# **Machine-Level Programming: Basics**

CSCI 237: Computer Organization 7<sup>th</sup> Lecture, Feb 24, 2025

Jeannie Albrecht

### Administrative Details

- All puzzles due Tue/Wed at midnight
- Come see me and/or TAs for help!
- Glow HW 2 due Friday
  - Covers signed/unsigned/floating point
- I will be in my office today from 12:45-1:30

# Last Time: Floating Point

- Numerical Form: (-1)<sup>s</sup> \* M \* 2<sup>E</sup>
  - Sign bit s determines whether number is negative or positive
  - Significand (mantissa) M normally a fractional value in range [1.0,2.0).
  - Exponent E weights value by power of two
  - Normalized, denormalized, and special cases
- Encoding
  - s field encodes sign bit s (0 for positive, 1 for negative value)
  - exp field encodes E (but is not equal to E)
  - frac field encodes M (but is not equal to M)

Single precision: 32 bits							
	s exp frac						
	1	8-bits	23-bits				
Double precision: 64 bits							
	s exp frac						
	1	11-bits	52-bits				

# **Today: Machine Programming: Basics**

- Floating point wrapup
- History of Intel processors and architectures
- Assembly Basics: Registers, operands, move
- Arithmetic & logical operations
- C, assembly, machine code



#### 

# C float Decoding Example

float: 0xC0A00000

 $v = (-1)^s \times M \times 2^E$ E = exp – bias

Bias = 2<sup>k-1</sup> – 1 = 127

	1	1000	0001	010	0000	0000	0000	0000	000	0	
	1	8-k	oits			23	-bits			L (	imal
E =									<b>he</b> 0		<b>B</b> ill
<mark>S</mark> =									2	23	001
M =									4 5 6	4 5 6	010 010 011
									7 8 9	7 8 9	011 100 100
<b>v</b> = (-1	) <mark>s</mark> x	M x 2 <sup>E</sup>	:						A B C	10 11 12	101 101 110
- ( '									D E F	13 14 15	110     111     111

# C float Decoding Example

float: 0xC0A00000

 $v = (-1)^s \times M \times 2^E$ E = exp - bias

Bias =  $2^{k-1} - 1 = 127$ 

Ε

F

14

15

1110

1111

	_								
1	1000	0001	010	0000	0000	0000	0000	0000	
1	8-k	oits			23	-bits		.et	ecimalary
E = exp - bias = 129 - 127 = 2									<b>) 0000</b> 0001
$S = 1 \text{ (negative number)} \qquad \qquad \frac{2}{3} \frac{2}{3}$								0010 0011 0100	
$M = 1.010\ 0000\ 0000\ 0000\ 0000\ 0000\ \frac{5}{7}$							5 5 6 6 7 7	0101 0110 0111	
= 1 +	1/4	= 1.2	5					8 8 9 9	1000 1001 1010
<b>v = (-1)</b> <sup>s</sup> :	× M x 2 <sup>₽</sup>	= (-1) <sup>1</sup> *	* 1.25	5 * 2 <sup>2</sup> =	: -5			A         1           B         1           C         1           D         1	$ \begin{array}{c ccccccccccccccccccccccccccccccccccc$

# Floating Point in C

#### C guarantees two levels

- •float single precision
- •double double precision
- Conversions/Casting
  - Casting between int, float, and double changes bit representation
  - double/float  $\rightarrow$  int
    - Truncates fractional part
    - Like rounding toward zero
    - Not defined when out of range or NaN: Generally sets to TMin
  - int  $\rightarrow$  double
    - Exact conversion, as long as int has ≤ 53 bit word size
  - int ightarrow float
    - Will round according to "rounding mode"

# 

### **Floating Point Puzzles**

For each of the following C expressions, either:

- Argue that it is true for all argument values
- Explain why not true

int x = ...;float f = ...; double d = ...;

> Assume neither d nor f is NaN

\* x == (int) (double) x
\* f == (float) (double) f
\* d == (double) (float) d
\* f == -(-f);
\* 2/3 == 2/3.0
\* d < 0.0 ⇒ ((d\*2) < 0.0)
\* d > f ⇒ -f > -d
\* d > f ⇒ -f > -d
\* d \* d >= 0.0
\* (d + f) - d == f

• x == (int) (float) x

### **FP** Summary

- IEEE Floating Point has clear mathematical properties
- Represents numbers of form M x 2<sup>E</sup>
- One can reason about operations independent of implementation
  - As if computed with perfect precision and then rounded
- Not exactly the same as real arithmetic
  - Violates associativity/distributivity
  - Makes life difficult for compilers & serious numerical applications programmers
- Next up: Chapter 3!

Single precision: 32 bits						
s exp	frac					
1 8-bits 23-bits						
Double precision: 64 bits						
s exp	frac					
1 11-bits	52-bits					

# **Today: Machine Programming: Basics**

#### Floating point wrapup

- History of Intel processors and architectures
- Assembly Basics: Registers, operands, move
- Arithmetic & logical operations
- C, assembly, machine code

# Definitions

- Architecture: The parts of a processor design that one needs to understand for writing assembly/machine code.
  - Examples: instruction set specification, registers
  - Example ISA (instruction set architecture) implementations:
    - Intel: x86, IA32, Itanium, x86-64
    - ARM: Used in almost all mobile phones
- Microarchitecture: Implementation of the architecture
  - Examples: cache sizes and core frequency
- Code Forms:
  - Machine Code: The byte-level programs that a processor executes
  - Assembly Code: A text representation of machine code (programming language that is assembled into machine code)

### Intel x86 Processors

- Dominate laptop/desktop/server market (for now)
- Evolutionary design
  - Backwards compatible up until 8086, introduced in 1978
  - Added more features as time goes on
- Complex instruction set computer (CISC)
  - Many different instructions with many different formats
    - But, only small subset encountered with Linux programs
  - Hard to match performance of Reduced Instruction Set Computers (RISC)
  - But, Intel has done just that!
    - In terms of speed. Less so for low power.
- An Aside: Apple's M1/M2 processors use a RISC design
  - Designed using ARM's instruction set (commonly found in smartphones and tablets rather than computers)

#### Transistors

# Intel x86 Evolution: Milestones



Name	Date	Transistors	MHz
<b>8086</b>	1978	29K	5-10
First 16-bit I	ntel processor. Ba	asis for IBM PC & DO	S
1MB addres	s space		
<b>386</b>	1985	275K	16-33
First 32 bit I	ntel processor , re	ferred to as IA32	
Added "flat	addressing", capa	ble of running Unix	
Pentium 4E	2004	125M	2800-3800
First 64-bit I	ntel x86 processo	r, referred to as x86-	64
Core 2	2006	291M	1060-3500
First multi-c	ore Intel processo	r (core = CPU)	
Core i7	2008	731M	1700-3900
Four cores			

### Moore's Law: The number of transistors on microchips doubles every two years Our World in Data

Moore's law describes the empirical regularity that the number of transistors on integrated circuits doubles approximately every two years. This advancement is important for other aspects of technological progress in computing – such as processing speed or the price of computers.



OurWorldinData.org – Research and data to make progress against the world's largest problems.

Licensed under CC-BY by the authors Hannah Ritchie and Max Roser.

### Intel x86 Processors, cont.

Past	Generations	Proce	ess techno	logy	
<b>1</b>	<sup>st</sup> Pentium Pro	1995	60	00 nm	
<b>1</b>	<sup>st</sup> Pentium III	1999	25	50 nm	
<b>1</b>	<sup>st</sup> Pentium 4	2000	18	30 nm	
<b>1</b>	<sup>st</sup> Core 2 Duo	2006	e	65 nm	
Rece	ent Generatio	ons			
1.	Nehalem	2008	Z	l5 nm	
2.	Sandy Bridge	2011	3	32 nm	
3.	lvy Bridge	2012	2	22 nm	
4.	Haswell	2013	2	22 nm	N
5.	Broadwell	2014	1	L4 nm	N
6.	Skylake	2015	1	L4 nm	Ì
7.	Kaby Lake	2016/7	1	L4 nm	
8.	Coffee Lake	2017	1	L4nm	
9.	C.L. Refreshed	2018	1	L4nm	
10.	Ice Lake	2019		10nm	1

Process technology dimension = width of narrowest wires (10 nm ≈ 100 atoms wide)



# 2019 State of the Art: Ice Lake

- Mobile Device: Core i7
  - 1-2.3 GHz
    - Turbo (3.8-4.1 GHz)
  - 9-28 W
  - Integrated Intel Gen 11 GPU
  - 2-4 CPUs
  - 2023 State of the Art is called Meteor Lake and uses a more complex (but similar) design involving tiles



https://www.servethehome.com/intel-ice-lake-era-with-microarchitecture-and-gen11-gpu-improvements/

### **Our Coverage**

#### IA32

The traditional x86 (32 bit)

#### x86-64 (64 bit)

- The standard
- > gcc hello.c
- > gcc -m64 hello.c

#### Presentation

- Book covers x86-64
- "Web aside" on IA32
- We will only cover x86-64

### Levels of Abstraction



Computer scientists love layers and abstraction...

**Computer Designer** 

Caches, clock freq, layout, ...

# Assembly/Machine Code View



**Programmer-Visible State** 

- PC: Program counter
  - Address of next instruction
  - Called "RIP" (x86-64)
- Register file
  - Heavily used program data
- Condition codes
  - Store status information about most recent arithmetic or logical operation
  - Used for conditional branching

#### Memory

- Byte addressable array
- Code and user data
- Stack to support procedures

### Machine Programming: Basics & Ops

- Floating point summary
- History of Intel processors and architectures
- Assembly Basics: Registers, operands, move
- Arithmetic & logical operations
- C, assembly, machine code

### Assembly Characteristics: Data Types

- "Integer" data of 1, 2, 4, or 8 bytes
  - Data values
  - Addresses (untyped pointers)
- Floating point data of 4, 8, or 10 bytes
- Code: Byte sequences encoding series of instructions
- No aggregate types such as arrays or structures



## x86-64 Integer Registers

%rax	<sup>⊗</sup> eax	% <b>r8</b>	%r8d
%rbx	%ebx	% <b>r9</b>	%r9d
%rcx	%ecx	% <b>r10</b>	%r10d
%rdx	%edx	% <b>r11</b>	% <b>r11d</b>
% <b>rsi</b>	%esi	% <b>r12</b>	%r12d
%rdi	%edi	% <b>r13</b>	%r13d
%rsp	%esp	8 <b>r14</b>	%r14d
%rbp	%ebp	% <b>r15</b>	%r15d

- Can reference low-order 4 bytes (also low-order 1 & 2 bytes)
- Not part of memory (or cache); part of CPU

### Some History: IA32 Registers

#### Origin (mostly obsolete)



### Assembly Characteristics: Operations

- Transfer data between memory and register
  - Load data from memory into register
  - Store register data into memory
  - Data stored in registers is much faster to access than memory
- Perform arithmetic functions on register or memory data

#### Transfer control

- Unconditional jumps to/from procedures
- Conditional branches



# Moving Data (Ch 3.4)

# Moving Data movq Source, Dest

- Operand Types for source and dest
  - Immediate: Constant integer data
    - Example: \$0x400, \$-533
    - Like C constant, but prefixed with `\$'
    - Encoded with 1, 2, or 4 bytes
  - Register: One of 16 integer registers
    - Example: %rax, %r13
    - But %rsp reserved for special use
    - Others have special uses for particular instructions
  - Memory: 8 consecutive bytes of memory at address given by register
    - Simplest example: (%rax)
    - Various other "addressing modes"
    - Note the parentheses

2	erax
2	ercx
2	&rdx
!	%rbx
	8 <b>rsi</b>
9	& <b>rdi</b>
	% <b>rsp</b>
	%rbp





# Moving Data (Ch 3.4)

Moving Data
 movqSource, Dest

Q = quad word (8 bytes) L = double word (4 bytes) W = word (2 bytes – historical!)

Operand Types for s B = byte (1 byte)

Immediate: Constant integer data

- Example: \$0x400, \$-533
- Like C constant, but prefixed with `\$'
- Encoded with 1, 2, or 4 bytes

Register: One of 16 integer registers

- Example: %rax, %r13
- But %rsp reserved for special use
- Others have special uses for particular instructions
- Memory: 8 consecutive bytes of memory at address given by register
  - Simplest example: (%rax)
  - Various other "addressing modes"
  - Note the parentheses

	%rax
	%rcx
)	%rdx
,	%rbx
	% <b>rsi</b>
	%rdi
	%rsp
	%rbp





### **Pointer Recap**

- int \*p; //variable p is a pointer to an integer
- int i; // integer value
- You dereference a pointer to get value with \*:
- You find address of value with &:
- Worth spending a few minutes digesting the concept of pointers!
   Check K&R for more info









NOTE: Cannot do memory-memory transfer with a single instruction!

# Simple Memory Addressing Modes

- $\blacksquare Normal \qquad (Reg[R]) \longrightarrow Mem[Reg[R]]$ 
  - Register R specifies memory address
  - Like pointer dereferencing in C

#### movq (%rcx), %rax

- Displacement D(Reg[R]) → Mem[Reg[R]+D]
  - Register R specifies start of memory region
  - Constant displacement D specifies offset

```
movq 8(%rbp),%rdx
```