

MIPS Instruction Set Architecture (2)

Jinkyu Jeong (jinkyu@skku.edu)

Computer Systems Laboratory

Sungkyunkwan University

<http://csl.skku.edu>

Outline

Textbook: P&H 2.5-2.12 (Except 2.11) – for 4th Ed.

P&H 2.5-2.11 (Except 2.10) – for 5th Ed.

- Representing Instructions
- Logical Operations
- Instructions for Making Decisions
- Supporting Procedures in Computer Hardware
- Communicating with People
- MIPS Addressing for 32-Bit Immediates and Addresses
- Translating and Starting a Program
- Array vs. Pointers – Will not cover (2.14 in 4th Ed.)

The screenshot shows the PCSpim MIPS simulator interface. The main window displays assembly code for a 'Text Segment'. The code consists of MIPS instructions with their corresponding hexadecimal addresses and offsets. The instructions include load word, add immediate, add immediate, shift left logical, add immediate, jump and link, ori, syscall, ori, lui, syscall, ori, lui, ori, syscall, ori, lui, ori, syscall, ori, lui, and syscall. The simulator status bar at the bottom shows 'For Help, press F1' and 'Bare=0; Pseudo=1; Mapped=1; LoadTrap= NUM'.

```
PCSpim
File Simulator Window Help
[Icons]
Text Segment
[0x00400000] 0x8fa40000 lw $4, 0($29) ; 102: lw
[0x00400004] 0x27a50004 addiu $5, $29, 4 ; 103: addi
[0x00400008] 0x24a60004 addiu $6, $5, 4 ; 104: addi
[0x0040000c] 0x00041080 sll $2, $4, 2 ; 105: sll
[0x00400010] 0x00c23021 addu $6, $6, $2 ; 106: addu
[0x00400014] 0x0c100008 jal 0x00400020 [main] ; 107: jal
[0x00400018] 0x3402000a ori $2, $0, 10 ; 108: ori
[0x0040001c] 0x0000000c syscall ; 109: syscall
[0x00400020] 0x34020004 ori $2, $0, 4 ; 10: ori
[0x00400024] 0x3c041001 lui $4, 4097 [prompt] ; 11: lui
[0x00400028] 0x0000000c syscall ; 12: syscall
[0x0040002c] 0x34020008 ori $2, $0, 8 ; 13: ori
[0x00400030] 0x3c011001 lui $1, 4097 [buffer] ; 14: lui
[0x00400034] 0x34240010 ori $4, $1, 16 [buffer] ; 15: ori
[0x00400038] 0x3405004e ori $5, $0, 78 ; 15: ori
[0x0040003c] 0x0000000c syscall ; 16: syscall
[0x00400040] 0x34020004 ori $2, $0, 4 ; 17: ori
[0x00400044] 0x3c011001 lui $1, 4097 [greet] ; 18: lui
[0x00400048] 0x34240060 ori $4, $1, 96 [greet] ; 19: ori
[0x0040004c] 0x0000000c syscall ; 19: syscall
[0x00400064] 0x0000000c syscall ; 24: syscall
KERNEL
[0x80000080] 0x0001d821 addu $27, $0, $1 ; 57: addu
For Help, press F1 Bare=0; Pseudo=1; Mapped=1; LoadTrap= NUM
```

Representing Instructions

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



- **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

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

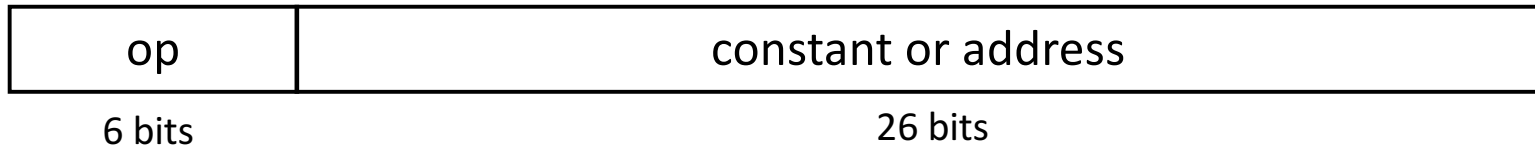
$$00000010\ 00110010\ 01000000\ 00100000_2 = 02324020_{16}$$

MIPS I-format Instructions



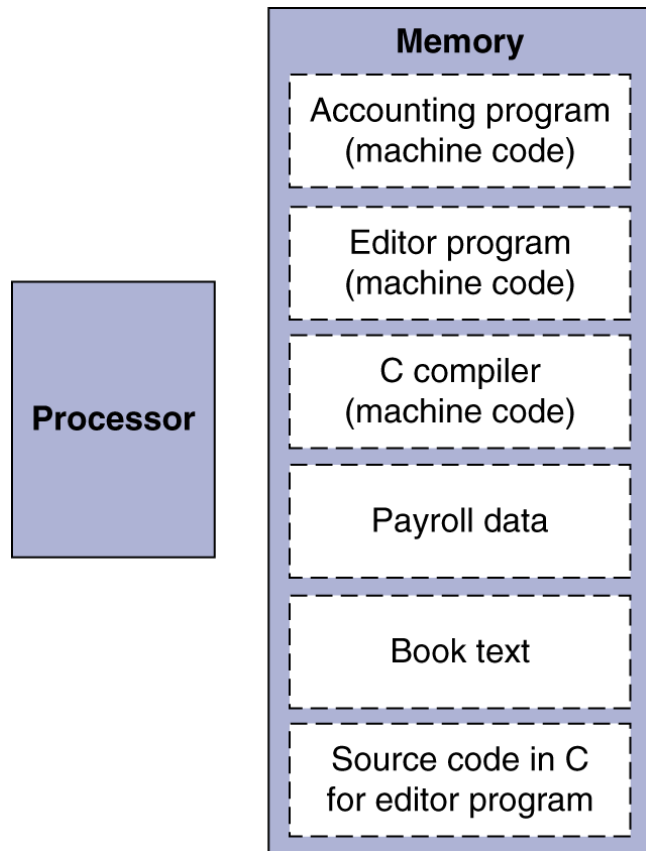
- Immediate arithmetic and load/store instructions
 - rt: destination or source register number
 - Constant: -2^{15} to $+2^{15} - 1$
 - Address: offset added to base address in rs
- *Design Principle 4: Good design demands good compromises*
 - Different formats complicate decoding, but allow 32-bit instructions uniformly
 - Keep formats as similar as possible

MIPS J-format Instructions



- Jump instructions (`j` and `jal`)
 - Address: encode 26-bit target address

Stored Program Computers



- Instructions represented in binary, just like data
- 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

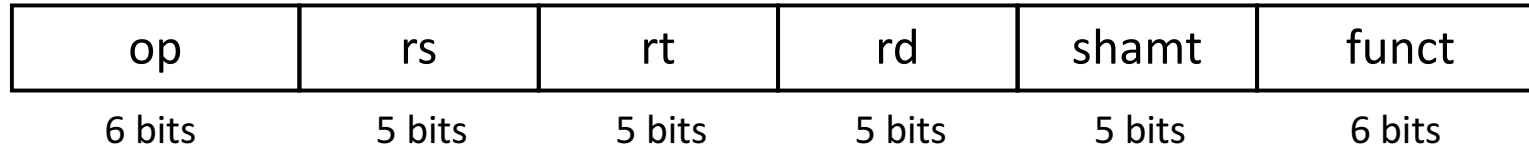
Logical Operations

- 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

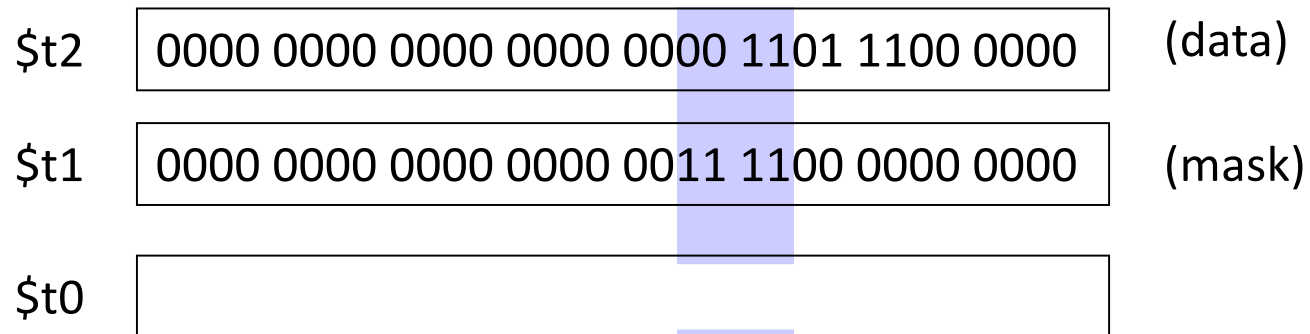


- **shamt**: how many positions to shift
- **Shift left logical**
 - Shift left and fill with 0 bits
 - sll by i bits multiplies by 2^i
- **Shift right logical**
 - Shift right and fill with 0 bits
 - srl by i bits divides by 2^i (unsigned only)
 - cf. sra : shift right arithmetic

AND Operations

- Useful to mask bits in a word
 - Select some bits, clear others to 0

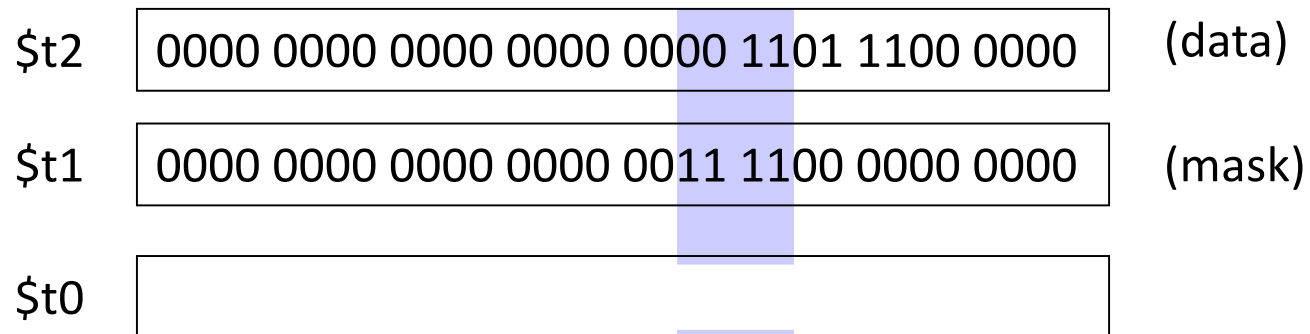
and \$t0, \$t1, \$t2



OR Operations

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

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



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
```

Instructions for Making Decisions

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

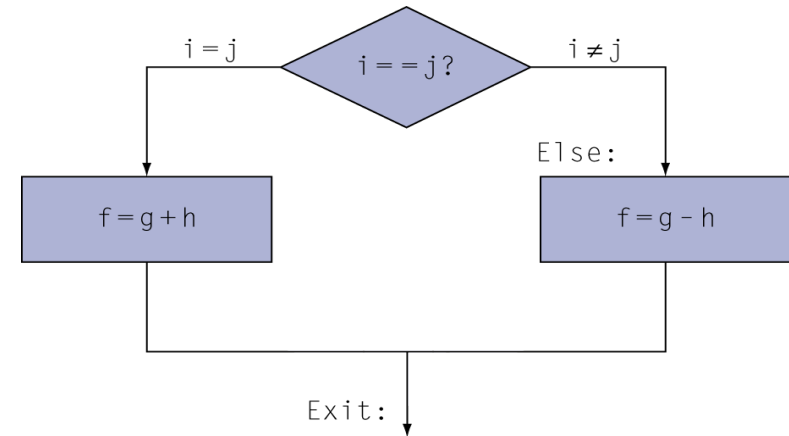
Compiling If Statements

- C code:

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

– f, g, ... in \$s0, \$s1, ...

- Compiled MIPS code:



Assembler calculates addresses

Compiling Loop Statements

- C code:

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

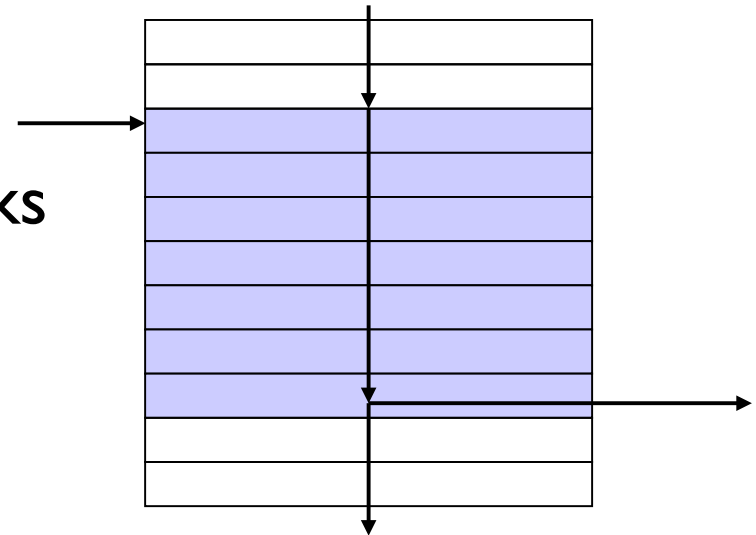
– i in \$s3, k in \$s5, address of save in \$s6

- Compiled MIPS code:

```
Loop:  sll    $t1, $s3, 2
       add   $t1, $t1, $s6
       lw    $t0, 0($t1)
       bne   $t0, $s5, Exit
       addi  $s3, $s3, 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



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 $rd = 0$;
- `slti rt, rs, constant`
 - if ($rs < \text{constant}$) $rt = 1$; 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

- 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: `slt`, `slti`
- Unsigned comparison: `sltu`, `sltiu`
- Example

`$s0 = 1111 1111 1111 1111 1111 1111 1111 1111`

`$s1 = 0000 0000 0000 0000 0000 0000 0000 0001`

– `slt $t0, $s0, $s1 # signed`

- $-1 < +1 \Rightarrow \$t0 = 1$

– `sltu $t0, $s0, $s1 # unsigned`

- $+4,294,967,295 > +1 \Rightarrow \$t0 = 0$

Supporting Procedures in Computer Hardware

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 (caller-save reg's)
 - Can be overwritten by callee
- **\$s0 – \$s7**: saved (callee-saved reg's)
 - 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 (Continued)

- MIPS code:

Leaf_example:	
addi \$sp, \$sp, -4 sw \$s0, 0(\$sp)	Save \$s0 on stack
add \$t0, \$a0, \$a1 add \$t1, \$a2, \$a3 sub \$s0, \$t0, \$t1	Procedure body
add \$v0, \$s0, \$zero	Result
lw \$s0, 0(\$sp) addi \$sp, \$sp, 4	Restore \$s0
jr \$ra	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

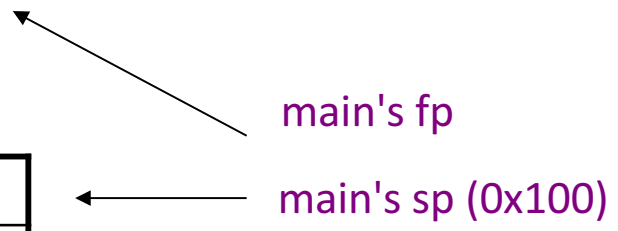
Non-Leaf Procedure Example

- MIPS code:

fact:		
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
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

fact() Animation:

0x100	
0x0FC	
0x0F8	
0x0F4	
0x0F0	
0x0EC	
0x0E8	
0x0E4	
0x0E0	
0x0DC	
0x0D8	
0x0D4	
0x0D0	
0x0CC	
0x0C8	
0x0C4	



fact() Animation:

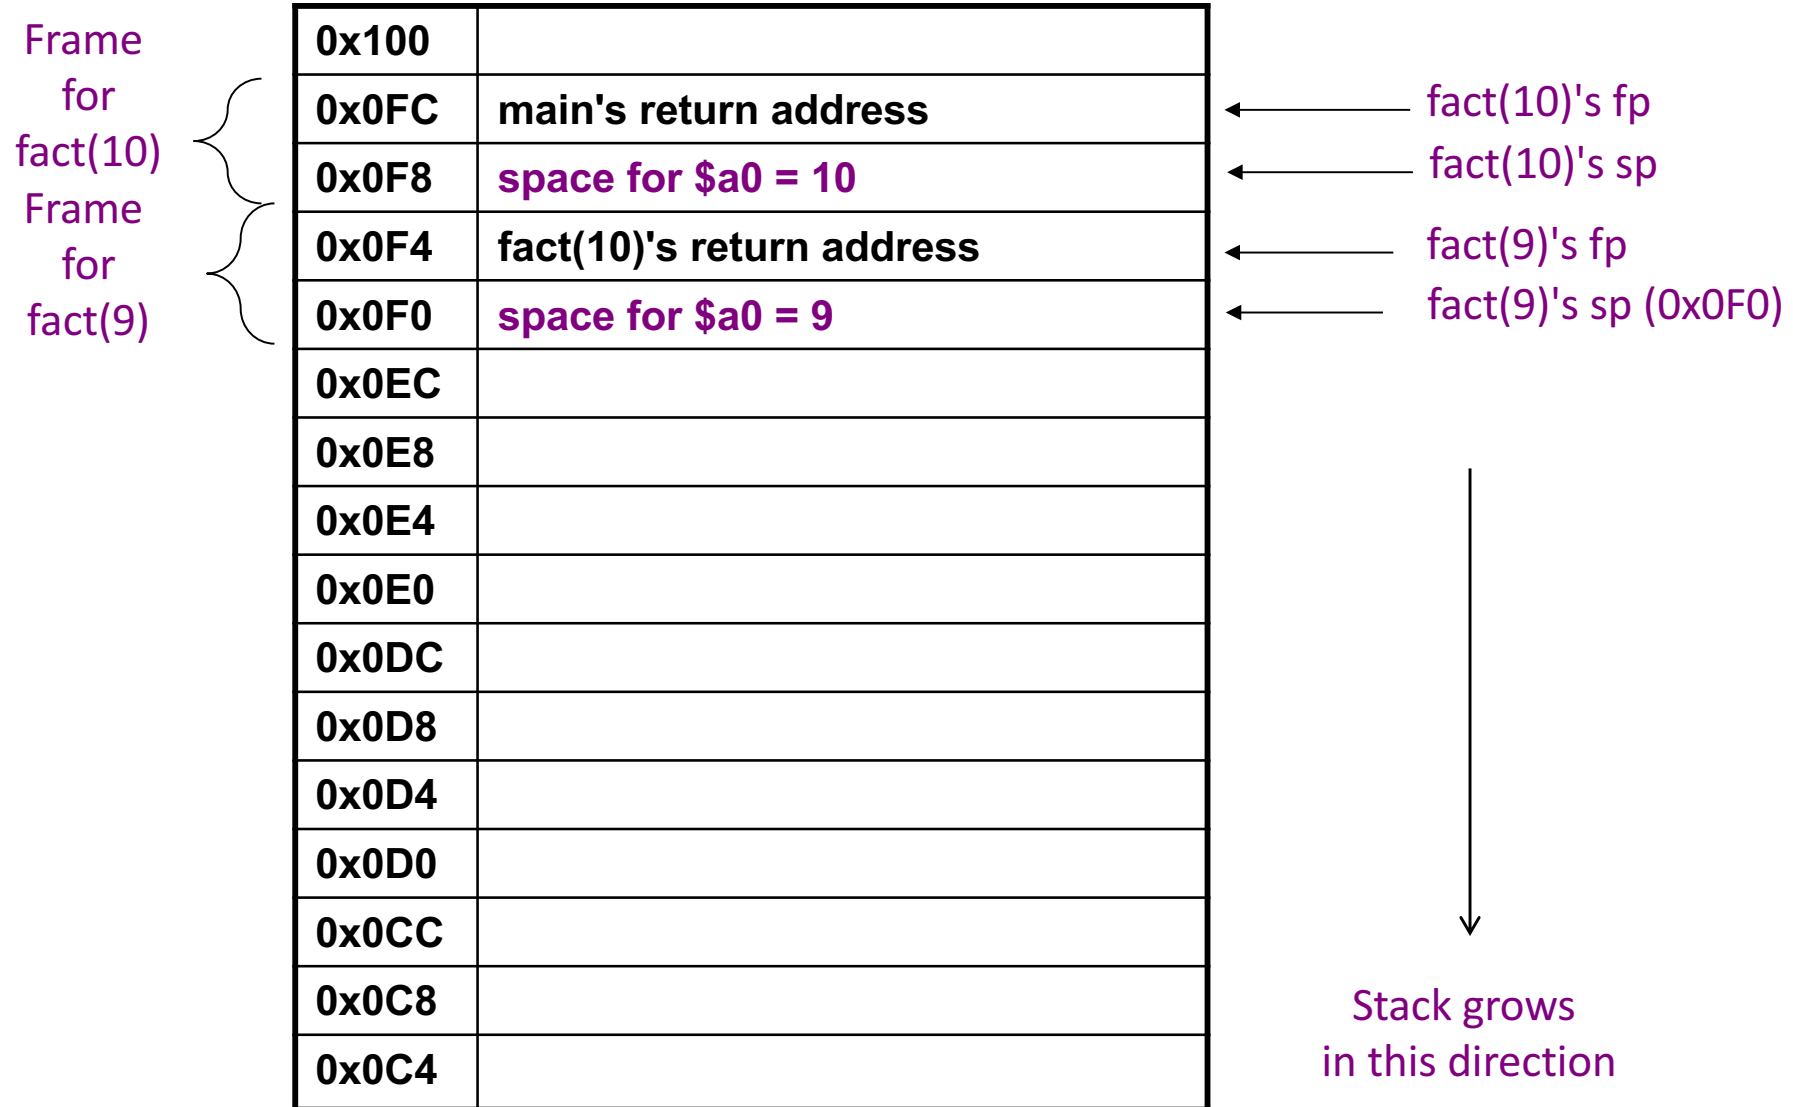
Frame
for
fact(10)



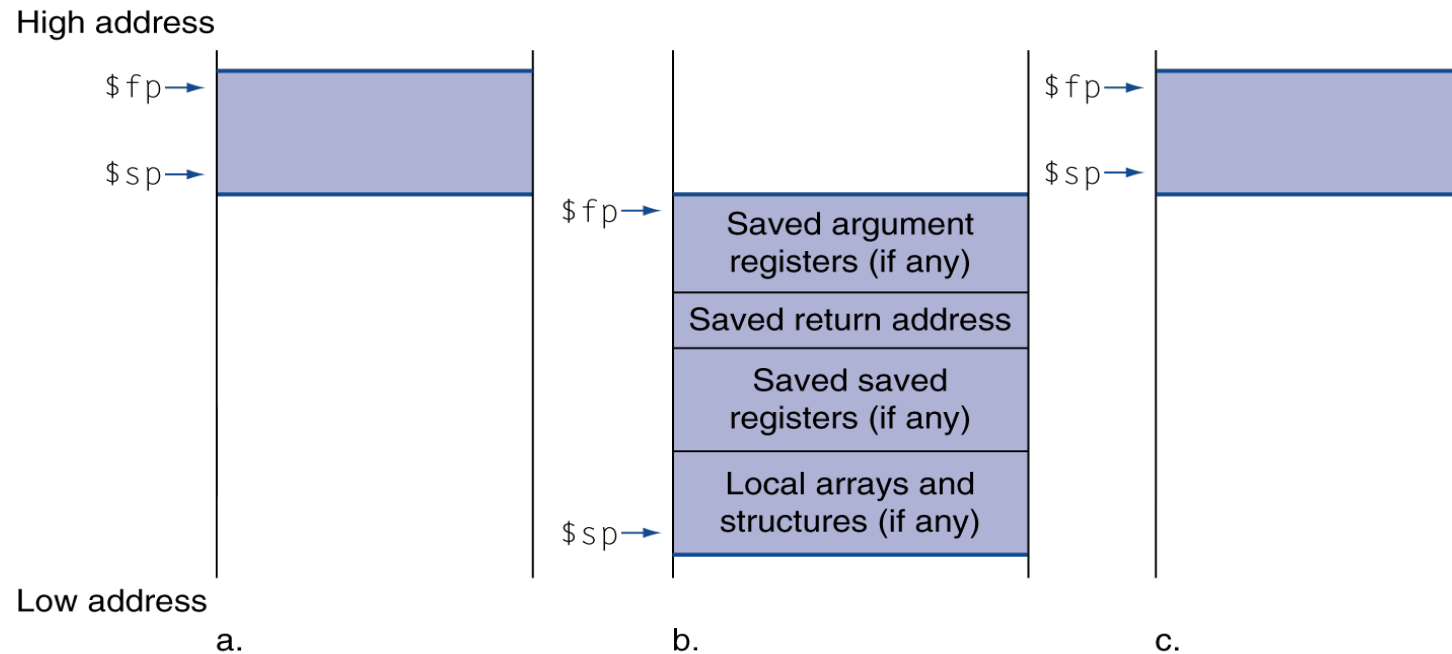
0x100	
0x0FC	main's return address
0x0F8	space for \$a0
0x0F4	
0x0F0	
0x0EC	
0x0E8	
0x0E4	
0x0E0	
0x0DC	
0x0D8	
0x0D4	
0x0D0	
0x0CC	
0x0C8	
0x0C4	

← fact(10)'s fp
← fact(10)'s sp

fact() Animation:



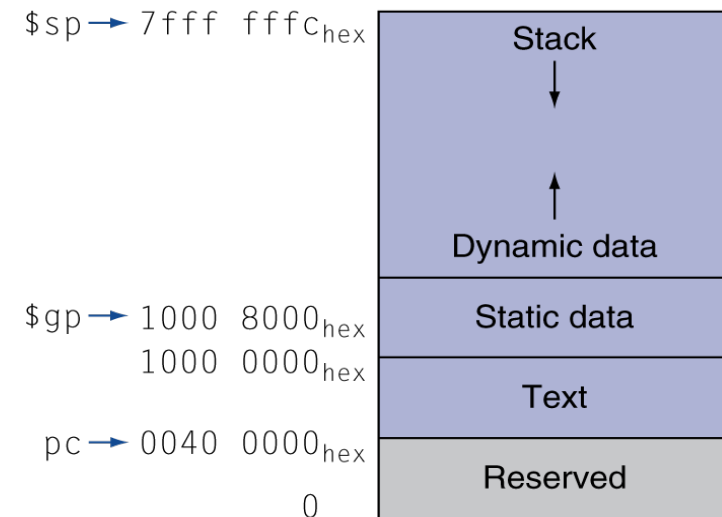
Local Data on the Stack



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

Memory Layout

- Text: program code
- Static data: global variables
 - e.g., static variables in C, constant arrays and strings
 - \$gp initialized to address allowing \pm offsets into this segment
- Dynamic data: heap
 - E.g., malloc in C, new in Java
- Stack: automatic storage



Communicating with People (Handling Characters)

Character Data

- **Byte-encoded character sets**
 - ASCII: 128 characters
 - 95 graphic, 33 control
 - Latin-1: 256 characters
 - ASCII, +96 more graphic characters
- **Unicode: 32-bit character set**
 - Used in Java, C++ wide characters, ...
 - Most of the world's alphabets, plus symbols
 - UTF-8, UTF-16: variable-length encodings

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

String Copy Example

- C code (naïve):

- Null-terminated string

```
void strcpy (char x[], char y[])  
{ int i;  
  i = 0;  
  while ((x[i]=y[i])!='\0')  
    i += 1;  
}
```

- Addresses of x, y in \$a0, \$a1
- i in \$s0

String Copy Example

- MIPS code:

strcpy:		
	addi \$sp, \$sp, -4	# adjust stack for 1 item
	sw \$s0, 0(\$sp)	# save \$s0
	add \$s0, \$zero, \$zero	# i = 0
L1:	add \$t1, \$s0, \$a1	# addr of y[i] in \$t1
	lbu \$t2, 0(\$t1)	# \$t2 = y[i]
	add \$t3, \$s0, \$a0	# addr of x[i] in \$t3
	sb \$t2, 0(\$t3)	# x[i] = y[i]
	beq \$t2, \$zero, L2	# exit loop if y[i] == 0
	addi \$s0, \$s0, 1	# i = i + 1
	j L1	# next iteration of loop
L2:	lw \$s0, 0(\$sp)	# restore saved \$s0
	addi \$sp, \$sp, 4	# pop 1 item from stack
	jr \$ra	# and return

MIPS Addressing for 32-Bit Immediates and Addresses

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

`lui $s0, 61`

0000 0000 0111 1101	0000 0000 0000 0000
---------------------	---------------------

`ori $s0, $s0, 2304`

0000 0000 0111 1101	0000 1001 0000 0000
---------------------	---------------------

Branch Addressing

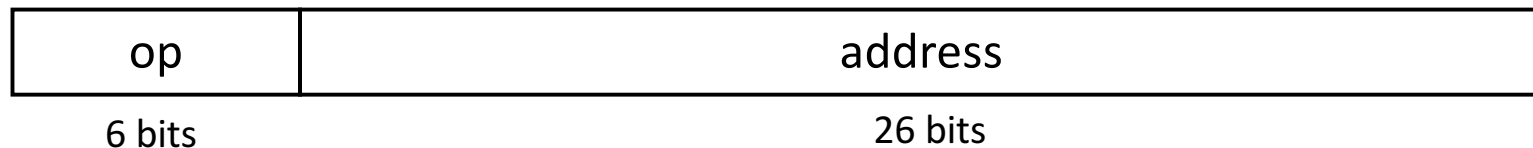
- Branch instructions specify
 - Opcode, two registers, target address



- Most branch targets are near branch
 - Forward or backward
- PC-relative addressing
 - Target address = PC + offset \times 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



- (Pseudo)Direct jump addressing
 - Target address = $PC_{31..28} : (\text{address} \times 4)$

Target Addressing Example

- Loop code from earlier example

– Assume Loop at location 80000

```

Loop: sll  $t1, $s3, 2      80000
      add  $t1, $t1, $s6    80004
      lw   $t0, 0($t1)     80008
      bne  $t0, $s5, Exit  80012
      addi $s3, $s3, 1     80016
      j    Loop            80020
Exit:  ...                80024
    
```

80000	0	0	19	9	4	0
80004	0	9	22	9	0	32
80008	35	9	8	0		
80012	5	8	21	2		
80016	8	19	19	1		
80020	2	20000				
80024						

Branching Far Away

- If branch target is too far to encode with 16-bit offset, assembler rewrites the code
- Example

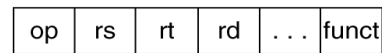
```
    beq $s0,$s1, L1
      ↓
    bne $s0,$s1, L2
    j  L1
L2:  ...
```


Addressing Mode Summary

1. Immediate addressing



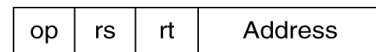
2. Register addressing



Registers

Register

3. Base addressing



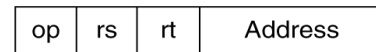
Register

+

Memory

Byte Halfword Word

4. PC-relative addressing



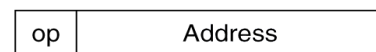
PC

+

Memory

Word

5. Pseudodirect addressing



PC

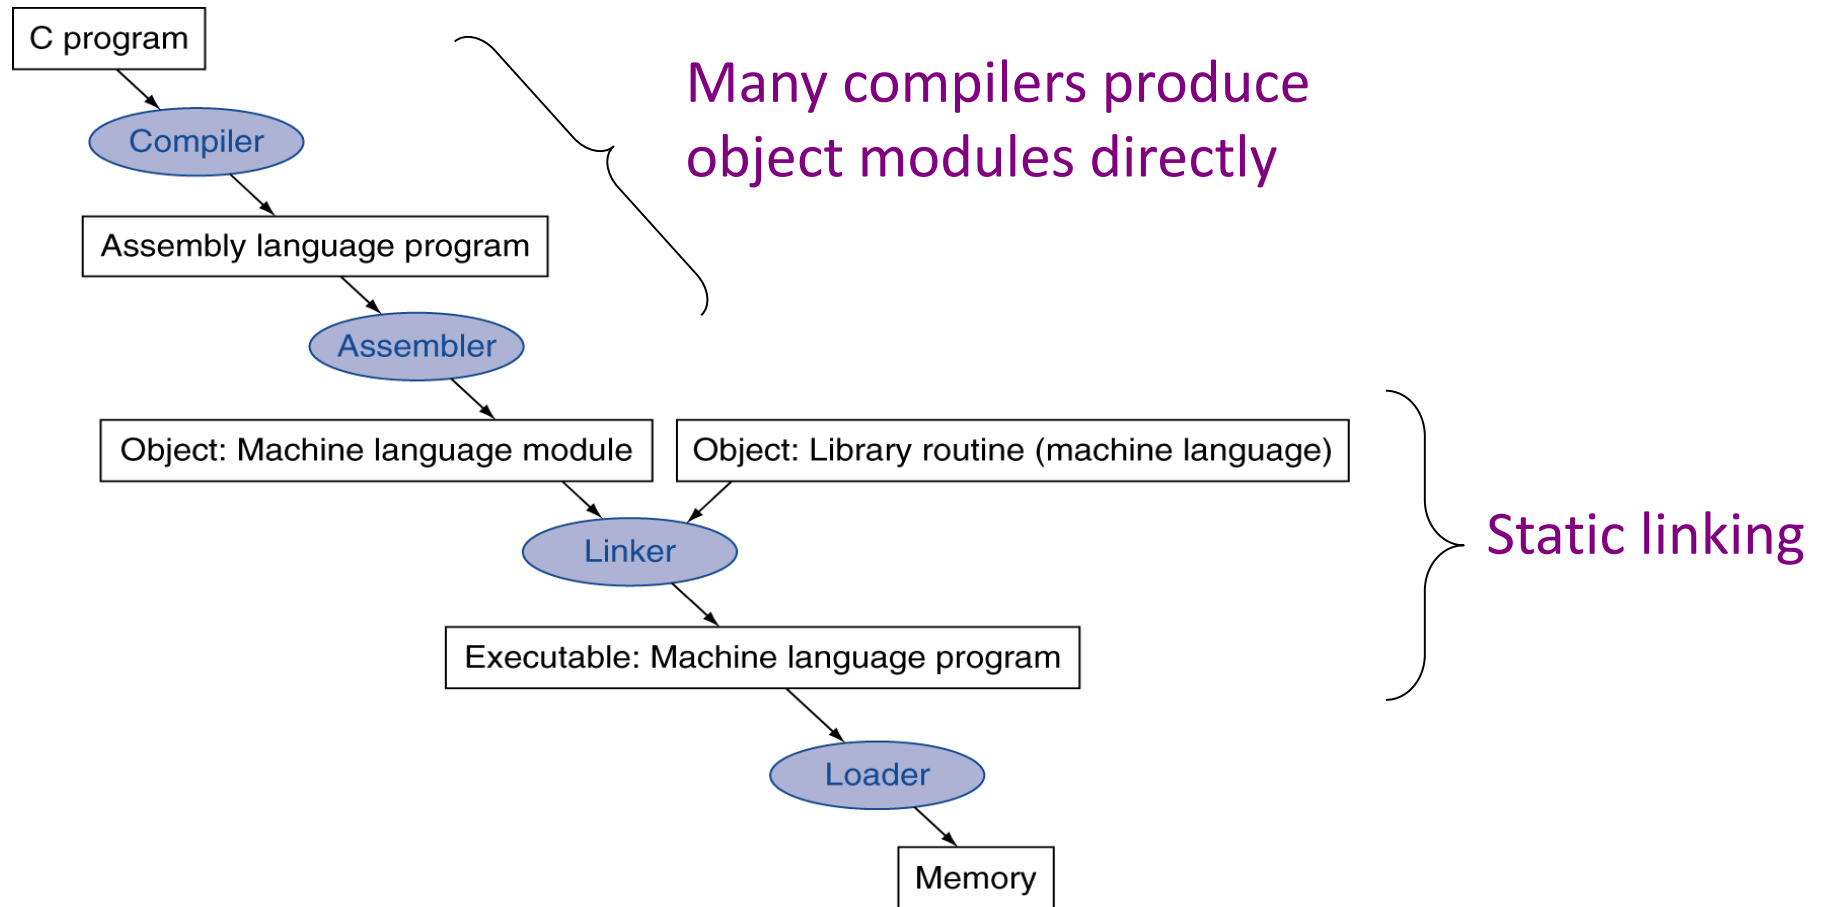
:

Memory

Word

Translating and Starting a Program

Translation and Startup



Assembler Pseudoinstructions

- Most assembler instructions represent machine instructions one-to-one
- Pseudoinstructions: figments of the assembler's imagination

`move $t0, $t1` → `add $t0, $zero, $t1`

`blt $t0, $t1, L` → `slt $at, $t0, $t1`
`bne $at, $zero, L`

– `$at` (register 1): assembler temporary

Producing an Object Module

- Assembler (or compiler) translates program into machine instructions
- Provides information for building a complete program from the pieces
 - Header: described contents of object module
 - Text segment: translated instructions
 - Static data segment: data allocated for the life of the program
 - Relocation info: for contents that depend on absolute location of loaded program
 - Symbol table: global definitions and external refs
 - Debug info: for associating with source code

Linking Object Modules

- Produces an executable image
 1. Merges segments
 2. Resolve labels (determine their addresses)
 3. Patch location-dependent and external refs
- Could leave location dependencies for fixing by a relocating loader
 - But with virtual memory, no need to do this
 - Program can be loaded into absolute location in virtual memory space

Loading a Program

- Load from image file on disk into memory
 1. Read header to determine segment sizes
 2. Create virtual address space
 3. Copy text and initialized data into memory
 - Or set page table entries so they can be faulted in
 4. Set up arguments on stack
 5. Initialize registers (including \$sp, \$fp, \$gp)
 6. Jump to startup routine
 - Copies arguments to \$a0, ... and calls main
 - When main returns, do exit syscall

Dynamic Linking

- Only link/load library procedure when it is called
 - Requires procedure code to be relocatable
 - Avoids image bloat caused by static linking of all (transitively) referenced libraries
 - Automatically picks up new library versions

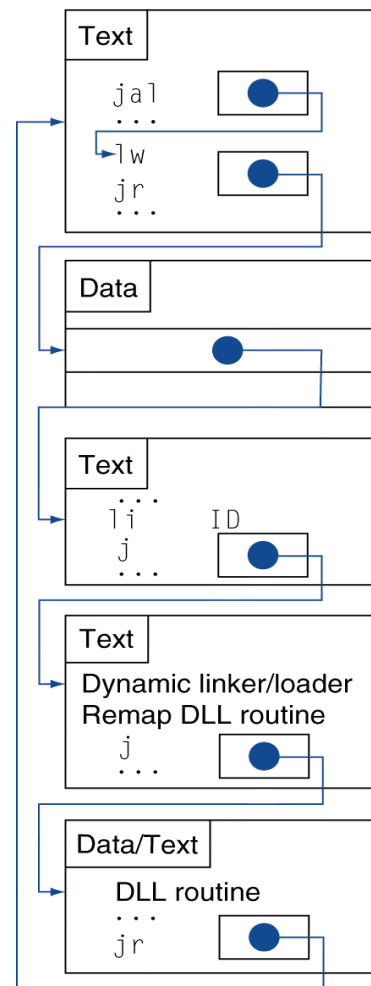
Lazy Linkage

Indirection table

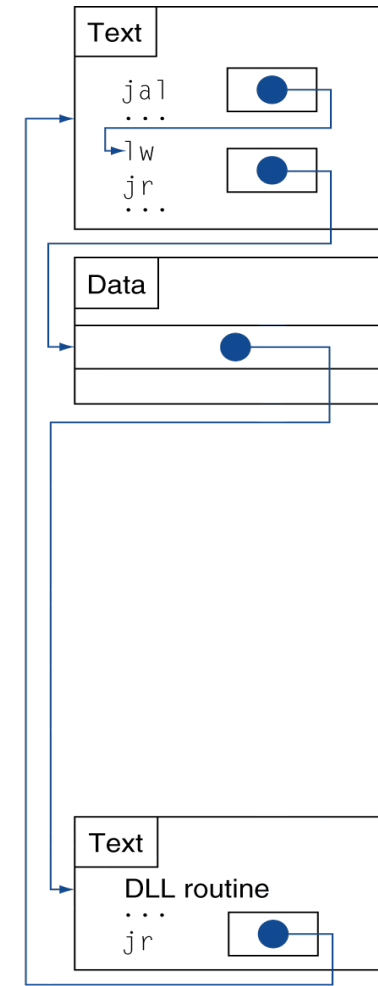
Stub: Loads routine ID,
Jump to linker/loader

Linker/loader code

Dynamically
mapped code

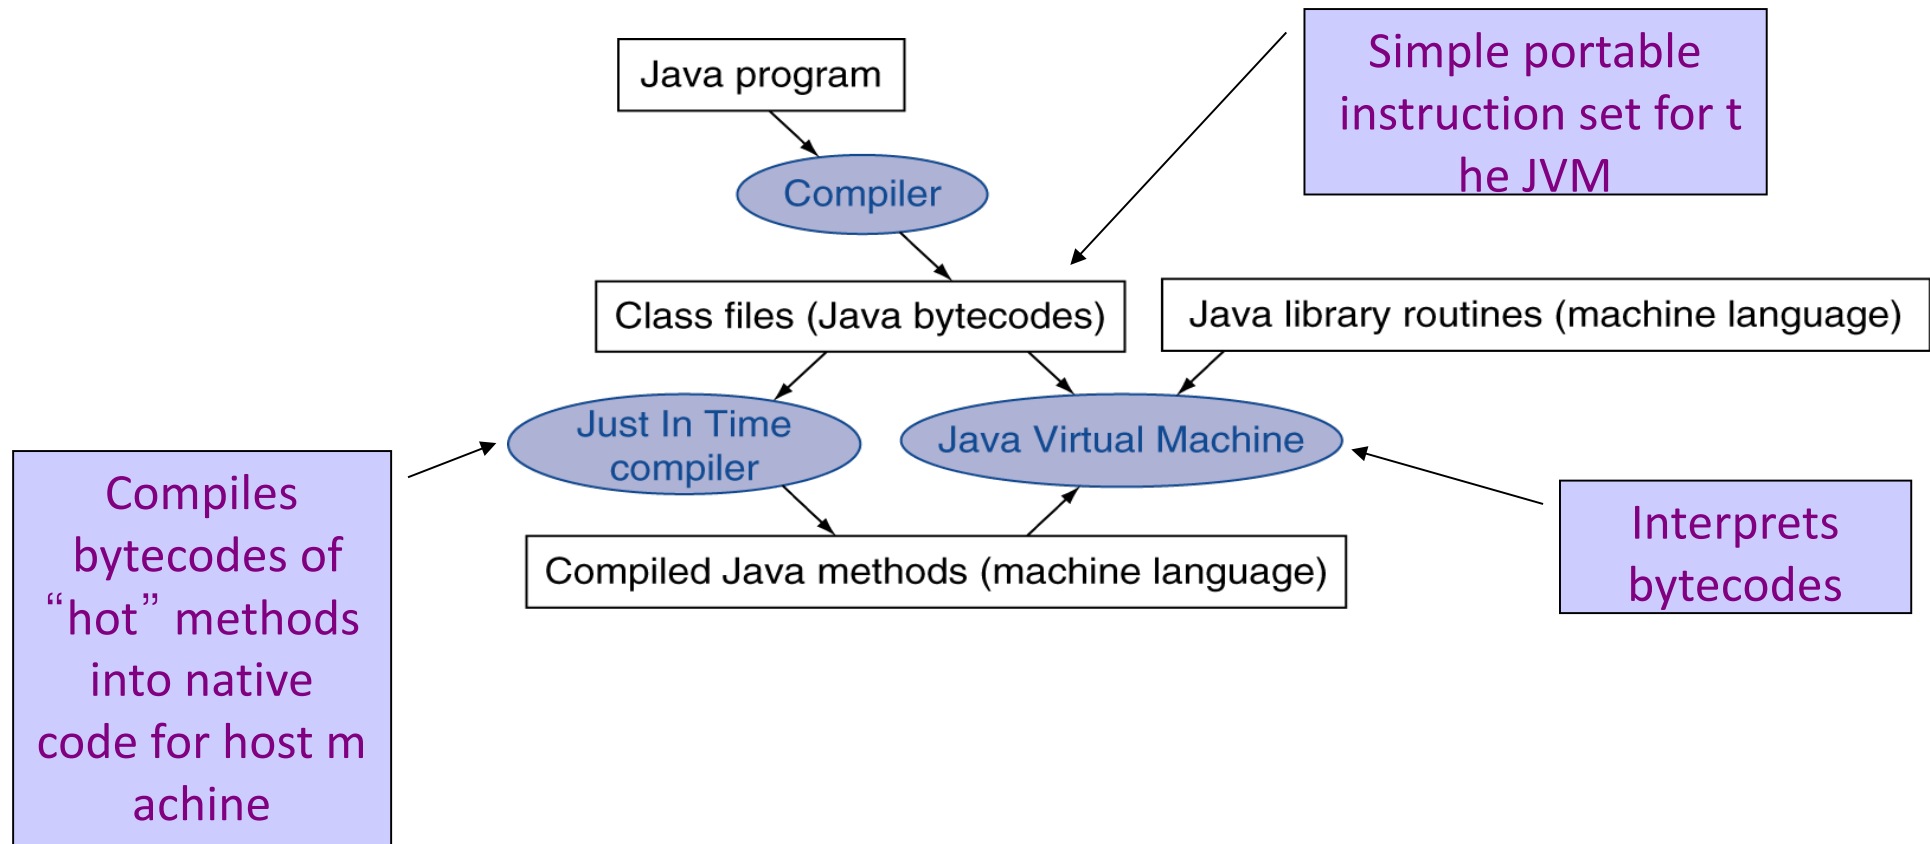


a. First call to DLL routine



b. Subsequent calls to DLL routine

Starting Java Applications



Array vs. Pointers

Arrays vs. Pointers

- **Array indexing involves**
 - Multiplying index by element size
 - Adding to array base address
- **Pointers correspond directly to memory addresses**
 - Can avoid indexing complexity

Example: Clearing and Array

<pre>clear1(int array[], int size) { int i; for (i = 0; i < size; i += 1) array[i] = 0; }</pre>	<pre>clear2(int *array, int size) { int *p; for (p = &array[0]; p < &array[size]; p = p + 1) *p = 0; }</pre>
<pre> move \$t0,\$zero # i = 0 loop1: sll \$t1,\$t0,2 # \$t1 = i * 4 add \$t2,\$a0,\$t1 # \$t2 = # &array[i] sw \$zero, 0(\$t2) # array[i] = 0 addi \$t0,\$t0,1 # i = i + 1 slt \$t3,\$t0,\$a1 # \$t3 = # (i < size) bne \$t3,\$zero,loop1 # if (...) # goto loop1</pre>	<pre> move \$t0,\$a0 # p = & array[0] sll \$t1,\$a1,2 # \$t1 = size * 4 add \$t2,\$a0,\$t1 # \$t2 = # &array[size] loop2: sw \$zero,0(\$t0) # Memory[p] = 0 addi \$t0,\$t0,4 # p = p + 4 slt \$t3,\$t0,\$t2 # \$t3 = # (p<&array[size]) bne \$t3,\$zero,loop2 # if (...) # goto loop2</pre>

Comparison of Array vs. Ptr

- Multiply “strength reduced” to shift
- Array version requires shift to be inside loop
 - Part of index calculation for incremented i
 - c.f. incrementing pointer
- Compiler can achieve same effect as manual use of pointers
 - Induction variable elimination
 - Better to make program clearer and safer