	IC	1	dst d
8	3.065	0.105	RR	CELL	1	dst

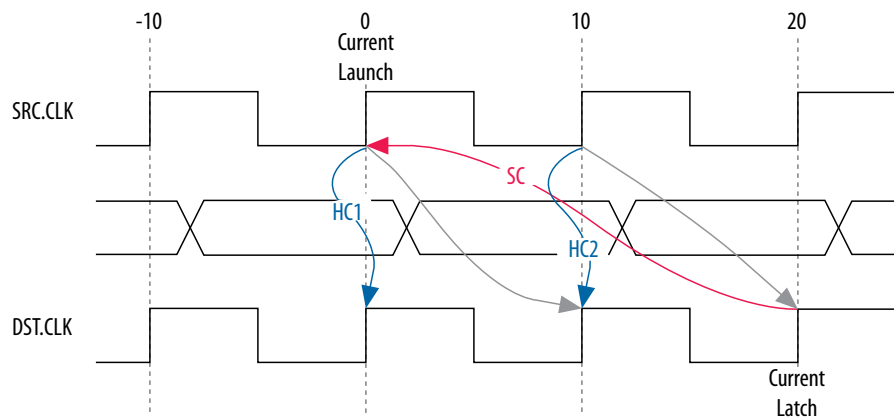
Data Required Path						
	Total	Incr	RF	Type	Fanout	Element
1	20.000	20.000				latch edge time
2	22.248	2.248	R			clock network delay
3	22.142	-0.106		uTsu	1	dst

Path #1: Setup slack is 19.077		
Path Summary   Statistics   Data Path   Waveform		
Property	Value	
1 From Node	src	
2 To Node	dst	
3 Launch Clock	clk_src	
4 Latch Clock	clk_dst	
5 Multicycle - Setup End	2	
6 Data Arrival Time	3.065	
7 Data Required Time	22.142	
8 Slack	19.077	

The timing diagram for the hold checks for this example. The hold checks are relative to the setup check.

Figure 7-30: Hold Timing Diagram



The calculation that the TimeQuest analyzer performs to determine the hold check. Both hold checks are equivalent.

Figure 7-31: Hold Check

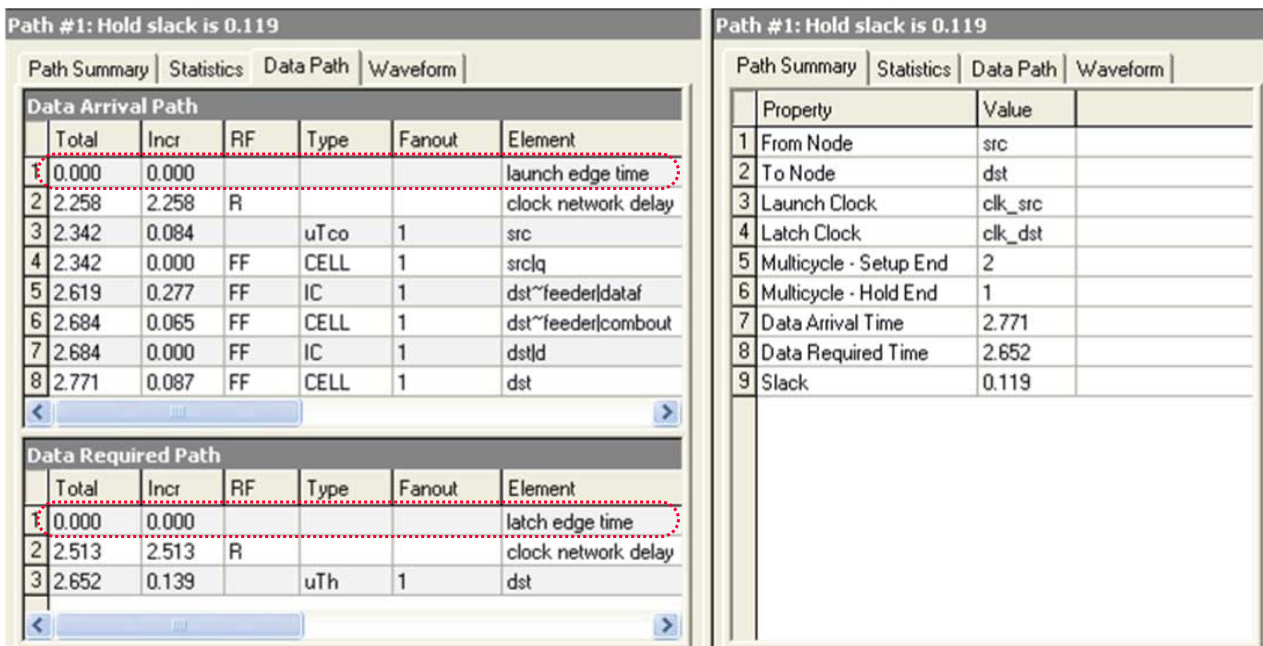
hold check 1 = current launch edge – previous latch edge  
 = 0 ns – 0 ns  
 = 0 ns

hold check 2 = next launch edge – current latch edge  
 = 10 ns – 10 ns  
 = 0 ns

The most restrictive hold relationship with an end multicycle setup assignment value of two and an end multicycle hold assignment value of one is 0 ns.

The hold report for this example in the TimeQuest analyzer with the launch and latch edges highlighted.

Figure 7-32: Hold Report




### Application of Multicycle Exceptions

This section shows the following examples of applications of multicycle exceptions. Each example explains how the multicycle exceptions affect the default setup and hold analysis in the TimeQuest analyzer. All of the examples are between related clock domains. If your design contains related clocks, such as PLL clocks, and paths between related clock domains, you can apply multicycle constraints.


#### Same Frequency Clocks with Destination Clock Offset

In this example, the source and destination clocks have the same frequency, but the destination clock is offset with a positive phase shift. Both the source and destination clocks have a period of 10 ns. The destination clock has a positive phase shift of 2 ns with respect to the source clock.

An example of a design with same frequency clocks and a destination clock offset.

**Figure 7-33: Same Frequency Clocks with Destination Clock Offset**

The timing diagram for default setup check analysis that the TimeQuest analyzer performs.

**Figure 7-34: Setup Timing Diagram**

The calculation that the TimeQuest analyzer performs to determine the setup check.

**Figure 7-35: Setup Check**

$$\begin{aligned}
 \text{setup check} &= \text{current latch edge} - \text{closest previous launch edge} \\
 &= 2 \text{ ns} - 0 \text{ ns} \\
 &= 2 \text{ ns}
 \end{aligned}$$


The setup relationship shown is too pessimistic and is not the setup relationship required for typical designs. To correct the default analysis, you must use an end multicycle setup exception of two. A multicycle exception used to correct the default analysis in this example in your SDC file.

### Multicycle Exceptions

```
set_multicycle_path -from [get_clocks clk_src] -to [get_clocks clk_dst] \
  -setup -end 2
```


The timing diagram for the preferred setup relationship for this example.

Figure 7-36: Preferred Setup Relationship



The timing diagram for default hold check analysis that the TimeQuest analyzer performs with an end multicycle setup value of two.

Figure 7-37: Default Hold Check



The calculation that the TimeQuest analyzer performs to determine the hold check.

Figure 7-38: Hold Check


$$\begin{aligned}
 \text{hold check 1} &= \text{current launch edge} - \text{previous latch edge} \\
 &= 0 \text{ ns} - 2 \text{ ns} \\
 &= -2 \text{ ns}
 \end{aligned}$$

$$\begin{aligned}
 \text{hold check 2} &= \text{next launch edge} - \text{current latch edge} \\
 &= 10 \text{ ns} - 12 \text{ ns} \\
 &= -2 \text{ ns}
 \end{aligned}$$

In this example, the default hold analysis returns the preferred hold requirements and no multicycle hold exceptions are required.

The associated setup and hold analysis if the phase shift is  $-2$  ns. In this example, the default hold analysis is correct for the negative phase shift of 2 ns, and no multicycle exceptions are required.

**Figure 7-39: Negative Phase Shift**




### Destination Clock Frequency is a Multiple of the Source Clock Frequency

In this example, the destination clock frequency value of 5 ns is an integer multiple of the source clock frequency of 10 ns. The destination clock frequency can be an integer multiple of the source clock frequency when a PLL is used to generate both clocks with a phase shift applied to the destination clock.


An example of a design where the destination clock frequency is a multiple of the source clock frequency.

**Figure 7-40: Destination Clock is Multiple of Source Clock**



The timing diagram for default setup check analysis that the TimeQuest analyzer performs.

**Figure 7-41: Setup Timing Diagram**



The calculation that the TimeQuest analyzer performs to determine the setup check.

**Figure 7-42: Setup Check**

$$\begin{aligned} \text{setup check} &= \text{current latch edge} - \text{closest previous launch edge} \\ &= 5 \text{ ns} - 0 \text{ ns} \\ &= 5 \text{ ns} \end{aligned}$$


The setup relationship demonstrates that the data does not need to be captured at edge one, but can be captured at edge two; therefore, you can relax the setup requirement. To correct the default analysis, you must shift the latch edge by one clock period with an end multicycle setup exception of two. The multicycle exception assignment used to correct the default analysis in this example.

**Multicycle Exceptions**

```
set_multicycle_path -from [get_clocks clk_src] -to [get_clocks clk_dst] \
    -setup -end 2
```


The timing diagram for the preferred setup relationship for this example.

**Figure 7-43: Preferred Setup Analysis**



The timing diagram for default hold check analysis performed by the TimeQuest analyzer with an end multicycle setup value of two.

**Figure 7-44: Default Hold Check**





The calculation that the TimeQuest analyzer performs to determine the hold check.

**Figure 7-45: Hold Check**


$$\begin{aligned}
 \text{hold check 1} &= \text{current launch edge} - \text{previous latch edge} \\
 &= 0 \text{ ns} - 5 \text{ ns} \\
 &= -5 \text{ ns} \\
 \\ 
 \text{hold check 2} &= \text{next launch edge} - \text{current latch edge} \\
 &= 10 \text{ ns} - 10 \text{ ns} \\
 &= 0 \text{ ns}
 \end{aligned}$$

In this example, hold check one is too restrictive. The data is launched by the edge at 0 ns and should check against the data captured by the previous latch edge at 0 ns, which does not occur in hold check one. To correct the default analysis, you must use an end multicycle hold exception of one.

### Destination Clock Frequency is a Multiple of the Source Clock Frequency with an Offset


This example is a combination of the previous two examples. The destination clock frequency is an integer multiple of the source clock frequency and the destination clock has a positive phase shift. The destination clock frequency is 5 ns and the source clock frequency is 10 ns. The destination clock also has a positive offset of 2 ns with respect to the source clock. The destination clock frequency can be an integer multiple of the source clock frequency with an offset when a PLL is used to generate both clocks with a phase shift applied to the destination clock. The following example shows a design in which the destination clock frequency is a multiple of the source clock frequency with an offset.

**Figure 7-46: Destination Clock is Multiple of Source Clock with Offset**



The timing diagram for the default setup check analysis the TimeQuest analyzer performs.

Figure 7-47: Setup Timing Diagram



The calculation that the TimeQuest analyzer performs to determine the setup check.

Figure 7-48: Hold Check

$$\begin{aligned}
 \text{setup check} &= \text{current latch edge} - \text{closest previous launch edge} \\
 &= 2 \text{ ns} - 0 \text{ ns} \\
 &= 2 \text{ ns}
 \end{aligned}$$

The setup relationship in this example demonstrates that the data does not need to be captured at edge one, but can be captured at edge two; therefore, you can relax the setup requirement. To correct the default analysis, you must shift the latch edge by one clock period with an end multicycle setup exception of three.

The multicycle exception code you can use to correct the default analysis in this example.


### Multicycle Exceptions

```

set_multicycle_path -from [get_clocks clk_src] -to [get_clocks clk_dst] \
    -setup -end 3
    
```


The timing diagram for the preferred setup relationship for this example.

Figure 7-49: Preferred Setup Analysis



The timing diagram for default hold check analysis the TimeQuest analyzer performs with an end multicycle setup value of three.

**Figure 7-50: Default Hold Check**



The calculation that the TimeQuest analyzer performs to determine the hold check.

**Figure 7-51: Hold Check**

$$\begin{aligned}
 \text{hold check 1} &= \text{current launch edge} - \text{previous latch edge} \\
 &= 0 \text{ ns} - 5 \text{ ns} \\
 &= -5 \text{ ns} \\
 \\ 
 \text{hold check 2} &= \text{next launch edge} - \text{current latch edge} \\
 &= 10 \text{ ns} - 10 \text{ ns} \\
 &= 0 \text{ ns}
 \end{aligned}$$


In this example, hold check one is too restrictive. The data is launched by the edge at 0 ns and should check against the data captured by the previous latch edge at 2 ns, which does not occur in hold check one. To correct the default analysis, you must use an end multicycle hold exception of one.

### Source Clock Frequency is a Multiple of the Destination Clock Frequency

In this example, the source clock frequency value of 5 ns is an integer multiple of the destination clock frequency of 10 ns. The source clock frequency can be an integer multiple of the destination clock frequency when a PLL is used to generate both clocks and different multiplication and division factors are used.


An example of a design where the source clock frequency is a multiple of the destination clock frequency.

**Figure 7-52: Source Clock Frequency is Multiple of Destination Clock Frequency**



The timing diagram for default setup check analysis performed by the TimeQuest analyzer.

**Figure 7-53: Default Setup Check Analysis**



The calculation that the TimeQuest analyzer performs to determine the setup check.

**Figure 7-54: Setup Check**

$$\begin{aligned}
 \text{setup check} &= \text{current latch edge} - \text{closest previous launch edge} \\
 &= 10 \text{ ns} - 5 \text{ ns} \\
 &= 5 \text{ ns}
 \end{aligned}$$

The setup relationship shown demonstrates that the data launched at edge one does not need to be captured, and the data launched at edge two must be captured; therefore, you can relax the setup requirement. To correct the default analysis, you must shift the launch edge by one clock period with a start multicycle setup exception of two.

The multicycle exception code you can use to correct the default analysis in this example.


### Multicycle Exceptions

```

set_multicycle_path -from [get_clocks clk_src] -to [get_clocks clk_dst] \
    -setup -start 2
    
```


The timing diagram for the preferred setup relationship for this example.

Figure 7-55: Preferred Setup Check Analysis



The timing diagram for default hold check analysis the TimeQuest analyzer performs for a start multicycle setup value of two.

Figure 7-56: Default Hold Check



The calculation that the TimeQuest analyzer performs to determine the hold check.

Figure 7-57: Hold Check

$$\begin{aligned} \text{hold check 1} &= \text{current launch edge} - \text{previous latch edge} \\ &= 0 \text{ ns} - 0 \text{ ns} \\ &= 0 \text{ ns} \end{aligned}$$


$$\begin{aligned} \text{hold check 2} &= \text{next launch edge} - \text{current latch edge} \\ &= 5 \text{ ns} - 10 \text{ ns} \\ &= -5 \text{ ns} \end{aligned}$$

In this example, hold check two is too restrictive. The data is launched next by the edge at 10 ns and should check against the data captured by the current latch edge at 10 ns, which does not occur in hold check two. To correct the default analysis, you must use a start multicycle hold exception of one.

### Source Clock Frequency is a Multiple of the Destination Clock Frequency with an Offset


In this example, the source clock frequency is an integer multiple of the destination clock frequency and the destination clock has a positive phase offset. The source clock frequency is 5 ns and destination clock frequency is 10 ns. The destination clock also has a positive offset of 2 ns with respect to the source clock. The source clock frequency can be an integer multiple of the destination clock frequency with an offset when a PLL is used to generate both clocks, different multiplication.

Figure 7-58: Source Clock Frequency is Multiple of Destination Clock Frequency with Offset



Timing diagram for default setup check analysis the TimeQuest analyzer performs.

Figure 7-59: Setup Timing Diagram



The calculation that the TimeQuest analyzer performs to determine the setup check.

Figure 7-60: Setup Check

$$\begin{aligned}
 \text{setup check} &= \text{current latch edge} - \text{closest previous launch edge} \\
 &= 12 \text{ ns} - 10 \text{ ns} \\
 &= 2 \text{ ns}
 \end{aligned}$$


The setup relationship in this example demonstrates that the data is not launched at edge one, and the data that is launched at edge three must be captured; therefore, you can relax the setup requirement. To correct the default analysis, you must shift the launch edge by two clock periods with a start multicyle setup exception of three.

The multicyle exception used to correct the default analysis in this example.

```
set_multicycle_path -from [get_clocks clk_src] -to [get_clocks clk_dst] \
  -setup -start 3
```


The timing diagram for the preferred setup relationship for this example.

**Figure 7-61: Preferred Setup Check Analysis**



The timing diagram for default hold check analysis the TimeQuest analyzer performs for a start multicycle setup value of three.

**Figure 7-62: Default Hold Check Analysis**



The calculation that the TimeQuest analyzer performs to determine the hold check.

**Figure 7-63: Hold Check**

hold check 1 = current launch edge – previous latch edge  
 = 0 ns – 2 ns  
 = -2 ns


hold check 2 = next launch edge – current latch edge  
 = 5 ns – 12 ns  
 = -7 ns

In this example, hold check two is too restrictive. The data is launched next by the edge at 10 ns and should check against the data captured by the current latch edge at 12 ns, which does not occur in hold check two. To correct the default analysis, you must use a start multicycle hold exception of one.

## A Sample Design with SDC File

An example circuit that includes two clocks, a phase-locked loop (PLL), and other common synchronous design elements helps demonstrate how to constrain your design with an SDC file.

**Figure 7-64: TimeQuest Constraint Example**



The following SDC file contains basic constraints for the example circuit.

### Example 7-5: Basic SDC Constraints

```
# Create clock constraints
create_clock -name clockone -period 10.000 [get_ports {clk1}]
create_clock -name clocktwo -period 10.000 [get_ports {clk2}]
# Create virtual clocks for input and output delay constraints
create_clock -name clockone_ext -period 10.000
create_clock -name clocktwo_ext -period 10.000
derive_pll_clocks
```



```
# derive clock uncertainty
derive_clock_uncertainty
# Specify that clockone and clocktwo are unrelated by assinging
# them to seperate asynchronous groups
set_clock_groups \
  -asynchronous \
  -group {clockone} \
  -group {clocktwo altp110|altp11_component|auto_generated|p111|clk[0]}

# set input and output delays
set_input_delay -clock { clockone_ext } -max 4 [get_ports
{data1}]set_input_delay -clock { clockone_ext } -min -1 [get_ports {data1}]
set_input_delay -clock { clockone_ext } -max 4 [get_ports
{data2}]set_input_delay -clock { clockone_ext } -min -1 [get_ports {data2}]
set_output_delay -clock { clocktwo_ext } -max 6 [get_ports {dataout}]
set_output_delay -clock { clocktwo_ext } -min -3 [get_ports {dataout}]
```

The SDC file contains the following basic constraints you should include for most designs:

- Definitions of `clockone` and `clocktwo` as base clocks, and assignment of those settings to nodes in the design.
- Definitions of `clockone_ext` and `clocktwo_ext` as virtual clocks, which represent clocks driving external devices interfacing with the FPGA.
- Automated derivation of generated clocks on PLL outputs.
- Derivation of clock uncertainty.
- Specification of two clock groups, the first containing `clockone` and its related clocks, the second containing `clocktwo` and its related clocks, and the third group containing the output of the PLL. This specification overrides the default analysis of all clocks in the design as related to each other.
- Specification of input and output delays for the design.

#### Related Information

[Asynchronous Clock Groups](#) on page 7-20

For more information about asynchronous clock groups.

## Running the TimeQuest Analyzer

When you compile a design, the TimeQuest timing analyzer automatically performs multi-corner signoff timing analysis after the Fitter has finished.

- To open the TimeQuest analyzer GUI directly from the Quartus II software GUI, click **TimeQuest Timing Analyzer** on the Tools menu.
- To perform or repeat multi-corner timing analysis from the Quartus II GUI, click **Processing > Start > Start TimeQuest Timing Analyzer**.
- To perform multi-corner timing analysis from a system command prompt, type `quartus_sta <options><project_name>`.
- To run the TimeQuest analyzer as a stand-alone GUI application, type the following command at the command prompt:`quartus_staw`.
- To run the TimeQuest analyzer in interactive command-shell mode, type the following command at a system command prompt: `quartus_sta -s <options><project_name>`.

The following table lists the command-line options available for the `quartus_sta` executable.

**Table 7-7: Summary of Command-Line Options**

Command-Line Option	Description
-h   --help	Provides help information on quartus_sta.
-t <script file>   --script=<script file>	Sources the <script file>.
-s   --shell	Enters shell mode.
--tcl_eval <tcl command>	Evaluates the Tcl command <tcl command>.
--do_report_timing	For all clocks in the design, run the following commands:  <pre>report_timing -npaths 1 -to_clock \$clock report_timing -setup -npaths 1 -to_ clock \$clock report_timing -hold -npaths 1 -to_clock \$clock report_timing -recovery -npaths 1 -to_ clock \$clock report_timing -removal -npaths 1 -to_ clock \$clock</pre>
--force_dat	Forces an update of the project database with new delay information.
--lower_priority	Lowest the computing priority of the quartus_sta process.
--post_map	Uses the post-map database results.
--sdc=<SDC file>	Specifies the SDC file to use.
--report_script=<script>	Specifies a custom report script to call.
--speed=<value>	Specifies the device speed grade used for timing analysis.
--tq2pt	Generates temporary files to convert the TimeQuest analyzer SDC file(s) to a PrimeTime SDC file.
-f <argument file>	Specifies a file containing additional command-line arguments.
-c <revision name>   --rev=<revision_name>	Specifies which revision and its associated Quartus II Settings File(.qsf) to use.
--multicorner	Specifies that all slack summary reports be generated for both slow- and fast-corners.
--multicorner [=on off]	Turns off multicorner timing analysis.
--voltage=<value_in_mV>	Specifies the device voltage, in mV used for timing analysis.
--temperature= <value_in_C>	Specifies the device temperature in degrees Celsius, used for timing analysis.
--parallel [=<num_processors>]	Specifies the number of computer processors to use on a multiprocessor system.

Command-Line Option	Description
--64bit	Enables 64-bit version of the executable.

#### Related Information

- [Constraining and Analyzing with Tcl Commands](#) on page 7-60  
For more information about using Tcl commands to constrain and analyze your design
- [Recommended Flow for First Time Users](#) on page 7-2  
For more information about steps to perform before opening the TimeQuest analyzer.
- [About TimeQuest Timing Analysis](#)  
For more information about the TimeQuest analyzer GUI, refer to Quartus II Help.

## Quartus II Settings

Within the Quartus II software, there are a number of quick steps for setting up your design with TimeQuest. You can modify the appropriate settings in **Assignments > Settings**.

In the **Settings** dialog box, select **TimeQuest Timing Analyzer** in the **Category** list.

The **TimeQuest Timing Analyzer** settings page is where you specify the title and location for a Synopsis Design Constraint (SDC) file. The SDC file is an industry standard format for specifying timing constraints. If no SDC file exists, you can create one based on the instructions in this document. The Quartus II software provides an SDC template you can use to create your own.

The following TimeQuest options should be on by default:

- Enable multicorner timing analysis—Directs the TimeQuest analyzer to analyze all the timing models of your FPGA against your constraints. This is required for final timing sign-off. Unchecked, only the slow timing model is analyzed.
- Enable common clock path pessimism removal— Prevents timing analysis from over-calculating the effects of **On-Die Variation**. This makes timing better, and there really is no reason for this to be disabled.
- Report worst-case paths during compilation—This optional setting displays summary of the worst paths in your timing report. This type of path analysis is covered in more detail later in this document. While useful, this summary can increase the size of the `<project>.sta.rpt` with all of these paths.
- Tcl script file for custom reports—This optional setting should prove useful later, allowing you to add custom reports to create a custom analysis. For example, if you are only working on a portion of the full FPGA, you may want additional timing reports that cover that hierarchy.

**Note:** In addition, certain values are set by default. The default duty-cycle is 50% and the default clock frequency is 1Ghz.

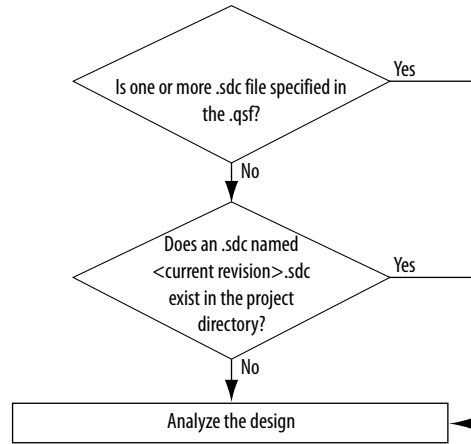
## SDC File Precedence

The Fitter and the TimeQuest analyzer process SDC files in the order you specify in the Quartus II Settings File (`.qsf`). You can add and remove SDC files to process and specify the order they are processed from the **Assignments** menu.

Click **Settings**, then **TimeQuest Timing Analyzer** and add or remove SDC files, or specify a processing order in the **SDC files to include in the project** box. When you create a new SDC file for a project, you must add it to the project for it to be read during fitting and timing analysis. If you use the Quartus II Text Editor to create an SDC file, the option to add it to the project is enabled by default when you save the file. If you use any other editor to create an SDC file, you must remember to add it to the project. If no SDC files are listed in the `.qsf`, the Quartus II software looks for an SDC named `<current revision>.sdc` in the project directory. When you use IP from Altera, and some third-parties, the SDC files are often included

in a project through an intermediate file called a Quartus II IP File (**.qip**). A **.qip** file points to all source files and constraints for a particular IP. If SDC files for IP blocks in your design are included through with a **.qip**, do not re-add them manually. An SDC file can also be added from a Quartus II IP File (**.qip**) included in the **.qsf**.

**Figure 7-65: .sdc File Order of Precedence**



**Note:** If you type the `read_sdc` command at the command line without any arguments, the TimeQuest analyzer reads constraints embedded in HDL files, then follows the SDC file precedence order.

The SDC file must contain only SDC commands that specify timing constraints. There are some techniques to control which constraints are applied during different parts of the compilation flow. Tcl commands to manipulate the timing netlist or control the compilation must be in a separate Tcl script.

## Understanding Results

Knowing how your constraints are displayed when analyzing a path is one of the most important skills of timing analysis. This information completes your understanding of timing analysis and lets you correlate the SDC input to the back-end analysis, and determine how the delays in the FPGA affect timing.

### Iterative Constraint Modification

Sometimes it is useful to change an SDC constraint and reanalyze the timing results. This flow is particularly common when you are creating timing constraints and want to ensure that they will be applied appropriately during compilation and timing analysis.

Use the following steps when you iteratively modify constraints:

1. Open the TimeQuest Timing Analyzer
2. Generate the appropriate reports.
3. Analyze your results
4. Edit your SDC file and save
5. Double-click **Reset Design**
6. Generate the appropriate reports.
7. Analyze your results
8. Repeat steps 4-7 as necessary.

**Open the TimeQuest Timing Analyzer**—It is most common to use this interactive approach in the TimeQuest GUI. You can also use the command-line shell mode, but it does not include some of the time-saving automatic features in the GUI.

**Generate the appropriate reports** —Use the **Report All Summaries** task under **Macros** to generate setup, hold, recovery, and removal summaries, as well as minimum pulse width checks, and a list of all the defined clocks. These summaries cover all constrained paths in your design. Especially when you are modifying or correcting constraints, you should also perform the Diagnostic task to create reports to identify unconstrained parts of your design, or ignored constraints. Double-click on any of the report tasks to automatically run the three tasks under **Netlist Setup** if they haven't already run. One of those tasks reads all SDC files.

**Analyze your results**—When you are modifying or correcting constraints, review the reports to find any unexpected results. For example, a cross-domain path might indicate that you forgot to cut a transfer by including a clock in a clock group.

**Edit your SDC file and save it**—Create or edit the appropriate constraints in your SDC files. If you edit your SDC file in the Quartus II Text Editor, you can benefit from tooltips showing constraint options, and dialog boxes that guide you when creating constraints.

**Reset the design**—Double click **Reset Design** task to remove all constraints from your design. Removing all constraints from your design prepares it to reread the SDC files, including your changes.

Be aware that this method just performs timing analysis using new constraints, but the fit being analyzed has not changed. The place-and-route was performed with the old constraints, but you are analyzing with new constraints, so if something is failing timing against these new constraints, you may need to run place-and-route again.

For example, the Fitter may concentrate on a very long path in your design, trying to close timing. For example, you may realize that a path runs at a lower rate, and so have added `set_multicycle_path` assignments to relax the relationship (open the window when data is valid). When you perform TimeQuest analysis iteratively with these new multicycles, new paths replace the old. The new paths may have sub-optimal placement since the Fitter was concentrating on the previous paths when it ran, because they were more critical. The iterative method is recommend for getting your SDC files correct, but you should perform a full compilation to see what the Quartus II software can do with those constraints.

#### Related Information

- [Relaxing Setup with `set\_multicycle\_path`](#) on page 7-27
- [Editing Quartus II Text Editor Files](#)

For information about editing files in the Quartus II Text Editor, refer to Quartus II Help.

## Report Timing (Dialog Box)

Once you are comfortable with the **Report All Summaries** command, the next tool in the TimeQuest analyzer toolbox is **Report Timing...**

The TimeQuest analyzer displays reports in the **Report** pane, and is similar to a table of contents for all the reports created. Selecting any name in the **Report** panel displays that report in the main viewing pane. Below is a design with the **Summary (Setup)** report highlighted:

The main viewing pane shows the Slack for every clock domain. Positive slack is good, saying these paths meet timing by that much. The End Point TNS stands for Total Negative Slack, and is the sum of all slacks for each destination and can be used as a relative assessment of how much a domain is failing.

However, this is just a summary. To get details on any domain, you can right-click that row and select **Report Timing...**

The **Report Timing** dialog box appears, auto-filled with the **Setup** radio button selected and the **To Clock** box filled with the selected clock. This occurs because you were viewing the **Setup Summary** report, and right-clicked on that particular clock. As such, the worst 10 paths where that is the destination clock were reported. You can modify the settings in various ways, such as increasing the number of paths to report, adding a **Target** filter, adding a **From Clock**, writing the report to a text file, etc.

Note that any `report_timing` command can be copied from the **Console** at the bottom into a user-created Tcl file, so that you can analyze specific paths again in the future without having to negotiate the TimeQuest analyzer UI. This is often done as users become more comfortable with TimeQuest and find themselves analyzing the same problematic parts of their design over and over, but is not required. Many complex designs successfully use TimeQuest as a diving tool, i.e. just starting with summaries and diving down into the failing paths after each compile.

## Analyzing Results with Report Timing

**Report Timing** is one of the most useful analysis tools in TimeQuest. Many designs require nothing but this command. In the TimeQuest analyzer, this command can be accessed from the **Tasks** menu, from the **Reports > Custom Reports** menu, or by right-clicking on nodes or assignments in TimeQuest.

You can review all of the options for **Report Timing** by typing `report_timing -long_help` in the TimeQuest console.

### Clocks

The **From Clock** and **To Clock** in the **Clocks** box are used to filter paths where the selected clock is used as the launch or latch. The pull-down menu allows you to choose from existing clocks (although admittedly has a "limited view" for long clock names).

### Targets

The boxes in the **Targets** box are targets for the **From Clock** and **To Clock** settings, and allow you to report paths with only particular endpoints. These are usually filled with register names or I/O ports, and can be wildcarded. For example, you might use the following to only report paths within a hierarchy of interest:

```
report_timing -from *|egress:egress_inst|* -to *|egress:egress_inst|* -(other options)
```

If the **From**, **To**, or **Through** boxes are empty, then the TimeQuest analyzer assumes you are referring to all possible targets in the device, which can also be represented with a wildcard (\*). The **From** and **To** options cover the majority of situations. The **Through** option is used to limit the report for paths that pass through combinatorial logic, or a particular pin on a cell. This is seldom used, and may not be very reliable due to combinatorial node name changes during synthesis. Clicking the browse **Browse** box after each target opens the **Name Finder** dialog box to search for specific names. This is especially useful to make sure the name being entered matches nodes in the design, since the **Name Finder** can immediately show what matches a user's wildcard.

### Analysis type

The **Analysis type** options are **Setup**, **Hold**, **Recovery**, or **Removal**. These will be explained in more detail later, as understanding them is the underpinning of timing analysis.

### Output

The **Detail** level, is an option often glanced over that should be understood. It has four options, but I will only discuss three.

The first level is called **Summary**, and produces a report which only displays Summary information such as

- **Slack**
- **From Node**
- **To Node**
- **Launch Clock**
- **Latch Clock**
- **Relationship**
- **Clock Skew**
- **Data Delay**

The **Summary** report is always reported with more detailed reports, so the user would choose this if they want less info. A good use for summary detail is when writing the report to a text file, where **Summary** can be quite brief.

The next level is **Path only**. This report displays all the detailed information, except the **Data Path** tab displays the clock tree as one line item. This is useful when you know the clock tree is correct, details are not relevant. This is common for most paths within the FPGA. A useful data point is to look at the **Clock Skew** column in the **Summary** report, and if it's a small number, say less than +/-150ps, then the clock tree is well balanced between source and destination.

If there is clock skew, you should select the **Full path** option.. This breaks the clock tree out into explicit detail, showing every cell it goes through, including such things as the input buffer, PLL, global buffer (called `CLKCTRL_`), and any logic. If there is clock skew, this is where you can determine what is causing the clock skew in your design. The **Full path** option is also recommended for I/O analysis, since only the source clock or destination clock is inside the FPGA, and therefore its delay plays a critical role in meeting timing.

The Data Path tab of a detailed report gives the delay break-downs, but there is also useful information in the **Path Summary** and **Statistics** tabs, while the **Waveform** tab is useful to help visualize the **Data Path** analysis. I would suggest taking a few minutes to look at these in the user's design. The whole analysis takes some time to get comfortable with, but hopefully is clear in what it's doing.

**Report Timing** also has the **Report panel name**, which displays the name used in TimeQuest's **Report** section. There is also an optional **File name** switch, which allows you to write the information to a file. If you append .htm as a suffix, the TimeQuest analyzer produces the report as HTML. The **File options** radio buttons allow you to choose between **Overwrite** and **Append** when saving the file.

## Paths

The default value for **Report number of paths** is 10. Two endpoints may have a lot of combinatorial logic between them and might have many different paths. Likewise, a single destination may have hundreds of paths leading to it. Because of this, you might list hundreds of paths, many of which have the same destination and might have the same source. By turning on **Pairs only** you can list only one path for each pair of source and destination. An even more powerful way to filter the report is limit the Maximum number of paths per endpoints. You can also filter paths by entering a value in the **Maximum slack limit** field.

## Tcl command

Finally, at the bottom is the **Tcl command** field, which displays the Tcl syntax of what is run in TimeQuest. You can edit this directly before running the **Report Timing** command.

**Note:**

A useful addition is to add the `-false_path` option to the command line string. With this option, only false paths are listed. A false path is any path where the launch and latch clock have been defined, but the path was cut with either a `set_false_path` assignment or `set_clock_groups_assignment`. Paths where the launch or latch clock was never constrained are not considered false paths. This option is useful to see if a false path assignment worked and what paths it covers, or to look for paths between clock domains that should not exist. The **Task** window's **Report False Path** custom report is nothing more than **Report Timing** with the `-false_path` flag enabled.

## Correlating Constraints to the Timing Report

A critical part of timing analysis is how timing constraints appear in the **Report Timing** analysis. Most constraints only affect the launch and latch edges. Specifically, `create_clock` and `create_generated_clock` create clocks with default relationships. The command `set_multicycle_path` will modify those default relationships, while `set_max_delay` and `set_min_delay` are low-level overrides that explicitly tell TimeQuest what the launch and latch edges should be.

The following figures are from an example of the output of **Report Timing** on a particular path.

Initially, the design features a clock driving the source and destination registers with a period of 10ns. This results in a setup relationship of 10ns (launch edge = 0ns, latch edge = 10ns) and hold relationship of 0ns (launch edge = 0ns, latch edge = 0ns) from the command:

```
create_clock -name clocktwo -period 10.000 [get_ports {clk2}]
```

Figure 7-66: Setup Relationship 10ns, Hold Relationship 0ns

**Path #1: Setup slack is 6.429**

Path Summary | Statistics | Data Path | Waveform | Extra Fitter Information

**Data Arrival Path**

	Total	Incr	RF	Type	Fanout	Location
1	0.000	0.000				launch edge time
2	4.578	4.578				clock path
1	0.000	0.000				source latency
2	0.000	0.000			1	PIN_H13 clk2

**Data Required Path**

	Total	Incr	RF	Type	Fanout	Location
1	10.000	10.000				latch edge time
2	13.876	3.876				clock path
1	10.000	0.000				source latency
2	10.000	0.000			1	PIN_H13 clk2
3	10.000	0.000	RR	IC	1	IOIBUF_X56_Y81_N1 clk2~input j



Path #1: Hold slack is 0.468															
Path Summary		Statistics	Data Path	Waveform	Extra Fitter Information										
<b>Data Arrival Path</b>															
	Total	Incr	RF	Type	Fanout	Location									
1	0.000	0.000				launch edge time									
2	4.397	4.397				clock path									
1	0.000	0.000				source latency									
2	0.000	0.000			1	PIN_N16	clk1								
3	0.000	0.000	DD	IC	1	TOTRIE_Y80_Y35_N44	clk1~input[i]								
<table border="1"> <tr> <td colspan="8">   </td> </tr> </table>															
<b>Data Required Path</b>															
	Total	Incr	RF	Type	Fanout	Location									
1	0.000	0.000				latch edge time									
2	4.539	4.539				clock path									
1	0.000	0.000				source latency									
2	0.000	0.000			1	PIN_N16	clk1								
3	0.000	0.000	DD	IC	1	TOTRIE_Y80_Y35_N44	clk1~input[i]								
<table border="1"> <tr> <td colspan="8">   </td> </tr> </table>															

In the next figure, using `set_multicycle_path` adds multicycles to relax the setup relationship, or open the window, making the setup relationship 20ns while the hold relationship is still 0ns:

```
set_multicycle_path -from clocktwo -to clocktwo -setup -end 2
set_multicycle_path -from clocktwo -to clocktwo -hold -end 1
```

Figure 7-67: Setup Relationship 20ns

Path #1: Setup slack is 16.429															
Path Summary		Statistics	Data Path	Waveform	Extra Fitter Information										
<b>Data Arrival Path</b>															
	Total	Incr	RF	Type	Fanout	Location									
1	0.000	0.000				launch edge time									
2	4.578	4.578				clock path									
1	0.000	0.000				source latency									
2	0.000	0.000			1	PIN_H13	clk2								
<table border="1"> <tr> <td colspan="8">   </td> </tr> </table>															
<b>Data Required Path</b>															
	Total	Incr	RF	Type	Fanout	Location									
1	20.000	20.000				latch edge time									
2	23.876	3.876				clock path									
1	20.000	0.000				source latency									
2	20.000	0.000			1	PIN_H13	clk2								
3	20.000	0.000	RR	IC	1	IOIBUF_X56_Y81_N1	clk2~input[i]								
<table border="1"> <tr> <td colspan="8">   </td> </tr> </table>															

In the last figure, using the `set_max_delay` and `set_min_delay` constraints lets you explicitly override the relationships. Note that the only thing changing for these different constraints is the Launch Edge Time and Latch Edge Times for setup and hold analysis. Every other line item comes from delays inside the FPGA and are static for a given fit. Whenever analyzing how your constraints affect the timing requirements, this is the place to look.

Figure 7-68: Using `set_max_delay` and `set_min_delay`

**Path #1: Setup slack is 11.429**

Path Summary | Statistics | Data Path | Waveform | Extra Fitter Information

**Data Arrival Path**

	Total	Incr	RF	Type	Fanout	Location
1	0.000	0.000				launch edge time
2	4.578	4.578				clock path
1	0.000	0.000				source latency
2	0.000	0.000			1	PIN_H13 clk2

**Data Required Path**

	Total	Incr	RF	Type	Fanout	Location
1	15.000	15.000				latch edge time
2	18.876	3.876				clock path
1	15.000	0.000				source latency
2	15.000	0.000			1	PIN_H13 clk2
3	15.000	0.000	RR	IC	1	IOIBUF_X56_Y81_N1 clk2~input i

**Path #1: Hold slack is -9.574 (VIOLATED)**

Path Summary | Statistics | Data Path | Waveform | Extra Fitter Information

**Data Arrival Path**

	Total	Incr	RF	Type	Fanout	Location
1	0.000	0.000				launch edge time
2	4.137	4.137				clock path
1	0.000	0.000				source latency
2	0.000	0.000			1	PIN_H13 clk2

**Data Required Path**

	Total	Incr	RF	Type	Fanout	Location
1	10.000	10.000				latch edge time
2	14.249	4.249				clock path
1	10.000	0.000				source latency
2	10.000	0.000			1	PIN_H13 clk2
3	10.000	0.000	RR	IC	1	IOIBUF_X56_Y81_N1 clk2~input i

For I/O, this all holds true except we must add in the `-max` and `-min` values. They are displayed as **iExt** or **oExt** in the **Type** column. An example would be an output port with a `set_output_delay -max 1.0` and `set_output_delay -min -0.5`:

Once again, the launch and latch edge times are determined by the clock relationships, multicycles and possibly `set_max_delay` or `set_min_delay` constraints. The value of `set_output_delay` is also added in as an **oExt** value. For outputs this value is part of the **Data Required Path**, since this is the external part of the analysis. The setup report on the left will subtract the `-max` value, making the setup relationship harder to meet, since we want the **Data Arrival Path** to be shorter than the **Data Required Path**. The `-min` value is also subtracted, which is why a negative number makes hold timing more restrictive, since we want the **Data Arrival Path** to be longer than the **Data Required Path**.

#### Related Information

[Relaxing Setup with `set\_multicycle\_path`](#) on page 7-27

## Constraining and Analyzing with Tcl Commands

You can use Tcl commands from the Quartus II software Tcl Application Programming Interface (API) to constrain, analyze, and collect information for your design. This section focuses on executing timing analysis tasks with Tcl commands; however, you can perform many of the same functions in the TimeQuest analyzer GUI. SDC commands are Tcl commands for constraining a design. SDC extension commands provide additional constraint methods and are specific to the TimeQuest analyzer. Additional TimeQuest analyzer commands are available for controlling timing analysis and reporting. These commands are contained in the following Tcl packages available in the Quartus II software:

- `::quartus::sta`
- `::quartus::sdc`
- `::quartus::sdc_ext`

#### Related Information

- [::quartus::sta](#)  
For more information about TimeQuest analyzer Tcl commands and a complete list of commands, refer to Quartus II Help.
- [::quartus::sdc](#)  
For more information about standard SDC commands and a complete list of commands, refer to Quartus II Help.
- [::quartus::sdc\\_ext](#)  
For more information about Altera extensions of SDC commands and a complete list of commands, refer to Quartus II Help.

## Collection Commands

The TimeQuest analyzer Tcl commands often return data in an object called a collection. In your Tcl scripts you can iterate over the values in collections to access data contained in them. The software returns collections instead of Tcl lists because collections are more efficient than lists for large sets of data.

The TimeQuest analyzer supports collection commands that provide easy access to ports, pins, cells, or nodes in the design. Use collection commands with any constraints or Tcl commands specified in the TimeQuest analyzer.

**Table 7-8: SDC Collection Commands**

Command	Description of the collection returned
<code>all_clocks</code>	All clocks in the design.
<code>all_inputs</code>	All input ports in the design.
<code>all_outputs</code>	All output ports in the design.
<code>all_registers</code>	All registers in the design.
<code>get_cells</code>	Cells in the design. All cell names in the collection match the specified pattern. Wildcards can be used to select multiple cells at the same time.
<code>get_clocks</code>	Lists clocks in the design. When used as an argument to another command, such as the <code>-from</code> or <code>-to</code> of <code>set_multicycle_path</code> , each node in the clock represents all nodes clocked by the clocks in the collection. The default uses the specific node (even if it is a clock) as the target of a command.
<code>get_nets</code>	Nets in the design. All net names in the collection match the specified pattern. You can use wildcards to select multiple nets at the same time.
<code>get_pins</code>	Pins in the design. All pin names in the collection match the specified pattern. You can use wildcards to select multiple pins at the same time.
<code>get_ports</code>	Ports (design inputs and outputs) in the design.

You can also examine collections and experiment with collections using wildcards in the TimeQuest analyzer by clicking **Name Finder** from the **View** menu.

## Wildcard Characters

To apply constraints to many nodes in a design, use the “\*” and “?” wildcard characters. The “\*” wildcard character matches any string; the “?” wildcard character matches any single character.

If you make an assignment to node `reg*`, the TimeQuest analyzer searches for and applies the assignment to all design nodes that match the prefix `reg` with any number of following characters, such as `reg`, `reg1`, `reg[2]`, `regbank`, and `reg12bank`.

If you make an assignment to a node specified as `reg?`, the TimeQuest analyzer searches and applies the assignment to all design nodes that match the prefix `reg` and any single character following; for example, `reg1`, `rega`, and `reg4`.

## Adding and Removing Collection Items

Wildcards used with collection commands define collection items identified by the command. For example, if a design contains registers named `src0`, `src1`, `src2`, and `dst0`, the collection command `[get_registers src*]` identifies registers `src0`, `src1`, and `src2`, but not register `dst0`. To identify register `dst0`, you must use an additional command, `[get_registers dst*]`. To include `dst0`, you could also specify a collection command `[get_registers {src* dst*}]`.

To modify collections, use the `add_to_collection` and `remove_from_collection` commands. The `add_to_collection` command allows you to add additional items to an existing collection.

### add\_to\_collection Command

```
add_to_collection <first collection> <second collection>
```

**Note:** The `add_to_collection` command creates a new collection that is the union of the two specified collections.

The `remove_from_collection` command allows you to remove items from an existing collection.

### remove\_from\_collection Command

```
remove_from_collection <first collection> <second collection>
```

You can use the following code as an example for using `add_to_collection` for adding items to a collection.

### Adding Items to a Collection

```
#Setting up initial collection of registers
set regsl [get_registers a*]
#Setting up initial collection of keepers
set kprsl [get_keepers b*]
#Creating a new set of registers of $regsl and $kprsl
set regs_union [add_to_collection $kprsl $regsl]
#OR
#Creating a new set of registers of $regsl and b*
#Note that the new collection appends only registers with name b*
# not all keepers
set regs_union [add_to_collection $regsl b*]
```

In the Quartus II software, keepers are I/O ports or registers. A SDC file that includes `get_keepers` can only be processed as part of the TimeQuest analyzer flow and is not compatible with third-party timing analysis flows.

#### Related Information

- [add\\_to\\_collection](#)
- [remove\\_from\\_collection](#)

For more information about the `add_to_collection` and `remove_from_collection` commands, refer to Quartus II Help.

### Getting Other Information about Collections

You can display the contents of a collection with the `query_collection` command. Use the `-report_format` option to return the contents in a format of one element per line. The `-list_format` option returns the contents in a Tcl list.

```
query_collection -report_format -all $regs_union
```

Use the `get_collection_size` command to return the size of a collection; the number of items it contains. If your collection is in a variable named `col`, it is more efficient to use `set num_items [get_collection_size $col]` than `set num_items [llength [query_collection -list_format $col]]`

### Using the get\_pins Command

The `get_pins` command supports options that control the matching behavior of the wildcard character (\*). Depending on the combination of options you use, you can make the wildcard character (\*) respect or

ignore individual levels of hierarchy, which are indicated by the pipe character (|). By default, the wildcard character (\*) matches only a single level of hierarchy.

These examples filter the following node and pin names to illustrate function:

- foo (a hierarchy level named foo)
- foo|dataa (an input pin in the instance foo)
- foo|datab (an input pin in the instance foo)
- foo|bar (a combinational node named bar in the foo instance)
- foo|bar|datac (an input pin to the combinational node named bar)
- foo|bar|datad (an input pin to the combinational node bar)

**Table 7-9: Sample Search Strings and Search Results**

Search String	Search Result
get_pins * dataa	foo dataa
get_pins * datac	<empty> <sup>(1)</sup>
get_pins * * datac	foo bar datac
get_pins foo* *	foo dataa, foo datab
get_pins -hierarchical * * datac	<empty> <sup>(1)</sup>
get_pins -hierarchical foo *	foo dataa, foo datab
get_pins -hierarchical * datac	foo bar datac
get_pins -hierarchical foo * datac	<empty> <sup>(1)</sup>
get_pins -compatibility_mode * datac	foo bar datac <sup>(2)</sup>
get_pins -compatibility_mode * * datac	foo bar datac

The default method separates hierarchy levels of instances from nodes and pins with the pipe character (|). A match occurs when the levels of hierarchy match, and the string values including wildcards match the instance and/or pin names. For example, the command `get_pins <instance_name>|*|datac` returns all the `datac` pins for registers in a given instance. However, the command `get_pins *|datac` returns an empty collection because the levels of hierarchy do not match.

Use the `-hierarchical` matching scheme to return a collection of cells or pins in all hierarchies of your design.

For example, the command `get_pins -hierarchical *|datac` returns all the `datac` pins for all registers in your design. However, the command `get_pins -hierarchical *|*|datac` returns an empty collection because more than one pipe character (|) is not supported.

The `-compatibility_mode` option returns collections matching wildcard strings through any number of hierarchy levels. For example, an asterisk can match a pipe character when using `-compatibility_mode`.

<sup>(1)</sup> The search result is `<empty>` because the wildcard character (\*) does not match more than one hierarchy level, indicated by a pipe character (|), by default. This command would match any pin named `datac` in instances at the top level of the design.

<sup>(2)</sup> When you use `-compatibility_mode`, pipe characters (|) are not treated as special characters when used with wildcards.

## Identifying the Quartus II Software Executable from the SDC File

To identify which Quartus II software executable is currently running you can use the `$$::TimeQuestInfo(nameofexecutable)` variable from within an SDC file. This technique is most commonly used when you want to use an overconstraint to cause the Fitter to work harder on a particular path or set of paths in the design.

### Identifying the Quartus II Executable

```
#Identify which executable is running:
set current_exe $$::TimeQuestInfo(nameofexecutable)
if { [string equal $current_exe "quartus_fit"] } {
  #Apply .sdc assignments for Fitter executable here
} else {
  #Apply .sdc assignments for non-Fitter executables here
}
if { ![string equal "quartus_sta" $$::TimeQuestInfo(nameofexecutable)] } {
  #Apply .sdc assignments for non-TimeQuest executables here
} else {
  #Apply .sdc assignments for TimeQuest executable here
}
```

Examples of different executable names are `quartus_map` for Analysis & Synthesis, `quartus_fit` for Fitter, and `quartus_sta` for the TimeQuest analyzer.

## Locating Timing Paths in Other Tools

You can locate paths and elements from the TimeQuest analyzer to other tools in the Quartus II software.

Use the **Locate** or `Locate Path` command in the TimeQuest analyzer GUI or the `locate` command in the Tcl console in the TimeQuest analyzer GUI. Right-click on most paths or node names in the TimeQuest analyzer GUI to access the **Locate** or **Locate Path** options.

The following commands are examples of how to locate the ten paths with the worst timing slack from TimeQuest analyzer to the **Technology Map Veiwier** and locate all ports matching `data*` in the **Chip Planner**.

### Example 7-6: Locating from the TimeQuest Analyzer

```
# Locate in the Technology Map Viewer the ten paths with the worst slack
locate [get_timing_paths -npaths 10] -tmv
# locate all ports that begin with data in the Chip Planner
locate [get_ports data*] -chip
```

#### Related Information

- [Viewing Timing Analysis Results](#)  
For more information about locating paths from the TimeQuest analyzer, refer to Quartus II Help.
- [locate](#)  
For more information on this command, refer to Quartus II Help.

## Generating Timing Reports

The TimeQuest analyzer provides real-time static timing analysis result reports. The TimeQuest analyzer does not automatically generate most reports; you must create each report individually in the TimeQuest

analyzer GUI or with command-line commands. You can customize in which report to display specific timing information, excluding fields that are not required.

Some of the different command-line commands you can use to generate reports in the TimeQuest analyzer and the equivalent reports shown in the TimeQuest analyzer GUI.

**Table 7-10: TimeQuest Analyzer Reports**

Command-Line Command	Report
report_timing	Timing report
report_exceptions	Exceptions report
report_clock_transfers	Clock Transfers report
report_min_pulse_width	Minimum Pulse Width report
report_ucp	Unconstrained Paths report

During compilation, the Quartus II software generates timing reports on different timing areas in the design. You can configure various options for the TimeQuest analyzer reports generated during compilation.

You can also use the `TIMEQUEST_REPORT_WORST_CASE_TIMING_PATHS` assignment to generate a report of the worst-case timing paths for each clock domain. This report contains worst-case timing data for setup, hold, recovery, removal, and minimum pulse width checks.

Use the `TIMEQUEST_REPORT_NUM_WORST_CASE_TIMING_PATHS` assignment to specify the number of paths to report for each clock domain.

An example of how to use the `TIMEQUEST_REPORT_WORST_CASE_TIMING_PATHS` and `TIMEQUEST_REPORT_NUM_WORST_CASE_TIMING_PATHS` assignments in the `.qsf` to generate reports.

### Generating Worst-Case Timing Reports

```
#Enable Worst-Case Timing Report
set_global_assignment -name TIMEQUEST_REPORT_WORST_CASE_TIMING_PATHS ON
#Report 10 paths per clock domain
set_global_assignment -name TIMEQUEST_REPORT_NUM_WORST_CASE_TIMING_PATHS 10
```

### Fmax Summary Report panel

Fmax Summary Report panel lists the maximum frequency of each clock in your design. In some designs you may see a note indicating "Limit due to hold check. Typically, Fmax is not limited by hold checks, because they are often same-edge relationships, and therefore independent of clock frequency, for example, launch = 0, latch = 0. However, if you have an inverted clock transfer, or a multicycle transfer such as setup=2, hold=0, then the hold relationship is no longer a same-edge transfer and changes as the clock frequency changes. The value in the **Restricted Fmax** column incorporates limits due to hold time checks in the situations described previously, as well as minimum period and pulse width checks. If hold checks limit the Fmax more than setup checks, that is indicated in the **Note:** column as "Limit due to hold check".

#### Related Information

- [::quartus::sta](#)  
For more information on this command, refer to Quartus II Help.



- [TimeQuest Timing Analyzer Page](#)  
For more information about the options you can set to customize TimeQuest analyzer reports.
- [Area and Timing Optimization](#)  
For more information about timing closure recommendations.

## Document Revision History

Table 7-11: Document Revision History

Date	Version	Changes
2014.12.15	14.1.0	Major reorganization. Revised and added content to the following topic areas: <ul style="list-style-type: none"> <li>• Timing Constraints</li> <li>• Create Clocks and Clock Constraints</li> <li>• Creating Generated Clocks</li> <li>• Creating Clock Groups</li> <li>• Clock Uncertainty</li> <li>• Running the TimeQuest Analyzer</li> <li>• Generating Timing Reports</li> <li>• Understanding Results</li> <li>• Constraining and Analyzing with Tcl Commands</li> </ul>
August 2014	14.0a10.0	Added command line compilation requirements for Arria 10 devices.
June 2014	14.0.0	<ul style="list-style-type: none"> <li>• Minor updates.</li> <li>• Updated format.</li> </ul>
November 2013	13.1.0	<ul style="list-style-type: none"> <li>• Removed HardCopy device information.</li> </ul>
June 2012	12.0.0	<ul style="list-style-type: none"> <li>• Reorganized chapter.</li> <li>• Added “Creating a Constraint File from Quartus II Templates with the Quartus II Text Editor” section on creating an SDC constraints file with the <b>Insert Template</b> dialog box.</li> <li>• Added “Identifying the Quartus II Software Executable from the SDC File” section.</li> <li>• Revised multicycle exceptions section.</li> </ul>
November 2011	11.1.0	<ul style="list-style-type: none"> <li>• Consolidated content from the Best Practices for the Quartus II TimeQuest Timing Analyzer chapter.</li> <li>• Changed to new document template.</li> </ul>
May 2011	11.0.0	<ul style="list-style-type: none"> <li>• Updated to improve flow. Minor editorial updates.</li> </ul>

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> <li>• Changed to new document template.</li> <li>• Revised and reorganized entire chapter.</li> <li>• Linked to Quartus II Help.</li> </ul>
July 2010	10.0.0	Updated to link to content on SDC commands and the TimeQuest analyzer GUI in Quartus II Help.
November 2009	9.1.0	Updated for the Quartus II software version 9.1, including: <ul style="list-style-type: none"> <li>• Added information about commands for adding and removing items from collections</li> <li>• Added information about the <code>set_timing_derate</code> and <code>report_skew</code> commands</li> <li>• Added information about worst-case timing reporting</li> <li>• Minor editorial updates</li> </ul>
November 2008	8.1.0	Updated for the Quartus II software version 8.1, including: <ul style="list-style-type: none"> <li>• Added the following sections:               <ul style="list-style-type: none"> <li>“set_net_delay” on page 7-42</li> <li>“Annotated Delay” on page 7-49</li> <li>“report_net_delay” on page 7-66</li> </ul> </li> <li>• Updated the descriptions of the <code>-append</code> and <code>-file &lt;name&gt;</code> options in tables throughout the chapter</li> <li>• Updated entire chapter using 8½” × 11” chapter template</li> <li>• Minor editorial updates</li> </ul>

**Related Information**

[Quartus II Handbook Archive](#)

For previous versions of the Quartus II Handbook

2014.12.15

QI15V3



Subscribe



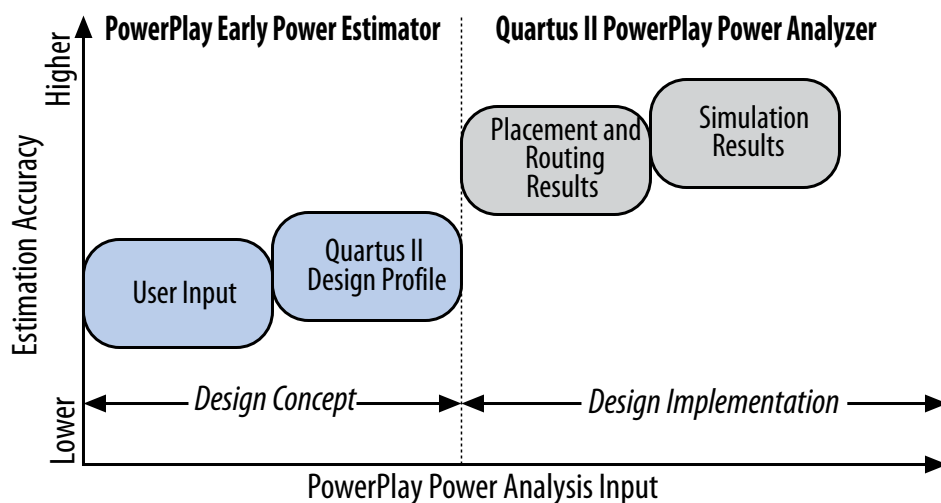
Send Feedback

The PowerPlay Power Analysis tools allow you to estimate device power consumption accurately.

As designs grow larger and process technology continues to shrink, power becomes an increasingly important design consideration. When designing a PCB, you must estimate the power consumption of a device accurately to develop an appropriate power budget, and to design the power supplies, voltage regulators, heat sink, and cooling system.

The following figure shows the PowerPlay Power Analysis tools ability to estimate power consumption from early design concept through design implementation.

**Figure 8-1: PowerPlay Power Analysis From Design Concept Through Design Implementation**



For the majority of the designs, the PowerPlay Power Analyzer and the PowerPlay EPE spreadsheet have the following accuracy after the power models are final:

- PowerPlay Power Analyzer— $\pm 20\%$  from silicon, assuming that the PowerPlay Power Analyzer uses the Value Change Dump File (.vcd) generated toggle rates.
- PowerPlay EPE spreadsheet—  $\pm 20\%$  from the PowerPlay Power Analyzer results using .vcd generated toggle rates. 90% of EPE designs (using .vcd generated toggle rates exported from PPPA) are within  $\pm 30\%$  silicon.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered

**ALTERA**®

The toggle rates are derived using the PowerPlay Power Analyzer with a .vcd file generated from a gate level simulation representative of the system operation.

#### Related Information

- [About Power Estimation and Analysis](#)
- [PowerPlay Early Power Estimators \(EPE\) and Power Analyzer](#)

## Types of Power Analyses

Understanding the uses of power analysis and the factors affecting power consumption helps you to use the PowerPlay Power Analyzer effectively. Power analysis meets the following significant planning requirements:

- **Thermal planning**—Thermal power is the power that dissipates as heat from the FPGA. You must use a heatsink or fan to act as a cooling solution for your device. The cooling solution must be sufficient to dissipate the heat that the device generates. The computed junction temperature must fall within normal device specifications.
- **Power supply planning**—Power supply is the power needed to run your device. Power supplies must provide adequate current to support device operation.

**Note:** For power supply planning, use the PowerPlay EPE at the early stages of your design cycle. Use the PowerPlay Power Analyzer reports when your design is complete to get an estimate of your design power requirement.

The two types of analyses are closely related because much of the power supplied to the device dissipates as heat from the device; however, in some situations, the two types of analyses are not identical. For example, if you use terminated I/O standards, some of the power drawn from the power supply of the device dissipates in termination resistors rather than in the device.

Power analysis also addresses the activity of your design over time as a factor that impacts the power consumption of the device. The static power ( $P_{\text{STATIC}}$ ) is the thermal power dissipated on chip, independent of user clocks.  $P_{\text{STATIC}}$  includes the leakage power from all FPGA functional blocks, except for I/O DC bias power and transceiver DC bias power, which are accounted for in the I/O and transceiver sections. Dynamic power is the additional power consumption of the device due to signal activity or toggling.

#### Related Information

- [PowerPlay Early Power Estimator \(EPE\) User Guide](#)

## Differences between the PowerPlay EPE and the Quartus II PowerPlay Power Analyzer

The following table lists the differences between the PowerPlay EPE and the Quartus II PowerPlay Power Analyzer.

**Table 8-1: Comparison of the PowerPlay EPE and Quartus II PowerPlay Power Analyzer**

Characteristic	PowerPlay EPE	Quartus II PowerPlay Power Analyzer
Phase in the design cycle	Any time, but it is recommended to use Quartus II PowerPlay Power Analyzer for post-fit power analysis.	Post-fit
Tool requirements	Spreadsheet program	The Quartus II software
Accuracy	Medium	Medium to very high
Data inputs	<ul style="list-style-type: none"> <li>Resource usage estimates</li> <li>Clock requirements</li> <li>Environmental conditions</li> <li>Toggle rate</li> </ul>	<ul style="list-style-type: none"> <li>Post-fit design</li> <li>Clock requirements</li> <li>Signal activity defaults</li> <li>Environmental conditions</li> <li>Register transfer level (RTL) simulation results (optional)</li> <li>Post-fit simulation results (optional)</li> <li>Signal activities per node or entity (optional)</li> </ul>
Data outputs (1)	<ul style="list-style-type: none"> <li>Total thermal power dissipation</li> <li>Thermal static power</li> <li>Thermal dynamic power</li> <li>Off-chip power dissipation</li> <li>Current drawn from voltage supplies</li> </ul>	<ul style="list-style-type: none"> <li>Total thermal power</li> <li>Thermal static power</li> <li>Thermal dynamic power</li> <li>Thermal I/O power</li> <li>Thermal power by design hierarchy</li> <li>Thermal power by block type</li> <li>Thermal power dissipation by clock domain</li> <li>Off-chip (non-thermal) power dissipation</li> <li>Device supply currents</li> </ul>

The result of the PowerPlay Power Analyzer is only an estimation of power. Altera does not recommend using the result as a specification. The purpose of the estimation is to help you establish guidelines for the power budget of your design. It is important that you verify the actual power during device operation as the information is sensitive to the actual device design and the environmental operating conditions.

<sup>(3)</sup> PowerPlay EPE and PowerPlay Power Analyzer outputs vary by device family. For more information, refer to the device-specific PowerPlay Early Power Estimators (EPE) and Power Analyzer Page and PowerPlay Power Analyzer Reports in the Quartus II Help.

**Note:** The PowerPlay Power Analyzer does not include the transceiver power for features that can only be enabled through dynamic reconfiguration (DFE, ADCE/AEQ, EyeQ). Use the EPE to estimate the incremental power consumption by these features.

#### Related Information

- [PowerPlay Early Power Estimators \(EPE\) and Power Analyzer Page](#)  
For more information, refer to the device-specific PowerPlay Early Power Estimators (EPE) page on the Altera website.
- [PowerPlay Power Analyzer Reports](#)  
For more information, refer to this page for device-specific information about the PowerPlay Early Power Estimator.

## Factors Affecting Power Consumption

Understanding the following factors that affect power consumption allows you to use the PowerPlay Power Analyzer and interpret its results effectively:

- [Device Selection](#)
- [Environmental Conditions](#)
- [Device Resource Usage](#)
- [Signal Activities](#)

### Device Selection

Device families have different power characteristics. Many parameters affect the device family power consumption, including choice of process technology, supply voltage, electrical design, and device architecture.

Power consumption also varies in a single device family. A larger device consumes more static power than a smaller device in the same family because of its larger transistor count. Dynamic power can also increase with device size in devices that employ global routing architectures.

The choice of device package also affects the ability of the device to dissipate heat. This choice can impact your required cooling solution choice to comply to junction temperature constraints.

Process variation can affect power consumption. Process variation primarily impacts static power because sub-threshold leakage current varies exponentially with changes in transistor threshold voltage. Therefore, you must consult device specifications for static power and not rely on empirical observation. Process variation has a weak effect on dynamic power.

### Environmental Conditions

Operating temperature primarily affects device static power consumption. Higher junction temperatures result in higher static power consumption. The device thermal power and cooling solution that you use must result in the device junction temperature remaining within the maximum operating range for the device. The main environmental parameters affecting junction temperature are the cooling solution and ambient temperature.

The following table lists the environmental conditions that could affect power consumption.

**Table 8-2: Environmental Conditions that Could Affect Power Consumption**

Environmental Conditions	Description
Airflow	<p>A measure of how quickly the device removes heated air from the vicinity of the device and replaces it with air at ambient temperature.</p> <p>You can either specify airflow as “still air” when you are not using a fan, or as the linear feet per minute rating of the fan in the system. Higher airflow decreases thermal resistance.</p>
Heat Sink and Thermal Compound	<p>A heat sink allows more efficient heat transfer from the device to the surrounding area because of its large surface area exposed to the air. The thermal compound that interfaces the heat sink to the device also influences the rate of heat dissipation. The case-to-ambient thermal resistance (<math>\theta_{CA}</math>) parameter describes the cooling capacity of the heat sink and thermal compound employed at a given airflow. Larger heat sinks and more effective thermal compounds reduce <math>\theta_{CA}</math>.</p>
Junction Temperature	<p>The junction temperature of a device is equal to:</p> $T_{\text{Junction}} = T_{\text{Ambient}} + P_{\text{Thermal}} \cdot \theta_{\text{JA}}$ <p>in which <math>\theta_{\text{JA}}</math> is the total thermal resistance from the device transistors to the environment, having units of degrees Celsius per watt. The value <math>\theta_{\text{JA}}</math> is equal to the sum of the junction-to-case (package) thermal resistance (<math>\theta_{\text{JC}}</math>), and the case-to-ambient thermal resistance (<math>\theta_{\text{CA}}</math>) of your cooling solution.</p>
Board Thermal Model	<p>The junction-to-board thermal resistance (<math>\theta_{\text{JB}}</math>) is the thermal resistance of the path through the board, having units of degrees Celsius per watt. To compute junction temperature, you can use this board thermal model along with the board temperature, the top-of-chip <math>\theta_{\text{JA}}</math> and ambient temperatures.</p>

## Device Resource Usage

The number and types of device resources used greatly affects power consumption.

- **Number, Type, and Loading of I/O Pins**—Output pins drive off-chip components, resulting in high-load capacitance that leads to a high-dynamic power per transition. Terminated I/O standards require external resistors that draw constant (static) power from the output pin.
- **Number and Type of Hard Logic Blocks**—A design with more logic elements (LEs), multiplier elements, memory blocks, transceiver blocks or HPS system tends to consume more power than a design with fewer circuit elements. The operating mode of each circuit element also affects its power consumption. For example, a DSP block performing  $18 \times 18$  multiplications and a DSP block performing multiply-accumulate operations consume different amounts of dynamic power because of different amounts of charging internal capacitance on each transition. The operating mode of a circuit element also affects static power.
- **Number and Type of Global Signals**—Global signal networks span large portions of the device and have high capacitance, resulting in significant dynamic power consumption. The type of global signal is important as well. For example, Stratix V devices support global clocks and quadrant (regional) clocks. Global clocks cover the entire device, whereas quadrant clocks only span one-fourth of the device. Clock networks that span smaller regions have lower capacitance and tend to consume less power. The location of the logic array blocks (LABs) driven by the clock network can also have an impact because the Quartus II software automatically disables unused branches of a clock.

## Signal Activities

The behavior of each signal in your design is an important factor in estimating power consumption. The following table lists the two vital behaviors of a signal, which are toggle rate and static probability:

**Table 8-3: Signal Behavior**

Signal Behavior	Description
Toggle rate	<ul style="list-style-type: none"> <li>• The toggle rate of a signal is the average number of times that the signal changes value per unit of time. The units for toggle rate are transitions per second and a transition is a change from 1 to 0, or 0 to 1.</li> <li>• Dynamic power increases linearly with the toggle rate as you charge the board trace model more frequently for logic and routing. The Quartus II software models full rail-to-rail switching. For high toggle rates, especially on circuit output I/O pins, the circuit can transition before fully charging the downstream capacitance. The result is a slightly conservative prediction of power by the PowerPlay Power Analyzer.</li> </ul>
Static probability	<ul style="list-style-type: none"> <li>• The static probability of a signal is the fraction of time that the signal is logic 1 during the period of device operation that is being analyzed. Static probability ranges from 0 (always at ground) to 1 (always at logic-high).</li> <li>• Static probabilities of their input signals can sometimes affect the static power that routing and logic consume. This effect is due to state-dependent leakage and has a larger effect on smaller process geometries. The Quartus II software models this effect on devices at 90 nm or smaller if it is important to the power estimate. The static power also varies with the static probability of a logic 1 or 0 on the I/O pin when output I/O standards drive termination resistors.</li> </ul>

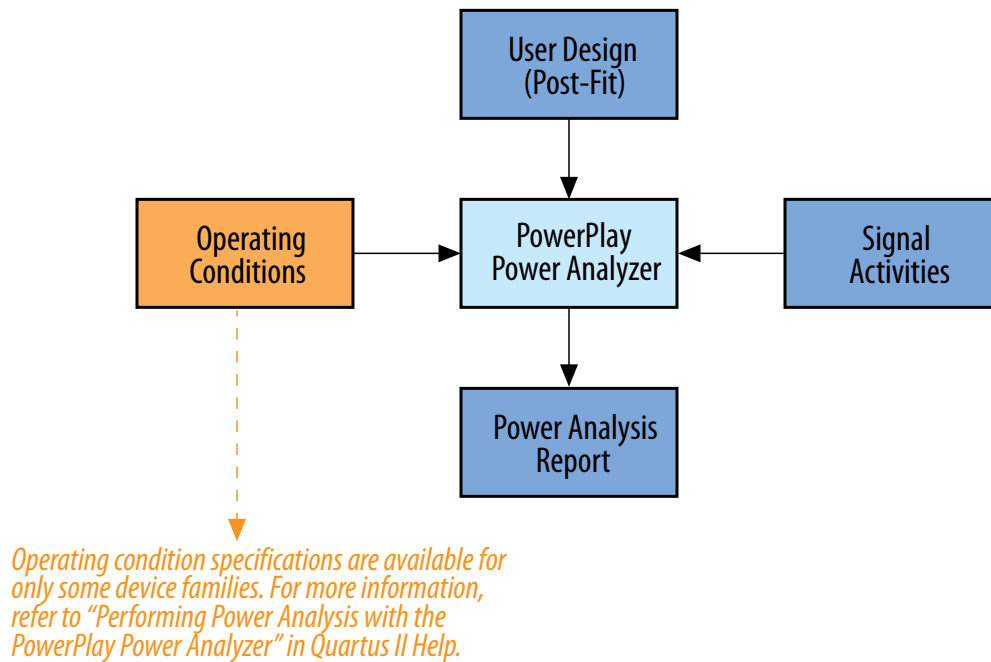


**Note:** To get accurate results from the power analysis, the signal activities for analysis must represent the actual operating behavior of your design. Inaccurate signal toggle rate data is the largest source of power estimation error.

## PowerPlay Power Analyzer Flow

The PowerPlay Power Analyzer supports accurate power estimations by allowing you to specify the important design factors affecting power consumption. The following figure shows the high-level PowerPlay Power Analyzer flow.

**Figure 8-2: PowerPlay Power Analyzer High-Level Flow**



To obtain accurate I/O power estimates, the PowerPlay Power Analyzer requires you to synthesize your design and then fit your design to the target device. You must specify the electrical standard on each I/O cell and the board trace model on each I/O standard in your design.

### Related Information

- [Performing Power Analysis with the PowerPlay Power Analyzer](#)

## Operating Settings and Conditions

You can specify device power characteristics, operating voltage conditions, and operating temperature conditions for power analysis in the Quartus II software.

On the **Operating Settings and Conditions** page of the **Settings** dialog box, you can specify whether the device has typical power consumption characteristics or maximum power consumption characteristics.

On the **Voltage** page of the **Settings** dialog box, you can view the operating voltage conditions for each power rail in the device, and specify supply voltages for power rails with selectable supply voltages.

**Note:** The Quartus II Fitter may override some of the supply voltages settings specified in this chapter. For example, supply voltages for several Stratix V transceiver power supplies depend on the data rate used. If the Fitter detects that voltage required is different from the one specified in the **Voltage** page, it will automatically set the correct voltage for relevant rails. The Quartus II PowerPlay Power Analyzer uses voltages selected by the Fitter if they conflict with the settings specified in the **Voltage** page.

On the **Temperature** page of the **Settings** dialog box, you can specify the thermal operating conditions of the device.

#### Related Information

- [Operating Settings and Conditions Page \(Settings Dialog Box\)](#)
- [Voltage Page \(Settings Dialog Box\)](#)
- [Temperature Page \(Settings Dialog Box\)](#)

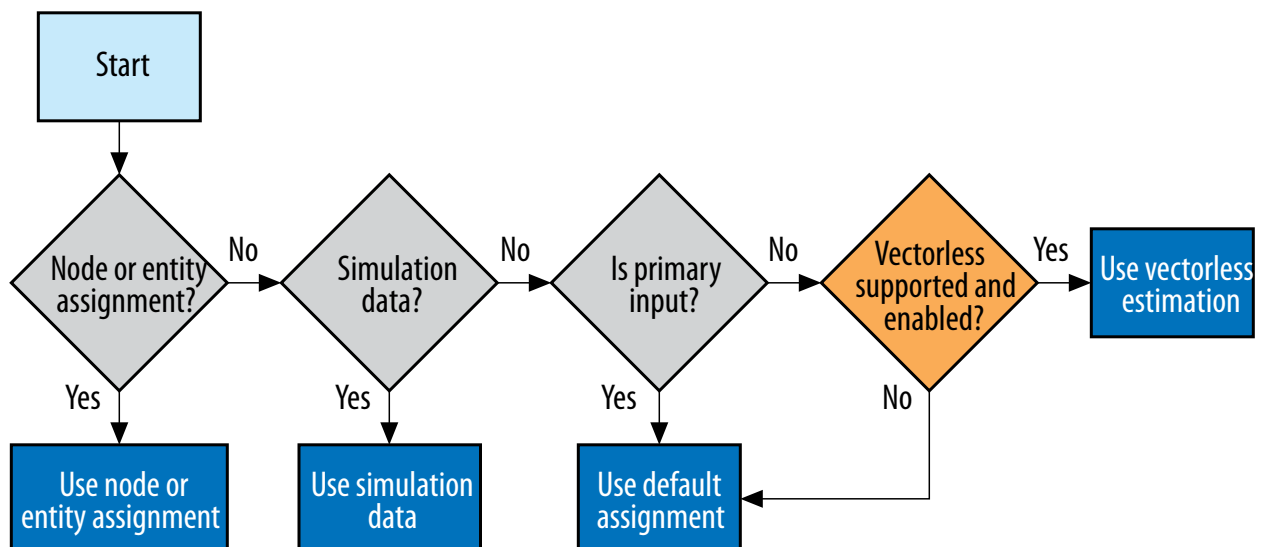
## Signal Activities Data Sources

The PowerPlay Power Analyzer provides a flexible framework for specifying signal activities. The framework reflects the importance of using representative signal-activity data during power analysis. Use the following sources to provide information about signal activity:

- Simulation results
- User-entered node, entity, and clock assignments
- User-entered default toggle rate assignment
- Vectorless estimation

The PowerPlay Power Analyzer allows you to mix and match the signal-activity data sources on a signal-by-signal basis. The following figure shows the priority scheme applied to each signal.

Figure 8-3: Signal-Activity Data Source Priority Scheme



### Related Information

- [Performing Power Analysis with the PowerPlay Power Analyzer](#)

## Simulation Results

The PowerPlay Power Analyzer directly reads the waveforms generated by a design simulation. Static probability and toggle rate can be calculated for each signal from the simulation waveform. Power analysis is most accurate when you use representative input stimuli to generate simulations.

The PowerPlay Power Analyzer reads results generated by the following simulators:

- ModelSim®
- ModelSim-Altera
- QuestaSim
- Active-HDL
- NCSim
- VCS
- VCS MX
- Riviera-PRO

Signal activity and static probability information are derived from a Verilog Value Change Dump File (.vcd). For more information, refer to [Signal Activities](#) on page 8-6.

For third-party simulators, use the **EDA Tool Settings** to specify the Generate Value Change Dump (VCD) file script option in the Simulation page of the Settings dialog box. These scripts instruct the third-party simulators to generate a .vcd that encodes the simulated waveforms. The Quartus II PowerPlay Power Analyzer reads this file directly to derive the toggle rate and static probability data for each signal.

Third-party EDA simulators, other than those listed, can generate a .vcd that you can use with the PowerPlay Power Analyzer. For those simulators, you must manually create a simulation script to generate the appropriate .vcd.

**Note:** You can use a .vcd created for power analysis to optimize your design for power during fitting by utilizing the appropriate settings in the PowerPlay power optimization list, available from **Assignments > Settings > Compiler Settings > Advanced Settings (Fitter)**.

### Related Information


#### [Power Optimization](#)

#### [Section I. Simulation](#)

## Using Simulation Files in Modular Design Flows

A common design practice is to create modular or hierarchical designs in which you develop each design entity separately, and then instantiate these modules in a higher-level entity to form a complete design. You can perform simulation on a complete design or on each module for verification. The PowerPlay Power Analyzer supports modular design flows when reading the signal activities from simulation files. The following figure shows an example of a modular design flow.

Figure 8-4: Modular Simulation Flow




When specifying a simulation file (a **.vcd**), the software provides support to specify an associated design entity name, such that the PowerPlay Power Analyzer imports the signal activities derived from that file for the specified design entity. The PowerPlay Power Analyzer also supports the specification of multiple **.vcd** files for power analysis, with each having an associated design entity name to enable the integration of partial design simulations into a complete design power analysis. When specifying multiple **.vcd** files for your design, more than one simulation file can contain signal-activity information for the same signal.

**Note:** When you apply multiple **.vcd** files to the same design entity, the signal activity used in the power analysis is the equal-weight arithmetic average of each **.vcd**.

**Note:** When you apply multiple simulation files to design entities at different levels in your design hierarchy, the signal activity in the power analysis derives from the simulation file that applies to the most specific design entity.

The following figure shows an example of a hierarchical design. The top-level module of your design, called **Top**, consists of three 8b/10b decoders, followed by a **mux**. The software then encodes the output of the **mux** to produce the final output of the top-level module. An error-handling module handles any 8b/10b decoding errors. The **Top** module contains the top-level entity of your design and any logic not defined as part of another module. The design file for the top-level module might be a wrapper for the hierarchical entities below it, or it might contain its own logic. The following usage scenarios show common ways that you can simulate your design and import the **.vcd** into the PowerPlay Power Analyzer.

Figure 8-5: Example Hierarchical Design



## Complete Design Simulation

You can simulate the entire design and generate a **.vcd** from a third-party simulator. The PowerPlay Power Analyzer can then import the **.vcd** (specifying the top-level design). The resulting power analysis uses the signal activities information from the generated **.vcd**, including those that apply to submodules, such as `decode [1-3]`, `err1`, `mux1`, and `encode1`.

## Modular Design Simulation

You can independently simulate of the top-level design, and then import all the resulting **.vcd** files into the PowerPlay Power Analyzer. For example, you can simulate the `8b10b_dec` independent of the entire design and `mux`, `8b10b_rxerr`, and `8b10b_enc`. You can then import the **.vcd** files generated from each simulation by specifying the appropriate instance name. For example, if the files produced by the simulations are **8b10b\_dec.vcd**, **8b10b\_enc.vcd**, **8b10b\_rxerr.vcd**, and **mux.vcd**, you can use the import specifications in the following table:

Table 8-4: Import Specifications

File Name	Entity
<b>8b10b_dec.vcd</b>	Top   8b10b_dec:decode1
<b>8b10b_dec.vcd</b>	Top   8b10b_dec:decode2
<b>8b10b_dec.vcd</b>	Top   8b10b_dec:decode3
<b>8b10b_rxerr.vcd</b>	Top   8b10b_rxerr:err1
<b>8b10b_enc.vcd</b>	Top   8b10b_enc:encode1
<b>mux.vcd</b>	Top   mux:mux1

The resulting power analysis applies the simulation vectors in each file to the assigned entity. Simulation provides signal activities for the pins and for the outputs of functional blocks. If the inputs to an entity instance are input pins for the entire design, the simulation file associated with that instance does not

provide signal activities for the inputs of that instance. For example, an input to an entity such as `mux1` has its signal activity specified at the output of one of the decode entities.

## Multiple Simulations on the Same Entity

You can perform multiple simulations of an entire design or specific modules of a design. For example, in the process of verifying the top-level design, you can have three different simulation testbenches: one for normal operation, and two for corner cases. Each of these simulations produces a separate `.vcd`. In this case, apply the different `.vcd` file names to the same top-level entity, as shown in the following table.

**Table 8-5: Multiple Simulation File Names and Entities**

File Name	Entity
<code>normal.vcd</code>	Top
<code>corner1.vcd</code>	Top
<code>corner2.vcd</code>	Top

The resulting power analysis uses an arithmetic average of the signal activities calculated from each simulation file to obtain the final signal activities used. If a signal `err_out` has a toggle rate of zero transition per second in `normal.vcd`, 50 transitions per second in `corner1.vcd`, and 70 transitions per second in `corner2.vcd`, the final toggle rate in the power analysis is 40 transitions per second.

If you do not want the PowerPlay Power Analyzer to read information from multiple instances and take an arithmetic average of the signal activities, use a `.vcd` that includes only signals from the instance that you care about.

## Overlapping Simulations

You can perform a simulation on the entire design, and more exhaustive simulations on a submodule, such as `8b10b_rxerr`. The following table lists the import specification for overlapping simulations.

**Table 8-6: Overlapping Simulation Import Specifications**

File Name	Entity
<code>full_design.vcd</code>	Top
<code>error_cases.vcd</code>	Top   8b10b_rxerr:err1

In this case, the software uses signal activities from `error_cases.vcd` for all the nodes in the generated `.vcd` and uses signal activities from `full_design.vcd` for only those nodes that do not overlap with nodes in `error_cases.vcd`. In general, the more specific hierarchy (the most bottom-level module) derives signal activities for overlapping nodes.

## Partial Simulations

You can perform a simulation in which the entire simulation time is not applicable to signal-activity calculation. For example, if you run a simulation for 10,000 clock cycles and reset the chip for the first 2,000 clock cycles. If the PowerPlay Power Analyzer performs the signal-activity calculation over all 10,000 cycles, the toggle rates are only 80% of their steady state value (because the chip is in reset for the

first 20% of the simulation). In this case, you must specify the useful parts of the `.vcd` for power analysis. The **Limit VCD Period** option enables you to specify a start and end time when performing signal-activity calculations.

## Specifying Start and End Time when Performing Signal-Activity Calculations using the Limit VCD Period Option

To specify a start and end time when performing signal-activity calculations using the **Limit VCD period** option, follow these steps:

1. In the Quartus II software, on the Assignments menu, click **Settings**.
2. Under the Category list, click **PowerPlay Power Analyzer Settings**.
3. Turn on the **Use input file(s) to initialize toggle rates and static probabilities during power analysis** option.
4. Click **Add**.
5. In the **File name** and **Entity** fields, browse to the necessary files.
6. Under Simulation period, turn on **VCD file** and **Limit VCD period** options.
7. In the **Start time** and **End time** fields, specify the desired start and end time.
8. Click **OK**.

You can also use the following tcl or qsf assignment to specify `.vcd` files:

```
set_global_assignment -name POWER_INPUT_FILE_NAME "test.vcd" -section_id test.vcd
set_global_assignment -name POWER_INPUT_FILE_TYPE VCD -section_id test.vcd
set_global_assignment -name POWER_VCD_FILE_START_TIME "10 ns" -section_id test.vcd
set_global_assignment -name POWER_VCD_FILE_END_TIME "1000 ns" -section_id test.vcd
set_instance_assignment -name POWER_READ_INPUT_FILE test.vcd -to test_design
```

### Related Information

- [set\\_power\\_file\\_assignment](#)
- [Add/Edit Power Input File Dialog Box](#)

## Node Name Matching Considerations

Node name mismatches happen when you have `.vcd` applied to entities other than the top-level entity. In a modular design flow, the gate-level simulation files created in different Quartus II projects might not match their node names with the current Quartus II project.

For example, you may have a file named `8b10b_enc.vcd`, which the Quartus II software generates in a separate project called `8b10b_enc` while simulating the `8b10b` encoder. If you import the `.vcd` into another project called `Top`, you might encounter name mismatches when applying the `.vcd` to the `8b10b_enc` module in the `Top` project. This mismatch happens because the Quartus II software might name all the combinational nodes in the `8b10b_enc.vcd` differently than in the `Top` project.

You can avoid name mismatching with only RTL simulation data, in which register names do not change, or with an incremental compilation flow that preserves node names along with a gate-level simulation.

**Note:** To ensure accuracy, Altera recommends that you use an incremental compilation flow to preserve the node names of your design.

**Related Information****Quartus II Incremental Compilation for Hierarchical and Team-Based Design**

## Glitch Filtering

The PowerPlay Power Analyzer defines a glitch as two signal transitions so closely spaced in time that the pulse, or glitch, occurs faster than the logic and routing circuitry can respond. The output of a transport delay model simulator contains glitches for some signals. The logic and routing structures of the device form a low-pass filter that filters out glitches that are tens to hundreds of picoseconds long, depending on the device family.

Some third-party simulators use different models than the transport delay model as the default model. Different models cause differences in signal activity and power estimation. The inertial delay model, which is the ModelSim default model, filters out more glitches than the transport delay model and usually yields a lower power estimate.

**Note:** Altera recommends that you use the transport simulation model when using the Quartus II software glitch filtering support with third-party simulators. Simulation glitch filtering has little effect if you use the inertial simulation model.

Glitch filtering in a simulator can also filter a glitch on one logic element (LE) (or other circuit element) output from propagating to downstream circuit elements to ensure that the glitch does not affect simulated results. Glitch filtering prevents a glitch on one signal from producing non-physical glitches on all downstream logic, which can result in a signal toggle rate and a power estimate that are too high. Circuit elements in which every input transition produces an output transition, including multipliers and logic cells configured to implement XOR functions, are especially prone to glitches. Therefore, circuits with such functions can have power estimates that are too high when glitch filtering is not used.

**Note:** Altera recommends that you use the glitch filtering feature to obtain the most accurate power estimates. For .vcd files, the PowerPlay Power Analyzer flows support two levels of glitch filtering.

### Enabling First Level of Glitch Filtering

To enable the first level of glitch filtering in the Quartus II software for supported third-party simulators, follow these steps:

1. On the Assignments menu, click **Settings**.
2. In the **Category** list, select **Simulation under EDA Tool Settings**.
3. Select the **Tool name** to use for the simulation.
4. Turn on **Enable glitch filtering**.

### Enabling Second Level of Glitch Filtering

The second level of glitch filtering occurs while the PowerPlay Power Analyzer is reading the .vcd generated by a third-party simulator. To enable the second level of glitch filtering, follow these steps:

1. On the Assignments menu, click **Settings**.
2. In the **Category** list, select **PowerPlay Power Analyzer Settings**.
3. Under **Input File(s)**, turn on **Perform glitch filtering on VCD files**.

The .vcd file reader performs filtering complementary to the filtering performed during simulation and is often not as effective. While the .vcd file reader can remove glitches on logic blocks, the file reader cannot determine how a given glitch affects downstream logic and routing, and may eliminate the impact of the



glitch completely. Filtering the glitches during simulation avoids switching downstream routing and logic automatically.

**Note:** When running simulation for design verification (rather than to produce input to the PowerPlay Power Analyzer), Altera recommends that you turn off the glitch filtering option to produce the most rigorous and conservative simulation from a functionality viewpoint. When performing simulation to produce input for the PowerPlay Power Analyzer, Altera recommends that you turn on the glitch filtering to produce the most accurate power estimates.

## Node and Entity Assignments

You can assign toggle rates and static probabilities to individual nodes and entities in the design. These assignments have the highest priority, overriding data from all other signal-activity sources.

You must use the Assignment Editor or Tcl commands to create the **Power Toggle Rate** and **Power Static Probability** assignments. You can specify the power toggle rate as an absolute toggle rate in transitions per second using the **Power Toggle Rate** assignment, or you can use the **Power Toggle Rate Percentage** assignment to specify a toggle rate relative to the clock domain of the assigned node for a more specific assignment made in terms of hierarchy level.

**Note:** If you use the **Power Toggle Rate Percentage** assignment, and the node does not have a clock domain, the Quartus II software issues a warning and ignores the assignment.

Assigning toggle rates and static probabilities to individual nodes and entities is appropriate for signals in which you have knowledge of the signal or entity being analyzed. For example, if you know that a 100 MHz data bus or memory output produces data that is essentially random (uncorrelated in time), you can directly enter a 0.5 static probability and a toggle rate of 50 million transitions per second.

The PowerPlay Power Analyzer treats bidirectional I/O pins differently. The combinational input port and the output pad for a pin share the same name. However, those ports might not share the same signal activities. For reading signal-activity assignments, the PowerPlay Power Analyzer creates a distinct name `<node_name~output>` when configuring the bidirectional signal as an output and `<node_name~result>` when configuring the signal as an input. For example, if a design has a bidirectional pin named `MYPIN`, assignments for the combinational input use the name `MYPIN~result`, and the assignments for the output pad use the name `MYPIN~output`.

**Note:** When you create the logic assignment in the Assignment Editor, you cannot find the `MYPIN~result` and `MYPIN~output` node names in the Node Finder. Therefore, to create the logic assignment, you must manually enter the two differentiating node names to create the assignment for the input and output port of the bidirectional pin.

### Related Information

#### Constraining Designs

For more information about how to use the Assignment Editor in the Quartus II software, refer to this document.

## Timing Assignments to Clock Nodes

For clock nodes, the PowerPlay Power Analyzer uses timing requirements to derive the toggle rate when neither simulation data nor user-entered signal-activity data is available.  $f_{MAX}$  requirements specify full cycles per second, but each cycle represents a rising transition and a falling transition. For example, a clock  $f_{MAX}$  requirement of 100 MHz corresponds to 200 million transitions per second for the clock node.

## Default Toggle Rate Assignment

You can specify a default toggle rate for primary inputs and other nodes in your design. The PowerPlay Power Analyzer uses the default toggle rate when no other method specifies the signal-activity data.

The PowerPlay Power Analyzer specifies the toggle rate in absolute terms (transitions per second), or as a fraction of the clock rate in effect for each node. The toggle rate for a clock derives from the timing settings for the clock. For example, if the PowerPlay Power Analyzer specifies a clock with an  $f_{MAX}$  constraint of 100 MHz and a default relative toggle rate of 20%, nodes in this clock domain transition in 20% of the clock periods, or 20 million transitions occur per second. In some cases, the PowerPlay Power Analyzer cannot determine the clock domain for a node because either the PowerPlay Power Analyzer cannot determine a clock domain for the node, or the clock domain is ambiguous. For example, the PowerPlay Power Analyzer may not be able to determine a clock domain for a node if the user did not specify sufficient timing assignments. In these cases, the PowerPlay Power Analyzer substitutes and reports a toggle rate of zero.

## Vectorless Estimation

For some device families, the PowerPlay Power Analyzer automatically derives estimates for signal activity on nodes with no simulation or user-entered signal-activity data. Vectorless estimation statistically estimates the signal activity of a node based on the signal activities of nodes feeding that node, and on the actual logic function that the node implements. Vectorless estimation cannot derive signal activities for primary inputs. Vectorless estimation is accurate for combinational nodes, but not for registered nodes. Therefore, the PowerPlay Power Analyzer requires simulation data for at least the registered nodes and I/O nodes for accuracy.

The **PowerPlay Power Analyzer Settings** dialog box allows you to disable vectorless estimation. When turned on, vectorless estimation takes precedence over default toggle rates. Vectorless estimation does not override clock assignments.

To disable vectorless estimation, perform the following steps:

1. In the Quartus II software, on the Assignments menu, click **Settings**.
2. In the Category list, select **PowerPlay Power Analyzer Settings**.
3. Turn off the **Use vectorless estimation** option.

### Related Information

- [Performing Power Analysis with the PowerPlay Power Analyzer](#)

## Using the PowerPlay Power Analyzer

For flows that use the PowerPlay Power Analyzer, you must first synthesize your design, and then fit it to the target device. You must either provide timing assignments for all the clocks in your design, or use a simulation-based flow to generate activity data. You must specify the I/O standard on each device input and output and the board trace model on each output in your design.

### Related Information

- [Performing Power Analysis with the PowerPlay Power Analyzer](#)

## Common Analysis Flows

You can use the analysis flows in this section with the PowerPlay Power Analyzer. However, vectorless activity estimation is only available for some device families.

### Signal Activities from RTL (Functional) Simulation, Supplemented by Vectorless Estimation

In the functional simulation flow, simulation provides toggle rates and static probabilities for all pins and registers in your design. Vectorless estimation fills in the values for all the combinational nodes between pins and registers, giving good results. This flow usually provides a compilation time benefit when you use the third-party RTL simulator.

#### RTL Simulation Limitation

RTL simulation may not provide signal activities for all registers in the post-fitting netlist because synthesis loses some register names. For example, synthesis might automatically transform state machines and counters, thus changing the names of registers in those structures.

### Signal Activities from Vectorless Estimation and User-Supplied Input Pin Activities

The vectorless estimation flow provides a low level of accuracy, because vectorless estimation for registers is not entirely accurate.

### Signal Activities from User Defaults Only

The user defaults only flow provides the lowest degree of accuracy.

## Importance of .vcd

Altera recommends that you use a **.vcd** or a **.saf** generated by gate-level timing simulation for an accurate power estimation because gate-level timing simulation takes all the routing resources and the exact logic array resource usage into account.

### Generating a .vcd

In previous versions of the Quartus II software, you could use either the Quartus II simulator or an EDA simulator to perform your simulation. The Quartus II software no longer supports a built-in simulator, and you must use an EDA simulator to perform simulation. Use the **.vcd** as the input to the PowerPlay Power Analyzer to estimate power for your design.

To create a **.vcd** for your design, follow these steps:

1. On the Assignments menu, click **Settings**.
2. In the **Category** list, under **EDA Tool Settings**, click **Simulation**.
3. In the **Tool name** list, select your preferred EDA simulator.
4. In the **Format for output netlist** list, select **Verilog HDL**, or **SystemVerilog HDL**, or **VHDL**.
5. Turn on **Generate Value Change Dump (VCD) file script**.

This option turns on the **Map illegal HDL characters** and **Enable glitch filtering** options. The **Map illegal HDL characters** option ensures that all signals have legal names and that signal toggle rates are available later in the PowerPlay Power Analyzer. The **Enable glitch filtering** option directs the EDA Netlist Writer to perform glitch filtering when generating VHDL Output Files, Verilog Output Files, and the corresponding Standard Delay Format Output Files for use with other EDA simulation tools. This option is available regardless of whether or not you want to generate **.vcd** scripts.

**Note:** When performing simulation using ModelSim, the **+nospecify** option for the `vsim` command disables the **specify path delays and timing checks** option in ModelSim. By enabling glitch filtering on the **Simulation** page, the simulation models include specified path delays. Thus, ModelSim might fail to simulate a design if you enabled glitch filtering and specified the **+nospecify** option. Altera recommends that you remove the **+nospecify** option from the ModelSim `vsim` command to ensure accurate simulation for power estimation.

6. Click **Script Settings**. Select the signals that you want to output to the **.vcd**.  
With **All signals** selected, the generated script instructs the third-party simulator to write all connected output signals to the **.vcd**. With **All signals except combinational lcell outputs** selected, the generated script tells the third-party simulator to write all connected output signals to the **.vcd**, except logic cell combinational outputs.  
  
**Note:** The file can become extremely large if you write all output signals to the file because the file size depends on the number of output signals being monitored and the number of transitions that occur.
7. Click **OK**.
8. In the **Design instance name** box, type a name for your testbench.
9. Compile your design with the Quartus II software and generate the necessary EDA netlist and script that instructs the third-party simulator to generate a **.vcd**.
10. Perform a simulation with the third-party EDA simulation tool. Call the generated script in the simulation tool before running the simulation. The simulation tool generates the **.vcd** and places it in the project directory.

#### Related Information

[Simulation Results](#) on page 8-9

[Glitch Filtering](#) on page 8-14

#### Section I. Simulation

### Generating a .vcd from ModelSim Software

To generate a **.vcd** with the ModelSim software, follow these steps:

1. In the Quartus II software, on the Assignments menu, click **Settings**.
2. In the **Category** list, under **EDA Tool Settings**, click **Simulation**.
3. In the **Tool name** list, select your preferred EDA simulator.
4. In the **Format for output netlist** list, select **Verilog HDL**, or **SystemVerilog HDL**, or **VHDL**.
5. Turn on **Generate Value Change Dump (VCD) file script**.
6. To generate the **.vcd**, perform a full compilation.
7. In the ModelSim software, compile the files necessary for simulation.
8. Load your design by clicking **Start Simulation** on the Tools menu, or use the `vsim` command.
9. Use the **.vcd** script created in [step 6](#) using the following command:  

```
source <design>_dump_all_vcd_nodes.tcl
```
10. Run the simulation (for example, run 2000ns or run -all).
11. Quit the simulation using the `quit -sim` command, if required.
12. Exit the ModelSim software.  
If you do not exit the software, the ModelSim software might end the writing process of the **.vcd** improperly, resulting in a corrupt **.vcd**.

## Generating a .vcd from Full Post-Fit Netlist (Zero Delay) Simulation

To successfully generate a .vcd from the full post-fit Netlist (zero delay) simulation, follow these steps:

1. Compile your design in the Quartus II software to generate the Netlist <project\_name>.vo.
2. In <project\_name>.vo, search for the include statement for <project\_name>.sdo, comment the statement out, and save the file.

Altera recommends that you use the Standard Delay Format Output File (.sdo) for gate-level timing simulation. The .sdo contains the delay information of each architecture primitive and routing element in your design; however, you must exclude the .sdo for zero delay simulation.

3. Generate a .vcd for power estimation by performing the steps in [Generating a .vcd](#) on page 8-17.

### Related Information

#### [Section I. Simulation](#)

## PowerPlay Power Analyzer Compilation Report

The following table list the items in the Compilation Report of the PowerPlay Power Analyzer section.

Section	Description
Summary	The Summary section of the report shows the estimated total thermal power consumption of your design. This includes dynamic, static, and I/O thermal power consumption. The I/O thermal power includes the total I/O power drawn from the $V_{CCIO}$ and $V_{CCPD}$ power supplies and the power drawn from $V_{CCINT}$ in the I/O subsystem including I/O buffers and I/O registers. The report also includes a confidence metric that reflects the overall quality of the data sources for the signal activities. For example, a <b>Low</b> power estimation confidence value reflects that you have provided insufficient toggle rate data, or most of the signal-activity information used for power estimation is from default or vectorless estimation settings. For more information about the input data, refer to the PowerPlay Power Analyzer Confidence Metric report.
Settings	The Settings section of the report shows the PowerPlay Power Analyzer settings information of your design, including the default input toggle rates, operating conditions, and other relevant setting information.
Simulation Files Read	The Simulation Files Read section of the report lists the simulation output file that the .vcd used for power estimation. This section also includes the file ID, file type, entity, VCD start time, VCD end time, the unknown percentage, and the toggle percentage. The unknown percentage indicates the portion of the design module unused by the simulation vectors.
Operating Conditions Used	The Operating Conditions Used section of the report shows device characteristics, voltages, temperature, and cooling solution, if any, during the power estimation. This section also shows the entered junction temperature or auto-computed junction temperature during the power analysis.

Section	Description
Thermal Power Dissipated by Block	<p>The Thermal Power Dissipated by Block section of the report shows estimated thermal dynamic power and thermal static power consumption categorized by atoms. This information provides you with estimated power consumption for each atom in your design.</p> <p>By default, this section does not contain any data, but you can turn on the report with the <b>Write power dissipation by block to report file</b> option on the <b>PowerPlay Power Analyzer Settings</b> page.</p>
Thermal Power Dissipation by Block Type (Device Resource Type)	<p>This Thermal Power Dissipation by Block Type (Device Resource Type) section of the report shows the estimated thermal dynamic power and thermal static power consumption categorized by block types. This information is further categorized by estimated dynamic and static power and provides an average toggle rate by block type. Thermal power is the power dissipated as heat from the FPGA device.</p>
Thermal Power Dissipation by Hierarchy	<p>This Thermal Power Dissipation by Hierarchy section of the report shows estimated thermal dynamic power and thermal static power consumption categorized by design hierarchy. This information is further categorized by the dynamic and static power that was used by the blocks and routing in that hierarchy. This information is useful when locating modules with high power consumption in your design.</p>
Core Dynamic Thermal Power Dissipation by Clock Domain	<p>The Core Dynamic Thermal Power Dissipation by Clock Domain section of the report shows the estimated total core dynamic power dissipation by each clock domain, which provides designs with estimated power consumption for each clock domain in the design. If the clock frequency for a domain is unspecified by a constraint, the clock frequency is listed as “unspecified.” For all the combinational logic, the clock domain is listed as no clock with zero MHz.</p>

Section	Description
<p>Current Drawn from Voltage Supplies</p>	<p>The Current Drawn from Voltage Supplies section of the report lists the current drawn from each voltage supply. The <math>V_{CCIO}</math> and <math>V_{CCPD}</math> voltage supplies are further categorized by I/O bank and by voltage. This section also lists the minimum safe power supply size (current supply ability) for each supply voltage. Minimum current requirement can be higher than user mode current requirement in cases in which the supply has a specific power up current requirement that goes beyond user mode requirement, such as the <math>V_{CCPD}</math> power rail in Stratix III and Stratix IV devices, and the <math>V_{CCIO}</math> power rail in Stratix IV devices.</p> <p>The I/O thermal power dissipation on the summary page does not correlate directly to the power drawn from the <math>V_{CCIO}</math> and <math>V_{CCPD}</math> voltage supplies listed in this report. This is because the I/O thermal power dissipation value also includes portions of the <math>V_{CCINT}</math> power, such as the I/O element (IOE) registers, which are modeled as I/O power, but do not draw from the <math>V_{CCIO}</math> and <math>V_{CCPD}</math> supplies.</p> <p>The reported current drawn from the I/O Voltage Supplies (ICCI0 and ICCPD) as reported in the PowerPlay Power Analyzer report includes any current drawn through the I/O into off-chip termination resistors. This can result in ICCI0 and ICCPD values that are higher than the reported I/O thermal power, because this off-chip current dissipates as heat elsewhere and does not factor in the calculation of device temperature. Therefore, total I/O thermal power does not equal the sum of current drawn from each <math>V_{CCIO}</math> and <math>V_{CCPD}</math> supply multiplied by <math>V_{CCIO}</math> and <math>V_{CCPD}</math> voltage.</p> <p>For SoC devices or for Arria V SoC and Cyclone V SoC devices, there is no standalone ICC_AUX_SHARED current drawn information. The ICC_AUX_SHARED is reported together with ICC_AUX.</p>
<p>Confidence Metric Details</p>	<p>The Confidence Metric is defined in terms of the total weight of signal activity data sources for both combinational and registered signals. Each signal has two data sources allocated to it; a toggle rate source and a static probability source.</p> <p>The Confidence Metric Details section also indicates the quality of the signal toggle rate data to compute a power estimate. The confidence metric is low if the signal toggle rate data comes from poor predictors of real signal toggle rates in the device during an operation. Toggle rate data that comes from simulation, user-entered assignments on specific signals or entities are reliable. Toggle rate data from default toggle rates (for example, 12.5% of the clock period) or vectorless estimation are relatively inaccurate. This section gives an overall confidence rating in the toggle rate data, from low to high. This section also summarizes how many pins, registers, and combinational nodes obtained their toggle rates from each of simulation, user entry, vectorless estimation, or default toggle rate estimations. This detailed information helps you understand how to increase the confidence metric, letting you determine your own confidence in the toggle rate data.</p>

Section	Description
Signal Activities	<p>The Signal Activities section lists toggle rates and static probabilities assumed by power analysis for all signals with fan-out and pins. This section also lists the signal type (pin, registered, or combinational) and the data source for the toggle rate and static probability. By default, this section does not contain any data, but you can turn on the report with the <b>Write signal activities to report file</b> option on the <b>PowerPlay Power Analyzer Settings</b> page.</p> <p>Altera recommends that you keep the <b>Write signal activities to report file</b> option turned off for a large design because of the large number of signals present. You can use the Assignment Editor to specify that activities for individual nodes or entities are reported by assigning an on value to those nodes for the <b>Power Report Signal Activities</b> assignment.</p>
Messages	The Messages section lists the messages that the Quartus II software generates during the analysis.

## Scripting Support

You can run procedures and create settings described in this chapter in a Tcl script. You can also run some procedures at a command prompt. For more information about scripting command options, refer to the Quartus II Command-Line and Tcl API Help browser. To run the Help browser, type the following command at the command prompt:

```
quartus_sh --qhelp
```

### Related Information

- [Tcl Scripting](#)
- [API Functions for Tcl](#)
- [Quartus II Settings File Reference Manual](#)
- [Command-Line Scripting](#)

## Running the PowerPlay Power Analyzer from the Command-Line

The executable to run the PowerPlay Power Analyzer is `quartus_pow`. For a complete listing of all command-line options supported by `quartus_pow`, type the following command at a system command prompt:

```
quartus_pow --help
```

or-

```
quartus_sh --qhelp
```

The following lists the examples of using the `quartus_pow` executable. Type the command listed in the following section at a system command prompt. These examples assume that operations are performed on Quartus II project called *sample*.



**To instruct the PowerPlay Power Analyzer to generate a PowerPlay EPE File:**

```
quartus_pow sample --output_epe=sample.csv
```

**To instruct the PowerPlay Power Analyzer to generate a PowerPlay EPE File without performing the power estimate:**

```
quartus_pow sample --output_epe=sample.csv --estimate_power=off
```

**To instruct the PowerPlay Power Analyzer to use a .vcd as input (sample.vcd):**

```
quartus_pow sample --input_vcd=sample.vcd
```

**To instruct the PowerPlay Power Analyzer to use two .vcd files as input files (sample1.vcd and sample2.vcd), perform glitch filtering on the .vcd and use a default input I/O toggle rate of 10,000 transitions per second:**

```
quartus_pow sample --input_vcd=sample1.vcd --input_vcd=sample2.vcd \  
--vcd_filter_glitches=on --\  
default_input_io_toggle_rate=10000transitions/s
```

**To instruct the PowerPlay Power Analyzer to not use an input file, a default input I/O toggle rate of 60%, no vectorless estimation, and a default toggle rate of 20% on all remaining signals:**

```
quartus_pow sample --no_input_file --default_input_io_toggle_rate=60% \  
--use_vectorless_estimation=off --default_toggle_rate=20%
```

**Note:** No command–line options are available to specify the information found on the **PowerPlay Power Analyzer Settings Operating Conditions** page. Use the Quartus II GUI to specify these options.

The `quartus_pow` executable creates a report file, `<revision name>.pow.rpt`. You can locate the report file in the main project directory. The report file contains the same information in [PowerPlay Power Analyzer Compilation Report](#) on page 8-19.

## Document Revision History

The following table lists the revision history for this chapter.

Date	Version	Changes
2014.12.15	14.1.0	<ul style="list-style-type: none"> <li>Removed Signal Activities from Full Post-Fit Netlist (Timing) Simulation and Signal Activities from Full Post-Fit Netlist (Zero Delay) Simulation sections as these are no longer supported.</li> <li>Updated location of Fitter Settings, Analysis &amp; Synthesis Settings, and Physical Synthesis Optimizations to Compiler Settings.</li> </ul>
2014.08.18	14.0a10.0	Updated "Current Drawn from Voltage Supplies" to clarify that for SoC devices or for Arria V SoC and Cyclone V SoC devices, there is no standalone ICC_AUX_SHARED current drawn information. The ICC_AUX_SHARED is reported together with ICC_AUX.

Date	Version	Changes
November 2012	12.1.0	<ul style="list-style-type: none"> <li>Updated “Types of Power Analyses” on page 8–2, and “Confidence Metric Details” on page 8–23.</li> <li>Added “Importance of .vcd” on page 8–20, and “Avoiding Power Estimation and Hardware Measurement Mismatch” on page 8–24</li> </ul>
June 2012	12.0.0	<ul style="list-style-type: none"> <li>Updated “Current Drawn from Voltage Supplies” on page 8–22.</li> <li>Added “Using the HPS Power Calculator” on page 8–7.</li> </ul>
November 2011	10.1.1	<ul style="list-style-type: none"> <li>Template update.</li> <li>Minor editorial updates.</li> </ul>
December 2010	10.1.0	<ul style="list-style-type: none"> <li>Added links to Quartus II Help, removed redundant material.</li> <li>Moved “Creating PowerPlay EPE Spreadsheets” to page 8–6.</li> <li>Minor edits.</li> </ul>
July 2010	10.0.0	<ul style="list-style-type: none"> <li>Removed references to the Quartus II Simulator.</li> <li>Updated Table 8–1 on page 8–6, Table 8–2 on page 8–13, and Table 8–3 on page 8–14.</li> <li>Updated Figure 8–3 on page 8–9, Figure 8–4 on page 8–10, and Figure 8–5 on page 8–12.</li> </ul>
November 2009	9.1.0	<ul style="list-style-type: none"> <li>Updated “Creating PowerPlay EPE Spreadsheets” on page 8–6 and “Simulation Results” on page 8–10.</li> <li>Added “Signal Activities from Full Post-Fit Netlist (Zero Delay) Simulation” on page 8–19 and “Generating a .vcd from Full Post-Fit Netlist (Zero Delay) Simulation” on page 8–21.</li> <li>Minor changes to “Generating a .vcd from ModelSim Software” on page 8–21.</li> <li>Updated Figure 11–8 on page 11–24.</li> </ul>
March 2009	9.0.0	<ul style="list-style-type: none"> <li>This chapter was chapter 11 in version 8.1.</li> <li>Removed Figures 11-10, 11-11, 11-13, 11-14, and 11-17 from 8.1 version.</li> </ul>
November 2008	8.1.0	<ul style="list-style-type: none"> <li>Updated for the Quartus II software version 8.1.</li> <li>Replaced Figure 11-3.</li> <li>Replaced Figure 11-14.</li> </ul>
May 2008	8.0.0	<ul style="list-style-type: none"> <li>Updated Figure 11–5.</li> <li>Updated “Types of Power Analyses” on page 11–5.</li> <li>Updated “Operating Conditions” on page 11–9.</li> <li>Updated “PowerPlay Power Analyzer Compilation Report” on page 11–31.</li> <li>Updated “Current Drawn from Voltage Supplies” on page 11–32.</li> </ul>

2014.06.30

QI15V3



Subscribe



Send Feedback

## About Altera System Debugging Tools

The Altera<sup>®</sup> system debugging tools help you verify your FPGA designs. As your product requirements continue to increase in complexity, the time you spend on design verification continues to rise. This manual provides a quick overview of the tools available in the system debugging suite and discusses the criteria for selecting the best tool for your design.

## System Debugging Tools Portfolio

The Quartus<sup>®</sup> II software provides a portfolio of system design debugging tools for real-time verification of your design. Each tool in the system debugging portfolio uses a combination of available memory, logic, and routing resources to assist in the debugging process. The tools provide visibility by routing (or “tapping”) signals in your design to debugging logic. The debugging logic is then compiled with your design and downloaded into the FPGA or CPLD for analysis. Because different designs can have different constraints and requirements, such as the number of spare pins available or the amount of logic or memory resources remaining in the physical device, you can choose a tool from the available debugging tools that matches the specific requirements for your design.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered

## System Debugging Tools Comparison

Table 9-1: Debugging Tools Portfolio

Tool	Description	Typical Usage
<p><b>System Console</b></p>	<p>Uses a Tcl interpreter to communicate with hardware modules instantiated in your design. You can use it with the Transceiver Toolkit to monitor or debug your design.</p> <p>System Console provides real-time in-system debugging capabilities. Using System Console, you can read from and write to Memory Mapped components in our system without the help of a processor or additional software.</p> <p>System Console uses Tcl as the fundamental infrastructure which means you can source scripts, set variables, write procedures, and take advantage of all the features of the Tcl scripting language.</p>	<p>You need to perform system-level debugging. For example, if you have an Avalon-MM slave or Avalon-ST interfaces, you can debug your design at a transaction level. The tool supports JTAG connectivity and TCP/IP connectivity to the target FPGA you wish to debug.</p>
<p><b>Transceiver Toolkit</b></p>	<p>The Transceiver Toolkit allows you to test and tune transceiver link signal quality. You can use a combination of bit error rate (BER), bathtub curve, and eye contour graphs as quality metrics. Auto Sweeping of physical medium attachment (PMA) settings allows you to quickly find an optimal solution.</p>	<p>You need to debug or optimize signal integrity of your board layout even before the actual design to be run on the FPGA is ready.</p>
<p><b>SignalTap<sup>®</sup> II Logic Analyzer</b></p>	<p>This logic analyzer uses FPGA resources to sample test nodes and outputs the information to the Quartus II software for display and analysis.</p>	<p>You have spare on-chip memory and you want functional verification of your design running in hardware.</p>

Tool	Description	Typical Usage
<b>SignalProbe</b>	This tool incrementally routes internal signals to I/O pins while preserving results from your last place-and-routed design.	You have spare I/O pins and you would like to check the operation of a small set of control pins using either an external logic analyzer or an oscilloscope.
<b>Logic Analyzer Interface (LAI)</b>	This tool multiplexes a larger set of signals to a smaller number of spare I/O pins. LAI allows you to select which signals are switched onto the I/O pins over a JTAG connection.	You have limited on-chip memory, and have a large set of internal data buses that you would like to verify using an external logic analyzer. Logic analyzer vendors, such as Tektronics and Agilent, provide integration with the tool to improve the usability of the tool.
<b>In-System Sources and Probes</b>	This tool provides an easy way to drive and sample logic values to and from internal nodes using the JTAG interface.	You want to prototype a front panel with virtual buttons for your FPGA design.
<b>In-System Memory Content Editor</b>	This tool displays and allows you to edit on-chip memory.	You would like to view and edit the contents of on-chip memory that is not connected to a Nios II processor. You can also use the tool when you do not want to have a Nios II debug core in your system.
<b>Virtual JTAG Interface</b>	This megafunction allows you to communicate with the JTAG interface so that you can develop your own custom applications.	You have custom signals in your design that you want to be able to communicate with.

## Altera JTAG Interface (AJI)

With the exception of SignalProbe, each of the on-chip debugging tools uses the JTAG port to control and read back data from debugging logic and signals under test. System Console uses JTAG and other interfaces as well. The JTAG resource is shared among all of the on-chip debugging tools.

## Required Arbitration Logic

For all system debugging tools except System Console, the Quartus II software compiles logic into your design automatically to distinguish between data and control information and each of the debugging logic blocks, when the JTAG resource is required. This arbitration logic, also known as the System-Level Debugging (SLD) infrastructure, is shown in the design hierarchy of your compiled project as

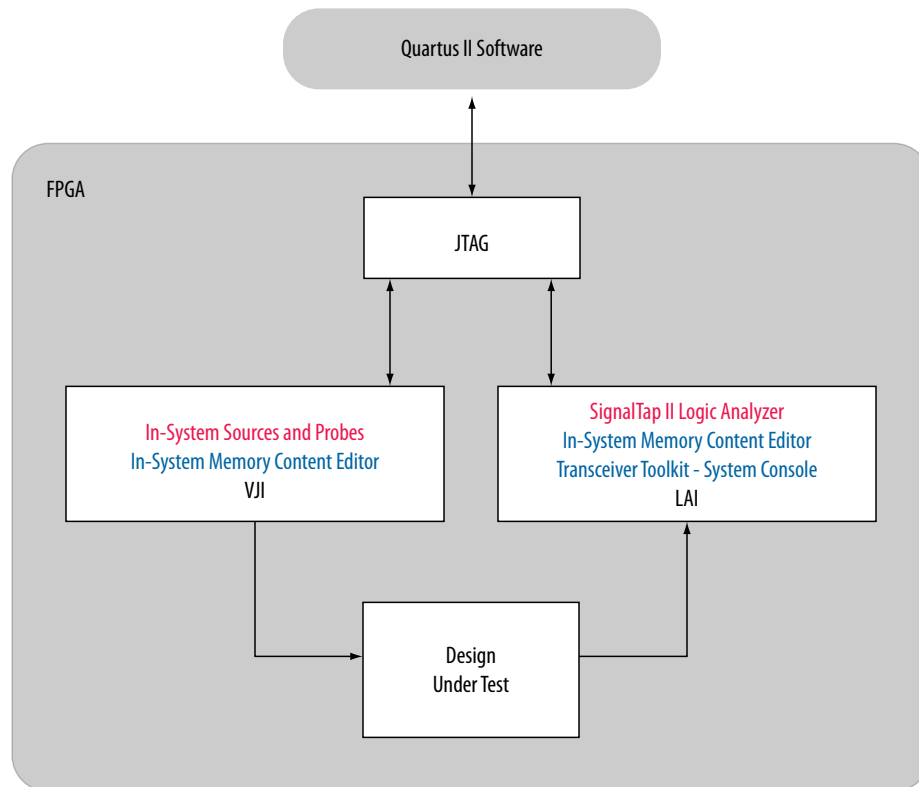
**sld\_hub:sld\_hub\_inst.** The SLD logic allows you to instantiate multiple debugging blocks into your design and run them simultaneously. For System Console, you must explicitly insert debug IP cores into your design to enable debugging.

## Debugging Ecosystem

To maximize debugging closure, the Quartus II software allows you to use a combination of the debugging tools in tandem to fully exercise and analyze the logic under test. All of the tools have basic analysis features built in; that is, all of the tools enable you to read back information collected from the design nodes that are connected to the debugging logic. Out of the set of debugging tools, the SignalTap II Logic Analyzer, the LAI, and the SignalProbe feature are general purpose debugging tools optimized for probing signals in your register transfer level (RTL) netlist. In-System Sources and Probes, the Virtual JTAG Interface, System Console, Transceiver Toolkit, and In-System Memory Content Editor, allow you to read back data from the debugging breakpoints, and to input values into your design during runtime.

Taken together, the set of on-chip debugging tools form a debugging ecosystem. The set of tools can generate a stimulus to and solicit a response from the logic under test, providing a complete debugging solution.

Figure 9-1: Debugging Ecosystem



## About Analysis Tools for RTL Nodes

The SignalTap II Logic Analyzer, SignalProbe, and LAI are designed specifically for probing and debugging RTL signals at system speed. They are general-purpose analysis tools that enable you to tap and analyze any routable node from the FPGA or CPLD. If you have spare logic and memory resources, the

SignalTap II Logic Analyzer is useful for providing fast functional verification of your design running on actual hardware.

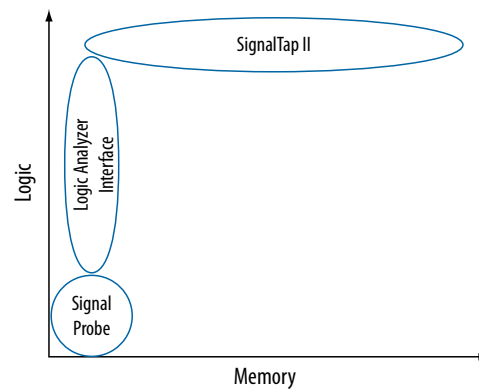
Conversely, if logic and memory resources are tight and you require the large sample depths associated with external logic analyzers, both the LAI and the SignalProbe make it easy to view internal design signals using external equipment.

**Note:** The SignalTap II Logic Analyzer is not supported on CPLDs, because there are no memory resources available on these devices.

## Resource Usage

The most important selection criteria for these three tools are the available resources remaining on your device after implementing your design and the number of spare pins available. You should evaluate your preferred debugging option early on in the design planning process to ensure that your board, your Quartus II project, and your design are all set up to support the appropriate options. Planning early can reduce time spent during debugging and eliminate the necessary late changes to accommodate your preferred debugging methodologies.

**Figure 9-2: Resource Usage per Debugging Tool**



## Overhead Logic

Any debugging tool that requires the use of a JTAG connection requires the SLD infrastructure logic, for communication with the JTAG interface and arbitration between any instantiated debugging modules. This overhead logic uses around 200 logic elements (LEs), a small fraction of the resources available in any of the supported devices. The overhead logic is shared between all available debugging modules in your design. Both the SignalTap II Logic Analyzer and the LAI use a JTAG connection.

### **For SignalProbe**

SignalProbe requires very few on-chip resources. Because it requires no JTAG connection, SignalProbe uses no logic or memory resources. SignalProbe uses only routing resources to route an internal signal to a debugging test point.

### **For Logic Analyzer Interface**

The LAI requires a small amount of logic to implement the multiplexing function between the signals under test, in addition to the SLD infrastructure logic. Because no data samples are stored on the chip, the LAI uses no memory resources.

### **For SignalTap II**


The SignalTap II Logic Analyzer requires both logic and memory resources. The number of logic resources used depends on the number of signals tapped and the complexity of the trigger logic. However,

the amount of logic resources that the SignalTap II Logic Analyzer uses is typically a small percentage of most designs. A baseline configuration consisting of the SLD arbitration logic and a single node with basic triggering logic contains approximately 300 to 400 Logic Elements (LEs). Each additional node you add to the baseline configuration adds about 11 LEs. Compared with logic resources, memory resources are a more important factor to consider for your design. Memory usage can be significant and depends on how you configure your SignalTap II Logic Analyzer instance to capture data and the sample depth that your design requires for debugging. For the SignalTap II Logic Analyzer, there is the added benefit of requiring no external equipment, as all of the triggering logic and storage is on the chip.

### Resource Estimation

The resource estimation feature for the SignalTap II Logic Analyzer and the LAI allows you to quickly judge if enough on-chip resources are available before compiling the tool with your design.

**Figure 9-3: Resource Estimator**



Instance	Status	LEs: 652	Memory: 524288	M512/MLAB: 0/94	M4K/M9K: 128/60
auto_signaltap_0	Not running	652 cells	524288 bits	0 blocks	Can't Fit 128 blocks

## Pin Usage

### For SignalProbe

The ratio of the number of pins used to the number of signals tapped for the SignalProbe feature is one-to-one. Because this feature can consume free pins quickly, a typical application for this feature is routing control signals to spare pins for debugging.

### For Logic Analyzer Interface

The ratio of the number of pins used to the number of signals tapped for the LAI is many-to-one. It can map up to 256 signals to each debugging pin, depending on available routing resources. The control of the active signals that are mapped to the spare I/O pins is performed via the JTAG port. The LAI is ideal for routing data buses to a set of test pins for analysis.

### For SignalTap II

Other than the JTAG test pins, the SignalTap II Logic Analyzer uses no additional pins. All data is buffered using on-chip memory and communicated to the SignalTap II Logic Analyzer GUI via the JTAG test port.

## Usability Enhancements

The SignalTap II Logic Analyzer, SignalProbe, and LAI tools can be added to your existing design with minimal effects. With the node finder, you can find signals to route to a debugging module without making any changes to your HDL files. SignalProbe inserts signals directly from your post-fit database. The SignalTap II Logic Analyzer and LAI support inserting signals from both pre-synthesis and post-fit netlists.

### Incremental Compilation

All three tools allow you to find and configure your debugging setup quickly. In addition, the Quartus II incremental compilation feature and the Quartus II incremental routing feature allow for a fast turnaround time for your programming file, increasing productivity and enabling fast debugging closure.



Both the LAI and SignalTap II Logic Analyzer support incremental compilation. With incremental compilation, you can add a SignalTap II Logic Analyzer instance or an LAI instance incrementally into your placed-and-routed design. This has the benefit of both preserving your timing and area optimizations from your existing design, and decreasing the overall compilation time when any changes are necessary during the debugging process. With incremental compilation, you can save up to 70% compile time of a full compilation.

### Incremental Routing

SignalProbe uses the incremental routing feature. The incremental routing feature runs only the Fitter stage of the compilation. This leaves your compiled design untouched, except for the newly routed node or nodes. With SignalProbe, you can save as much as 90% compile time of a full compilation.

### Automation Via Scripting

As another productivity enhancement, all tools in the on-chip debugging tool set support scripting via the `quartus_stp` Tcl package. For the SignalTap II Logic Analyzer and the LAI, scripting enables user-defined automation for data collection while debugging in the lab. The System Console includes a full Tcl interpreter for scripting.

### Remote Debugging

You can perform remote debugging of your system with the Quartus II software via the System Console. This feature allows you to debug equipment deployed in the field through an existing TCP/IP connection.

There are two Application Notes available to assist you.

- Application Note 624 describes how to set up your NIOS II system to use the System Console to perform remote debugging.
- Application Note 693 describes how to set up your Altera SoC to use the SLD tools to perform remote debugging.

### Related Information

- [Application Note 624: Debugging with System Console over TCP/IP](#)
- [Application Note 693: Remote Debugging over TCP/IP for Altera SoC](#)

## Suggested On-Chip Debugging Tools for Common Debugging Features

Table 9-2: Tools for Common Debugging Features <sup>(1)</sup>

Feature	SignalProbe	Logic Analyzer Interface (LAI)	SignalTap II Logic Analyzer	Description
<b>Large Sample Depth</b>	N/A	X	—	An external logic analyzer used with the LAI has a bigger buffer to store more captured data than the SignalTap II Logic Analyzer. No data is captured or stored with SignalProbe.
<b>Ease in Debugging Timing Issue</b>	X	X	—	External equipment, such as oscilloscopes and mixed signal oscilloscopes (MSOs), can be used with either LAI or SignalProbe. When used with the LAI, external equipment provides you with access to timing mode, which allows you to debug combined streams of data.

Feature	SignalProbe	Logic Analyzer Interface (LAI)	SignalTap II Logic Analyzer	Description
<b>Minimal Effect on Logic Design</b>	X	X <sup>(2)</sup>	X <sup>(2)</sup>	The LAI adds minimal logic to a design, requiring fewer device resources. The SignalTap II Logic Analyzer has little effect on the design, because it is set as a separate design partition. SignalProbe incrementally routes nodes to pins, not affecting the design at all.
<b>Short Compile and Recompile Time</b>	X	X <sup>(2)</sup>	X <sup>(2)</sup>	SignalProbe attaches incrementally routed signals to previously reserved pins, requiring very little recompilation time to make changes to source signal selections. The SignalTap II Logic Analyzer and the LAI can take advantage of incremental compilation to refit their own design partitions to decrease recompilation time.

Feature	SignalProbe	Logic Analyzer Interface (LAI)	SignalTap II Logic Analyzer	Description
<b>Triggering Capability</b>	N/A	N/A	X	The SignalTap II Logic Analyzer offers triggering capabilities that are comparable to commercial logic analyzers.
<b>I/O Usage</b>	—	—	X	No additional output pins are required with the SignalTap II Logic Analyzer. Both the LAI and SignalProbe require I/O pin assignments.
<b>Acquisition Speed</b>	N/A	—	X	The SignalTap II Logic Analyzer can acquire data at speeds of over 200 MHz. The same acquisition speeds are obtainable with an external logic analyzer used with the LAI, but might be limited by signal integrity issues.
<b>No JTAG Connection Required</b>	X	—	X	A FPGA design with the LAI requires an active JTAG connection to a host running the Quartus II software. SignalProbe and SignalTap II do not require a host for debugging purposes.

Feature	SignalProbe	Logic Analyzer Interface (LAI)	SignalTap II Logic Analyzer	Description
<b>No External Equipment Required</b>	—	—	X	The SignalTap II Logic Analyzer logic is completely internal to the programmed FPGA device. No extra equipment is required other than a JTAG connection from a host running the Quartus II software or the stand-alone SignalTap II Logic Analyzer software. SignalProbe and the LAI require the use of external debugging equipment, such as multimeters, oscilloscopes, or logic analyzers.

Notes to Table:

1. • X indicates the recommended tools for the feature.
  - — indicates that while the tool is available for that feature, that tool might not give the best results.
  - N/A indicates that the feature is not applicable for the selected tool.
2. When used with incremental compilation.

## About Stimulus-Capable Tools

The In-System Memory Content Editor, In-System Sources and Probes, and Virtual JTAG interface enable you to use the JTAG interface as a general-purpose communication port. Though all three tools can be used to achieve the same results, there are some considerations that make one tool easier to use in certain applications than others. In-System Sources and Probes is ideal for toggling control signals. The In-System Memory Content Editor is useful for inputting large sets of test data. Finally, the Virtual JTAG interface is well suited for more advanced users who want to develop their own customized JTAG solution.

System Console provides system-level debugging at a transaction level, such as with Avalon-MM slave or Avalon-ST interfaces. You can communicate to a chip through JTAG, and TCP/IP protocols. System

Console uses a Tcl interpreter to communicate with hardware modules that you have instantiated into your design.

## In-System Sources and Probes

In-System Sources and Probes is an easy way to access JTAG resources to both read and write to your design. You can start by instantiating a megafunction into your HDL code. The megafunction contains source ports and probe ports for driving values into and sampling values from the signals that are connected to the ports, respectively. Transaction details of the JTAG interface are abstracted away by the megafunction. During runtime, a GUI displays each source and probe port by instance and allows you to read from each probe port and drive to each source port. The GUI makes this tool ideal for toggling a set of control signals during the debugging process.

### Push Button Functionality

A good application of In-System Sources and Probes is to use the GUI as a replacement for the push buttons and LEDs used during the development phase of a project. Furthermore, In-System Sources and Probes supports a set of scripting commands for reading and writing using `quartus_stp`. When used with the Tk toolkit, you can build your own graphical interfaces. This feature is ideal for building a virtual front panel during the prototyping phase of the design.

## In-System Memory Content Editor

The In-System Memory Content Editor allows you to quickly view and modify memory content either through a GUI interface or through Tcl scripting commands. The In-System Memory Content Editor works by turning single-port RAM blocks into dual-port RAM blocks. One port is connected to your clock domain and data signals, and the other port is connected to the JTAG clock and data signals for editing or viewing.

### Generate Test Vectors

Because you can modify a large set of data easily, a useful application for the In-System Memory Content Editor is to generate test vectors for your design. For example, you can instantiate a free memory block, connect the output ports to the logic under test (using the same clock as your logic under test on the system side), and create the glue logic for the address generation and control of the memory. At runtime, you can modify the contents of the memory using either a script or the In-System Memory Content Editor GUI and perform a burst transaction of the data contents in the modified RAM block synchronous to the logic being tested.

## Virtual JTAG Interface Megafunction

The Virtual JTAG Interface megafunction provides the finest level of granularity for manipulating the JTAG resource. This megafunction allows you to build your own JTAG scan chain by exposing all of the JTAG control signals and configuring your JTAG Instruction Registers (IRs) and JTAG Data Registers (DRs). During runtime, you control the IR/DR chain through a Tcl API, or with System Console. This feature is meant for users who have a thorough understanding of the JTAG interface and want precise control over the number and type of resources used.

## System Console

System Console is a framework that you can launch from the Quartus II software to start services for performing various debugging tasks. System Console provides you with Tcl scripts and a GUI to access the Qsys system integration tool to perform low-level hardware debugging of your design, as well as identify a module by its path, and open and close a connection to a Qsys module. You can access your

design at a system level for purposes of loading, unloading, and transferring designs to multiple devices. Also, System Console supports the Tk toolkit for building graphical interfaces.

### Test Signal Integrity

System Console also allows you to access commands that allow you to control how you generate test patterns, as well as verify the accuracy of data generated by test patterns. You can use JTAG debug commands in System Console to verify the functionality and signal integrity of your JTAG chain. You can test clock and reset signals.

### Board Bring-Up and Verification

You can use System Console to access programmable logic devices on your development board, perform board bring-up, and perform verification. You can also access software running on a Nios II or Altera SoC processor, as well as access modules that produce or consume a stream of bytes.

### Test Link Signal Integrity with Transceiver Toolkit

Transceiver Toolkit runs from the System Console framework, and allows you to run automatic tests of your transceiver links for debugging and optimizing your transceiver designs. You can use the Transceiver Toolkit GUI to set up channel links in your transceiver devices, and then automatically run EyeQ and Auto Sweep testing to view a graphical representation of your test data.

## Document Revision History

Table 9-3: Document Revision History

Date	Version	Changes
June 2014	14.0.0	Added information that System Console supports the Tk toolkit.
November 2013	13.1.0	Dita conversion. Added link to Remote Debugging over TCP/IP for Altera SoC Application Note.
June 2012	12.0.0	Maintenance release.
November 2011	10.0.2	Maintenance release. Changed to new document template.
December 2010	10.0.1	Maintenance release. Changed to new document template.
July 2010	10.0.0	Initial release

For previous versions of the *Quartus II Handbook*, refer to the Quartus II Handbook Archive.

### Related Information

[Quartus II Handbook Archive](#)

# Analyzing and Debugging Designs with System Console 10

2014.12.15

QI15V3



Subscribe



Send Feedback

## About System Console

System Console provides visibility into your system. This visibility allows faster debugging and time to market for your FPGA. System Console is both a platform and an application for interacting with the debug-enabled portions of your design.

You can perform the following high-level tasks with System Console and the tools built on top of System Console:

- Perform board bring-up, with both finalized and partially complete designs.
- Remote debug from anywhere with internet access.
- Automate complex run-time verification solutions through scripting across multiple devices in your system.
- Test serial links with point-and-click configuration tuning in the Transceiver Toolkit.
- Debug memory interfaces with the External Memory Interface Toolkit.
- Integrate your own debug IP into the debugging platform.
- Test the performance of your ADC and analog chain on a MAX<sup>®</sup> 10 device using the ADC Toolkit.
- Use MATLAB/Simulink environment with System Console to perform system verification.

### Related Information

[System Console Online Training](#)

## Use Cases for System Console

You can leverage System Console for multiple debugging use cases. You can access tutorials, application notes, and design examples to learn more about debugging with System Console.

### Related Information

- [Board Bring-Up with System Console Tutorial](#) on page 10-9
- [About the ADC Toolkit](#) on page 10-26
- [External Memory Interface Documentation](#)
- [Debugging Transceiver Links Documentation](#) on page 11-1
- [Application Note 693: Remote Hardware Debugging over TCP/IP for Altera SoC](#)
- [Application Note 624: Debugging with System Console over TCP/IP](#)

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



- [White Paper 01208: Hardware in the Loop from the MATLAB/Simulink Environment](#)

## Using Debug Agents

System Console runs on your host computer and communicates with your running design through debug agents. These debug agents are soft-logic added to particular IP cores to enable debug communication with the host computer. Some debug agents have this single-purpose function. Other debug agents, such as the Nios II processor with debug enabled, are for both debugging the hardware in your design as well as doing software debugging of the code running on the Nios II processor.

By including debugging IP cores in your design, you can make large portions of a design debug-accessible. The IP allows reading of memory and altering peripheral registers from a host computer. For example, adding a JTAG to Avalon Master Bridge instance to a Qsys system enables you to read and write to memory-mapped slaves connected to the bridge. Other types of debug agents are also available.

You can instantiate debug IP cores using the IP Catalog.

**Note:** The following IP cores in the IP Catalog do not support VHDL simulation generation in the current version of the Quartus II software:

- JTAG Debug Link
- SLD Hub Controller System
- USB Debug Link

## System Console Flow

1. Add required component(s) to Qsys.
2. Generate and compile design.
3. Connect board and program FPGA.
4. Start System Console.
5. Locate and open service path.
6. Perform desired operation(s) with service.
7. Close the service.

## Application and Interfaces

Use the Tcl scripting language to interact with your running design in both the graphical and command-line interface modes. The System Console GUI provides additional panes to make important design information available.

System Console understands the particulars of the communication channel because of design information embedded in the programmable SRAM Object File (.sof). When System Console launches from the Quartus II software or Qsys while your design is open, any existing programmable file is automatically found and linked to the detected running device if they are compatible. In more complicated systems, the designs and devices may need to be linked manually.

### Related Information

- [API](#) on page 10-38
- [Quartus II Scripting Reference Manual](#)  
Information about Tcl scripting support
- [Introduction to Tcl Online Training](#)

## Starting System Console

There are several different ways to launch System Console.

### Starting System Console from Quartus II

- Click **Tools > System Debugging Tools > System Console**.

### Starting System Console from Qsys

- Click **Tools > System Console**.

### Starting System Console from Nios II Command Shell

1. On the Windows Start menu, click **All Programs > Altera > Nios II EDS <version> > Nios II <version> Command Shell**.
2. Type the following command:

```
system-console
```

**Note:** To get help information, type the command `system-console --help`

### Customizing Startup

You can customize your System Console environment by adding commands to the `system_console_rc` configuration file. You can locate this file in the following location:

- `<$HOME>/system_console/system_console_rc.tcl`, the file in this location is known as the user configuration file, which only affects the owner of that home directory.

You can alternatively specify your own design specific startup configuration file by using the command-line argument `--rc_script=<path_to_script>`, when you launch System Console from the Nios II command shell.

You can use the `system_console_rc.tcl` file in combination with your custom `rc_script.tcl` file. In this capacity, the `system_console_rc.tcl` file performs actions that System Console always needs and the local `rc_script.tcl` file performs actions for particular experiments.

On startup, System Console automatically runs any Tcl commands in these files. The commands in the `system_console_rc.tcl` file run first, then the commands in the `rc_script.tcl` file run.

### Command-Line Arguments

The `--cli` command-line argument runs System Console in command-line mode.

The `--project_dir=<project_dir>` command-line argument directs System Console to the location of your hardware project. Ensure that you are working with the project you intend. The JTAG chain details and other information depend on the specific project.

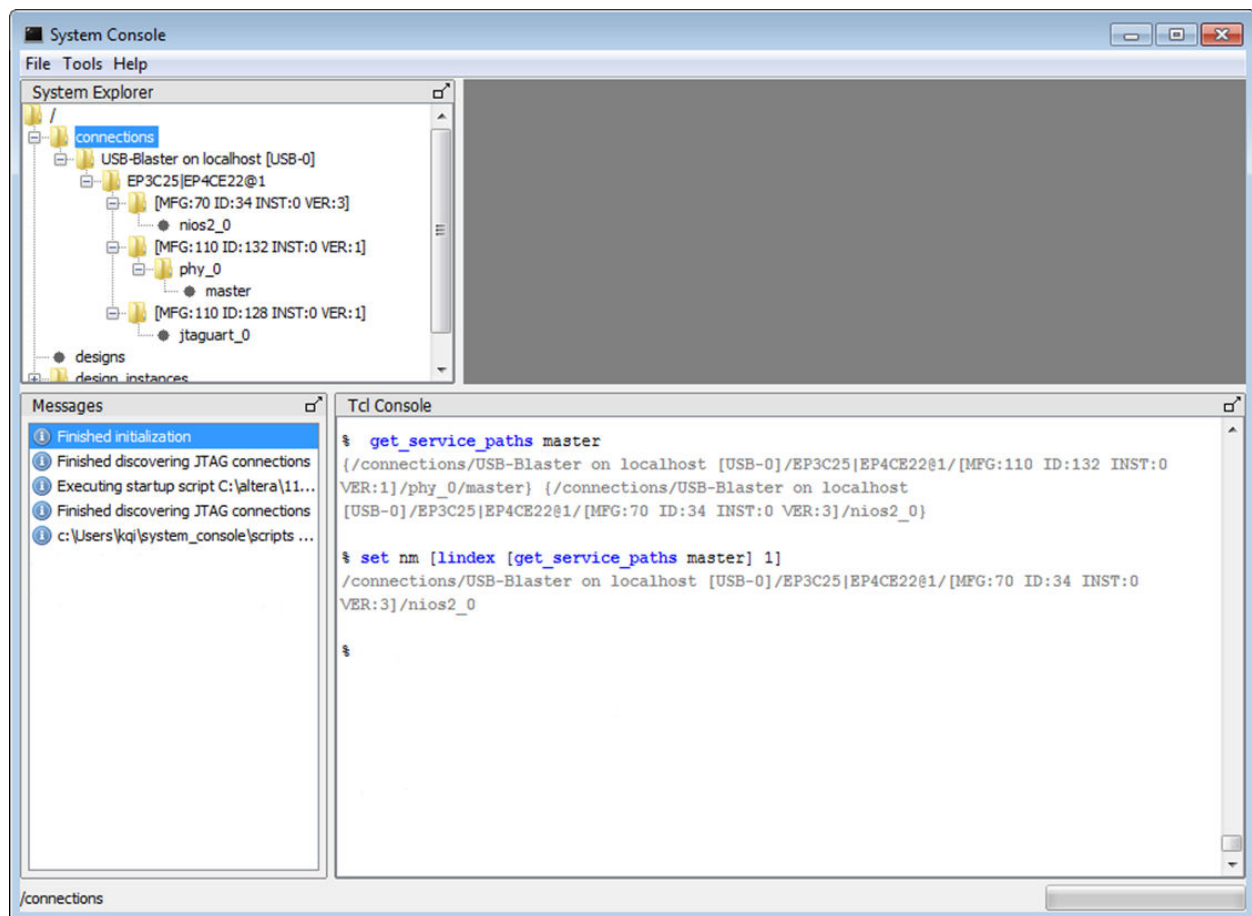
The `--script=<your_script>.tcl` command-line argument directs System Console to run your Tcl script.

## The System Console GUI

The System Console GUI consists of a main window with four separate panes.

- The **System Explorer** pane displays the hierarchy of the System Console virtual file system in your design, including board connections, devices, designs, and scripts.
- The **Tools** pane displays the ADC Toolkit, Transceiver Toolkit, Toolkits, GDB Server Control Panel, and Bus Analyzer. Click the **Tools** menu to launch the applications.
- The **Tcl Console** is where the design interactions take place. Common actions are sourcing scripts, writing procedures, and using the System Console API.
- The **Messages** pane displays status, warning, and error messages regarding connections and debug actions.

Figure 10-1: System Console GUI



### Related Information

[System Console Online Help](#)

### System Explorer Pane

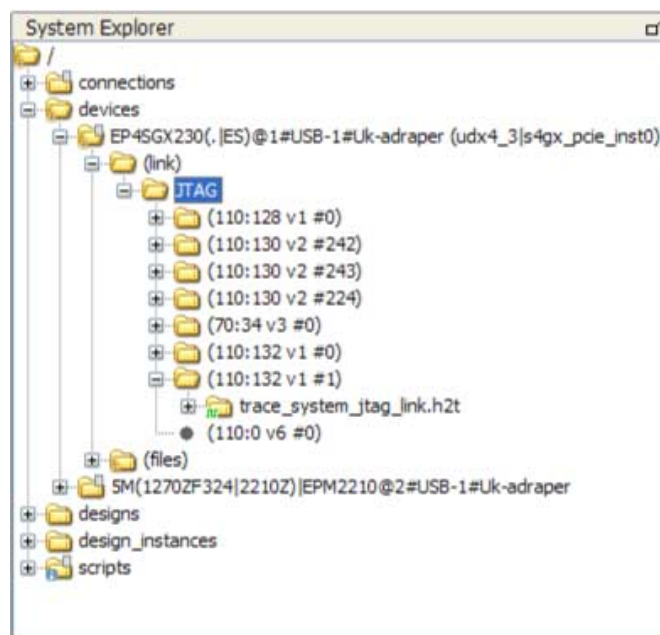
The System Explorer pane displays the virtual file system for all connected debugging components. This virtual file system contains the following information:

- The **devices** folder contains information about each device connected to System Console.
- The **scripts** folder stores scripts for easy execution.
- The **connections** folder displays information about the board connections which are visible to System Console, such as USB Blaster. Multiple connections are possible.
- The **designs** folder displays information about Quartus II project designs connected to System Console.

Within the **devices** folder is a folder for each device currently connected to System Console. Each device folder contains a **(link)** folder and sometimes contains a **(files)** folder.

The **(link)** folder shows debug agents (and other hardware) that System Console is able to access. The **(files)** folder is a copy of the tree under the designs folder for the project that is currently linked to this device.

Figure 10-2: System Explorer Pane



- The figure above shows that under the devices folder there is the **EP4SGX230** folder which contains a **(link)** folder. The **(link)** folder contains a **JTAG** folder. The **JTAG** folder contains folders that describe the debug pipes (i.e. JTAG, USB, Ethernet, etc) and agents that are connected to the EP4SGX230 device via a JTAG connection.
- The **(files)** folder contains information about the design files loaded from the Quartus II project for the device.
- Folders that have a context menu available show a small context menu icon. Right-click these folders to view a context menu. For example, the **connections** folder in **Figure 10-2** shows a context menu icon.
- Folders that have informational messages available display a small informational message icon. Hover over these folders to see the informational message. For example, the **scripts** folder in **Figure 10-2** shows an informational message icon.
- Debug agents that sense the clock and reset state of the target show an informational or error message with a clock status icon. The icon indicates whether the clock is running (info, green), stopped (error, red), or running but in reset (error, red). For example, the **trace\_system\_jtag\_link.h2t** folder in **Figure 10-2** has a running clock.

## Interactive Help

Typing `help help` into the Tcl Console lists all available commands. Typing `help <command name>` provides the syntax of commands. System Console provides command completion if you type the beginning letters of a command and then press the Tab key.

## Services

System Console services allow you to access different parts of your running design. For example, the master service provides access to memory-mapped slave interfaces and the processor service provides access to fine-grained processor controls. The services do not intermix, but a single IP core can provide multiple services. For example, the Nios II processor contains a debug core. It is a processor and it has a memory-mapped master interface that can connect to slaves. The master service can access the memory-mapped slaves that connect to the Nios II processor. Also, you can use the processor service to do software debugging.

## Common Services

Each common service exposes a separate API. By adding the appropriate debug agent to your design, System Console services can use the associated capabilities of the debug agent.

**Table 10-1: Common Services for System Console**

Service	Function	Debug Agent Providing Service
master	Access memory-mapped (Avalon-MM or AXI) slaves connected to the master interface.	<ul style="list-style-type: none"> <li>• Nios II with debug</li> <li>• JTAG to Avalon Master Bridge</li> <li>• USB Debug Master</li> </ul>

Service	Function	Debug Agent Providing Service
slave	Allows the host to access a single slave without needing to know the location of the slave in the host's memory map. Any slave that is accessible to a System Console master can provide this service.	<ul style="list-style-type: none"> <li>Nios II with debug</li> <li>JTAG to Avalon Master Bridge</li> </ul>
processor	<ul style="list-style-type: none"> <li>Start, stop, or step the processor.</li> <li>Read and write processor registers.</li> </ul>	Nios II with debug

**Related Information**

- [System Console Examples](#) on page 10-8
- [API](#) on page 10-38

## Locating Available Services

System Console uses a virtual file system to organize the available services, which is similar to the **/dev location** on Linux systems. Board connection, device type, and IP names are all part of a service path. Instances of services are referred to by their unique service path in the file system. You can retrieve service paths for a particular service with the command `get_service_paths <service-type>`.

### Example 10-1: Locating a Service Path

```
#We are interested in master services.
set service_type "master"

#Get all the paths as a list.
set master_service_paths [get_service_paths $service_type]

#We are interested in the first service in the list.
set master_index 0

#The path of the first master.
set master_path [lindex $master_service_paths $master_index]

#Or condense the above statements into one statement:
set master_path [lindex [get_service_paths master] 0]
```

System Console commands require service paths to identify the service instance you want to access. The paths for different components can change between runs of System Console and between versions. Use `get_service_paths` to obtain service paths rather than hard coding them into your Tcl scripts.

The string values of service paths change with different releases of the tool, so you should not infer meaning from the actual strings within the service path. Use `marker_node_info` to get information from the path.

System Console automatically discovers most services at startup. System Console automatically scans for all JTAG and USB-based service instances and retrieves their service paths. System Console does not automatically discover some services, such as TCP/IP. Use `add_service` to inform System Console about those services.

### Example 10-2: Marker\_node\_info

You can also use the `marker_node_info` command to get information about the discovered services so you can choose the right one.

```
foreach m [get_service_paths master] {
  array set minfo [marker_node_info $m]
  if {[string match {*myhpath} $minfo(full_hpath)]} {
    set master_path $m
    break
  }
}
```

## Opening and Closing Services

After you have a service path to a particular service instance, you can access the service for use.

The `claim_service` command directs System Console to start using a particular service instance. The `claim_service` command claims a service instance for exclusive use.

### Example 10-3: Opening a Service

```
set service_type "master"
set claim_path [claim_service $service_type $master_path mylib];#Claims
service.
```

You can pass additional arguments to the `claim_service` command to direct System Console to start accessing a particular portion of a service instance. For example, if you use the master service to access memory, then use `claim_service` to only access the address space between 0x0 and 0x1000. System Console then allows other users to access other memory ranges, and denies access to the claimed memory range. The `claim_service` command returns a newly created service path that you can use to access your claimed resources.

You can access a service after you open it. When you finish accessing a service instance, use the `close_service` command to direct System Console to make this resource available to other users.

### Example 10-4: Closing a Service

```
close_service master $claim_path; #Closes the service.
```

## System Console Examples

Altera provides examples for performing board bring-up, creating a simple dashboard, and programming a Nios II processor. The **System\_Console.zip** file contains design files for the board bring-up example. The Nios II Ethernet Standard **.zip** files contain the design files for the Nios II processor example.

**Note:** The instructions for these examples assume that you are familiar with the Quartus II software, Tcl commands, and Qsys.

**Related Information**

[On-Chip Debugging Design Examples Website](#)

Contains the design files for the example designs that you can download.

## Board Bring-Up with System Console Tutorial

You can perform low-level hardware debugging of Qsys systems with System Console. You can debug systems that include IP cores instantiated in your Qsys system or perform initial bring-up of your PCB. This board bring-up tutorial uses a Nios II Embedded Evaluation Kit (NEEK) board and USB cable. If you have a different development kit, you need to change the device and pin assignments to match your board and then recompile the design.

**Related Information**

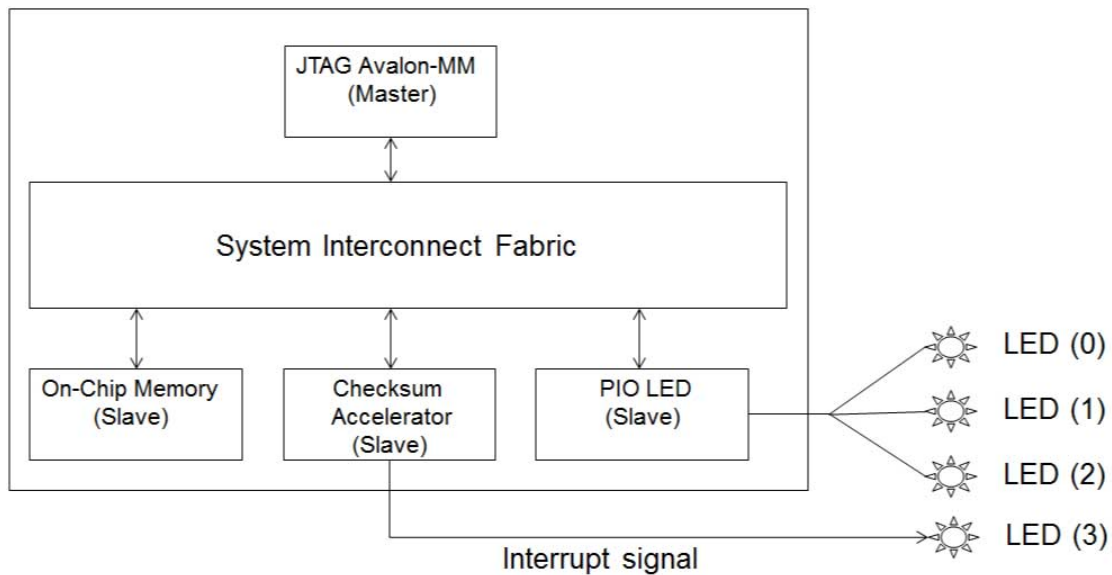
- [Use Cases for System Console](#) on page 10-1
- [Faster Board Bring-Up with System Console Demo Video](#)

### Board Bring-Up Flow

1. Set up the board bring-up example.
2. Verify clock and reset signals.
3. Verify memory and other peripheral interfaces.

### Qsys Modules

Figure 10-3: Qsys Modules for Board Bring-up Example





The Qsys design for this example includes the following modules:

- JTAG to Avalon Master Bridge—Provides System Console host access to the memory-mapped IP in the design via the JTAG interface.
- On-chip memory—Simplest type of memory for use in an FPGA-based embedded system. The memory is implemented on the FPGA; consequently, external connections on the circuit board are not necessary.
- Parallel I/O (PIO) module—Provides a memory-mapped interface for sampling and driving general I/O ports.
- Checksum Accelerator—Calculates the checksum of a data buffer in memory. The Checksum Accelerator consists of the following:
  - Checksum Calculator (**checksum\_transform.v**)
  - Read Master (**slave.v**)
  - Checksum Controller (**latency\_aware\_read\_master.v**)

### Checksum Accelerator Functionality

The base address of the memory buffer and data length passes to the Checksum Controller from a memory-mapped master. The Read Master continuously reads data from memory and passes the data to the Checksum Calculator. When the checksum calculations finish, the Checksum Calculator issues a valid signal along with the checksum result to the Checksum Controller. The Checksum Controller sets the DONE bit in the status register and also asserts the interrupt signal. You should only read the result from the Checksum Controller when the DONE bit and interrupt signal are asserted.

### Setting Up the Board Bring-Up Design Example

To load the design example into the Quartus II software and program your device, follow these steps:

1. Unzip the **System\_Console.zip** file to your local hard drive.
2. Click **File > Open Project** and select **Systemconsole\_design\_example.qpf** with the Quartus II software.
3. Change the device and pin assignments (LED, clock, and reset pins) in the **Systemconsole\_design\_example.qsf** file to match your board.
4. Click **Processing > Start Compilation**
5. To Program your device, follow these steps:
  - a. Click **Tools > Programmer**.
  - b. Click **Hardware Setup**.
  - c. Click the **Hardware Settings** tab.
  - d. Under **Currently selected hardware**, click **USB-Blaster**, and click **Close**.

**Note:** If you do not see the **USB-Blaster** option, then your device was not detected. Verify that the USB-Blaster driver is installed, your board is powered on, and the USB cable is intact.

This design example uses a USB-Blaster cable. If you do not have a USB-Blaster cable and you are using a different cable type, then select your cable from the **Currently selected hardware** options.

- e. Click **Auto Detect**, and then select your device.
- f. Double-click your device under **File**.
- g. Browse to your project folder and click **Systemconsole\_design\_example.sof** in the subdirectory **output\_files**.

- h. Turn on the **Program/Configure** option.
  - i. Click **Start**.
  - j. Close the Programmer.
6. Click **Tools > System Debugging Tools > System Console**.

**Related Information**

[System\\_Console.zip file](#)

Contains the design files for this tutorial.

## Verifying Clock and Reset Signals

You can use the System Explorer pane to verify clock and reset signals. Open the appropriate node and check for either a green clock icon or a red clock icon.

**Related Information**

[System Explorer Pane](#) on page 10-4

## Verifying Memory and Other Peripheral Interfaces

The Avalon-MM service accesses memory-mapped slaves via a suitable Avalon-MM master, which can be controlled by the host. You can use Tcl commands to read and write to memory with a master service.

### Locating and Opening the Master Service

```
#Select the master service type and check for available service paths.
set service_paths [get_service_paths master]

#Set the master service path.
set master_service_path [lindex $service_paths 0]


#Open the master service.
set claim_path [claim_service master $master_service_path mylib]
```

### Avalon-MM Slaves

The **Address Map** tab shows the address range for every Qsys component. The Avalon-MM master communicates with slaves using these addresses.

The register maps for all Altera components are in their respective Data Sheets.

**Figure 10-4: Address Map**



**Related Information**

[Data Sheets Website](#)

## Testing the PIO component

In this example design, the PIO connects to the LEDs of the board. Test if this component is operating properly and the LEDs are connected, by driving the outputs with the Avalon-MM master.

**Table 10-2: Register Map for the PIO Core**

Offset	Register Name		R/W	Fields				
				(n-1)	...	2	1	0
0	data	read access	R	Data value currently on PIO inputs.				
		write access	W	New value to drive on PIO outputs.				
1	direction		R/W	Individual direction control for each I/O port. A value of 0 sets the direction to input; 1 sets the direction to output.				
2	interruptmask		R/W	IRQ enable/disable for each input port. Setting a bit to 1 enables interrupts for the corresponding port.				
3	edgecapture		R/W	Edge detection for each input port.				

```
#Write the driver output values for the Parallel I/O component.
set offset 0x0; #Register address offset.
set value 0x7; #Only set bits 0, 1, and 2.
master_write_8 $claim_path $offset $value

#Read back the register value.
set offset 0x0
set count 0x1
master_read_8 $claim_path $offset $count

master_write_8 $claim_path 0x0 0x2; #Only set bit 1.
master_write_8 $claim_path 0x0 0xe; #Only set bits 1, 2, 3.
master_write_8 $claim_path 0x0 0x7; #Only set bits 0, 1, 2.

#Observe the LEDs turn on and off as you execute these Tcl commands.
#The LED is on if the register value is zero and off if the register value is one.
#LED 0, LED 1, and LED 2 connect to the PIO.
#LED 3 connects to the interrupt signal of the CheckSum Accelerator.
```

## Testing On-chip Memory

Test the memory with a recursive function that writes to incrementing memory addresses.

```
#Load the design example utility procedures for writing to memory.
source set_memory_values.tcl

#Write to the on-chip memory.
set base_address 0x80
set write_length 0x80
set value 0x5a5a5a5a
fill_memory $claim_path $base_address $write_length $value

#Verify the memory was written correctly.
#This utility proc returns 0 if the memory range is not uniform with this value.
verify_memory $claim_path $base_address $write_length $value

#Check that the memory is re-initialized when reset.
#Trigger reset then observe verify_memory returns 0.
set jtag_debug_path [lindex [get_service_paths jtag_debug] 0]
```

```
set claim_jtag_debug_path [claim_service jtag_debug $jtag_debug_path mylib]
jtag_debug_reset_system $claim_jtag_debug_path; #Reset the connected on-chip memory
#peripheral.
close_service jtag_debug $claim_jtag_debug_path
verify_memory $claim_path $base_address $write_length $value

#The on-chip memory component was parameterized to re-initialized to 0 on reset.
#Check the actual value.
master_read_8 $claim_path 0x0 0x1
```

### Testing the Checksum Accelerator

The Checksum Accelerator calculates the checksum of a data buffer in memory. It calculates the value for a specified memory buffer, sets the DONE bit in the status register, and asserts the interrupt signal. You should only read the result from the controller when both the DONE bit and the interrupt signal are asserted. The host should assert the interrupt enable control bit in order to check the interrupt signal.

**Table 10-3: Register Map for Checksum Component**

Offset (Bytes)	Hexadecimal value (after adding offset)	Register	Access	Bits (32 bits)							
				31-9	8	7-5	4	3	2	1	0
0	0x20	Status	Read/Write to clear							BUSY	DONE
4	0x24	Address	Read/Write	Read Address							
12	0x2C	Length	Read/Write	Length in bytes							
24	0x38	Control	Read/Write		Fixed Read Address Bit		Interrupt Enable	GO		INV	Clear
28	0x3C	Result	Read	Checksum result (upper 16 bits are zero)							

```
1. #Pass the base address of the memory buffer Checksum Accelerator.
set base_address 0x20
set offset 4
set address_reg [expr {$base_address + $offset}]
set memory_address 0x80
master_write_32 $claim_path $address_reg $memory_address

#Pass the memory buffer to the Checksum Accelerator.
set length_reg [expr {$base_address + 12}]
set length 0x20
master_write_32 $claim_path $length_reg $length

#Write clear to status and control registers.
#Status register:
set status_reg $base_address
master_write_32 $claim_path $status_reg 0x0
#Control register:
set clear 0x1
set control_reg [expr {$base_address + 24}]
master_write_32 $claim_path $control_reg $clear
```

```
#Write GO to the control register.
set go 0x8
master_write_32 $claim_path $control_reg $go

#Cross check if the checksum DONE bit is set.
master_read_32 $claim_path $status_reg 0x1

#Is the DONE bit set?
#If yes, check the result and you are finished with the board bring-up design
example.
set result_reg [expr {$base_address + 28}]
master_read_16 $claim_path $result_reg 0x1
```

2. If the result is zero and the JTAG chain works properly, the clock and reset signals work properly, and the memory works properly, then the problem is the Checksum Accelerator component.

```
#Confirm if the DONE bit in the status register (bit 0)
#and interrupt signal are asserted.
#Status register:
master_read_32 $claim_path $status_reg 0x1
#Check DONE bit should return a one.

#Enable interrupt and go:
set interrupt_and_go 0x18
master_write_32 $claim_path $control_reg $interrupt_and_go
```

3. Check the control enable to see the interrupt signal. LED 3 (MSB) should be off. This indicates the interrupt signal is asserted.
4. You have narrowed down the problem to the data path. View the RTL to check the data path.
5. Open the **Checksum\_transform.v** file from your project folder.
  - `<unzip dir>/System_Console/ip/checksum_accelerator/checksum_accelerator.v`
6. Notice that the `data_out` signal is grounded in **Figure 10-5** (uncommented line 87 and comment line 88). Fix the problem.
7. Save the file and regenerate the Qsys system.
8. Re-compile the design and reprogram your device.
9. Redo the above steps, starting with **Verifying Memory and Other Peripheral Interfaces** on page 10-11 or run the Tcl script included with this design example.

```
source set_memory_and_run_checksum.tcl
```

Figure 10-5: Checksum.v File

```
83 // first folding
84 assign first_folded_sum = (initial_sum [32] + initial_sum[31:16] + initial_sum[15:0]); // this result is at most 17 bits wide (16 bits with rollover)
85
86 // second folding and optional inversion, this result is at most 16 bits wide
87 assign data_out = (invert == 1)? ~(first_folded_sum[16] + first_folded_sum[15:0]) : (first_folded_sum[16] + first_folded_sum[15:0]);
88 // assign data_out = 16'h0000;
89
90 endmodule
```

## Dashboard Service

The dashboard service enables you to construct GUIs for visualizing and interacting with debug data. The dashboard service provides graphical widgets such as buttons and text fields. The dashboard is a graphical pane for the layout of your widgets. Widgets can pull data from other System Console services. Similarly, widgets can leverage user input to act on debug logic in your design through services.

## Properties

Widget properties can push information to the user interface and pull information from the user interface. Widgets have properties specific to their type. For example, the `onClick` property performs an action when the button is clicked. A label widget does not have the same property because it does not perform an action when clicked. However, both the button and label widgets have the `text` property for the string they display.

## Layout

The dashboard service creates a widget hierarchy where the dashboard is at the top-level. The dashboard service can implement group-type widgets that contain child widgets. Layout properties dictate layout actions performed by a parent on its children.

An example layout property is `expandable`: if true, the widget expands horizontally to encompass all the space available to it. Another property is `visible`: a widget is only laid out when this property is true.

## User Input

Some of the available widgets allow user interaction. For example, the `textField` widget is a box that allows you to type text. For this widget, the contents of the box are accessible through the `text` property. A Tcl script can either get or set the contents of the field by accessing this property.

## Callbacks

Some widgets can perform user-specified actions, referred to as callbacks, upon certain events. The `textField` widget has the `onChange` property, which is called when the text contents change. The `button` widget has the `onClick` property, which is called when the button is clicked. These callbacks may update widgets or interact with services based on the contents of the text field or the state of any other widget.

### Related Information

[Dashboard Commands](#) on page 10-45

## Dashboard Example

### Example 10-5: Adding the Service

The dashboard is not initialized by default. You must add the service before it can be used.

```
set dash [add_service dashboard dashboard_example "Dashboard Example" \  
"Tools/Example"]
```

### Example 10-6: Showing the Dashboard

Once instantiated, you must explicitly make the dashboard visible. Use the `dashboard_set_property` command to modify the `visible` property of the root dashboard:

```
dashboard_set_property $dash self visible true
```

In this command, `$dash` represents the dashboard service. `self` is the name of the root dashboard widget. `visible` is the property being set. `true` is the value to set. Executed as a single command, it causes the root dashboard to be made visible.

### Example 10-7: Adding Widgets

Use the `dashboard_add` command to add widgets:

```
set name "my_label"  
set widget_type "label"  
set parent "self"  
dashboard_add $dash $name $widget_type $parent
```

The following commands add a label widget named "my\_label" to the root dashboard. In the GUI, it appears as the text "label." Change the text:

```
set content "Text to display goes here"  
dashboard_set_property $dash $name text $content
```

This command sets the `text` property to that string. In the GUI, the displayed text changes to the new value. Add one more label:

```
dashboard_add $dash my_label_2 label self  
dashboard_set_property $dash my_label_2 text "Another label"
```

Notice the new label appears to the right of the first label. Cause the layout to put the label below instead:

```
dashboard_set_property $dash self itemsPerRow 1
```

### Example 10-8: Gathering Input

Incorporate user input into our dashboard:

```
set name "my_text_field"  
set widget_type "textField"  
set parent "self"  
dashboard_add $dash $name $widget_type $parent
```

The widget appears, but it is very small. Make the widget fill the horizontal space:

```
dashboard_set_property $dash my_text_field expandableX true
```

Now the text field is fully visible. Text can be typed into it once clicked. Type a sentence. Now, retrieve the contents of the field:

```
set content [dashboard_get_property $dash my_text_field text]  
puts $content
```

This prints the contents into the console.

### Example 10-9: Updating Widgets Upon User Events

You can make the dashboard perform actions without interactive typing. Use callbacks to accomplish this. Start by defining a procedure that updates the first label with the text field contents:

```
proc update_my_label_with_my_text_field {dash} {  
    set content [dashboard_get_property $dash my_text_field text]  
    dashboard_set_property $dash my_label text $content  
}
```

Run the `update_my_label_with_my_text_field $dash` command in the Tcl Console. Notice that the first label now matches the text field contents. Have the `update_my_label_with_my_text_field $dash` command called whenever the text field changes:

```
dashboard_set_property $dash my_text_field onChange \  
"update_my_label_with_my_text_field $dash"
```

The `onChange` property is executed each time the text field changes. The effect is the first field changes to match what is typed.

### Example 10-10: Buttons

You can use buttons to trigger actions. Create a button that changes the second label:

```
proc append_to_my_label_2 {dash suffix} {  
    set old_text [dashboard_get_property $dash my_label_2 text]  
    set new_text "${old_text}${suffix}"  
    dashboard_set_property $dash my_label_2 text $new_text  
}  
set text_to_append ", and more"  
dashboard_add $dash my_button button self  
dashboard_set_property $dash my_button onClick [list append_to_my_label_2 \  
$dash $text_to_append]
```

Click the button and the second label gets some text appended to it.

### Example 10-11: Groups

The property `itemsPerRow` dictates how widgets are laid out in a group. For more complicated layouts where the number of widgets per row is different per row, nested groups should be used. Add a new group with more widgets per row:

```
dashboard_add $dash my_inner_group group self  
dashboard_set_property $dash my_inner_group itemsPerRow 2  
dashboard_add $dash inner_button_1 button my_inner_group  
dashboard_add $dash inner_button_2 button my_inner_group
```

There is now a row with a group of two buttons. The border with the group name can be removed to make the nested group more seamless.

```
dashboard_set_property $dash inner_group title ""
```



The `title` property can be set to any other string to have the border and title text show up.

### Example 10-12: Tabs

GUIs do not require all the widgets to be visible at the same time. Tabs accomplish this.

```
dashboard_add $dash my_tabs tabbedGroup self
dashboard_set_property $dash my_tabs expandableX true
dashboard_add $dash my_tab_1 group my_tabs
dashboard_add $dash my_tab_2 group my_tabs
dashboard_add $dash tabbed_label_1 label my_tab_1
dashboard_add $dash tabbed_label_2 label my_tab_2
dashboard_set_property $dash tabbed_label_1 text "in the first tab"
dashboard_set_property $dash tabbed_label_2 text "in the second tab"
```

This adds a set of two tabs, each with a group containing a label. Clicking on the tabs changes the displayed group/label.

## Nios II Processor Example

This example programs the Nios II processor on your board to run the count binary software example included in the Nios II installation. This is a simple program that uses an 8-bit variable to repeatedly count from 0x00 to 0xFF. The output of this variable is displayed on the LEDs on your board. After programming the Nios II processor, you use System Console processor commands to start and stop the processor.

To run this example, perform the following steps:

1. Download the Nios II Ethernet Standard Design Example for your board from the Altera website.
2. Create a folder to extract the design. For this example, use **C:\Count\_binary**.
3. Unzip the Nios II Ethernet Standard Design Example into **C:\Count\_binary**.
4. In a Nios II command shell, change to the directory of your new project.
5. Program your board. In a Nios II command shell, type the following:

```
nios2-configure-sof niosii_ethernet_standard_<board_version>.sof
```

6. Using Nios II Software Build Tools for Eclipse, create a new Nios II Application and BSP from Template using the **Count Binary** template and targeting the Nios II Ethernet Standard Design Example.
7. To build the executable and linkable format (ELF) file (**.elf**) for this application, right-click the **Count Binary** project and select **Build Project**.
8. Download the **.elf** file to your board by right-clicking **Count Binary** project and selecting **Run As, Nios II Hardware**.
  - The LEDs on your board provide a new light show.
9. Type the following:

```
system-console; #Start System Console.

#Set the processor service path to the Nios II processor.
set niosii_proc [lindex [get_service_paths processor] 0]

set claimed_proc [claim_service processor $niosii_proc mylib]; #Open the service.
```

```
processor_stop $claimed_proc; #Stop the processor.  
#The LEDs on your board freeze.  
  
processor_run $claimed_proc; #Start the processor.  
#The LEDs on your board resume their previous activity.  
  
processor_stop $claimed_proc; #Stop the processor.  
  
close_service processor $claimed_proc; #Close the service.
```

- The `processor_step`, `processor_set_register`, and `processor_get_register` commands provide additional control over the Nios II processor.

#### Related Information

- [Processor Commands](#) on page 10-55
- [Nios II Ethernet Standard Design Example](#)
- [Nios II Software Build Tools User Guide](#)

## Additional Services

### Design Service

You can use design service commands to work with Quartus II design information.

#### Example 10-13: Load

When you open System Console from the Quartus II software or Qsys, the current project's debug information is sourced automatically if the `.sof` has been built. In other situations, you can load manually.

```
set sof_path [file join project_dir output_files project_name.sof]  
set design [design_load $sof_path]
```

System Console is now aware that this particular `.sof` has been loaded.

#### Example 10-14: Linking

Once a `.sof` is loaded, System Console automatically links design information to the connected device. The resultant link persists and you can choose to unlink or reuse the link on an equivalent device with the same `.sof`.

You can perform manual linking.

```
set device_index 0; # Device index for our target  
set device [lindex [get_service_paths device] $device_index]  
design_link $design $device
```

Manually linking fails if the target device does not match the design service.

Linking succeeds even if the `.sof` programmed to the target is not the same as the design `.sof`.

**Related Information**

[Design Service Commands](#) on page 10-41

## Device Service

The device service supports device-level actions.

### Example 10-15: Programming

You can use the device service with Tcl scripting to perform device programming.

```
set device_index 0 ; #Device index for target
set device [lindex [get_service_paths device] $device_index]
set sof_path [file join project_path output_files project_name.sof]
device_download_sof $device $sof_path
```

To program, all you need are the device service path and the file system path to a **.sof**. Ensure that no other service (e.g. master service) is open on the target device or else the command fails. Afterwards, you may do the following to check that the design linked to the device is the same one programmed:

```
device_get_design $device
```

**Related Information**

[Device Commands](#) on page 10-42

## Monitor Service

The monitor service builds on top of the master service to allow reads of Avalon-MM slaves at a regular interval. The service is fully software-based. The monitor service requires no extra soft-logic. This service streamlines the logic to do interval reads, and it offers better performance than exercising the master service manually for the reads.

### Example 10-16: Monitor Service

Start by determining a master and a memory address range that you are interested in polling continuously.

```
set master_index 0
set master [lindex [get_service_paths master] $master_index]
set address 0x2000
set bytes_to_read 100
set read_interval_ms 100
```

You can use the first master to read 100 bytes starting at address 0x2000 every 100 milliseconds. Open the monitor service:

```
set monitor [lindex [get_service_paths monitor] 0]
set claimed_monitor [claim_service monitor $monitor mylib]
```

Notice that the master service was not opened. The monitor service opens the master service automatically. Register the previously-defined address range and time interval with the monitor service:

```
monitor_add_range $claimed_monitor $master $address $bytes_to_read
monitor_set_interval $claimed_monitor $read_interval_ms
```

More ranges can be added. Define what happens at each interval:

```
global monitor_data_buffer
set monitor_data_buffer [list]
proc store_data {monitor master address bytes_to_read} {
    global monitor_data_buffer
    set data [monitor_read_data $claimed_monitor $master $address $bytes_to_read]
    lappend monitor_data_buffer $data
}
```

The code example above, gathers the data and appends it with a global variable. `monitor_read_data` returns the range of data polled from the running design as a list. In this example, data will be a 100-element list. This list is then appended as a single element in the `monitor_data_buffer` global list. If this procedure takes longer than the interval period, the monitor service may have to skip the next one or more calls to the procedure. In this case, `monitor_read_data` will return the latest data polled. Register this callback with the opened monitor service:

```
set callback [list store_data $claimed_monitor $master $address $bytes_to_read]
monitor_set_callback $claimed_monitor $callback
```

Use the callback variable to call when the monitor finishes an interval. Start monitoring:

```
monitor_set_enabled $claimed_monitor 1
```

Immediately, the monitor reads the specified ranges from the device and invokes the callback at the specified interval. Check the contents of `monitor_data_buffer` to verify this. To turn off the monitor, use 0 instead of 1 in the above command.

#### Related Information

[Monitor Commands](#) on page 10-53

## Bytestream Service

The bytestream service provides access to modules that produce or consume a stream of bytes. You can use the bytestream service to communicate directly to the IP core that provides bytestream interfaces, such as the Altera JTAG UART of the Avalon-ST JTAG interface.

### Example 10-17: Bytestream Service

The following code finds the bytestream service for your interface and opens it.

```
set bytestream_index 0
set bytestream [lindex [get_service_paths bytestream] $bytestream_index]
set claimed_bytestream [claim_service bytestream $bytestream mylib]
```

To specify the outgoing data as a list of bytes and send it through the opened service:

```
set payload [list 1 2 3 4 5 6 7 8]
bytestream_send $claimed_bytestream $payload
```

Incoming data also comes as a list of bytes.

```
set incoming_data [list]
while {[length $incoming_data] == 0} {
    set incoming_data [bytestream_receive $claimed_bytestream 8]
}
```

Close the service when done.

```
close_service bytestream $claimed_bytestream
```

### Related Information

[Bytestream Commands](#) on page 10-56

## SLD Service

The SLD Service shifts values into the instruction and data registers of SLD nodes and captures the previous value. When interacting with a SLD node, start by acquiring exclusive access to the node on an opened service.

### Example 10-18: SLD Service

```
set timeout_in_ms 1000
set lock_failed [sld_lock $sld_service_path $timeout_in_ms]
```

This code attempts to lock the selected SLD node. If it is already locked, `sld_lock` waits for the specified timeout. Confirm the procedure returns non-zero before proceeding. Set the instruction register and capture the previous one:

```
if {$lock_failed} {
    return
}
set instr 7
set delay_us 1000
set capture [sld_access_ir $sld_service_path $instr $delay_us]
```

The 1000 microsecond delay guarantees that the following SLD command executes least 1000 microseconds later. Data register access works the same way.

```
set data_bit_length 32
set delay_us 1000
set data_bytes [list 0xEF 0xBE 0xAD 0xDE]
set capture [sld_access_dr $sld_service_path $data_bit_length $delay_us \
    $data_bytes]
```

Shift count is specified in bits, but the data content is specified as a list of bytes. The capture return value is also a list of bytes. Always unlock the SLD node once finished with the SLD service.

```
sld_unlock $sld_service_path
```

## Related Information

- [SLD Commands](#) on page 10-40
- [Virtual JTAG Megafunction documentation](#)

## In-System Sources and Probes Service

The In-System Sources and Probes (ISSP) service provides scriptable access to the `altsource_probe` IP core in a similar manner to using the **In-System Sources and Probes Editor** in Quartus II.

### Example 10-19: ISSP Service

Before you use the ISSP service, ensure your design works in the **In-System Sources and Probes Editor**. In System Console, open the service for an ISSP instance.

```
set issp_index 0
set issp [lindex [get_service_paths issp] 0]
set claimed_issp [claim_service issp $issp mylib]
```

View information about this particular ISSP instance.

```
array set instance_info [issp_get_instance_info $claimed_issp]
set source_width $instance_info(source_width)
set probe_width $instance_info(probe_width)
```

Probe data is read as a single bitstring of length equal to the probe width.

```
set all_probe_data [issp_read_probe_data $claimed_issp]
```

As an example, you can define the following procedure to extract an individual probe line's data.

```
proc get_probe_line_data {all_probe_data index} {
    set line_data [expr { ($all_probe_data >> $index) & 1 }]
    return $line_data
}
set initial_all_probe_data [issp_read_probe_data $claim_issp]
set initial_line_0 [get_probe_line_data $initial_all_probe_data 0]
set initial_line_5 [get_probe_line_data $initial_all_probe_data 5]
# ...
set final_all_probe_data [issp_read_probe_data $claimed_issp]
set final_line_0 [get_probe_line_data $final_all_probe_data 0]
```

Similarly, source data is written as a single bitstring of length equal to the source width.

```
set source_data 0xDEADBEEF
issp_write_source_data $claimed_issp $source_data
```

The currently set source data can also be retrieved.

```
set current_source_data [issp_read_source_data $claimed_issp]
```

As an example, you can invert the data for a 32-bit wide source by doing the following:

```
set current_source_data [issp_read_source_data $claimed_issp]
set inverted_source_data [expr { $current_source_data ^ 0xFFFFFFFF }]
issp_write_source_data $claimed_issp $inverted_source_data
```

**Related Information**

[In-System Sources and Probes Commands](#) on page 10-57

## System Console Infrastructure

Services associated with debug agents in the running design can be directly opened and closed. Behind the scenes, System Console is responsible for determining and using the lower level protocol for communication with the debug agent. As part of this, the System Console infrastructure finds the best board connection to use for command and data transmission.

**Related Information**

[WP-01170 System-Level Debugging and Monitoring of FPGA Designs white paper](#)

Detailed information about the architecture for system level debugging.

## On-Board USB Blaster II Support

System Console supports an On-Board USB-Blaster™ II circuit via the USB Debug Master IP component. This IP core supports the master service.

Not all Stratix V boards support the On-Board USB-Blaster II. For example, the transceiver signal integrity board does not support the On-Board USB-Blaster II.

## About Using MATLAB and Simulink in a System Verification Flow

System Console can be used with MATLAB and Simulink to perform system development testing. You can use the Altera Hardware in the Loop (HIL) tools to set up a system verification flow. In this approach, the design is deployed to hardware and runs in real time. The surrounding components in your system are simulated in a software environment. The HIL approach allows you to use the flexibility of software tools with the real-world accuracy and speed of hardware. You can gradually introduce more hardware components to your system verification testbench. This gives you more control over the integration process as you tune and validate your system. When your full system is integrated, the HIL approach allows you to provide stimuli via software to test your system under a variety of scenarios.

### Advantages of HIL Approach

- Avoid long computational delays for algorithms with high processing rates
- API helps to control, debug, visualize, and verify FPGA designs all within the MATLAB environment
- FPGA results are read back by the MATLAB software for further analysis and display

### Required Tools and Components

- MATLAB software
- DSP Builder software
- Quartus II software
- Altera FPGA

**Note:** The System Console MATLAB API is included in the DSP Builder installation bundle.

Figure 10-6: Hardware in the Loop Host-Target Setup



### Supported MATLAB API Commands

You can perform your work from the MATLAB environment and leverage the capability of System Console to read and write to masters and slaves. By using the supported MATLAB API commands, you do not have to launch the System Console software. The supported commands are the following:

- `SystemConsole.refreshMasters;`
- `M = SystemConsole.openMaster(1);`
- `M.write (type, byte address, data);`
- `M.read (type, byte address, number of words);`
- `M.close`

### Example 10-20: MATLAB API Script Example

```
SystemConsole.refreshMasters; %Investigate available targets
M = SystemConsole.openMaster(1); %Creates connection with FPGA target
%%%%%%%% User Application %%%%%%%%%%%
....
M.write('uint32',write_address,data); %Send data to FPGA target
....
data = M.read('uint32',read_address,size); %Read data from FPGA target
....
%%%%%%%%%%
M.close; %Terminates connection to FPGA target
```



## High-Level Flow

1. Install the DSP Builder software so you have the necessary libraries to enable this flow
2. Build your design using Simulink and the DSP Builder libraries (DSP Builder helps to convert the Simulink design to HDL)
3. Include Avalon-MM components in your design (DSP Builder can port non-Avalon-MM components)
4. Include Signals and Control blocks in your design
5. Use boundary blocks to separate synthesizable and non-synthesizable logic
6. Integrate your DSP system in Qsys
7. Program your Altera FPGA
8. Use the supported MATLAB API commands to interact with your Altera FPGA

## Related Information

- [Hardware in the Loop from the MATLAB/Simulink Environment white paper](#)
- [System in the Loop - Enabling Real-Time FPGA Verification within MATLAB website](#)
- [DSP Builder website](#)

## About the ADC Toolkit

The ADC Toolkit is designed to work with MAX 10 devices and helps you understand the performance of the analog signal chain as seen by the on-board ADC hardware. The GUI displays the performance of the ADC using industry standard metrics. You can export the collected data to a **.csv** file and process this raw data yourself. The ADC Toolkit is built on the System Console framework and can only be operated using the GUI. There is no Tcl support for the tool.

## Prerequisites for Using the ADC Toolkit

- Altera Modular ADC IP core
  - **External Reference Voltage** if you select **External** in the Altera Modular ADC IP parameters
- Reference signal

The ADC Toolkit needs a sine wave signal to be fed to the analog inputs. You need the capability to precisely set the level and frequency of the reference signal. A high-precision sine wave is needed for accurate test results; however, there are useful things that can be read in **Scope** mode with any input signal.

To achieve the best testing results, the reference signal should have less distortions than the device ADC is able to resolve. If this is not the case, then you will be adding distortions from the source into the resulting ADC distortion measurements. The limiting factor is based on hardware precision.

**Note:** When applying a sine wave, the ADC should sample at 2x the fundamental sine wave frequency. There should be a low-pass filter, 3dB point set to the fundamental frequency.

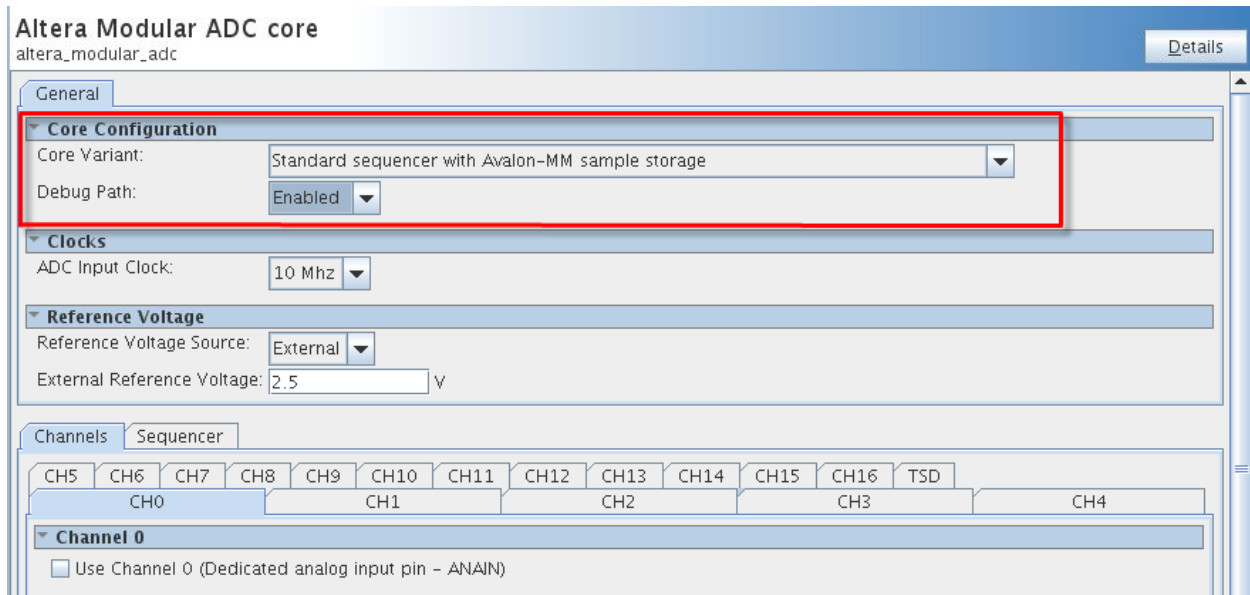
## Configuring the Altera Modular ADC IP Core

The Altera Modular ADC IP core needs to be included in your design. You can instantiate this IP core from the **IP Catalog**. When you configure this IP core in the **Parameter Editor**, you need to enable the **Debug Path** option located under **Core Configuration**.

There are two limitations in the Quartus II software v14.1 for the Altera Modular ADC IP core. The ADC Toolkit does not support the **ADC control core only** option under **Core Configuration**. You must select

a core variant that uses the standard sequencer in order for the Altera Modular ADC IP core to work with ADC Toolkit. Also, if an Avalon Master is not connected to the sequencer, you must manually start the sequencer before the ADC Toolkit will work.

**Figure 10-7: Altera Modular ADC Core**



### Starting the ADC Toolkit

The ADC Toolkit can be launched from System Console. Before starting the ADC toolkit, you need to verify that your board is programmed. You can then load your **.sof** by clicking **File > Load Design**. If System Console was started with an active project, your design is auto-loaded when you start System Console.

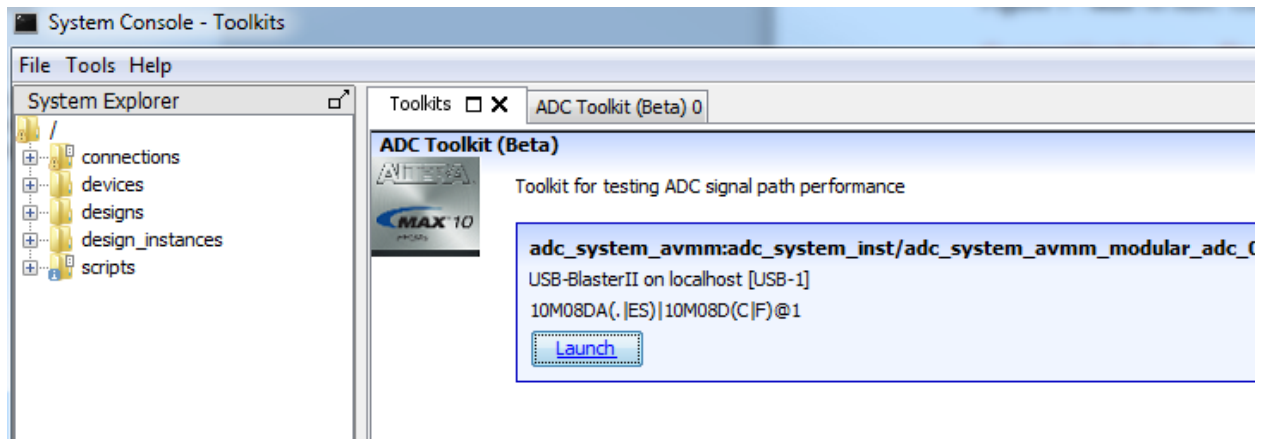
There are two methods to start the ADC Toolkit. Both methods require you to have a MAX 10 device connected, programmed with a project, and linked to this project. However, the **Launch** command only shows up if these requirements are met. You can always start the ADC Toolkit from the **Tools** menu, but if the above requirements are not met, no connection will be made.

- Click **Tools > ADC Toolkit**
- Alternatively, click **Launch** from the **Toolkits** tab. The path for the device is displayed above the **Launch** button.

**Note:** Only one ADC Toolkit enabled device can be connected at a time.

Upon starting the ADC Toolkit, an identifier path on the ADC Toolkit tab shows you which ADC on the device is being used for this instance of the ADC Toolkit.

Figure 10-8: Launching ADC Toolkit




### ADC Toolkit Flow

The ADC Toolkit GUI consists of four panels: **Frequency Selection**, **Scope**, **Signal Quality**, and **Linearity**.

1. Use the **Frequency Selection** panel to calculate the required sine wave frequency for proper signal quality testing. The ADC Toolkit will give you the nearest ideal frequency based on your desired reference signal frequency.
2. Use the **Scope** panel to tune your signal generator or inspect input signal characteristics.
3. Use the **Signal Quality** panel to test the performance of your ADC using industry standard metrics.
4. Use the **Linearity** panel to test the linearity performance of your ADC and display differential and integral non-linearity results.

Figure 10-9: ADC Toolkit GUI



#### Related Information

- [MAX 10 FPGA Device Overview](#)
- [MAX 10 FPGA Device Datasheet](#)
- [MAX 10 FPGA Design Guidelines](#)
- [MAX 10 Analog to Digital Converter User Guide](#)
- [Additional information about sampling frequency](#)

Nyquist sampling theorem and how it relates to the nominal sampling interval required to avoid aliasing.

## Frequency Selection

You use the **Frequency Selection** panel to compute the required reference signal frequency to run the ADC performance tests. The sine wave frequency is critical and affects the validity of your test results.

Figure 10-10: Frequency Selection Panel

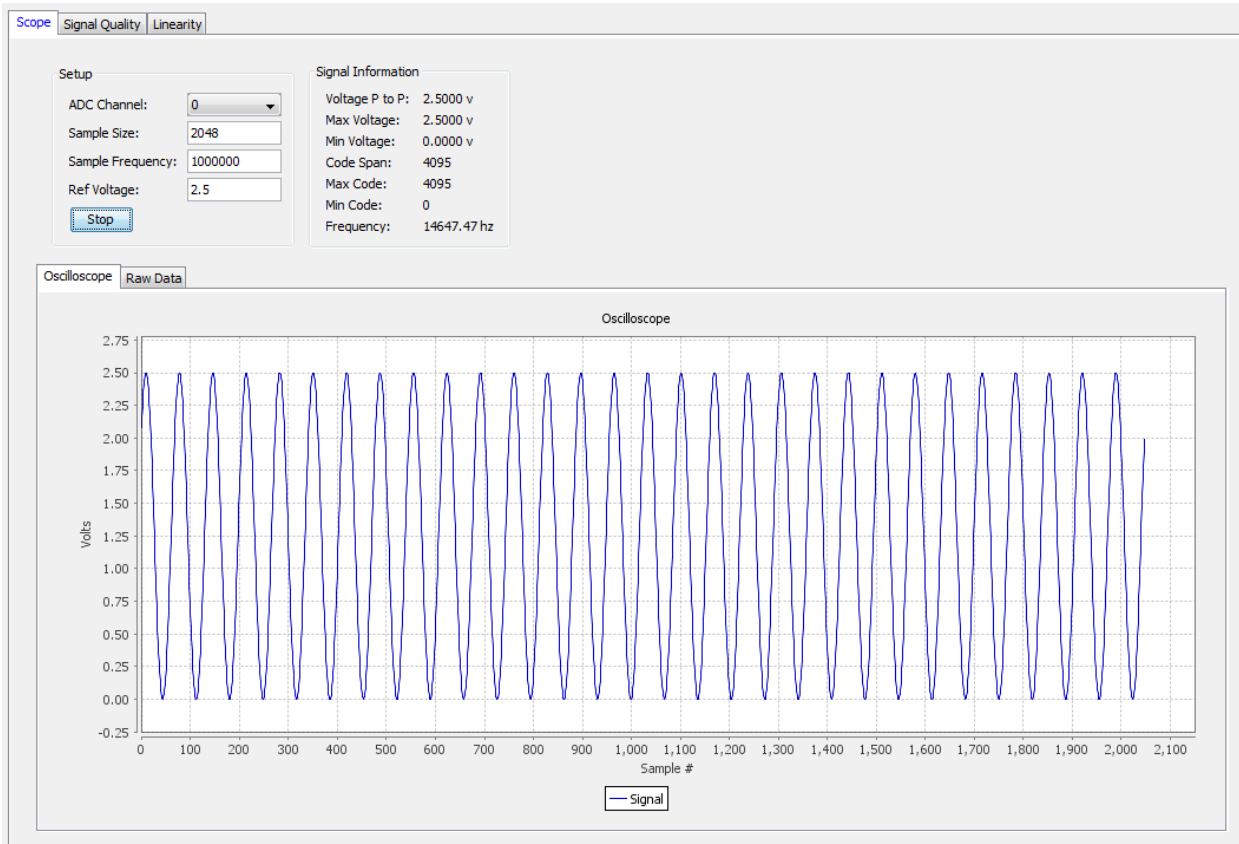
To set the frequency of the reference signal:

1. On **ADC Channel**, select the ADC channel that you plan to test.  
The tool populates the **Sample Size** and **Sample Frequency** fields.
2. Enter the **Desired Frequency**. This is your desired frequency for testing. You need to complete this procedure to calculate the frequency that you set your signal generator to, which will differ depending on the type of test you want to do with the ADC Toolkit.
3. Click **Calculate**.
  - The closest frequency for valid testing near your desired frequency displays under both **Signal Quality Test** and **Linearity Test**.
  - The nearest required sine wave frequencies are different for the signal quality test and linearity test.
4. Set your signal generator to the precise frequency given by the tool based on the type of test you want to run.

## Scope Mode

You use the **Scope** panel to tune your signal generator in order to achieve the best possible performance from the ADC.

Figure 10-11: Scope Mode Panel



To tune your signal generator:

1. On **ADC Channel**, select the ADC channel that you plan to test.
2. Enter your reference **Sample Frequency** (unless the tool can extract this value from your IP).
3. Enter your **Ref Voltage** (unless the tool can extract this value from your IP).
4. Click **Run**.  
The tool will repeatedly capture a buffer worth of data and display the data as a waveform and display additional information under **Signal Information**.
5. Tune your signal generator to use the maximum dynamic range of the ADC without clipping. Avoid hitting 0 or 4095 because your signal will likely be clipping. Look at the displayed sine wave under **Oscilloscope** to see that the top and bottom peaks are evenly balanced to ensure you have selected the optimum value.
  - For MAX 10 devices, you want to get as close to **Min Code = 0** and **Max Code = 4095** without actually hitting those values.
  - The frequency should be set precisely to the value needed for testing such that coherent sampling is observed in the test window. Before moving forward, follow the suggested value for signal quality testing or linearity testing, which is displayed next to the actual frequency that is detected.
  - From the **Raw Data** tab, you can export your data as a **.csv** file.

**Related Information**[Additional information about coherent sampling vs window sampling](#)

## Signal Quality Test Mode

The available performance metrics in signal quality test mode are the following: signal to noise ratio (SNR), total harmonic distortion (THD), spurious free dynamic range (SFDR), signal to noise and distortion ratio (SINAD), effective number of bits (ENOB), and a frequency response graph.


The frequency response graph shows the signal, noise floor, and any spurs or harmonics.

The signal quality parameters are measurements relative to the carrier signal and not the full scale of the ADC.

### Before you begin

Before running a signal quality test, ensure that you have set up the frequency of the reference signal using **Scope** mode.

**Figure 10-12: Signal Quality Panel**



To run a signal quality test:

1. On **ADC Channel**, select the ADC channel that you plan to test.
2. Click **Run**.

From the **Raw Data** tab, you can export your data as a **.csv** file.

For signal quality tests, the signal must be coherently sampled. Based on the sampling rate and number of samples to test, specific input frequencies are required for coherent sampling.

The sample frequency for each channel is calculated based on the ADC sequencer configuration.

#### Related Information

[Additional information about dynamic parameters such as SNR, THD, etc](#)

## Linearity Test Mode

The linearity test determines the linearity of the step sizes of each ADC code. It uses a histogram testing method which requires sinusoidal inputs which are easier to source from signal generators and DACs than other test methods.

When using **Linearity** test mode, your reference signal must meet specific requirements.

- The full code range of the ADC is covered by the signal source. Results improve if the time spent at code ends is equivalent, by tuning the reference signal in **Scope** mode.
- You have to make sure if using code ends that you are not clipping the signal. Look at the signal in **Scope** mode to see that it does not look flat at the top or bottom. It may be desirable to back away from code ends and test a smaller range within the desired operating range of the ADC input signal.
- Choosing a frequency that is not an integer multiple of the sample rate and buffer size helps to ensure all code bins are filled relatively evenly to the probability density function of a sine wave. If an integer multiple is selected, some bins may be skipped entirely while others are over populated. This makes the tests results invalid. Use the frequency calculator feature to determine a good signal frequency near your desired frequency.

To run a linearity test:

1. On **ADC Channel**, select the ADC channel that you plan to test.
2. Enter the test sample size in **Burst Size**. Larger samples increase the confidence in the test results.
3. Click **Run**.
  - You can stop the test at anytime, as well as click **Run** again to continue adding to the aggregate data. To start fresh, click **Reset** after you stop a test. Anytime you change the input signal or channel, you should click **Reset** so your results are correct for a particular input.
  - There are three graphical views of the data: **Histogram** view, **DNL** view, and **INL** view.
  - From the **Raw Data** tab, you can export your data as a **.csv** file.

## Data Views


### Histogram View

The **Histogram** view shows how often each code shows up. The graph updates every few seconds as data is collected. You can use the **Histogram** view to quickly check if your test signal was set up appropriately.

The figure below shows the shape of a pure sine wave signal that was sent to the ADC. Your reference signal should look similar to the example below.




Figure 10-13: Example of Pure Sine Wave Histogram



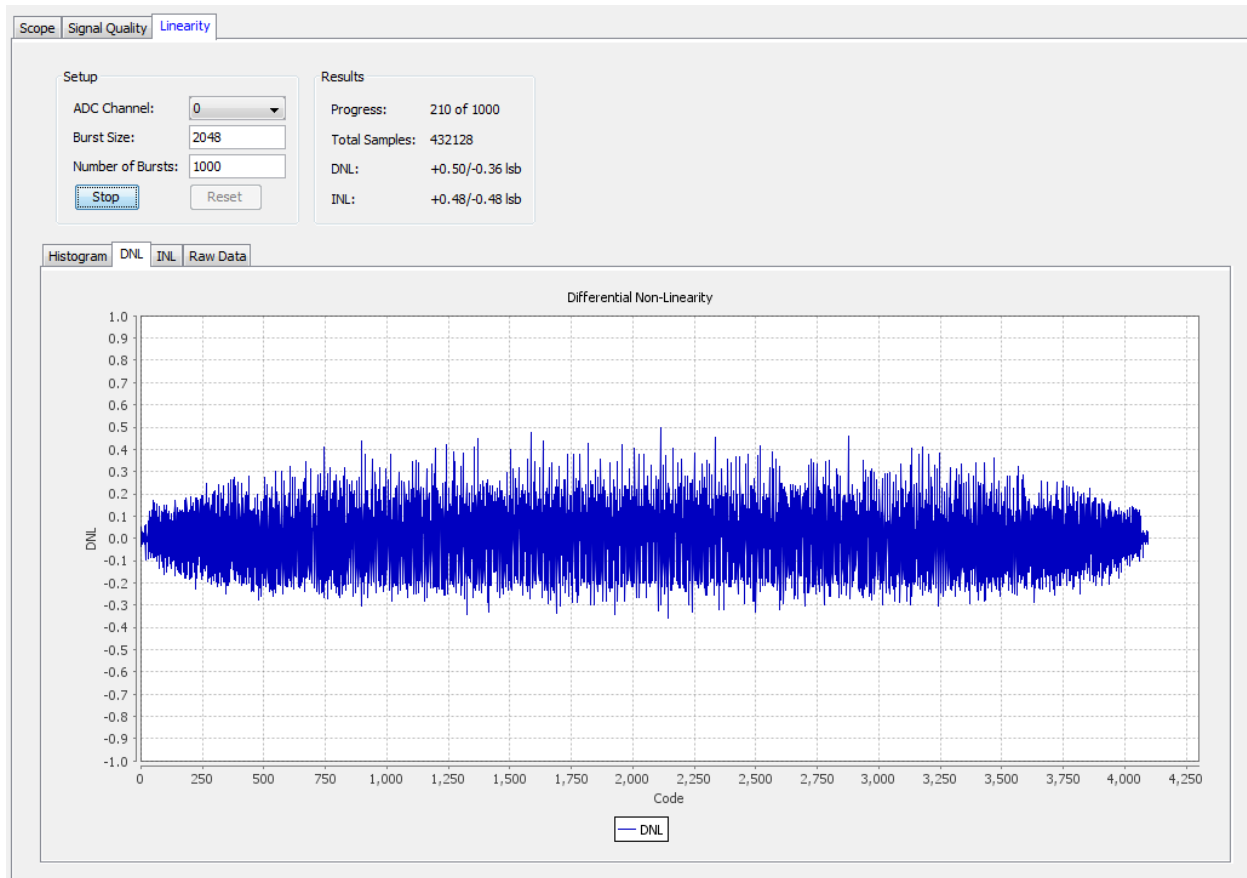
If your reference signal is not a relatively smooth line, but has jagged edges with some bins having a value of 0 while adjacent bins have a much higher value, then the test signal frequency was not adequate for this test. You can use **Scope** mode to help choose a good frequency for linearity testing.

Figure 10-14: Examples of (Left) Poor Frequency Choice vs (Right) Good Frequency Choice



## Differential Non-linearity View

Figure 10-15: Example of Good Differential Non-linearity




The **DNL** view shows all currently collected data. Ideally, you want your data to look like a straight line through the 0 on the x-axis. The line will look rough when there are not enough samples of data, but the line should look better as more data is collected and averaged out.

Each point in the graph represents how many LSB values a particular code differs from the ideal step size of 1 LSB. The highest positive and negative DNL values are listed in the **Results** box.

## Integral Non-linearity View

Figure 10-16: Example of Good Integral Non-linearity



The **INL** view shows all currently collected data. Ideally, with a perfect ADC and enough samples, the graph would look like a straight line through 0 on the x-axis.

Each point in the graph represents how many LSB values a particular code differs from its expected point in the voltage slope. The highest positive and negative INL values are listed in the **Results** box.

## Common ADC Terms

**SNR** is defined as the ratio of the output signal voltage level to the output noise level.

**THD** is defined as the ratio of the sum of powers of the harmonic frequency components to the power of the fundamental/original frequency component.

**SFDR** characterizes the ratio between the fundamental signal and the highest spurious in the spectrum.

**SINAD** is defined as the ratio of the RMS value of the signal amplitude to the RMS value of all other spectral components, including harmonics, but excluding DC.

**ENOB** is the number of bits with which the ADC behaves like a perfect ADC.

**DNL** is defined as the maximum and minimum difference in the step width between actual transfer function and the perfect transfer function.

*INL* is defined as the maximum vertical difference between the actual and the ideal curve. It indicates the amount of deviation of the actual curve from the ideal transfer curve.

## API

### Console Commands

The console commands enable testing. Use console commands to identify a module by its path, and to open and close a connection to it. The `path` that identifies a module is the first argument to most System Console commands.

**Table 10-4: Console Commands**

Command	Arguments	Function
<code>get_service_types</code>	N/A	Returns a list of service types that System Console manages. Examples of service types include master, bytestream, processor, sld, jtag_debug, device, and design.
<code>get_service_paths</code>	<code>&lt;service_type&gt;</code>	Returns a list of paths to nodes that implement the requested service type.
<code>claim_service</code>	<code>&lt;service-type&gt;</code> <code>&lt;service-path&gt;</code> <code>&lt;claim-group&gt;</code> <code>&lt;claims&gt;</code>	Provides finer control of the portion of a service you want to use.  The return value from <code>claim_service</code> is the path of the claimed service which should be used to access and finally close the service.  Run <code>help claim_service</code> to get a <code>&lt;service-type&gt;</code> list.  Then run <code>help claim_service &lt;service-type&gt;</code> to get specific help on that service.
<code>close_service</code>	<code>&lt;service_type&gt;</code> <code>&lt;service_path&gt;</code>	Closes the specified service type at the specified path.
<code>is_service_open</code>	<code>&lt;service_type&gt;</code> <code>&lt;service_path&gt;</code>	Returns 1 if the service type provided by the path is open, 0 if the service type is closed.
<code>get_services_to_add</code>	—	Returns a list of all services that are instantiable with the <code>add_service</code> command.

Command	Arguments	Function
add_service	<service-type> <instance-name> <optional-parameters>	Adds a service of the specified service type with the given instance name. Run <code>get_services_to_add</code> to retrieve a list of instantiable services. This command returns the path where the service was added.  Run <code>help add_service &lt;service-type&gt;</code> to get specific help about that service type, including any parameters that might be required for that service.
add_service dashboard	<name> <title> <menu>	Creates a new GUI dashboard in System Console.
add_service gdbserver	<Processor Service> <port number>	Instantiates a gdbserver.
add_service tcp	<instance_name> <ip_addr> <port number>	Instantiates a tcp service.
add_service transceiver_channel_rx	<data_pattern_checker path> <transceiver path> <transceiver channel address> <reconfig path> <reconfig channel address>	Instantiates a Transceiver Toolkit receiver channel.
add_service transceiver_channel_tx	<data_pattern_generator path> <transceiver path> <transceiver channel address> <reconfig path> <reconfig channel address>	Instantiates a Transceiver Toolkit transmitter channel.

Command	Arguments	Function
add_service transceiver_debug_ link	<transceiver_channel_ tx path>  <transceiver_channel_ rx path>	Instantiates a Transceiver Toolkit debug link.
get_version	—	Returns the current System Console version and build number.
get_claimed_services	<claim-group>	For the given claim group, returns a list of services claimed. The returned list consists of pairs of paths and service types. Each pair is one claimed service.
refresh_connections	—	Scans for available hardware and updates the available service paths if there have been any changes.
send_message	<level>  <message>	Sends a message of the given level to the message window. Available levels are info, warning, error, and debug.

**Related Information**

[Application and Interfaces](#) on page 10-2

## SLD Commands

Table 10-5: SLD Commands

Command	Arguments	Function
sld_access_ir	<service-path>  <ir-value>  <delay> (in $\mu$ s)	Shifts the instruction value into the instruction register of the specified node. Returns the previous value of the instruction.  If the <delay> parameter is non-zero, then the JTAG clock is paused for this length of time after the access.

Command	Arguments	Function
sld_access_dr	<code>&lt;service-path&gt;</code> <code>&lt;size_in_bits&gt;</code> <code>&lt;delay-in-μs&gt;</code> , <code>&lt;list_of_byte_values&gt;</code>	<p>Shifts the byte values into the data register of the SLD node up to the size in bits specified.</p> <p>If the <code>&lt;delay&gt;</code> parameter is non-zero, then the JTAG clock is paused for this length of time after the access.</p> <p>Returns the previous contents of the data register.</p>
sld_lock	<code>&lt;service-path&gt;</code> <code>&lt;timeout-in-milliseconds&gt;</code>	<p>Locks the SLD chain to guarantee exclusive access.</p> <p>Returns 0 if successful. If the SLD chain is already locked by another user, tries for <code>&lt;timeout&gt;</code> ms before throwing a Tcl error. You can use the catch command if you want to handle the error.</p>
sld_unlock	<code>&lt;service-path&gt;</code>	<p>Unlocks the SLD chain.</p>

**Related Information**

[SLD Service](#) on page 10-22

## Design Service Commands

Design service commands load and work with your design at a system level.

**Table 10-6: Design Service Commands**

Command	Arguments	Function
design_load	<code>&lt;quartus-project-path&gt;</code> , <code>&lt;sof-file-path&gt;</code> , or <code>&lt;qpf-file-path&gt;</code>	<p>Loads a model of a Quartus II design into System Console. Returns the design path.</p> <p>For example, if your Quartus II Project File (<b>.qpf</b>) is in <b>c:/projects/loopback</b>, type the following command: <code>design_load {c:\projects\loopback\}</code></p>



Command	Arguments	Function
design_link	<design-instance-path> <device-service-path>	Links a Quartus II logical design with a physical device.  For example, you can link a Quartus II design called <b>2c35_quartus_design</b> to a 2c35 device. After you create this link, System Console creates the appropriate correspondences between the logical and physical submodules of the Quartus II project.
design_extract_debug_files	<design-path> <zip-file-name>	Extracts debug files from a <b>.sof</b> to a zip file which can be emailed to <i>Altera Support</i> for analysis.
design_get_warnings	<design-path>	Gets the list of warnings for this design. If the design loads correctly, then an empty list returns.

**Related Information**

[Design Service](#) on page 10-19

## Device Commands

The device commands provide access to programmable logic devices on your board. Before you use these commands, identify the path to the programmable logic device on your board using the `get_service_paths`.

**Table 10-7: Device Commands**

Command	Arguments	Function
device_download_sof	<service_path> <sof-file-path>	Loads the specified <b>.sof</b> to the device specified by the path.
device_get_connections	<service_path>	Returns all connections which go to the device at the specified path.
device_get_design	<device_path>	Returns the design this device is currently linked to.

**Related Information**

[Device Service](#) on page 10-20

## Avalon-MM Commands

Using the 8, 16, or 32 versions of the `master_read` or `master_write` commands is less efficient than using the `master_write_memory` or `master_read_memory` commands. Master commands can also be used on slave services. If you are working on a slave service, the address field can be a register (if the slave defines register names).<sup>(4)</sup>

**Table 10-8: Avalon-MM Commands**

Command	Arguments	Function
<code>master_write_memory</code>	<code>&lt;service-path&gt;</code> <code>&lt;address&gt;</code> <code>&lt;list_of_byte_values&gt;</code>	Writes the list of byte values, starting at the specified base address.
<code>master_write_8</code>	<code>&lt;service-path&gt;</code> <code>&lt;address&gt;</code> <code>&lt;list_of_byte_values&gt;</code>	Writes the list of byte values, starting at the specified base address, using 8-bit accesses.
<code>master_write_16</code>	<code>&lt;service-path&gt;</code> <code>&lt;address&gt;</code> <code>&lt;list_of_16_bit_words&gt;</code>	Writes the list of 16-bit values, starting at the specified base address, using 16-bit accesses.
<code>master_write_from_file</code>	<code>&lt;service-path&gt;</code> <code>&lt;file-name&gt;</code> <code>&lt;address&gt;</code>	Writes the entire contents of the file through the master, starting at the specified address. The file is treated as a binary file containing a stream of bytes.
<code>master_write_32</code>	<code>&lt;service-path&gt;</code> <code>&lt;address&gt;</code> <code>&lt;list_of_32_bit_words&gt;</code>	Writes the list of 32-bit values, starting at the specified base address, using 32-bit accesses.
<code>master_read_memory</code>	<code>&lt;service-path&gt;</code> <code>&lt;address&gt;</code> <code>&lt;size_in_bytes&gt;</code>	Returns a list of <code>&lt;size&gt;</code> bytes. Read from memory starts at the specified base address.
<code>master_read_8</code>	<code>&lt;service-path&gt;</code> <code>&lt;address&gt;</code> <code>&lt;size_in_bytes&gt;</code>	Returns a list of <code>&lt;size&gt;</code> bytes. Read from memory starts at the specified base address, using 8-bit accesses.

<sup>(4)</sup> Transfers performed in 16- and 32-bit sizes are packed in little-endian format.

Command	Arguments	Function
master_read_16	<service-path> <address> <size_in_multiples_of_16_bits>	Returns a list of <size> 16-bit values. Read from memory starts at the specified base address, using 16-bit accesses.
master_read_32	<service-path> <address> <size_in_multiples_of_32_bits>	Returns a list of <size> 32-bit values. Read from memory starts at the specified base address, using 32-bit accesses.
master_read_to_file	<service-path> <file-name> <address> <count>	Reads the number of bytes specified by <count> from the memory address specified and creates (or overwrites) a file containing the values read. The file is written as a binary file.
master_get_register_names	<service-path>	When a register map is defined, returns a list of register names in the slave.

## JTAG Debug Commands

Table 10-9: JTAG Commands

Command	Arguments	Function
jtag_debug_loop	<service-path> <list_of_byte_values>	Loops the specified list of bytes through a loopback of <code>tdi</code> and <code>tdo</code> of a system-level debug (SLD) node. Returns the list of byte values in the order that they were received. Blocks until all bytes are received. Byte values are given with the 0x (hexadecimal) prefix and delineated by spaces.
jtag_debug_reset_system	<service-path>	Issues a reset request to the specified service. Connectivity within your device determines which part of the system is reset.

## Clock and Reset Signal Commands

Table 10-10: Clock and Reset Commands

Command	Argument	Function
<code>jtag_debug_sample_clock</code>	<code>&lt;service-path&gt;</code>	Returns the value of the clock signal of the system clock that drives the module's system interface. The clock value is sampled asynchronously; consequently, you may need to sample the clock several times to guarantee that it is toggling.
<code>jtag_debug_sample_reset</code>	<code>&lt;service-path&gt;</code>	Returns the value of the <code>reset_n</code> signal of the Avalon-ST JTAG Interface core. If <code>reset_n</code> is low (asserted), the value is 0 and if <code>reset_n</code> is high (deasserted), the value is 1.
<code>jtag_debug_sense_clock</code>	<code>&lt;service-path&gt;</code>	Returns the result of a sticky bit that monitors for system clock activity. If the clock has toggled since the last execution of this command, the bit is 1. Returns <code>true</code> if the bit has ever toggled and otherwise returns <code>false</code> . The sticky bit is reset to 0 on read.

## Dashboard Commands

Dashboard commands create graphical tools that seamlessly integrate into System Console. This section describes the supported dashboard Tcl commands and the properties that you can assign to the widgets on your dashboard. The dashboard allows you to create tools that interact with live instances of an IP core on your device.

Table 10-11: Dashboard Commands

Command	Arguments	Description
<code>dashboard_add</code>	<code>&lt;service-path&gt;</code> <code>&lt;id&gt;</code> <code>&lt;type&gt;</code> <code>&lt;group id&gt;</code>	Adds a specified widget to your GUI dashboard.

Command	Arguments	Description
dashboard_remove	<i>&lt;service-path&gt;</i> <i>&lt;id&gt;</i>	Removes a specified widget from your GUI dashboard.
dashboard_set_property	<i>&lt;service-path&gt;</i> <i>&lt;property&gt;</i> <i>&lt;id&gt;</i> <i>&lt;value&gt;</i>	Sets the specified properties of the specified widget that has been added to your GUI dashboard.
dashboard_get_property	<i>&lt;service-path&gt;</i> <i>&lt;id&gt;</i> <i>&lt;type&gt;</i>	Determines the existing properties of a widget added to your GUI dashboard.
dashboard_get_types	—	Returns a list of all possible widgets that you can add to your GUI dashboard.
dashboard_get_properties	<i>&lt;widget type&gt;</i>	Returns a list of all possible properties of the specified widgets in your GUI dashboard.

**Related Information**

[Dashboard Service](#) on page 10-14

## Specifying Widgets

You can specify the widgets that you add to your dashboard.

**Note:** `dashboard_add` performs a case-sensitive match against the widget type name.

**Table 10-12: Dashboard Widgets**

Widget	Description
group	Adds a collection of widgets and controls the general layout of the widgets.
button	Adds a button.
tabbedGroup	Allows you to group tabs together.
fileChooserButton	Defines button actions.
label	Adds a text string.
text	Displays text.

Widget	Description
textField	Adds a text field.
list	Adds a list.
table	Adds a table.
led	Adds a LED with a label.
dial	Adds the shape of an analog dial.
timeChart	Adds a chart of historic values, with the X-axis of the chart representing time.
barChart	Adds a bar chart.
checkBox	Adds a check box.
comboBox	Adds a combo box.
lineChart	Adds a line chart.
pieChart	Adds a pie chart.

## Customizing Widgets

You can change widget properties. Use `dashboard_set_property` to interact with the widgets you instantiate. This functionality is most useful when you change part of the execution of a callback.

## Assigning Dashboard Widget Properties

The following tables list the various properties that you can apply to the widgets on your dashboard.

**Table 10-13: Properties Common to All Widgets**

Property	Description
enabled	Enables or disables the widget.
expandable	Allows the widget to be expanded.
expandableX	Allows the widget to be resized horizontally if there is space available in the cell where it resides.
expandableY	Allows the widget to be resized vertically if there's space available in the cell where it resides.
maxHeight	If the widget's <code>expandableY</code> is set, this is the maximum height in pixels that the widget can take.

Property	Description
minHeight	If the widget's expandableY is set, this is the minimum height in pixels that the widget can take.
maxWidth	If the widget's expandableX is set, this is the maximum width in pixels that the widget can take.
minWidth	If the widget's expandableX is set, this is the minimum width in pixels that the widget can take.
preferredHeight	The height of the widget if expandableY is not set.
preferredWidth	The width of the widget if expandableX is not set.
toolTip	Implements a mouse-over tooltip.
selected	The value of the checkbox, whether it is selected or not.
visible	Displays the widget.
onChange	Registers a callback function to be called when the value of the box changes.
options	Allows you to list available options.

Table 10-14: button Properties

Property	Description
onClick	A Tcl command to run, usually a <code>proc</code> , every time the button is clicked.
text	The text on the button.

Table 10-15: fileChooserButton Properties

Property	Description
text	The text on the button.
onChoose	A Tcl command to run, usually a <code>proc</code> , every time the button is clicked.
title	The dialog box title.
chooserButtonText	The text of dialog box approval button. By default, it is "Open."

Property	Description
filter	The file filter based on extension. Only one extension is supported. By default, all file names are allowed. The filter is specified as [list filter_description file_extension], for example [list "Text Document (.txt)" ".txt"].
mode	Specifies what kind of files or directories can be selected. The default is "files_only." Possible options are "files_only" and "directories_only."
multiSelectionEnabled	Controls whether multiple files can be selected. False, by default.
paths	Returns a list of file paths selected in the file chooser dialog box. This property is read-only. It is most useful when used within the onclick script or a procedure when the result is freshly updated after the dialog box closes.

**Table 10-16: dial Properties**

Property	Description
max	The maximum value that the dial can show.
min	The minimum value that the dial can show.
tickSize	The space between the different tick marks of the dial.
title	The title of the dial.
value	The value that the dial's needle should mark. It must be between min and max.

**Table 10-17: group Properties**

Property	Description
itemsPerRow	The number of widgets the group can position in one row, from left to right, before moving to the next row.
title	The title of the group. Groups with a title can have a border around them, and setting an empty title removes the border.



**Table 10-18: label Properties**

Property	Description
text	The text to show in the label.

**Table 10-19: led Properties**

Property	Description
color	The color of the LED. The options are: red_off, red, yellow_off, yellow, green_off, green, blue_off, blue, and black.
text	The text to show next to the LED.

**Table 10-20: text Properties**

Property	Description
editable	Controls whether the text box is editable.
htmlCapable	Controls whether the text box can format HTML.
text	The text to show in the text box.

**Table 10-21: timeChart Properties**

Property	Description
labelX	The label for the X axis.
labelY	The label for the Y axis.
latest	The latest value in the series.
maximumItemCount	The number of sample points to display in the historic record.
title	The title of the chart.

**Table 10-22: table Properties**

Property	Description
<b>Table-wide Properties</b>	
columnCount	The number of columns (Mandatory) (0, by default).

Property	Description
rowCount	The number of rows (Mandatory) (0, by default).
headerReorderingAllowed	Controls whether you can drag the columns (false, by default).
headerResizingAllowed	Controls whether you can resize all column widths. (false, by default). Note, each column can be individually configured to be resized by using the columnWidthResizable property.
rowSorterEnabled	Controls whether you can sort the cell values in a column (false, by default).
showGrid	Controls whether to draw both horizontal and vertical lines (true, by default).
showHorizontalLines	Controls whether to draw horizontal line (true, by default).
showVerticalLines	Controls whether to draw vertical line (true, by default).
rowIndex	Current row index. Zero-based. This value affects some properties below (0, by default).
columnIndex	Current column index. Zero-based. This value affects all column specific properties below (0, by default).
cellText	Specifies the text to be filled in the cell specified in the current rowIndex and columnIndex (Empty, by default).
selectedRows	Control or retrieve row selection.
<b>Column-specific Properties</b>	
columnHeader	The text to be filled in the column header.
columnHorizontalAlignment	The cell text alignment in the specified column. Supported types are "leading"(default), "left", "center", "right", "trailing".
columnRowSorterType	The type of sorting method used. This is applicable only if rowSorterEnabled is true. Each column has its own sorting type. Supported types are "string" (default), "int", and "float".
columnWidth	The number of pixels used for the column width.
columnWidthResizable	Controls whether the column width is resizable by you (false, by default).

**Table 10-23: barChart Properties**

Property	Description
title	Chart title.
labelX	X axis label text.
labelY	Y axis label text.
range	Y axis value range. By default, it is auto range. Range is specified in a Tcl list, for example [list lower_numerical_value upper_numerical_value].
itemValue	Item value. Value is specified in a Tcl list, for example list bar_category_str numerical_value.

**Table 10-24: lineChart Properties**

Property	Description
title	Chart title.
labelX	Axis X label text.
labelY	Axis Y label text.
range	Axis Y value range. By default, it is auto range. Range is specified in a Tcl list, for example list lower_numerical_value upper_numerical_value.
itemValue	Item value. Value is specified in a Tcl list, for example list bar_category_str numerical_value.

**Table 10-25: pieChart Properties**

Property	Description
title	Chart title.
itemValue	Item value. Value is specified in a Tcl list, for example list bar_category_str numerical_value.

## Monitor Commands

You can use the Monitor commands to read many Avalon-MM slave memory locations at a regular interval.

Under normal load, the monitor service reads the data after each interval and then calls the callback. If the value you read is timing sensitive, you can use the `monitor_get_read_interval` command to read the exact time between the intervals at which the data was read.

Under heavy load, or with a callback that takes a long time to execute, the monitor service skips some callbacks. If the registers you read do not have side effects (for example, they read the total number of events since reset), skipping callbacks has no effect on your code. The `monitor_read_data` command and `monitor_get_read_interval` command are adequate for this scenario.

If the registers you read have side effects (for example, they return the number of events since the last read), you must have access to the data that was read, but for which the callback was skipped. The `monitor_read_all_data` and `monitor_get_all_read_intervals` commands provide access to this data.

**Table 10-26: Main Monitoring Commands**

Command	Arguments	Function
<code>monitor_add_range</code>	<code>&lt;service-path&gt;</code> <code>&lt;target-path&gt;</code> <code>&lt;address&gt;</code> <code>&lt;size&gt;</code>	<p>Adds a contiguous memory address into the monitored memory list.</p> <p><code>&lt;service path&gt;</code> is the value returned when you opened the service.</p> <p><code>&lt;target-path&gt;</code> argument is the name of a master service to read. The address is within the address space of this service. <code>&lt;target-path&gt;</code> is returned from <code>[lindex [get_service_paths master] n]</code> where <code>n</code> is the number of the master service.</p> <p><code>&lt;address&gt;</code> and <code>&lt;size&gt;</code> are relative to the master service.</p>
<code>monitor_set_callback</code>	<code>&lt;service-path&gt;</code> <code>&lt;Tcl-expression&gt;</code>	<p>Defines a Tcl expression in a single string that will be evaluated after all the memories monitored by this service are read. Typically, this expression should be specified as a Tcl procedure call with necessary argument passed in.</p>

Command	Arguments	Function
monitor_set_interval	<i>&lt;service-path&gt;</i> <i>&lt;interval&gt;</i>	Specifies the frequency of the polling action by specifying the interval between two memory reads. The actual polling frequency varies depending on the system activity. The monitor service will try to keep it as close to this specification as possible.
monitor_get_interval	<i>&lt;service-path&gt;</i>	Returns the current interval set which specifies the frequency of the polling action.
monitor_set_enabled	<i>&lt;service-path&gt;</i> <i>&lt;enable(1)/disable(0)&gt;</i>	Enables and disables monitoring. Memory read starts after this is enabled, and Tcl callback is evaluated after data is read.

Table 10-27: Monitor Callback Commands

Command	Arguments	Function
monitor_add_range	<i>&lt;service-path&gt;</i> <i>&lt;target-path&gt;</i> <i>&lt;address&gt;</i> <i>&lt;size&gt;</i>	Adds contiguous memory addresses into the monitored memory list.  The <i>&lt;target-path&gt;</i> argument is the name of a master service to read. The address is within the address space of this service.
monitor_set_callback	<i>&lt;service-path&gt;</i> <i>&lt;Tcl-expression&gt;</i>	Defines a Tcl expression in a single string that will be evaluated after all the memories monitored by this service are read. Typically, this expression should be specified as a Tcl procedure call with necessary argument passed in.
monitor_read_data	<i>&lt;service-path&gt;</i> <i>&lt;target-path&gt;</i> <i>&lt;address&gt;</i> <i>&lt;size&gt;</i>	Returns a list of 8-bit values read from the most recent values read from device. The memory range specified must be the same as the monitored memory range as defined by <code>monitor_add_range</code> .

Command	Arguments	Function
monitor_read_all_data	<service-path> <target-path> <address> <size>	Returns a list of 8-bit values read from all recent values read from device since last Tcl callback. The memory range specified must be within the monitored memory range as defined by monitor_add_range.
monitor_get_read_interval	<service-path> <target-path> <address> <size>	Returns the number of milliseconds between last two data reads returned by monitor_read_data.
monitor_get_all_read_intervals	<service-path> <target-path> <address> <size>	Returns a list of intervals in milliseconds between two reads within the data returned by monitor_read_all_data.
monitor_get_missing_event_count	<service-path>	Returns the number of callback events missed during the evaluation of last Tcl callback expression.

**Related Information**

[Monitor Service](#) on page 10-20

## Processor Commands

Table 10-28: Processor Commands

Command <sup>(5)</sup>	Arguments	Function
processor_download_elf	<service-path> <elf-file-path>	Downloads the given Executable and Linking Format File (.elf) to memory using the master service associated with the processor. Sets the processor's program counter to the .elf entry point.
processor_in_debug_mode	<service-path>	Returns a non-zero value if the processor is in debug mode.

<sup>(5)</sup> If your system includes a Nios II/f core with a data cache, it may complicate the debugging process. If you suspect the Nios II/f core writes to memory from the data cache at nondeterministic intervals; thereby, overwriting data written by the System Console, you can disable the cache of the Nios II/f core while debugging.

Command <sup>(5)</sup>	Arguments	Function
processor_reset	<service-path>	Resets the processor and places it in debug mode.
processor_run	<service-path>	Puts the processor into run mode.
processor_stop	<service-path>	Puts the processor into debug mode.
processor_step	<service-path>	Executes one assembly instruction.
processor_get_register_names	<service-path>	Returns a list with the names of all of the processor's accessible registers.
processor_get_register	<service-path> <register_name>	Returns the value of the specified register.
processor_set_register	<service-path> <register_name> <value>	Sets the value of the specified register.

**Related Information**

[Nios II Processor Example](#) on page 10-18

## Bytestream Commands

Table 10-29: Bytestream Commands

Command	Arguments	Function
bytestream_send	<service-path> <values>	Sends the list of bytes to the specified bytestream service. Values argument is the list of bytes to send.

<sup>(5)</sup> If your system includes a Nios II/f core with a data cache, it may complicate the debugging process. If you suspect the Nios II/f core writes to memory from the data cache at nondeterministic intervals; thereby, overwriting data written by the System Console, you can disable the cache of the Nios II/f core while debugging.

Command	Arguments	Function
bytestream_receive	<i>&lt;service-path&gt;</i> <i>&lt;length&gt;</i>	Returns a list of bytes currently available in the specified services receive queue, up to the specified limit. Length argument is the maximum number of bytes to receive.

**Related Information**

[Bytestream Service](#) on page 10-21

## In-System Sources and Probes Commands

**Note:** The valid values for probe claims include `read_only`, `normal`, and `exclusive`.

**Table 10-30: In-System Sources and Probes Commands**

Command	Arguments	Function
issp_get_instance_info	<i>&lt;service-path&gt;</i>	Returns a list of the configurations of the In-System Sources and Probes instance, including:  instance_index instance_name source_width probe_width
issp_read_probe_data	<i>&lt;service-path&gt;</i>	Retrieves the current value of the probe input. A hex string is returned representing the probe port value.
issp_read_source_data	<i>&lt;service-path&gt;</i>	Retrieves the current value of the source output port. A hex string is returned representing the source port value.
issp_write_source_data	<i>&lt;service-path&gt;</i> <i>&lt;source-value&gt;</i>	Sets values for the source output port. The value can be either a hex string or a decimal value supported by the System Console Tcl interpreter.

**Related Information**

[In-System Sources and Probes Service](#) on page 10-23



## Deprecated Commands

The table lists commands that have been deprecated. These commands are currently supported, but are targeted for removal from System Console.

**Table 10-31: Deprecated Commands**

Command	Arguments	Function
<code>open_service</code>	<code>&lt;service_type&gt;</code> <code>&lt;service_path&gt;</code>	<p>Opens the specified service type at the specified path.</p> <p>Calls to <code>open_service</code> may be replaced with calls to <code>claim_service</code> providing that the return value from <code>claim_service</code> is stored and used to access and close the service.</p>

## Document Revision History

**Table 10-32: Document Revision History**

Date	Version	Changes
December 2014	14.1.0	<ul style="list-style-type: none"> <li>Added overview and procedures for using ADC Toolkit on MAX 10 devices.</li> <li>Added overview for using MATLABS/Simulink Environment with System Console for system verification.</li> </ul>
June 2014	14.0.0	Updated design examples for the following: board bring-up, dashboard service, Nios II processor, design service, device service, monitor service, bytestream service, SLD service, and ISSP service.
November 2013	13.1.0	Re-organization of sections. Added high-level information with block diagram, workflow, SLD overview, use cases, and example Tcl scripts.
June 2013	13.0.0	Updated Tcl command tables. Added board bring-up design example. Removed SOPC Builder content.
November 2012	12.1.0	Re-organization of content.
August 2012	12.0.1	Moved Transceiver Toolkit commands to Transceiver Toolkit chapter.
June 2012	12.0.0	Maintenance release. This chapter adds new System Console features.

Date	Version	Changes
November 2011	11.1.0	Maintenance release. This chapter adds new System Console features.
May 2011	11.0.0	Maintenance release. This chapter adds new System Console features.
December 2010	10.1.0	Maintenance release. This chapter adds new commands and references for Qsys.
July 2010	10.0.0	Initial release. Previously released as the System Console User Guide, which is being obsoleted. This new chapter adds new commands.

For previous versions of the *Quartus II Handbook* , refer to the Quartus II Handbook Archive.

**Related Information**

[Quartus II Handbook Archive](#)

# Debugging Transceiver Links 11

2014.12.15

QI15V3



Subscribe



Send Feedback

This chapter describes using the Transceiver Toolkit to optimize high-speed serial links in your board design. The Transceiver Toolkit provides real-time control, monitoring, and debugging of the transceiver links running on your board.

You can control the transmitter or receiver channels to optimize transceiver settings and hardware features. The toolkit tests bit-error rate (BER) while running multiple links at the target data rate. You can also run auto sweep tests to identify the best physical media attachment (PMA) settings for each link. An EyeQ graph displays the receiver horizontal and vertical eye margin during testing. The toolkit supports testing of multiple devices across one or more boards simultaneously.

**Note:** The Transceiver Toolkit is not a stand-alone application.

## Starting the Transceiver Toolkit

The Transceiver Toolkit GUI helps you to easily visualize and debug transceiver links in your design. To launch the GUI, click **Tools > System Debugging Tools > Transceiver Toolkit**.

Alternatively, you can run Tcl scripts from the command-line.

```
system-console --script=<name of script>
```

## Transceiver Toolkit GUI

- The **System Explorer** displays the main components and hardware connections for your design.
- The **Channel Manager** contains three tabs for the transmitter channels, receiver channels, and transceiver links. From the **Channel Manager**, you can control and view the status of multiple channels simultaneously.
- The **Tcl Console** supports scripting control of Transceiver Toolkit.
- The **Messages** pane displays informational, warning, and error messages about Transceiver Toolkit processes.

Additional information about the **System Explorer** and **Messages** panes can be found in the [Analyzing and Debugging Designs with System Console](#).

## Quick Start

Get started quickly by downloading Transceiver Toolkit design examples from the [On-Chip Debugging Design Examples](#) website. For an online demonstration of how to use the Transceiver Toolkit to run a high-speed link test with one of the design examples, refer to the [Transceiver Toolkit Online Demo](#) on the Altera website.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



## Transceiver Debugging Overview

Testing transceiver links involves configuring your system for debug, and then running various link tests.

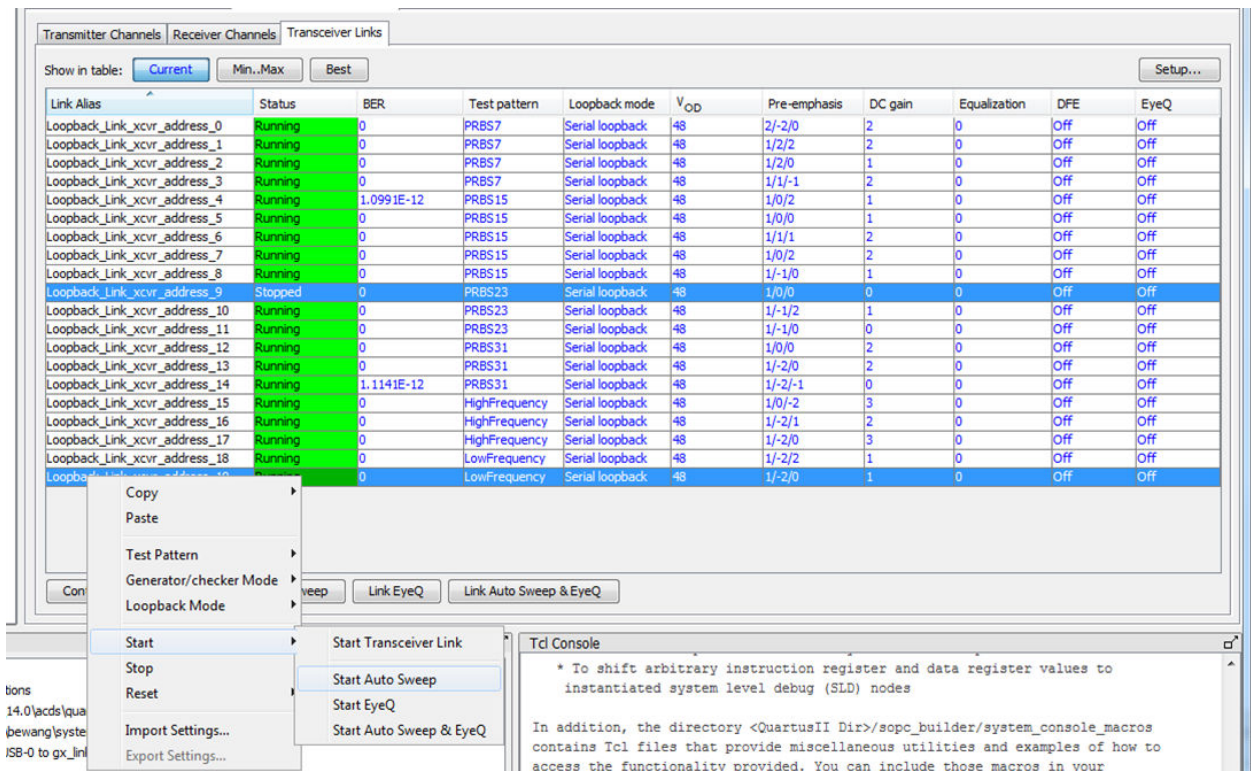
**Table 11-1: Transceiver Link Debugging Flow**

	Flow Description
System Configuration Steps	<ol style="list-style-type: none"> <li>1. Use one of the following methods to define a system that includes necessary transceiver debugging components: <ul style="list-style-type: none"> <li>• Click <b>Tools</b> &gt; <b>Qsys</b> and modify Altera design examples.</li> <li>• Use the <b>IP Catalog</b> and Parameter Editor to define and integrate debugging components into your own design.</li> </ul> </li> <li>2. Click <b>Assignments</b> &gt; <b>Pin Planner</b> to assign device I/O pins to match your device and board.</li> <li>3. Click <b>Processing</b> &gt; <b>Start Compilation</b> to compile your design.</li> <li>4. Connect your target device to Altera programming hardware.</li> <li>5. Click <b>Tools</b> &gt; <b>Programmer</b> to program your target device.</li> </ol>
Link Debugging Steps	<ol style="list-style-type: none"> <li>1. Click <b>File</b> &gt; <b>Load Design</b>, and select the SRAM Object File (<b>.sof</b>) generated for your transceiver design. <ul style="list-style-type: none"> <li>• If you start the toolkit from Quartus II while a project is open, your project will auto-load into the toolkit.</li> </ul> </li> <li>2. (Optional) Create additional links between transmitter and receiver channels.</li> <li>3. Run any of the following tests: <ul style="list-style-type: none"> <li>• Run BER with various combinations of PMA settings.</li> <li>• Run PRBS Signal eye tests.</li> <li>• Run custom traffic tests.</li> <li>• Run link optimization tests.</li> <li>• Directly control PMA analog settings to experiment with settings while the link is running.</li> </ul> </li> </ol>

## Channel Manager

The **Channel Manager** allows you to configure and control large numbers of channels using a table view. You can view all the PMA and sweep settings for all channels. From the **Channel Manager**, you can copy/paste and import/export settings to and from channels. You can also start and stop sweeps for any or all channels. Right-click in the **Channel Manager** to view pop-up menus with additional commands for interacting with channels. The columns in the **Channel Manager** are movable, resizable, and sortable.

Figure 11-1: Channel Manager GUI



### Copying and Pasting Settings

You can copy PMA and/or sweep settings from a selected row. You can paste PMA and/or sweep settings to one or more rows.

### Importing and Exporting Settings

You can select a row in the **Channel Manager** to export your PMA settings to a text file. You can then select one or more rows in the **Channel Manager** to apply the PMA settings from a text file. The PMA settings in the text file apply to a single channel. When you import the PMA settings from a text file, you are duplicating one set of PMA settings for all selected channels.

### Starting and Stopping Tests

The **Channel Manager** gives you the flexibility to start and stop test from the right-click menus. You can select several rows in the **Channel Manager** to start or stop test for multiple channels.

## Display Modes

The three display modes are **Current**, **Min/Max**, and **Best**. The default display mode is **Current**.

- **Current**—shows the current values from the device. The blue color text indicates that the settings are live.
- **Min/Max**—shows the minimum and maximum values to be used in the auto sweep.
- **Best**—shows the best tested values from the last completed auto sweep run.

**Note:** The **Transmitter Channels** tab only shows the **Current** display mode. Auto sweep cannot be performed on only a transmitter channel; a receiver channel is required to perform an auto sweep test.

## Creating Links

The toolkit automatically creates links when a receiver and transmitter share a transceiver channel. You can also manually create and delete links between transmitter and receiver channels. You create links in the **Setup** dialog.

### Setup Dialog

Click **Setup** from the **Channel Manager** to open the **Setup** dialog box.

**Table 11-2: Setup Dialog Popup Menu**

Command Name	Action When Clicked	Enabled If
<b>Edit Transmitter Alias</b>	Starts the inline edit of the alias of the selected row.	Only enabled if one row is selected.
<b>Edit Receiver Alias</b>	Starts the inline edit of the alias of the selected row.	Only enabled if one row is selected.
<b>Edit Transceiver Link Alias</b>	Starts the inline edit of the alias of the selected row.	Only enabled if one row is selected.
<b>Copy</b>	Copies the text of the selected row(s) to the clipboard. The text copied depends on the column clicked on. The text copied to the clipboard is newline delimited.	Enabled if one or more rows are selected.

## Controlling Transceiver Channels

You can directly control and monitor transmitters, receivers, and links running on the board in real time. You can transmit a data pattern across the transceiver link, and then report the signal quality of the received data in terms of bit error rate or eye margin with EyeQ.

Click **Control Transmitter Channel** (**Transmitter Channels** tab), **Control Receiver Channel** (**Receiver Channels** tab), or **Control Transceiver Link** (**Transceiver Links** tab) to adjust transmitter or receiver settings while the channels are running.

You can use the GUI or the right-click popup menus to execute commands.

## Auto Sweep Testing

Use the auto sweep feature to automatically sweep ranges for the best transceiver PMA settings. You can store a history of the test runs and keep a record of the best PMA settings. You can use the best found settings in your final design for improved signal integrity compared with the default settings.

## Adaptive Equalization Control

Adaptive equalization (AEQ) automatically evaluates and selects the best combination of reconfiguration equalizer settings for the receiver. AEQ continuously evaluates and changes the settings for current conditions. You can use AEQ for multiple, independently controlled receiver channels.

Enable this feature by selecting **One-time adaptation** for the **Equalization mode** receiver setting.

## Signal Eye Margin Testing

Some Altera devices include EyeQ circuitry that allows visualization of the horizontal and vertical eye margin at the receiver. For supported devices, use signal eye tests to tune the PMA settings of your transceiver. This results in the best eye margin and BER at high data rates. This GUI is disabled for unsupported devices.

The EyeQ graph can display a bathtub curve, eye diagram representing eye margin, or heat map display. The run list displays the statistics of each EyeQ test. When PMA settings are suitable, the bathtub curve is wide, with sharp slopes near the edges. The curve is up to 30 units wide. If the bathtub is narrow, then the signal quality is poor. The wider the bathtub curve, the wider the eye. The smaller the bathtub curve, the smaller the eye. The eye contour shows the estimated horizontal and vertical eye opening at the receiver.

You can right-click any of the test runs in the list, and then click **Apply Settings to Device** to quickly apply those PMA setting to your device. You can also click **Export**, **Import**, or **Create Report**.

Figure 11-2: EyeQ Settings and Status Showing Results of Two Test Runs

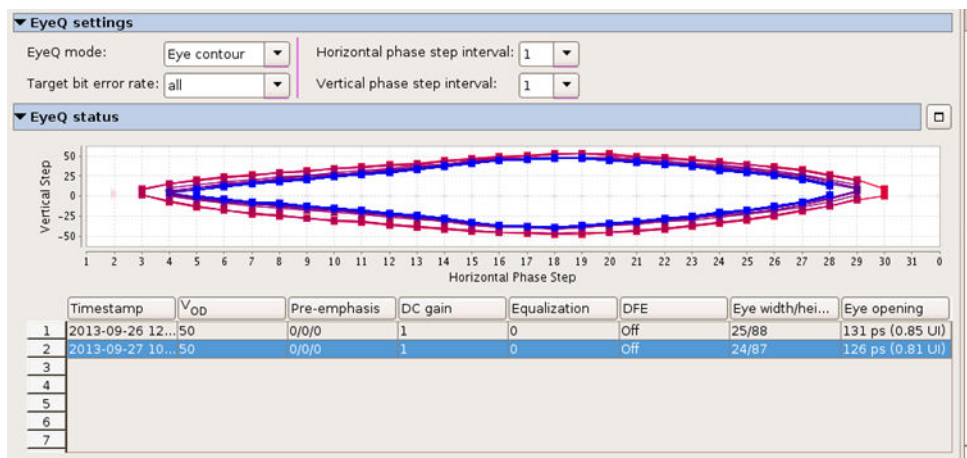
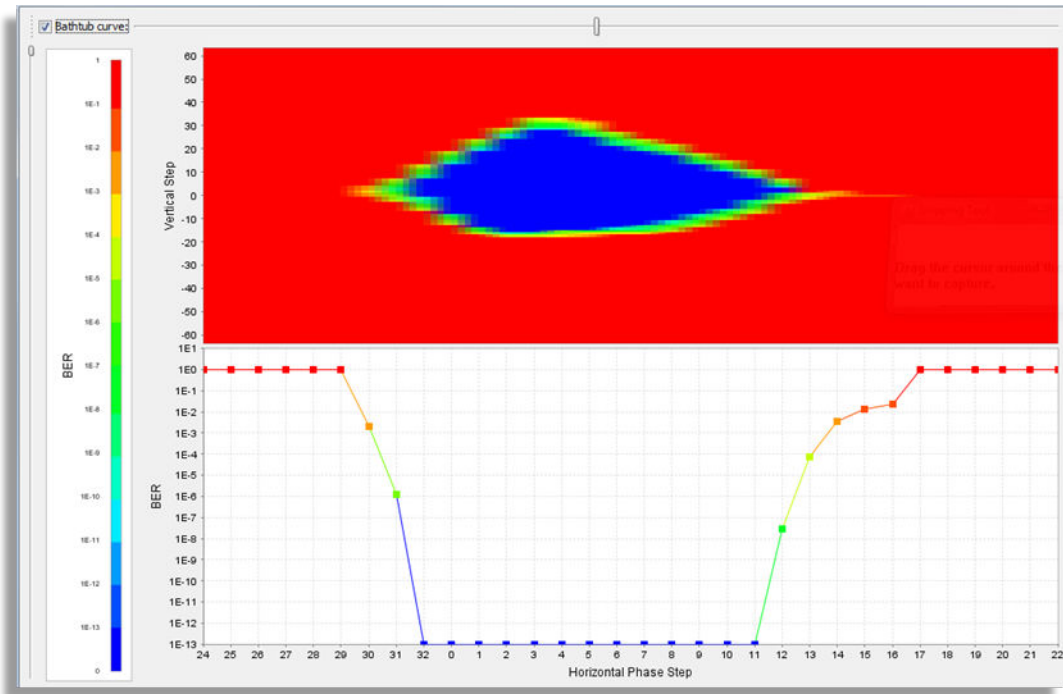


Figure 11-3: Heat Map Display and Bathtub Curve Through Eye



## Serial Bit Comparator Mode

**Serial bit comparator** mode allows you to run EyeQ diagnostic features with any PRBS patterns or user-design data, without disrupting the data path.

To enable this mode you must enable the following debugging component options when configuring the debugging system:

**Table 11-3: Component Settings for Serial Bit Comparator Mode**

Debugging Component	Setting for Serial Bit Mode <sup>(6)</sup>
Transceiver Reconfiguration Controller	Turn on <b>Enable EyeQ block</b> and <b>Enable Bit Error Rate Block</b>
Data Pattern Generator <sup>(7)</sup>	Turn on <b>Enable Bypass interface</b>

**Serial bit comparator** mode is less accurate than **Data pattern checker** mode for single bit error checking. Do not use **Serial bit comparator** mode if you require an exact error rate. Use the **Serial bit comparator** mode for checking a large window of error.

The bit error counter is not read in real-time because it is read through the memory-mapped interface.

<sup>(6)</sup> Settings in [Table 11-3](#) are supported in Stratix V devices only.

<sup>(7)</sup> Limited support for Data Pattern Generator or data pattern in Serial Bit Mode.



Serial bit comparator mode has the following hardware limitations for Stratix V devices:

- The serial bit checker can only be used on a single channel per reconfiguration controller at a time.
- When the serial bit checker is running on channel  $n$ , only the  $V_{OD}$ , pre-emphasis, DC gain, and EyeQ settings on that channel can be changed. Changing or enabling DFE or CTLE can cause corruption of the serial bit checker results.
- When the serial bit checker is running on a channel, no settings on any other channel on the same reconfiguration controller can be changed.
- When the serial bit checker is running on a channel, no other channel should be opened in the Transceiver Toolkit.
- When the serial bit checker is running on a channel, copying PMA settings from any channel on the same reconfiguration controller is unsupported.

## Scripting Support

You can alternatively use Tcl commands to access Transceiver Toolkit functions, rather than using the GUI. You can script various tasks, such as loading a project, creating design instances, linking device resources, and identifying high-speed serial links. You can save your project setup in a Tcl script for use in subsequent testing sessions. You can also build a custom test routine script.

After you set up and define links that describe the entire physical system, you can click Save Tcl Script to save the setup for future use. To run the scripts, double-click script names in the **System Explorer** scripts folder.

### Related Information

[Scripting API](#) on page 11-31

## Arria 10 Support and Limitations

The following are details about support and limitations for using Arria 10 devices with the Transceiver Toolkit.

Quartus II software version 14.1

- The JTAG Debug Link no longer needs to be instantiated in Qsys project. It will be auto-instantiated.
- The data rate may sometimes show the value 0 Mbps. You can click the (refresh) button next to the displayed data rate to override this.

Quartus II software version 14.1a10s

- All the features in the Transceiver Toolkit support Arria 10, with the exception of EyeQ.
- Full support for the Arria 10 hardended PRBS generator and checker. You must enable accumulators and capability registers.
- Arria 10 DFE supports two modes, a 3-tap mode and a 7-tap mode. You can switch between the two modes via a checkbox.
- Support for setting the pre-emphasis second pre-tap.

## Configuring Systems for Debug

To debug transceivers, you must first configure a system that includes the appropriate Altera IP core(s) that supports each debugging operation. You can create such a system by either modifying an Altera design example, or by integrating debugging components into your own design.

You can quickly parameterize the debugging components by using the **IP Catalog** and Parameter Editor.

**Table 11-4: Transceiver Toolkit IP Core Configuration**

Component	Debugging Functions	Parameterization Notes
Transceiver Native PHY	Supports all debugging functions	<ul style="list-style-type: none"> <li>If <b>Enable 10G PCS</b> is enabled, <b>10G PCS protocol mode</b> must be set to <b>basic</b> on the <b>10G PCS</b> tab.</li> </ul>
Custom PHY	Test all possible transceiver parallel data widths	<ul style="list-style-type: none"> <li>Set lanes, group size, serialization factor, data rate, and input clock frequency to match your application.</li> <li>Turn on <b>Avalon data interfaces</b>.</li> <li>Disable <b>8B/10B</b>.</li> <li>Set <b>Word alignment mode</b> to <b>manual</b>.</li> <li>Disable <b>rate match FIFO</b>.</li> <li>Disable <b>byte ordering block</b>.</li> </ul>
Low Latency PHY	Test at more than 8.5 Gbps in GT devices or use of PMA direct mode (such as when using six channels in one quad)	<ul style="list-style-type: none"> <li>Set <b>Phase compensation FIFO mode</b> to <b>EMBEDDED</b> above certain data rates and set to <b>NONE</b> for PMA direct mode (Stratix IV designs only).</li> <li>Turn on <b>Avalon data interfaces</b>.</li> <li>Set serial loopback mode to enable serial loopback controls in the toolkit.</li> </ul>
Avalon-ST Data Pattern Generator	Generates standard data test patterns at Avalon-ST source ports	<ul style="list-style-type: none"> <li>Select PRBS7, PRBS15, PRBS23, PRBS31, high frequency, or low frequency patterns.</li> </ul>
Avalon-ST Data Pattern Checker	Validates incoming data stream against test patterns accepted on Avalon streaming sink ports	<ul style="list-style-type: none"> <li>Specify a value for <code>ST_DATA_W</code> that matches the FPGA-fabric interface width.</li> </ul>

Component	Debugging Functions	Parameterization Notes
Reconfiguration Controller	Supports PMA control and other transceiver settings	<ul style="list-style-type: none"> <li>• Connect the reconfiguration controller to all PHYs that you want controlled by the toolkit.</li> <li>• Connect <code>reconfig_from_xcvr</code> to <code>reconfig_to_xcvr</code>.</li> <li>• Enable Analog controls.</li> <li>• Enable EyeQ block (Stratix V devices only).</li> <li>• Enable AEQ block (Stratix V devices only).</li> <li>• Enable DFE block (Stratix V devices only).</li> </ul>
JTAG to Avalon Master Bridge	Accepts encoded streams of bytes with transaction data and initiates Avalon-MM transactions	N/A
<b>Arria 10 Implementation</b>		
Altera Debug Master Endpoint (ADME)	<ul style="list-style-type: none"> <li>• Supports control of PMA analog settings, ADCE settings, DFE settings, and EyeQ.</li> <li>• Can discover PHY and use toolkit on designs not using Qsys.</li> <li>• RBC support, but invalid settings are not blocked from being applied.</li> <li>• Optionally, you can turn off to save resource count.</li> </ul>	<ul style="list-style-type: none"> <li>• Enabled from the Arria 10 Transceiver Native PHY Parameter Editor, Dynamic Reconfiguration tab.</li> <li>• Enable Shared reconfiguration interface.</li> </ul>
JTAG Debug Link	Required for Arria 10	For Qsys projects, the JTAG Debug Link is auto-instantiated.
Transceiver PHY Reset Controller	Required for Arria 10	N/A

**Related Information**

[Integrating Debug Components In Your Design](#) on page 11-12

## Adapting Altera Design Examples

Altera provides design examples to help you quickly test your own design. You can experiment with these designs and modify them for your own application. Refer to the **readme.txt** of each design example for more information. Download the Transceiver Toolkit design examples from the On-Chip Debugging Design Examples page of the Altera website.

You can use the design examples as a starting point to work with a particular signal integrity development board. The design examples provide the components to quickly test the functionality of the receiver and transmitter channels in your design. You can easily change the transceiver settings in the design examples to see how they affect your transceiver link performance. You can isolate and verify the high-speed serial links without debugging other logic in your design. You can modify and customize the design examples to match your intended transceiver design.

Once you have downloaded the design examples, open the Quartus II software and restore the design example project archive. If you have access to the same development board with the same device as mentioned in the **readme.txt** file of the example, you can directly program the device with the provided programming file in that example. If you want to recompile the design, you must make your modifications to the configuration in Qsys, regenerate in Qsys, and recompile the design in the Quartus II software to generate a new programming file.

If you have the same board as mentioned in the **readme.txt** file, but a different device on your board, you must choose the appropriate device and recompile the design. For example, some early development boards are shipped with engineering sample devices.

You can make changes to the design examples so that you can use a different development board or a different device. If you have a different board, you must edit the necessary pin assignments and recompile the design examples.

#### Related Information

- [On-Chip Debugging Design Examples](#)

## Modifying Design Examples

You can adapt an Altera design example to experiment with various configurations that match your own design. For example, you can change data rate, number of lanes, PCS-PMA width, FPGA-fabric interface width, or input reference clock frequency. To modify the design examples, you modify the IP core parameters and regenerate the system in Qsys. Next, you modify the top-level design file, and reassign device I/O pins as necessary.

To modify a design example PHY block to match your design, follow these steps:

1. Determine the number of channels required by your design.
2. Open the *<project name>.qpf* for the design example in the Quartus II software.
3. Click **Tools** > **Qsys**.
4. On the **System Contents** tab, right-click the PHY block and click **Edit**. Specify options for the PHY block to match your design requirement for number of lanes, data rate, PCS-PMA width, FPGA-fabric interface width, and input reference clock frequency.
5. Specify a multiple of the FPGA-fabric interface data width for **Avalon Data Symbol Size**. The available values are 8 or 10. Click **Finish**.
6. Delete any timing adapter from the design. The timing adapters are not required.
7. From the **IP Catalog**, add one **data pattern generator** and **data pattern checker** for each transmitter and receiver lane.
8. Right-click **data pattern generator** and click **Edit**. Specify a value for ST\_DATA\_W that matches the FPGA-fabric interface width.
9. Right-click **data pattern checker** and click **Edit**. Specify a value for ST\_DATA\_W that matches the FPGA-fabric interface width.
10. From the **IP Catalog**, add a **Transceiver Reconfiguration Controller**.

11. Right-click **Transceiver Reconfiguration Controller** and click **Edit**. Specify 2\* number of lanes for the number of reconfigurations interfaces. Click **finish**.
12. Create connections for the data pattern generator and data pattern checker components. Right-click the net name in the **System Contents** tab and specify the following connections.

From		To	
Block Name	Net Name	Block Name	Net Name
clk_100	clk	data_pattern_generator	csr_clk
clk_100	clk_reset	data_pattern_generator	csr_clk_reset
master_0	master	data_pattern_generator	csr_slave
xcvr_*_phy_0	tx_clk_out0	data_pattern_generator	pattern_out_clk
xcvr_*_phy_0	tx_parallel_data0	data_pattern_generator	pattern_out
clk_100	clk	data_pattern_checker	csr_clk
clk_100	clk_reset	data_pattern_checker	csr_clk_reset
master_0	master	data_pattern_checker	csr_slave
xcvr_*_phy_0	rx_clk_out0	data_pattern_checker	pattern_in_clk
xcvr_*_phy_0	rx_parallel_data0	data_pattern_checker	pattern_in

13. Click **System** > **Assign Base Addresses**.

14. Connect the reset port of timing adapters to `clk_reset` of `clk_100`.

15. To implement the changes to the system, click **Generate** > **Generate HDL**.

16. If you modify the number of lanes in the PHY, you must update the top-level file accordingly. The following example shows Verilog HDL code for a two-channel design that declares input and output ports in the top-level design. The example design includes the low latency PHY IP core. If you modify the PHY parameters, you must modify the top-level design with the correct port names. Qsys displays an example of the PHY, click **Generate** > **HDL Example**.

```

module low_latency_10g_1ch DUT (
    input  wire GXB_RXL11,
    input  wire GXB_RXL12,
    output wire GXB_TXL11,
    output wire GXB_TXL12
);
.....

low_latency_10g_1ch DUT (
    .....
    .xcvr_low_latency_phy_0_tx_serial_data_export    ({GXB_TXL11,
GXB_TXL12}),
    .xcvr_low_latency_phy_0_rx_serial_data_export    ({GXB_RXL11,
GXB_TXL12}),
    .....
);

```

17. From Quartus II, click **Assignments** > **Pin Planner** and update pin assignments to match your board.

18. Edit the design's Synopsys Design Constraints (**.sdc**) to reflect the reference clock change. Ignore the reset warning messages.

19. Recompile the design.

## How to Use Internal PLL to Generate the `reconfig_clk`

You can use an internal PLL to generate the `reconfig_clk` by changing the Qsys connections so that offset cancellation is delayed until the generated clock is stable.

- If you do not have a free running clock that is within the required frequency range of the reconfiguration clock, a PLL can be added to the top-level of the design example. The frequency range varies depending on the device family. See the data sheet for your device.
- When an internal PLL is used, you need to hold off offset cancellation until the generated clock is stable. You do this by connecting the `pll_locked` signal of the internal PLL to the `.clk_clk_in_reset_n` port of the Qsys system, instead of the `system_reset` signal.
- The filter logic, inverter, and synchronization to the `reconfig_clk` should be implemented outside the Qsys system using user-created logic.

You can find the [support solution](#) in the Altera Knowledge Base. The solution applies to only Arria V, Cyclone V, Stratix IV GX/GT, and Stratix V.

## Integrating Debug Components In Your Design

This section describes integrating debugging IP components into your own design. You can integrate transceiver debugging components into your design, rather than modifying the Altera Debugging Design Examples.

### Configuring BER Tests

To integrate components with your design for BER testing, follow these steps:

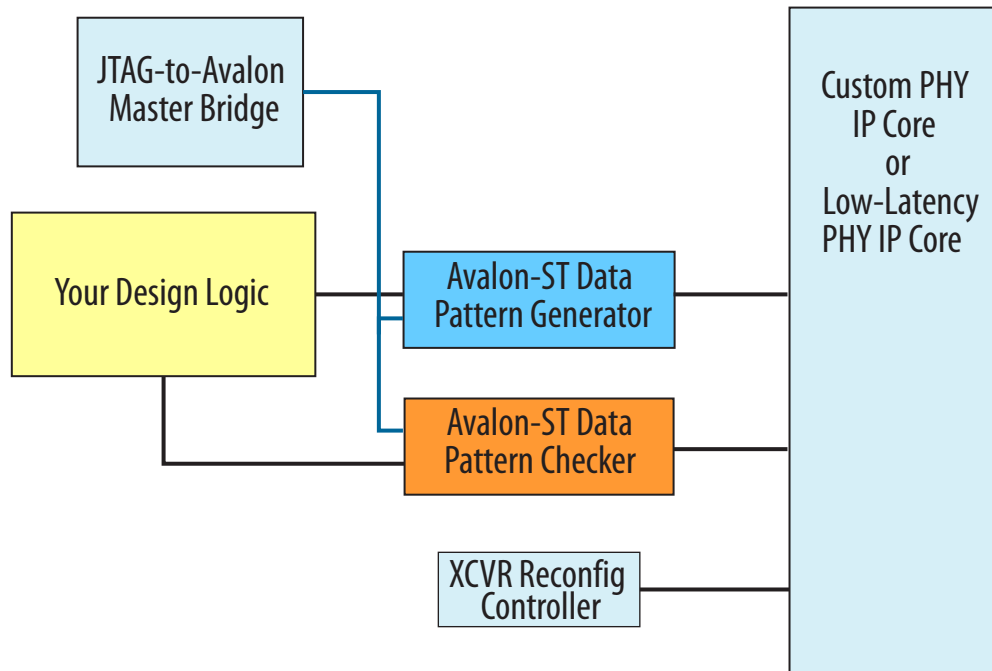
1. To add and connect debugging components to your system, click **Tools > Qsys**.
2. Define and instantiate the following from the **IP Catalog**:
  - **Altera Avalon Data Pattern Generator**. Turn on **Enable Bypass interface** for connection to design logic.
  - **Altera Avalon Data Pattern Checker**. Turn on **Enable Bypass interface** for connection to design logic.
  - **JTAG to Avalon Master Bridge**.
  - **Transceiver Reconfiguration Controller**.
3. Make the following connections between components in your system:

From	To
Your Design Logic	Data Pattern Generator bypass port
Data Pattern Generator	PHY input port
JTAG to Avalon Master Bridge	Altera Avalon Data Pattern Generator
JTAG to Avalon Master Bridge	Altera Avalon Data Pattern Checker
JTAG to Avalon Master Bridge	PHY input port
Data Pattern Checker	PHY output port

From	To
Transceiver Reconfiguration Controller	PHY input port

4. To generate the system, click **Generate** > **Generate HDL**.
5. To compile the design and generate configuration files, click **Processing** > **Start Compilation**.
6. Click **Tools** > **Programmer** and configure the target device with your debugging design.

**Figure 11-4: BER Test Configuration**



**Related Information**

[Running BER Tests](#) on page 11-23

**Configuring PRBS Signal Eye Tests**

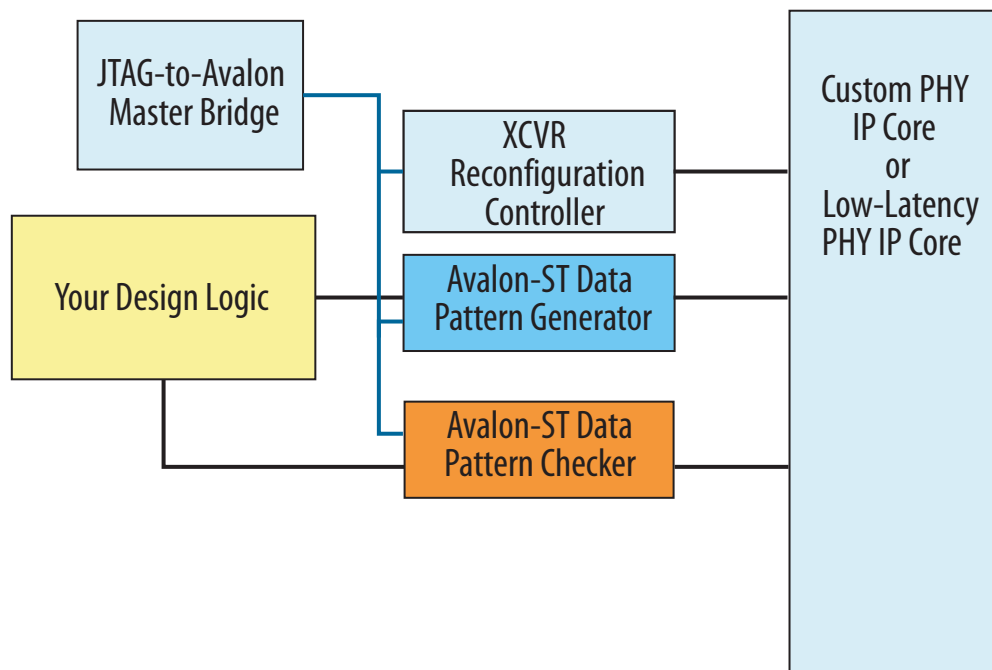
To integrate components with your design for testing PRBS signal eye, follow these steps.

1. To add and connect debugging components to your system, click **Tools** > **Qsys**.
2. Define and instantiate the following from the **IP Catalog**:
  - **Altera Avalon Data Pattern Generator**. Turn on **Enable Bypass interface** for connection to design logic.
  - **Altera Avalon Data Pattern Checker**. Turn on **Enable Bypass interface** for connection to design logic.
  - **JTAG to Avalon Master Bridge**.
  - **Transceiver Reconfiguration Controller**. Turn on **Enable EyeQ block** to enable signal eye analysis.
3. Make the following connections between components in your system:

From	To
Your Design Logic	Data Pattern Generator bypass port
Data Pattern Generator	PHY input port
JTAG to Avalon Master Bridge	Altera Avalon Data Pattern Generator
JTAG to Avalon Master Bridge	Altera Avalon Data Pattern Checker
Data Pattern Checker	PHY output port
JTAG to Avalon Master Bridge	Transceiver Reconfiguration Controller
JTAG to Avalon Master Bridge	PHY input port
Transceiver Reconfiguration Controller	PHY input port

4. To generate the system, click **Generate** > **Generate HDL**.
5. To compile the design and generate configuration files, click **Processing** > **Start Compilation**.
6. Click **Tools** > **Programmer** and configure the target device with your debugging design.

**Figure 11-5: PRBS Signal Eye Test Configuration**



#### Related Information

[Running PRBS Signal Eye Tests](#) on page 11-23



## Configuring Custom Traffic Signal Eye Tests

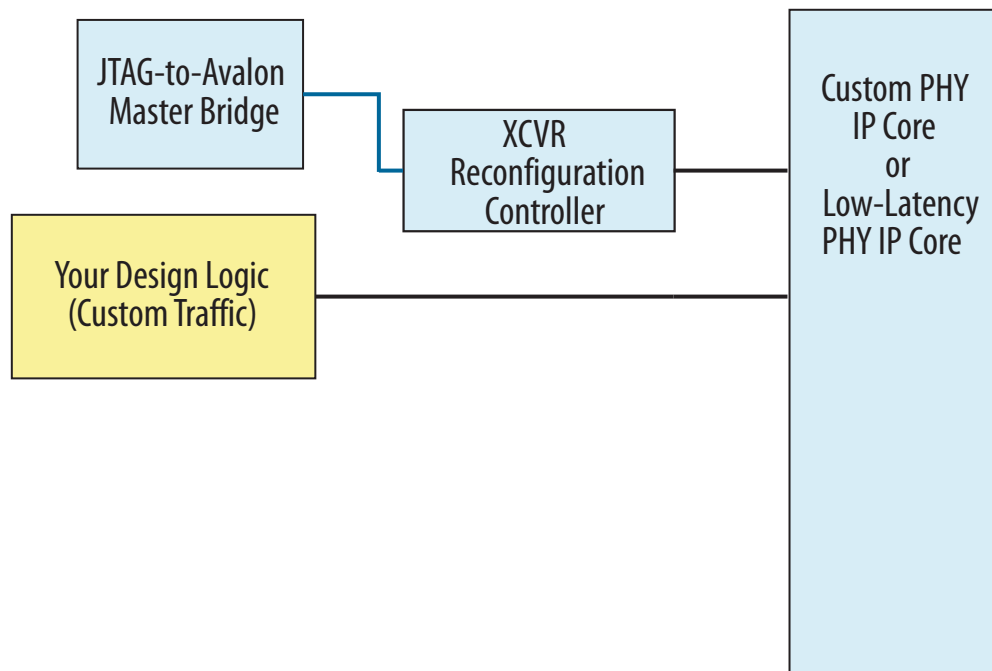
To integrate components with your design for testing custom traffic signal eye, follow these steps.

1. To add and connect debugging components to your system, click **Tools > Qsys**.
2. Define and instantiate the following from the **IP Catalog**:
  - **JTAG to Avalon Master Bridge**.
  - **Transceiver Reconfiguration Controller**. Turn on **Enable EyeQ block** to enable signal eye analysis. Turn on **Enable Bit Error Rate Block** to perform BER testing.
3. Make the following connections between components in your system:

From	To
Your design logic with custom traffic	PHY input port
JTAG to Avalon Master Bridge	Transceiver Reconfiguration Controller
JTAG to Avalon Master Bridge	PHY input port
Transceiver Reconfiguration Controller	PHY input port

4. To generate the system, click **Generate > Generate HDL**.
5. To compile the design and generate configuration files, click **Processing > Start Compilation**.
6. Click **Tools > Programmer** and configure the target device with your debugging design.

**Figure 11-6: Custom Traffic Signal Eye Test Configuration**



### Related Information

[Running Custom Traffic Tests](#) on page 11-24

## Configuring Link Optimization Tests

To integrate components with your design for link optimization tests, follow these steps.

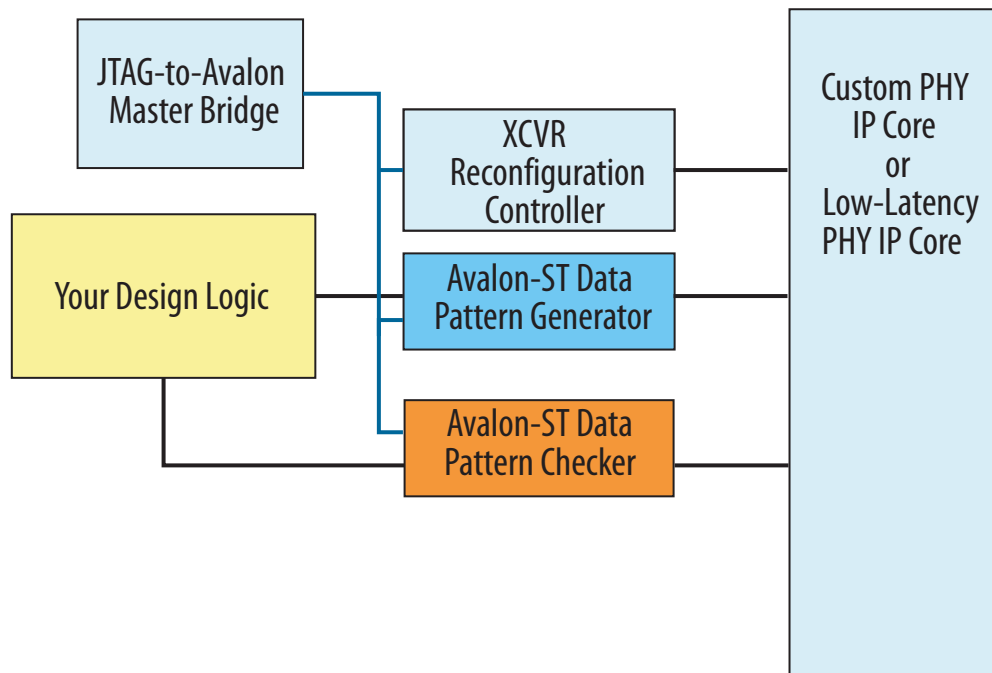
1. To add and connect debugging components to your system, click **Tools > Qsys**.
2. Define and instantiate the following from the **IP Catalog**:
  - **Altera Avalon Data Pattern Generator**. Turn on **Enable Bypass interface** for connection to design logic.
  - **Altera Avalon Data Pattern Checker**. Turn on **Enable Bypass interface** for connection to design logic.
  - **JTAG to Avalon Master Bridge**.
  - **Transceiver Reconfiguration Controller**. Turn on **Enable EyeQ block**, **Enable Analog controls**, **Enable decision feedback equalizer (DFE) block**, and **Enable adaptive equalization (AEQ) block** to enable all types of link analysis.
3. Make the following connections between components in your system:

From	To
Your Design Logic	Data Pattern Generator bypass port
Data Pattern Generator	PHY input port
JTAG to Avalon Master Bridge	Altera Avalon Data Pattern Generator
JTAG to Avalon Master Bridge	Altera Avalon Data Pattern Checker
Data Pattern Checker	PHY output port
JTAG to Avalon Master Bridge	Transceiver Reconfiguration Controller
JTAG to Avalon Master Bridge	PHY input port
Transceiver Reconfiguration Controller	PHY input port

4. To generate the system, click **Generate > Generate HDL**.
5. To compile the design and generate configuration files, click **Processing > Start Compilation**.
6. Click **Tools > Programmer** and configure the target device with your debugging design.



Figure 11-7: Link Optimization Test Configuration



**Related Information**

[Running Link Optimization Tests](#) on page 11-25

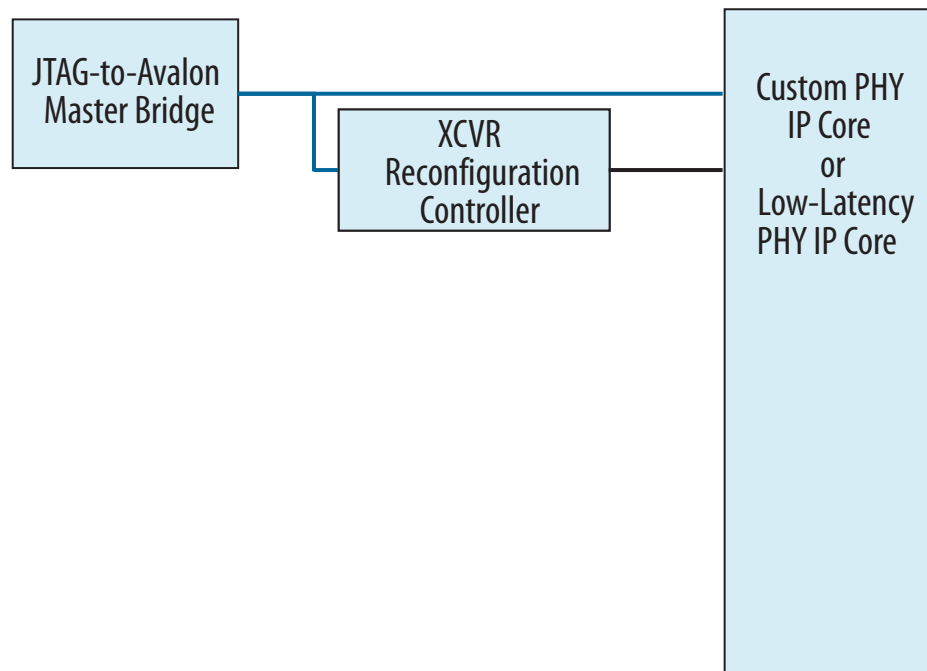
**Configuring PMA Analog Setting Control**

To integrate components for PMA analog setting control, follow these steps.

1. To add and connect debugging components to your system, click **Tools** > **Qsys**.
2. Define and instantiate the following from the **IP Catalog**:
  - **JTAG to Avalon Master Bridge**.
  - **Transceiver Reconfiguration Controller**. Turn on **Enable Analog controls**. You can optionally turn on **Enable EyeQ block**, **Enable decision feedback equalizer (DFE) block**, and **Enable adaptive equalization (AEQ) block** to enable these types of link analysis.
3. Make the following connections between components in your system:

From	To
JTAG to Avalon Master Bridge	Transceiver Reconfiguration Controller
JTAG to Avalon Master Bridge	PHY input port
Transceiver Reconfiguration Controller	PHY input port

4. To generate the system, click **Generate** > **Generate HDL**.
5. To compile the design and generate configuration files, click **Processing** > **Start Compilation**.
6. Click **Tools** > **Programmer** and configure the target device with your debugging design.

**Figure 11-8: PMA Analog Setting Control Configuration****Related Information**

[Controlling PMA Analog Settings](#) on page 11-25

## Debugging Transceiver Links

The Transceiver Toolkit allows you to control and monitor the performance of high-speed serial links running on your board in real-time. You can identify the transceiver links in your design, transmit a data pattern across the transceiver link, and report the signal quality of the received data in terms of bit error rate, bathtub curve, heat map, or EyeQ graph (for supported families).

The toolkit automatically identifies the transceiver links in your design, or you can manually create transceiver links. You can then run auto sweep to help you quickly identify the best PMA settings for each link. You can directly control the transmitter/receiver channels to experiment with various settings suggested by auto sweep. The EyeQ graph allows you to visualize the estimated horizontal and vertical eye opening at the receiver.

The Transceiver Toolkit supports various transceiver link testing configurations. You can identify and test the transceiver link between two Altera devices, or you can transmit a test pattern with a third-party device and monitor the data on an Altera device receiver channel. If a third-party chip includes self-test capability, then you can send the test pattern from the Altera device and monitor the signal integrity at the third-party device receiver channel. If the third-party device supports reverse serial loopback, you can run the test entirely within the Transceiver Toolkit.

Before you can monitor transceiver channels, you must configure a system with debugging components, and program the design into an FPGA. Once those steps are complete, use the following flow to test the channels:

1. Load the design in Transceiver Toolkit
2. Link hardware resources
3. Verify hardware connections
4. Identify transceiver channels
5. Run link tests or control PMA analog settings
6. View results

## Step 1: Load Your Design

The Transceiver Toolkit automatically loads the last compiled design upon opening. To load any design into the toolkit, click **File > Load Design** and select the **.sof** programming file generated for your transceiver design. Loading the **.sof** automatically links the design to the target hardware in the toolkit. The toolkit automatically discovers links between transmitter and receiver of the same channel. The **System Explorer** displays information about the loaded design.

## Step 2: Link Hardware Resources

The toolkit automatically discovers connected hardware and designs. You can also manually link a design to connected hardware resources in the **System Explorer**.

If you are using more than one Altera board, you can set up a test with multiple devices linked to the same design. This setup is useful when you want to perform a link test between a transmitter and receiver on two separate devices. You can also load multiple Quartus II projects and make links between different systems. You can perform tests on completely separate and unrelated systems in a single tool instance.

**Note:** Prior to the Transceiver Toolkit version 11.1, you must manually load and link your design to hardware. In version 11.1 and later, the Transceiver Toolkit automatically links any device programmed with a project.

Figure 11-9: One Channel Loopback Mode

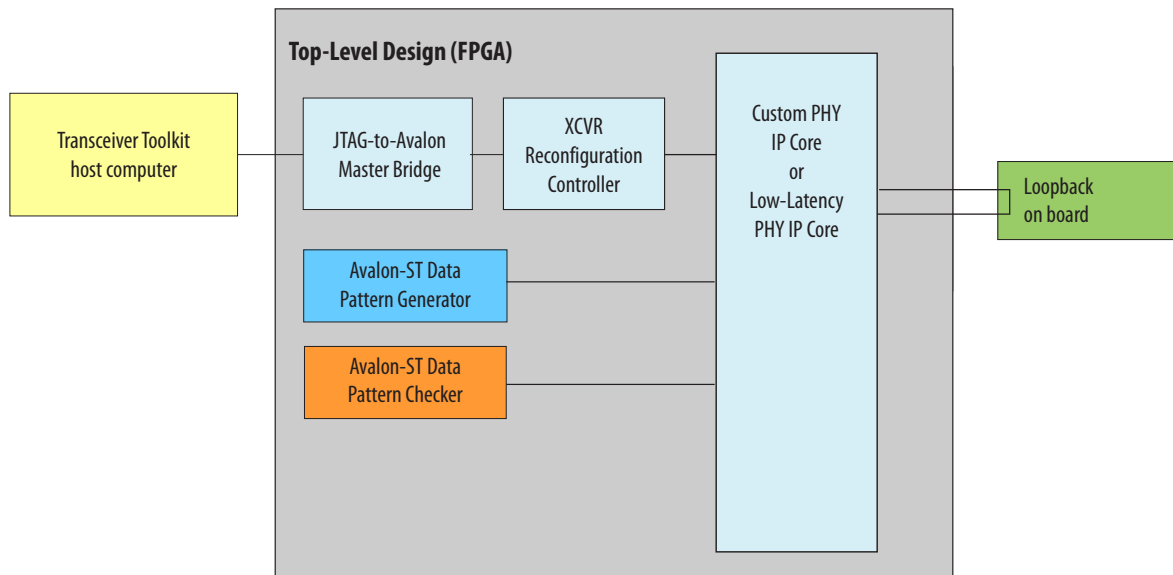
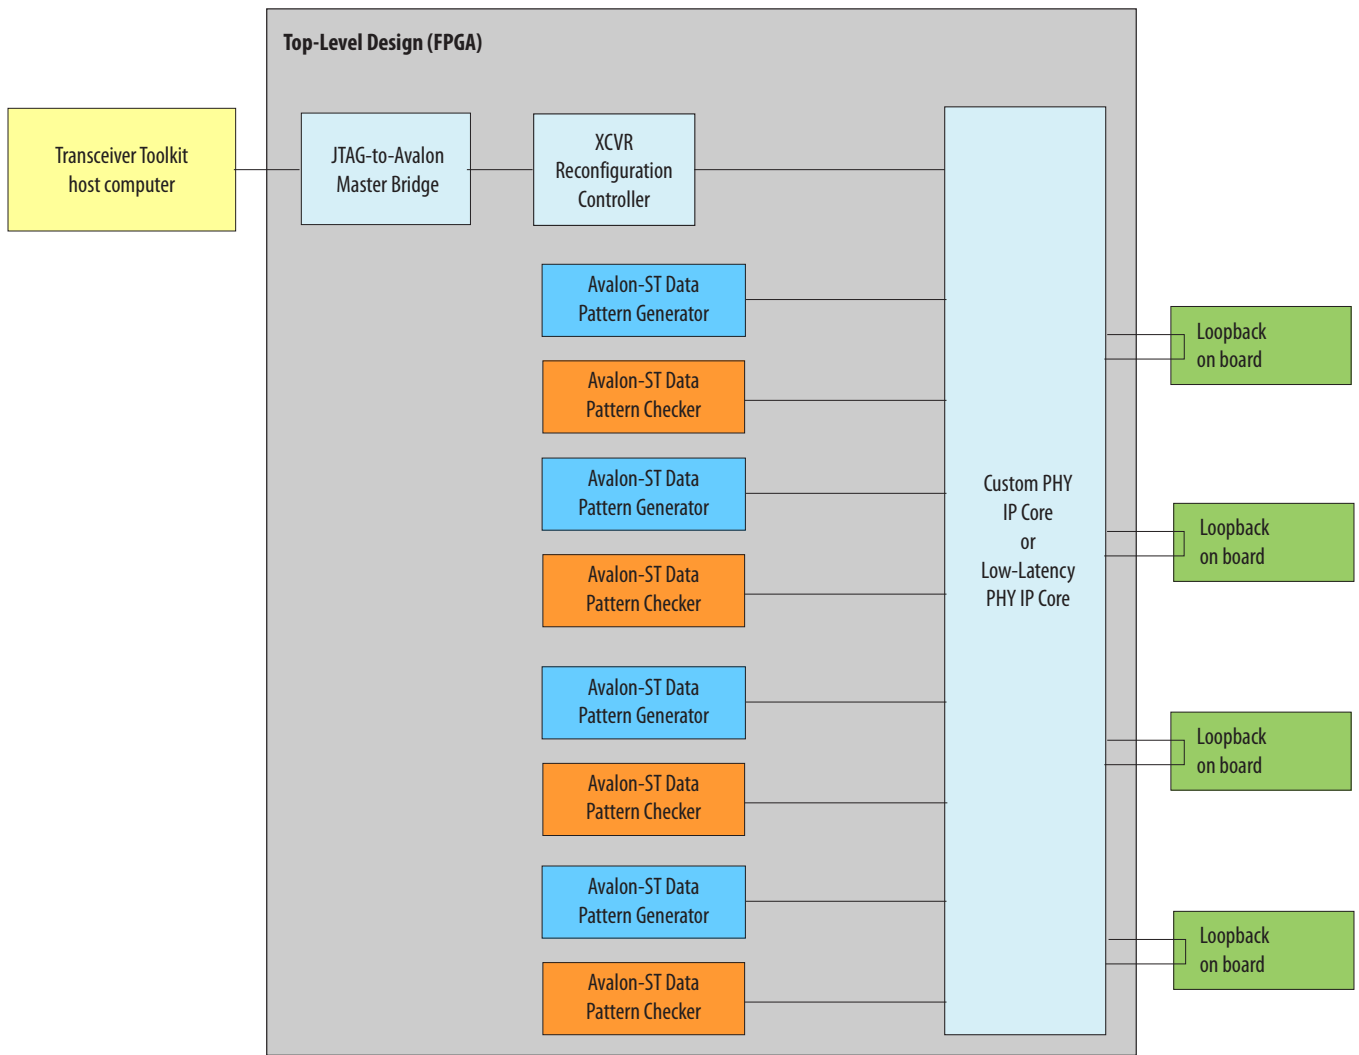


Figure 11-10: Four Channel Loopback Mode



### Linking One Design to One Device

To link one design to one device by one USB-Blaster download cable, follow these steps:

1. Load the design for your Quartus II project.
2. Link each device to an appropriate design if the design has not auto-linked.
3. Create the link between channels on the device to test.

### Linking Two Designs to Two Devices

To link two designs to two separate devices on the same board, connected by one USB-Blaster download cable, follow these steps:

1. Load the design for all the Quartus II project files you might need.
2. Link each device to an appropriate design if the design has not auto-linked.
3. Open the project for the second device.
4. Link the second device on the JTAG chain to the second design (unless the design auto-links).
5. Create a link between the channels on the devices you want to test.

### Linking Designs and Devices on Separate Boards

To link two designs to two separate devices on separate boards, connected to separate USB-Blaster download cables, follow these steps:

1. Load the design for all the Quartus II project files you might need.
2. Link each device to an appropriate design if the design has not auto-linked.
3. Create the link between channels on the device to test.
4. Link the device you connected to the second USB-Blaster download cable to the second design.
5. Create a link between the channels on the devices you want to test.

### Linking One Design on Two Devices

To link the same design on two separate devices, follow these steps:

1. In the Transceiver Toolkit, open the **.sof** you are using on both devices.
2. Link the first device to this design instance.
3. Link the second device to the design.
4. Create a link between the channels on the devices you want to test.

### Step 3: Verify Hardware Connections

After you load your design and link your hardware, verify that the channels are connected correctly and looped back properly on the hardware. Use the toolkit to send data patterns and receive them correctly. Verifying your link and correct channel before you perform Auto Sweep or EyeQ tests can save time in the work flow.

After you have verified that the transmitter and receiver are communicating with each other, you can create a link between the two transceivers so that you can perform Auto Sweep and EyeQ tests with this pair.

### Step 4: Identify Transceiver Channels

The Transceiver Toolkit automatically displays recognized transmitter and receiver channels. The toolkit identifies a channel automatically whenever a receiver and transmitter share a transceiver channel. You can also manually identify the transmitter and receiver in a transceiver channel and create a link between the two for testing.

When you run link tests, channel color highlights indicate the test status:

**Table 11-5: Channel Color Highlights**

Color	Transmitter Channel	Receiver Channel
Red	Channel is closed or generator clock is not running	Channel is closed or checker clock is not running



Color	Transmitter Channel	Receiver Channel
Green	Generator is sending a pattern	Checker is checking and data pattern is locked
Neutral	Channel is open, generator clock is running, and generator is not sending a pattern, or generator is not sending a pattern	Channel is open, checker clock is running, and checker is not checking, or checker is not checking
Yellow	N/A	Checker is checking and data pattern is not locked

## Step 5: Run Link Tests

Once you identify the transceiver channels for debugging, you can run various link tests in the toolkit.

Use the **Transceiver Links** tab to control link tests. For example, use the Auto Sweep feature to sweep transceiver settings to determine the parameters that support the best BER value. Click **Link Auto Sweep**, **Link EyeQ** or **Link Auto Sweep & EyeQ** to adjust the PMA settings and run tests.

### Running BER Tests

You can run BER tests across your transceiver link. After programming the FPGA with your debugging design, loading the design in the toolkit, and linking hardware, follow these steps to run BER tests:

1. Click **Setup**.
  - a. Select the generator and checker you want to control.
  - b. Select the transmitter and receiver pair you want to control.
  - c. Click **Create Transceiver Link** and click **Close**
2. Click **Control Transceiver Link**, and specify a PRBS **Test pattern** and **Data pattern checker** for **Checker mode**. The checker mode option is only available after you turn on **Enable EyeQ block** and **Enable Bit Error Rate Block** in the Reconfiguration Controller component.
 

If you select **Bypass** for the **Test pattern**, the toolkit bypasses the PRBS generator and runs your design through the link. The bypass option is only available after you turn on **Enable Bypass interface** in the Reconfiguration Controller component.
3. Experiment with **Reconfiguration**, **Generator**, or **Checker** settings.
4. Click **Start** to run the pattern with your settings. You can then click **Inject Error** to inject error bits, **Reset** the counter, or **Stop** the test.

#### Related Information

[Configuring BER Tests](#) on page 11-12

### Running PRBS Signal Eye Tests

You can run PRBS signal eye tests to visualize the estimated horizontal and vertical eye opening at the receiver. After programming the FPGA with your debugging design, loading the design in the toolkit, and linking hardware, follow these steps to run PRBS signal eye tests:

1. Click **Setup**.

- a. Select the generator and checker you want to control.
  - b. Select the transmitter and receiver pair you want to control.
  - c. Click **Create Transceiver Link** and click **Close**.
2. Click **Link EyeQ**, and select **EyeQ** as the **Test mode**. The **EyeQ** mode displays test results as a bathtub curve, heat map, or eye contour representing bit error and phase offset data.
  3. Specify the PRBS **Test pattern** and the **Checker mode**. Use **Serial bit comparator** checker mode only for checking a large window of error with custom traffic.  
The checker mode option is only available after you turn on **Enable EyeQ block** and **Enable Bit Error Rate Block** in the Reconfiguration Controller component.
  4. Specify **Run length** and **EyeQ settings** to control the test coverage and type of EyeQ results displayed, respectively.
  5. Click **Start** to run the pattern with your settings. EyeQ uses the current channel settings to start a phase sweep of the channel. The phase sweep runs 32 iterations. As the run progresses, view the status under **EyeQ status**. Use this diagram to compare PMA settings for the same channel and to choose the best combination of PMA settings for a particular channel.
  6. When the run completes, the chart is displayed and the characteristics of each run are listed in the run list. You can click **Stop** to halt the test, change the PMA settings, and re-start the test. Click **Create Report** to export data to a table format for further viewing.

#### Related Information

[Configuring PRBS Signal Eye Tests](#) on page 11-13

## Running Custom Traffic Tests

After programming the FPGA with your debugging design, loading the design in the toolkit, and linking hardware, follow these steps to run custom traffic tests:

1. Click **Setup**.
  - a. Select the associated reconfiguration controller
  - b. Click **Create Transceiver Link** and click **Close**
2. Click the **Receiver EyeQ**, and select **EyeQ** as the **Test mode**. The **EyeQ** mode displays test results as a bathtub curve, heat map, or eye contour representing bit error and phase offset data.
3. Specify the PRBS **Test pattern**
4. For **Checker mode**, select **Serial bit comparator**.  
The checker mode option is only available after you turn on **Enable EyeQ block** and **Enable Bit Error Rate Block** in the Reconfiguration Controller component.
5. Specify **Run length** and **EyeQ settings** to control the test coverage and type of EyeQ results displayed, respectively.
6. Click **Start** to run the pattern with your settings. EyeQ uses the current channel settings to start a phase sweep of the channel. The phase sweep runs 32 iterations. As the run progresses, view the status under **EyeQ status**.
7. When the run completes, the chart is displayed and the characteristics of each run are listed in the run list. You can click **Stop** to halt the test, change the PMA settings, and re-start the test. Click **Create Report** to export data to a table format for further viewing.

#### Related Information

[Configuring Custom Traffic Signal Eye Tests](#) on page 11-15

## Running Link Optimization Tests

After programming the FPGA with your debugging design, loading the design in the toolkit, and linking hardware, follow these steps to run link optimization tests:

1. Click the **Transceiver Links** tab, and select the channel you want to control.
2. Click **Link Auto Sweep**. The **Advanced** tab appears with **Auto sweep** as **Test mode**.
3. Specify the PRBS **Test pattern**.
4. Specify **Run length**, experiment with the **Transmitter settings**, and **Receiver settings** to control the test coverage and PMA settings, respectively.
5. Click **Start** to run all combinations of tests meeting the PMA parameter limits.
6. When the run completes the chart is displayed and the characteristics of each run are listed in the run list. You can click **Stop** to halt the test, change the PMA settings, and re-start the test. Click **Create Report** to export data to a table format for further viewing.
7. To use decision feedback equalization (DFE) to determine the best tap settings, follow these steps:
  - a. Use Auto Sweep to find optimal PMA settings while leaving the **DFE mode** set to **Off**.
  - b. If BER = 0, use the best PMA settings achieved.
  - c. If BER > 0, use this PMA setting, and set the minimum and maximum values obtained from Auto Sweep to match this setting. Set the maximum DFE range to limits for each of the three DFE settings.
  - d. Run **Create Report** to view the results and determine which DFE setting has the best BER. Use these settings in conjunction with the PMA settings for the best results.

### Related Information

[Configuring Link Optimization Tests](#) on page 11-16

## Controlling PMA Analog Settings

You can directly control PMA analog settings to experiment with settings while the link is running. To control PMA analog settings, follow these steps:

1. Click **Setup**.
  - a. Click the **Transmitter Channels** tab, define a transmitter without a generator, and click **Create Transmitter Channel**.
  - b. Click the **Receiver Channels** tab, define a receiver without a generator, and click **Create Receiver Channel**.
  - c. Click the **Transceiver Links** tab, select the transmitter and receivers you want to control, and click **Create Transceiver Link**.
  - d. Click **Close**.
2. Click **Control Receiver Channel**, **Control Transmitter Channel**, or **Control Transceiver Link** to directly control the PMA settings while running.

### Related Information

[Configuring PMA Analog Setting Control](#) on page 11-17

## Toolkit GUI Setting Reference

The following settings are available for interaction with the transmitter channels or receiver channels or transceiver links in the Transceiver Toolkit GUI.

**Table 11-6: Transceiver Toolkit Control Panel Settings**

Setting	Description	Control Panel
1-D EyeQ mode	This feature should be used when <b>DFE</b> is on and the <b>EyeQ mode</b> is set to <b>Bathtub curve</b> . Ignores the vertical step settings for EyeQ.	Receiver Transceiver Link
Alias	Name you choose for the channel.	Transmitter Receiver Transceiver Link
Auto sweep status	Reports the current and best tested bits, errors, bit error rate, and case count for the current auto sweep test.	Receiver Transceiver Link
Bit error rate (BER)	Specifies errors divided by bits tested since the last reset of the checker.	Receiver Transceiver Link
Channel address	Logical address number of the transceiver channel.	Transmitter Receiver Transceiver Link
Checker mode	Specify <b>Data pattern checker</b> or <b>Serial bit comparator</b> for BER tests.  If you enable <b>Serial bit comparator</b> the Data Pattern Generator sends the PRBS pattern, but the pattern is checked by the serial bit comparator.  In <b>Bypass mode</b> , clicking <b>Start</b> begins counting on the Serial bit comparator.	Receiver Transceiver Link
Data rate	Data rate of the channel as read from the project file or data rate as measured by the frequency detector.  To use the frequency detector, turn on <b>Enable Frequency Counter</b> in the Data Pattern Checker IP core and/or Data Pattern Generator IP core, regenerate the IP cores, and recompile the design.  The measured data rate depends on the Avalon management clock frequency as read from the project file.  Click the refresh button next to the measured <b>Data rate</b> if you make changes to your settings and want to sample the data rate again.	Transmitter Receiver Transceiver Link

Setting	Description	Control Panel
DC gain	Circuitry that provides an equal boost to the incoming signal across the frequency spectrum.	Receiver Transceiver Link
DFE mode and tap values 1-5	Decision feedback equalization (DFE) for improving signal quality. One-time mode DFE determines the best tap settings and stops searching. There is also a one-time adaptive mode button that automatically turns on one-time mode and immediately populates converged values into the manual settings lists. Adaptive mode DFE automatically tries to find the best tap values.	Receiver Transceiver Link
Enable word aligner	Forces the transceiver channel to align to the word you specify.	Receiver Transceiver Link
Equalization control	Boosts the high-frequency gain of the incoming signal, thereby compensating for the low-pass filter effects of the physical medium. AEQ one-time adaptation is supported in Auto Sweep. When used with DFE, you need to use DFE triggered mode or DFE continuous.	Receiver Transceiver Link
Equalization mode	Adaptive equalization (AEQ) automatically evaluates and selects the best combination of equalizer settings. The setting applies only to Stratix V devices. When turned on, it automatically turns off Equalization Control. The one-time selection determines the best setting and stops searching. You can use AEQ for multiple, independently controlled receiver channels.	Receiver Transceiver Link
Error rate limit	Turns on or off error rate limits. <b>Start checking after</b> waits until the set number of bits are satisfied until it starts looking at the bit error rate (BER) for the next two checks.  <b>Bit error rate achieves below</b> sets upper bit error rate limits. If the error rate is better than the set error rate, the test ends.  <b>Bit error rate exceeds</b> Sets lower bit error rate limits. If the error rate is worse than the set error rate, the test ends.	Receiver Transceiver Link

Setting	Description	Control Panel
EyeQ mode	Allows you to specify Eye contour or Bathtub curve as the type of EyeQ graph generated by the test.	Transmitter Receiver Transceiver Link
EyeQ phase step	Sets the phase step for sampling the data from an offset of the CDR (clock data recovery) data path; set to Off to use the regular clock data recovery (CDR) data path.	Receiver Transceiver Link
EyeQ status	Displays a graphical representation of signal integrity as an eye contour, bathtub curve plot, or heat map.	Transmitter Receiver Transceiver Link
EyeQ vertical step	Sets the voltage threshold of the sampler to report the height of the eye. Negative numbers are allowed for vertical steps to capture asymmetric eye.	Receiver Transceiver Link
Horizontal phase step interval	Specify the number of horizontal steps to increment when performing a sweep. Increasing the value increases the speed of the test but at a lower resolution. This option only applies to eye contour.	Transmitter Receiver Transceiver Link
Increase test range	Right-click in the <b>Advanced</b> panel to use the span capabilities of Auto Sweep to automatically increase the span of tests by one unit down for the minimum and one unit up for the maximum, for the selected set of controls. You can span either PMA Analog controls (non-DFE controls), or the DFE controls. You can quickly set up a test to check if any PMA setting combinations near your current best could yield better results.	Receiver Transceiver Link
Inject Error	Flips one bit to the output of the data pattern generator to introduce an artificial error.	Transmitter Transceiver Link
Maximum tested bits	Sets the maximum number of tested bits for each test iteration.	Receiver Transceiver Link
Number of bits tested	Specifies the number of bits tested since the last reset of the checker.	Receiver Transceiver Link

Setting	Description	Control Panel
Number of error bits	Specifies the number of error bits encountered since the last reset of the checker.	Receiver Transceiver Link
Number of preamble beats	The number of clock cycles to which the preamble word is sent before the test pattern begins.	Transmitter Transceiver Link
PLL refclk freq	Channel reference clock frequency as read from the project file or measured reference clock frequency as calculated from the measured data rate.	Transmitter Receiver Transceiver Link
Populate with	Right-click in the <b>Advanced</b> panel to load current values on the device as a starting point, or initially load the best settings determined through auto sweep. The Quartus II software automatically applies the values you specify in the drop-down lists for the Transmitter settings and Receiver settings.	Receiver Transceiver Link
Preamble word	Word to send out if the preamble mode is used.	Transmitter Transceiver Link
Pre-emphasis	The programmable pre-emphasis module in each transmit buffer boosts high frequencies in the transmit data signal, which may be attenuated in the transmission media. Using pre-emphasis can maximize the data eye opening at the far-end receiver.	Transmitter Transceiver Link
Receiver channel	Specifies the name of the selected receiver channel.	Receiver Transceiver Link
Reset	Resets the current test.	Receiver Transceiver Link
Rules Based Configuration (RBC) validity checking	Displays invalid combination of settings in red in each list under <b>Transmitter settings</b> and <b>Receiver settings</b> , based on previous settings. If selected, the settings remain in red to indicate the currently selected combination is invalid. This feature helps you to avoid manually testing invalid settings that cannot be compiled into your design. This feature helps to prevent you from setting the device into an invalid mode for extended periods of time and potentially damaging the circuits.	Receiver Transceiver Link

Setting	Description	Control Panel
Run length	Sets coverage parameters for test runs.	Transmitter Receiver Transceiver Link
Run results table	Lists the statistics of each EyeQ test run. The run results table is sortable. You can right-click any of the tests in the table and then click <b>Apply Settings to Device</b> to quickly apply the chosen PMA settings to your device. You can click <b>Import</b> to load reports from previously generated EyeQ runs into the run results table. You can click <b>Export</b> to export single or multiple runs from the run results table to a report.	Transmitter Receiver Transceiver Link
RX CDR PLL status <sup>(8)</sup>	Shows the receiver in lock-to-reference (LTR) mode. When in auto-mode, if data cannot be locked, this signal alternates in LTD mode if the CDR is locked to data.	Receiver Transceiver Link
RX CDR data status	Shows the receiver in lock-to-data (LTD) mode. When in auto-mode, if data cannot be locked, the signal stays high when locked to data and never toggles.	Receiver Transceiver Link
Serial loopback enabled	Inserts a serial loopback before the buffers, allowing you to form a link on a transmitter and receiver pair on the same physical channel of the device.	Transmitter Receiver Transceiver Link
Start	Starts the pattern generator or checker on the channel to verify incoming data.	Transmitter Receiver Transceiver Link
Stop	Stops generating patterns and testing the channel.	Transmitter Receiver Transceiver Link
Target bit error rate	Finds the contour edge of the bit error rate that you select. This option only applies to eye contour mode.	Transmitter Receiver Transceiver Link

<sup>(8)</sup> For Stratix V, the Phase Frequency Detector (PFD) is inactive in LTD mode. The `rx_is_lockedtoref` status signal toggles randomly and is not significant in LTD mode.



Setting	Description	Control Panel
Test mode	Allows you to specify the <b>Auto sweep</b> , <b>EyeQ</b> , or <b>Auto sweep and EyeQ</b> test mode.	Receiver Transceiver Link
Test pattern	Test pattern sent by the transmitter channel. Options include <b>PRBS7</b> , <b>PRBS15</b> , <b>PRBS23</b> , <b>PRBS31</b> , <b>LowFrequency</b> , and <b>HighFrequency</b> and <b>Bypass mode</b> . The Data Pattern Checker self-aligns both high and low frequency patterns. Use <b>Bypass mode</b> to send user-design data.	Transmitter Receiver Transceiver Link
Time limit	Specifies the time limit unit and value to have a maximum bounds time limit for each test iteration	Receiver Transceiver Link
Transmitter channel	Specifies the name of the selected transmitter channel.	Transmitter Transceiver Link
TX/CMU PLL status	Provides status of whether the transmitter channel PLL is locked to the reference clock.	Transmitter Transceiver Link
Use preamble upon start	If turned on, sends the preamble word before the test pattern. If turned off, starts sending the test pattern immediately.	Transmitter Transceiver Link
Vertical phase step interval	Specify the number of vertical steps to increment when performing a sweep. Increasing the value increases the speed of the test but at a lower resolution. This option only applies to the eye contour.	Transmitter Receiver Transceiver Link
V <sub>OD</sub> control	Programmable transmitter differential output voltage.	Transmitter Transceiver Link

## Scripting API

You can alternatively use Tcl commands to access Transceiver Toolkit functions, rather than using the GUI. You can script various tasks, such as loading a project, creating design instances, linking device resources, and identifying high-speed serial links. You can save your project setup in a Tcl script for use in subsequent testing sessions. You can also build a custom test routine script.

After you set up and define links that describe the entire physical system, you can click Save Tcl Script to save the setup for future use. To run the scripts, double-click script names in the System Explorer scripts folder.

View a list of the available Tcl commands in the Tcl Console window. Select Tcl commands in the list to view descriptions, including example usage.

To view Tcl command descriptions from the Tcl Console window:

1. Type `help help`. The Console displays all Transceiver Toolkit Tcl commands.
2. Type `help <command name>`. The Console displays the command description.

## Transceiver Toolkit Commands

The following tables list the available Transceiver Toolkit scripting commands.

**Table 11-7: Transceiver Toolkit Channel\_rx Commands**

Command	Arguments	Function
<code>transceiver_channel_rx_get_data</code>	<code>&lt;service-path&gt;</code>	Returns a list of the current checker data. The results are in the order of number of bits, number of errors, and bit error rate.
<code>transceiver_channel_rx_get_dcgain</code>	<code>&lt;service-path&gt;</code>	Gets the DC gain value on the receiver channel.
<code>transceiver_channel_rx_get_dfe_tap_value</code>	<code>&lt;service-path&gt; &lt;tap position&gt;</code>	Gets the current tap value of the specified channel at the specified tap position.
<code>transceiver_channel_rx_get_eqctrl</code>	<code>&lt;service-path&gt;</code>	Gets the equalization control value on the receiver channel.
<code>transceiver_channel_rx_get_pattern</code>	<code>&lt;service-path&gt;</code>	Returns the current data checker pattern by name.
<code>transceiver_channel_rx_has_dfe</code>	<code>&lt;service-path&gt;</code>	Gets whether this channel has the DFE feature available.
<code>transceiver_channel_rx_has_eyeq</code>	<code>&lt;service-path&gt;</code>	Gets whether the EyeQ feature is available for the specified channel.
<code>transceiver_channel_rx_is_checking</code>	<code>&lt;service-path&gt;</code>	Returns non-zero if the checker is running.
<code>transceiver_channel_rx_is_dfe_enabled</code>	<code>&lt;service-path&gt;</code>	Gets whether the DFE feature is enabled on the specified channel.
<code>transceiver_channel_rx_is_locked</code>	<code>&lt;service-path&gt;</code>	Returns non-zero if the checker is locked onto the incoming data.
<code>transceiver_channel_rx_reset_counters</code>	<code>&lt;service-path&gt;</code>	Resets the bit and error counters inside the checker.

Command	Arguments	Function
transceiver_channel_rx_reset	<service-path>	Resets the specified channel.
transceiver_channel_rx_set_dcgain	<service-path> <value>	Sets the DC gain value on the receiver channel.
transceiver_channel_rx_set_dfe_enabled	<service-path><disable(0)/enable(1)>	Enables or disables the DFE feature on the specified channel.
transceiver_channel_rx_set_dfe_tap_value	<service-path> <tap position> <tap value>	Sets the current tap value of the specified channel at the specified tap position to the specified value.
transceiver_channel_rx_set_dfe_adaptive	<service-path>	Sets the mode of DFE adaptation. 0=off, 1=adaptive, 2= one-time adaptive
transceiver_channel_rx_set_eqctrl	<service-path> <value>	Sets the equalization control value on the receiver channel.
transceiver_channel_rx_start_checking	<service-path>	Starts the checker.
transceiver_channel_rx_stop_checking	<service-path>	Stops the checker.
transceiver_channel_rx_get_eyeq_phase_step	<service-path>	Gets the current phase step of the specified channel.
transceiver_channel_rx_set_pattern	<service-path> <pattern-name>	Sets the expected pattern to the one specified by the pattern name.
transceiver_channel_rx_is_eyeq_enabled	<service-path>	Gets whether the EyeQ feature is enabled on the specified channel.
transceiver_channel_rx_set_eyeq_enabled	<service-path> <disable(0)/enable(1)>	Enables or disables the EyeQ feature on the specified channel.
transceiver_channel_rx_set_eyeq_phase_step	<service-path><phase step>	Sets the phase step of the specified channel.
transceiver_channel_rx_set_word_aligner_enabled	<service-path><disable(0)/enable(1)>	Enables or disables the word aligner of the specified channel.
transceiver_channel_rx_is_word_aligner_enabled	<service-path><disable(0)/enable(1)>	Gets whether the word aligner feature is enabled on the specified channel.

Command	Arguments	Function
transceiver_channel_rx_is_locked	<service-path>	Returns non-zero if the checker is locked onto the incoming signal.
transceiver_channel_rx_is_rx_locked_to_data	<service-path>	Returns 1 if transceiver is in lock to data (LTD) mode. Otherwise 0.
transceiver_channel_rx_is_rx_locked_to_ref	<service-path>	Returns 1 if transceiver is in lock to reference (LTR) mode. Otherwise 0.
transceiver_channel_rx_has_eyeq_1d	<service-path>	Detects whether the eye viewer pointed to by <service-path> supports 1D-EyeQ mode.
transceiver_channel_rx_set_1deye_mode	<service-path><disable(0)/enable(1)>	Enables or disables 1D-EyeQ mode.
transceiver_channel_rx_get_1deye_mode	<service-path>	Returns the current on or off status of 1D-EyeQ mode.

Table 11-8: Transceiver Toolkit Channel\_tx Commands

Command	Arguments	Function
transceiver_channel_tx_disable_preamble	<service-path>	Disables the preamble mode at the beginning of generation.
transceiver_channel_tx_enable_preamble	<service-path>	Enables the preamble mode at the beginning of generation.
transceiver_channel_tx_get_number_of_preamble_beats	<service-path>	Returns the currently set number of beats to send out the preamble word.
transceiver_channel_tx_get_pattern	<service-path>	Returns the currently set pattern.
transceiver_channel_tx_get_preamble_word	<service-path>	Returns the currently set preamble word.
transceiver_channel_tx_get_preemph0t	<service-path>	Gets the pre-emphasis pre-tap value on the transmitter channel.
transceiver_channel_tx_get_preemph1t	<service-path>	Gets the pre-emphasis first post-tap value on the transmitter channel.

Command	Arguments	Function
transceiver_channel_tx_get_preemph2t	<service-path>	Gets the pre-emphasis second post-tap value on the transmitter channel.
transceiver_channel_tx_get_vodctrl	<service-path>	Gets the V <sub>OD</sub> control value on the transmitter channel.
transceiver_channel_tx_inject_error	<service-path>	Injects a 1-bit error into the generator's output.
transceiver_channel_tx_is_generating	<service-path>	Returns non-zero if the generator is running.
transceiver_channel_tx_is_preamble_enabled	<service-path>	Returns non-zero if preamble mode is enabled.
transceiver_channel_tx_set_number_of_preamble_beats	<service-path><number-of-preamble-beats>	Sets the number of beats to send out the preamble word.
transceiver_channel_tx_set_pattern	<service-path><pattern-name>	Sets the output pattern to the one specified by the pattern name.
transceiver_channel_tx_set_preamble_word	<service-path><preamble-word>	Sets the preamble word to be sent out.
transceiver_channel_tx_set_preemph0t	<service-path><preemph0t value>	Sets the pre-emphasis pre-tap value on the transmitter channel.
transceiver_channel_tx_set_preemph1t	<service-path><preemph1t value>	Sets the pre-emphasis first post-tap value on the transmitter channel.
transceiver_channel_tx_set_preemph2t	<service-path><preemph2t value>	Sets the pre-emphasis second post-tap value on the transmitter channel.
transceiver_channel_tx_set_vodctrl	<service-path><vodctrl value>	Sets the V <sub>OD</sub> control value on the transmitter channel.
transceiver_channel_tx_start_generation	<service-path>	Starts the generator.
transceiver_channel_tx_stop_generation	<service-path>	Stops the generator.

Table 11-9: Transceiver Toolkit Transceiver Toolkit Debug\_Link Commands

Command	Arguments	Function
transceiver_debug_link_get_pattern	<service-path>	Gets the currently set pattern the link uses to run the test.
transceiver_debug_link_is_running	<service-path>	Returns non-zero if the test is running on the link.
transceiver_debug_link_set_pattern	<service-path> <data pattern>	Sets the pattern the link uses to run the test.
transceiver_debug_link_start_running	<service-path>	Starts running a test with the currently selected test pattern.
transceiver_debug_link_stop_running	<service-path>	Stops running the test.

Table 11-10: Transceiver Toolkit Reconfig\_Analog Commands

Command	Arguments	Function
transceiver_reconfig_analog_get_logical_channel_address	<service-path>	Gets the transceiver logical channel address currently set.
transceiver_reconfig_analog_get_rx_dcgain	<service-path>	Gets the DC gain value on the receiver channel specified by the current logical channel address.
transceiver_reconfig_analog_get_rx_eqctrl	<service-path>	Gets the equalization control value on the receiver channel specified by the current logical channel address.
transceiver_reconfig_analog_get_tx_preemph0t	<service-path>	Gets the pre-emphasis pre-tap value on the transmitter channel specified by the current logical channel address.
transceiver_reconfig_analog_get_tx_preemph1t	<service-path>	Gets the pre-emphasis first post-tap value on the transmitter channel specified by the current logical channel address.

Command	Arguments	Function
transceiver_reconfig_analog_get_tx_preemph2t	<service-path>	Gets the pre-emphasis second post-tap value on the transmitter channel specified by the current logical channel address.
transceiver_reconfig_analog_get_tx_vodctrl	<service-path>	Gets the V <sub>OD</sub> control value on the transmitter channel specified by the current logical channel address.
transceiver_reconfig_analog_set_logical_channel_address	<service-path><logical channel address>	Sets the transceiver logical channel address.
transceiver_reconfig_analog_set_rx_dcgain	<service-path><dc_gain value>	Sets the DC gain value on the receiver channel specified by the current logical channel address
transceiver_reconfig_analog_set_rx_eqctrl	<service-path> <eqctrl value>	Sets the equalization control value on the receiver channel specified by the current logical channel address.
transceiver_reconfig_analog_set_tx_preemph0t	<service-path><preemph0t value>	Sets the pre-emphasis pre-tap value on the transmitter channel specified by the current logical channel address.
transceiver_reconfig_analog_set_tx_preemph1t	<service-path><preemph1t value>	Sets the pre-emphasis first post-tap value on the transmitter channel specified by the current logical channel address.
transceiver_reconfig_analog_set_tx_preemph2t	<service-path> <preemph2t value>	Sets the pre-emphasis second post-tap value on the transmitter channel specified by the current logical channel address.
transceiver_reconfig_analog_set_tx_vodctrl	<service-path> <vodctrl value>	Sets the V <sub>OD</sub> control value on the transmitter channel specified by the current logical channel address.

Table 11-11: Transceiver Toolkit Decision Feedback Equalization (DFE) Commands

Command	Arguments	Function
alt_xcvr_reconfig_dfe_get_logical_channel_address	<service-path>	Gets the logical channel address that other alt_xcvr_reconfig_dfe commands use to apply.
alt_xcvr_reconfig_dfe_is_enabled	<service-path>	Gets whether the DFE feature is enabled on the previously specified channel.
alt_xcvr_reconfig_dfe_set_enabled	<service-path> <disable(0)/enable(1)>	Enables or disables the DFE feature on the previously specified channel.
alt_xcvr_reconfig_dfe_set_logical_channel_address	<service-path> <logical channel address>	Sets the logical channel address that other alt_xcvr_reconfig_eye_viewer commands use.
alt_xcvr_reconfig_dfe_set_tap_value	<service-path> <tap position> <tap value>	Sets the tap value at the previously specified channel at specified tap position and value.

Table 11-12: Transceiver Toolkit Eye Monitor Commands

Command	Arguments	Function
alt_xcvr_custom_is_word_aligner_enabled	<service-path> <disable(0)/enable(1)>	Gets whether the word aligner feature is enabled on the previously specified channel.
alt_xcvr_custom_set_word_aligner_enabled	<service-path> <disable(0)/enable(1)>	Enables or disables the word aligner of the previously specified channel.
alt_xcvr_custom_is_rx_locked_to_data	<service-path>	Returns whether the receiver CDR is locked to data.
alt_xcvr_custom_is_rx_locked_to_ref	<service-path>	Returns whether the receiver CDR PLL is locked to the reference clock.
alt_xcvr_custom_is_serial_loopback_enabled	<service-path>	Returns whether the serial loopback mode of the previously specified channel is enabled.



Command	Arguments	Function
alt_xcvr_custom_set_serial_loopback_enabled	<service-path> <disable(0)/enable(1)>	Enables or disables the serial loopback mode of the previously specified channel.
alt_xcvr_custom_is_tx_pll_locked	<service-path>	Returns whether the transmitter PLL is locked to the reference clock.
alt_xcvr_reconfig_eye_viewer_get_logical_channel_address	<service-path>	Gets the logical channel address on which other alt_reconfig_eye_viewer commands will use to apply.
alt_xcvr_reconfig_eye_viewer_get_phase_step	<service-path>	Gets the current phase step of the previously specified channel.
alt_xcvr_reconfig_eye_viewer_is_enabled	<service-path>	Gets whether the EyeQ feature is enabled on the previously specified channel.
alt_xcvr_reconfig_eye_viewer_set_enabled	<service-path> <disable(0)/enable(1)>	Enables or disables the EyeQ feature on the previously specified channel.  Setting a value of 2 enables both EyeQ and the Serial Bit Comparator.
alt_xcvr_reconfig_eye_viewer_set_logical_channel_address	<service-path> <logical channel address>	Sets the logical channel address on which other alt_reconfig_eye_viewer commands will use to apply.
alt_xcvr_reconfig_eye_viewer_set_phase_step	<service-path> <phase step>	Sets the phase step of the previously specified channel.
alt_xcvr_reconfig_eye_viewer_has_ber_checker	<service-path>	Detects whether the eye viewer pointed to by <service-path> supports the Serial Bit Comparator.
alt_xcvr_reconfig_eye_viewer_ber_checker_is_enabled	<service-path>	Detects whether the Serial Bit Comparator is enabled.

Command	Arguments	Function
alt_xcvr_reconfig_eye_viewer_ber_checker_start	<service-path>	Starts the Serial Bit Comparator counters.
alt_xcvr_reconfig_eye_viewer_ber_checker_stop	<service-path>	Stops the Serial Bit Comparator counters.
alt_xcvr_reconfig_eye_viewer_ber_checker_reset_counters	<service-path>	Resets the Serial Bit Comparator counters.
alt_xcvr_reconfig_eye_viewer_ber_checker_is_running	<service-path>	Gets whether the Serial Bit Comparator counters are currently running or not.
alt_xcvr_reconfig_eye_viewer_ber_checker_get_data	<service-path>	Gets the current total bit, error bit, and exception counts for the Serial Bit Comparator.
alt_xcvr_reconfig_eye_viewer_has_ldeye	<service-path>	Detects whether the eye viewer pointed to by <service-path> supports 1D-EyeQ mode.
alt_xcvr_reconfig_eye_viewer_set_ldeye_mode	<service-path> <disable(0)/enable(1)>	Enables or disables 1D-EyeQ mode.
alt_xcvr_reconfig_eye_viewer_get_ldeye_mode	<service-path>	Gets the enable or disabled state of 1D-EyeQ mode.

Table 11-13: Channel Type Commands

Command	Arguments	Function
get_channel_type	<service-path><logical-channel-num>	Reports the detected type (GX/GT) of channel <logical-channel-num > for the reconfiguration block located at <service-path>.
set_channel_type	<service-path><logical-channel-num> <channel-type>	Overrides the detected channel type of channel <logical-channel-num > for the reconfiguration block located at <service-path> to the type specified (0:GX, 1:GT).

**Table 11-14: Loopback Commands**

Command	Arguments	Function
loopback_get	<service-path>	Returns the value of a setting or result on the loopback channel. Available results include: <ul style="list-style-type: none"> <li>• Status—running or stopped.</li> <li>• Bytes—number of bytes sent through the loopback channel.</li> <li>• Errors—number of errors reported by the loopback channel.</li> <li>• Seconds—number of seconds since the loopback channel was started.</li> </ul>
loopback_set	<service-path>	Sets the value of a setting controlling the loopback channel. Some settings are only supported by particular channel types. Available settings include: <ul style="list-style-type: none"> <li>• Timer—number of seconds for the test run.</li> <li>• Size—size of the test data.</li> <li>• Mode—mode of the test.</li> </ul>
loopback_start	<service-path>	Starts sending data through the loopback channel.
loopback_stop	<service-path>	Stops sending data through the loopback channel.

## Data Pattern Generator Commands

You can use Data Pattern Generator commands to control data patterns for debugging transceiver channels. You must instantiate the Data Pattern Generator component to support these commands.

**Table 11-15: Data Pattern Generator Commands**

Command	Arguments	Function
data_pattern_generator_start	<service-path>	Starts the data pattern generator.
data_pattern_generator_stop	<service-path>	Stops the data pattern generator.

Command	Arguments	Function
data_pattern_generator_is_generating	<service-path>	Returns non-zero if the generator is running.
data_pattern_generator_inject_error	<service-path>	Injects a 1-bit error into the generator output.
data_pattern_generator_set_pattern	<service-path> <pattern-name>	<p>Sets the output pattern specified by the &lt;pattern-name&gt;. In all, 6 patterns are available, 4 are pseudo-random binary sequences (PRBS), 1 is high frequency and 1 is low frequency.</p> <p>The PRBS7, PRBS15, PRBS23, PRBS31, HF (outputs high frequency, constant pattern of alternating 0s and 1s), and LF (outputs low frequency, constant pattern of 10b'1111100000 for 10-bit symbols and 8b'11110000 for 8-bit symbols) pattern names are defined.</p> <p>PRBS files are clear text and you can modify the PRBS files.</p>
data_pattern_generator_get_pattern	<service-path>	Returns currently selected output pattern.
data_pattern_generator_get_available_patterns	<service-path>	Returns a list of available data patterns by name.
data_pattern_generator_enable_preamble	<service-path>	Enables the preamble mode at the beginning of generation.
data_pattern_generator_disable_preamble	<service-path>	Disables the preamble mode at the beginning of generation.
data_pattern_generator_is_preamble_enabled	<service-path>	Returns a non-zero value if preamble mode is enabled.
data_pattern_generator_set_preamble_word	<preamble-word>	Sets the preamble word (could be 32-bit or 40-bit).
data_pattern_generator_get_preamble_word	<service-path>	Gets the preamble word.

Command	Arguments	Function
data_pattern_generator_set_preamble_beats	<service-path><number-of-preamble-beats>	Sets the number of beats to send out in the preamble word.
data_pattern_generator_get_preamble_beats	<service-path>	Returns the currently set number of beats to send out in the preamble word.
data_pattern_generator_fcnter_start	<service-path><max-cycles>	Sets the max cycle count and starts the frequency counter.
data_pattern_generator_check_status	<service-path>	Queries the data pattern generator for current status. Returns a bitmap indicating the status, with bits defined as follows: [0]-enabled, [1]-bypass enabled, [2]-avalon, [3]-sink ready, [4]-source valid, and [5]-frequency counter enabled.
data_pattern_generator_fcnter_report	<service-path><force-stop>	Reports the current measured clock ratio, stopping the counting first depending on <force-stop>.

## Data Pattern Checker Commands

You can use Data Pattern Checker commands to verify your generated data patterns. You must instantiate the Data Pattern Checker component to support these commands.

Table 11-16: Data Pattern Checker Commands

Command	Arguments	Function
data_pattern_checker_start	<service-path>	Starts the data pattern checker.
data_pattern_checker_stop	<service-path>	Stops the data pattern checker.
data_pattern_checker_is_checking	<service-path>	Returns a non-zero value if the checker is running.
data_pattern_checker_is_locked	<service-path>	Returns non-zero if the checker is locked onto the incoming data.

Command	Arguments	Function
<code>data_pattern_checker_set_pattern</code>	<code>&lt;service-path&gt; &lt;pattern-name&gt;</code>	Sets the expected pattern to the one specified by the <code>&lt;pattern-name&gt;</code> .
<code>data_pattern_checker_get_pattern</code>	<code>&lt;service-path&gt;</code>	Returns the currently selected expected pattern by name.
<code>data_pattern_checker_get_available_patterns</code>	<code>&lt;service-path&gt;</code>	Returns a list of available data patterns by name.
<code>data_pattern_checker_get_data</code>	<code>&lt;service-path&gt;</code>	Returns a list of the current checker data. The results are in the following order: number of bits, number of errors, and bit error rate.
<code>data_pattern_checker_reset_counters</code>	<code>&lt;service-path&gt;</code>	Resets the bit and error counters inside the checker.
<code>data_pattern_checker_fcnter_start</code>	<code>&lt;service-path&gt;&lt;max-cycles&gt;</code>	Sets the max cycle count and starts the frequency counter.
<code>data_pattern_checker_check_status</code>	<code>&lt;service-path&gt;</code> <code>&lt;service-path&gt;</code>	Queries the data pattern checker for current status. Returns a bitmap indicating status, with bits defined as follows: [0]-enabled, [1]-locked, [2]-bypass enabled, [3]-avalon, [4]-sink ready, [5]-source valid, and [6]-frequency counter enabled.
<code>data_pattern_checker_fcnter_report</code>	<code>&lt;service-path&gt;&lt;force-stop&gt;</code>	Reports the current measured clock ratio, stopping the counting first depending on <code>&lt;force-stop&gt;</code> .

## Revision History

Table 11-17: Document Revision History

Date	Version	Changes
December, 2014	14.1.0	<ul style="list-style-type: none"> <li>Added section about Arria 10 support and limitations.</li> </ul>

Date	Version	Changes
June, 2014	14.0.0	<ul style="list-style-type: none"> <li>Updated GUI changes for Channel Manager with popup menus, IP Catalog, Quartus II, and Qsys.</li> <li>Added ADME and JTAG debug link info for Arria 10.</li> <li>Added instructions to run Tcl script from command line.</li> <li>Added heat map display option.</li> <li>Added procedure to use internal PLL to generate reconfig_clk.</li> <li>Added note stating RX CDR PLL status can toggle in LTD mode.</li> </ul>
November, 2013	13.1.0	<ul style="list-style-type: none"> <li>Reorganization and conversion to DITA.</li> </ul>
May, 2013	13.0.0	<ul style="list-style-type: none"> <li>Added Conduit Mode Support, Serial Bit Comparator, Required Files and Tcl command tables.</li> </ul>
November, 2012	12.1.0	<ul style="list-style-type: none"> <li>Minor editorial updates. Added Tcl help information and removed Tcl command tables. Added 28-Gbps Transceiver support section.</li> </ul>
August, 2012	12.0.1	<ul style="list-style-type: none"> <li>General reorganization and revised steps in modifying Altera example designs.</li> </ul>
June, 2012	12.0.0	<ul style="list-style-type: none"> <li>Maintenance release for update of Transceiver Toolkit features.</li> </ul>
November, 2011	11.1.0	<ul style="list-style-type: none"> <li>Maintenance release for update of Transceiver Toolkit features.</li> </ul>
May, 2011	11.0.0	<ul style="list-style-type: none"> <li>Added new Tcl scenario.</li> </ul>
December, 2010	10.1.0	<ul style="list-style-type: none"> <li>Changed to new document template. Added new 10.1 release features.</li> </ul>
August, 2010	10.0.1	<ul style="list-style-type: none"> <li>Corrected links.</li> </ul>
July 2010	10.0.0	<ul style="list-style-type: none"> <li>Initial release.</li> </ul>

**Related Information**

[Quartus II Handbook Archive](#)

# Quick Design Debugging Using SignalProbe 12

2014.06.30

QII5V3



Subscribe



Send Feedback

## Quick Design Debugging Using SignalProbe

The SignalProbe incremental routing feature helps reduce the hardware verification process and time-to-market for system-on-a-programmable-chip (SOPC) designs. Easy access to internal device signals is important in the design or debugging process. The SignalProbe feature makes design verification more efficient by routing internal signals to I/O pins quickly without affecting the design. When you start with a fully routed design, you can select and route signals for debugging to either previously reserved or currently unused I/O pins.

The SignalProbe feature supports the Arria<sup>®</sup> series, Cyclone<sup>®</sup> series, MAX<sup>®</sup> II, and Stratix<sup>®</sup> series device families.

### Related Information

- [System Debugging Tools Overview documentation](#) on page 9-1  
Overview and comparison of all the tools available in the Quartus II software

## Design Flow Using SignalProbe

The SignalProbe feature allows you to reserve available pins and route internal signals to those reserved pins, while preserving the behavior of your design. SignalProbe is an effective debugging tool that provides visibility into your FPGA.

You can reserve pins for SignalProbe and assign I/O standards after a full compilation. Each SignalProbe-source to SignalProbe-pin connection is implemented as an engineering change order (ECO) that is applied to your netlist after a full compilation.

To route the internal signals to the device's reserved pins for SignalProbe, perform the following tasks:

1. Perform a full compilation.
2. Reserve SignalProbe Pins.
3. Assign SignalProbe sources.
4. Add registers between pipeline paths and Signalprobe pins.
5. Perform a SignalProbe compilation.
6. Analyze the results of a SignalProbe compilation.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered





## Perform a Full Compilation

You must complete a full compilation to generate an internal netlist containing a list of internal nodes to probe.

To perform a full compilation, on the Processing menu, click **Start Compilation**.

## Reserve SignalProbe Pins

SignalProbe pins can only be reserved after a full compilation. You can also probe any unused I/Os of the device. Assigning sources is a simple process after reserving SignalProbe pins. The sources for SignalProbe pins are the internal nodes and registers in the post-compilation netlist that you want to probe.

**Note:** Although you can reserve SignalProbe pins using many features within the Quartus II software, including the Pin Planner and the Tcl interface, you should use the **SignalProbe Pins** dialog box to create and edit your SignalProbe pins.

### Related Information

[SignalProbe online help](#)

## Assign SignalProbe Sources

A SignalProbe source can be any combinational node, register, or pin in your post-compilation netlist. To find a SignalProbe source, in the Node Finder, use the SignalProbe filter to remove all sources that cannot be probed. You might not be able to find a particular internal node because the node can be optimized away during synthesis, or the node cannot be routed to the SignalProbe pin. For example, you cannot probe nodes and registers within Gigabit transceivers in Stratix IV devices because there are no physical routes available to the pins.

**Note:** To probe virtual I/O pins generated in low-level partitions in an incremental compilation flow, select the source of the logic that feeds the virtual pin as your SignalProbe source pin.

Because SignalProbe pins are implemented and routed as ECOs, turning the **SignalProbe enable** option on or off is the same as selecting **Apply Selected Change** or **Restore Selected Change** in the Change Manager window. If the Change Manager window is not visible at the bottom of your screen, on the View menu, point to **Utility Windows** and click **Change Manager**.

### Related Information

- [SignalProbe Pins Dialog Box online help](#)
- [Add SignalProbe Pins Dialog Box online help](#)
- [Engineering Change Management with the Chip Planner documentation](#)

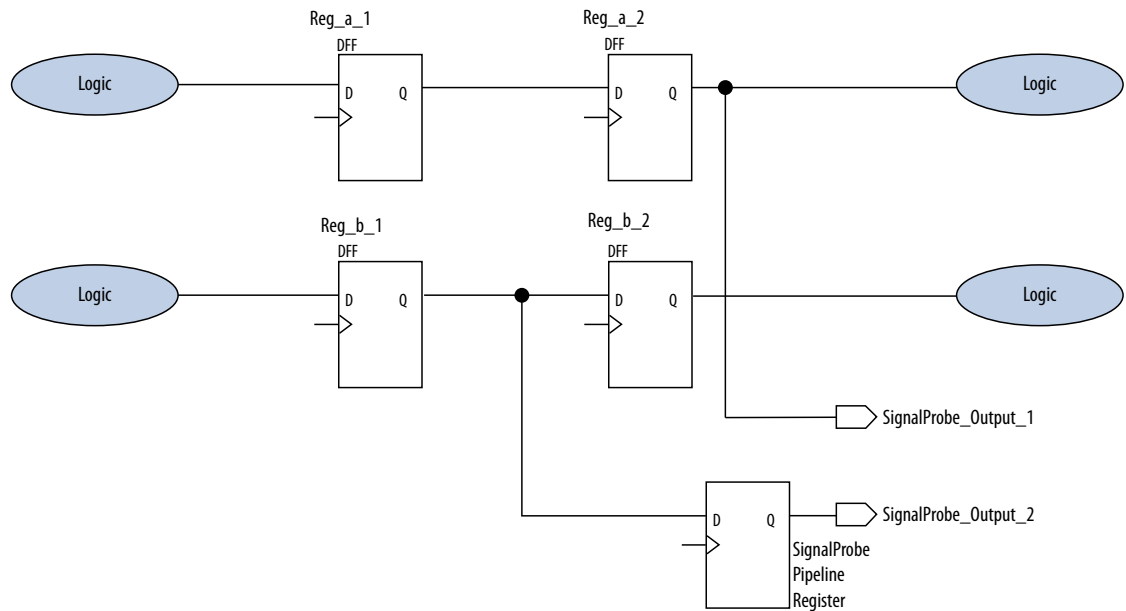
## Add Registers Between Pipeline Paths and SignalProbe Pins

You can specify the number of registers placed between a SignalProbe source and a SignalProbe pin. The registers synchronize data to a clock and control the latency of the SignalProbe outputs. The SignalProbe feature automatically inserts the number of registers specified into the SignalProbe path.

The figure shows a single register between the SignalProbe source `Reg_b_1` and SignalProbe `SignalProbe_Output_2` output pin added to synchronize the data between the two SignalProbe output pins.

**Note:** When you add a register to a SignalProbe pin, the SignalProbe compilation attempts to place the register to best meet timing requirements. You can place SignalProbe registers either near the SignalProbe source to meet  $f_{MAX}$  requirements, or near the I/O to meet  $t_{CO}$  requirements.

Figure 12-1: Synchronizing SignalProbe Outputs with a SignalProbe Register



In addition to clock input for pipeline registers, you can also specify a reset signal pin for pipeline registers. To specify a reset pin for pipeline registers, use the Tcl command `make_sp`.

#### Related Information

[Add SignalProbe Pins Dialog Box online help](#)

Information about how to pipeline an existing SignalProbe connection

## Perform a SignalProbe Compilation

Perform a SignalProbe compilation to route your SignalProbe pins. A SignalProbe compilation saves and checks all netlist changes without recompiling the other parts of the design. A SignalProbe compilation takes a fraction of the time of a full compilation to finish. The design's current placement and routing are preserved.

To perform a SignalProbe compilation, on the Processing menu, point to **Start** and click **Start SignalProbe Compilation**.

## Analyze the Results of a SignalProbe Compilation

After a SignalProbe compilation, the results are available in the compilation report file. Each SignalProbe pin is displayed in the **SignalProbe Fitting Result** page in the **Fitter** section of the Compilation Report. To view the status of each SignalProbe pin in the **SignalProbe Pins** dialog box, on the Tools menu, click **SignalProbe Pins**.

The status of each SignalProbe pin appears in the Change Manager window. If the Change Manager window is not visible at the bottom of your GUI, from the View menu, point to **Utility Windows** and click **Change Manager**.

Figure 12-2: Change Manager Window with SignalProbe Pins

Index	Node Name	Change Type	Old Value	Target Value	Current Value	Disk Value
1	signalprobe_1	SignalProbe	Disconnected	filterstate_m_inst1filter.idle	filterstate_m_inst1filter.idle	filterstate_m_inst1filter.idle
2	signalprobe_2	SignalProbe	Disconnected	filterstate_m_inst1filter.tap1	filterstate_m_inst1filter.tap1	filterstate_m_inst1filter.tap1
3	signalprobe_3	SignalProbe	Disconnected	filterstate_m_inst1filter.tap2	filterstate_m_inst1filter.tap2	filterstate_m_inst1filter.tap2
4	signalprobe_4	SignalProbe	Disconnected	filterstate_m_inst1filter.tap3	filterstate_m_inst1filter.tap3	filterstate_m_inst1filter.tap3
5	signalprobe_5	SignalProbe	Disconnected	filterstate_m_inst1filter.tap4	filterstate_m_inst1filter.tap4	filterstate_m_inst1filter.tap4

To view the timing results of each successfully routed SignalProbe pin, on the Processing menu, point to **Start** and click **Start Timing Analysis**.

#### Related Information

[Engineering Change Management with the Chip Planner documentation](#)

## What a SignalProbe Compilation Does

After a full compilation, you can start a SignalProbe compilation either manually or automatically. A SignalProbe compilation performs the following functions:

- Validates SignalProbe pins
- Validates your specified SignalProbe sources
- Adds registers into SignalProbe paths, if applicable
- Attempts to route from SignalProbe sources through registers to SignalProbe pins

To run the SignalProbe compilation immediately after a full compilation, on the Tools menu, click **SignalProbe Pins**. In the **SignalProbe Pins** dialog box, click **Start Check & Save All Netlist Changes**.

To run a SignalProbe compilation manually after a full compilation, on the Processing menu, point to **Start** and click **Start SignalProbe Compilation**.

**Note:** You must run the Fitter before a SignalProbe compilation. The Fitter generates a list of all internal nodes that can serve as SignalProbe sources.

Turn the **SignalProbe enable** option on or off in the **SignalProbe Pins** dialog box to enable or disable each SignalProbe pin.

## Understanding the Results of a SignalProbe Compilation


After a SignalProbe compilation, the results appear in two sections of the compilation report file. The fitting results and status of each SignalProbe pin appears in the **SignalProbe Fitting Result** screen in the Fitter section of the Compilation Report.

Table 12-1: Status Values

Status	Description
Routed	Connected and routed successfully
Not Routed	Not enabled


Status	Description
Failed to Route	Failed routing during last SignalProbe compilation
Need to Compile	Assignment changed since last SignalProbe compilation

Figure 12-3: SignalProbe Fitting Results Page in the Compilation Report Window



The **SignalProbe source to output delays** screen in the Timing Analysis section of the Compilation Report displays the timing results of each successfully routed SignalProbe pin.

Figure 12-4: SignalProbe Source to Output Delays Page in the Compilation Report Window



**Note:** After a SignalProbe compilation, the processing screen of the Messages window also provides the results for each SignalProbe pin and displays slack information for each successfully routed SignalProbe pin.

## Analyzing SignalProbe Routing Failures

A SignalProbe compilation can fail for any of the following reasons:

- **Route unavailable**—the SignalProbe compilation failed to find a route from the SignalProbe source to the SignalProbe pin because of routing congestion.
- **Invalid or nonexistent SignalProbe source**—you entered a SignalProbe source that does not exist or is invalid.
- **Unusable output pin**—the output pin selected is found to be unusable.

Routing failures can occur if the SignalProbe pin's I/O standard conflicts with other I/O standards in the same I/O bank.

If routing congestion prevents a successful SignalProbe compilation, you can allow the compiler to modify routing to the specified SignalProbe source. On the Tools menu, click **SignalProbe Pins** and turn on **Modify latest fitting results during SignalProbe compilation**. This setting allows the Fitter to modify existing routing channels used by your design.

**Note:** Turning on **Modify latest fitting results during SignalProbe compilation** can change the performance of your design.

## Scripting Support

You can also run some procedures at a command prompt. For detailed information about scripting command options, refer to the Quartus II command-line and Tcl API Help browser. To run the Help browser, type the following command at the command prompt:

```
quartus_sh --qhelp
```

**Note:** The Tcl commands in this section are part of the `::quartus::chip_planner` Quartus II Tcl API. Source or include the `::quartus::chip_planner` Tcl package in your scripts to make these commands available.

### Related Information

- [Tcl Scripting documentation](#)
- [Quartus II Settings File Reference Manual](#)  
Information about all settings and constraints in the Quartus II software
- [Command-Line Scripting documentation](#)

## Making a SignalProbe Pin

To make a SignalProbe pin, type the following command:

```
make_sp [-h | -help] [-long_help] [-clk <clk>] [-io_std <io_std>] \  
-loc <loc> -pin_name <pin name> [-regs <regs>] [-reset <reset>] \  
-src_name <source name>
```

## Deleting a SignalProbe Pin

To delete a SignalProbe pin, type the following Tcl command:

```
delete_sp [-h | -help] [-long_help] -pin_name <pin name>
```

## Enabling a SignalProbe Pin

To enable a SignalProbe pin, type the following Tcl command:

```
enable_sp [-h | -help] [-long_help] -pin_name <pin name>
```

## Disabling a SignalProbe Pin

To disable a SignalProbe pin, type the following Tcl command:

```
disable_sp [-h | -help] [-long_help] -pin_name <pin name>
```

## Performing a SignalProbe Compilation

To perform a SignalProbe compilation, type the following command:

```
quartus_sh --flow signalprobe <project name>
```

## Script Example

The example shows a script that creates a SignalProbe pin called `sp1` and connects the `sp1` pin to source node `reg1` in a project that was already compiled.

### Creating a SignalProbe Pin Called `sp1`

```
package require ::quartus::chip_planner
project_open project
read_netlist
make_sp -pin_name sp1 -src_name reg1
check_netlist_and_save
project_close
```

## Reserving SignalProbe Pins

To reserve a SignalProbe pin, add the commands shown in the example to the Quartus II Settings File (`.qsf`) for your project.

### Reserving a SignalProbe Pin

```
set_location_assignment <location> -to <SignalProbe pin name>
set_instance_assignment -name RESERVE_PIN \
"AS SIGNALPROBE OUTPUT" -to <SignalProbe pin name>
```

Valid locations are pin location names, such as `Pin_A3`.

## Common Problems When Reserving a SignalProbe Pin

If you cannot reserve a SignalProbe pin in the Quartus II software, it is likely that one of the following is true:

- You have selected multiple pins.
- A compilation is running in the background. Wait until the compilation is complete before reserving the pin.
- You have the Quartus II Web Edition software, in which the SignalProbe feature is not enabled by default. You must turn on TalkBack to enable the SignalProbe feature in the Quartus II Web Edition software.
- You have not set the pin reserve type to **As Signal Probe Output**. To reserve a pin, on the Assignments menu, in the **Assign Pins** dialog box, select **As SignalProbe Output**.
- The pin is reserved from a previous compilation. During a compilation, the Quartus II software reserves each pin on the targeted device. If you end the Quartus II process during a compilation, for example, with the **Windows Task Manager End Process** command or the UNIX `kill` command, perform a full recompilation before reserving pins as SignalProbe outputs.
- The pin does not support the SignalProbe feature. Select another pin.
- The current device family does not support the SignalProbe feature.

## Adding SignalProbe Sources

To assign the node name to a SignalProbe pin, type the following Tcl command:

```
set_instance_assignment -name SIGNALPROBE_SOURCE <node name> \
-to <SignalProbe pin name>
```

The next command turns on SignalProbe routing. To turn off individual SignalProbe pins, specify `OFF` instead of `ON` with the following command:

```
set_instance_assignment -name SIGNALPROBE_ENABLE ON \
-to <SignalProbe pin name>
```

### Related Information

- [SignalProbe Pins Dialog Box online help](#)
- [Add SignalProbe Pins Dialog Box online help](#)  
Information about how to pipeline an existing SignalProbe connection

## Assigning I/O Standards

To assign an I/O standard to a pin, type the following Tcl command:

```
set_instance_assignment -name IO_STANDARD <I/O standard> -to <SignalProbe pin name>
```

### Related Information

[I/O Standards online help](#)

## Adding Registers for Pipelining

To add registers for pipelining, type the following Tcl command:

```
set_instance_assignment -name SIGNALPROBE_CLOCK <clock name> \
-to <SignalProbe pin name>

set_instance_assignment \
-name SIGNALPROBE_NUM_REGISTERS <number of registers> -to <SignalProbe pin name>
```

## Running SignalProbe Immediately After a Full Compilation

To run SignalProbe immediately after a full compilation, type the following Tcl command:

```
set_global_assignment -name SIGNALPROBE_DURING_NORMAL_COMPILATION ON
```

## Running SignalProbe Manually

To run SignalProbe as part of a scripted flow using Tcl, use the following in your script:

```
execute_flow -signalprobe
```

To perform a Signal Probe compilation interactively at a command prompt, type the following command:

```
quartus_sh_fit --flow signalprobe <project name>
```

## Enabling or Disabling All SignalProbe Routing

Use the Tcl command in the example to turn on or turn off SignalProbe routing. When using this command, to turn SignalProbe routing on, specify `ON`. To turn SignalProbe routing off, specify `OFF`.

### Turning SignalProbe On or Off with Tcl Commands

```
set spe [get_all_assignments -name SIGNALPROBE_ENABLE] \
foreach_in_collection asgn $spe {
    set signalprobe_pin_name [lindex $asgn 2]
    set_instance_assignment -name SIGNALPROBE_ENABLE \
        -to $signalprobe_pin_name <ON/OFF> }
```

## Allowing SignalProbe to Modify Fitting Results

To turn on **Modify latest fitting results**, type the following Tcl command:

```
set_global_assignment -name SIGNALPROBE_ALLOW_OVERUSE ON
```

## Document Revision History

Table 12-2: Document Revision History

Date	Version	Changes
June 2014	14.0.0	Dita conversion.
May 2013	13.0.0	Changed sequence of flow to clarify that you need to perform a full compilation before reserving SignalProbe pins. Affected sections are “Debugging Using the SignalProbe Feature” on page 12–1 and “Reserving SignalProbe Pins” on page 12–2. Moved “Performing a Full Compilation” on page 12–2 before “Reserving SignalProbe Pins” on page 12–2.
June 2012	12.0.0	Removed survey link.
November 2011	10.0.2	Template update.



Date	Version	Changes
December 2010	10.0.1	Changed to new document template.
July 2010	10.0.0	<ul style="list-style-type: none"> <li>Revised for new UI.</li> <li>Removed section SignalProbe ECO flows</li> <li>Removed support for SignalProbe pin preservation when recompiling with incremental compilation turned on.</li> <li>Removed outdated FAQ section.</li> <li>Added links to Quartus II Help for procedural content.</li> </ul>
November 2009	9.1.0	<ul style="list-style-type: none"> <li>Removed all references and procedures for APEX devices.</li> <li>Style changes.</li> </ul>
March 2009	9.0.0	<ul style="list-style-type: none"> <li>Removed the “Generate the Programming File” section</li> <li>Removed unnecessary screenshots</li> <li>Minor editorial updates</li> </ul>
November 2008	8.1.0	<ul style="list-style-type: none"> <li>Modified description for preserving SignalProbe connections when using Incremental Compilation</li> <li>Added plausible scenarios where SignalProbe connections are not reserved in the design</li> </ul>
May 2008	8.0.0	<ul style="list-style-type: none"> <li>Added “Arria GX” to the list of supported devices</li> <li>Removed the “On-Chip Debugging Tool Comparison” and replaced with a reference to the Section V Overview on page 13–1</li> <li>Added hyperlinks to referenced documents throughout the chapter</li> <li>Minor editorial updates</li> </ul>

**Related Information**[Quartus II Handbook Archive](#)

For previous versions of the Quartus II Handbook

# Design Debugging Using the SignalTap II Logic Analyzer 13

2014.12.15

QII5V3



Subscribe



Send Feedback

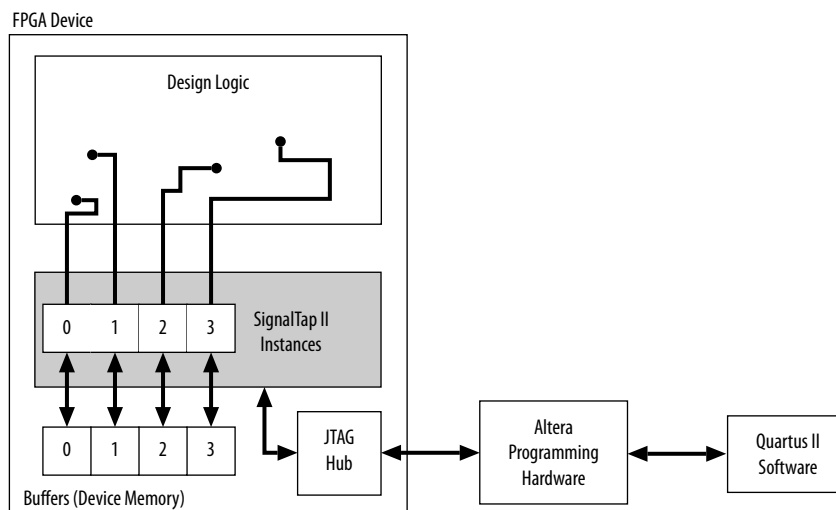
## About the SignalTap II Logic Analyzer

Altera provides the SignalTap<sup>®</sup> II Logic Analyzer to help with design debugging. This logic analyzer allows you to examine the behavior of internal signals, without using extra I/O pins, while the design is running at full speed on an FPGA device.

The SignalTap II Logic Analyzer is scalable, easy to use, and available as a stand-alone package or included with the Quartus<sup>®</sup> II software subscription. This logic analyzer helps debug an FPGA design by probing the state of the internal signals in the design without the use of external equipment. Defining custom trigger-condition logic provides greater accuracy and improves the ability to isolate problems. The SignalTap II Logic Analyzer does not require external probes or changes to the design files to capture the state of the internal nodes or I/O pins in the design. All captured signal data is conveniently stored in device memory until you are ready to read and analyze the data.

The SignalTap II Logic Analyzer is a next-generation, system-level debugging tool that captures and displays real-time signal behavior in a system-on-a-programmable-chip (SOPC) or any FPGA design. The SignalTap II Logic Analyzer supports the highest number of channels, largest sample depth, and fastest clock speeds of any logic analyzer in the programmable logic market.

Figure 13-1: SignalTap II Logic Analyzer Block Diagram



© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered

ALTERA<sup>®</sup>

Note to figure :

1. This diagram assumes that you compiled the SignalTap II Logic Analyzer with the design as a separate design partition using the Quartus II incremental compilation feature. This is the default setting for new projects in the Quartus II software. If incremental compilation is disabled or not used, the SignalTap II logic is integrated with the design.

This chapter is intended for any designer who wants to debug an FPGA design during normal device operation without the need for external lab equipment. Because the SignalTap II Logic Analyzer is similar to traditional external logic analyzers, familiarity with external logic analyzer operations is helpful, but not necessary. To take advantage of faster compile times when making changes to the SignalTap II Logic Analyzer, knowledge of the Quartus II incremental compilation feature is helpful.

#### Related Information

[Quartus II Incremental Compilation for Hierarchical and Team-Based Design documentation](#)

## Hardware and Software Requirements

You need the following components to perform logic analysis with the SignalTap II Logic Analyzer:

- Software:
  - Quartus II design software
    - or
  - Quartus II Web Edition (with the TalkBack feature enabled)
    - or
  - SignalTap II Logic Analyzer standalone software and standalone Programmer software.
- Download/upload cable
- Altera® development kit or your design board with JTAG connection to device under test

**Note:** The Quartus II software Web Edition does not support the SignalTap II Logic Analyzer with the incremental compilation feature.

The memory blocks of the device store captured data and transfers the data to the Quartus II software waveform display with a JTAG communication cable, such as EthernetBlaster or USB-Blaster™.

**Table 13-1: SignalTap II Logic Analyzer Features and Benefits**

Feature	Benefit
Toolbar with commonly used menu items.	Single-click operation of commonly used menu items. Hover over the icons to see tool tips.
Multiple logic analyzers in a single device	Captures data from multiple clock domains in a design at the same time.
Multiple logic analyzers in multiple devices in a single JTAG chain	Simultaneously captures data from multiple devices in a JTAG chain.
Plug-In Support	Easily specifies nodes, triggers, and signal mnemonics for IP, such as the Nios® II processor.

Feature	Benefit
Up to 10 basic or advanced trigger conditions for each analyzer instance	Enables sending more complex data capture commands to the logic analyzer, providing greater accuracy and problem isolation.
Power-Up Trigger	Captures signal data for triggers that occur after device programming, but before manually starting the logic analyzer.
Custom trigger HDL object	You can code your own trigger in Verilog or VHDL and tap specific instances of modules located anywhere in the hierarchy of your design without needing to manually route all the necessary connections. This simplifies the process of tapping nodes spread out across your design.
State-based Triggering Flow	Enables you to organize your triggering conditions to precisely define what your logic analyzer captures.
Incremental Compilation	Modifies the SignalTap II Logic Analyzer monitored signals and triggers without performing a full compilation, saving time.
Incremental Route with Rapid Recompile	Manually allocate trigger input, data input, storage filter input node count, and perform a full compilation to include the SignalTap II Logic Analyzer in your design, then you can selectively connect, disconnect, and swap to different nodes in your design. Use Rapid Recompile to perform incremental routing and gain a 2-4x speedup over the initial full compilation.
Flexible buffer acquisition modes	The buffer acquisition control allows you to precisely control the data that is written into the acquisition buffer. Both segmented buffers and non-segmented buffers with storage qualification allow you to discard data samples that are not relevant to the debugging of your design.
MATLAB integration with included MEX function	Collects the SignalTap II Logic Analyzer captured data into a MATLAB integer matrix.
Up to 2,048 channels per logic analyzer instance	Samples many signals and wide bus structures.
Up to 128K samples in each device	Captures a large sample set for each channel.
Fast clock frequencies	Synchronous sampling of data nodes using the same clock tree driving the logic under test.
Resource usage estimator	Provides estimate of logic and memory device resources used by SignalTap II Logic Analyzer configurations.
No additional cost	The SignalTap II Logic Analyzer is included with a Quartus II subscription and with the Quartus II Web Edition (with TalkBack enabled).
Compatibility with other on-chip debugging utilities	You can use the SignalTap II Logic Analyzer in tandem with any JTAG-based on-chip debugging tool, such as an In-System Memory Content editor, allowing you to change signal values in real-time while you are running an analysis with the SignalTap II Logic Analyzer.

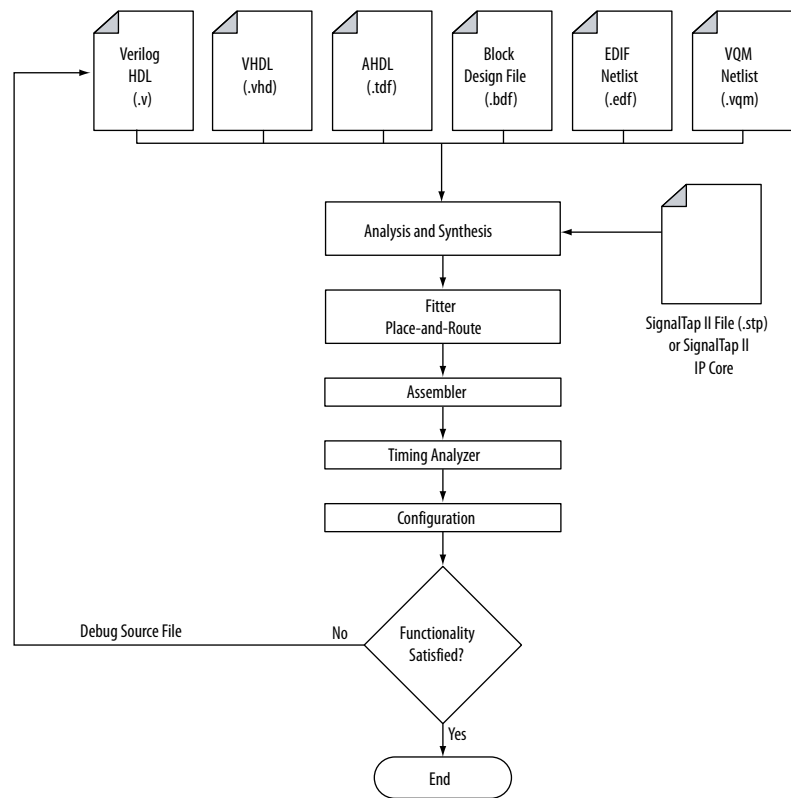
### Related Information

- [System Debugging Tools Overview documentation](#) on page 9-1  
Overview and comparison of all tools available in the In-System Verification Tool set

## Design Flow Using the SignalTap II Logic Analyzer

**Figure 13-2** shows a typical overall FPGA design flow using the SignalTap II Logic Analyzer. A SignalTap II file (.stp) is added to and enabled in your project, or a SignalTap II IP core, created with the IP Catalog, is instantiated in your design. **Figure 13-2** shows the flow of operations from initially adding the SignalTap II Logic Analyzer to your design to final device configuration, testing, and debugging.

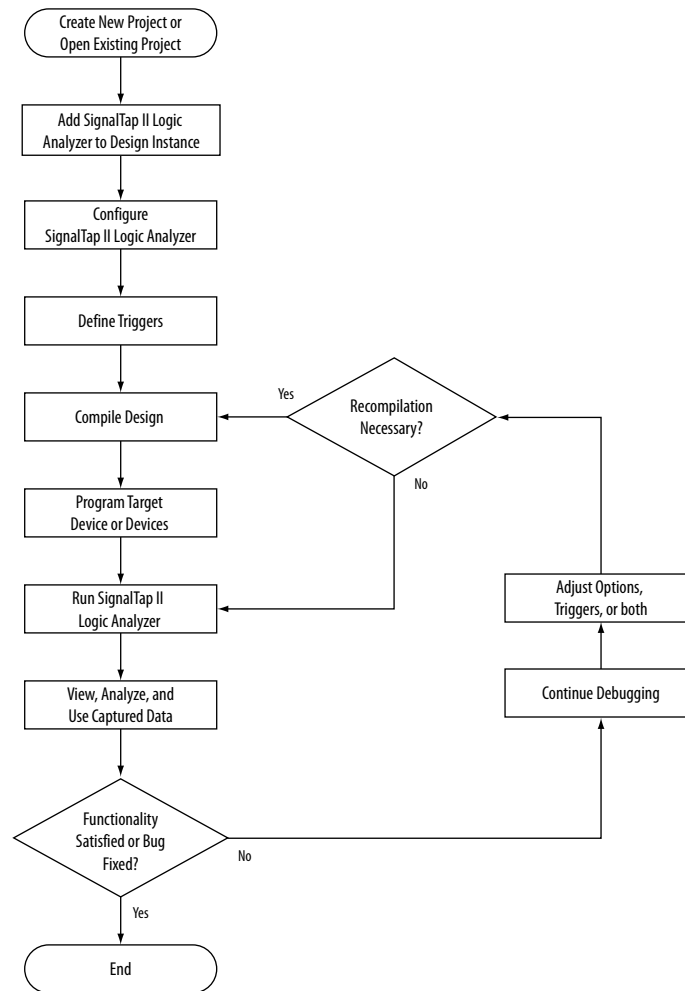
**Figure 13-2: SignalTap II FPGA Design and Debugging Flow**



## SignalTap II Logic Analyzer Task Flow

To use the SignalTap II Logic Analyzer to debug your design, you perform a number of tasks to add, configure, and run the logic analyzer. **Figure 13-3** shows a typical flow of the tasks you complete to debug your design.

Figure 13-3: SignalTap II Logic Analyzer Task Flow



## Add the SignalTap II Logic Analyzer to Your Design

Create an **.stp** or create a parameterized HDL instance representation of the logic analyzer using the IP Catalog. If you want to monitor multiple clock domains simultaneously, add additional instances of the logic analyzer to your design, limited only by the available resources in your device.

### Related Information

[Setting Up the SignalTap II Logic Analyzer online help](#)

## Configure the SignalTap II Logic Analyzer

After you add the SignalTap II Logic Analyzer to your design, configure the logic analyzer to monitor the signals you want. You can manually add signals or use a plug-in, such as the Nios II processor plug-in, to quickly add entire sets of associated signals for a particular intellectual property (IP). You can also specify settings for the data capture buffer, such as its size, the method in which data is captured and stored, and the device memory type to use for the buffer in devices that support memory type selection.

**Related Information**

[Creating a Power-Up Trigger](#) on page 13-44

## Define Trigger Conditions

The SignalTap II Logic Analyzer captures data continuously while the logic analyzer is running. To capture and store specific signal data, set up triggers that tell the logic analyzer under what conditions to stop capturing data. The SignalTap II Logic Analyzer allows you to define trigger conditions that range from very simple, such as the rising edge of a single signal, to very complex, involving groups of signals, extra logic, and multiple conditions. Power-Up Triggers allow you to capture data from trigger events occurring immediately after the device enters user-mode after configuration.

## Compile the Design

With the **.stp** configured and trigger conditions defined, compile your project as usual to include the logic analyzer in your design. Because you may need to change monitored signal nodes or adjust trigger settings frequently during debugging, Altera recommends that you use the incremental compilation feature built into the SignalTap II Logic Analyzer, along with Quartus II incremental compilation, to reduce recompile times. You can also use Incremental Route with Rapid Recompile to reduce recompile times.

**Related Information**

[Compiling a Design that Contains a SignalTap II Logic Analyzer online help](#)

## Program the Target Device or Devices

When you debug a design with the SignalTap II Logic Analyzer, you can program a target device directly from the **.stp** without using the Quartus II Programmer. You can also program multiple devices with different designs and simultaneously debug them.

**Note:** The SignalTap II Logic Analyzer supports all current Altera FPGA device families including Arria<sup>®</sup>, Cyclone<sup>®</sup>, MAX<sup>®</sup> 10 and Stratix<sup>®</sup> devices.

**Related Information**

- [Managing Multiple SignalTap II Files and Configurations](#) on page 13-22
- [Running the SignalTap II Logic Analyzer online help](#)

## Run the SignalTap II Logic Analyzer

In normal device operation, you control the logic analyzer through the JTAG connection, specifying when to start looking for trigger conditions to begin capturing data. With Runtime or Power-Up Triggers, read and transfer the captured data from the on-chip buffer to the **.stp** for analysis.

**Related Information**

[Analyzing Data in the SignalTap II Logic Analyzer online help](#)

## View, Analyze, and Use Captured Data

After you have captured data and read it into the **.stp**, that data is available for analysis and debugging. Set up mnemonic tables, either manually or with a plug-in, to simplify reading and interpreting the captured signal data. To speed up debugging, use the **Locate** feature in the **SignalTap II node** list to find the locations of problem nodes in other tools in the Quartus II software. Save the captured data for later analysis, or convert the data to other formats for sharing and further study.

## Embedding Multiple Analyzers in One FPGA

The SignalTap II Logic Analyzer Editor includes support for adding multiple logic analyzers by creating instances in the `.stp`. You can create a unique logic analyzer for each clock domain in the design.

## Monitoring FPGA Resources Used by the SignalTap II Logic Analyzer

The SignalTap II Logic Analyzer has a built-in resource estimator that calculates the logic resources and amount of memory that each logic analyzer instance uses. Furthermore, because the most demanding on-chip resource for the logic analyzer is memory usage, the resource estimator reports the ratio of total RAM usage in your design to the total amount of RAM available, given the results of the last compilation. The resource estimator provides a warning if a potential for a “no-fit” occurs.

You can see resource usage of each logic analyzer instance and total resources used in the columns of the **Instance Manager** pane of the SignalTap II Logic Analyzer Editor. Use this feature when you know that your design is running low on resources.

The logic element value reported in the resource usage estimator may vary by as much as 10% from the actual resource usage.

## Using the IP Catalog to Create Your Logic Analyzer

You can create a SignalTap II Logic Analyzer instance by using the IP Catalog. The IP Catalog generates an HDL file that you instantiate in your design.

**Note:** The State-based trigger flow, the state machine debugging feature, and the storage qualification feature are not supported when using the IP Catalog to create the logic analyzer.

## Configure the SignalTap II Logic Analyzer

There are many ways to configure instances of the SignalTap II Logic Analyzer. Some of the settings are similar to those found on traditional external logic analyzers. Other settings are unique to the SignalTap II Logic Analyzer because of the requirements for configuring a logic analyzer. All settings allow you to configure the logic analyzer the way you want to help debug your design.

**Note:** Some settings can only be adjusted when you are viewing Run-Time Trigger conditions instead of Power-Up Trigger conditions.

## Assigning an Acquisition Clock

Assign a clock signal to control the acquisition of data by the SignalTap II Logic Analyzer. The logic analyzer samples data on every positive (rising) edge of the acquisition clock. The logic analyzer does not support sampling on the negative (falling) edge of the acquisition clock. You can use any signal in your design as the acquisition clock. However, for best results, Altera recommends that you use a global, non-gated clock synchronous to the signals under test for data acquisition. Using a gated clock as your acquisition clock can result in unexpected data that does not accurately reflect the behavior of your design. The Quartus II static timing analysis tools show the maximum acquisition clock frequency at which you can run your design. Refer to the Timing Analysis section of the Compilation Report to find the maximum frequency of the logic analyzer clock.

**Note:** Altera recommends that you exercise caution when using a recovered clock from a transceiver as an acquisition clock for the SignalTap II Logic Analyzer. Incorrect or unexpected behavior has



been noted, particularly when a recovered clock from a transceiver is used as an acquisition clock with the power-up trigger feature.

If you do not assign an acquisition clock in the SignalTap II Logic Analyzer Editor, the Quartus II software automatically creates a clock pin called `auto_stp_external_clk`.

You must make a pin assignment to this pin independently from the design. Ensure that a clock signal in your design drives the acquisition clock.

#### Related Information

- [Working with Nodes in the SignalTap II Logic Analyzer online help](#)  
Information about assigning an acquisition clock
- [Managing Device I/O Pins documentation](#)  
Information about assigning signals to pins

## Adding Signals to the SignalTap II File

While configuring the logic analyzer, add signals to the node list in the `.stp` to select which signals in your design you want to monitor. You can also select signals to define triggers. You can assign the following two types of signals to your `.stp` file:

- **Pre-synthesis**—These signals exist after design elaboration, but before any synthesis optimizations are done. This set of signals should reflect your Register Transfer Level (RTL) signals.
- **Post-fitting**—This signal exists after physical synthesis optimizations and place-and-route.

**Note:** If you are not using incremental compilation, add only pre-synthesis signals to the `.stp`. Using pre-synthesis helps when you want to add a new node after you change a design. Source file changes appear in the Node Finder after you perform an Analysis and Elaboration. On the Processing Menu, point to **Start** and click **Start Analysis & Elaboration**.

The Quartus II software does not limit the number of signals available for monitoring in the SignalTap II window waveform display. However, the number of channels available is directly proportional to the number of logic elements (LEs) or adaptive logic modules (ALMs) in the device. Therefore, there is a physical restriction on the number of channels that are available for monitoring. Signals shown in blue text are post-fit node names. Signals shown in black text are pre-synthesis node names.

After successful Analysis and Elaboration, invalid signals are displayed in red. Unless you are certain that these signals are valid, remove them from the `.stp` for correct operation. The SignalTap II Status Indicator also indicates if an invalid node name exists in the `.stp`.

You can tap signals if a routing resource (row or column interconnects) exists to route the connection to the SignalTap II instance. For example, signals that exist in the I/O element (IOE) cannot be directly tapped because there are no direct routing resources from the signal in an IOE to a core logic element. For input pins, you can tap the signal that is driving a logic array block (LAB) from an IOE, or, for output pins, you can tap the signal from the LAB that is driving an IOE.

When adding pre-synthesis signals, make all connections to the SignalTap II Logic Analyzer before synthesis. Logic and routing resources are allocated during recompilation to make the connection as if a change in your design files had been made. Pre-synthesis signal names for signals driving to and from IOEs coincide with the signal names assigned to the pin.

In the case of post-fit signals, connections that you make to the SignalTap II Logic Analyzer are the signal names from the actual atoms in your post-fit netlist. You can only make a connection if the signals are

part of the existing post-fit netlist and existing routing resources are available from the signal of interest to the SignalTap II Logic Analyzer. In the case of post-fit output signals, tap the `COMBOUT` or `REGOUT` signal that drives the IOE block. For post-fit input signals, signals driving into the core logic coincide with the signal name assigned to the pin.

**Note:** Because NOT-gate push back applies to any register that you tap, the signal from the atom may be inverted. You can check this by locating the signal in either the Resource Property Editor or the Technology Map Viewer. The Technology Map viewer and the Resource Property Editor can also be used to help you find post-fit node names.

#### Related Information

- [Faster Compilations with Quartus II Incremental Compilation](#) on page 13-47
- [Incremental Compilation online help](#)
- [Analyzing Designs with Quartus II Netlist Viewers documentation](#)  
Information about cross-probing to source design files and other Quartus II windows

## Signal Preservation

Many of the RTL signals are optimized during the process of synthesis and place-and-route. RTL signal names frequently may not appear in the post-fit netlist after optimizations. For example, the compilation process can add tildes (“~”) to nets that fan-out from a node, making it difficult to decipher which signal nets they actually represent. These process results can cause problems when you use the incremental compilation flow with the SignalTap II Logic Analyzer. Because you can only add post-fitting signals to the SignalTap II Logic Analyzer in partitions of type **post-fit**, RTL signals that you want to monitor may not be available, preventing their use. To avoid this issue, use synthesis attributes to preserve signals during synthesis and place-and-route. When the Quartus II software encounters these synthesis attributes, it does not perform any optimization on the specified signals, forcing them to continue to exist in the post-fit netlist. However, if you do this, you could see an increase in resource utilization or a decrease in timing performance. The two attributes you can use are:

- `keep`—Ensures that combinational signals are not removed
- `preserve`—Ensures that registers are not removed

If you are debugging an IP core, such as the Nios II CPU or other encrypted IP, you might need to preserve nodes from the core to make them available for debugging with the SignalTap II Logic Analyzer. Preserving nodes is often necessary when a plug-in is used to add a group of signals for a particular IP.

If you use incremental compilation flow with the SignalTap II Logic Analyzer, pre-synthesis nodes may not be connected to the SignalTap II Logic Analyzer if the affected partition is of the post-fit type. A critical warning is issued for all pre-synthesis node names that are not found in the post-fit netlist.

#### Related Information

- [Quartus II Integrated Synthesis documentation](#)  
Information about using signal preservation attributes
- [Working with Nodes in the SignalTap II Logic Analyzer online help](#)

## Assigning Data Signals Using the Technology Map Viewer

You can easily add post-fit signal names that you find in the Technology map viewer. To do so, launch the Technology map viewer (post-fitting) after compiling your design. When you find the desired node, copy the node to either the active `.stp` for your design or a new `.stp`

## Node List Signal Use Options

When a signal is added to the node list, you can select options that specify how the signal is used with the logic analyzer. You can turn off the ability of a signal to trigger the analyzer by disabling the **Trigger Enable** option for that signal in the node list in the **.stp**. This option is useful when you want to see only the captured data for a signal and you are not using that signal as part of a trigger.

You can turn off the ability to view data for a signal by disabling the **Data Enable** column. This option is useful when you want to trigger on a signal, but have no interest in viewing data for that signal.

### Related Information

[Define Triggers](#) on page 13-23

### Disabling and Enabling a SignalTap II Instance

From the **Instance Manager** pane, you can disable and enable SignalTap II instances. Physically adding or removing instances requires recompilation after disabling and enabling a SignalTap II instance.

### Untappable Signals

Not all of the post-fitting signals in your design are available in the **SignalTap II : post-fitting filter** in the **Node Finder** dialog box. The following signal types cannot be tapped:

- **Post-fit output pins**—You cannot tap a post-fit output pin directly. To make an output signal visible, tap the register or buffer that drives the output pin. This includes pins defined as bidirectional.
- **Signals that are part of a carry chain**—You cannot tap the carry out (`cout0` or `cout1`) signal of a logic element. Due to architectural restrictions, the carry out signal can only feed the carry in of another LE.
- **JTAG Signals**—You cannot tap the JTAG control (`TCK`, `TDI`, `TDO`, and `TMS`) signals.
- **ALTGXB IP core**—You cannot directly tap any ports of an ALTGXB instantiation.
- **LVDS**—You cannot tap the data output from a serializer/deserializer (SERDES) block.
- **DQ, DQS Signals**—You cannot directly tap the `DQ` or `DQS` signals in a DDR/DDR II design.

### Adding Signals with a Plug-In

Instead of adding individual or grouped signals through the **Node Finder**, you can add groups of relevant signals of a particular type of IP with a plug-in. The SignalTap II Logic Analyzer comes with one plug-in already installed for the Nios II processor. Besides easy signal addition, plug-ins also provide features such as pre-designed mnemonic tables, useful for trigger creation and data viewing, as well as the ability to disassemble code in captured data.

The Nios II plug-in, for example, creates one mnemonic table in the **Setup** tab and two tables in the **Data** tab:

- **Nios II Instruction (Setup tab)**—Capture all the required signals for triggering on a selected instruction address.
- **Nios II Instance Address (Data tab)**—Display address of executed instructions in hexadecimal format or as a programming symbol name if defined in an optional Executable and Linking Format (**.elf**) file.
- **Nios II Disassembly (Data tab)**—Displays disassembled code from the corresponding address.

To add signals to the **.stp** using a plug-in, perform the following steps after running Analysis and Elaboration on your design:

1. Right-click in the node list. On the Add Nodes with Plug-In submenu, choose the plug-in you want to use, such as the included plug-in named **Nios II**.

**Note:** If the IP for the selected plug-in does not exist in your design, a message informs you that you cannot use the selected plug-in.

2. The **Select Hierarchy Level** dialog box appears showing the IP hierarchy of your design. Select the IP that contains the signals you want to monitor with the plug-in and click **OK**.
3. If all the signals in the plug-in are available, a dialog box might appear, depending on the plug-in selected, where you can specify options for the plug-in. With the Nios II plug-in, you can optionally select an **.elf** containing program symbols from your Nios II Integrated Development Environment (IDE) software design. Specify options for the selected plug-in as desired and click **OK**.

**Note:** To make sure all the required signals are available, in the Quartus II **Analysis & Synthesis** settings, click **Assignments > Settings > Compiler Settings > Advanced Settings (Synthesis)**. Turn on **Create debugging nodes for IP cores**.

All the signals included in the plug-in are added to the node list.

**Related Information**

- [Define Triggers](#) on page 13-23
- [View, Analyze, and Use Captured Data](#) on page 13-6


## Adding Finite State Machine State Encoding Registers

Finding the signals to debug Finite State Machines (FSM) can be challenging. Finding nodes from the post-fit netlist may be impossible, as FSM encoding signals may be changed or optimized away during synthesis and place-and-route. If you can find all of the relevant nodes in the post-fit netlist or you used the nodes from the pre-synthesis netlist, an additional step is required to find and map FSM signal values to the state names that you specified in your HDL.

The SignalTap II Logic Analyzer GUI can detect FSMs in your compiled design. The SignalTap II Logic Analyzer configuration automatically tracks the FSM state signals as well as state encoding through the compilation process. Shortcut menu commands from the SignalTap II Logic Analyzer GUI allow you to add all of the FSM state signals to your logic analyzer with a single command. For each FSM added to your SignalTap II configuration, the FSM debugging feature adds a mnemonic table to map the signal values to the state enumeration that you provided in your source code. The mnemonic tables enable you to visualize state machine transitions in the waveform viewer. The FSM debugging feature supports adding FSM signals from both the pre-synthesis and post-fit netlists.

**Figure 13-4: Decoded FSM Mnemonics**

The waveform viewer with decoded signal values from a state machine added with the FSM debugging feature.



**Related Information**

- [Recommended HDL Coding Styles documentation](#)  
Coding guidelines for specifying FSM in Verilog and VHDL
- [Setting Up the SignalTap II Logic Analyzer online help](#)  
Information about adding FSM signals to the configuration file

## Modifying and Restoring Mnemonic Tables for State Machines

When you add FSM state signals via the FSM debugging feature, the SignalTap II Logic Analyzer GUI creates a mnemonic table using the format `<StateSignalName>_table`, where **StateSignalName** is the name of the state signals that you have declared in your RTL. You can edit any mnemonic table using the **Mnemonic Table Setup** dialog box.

If you want to restore a mnemonic table that was modified, right-click anywhere in the node list window and select **Recreate State Machine Mnemonics**. By default, restoring a mnemonic table overwrites the existing mnemonic table that you modified. To restore a FSM mnemonic table to a new record, turn off **Overwrite existing mnemonic table** in the **Recreate State Machine Mnemonics** dialog box.

**Note:** If you have added or deleted a signal from the FSM state signal group from within the setup tab, delete the modified register group and add the FSM signals back again.

### Related Information

[Creating Mnemonics for Bit Patterns](#) on page 13-60

## Additional Considerations

The SignalTap II configuration GUI recognizes state machines from your design only if you use Quartus II Integrated Synthesis (QIS). The state machine debugging feature is not able to track the FSM signals or state encoding if you use other EDA synthesis tools.

If you add post-fit FSM signals, the SignalTap II Logic Analyzer FSM debug feature may not track all optimization changes that are a part of the compilation process. If the following two specific optimizations are enabled, the SignalTap II FSM debug feature may not list mnemonic tables for state machines in the design:

- If you have physical synthesis turned on, state registers may be resource balanced (register retiming) to improve  $f_{MAX}$ . The FSM debug feature does not list post-fit FSM state registers if register retiming occurs.
- The FSM debugging feature does not list state signals that have been packed into RAM and DSP blocks during QIS or Fitter optimizations.

You can still use the FSM debugging feature to add pre-synthesis state signals.

## Specifying the Sample Depth

The sample depth specifies the number of samples that are captured and stored for each signal in the captured data buffer. To specify the sample depth, select the desired number of samples to store in the **Sample Depth** list. The sample depth ranges from 0 to 128K.

If device memory resources are limited, you may not be able to successfully compile your design with the sample buffer size you have selected. Try reducing the sample depth to reduce resource usage.

## Capturing Data to a Specific RAM Type

When you use the SignalTap II Logic Analyzer with some devices, you have the option to select the RAM type where acquisition data is stored. Once SignalTap II Logic Analyzer is allocated to a particular RAM block, the entire RAM block becomes a dedicated resource for the logic analyzer. RAM selection allows you to preserve a specific memory block for your design and allocate another portion of memory for SignalTap II Logic Analyzer data acquisition. For example, if your design has an application that requires a large block of memory resources, such a large instruction or data cache, you would choose to use MLAB, M512, or M4k blocks for data acquisition and leave the M9k blocks for the rest of your design.

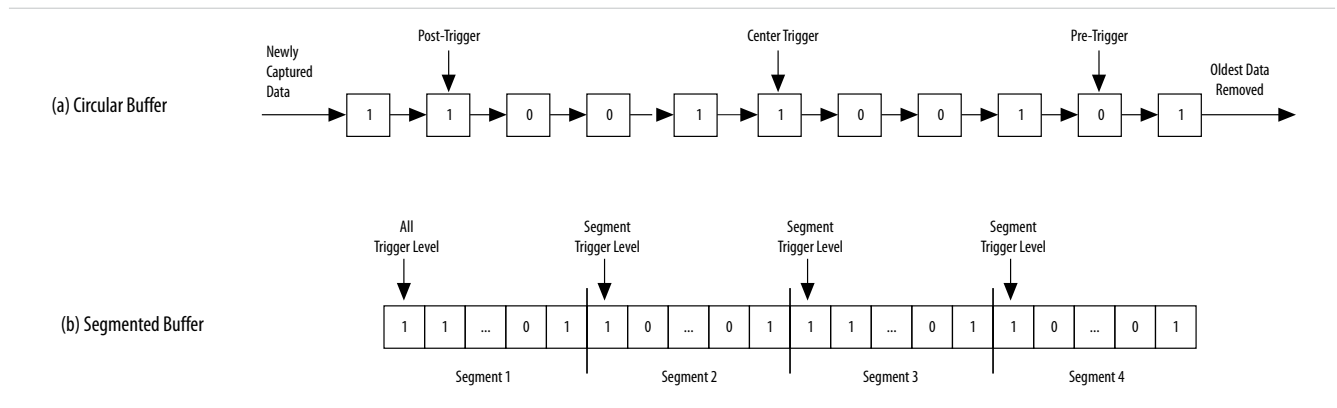
To select the RAM type to use for the SignalTap II Logic Analyzer buffer, select it from the RAM type list. Use this feature when the acquired data (as reported by the SignalTap II resource estimator) is not larger than the available memory of the memory type that you have selected in the FPGA.

## Choosing the Buffer Acquisition Mode

The Buffer Acquisition Type Selection feature in the SignalTap II Logic Analyzer lets you choose how the captured data buffer is organized and can potentially reduce the amount of memory that is required for SignalTap II data acquisition. There are two types of acquisition buffer within the SignalTap II Logic Analyzer—a non-segmented buffer and a segmented buffer. With a non-segmented buffer, the SignalTap II Logic Analyzer treats entire memory space as a single FIFO, continuously filling the buffer until the logic analyzer reaches a defined set of trigger conditions. With a segmented buffer, the memory space is split into a number of separate buffers. Each buffer acts as a separate FIFO with its own set of trigger conditions. Only a single buffer is active during an acquisition. The SignalTap II Logic Analyzer advances to the next segment after the trigger condition or conditions for the active segment has been reached.

When using a non-segmented buffer, you can use the storage qualification feature to determine which samples are written into the acquisition buffer. Both the segmented buffers and the non-segmented buffer with the storage qualification feature help you maximize the use of the available memory space. [Figure 13-5](#) illustrates the differences between the two buffer types.

**Figure 13-5: Buffer Type Comparison in the SignalTap II Logic Analyzer**



Notes to figure :

1. Both non-segmented and segmented buffers can use a predefined trigger (Pre-Trigger, Center Trigger, Post-Trigger) position or define a custom trigger position using the **State-Based Triggering** tab. Refer to [Specifying the Trigger Position](#) for more details.
2. Each segment is treated like a FIFO, and behaves as the non-segmented buffer shown in (a).

### Related Information

[Using the Storage Qualifier Feature](#) on page 13-15

## Non-Segmented Buffer

The non-segmented buffer (also known as a circular buffer) shown in [Figure 13-5](#) (a) is the default buffer type used by the SignalTap II Logic Analyzer. While the logic analyzer is running, data is stored in the buffer until it fills up, at which point new data replaces the oldest data. This continues until a specified

trigger event, consisting of a set of trigger conditions, occurs. When the trigger event happens, the logic analyzer continues to capture data after the trigger event until the buffer is full, based on the trigger position setting in the **Signal Configuration** pane or the **.stp**. To capture the majority of the data before the trigger occurs, select **Post trigger position** from the list. To capture the majority of the data after the trigger, select **Pre-trigger position**. To center the trigger position in the data, select **Center trigger position**. Alternatively, use the custom State-based triggering flow to define a custom trigger position within the capture buffer.

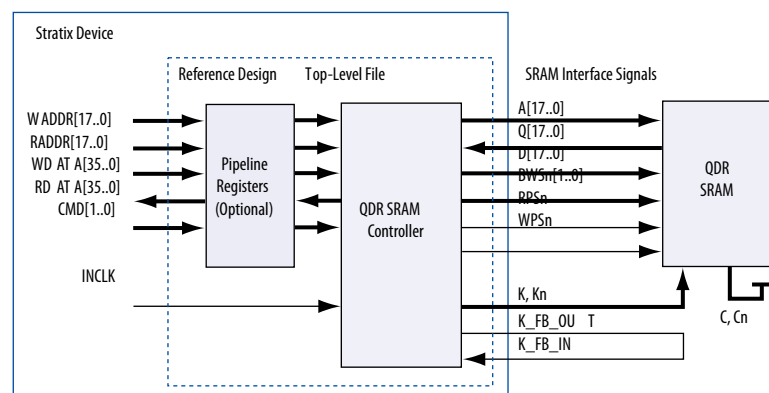
#### Related Information

[Specifying the Trigger Position](#) on page 13-43

## Segmented Buffer

A segmented buffer allows you to debug systems that contain relatively infrequent recurring events. The acquisition memory is split into evenly sized segments, with a set of trigger conditions defined for each segment. Each segment acts as a non-segmented buffer. If you want to have separate trigger conditions for each of the segmented buffers, you must use the state-based trigger flow. **Figure 13-6** shows an example of a segmented buffer system. If you want to have separate trigger conditions for each of the segmented buffers, you must use the state-based trigger flow.

**Figure 13-6: Example System that Generates Recurring Events**



The SignalTap II Logic Analyzer verifies the functionality of the design shown in **Figure 13-6** to ensure that the correct data is written to the SRAM controller. Buffer acquisition in the SignalTap II Logic Analyzer allows you to monitor the `RDATA` port when `H'0F0F0F0F` is sent into the `RADDR` port. You can monitor multiple read transactions from the SRAM device without running the SignalTap II Logic Analyzer again. The buffer acquisition feature allows you to segment the memory so you can capture the same event multiple times without wasting allocated memory. The number of cycles that are captured depends on the number of segments specified under the **Data** settings.

To enable and configure buffer acquisition, select **Segmented** in the SignalTap II Logic Analyzer Editor and select the number of segments to use. In the example in **Figure 13-6**, selecting sixty-four 64-sample segments allows you to capture 64 read cycles when the `RADDR` signal is `H'0F0F0F0F`.

#### Related Information

[Configuring the Trigger Flow in the SignalTap II Logic Analyzer online help](#)

Information about buffer acquisition mode

## Using the Storage Qualifier Feature

Both non-segmented and segmented buffers described in the previous section offer a snapshot in time of the data stream being analyzed. The default behavior for writing into acquisition memory with the SignalTap II Logic Analyzer is to sample data on every clock cycle. With a non-segmented buffer, there is one data window that represents a comprehensive snapshot of the datastream. Similarly, segmented buffers use several smaller sampling windows spread out over more time, with each sampling window representing a contiguous data set.

With carefully chosen trigger conditions and a generous sample depth for the acquisition buffer, analysis using segmented and non-segmented buffers captures a majority of functional errors in a chosen signal set. However, each data window can have a considerable amount of redundancy associated with it; for example, a capture of a data stream containing long periods of idle signals between data bursts. With default behavior using the SignalTap II Logic Analyzer, you cannot discard the redundant sample bits.

The Storage Qualification feature allows you to filter out individual samples not relevant to debugging the design. With this feature, a condition acts as a write enable to the buffer during each clock cycle of data acquisition. Through fine tuning the data that is actually stored in acquisition memory, the Storage Qualification feature allows for a more efficient use of acquisition memory in the specified number of samples over a longer period of analysis.

Use of the Storage Qualification feature is similar to an acquisition using a segmented buffer, in that you can create a discontinuity in the capture buffer. Because you can create a discontinuity between any two samples in the buffer, the Storage Qualification feature is equivalent to being able to create a customized segmented buffer in which the number and size of segment boundaries are adjustable.

**Note:** You can only use the Storage Qualification feature with a non-segmented buffer. The IP Catalog flow only supports the Input Port mode for the Storage Qualification feature.



**Figure 13-7: Data Acquisition Using Different Modes of Controlling the Acquisition Buffer****Non-segmented Buffer (1)****Segmented Buffer (2)****Non-segmented Buffer with Storage Qualifier (3)**

Notes to figure :

1. Non-segmented Buffers capture a fixed sample window of contiguous data.
2. Segmented buffers divide the buffer into fixed sized segments, with each segment having an equal sample depth.
3. Storage Qualification allows you to define a custom sampling window for each segment you create with a qualifying condition. Storage qualification potentially allows for a larger time scale of coverage.

There are six storage qualifier types available under the Storage Qualification feature:

- **Continuous**
- **Input port**
- **Transitional**
- **Conditional**
- **Start/Stop**
- **State-based**

Continuous (the default mode selected) turns the Storage Qualification feature off.

Each selected storage qualifier type is active when an acquisition starts. Upon the start of an acquisition, the SignalTap II Logic Analyzer examines each clock cycle and writes the data into the acquisition buffer

based upon storage qualifier type and condition. The acquisition stops when a defined set of trigger conditions occur.

**Note:** Trigger conditions are evaluated independently of storage qualifier conditions. The SignalTap II Logic Analyzer evaluates the data stream for trigger conditions on every clock cycle after the acquisition begins.

The storage qualifier operates independently of the trigger conditions.

**Related Information**

[Define Trigger Conditions](#) on page 13-6

**Input Port Mode**


When using the Input port mode, the SignalTap II Logic Analyzer takes any signal from your design as an input. When the design is running, if the signal is high on the clock edge, the SignalTap II Logic Analyzer stores the data in the buffer. If the signal is low on the clock edge, the data sample is ignored. A pin is created and connected to this input port by default if no internal node is specified.

If you are using an **.stp** to create a SignalTap II Logic Analyzer instance, specify the storage qualifier signal using the input port field located on the **Setup** tab. You must specify this port for your project to compile.


If you use the IP Catalog flow, the storage qualification input port, if specified, appears in the generated instantiation template. You can then connect this port to a signal in your RTL.

**Figure 13-8** shows a data pattern captured with a segmented buffer. **Figure 13-9** shows a capture of the same data pattern with the storage qualification feature enabled.

**Figure 13-8: Data Acquisition of a Recurring Data Pattern in Continuous Capture Mode (to illustrate Input port mode)**



**Figure 13-9: Data Acquisition of a Recurring Data Pattern Using an Input Signal as a Storage Qualifier**



(1) Markers display samples when the logic analyzer paused a write into acquisition memory. These markers are enabled with the option "Record data discontinuities."

**Transitional Mode**

In Transitional mode, you choose a set of signals for inspection using the node list check boxes in the **Storage Qualifier** column. During acquisition, if any of the signals marked for inspection have changed since the previous clock cycle, new data is written to the acquisition buffer. If none of the signals marked have changed since the previous clock cycle, no data is stored. **Figure 13-10** shows the transitional storage

qualifier setup. [Figure 13-11](#) and [Figure 13-12](#) show captures of a data pattern in continuous capture mode and a data pattern using the Transitional mode for storage qualification.

Figure 13-10: Transitional Storage Qualifier Setup

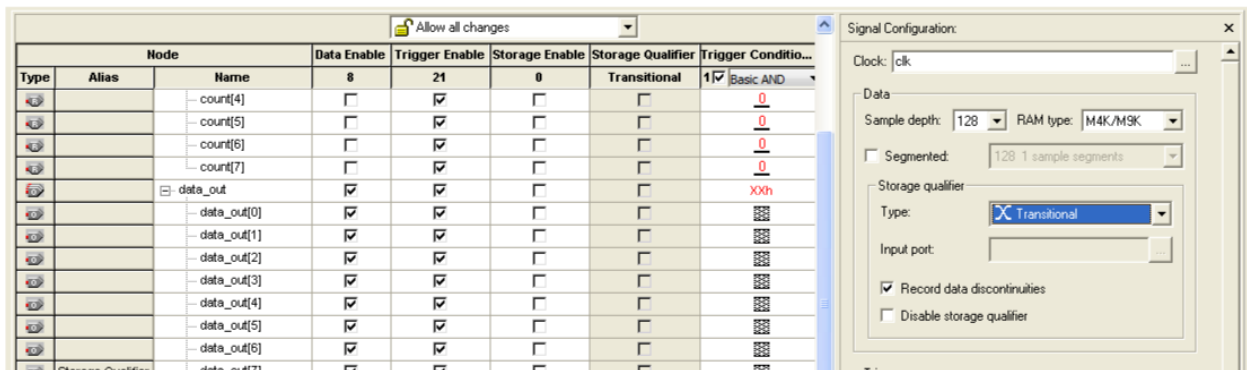


Figure 13-11: Data Acquisition of a Recurring Data Pattern in Continuous Capture Mode (to illustrate Transitional mode)





Figure 13-12: Data Acquisition of Recurring Data Pattern Using a Transitional Mode as a Storage Qualifier



## Conditional Mode

In Conditional mode, the SignalTap II Logic Analyzer evaluates a combinational function of storage qualifier enabled signals within the node list to determine whether a sample is stored. The SignalTap II Logic Analyzer writes into the buffer during the clock cycles in which the condition you specify evaluates TRUE.

You can select either **Basic AND**, **Basic OR**, or **Advanced** storage qualifier conditions. A **Basic AND** or **Basic OR** storage qualifier condition matches each signal to one of the following:

- **Don't Care**
- **Low**
- **High**
- **Falling Edge**
- **Rising Edge**
- **Either Edge**

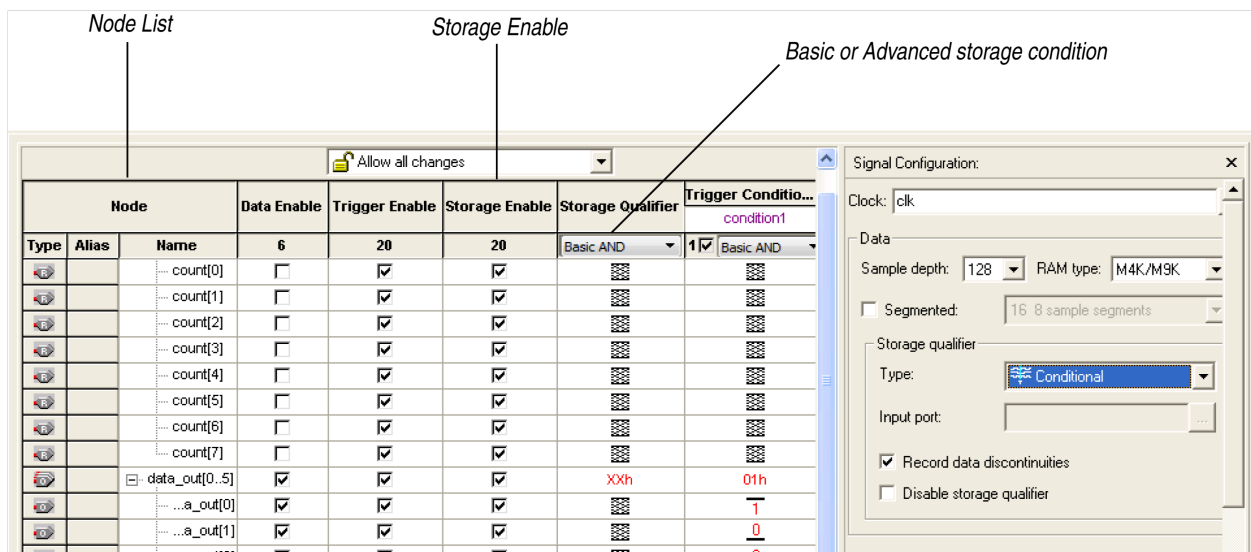
If you specify a **Basic AND** storage qualifier condition for more than one signal, the SignalTap II Logic Analyzer evaluates the logical AND of the conditions.

Any other combinational or relational operators that you may want to specify with the enabled signal set for storage qualification can be done with an advanced storage condition. **Figure 13-13** details the conditional storage qualifier setup in the **.stp**.


You can specify storage qualification conditions similar to the manner in which trigger conditions are specified.

**Figure 13-14** and **Figure 13-15** show a data capture with continuous sampling, and the same data pattern using the conditional mode for analysis, respectively.

**Figure 13-13: Conditional Storage Qualifier Setup**




**Figure 13-14: Data Acquisition of a Recurring Data Pattern in Continuous Capture Mode (to illustrate Conditional capture)**



(1) Storage Qualifier condition is set up to evaluate data\_out[6] AND data\_out[7].

**Figure 13-15: Data Acquisition of a Recurring Data Pattern in Conditional Capture Mode**



#### Related Information

- [Creating Basic Trigger Conditions](#) on page 13-24
- [Creating Advanced Trigger Conditions](#) on page 13-26

#### Start/Stop Mode

The Start/Stop mode is similar to the Conditional mode for storage qualification. However, in this mode there are two sets of conditions, one for start and one for stop. If the start condition evaluates to TRUE, data is stored in the buffer every clock cycle until the stop condition evaluates to TRUE, which then pauses the data capture. Additional start signals received after the data capture has started are ignored. If both start and stop evaluate to TRUE at the same time, a single cycle is captured.

**Note:** You can force a trigger by pressing the **Stop** button if the buffer fails to fill to completion due to a stop condition.

**Figure 13-16** shows the Start/Stop mode storage qualifier setup. **Figure 13-17** and **Figure 13-18** show captures data pattern in continuous capture mode and a data pattern in using the Start/Stop mode for storage qualification.

Figure 13-16: Start/Stop Mode Storage Qualifier Setup

The screenshot shows the Signal Configuration window for a logic analyzer. The main table lists nodes and their configuration:

Node	Alias	Name	Data Enable	Trigger Enable	Storage Enable	Storage Qualifier	Trig...
count		count[0]	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>		
		count[1]	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>		
		count[2]	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>		
		count[3]	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>		
		count[4]	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>		
		count[5]	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>		
		count[6]	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>		
		count[7]	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>		
data_out[0..5]		data_out[6]	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Start Basic AND	1
Storage Qualifier 1		data_out[7]	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Stop Basic AND	0
Storage Qualifier 2		start_sig_pulse	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>		

The right-hand pane shows the Signal Configuration for the selected node:

- Clock: clk
- Data: Sample depth: 128, RAM type: M4K/M9K
- Segmented: 128 1 sample segments
- Storage qualifier Type: Start/Stop
- Input port: (empty)
- Record data discontinuities:
- Disable storage qualifier:

Labels below the table point to specific cells:

- Storage Qualifier: points to the 'Storage Qualifier' column header.
- Storage Qualifier Start Condition: points to the 'Start' cell for data\_out[6].
- Storage Qualifier Stop Condition: points to the 'Stop' cell for data\_out[7].

Figure 13-17: Data Acquisition of a Recurring Data Pattern in Continuous Mode (to illustrate Start/Stop mode)





Figure 13-18: Data Acquisition of a Recurring Data Pattern with Start/Stop Storage Qualifier Enabled



### State-Based

The State-based storage qualification mode is used with the State-based triggering flow. The state based triggering flow evaluates an if-else based language to define how data is written into the buffer. With the State-based trigger flow, you have command over boolean and relational operators to guide the execution flow for the target acquisition buffer. When the storage qualifier feature is enabled for the State-based flow, two additional commands are available, the `start_store` and `stop_store` commands. These commands operate similarly to the Start/Stop capture conditions described in the previous section. Upon the start of acquisition, data is not written into the buffer until a `start_store` action is performed. The `stop_store` command pauses the acquisition. If both `start_store` and `stop_store` actions are performed within the same clock cycle, a single sample is stored into the acquisition buffer.

**Related Information**[State-Based Triggering](#) on page 13-33**Showing Data Discontinuities**

When you turn on **Record data discontinuities**, the SignalTap II Logic Analyzer marks the samples during which the acquisition paused from a storage qualifier. This marker is displayed in the waveform viewer after acquisition completes.

**Disable Storage Qualifier**


You can turn off the storage qualifier quickly with the **Disable Storage Qualifier** option, and perform a continuous capture. This option is run-time reconfigurable; that is, the setting can be changed without recompiling the project. Changing storage qualifier mode from the **Type** field requires a recompilation of the project.

**Related Information**[Runtime Reconfigurable Options](#) on page 13-54**Managing Multiple SignalTap II Files and Configurations**

You may have more than one **.stp** in one design. Each file potentially has a different group of monitored signals. These signal groups make it possible to debug different blocks in your design. In turn, each group of signals can also be used to define different sets of trigger conditions. Along with each **.stp**, there is also an associated programming file (SRAM Object File [**.sof**]). The settings in a selected SignalTap II file must match the SignalTap II logic design in the associated **.sof** for the logic analyzer to run properly when the device is programmed. Use the Data Log feature and the SOF Manager to manage all of the **.stp** files and their associated settings and programming files.

The Data Log allows you to store multiple SignalTap II configurations within a single **.stp**. **Figure 13-19** shows two signal set configurations with multiple trigger conditions in one **.stp**. To toggle between the active configurations, double-click on an entry in the Data Log. As you toggle between the different configurations, the signal list and trigger conditions change in the **Setup** tab of the **.stp**. The active configuration displayed in the **.stp** is indicated by the blue square around the signal specified in the Data Log. Enable the Data Log by clicking the check box next to **Data Log**. To store a configuration in the Data Log, on the Edit menu, click **Save to Data Log** or click the **Save to Data Log** icon at the top of the Data Log. The time stamping for the Data Log entries display the wall-clock time when SignalTap II triggered and the elapsed time from when acquisition started to when the device triggered.

Figure 13-19: Data Log



The SOF Manager allows you to embed multiple SOFs into one **.stp**. Embedding an SOF in an **.stp** lets you move the **.stp** to a different location, either on the same computer or across a network, without the need to include the associated **.sof** separately. To embed a new SOF in the **.stp**, right-click in the SOF Manager, and click **Attach SOF File**.

Figure 13-20: SOF Manager



As you switch between configurations in the Data Log, you can extract the SOF that is compatible with that particular configuration. You can use the programmer in the SignalTap II Logic Analyzer to download the new SOF to the FPGA, ensuring that the configuration of your **.stp** always matches the design programmed into the target device.

## Define Triggers

When you start the SignalTap II Logic Analyzer, it samples activity continuously from the monitored signals. The SignalTap II Logic Analyzer “triggers”—that is, the logic analyzer stops and displays the data—when a condition or set of conditions that you specified has been reached. This section describes the various types of trigger conditions that you can specify using the SignalTap II Logic Analyzer on the **Signal Configuration** pane.



## Creating Basic Trigger Conditions

The simplest kind of trigger condition is a basic trigger. Select this from the list at the top of the **Trigger Conditions** column in the node list in the SignalTap II Logic Analyzer Editor. If you select the **Basic AND** or **Basic OR** trigger type, you must specify the trigger pattern for each signal you have added in the **.stp**. To specify the trigger pattern, right-click in the **Trigger Conditions** column and click the desired pattern. Set the trigger pattern to any of the following conditions:

- **Don't Care**
- **Low**
- **High**
- **Falling Edge**
- **Rising Edge**
- **Either Edge**

For buses, type a pattern in binary, or right-click and select **Insert Value** to enter the pattern in other number formats. Note that you can enter x to specify a set of “don't care” values in either your hexadecimal or your binary string. For signals added to the **.stp** that have an associated mnemonic table, you can right-click and select an entry from the table to specify pre-defined conditions for the trigger.

For more information about creating and using mnemonic tables, refer to [View, Analyze, and Use Captured Data](#), and to the Quartus II Help.

For signals added with certain plug-ins, you can create basic triggers easily using predefined mnemonic table entries. For example, with the Nios II plug-in, if you have specified an **.elf** from your Nios II IDE design, you can type the name of a function from your Nios II code. The logic analyzer triggers when the Nios II instruction address matches the address of the specified code function name.

Data capture stops and the data is stored in the buffer when the logical AND of all the signals for a given trigger condition evaluates to **TRUE**.

### Using the Basic OR Triggering Condition with Nested Groups

When you specify a set of signals as a nested group (group of groups) with the **Basic OR** trigger type, an advanced trigger condition is generated. This advanced trigger condition sorts signals within groups to minimize the need to recompile your design. As long as the parent-child relationships of nodes are kept constant, the generated advanced trigger condition does not change. You can modify the sibling relationships of nodes and not need to recompile your design. The precedence of how this trigger condition is evaluated starts at the bottom-level with the leaf-groups first, then their resulting logic-1 or logic-0 value is used to compute the result of their parent group's logic value. Specifying a value of **TRUE** for a group sets that group's logical result to logic-1 and effectively eliminates all members beneath it from affecting the result of the group trigger. Specifying a value of **FALSE** for a group sets that group's logical result to logic-0 and effectively eliminates all members beneath it from affecting the result of the group trigger.

1. Select **Basic OR** under **Trigger Conditions**.
2. In the **Setup** tab, select nodes including groups.
3. Right-click in the **Setup** tab and select **Group**.
4. Select your signal(s) and right-click to set a group trigger condition that applies the reduction **AND**, **OR**, **NAND**, **NOR**, **XOR**, **XNOR**, or logical **TRUE** or **FALSE**.

**Note:** The OR and AND group trigger conditions are only selectable for groups with no groups as children (bottom-level groups).

Figure 13-21: Creating Nested Groups

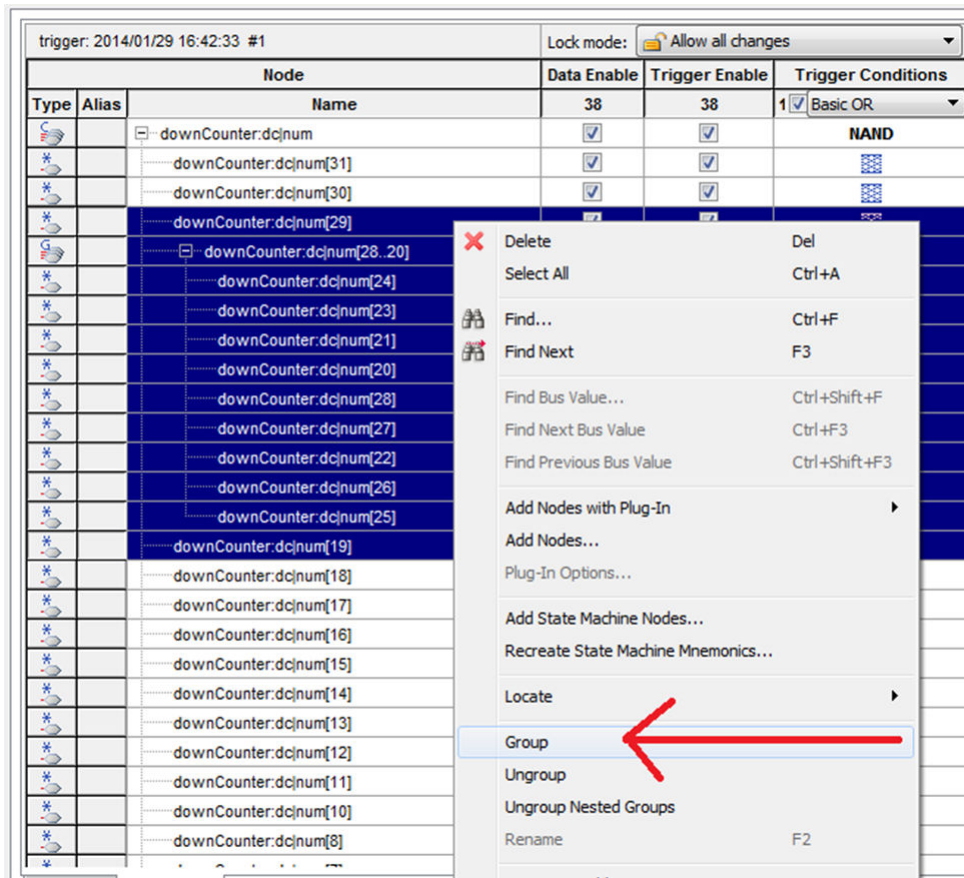
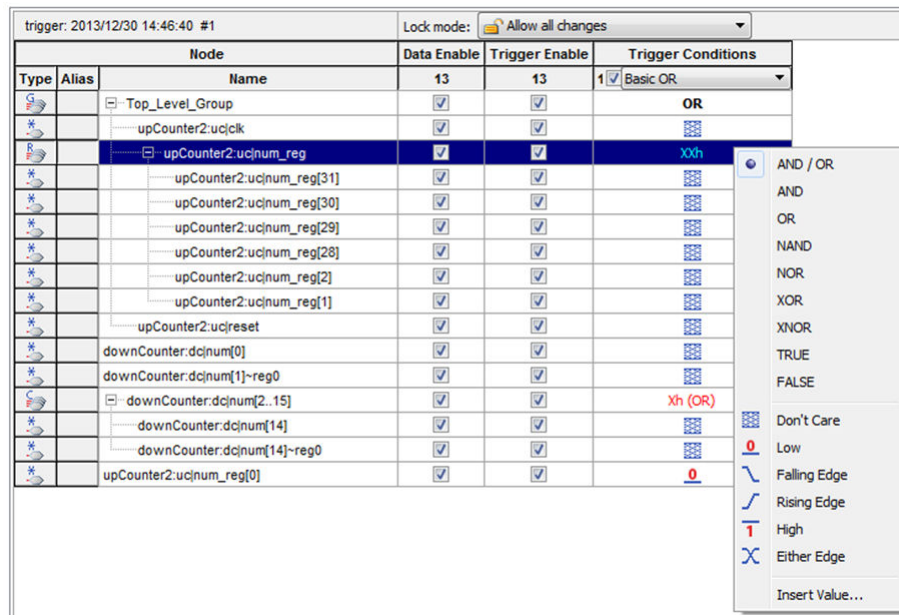


Figure 13-22: Applying Group Trigger Condition



## Creating Advanced Trigger Conditions

With the basic triggering capabilities of the SignalTap II Logic Analyzer, you can build more complex triggers with extra logic that enables you to capture data when a combination of conditions exist. If you select the **Advanced** trigger type at the top of the **Trigger Conditions** column in the node list of the SignalTap II Logic Analyzer Editor, a new tab named **Advanced Trigger** appears where you can build a complex trigger expression using a simple GUI. Drag-and-drop operators into the Advanced Trigger Configuration Editor window to build the complex trigger condition in an expression tree. To configure the operators' settings, double-click or right-click the operators that you have placed and select **Properties**.

Table 13-2: Advanced Triggering Operators

Name of Operator	Type
Less Than	Comparison
Less Than or Equal To	Comparison
Equality	Comparison
Inequality	Comparison
Greater Than	Comparison
Greater Than or Equal To	Comparison
Logical NOT	Logical
Logical AND	Logical
Logical OR	Logical

Name of Operator	Type
Logical XOR	Logical
Reduction AND	Reduction
Reduction OR	Reduction
Reduction XOR	Reduction
Left Shift	Shift
Right Shift	Shift
Bitwise Complement	Bitwise
Bitwise AND	Bitwise
Bitwise OR	Bitwise
Bitwise XOR	Bitwise
Edge and Level Detector	Signal Detection

Note to table :

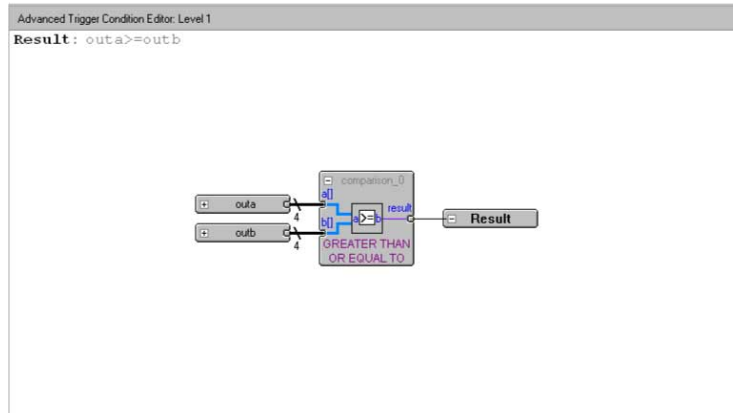
1. For more information about each of these operators, refer to the Quartus II Help.

Adding many objects to the Advanced Trigger Condition Editor can make the work space cluttered and difficult to read. To keep objects organized while you build your advanced trigger condition, use the shortcut menu and select **Arrange All Objects**. You can also use the **Zoom-Out** command to fit more objects into the Advanced Trigger Condition Editor window.

## Examples of Advanced Triggering Expressions

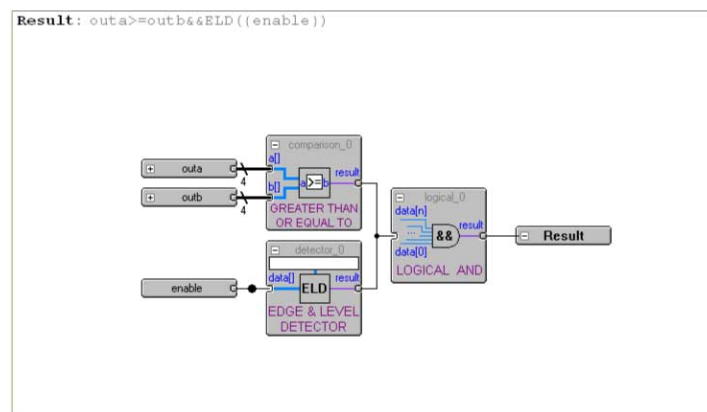
The following examples show how to use Advanced Triggering:

- Trigger when bus `outa` is greater than or equal to `outb`.

Figure 13-23: Bus `outa` Is Greater Than or Equal to Bus `outb`

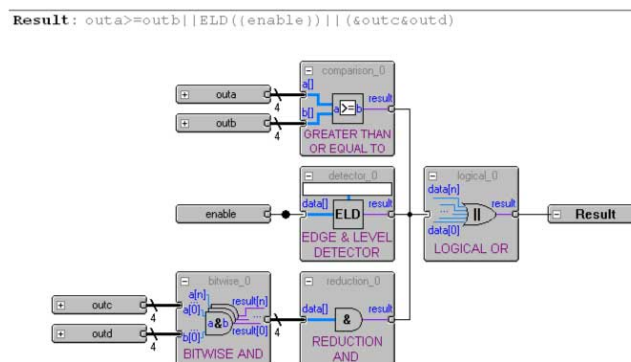
- Trigger when bus `outa` is greater than or equal to bus `outb`, and when the enable signal has a rising edge (Figure 13-24).

Figure 13-24: Enable Signal Has a Rising Edge



- Trigger when bus `outa` is greater than or equal to bus `outb`, or when the enable signal has a rising edge. Or, when a bitwise AND operation has been performed between bus `outc` and bus `outd`, and all bits of the result of that operation are equal to 1 (Figure 13-25).

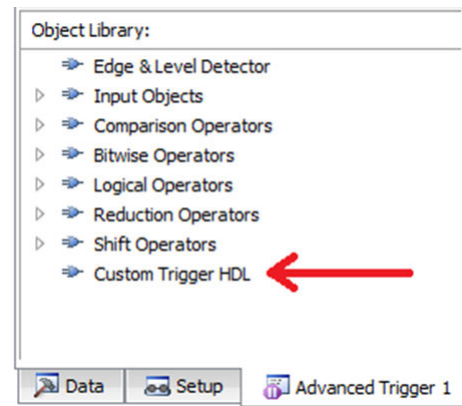
Figure 13-25: Bitwise AND Operation



## Custom Trigger HDL Object

The Custom Trigger HDL object found in the **Advanced Trigger** editor allows you to create a customized trigger condition with your own HDL module in either Verilog or VHDL. You can use this object to simulate the behavior of your triggering logic to make sure that the logic itself is not faulty. You can tap specific instances of modules located anywhere in the hierarchy of your design without having to manually route all the necessary connections.


Figure 13-26: Object Library



## Custom Trigger Flow


1. Select Advanced for a given trigger-level to make the Advanced Trigger editor active.
2. Prepare your Custom Trigger HDL module. You can either add a new source file to Quartus II that contains the trigger module or append the HDL for your trigger module to a source file already included in Quartus II **Files** under the **Project Navigator**.

Figure 13-27: Project Navigator



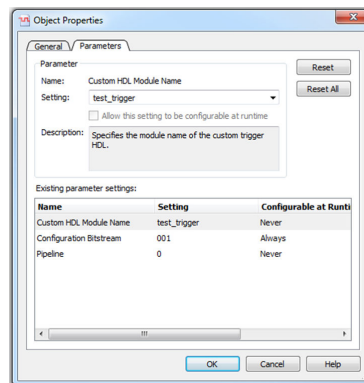
3. Implement the required inputs and outputs for your Custom Trigger HDL module, see [Table 13-3](#).
4. Drag in your Custom Trigger HDL object and connect the object's data input bus and result output bit to the final trigger result.

Figure 13-28: Custom Trigger HDL Object



- Right-click your Custom Trigger HDL object and configure the object's properties, see [Table 13-4](#).

Figure 13-29: Configure Object Properties



- Compile your design.
- Acquire data with SignalTap II using your custom Trigger HDL object.

Table 13-3: Custom Trigger HDL Module Required Inputs and Outputs

Name	Description	Input/Output	Required/Optional
acq_clk	Acquisition clock used by SignalTap II	Input	Required
reset	Reset signal used by SignalTap II when restarting a capture.	Input	Required
data_in	<ul style="list-style-type: none"> <li>Data input to be connected in the Advanced Trigger editor.</li> <li>Data your module will use to trigger.</li> </ul>	Input	Required
pattern_in	<ul style="list-style-type: none"> <li>Module's input for the configuration bitstream property.</li> <li>Runtime configurable property that can be set from SignalTap II GUI to change the behavior of your trigger logic.</li> </ul>	Input	Optional
trigger_out	Output signal of your module to be asserted when triggering conditions have been met.	Output	Required

**Table 13-4: Custom Trigger HDL Module Properties**

Property	Description
Custom HDL Module Name	Module name of your triggering logic i.e. module <b>trigger_foo</b> (input x, y ...);
Configuration Bitstream	<ul style="list-style-type: none"> <li>Allows you to create runtime-configurable trigger logic which can change its behavior based upon the value of the configuration bitstream.</li> <li>Configuration bitstream property is interpreted as binary and should only contain the characters 1 and 0. The bit-width (number of 1s and 0s) should match the <code>pattern_in</code> bit width in <a href="#">Table 13-3</a>.</li> <li>A blank configuration bitstream implies that your module does not have a <code>pattern_in</code> input.</li> </ul>
Pipeline	<ul style="list-style-type: none"> <li>Tells the advanced trigger editor how many stages of pipeline your triggering logic has.</li> <li>If it takes three clock cycles after a triggering input is received for the trigger output to be asserted, you can denote a pipeline value of three.</li> </ul>

**Figure 13-30: Example of Verilog Trigger Using Configuration Bitstream**

```

module test_trigger(input acq_clk, reset, input [3:0] data_in, input
[1:0] pattern_in, output reg trigger_out);

    always @ (pattern_in) begin
        case (pattern_in)
            2'b00:
                trigger_out = &data_in;
            2'b01:
                trigger_out = |data_in;
            2'b10:
                trigger_out = 1'b0;
            2'b11:
                trigger_out = 1'b1;
        endcase
    end

endmodule

```



Figure 13-31: Example of Verilog Trigger with No Configuration Bitstream

```
module test_trigger_no_bs(input acq_clk, reset, input [3:0]
data_in, output trigger_out);

    assign trigger_out = &data_in;

endmodule
```

## Trigger Condition Flow Control

The SignalTap II Logic Analyzer offers multiple triggering conditions to give you precise control of the method in which data is captured into the acquisition buffers. Trigger Condition Flow allows you to define the relationship between a set of triggering conditions. The SignalTap II Logic Analyzer **Signal Configuration** pane offers two flow control mechanisms for organizing trigger conditions:

- **Sequential Triggering**—The default triggering flow. Sequential triggering allows you to define up to 10 triggering levels that must be satisfied before the acquisition buffer finishes capturing.
- **State-Based Triggering**—Allows you the greatest control over your acquisition buffer. Custom-based triggering allows you to organize trigger conditions into states based on a conditional flow that you define.

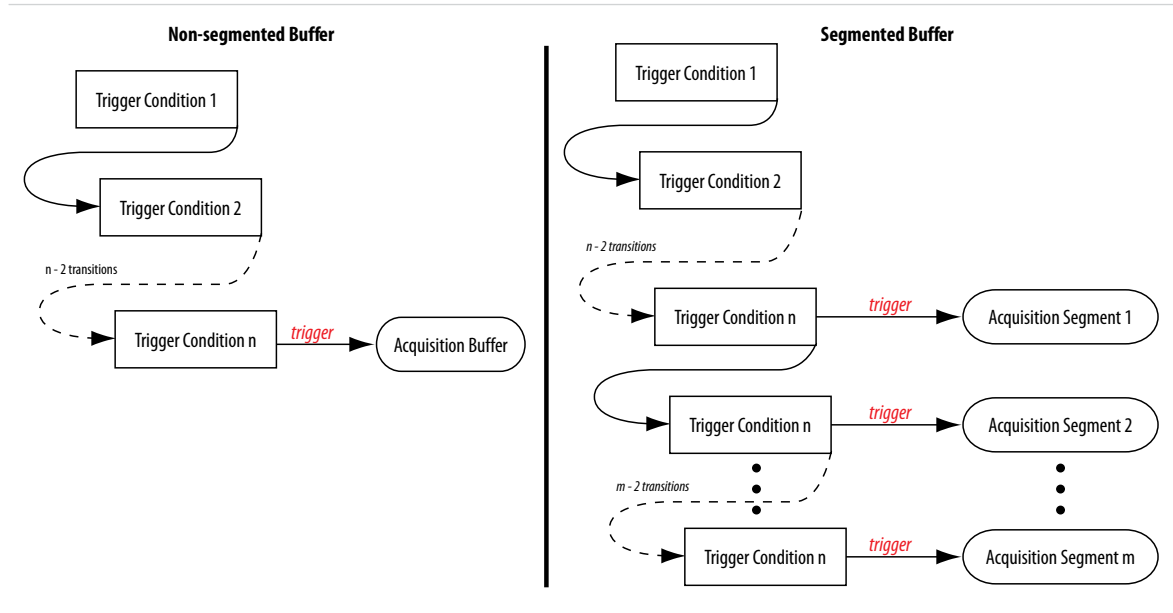
You can use sequential or state based triggering with either a segmented or a non-segmented buffer.

### Sequential Triggering

Sequential triggering flow allows you to cascade up to 10 levels of triggering conditions. The SignalTap II Logic Analyzer sequentially evaluates each of the triggering conditions. When the last triggering condition evaluates to `TRUE`, the SignalTap II Logic Analyzer triggers the acquisition buffer. For segmented buffers, every acquisition segment after the first segment triggers on the last triggering condition that you have specified. Use the Simple Sequential Triggering feature with basic triggers, advanced triggers, or a mix of both. [Figure 13-32](#) illustrates the simple sequential triggering flow for non-segmented and segmented buffers.

**Note:** The external trigger is considered as trigger level 0. The external trigger must be evaluated before the main trigger levels are evaluated.

Figure 13-32: Sequential Triggering Flow



Notes to figure :

1. The acquisition buffer stops capture when all  $n$  triggering levels are satisfied, where  $n \leq 10$ .
2. An external trigger input, if defined, is evaluated before all other defined trigger conditions are evaluated.

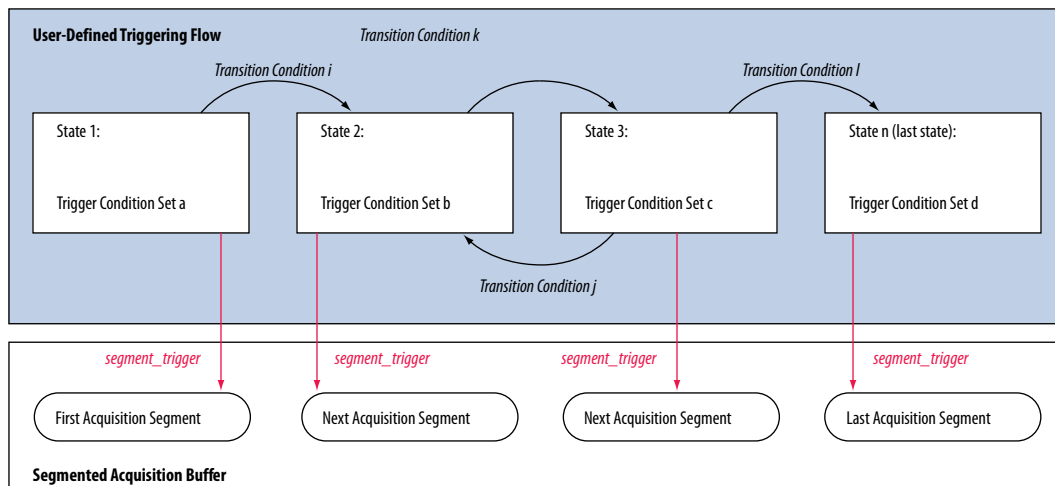
To configure the SignalTap II Logic Analyzer for Sequential triggering, in the SignalTap II editor on the **Trigger flow control** list, select **Sequential**. Select the desired number of trigger conditions from the **Trigger Conditions** list. After you select the desired number of trigger conditions, configure each trigger condition in the node list. To disable any trigger condition, turn on the trigger condition at the top of the column in the node list.

## State-Based Triggering

Custom State-based triggering provides the most control over triggering condition arrangement. The State-Based Triggering flow allows you to describe the relationship between triggering conditions precisely, using an intuitive GUI and the SignalTap II Trigger Flow Description Language, a simple description language based upon conditional expressions. Tooltips within the custom triggering flow GUI allow you to describe your desired flow quickly. The custom State-based triggering flow allows for more efficient use of the space available in the acquisition buffer because only specific samples of interest are captured.

Events that trigger the acquisition buffer are organized by a state diagram that you define. All actions performed by the acquisition buffer are captured by the states and all transition conditions between the states are defined by the conditional expressions that you specify within each state.

Figure 13-33: State-Based Triggering Flow



Notes to figure:

1. You are allowed up to 20 different states.
2. An external trigger input, if defined, is evaluated before any conditions in the custom State-based triggering flow are evaluated.

Each state allows you to define a set of conditional expressions. Each conditional expression is a Boolean expression dependent on a combination of triggering conditions (configured within the **Setup** tab), counters, and status flags. Counters and status flags are resources provided by the SignalTap II Logic Analyzer custom-based triggering flow.

Within each conditional expression you define a set of actions. Actions include triggering the acquisition buffer to stop capture, a modification to either a counter or status flag, or a state transition.

Trigger actions can apply to either a single segment of a segmented acquisition buffer or to the entire non-segmented acquisition buffer. Each trigger action provides you with an optional count that specifies the number of samples captured before stopping acquisition of the current segment. The count argument allows you to control the amount of data captured precisely before and after triggering event.

Resource manipulation actions allow you to increment and decrement counters or set and clear status flags. The counter and status flag resources are used as optional inputs in conditional expressions. Counters and status flags are useful for counting the number of occurrences of particular events and for aiding in triggering flow control.

This SignalTap II custom State-based triggering flow allows you to capture a sequence of events that may not necessarily be contiguous in time; for example, capturing a communication transaction between two devices that includes a handshaking protocol containing a sequence of acknowledgements.

The **State-Based Trigger Flow** tab is the control interface for the custom state-based triggering flow. To enable this tab, select **State-based** on the **Trigger Flow Control** list. (Note that when **Trigger Flow Control** is specified as **Sequential**, the **State-Based Trigger Flow** tab is hidden.)

The **State-Based Trigger Flow** tab is partitioned into the following three panes:

- **State Diagram** pane
- **Resources** pane
- **State Machine** pane

### State Diagram Pane

The **State Diagram** pane provides a graphical overview of the triggering flow that you define. It shows the number of states available and the state transitions between the states. You can adjust the number of available states by using the menu above the graphical overview.

### State Machine Pane

The **State Machine** pane contains the text entry boxes where you can define the triggering flow and actions associated with each state. You can define the triggering flow using the SignalTap II Trigger Flow Description Language, a simple language based on “if-else” conditional statements. Tooltips appear when you move the mouse over the cursor, to guide command entry into the state boxes. The GUI provides a syntax check on your flow description in real-time and highlights any errors in the text flow.

The State Machine description text boxes default to show one text box per state. You can also have the entire flow description shown in a single text field. This option can be useful when copying and pasting a flow description from a template or an external text editor. To toggle between one window per state, or all states in one window, select the appropriate option under **State Display mode**.

#### Related Information

- [SignalTap II Trigger Flow Description Language](#) on page 13-36
- [SignalTap II Trigger Flow Description Language online help](#)

### Resources Pane

The **Resources** pane allows you to declare Status Flags and Counters for use in the conditional expressions in the Custom Triggering Flow. Actions to decrement and increment counters or to set and clear status flags are performed within the triggering flow that you define.

You can specify up to 20 counters and 20 status flags. Counter and status flags values may be initialized by right-clicking the status flag or counter name after selecting a number of them from the respective pull-down list, and selecting **Set Initial Value**. To specify a counter width, right-click the counter name and select **Set Width**. Counters and flag values are updated dynamically after acquisition has started to assist in debugging your trigger flow specification.

The **configurable at runtime** options in the **Resources** pane allows you to configure the custom-flow control options that can be changed at runtime without requiring a recompilation.

**Table 13-5: Runtime Reconfigurable Settings, State-Based Triggering Flow**

Setting	Description
Destination of goto action	Allows you to modify the destination of the state transition at runtime.
Comparison values	Allows you to modify comparison values in Boolean expressions at runtime. In addition, you can modify the <code>segment_trigger</code> and trigger action post-fill count argument at runtime.
Comparison operators	Allows you to modify the operators in Boolean expressions at runtime.

Setting	Description
Logical operators	Allows you to modify the logical operators in Boolean expressions at runtime.

You can restrict changes to your SignalTap II configuration to include only the options that do not require a recompilation. Trigger lock-mode allows you to make changes that can be immediately reflected in the device.

1. On the **Setup** tab, select **Allow trigger condition changes only**.
2. Modify the Trigger Flow conditions in the **Custom Trigger Flow** tab.
3. Click the desired parameter in the text box and select a new parameter from the menu that appears.

**Note:** Trigger lock mode restricts changes to the configuration settings that have **configurable at runtime** specified. The runtime configurable settings for the **Custom Trigger Flow** tab are on by default. You may get some performance advantages by disabling some of the runtime configurable options.

Incremental Route lock-mode restricts the GUI to only allow changes that require an Incremental Route compilation using Rapid Recompile. Use Rapid Recompile to perform incremental routing and gain a 2-4x speedup over the initial full compilation. Refer to [Incremental Route with Rapid Recompile](#) on page 13-48.

#### Related Information

- [Performance and Resource Considerations](#) on page 13-51
- [Runtime Reconfigurable Options](#) on page 13-54

## SignalTap II Trigger Flow Description Language

The Trigger Flow Description Language is based on a list of conditional expressions per state to define a set of actions. Each line in the example shows a language format. Keywords are shown in bold. Non-terminals are delimited by “<>” and are further explained in the following sections. Optional arguments are delimited by “[ ]”.

```
state <State_label>:
<action_list>

if( <Boolean_expression> )
<action_list>
[else if ( <boolean_expression> )
<action_list>]
[else
<action_list>]
```

Note to example :

1. Multiple `else if` conditions are allowed.

The priority for evaluation of conditional statements is assigned from top to bottom. The `<boolean_expression>` in an `if` statement can contain a single event, or it can contain multiple event conditions. The `action_list` within an `if` or an `else if` clause must be delimited by the begin and end tokens when the action list contains multiple statements. When the boolean expression is evaluated `TRUE`,

the logic analyzer analyzes all of the commands in the action list concurrently. The possible actions include:

- Triggering the acquisition buffer
- Manipulating a counter or status flag resource
- Defining a state transition

**Related Information**

[Custom Triggering Flow Application Examples](#) on page 13-67

**State Labels**

State labels are identifiers that can be used in the action `goto`.

`state <state_label>`: begins the description of the actions evaluated when this state is reached.

The description of a state ends with the beginning of another state or the end of the whole trigger flow description.

**Boolean\_expression**

`Boolean_expression` is a collection of logical operators, relational operators, and their operands that evaluate into a Boolean result. Depending on the operator, the operand can be a reference to a trigger condition, a counter and a register, or a numeric value. Within an expression, parentheses can be used to group a set of operands.

Logical operators accept any boolean expression as an operand.

**Table 13-6: Logical Operators**

Operator	Description	Syntax
!	NOT operator	! expr1
&&	AND operator	expr1 && expr2
	OR operator	expr1    expr2

Relational operators are performed on counters or status flags. The comparison value, the right operator, must be a numerical value.

**Table 13-7: Relational Operators**

Operator	Description	Syntax
>	Greater than	<identifier> > <numerical_value>
>=	Greater than or Equal to	<identifier> >= <numerical_value>
==	Equals	<identifier> == <numerical_value>
!=	Does not equal	<identifier> != <numerical_value>

Operator	Description	Syntax
<=	Less than or equal to	<code>&lt;identifier&gt; &lt;= &lt;numerical_value&gt;</code>
<	Less than	<code>&lt;identifier&gt; &lt; &lt;numerical_value&gt;</code>

Notes to table :

1. `<identifier>` indicates a counter or status flag.
2. `<numerical_value>` indicates an integer.

## Action\_list

**Action\_list** is a list of actions that can be performed when a state is reached and a condition is also satisfied. If more than one action is specified, they must be enclosed by `begin` and `end`. The actions can be categorized as resource manipulation actions, buffer control actions, and state transition actions. Each action is terminated by a semicolon (;).

## Resource Manipulation Action

The resources used in the trigger flow description can be either counters or status flags.

**Table 13-8: Resource Manipulation Action**

Action	Description	Syntax
increment	Increments a counter resource by 1	<code>increment &lt;counter_identifier&gt;;</code>
decrement	Decrements a counter resource by 1	<code>decrement &lt;counter_identifier&gt;;</code>
reset	Resets counter resource to initial value	<code>reset &lt;counter_identifier&gt;;</code>
set	Sets a status Flag to 1	<code>set &lt;register_flag_identifier&gt;;</code>
clear	Sets a status Flag to 0	<code>clear &lt;register_flag_identifier&gt;;</code>

## Buffer Control Action

Buffer control actions specify an action to control the acquisition buffer.

**Table 13-9: Buffer Control Action**

Action	Description	Syntax
trigger	Stops the acquisition for the current buffer and ends analysis. This command is required in every flow definition.	<code>trigger &lt;post-fill_count&gt;;</code>

Action	Description	Syntax
segment_trigger	<p>Ends the acquisition of the current segment. The SignalTap II Logic Analyzer starts acquiring from the next segment on evaluating this command. If all segments are filled, the oldest segment is overwritten with the latest sample. The acquisition stops when a trigger action is evaluated.</p> <p>This action cannot be used in non-segmented acquisition mode.</p>	segment_trigger <post-fill_count>;
start_store	<p>Asserts the write_enable to the SignalTap II acquisition buffer. This command is active only when the State-based storage qualifier mode is enabled.</p>	start_store
stop_store	<p>De-asserts the write_enable signal to the SignalTap II acquisition buffer. This command is active only when the State-based storage qualifier mode is enabled.</p>	stop_store

Both trigger and segment\_trigger actions accept an optional post-fill count argument. If provided, the current acquisition acquires the number of samples provided by post-fill count and then stops acquisition. If no post-count value is specified, the trigger position for the affected buffer defaults to the trigger position specified in the **Setup** tab.

**Note:** In the case of segment\_trigger, acquisition of the current buffer stops immediately if a subsequent triggering action is issued in the next state, regardless of whether or not the post-fill count has been satisfied for the current buffer. The remaining unfilled post-count acquisitions in the current buffer are discarded and displayed as grayed-out samples in the data window.

## State Transition Action

The State Transition action specifies the next state in the custom state control flow. It is specified by the goto command. The syntax is as follows:

```
goto <state_label>;
```

## Using the State-Based Storage Qualifier Feature

When you select State-based for the storage qualifier type, the start\_store and stop\_store actions are enabled in the State-based trigger flow. These commands, when used in conjunction with the expressions of the State-based trigger flow, give you maximum flexibility to control data written into the acquisition buffer.

**Note:** The start\_store and stop\_store commands can only be applied to a non-segmented buffer.

The start\_store and stop\_store commands function similar to the start and stop conditions when using the **start/stop** storage qualifier mode conditions. If storage qualification is enabled, the start\_store command must be issued for SignalTap II to write data into the acquisition buffer. No data is acquired until the start\_store command is performed. Also, a trigger command must be included as part of the trigger flow description. The trigger command is necessary to complete the acquisition and display the results on the waveform display.



The following example illustrates the behavior of the State-based trigger flow with the storage qualification commands.

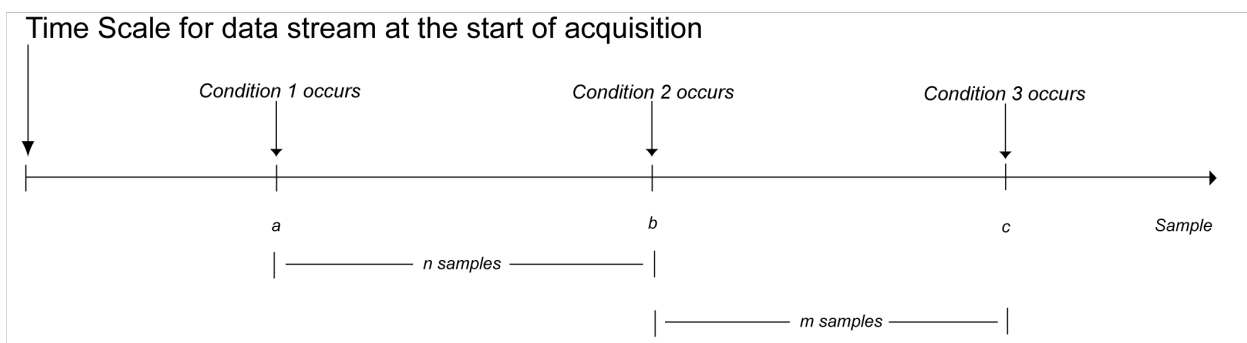
```

State 1: ST1:
if ( condition1 )
  start_store;
else if ( condition2 )
  trigger value;
else if ( condition3 )
  stop_store;

```

**Figure 13-34** shows a hypothetical scenario with three trigger conditions that happen at different times after you click **Start Analysis**. The trigger flow description in the example above, when applied to the scenario shown in **Figure 13-34**, illustrates the functionality of the storage qualification feature for the state-based trigger flow.

**Figure 13-34: Capture Scenario for Storage Qualification with the State-Based Trigger Flow**



In this example, the SignalTap II Logic Analyzer does not write into the acquisition buffer until sample a, when Condition 1 occurs. Once sample b is reached, the `trigger value` command is evaluated. The logic analyzer continues to write into the buffer to finish the acquisition. The trigger flow specifies a `stop_store` command at sample c, m samples after the trigger point occurs.

The logic analyzer finishes the acquisition and displays the contents of the waveform if it can successfully finish the post-fill acquisition samples before Condition 3 occurs. In this specific case, the capture ends if the post-fill count value is less than *m*.

If the post-fill count value specified in Trigger Flow description 1 is greater than *m* samples, the buffer pauses acquisition indefinitely, provided there is no recurrence of Condition 1 to trigger the logic analyzer to start capturing data again. The SignalTap II Logic Analyzer continues to evaluate the `stop_store` and `start_store` commands even after the `trigger` command is evaluated. If the acquisition has paused, you can click **Stop Analysis** to manually stop and force the acquisition to trigger. You can use counter values, flags, and the State diagram to help you perform the trigger flow. The counter values, flags, and the current state are updated in real-time during a data acquisition.

**Figure 13-35** and **Figure 13-36** show a real data acquisition of the scenario. **Figure 13-35** illustrates a scenario where the data capture finishes successfully. It uses a buffer with a sample depth of 64,  $m = n = 10$ , and the post-fill count value = 5. **Figure 13-36** illustrates a scenario where the logic analyzer pauses indefinitely even after a trigger condition occurs due to a `stop_store` condition. This scenario uses a sample depth of 64, with  $m = n = 10$  and post-fill count = 15.

Figure 13-35: Storage Qualification with Post-Fill Count Value Less than  $m$  (Acquisition Successfully Completes)




Figure 13-36: Storage Qualification with Post-Fill Count Value Greater than  $m$  (Acquisition Indefinitely Paused)

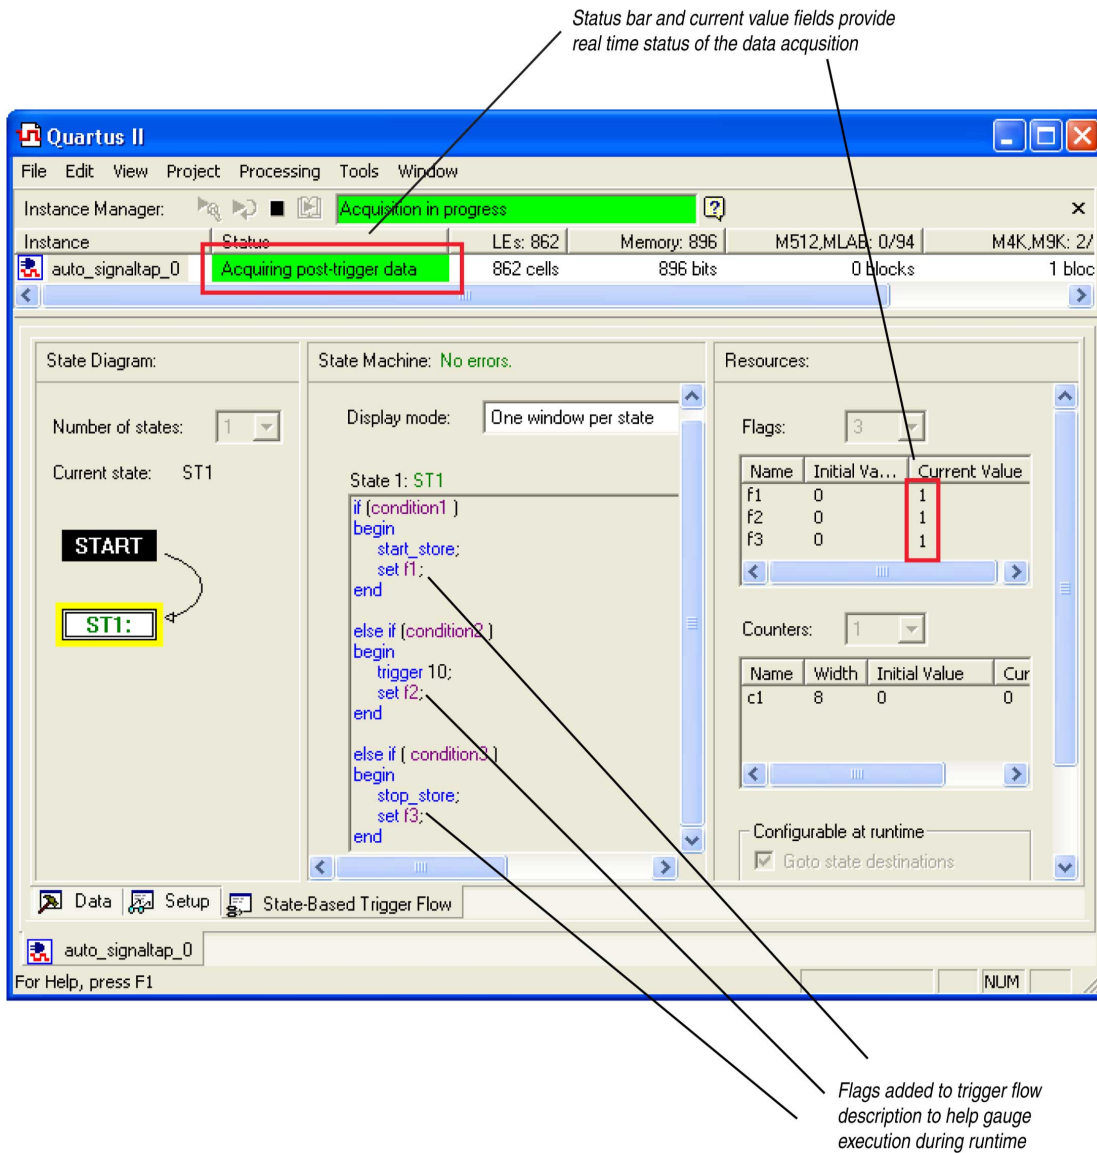



Figure 13-37: Waveform After Forcing the Analysis to Stop



The combination of using counters, Boolean and relational operators in conjunction with the `start_store` and `stop_store` commands can give a clock-cycle level of resolution to controlling the samples that are written into the acquisition buffer. The code example below shows a trigger flow description that skips three clock cycles of samples after hitting condition 1. **Figure 13-38** shows the data transaction on a continuous capture and **Figure 13-40** shows the data capture with the Trigger flow description applied, in the example below.

```

State 1: ST1
start_store
if ( condition1 )
begin
    stop_store;
    goto ST2;
end

State 2: ST2
if (c1 < 3)
    increment c1; //skip three clock cycles; c1 initialized to 0
else if (c1 == 3)
begin
    start_store; //start_store necessary to enable writing to finish
                //acquisition
    trigger;
end
    
```

Figure 13-38: Continuous Capture of Data Transaction for Example 2





Figure 13-39: Capture of Data Transaction with Trigger Flow Description Applied



## Specifying the Trigger Position

The SignalTap II Logic Analyzer allows you to specify the amount of data that is acquired before and after a trigger event. You can specify the trigger position independently between a Runtime and Power-Up Trigger. Select the desired ratio of pre-trigger data to post-trigger data by choosing one of the following ratios:

- **Pre**—Saves signal activity that occurred after the trigger (12% pre-trigger, 88% post-trigger).
- **Center**—Saves 50% pre-trigger and 50% post-trigger data.
- **Post**—Saves signal activity that occurred before the trigger (88% pre-trigger, 12% post-trigger).

These pre-defined ratios apply to both non-segmented buffers and segmented buffers.

If you use the custom-state based triggering flow, you can specify a custom trigger position. The `segment_trigger` and `trigger` actions accept a post-fill count argument. The post-fill count specifies the number of samples to capture before stopping data acquisition for the non-segmented buffer or a data segment when using the `trigger` and `segment_trigger` commands, respectively. When the captured data is displayed in the SignalTap II data window, the trigger position appears as the number of post-count samples from the end of the acquisition segment or buffer.

Sample Number of Trigger Position =  $(N - \text{Post-Fill Count})$

In this case,  $N$  is the sample depth of either the acquisition segment or non-segmented buffer.

For segmented buffers, the acquisition segments that have a post-count argument define use of the post-count setting. Segments that do not have a post-count setting default to the trigger position ratios defined in the **Setup** tab.

#### Related Information

[State-Based Triggering](#) on page 13-33

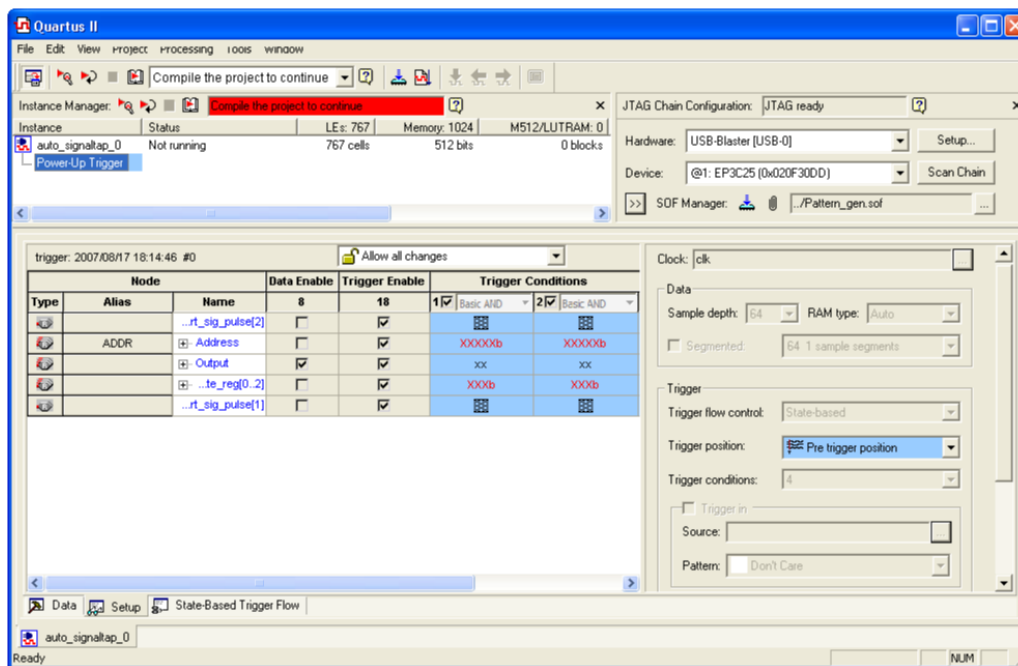
## Creating a Power-Up Trigger

Typically, the SignalTap II Logic Analyzer is used to trigger on events that occur during normal device operation. You start an analysis manually once the target device is fully powered on and the JTAG connection for the device is available. However, there may be cases when you would like to capture trigger events that occur during device initialization, immediately after the FPGA is powered on or reset. With the SignalTap II Power-Up Trigger feature, you arm the SignalTap II Logic Analyzer and capture data immediately after device programming.

## Enabling a Power-Up Trigger

You can add a different Power-Up Trigger to each logic analyzer instance in the **SignalTap II Instance Manager** pane. To enable the Power-Up Trigger for a logic analyzer instance, right-click the instance and click **Enable Power-Up Trigger**, or select the instance, and on the **Edit** menu, click **Enable Power-Up Trigger**. To disable a Power-Up Trigger, click **Disable Power-Up Trigger** in the same locations. Power-Up Trigger is shown as a child instance below the name of the selected instance with the default trigger conditions specified in the node list.

Figure 13-40: SignalTap II Logic Analyzer Editor with Power-Up Trigger Enabled



## Managing and Configuring Power-Up and Runtime Trigger Conditions

When the Power-Up Trigger is enabled for a logic analyzer instance, you can create basic and advanced trigger conditions for the trigger as you do with a Run-Time Trigger. Power-Up Trigger conditions that you can adjust are color coded light blue, while Run-Time Trigger conditions you cannot adjust remain white. Since each instance now has two sets of trigger conditions—the Power-Up Trigger and the Run-Time Trigger—you can differentiate between the two with color coding. To switch between the trigger conditions of the Power-Up Trigger and the Run-Time Trigger, double-click the instance name or the Power-Up Trigger name in the **Instance Manager**.

You cannot make changes to Power-Up Trigger conditions that would normally require a full recompile with Runtime Trigger conditions, such as adding signals, deleting signals, or changing between basic and advanced triggers. To apply these changes to the Power-Up Trigger conditions, first make the changes using the Runtime Trigger conditions.

**Note:** Any change made to the Power-Up Trigger conditions requires that you recompile the SignalTap II Logic Analyzer instance, even if a similar change to the Runtime Trigger conditions does not require a recompilation.

While creating or making changes to the trigger conditions for the Run-Time Trigger or the Power-Up Trigger, you may want to copy these conditions to the other trigger. This enables you to look for the same trigger during both power-up and runtime. To do this, right-click the instance name or the Power-Up Trigger name in the **Instance Manager** and click **Duplicate Trigger**, or select the instance name or the Power-Up Trigger name and on the **Edit** menu, click **Duplicate Trigger**.

You can also use In-System Sources and Probes in conjunction with the SignalTap II Logic Analyzer to force trigger conditions. The In-System Sources and Probes feature allows you to drive and sample values on to selected nets over the JTAG chain.

### Related Information

- [Design Debugging Using In-System Sources and Probes documentation](#) on page 17-1

## Using External Triggers

You can create a trigger input that allows you to trigger the SignalTap II Logic Analyzer from an external source. The external trigger input behaves like trigger condition 1, is evaluated, and must be `TRUE` before any other configured trigger conditions are evaluated. The logic analyzer supplies a signal to trigger external devices or other SignalTap II Logic Analyzer instances. These features allow you to synchronize external logic analysis equipment with the internal logic analyzer. Power-Up Triggers can use the external triggers feature, but they must use the same source or target signal as their associated Run-Time Trigger.

You can use external triggers to perform cross-triggering on a hard processor system (HPS). Use your processor debugger to configure the HPS to obey or disregard cross-trigger request from the FPGA, and to issue or not issue cross-trigger requests to the FPGA. Use your processor debugger in combination with the SignalTap II external trigger feature to develop a dynamic combination of cross-trigger behaviors. You can use the cross-triggering feature with the ARM Development Studio 5 (DS-5) software to implement a system-level debugging solution for your Altera SoC.

### Related Information

- [FPGA-Adaptive Software Debug and Performance Analysis white paper](#)  
Information about the ARM DS-5 debugging solution
- [Signal Configuration Pane online help](#)  
Information about setting up external triggers

## Using the Trigger Out of One Analyzer as the Trigger In of Another Analyzer

An advanced feature of the SignalTap II Logic Analyzer is the ability to use the **Trigger out** of one analyzer as the **Trigger in** to another analyzer. This feature allows you to synchronize and debug events that occur across multiple clock domains.

To perform this operation, first turn on **Trigger out** for the source logic analyzer instance. On the **Instance** list of the **Trigger out** trigger, select the targeted logic analyzer instance. For example, if the instance named `auto_signaltap_0` should trigger `auto_signaltap_1`, select `auto_signaltap_1 | trigger_in`.

Turning on **Trigger out** automatically enables the **Trigger in** of the targeted logic analyzer instance and fills in the **Instance** field of the **Trigger in** trigger with the **Trigger out** signal from the source logic analyzer instance. In this example, `auto_signaltap_0` is targeting `auto_signaltap_1`. The **Trigger In Instance** field of `auto_signaltap_1` is automatically filled in with `auto_signaltap_0 | trigger_out`.

## Compile the Design

When you add an `.stp` to your project, the SignalTap II Logic Analyzer becomes part of your design. You must compile your project to incorporate the SignalTap II logic and enable the JTAG connection you use to control the logic analyzer. When you are debugging with a traditional external logic analyzer, you must often make changes to the signals monitored as well as the trigger conditions. Because these adjustments require that you recompile your design when using the SignalTap II Logic Analyzer, use the SignalTap II Logic Analyzer feature along with incremental compilation in the Quartus II software to reduce recompilation time.

### Related Information

[Using the Incremental Compilation Design Flow online help](#)

## Faster Compilations with Quartus II Incremental Compilation

When you compile your design with an **.stp**, the `sld_signaltap` and `sld_hub` entities are automatically added to the compilation hierarchy. These two entities are the main components of the SignalTap II Logic Analyzer, providing the trigger logic and JTAG interface required for operation.

Incremental compilation enables you to preserve the synthesis and fitting results of your original design and add the SignalTap II Logic Analyzer to your design without recompiling your original source code. Incremental compilation is also useful when you want to modify the configuration of the **.stp**. For example, you can modify the buffer sample depth or memory type without performing a full compilation after the change is made. Only the SignalTap II Logic Analyzer, configured as its own design partition, must be recompiled to reflect the changes.

### Enabling Incremental Compilation for Your Design

When enabled for your design, the SignalTap II Logic Analyzer is always a separate partition. After the first compilation, you can use the SignalTap II Logic Analyzer to analyze signals from the post-fit netlist. If your partitions are designed correctly, subsequent compilations due to SignalTap II Logic Analyzer settings take less time.

The netlist type for the top-level partition defaults to **source**. To take advantage of incremental compilation, specify the Netlist types for the partitions you wish to tap as **Post-fit**.

### Related Information

[Quartus II Incremental Compilation for Hierarchical and Team-Based Design documentation](#)

### Using Incremental Compilation with the SignalTap II Logic Analyzer

The SignalTap II Logic Analyzer is automatically configured to work with the incremental compilation flow. For all signals that you want to connect to the SignalTap II Logic Analyzer from the post-fit netlist, set the netlist type of the partition containing the desired signals to **Post-Fit** with a Fitter Preservation Level of **Placement and Routing** using the Design Partitions window. Use the **SignalTap II: post-fitting filter** in the **Node Finder** to add the signals of interest to your SignalTap II configuration file. If you want to add signals from the pre-synthesis netlist, set the netlist type to **Source File** and use the **SignalTap II: pre-synthesis filter** in the **Node Finder**. Do not use the netlist type **Post-Synthesis** with the SignalTap II Logic Analyzer.

**Caution:** Be sure to conform to the following guidelines when using post-fit and pre-synthesis nodes:

- Read all incremental compilation guidelines to ensure the proper partitioning of a project.
- To speed up compile time, use only post-fit nodes for partitions specified as preservation-level post-fit.
- Do not mix pre-synthesis and post-fit nodes in any partition. If you must tap pre-synthesis nodes for a particular partition, make all tapped nodes in that partition pre-synthesis nodes and change the netlist type to **source** in the design partitions window.

Node names may be different between a pre-synthesis netlist and a post-fit netlist. In general, registers and user input signals share common names between the two netlists. During compilation, certain optimizations change the names of combinational signals in your RTL. If the type of node name chosen does not match the netlist type, the compiler may not be able to find the signal to connect to your



SignalTap II Logic Analyzer instance for analysis. The compiler issues a critical warning to alert you of this scenario. The signal that is not connected is tied to ground in the **SignalTap II data** tab.

If you do use incremental compilation flow with the SignalTap II Logic Analyzer and source file changes are necessary, be aware that you may have to remove compiler-generated post-fit net names. Source code changes force the affected partition to go through resynthesis. During synthesis, the compiler cannot find compiler-generated net names from a previous compilation.

**Note:** Altera recommends using only registered and user-input signals as debugging taps in your **.stp** whenever possible.

Both registered and user-supplied input signals share common node names in the pre-synthesis and post-fit netlist. As a result, using only registered and user-supplied input signals in your **.stp** limits the changes you need to make to your SignalTap II Logic Analyzer configuration.

You can check the nodes that are connected to each SignalTap II instance using the In-System Debugging compilation reports. These reports list each node name you selected to connect to a SignalTap II instance, the netlist type used for the particular connection, and the actual node name used after compilation. If the incremental compilation flow is not used, the In-System Debugging reports are located in the Analysis & Synthesis folder. If the incremental compilation flow is used, this report is located in the Partition Merge folder.

To verify that your original design was not modified, examine the messages in the **Partition Merge** section of the Compilation Report.

Unless you make changes to your design partitions that require recompilation, only the SignalTap II design partition is recompiled. If you make subsequent changes to only the **.stp**, only the SignalTap II design partition must be recompiled, reducing your recompilation time.

## Preventing Changes Requiring Recompilation

You can configure the **.stp** to prevent changes that normally require recompilation. To do this, select a lock mode from above the node list in the **Setup** tab. To lock your configuration, choose to allow only trigger condition changes, regardless of whether you use incremental compilation.

### Related Information

[Setup Tab \(SignalTap II Logic Analyzer\) online help](#)

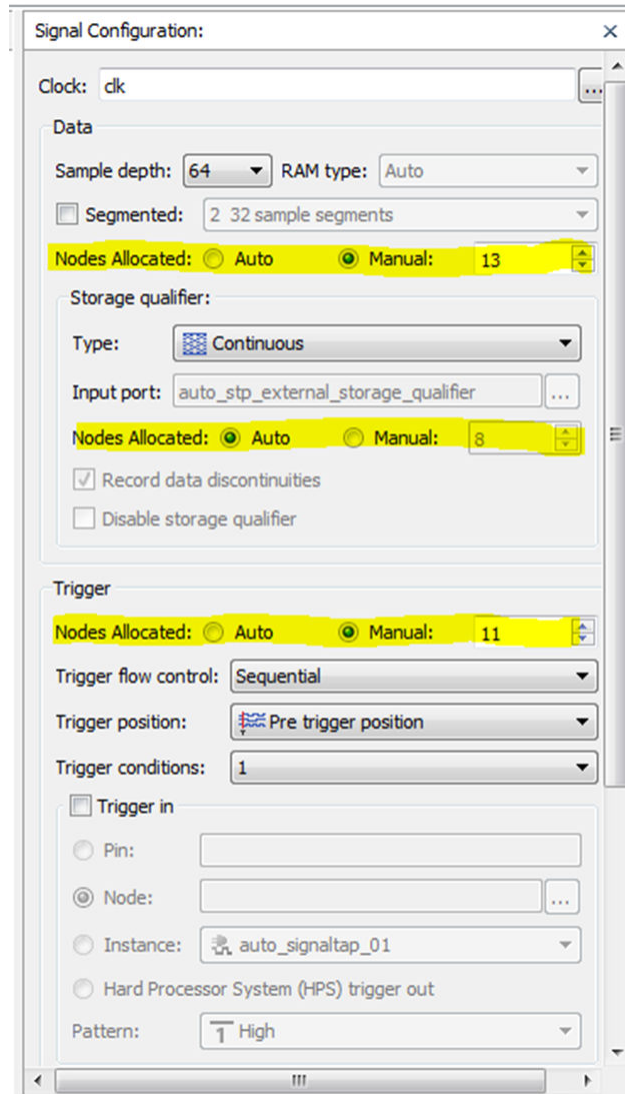
## Incremental Route with Rapid Recompile

You can use Incremental Route with Rapid Recompile to decrease compilation times. After performing a full compilation on your design, you can use the Incremental Route flow to achieve a 2-4x speedup over a flat compile. The Incremental Route flow is not compatible with Partial Reconfiguration.

Device support in Quartus II software v14.0 for Incremental Route with Rapid Recompile is limited to Arria V, Cyclone V, and Stratix V.

## Incremental Route Flow

Figure 13-41: Manually Allocate Nodes




1. Open your design, and run **Analysis & Elaboration** (or a full compilation) to give node visibility in SignalTap II.
2. Add SignalTap II to your design and specify manual allocation for Trigger and Data (Storage Qualifier, if used) nodes in the SignalTap II **Signal Configuration** pane.

**Note:** Selecting **Manual** allows you to control the number of nodes compiled into the design. This is critical for the Incremental Route flow. If you select **Auto**, the number of nodes compiled into the design will directly reflect the number of nodes (not including groups, which are not signals) currently in the **Setup** tab. If you then add a node, the number of nodes required on the

device will not match what has already been compiled, and you will then need to perform a full compilation.

3. Specify the number of nodes that you estimate will be needed for the debugging process. You can increase the number of nodes later, but this will require more compilation time.
4. Connect nodes you are interested in tapping.
5. Run a full compilation, if you have not already done a full compile on your project. Otherwise, you can start incremental compile using Rapid Recompile.
6. Debug and determine additional signals of interest.
7. (Optional) Turn on **Allow incremental route changes only** lock-mode.


Figure 13-42: Incremental Route Lock-Mode



8. Add additional nodes in the SignalTap II **Setup** tab without exceeding the number of manually allocated nodes in step 2. Avoid making changes to non-runtime configurable settings.
9. Click the Rapid Recompile icon from the toolbar (or from the Processing menu, click **Start Rapid Recompile**).

**Note:** The previous steps set up your design for Incremental Route, but the actual Incremental Route process begins when you perform a Rapid Recompile.

Figure 13-43: Rapid Recompile Icon



## Tips to Achieve Maximum Speedup

- Basic AND (which applies to Storage Qualifier as well as trigger input) is the fastest for the Incremental Route flow.
- Basic OR is slower for the Incremental Route flow, but if you avoid changing the parent-child relationship of nodes within groups, you can minimize the impact on compile time. You can change the sibling relationships of nodes.
  - Basic OR and advanced triggers require re-synthesis when the number/names of tapped nodes are changed.
- Use the Incremental Route lock-mode to avoid inadvertent changes requiring a full compilation.

## Timing Preservation with the SignalTap II Logic Analyzer

In addition to verifying functionality, timing closure is one of the most crucial processes in successfully completing a design. When you compile a project with a SignalTap II Logic Analyzer without the use of

incremental compilation, you add IP to your existing design. Therefore, you can affect the existing placement, routing, and timing of your design. To minimize the effect that the SignalTap II Logic Analyzer has on your design, Altera recommends that you use incremental compilation for your project. Incremental compilation is the default setting in new designs and can be easily enabled and configured in existing designs. With the SignalTap II Logic Analyzer instance in its own design partition, it has little to no effect on your design.

In addition to using the incremental compilation flow for your design, you can use the following techniques to help maintain timing:

- Avoid adding critical path signals to your **.stp**.
- Minimize the number of combinational signals you add to your **.stp** and add registers whenever possible.
- Specify an  $f_{MAX}$  constraint for each clock in your design.

#### Related Information

[Timing Closure and Optimization documentation](#)

## Performance and Resource Considerations

There is a necessary trade-off between the runtime flexibility of the SignalTap II Logic Analyzer, the timing performance of the SignalTap II Logic Analyzer, and resource usage. The SignalTap II Logic Analyzer allows you to select the runtime configurable parameters to balance the need for runtime flexibility, speed, and area. The default values have been chosen to provide maximum flexibility so you can complete debugging as quickly as possible; however, you can adjust these settings to determine whether there is a more optimal configuration for your design.

The following tips provide extra timing slack if you have determined that the SignalTap II logic is in your critical path, or to alleviate the resource requirements that the SignalTap II Logic Analyzer consumes if your design is resource-constrained.

If SignalTap II logic is part of your critical path, follow these tips to speed up the performance of the SignalTap II Logic Analyzer:

- **Disable runtime configurable options**—Certain resources are allocated to accommodate for runtime flexibility. If you use either advanced triggers or State-based triggering flow, disable runtime configurable parameters for a boost in  $f_{MAX}$  of the SignalTap II logic. If you are using State-based triggering flow, try disabling the **Goto state destination** option and performing a recompilation before disabling the other runtime configurable options. The **Goto state destination** option has the greatest impact on  $f_{MAX}$ , as compared to the other runtime configurable options.
- **Minimize the number of signals that have Trigger Enable selected**—All signals that you add to the **.stp** have **Trigger Enable** turned on. Turn off **Trigger Enable** for signals that you do not plan to use as triggers.
- **Turn on Physical Synthesis for register retiming**—If you have a large number of triggering signals enabled (greater than the number of inputs that would fit in a LAB) that fan-in logic to a gate-based triggering condition, such as a basic trigger condition or a logical reduction operator in the advanced trigger tab, turn on **Perform register retiming**. This can help balance combinational logic across LABs.

If your design is resource constrained, follow these tips to reduce the amount of logic or memory used by the SignalTap II Logic Analyzer:

- **Disable runtime configurable options**—Disabling runtime configurability for advanced trigger conditions or runtime configurable options in the State-based triggering flow results in using fewer LEs.
- **Minimize the number of segments in the acquisition buffer**—You can reduce the number of logic resources used for the SignalTap II Logic Analyzer by limiting the number of segments in your sampling buffer to only those required.
- **Disable the Data Enable for signals that are used for triggering only**—By default, both the **data enable** and **trigger enable** options are selected for all signals. Turning off the **data enable** option for signals used as trigger inputs only saves on memory resources used by the SignalTap II Logic Analyzer.

Because performance results are design-dependent, try these options in different combinations until you achieve the desired balance between functionality, performance, and utilization.

## Program the Target Device or Devices

After you compile your project, including the SignalTap II Logic Analyzer, configure the FPGA target device. When you are using the SignalTap II Logic Analyzer for debugging, configure the device from the **.stp** instead of the Quartus II Programmer. Because you configure from the **.stp**, you can open more than one **.stp** and program multiple devices to debug multiple designs simultaneously.

The settings in an **.stp** must be compatible with the programming **.sof** used to program the device. An **.stp** is considered compatible with an **.sof** when the settings for the logic analyzer, such as the size of the capture buffer and the signals selected for monitoring or triggering, match the way the target device is programmed. If the files are not compatible, you can still program the device, but you cannot run or control the logic analyzer from the SignalTap II Logic Analyzer Editor.

**Note:** When the SignalTap II Logic Analyzer detects incompatibility after analysis is started, a system error message is generated containing two CRC values, the expected value and the value retrieved from the **.stp** instance on the device. The CRC values are calculated based on all SignalTap II settings that affect the compilation.

To ensure programming compatibility, make sure to program your device with the latest **.sof** created from the most recent compilation. Checking whether or not a particular **.sof** is compatible with the current SignalTap II configuration is achieved quickly by attaching the **.sof** to the SOF manager.

Before starting a debugging session, do not make any changes to the **.stp** settings that would require recompiling the project. You can check the SignalTap II status display at the top of the **Instance Manager** pane to verify whether a change you made requires recompiling the project, producing a new **.sof**. This gives you the opportunity to undo the change, so that you do not need to recompile your project. To prevent any such changes, select **Allow trigger condition changes only** to lock the **.stp**. The Incremental Route lock mode, **Allow incremental route changes only**, limits changes which will only require an Incremental Route using Rapid Recompile, and not a full compile.

Although the Quartus II project is not required when using an **.stp**, it is recommended. The project database contains information about the integrity of the current SignalTap II Logic Analyzer session. Without the project database, there is no way to verify that the current **.stp** matches the **.sof** that is downloaded to the device. If you have an **.stp** that does not match the **.sof**, incorrect data is captured in the SignalTap II Logic Analyzer.

Related Information

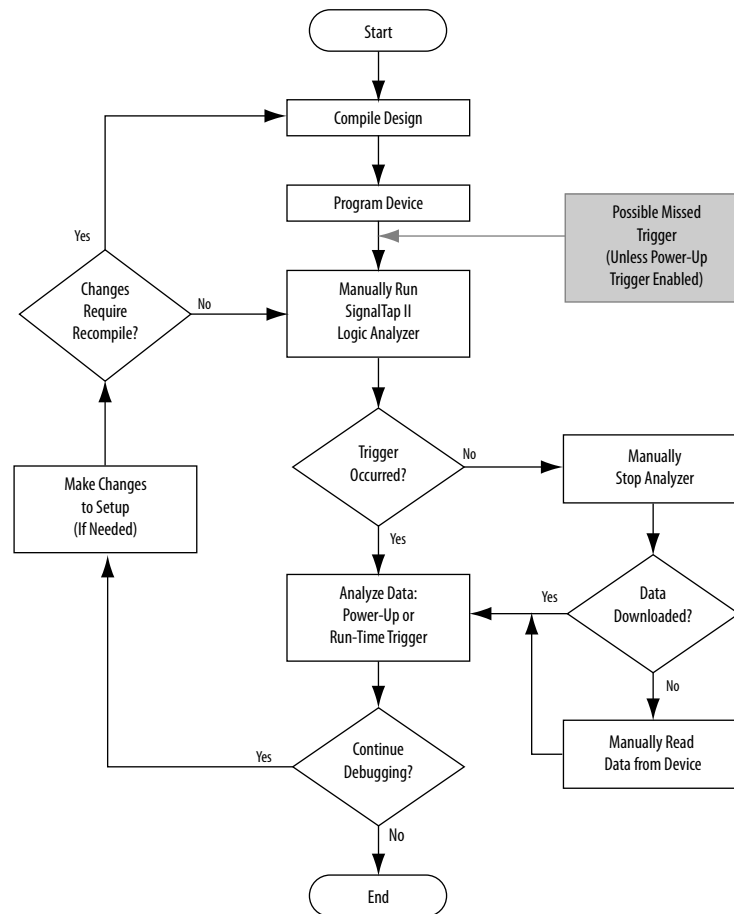
[Running the SignalTap II Logic Analyzer online help](#)

## Run the SignalTap II Logic Analyzer

After the device is configured with your design that includes the SignalTap II Logic Analyzer, perform debugging operations in a manner similar to when you use an external logic analyzer. You initialize the logic analyzer by starting an analysis. When your trigger event occurs, the captured data is stored in the memory buffer on the device and then transferred to the .stp with the JTAG connection.

You can also perform the equivalent of a force trigger instruction that lets you view the captured data currently in the buffer without a trigger event occurring. [Figure 13-44](#) illustrates a flow that shows how you operate the SignalTap II Logic Analyzer. The flowchart indicates where Power-Up and Runtime Trigger events occur and when captured data from these events is available for analysis.

**Figure 13-44: Power-Up and Runtime Trigger Events Flowchart**



You can also use In-System Sources and Probes in conjunction with the SignalTap II Logic Analyzer to force trigger conditions. The In-System Sources and Probes feature allows you to drive and sample values on to selected signals over the JTAG chain.

**Related Information**[Running the SignalTap II Logic Analyzer online help](#)

Information on running the analyzer from the **Instance Manager** pane

[Design Debugging Using In-System Sources and Probes documentation](#) on page 17-1

## Runtime Reconfigurable Options

Certain settings in the **.stp** are changeable without recompiling your design when you use Runtime Trigger mode.

**Table 13-10: Runtime Reconfigurable Features**

Runtime Reconfigurable Setting	Description
Basic Trigger Conditions and Basic Storage Qualifier Conditions	All signals that have the Trigger condition turned on can be changed to any basic trigger condition value without recompiling.
Advanced Trigger Conditions and Advanced Storage Qualifier Conditions	Many operators include runtime configurable settings. For example, all comparison operators are runtime-configurable. Configurable settings are shown with a white background in the block representation. This runtime reconfigurable option is turned on in the <b>Object Properties</b> dialog box.
Switching between a storage-qualified and a continuous acquisition	Within any storage-qualified mode, you can switch to continuous capture mode without recompiling the design. To enable this feature, turn on <b>disable storage qualifier</b> .
State-based trigger flow parameters	<a href="#">Table 13-5</a> lists Reconfigurable State-based trigger flow options.

Runtime Reconfigurable options can potentially save time during the debugging cycle by allowing you to cover a wider possible scenario of events without the need to recompile the design. You may experience a slight impact to the performance and logic utilization. You can turn off Runtime re-configurability for Advanced Trigger Conditions and the State-based trigger flow parameters, boosting performance and decreasing area utilization.

You can configure the **.stp** to prevent changes that normally require recompilation. To do this, in the **Setup** tab, select **Allow Trigger Condition changes only** above the node list.

Incremental Route lock mode, **Allow incremental route changes only**, limits changes which will only require an Incremental Route compilation, and not a full compile.

The example below illustrates a potential use case for Runtime Reconfigurable features. This example provides a storage qualified enabled State-based trigger flow description and shows how you can modify the size of a capture window at runtime without a recompile. This example gives you equivalent functionality to a segmented buffer with a single trigger condition where the segment sizes are runtime reconfigurable.

```
state ST1:
if ( condition1 && (c1 <= m) ) // each "segment" triggers on condition
//1
begin // m = number of total "segments"
start_store;
increment c1;
goto ST2;
End

else (c1 > m) //This else condition handles the last
```

```

//segment.
begin
    start_store
    Trigger (n-1)
end

state ST2:
if ( c2 >= n) //n = number of samples to capture in each
//segment.
begin
    reset c2;
    stop_store;
    goto ST1;
end

else (c2 < n)
begin
    increment c2;
    goto ST2;
end
end

```

Note to example :

1.  $m \times n$  must equal the sample depth to efficiently use the space in the sample buffer.

Figure 13-45 shows a segmented buffer described by the trigger flow in example above.

During runtime, the values  $m$  and  $n$  are runtime reconfigurable. By changing the  $m$  and  $n$  values in the preceding trigger flow description, you can dynamically adjust the segment boundaries without incurring a recompile.

**Figure 13-45: Segmented Buffer Created with Storage Qualifier and State-Based Trigger**



Note to figure :

1. Total sample depth is fixed, where  $m \times n$  must equal sample depth.

You can add states into the trigger flow description and selectively mask out specific states and enable other ones at runtime with status flags.

The example below shows a modified description of the example above with an additional state inserted. You use this extra state to specify a different trigger condition that does not use the storage qualifier feature. You insert status flags into the conditional statements to control the execution of the trigger flow.

```

state ST1 :
if (condition2 && f1) //additional state added for a non-
segmented //acquisition Set f1 to enable state
begin
    start_store;
    trigger
end
else if (! f1)
goto ST2;
state ST2:
if ( (condition1 && (c1 <= m) && f2) //f2 status flag used to mask state. Set
f2

```



```

//to enable.
begin
    start_store;
    increment c1;
    goto ST3;
end
else (c1 > m )
    start_store
    Trigger (n-1)
end
state ST3:
if ( c2 >= n)
begin
    reset c2;
    stop_store;
    goto ST1;
end
else (c2 < n)
begin
    increment c2;
    goto ST2;
end
end

```

## SignalTap II Status Messages

**Table 13-11** describes the text messages that might appear in the SignalTap II Status Indicator in the **Instance Manager** pane before, during, and after a data acquisition. Use these messages to monitor the state of the logic analyzer or what operation it is performing.

**Table 13-11: Text Messages in the SignalTap II Status Indicator**

Message	Message Description
Not running	The SignalTap II Logic Analyzer is not running. There is no connection to a device or the device is not configured.
(Power-Up Trigger) Waiting for clock <b>(1)</b>	The SignalTap II Logic Analyzer is performing a Runtime or Power-Up Trigger acquisition and is waiting for the clock signal to transition.
Acquiring (Power-Up) pre-trigger data <b>(1)</b>	The trigger condition has not been evaluated yet. A full buffer of data is collected if using the non-segmented buffer acquisition mode and storage qualifier type is continuous.
Trigger In conditions met	Trigger In condition has occurred. The SignalTap II Logic Analyzer is waiting for the condition of the first trigger condition to occur. This can appear if Trigger In is specified.
Waiting for (Power-up) trigger <b>(1)</b>	The SignalTap II Logic Analyzer is now waiting for the trigger event to occur.
Trigger level <x> met	The condition of trigger condition $x$ has occurred. The SignalTap II Logic Analyzer is waiting for the condition specified in condition $x + 1$ to occur.
Acquiring (power-up) post-trigger data <b>(1)</b>	The entire trigger event has occurred. The SignalTap II Logic Analyzer is acquiring the post-trigger data. The amount of post-trigger data collected is you define between 12%, 50%, and 88% when the non-segmented buffer acquisition mode is selected.

Message	Message Description
Offload acquired (Power-Up) data (1)	Data is being transmitted to the Quartus II software through the JTAG chain.
Ready to acquire	The SignalTap II Logic Analyzer is waiting for you to initialize the analyzer.

Note to [Table 13-11](#) :

1. This message can appear for both Runtime and Power-Up Trigger events. When referring to a Power-Up Trigger, the text in parentheses is added.

**Note:** In segmented acquisition mode, pre-trigger and post-trigger do not apply.

## View, Analyze, and Use Captured Data

Once a trigger event has occurred or you capture data manually, you can use the SignalTap II interface to examine the data, and use your findings to help debug your design.

When in the Data view, you can use the drag-to-zoom feature by left-clicking to isolate the data of interest.

### Related Information


[Analyzing Data in the SignalTap II Logic Analyzer online help](#)

Information about what you can do with captured data

## Capturing Data Using Segmented Buffers

Segmented Acquisition buffers allow you to perform multiple captures with a separate trigger condition for each acquisition segment. This feature allows you to capture a recurring event or sequence of events that span over a long period time efficiently. Each acquisition segment acts as a non-segmented buffer, continuously capturing data when it is activated. When you run an analysis with the **segmented buffer** option enabled, the SignalTap II Logic Analyzer performs back-to-back data captures for each acquisition segment within your data buffer. The trigger flow, or the type and order in which the trigger conditions evaluate for each buffer, is defined by either the Sequential trigger flow control or the Custom State-based trigger flow control. [Figure 13-46](#) shows a segmented acquisition buffer with four segments represented as four separate non-segmented buffers.

**Figure 13-46: Segmented Acquisition Buffer**




The SignalTap II Logic Analyzer finishes an acquisition with a segment, and advances to the next segment to start a new acquisition. Depending on when a trigger condition occurs, it may affect the way the data capture appears in the waveform viewer. [Figure 13-46](#) illustrates the method in which data is captured.

The Trigger markers in **Figure 13-46**—Trigger 1, Trigger 2, Trigger 3 and Trigger 4—refer to the evaluation of the `segment_trigger` and `trigger` commands in the Custom State-based trigger flow. If you use a sequential flow, the Trigger markers refer to trigger conditions specified within the **Setup** tab.

If the Segment 1 Buffer is the active segment and Trigger 1 occurs, the SignalTap II Logic Analyzer starts evaluating Trigger 2 immediately. Data Acquisition for Segment 2 buffer starts when either Segment Buffer 1 finishes its post-fill count, or when Trigger 2 evaluates as `TRUE`, whichever condition occurs first. Thus, trigger conditions associated with the next buffer in the data capture sequence can preempt the post-fill count of the current active buffer. This allows the SignalTap II Logic Analyzer to accurately capture all of the trigger conditions that have occurred. Samples that have not been used appear as a blank space in the waveform viewer.

**Figure 13-47** shows an example of a capture using sequential flow control with the trigger condition for each segment specified as **Don't Care**. Each segment before the last captures only one sample, because the next trigger condition immediately preempts capture of the current buffer. The trigger position for all segments is specified as pre-trigger (10% of the data is before the trigger condition and 90% of the data is after the trigger position). Because the last segment starts immediately with the trigger condition, the segment contains only post-trigger data. The three empty samples in the last segment are left over from the pre-trigger samples that the SignalTap II Logic Analyzer allocated to the buffer.

**Figure 13-47: Segmented Capture with Preemption of Acquisition Segments**



Note to **Figure 13-47** :

1. A segmented acquisition buffer using the sequential trigger flow with a trigger condition specified as **Don't Care**. All segments, with the exception of the last segment, capture only one sample because the next trigger condition preempts the current buffer from filling to completion.

For the sequential trigger flow, the **Trigger Position** option applies to every segment in the buffer. For maximum flexibility on how the trigger position is defined, use the custom state-based trigger flow. By adjusting the trigger position specific to your debugging requirements, you can help maximize the use of the allocated buffer space.

## Differences in Pre-fill Write Behavior Between Different Acquisition Modes

The SignalTap II Logic Analyzer uses one of the following three modes when writing into the acquisition memory:

- **Non-segmented buffer**
- **Non-segmented buffer with a storage qualifier**
- **Segmented buffer**

There are subtle differences in the amount of data captured immediately after running the SignalTap II Logic Analyzer and before any trigger conditions occur. A non-segmented buffer, running in continuous mode, completely fills the buffer with sampled data before evaluating any trigger conditions. Thus, a non-segmented capture without any storage qualification enabled always shows a waveform with a full buffer's worth of data captured.

Filling the buffer provides you with as much data as possible within the capture window. The buffer gets pre-filled with data samples prior to evaluating the trigger condition. As such, SignalTap requires that the buffer be filled at least once before any data can be retrieved through the JTAG connection and prevents the buffer from being dumped during the first acquisition prior to a trigger condition when you perform a **Stop Analysis**.

For segmented buffers and non-segmented buffers using any storage qualification mode, the SignalTap II Logic Analyzer immediately evaluates all trigger conditions while writing samples into the acquisition memory. The logic analyzer evaluates each trigger condition before acquiring a full buffer's worth of samples. This evaluation is especially important when using any storage qualification on the data set. The logic analyzer may miss a trigger condition if it waits until a full buffer's worth of data is captured before evaluating any trigger conditions.

If the trigger event occurs on any data sample before the specified amount of pre-trigger data has occurred, then the SignalTap II Logic Analyzer triggers and begins filling memory with post-trigger data, regardless of the amount of pre-trigger data you specify. For example, if you set the trigger position to 50% and set the logic analyzer to trigger on a processor reset, start the logic analyzer, and then power on your target system, the logic analyzer triggers. However, the logic analyzer memory is filled only with post-trigger data, and not any pre-trigger data, because the trigger event, which has higher precedence than the capture of pre-trigger data, occurred before the pre-trigger condition was satisfied.

**Figure 13-48** and **Figure 13-49** show the difference between a non-segmented buffer in continuous mode and a non-segmented buffer using a storage qualifier. The logic analyzer for the waveforms below is configured with a sample depth of 64 bits, with a trigger position specified as **Post trigger position**.

**Figure 13-48: SignalTap II Logic Analyzer Continuous Data Capture**




Note to **Figure 13-48** :

- 1. Continuous capture mode with post-trigger position.
- 2. Capture of a recurring pattern using a non-segmented buffer in continuous mode. The SignalTap II Logic Analyzer is configured with a basic trigger condition (shown in the figure as "Trig1") with a sample depth of 64 bits.

Notice in **Figure 13-48** that Trig1 occurs several times in the data buffer before the SignalTap II Logic Analyzer actually triggers. A full buffer's worth of data is captured before the logic analyzer evaluates any trigger conditions. After the trigger condition occurs, the logic analyzer continues acquisition until it captures eight additional samples (12% of the buffer, as defined by the "post-trigger" position).

Figure 13-49: SignalTap II Logic Analyzer Conditional Data Capture



Note to [Figure 13-49](#) :

1. Conditional capture, storage always enabled, post-fill count.
2. SignalTap II Logic Analyzer capture of a recurring pattern using a non-segmented buffer in conditional mode. The logic analyzer is configured with a basic trigger condition (shown in the figure as "Trig1"), with a sample depth of 64 bits. The "Trigger in" condition is specified as "Don't care", which means that every sample is captured.

Notice in [Figure 13-49](#) that the logic analyzer triggers immediately. As in [Figure 13-48](#), the logic analyzer completes the acquisition with eight samples, or 12% of 64, the sample capacity of the acquisition buffer.

## Creating Mnemonics for Bit Patterns

The mnemonic table feature allows you to assign a meaningful name to a set of bit patterns, such as a bus. To create a mnemonic table, right-click in the **Setup** or **Data** tab of an **.stp** and click **Mnemonic Table Setup**. You create a mnemonic table by entering sets of bit patterns and specifying a label to represent each pattern. Once you have created a mnemonic table, assign the table to a group of signals. To assign a mnemonic table, right-click on the group, click **Bus Display Format** and select the desired mnemonic table.

You use the labels you create in a table in different ways on the **Setup** and **Data** tabs. On the **Setup** tab, you can create basic triggers with meaningful names by right-clicking an entry in the **Trigger Conditions** column and selecting a label from the table you assigned to the signal group. On the **Data** tab, if any captured data matches a bit pattern contained in an assigned mnemonic table, the signal group data is replaced with the appropriate label, making it easy to see when expected data patterns occur.

## Automatic Mnemonics with a Plug-In

When you use a plug-in to add signals to an **.stp**, mnemonic tables for the added signals are automatically created and assigned to the signals defined in the plug-in. To enable these mnemonic tables manually, right-click on the name of the signal or signal group. On the **Bus Display Format** shortcut menu, then click the name of the mnemonic table that matches the plug-in.

As an example, the Nios II plug-in helps you to monitor signal activity for your design as the code is executed. If you set up the logic analyzer to trigger on a function name in your Nios II code based on data from an **.elf**, you can see the function name in the **Instance Address** signal group at the trigger sample, along with the corresponding disassembled code in the **Disassembly** signal group, as shown in [Figure 13-50](#). Captured data samples around the trigger are referenced as offset addresses from the trigger function name.

Figure 13-50: Data Tab when the Nios II Plug-In is Used

Type	Alias	Name	37	Value	38	48	49	50	51	52
PC	...	...Nios II Inst Address	alt_main+0x8		<empty>		alt_main+0xc		<empty>	<empty>
DIS	...	...Nios II Disassembly	mov fp, sp		<empty>		movi r2, 2		<empty>	<empty>

## Locating a Node in the Design

When you find the source of an error in your design using the SignalTap II Logic Analyzer, you can use the node locate feature to locate that signal in many of the tools found in the Quartus II software, as well as in your design files. This lets you find the source of the problem quickly so you can modify your design to correct the flaw. To locate a signal from the SignalTap II Logic Analyzer in one of the Quartus II software tools or your design files, right-click on the signal in the **.stp**, and click **Locate in <tool name>**.

You can locate a signal from the node list with the following tools:

- Assignment Editor
- Pin Planner
- Timing Closure Floorplan
- Chip Planner
- Resource Property Editor
- Technology Map Viewer
- RTL Viewer
- Design File

## Saving Captured Data

The data log shows the history of captured data and the triggers used to capture the data. The SignalTap II Logic Analyzer acquires data, stores it in a log, and displays it as waveforms. When the logic analyzer is in auto-run mode and a trigger event occurs more than once, captured data for each time the trigger occurred is stored as a separate entry in the data log. This allows you to review the captured data for each trigger event. The default name for a log is based on the time when the data was acquired. Altera recommends that you rename the data log with a more meaningful name.

The logs are organized in a hierarchical manner; similar logs of captured data are grouped together in trigger sets. To open the **Data Log** pane, on the View menu, select **Data Log**. To turn on data logging, turn on **Enable data log** in the **Data Log** (Figure 13-19). To recall and activate a data log for a given trigger set, double-click the name of the data log in the list. The time stamping for the Data Log entries display the wall-clock time when SignalTap II triggered and the elapsed time from when acquisition started to when the device triggered.

### Related Information

[Managing Multiple SignalTap II Files and Configurations](#) on page 13-22

You can use the Data Log feature for organizing different sets of trigger conditions and different sets of signal configurations.

## Exporting Captured Data to Other File Formats

You can export captured data to the following file formats, for use with other EDA simulation tools:

- Comma Separated Values File (**.csv**)
- Table File (**.tbl**)
- Value Change Dump File (**.vcd**)
- Vector Waveform File (**.vwf**)
- Graphics format files (**.jpg**, **.bmp**)

To export the captured data from SignalTap II Logic Analyzer, on the File menu, click **Export** and specify the **File Name**, **Export Format**, and **Clock Period**.

## Creating a SignalTap II List File

Captured data can also be viewed in an **.stp** list file. An **.stp** list file is a text file that lists all the data captured by the logic analyzer for a trigger event. Each row of the list file corresponds to one captured sample in the buffer. Columns correspond to the value of each of the captured signals or signal groups for that sample. If a mnemonic table was created for the captured data, the numerical values in the list are replaced with a matching entry from the table. This is especially useful with the use of a plug-in that includes instruction code disassembly. You can immediately see the order in which the instruction code was executed during the same time period of the trigger event. To create an **.stp** list file in the Quartus II software, on the File menu, select **Create/Update** and click **Create SignalTap II List File**.

## Other Features

The SignalTap II Logic Analyzer has other features that do not necessarily belong to a particular task in the task flow.

## Using the SignalTap II MATLAB MEX Function to Capture Data

If you use MATLAB for DSP design, you can call the MATLAB MEX function `alt_signaltap_run`, built into the Quartus II software, to acquire data from the SignalTap II Logic Analyzer directly into a matrix in the MATLAB environment. If you use the MATLAB MEX function in a loop, you can perform as many acquisitions in the same amount of time as you can when using SignalTap II in the Quartus II software environment.

**Note:** The SignalTap II MATLAB MEX function is available in the Windows version and Linux version of the Quartus II software. It is compatible with MATLAB Release 14 Original Release Version 7 and later.

To set up the Quartus II software and the MATLAB environment to perform SignalTap II acquisitions, perform the following steps:

1. In the Quartus II software, create an **.stp** file.
2. In the node list in the **Data** tab of the SignalTap II Logic Analyzer Editor, organize the signals and groups of signals into the order in which you want them to appear in the MATLAB matrix. Each

column of the imported matrix represents a single SignalTap II acquisition sample, while each row represents a signal or group of signals in the order they are organized in the **Data** tab.

**Note:** Signal groups acquired from the SignalTap II Logic Analyzer and transferred into the MATLAB MEX function are limited to a width of 32 signals. If you want to use the MATLAB MEX function with a bus or signal group that contains more than 32 signals, split the group into smaller groups that do not exceed the 32-signal limit.

3. Save the **.stp** and compile your design. Program your device and run the SignalTap II Logic Analyzer to ensure your trigger conditions and signal acquisition work correctly.
4. In the MATLAB environment, add the Quartus II binary directory to your path with the following command:

```
addpath <Quartus install directory>\win
```

You can view the help file for the MEX function by entering the following command in MATLAB without any operators:

```
alt_signaltap_run
```

Use the MATLAB MEX function to open the JTAG connection to the device and run the SignalTap II Logic Analyzer to acquire data. When you finish acquiring data, close the JTAG connection.

To open the JTAG connection and begin acquiring captured data directly into a MATLAB matrix called `stp`, use the following command:

```
stp = alt_signaltap_run \
(' <stp filename>' [, ('signed' | 'unsigned')], '<instance names>' [, \
'<signalset name>' [, '<trigger name>'] ] ] ] );
```

When capturing data you must assign a filename, for example, `<stp filename>` as a requirement of the MATLAB MEX function. Other MATLAB MEX function options are described in [Table 13-12](#).

**Table 13-12: SignalTap II MATLAB MEX Function Options**

Option	Usage	Description
signed	'signed'	The <b>signed</b> option turns signal group data into 32-bit two's-complement signed integers. The MSB of the group as defined in the SignalTap II <b>Data</b> tab is the sign bit. The <b>unsigned</b> option keeps the data as an unsigned integer. The default is <b>signed</b> .
unsigned	'unsigned'	
<instance name>	'auto_signaltap_0'	Specify a SignalTap II instance if more than one instance is defined. The default is the first instance in the <b>.stp</b> , <code>auto_signaltap_0</code> .
<signal set name>	'my_signalset'	Specify the signal set and trigger from the SignalTap II data log if multiple configurations are present in the <b>.stp</b> . The default is the active signal set and trigger in the file.
<trigger name>	'my_trigger'	



You can enable or disable verbose mode to see the status of the logic analyzer while it is acquiring data. To enable or disable verbose mode, use the following commands:

```
alt_signaltap_run('VERBOSE_ON');  
alt_signaltap_run('VERBOSE_OFF');
```

When you finish acquiring data, close the JTAG connection with the following command:

```
alt_signaltap_run('END_CONNECTION');
```

For more information about the use of MATLAB MEX functions in MATLAB, refer to the MATLAB Help.

## Using SignalTap II in a Lab Environment

You can install a stand-alone version of the SignalTap II Logic Analyzer. This version is particularly useful in a lab environment in which you do not have a workstation that meets the requirements for a complete Quartus II installation, or if you do not have a license for a full installation of the Quartus II software. The standalone version of the SignalTap II Logic Analyzer is included with and requires the Quartus II stand-alone Programmer which is available from the Downloads page of the [Altera website](#).

## Remote Debugging Using the SignalTap II Logic Analyzer

### Debugging Using a Local PC and an Altera SoC

You can use the System Console with SignalTap II Logic Analyzer to remote debug your Altera SoC. This method requires one local PC, an existing TCP/IP connection, a programming device at the remote location, and an Altera SoC.

#### Related Information

[Remote Hardware Debugging over TCP/IP for Altera SoC application note](#)

### Debugging Using a Local PC and a Remote PC

You can use the SignalTap II Logic Analyzer to debug a design that is running on a device attached to a PC in a remote location.

To perform a remote debugging session, you must have the following setup:

- The Quartus II software installed on the local PC
- Stand-alone SignalTap II Logic Analyzer or the full version of the Quartus II software installed on the remote PC
- Programming hardware connected to the device on the PCB at the remote location
- TCP/IP protocol connection

### Equipment Setup

On the PC in the remote location, install the standalone version of the SignalTap II Logic Analyzer, included in the Quartus II standalone Programmer, or the full version of the Quartus II software. This remote computer must have Altera programming hardware connected, such as the EthernetBlaster or USB-Blaster.

On the local PC, install the full version of the Quartus II software. This local PC must be connected to the remote PC across a LAN with the TCP/IP protocol.

**Related Information**

[Using the JTAG Server online help](#)

Information about enabling remote access to a JTAG server

## Using the SignalTap II Logic Analyzer in Devices with Configuration Bitstream Security

Certain device families support bitstream decryption during configuration using an on-device AES decryption engine. You can still use the SignalTap II Logic Analyzer to analyze functional data within the FPGA. However, note that JTAG configuration is not possible after the security key has been programmed into the device.

Altera recommends that you use an unencrypted bitstream during the prototype and debugging phases of the design. Using an unencrypted bitstream allows you to generate new programming files and reconfigure the device over the JTAG connection during the debugging cycle.

If you must use the SignalTap II Logic Analyzer with an encrypted bitstream, first configure the device with an encrypted configuration file using Passive Serial (PS), Fast Passive Parallel (FPP), or Active Serial (AS) configuration modes. The design must contain at least one instance of the SignalTap II Logic Analyzer. After the FPGA is configured with a SignalTap II Logic Analyzer instance in the design, when you open the SignalTap II Logic Analyzer in the Quartus II software, you then scan the chain and are ready to acquire data with the JTAG connection.

## Backward Compatibility with Previous Versions of Quartus II Software

You can open an **.stp** created in a previous version in a current version of the Quartus II software. However, opening an **.stp** modifies it so that it cannot be opened in a previous version of the Quartus II software.

If you have a Quartus II project file from a previous version of the software, you may have to update the **.stp** configuration file to recompile the project. You can update the configuration file by opening the SignalTap II Logic Analyzer. If you need to update your configuration, a prompt appears asking if you would like to update the **.stp** to match the current version of the Quartus II software.

## SignalTap II Command-Line Options

To compile your design with the SignalTap II Logic Analyzer using the command prompt, use the `quartus_stp` command. **Table 13-13** shows the options that help you use the `quartus_stp` executable.

**Table 13-13: SignalTap II Command-Line Options**

Option	Usage	Description
<code>stp_file</code>	<code>quartus_stp --stp_file &lt;stp_filename&gt;</code>	Assigns the specified <b>.stp</b> to the <code>USE_SIGNALTAP_FILE</code> in the <b>.qsf</b> .

Option	Usage	Description
enable	<code>quartus_stp --enable</code>	Creates assignments to the specified <b>.stp</b> in the <b>.qsf</b> and changes <code>ENABLE_SIGNALTAP</code> to ON. The SignalTap II Logic Analyzer is included in your design the next time the project is compiled. If no <b>.stp</b> is specified in the <b>.qsf</b> , the <code>--stp_file</code> option must be used. If the <code>--enable</code> option is omitted, the current value of <code>ENABLE_SIGNALTAP</code> in the <b>.qsf</b> is used.
disable	<code>quartus_stp --disable</code>	Removes the <b>.stp</b> reference from the <b>.qsf</b> and changes <code>ENABLE_SIGNALTAP</code> to OFF. The SignalTap II Logic Analyzer is removed from the design database the next time you compile your design. If the <code>--disable</code> option is omitted, the current value of <code>ENABLE_SIGNALTAP</code> in the <b>.qsf</b> is used.
<code>create_signaltap_hdl_file</code>	<code>quartus_stp --create_signaltap_hdl_file</code>	Creates an <b>.stp</b> representing the SignalTap II instance. The file is based on the last compilation. You must use the <code>--stp_file</code> option to create an <b>.stp</b> properly. Analogous to the <b>Create SignalTap II File from Design Instance(s)</b> command in the Quartus II software.

The first example illustrates how to compile a design with the SignalTap II Logic Analyzer at the command line.

```
quartus_stp filtref --stp_file stp1.stp --enable
quartus_map filtref --source=filtref.bdf --family=CYCLONE
quartus_fit filtref --part=EP1C12Q240C6 --fmax=80MHz --tsu=8ns
quartus_asm filtref
```

The `quartus_stp --stp_file stp1.stp --enable` command creates the QSF variable and instructs the Quartus II software to compile the **stp1.stp** file with your design. The `--enable` option must be applied for the SignalTap II Logic Analyzer to compile properly into your design.

The example below shows how to create a new **.stp** after building the SignalTap II Logic Analyzer instance with the IP Catalog.

```
quartus_stp filtref --create_signaltap_hdl_file --stp_file stp1.stp
```

#### Related Information

#### [Command-Line Scripting documentation](#)

Information about the other command line executables and options

## SignalTap II Tcl Commands

The `quartus_stp` executable supports a Tcl interface that allows you to capture data without running the Quartus II GUI. You cannot execute SignalTap II Tcl commands from within the Tcl console in the Quartus II software. They must be executed from the command line with the `quartus_stp` executable. To execute a Tcl file that has SignalTap II Logic Analyzer Tcl commands, use the following command:

```
quartus_stp -t <Tcl file>
```

The example is an excerpt from a script you can use to continuously capture data. Once the trigger condition is met, the data is captured and stored in the data log.

```
#opens signaltap session
open_session -name stpl.stp
#start acquisition of instance auto_signaltap_0 and
#auto_signaltap_1 at the same time
#calling run_multiple_end will start all instances
#run after run_multiple_start call
run_multiple_start
run -instance auto_signaltap_0 -signal_set signal_set_1 -trigger \
trigger_1 -data_log log_1 -timeout 5
run -instance auto_signaltap_1 -signal_set signal_set_1 -trigger \
trigger_1 -data_log log_1 -timeout 5
run_multiple_end
#close signaltap session
close_session
```

When the script is completed, open the **.stp** that you used to capture data to examine the contents of the Data Log.

#### Related Information

[::quartus::stp online help](#)

Information about Tcl commands that you can use with the SignalTap II Logic Analyzer Tcl package

## Design Example: Using SignalTap II Logic Analyzers

The system in this example contains many components, including a Nios processor, a direct memory access (DMA) controller, on-chip memory, and an interface to external SDRAM memory. In this example, the Nios processor executes a simple C program from on-chip memory and waits for you to press a button. After you press a button, the processor initiates a DMA transfer, which you analyze using the SignalTap II Logic Analyzer.

#### Related Information

[AN 446: Debugging Nios II Systems with the SignalTap II Embedded Logic Analyzer application note](#)

## Custom Triggering Flow Application Examples

The custom triggering flow in the SignalTap II Logic Analyzer is most useful for organizing a number of triggering conditions and for precise control over the acquisition buffer. This section provides two application examples for defining a custom triggering flow within the SignalTap II Logic Analyzer. Both examples can be easily copied and pasted directly into the state machine description box by using the state display mode **All states in one window**.

#### Related Information

[On-chip Debugging Design Examples website](#)

## Design Example 1: Specifying a Custom Trigger Position

Actions to the acquisition buffer can accept an optional post-count argument. This post-count argument enables you to define a custom triggering position for each segment in the acquisition buffer. The example shows how to apply a trigger position to all segments in the acquisition buffer. The example describes a triggering flow for an acquisition buffer split into four segments. If each acquisition segment is 64 samples

### Design Example 2: Trigger When triggercond1 Occurs Ten Times between triggercond2 and triggercond3

in depth, the trigger position for each buffer will be at sample #34. The acquisition stops after all four segments are filled once.


```

if (c1 == 3 && condition1)
    trigger 30;
else if ( condition1 )
begin
    segment_trigger 30;
    increment c1;
end

```

Each segment acts as a non-segmented buffer that continuously updates the memory contents with the signal values. The last acquisition before stopping the buffer is displayed on the **Data** tab as the last sample number in the affected segment. The trigger position in the affected segment is then defined by  $N - \text{post count fill}$ , where  $N$  is the number of samples per segment. **Figure 13-51** illustrates the triggering position.

**Figure 13-51: Specifying a Custom Trigger Position**



### Design Example 2: Trigger When triggercond1 Occurs Ten Times between triggercond2 and triggercond3

The custom trigger flow description is often useful to count a sequence of events before triggering the acquisition buffer. The example shows such a sample flow. This example uses three basic triggering conditions configured in the SignalTap II **Setup** tab.

This example triggers the acquisition buffer when `condition1` occurs after `condition3` and occurs ten times prior to `condition3`. If `condition3` occurs prior to ten repetitions of `condition1`, the state machine transitions to a permanent wait state.

```

state ST1:
if ( condition2 )
begin
    reset c1;
    goto ST2;
end
State ST2 :

```

```

if ( condition1 )
    increment c1;
else if (condition3 && c1 < 10)
    goto ST3;
else if ( condition3 && c1 >= 10)
    trigger;
ST3:
goto ST3;

```

## SignalTap II Scripting Support

You can run procedures and make settings described in this chapter in a Tcl script. You can also run some procedures at a command prompt. For detailed information about scripting command options, refer to the Quartus II Command-Line and Tcl API Help browser. To run the Help browser, type the following at the command prompt:

```
quartus_sh --qhelp
```

### Related Information

- [Tcl Scripting documentation](#)
- [Quartus II Tcl Scripting online help](#)

## Document Revision History

Table 13-14: Document Revision History

Date	Version	Changes Made
2014.12.15	14.1.0	Updated location of Fitter Settings, Analysis & Synthesis Settings, and Physical Synthesis Optimizations to Compiler Settings.
December 2014	14.1.0	<ul style="list-style-type: none"> <li>• Added MAX 10 as supported device.</li> <li>• Removed Full Incremental Compilation setting and Post-Fit (Strict) netlist type setting information.</li> <li>• Removed outdated GUI images from "Using Incremental Compilation with the SignalTap II Logic Analyzer" section.</li> </ul>
June 2014	14.0.0	<ul style="list-style-type: none"> <li>• DITA conversion.</li> <li>• Replaced MegaWizard Plug-In Manager and Megafunction content with IP Catalog and parameter editor content.</li> <li>• Added flows for custom trigger HDL object, Incremental Route with Rapid Recompile, and nested groups with Basic OR.</li> <li>• GUI changes: toolbar, drag to zoom, disable/enable instance, trigger log time-stamping.</li> </ul>

Date	Version	Changes Made
November 2013	13.1.0	Removed HardCopy material. Added section on using cross-triggering with DS-5 tool and added link to white paper 01198. Added section on remote debugging an Altera SoC and added link to application note 693. Updated support for MEX function.
May 2013	13.0.0	<ul style="list-style-type: none"> <li>Added recommendation to use the state-based flow for segmented buffers with separate trigger conditions, information about Basic OR trigger condition, and hard processor system (HPS) external triggers.</li> <li>Updated “Segmented Buffer” on page 13-17, Conditional Mode on page 13-21, Creating Basic Trigger Conditions on page 13-16, and Using External Triggers on page 13-48.</li> </ul>
June 2012	12.0.0	Updated Figure 13-5 on page 13-16 and “Adding Signals to the SignalTap II File” on page 13-10.
November 2011	11.0.1	<p>Template update.</p> <p>Minor editorial updates.</p>
May 2011	11.0.0	Updated the requirement for the standalone SignalTap II software.
December 2010	10.0.1	Changed to new document template.
July 2010	10.0.0	<ul style="list-style-type: none"> <li>Add new acquisition buffer content to the “View, Analyze, and Use Captured Data” section.</li> <li>Added script sample for generating hexadecimal CRC values in programmed devices.</li> <li>Created cross references to Quartus II Help for duplicated procedural content.</li> </ul>
November 2009	9.1.0	No change to content.
March 2009	9.0.0	<ul style="list-style-type: none"> <li>Updated Table 13-1</li> <li>Updated “Using Incremental Compilation with the SignalTap II Logic Analyzer” on page 13-45</li> <li>Added new Figure 13-33</li> <li>Made minor editorial updates</li> </ul>
November 2008	8.1.0	<p>Updated for the Quartus II software version 8.1 release:</p> <ul style="list-style-type: none"> <li>Added new section “Using the Storage Qualifier Feature” on page 14-25</li> <li>Added description of <code>start_store</code> and <code>stop_store</code> commands in section “Trigger Condition Flow Control” on page 14-36</li> <li>Added new section “Runtime Reconfigurable Options” on page 14-63</li> </ul>

Date	Version	Changes Made
May 2008	8.0.0	Updated for the Quartus II software version 8.0: <ul style="list-style-type: none"><li>• Added “Debugging Finite State machines” on page 14-24</li><li>• Documented various GUI usability enhancements, including improvements to the resource estimator, the bus find feature, and the dynamic display updates to the counter and flag resources in the State-based trigger flow control tab</li><li>• Added “Capturing Data Using Segmented Buffers” on page 14–16</li><li>• Added hyperlinks to referenced documents throughout the chapter</li><li>• Minor editorial updates</li></ul>

**Related Information**

[Quartus II Handbook Archive](#)

For previous versions of the Quartus II Handbook



# Debugging Single Event Upsets Using the Fault Injection Debugger 14

2014.06.30

QI15V3



Subscribe



Send Feedback

The Fault Injection Debugger is a Quartus II Programmer tool that uses the Altera Fault Injection IP core to inject, detect, and scrub soft errors in your device.

The Fault Injection Debugger is available for Stratix V family devices. For assistance with support for Arria V or Cyclone V family devices, file a service request using [mySupport](#).

The Fault Injection Debugger provides the following benefits:

- Allows you to perform single event functional interrupts (SEFI) characterization in-house. This eliminates the requirement for entire system beam testing. Instead, you can limit the beam testing to Failures in time (FIT)/Mb measurement at the device level.
- Scale FIT rates according to the SEFI characterization that is relevant to your design architecture.
- Optimize your design to reduce disruption caused by a single event upsets (SEU).

## Related Information

[Single Event Upsets](#)

## Single Event Upset Mitigation

Integrated circuits and programmable logic devices such as FPGAs in particular are susceptible to SEUs. SEU are random, nondestructive events, caused by two major sources: alpha particles and neutrons from cosmic rays. Radiation can cause either the logic register, embedded memory bit, or a configuration RAM (CRAM) bit to flip its state.

Arria V, Cyclone V, Stratix V and newer devices have the following CRAM capabilities:

- Error Detection Cyclical Redundance Checking (EDCRC)
- Automatic correction of an upset CRAM (scrubbing)
- Ability to create an upset CRAM condition (fault injection)

For more information about SEU mitigation in the Arria V, Cyclone V and Stratix V family devices, refer to the *SEU Mitigation* chapter in the respective device handbook.

## Related Information

[Single Event Upsets](#)

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



## Hardware and Software Requirements

The following hardware and software is required to use the Fault Injection Debugger:

- Quartus II software version 14.0 or later.
- License FEATURE line enabling the use of the Fault Injection Debugger. For more information, contact your local Altera sales representative.
- Download cable (USB-Blaster, USB-Blaster II, EthernetBlaster or EthernetBlaster II cable)
- Altera development kit or user design board with a JTAG connection to the device under test

### Related Information

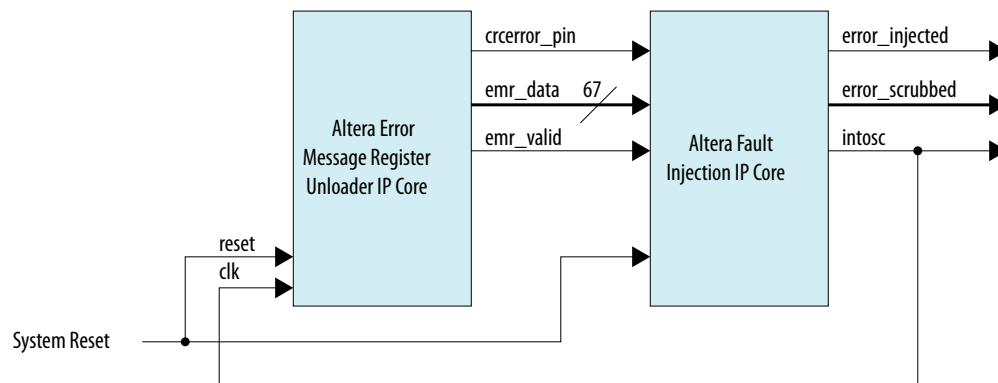
[Contact Altera](#)

## Using the Fault Injection Debugger

You instantiate the Altera Fault Injection IP core into your design to use the Fault Injection Debugger. You can use the Fault Injection Debugger in the Quartus II software GUI or at the command line interface (CLI) to inject and scrub errors from the FPGA's CRAM.

The Altera Fault Injection IP core is used with the Error Message Register (EMR) Unloader IP core.

**Figure 14-1: Example of Altera Fault Injection IP Core and EMR Unloader IP Core**



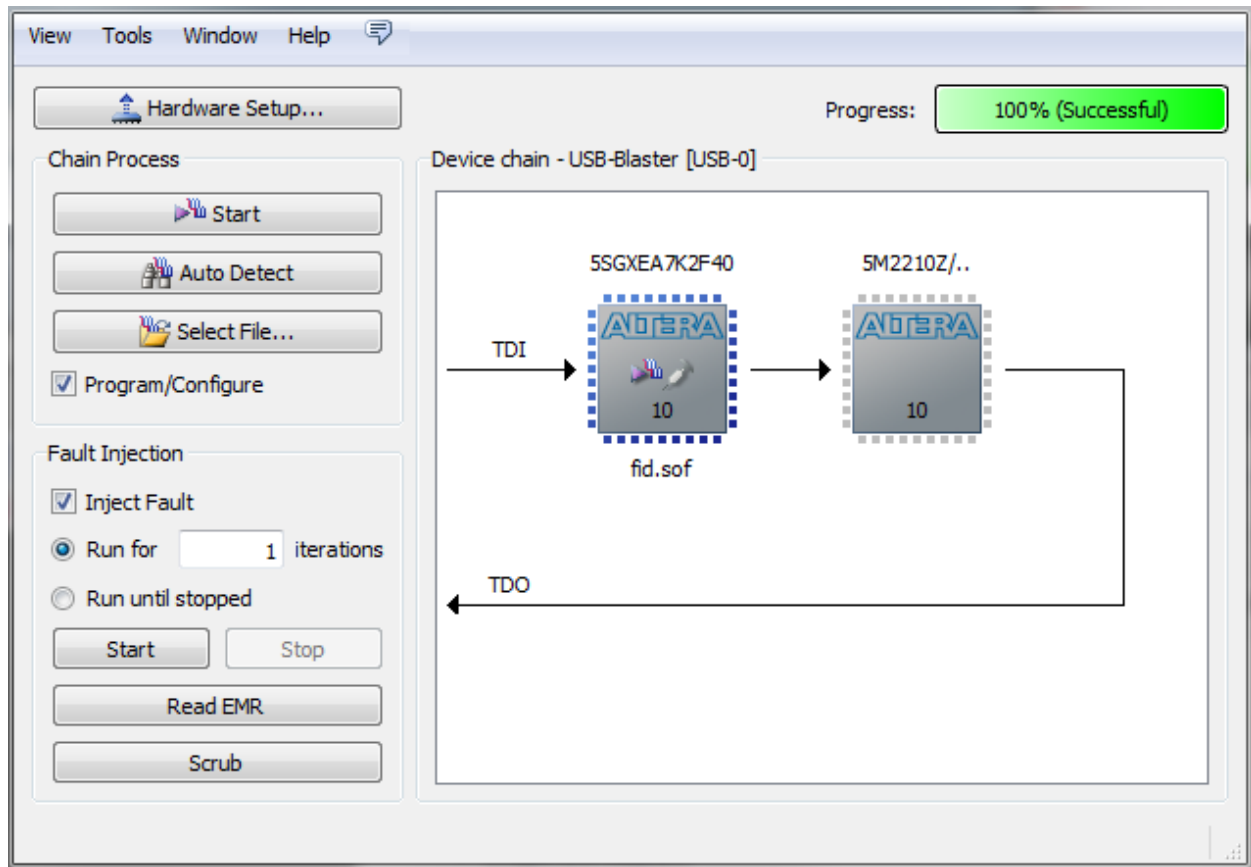
### Related Information

- [Development Kits, Daughter Cards and Programming Hardware Support](#)
- [Download Center](#)
- [AN 539: Test Methodology or Error Detection and Recovery using CRC in Altera FPGA Devices](#)

## Fault Injection Debugger GUI

To launch the Fault Injection Debugger, click **Tools > Fault Injection Debugger**. The GUI allows you to set up hardware, program and configure your device, and perform fault injection.

Figure 14-2: Fault Injection Debugger



## Injecting Errors and Performing Scrubbing with the Fault Injection Debugger

To set up, inject errors, and perform scrubbing with the Fault Injection Debugger in the Quartus II software, follow these steps:

1. Create or open a project that is targeted to an Arria V, Cyclone V, or Stratix V family device.
2. Use the IP Catalog and parameter editor to instantiate and specify your parameters for the Altera Fault Injection IP.
3. To compile your design, click **Processing > Start Compilation**.
4. Click **Tools > Fault Injection Debugger**.
5. To set up the programming hardware, click **Hardware Setup** and select the desired hardware.
6. To automatically add devices to the device chain to the programming list, click **Auto Detect**.
7. To add a programming file to the targeted device, click **Select File** and select the SRAM Object File (.sof) you want to open.

**Note:** After you add the .sof to the device, the **Program/Configure** and **Inject Fault** boxes are enabled.

8. Program the device with the Quartus II Programmer.
9. To inject errors into the device:

- a. Turn on **Inject Fault**.
  - b. Select **Run for <number> iterations** and type the number of iterations you want to run or select **Run until stopped**. You can enter up to 50 iterations.
  - c. Click **Start**.
  - d. Click **Stop** to manually stop injecting errors into the device.
10. To read the error message register, which displays information about detected soft errors in the Quartus II Messages window, click **Read EMR**.
11. To perform external scrubbing on the device, click **Scrub**.

#### Related Information

- [About Programming](#)
- [About Compilation](#)
- [Managing Quartus II Projects](#)

## Command-Line Interface

You can run the Fault Injection Debugger at the command line with the **quartus\_fid** executable.

**Table 14-1: Command line Arguments for Fault Injection**

Short Argument	Long Argument	Description
c	cable	Specify programming hardware or cable. (Required)
i	index	Specify the active device to inject fault. (Required)
n	number	Specify the number of errors to inject. The default value is 1. (Not required)
t	time	Interval time between injections. (Not required)

## Targeted Fault Injection Feature

The Fault Injection Debugger injects faults into the FPGA randomly. However, the Targeted Fault Injection feature allows you to inject faults into targeted locations in the CRAM. This may be useful, for example, if you noted an SEU event and want to test the FPGA or system response to the same event after modifying a recovery strategy.

**Note:** The Targeted Fault Injection feature is available only from the command line interface, not the GUI.

You can direct the Fault Injection Debugger to inject these types of errors within a CRAM frame:

- single Bit Errors (correctable with EDCRC scrubbing)
- adjacent Double-bit Errors (correctable with EDCRC scrubbing)
- multi-bit, uncorrectable errors

You can specify that errors are injected from the command line, or in prompt mode.

**Related Information**

[AN 539: Test Methodology or Error Detection and Recovery using CRC in Altera FPGA Devices](#)

**Specifying an Error List From the Command Line**

The Targeted Fault Injection feature allows you to specify an error list from the command line, as shown in the following example.

```
c:\Users\sng> quartus_fid -c 1 -i "@1= svgx.sof#i " -n 2 -user="@1= 0x2274 0x05EF 0x2264 0x0500 "
```

`c 1` indicates that the fpga is controlled by the first cable on your computer.

`i "@1= svgx.sof#i "` indicates that the first device in the chain is loaded with the object file `svgx.sof` and will be injected with faults.

`n 2` indicates that two faults will be injected.

`user="@1= 0x2274 0x05EF 0x2264 0x0500 "` is a user-specified list of faults to be injected: in this case, Device 1 has two faults at frame 0x2274, bit 0x05EF and at frame 0x2264, bit 0x0500.

**Specifying an Error List From Prompt Mode**

You can operate the Targeted Fault Injection feature interactively by specifying the number of faults to be 0 (-n 0). The Fault Injection Debugger presents prompt mode commands and their descriptions.

Prompt Mode Command	Description
F	to inject fault
E	to read EMR
S	to scrub errors
Q	to quit

In the prompt mode, the F command can be issued alone, resulting in a single fault in a random location in the device. In the following examples of the F command in prompt mode, three errors are injected.

```
F #3 0x12 0x34 0x56 0x78 * 0x9A 0xBC +
```

- Error 1 – Single Bit Error at frame 0x12, bit 0x34
- Error 2 – Uncorrectable Error at frame 0x56, bit 0x78 (\* indicates a multiple bit error)
- Error 3 – Double-adjacent Error at frame 0x9A, bit 0xBC (+ indicates a double bit error)

```
F 0x12 0x34 0x56 0x78 *
```

One (default) error is injected:

Error 1 – Single Bit Error at frame 0x12, bit 0x34. Locations after the first frame/bit location will be ignored.

```
F #3 0x12 0x34 0x56 0x78 * 0x9A 0xBC + 0xDE 0x00
```

Three errors are injected:

- Error 1 – Single Bit Error at frame 0x12, bit 0x34
- Error 2 – Uncorrectable Error at frame 0x56, bit 0x78
- Error 3 – Double-adjacent Error at frame 0x9A, bit 0xBC
- Locations after the first 3 frame/bit pairs will be ignored.

### Determining CRAM Bit Locations

When the Fault Injection Debugger detects a CRAM EDCRC error, the Error Message Register (EMR) contains the syndrome, frame number, bit location, and error type (single, double, multi-bit) of the detected CRAM error.

During system testing, you should save the EMR contents when an EDCRC fault is detected.

**Note:** With the recorded EMR contents, you can supply the frame and bit numbers to the Fault Injection Debugger program to re-play the errors noted during system testing, to further design and characterize a system recovery response to that error.

#### Related Information

[AN 539: Test Methodology or Error Detection and Recovery using CRC in Altera FPGA Devices](#)

## Altera Fault Injection IP Pin Description

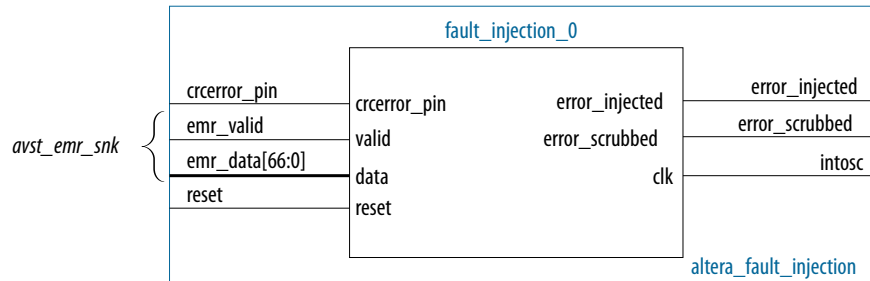
The Altera Fault Injection IP core includes the following I/O pins.

**Table 14-2: Altera Fault Injection IP Core I/O Pins**

Pin Name	Pin Direction	Pin Description
crcerror_pin	input	Input from Altera Error Message Register Unloader IP. This signal is asserted when a CRC error has been detected by the device's EDCRC.
emr_data	input	<p>Error Message Register (EMR) contents. For Arria V, Cyclone V, and Stratix V devices, the EMR fields are:</p> <p>EMR [66:35] - Syndrome (32 bits)</p> <p>EMR [34:19] - Frame Address (14 bits)</p> <p>EMR [18:5] - Byte Location (11 bits)</p> <p>EMR [1:0] - Error Type (2 bits)</p> <p>For other device families, see the appropriate device Handbook.</p> <p>This input complies with the Avalon Streaming data interface signal.</p>
emr_valid	input	Indicates the emr_data inputs contain valid data. This is an Avalon Streaming valid interface signal.

Pin Name	Pin Direction	Pin Description
Reset	input	Module reset input.
error_injected	output	Indicates an error was injected into CRAM as commanded via the JTAG interface.
error_scrubbed	output	Indicates the device scrubbing is complete as commanded via the JTAG interface.
intosc	output	Optional output. This is the clock used by the Altera Fault Injection IP. It can be used, for example, to clock the EMR_unloader block.

Figure 14-3: Altera Fault Injection IP Pin Diagram



## Document Revision History

Table 14-3: Document Revision History

Date	Version	Changes
June 2014	2014.06.30	<ul style="list-style-type: none"> <li>Removed “Modifying the Quartus INI File” section.</li> <li>Added “Targeted Fault Injection Feature” section.</li> <li>Updated “Hardware and Software Requirements” section.</li> </ul>
December 2012	2012.12.01	Preliminary release.

# In-System Debugging Using External Logic Analyzers 15

2014.06.30

QI15V3



Subscribe



Send Feedback

## About the Quartus II Logic Analyzer Interface

The Quartus II Logic Analyzer Interface (LAI) allows you to use an external logic analyzer and a minimal number of Altera-supported device I/O pins to examine the behavior of internal signals while your design is running at full speed on your Altera® - supported device.

The LAI connects a large set of internal device signals to a small number of output pins. You can connect these output pins to an external logic analyzer for debugging purposes. In the Quartus II LAI, the internal signals are grouped together, distributed to a user-configurable multiplexer, and then output to available I/O pins on your Altera-supported device. Instead of having a one-to-one relationship between internal signals and output pins, the Quartus II LAI enables you to map many internal signals to a smaller number of output pins. The exact number of internal signals that you can map to an output pin varies based on the multiplexer settings in the Quartus II LAI.

**Note:** The term “logic analyzer” when used in this document includes both logic analyzers and oscilloscopes equipped with digital channels, commonly referred to as mixed signal analyzers or MSOs.

The LAI does not support Hard Processor System (HPS) I/Os.

### Related Information

[Device Support website](#)

## Choosing a Logic Analyzer

The Quartus II software offers the following two general purpose on-chip debugging tools for debugging a large set of RTL signals from your design:

- The SignalTap® II Logic Analyzer
- An external logic analyzer, which connects to internal signals in your Altera-supported device by using the Quartus II LAI

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered





**Table 15-1: Comparing the SignalTap II Logic Analyzer with the Logic Analyzer Interface**

Feature	Description	Recommended Logic Analyzer
Sample Depth	You have access to a wider sample depth with an external logic analyzer. In the SignalTap II Logic Analyzer, the maximum sample depth is set to 128 Kb, which is a device constraint. However, with an external logic analyzer, there are no device constraints, providing you a wider sample depth.	LAI
Debugging Timing Issues	Using an external logic analyzer provides you with access to a “timing” mode, which enables you to debug combined streams of data.	LAI
Performance	You frequently have limited routing resources available to place and route when you use the SignalTap II Logic Analyzer with your design. An external logic analyzer adds minimal logic, which removes resource limits on place-and-route.	LAI
Triggering Capability	The SignalTap II Logic Analyzer offers triggering capabilities that are comparable to external logic analyzers.	LAI or SignalTap II
Use of Output Pins	Using the SignalTap II Logic Analyzer, no additional output pins are required. Using an external logic analyzer requires the use of additional output pins.	SignalTap II
Acquisition Speed	With the SignalTap II Logic Analyzer, you can acquire data at speeds of over 200 MHz. You can achieve the same acquisition speeds with an external logic analyzer; however, you must consider signal integrity issues.	SignalTap II

**Related Information**

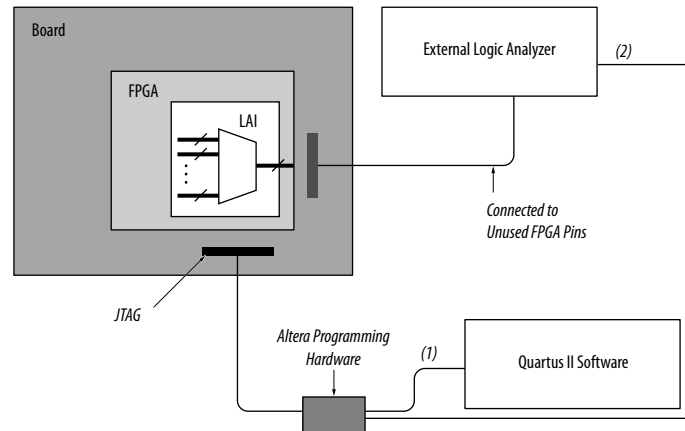
- [System Debugging Tools Overview](#) on page 9-1  
Overview and comparison of all tools available in the Quartus II software on-chip debugging tool suite

**Required Components**

You must have the following components to perform analysis using the LAI:

- The Quartus II software starting with version 5.1 and later
- The device under test
- An external logic analyzer
- An Altera communications cable
- A cable to connect the Altera-supported device to the external logic analyzer

**Figure 15-1: LAI and Hardware Setup**

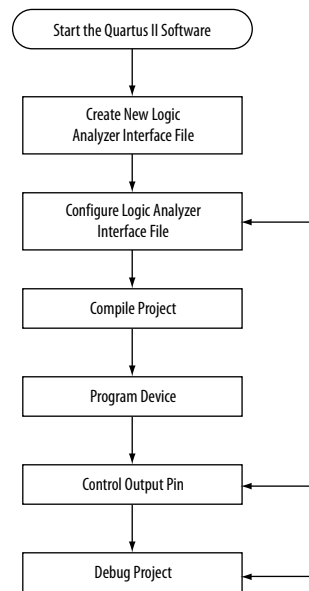


Notes to figure:

1. Configuration and control of the LAI using a computer loaded with the Quartus II software via the JTAG port.
2. Configuration and control of the LAI using a third-party vendor logic analyzer via the JTAG port. Support varies by vendor.

## Flow for Using the LAI

Figure 15-2: LAI Workflow



Notes to figure:

1. Configuration and control of the LAI using a computer loaded with the Quartus II software via the JTAG port.
2. Configuration and control of the LAI using a third-party vendor logic analyzer via the JTAG port. Support varies by vendor.

## Working with LAI Files

The **.lai** file stores the configuration of an LAI instance. The **.lai** file opens in the LAI editor. The editor allows you to group multiple internal signals to a set of external pins.

### Related Information

[Setting Up the Logic Analyzer Interface online help](#)

## Configuring the File Core Parameters

After you create the **.lai** file, you must configure the **.lai** file core parameters by clicking on the **Setup View** list, and then selecting **Core Parameters**. The table below lists the **.lai** file core parameters.

Table 15-2: LAI File Core Parameters

Parameter	Description
<b>Pin Count</b>	<p>The <b>Pin Count</b> parameter signifies the number of pins you want dedicated to your LAI. The pins must be connected to a debug header on your board. Within the Altera-supported device, each pin is mapped to a user-configurable number of internal signals.</p> <p>The <b>Pin Count</b> parameter can range from 1 to 255 pins.</p>
<b>Bank Count</b>	<p>The <b>Bank Count</b> parameter signifies the number of internal signals that you want to map to each pin. For example, a <b>Bank Count</b> of 8 implies that you will connect eight internal signals to each pin.</p> <p>The <b>Bank Count</b> parameter can range from 1 to 255 banks.</p>
<b>Output/Capture Mode</b>	<p>The <b>Output/Capture Mode</b> parameter signifies the type of acquisition you perform. There are two options that you can select:</p> <p><b>Combinational/Timing</b>—This acquisition uses your external logic analyzer’s internal clock to determine when to sample data. Because <b>Combinational/Timing</b> acquisition samples data asynchronously to your Altera-supported device, you must determine the sample frequency you should use to debug and verify your system. This mode is effective if you want to measure timing information, such as channel-to-channel skew. For more information about the sampling frequency and the speeds at which it can run, refer to the data sheet for your external logic analyzer.</p> <p><b>Registered/State</b>—This acquisition uses a signal from your system under test to determine when to sample. Because <b>Registered/State</b> acquisition samples data synchronously with your Altera-supported device, it provides you with a functional view of your Altera-supported device while it is running. This mode is effective when you verify the functionality of your design.</p>
<b>Clock</b>	<p>The <b>Clock</b> parameter is available only when <b>Output/Capture Mode</b> is set to <b>Registered State</b>. You must specify the sample clock in the <b>Core Parameters</b> view. The sample clock can be any signal in your design. However, for best results, Altera recommends that you use a clock with an operating frequency fast enough to sample the data you would like to acquire.</p>
<b>Power-Up State</b>	<p>The <b>Power-Up State</b> parameter specifies the power-up state of the pins you have designated for use with the LAI. You have the option of selecting tri-stated for all pins, or selecting a particular bank that you have enabled.</p>

## Mapping the LAI File Pins to Available I/O Pins

To configure the .jai file I/O pin parameters, select **Pins** in the **Setup View** list. To assign pin locations for the LAI, double-click the **Location** column next to the reserved pins in the **Name** column, and the Pin Planner opens.

### Related Information

[Managing Device I/O Pins documentation](#)

Information about how to use the Pin Planner

## Mapping Internal Signals to the LAI Banks

After you have specified the number of banks to use in the **Core Parameters** settings page, you must assign internal signals for each bank in the LAI. Click the **Setup View** arrow and select **Bank n** or **All Banks**.

To view all of your bank connections, click **Setup View** and select **All Banks**.

## Using the Node Finder

Before making bank assignments, on the View menu, point to **Utility Windows** and click **Node Finder**. Find the signals that you want to acquire, then drag and drop the signals from the **Node Finder** dialog box into the bank **Setup View**. When adding signals, use **SignalTap II: pre-synthesis** for non-incrementally routed instances and **SignalTap II: post-fitting** for incrementally routed instances.

As you continue to make assignments in the bank **Setup View**, the schematic of your LAI in the **Logical View** of your **.lai** file begins to reflect your assignments. Continue making assignments for each bank in the **Setup View** until you have added all of the internal signals for which you wish to acquire data.

## Compiling Your Quartus II Project

When you save your **.lai** file, a dialog box prompts you to enable the LAI instance for the active project. Alternatively, you can specify the **.lai** file your project uses in the **Global Project Settings** dialog box.

After you specify the name of your **.lai** file, you must compile your project. To compile your project, on the Processing menu, click **Start Compilation**.

To ensure that the LAI is properly compiled with your project, expand the entity hierarchy in the Project Navigator. (To display the Project Navigator, on the View menu, point to **Utility Windows** and click **Project Navigator**.) If the LAI is compiled with your design, the `sld_hub` and `sld_multitap` entities are shown in the Project Navigator.

Figure 15-3: Project Navigator

Entity	Logic Cells	LC Registers
Stratix: EP1S10B672C7		
test	136 (1)	81
sld_multitap:auto_lai_0	35 (11)	15
sld_hub:sld_hub_inst	100 (25)	65

## Programming Your Altera-Supported Device Using the LAI

After compilation completes, you must configure your Altera-supported device before using the LAI.

You can use the LAI with multiple devices in your JTAG chain. Your JTAG chain can also consist of devices that do not support the LAI or non-Altera, JTAG-compliant devices. To use the LAI in more than one Altera-supported device, create an **.lai** file and configure an **.lai** file for each Altera-supported device that you want to analyze.

### Related Information

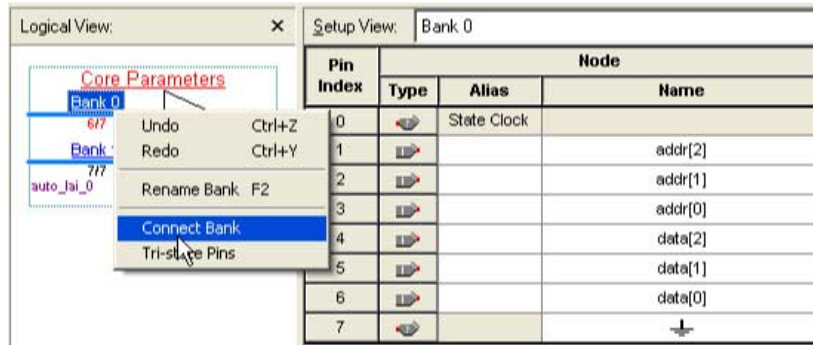
#### [Enabling the Logic Analyzer Interface online help](#)

Information to configure a device or a set of devices for use with LAI

## Controlling the Active Bank During Runtime

When you have programmed your Altera-supported device, you can control which bank you map to the reserved .jai file output pins. To control which bank you map, in the schematic in the Logical View, right-click the bank and click **Connect Bank**.

Figure 15-4: Configuring Banks



## Acquiring Data on Your Logic Analyzer

To acquire data on your logic analyzer, you must establish a connection between your device and the external logic analyzer. For more information about this process and for guidelines about how to establish connections between debugging headers and logic analyzers, refer to the documentation for your logic analyzer.

## Using the LAI with Incremental Compilation

The Incremental Compilation feature in the Quartus II software allows you to preserve the synthesis and fitting results of your design. This is an effective feature for reducing compilation times if you only modify a portion of a design or you wish to preserve the optimization results from a previous compilation.

The Incremental Compilation feature is well suited for use with LAI since LAI comprises a small portion of most designs. Because LAI consists of only a small portion of your design, incremental compilation helps to minimize your compilation time. Incremental compilation works best when you are only changing a small portion of your design. Incremental compilation yields an accurate representation of your design behavior when changing the .jai file through multiple compilations.

### Related Information

[Enabling the Logic Analyzer Interface online help](#)

## Document Revision History

Table 15-3: Document Revision History

Date	Version	Changes
June 2014	14.0.0	<ul style="list-style-type: none"> <li>• Dita conversion</li> <li>• Added limitation about HPS I/O support</li> </ul>
June 2012	12.0.0	Removed survey link
November 2011	10.1.1	Changed to new document template
December 2010	10.1.0	<ul style="list-style-type: none"> <li>• Minor editorial updates</li> <li>• Changed to new document template</li> </ul>
August 2010	10.0.1	Corrected links
July 2010	10.0.0	<ul style="list-style-type: none"> <li>• Created links to the Quartus II Help</li> <li>• Editorial updates</li> <li>• Removed Referenced Documents section</li> </ul>
November 2009	9.1.0	<ul style="list-style-type: none"> <li>• Removed references to APEX devices</li> <li>• Editorial updates</li> </ul>
March 2009	9.0.0	<ul style="list-style-type: none"> <li>• Minor editorial updates</li> <li>• Removed Figures 15-4, 15-5, and 15-11 from 8.1 version</li> </ul>
November 2008	8.1.0	Changed to 8-1/2 x 11 page size. No change to content
May 2008	8.0.0	<ul style="list-style-type: none"> <li>• Updated device support list on page 15-3</li> <li>• Added links to referenced documents throughout the chapter</li> <li>• Added “Referenced Documents”</li> <li>• Added reference to <i>Section V. In-System Debugging</i></li> <li>• Minor editorial updates</li> </ul>

### Related Information

#### [Quartus II Handbook Archive](#)

For previous versions of the Quartus II Handbook

# In-System Modification of Memory and Constants 16

2014.06.30

QI15V3



Subscribe



Send Feedback

## About the In-System Memory Content Editor

The Quartus® II In-System Memory Content Editor allows you to view and update memories and constants with the JTAG port connection.

The In-System Memory Content Editor allows access to dense and complex FPGA designs. When you program devices, you have read and write access to the memories and constants through the JTAG interface. You can then identify, test, and resolve issues with your design by testing changes to memory contents in the FPGA while your design is running.

When you use the In-System Memory Content Editor in conjunction with the SignalTap II Logic Analyzer, you can more easily view and debug your design in the hardware lab.

The ability to read data from memories and constants allows you to quickly identify the source of problems. The write capability allows you to bypass functional issues by writing expected data. For example, if a parity bit in your memory is incorrect, you can use the In-System Memory Content Editor to write the correct parity bit values into your RAM, allowing your system to continue functioning. You can also intentionally write incorrect parity bit values into your RAM to check the error handling functionality of your design.

### Related Information

[System Debugging Tools Overview](#) on page 9-1

Overview and comparison of all tools available in the Quartus II software on-chip debugging tool suite

[Design Debugging Using the SignalTap II Logic Analyzer documentation](#) on page 13-1

[Library of Parameterized Modules online help](#)

List of the types of memories and constants currently supported by the Quartus II software

## Design Flow Using the In-System Memory Content Editor

To use the In-System Memory Content Editor, perform the following steps:

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



1. Identify the memories and constants that you want to access.
2. Edit the memories and constants to be run-time modifiable.
3. Perform a full compilation.
4. Program your device.
5. Launch the In-System Memory Content Editor.

## Creating In-System Modifiable Memories and Constants

When you specify that a memory or constant is run-time modifiable, the Quartus II software changes the default implementation. A single-port RAM is converted to a dual-port RAM, and a constant is implemented in registers instead of look-up tables (LUTs). These changes enable run-time modification without changing the functionality of your design.

If you instantiate a memory or constant IP core directly with ports and parameters in VHDL or Verilog HDL, add or modify the `lpm_hint` parameter as follows:

In VHDL code, add the following:

```
lpm_hint => "ENABLE_RUNTIME_MOD = YES,  
           INSTANCE_NAME = <instantiation name>";
```

In Verilog HDL code, add the following:

```
defparam <megafunction instance name>.lpm_hint =  
    "ENABLE_RUNTIME_MOD = YES,  
    INSTANCE_NAME = <instantiation name>";
```

### Related Information

[Setting up the In-System Memory Content Editor online help](#)

## Running the In-System Memory Content Editor

The In-System Memory Content Editor has three separate panes: the **Instance Manager**, the **JTAG Chain Configuration**, and the **Hex Editor**.

The **Instance Manager** pane displays all available run-time modifiable memories and constants in your FPGA device. The **JTAG Chain Configuration** pane allows you to program your FPGA and select the Altera® device in the chain to update.

Using the In-System Memory Content Editor does not require that you open a project. The In-System Memory Content Editor retrieves all instances of run-time configurable memories and constants by scanning the JTAG chain and sending a query to the specific device selected in the **JTAG Chain Configuration** pane.

If you have more than one device with in-system configurable memories or constants in a JTAG chain, you can launch multiple In-System Memory Content Editors within the Quartus II software to access the memories and constants in each of the devices. Each In-System Memory Content Editor can access the in-system memories and constants in a single device.

## Instance Manager

When you scan the JTAG chain to update the **Instance Manager** pane, you can view a list of all run-time modifiable memories and constants in the design. The **Instance Manager** pane displays the Index, Instance, Status, Width, Depth, Type, and Mode of each element in the list.

You can read and write to in-system memory with the **Instance Manager** pane.

**Note:** In addition to the buttons available in the **Instance Manager** pane, you can read and write data by selecting commands from the Processing menu, or the right-click menu in the **Instance Manager** pane or **Hex Editor** pane.

The status of each instance is also displayed beside each entry in the **Instance Manager** pane. The status indicates if the instance is **Not running**, **Offloading data**, or **Updating data**. The health monitor provides information about the status of the editor.

The Quartus II software assigns a different index number to each in-system memory and constant to distinguish between multiple instances of the same memory or constant function. View the **In-System Memory Content Editor Settings** section of the Compilation Report to match an index number with the corresponding instance ID.

### Related Information

[Instance Manager Pane online help](#)

## Editing Data Displayed in the Hex Editor Pane

You can edit data read from your in-system memories and constants displayed in the **Hex Editor** pane by typing values directly into the editor or by importing memory files.

### Related Information

[Working with In-System Memory Content Editor Data online help](#)

## Importing and Exporting Memory Files

The In-System Memory Content Editor allows you to import and export data values for memories that have the In-System Updating feature enabled. Importing from a data file enables you to quickly load an entire memory image. Exporting to a data file enables you to save the contents of the memory for future use.

## Scripting Support

The In-System Memory Content Editor supports reading and writing of memory contents via a Tcl script or Tcl commands entered at a command prompt. For detailed information about scripting command options, refer to the Quartus II command-line and Tcl API Help browser.

To run the Help browser, type the following command at the command prompt:

```
quartus_sh --qhelp
```

The commonly used commands for the In-System Memory Content Editor are as follows:

- Reading from memory:

```
read_content_from_memory  
[-content_in_hex]  
-instance_index <instance index>  
-start_address <starting address>  
-word_count <word count>
```

- Writing to memory:

```
write_content_to_memory
```

- Saving memory contents to a file:

```
save_content_from_memory_to_file
```

- Updating memory contents from a file:

```
update_content_to_memory_from_file
```

#### Related Information

- [Tcl Scripting documentation](#)
- [Command-Line Scripting documentation](#)
- [API Functions for Tcl online help](#)

Descriptions of the command options and scripting examples

## Programming the Device with the In-System Memory Content Editor

If you make changes to your design, you can program the device from within the In-System Memory Content Editor.

#### Related Information

[Setting up the In-System Memory Content Editor online help](#)

## Example: Using the In-System Memory Content Editor with the SignalTap II Logic Analyzer

The following scenario describes how you can use the In-System Updating of Memory and Constants feature with the SignalTap II Logic Analyzer to efficiently debug your design. You can use the In-System Memory Content Editor and the SignalTap II Logic Analyzer simultaneously with the JTAG interface.

Scenario: After completing your FPGA design, you find that the characteristics of your FIR filter design are not as expected.

1. To locate the source of the problem, change all your FIR filter coefficients to be in-system modifiable and instantiate the SignalTap II Logic Analyzer.
2. Using the SignalTap II Logic Analyzer to tap and trigger on internal design nodes, you find the FIR filter to be functioning outside of the expected cutoff frequency.
3. Using the **In-System Memory Content Editor**, you check the correctness of the FIR filter coefficients. Upon reading each coefficient, you discover that one of the coefficients is incorrect.
4. Because your coefficients are in-system modifiable, you update the coefficients with the correct data with the **In-System Memory Content Editor**.

In this scenario, you can quickly locate the source of the problem using both the In-System Memory Content Editor and the SignalTap II Logic Analyzer. You can also verify the functionality of your device by changing the coefficient values before modifying the design source files.

You can also modify the coefficients with the In-System Memory Content Editor to vary the characteristics of the FIR filter, for example, filter attenuation, transition bandwidth, cut-off frequency, and windowing function.

## Document Revision History

**Table 16-1: Document Revision History**

Date	Version	Changes
June 2014	14.0.0	<ul style="list-style-type: none"> <li>• Dita conversion.</li> <li>• Removed references to megafunction and replaced with IP core.</li> </ul>
June 2012	12.0.0	Removed survey link.
November 2011	10.0.3	Template update.
December 2010	10.0.2	Changed to new document template. No change to content.
August 2010	10.0.1	Corrected links
July 2010	10.0.0	<ul style="list-style-type: none"> <li>• Inserted links to Quartus II Help</li> <li>• Removed Reference Documents section</li> </ul>
November 2009	9.1.0	<ul style="list-style-type: none"> <li>• Delete references to APEX devices</li> <li>• Style changes</li> </ul>
March 2009	9.0.0	No change to content
November 2008	8.1.0	Changed to 8-1/2 x 11 page size. No change to content.
May 2008	8.0.0	<ul style="list-style-type: none"> <li>• Added reference to Section V. In-System Debugging in volume 3 of the Quartus II Handbook on page 16-1</li> <li>• Removed references to the Mercury device, as it is now considered to be a “Mature” device</li> <li>• Added links to referenced documents throughout document</li> <li>• Minor editorial updates</li> </ul>

### Related Information

#### [Quartus II Handbook Archive](#)

For previous versions of the Quartus II Handbook

# Design Debugging Using In-System Sources and Probes 17

2014.06.30

QI15V3



Subscribe



Send Feedback

Traditional debugging techniques often involve using an external pattern generator to exercise the logic and a logic analyzer to study the output waveforms during run time. The SignalTap® II Logic Analyzer and SignalProbe allow you to read or “tap” internal logic signals during run time as a way to debug your logic design.

You can make the debugging cycle more efficient when you can drive any internal signal manually within your design, which allows you to perform the following actions:

- Force the occurrence of trigger conditions set up in the SignalTap II Logic Analyzer
- Create simple test vectors to exercise your design without using external test equipment
- Dynamically control run time control signals with the JTAG chain

The In-System Sources and Probes Editor in the Quartus II® software extends the portfolio of verification tools, and allows you to easily control any internal signal and provides you with a completely dynamic debugging environment. Coupled with either the SignalTap II Logic Analyzer or SignalProbe, the In-System Sources and Probes Editor gives you a powerful debugging environment in which to generate stimuli and solicit responses from your logic design.

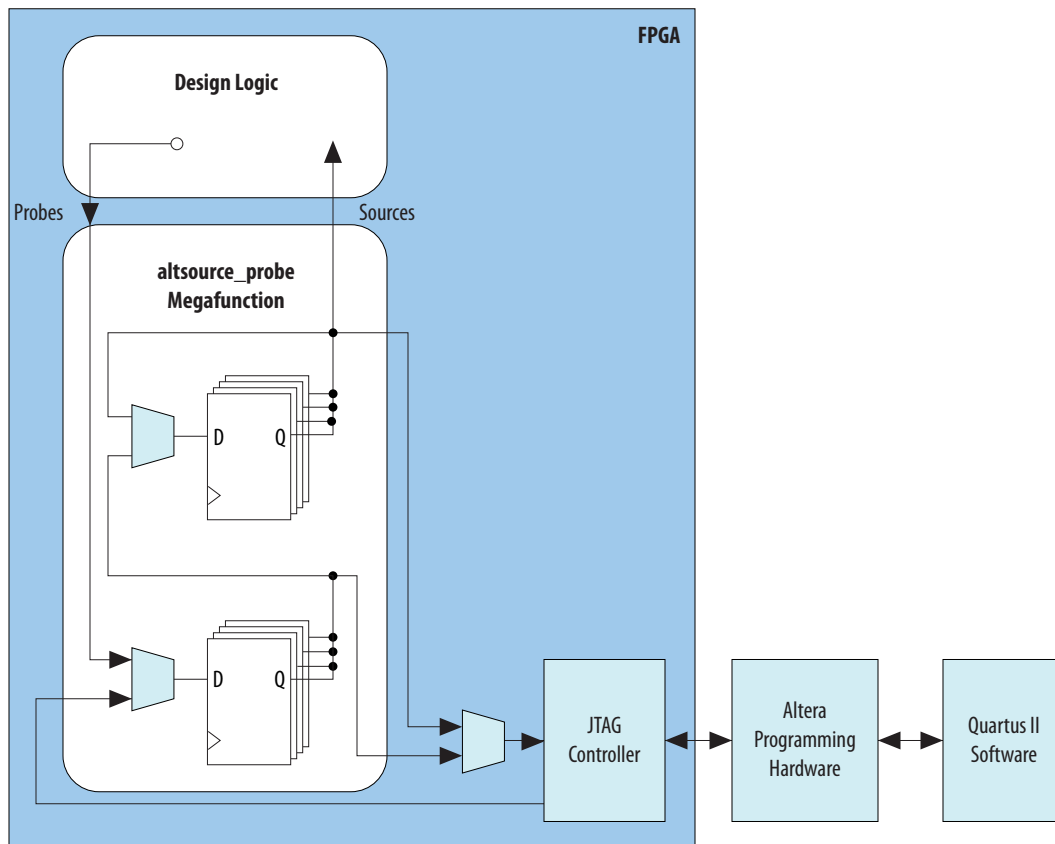
The Virtual JTAG IP core and the In-System Memory Content Editor also give you the capability to drive virtual inputs into your design. The Quartus II software offers a variety of on-chip debugging tools.

The In-System Sources and Probes Editor consists of the ALTSOURCE\_PROBE IP core and an interface to control the ALTSOURCE\_PROBE IP core instances during run time. Each ALTSOURCE\_PROBE IP core instance provides you with source output ports and probe input ports, where source ports drive selected signals and probe ports sample selected signals. When you compile your design, the ALTSOURCE\_PROBE IP core sets up a register chain to either drive or sample the selected nodes in your logic design. During run time, the In-System Sources and Probes Editor uses a JTAG connection to shift data to and from the ALTSOURCE\_PROBE IP core instances. The figure shows a block diagram of the components that make up the In-System Sources and Probes Editor.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered

Figure 17-1: In-System Sources and Probes Editor Block Diagram



The ALTSOURCE\_PROBE IP core hides the detailed transactions between the JTAG controller and the registers instrumented in your design to give you a basic building block for stimulating and probing your design. Additionally, the In-System Sources and Probes Editor provides single-cycle samples and single-cycle writes to selected logic nodes. You can use this feature to input simple virtual stimuli and to capture the current value on instrumented nodes. Because the In-System Sources and Probes Editor gives you access to logic nodes in your design, you can toggle the inputs of low-level components during the debugging process. If used in conjunction with the SignalTap II Logic Analyzer, you can force trigger conditions to help isolate your problem and shorten your debugging process.

The In-System Sources and Probes Editor allows you to easily implement control signals in your design as virtual stimuli. This feature can be especially helpful for prototyping your design, such as in the following operations:

- Creating virtual push buttons
- Creating a virtual front panel to interface with your design
- Emulating external sensor data
- Monitoring and changing run time constants on the fly

The In-System Sources and Probes Editor supports Tcl commands that interface with all your ALTSOURCE\_PROBE IP core instances to increase the level of automation.

### Related Information

#### [System Debugging Tools](#)

For an overview and comparison of all the tools available in the Quartus II software on-chip debugging tool suite

## Hardware and Software Requirements

The following components are required to use the In-System Sources and Probes Editor:

- Quartus II software

or

- Quartus II Web Edition (with the TalkBack feature turned on)
- Download Cable (USB-Blaster™ download cable or ByteBlaster™ cable)
- Altera® development kit or user design board with a JTAG connection to device under test

The In-System Sources and Probes Editor supports the following device families:

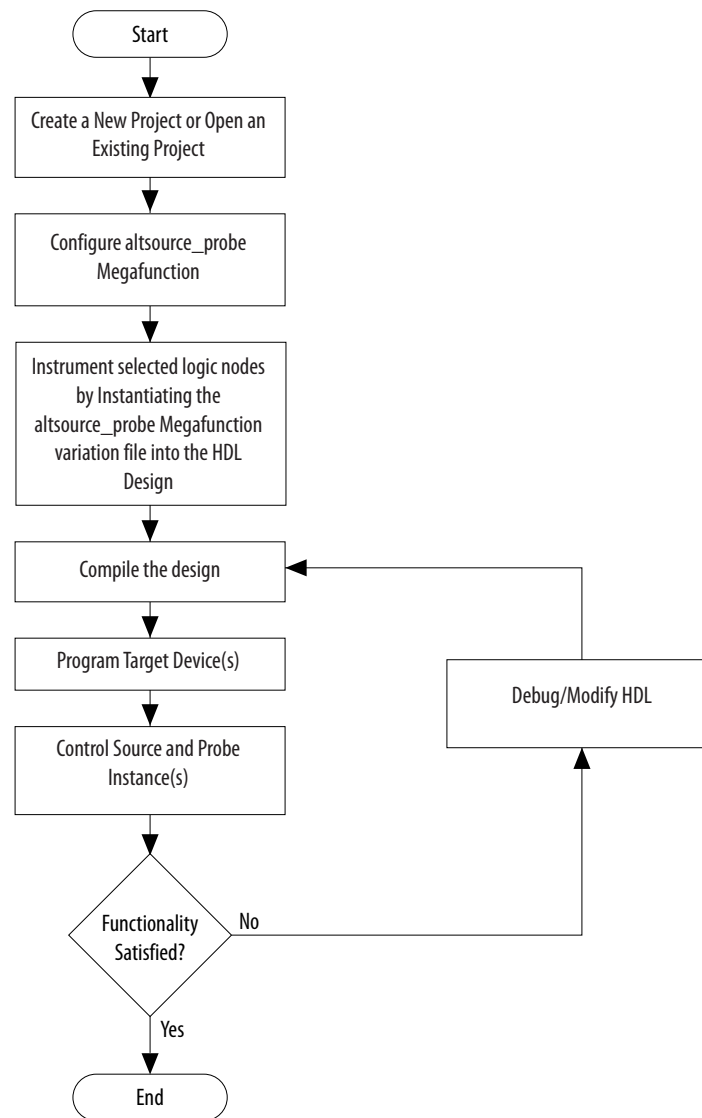
- Arria® series
- Stratix® series
- Cyclone® series
- MAX® series

## Design Flow Using the In-System Sources and Probes Editor

The In-System Sources and Probes Editor supports an RTL flow. Signals that you want to view in the In-System Sources and Probes editor are connected to an instance of the In-System Sources and Probes IP core.

After you compile the design, you can control each instance via the **In-System Sources and Probes Editor** pane or via a Tcl interface.

Figure 17-2: FPGA Design Flow Using the In-System Sources and Probes Editor



## Instantiating the In-System Sources and Probes IP Core

You must instantiate the In-System Sources and Probes IP core before you can use the In-System Sources and Probes editor. Use the IP Catalog and parameter editor to instantiate a custom variation of the In-System Sources and Probes IP core.



To configure the In-System Sources and Probes IP core, perform the following steps::

1. On the Tools menu, click **Tools > IP Catalog**.
2. Locate and double-click the In-System Sources and Probes IP core. The parameter editor appears.
3. Specify a name for your custom IP variation.
4. Specify the desired parameters for your custom IP variation. You can specify up to up to 256 bits for each source. Your design may include up to 128 instances of this IP core.
5. Click **Generate** or **Finish** to generate IP core synthesis and simulation files matching your specifications. The parameter editor generates the necessary variation files and the instantiation template based on your specification. Use the generated template to instantiate the In-System Sources and Probes IP core in your design.

**Note:** The In-System Sources and Probes Editor does not support simulation. You must remove the In-System Sources and Probes IP core before you create a simulation netlist.

## In-System Sources and Probes IP Core Parameters

Use the template to instantiate the variation file in your design.

**Table 17-1: In-System Sources and Probes IP Port Information**

Port Name	Required?	Direction	Comments
probe[]	No	Input	The outputs from your design.
source_clk	No	Input	Source Data is written synchronously to this clock. This input is required if you turn on <b>Source Clock</b> in the <b>Advanced Options</b> box in the parameter editor.
source_ena	No	Input	Clock enable signal for source_clk. This input is required if specified in the <b>Advanced Options</b> box in the parameter editor.
source[]	No	Output	Used to drive inputs to user design.

You can include up to 128 instances of the in-system sources and probes IP core in your design, if your device has available resources. Each instance of the IP core uses a pair of registers per signal for the width of the widest port in the IP core. Additionally, there is some fixed overhead logic to accommodate communication between the IP core instances and the JTAG controller. You can also specify an additional pair of registers per source port for synchronization.

When you compile your design that includes the In-System Sources and Probes IP core, the In-System Sources and Probes and SLD Hub Controller IP core are added to your compilation hierarchy automatically. These IP cores provide communication between the JTAG controller and your instrumented logic.

You can modify the number of connections to your design by editing the In-System Sources and Probes IP core. To open the design instance you want to modify in the parameter editor, double-click the instance in the Project Navigator. You can then modify the connections in the HDL source file. You must recompile your design after you make changes.

You can use the Quartus II incremental compilation feature to reduce compilation time. Incremental compilation allows you to organize your design into logical partitions. During recompilation of a design, incremental compilation preserves the compilation results and performance of unchanged partitions and reduces design iteration time by compiling only modified design partitions.

## Compiling the Design

When you compile your design that includes the In-System Sources and ProbesIP core, the In-System Sources and Probes and SLD Hub Controller IP core are added to your compilation hierarchy automatically. These IP cores provide communication between the JTAG controller and your instrumented logic.

You can modify the number of connections to your design by editing the In-System Sources and Probes IP core. To open the design instance you want to modify in the parameter editor, double-click the instance in the Project Navigator. You can then modify the connections in the HDL source file. You must recompile your design after you make changes.

You can use the Quartus II incremental compilation feature to reduce compilation design into logical partitions. During recompilation of a design, incremental compilation preserves the compilation results and performance of unchanged partitions and reduces design iteration time by compiling only modified design partitions.

### Related Information

[Quartus II Incremental Compilation for Hierarchical and Team-Based Design](#)

## Running the In-System Sources and Probes Editor

The In-System Sources and Probes Editor gives you control over all ALTSOURCE\_PROBE IP core instances within your design. The editor allows you to view all available run time controllable instances of the ALTSOURCE\_PROBE IP core in your design, provides a push-button interface to drive all your source nodes, and provides a logging feature to store your probe and source data.

To run the In-System Sources and Probes Editor:

- On the **Tools** menu, click **In-System Sources and Probes Editor**.

## In-System Sources and Probes Editor GUI

The In-System Sources and Probes Editor contains three panes:

- **JTAG Chain Configuration**—Allows you to specify programming hardware, device, and file settings that the In-System Sources and Probes Editor uses to program and acquire data from a device.
- **Instance Manager**—Displays information about the instances generated when you compile a design, and allows you to control data that the In-System Sources and Probes Editor acquires.
- **In-System Sources and Probes Editor**—Logs all data read from the selected instance and allows you to modify source data that is written to your device.

When you use the In-System Sources and Probes Editor, you do not need to open a Quartus II software project. The In-System Sources and Probes Editor retrieves all instances of the ALTSOURCE\_PROBE IP core by scanning the JTAG chain and sending a query to the device selected in the **JTAG Chain Configuration** pane. You can also use a previously saved configuration to run the In-System Sources and Probes Editor.

Each **In-System Sources and Probes Editor** pane can access the ALTSOURCE\_PROBE IP core instances in a single device. If you have more than one device containing IP core instances in a JTAG chain, you can launch multiple **In-System Sources and Probes Editor** panes to access the IP core instances in each device.

## Programming Your Device With JTAG Chain Configuration

After you compile your project, you must configure your FPGA before you use the In-System Sources and Probes Editor.

To configure a device to use with the In-System Sources and Probes Editor, perform the following steps:

1. Open the In-System Sources and Probes Editor.
2. In the **JTAG Chain Configuration** pane, point to **Hardware**, and then select the hardware communications device. You may be prompted to configure your hardware; in this case, click **Setup**.
3. From the **Device** list, select the FPGA device to which you want to download the design (the device may be automatically detected). You may need to click **Scan Chain** to detect your target device.
4. In the **JTAG Chain Configuration** pane, click to browse for the SRAM Object File (.sof) that includes the In-System Sources and Probes instance or instances. (The .sof may be automatically detected).
5. Click **Program Device** to program the target device.

## Instance Manager

The **Instance Manager** pane provides a list of all ALTSOURCE\_PROBE instances in the design and allows you to configure how data is acquired from or written to those instances.

The following buttons and sub-panes are provided in the **Instance Manager** pane:

- **Read Probe Data**—Samples the probe data in the selected instance and displays the probe data in the **In-System Sources and Probes Editor** pane.
- **Continuously Read Probe Data**—Continuously samples the probe data of the selected instance and displays the probe data in the **In-System Sources and Probes Editor** pane; you can modify the sample rate via the **Probe read interval** setting.
- **Stop Continuously Reading Probe Data**—Cancels continuous sampling of the probe of the selected instance.
- **Write Source Data**—Writes data to all source nodes of the selected instance.
- **Probe Read Interval**—Displays the sample interval of all the In-System Sources and Probe instances in your design; you can modify the sample interval by clicking **Manual**.
- **Event Log**—Controls the event log in the **In-System Sources and Probes Editor** pane.
- **Write Source Data**—Allows you to manually or continuously write data to the system.

The status of each instance is also displayed beside each entry in the **Instance Manager** pane. The status indicates if the instance is **Not running Offloading data**, **Updating data**, or if an **Unexpected JTAG communication error** occurs. This status indicator provides information about the sources and probes instances in your design.

## In-System Sources and Probes Editor Pane

The **In-System Sources and Probes Editor** pane allows you to view data from all sources and probes in your design.

The data is organized according to the index number of the instance. The editor provides an easy way to manage your signals, and allows you to rename signals or group them into buses. All data collected from in-system source and probe nodes is recorded in the event log and you can view the data as a timing diagram.

## Reading Probe Data

You can read data by selecting the ALTSOURCE\_PROBE instance in the **Instance Manager** pane and clicking **Read Probe Data**.

This action produces a single sample of the probe data and updates the data column of the selected index in the **In-System Sources and Probes Editor** pane. You can save the data to an event log by turning on the **Save data to event log** option in the **Instance Manager** pane.

If you want to sample data from your probe instance continuously, in the **Instance Manager** pane, click the instance you want to read, and then click **Continuously read probe data**. While reading, the status of the active instance shows **Unloading**. You can read continuously from multiple instances.

You can access read data with the shortcut menus in the **Instance Manager** pane.

To adjust the probe read interval, in the **Instance Manager** pane, turn on the **Manual** option in the **Probe read interval** sub-pane, and specify the sample rate in the text field next to the **Manual** option. The maximum sample rate depends on your computer setup. The actual sample rate is shown in the **Current interval** box. You can adjust the event log window buffer size in the **Maximum Size** box.

## Writing Data

To modify the source data you want to write into the ALTSOURCE\_PROBE instance, click the name field of the signal you want to change. For buses of signals, you can double-click the data field and type the value you want to drive out to the ALTSOURCE\_PROBE instance. The In-System Sources and Probes Editor stores the modified source data values in a temporary buffer.

Modified values that are not written out to the ALTSOURCE\_PROBE instances appear in red. To update the ALTSOURCE\_PROBE instance, highlight the instance in the **Instance Manager** pane and click **Write source data**. The **Write source data** function is also available via the shortcut menus in the **Instance Manager** pane.

The In-System Sources and Probes Editor provides the option to continuously update each ALTSOURCE\_PROBE instance. Continuous updating allows any modifications you make to the source data buffer to also write immediately to the ALTSOURCE\_PROBE instances. To continuously update the ALTSOURCE\_PROBE instances, change the **Write source data** field from **Manually** to **Continuously**.

## Organizing Data

The **In-System Sources and Probes Editor** pane allows you to group signals into buses, and also allows you to modify the display options of the data buffer.

To create a group of signals, select the node names you want to group, right-click and select **Group**. You can modify the display format in the Bus Display Format and the Bus Bit order shortcut menus.

The **In-System Sources and Probes Editor** pane allows you to rename any signal. To rename a signal, double-click the name of the signal and type the new name.

The event log contains a record of the most recent samples. The buffer size is adjustable up to 128k samples. The time stamp for each sample is logged and is displayed above the event log of the active instance as you move your pointer over the data samples.

You can save the changes that you make and the recorded data to a Sources and Probes File (**.spf**). To save changes, on the File menu, click **Save**. The file contains all the modifications you made to the signal groups, as well as the current data event log.

## Tcl interface for the In-System Sources and Probes Editor

To support automation, the In-System Sources and Probes Editor supports the procedures described in this chapter in the form of Tcl commands. The Tcl package for the In-System Sources and Probes Editor is included by default when you run **quartus\_stp**.

The Tcl interface for the In-System Sources and Probes Editor provides a powerful platform to help you debug your design. The Tcl interface is especially helpful for debugging designs that require toggling multiple sets of control inputs. You can combine multiple commands with a Tcl script to define a custom command set.


**Table 17-2: In-System Sources and Probes Tcl Commands**

Command	Argument	Description
start_insystem_source_probe	-device_name <device name> -hardware_name <hardware name>	Opens a handle to a device with the specified hardware.  Call this command before starting any transactions.
get_insystem_source_probe_instance_info	-device_name <device name> -hardware_name <hardware name>	Returns a list of all ALTSOURCE_PROBE instances in your design. Each record returned is in the following format:  {<instance Index>, <source width>, <probe width>, <instance name>}
read_probe_data	-instance_index <instance_index> -value_in_hex (optional)	Retrieves the current value of the probe.  A string is returned that specifies the status of each probe, with the MSB as the left-most bit.
read_source_data	-instance_index <instance_index> -value_in_hex (optional)	Retrieves the current value of the sources.  A string is returned that specifies the status of each source, with the MSB as the left-most bit.
write_source_data	-instance_index <instance_index> -value <value> -value_in_hex (optional)	Sets the value of the sources.  A binary string is sent to the source ports, with the MSB as the left-most bit.
end_interactive_probe	None	Releases the JTAG chain.  Issue this command when all transactions are finished.

The example shows an excerpt from a Tcl script with procedures that control the ALTSOURCE\_PROBE instances of the design as shown in the figure below. The example design contains a DCFIFO with ALTSOURCE\_PROBE instances to read from and write to the DCFIFO. A set of control muxes are added to the design to control the flow of data to the DCFIFO between the input pins and the ALTSOURCE\_PROBE instances. A pulse generator is added to the read request and write request control lines to guarantee a single sample read or write. The ALTSOURCE\_PROBE instances, when used with the script in the example below, provide visibility into the contents of the FIFO by performing single sample write and read operations and reporting the state of the full and empty status flags.

Use the Tcl script in debugging situations to either empty or preload the FIFO in your design. For example, you can use this feature to preload the FIFO to match a trigger condition you have set up within the SignalTap II Logic Analyzer.

Figure 17-3: DCFIFO Example Design Controlled by Tcl Script



```
## Setup USB hardware - assumes only USB Blaster is installed and
## an FPGA is the only device in the JTAG chain
set usb [lindex [get_hardware_names] 0]
set device_name [lindex [get_device_names -hardware_name $usb] 0]
## write procedure : argument value is integer
proc write {value} {
  global device_name usb
  variable full
  start_insystem_source_probe -device_name $device_name -hardware_name $usb
  #read full flag
```

```
set full [read_probe_data -instance_index 0]
if {$full == 1} {end_insystem_source_probe
return "Write Buffer Full"
}
##toggle select line, drive value onto port, toggle enable
##bits 7:0 of instance 0 is S_data[7:0]; bit 8 = S_write_req;
##bit 9 = Source_write_sel
##int2bits is custom procedure that returns a bitstring from an integer
## argument
write_source_data -instance_index 0 -value /[int2bits [expr 0x200 | $value]]
write_source_data -instance_index 0 -value [int2bits [expr 0x300 | $value]]
##clear transaction
write_source_data -instance_index 0 -value 0
end_insystem_source_probe
}
proc read {} {
global device_name usb
variable empty
start_insystem_source_probe -device_name $device_name -hardware_name $usb
##read empty flag : probe port[7:0] reads FIFO output; bit 8 reads empty_flag
set empty [read_probe_data -instance_index 1]
if {[regexp {1.....} $empty]} { end_insystem_source_probe
return "FIFO empty" }
## toggle select line for read transaction
## Source_read_sel = bit 0; s_read_reg = bit 1
## pulse read enable on DC FIFO
write_source_data -instance_index 1 -value 0x1 -value_in_hex
write_source_data -instance_index 1 -value 0x3 -value_in_hex
set x [read_probe_data -instance_index 1 ]
end_insystem_source_probe
return $x
}
```

#### Related Information


- [Tcl Scripting](#)
- [Quartus II Settings File Manual](#)
- [Command Line Scripting](#)

## Design Example: Dynamic PLL Reconfiguration

The In-System Sources and Probes Editor can help you create a virtual front panel during the prototyping phase of your design. You can create relatively simple, high functioning designs of in a short amount of time. The following PLL reconfiguration example demonstrates how to use the In-System Sources and Probes Editor to provide a GUI to dynamically reconfigure a Stratix PLL.

Stratix PLLs allow you to dynamically update PLL coefficients during run time. Each enhanced PLL within the Stratix device contains a register chain that allows you to modify the pre-scale counters (m and n values), output divide counters, and delay counters. In addition, the ALTPLL\_RECONFIG IP core provides an easy interface to access the register chain counters. The ALTPLL\_RECONFIG IP core provides a cache that contains all modifiable PLL parameters. After you update all the PLL parameters in the cache, the ALTPLL\_RECONFIG IP core drives the PLL register chain to update the PLL with the updated parameters. The figure shows a Stratix-enhanced PLL with reconfigurable coefficients.


Figure 17-4: Stratix-Enhanced PLL with Reconfigurable Coefficients



The following design example uses an `ALTSOURCE_PROBE` instance to update the PLL parameters in the `ALTPLL_RECONFIG` IP core cache. The `ALTPLL_RECONFIG` IP core connects to an enhanced PLL in a Stratix FPGA to drive the register chain containing the PLL reconfigurable coefficients. This design example uses a Tcl/Tk script to generate a GUI where you can enter in new  $m$  and  $n$  values for the enhanced PLL. The Tcl script extracts the  $m$  and  $n$  values from the GUI, shifts the values out to the `ALTSOURCE_PROBE` instances to update the values in the `ALTPLL_RECONFIG` IP core cache, and asserts the reconfiguration signal on the `ALTPLL_RECONFIG` IP core. The reconfiguration signal on the `ALTPLL_RECONFIG` IP core starts the register chain transaction to update all PLL reconfigurable coefficients.




Figure 17-5: Block Diagram of Dynamic PLL Reconfiguration Design Example



This design example was created using a Nios<sup>®</sup> II Development Kit, Stratix Edition. The file **sourceprobe\_DE\_dynamic\_pll.zip** contains all the necessary files for running this design example, including the following:

- **Readme.txt**—A text file that describes the files contained in the design example and provides instructions about running the Tk GUI shown in the figure below.
- **Interactive\_Reconfig.qar**—The archived Quartus II project for this design example.

Figure 17-6: Interactive PLL Reconfiguration GUI Created with Tk and In-System Sources and Probes Tcl Package



### Related Information

#### On-chip Debugging Design Examples

to download the In-System Sources and Probes Example

## Document Revision History

Table 17-3: Document Revision History

Date	Version	Changes
June 2014	14.0.0	Updated formatting.
June 2012	12.0.0	Removed survey link.

Date	Version	Changes
November 2011	10.1.1	Template update.
December 2010	10.1.0	Minor corrections. Changed to new document template.
July 2010	10.0.0	Minor corrections.
November 2009	9.1.0	<ul style="list-style-type: none"> <li>• Removed references to obsolete devices.</li> <li>• Style changes.</li> </ul>
March 2009	9.0.0	No change to content.
November 2008	8.1.0	Changed to 8-1/2 x 11 page size. No change to content.
May 2008	8.0.0	<ul style="list-style-type: none"> <li>• Documented that this feature does not support simulation on page 17–5</li> <li>• Updated Figure 17–8 for Interactive PLL reconfiguration manager</li> <li>• Added hyperlinks to referenced documents throughout the chapter</li> <li>• Minor editorial updates</li> </ul>

**Related Information****[Quartus II Handbook Archive](#)**

For previous versions of the Quartus II Handbook

2014.12.15

QI15V3



Subscribe



Send Feedback

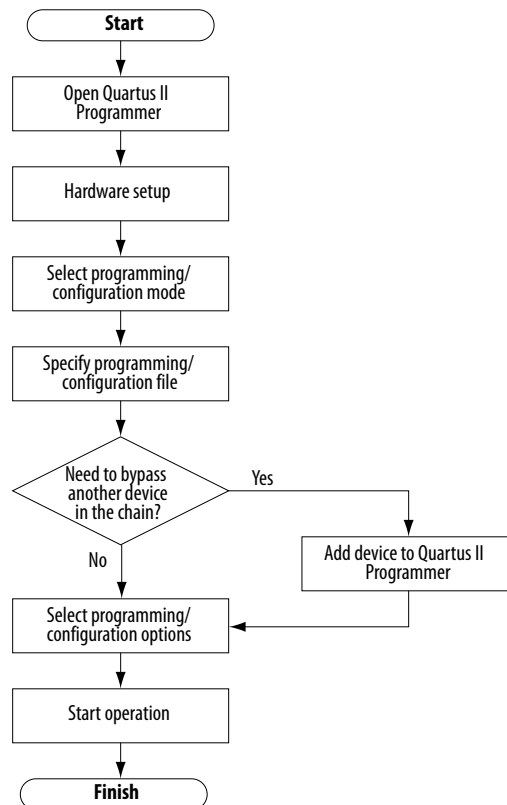
The Quartus II Programmer allows you to program and configure Altera® CPLD, FPGA, and configuration devices. After compiling your design, use the Quartus II Programmer to program or configure your device, to test the functionality of the design on a circuit board.

## Related Information

- [Programming Devices](#)

## Programming Flow

Figure 18-1: Programming Flow



© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at [www.altera.com/common/legal.html](http://www.altera.com/common/legal.html). Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO  
9001:2008  
Registered



The following steps describe the programming flow:

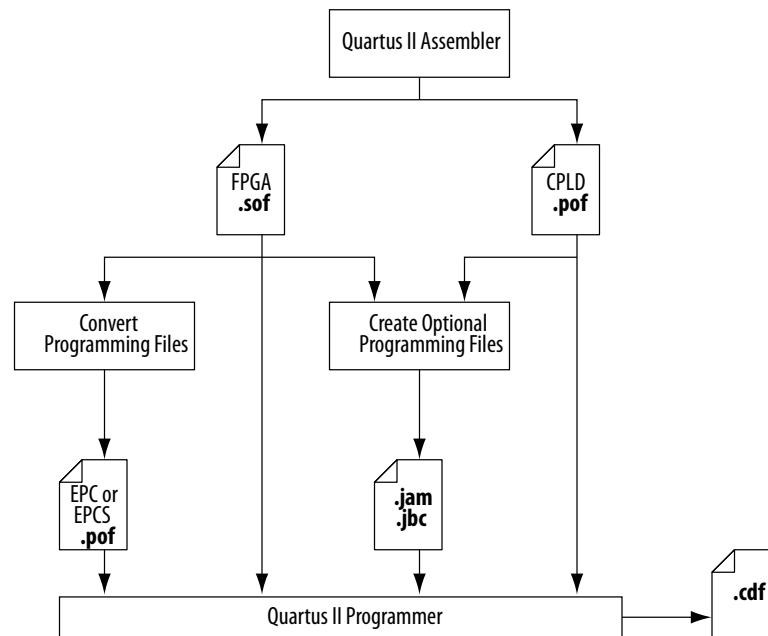
1. Compile your design, such that the Quartus II Assembler generates the programming or configuration file.
2. Convert the programming or configuration file to target your configuration device and, optionally, create secondary programming files.

**Table 18-1: Programming and Configuration File Format**

File Format	FPGA	CPLD	Configuration Device	Serial Configuration Device
SRAM Object File (.sof)	Yes	—	—	—
Programmer Object File (.pof)	—	Yes	Yes	Yes
JEDEC JESD71 STAPL Format File (.jam)	Yes	Yes	Yes	—
Jam Byte Code File (.jbc)	Yes	Yes	Yes	—

3. Program and configure the FPGA, CPLD, or configuration device using the programming or configuration file with the Quartus II Programmer.

**Figure 18-2: Programming File Generation Flow**



#### Related Information

- [About Programming](#)  
Provides information about Chain Description Files (.cdf).

## Optional Programming or Configuration Files

The Quartus II software can generate optional programming or configuration files in various formats that you can use with programming tools other than the Quartus II Programmer. When you compile a design in the Quartus II software, the Assembler automatically generates either a **.sof** or **.pof**. The Assembler also allows you to convert FPGA configuration files to programming files for configuration devices.

### Related Information

- [About Optional Programming Files](#)
- [AN 425: Using Command-Line Jam STAPL Solution for Device Programming](#)

Describes how to use the **.jam** and **.jbc** programming files with the Jam STAPL Player, Jam STAPL Byte-Code Player, and the `quartus_jli` command-line executable.

## Secondary Programming Files

The Quartus II software generates programming files in various formats for use with different programming tools.

**Table 18-2: File Types Generated by the Quartus II Software and Supported by the Quartus II Programmer**

File Type	Generated by the Quartus II Software	Supported by the Quartus II Programmer
<b>.sof</b>	Yes	Yes
<b>.pof</b>	Yes	Yes
<b>.jam</b>	Yes	Yes
<b>.jbc</b>	Yes	Yes
JTAG Indirect Configuration File ( <b>.jic</b> )	Yes	Yes
Serial Vector Format File ( <b>.svf</b> )	Yes	—
In System Configuration File ( <b>.isc</b> )	Yes	—
Hexadecimal (Intel-Format) Output File ( <b>.hexout</b> )	Yes	—
Raw Binary File ( <b>.rbf</b> )	Yes	Yes <sup>(9)</sup>
Raw Binary File for Partial Reconfiguration ( <b>.rbf</b> )	Yes	Yes <sup>(10)</sup>
Tabular Text File ( <b>.ttf</b> )	Yes	—
Raw Programming Data File ( <b>.rpd</b> )	Yes	—

<sup>(9)</sup> Raw Binary File (**.rbf**) is supported by the Quartus II Programmer in Passive Serial (PS) configuration mode.

<sup>(10)</sup> Raw Binary File for Partial Reconfiguration (**.rbf**) is supported by the Quartus II Programmer in JTAG debug mode.

### Related Information

- [Generating Secondary Programming Files](#)

## Quartus II Programmer GUI

The Quartus II Programmer GUI is a window that allows you to perform the following tasks:

- Adding your programming and configuration files.
- Specifying programming options and hardware.
- Starting the programming or configuration of the device.

To open the Programmer window, on the Tools menu, click **Programmer**. As you proceed through the programming flow, the Quartus II Message window reports the status of each operation.

### Related Information

- [Programmer Window](#)  
Describes the Programmer window.
- [Programmer Page \(Options Dialog Box\)](#)  
Describes the options in the **Tools** menu.

## Editing the Device Details of an Unknown Device

If the Quartus II Programmer automatically detects devices with shared JTAG IDs, the Programmer prompts you to specify the correct device in the JTAG chain.

If the Programmer does not prompt you to specify the correct device in the JTAG chain, then you must add a user defined device in the Quartus II software for each unknown device in the JTAG chain and specify the instruction register length for each device.

To edit the device details of an unknown device, follow these steps:

1. Double-click on the unknown device listed under the device column.
2. Click **Edit**.
3. Change the device **Name**.
4. Enter the **Instruction register Length**.
5. Click **OK**.
6. Save the **.cdf**.

## Setting Up Your Hardware

The Quartus II Programmer provides the flexibility to choose a download cable or programming hardware. Before you can program or configure your device, you must have the correct hardware setup.

### Related Information

- [Setting Up Programming Hardware](#)  
Describes the steps to set up your hardware.
- [Setting up Programming Hardware in Quartus II Software](#)  
Describes the programming hardware driver installation.

## Setting the JTAG Hardware

The JTAG server allows the Quartus II Programmer to access the JTAG hardware. You can also access the JTAG download cable or programming hardware connected to a remote computer through the JTAG server of that computer. With the JTAG server, you can control the programming or configuration of devices from a single computer through other computers at remote locations. The JTAG server uses the TCP/IP communications protocol.

### Related Information

- [Using the JTAG Server](#)  
Lists how to use the JTAG Server

## Running JTAG Daemon with Linux

The JTAGD daemon is the Linux version of a JTAG server. The JTAGD daemon allows a board which is connected to a Linux host to be programmed or debugged over the network from a remote machine. The JTAGD daemon also allows multiple programs to use JTAG resources at the same time.

Run the JTAGD daemon to avoid:

- the JTAGD server from exiting after two minutes of idleness.
- the JTAGD server from not accepting connections from remote machines, which might lead to an intermittent failure.

To run JTAGD as a daemon, follow these steps:

1. Create an `/etc/jtagd` directory.
2. Set the permissions of this directory and the files in the directory to allow you to have the read/write access.
3. Run `jtagd` (with no arguments) from your `quartus/bin` directory.

The JTAGD daemon is now running and does not terminate when you log off.

## Using the JTAG Chain Debugger Tool

The JTAG Chain Debugger tool allows you to test the JTAG chain integrity and detect intermittent failures of the JTAG chain. In addition, the tool allows you to shift in JTAG instructions and data through the JTAG interface and step through the test access port (TAP) controller state machine for debugging purposes. You access the tool from the Tools menu on the main menu of the Quartus II software.

### Related Information

- [Using the JTAG Chain Debugger](#)

## Stand-Alone Quartus II Programmer

Altera offers the free stand-alone Quartus II Programmer, which has the same full functionality as the Quartus II Programmer in the Quartus II software. The stand-alone Quartus II Programmer is useful when programming your devices with another workstation, so you do not need two full licenses. You can download the stand-alone Quartus II Programmer from the Download Center on the Altera website.

**Related Information**

- [Download Center](#)

You can download the stand-alone Quartus II Programmer from this page.

## Programming and Configuration Modes

The following table lists the programming and configuration modes supported by Altera devices.

**Table 18-3: Programming and Configuration Modes**

Configuration Mode Supported by the Quartus II Programmer	FPGA	CPLD	Configuration Device	Serial Configuration Device
JTAG	Yes	Yes	Yes	—
Passive Serial (PS)	Yes	—	—	—
Active Serial (AS) Programming	—	—	—	Yes
Configuration via Protocol (CvP)	Yes	—	—	—
In-Socket Programming	—	Yes (except for MAX II CPLDs)	Yes	Yes

**Related Information**

- [About Programming](#)
- [Configuration via Protocol \(CvP\) Implementation in Altera FPGAs User Guide](#)  
Describes the CvP configuration mode.
- [Programming Adapters](#)  
Contains a list of programming adapters available for Altera devices.

## Design Security Keys

The Quartus II Programmer supports the generation of encryption key programming files and encrypted configuration files for Altera FPGAs that support the design security feature. You can also use the Quartus II Programmer to program the encryption key into the FPGA.

**Related Information**

- [AN 556: Using the Design Security Features in Altera FPGAs](#)

## Convert Programming Files Dialog Box

The **Convert Programming Files** dialog box in the Programmer allows you to convert programming files from one file format to another. For example, to store the FPGA data in configuration devices, you can convert the **.sof** data to another format, such as **.pof**, **.hexout**, **.rbf**, **.rpd**, or **.jic**, and then program the configuration device.



You can also configure multiple devices with an external host, such as a microprocessor or CPLD. For example, you can combine multiple **.sof** files into one **.pof**.

To access the **Convert Programming Files** dialog box, on the main menu of the Quartus II software, click **File**, and then click **Convert Programming Files**.

**Related Information**

- [Convert Programming Files Dialog Box](#)

## Debugging Your Configuration

Use the **Advanced** option in the **Convert Programming Files** dialog box to debug your configuration. You must choose the advanced settings that apply to your Altera device. You can direct the Quartus II software to enable or disable an advanced option by turning the option on or off in the **Advanced Options** dialog box.

When you change settings in the **Advanced Options** dialog box, the change affects **.pof**, **.jic**, **.rpd**, and **.rbf** files.

The following table lists the **Advanced Options** settings in more detail.

**Table 18-4: Advanced Options Settings**

Option Setting	Description
Disable EPCS ID check	<p>FPGA skips the EPCS silicon ID verification.</p> <p>Default setting is unavailable (EPCS ID check is enabled).</p> <p>Applies to the single- and multi-device AS configuration modes on all FPGA devices.</p>
Disable AS mode CONF_DONE error check	<p>FPGA skips the CONF_DONE error check.</p> <p>Default setting is unavailable (AS mode CONF_DONE error check is enabled).</p> <p>Applies to single- and multi-device (AS) configuration modes on all FPGA devices.</p> <p>The CONF_DONE error check is disabled by default for Stratix V, Arria V, and Cyclone V devices for AS-PS multi device configuration mode.</p>
Program Length Count adjustment	<p>Specifies the offset you can apply to the computed PLC of the entire bitstream.</p> <p>Default setting is 0. The value must be an integer.</p> <p>Applies to single- and multi-device (AS) configuration modes on all FPGA devices.</p>

Option Setting	Description
Post-chain bitstream pad bytes	<p>Specifies the number of pad bytes appended to the end of an entire bitstream.</p> <p>Default value is set to 0 if the bitstream of the last device is uncompressed. Set to 2 if the bitstream of the last device is compressed.</p>
Post-device bitstream pad bytes	<p>Specifies the number of pad bytes appended to the end of the bitstream of a device.</p> <p>Default value is 0. No negative integer.</p> <p>Applies to all single-device configuration modes on all FPGA devices.</p>
Bitslice padding value	<p>Specifies the padding value used to prepare bitslice configuration bitstreams, such that all bitslice configuration chains simultaneously receive their final configuration data bit.</p> <p>Default value is 1. Valid setting is 0 or 1.</p> <p>Use only in 2, 4, and 8-bit PS configuration mode, when you use an EPC device with the decompression feature enabled.</p> <p>Applies to all FPGA devices that support enhanced configuration devices.</p>

The following table lists the symptoms you may encounter if a configuration fails, and describes the advanced options you must use to debug your configuration.

Failure Symptoms	Disable EPCS ID Check	Disable AS Mode CONF_DONE Error Check	PLC Settings	Post-Chain Bitstream Pad Bytes	Post-Device Bitstream Pad Bytes	Bitslice Padding Value
Configuration failure occurs after a configuration cycle.	—	Yes	Yes	Yes <sup>(11)</sup>	Yes <sup>(12)</sup>	—
Decompression feature is enabled.	—	Yes	Yes	Yes <sup>(11)</sup>	Yes <sup>(12)</sup>	—

<sup>(11)</sup> Use only for multi-device chain

<sup>(12)</sup> Use only for single-device chain

Failure Symptoms	Disable EPCS ID Check	Disable AS Mode CONF_DONE Error Check	PLC Settings	Post-Chain Bitstream Pad Bytes	Post-Device Bitstream Pad Bytes	Bitslice Padding Value
Encryption feature is enabled.	—	Yes	Yes	Yes <sup>(11)</sup>	Yes <sup>(12)</sup>	—
CONF_DONE stays low after a configuration cycle.	—	Yes	Yes <sup>(13)</sup>	Yes <sup>(11)</sup>	Yes <sup>(12)</sup>	—
CONF_DONE goes high momentarily after a configuration cycle.	—	Yes	Yes <sup>(14)</sup>	—	—	—
FPGA does not enter user mode even though CONF_DONE goes high.	—	—	—	Yes <sup>(11)</sup>	Yes <sup>(12)</sup>	—
Configuration failure occurs at the beginning of a configuration cycle.	Yes	—	—	—	—	—
Newly introduced EPCS, such as EPCS128.	Yes	—	—	—	—	—

<sup>(13)</sup> Start with positive offset to the PLC settings

<sup>(14)</sup> Start with negative offset to the PLC settings

Failure Symptoms	Disable EPCS ID Check	Disable AS Mode CONF_DONE Error Check	PLC Settings	Post-Chain Bitstream Pad Bytes	Post-Device Bitstream Pad Bytes	Bitslice Padding Value
Failure in <b>.pof</b> generation for EPC device using Quartus II Convert Programming File Utility when the decompression feature is enabled.	—	—	—	—	—	Yes

## Converting Programming Files for Partial Reconfiguration

The **Convert Programming File** dialog box supports the following programming file generation and option for Partial Reconfiguration:

- Partial-Masked SRAM Object File (**.pmsf**) output file generation, with **.msf** and **.sof** as input files.
- **.rbf** for Partial Reconfiguration output file generation, with a **.pmsf** as the input file.

**Note:** The **.rbf** for Partial Reconfiguration file is only for Partial Reconfiguration.

- Providing the **Enable decompression during Partial Reconfiguration** option to enable the option bit for bitstream decompression during Partial Reconfiguration, when converting a full design **.sof** to any supported file type.

### Related Information

[Design Planning for Partial Reconfiguration](#)

### Generating **.pmsf** using a **.msf** and a **.sof**

To generate the **.pmsf** in the **Convert Programming Files** dialog box, follow these steps:

1. In the **Convert Programming Files** dialog box, under the **Programming file type** field, select **Partial-Masked SRAM Object File (.pmsf)**.
2. In the **File name field**, specify the necessary output file name.
3. In the **Input files to convert** field, add necessary input files to convert. You can add only a **.msf** and **.sof**.
4. Click **Generate**.

### Generating **.rbf** for Partial Reconfiguration Using a **.pmsf**

After generating the **.pmsf**, convert the **.pmsf** to a **.rbf** for Partial Reconfiguration in the **Convert Programming Files** dialog box.

To generate the **.rbf** for Partial Reconfiguration, follow these steps:

1. In the **Convert Programming Files** dialog box, in the **Programming file type** field, select **Raw Binary File for Partial Reconfiguration (.rbf)**.
2. In the **File name** field, specify the output file name.
3. In the **Input files to convert** field, add input files to convert. You can add only a **.pmsf**.
4. After adding the **.pmsf**, select the **.pmsf** and click **Properties**. The **PMSF File Properties** dialog box appears.
5. Make your selection either by turning on or turning off the following options:
  - **Compression option**—This option enables compression on Partial Reconfiguration bitstream. If you turn on this option, then you must turn on the **Enable decompression during Partial Reconfiguration** option.
  - **Enable SCRUB mode option**—The default of this option is based on AND/OR mode. This option is valid only when Partial Reconfiguration masks in your design are not overlapped vertically. Otherwise, you cannot generate the **.rbf** for Partial Reconfiguration.
  - **Write memory contents option**—This option is a workaround for initialized RAM/ROM in a Partial Reconfiguration region.

For more information about these option, refer to the [Design Planning for Partial Reconfiguration](#).

6. Click **OK**.
7. Click **Generate**.

## Enable Decompression during Partial Reconfiguration Option

You can turn on the **Enable decompression during Partial Reconfiguration** option in the **SOF File Properties: Bitstream Encryption** dialog box, which can be accessed from the **Convert Programming File** dialog box. This option is available when converting a **.sof** to any supported programming file types listed in [Table 18-2](#).

This option is hidden for other targeted devices that do not support Partial Reconfiguration. To view this option in the **SOF File Properties: Bitstream Encryption** dialog box, the **.sof** must be targeted on an Altera device that supports Partial Reconfiguration.

If you turn on the **Compression** option when generating the **.rbf** for Partial Reconfiguration, then you must turn on the **Enable decompression during Partial Reconfiguration** option.

## Flash Loaders

Parallel and serial configuration devices do not support the JTAG interface. However, you can use a flash loader to program configuration devices in-system via the JTAG interface. You can use an FPGA as a bridge between the JTAG interface and the configuration device. The Quartus II software supports parallel and serial flash loaders.

### Related Information

- [About Flash Loaders](#)

## JTAG Debug Mode for Partial Reconfiguration

The JTAG debug mode allows you to configure partial reconfiguration bitstream through the JTAG interface. Use this feature to debug PR bitstream and eventually helping you in your PR design prototyping. This feature is available for internal and external host.

During JTAG debug operation, the JTAG command sent from the Quartus II Programmer ignores and overrides most of the Partial Reconfiguration IP core interface signals (`clk`, `pr_start`, `double_pr`, `data[]`, `data_valid`, and `data_read`).

**Note:** The `TCK` is the main clock source for PR IP core during this operation.

You can view the status of Partial Reconfiguration operation in the messages box and the Progress bar in the Quartus II Programmer. The `PR_DONE`, `PR_ERROR`, and `CRC_ERROR` signals will be monitored during PR operation and reported in the Messages box at the end of the operation.

The Quartus II Programmer can detect the number of `PR_DONE` instruction(s) in plain or compressed PR bitstream and, therefore, can handle single or double PR cycle accordingly. However, only single PR cycle is supported for encrypted Partial Reconfiguration bitstream in JTAG debug mode (provided that the specified device is configured with the encrypted base bitstream which contains the PR IP core in the design).

**Note:** Configuring an incompatible PR bitstream to the specified device may corrupt your design, including the routing path and the PR IP core placed in the static region. When this issue occurs, the PR IP core stays in an undefined state, and the Quartus II Programmer is unable to reset the IP core. As a result, the Quartus II Programmer generates the following error when you try to configure a new PR bitstream:

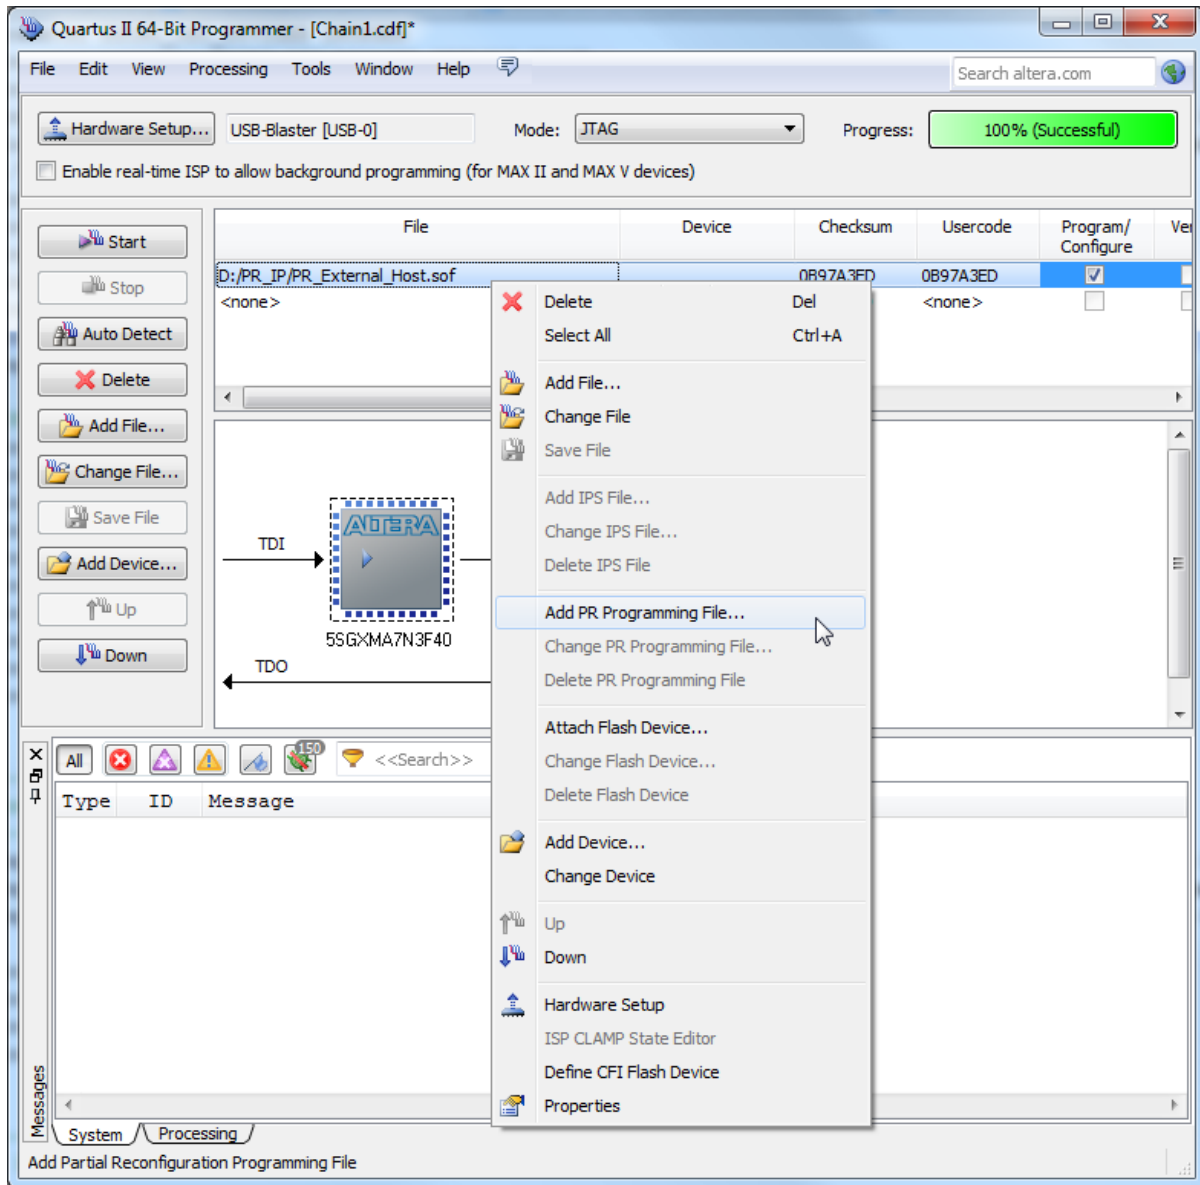
```
Error (12897): Partial Reconfiguration status: Can't reset the PR megafunction.  
This issue occurred because the design was corrupted by an incompatible PR  
bitstream in the previous PR operation. You must reconfigure the device with a good  
design.
```

## Configuring Partial Reconfiguration Bitstream in JTAG Debug Mode

To configure the Partial Reconfiguration bitstream in JTAG debug mode, follow these steps:


1. In the Quartus II Programmer GUI, right click on a highlighted base bitstream (in `.sof`) and then click **Add PR Programming File** to add the PR bitstream (`.rbf`).

Figure 18-3: Adding PR Programming File



2. After adding the PR bitstream, you can change or delete the Partial Reconfiguration programming file by clicking **Change PR Programming File** or **Delete PR Programming File**.

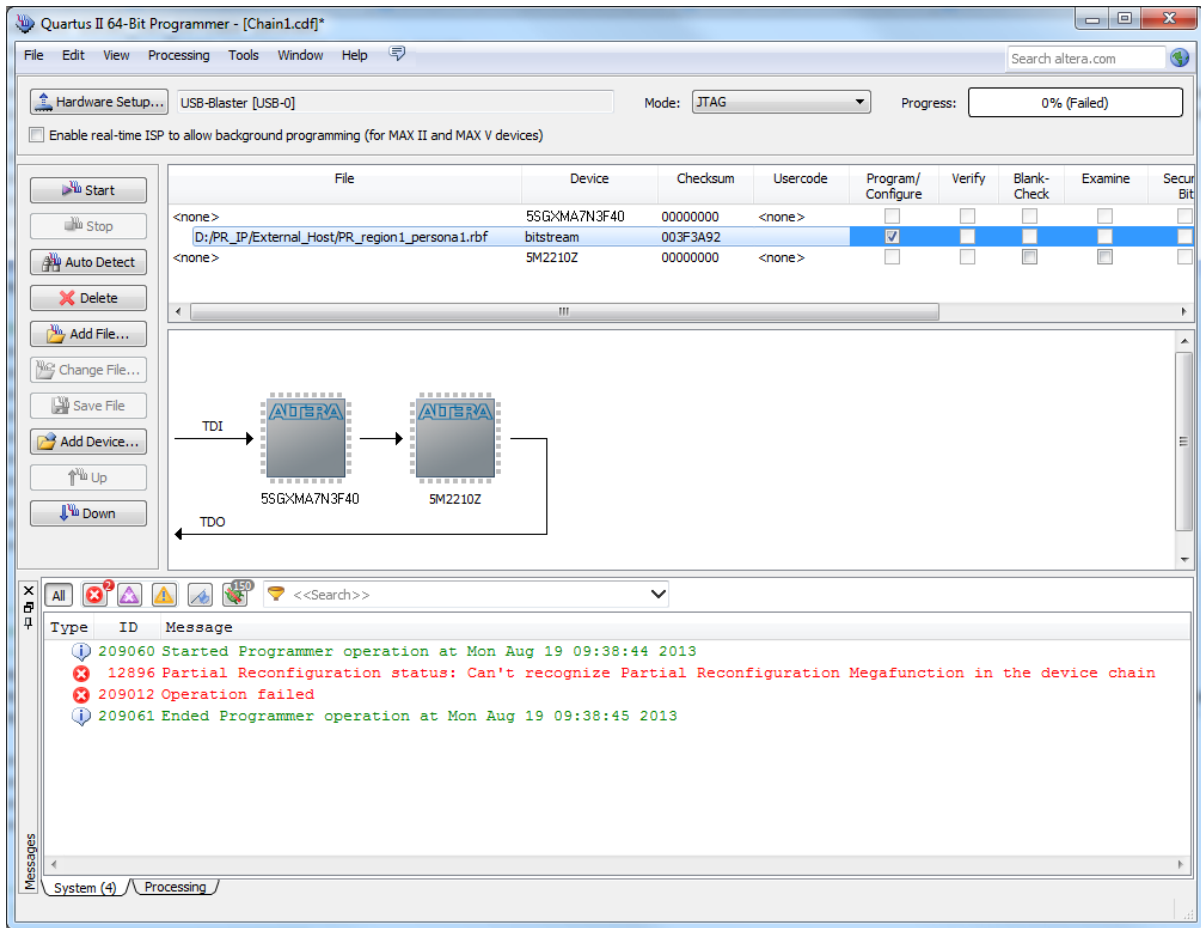
Figure 18-4: Change PR Programming File or Delete PR Programming File



3. Click **Start** to configure the PR bitstream. The Quartus II Programmer generates an error message if the specified device does not contain the PR IP core in the design (you must instantiate the Partial Reconfiguration IP core in your design to use the JTAG debug mode).

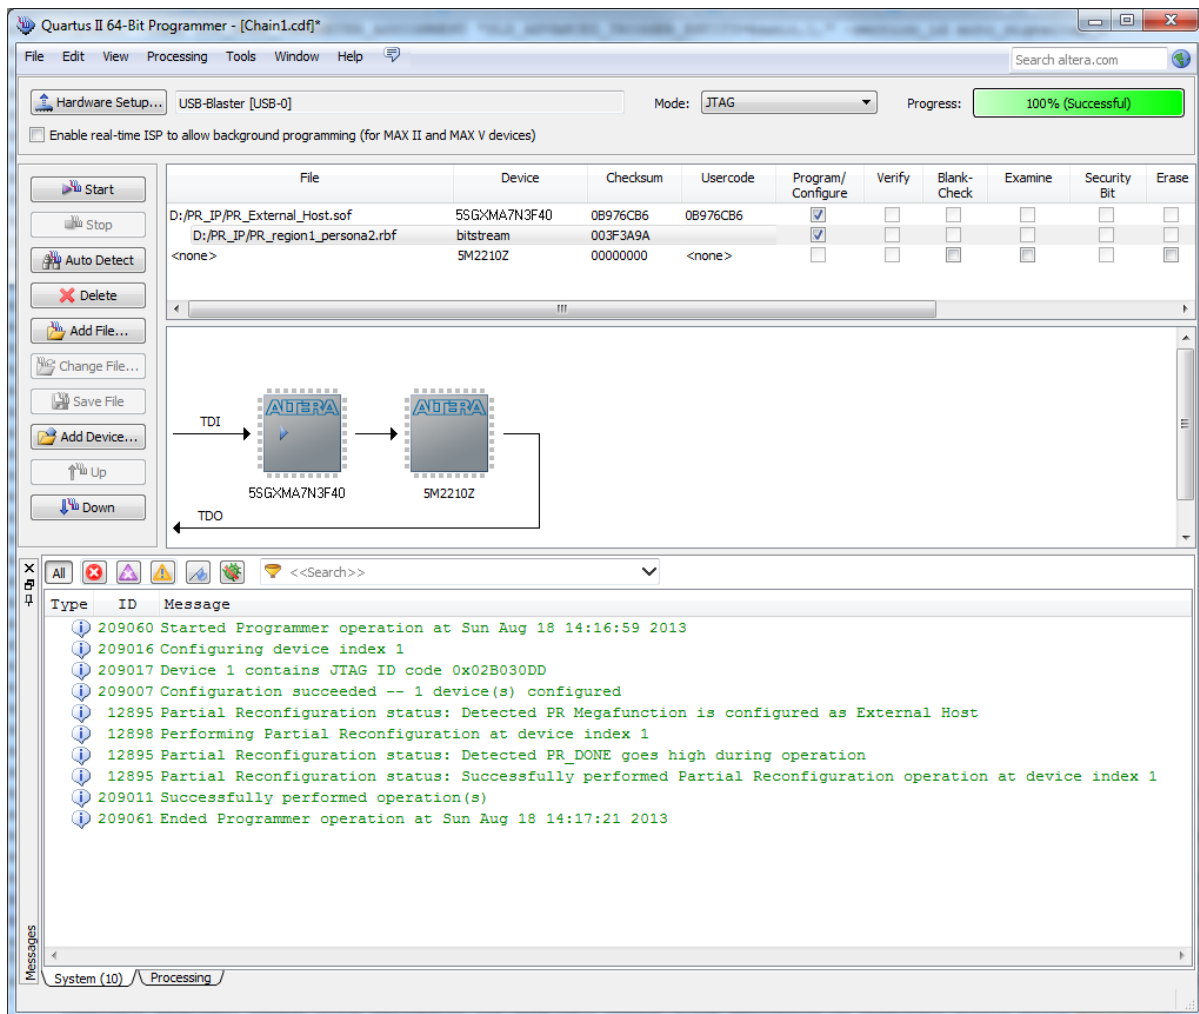


Figure 18-5: Starting PR Bitstream Configuration



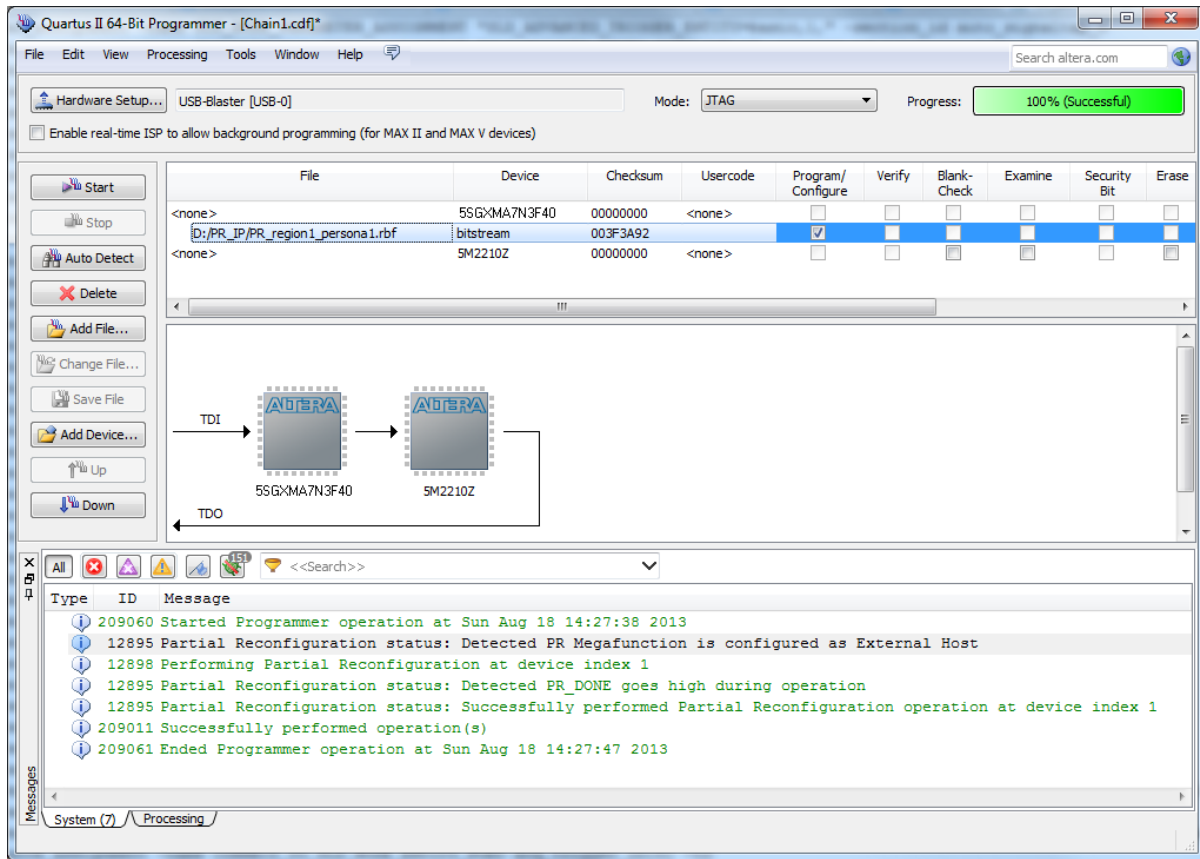
4. Configure the valid .rbf in JTAG debug mode with the Quartus II Programmer.

Figure 18-6: Configuring Valid .rbf



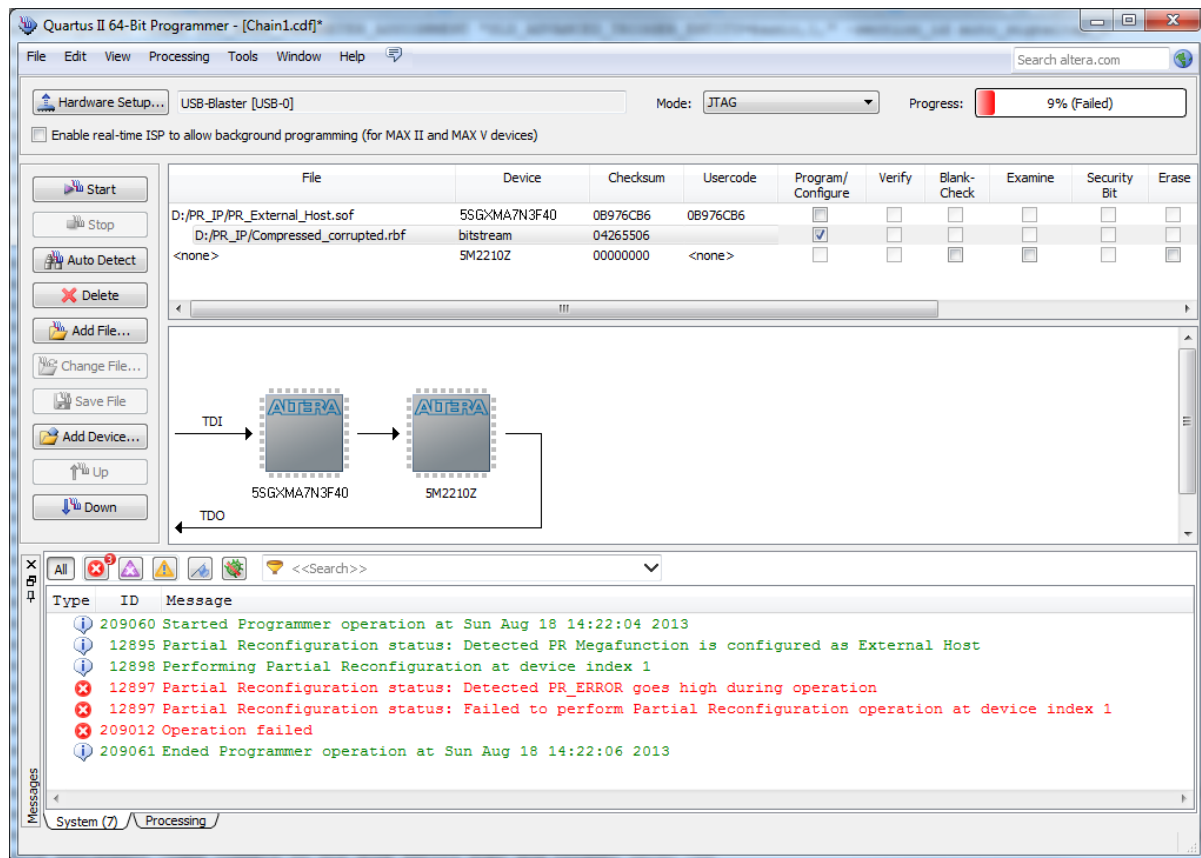
- The JTAG debug mode is also supported if the PR IP core is pre-programmed on the specified device.

Figure 18-7: Partial Reconfiguration IP Core Successfully Pre-programmed



- The Quartus II Programmer reports error when you try to configure the corrupted .rbf in JTAG debug mode.

Figure 18-8: Configuring Corrupted .rbf



## Scripting Support

In addition to the Quartus II Programmer GUI, you can use the Quartus II command-line executable `quartus_pgm.exe` to access programmer functionality from the command line and from scripts. The programmer accepts `.pof`, `.sof`, and `.jic` programming or configuration files and `.cdf`.

The following example shows a command that programs a device:

```
quartus_pgm -c byteblasterII -m jtag -o bpv:design.pof
```

Where:

- `-c byteblasterII` specifies the ByteBlaster II download cable
- `-m jtag` specifies the JTAG programming mode
- `-o bpv` represents the blank-check, program, and verify operations
- `design.pof` represents the `.pof` used for the programming

The Programmer automatically executes the erase operation before programming the device.

**Note:** For linux terminal, use the following command:

```
quartus_pgm -c byteblasterII -m jtag -o bpv\;design.pof
```

**Related Information**

- [About Quartus II Scripting](#)

## The jtagconfig Debugging Tool

You can use the `jtagconfig` command-line utility (which is similar to the auto detect operation in the Quartus II Programmer) to check the devices in a JTAG chain and the user-defined devices.

For more information about the `jtagconfig` utility, type one of the following commands at the command prompt:

```
jtagconfig -h
```

```
jtagconfig --help
```

**Note:** The help switch does not reference the `-n` switch. The `jtagconfig -n` command shows each node for each JTAG device.

**Related Information**

- [Command-Line Scripting](#)

## Generating .pmsf using a .msf and a .sof

You can generate a `.pmsf` with the `quartus_cpf` command by typing the following command:

```
quartus_cpf -p <pr_revision.msf> <pr_revision.sof> <new_filename.pmsf>
```

## Generating .rbf for Partial Reconfiguration using a .pmsf

You can generate a `.rbf` for Partial Reconfiguration with the `quartus_cpf` command by typing the following command:

```
quartus_cpf -o foo.txt -c <pr_revision.pmsf> <pr_revision.rbf>
```

**Note:** You must run this command in the same directory where the files are located.

## Document Revision History

**Table 18-5: Document Revision History**

Date	Version	Chages
December 2014	14.1.0	Updated the Scripting Support section to include a Linux command to program a device.

Date	Version	Changes
June 2014	14.0.0	<ul style="list-style-type: none"> <li>Added Running JTAG Daemon.</li> <li>Removed Cyclone III and Stratix III devices references.</li> <li>Removed MegaWizard Plug-In Manager references.</li> <li>Updated Secondary Programming Files section to add notes about the Quartus II Programmer support for <b>.rbf</b> files.</li> </ul>
November 2013	13.1.0	<ul style="list-style-type: none"> <li>Converted to DITA format.</li> <li>Added JTAG Debug Mode for Partial Reconfiguration and Configuring Partial Reconfiguration Bitstream in JTAG Debug Mode sections.</li> </ul>
November 2012	12.1.0	<ul style="list-style-type: none"> <li>Updated Table 18–3 on page 18–6, and Table 18–4 on page 18–8.</li> <li>Added “Converting Programming Files for Partial Reconfiguration” on page 18–10, “Generating .pmsf using a .msf and a .sof” on page 18–10, “Generating .rbf for Partial Reconfiguration Using a .pmsf” on page 18–12, “Enable Decompression during Partial Reconfiguration Option” on page 18–14</li> <li>Updated “Scripting Support” on page 18–15.</li> </ul>
June 2012	12.0.0	<ul style="list-style-type: none"> <li>Updated Table 18–5 on page 18–8.</li> <li>Updated “Quartus II Programmer GUI” on page 18–3.</li> </ul>
November 2011	11.1.0	<ul style="list-style-type: none"> <li>Updated “Configuration Modes” on page 18–5.</li> <li>Added “Optional Programming or Configuration Files” on page 18–6.</li> <li>Updated Table 18–2 on page 18–5.</li> </ul>
May 2011	11.0.0	<ul style="list-style-type: none"> <li>Added links to Quartus II Help.</li> <li>Updated “Hardware Setup” on page 21–4 and “JTAG Chain Debugger Tool” on page 21–4.</li> </ul>
December 2010	10.1.0	<ul style="list-style-type: none"> <li>Changed to new document template.</li> <li>Updated “JTAG Chain Debugger Example” on page 20–4.</li> <li>Added links to Quartus II Help.</li> <li>Reorganized chapter.</li> </ul>
July 2010	10.0.0	<ul style="list-style-type: none"> <li>Added links to Quartus II Help.</li> <li>Deleted screen shots.</li> </ul>
November 2009	9.1.0	No change to content.

Date	Version	Chages
March 2009	9.0.0	<ul style="list-style-type: none"><li>• Added a row to Table 21-4.</li><li>• Changed references from “JTAG Chain Debug” to “JTAG Chain Debugger”.</li><li>• Updated figures.</li></ul>

**Related Information**

[Quartus II Handbook Archive](#)

For previous versions of the Quartus II Handbook