

CS 61C: Great Ideas in Computer Architecture *Assemblers, Linkers, Compilers*

Instructors:
Krste Asanovic, Randy H. Katz
<http://inst.eecs.Berkeley.edu/~cs61c/fa12>

9/21/12 Fall 2012 -- Lecture #12 1

New-School Machine Structures (It's a bit more complicated!)

- Parallel Requests**
Assigned to computer
e.g., Search "Katz"
- Parallel Threads**
Assigned to core
e.g., Lookup, Ads
- Parallel Instructions**
>1 instruction @ one time
e.g., 5 pipelined instructions
- Parallel Data**
>1 data item @ one time
e.g., Add of 4 pairs of words
- Hardware descriptions**
All gates @ one time
- Programming Languages**

Software | **Hardware**

Harness Parallelism & Achieve High Performance

Today's Lecture

9/21/12 Fall 2012 -- Lecture #12 2

Big Idea #1: Levels of Representation/ Interpretation

High Level Language Program (e.g., C)

```
temp = v[k];
v[k] = v[k+1];
v[k+1] = temp;
```

Compiler

Assembly Language Program (e.g., MIPS)

```
lw $t0, 0($2)
lw $t1, 4($2)
sw $t1, 0($2)
sw $t0, 4($2)
```

Assembler

Machine Language Program (MIPS)

```
0000 1001 1100 0110 1010 1010 1111 0101 1000
1010 1111 0101 1000 0000 1001 1100 0110
1100 0110 1010 1111 0101 1000 0000 1001
0101 1000 0000 1001 1100 0110 1010 1111
```

Anything can be represented as a number, i.e., data or instructions

Machine Interpretation

Hardware Architecture Description (e.g., block diagrams)

Architecture Implementation

Logic Circuit Description (Circuit Schematic Diagrams)

9/21/12 Fall 2012 -- Lecture #12 3

Agenda

- Review: Performance
- Assemblers
- Administrivia
- Linkers
- Compilers vs. Interpreters
- And in Conclusion, ...

9/21/12 Fall 2012 -- Lecture #12 4

Review: Defining Relative CPU Performance

- Performance_x = 1/Program Execution Time_x
- Performance_x > Performance_y => 1/Execution Time_x > 1/Execution Time_y => Execution Time_y > Execution Time_x
- *Computer X is N times faster than Computer Y*
Performance_x / Performance_y = N or Execution Time_y / Execution Time_x = N

9/21/12 Fall 2012 -- Lecture #12 5

Review: Performance Equation

- Time = $\frac{\text{Seconds}}{\text{Program}}$
= $\frac{\text{Instructions}}{\text{Program}} \times \frac{\text{Clock cycles}}{\text{Instruction}} \times \frac{\text{Seconds}}{\text{Clock Cycle}}$

9/21/12 Fall 2012 -- Lecture #102 6

Performance Comparison

MacAir 3.1: 1.6 Ghz
MacAir 5.1: 2.0 Ghz

	Late 2010 11"	Mid-2011 11"
Processor Speed:	1.4 GHz*	1.6 GHz**
Processor Type:	Core 2 Duo (SU9400)*	Core i5 (I5-2467M)**
L2/L3 Cache:	3 MB	256k x2, 3 MB
System Bus/DMI:	800 Mhz	5 GT/s
Standard RAM:	2 GB	2 GB, 4 GB
Maximum RAM:	4 GB*	4 GB*
Onboard RAM Type:	1066 Mhz DDR3	1333 Mhz DDR3
Internal Storage:	64 GB, 128 GB	64 GB, 128 GB
Storage Type:	SSD (Mini SATA)	SSD (Mini SATA)
Video Processor:	GeForce 320M	HD Graphics 3000
Shared VRAM:	256 MB	256 MB, 384 MB
Thunderbolt:	No	Yes
SD Card Slot:	No	No
Bluetooth:	2.1+EDR	4.0
Display Size:	11.6" Widescreen	11.6" Widescreen
Display Resolution:	1366x768	1366x768
Backlit Keyboard:	No	Yes
Battery Life:	5 Hours	5 Hours
Dimensions:	11. - 68 x 11.8 x 7.56	11. - 68 x 11.8 x 7.56
Weight:	2.3 Pounds	2.38 Pounds
Order Numbers:	MC505LL/A, MC506LL/A	MC968LL/A, MC969LL/A
EMC Number:	2393	2471
Model Identifier:	MacBookAir3,1	MacBookAir4,1

9/21/12

Performance Comparison

MacAir 3.1: 1.6 Ghz
MacAir 5.1: 2.0 Ghz

	MacBook Air
Model Name:	MacBook Air
Model Identifier:	MacBookAir3,1
Processor Name:	Intel Core 2 Duo
Processor Speed:	1.6 GHz
Number of Processors:	1
Total Number of Cores:	2
L2 Cache:	3 MB
Memory:	4 GB
Bus Speed:	800 Mhz
Boot ROM Version:	MBA31.0061.B01
SMC Version (system):	1.6714
Serial Number (system):	C02DW245DDR0
Hardware UUID:	35273131-BD15-5DEB-A919-701640168C81

	MacBook Air
Model Name:	MacBook Air
Model Identifier:	MacBookAir5,1
Processor Name:	Intel Core i7
Processor Speed:	2 GHz
Number of Processors:	1
Total Number of Cores:	2
L2 Cache (per Core):	256 KB
L3 Cache:	4 MB
Memory:	8 GB
Boot ROM Version:	MBA51.00EF.B00
SMC Version (system):	2.4117
Serial Number (system):	C02H4248DRV9
Hardware UUID:	7CEA9B1F-172E-59ED-ACCA-F8928D0E0A98

9/21/12

Review: Performance Equation

- Time = $\frac{\text{Seconds}}{\text{Program}}$
- $$= \frac{\text{Instructions}}{\text{Program}} \times \frac{\text{Clock cycles}}{\text{Instruction}} \times \frac{\text{Seconds}}{\text{Clock Cycle}}$$

Time_{3,1} = 1/1.6Ghz = 625 ns
Time_{5,1} = 1/2.0Ghz = 500 ns

$$\frac{\text{Performance}_{5,1}}{\text{Performance}_{3,1}} = \frac{\text{Time}_{3,1}}{\text{Time}_{5,1}} = \frac{625}{500} = 1.25$$

9/21/12 Fall 2012 - Lecture #12 9

gcc Optimization Experiment

	BubbleSort.c	Dhrystone.c
MacAir 3.1		
No Opt		
-O2		
MacAir 5.1		
No Opt		
-O2		

9/21/12 Fall 2012 - Lecture #12 10

gcc Optimization Experiment

	BubbleSort.c	Dhrystone.c
MacAir 3.1	3.2 s	7.4 s
No Opt		3125000 dhrys/s
-O2	1.5 s	2.7 s
		8333333 dhrys/s
MacAir 5.1	1.9 s (1.7x)	3.0 s (2.4x)
No Opt		8333333 dhrys/s (2.7x)
-O2	0.8 s (1.9x)	0.7 s (3.8x)
		25000000 dhrys/s (3x)

9/21/12 Fall 2012 - Lecture #12 11

Agenda

- Review: Performance
- Assemblers
- Administrivia
- Linkers
- Compilers vs. Interpreters
- And in Conclusion, ...

9/21/12 Fall 2012 - Lecture #12 12

Converting C to MIPS Machine code

\$t0 (reg 8), &A in \$t1 (reg 9), h=\$s2 (reg 18)
 A[300] = h + A[300];

Format? Student Roulette?

lw \$t0,1200(\$t1)

--	--	--	--	--	--	--	--	--

addu \$t0,\$s2,\$t0

--	--	--	--	--	--	--	--	--

sw \$t0,1200(\$t1)

--	--	--	--	--	--	--	--	--

Instruction	Format	op	rs	rt	rd	shamt	funct	address
addu	R	0	reg	reg	reg	0	33 _{ten}	n.a.
lw (load word)	I	35 _{ten}	reg	reg	n.a.	n.a.	n.a.	address
sw (store word)	I	43 _{ten}	reg	reg	n.a.	n.a.	n.a.	address
R-type		op	rs	rt	rd	shamt	funct	
I-type		op	rs	rt	address or constant			
J-type		op	address					

9/21/12 Fall 2012 – Lecture #12 13

Converting C to MIPS Machine code

\$t0 (reg 8), &A in \$t1 (reg 9), h=\$s2 (reg 18)
 A[300] = h + A[300];

Format? Student Roulette?

lw \$t0,1200(\$t1)

35	9	8						1200
----	---	---	--	--	--	--	--	------

addu \$t0,\$s2,\$t0

0	18	8	8					0 33
---	----	---	---	--	--	--	--	---------

sw \$t0,1200(\$t1)

43	9	8						1200
----	---	---	--	--	--	--	--	------

Instruction	Format	op	rs	rt	rd	shamt	funct	address
addu	R	0	reg	reg	reg	0	33 _{ten}	n.a.
lw (load word)	I	35 _{ten}	reg	reg	n.a.	n.a.	n.a.	address
sw (store word)	I	43 _{ten}	reg	reg	n.a.	n.a.	n.a.	address
R-type		op	rs	rt	rd	shamt	funct	
I-type		op	rs	rt	address or constant			
J-type		op	address					

9/21/12 Fall 2012 – Lecture #12 14

Converting to MIPS Machine code

Address Loop: Format? Student Roulette?

800 sll \$t1,\$s3,2 R

--	--	--	--	--	--	--	--	--

804 addu \$t1,\$t1,\$s6 R

--	--	--	--	--	--	--	--	--

808 lw \$t0,0(\$t1) I

--	--	--	--	--	--	--	--	--

812 bne \$t0,\$s5, Exit I

--	--	--	--	--	--	--	--	--

816 addiu \$s3,\$s3,1 I

--	--	--	--	--	--	--	--	--

820 j Loop J

--	--	--	--	--	--	--	--	--

Exit:

R-type	op	rs	rt	rd	shamt	funct	
I-type	op	rs	rt	address or constant			
J-type	op	address					

9/21/12 Fall 2012 – Lecture #12 15

Converting to MIPS Machine code

Address Loop: Format? Student Roulette?

800 sll \$t1,\$s3,2 R

0	0	19	9	2	0
---	---	----	---	---	---

804 addu \$t1,\$t1,\$s6 R

0	9	22	9	0	33
---	---	----	---	---	----

808 lw \$t0,0(\$t1) I

35	9	8			0
----	---	---	--	--	---

812 bne \$t0,\$s5, Exit I

5	8	21			2
---	---	----	--	--	---

816 addiu \$s3,\$s3,1 I

8	19	19			1
---	----	----	--	--	---

820 j Loop J

2					200
---	--	--	--	--	-----

Exit:

R-type	op	rs	rt	rd	shamt	funct	
I-type	op	rs	rt	address or constant			
J-type	op	address					

9/21/12 Fall 2012 – Lecture #12 16

32-bit Constants in MIPS

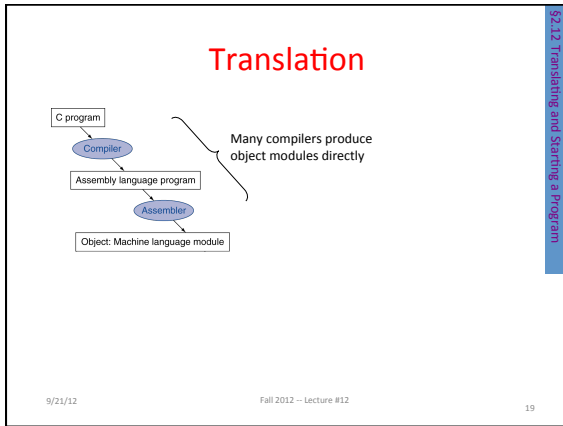
- Can create a 32-bit constant from two 32-bit MIPS instructions
- *Load Upper Immediate (lui or “Louie”)* puts 16 bits into upper 16 bits of destination register
- MIPS to load 32-bit constant into register \$s0?
 0000 0000 0011 1101 0000 1001 0000 0000_{two}
 lui \$s0, 61 # 61 = 0000 0000 0011 1101_{two}
 ori \$s0, \$s0, 2304 # 2304 = 0000 1001 0000 0000_{two}

9/21/12 Fall 2012 – Lecture #12 17

Assembler

- Input: Assembly Language Code (e.g., `foo.s` for MIPS)
- Output: Object Code, information tables (e.g., `foo.o` for MIPS)
- Reads and Uses **Directives**
- Replace **Pseudo-instructions**
- Produce Machine Language
- Creates **Object File**

9/21/12 Fall 2012 – Lecture #12 18



- ### Assembly and Pseudo-instructions
- Turning textual MIPS instructions into machine code called *assembly*, program called *assembler*
 - Calculates addresses, maps register names to numbers, produces binary machine language
 - Textual language called *assembly language*
 - Can also accept instructions convenient for programmer but not in hardware
 - Load immediate (li)* allows 32-bit constants, assembler turns into `lui + ori` (if needed)
 - Load double (ld)* uses two `lwc1` instructions to load a pair of 32-bit floating point registers
 - Called *Pseudo-Instructions*
- 9/21/12 Fall 2012 -- Lecture #12 20

- ### Assembler Directives (P&H B-5 to B-7)
- Give directions to assembler, but do not produce machine instructions
 - `.text`: Subsequent items put in user text segment
 - `.data`: Subsequent items put in user data segment
 - `.globl sym`: declares `sym` global and can be referenced from other files
 - `.ascii str`: Store the string `str` in memory and null-terminate it
 - `.word w1...wn`: Store the `n` 32-bit quantities in successive memory words
- 9/21/12 Fall 2012 -- Lecture #12 21

- ### Assembler Pseudo-instructions
- Most assembler instructions represent machine instructions one-to-one
 - Pseudo-instructions: 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
- 9/21/12 Fall 2012 -- Lecture #12 22

- ### More Pseudo-instructions
- Asm. treats convenient variations of machine language instructions as if real instructions
- | Pseudo: | Real: |
|--------------------------------|-------|
| <code>addu \$t0,\$t6,1</code> | _____ |
| <code>subu \$sp,\$sp,32</code> | _____ |
| <code>sd \$a0,32(\$sp)</code> | _____ |
| <code>la \$a0,str</code> | _____ |
- 9/21/12 Fall 2012 -- Lecture #12 [Student Roulette?](#) 23

- ### More Pseudoinstructions
- Asm. treats convenient variations of machine language instructions as if real instructions
- | Pseudo: | Real: |
|--------------------------------|--------------------------------|
| <code>addu \$t0,\$t6,1</code> | <code>addiu \$t0,\$t6,1</code> |
| <code>subu \$sp,\$sp,32</code> | _____ |
| <code>sd \$a0,32(\$sp)</code> | _____ |
| <code>la \$a0,str</code> | _____ |
- 9/21/12 Fall 2012 -- Lecture #12 [Student Roulette?](#) 24

More Pseudoinstructions

- Asm. treats convenient variations of machine language instructions as if real instructions

Pseudo:	Real:
<code>addu \$t0,\$t6,1</code>	<code>addiu \$t0,\$t6,1</code>
<code>subu \$sp,\$sp,32</code>	<code>addiu \$sp,\$sp,-32</code>
<code>sd \$a0,32(\$sp)</code>	_____

<code>la \$a0,str</code>	_____

9/21/12 Fall 2012 -- Lecture #12 Student Roulette? 25

More Pseudoinstructions

- Asm. treats convenient variations of machine language instructions as if real instructions

Pseudo:	Real:
<code>addu \$t0,\$t6,1</code>	<code>addiu \$t0,\$t6,1</code>
<code>subu \$sp,\$sp,32</code>	<code>addiu \$sp,\$sp,-32</code>
<code>sd \$a0,32(\$sp)</code>	<code>sw \$a0,32(\$sp)</code>
	<code>sw \$a1,36(\$sp)</code>
<code>la \$a0,str</code>	<code>lui \$at,left(str)</code>
	<code>ori \$a0,\$at,right(str)</code>

9/21/12 Fall 2012 -- Lecture #12 Student Roulette? 26

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

9/21/12 Fall 2012 -- Lecture #12 27

Agenda

- Review
- Assemblers
- Administrivia
- Linkers
- Compilers vs. Interpreters
- And in Conclusion, ...

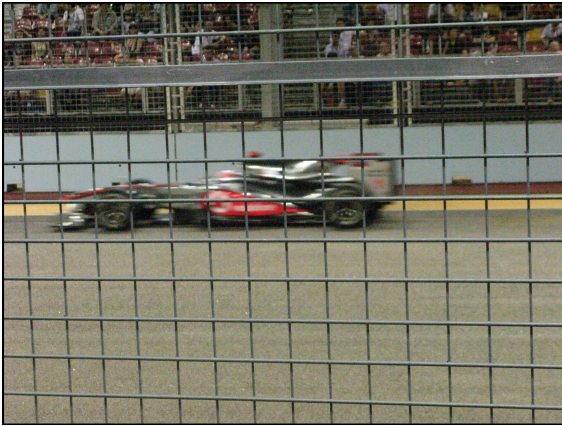
9/21/12 Fall 2012 -- Lecture #12 28

Administrivia

- Midterm! October 9th, Evening, 8 PM – 10 PM
 - 1 hour exam with 2 hours to complete it
 - Closed notes, book; one 8.5x11" crib sheet, MIPS green card provided
 - Comprehensive from course start
 - Lectures, Labs, Projects
 - There will be a TA-led review session
 - Special accommodations, contact instructors

9/21/12 Fall 2012 -- Lecture #7 29

The screenshot shows a Piazza Q&A forum for a course. The interface includes a search bar, a list of questions with timestamps and view counts, and a detailed view of a question titled "Finally got project working correctly, recommit?". The question asks if students should recommit the project even if it's past the 48-hour late submission deadline. The instructor's answer states that it's probably not a good idea to turn in submissions after the deadline. The forum also shows a "followup discussions" section and a "Special Mentions" section highlighting a user named Alan Christopher.



Agenda

- Review
- Assemblers
- **Linkers**
- Administrivia
- Compilers vs. Interpreters
- And in conclusion, ...

9/21/12 Fall 2012 -- Lecture #12 32

Separate Compilation and Assembly

- No need to compile all code at once
- How to put pieces together?

FIGURE B.1.1 The process that produces an executable file. An assembler translates a file of assembly language into an object file, which is linked with other files and libraries into an executable file. Copyright © 2009 Elsevier, Inc. All rights reserved.

9/21/12 Fall 2012 -- Lecture #12 33

Translation and Startup

§2.12 Translating and Starting a Program

9/21/12 Fall 2012 -- Lecture #12 34

Linker Stitches Files Together

FIGURE B.3.1 The linker searches a collection of object files and program libraries to find nonlocal routines used in a program, combines them into a single executable file, and resolves references between routines in different files. Copyright © 2009 Elsevier, Inc. All rights reserved.

9/21/12 Fall 2012 -- Lecture #12 35

Linking Object Modules

- Produces an executable image
 1. Merges segments
 2. Resolve labels (determine their addresses)
 3. Patch location-dependent and external refs
- Often slower than compiling
 - All the machine code files must be read into memory and linked together

9/21/12 Fall 2012 -- Lecture #12 36

Loading a Program

- Load from image file on disk into memory
 1. Read header to determine segment sizes
 2. Create virtual address space (covered later in semester)
 3. Copy text and initialized data into memory
 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" systems call

9/21/12

Fall 2012 -- Lecture #12

37

Agenda

- Review
- Assemblers
- Administrivia
- Linkers
- Compilers vs. Interpreters
- And in Conclusion, ...

9/21/12

Fall 2012 -- Lecture #12

38

What's a Compiler?

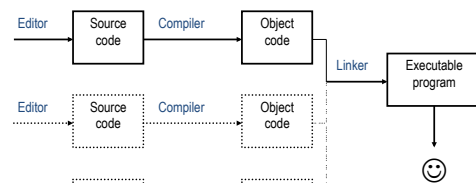
- *Compiler*: a program that accepts as input a program text in a certain language and produces as output a program text in another language, *while preserving the meaning of that text*.
- Text must comply with the syntax rules of whichever programming language it is written in.
- Compiler's complexity depends on the syntax of the language and how much abstraction that programming language provides.
 - A C compiler is much simpler than C++ Compiler
- Compiler executes *before* compiled program runs

9/21/12

Fall 2012 -- Lecture #12

39

Compiled Languages: Edit-Compile-Link-Run



9/21/12

Fall 2012 -- Lecture #12

2-40

Compiler Optimization

- gcc compiler options
- O1: the compiler tries to reduce code size and execution time, without performing any optimizations that take a great deal of compilation time
- O2: Optimize even more. GCC performs nearly all supported optimizations that do not involve a space-speed tradeoff. As compared to -O, this option increases both compilation time and the performance of the generated code
- O3: Optimize yet more. All -O2 optimizations and also turns on the `-finline-functions`, ...

9/21/12

Fall 2012 -- Lecture #12

41

What is Typical Benefit of Compiler Optimization?

- What is a typical program?
- For now, try a toy program:
BubbleSort.c

```

#define ARRAY_SIZE 20000
int main() {
    int iarray[ARRAY_SIZE], x, y, holder;
    for(x = 0; x < ARRAY_SIZE; x++)
        for(y = 0; y < ARRAY_SIZE-1; y++)
            if(iarray[y] > iarray[y+1]) {
                holder = iarray[y+1];
                iarray[y+1] = iarray[y];
                iarray[y] = holder;
            }
}
  
```

9/21/12

Fall 2012 -- Lecture #12

42

Unoptimized MIPS Code

```

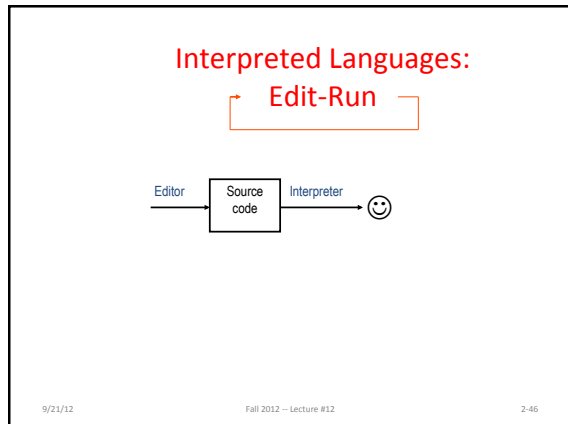
SL3:      addu $2,$3,$2      lw $3,80020($sp)      SL11:
lw $2,80016($sp)  lw $4,80020($sp)  addu $2,$3,1      SL9:
slt $3,$2,20000  addu $3,$4,1      move $3,$2      lw $2,80020($sp)
bne $3,$0,$16    move $4,$3      sll $2,$3,2      addu $3,$2,1
j $L4           sll $3,$4,2      addu $3,$sp,16    sw $3,80020($sp)
SL6:      addu $4,$sp,16  addu $2,$3,$2      j $L7
.set noreorder   addu $3,$4,$3  lw $3,80020($sp)  SL8:
nop            lw $2,0($2)      move $4,$3      SL5:
.set reorder     lw $3,0($3)  sll $3,$4,2      lw $2,80016($sp)
sw $0,80020($sp)  slt $2,$3,$2      addu $4,$sp,16  addu $3,$2,1
SL7:      beq $2,$0,$L9    addu $3,$4,$3      lw $3,80016($sp)  sw $3,80016($sp)
lw $2,80020($sp)  lw $3,80020($sp)  lw $4,0($3)      j $L3
slt $3,$2,19999  addu $2,$3,1      sw $4,0($2)      SL4:
bne $3,$0,$L10  move $3,$2      lw $2,80020($sp)  SL2:
j $L5          sll $2,$3,2      move $3,$2
SL10:     addu $3,$sp,16    sll $2,$3,2      li $12,65536
lw $2,80020($sp)  addu $2,$3,$2      addu $3,$sp,16  ori
move $3,$2      lw $3,0($2)      addu $2,$3,$2  $12,$12,0x38b0
sll $2,$3,2      sw $3,80024($sp)  lw $3,80024($sp)  addu $13,$12,$sp
addu $3,$sp,16   sw $3,0($2)      addu $sp,$sp,$12  j $31
    
```

-O2 optimized MIPS Code

```

li $13,65536      slt $2,$4,$3
ori $13,$13,0x3890  beq $2,$0,$L9
addu $13,$13,$sp  sw $3,0($5)
sw $28,0($13)     sw $4,0($6)
move $4,$0        SL9:
addu $8,$sp,16   move $3,$7
SL6:             slt $2,$3,19999
move $3,$0       bne $2,$0,$L10
addu $9,$4,1     move $4,$9
.p2align 3       slt $2,$4,20000
SL10:            bne $2,$0,$L6
sll $2,$3,2      li $12,65536
addu $6,$8,$2    ori $12,$12,0x38a0
addu $7,$3,1     addu $13,$12,$sp
sll $2,$7,2      addu $sp,$sp,$12
addu $5,$8,$2    j $31
lw $3,0($6)      .
lw $4,0($5)      .
    
```

- ### What's an Interpreter?
- Reads and executes source statements executed one at a time
 - No linking
 - No machine code generation, so more portable
 - Starts executing quicker, but runs much more slowly than compiled code
 - Performing the actions straight from the text allows better error checking and reporting to be done
 - Interpreter stays around during execution
 - Unlike compiler, some work is done *after* program starts
 - Writing an interpreter is much less work than writing a compiler



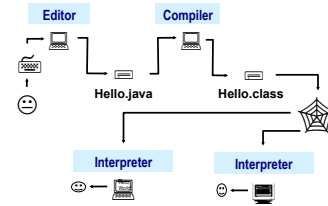
- ### Compiler vs. Interpreter Advantages
- | | |
|--|---|
| <p>Compilation:</p> <ul style="list-style-type: none"> • Faster Execution • Single file to execute • Compiler can do better diagnosis of syntax and semantic errors, since it has more info than an interpreter (Interpreter only sees one line at a time) • Can find syntax errors <i>before</i> run program • Compiler can optimize code | <p>Interpreter:</p> <ul style="list-style-type: none"> • Easier to debug program • Faster development time |
|--|---|

- ### Compiler vs. Interpreter Disadvantages
- | | |
|---|---|
| <p>Compilation:</p> <ul style="list-style-type: none"> • Harder to debug program • Takes longer to change source code, recompile, and relink | <p>Interpreter:</p> <ul style="list-style-type: none"> • Slower execution times • No optimization • Need all of source code available • Source code larger than executable for large systems • Interpreter must remain installed while the program is interpreted |
|---|---|

Java's Hybrid Approach: Compiler + Interpreter

- A Java compiler converts Java source code into instructions for the *Java Virtual Machine (JVM)*
- These instructions, called *bytecodes*, are same for any computer / OS
- A CPU-specific Java interpreter interprets bytecodes on a particular computer

Java's Compiler + Interpreter



Why Bytecodes?

- Platform-independent
- Load from the Internet faster than source code
- Interpreter is faster and smaller than it would be for Java source
- Source code is not revealed to end users
- Interpreter performs additional security checks, screens out malicious code

JVM uses Stack vs. Registers

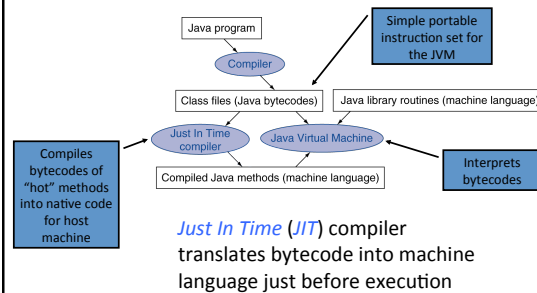
```

a = b + c;
=>
iload b ; push b onto Top Of Stack (TOS)
iload c ; push c onto Top Of Stack (TOS)
iadd ; Next to top Of Stack (NOS) = ; Top Of Stack (TOS) + NOS
istore a ; store TOS into a and pop stack
  
```

Java Bytecodes (Stack) vs. MIPS (Reg.)

Category	Operation	Java bytecode	Size (bits)	MIPS instr.	Meaning
Arithmetic	add	iadd	8	add	NOS←TOS+NOS; pop
	subtract	isub	8	sub	NOS←TOS-NOS; pop
Data transfer	increment	inc iB/iBb	8	addi	Frame[<i>iB</i>]= Frame[<i>iB</i>] + iBb
	load local integer/address	iload iB/aload iB	16	lw	TOS←Frame[<i>iB</i>]
	load local integer/address	iload ₂ /aload ₂	8	lw	TOS←Frame[0,1,2,3]
	store local integer/address	istore iB/astore iB	16	sw	Frame[<i>iB</i>]=TOS; pop
	load integer/address from array	iaload/aaload	8	lw	NOS←NOS[TOS]; pop
	store integer/address into array	istore/astore	8	sw	*NOS[NOS]=TOS; pop2
	load half from array	saload	8	lh	NOS←NOS[TOS]; pop
	store half into array	sastore	8	sh	*NOS[NOS]=TOS; pop2
	load byte from array	baload	8	lb	NOS←NOS[TOS]; pop
	store byte into array	bastore	8	sb	*NOS[NOS]=TOS; pop2
Logical	load immediate	bipush iB, sipush i16	16, 24	addi	push: TOS←iB or i16
	load immediate	iconst: [-1,0,1,2,3,4,5]	8	addi	push: TOS←[-1,0,1,2,3,4,5]
Conditional branch	and	iand	8	and	NOS←TOS&NOS; pop
	or	ior	8	or	NOS←TOS NOS; pop
	shift left	lshl	8	sll	NOS←NOS << TOS; pop
	shift right	lshr	8	srl	NOS←NOS >> TOS; pop
Unconditional jump	branch on equal	if_icmpeq i16	24	beq	if TOS == NOS, go to i16; pop2
	branch on not equal	if_icmpne i16	24	bne	if TOS != NOS, go to i16; pop2
Unconditional jump	compare	if_comp(i16,gt,ge) i16	24	slt	if TOS {<,≤,≥,}> NOS, go to i16; pop2
	jump	goto i16	24	j	go to i16
	return	ret, ireturn	8	jr	
Unconditional jump	jump to subroutine	jar i16	24	jail	go to i16; push; TOS←PC+3

Starting Java Applications



And, in Conclusion, ...

- Assemblers can enhance machine instruction set to help assembly-language programmer
- Translate from text that easy for programmers to understand into code that machine executes efficiently: Compilers, Assemblers
- Linkers allow separate translation of modules
- Interpreters for debugging, but slow execution
- Hybrid (Java): Compiler + Interpreter to try to get best of both
- Compiler Optimization to relieve programmer

9/21/12

Fall 2012 -- Lecture #12

55