

Chapter 2A

Instructions: Language of the Computer

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

MIPS Design Principles

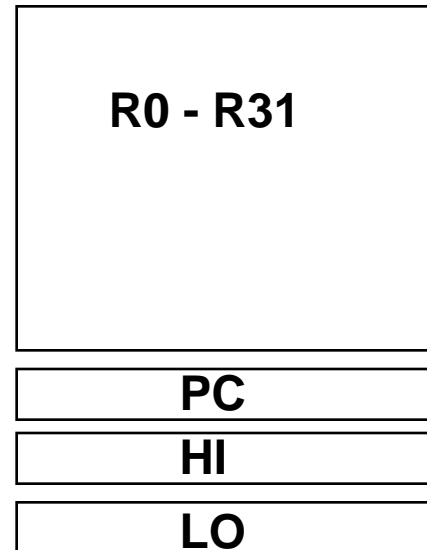
- **Simplicity favors regularity**
 - fixed size instructions
 - small number of instruction formats
 - opcode always the first 6 bits
- **Smaller is faster**
 - limited instruction set
 - limited number of registers in register file
 - limited number of addressing modes
- **Make the common case fast**
 - arithmetic operands from the register file (load-store machine)
 - allow instructions to contain immediate operands
- **Good design demands good compromises**
 - three instruction formats

MIPS-32 ISA

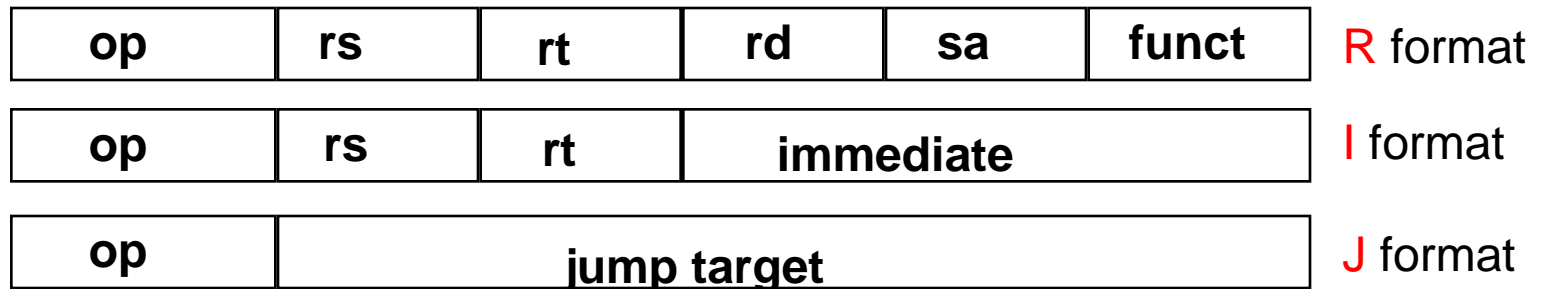
■ Instruction Categories

- Computational
- Load/Store
- Jump and Branch
- Floating Point: coprocessor
- Memory Management
- Special

Registers



3 Instruction Formats: all 32 bits wide



The MIPS Instruction Set



- Used as the example throughout the book
- Stanford MIPS commercialized by MIPS Technologies (www.mips.com)
- **Large (?)** share of embedded core market
 - Applications in consumer electronics, network/storage equipment, cameras, printers, ...
- See MIPS Reference Data tear-out card,
- Read **Appendix B: Assemblers, linkers, and the SPIM Simulator**

Arithmetic Operations

- Add and subtract, three operands
 - Two sources and one destination

add a, b, c $\#$ a = b + c
- All arithmetic operations have this form
- *Design Principle 1: Simplicity favors regularity*
 - Regularity makes implementation simpler
 - Simplicity enables higher performance at lower cost

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 (memory) is faster*
 - c.f. main memory: millions of locations

Arithmetic Example

- C code:

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

- f, i, j, g, h: assigned to \$s0, \$s1, \$s2, \$3, \$s4
- Compiled MIPS code:

```
add $t0, $s1, $s2 # temp t0 = g + h
add $t1, $s3, $s4 # temp t1 = i + j
sub $s0, $t0, $t1 # f = t0 - t1
```


MIPS Register File

- Holds thirty-two 32-bit registers

- Two **read ports** and
- One **write port**

- Registers are

- ┆ Faster than main memory

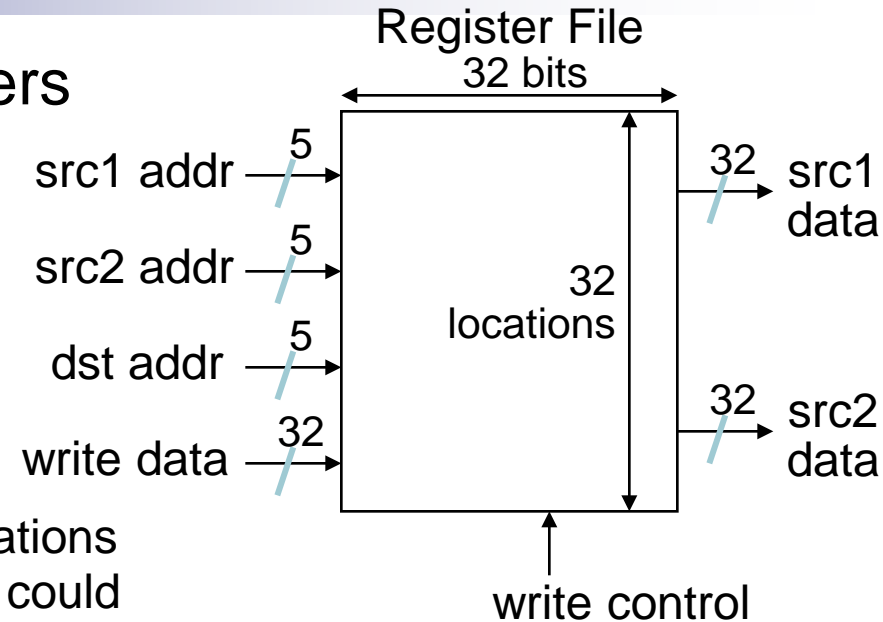
- But register files with more locations are slower (e.g., a 64 word file could be as much as 50% slower than a 32 word file)
- **Read/write port increase impacts speed quadratically**

- ┆ Easier for a compiler to use

- e.g., $(A*B) - (C*D) - (E*F)$ can do multiplies in any order

- ┆ Can hold variables so that

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



MIPS Assembly language

MIPS operands

Name	Example	Comments
32 registers	<code>\$s0-\$s7, \$t0-\$t9, \$zero, \$a0-\$a3, \$v0-\$v1, \$gp, \$fp, \$sp, \$ra, \$at</code>	Fast locations for data. In MIPS, data must be in registers to perform arithmetic, register <code>\$zero</code> always equals 0 and register <code>\$at</code> is reserved by the assembler to handle large constants.
2 ³⁰ memory words	<code>Memory[0], Memory[4], Memory[4294967292]</code>	Accessed only by data transfer instructions. MIPS uses byte addresses so sequential word addresses differ by 4. Memory holds data structures, arrays, and spilled registers.

MIPS assembly language

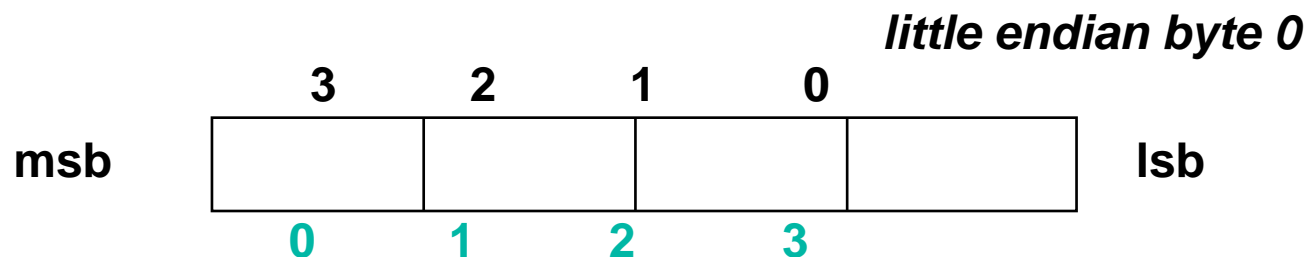
Category	Instruction	Example	Meaning	Comments
Arithmetic	<code>add</code>	<code>add \$s1,\$s2,\$s3</code>	$\$s1 = \$s2 + \$s3$	Three register operands
	<code>subtract</code>	<code>sub \$s1,\$s2,\$s3</code>	$\$s1 = \$s2 - \$s3$	Three register operands
	<code>add immediate</code>	<code>addi \$s1,\$s2,20</code>	$\$s1 = \$s2 + 20$	Used to add constants
Data transfer	<code>load word</code>	<code>lw \$s1,20(\$s2)</code>	$\$s1 = \text{Memory}[\$s2 + 20]$	Word from memory to register
	<code>store word</code>	<code>sw \$s1,20(\$s2)</code>	$\text{Memory}[\$s2 + 20] = \$s1$	Word from register to memory
	<code>load half</code>	<code>lh \$s1,20(\$s2)</code>	$\$s1 = \text{Memory}[\$s2 + 20]$	Halfword memory to register
	<code>load half unsigned</code>	<code>lhu \$s1,20(\$s2)</code>	$\$s1 = \text{Memory}[\$s2 + 20]$	Halfword memory to register
	<code>store half</code>	<code>sh \$s1,20(\$s2)</code>	$\text{Memory}[\$s2 + 20] = \$s1$	Halfword register to memory
	<code>load byte</code>	<code>lb \$s1,20(\$s2)</code>	$\$s1 = \text{Memory}[\$s2 + 20]$	Byte from memory to register
	<code>load byte unsigned</code>	<code>lbu \$s1,20(\$s2)</code>	$\$s1 = \text{Memory}[\$s2 + 20]$	Byte from memory to register
	<code>store byte</code>	<code>sb \$s1,20(\$s2)</code>	$\text{Memory}[\$s2 + 20] = \$s1$	Byte from register to memory
	<code>load linked word</code>	<code>ll \$s1,20(\$s2)</code>	$\$s1 = \text{Memory}[\$s2 + 20]$	Load word as 1st half of atomic swap
	<code>store condition. word</code>	<code>sc \$s1,20(\$s2)</code>	$\text{Memory}[\$s2+20]=\$s1; \$s1=0 \text{ or } 1$	Store word as 2nd half of atomic swap
<code>load upper immed.</code>	<code>lui \$s1,20</code>	$\$s1 = 20 * 2^{16}$	Loads constant in upper 16 bits	
Logical	<code>and</code>	<code>and \$s1,\$s2,\$s3</code>	$\$s1 = \$s2 \& \$s3$	Three reg. operands; bit-by-bit AND
	<code>or</code>	<code>or \$s1,\$s2,\$s3</code>	$\$s1 = \$s2 \$s3$	Three reg. operands; bit-by-bit OR
	<code>nor</code>	<code>nor \$s1,\$s2,\$s3</code>	$\$s1 = \sim(\$s2 \$s3)$	Three reg. operands; bit-by-bit NOR
	<code>and immediate</code>	<code>andi \$s1,\$s2,20</code>	$\$s1 = \$s2 \& 20$	Bit-by-bit AND reg with constant
	<code>or immediate</code>	<code>ori \$s1,\$s2,20</code>	$\$s1 = \$s2 20$	Bit-by-bit OR reg with constant
	<code>shift left logical</code>	<code>sll \$s1,\$s2,10</code>	$\$s1 = \$s2 \ll 10$	Shift left by constant
	<code>shift right logical</code>	<code>srl \$s1,\$s2,10</code>	$\$s1 = \$s2 \gg 10$	Shift right by constant
Conditional branch	<code>branch on equal</code>	<code>beq \$s1,\$s2,25</code>	if ($\$s1 == \$s2$) go to PC + 4 + 100	Equal test; PC-relative branch
	<code>branch on not equal</code>	<code>bne \$s1,\$s2,25</code>	if ($\$s1 \neq \$s2$) go to PC + 4 + 100	Not equal test; PC-relative
	<code>set on less than</code>	<code>slt \$s1,\$s2,\$s3</code>	if ($\$s2 < \$s3$) $\$s1 = 1$; else $\$s1 = 0$	Compare less than; for beq, bne
	<code>set on less than unsigned</code>	<code>sltu \$s1,\$s2,\$s3</code>	if ($\$s2 < \$s3$) $\$s1 = 1$; else $\$s1 = 0$	Compare less than unsigned
	<code>set less than immediate</code>	<code>slti \$s1,\$s2,20</code>	if ($\$s2 < 20$) $\$s1 = 1$; else $\$s1 = 0$	Compare less than constant
	<code>set less than immediate unsigned</code>	<code>sltiu \$s1,\$s2,20</code>	if ($\$s2 < 20$) $\$s1 = 1$; else $\$s1 = 0$	Compare less than constant unsigned
Unconditional jump	<code>jump</code>	<code>j 2500</code>	go to 10000	Jump to target address
	<code>jump register</code>	<code>jr \$ra</code>	go to <code>\$ra</code>	For switch, procedure return
	<code>jump and link</code>	<code>jal 2500</code>	$\$ra = \text{PC} + 4$; go to 10000	For procedure call

Memory Operands

- Main memory used for composite data
 - Arrays, structures, dynamic data
- To apply arithmetic operations
 - **Load values** from memory into registers
 - Execute operations
 - **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

Byte Addressing

- 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)



big endian byte 0

Other way to show the endian

- As an example, suppose we have the hexadecimal number **12345678**.
- The big endian and little endian arrangements of the bytes are shown below.

	Word address			
Address →	00	01	10	11
Big Endian	12	34	56	78
Little Endian	78	56	34	12

Memory Operand Example 1

- C code:

```
g = h + A[8];
```

- g in \$s1, h in \$s2, base address of A in \$s3

- Compiled MIPS code:

- Index 8 requires offset of 32
 - 4 bytes per word

```
lw    $t0, 32($s3)    # load word A[8]
add   $s1, $s2, $t0   # g = h + A[8]
```

offset



base register



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 (8×4)

```
lw    $t0, 32($s3)    # load word A[8]
add   $t0, $s2, $t0
sw    $t0, 48($s3)    # store word A[12]
```

Registers vs. Memory

- Registers are faster to access than memory: **why?**
- Operating on memory data requires loads and stores: **in RISC**
 - 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 (**usage**) optimization is important!
 - graph coloring problem (NP-complete): no two adjacent nodes in the **dependency graph** have the same color.
 - dependency graph
 - nodes: variables, edges: dependencies

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 (hard-wired)
- Useful for common operations
 - E.g., move between registers
`add $t2, $s1, $zero`

Unsigned Binary Integers

- Given an n-bit number

$$x = x_{n-1}2^{n-1} + x_{n-2}2^{n-2} + \dots + x_12^1 + x_02^0$$

- Range: 0 to $+2^n - 1$

- Example

- $0000\ 0000\ 0000\ 0000\ 0000\ 0000\ 0000\ 1011_2$
 $= 0 + \dots + 1 \times 2^3 + 0 \times 2^2 + 1 \times 2^1 + 1 \times 2^0$
 $= 0 + \dots + 8 + 0 + 2 + 1 = 11_{10}$

- Using 32 bits

- 0 to $+4,294,967,295$

2's-Complement Signed Integers

- Given an n-bit number

$$x = -x_{n-1}2^{n-1} + x_{n-2}2^{n-2} + \dots + x_12^1 + x_02^0$$

- Range: -2^{n-1} to $+2^{n-1} - 1$

- Example

- $1111\ 1111\ 1111\ 1111\ 1111\ 1111\ 1111\ 1100_2$
 $= -1 \times 2^{31} + 1 \times 2^{30} + \dots + 1 \times 2^2 + 0 \times 2^1 + 0 \times 2^0$
 $= -2,147,483,648 + 2,147,483,644 = -4_{10}$

- Using 32 bits

- $-2,147,483,648$ to $+2,147,483,647$

2's-Complement Signed Integers

- **Bit 31 (left-most)** is sign bit
 - 1 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

Signed Negation

- Complement and add 1
 - Complement means $1 \rightarrow 0, 0 \rightarrow 1$

$$x + \bar{x} = 1111\dots111_2 = -1$$

$$\bar{x} + 1 = -x$$

- Example: negate +2
 - $+2 = 0000\ 0000 \dots 0010_2$
 - $-2 = 1111\ 1111 \dots 1101_2 + 1$
 $= 1111\ 1111 \dots 1110_2$

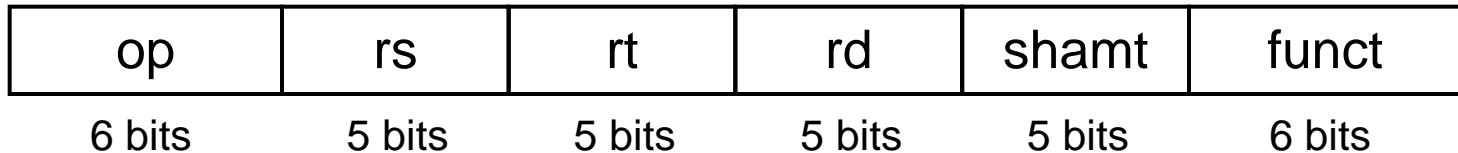
Sign Extension

- Representing a number using more bits
 - Preserve the numeric value
- In MIPS instruction set
 - `addi` : extend immediate value
 - `l b`, `l h`: 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
 - \$s0 – \$s7 are reg's 16 – 23
 - \$t8 – \$t9 are reg's 24 – 25

MIPS R-format Instructions



■ Instruction fields

- op: operation code (opcode)
- rs: first **s**ource register number
- rt: second source register number
- rd: **d**estination register number
- **shamt**: shift amount (00000 for now)
- **funct**: function code (**extends opcode**)

R-format Example

op	rs	rt	rd	shamt	funct
6 bits	5 bits	5 bits	5 bits	5 bits	6 bits

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

special	\$s1	\$s2	\$t0	0	add
---------	------	------	-------------	---	-----

0	17	18	8	0	32
---	----	----	---	---	----

000000	10001	10010	01000	00000	100000
--------	-------	-------	-------	-------	--------

$$00000010001100100100000000100000_2 = 02324020_{16}$$

Hexadecimal

- Base 16
 - Compact representation of bit strings
 - 4 bits per hex digit

0	0000	4	0100	8	1000	c	1100
1	0001	5	0101	9	1001	d	1101
2	0010	6	0110	a	1010	e	1110
3	0011	7	0111	b	1011	f	1111

- Example: eca8 6420
 - 1110 1100 1010 1000 0110 0100 0010 0000

MIPS I-format Instructions



- Immediate arithmetic and load/store instructions
 - rt: destination register (addi, lw)
source register (sw)
 - Constant or offset address: -2^{15} to $+2^{15} - 1$
 - 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

Memory Operand Example 2

- C code:

`A[300] = 200 + A[300];`

- `h` in `$s2`, base address of `A` in `$s3`

- Compiled MIPS code:

- Index 300 requires offset of 1200

```
lw    $t0, 1200($s3) #load word A[300]
```

```
addi  $t0, $t0, 200
```

```
sw    $t0, 1200($s3) #store word A[300]
```

I-format Examples

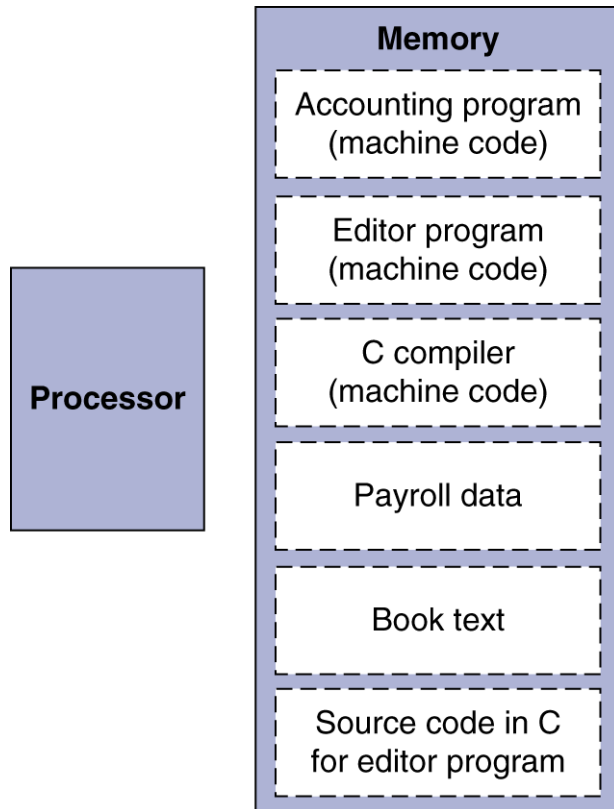
$A[300] = 200 + A[300];$

```
lw $t0, 1200($t1)
addi $t0, $t0, 200
sw $t0, 1200($t1)
```

op	rs	rt	constant or address
6 bits	5 bits	5 bits	16 bits
35	9	8	1200
8	8	8	200
43	9	8	1200

Stored Program Computers

The BIG Picture



- Instructions represented in binary, just like data
- Both instructions and data stored in memory
- Programs can operate on programs
 - e.g., compilers, linkers, ...
- Binary compatibility allows compiled programs to work on different computers
 - Standardized ISAs

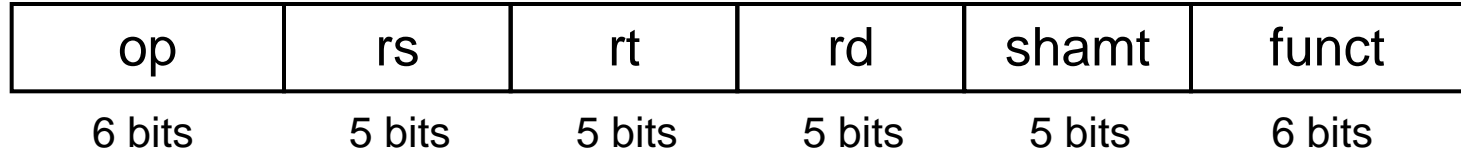
Logical Operations

- Instructions for bitwise manipulation

Operation	C	Java	MIPS
Shift left	<<	<<	sll
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



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

AND Operations

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

\$t2	0000 0000 0000 0000 0000 1101 1100 0000
\$t1	0000 0000 0000 0000 0011 1100 0000 0000
\$t0	0000 0000 0000 0000 0000 1100 0000 0000

OR Operations

- Useful to include bits in a word
 - Set some bits to 1, leave others unchanged
- or \$t0, \$t1, \$t2

\$t2 0000 0000 0000 0000 0000 1101 1100 0000

\$t1 0000 0000 0000 0000 0011 1100 0000 0000

\$t0 0000 0000 0000 0000 0011 1101 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
 - $a \text{ NOR } b == \text{NOT} (a \text{ OR } b)$

`nor $t0, $t1, $zero` ←

Register 0: always read as zero

\$t1 0000 0000 0000 0000 0011 1100 0000 0000

\$t0 1111 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 # I-type**
 - if (rs == rt) branch to instruction labeled L1;
- **bne rs, rt, L1 # I-type**
 - if (rs != rt) branch to instruction labeled L1;
- **j L1 # J-type**
 - unconditional jump to instruction labeled L1

Compiling If Statements

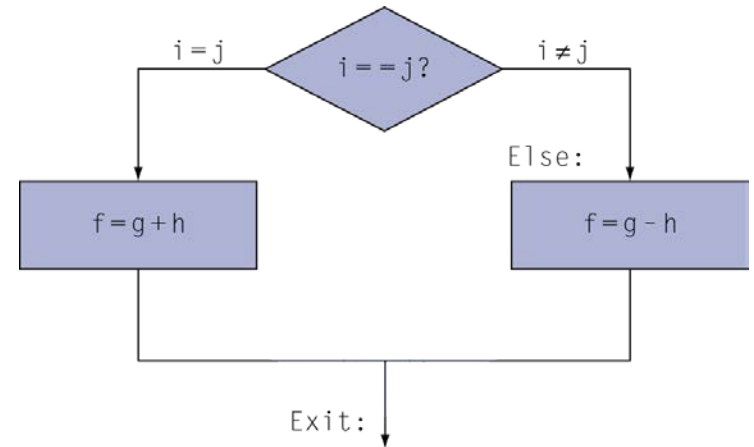
- C code:

```
if (i == j) f = g+h;
else f = g-h;
```

- f, g, h: in \$s0, \$s1, \$s2

- Compiled MIPS code:

```
                bne $s3, $s4, Else
                add $s0, $s1, $s2
                j   Exit
Else:           sub $s0, $s1, $s2
Exit:          ...
```



Assembler calculates addresses

Compiling Loop Statements

- C code:

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

- i in \$s3, k in \$s5, address of save in \$s6

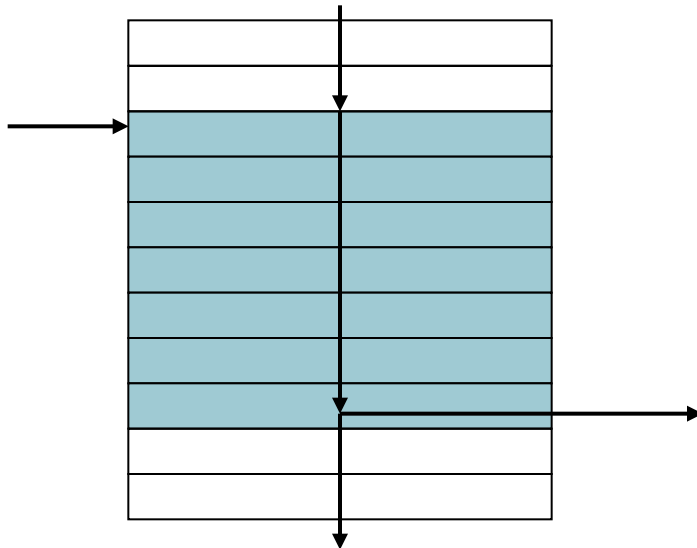
- Compiled MIPS code:

```
Loop:  sll    $t1, $s3, 2    # offset=4*i
        add   $t1, $t1, $s6 # base+offset
        lw    $t0, 0($t1)   # save[i]
        bne   $t0, $s5, Exit
        addi  $s3, $s3, 1    # i +=1
        j     Loop
```

```
Exit:  ...
```

Basic Blocks

- A **basic block** is a sequence of instructions with
 - No embedded branches (except at end)
 - No branch targets (except at beginning)



- A compiler identifies basic blocks for **optimization**
- An advanced processor can **accelerate** execution of basic blocks

Set on less than

- Set result to 1 if a condition is true
 - Otherwise, set to 0

0	rs	rt	rd	0	0x2a
---	----	----	----	---	------

- `slt rd, rs, rt`
 - if ($rs < rt$) $rd = 1$; else $rd = 0$;
- `slti rt, rs, constant`
 - if ($rs < constant$) $rt = 1$; else $rt = 0$;
- Use in combination with `beq`, `bne`

```
slt $t0, $s1, $s2 # if ($s1 < $s2)
bne $t0, $zero, L # branch to L
```

Branch Instruction Design

- MIPS compilers use `slt`, `slti`, `beq`, `bne` and the fixed value of 0 to create all relative conditions.
- Why not `blt`, `bge`, etc?
- Hardware for $<$, \geq , ... 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: `sl t, sl ti`
- Unsigned comparison: `sl tu, sl tui`
- Example
 - `$s0 = 1111 1111 1111 1111 1111 1111 1111 1111`
 - `$s1 = 0000 0000 0000 0000 0000 0000 0000 0001`
 - `sl t $t0, $s0, $s1 # signed`
 - $-1 < +1 \Rightarrow \$t0 = 1$
 - `sl tu $t0, $s0, $s1 # unsigned`
 - $+4,294,967,295 > +1 \Rightarrow \$t0 = 0$

Note

- More decision statements make code easier to read and understand
- Fewer decision statements simplify the task of the underlying layer that is responsible for execution
- More decision statements mean fewer lines of code, which generally reduces coding time and results in the execution of fewer operations.

Procedure

- A stored subroutine that performs a specific task based on the parameters with which it is provided.
- Make them easier to understand
- Allow code to be reused
- You can think of a procedure like a spy
- A spy operates on only a “need to know” basis, so the spy can’t make assumption about his employer

Six Steps in Procedure Calling

1. Main routine (**caller**) places parameters in a place where the procedure (**callee**) can access them.

\$a0 - \$a3: four **argument** registers

2. Transfers control to the procedure.

3. Acquires the storage resources needed.

4. Performs the desired task.

5. Places the result value in a place where the **caller** can access it.

\$v0 - \$v1: two **value** registers for result value

6. Return control to the **caller**.

\$ra: one **return address** register

MIPS Register Convention

Name	Register Number	Usage	Preserve on call by callee?
\$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
\$ra	31	return address	yes

MIPS registers and usage convention

Register name	Number	Usage
\$zero	0	constant 0
\$at	1	reserved for assembler
\$v0	2	expression evaluation and results of a function
\$v1	3	expression evaluation and results of a function
\$a0	4	argument 1
\$a1	5	argument 2
\$a2	6	argument 3
\$a3	7	argument 4
\$t0	8	temporary (not preserved across call)
\$t1	9	temporary (not preserved across call)
\$t2	10	temporary (not preserved across call)
\$t3	11	temporary (not preserved across call)
\$t4	12	temporary (not preserved across call)
\$t5	13	temporary (not preserved across call)
\$t6	14	temporary (not preserved across call)
\$t7	15	temporary (not preserved across call)
\$s0	16	saved temporary (preserved across call)
\$s1	17	saved temporary (preserved across call)
\$s2	18	saved temporary (preserved across call)
\$s3	19	saved temporary (preserved across call)
\$s4	20	saved temporary (preserved across call)
\$s5	21	saved temporary (preserved across call)
\$s6	22	saved temporary (preserved across call)
\$s7	23	saved temporary (preserved across call)
\$t8	24	temporary (not preserved across call)
\$t9	25	temporary (not preserved across call)
\$k0	26	reserved for OS kernel
\$k1	27	reserved for OS kernel
\$gp	28	pointer to global area
\$sp	29	stack pointer
\$fp	30	frame pointer
\$ra	31	return address (used by function call)

Caller, Callee

Caller-**saved** registers:

\$t0 - \$t7: temporary

\$t8 - \$t9: temporary

Callee-**saved** registers:

\$s0 - \$s7 : long-lived

Procedure call frame

additional arguments

callee-**saved** registers

\$sp, \$fp, \$ra

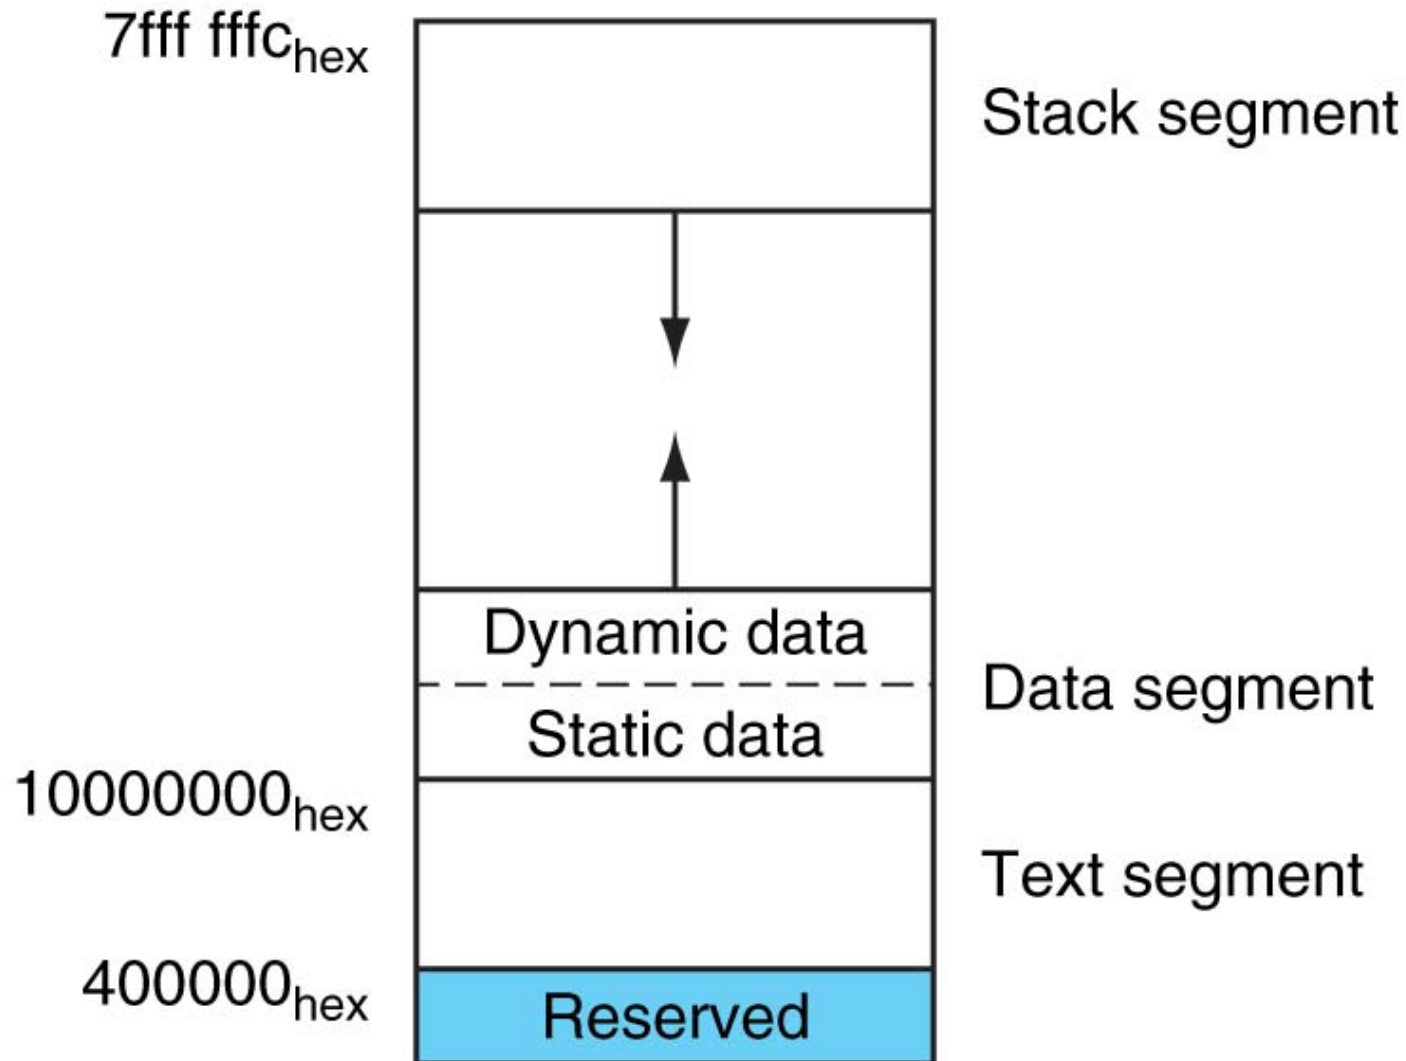
local variables

Stack: for spilling registers

last-in first-out (LIFO)

push, pop

Layout of Memory



Procedure Call Instructions

- Procedure call: jump and link

jal ProcedureLabel

- Address of following instruction (**PC+4**) 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:

leaf_example:

```
addi    $sp, $sp, -12    # push stack
sw      $t1, 8($sp)     # can drop greens
sw      $t0, 4($sp)
sw      $s0, 0($sp)
add     $t0, $a0, $a1    #body
add     $t1, $a2, $a3
sub     $s0, $t0, $t1
add     $v0, $s0, $zero # return f; move
lw      $s0, 0($sp)     # restore variables
lw      $t0, 4($sp)
lw      $t1, 8($sp)
addi    $sp, $sp, 12   # pop stack
jr      $ra             # jump to caller
```

Stack push and pop

Stack pointer : empty or **full** ?

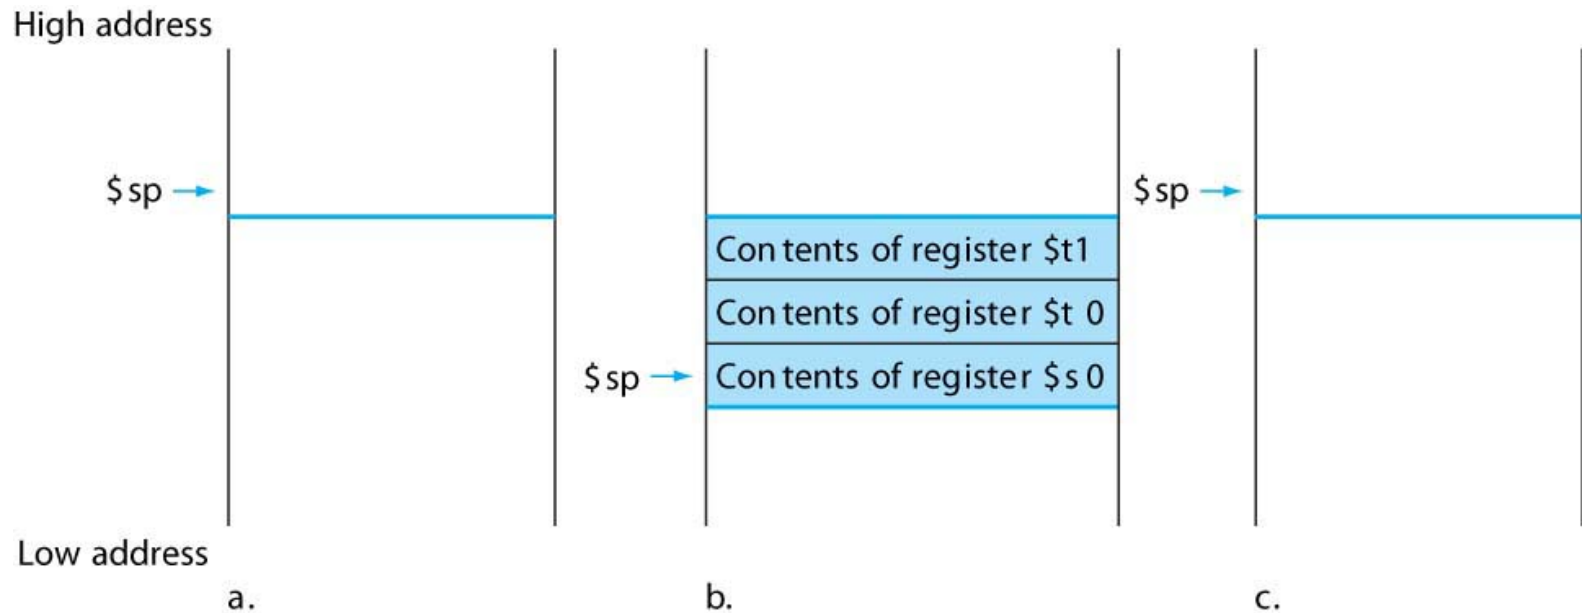


FIGURE 2.10 The values of the stack pointer and the stack (a) before, (b) during, and (c) after the procedure call. The stack pointer always points to the “top” of the stack, or the last word in the stack in this drawing.

Leaf Procedure Example

- MIPS code:

leaf_example:			
addi	\$sp,	\$sp,	- 4
sw	\$s0,	0(\$sp)	
add	\$t0,	\$a0,	\$a1
add	\$t1,	\$a2,	\$a3
sub	\$s0,	\$t0,	\$t1
add	\$v0,	\$s0,	\$zero
lw	\$s0,	0(\$sp)	
addi	\$sp,	\$sp,	4
jr	\$ra		

Save \$s0 on stack to use it

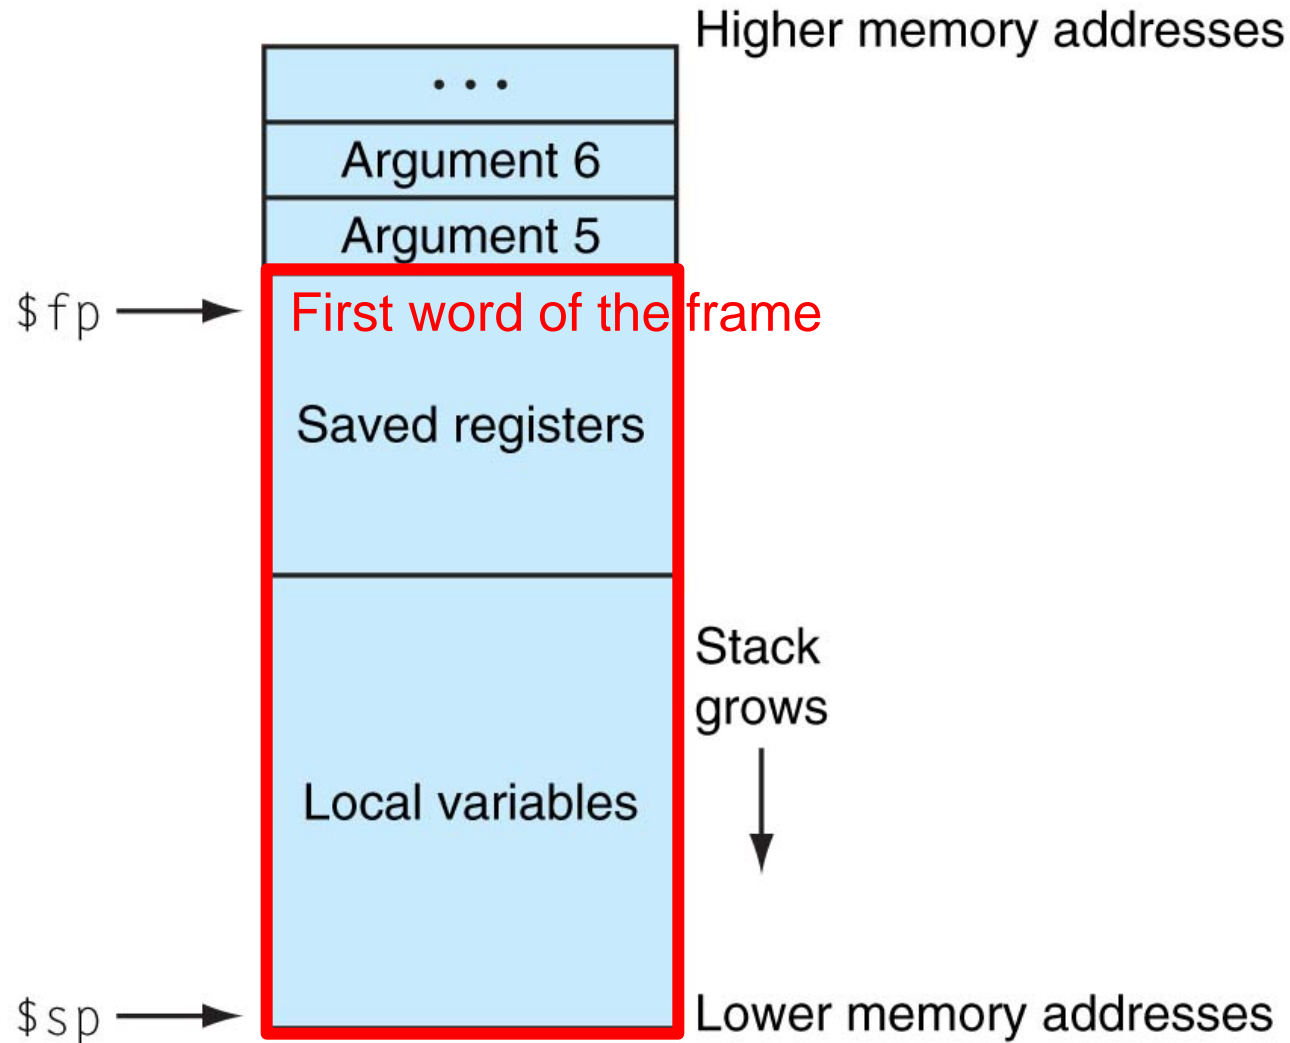
Procedure body

Move result to \$v0

Restore \$s0

Return

Layout of a stack memory



Non-Leaf (Nested) 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
- Stack operations: push & pop

Non-Leaf Procedure Example

- C code:

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

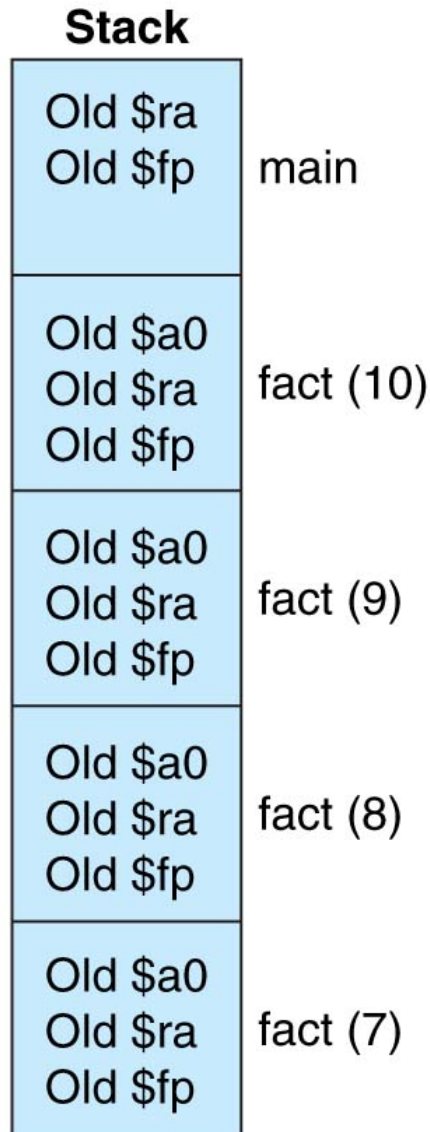
- Argument n in \$a0
- Result in \$v0

Non-Leaf Procedure Example

- MIPS code:

fact:		
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	# if n >= 1, go to L1
beq	\$t0, \$zero, L1	
addi	\$v0, \$zero, 1	# if n=0, 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
lw	\$a0, 0(\$sp)	# restore original n
lw	\$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

Stack during the call of fact(7)



Home work: Read Appendix B-27 ~B-30 (factorial example)

A stack frame : built upon the MIPS convention

1. Minimum stack frame: 24 bytes
4 arguments + return address
= 5 → 6

Stack pointer should be doubleword aligned

2. Main needs to save \$fp : -> 8

Stack grows