

COMPUTER ORGANIZATION AND DESIGN

The Hardware/Software Interface



CSE 341 Review

ISA and Programming

Instruction Set

- The repertoire of instructions of a computer
- Different computers have different instruction sets
 - But with many aspects in common
- Early computers had very simple instruction sets
 - Simplified implementation
- Many modern computers also have simple instruction sets



The MIPS Instruction Set

- Used as the example throughout the book
- Stanford MIPS commercialized by MIPS Technologies (<u>www.mips.com</u>)
- Large share of embedded core market
 - Applications in consumer electronics, network/storage equipment, cameras, printers, ...
- Typical of many modern ISAs
 - See MIPS Reference Data tear-out card, and Appendixes B and E



Arithmetic Operations

- Add and subtract, three operands
 Two sources and one destination
 - add a, b, c # a gets b + c
- All arithmetic operations have this form
 - Design Principle 1: Simplicity favours regularity
 - Regularity makes implementation simpler
 - Simplicity enables higher performance at lower cost



Arithmetic Example

C code:

$$f = (g + h) - (i + j);$$

Compiled MIPS code:

add t0, g, h # temp t0 = g + h add t1, i, j # temp t1 = i + j sub f, t0, t1 # f = t0 - t1



Register Operands

- Arithmetic instructions use register operands
- MIPS has a 32 × 32-bit register file
 - Use for frequently accessed data
 - Numbered 0 to 31
 - 32-bit data called a "word"
- Assembler names
 - \$t0, \$t1, ..., \$t9 for temporary values
 - \$s0, \$s1, ..., \$s7 for saved variables
 - Design Principle 2: Smaller is faster
 - c.f. main memory: millions of locations



Register Operand Example

C code: f = (g + h) - (i + j);
f, ..., j in \$s0, ..., \$s4
Compiled MIPS code: add \$t0, \$s1, \$s2 add \$t1, \$s3, \$s4 sub \$s0, \$t0, \$t1



Memory Operands

- Main memory used for composite data
 - Arrays, structures, dynamic data
- To apply arithmetic operations
 - Load values from memory into registers
 - Store result from register to memory
- Memory is byte addressed
 - Each address identifies an 8-bit byte
- Words are aligned in memory
 - Address must be a multiple of 4
- MIPS is Big Endian
 - Most-significant byte at least address of a word
 - *c.f.* Little Endian: least-significant byte at least address



Memory Operand Example 1

- C code:
 - g = h + A[8];
 - g in \$\$1, h in \$\$2, base address of A in \$\$3
- Compiled MIPS code:
 - Index 8 requires offset of 32
 - 4 bytes per word



Memory Operand Example 2

C code: A[12] = h + A[8];h in \$s2, base address of A in \$s3 Compiled MIPS code: Index 8 requires offset of 32 lw \$t0, 32(\$s3) # load word add \$t0, \$s2, \$t0 sw \$t0, 48(\$s3) # store word



Registers vs. Memory

- Registers are faster to access than memory
- Operating on memory data requires loads and stores
 - More instructions to be executed
- Compiler must use registers for variables as much as possible
 - Only spill to memory for less frequently used variables
 - Register optimization is important!



Immediate Operands

- Constant data specified in an instruction addi \$s3, \$s3, 4
- No subtract immediate instruction
 - Just use a negative constant addi \$s2, \$s1, -1
 - Design Principle 3: Make the common case fast
 - Small constants are common
 - Immediate operand avoids a load instruction



The Constant Zero

- MIPS register 0 (\$zero) is the constant 0
 - Cannot be overwritten
- Useful for common operations
 - E.g., move between registers add \$t2, \$s1, \$zero



2s-Complement Signed Integers

- Bit 31 is sign bit
 - I for negative numbers
 - 0 for non-negative numbers
- –(–2^{n 1}) can't be represented
- Non-negative numbers have the same unsigned and 2s-complement representation
- Some specific numbers
 - 0: 0000 0000 ... 0000
 - −1: 1111 1111 ... 1111
 - Most-negative: 1000 0000 ... 0000
 - Most-positive: 0111 1111 ... 1111



Sign Extension

- Representing a number using more bits
 - Preserve the numeric value
- In MIPS instruction set
 - addi: extend immediate value
 - Ib, Ih: extend loaded byte/halfword
 - beq, bne: extend the displacement
- Replicate the sign bit to the left
 - c.f. unsigned values: extend with 0s
- Examples: 8-bit to 16-bit
 - +2: 0000 0010 => 0000 0000 0000 0010
 - -2: 1111 1110 => 1111 1111 1111 1110



Representing Instructions

- Instructions are encoded in binary
 - Called machine code
- MIPS instructions
 - Encoded as 32-bit instruction words
 - Small number of formats encoding operation code (opcode), register numbers, ...
 - Regularity!
- Register numbers
 - \$t0 \$t7 are reg's 8 15
 - \$t8 \$t9 are reg's 24 25
 - \$s0 \$s7 are reg's 16 23



MIPS R-format Instructions

| ор | rs | rt | rd | shamt | funct |
|--------|--------|--------|--------|--------|--------|
| 6 bits | 5 bits | 5 bits | 5 bits | 5 bits | 6 bits |

- Instruction fields
 - op: operation code (opcode)
 - rs: first source register number
 - rt: second source register number
 - rd: destination register number
 - shamt: shift amount (00000 for now)
 - funct: function code (extends opcode)



R-format Example

| ор | rs | rt | rd | shamt | funct |
|--------|--------|--------|--------|--------|--------|
| 6 bits | 5 bits | 5 bits | 5 bits | 5 bits | 6 bits |

add \$t0, \$s1, \$s2

| special | \$s1 | \$s2 | \$tO | 0 | add |
|---------|-------|-------|-------|-------|--------|
| 0 | 17 | 18 | 8 | 0 | 32 |
| 000000 | 10001 | 10010 | 01000 | 00000 | 100000 |

$000001000110010010000000100000_2 = 02324020_{16}$



MIPS I-format Instructions

| ор | rs | rt | constant or address |
|--------|--------|--------|---------------------|
| 6 bits | 5 bits | 5 bits | 16 bits |

- Immediate arithmetic and load/store instructions
 - rt: destination or source register number
 - Constant: -2¹⁵ to +2¹⁵ 1
 - Address: offset added to base address in rs
- Design Principle 4: Good design demands good compromises
 - Different formats complicate decoding, but allow 32-bit instructions uniformly
 - Keep formats as similar as possible



Logical Operations

Instructions for bitwise manipulation

| Operation | С | Java | MIPS |
|-------------|----|------|-----------|
| Shift left | << | << | s]] |
| Shift right | >> | >>> | srl |
| Bitwise AND | & | & | and, andi |
| Bitwise OR | | | or, ori |
| Bitwise NOT | ~ | ~ | nor |

Useful for extracting and inserting groups of bits in a word



Shift Operations

| ор | rs | rt | rd | shamt | funct |
|--------|--------|--------|--------|--------|--------|
| 6 bits | 5 bits | 5 bits | 5 bits | 5 bits | 6 bits |

- shamt: how many positions to shiftShift left logical
 - Shift left and fill with 0 bits
 - s11 by *i* bits multiplies by 2ⁱ
- Shift right logical
 - Shift right and fill with 0 bits
 - srl by *i* bits divides by 2ⁱ (unsigned only)



AND Operations

Useful to mask bits in a word
Select some bits, clear others to 0 and \$t0, \$t1, \$t2





OR Operations

Useful to include bits in a word Set some bits to 1, leave others unchanged

or \$t0, \$t1, \$t2



\$t0 0000 0000 0000 000 00<mark>11 11</mark>01 1100 0000



NOT Operations

Useful to invert bits in a word
Change 0 to 1, and 1 to 0
MIPS has NOR 3-operand instruction
a NOR b == NOT (a OR b)
nor \$t0, \$t1, \$zero Register

Register 0: always read as zero

\$t1 0000 0000 0000 00011 1100 0000 0000

\$t0 | 1111 1111 1111 1100 0011 1111 1111



Conditional Operations

- Branch to a labeled instruction if a condition is true
 - Otherwise, continue sequentially
- beq rs, rt, L1
 - if (rs == rt) branch to instruction labeled L1;
- bne rs, rt, L1
 - if (rs != rt) branch to instruction labeled L1;
- ∎j L1
 - unconditional jump to instruction labeled L1



Compiling If Statements



Compiling Loop Statements

C code:

while (save[i] == k) i += 1;

i in \$s3, k in \$s5, address of save in \$s6
 Compiled MIPS code:





More Conditional Operations

Set result to 1 if a condition is true Otherwise, set to 0 slt rd, rs, rt if (rs < rt) rd = 1; else rd = 0;</p> slti rt, rs, constant • if (rs < constant) rt = 1; else rt = 0; Use in combination with beg, bne slt \$t0, \$s1, \$s2 # if (\$s1 < \$s2) bne \$t0, \$zero, L # branch to L



Branch Instruction Design

- Why not blt, bge, etc?
- Hardware for $<, \geq, \dots$ slower than $=, \neq$
 - Combining with branch involves more work per instruction, requiring a slower clock
 - All instructions penalized!
 - beq and bne are the common case
- This is a good design compromise



Signed vs. Unsigned

- Signed comparison: slt, slti
- Unsigned comparison: sltu, sltui
- Example
 - \$s0 = 1111 1111 1111 1111 1111 1111 1111

 - slt \$t0, \$s0, \$s1 # signed __1 < +1 ⇒ \$t0 = 1</pre>
 - sltu \$t0, \$s0, \$s1 # unsigned +4,294,967,295 > +1 ⇒ \$t0 = 0



Procedure Calling

Steps required

- 1. Place parameters in registers
- 2. Transfer control to procedure
- 3. Acquire storage for procedure
- 4. Perform procedure's operations
- 5. Place result in register for caller
- 6. Return to place of call



Register Usage

- \$a0 \$a3: arguments (reg's 4 7)
- \$v0, \$v1: result values (reg's 2 and 3)
- \$t0 \$t9: temporaries
 - Can be overwritten by callee
- \$s0 \$s7: saved
 - Must be saved/restored by callee
- \$gp: global pointer for static data (reg 28)
- \$sp: stack pointer (reg 29)
- \$fp: frame pointer (reg 30)
- \$ra: return address (reg 31)



Procedure Call Instructions

- Procedure call: jump and link
- jal ProcedureLabel
 - Address of following instruction put in \$ra
 - Jumps to target address
- Procedure return: jump register
 - jr \$ra
 - Copies \$ra to program counter
 - Can also be used for computed jumps
 e.g., for case/switch statements



Leaf Procedure Example

- C code: int leaf_example (int g, h, i, j) { int f; f = (g + h) - (i + j); return f; }
 - Arguments g, ..., j in \$a0, ..., \$a3
 - f in \$s0 (hence, need to save \$s0 on stack)
 - Result in \$v0



Leaf Procedure Example

MIPS code:

| <pre>leaf_example:</pre> | | | | | |
|--------------------------|-------|--------|--------|--|--|
| addi | \$sp, | \$sp, | -4 | | |
| SW | \$s0, | 0(\$sp | o) | | |
| add | \$t0, | \$a0, | \$a1 | | |
| add | \$t1, | \$a2, | \$a3 | | |
| sub | \$s0, | \$t0, | \$t1 | | |
| add | \$v0, | \$s0, | \$zero | | |
| ٦w | \$s0, | 0(\$s | o) | | |
| addi | \$sp, | \$sp, | 4 | | |
| jr | \$ra | | | | |

Save \$s0 on stack

Procedure body

Result

Restore \$s0

Return



Non-Leaf Procedures

- Procedures that call other procedures
- For nested call, caller needs to save on the stack:
 - Its return address
 - Any arguments and temporaries needed after the call
- Restore from the stack after the call


Non-Leaf Procedure Example

```
• C code:
int fact (int n)
{
  if (n < 1) return f;
  else return n * fact(n - 1);
}
```

- Argument n in \$a0
- Result in \$v0



Non-Leaf Procedure Example

MIPS code:

| fact | t: | | | | |
|------|------|-------|------------|---|--------------------------|
| | addi | \$sp, | \$sp, -8 | # | adjust stack for 2 items |
| | SW | \$ra, | 4(\$sp) | # | save return address |
| | SW | \$a0, | 0(\$sp) | # | save argument |
| | slti | \$t0, | \$a0, 1 | # | test for n < 1 |
| | beq | \$t0, | \$zero, L1 | | |
| | addi | \$v0, | \$zero, 1 | # | if so, result is 1 |
| | addi | \$sp, | \$sp, 8 | # | pop 2 items from stack |
| | jr | \$ra | | # | and return |
| L1: | addi | \$a0, | \$a0, -1 | # | else decrement n |
| | jal | fact | | # | recursive call |
| | ٦w | \$a0, | 0(\$sp) | # | restore original n |
| | ٦w | \$ra, | 4(\$sp) | # | and return address |
| | addi | \$sp, | \$sp, 8 | # | pop 2 items from stack |
| | mul | \$v0, | \$a0, \$v0 | # | multiply to get result |
| | jr | \$ra | | # | and return |



Local Data on the Stack



- Local data allocated by callee
 - e.g., C automatic variables
- Procedure frame (activation record)
 - Used by some compilers to manage stack storage



Byte/Halfword Operations

Could use bitwise operations MIPS byte/halfword load/store String processing is a common case lb rt, offset(rs) lh rt, offset(rs) Sign extend to 32 bits in rt lbu rt, offset(rs) lhu rt, offset(rs) Zero extend to 32 bits in rt sb rt, offset(rs) sh rt, offset(rs) Store just rightmost byte/halfword



32-bit Constants

- Most constants are small
 - 16-bit immediate is sufficient
- For the occasional 32-bit constant
 - lui rt, constant
 - Copies 16-bit constant to left 16 bits of rt
 - Clears right 16 bits of rt to 0

lhi \$s0, 61

0000 0000 0111 1101 0000 0000 0000 0000





Branch Addressing

Branch instructions specify

- Opcode, two registers, target address
- Most branch targets are near branch
 - Forward or backward

| ор | rs | rt | constant or address |
|--------|--------|--------|---------------------|
| 6 bits | 5 bits | 5 bits | 16 bits |

- PC-relative addressing
 - Target address = PC + offset × 4
 - PC already incremented by 4 by this time



Jump Addressing

- Jump (j and jal) targets could be anywhere in text segment
 - Encode full address in instruction

| ор | address |
|--------|---------|
| 6 bits | 26 bits |

- (Pseudo)Direct jump addressing
 Target address PC
 (address ×
 - Target address = PC_{31...28} : (address × 4)



Target Addressing Example

Loop code from earlier exampleAssume Loop at location 80000





Branching Far Away

If branch target is too far to encode with 16-bit offset, assembler rewrites the code
Example beq \$\$0,\$\$1, L1

```
↓
bne $s0,$s1, L2
j L1
L2: …
```



Addressing Mode Summary

1. Immediate addressing

op rs rt Immediate

2. Register addressing



3. Base addressing



4. PC-relative addressing



5. Pseudodirect addressing







Review

<u>_MIPS-32_IS</u>A

| 🗆 Ins | struction Categories |
|-------|----------------------|
| • | Computational |
| • | Load/Store |
| • | Jump and Branch |
| • | Floating Point |
| | - coprocessor |
| • | Memory Management |

Special



3 Instruction Formats: all 32 bits wide



<u>MIPS Arithmetic Instructions</u>

□ MIPS assembly language arithmetic statement

add \$t0, \$s1, \$s2 sub \$t0, \$s1, \$s2

Each arithmetic instruction performs one operation

Each specifies exactly three operands that are all contained in the datapath's register file (\$t0,\$s1,\$s2)

destination \leftarrow source1 op source2

□ Instruction Format (R format)



<u>MIPS Arithmetic Instructions</u>



MIPS Register File



- code density improves (since register are named with fewer bits than a memory location)

<u>Aside: MIPS RegisterConvention</u>

| Name | Registe | Usage | Preserv |
|--------------------|---------|------------------------|-----------------------------|
| | r | | e on |
| | Number | | call? |
| \$zero | 0 | constant 0 (hardware) | n.a. |
| \$at | 1 | reserved for assembler | n.a. |
| \$v0 - \$v1 | 2-3 | returned values | no |
| \$a0 - \$a3 | 4-7 | arguments | yes |
| \$t0 - \$t7 | 8-15 | temporaries | no |
| \$s0 - \$s7 | 16-23 | saved values | yes |
| \$t8 - \$t9 | 24-25 | temporaries | no |
| \$gp | 28 | global pointer | yes |
| \$sp | 29 | stack pointer | yes |
| \$fp | 30 | frame pointer | yes |
| SE451 Chapter 2.52 | 31 | return addr (hardware) | Yes Irwin, PSU, 2 |

800

<u>MIPS Memory Access Instructions</u>

MIPS has two basic data transfer instructions for accessing memory

lw \$t0, 4(\$s3) #load word from memory
sw \$t0, 8(\$s3) #store word to memory

The data is loaded into (lw) or stored from (sw) a register in the register file – a 5 bit address

- The memory address a 32 bit address is formed by adding the contents of the base address register to the offset value
 - A 16-bit field meaning access is limited to memory locations within a region of ±2¹³ or 8,192 words (±2¹⁵ or 32,768 bytes) of the address in the base register

<u>_Machine Language – Load Instruction</u>

□ Load/Store Instruction Format (I format):





<u>ByteAddresses</u>

Since 8-bit bytes are so useful, most architectures address individual bytes in memory

 Alignment restriction - the memory address of a word must be on natural word boundaries (a multiple of 4 in MIPS-32)

□ Big Endian: leftmost byte is word address IBM 360/370, Motorola 68k, MIPS, Sparc, HP PA

□ Little Endian: rightmost byte is word address Intel 80x86, DEC Vax, DEC Alpha (Windows NT)

little endian byte 0



<u>Aside</u>: Loading and Storing Bytes

MIPS provides special instructions to move bytes

- lb \$t0, 1(\$s3) #load byte from memory
- sb \$t0, 6(\$s3) #store byte to memory

| 0x28 19 | 8 | 16 bit offset |
|---------|---|---------------|
|---------|---|---------------|

□ What 8 bits get loaded and stored?

- load byte places the byte from memory in the rightmost 8 bits of the destination register
 - what happens to the other bits in the register?
- store byte takes the byte from the rightmost 8 bits of a register and writes it to a byte in memory
 - what happens to the other bits in the memory word?

<u>_MIPS Immediate Instructions</u>

□ Small constants are used often in typical code

Possible approaches?

- put "typical constants" in memory and load them
- create hard-wired registers (like \$zero) for constants like 1
- have special instructions that contain constants !

addi \$sp, \$sp, 4 #\$sp = \$sp + 4

slti \$t0, \$s2, 15 #\$t0 = 1 if \$s2<15

□ Machine format (I format):

| 0x0A 18 8 0x0F | |
|----------------|--|
|----------------|--|

- □ The constant is kept inside the instruction itself!
 - Immediate format limits values to the range +2¹⁵-1 to -2¹⁵

<u>Review: Unsigned Binary Representation</u>

| Hex | Binary | Decimal |
|------------|--------|---------------------|
| 0x0000000 | 00000 | 0 |
| 0x0000001 | 00001 | 1 |
| 0x0000002 | 00010 | 2 |
| 0x0000003 | 00011 | 3 |
| 0x0000004 | 00100 | 4 |
| 0x0000005 | 00101 | 5 |
| 0x0000006 | 00110 | 6 |
| 0x0000007 | 00111 | 7 |
| 0x0000008 | 01000 | 8 |
| 0x0000009 | 01001 | 9 |
| | | |
| 0xFFFFFFFC | 11100 | 2 ³² - 4 |
| 0xFFFFFFD | 11101 | 2 ³² - 3 |
| 0xFFFFFFF | 11110 | 2 ³² - 2 |
| 0xFFFFFFFF | 11111 | 2 ³² - 1 |



CSE431 Chapter 2.58



<u>MIPS Shift Operations</u>

- Need operations to pack and unpack 8-bit characters into 32-bit words
- Shifts move all the hits in a word left or right

sll \$t2, \$s0, 8 #\$t2 = \$s0 << 8 bits
srl \$t2, \$s0, 8 #\$t2 = \$s0 >> 8 bits

□ Instruction Format (R format)

| 0 | 16 | 10 | 8 | 0x00 |
|---|----|----|---|------|
|---|----|----|---|------|

Such shifts are called logical because they fill with zeros

Notice that a 5-bit shamt field is enough to shift a 32-bit value 2⁵ – 1 or 31 bit positions

MIPS Logical Operations

There are a number of bit-wise logical operations in the MIPS ISA

and \$t0, \$t1, \$t2 #\$t0 = \$t1 & \$t2

or \$t0, \$t1, \$t2 #\$t0 = \$t1 | \$t2

nor \$t0, \$t1, \$t2 #\$t0 = not(\$t1 | \$t2)

□ Instruction Format (R format)



ori \$t0, \$t1, 0xFF00 #\$t0 = \$t1 | ff00

□ Instruction Format (I format)

<u>_MIPS Control Flow Instructions</u>

□ MIPS conditional branch instructions:

```
bne $s0, $s1, Lbl #go to Lbl if $s0≠$s1
beg $s0, $s1, Lbl #go to Lbl if $s0=$s1
```

• Ex: if (i==j) h = i + j;

bne \$s0, \$s1, Lbl1 add \$s3, \$s0, \$s1 Lbl1: ...

□ Instruction Format (I format):

| 0x05 16 17 16 bit offset | 0x05 | 16 | 17 | 16 bit offset |
|--------------------------------|------|----|----|---------------|
|--------------------------------|------|----|----|---------------|

How is the branch destination address specified?

<u>Specifying Branch Destinations</u>

□ Use a register (like in lw and sw) added to the 16-bit offset

- which register? Instruction Address Register (the PC)
 - its use is automatically implied by instruction
 - PC gets updated (PC+4) during the fetch cycle so that it holds the address of the next instruction
- limits the branch distance to -2¹⁵ to +2¹⁵-1 (word) instructions from the (instruction after the) branch instruction, but most branches are local anyway

from the low order 16 bits of the branch instruction



In <u>Support of Branch Instructions</u>

- We have beq, bne, but what about other kinds of branches (e.g., branch-if-less-than)? For this, we need yet another instruction, slt
- □ Set on less than instruction:

□ Instruction format (R format):



Alternate versions of slt

slti \$t0, \$s0, 25 # if \$s0 < 25 then \$t0=1 ...
sltu \$t0, \$s0, \$s1 # if \$s0 < \$s1 then \$t0=1 ...
sltiu \$t0, \$s0, 25 # if \$s0 < 25 then \$t0=1 ...</pre>

<u>Aside: More Branch Instructions</u>

- Can use slt, beq, bne, and the fixed value of 0 in register \$zero to create other conditions
 - less than blt \$s1, \$s2, Label

slt \$at, \$s1, \$s2 #\$at set to 1 if
bne \$at, \$zero, Label #\$s1 < \$s2</pre>

- less than or equal to ble \$s1, \$s2, Label
- greater than bgt \$s1, \$s2, Label
- great than or equal to bge \$s1, \$s2, Label
- Such branches are included in the instruction set as pseudo instructions - recognized (and expanded) by the assembler
 - Its why the assembler needs a reserved register (\$at)

<u>Other Control Flow Instructions</u>

MIPS also has an unconditional branch instruction or jump instruction:

□ Instruction Format (J Format):

| 0x02 26-bit address |
|---------------------|
|---------------------|

from the low order 26 bits of the jump instruction



Six Steps in Execution of a Procecure

- 1. Main routine (caller) places parameters in a place where the procedure (callee) can access them
 - \$a0 \$a3: four argument registers
- 2. Caller transfers control to the callee
- 3. Callee acquires the storage resources needed
- 4. Callee performs the desired task
- 5. Callee places the result value in a place where the caller can access it
 - \$v0 \$v1: two value registers for result values
- 6. Callee returns control to the caller
 - \$ra: one return address register to return to the point of origin

<u>Aside</u>: <u>Spilling</u>Registers

What if the callee needs to use more registers than allocated to argument and return values?

callee uses a stack – a last-in-first-out queue



One of the general registers, \$sp (\$29), is used to address the stack (which "grows" from high address to low address)

add data onto the stack – push

\$sp = \$sp - 4
data on stack at new \$sp

remove data from the stack – pop

data from stack at \$sp \$sp = \$sp + 4

<u>Aside: Allocating Space on the stack</u>



The segment of the stack containing a procedure's saved registers and local variables is its procedure frame (aka activation record)

> The frame pointer (\$fp) points to the first word of the frame of a

procedure – providing a stable "base" register for the procedure

\$fp is initialized using \$sp on a call and \$sp is restored using \$fp on a return

MIPS Organization So Far



Number Representations

32-bit signed numbers (2's complement):



Converting <32-bit values into 32-bit values</p>

copy the most significant bit (the sign bit) into the "empty" bits 0010 -> 0000 0010 1010 -> 1111 1010

□ sign extend versus zero extend (lb vs. lbu)

MIPS Arithmetic Logic Unit (ALU)



With special handling for

- 🗆 sign extend addi, addiu, slti, sltiu
- zero extend andi, ori, xori
- overflow detection add, addi, sub
ARM & MIPS Similarities

ARM: the most popular embedded core
Similar basic set of instructions to MIPS

| | ARM | MIPS |
|-----------------------|------------------|------------------|
| Date announced | 1985 | 1985 |
| Instruction size | 32 bits | 32 bits |
| Address space | 32-bit flat | 32-bit flat |
| Data alignment | Aligned | Aligned |
| Data addressing modes | 9 | 3 |
| Registers | 15 × 32-bit | 31 × 32-bit |
| Input/output | Memory mapped | Memory mapped |



Compare and Branch in ARM

- Uses condition codes for result of an arithmetic/logical instruction
 - Negative, zero, carry, overflow
 - Compare instructions to set condition codes without keeping the result
- Each instruction can be conditional
 - Top 4 bits of instruction word: condition value
 - Can avoid branches over single instructions



Instruction Encoding

| | ARM | 31 28 Opx ⁴ | 27 Op ⁸ | 20 | 0 19 1 Rs1 ⁴ | 6 15 R | 12 11 d ⁴ | Opx ⁸ | 4 | 3 0 Rs2 ⁴ |
|-------------------|--------------------------|---------------------------|-----------------------------|------------|---|-----------|-------------------------|--------------------------|-------------------|-------------------------|
| Register-register | MIPS | 31 Op ⁶ | 26 25 Rs1 | 21 24 5 | 0 1 Rs2 ⁵ | 6 15 | 11 Rd⁵ | 10 Const ⁵ | 65 C | 0 Dpx ⁶ |
| Data transfer | ARM | 31 28 Opx ⁴ | 27 Op ⁸ | 21 2 | 0 19 1 Rs1 ⁴ | 6 15 R | 12 11 d ⁴ | Cor | nst ¹² | 0 |
| M | MIPS | Op ⁶ | Rs1 | 5 | Rd ⁵ | | | Const ¹⁶ | | |
| Propoh | ARM | 31 28 Opx ⁴ | 27 24 23 Op ⁴ | 3 | | | Const ²⁴ | | | 0 |
| Branch | MIPS | 31 Op ⁶ | 26 25 | 21 20 | 0 1 Opx ⁵ /Rs2 ⁵ | 6 15 | | Const ¹⁶ | | 0 |
| lumn/Call | ARM | 31 28 Opx ⁴ | 27 24 23 Op ⁴ | 3 | | | Const ²⁴ | | | 0 |
| ounp/oun | MIPS | 31 Op ⁶ | 26 25 | | | Co | onst ²⁶ | | | 0 |
| | Opcode Register Constant | | | | | | | | | |



ARM v8 Instructions

- In moving to 64-bit, ARM did a complete overhaul
- ARM v8 resembles MIPS
 - Changes from v7:
 - No conditional execution field
 - Immediate field is 12-bit constant
 - Dropped load/store multiple
 - PC is no longer a GPR
 - GPR set expanded to 32
 - Addressing modes work for all word sizes
 - Divide instruction
 - Branch if equal/branch if not equal instructions



Fallacies

- Powerful instruction \Rightarrow higher performance
 - Fewer instructions required
 - But complex instructions are hard to implement
 - May slow down all instructions, including simple ones
 - Compilers are good at making fast code from simple instructions
- Use assembly code for high performance
 - But modern compilers are better at dealing with modern processors
 - More lines of code ⇒ more errors and less productivity



Concluding Remarks

- Measure MIPS instruction executions in benchmark programs
 - Consider making the common case fast
 - Consider compromises

| Instruction class | MIPS examples | SPEC2006 Int | SPEC2006 FP |
|-------------------|--------------------------------------|--------------|-------------|
| Arithmetic | add, sub, addi | 16% | 48% |
| Data transfer | lw, sw, lb, lbu, lh, lhu, sb, lui | 35% | 36% |
| Logical | and, or, nor, andi, ori, sll, srl | 12% | 4% |
| Cond. Branch | beq, bne, slt, slti, sltiu | 34% | 8% |
| Jump | j, jr, jal | 2% | 0% |



Arithmetic for Multimedia

- Graphics and media processing operates on vectors of 8-bit and 16-bit data
 - Use 64-bit adder, with partitioned carry chain
 - Operate on 8×8-bit, 4×16-bit, or 2×32-bit vectors
 - SIMD (single-instruction, multiple-data)
- Saturating operations
 - On overflow, result is largest representable value
 - c.f. 2s-complement modulo arithmetic
 - E.g., clipping in audio, saturation in video



Arithmetic for Multimedia

- Graphics and media processing operates on vectors of 8-bit and 16-bit data
 - Use 64-bit adder, with partitioned carry chain
 - Operate on 8×8-bit, 4×16-bit, or 2×32-bit vectors
 - SIMD (single-instruction, multiple-data)
- Saturating operations
 - On overflow, result is largest representable value
 - c.f. 2s-complement modulo arithmetic
 - E.g., clipping in audio, saturation in video

