

CS 61C: Great Ideas in Computer Architecture *Compilers and Floating Point*

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

9/24/12 Fall 2012 - Lecture #13 1

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

Software

- 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**

Hardware

Warehouse Scale Computer

Smart Phone

Harness Parallelism & Achieve High Performance

Computer

Core ... Core

Memory (Cache)

Input/Output

Instruction Unit(s)

Functional Unit(s)

Cache Memory

Logic Gates

9/24/12 Fall 2012 - Lecture #13 2

Big Idea #1: Levels of Representation/ Interpretation Today's Lecture

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

Compiler

Assembly Language Program (e.g., MIPS)

Assembler

Machine Language Program (MIPS)

Machine Interpretation

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

Architecture Implementation

Logic Circuit Description (Circuit Schematic Diagrams)

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

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

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

```
0000 1001 1100 0110 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
```

9/24/12 Fall 2012 - Lecture #13 3

Agenda

- Review
- Compilers
- Administrivia
- Floating Point Revisited
- And in Conclusion, ...

9/24/12 Fall 2012 - Lecture #13 4

Translation and Startup

Many compilers produce object modules directly

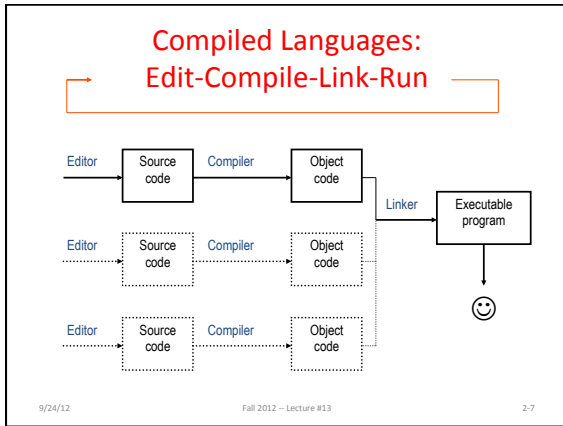
Static linking

9/24/12 Fall 2012 - Lecture #13 5

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

9/24/12 Fall 2012 - Lecture #13 6



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, ...

9/24/12 Fall 2012 - Lecture #13 8

gcc Optimization Experiment

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

9/24/12 Fall 2012 - Lecture #13 9

Performance Equation

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

Compiler affects this!

9/24/12 Fall 2012 - Lecture #13 10

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;
            }
}
  
```

9/24/12 Fall 2012 - Lecture #13 11

Unoptimized MIPS Code

```

SL3:
lw $2,80016($sp)      addu $2,$3,$2      lw $3,80020($sp)      $L11:
sll $3,$2,20000      addu $3,$4,1        addu $2,$3,1        $L9:
bne $3,$0,$L6        move $3,$2          lw $2,80020($sp)    lw $2,80020($sp)
j $L4                sll $3,$4,2        addu $3,$2,1        addu $3,$2,1
SL6:
.set noreorder      addu $4,$sp,16     addu $2,$3,$2      sw $3,80020($sp)
nop                 addu $3,$4,$3      lw $3,80020($sp)   j $L7
.set reorder        lw $3,$3          sll $3,$4,2        $L5:
sw $0,80020($sp)    st $2,$3,$2       addu $4,$sp,16     lw $2,80016($sp)
SL7:
lw $2,80020($sp)    lw $2,$0($2)      beq $2,$0,$L9     addu $3,$2,1
sll $3,$2,19999    addu $3,$4,$3      addu $3,$4,$3      sw $3,80016($sp)
bne $3,$0,$L10     move $3,$2        sw $4,$0($2)       j $L3
j $L5                sll $2,$3,2       move $3,$2        $L4:
SL10:
lw $2,80020($sp)    addu $3,$sp,16     sll $2,$3,2        li $12,65536
move $3,$2          addu $2,$3,$2     addu $3,$sp,16     ori
sll $2,$3,2         lw $3,$0($2)     addu $2,$3,$2      addu $12,$0x3880
addu $3,$sp,16     sw $3,80024($sp) lw $3,80024($sp)  addu $13,$12,$sp
                                                            addu $sp,$sp,$12
                                                            j $31
  
```

9/24/12 Fall 2012 - Lecture #13 12

-O2 optimized MIPS Code

```

li $13,65536          slt $2,$4,$3
ori $13,$13,0x3890   beq $2,$0,$19
addu $13,$13,$5p     sw $3,0($5)
sw $28,0($13)        sw $4,0($6)
move $4,$0           $L9:
addu $8,$5p,16       move $3,$7
$SL6:
move $3,$0           slt $2,$3,19999
addu $9,$4,1         bne $2,$0,$L10
.p2align 3           move $4,$9
$SL10:
sll $2,$3,2          li $12,65536
addu $6,$8,$2        ori $12,$12,0x38a0
addu $7,$3,1         addu $13,$12,$5p
sll $2,$7,2          addu $5p,$5p,$12
addu $5,$8,$2        j $31
lw $3,0($6)          .
lw $4,0($5)
    
```

9/24/12

Fall 2012 -- Lecture #13

13

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

9/24/12

Fall 2012 -- Lecture #13

14

Interpreted Languages:



9/24/12

Fall 2012 -- Lecture #13

2-15

Compiler vs. Interpreter Advantages

Compilation:

- 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 *before* run program
- Compiler can optimize code

Interpreter:

- Easier to debug program
- Faster development time

9/24/12

Fall 2012 -- Lecture #13

16

Compiler vs. Interpreter Disadvantages

Compilation:

- Harder to debug program
- Takes longer to change source code, recompile, and relink

Interpreter:

- 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

9/24/12

Fall 2012 -- Lecture #13

17

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

9/24/12

Fall 2012 -- Lecture #13

18

Java's Compiler + Interpreter



9/24/12

Fall 2012 - Lecture #13

19

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

9/24/12

Fall 2012 - Lecture #13

20

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
  
```

9/24/12

Fall 2012 - Lecture #13

21

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

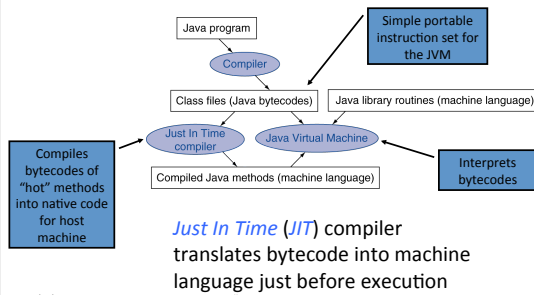
Category	Operation	Java bytecode	Size (bits)	MIPS Instr.	Meaning
Arithmetic	add	iadd	8	add	NOS=TOS+NOS; pop
	subtract	isub	8	sub	NOS=TOS-NOS; pop
	increment	inc I8a I8b	8	addl	Frame[8a]-Frame[8a] + I8b
Data transfer	load local integer/address	iload I8/aload I8	16	lw	TOS=Frame[I8]
	load local integer/address	iload I8/aload I8	8	lw	TOS=Frame[I0,I2,3]
	store local integer/address	istore I8/astore I8	16	sw	Frame[8]=TOS; pop
	load integer/address from array	iaload/aaload	8	lw	NOS=*NOS[TOS]; pop
	store integer/address into array	iastore/aastore	8	sw	*NOS[TOS]=TOS; pop2
	load half from array	lshld	8	lh	NOS=*NOS[TOS]; pop
	store half into array	sastore	8	sh	*NOS[TOS]=TOS; pop2
	load byte from array	baload	8	lb	NOS=*NOS[TOS]; pop
	store byte into array	bastore	8	sb	*NOS[TOS]=TOS; pop2
	load immediate	bipush I8, sipush I16	16, 24	addi	push; TOS=I8 or I16
Logical	load immediate	iconst [-1,0,1,2,3,4,5]	8	addi	push; TOS=[-1,0,1,2,3,4,5]
	and	iand	8	and	NOS=TOS&NOS; pop
	or	ior	8	or	NOS=TOS NOS; pop
	shift left	ishl	8	sl	NOS=NOS << TOS; pop
	shift right	ishr	8	srl	NOS=NOS >> TOS; pop
Conditional branch	branch on equal	if_icmpeq I16	24	beq	if TOS == NOS, go to I16; pop2
	branch on not equal	if_icmpneq I16	24	bne	if TOS != NOS, go to I16; pop2
	compare	if_icmpgt,le,gt,ge I16	24	slt	if TOS [< < > >] NOS, go to I16; pop2
Unconditional jump	jump	goto I16	24	j	go to I16
	return	ret, ireturn	8	jr	
	jump to subroutine	jsr I16	24	jal	go to I16; push; TOS=PC+3

9/24/12

Fall 2012 - Lecture #13

22

Starting Java Applications



9/24/12

Fall 2012 - Lecture #13

23

Agenda

- Review
- Compilers
- Administrivia
- Floating Point Revisited
- And in Conclusion, ...

9/24/12

Fall 2012 - Lecture #13

24

Administrivia

- Lab #5: MIPS Assembly
- HW #4 (of six), due Sunday
- Project 2a: MIPS Emulator, due Sunday
- Midterm, two weeks from Tuesday

9/24/12 Fall 2012 -- Lecture #13 25

THE CLOUD FACTORIES Power, Pollution and the Internet



Data centers are filled with servers, which are like bulked-up desktop computers, minus screens and keyboards, that contain chips to process data. (Ethian Prine for The New York Times)

By JAMES GLANZ
Published: September 22, 2012 | 273 Comments

SANTA CLARA, Calif. — Jeff Rothschild’s machines at Facebook had a problem he knew he had to solve immediately. They were about to melt.

[FACEBOOK](#)
[TWITTER](#)
[GOOGLE+](#)

9/24/12 26

CS61c in the News

“Most data centers, by design, consume vast amounts of energy in an incongruously wasteful manner, interviews and documents show. Online companies typically run their facilities at maximum capacity around the clock, whatever the demand. As a result, data centers can waste 90 percent or more of the electricity they pull off the grid, The Times found.”

“Worldwide, the digital warehouses use about 30 billion watts of electricity, roughly equivalent to the output of 30 nuclear power plants, according to estimates industry experts compiled for The Times.”

“The consulting firm McKinsey & Company analyzed energy use by data centers and found that, on average, they were using only 6 percent to 12 percent of the electricity powering their servers to perform computations. The rest was essentially used to keep servers idling and ready in case of a surge in activity that could slow or crash their operations.”

9/24/12 Fall 2012 -- Lecture #13 27

THE CLOUD FACTORIES Data Barns in a Farm Town, Gobbling Power and Flexing Muscle



Relatively low-cost hydroelectric power has lured technology giants to central Washington State.

By JAMES GLANZ
Published: September 23, 2012 | 72 Comments

QUINCY, Wash. — Set in the dry hills and irrigated farmland of Central Washington, Grant County is known for its robust harvest of apples, potatoes, cherries and beans. But for Microsoft, a prime lure was the region’s other valuable resource: cheap electrical power.

[FACEBOOK](#)
[TWITTER](#)
[GOOGLE+](#)
[E-MAIL](#)

9/24/12 28



Agenda

- Review
- Compilers
- Administrivia
- Floating Point Revisited
- And in Conclusion, ...

9/24/12 Fall 2012 -- Lecture #13 30

Goals for Floating Point

- Standard arithmetic for reals for all computers
 - Like two's complement
- Keep as much precision as possible in formats
- Help programmer with errors in real arithmetic
 - $+\infty$, $-\infty$, Not-A-Number (NaN), exponent overflow, exponent underflow
- Keep encoding that is somewhat compatible with two's complement
 - E.g., 0 in Fl. Pt. is 0 in two's complement
 - Make it possible to sort without needing to do floating point comparison

9/24/12

Fall 2012 – Lecture #13

31

Floating Point: Representing Very Small Numbers

- Zero: Bit pattern of all 0s is encoding for 0.000
 - ⇒ But 0 in exponent should mean most negative exponent (want 0 to be next to smallest real)
 - ⇒ Can't use two's complement ($1000\ 0000_{\text{two}}$)
- *Bias notation*: subtract bias from exponent
 - Single precision uses bias of 127; DP uses 1023
- 0 uses $0000\ 0000_{\text{two}} \Rightarrow 0-127 = -127$;
 ∞ , NaN uses $1111\ 1111_{\text{two}} \Rightarrow 255-127 = +128$
 - Smallest SP real can represent: $1.00\dots00 \times 2^{-126}$
 - Largest SP real can represent: $1.11\dots11 \times 2^{+127}$

9/24/12

Fall 2012 – Lecture #13

32

What If Operation Result Doesn't Fit in 32 Bits?

- *Overflow*: calculate too big a number to represent within a word
- Unsigned numbers: $1 + 4,294,967,295$ ($2^{32}-1$)
- Signed numbers: $1 + 2,147,483,647$ ($2^{31}-1$)
- How to handle depends on the programming language
 - C signed number arithmetic ignores overflow
 - Other languages want overflow signal on signed numbers (e.g., Fortran)
 - What's a computer architect to do?

9/24/12

Fall 2012 – Lecture #13

33

MIPS Solution: Offer Both

- Instructions that can trigger overflow:
 - add, sub, mult, div, addi, multi, divi
- Instructions that don't overflow are called "unsigned" (really means "no overflow"):
 - addu, subu, multu, divu, addiu, multiu, diviu
- Given semantics of C, always use unsigned versions
- Note: `slt` and `slti` do signed comparisons, while `sltu` and `sltiu` do unsigned comparisons
 - Nothing to do with overflow
 - When would get different answer for `slt` vs. `sltu`?

9/24/12

Fall 2012 – Lecture #13

Student Roulette? 34

What About *Real* Numbers in Base 2?

- $r \times 2^i$, E where i is exponent (2), i is a positive or negative integer, r is a real number ≥ 1.0 , < 2
- Computers version of normalized scientific notation called *Floating Point* notation

9/24/12

Fall 2012 – Lecture #13

35

Floating Point Numbers

- 32-bit word has 2^{32} patterns, so must be approximation of real numbers ≥ 1.0 , < 2
- IEEE 754 Floating Point Standard:
 - 1 bit for *sign* (s) of floating point number
 - 8 bits for *exponent* (E)
 - 23 bits for *fraction* (F)
(get 1 extra bit of precision if leading 1 is implicit)
- $(-1)^s \times (1 + F) \times 2^E$
- Can represent from 2.0×10^{-38} to 2.0×10^{38}

9/24/12

Fall 2012 – Lecture #13

36

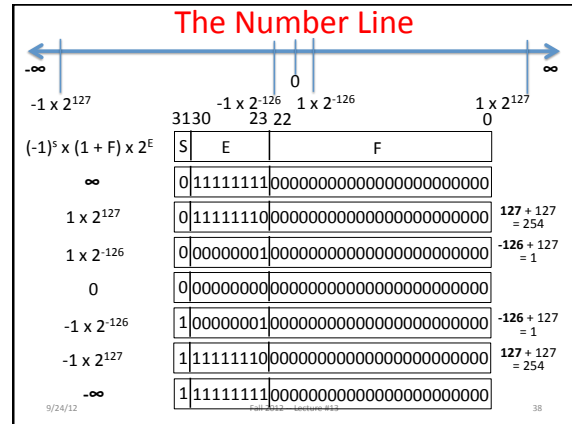
Bias Notation (+127)

	How it is interpreted		How it is encoded	
	Decimal Exponent	signed 2's complement	Biased Notation	Decimal Value of Biased Notation
∞, NaN	For infinities		11111111	255
	127	01111111	11111110	254

	2	00000010	10000001	129
	1	00000001	10000000	128
	0	00000000	01111111	127
	-1	11111111	01111110	126
	-2	11111110	01111101	125

	-126	10000010	00000001	1
Zero	For Denorms	10000001	00000000	0

9/24/12 Fall 2012 - Lecture #13 37



- ### Floating Point Numbers
- What about bigger or smaller numbers?
 - IEEE 754 Floating Point Standard: *Double Precision* (64 bits)
 - 1 bit for *sign* (s) of floating point number
 - 11 bits for *exponent* (E)
 - 52 bits for *fraction* (F)
 - (get 1 extra bit of precision if leading 1 is implicit)
- $(-1)^s \times (1 + F) \times 2^E$
- Can represent from 2.0×10^{-308} to 2.0×10^{308}
 - 32 bit format called *Single Precision*
- 9/24/12 Fall 2012 - Lecture #13 39

- ### More Floating Point
- What about 0?
 - Bit pattern all 0s means 0, so no implicit leading 1
 - What if divide 1 by 0?
 - Can get infinity symbols $+\infty, -\infty$
 - Sign bit 0 or 1, largest exponent, 0 in fraction
 - What if do something stupid? ($\infty - \infty, 0 \div 0$)
 - Can get special symbols NaN for Not-a-Number
 - Sign bit 0 or 1, largest exponent, not zero in fraction
 - What if result is too big? ($2 \times 10^{308} \times 2 \times 10^2$)
 - Get *overflow* in exponent, alert programmer!
 - What if result is too small? ($2 \times 10^{-308} \div 2 \times 10^2$)
 - Get *underflow* in exponent, alert programmer!
- 9/24/12 Fall 2012 - Lecture #13 40

- ### Floating Point Add Associativity?
- $A = (1000000.0 + 0.000001) - 1000000.0$
 - $B = (1000000.0 - 1000000.0) + 0.000001$
 - In single precision floating point arithmetic, A does not equal B
 - $A = 0.000000, B = 0.000001$
 - Floating Point Addition is not Associative!
 - Integer addition is associative
 - When does this matter?
- 9/24/12 Fall 2012 - Lecture #13 41

- ### MIPS Floating Point Instructions
- C, Java has single precision (`float`) and double precision (`double`) types
 - MIPS instructions: `.s` for single, `.d` for double
 - Fl. Pt. Addition single precision: `add.s`
 - Fl. Pt. Addition double precision: `add.d`
 - Fl. Pt. Subtraction single precision: `sub.s`
 - Fl. Pt. Subtraction double precision: `sub.d`
 - Fl. Pt. Multiplication single precision: `mul.s`
 - Fl. Pt. Multiplication double precision: `mul.d`
 - Fl. Pt. Divide single precision: `div.s`
 - Fl. Pt. Divide double precision: `div.d`
- 9/24/12 Fall 2012 - Lecture #13 42

MIPS Floating Point Instructions

- C, Java have single precision (float) and double precision (double) types
- MIPS instructions: .s for single, .d for double
 - Fl. Pt. Comparison single precision:
 - Fl. Pt. Comparison double precision:
 - Fl. Pt. branch:
- Since rarely mix integers and Floating Point, MIPS has separate registers for floating-point operations: \$f0, \$f1, ..., \$f31
 - Double precision uses adjacent even-odd pairs of registers:
 - \$f0 and \$f1, \$f2 and \$f3, \$f4 and \$f5, ..., \$f30 and \$f31
- Need data transfer instructions for these new registers
 - lwc1 (load word), swc1 (store word)
 - Double precision uses two lwc1 instructions, two swc1 instructions

9/24/12 Fall 2012 – Lecture #13 43

Who is Bigger than Whom? What's the Representation?

```

0000 0000 0000 0000 0000 0000 0000 0000two
> 1111 1111 1111 1111 1111 1111 1111 1111two
1111 1111 1111 1111 1111 1111 1111 1110two
> 1111 1111 1111 1111 1111 1111 1111 1111two
1000 0000 0000 0000 0000 0000 0000 0000two
> 1111 1111 0111 1111 1111 1111 1111 1111two
    
```

- Ones Complement
- Twos Complement
- Sign and Magnitude
- IEEE FP Standard
- More than one can apply

9/24/12 Fall 2012 – Lecture #13 44

Who is Bigger than Whom? What's the Representation?

```

0000 0000 0000 0000 0000 0000 0000 0000two
> 1111 1111 1111 1111 1111 1111 1111 1111two
1111 1111 1111 1111 1111 1111 1111 1110two
> 1111 1111 1111 1111 1111 1111 1111 1111two
1000 0000 0000 0000 0000 0000 0000 0000two
> 1111 1111 0111 1111 1111 1111 1111 1111two
    
```

- Ones Complement
- Twos Complement
- Sign and Magnitude
- IEEE FP Standard
- More than one or none can apply

9/24/12 Fall 2012 – Lecture #13 45

Peer Instruction Question

Suppose Big, Tiny, and BigNegative are floats in C, with Big initialized to a big number (e.g., age of universe in seconds or 4.32×10^{17}), Tiny to a small number (e.g., seconds/femtosecond or 1.0×10^{-15}), BigNegative = - Big. Here are two conditionals:

I. $(\text{Big} * \text{Tiny}) * \text{BigNegative} == (\text{Big} * \text{BigNegative}) * \text{Tiny}$
 II. $(\text{Big} + \text{Tiny}) + \text{BigNegative} == (\text{Big} + \text{BigNegative}) + \text{Tiny}$

Which statement about these is correct?

Orange. I. is false and II. is false
 Green. I. is false and II. is true
 Pink. I. is true and II. is false
 Yellow. I. is true and II. is true

9/24/12 Fall 2012 – Lecture #13 46

Peer Instruction Answer

Suppose Big, Tiny, and BigNegative are floats in C, with Big initialized to a big number (e.g., age of universe in seconds or 4.32×10^{17}), Tiny to a small number (e.g., seconds/femtosecond or 1.0×10^{-15}), BigNegative = - Big. Here are two conditionals:

I. $(\text{Big} * \text{Tiny}) * \text{BigNegative} == (\text{Big} * \text{BigNegative}) * \text{Tiny}$
 II. $(\text{Big} + \text{Tiny}) + \text{BigNegative} == (\text{Big} + \text{BigNegative}) + \text{Tiny}$

Which statement about these is correct?

Pink. I. is true and II. is false (if we don't consider overflow) — but there are cases where one side overflows while the other does not!

I. Works ok if no overflow, but because exponents add, if Big * BigNeg overflows, then result is overflow, not -1
 II. Left hand side is 0, right hand side is tiny

9/24/12 Fall 2012 – Lecture #13 47

Pitfalls

- Floating point addition is NOT associative
- Some optimizations can change order of floating point computations, which can change results
- Need to ensure that floating point algorithm is correct even with optimizations

9/24/12 Fall 2012 – Lecture #13 48

And, in Conclusion, ...

- Interpreters for debugging, but slow execution
- Hybrid (Java): Compiler + Interpreter to try to get best of both
- Compiler Optimization to relieve programmer
- Floating point is an approximation of reals