

Introduction to the MIPS ISA

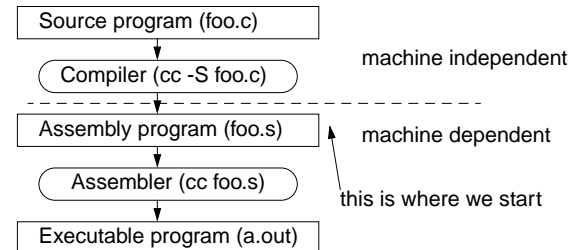
CSE378

WINTER, 2001

20

Overview

- Remember that the machine only understands very basic instructions (machine instructions)
- It is the compiler's job to translate your high-level (e.g. C program) into machine instructions. In more detail (forgetting linking):



- Assembly language is a thin veneer over machine language.

CSE378

WINTER, 2001

21

Overview (2)

- Think about a simple C program...

```
int array[100];
void main () {
    int i;
    for (i=0; i<100; i++)
        array[i] = i;
}
```

- What set of instructions (ISA) should the machine provide to execute it?
- What does your intuition tell you about trade-offs between the ISA and the size (length) of the resulting machine program?
- What kind of trade-offs exist between the ISA and the speed, cost, complexity of the hardware needed to execute the program?
- Tensions and contributing factors: ease of programming, ease of hardware design, program/memory size, compiler technology

CSE378

WINTER, 2001

22

MIPS ISA Overview

- MIPS is a "computer family": R2000/3000 (32-bit), R4000/4400 (64-bit)
- New entries include R8000 (scientific/graphics) and R10000
- MIPS originated as a Stanford project: **M**icroprocessor without **I**nterlocked **P**ipe **S**tages
- H+P posit 4 principles of hardware design. Try to keep them in mind during our discussion of the MIPS ISA:
 1. Simplicity favors regularity
 2. Smaller is faster
 3. Compromise
 4. Make the common case fast

CSE378

WINTER, 2001

23

MIPS is a RISC

- RISC = Reduced (Regular/Restricted) Instruction Set Computer
- All arithmetic operations are of the form:

$$R_d \leftarrow R_s \text{ OP } R_t \quad \# \text{ the } R_s \text{ are registers}$$

- Important restriction: MIPS is a load store architecture: the ALU can only operate on registers Why?
- Basic operations (really only a few kinds)
 1. Arithmetic (addition, subtraction, etc)
 2. Logical (and, or, xor, etc)
 3. Comparison (less-than, greater-than, etc)
 4. Control (branches, jumps, etc)
 5. Memory access (load and store)
- All MIPS instructions are 32 bits long

CSE378

WINTER, 2001

24

MIPS is a Load-Store Architecture

- Every operand of a MIPS instruction must be in a register (with some exceptions)
- Variables must be loaded into registers
- Results have to be stored back into memory
- Example C fragment...

```
a = b + c;
d = a + b;
```

- ... would be "translated" into something like:

```
Load b into register Rx
Load c into register Ry
Rz <- Rx + Ry
Store Rz into a
Rz <- Rz + Rx
Store Rz into d
```

CSE378

WINTER, 2001

25

MIPS Registers

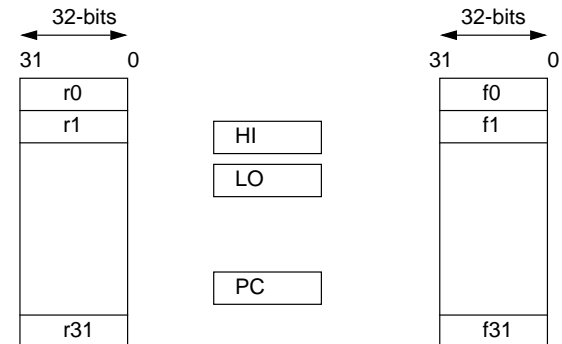
- Provides thirty-two, 32-bit registers, named \$0, \$1, \$2 .. \$31 used for:
 - integer arithmetic
 - address calculations
 - special-purpose functions defined by convention
 - temporaries
- A 32-bit program counter (PC)
- Two 32-bit registers HI and LO used specifically for multiplication and division
- Thirty-two 32-bit registers \$f0, \$f1, \$f2 .. \$f31 used for floating point arithmetic
- Other special-purpose registers (see later)

CSE378

WINTER, 2001

26

Registers are part of the process "state"



CSE378

WINTER, 2001

27

MIPS Register Names and Conventions

Register	Name	Function	Comment
\$0	zero	Always 0	No-op on write don't use it!
\$1	\$at	reserved for assembler	
\$2-3	\$v0-v1	expression eval./function return	
\$4-7	\$a0-a3	proc/func call parameters	
\$8-15	\$t0-t7	volatile temporaries	not saved on call
\$16-23	\$s0-s7	temporaries (saved across calls)	saved on call
\$24-25	\$t8-t9	volatile temporaries	not saved on call
\$26-27	\$k0-k1	reserved kernel/OS	don't use them
\$28	\$gp	pointer to global data area	
\$29	\$sp	stack pointer	
\$30	\$fp	frame pointer	
\$31	\$ra	proc/func return address	

CSE378

WINTER, 2001

28

MIPS Information Units

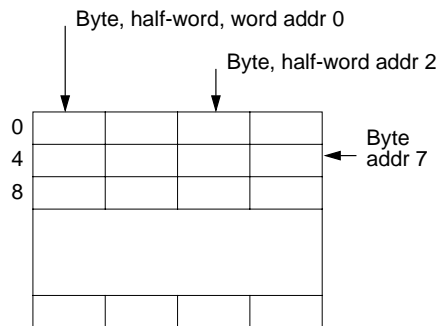
- Data types and size:
 - Byte
 - Half-word (2 bytes)
 - Word (4 bytes)
 - Float (4 bytes, single precision format)
 - Double (8 bytes, double precision format)
- Memory is byte addressable.
- A data type must start on an address divisible by its size (in bytes)
- The address of the data type is the address of its lowest byte (MIPS on DEC is little endian)

CSE378

WINTER, 2001

29

MIPS Addressing



- In MIPS (and most byte addressable machines) every word should start at an address divisible by 4.
- Why?

CSE378

WINTER, 2001

30

MIPS Instruction Types

- As we said earlier, there are very few basic operations :
 1. Memory access (load and store)
 2. Arithmetic (addition, subtraction, etc)
 3. Logical (and, or, xor, etc)
 4. Comparison (less-than, greater-than, etc)
 5. Control (branches, jumps, etc)
- We'll use the following notation when describing instructions:

rd: destination register (modified by instruction)
rs: source register (read by instruction)
rt: source/destination register (read or read+modified)
immed: a 16-bit value

CSE378

WINTER, 2001

31

Running Example

- Let's translate this simple C program into MIPS assembly code:

```
int x, y;

void main() {
    ...
    x = x + y;
    if (x==y) {
        x = x + 3;
    }
    x = x + y + 42;
    ...
}
```

CSE378

WINTER, 2001

32

Load and Store Instructions

- Data is explicitly moved between memory and registers through load and store instructions.
- Each load or store must specify the memory address of the memory data to be read or written.
- Think of a MIPS address as a 32-bit, unsigned integer.
- Because a MIPS instruction is always 32 bits long, the address must be specified in a more compact way.
- We always use a *base register* to address memory
- The base register points somewhere in memory, and the instruction specifies the register number, and a 16-bit, *signed offset*
- A single base register can be used to access any byte within ??? bytes from where it points in memory.

CSE378

WINTER, 2001

33

Load and Store Examples

- Load a word from memory:

```
lw rt, offset(base) # rt <- memory[base+offset]
```

- Store a word into memory:

```
sw rt, offset(base) # memory[base+offset] <- rt
```

- For smaller units (bytes, half-words) only the lower bits of a register are accessible. Also, for loads, you need to specify whether to sign or zero extend the data.

```
lb rt, offset(base) # rt <- sign-extended byte
lbu rt, offset(base) # rt <- zero-extended byte
sb rt, offset(base) # store low order byte of rt
```

CSE378

WINTER, 2001

34

Arithmetic Instructions

Opcode	Operands	Comments
ADD	rd, rs, rt	# rd <- rs + rt
ADDI	rt, rs, immed	# rt <- rs + immed
SUB	rd, rs, rt	# rd <- rs - rt

Examples:

ADD	\$8, \$8, \$10	# r8 <- r9 + r10
ADD	\$t0, \$t1, \$t2	# t0 <- t1 + t2
SUB	\$s0, \$s0, \$s1	# s0 <- s0 - s1
ADDI	\$t3, \$t4, 5	# t3 <- t4 + 5

CSE378

WINTER, 2001

35

Multiply and Divide Instructions

- Multiplying two 32-bit numbers can yield a 64 bit number. Hence the use of HI and LO registers.
- Dividing two numbers yields a quotient and a remainder.

Opcode	Operands	Comments
MULT	rs, rt	# HI/LO <- rs * rt
MULTU	rs, rt	# HI/LO <- rs * rt
DIV	rs, rt	# LO <- rs/rt # HI <- rs rem rt
DIVU	rs, rt	# LO <- rs/rt # HI <- rs rem rt

- If an operand is negative, the remainder is not specified by the MIPS architecture.

CSE378

WINTER, 2001

36

Multiply and Divide Instructions (2)

- There are instructions to move between HI/LO registers.

Opcode	Operands	Comments
MFHI	rd	# rd <- HI
MTHI	rs	# HI <- rs
MFLO	rd	# rd <- LO
MTLO	rs	# LO <- rs

CSE378

WINTER, 2001

37

Integer Arithmetic

- Numbers can be either *signed* or *unsigned*
- The above instructions all check for, and signal overflow should it occur.
- MIPS ISA provides instructions that don't care about overflows:
 - ADDU
 - ADDIU
 - SUBU, etc.
- For add and subtract, the computation is the same for both, but the machine will signal an overflow when one occurs for signed numbers.

CSE378

WINTER, 2001

38

Overflows in 2's Complement

- Overflow occurs when the addition of two numbers of the same sign results in a sum of the opposite sign
- Overflow cannot occur when adding operands of different signs

- Example 1: Assume a 4-bit machine. Register 9 contains 7 and register 10 contains 3
 - What happens when we use ADD? ADDU?

- Example 2: Assume a 4-bit machine. Register 9 contains 7 and register 10 contains -3
 - What happens when we use ADD? ADDU?

CSE378

WINTER, 2001

39

Flow of Control: Conditional Branches

- You can compare on...
 - equality or inequality of two registers
 - comparison of register to zero (>, <, <=, >=)
- ... and branch to a target that is a signed displacement (expressed in number of instructions [words not bytes!]) from the instruction following the branch.

CSE378

WINTER, 2001

40

Branches (2)

- In assembly language, it's easiest to just use the target address (from the label), rather than trying to figure out the number of instructions.

```
BEQ    rs, rt, target # branch if rs == rt
BNE    rs, rt, target # branch if rs != rt
BGTZ   rs, target     # branch if rs > 0
BGEZ   rs, target     # branch if rs >= 0
BLTZ   rs, target     # branch if rs < 0
BLEZ   rs, target     # branch if rs <= 0
```

CSE378

WINTER, 2001

41

Comparison Between Registers

- What if you want to branch if R6 is greater than R7?
- We can use the SLT instruction:

```
SLT    rd, rs, rt     # if rs<rt then rd <- 1
                          # else rd <- 0
SLTU   rd, rs, rt     # same, but rs,rt unsigned
```

- Example: Branch to L1 if \$5 > \$6

```
SLT    $7, $6, $5     # $7 = 1, if $6 < $5
BNE    $7, $0, L1
```

CSE378

WINTER, 2001

42

Jump Instructions

- Jump instructions allow for unconditional transfer of control:

```
J      target         # go to specified target
JR     rs              # jump to addr stored in rs
```

- Jump and link is used for procedure calls:

```
JAL    target         # jump to target, $31 <- PC
JALR   rs, rd         # jump to addr in rs
                          # rd <- PC
```

- When calling a procedure, use JAL; to return, use JR \$31

CSE378

WINTER, 2001

43

Logic Instructions

- Used to manipulate bits within words, set up masks, etc.

Opcode	Operands	Comments
AND	rd, rs, rt	# rd <- AND(rs, rt)
ANDI	rt, rs, immed	# rt <- AND(rs, immed)
OR	rd, rs, rt	
ORI	rt, rs, immed	
XOR	rd, rs, rt	
XORI	rt, rs, immed	

- The immediate constant is limited to 16 bits
- To load a constant in the 16 upper bits of a register we use LUI:

Opcode	Operands	Comments
LUI	rt, immed	# rt<31,16> <- immed # rt<15,0> <- 0

CSE378

WINTER, 2001

44

Logic Instruction Examples

- Turn *on* the bits in the low order byte of R6:

```
ORI    $6, $6, 0x00ff    # set r6<7,0> to 1s
```

- Turn *off* the bits in the low order byte of R6:

```
LUI    $5, 0xffff      # set r5<31,16> to 1s
ORI    $5, 0xff00      # zero low order byte
AND    $6, $6, $5      # zap low order byte in R6
```

- Flip the the bits in the high order byte of R6: (check this one)

```
LUI    $5, 0xff00      # 1s in upper byte
ANDI   $5, 0x0000      # 0s every where else
XOR    $6, $6, $5      # flip upper bits...
```

CSE378

WINTER, 2001

45

Shift Instructions

- Used to move bits around within registers.
- Logical shifts (zeros are shifted in from end).

SLL	rd, rt, shamt	# rd = rt shifted left by # shamt
SRL	rd, rt, shamt	# right shift

- Arithmetic shift right (sign extend from left)

SRA	rd, rt, shamt	# rd = rt shifted right by # shamt, and sign extended
-----	---------------	--

- shamt is a 5-bit shift amount

CSE378

WINTER, 2001

46

Back to our example

```
.data          # start of data segment
x: .word       # data layout directive
y: .word       # allocate two words
.text         # start of text segment
.globl main
main:
la    $t0, x           # t0 holds &x
lw    $t1, 0($t0)      # t1 holds x
la    $t2, y           # t2 holds &y
lw    $t3, 0($t2)      # t3 holds y
add   $t1, $t1, $t3    # x = x+y
sw    $t1, 0($t0)
bne   $t1, $t3, L1     # if x == y
add   $t1, $t1, 3      # x = x+3
sw    $t1, 0($t0)
L1:  addi $t4, $t3, 17   # t4 = y + 17
add   $t1, $t1, $t4
sw    $t1, 0($t0)
```

CSE378

WINTER, 2001

47

Discussion

- Note that we're going to great lengths to preserve the semantics of the original C program.
- We're storing back values to their memory locations immediately after computing them.
- Why might this be a good idea?
- Why might this be a bad idea?

CSE378

WINTER, 2001

48

An optimized example

- We eliminated “unnecessary” stores..

```
.data          # start of data segment
x: .word       # data layout directive
y: .word       # allocate two words
.text         # start of text segment
.globl main
main:
    la    $t0, x          # t0 holds &x
    lw    $t1, 0($t0)     # t1 holds x
    la    $t2, y          # t2 holds &y
    lw    $t3, 0($t2)     # t3 holds y
    add   $t1, $t1, $t3   # x = x+y
    bne  $t1, $t3, L1     # if x == y
    add   $t1, $t1, 3      # x = x+3
L1:  addi $t4, $t3, 17     # t4 = y + 17
    add   $t1, $t1, $t4
    sw    $t1, 0($t0)
```

CSE378

WINTER, 2001

49

Example C Program

```
#include <stdio.h>
int array[100];

void main ()
{
    int i;
    for (i=0; i<100; i++)
        array[i] = i;
}
```

CSE378

WINTER, 2001

50

Assembly Version (Hand coded)

```
.data          # begin data segment
array: .space 400 # allocate 400 bytes

.text         # begin code segment
.globl main   # entry point must be global

main:  move $t0, $0 # $t0 is used as counter
    la  $t1, array # $t1 is pointer into array
start: bge $t0, 100, exit # more than 99 iterations?
    sw  $t0, 0($t1) # store zero into array
    addi $t0, $t0, 1 # increment counter
    addi $t1, $t1, 4 # increment pointer into array
    j   start      # goto top of loop
exit:  j   $ra      # return to caller of main...
```

CSE378

WINTER, 2001

51

Assembly Version (Compiler Generated)

```
.data
array: .space 400      # the comments are obviously
       .text          # NOT generated by the compiler!
       .globl main

main:
subu   $sp, 8          # make room on stack
sw     $0, 4($sp)     # i lives at 4($sp)...
$32:   # ...initialize it to zero
lw     $14, 4($sp)    # load i into $14
mul    $15, $14, 4    # $15 is used as base reg
sw     $14, array($15) # store i into array[i]
lw     $24, 4($sp)    # load i into $24
addu   $25, $24, 1    # increment $24
sw     $25, 4($sp)    # store new val into i
blt    $25, 100, $32  # if i<100 goto top
move   $2, $0         # set $2 to 0
addu   $sp, 8         # reset stack
j      $31
       .end main
```