# Vortex RTL Textual Description Grammar

The Cecil Group

Department of Computer Science and Engineering
University of Washington
Box 352350, Seattle, Washington 98195-2350
cecil@cs.washington.edu

February 2, 1997 4:28 pm

This document describes the grammar for the textual form of the Vortex compiler's RTL. This is the language used to communicate between the C++ and Modula-3 front-ends and Vortex, and is also used to describe some Cecil primitives. These grammar rules are considered to extend those in the Cecil language reference manual.

Warning: the RTL language has not been designed from scratch, but rather via accretion of features. It is slowly being reworked to make the operators more general and language-independent, but this process is not complete.

## 1 Top-Level Entry Points

For Cecil RTL primitives, the syntax of the `code` string part of the `prim_body` is as follows:

```
rtl_prim        ::= rtl_stmts
```

For languages that have separate front-ends, such as C++ and Modula-3, source files are translated by the front-end into RTL files. An RTL file is described by the following grammar:

```
rtl_file        ::= { rtl_decl | pragma }
rtl_decl        ::= rtl_object_decl
                |   rtl_method_decl
                |   rtl_extern_method_decl
                |   rtl_field_decl
                |   rtl_var_decl
                |   rtl_assoc_decl
                |   rtl_include_decl
                |   rtl_top_level_stmts
```

## 2 Global Declarations

Declarations at the top level can appear in any order, and can be mutually recursive. Evaluation of global variable initializers is done top-down in the order given.

An object declaration introduces a new class. The `isa` list names its immediate superclasses. (There is a root class implicitly defined, and all classes inherit (directly or indirectly) from this root class implicitly.) Inheritance defines a partial order on objects; redundant or transitively-implied inheritance links are ignored. In the presence of diamond-shaped multiple inheritance, shared ancestor classes are including only once in the hierarchy, as if all classes were virtual base classes in C++ terminology.

```
rtl_object_decl ::= "object" rtl_name ["isa" rtl_names] [object_id_pragma] ";"
rtl_names       ::= rtl_name {"," rtl_name}
```

An `object_id_pragma` is attached to an `object_decl`, and specifies the names of the symbols to use as class identifiers for that class, when it is viewed from particular static types.

```
object_id_pragma ::= "(**" "ID" "(" "[" object_id_pair {"," object_id_pair } "]"
                     "**)"
object_id_pair   ::= "[" static_type "," rtl_name "]"
```

A method declaration defines a procedure in the intermediate program.

```
rtl_method_decl  ::= "method" [rtl_visibility] rtl_name "(" {rtl_formals} ")"
                       ":" representation [pragma] "{" rtl_stmts "}" [";"]
rtl_formals      ::= rtl_formal {"," rtl_formal}
```

Each formal has a name (can be omitted if it's not used within the body of the method), a specializer class (if omitted or if given as `@any`, then the argument is unspecialized), and a representation (if omitted, then it defaults to `OOP`). If `formals_are_immutable` is defined, then formal variables cannot be assigned, otherwise they can be. If `uses_PIC_based_runtime` is set, then only `OOP` formals can be specialized to something other than `any`. Formal parameters and results are passed by value (i.e. shallow-copied), including formals represented using `bytes[`$n$`]`, unlike C which passes arrays by reference.

```
rtl_formal       ::= [rtl_name] ["@" rtl_name] [":" representation]
```

A method (or field or global variable declaration) can be marked `private` (meaning that it is scoped within the current file, as in C; `static` has the same meaning), or `extern` (meaning that it is visible to other files, as in C, and that it has been declared already in some other file).

```
rtl_visibility   ::= "private" | "static" | "extern"
```

An extern method declaration declares a method and says that its definition may be external to the part of the program that the compiler sees. It is allowable to provide an extern method declaration in one source file and the definition of the method in another source file. In this case, the actual definition of the method takes precedence over the extern declaration.

```
rtl_extern_method_decl::= "extern" "method" rtl_name "(" {rtl_formals} ")"
                        ":" representation ";"
```

A field declaration declares an instance variable on a class. If var is present, the field is mutable, otherwise the field is considered immutable. The formal parameter specifies which class the instance variable is part of. If `front_end_does_field_layout` is set, then field declarations are unnecessary, and only low-level object allocation (with precomputed size) and pointer load & store operations are supported, otherwise field declarations should be provided and high-level object allocation operations and field load and store operations are provided. If `generate_field_accessors` is set, then a field declaration implicitly generates a get accessor method (and a set accessor method, if mutable).

```
rtl_field_decl   ::= ["var"] "field" [rtl_visibility] rtl_name
                       "(" rtl_formal ")" ":" representation [pragma]
                       [":=" literal] ";"
```

A global variable declaration introduces a variable in the global (current) scope, with the given visibility and representation. Variables (global and local) are assumed to be assignable.

```
rtl_var_decl     ::= "decl" [rtl_visibility] representation rtl_name
                       [":=" (literal | static_init)] ";"
```

A static initializer provides the mechanism for specifying the static initial value of a global variable whose representation is an array of bytes. It provides a sequence of initializers, each of which specifies an offset, a representation, and an initial value of that representation placed at that offset. The different kinds of initial values are pad (uninitialized data), the address of some global variable plus the given offset, a procedure

address (also specifying the return type of the procedure), an integer constant (possibly given in a string literal), a float constant (possibly given in a string literal), and a byte sequence.

```
static_init      ::= "static_init" "{" {static_init_entry} "}" ";"
static_init_entry::= "[" int "," representation "," static_init_value "]"
static_init_value::= "pad"
                   |   "&" rtl_name ["+" integer]
                   |   "&_proc" rtl_name representation
                   |   integer | "istr" "(" string ")"
                   |   float | double | "fstr" "(" string ")"
                   |   "str" "(" string ")"
```

An associate declaration treats the named method as a method in the given message set with the listed specializers.

```
rtl_assoc_decl   ::= "associate" rtl_name "with"
                       rtl_name "(" rtl_formals ")" ";"
```

For example:

```
   associate wrapper__3FooFi with wrapper__Fi(@Foo, @any);
```

This says consider the routine named `wrapper__3FooFi` to be an implementation of the message `wrapper__Fi` attached at class `Foo`.

An include declaration names another file of RTL code to include as part of the program.

```
rtl_include_decl  ::= "include" string ";"
```

A top-level statement block is a list of RTL statements that are to be executed upon program start-up. These expressions are evaluated in the order encountered during reading of the input files.

```
rtl_top_level_stmts ::= "rtl" "{" rtl_stmts "}" [";"]
```

## 3 Data Representations

Representations are specified in declarations and in statements. They indicate the representation (format, bit encoding, size, etc.) of global variables, formals, method results, locals, and so on. A representation is an unstructured sequence of bytes (whose components are accessed through pointer loads & stores), optionally with some alignment, something that is roughly equivalent to one of C's atomic data types (perhaps with a specified size, otherwise with a "natural" size), or something that is a special tagged-pointer-sized integer, float, tagged pointer, NLR return code, or return pair.

```
representation   ::= "bytes" "[" integer "]" [ "align" integer ]
                 |   base_rep {"*"}
base_rep         ::= "char" | "u_char" | "short" | "u_short"
                 |   "int"  | "u_int"  | "long"  | "u_long"
                 |   "float" | "double"
                 |   "int1" | "u_int1" | "int2" | "u_int2"
                 |   "int4"  | "u_int4"  | "int8"  | "u_int8"
                 |   "float4" | "float8"
                 |   "void"
                 |   "wordInt" | "unsignedWordInt" | "wordFloat" | "OOP"
                 |   "ReturnCode" | "PairType"
```

## 4 Statements

RTL statements are the following:

```
rtl_stmts        ::= { rtl_stmt ";" }
```

```
rtl_stmt        ::= alu
                  | vector_op
                  | field_op
                  | ptr_op
                  | allocation
                  | exception_scope
                  | label
                  | goto
                  | fn_call
                  | return
                  | fatal
                  | static_info
```

Many RTL statements may be prefaced with a declaration. These declarations are used to pass along type information that will be needed to declare variables in the C++ code we generate as output. A given variable should be declared at the first textual place it is assigned or read, but the scope of the declaration is the entire method. References

```
decl            ::= "decl" representation
```

Local variable declarations, assignments, unary and binary operations, and conversions (which specify the target representation to convert to).

```
alu             ::= local_decl | assign | unop | binop | conversion
local_decl      ::= [decl] rtl_name
assign          ::= [decl] rtl_name ":=" value
unop            ::= [decl] rtl_name ":=" unop_op value
binop           ::= [decl] rtl_name ":=" value binop_op value
conversion      ::= [decl] rtl_name ":=" "convert" value representation
```

Arguments to RTL operators can name RTL variables, can name undeclared variables (in which case they are assumed to name global variables of type `void*` that are defined by the run-time system and the header files included with the generated code), can mention integer, float, double, char, or string literals, and can mention language-specific void, true, false, and null-pointer values.

```
value           ::= rtl_name | literal | "void" | "true" | "false" | "null"
literal         ::= integer | float | double | character | string
```

Vector operations. Stores or fetches from a vector, array, or string object. The representation of the element is specified as part of the operation. It is possible to mark a vector location as immutable, meaning that the location won't be modified in the future (except possibly for explicit stores later in this procedure).

```
vector_op       ::= vector_load | vector_store
vector_load     ::= [decl] rtl_name ":=" vector_location
vector_store    ::= [decl] vector_location ":=" value
vector_location ::= rtl_name "[" index "]" representation ["is_immutable"]
index           ::= rtl_name | integer
```

Field (instance variable) operations. Stores or fetches from a declared field of an object. Field contents are considered immutable when the corresponding field declaration specified an immutable field.

```
field_op        ::= field_load | field_store
field_load      ::= [decl] rtl_name ":=" field_location
field_store     ::= [decl] field_location ":=" value
field_location  ::= rtl_name "." rtl_name "@" rtl_name
```

Simple pointer operations. Used to access language-level arrays, structs, and so on, if they've been translated down into this low level. The representation specifies the representation of the pointer target value being loaded or stored, and the index is a byte offset from the base pointer. Pointer target locations can be marked immutable, as with vector locations.

```
ptr_op          ::= ptr_load | ptr_store | addr_op
ptr_load        ::= [decl] rtl_name ":=" ptr_location
ptr_store       ::= ptr_location ":=" value
ptr_location    ::= "*" "(" rtl_name ["+" index] ")" representation
                        ["is_immutable"]
```

Pointers can be created by taking the address of a variable or of a procedure.

```
addr_op         ::= [decl] rtl_name ":=" "&" rtl_name
                |   [decl] rtl_name ":=" "&_proc" rtl_name
```

High-level allocation of new data structures. Object allocation can be used for classes with associated field declarations specifying the class's structure and `front_end_does_field_layout` turned off; various class-testing operations and field access operations can be applied to these objects. The vector allocator operations are Cecil-specific. The array allocator is used to allocate array-like objects, given the name of the array class, the representation of the elements, and the length of the array. The array can be defined as mutable or immutable. Vector and array objects are indexed using the vector load and store operations above, and the num_elems unary operator can be applied to vectors and arrays.

```
allocation      ::= [decl] rtl_name ":=" obj_allocator
                |   [decl] rtl_name ":=" vec_allocator
                |   [decl] rtl_name ":=" array_allocator
obj_allocator   ::= "new" rtl_name
vec_allocator   ::= vec_alloc_op value
vec_alloc_op    ::= "new_m_vector" | "new_i_vector"
                |   "new_m_string" | "new_i_string"
                |   "new_m_float_vector" | "new_i_float_vector"
                |   "new_m_word_vector" | "new_i_word_vector"
array_allocator ::= array_alloc_op rtl_name representation value
array_alloc_op  ::= "new_m_array" | "new_i_array"
```

Exceptions can be raised by calls, sends, and explicit raises. The scope of a handler for some exceptions is bracketed by `enter_` and `exit_exn_scope` RTL statements. These statements name the label where the innermost-enclosing exception handler for this scope begins. For languages where `exceptions_through_longjmp` is set, the runtime system is assumed to select the right handler if an exception is raised; the enter_exn_scope provides enough information to the runtime system to make this possible. Otherwise, exceptions propagate outwards through exception scopes incrementally, and each exception handling scope is responsible for forwarding control to the enclosing exception handling scope if it is not the right handler.

```
exception_scope ::= enter_exn_scope | exit_exn_scope
enter_exn_scope ::= "enter_exn_scope" label_name rtl_name ["(" rtl_name ")"]
                      ["[" value "," value "," value "]"]
exit_exn_scope  ::= "exit_exn_scope" label_name
```

The `enter_exn_scope` operator implicitly declares the variable to which the "name" of the exception being raised will be bound, and optionally declares the name which will be bound to the argument of the raised exception; these names are expected to be used in the exception handling code. For Modula-3, three

additional parameters are needed, giving the address of the exception record, the address of the module's exception handler list, and the class of the exception. (Other languages will presumably need extensions to support arguments appropriate for their runtime system implementation of exceptions.)

Labels that start exception handlers must be flagged as such.

```
label           ::= ("label" | "exn_handler_label") label_name
label_name      ::= rtl_name
```

Various kinds of unconditional and conditional branches. `is_class` is true iff the value is an instance of the named class; `inherits_from` extends this test to include subclasses as well. `inherits_from` can take a variable name (a computed `CecilMap*`) as an argument instead of a class name constant.

```
goto            ::= cond_goto | uncond_goto | case_goto | raise
cond_goto       ::= "if" test "goto" label_name
test            ::= unop_op value
                |   value binop_op value
                |   value "inherits_from" rtl_name
                |   value "is_class" rtl_name
uncond_goto     ::= "goto" label_name
```

`case_goto` implements an N-way branch. The argument value should be an integer in the range [0..N-1], and one of the N target labels is jumped to based on this value.

```
case_goto       ::= "case" value "goto" "[" label_names "]"
label_names     ::= label_name {"," label_name}
```

Exceptions are raised using the `raise` command. The first operand evaluates to the exception "name" (or exception object, for e.g. C++ and Java), and the second optional value is the argument to the exception. Control transfers unconditionally to the appropriate enclosing exception handler when the raise is executed.

```
raise           ::= "raise_exn" value ["(" value ")"]
```

Various kinds of procedure calls. The result of a call can be omitted (e.g. if the callee returns no value). All kinds of calls allow a marker blocking inlining of the call as well as an indication of whether the call can raise exceptions (a.k.a. produce a non-local return). (If `sends_can_raise` is set, then sends and table sends default to raising exceptions, otherwise calls default to not raising exceptions.)

```
fn_call         ::= [[decl] rtl_name ":="] (call | send | table_send)
                      ["no_inline"] ["can_raise_exns" | "no_exns"]
```

Calls are regular direct calls. The callee is either a regular named procedure (potentially one defined in RTL with a method decl, in which case it can be inlined; `call_external` blocks this check), or a value that evaluates to the address of the procedure to call.

```
call            ::= call_desc "(" [values] ")"
call_desc       ::= "call" rtl_name
                |   "call_external" rtl_name
                |   "call_indirect" value
values          ::= value {"," value}
```

Sends are messages, and search (using Cecil's inheritance rules) the named method set for the most specific applicable method, based on the run-time classes of the send arguments. Send RTLs can only be used if `uses_PIC_based_runtime` is set.

```
send            ::= "send" rtl_name "(" [values] ")"
```

Table sends are similar to sends, except that they get translated into the appropriate table lookup sequence (precomputed by the language front-ends) if they cannot be statically bound.

```
table_send       ::= "table_send" rtl_name "(" [values] ")"
                       "[" table_index "]" [ "id_offset" "(" int_literal ")" ]
                       [static_type_info]
table_index      ::= int_literal
```

The address of a lookup table is fetched from the object specified as the first argument of the send, at the byte offset specified by the `id_offset` clause. If no `id_offset` is specified, the table address is assumed to be in the first word of the object. The `table_index` specifies the index into the table, giving a table entry to use for dispatching this send. For example:

```
t1 := table_send my_message(t2, t3) [2] id_offset(4)
```

This means: fetch the address of the table from `table_base := *(t2+4)`, then use the table entry at `table_base+(2*sizeof(table_entry))` to dispatch this send. The precise way in which the table is used to dispatch the send is language dependent. Currently, Vortex supports two kinds of tables: a simple table with address-sized entries that contain the address of a routine to call, and a more complicated, C++-style table format that contains both the address of a routine and an adjustment value that should be added to the first argument before entering the callee.

The static type of the receiver argument of a table_send can be specified by naming the class that the receiver is known to inherit from.

```
static_type_info ::= "static_type" "(" static_type ")"
static_type      ::= rtl_name
```

Return statements.

```
return           ::= "return" [value]
```

Fatal errors: abort execution with a string message.

```
fatal            ::= "fatal" "(" string ")"
```

Assertions about the static properties of a value.

```
static_info      ::= "assert_type" rtl_name static_type_info
```

Currently, static information specifies either an exact class for the variable, or specifies a cone of possible classes (and the value could be an instance of the class or any subclass of the class), or gives another variable whose type to copy. The optimizer uses this information to reason about the possible values that could be stored in the given variable.

```
static_type_info ::= "map" "(" rtl_name ")"
                   |   "cone" "(" rtl_name ")"
                   |   "same_as" "(" rtl_name ")"
```

## 5 Tokens

The token syntax for RTL is the same as for Cecil, with the addition of `rtl_name`, which has the same syntax as a C identifier:

```
rtl_name         ::= (letter | "_") {letter | digit | "_"}
```

## 6 Unary and Binary Operators

These are the unary and binary operators currently defined. **[Need to add int8 versions of these.]**

| unop_op name | description |
|---|---|
| unary_-_int | integer negation |
| unary_-_float | float negation |
| unary_-_double | double negation |
| ~_int | bitwise complement |
| num_elems_int | extract the number of elements of an array, vector, or string object. |
| box_int | convert an untagged machine integer into a tagged Cecil integer OOP |
| box_float | convert an untagged machine float into a boxed, heap-allocated Cecil float OOP |
| box_double | convert an untagged machine double into a boxed, heap-allocated Cecil double OOP |
| box_char | convert an untagged machine integer into a boxed, heap-allocated Cecil char OOP |
| unbox_int | inverse of box_int |
| unbox_float | inverse of box_float |
| unbox_double | inverse of box_double |
| unbox_char | inverse of box_char |
| map_of | returns the `CecilMap*` pointer representing the run-time class of the argument; used as the argument to inherits_from |

| binop_op name | description |
|---|---|
| +_int | Add two signed integers |
| -_int | Subtract ... |
| *_int | Multiply ... |
| /_int | Divide ... |
| %_int | Modulus ... |
| &_int | Bitwise and two integers |
| \|_int | ... or ... |
| ^_int | ... xor ... |
| xnor_int | ... xnor ... |
| <<_int | bitwise left shift one integer by another |
| >>_int | ... arithmetic right shift ... |
| >>_logical_int | ... logical right shift ... |
| +_unsigned_int | Add two unsigned integers |
| -_unsigned_int | Subtract ... |
| *_unsigned_int | Multiply ... |

| binop_op name | description |
|---|---|
| /_unsigned_int | Divide ... |
| %_unsigned_int | Modulus ... |
| &_unsigned_int | Bitwise and two unsigned integers |
| \|_unsigned_int | ... or ... |
| ^_unsigned_int | ... xor ... |
| xnor_unsigned_int | ... xnor ... |
| <<_unsigned_int | bitwise left shift one unsigned integer by another |
| >>_unsigned_int | ... (logical) right shift ... |
| +_ptr | Add two values, one of which is assumed to be a pointer and the other an integer byte count, to produce a new pointer |
| -_ptr | Subtract an integer byte count from a pointer to produce a new pointer |
| -_ptrdiff | Subtract two pointers to produce an integer difference, in bytes |
| +_float | Add two floats |
| -_float | Subtract ... |
| *_float | Multiply ... |
| /_float | Divide ... |
| +_double | Add two doubles |
| -_double | Subtract ... |
| *_double | Multiply ... |
| /_double | Divide ... |
| =_int_log | Test whether two integers are equal |
| !=_int_log | ... not equal |
| >_int_log | Test whether one signed integer is greater than another |
| >=_int_log | ... greater than or equal to ... |
| <_int_log | ... less than ... |
| <=_int_log | ... less than or equal to ... |
| >_unsigned_log | Test whether one unsigned integer is greater than another |
| >=_unsigned_log | ... greater than or equal to ... |
| <_unsigned_log | ... less than ... |
| <=_unsigned_log | ... less than or equal to ... |
| =_ptr_log | Test whether two pointers are equal |
| !=_ptr_log | ... not equal |
| >_ptr_log | Test whether one pointer is greater than another |
| >=_ptr_log | ... greater than or equal to ... |
| <_ptr_log | ... less than ... |

| binop_op name | description |
|---|---|
| <=_ptr_log | ... less than or equal to ... |
| =_float_log | Test whether two floats are equal |
| !=_float_log | ... not equal |
| >_float_log | Test whether one float is greater than another |
| >=_float_log | ... greater than or equal to ... |
| <_float_log | ... less than ... |
| <=_float_log | ... less than or equal to ... |
| =_double_log | Test whether two doubles are equal |
| !=_double_log | ... not equal |
| >_double_log | Test whether one double is greater than another |
| >=_double_log | ... greater than or equal to ... |
| <_double_log | ... less than ... |
| <=_double_log | ... less than or equal to ... |