

Assembly Programming I

CSE 351 Spring 2017

Instructor:

Ruth Anderson

Teaching Assistants:

Dylan Johnson

Kevin Bi

Linxing Preston Jiang

Cody Ohlsen

Yufang Sun

Joshua Curtis

Administrivia

- ❖ Lab 1 due Friday (4/14)
 - Prelim submission (3+ of `bits.c`) due on TONIGHT (4/10). Turn in whatever you have at that time (drop box closes at 11:59pm), no lates. Worth a small part (no more than 10%) of total points for lab 1.
- ❖ Homework 2 due next Wednesday (4/19)

Roadmap

C:

```
car *c = malloc(sizeof(car));
c->miles = 100;
c->gals = 17;
float mpg = get_mpg(c);
free(c);
```

Java:

```
Car c = new Car();
c.setMiles(100);
c.setGals(17);
float mpg =
    c.getMPG();
```

- Memory & data
- Integers & floats
- x86 assembly**
- Procedures & stacks
- Executables
- Arrays & structs
- Memory & caches
- Processes
- Virtual memory
- Memory allocation
- Java vs. C

Assembly language:

```
get_mpg:
    pushq    %rbp
    movq    %rsp, %rbp
    ...
    popq    %rbp
    ret
```


Machine code:

```
0111010000011000
100011010000010000000010
1000100111000010
110000011111101000011111
```

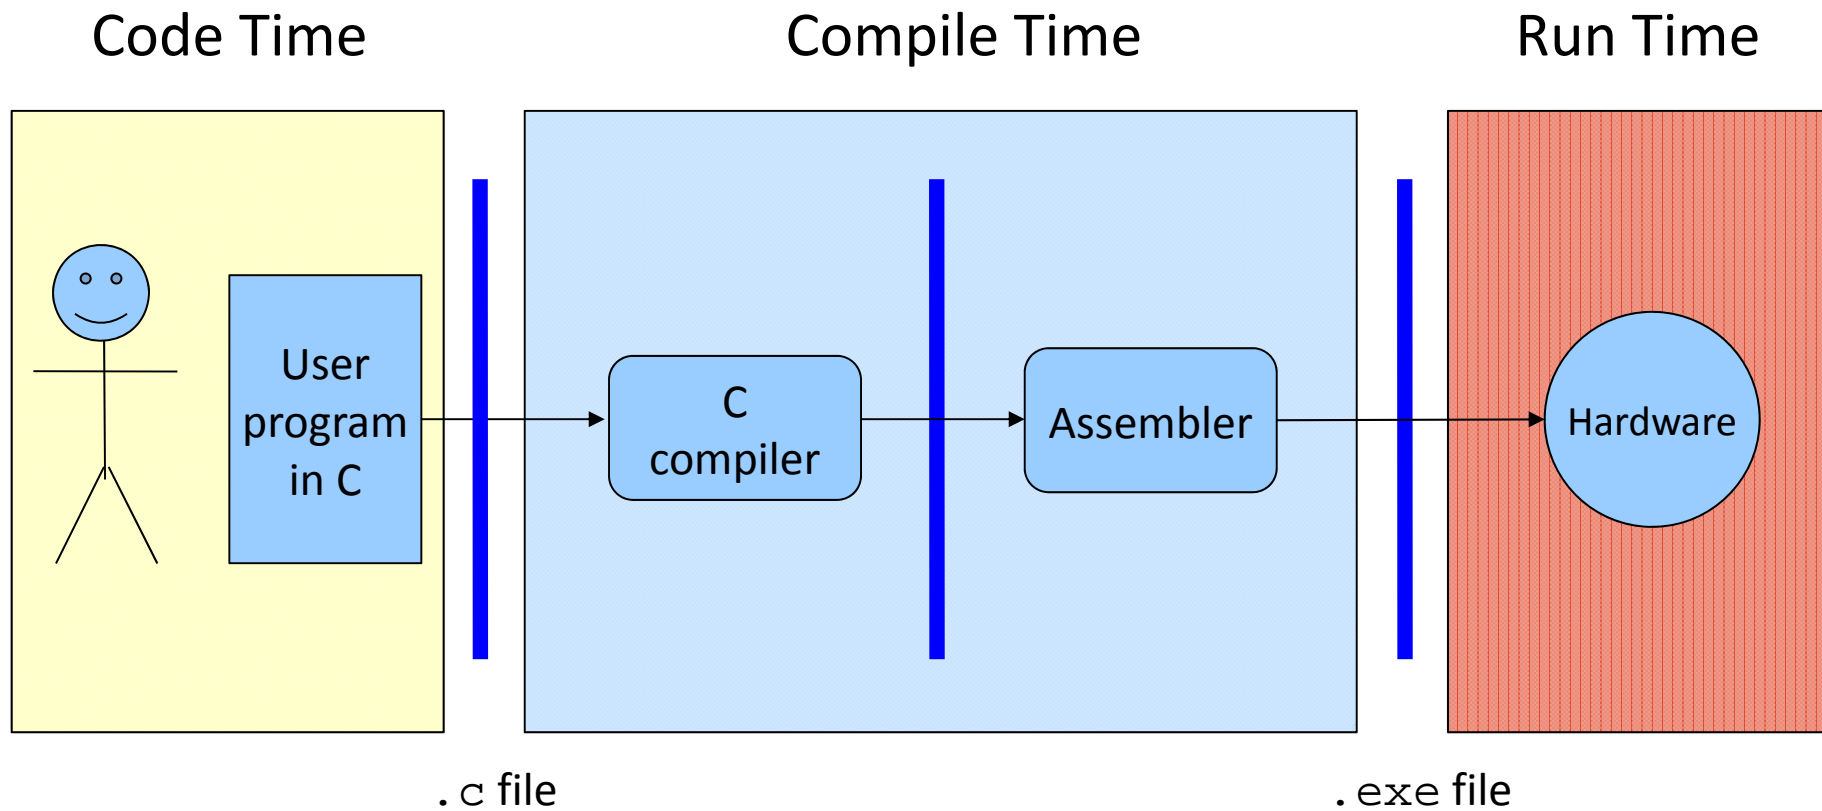
OS:



Computer system:



Translation



What makes programs run fast(er)?

HW Interface Affects Performance

Source code

Different applications or algorithms

Compiler

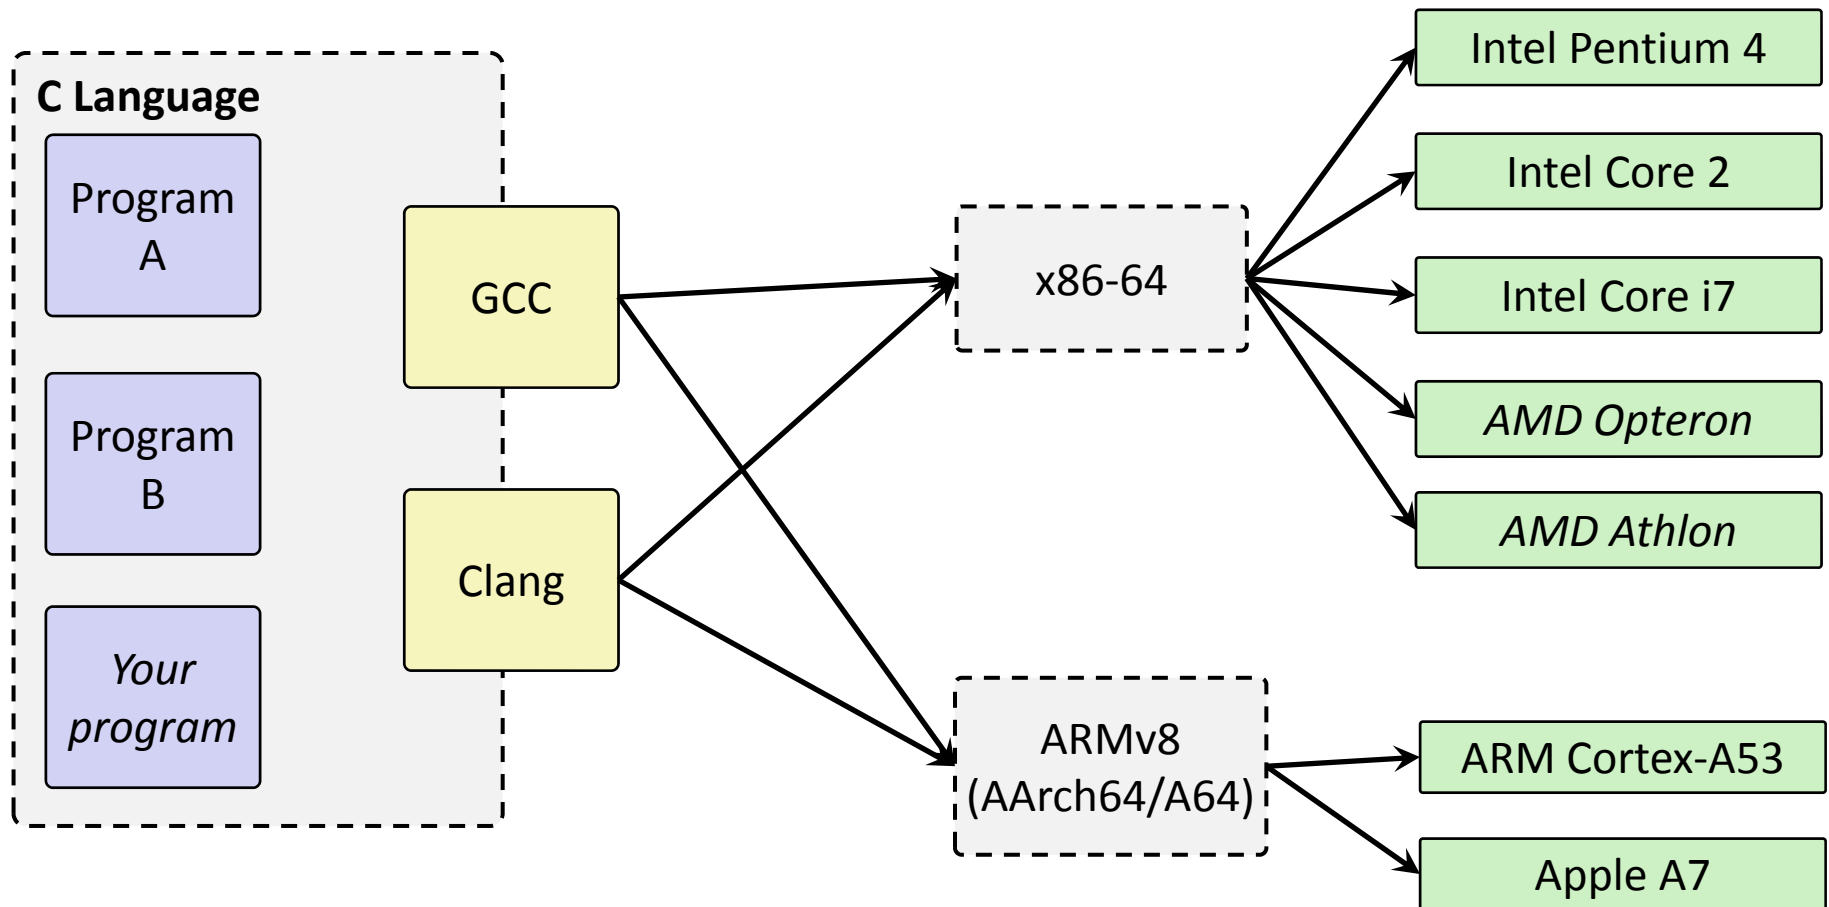
Perform optimizations, generate instructions

Architecture

Instruction set


Hardware

Different implementations



Instruction Set Architectures

- ❖ The ISA defines:
 - The system's state (*e.g.* registers, memory, program counter)
 - The instructions the CPU can execute
 - The effect that each of these instructions will have on the system state



Instruction Set Philosophies

- ❖ *Complex Instruction Set Computing (CISC)*: Add more and more elaborate and specialized instructions as needed
 - Lots of tools for programmers to use, but hardware must be able to handle all instructions
 - x86-64 is CISC, but only a small subset of instructions encountered with Linux programs
- ❖ *Reduced Instruction Set Computing (RISC)*: Keep instruction set small and regular
 - Easier to build fast hardware
 - Let software do the complicated operations by composing simpler ones

General ISA Design Decisions

- ❖ Instructions
 - What instructions are available? What do they do?
 - How are they encoded?

- ❖ Registers
 - How many registers are there?
 - How wide are they?

- ❖ Memory
 - How do you specify a memory location?

General ISA Design Decisions

- ❖ Instructions
 - What instructions are available? What do they do?
 - How are they encoded? **Instructions are data!**

- ❖ Registers
 - How many registers are there?
 - How wide are they? **Size of a word**

- ❖ Memory
 - How do you specify a memory location? **Different ways to build up an address**

Mainstream ISAs



x86

Designer	Intel, AMD
Bits	16-bit, 32-bit and 64-bit
Introduced	1978 (16-bit), 1985 (32-bit), 2003 (64-bit)
Design	CISC
Type	Register-memory
Encoding	Variable (1 to 15 bytes)
Endianness	Little

Macbooks & PCs
(Core i3, i5, i7, M)
[x86-64 Instruction Set](#)



ARM architectures

Designer	ARM Holdings
Bits	32-bit, 64-bit
Introduced	1985; 31 years ago
Design	RISC
Type	Register-Register
Encoding	AArch64/A64 and AArch32/A32 use 32-bit instructions, T32 (Thumb-2) uses mixed 16- and 32-bit instructions. ARMv7 user-space compatibility ^[1]
Endianness	Bi (little as default)

Smartphone-like devices
(iPhone, iPad, Raspberry Pi)
[ARM Instruction Set](#)



MIPS

Designer	MIPS Technologies, Inc.
Bits	64-bit (32→64)
Introduced	1981; 35 years ago
Design	RISC
Type	Register-Register
Encoding	Fixed
Endianness	Bi

Digital home & networking
equipment
(Blu-ray, PlayStation 2)
[MIPS Instruction Set](#)

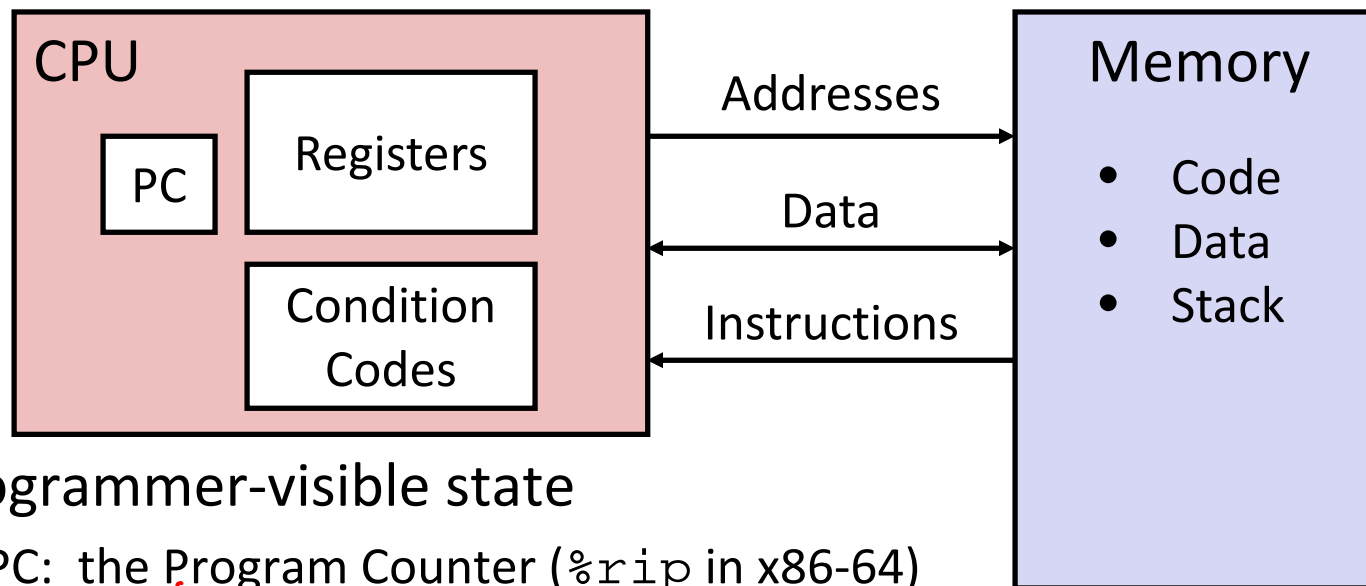
Definitions

- ❖ **Architecture (ISA):** The parts of a processor design that one needs to understand to write assembly code
 - “What is directly visible to software”
- ❖ **Microarchitecture:** Implementation of the architecture
 - CSE/EE 469, 470
- ❖ Are the following part of the architecture?
 - Number of registers?
 - How about CPU frequency?
 - Cache size? Memory size?

Definitions

- ❖ **Architecture (ISA):** The parts of a processor design that one needs to understand to write assembly code
 - “What is directly visible to software”
- ❖ **Microarchitecture:** Implementation of the architecture
 - CSE/EE 469, 470
- ❖ Are the following part of the architecture?
 - Number of registers? **Yes**
 - How about CPU frequency? **No**
 - Cache size? Memory size? **No**

Assembly Programmer's View



❖ Programmer-visible state

- PC: the Program Counter (`%rip` in x86-64)
 - Address of next instruction
- Named registers
 - Together in “register file”
 - Heavily used program data
- Condition codes
 - Store status information about most recent arithmetic operation
 - Used for conditional branching

❖ Memory

- Byte-addressable array
- Code and user data
- Includes *the Stack* (for supporting procedures)

x86-64 Assembly “Data Types”

- ❖ Integral data of 1, 2, 4, or 8 bytes
 - Data values
 - Addresses (untyped pointers)
 - ❖ Floating point data of 4, 8, 10 or 2x8 or 4x4 or 8x2
 - Different registers for those (e.g. `%xmm1`, `%ymm2`)
 - Come from *extensions to x86* (SSE, AVX, ...)
 - ❖ No aggregate types such as arrays or structures
 - Just contiguously allocated bytes in memory
 - ❖ Two common syntaxes
 - “AT&T”: used by our course, slides, textbook, gnu tools, ...
 - “Intel”: used by Intel documentation, Intel tools, ...
 - Must know which you’re reading
- } Not covered
In CSE 351

What is a Register?

- ❖ A location in the CPU that stores a small amount of data, which can be accessed very quickly (once every clock cycle)
- ❖ Registers have names, not *addresses*
 - In assembly, they start with % (e.g. %rsi)
- ❖ Registers are at the heart of assembly programming
 - They are a precious commodity in all architectures, but *especially* x86

x86-64 Integer Registers – 64 bits wide

%rax	%eax
%rbx	%ebx
%rcx	%ecx
%rdx	%edx
%rsi	%esi
%rdi	%edi
%rsp	%esp
%rbp	%ebp

%r8	%r8d
%r9	%r9d
%r10	%r10d
%r11	%r11d
%r12	%r12d
%r13	%r13d
%r14	%r14d
%r15	%r15d

- Can reference low-order 4 bytes (also low-order 2 & 1 bytes)

Some History: IA32 Registers – 32 bits wide



Memory

- ❖ Addresses
 - `0x7FFFD024C3DC`
- ❖ Big
 - `~ 8 GiB`
- ❖ Slow
 - `~50-100 ns`
- ❖ Dynamic
 - Can “grow” as needed while program runs

vs. Registers

- vs. Names
 - `%rdi`
- vs. Small
 - `(16 x 8 B)` = 128 B
- vs. Fast
 - sub-nanosecond timescale
- vs. Static
 - fixed number in hardware

Three Basic Kinds of Instructions

1) Transfer data between memory and register

- *Load* data from memory into register
 - `%reg = Mem[address]`
- *Store* register data into memory
 - `Mem[address] = %reg`

Remember: Memory is indexed just like an array of bytes!

2) Perform arithmetic operation on register or memory data

- `c = a + b;` `z = x << y;` `i = h & g;`

3) Control flow: what instruction to execute next

- Unconditional jumps to/from procedures
- Conditional branches

Operand types

- ❖ **Immediate:** Constant integer data
 - Examples: `$0x400`, `$-533`
 - Like C literal, but prefixed with ``$'`
 - Encoded with 1, 2, 4, or 8 bytes *depending on the instruction*
- ❖ **Register:** 1 of 16 integer registers
 - Examples: `%rax`, `%r13`
 - But `%rsp` reserved for special use
 - Others have special uses for particular instructions
- ❖ **Memory:** Consecutive bytes of memory at a computed address
 - Simplest example: `(%rax)`
 - Various other “address modes”

`%rax``%rcx``%rdx``%rbx``%rsi``%rdi``%rsp``%rbp``%rN`

Moving Data

- ❖ General form: `mov_ source, destination`
 - Missing letter (`_`) specifies size of operands
 - Note that due to backwards-compatible support for 8086 programs (16-bit machines!), “word” means 16 bits = 2 bytes in x86 instruction names
 - Lots of these in typical code
- ❖ `movb src, dst`
 - Move 1-byte “byte”
- ❖ `movw src, dst`
 - Move 2-byte “word”
- ❖ `movl src, dst`
 - Move 4-byte “long word”
- ❖ `movq src, dst`
 - Move 8-byte “quad word”

movq Operand Combinations

	Source	Dest	Src, Dest	C Analog
movq	Imm	Reg	movq \$0x4, %rax	var_a = 0x4;
		Mem	movq \$-147, (%rax)	*p_a = -147;
	Reg	Reg	movq %rax, %rdx	var_d = var_a;
		Mem	movq %rax, (%rdx)	*p_d = var_a;
	Mem	Reg	movq (%rax), %rdx	var_d = *p_a;

❖ *Cannot do memory-memory transfer with a single instruction*

- How would you do it?

Question

- ❖ Which of the following statements is TRUE?
 - A. For float f , $(f+2 == f+1+1)$ always returns TRUE
 - B. The width of a “word” is part of a system’s *architecture* (as opposed to *microarchitecture*)
 - C. Having more registers increases the performance of the hardware, but decreases the performance of the software
 - D. Mem to Mem (src to dst) is the only disallowed operand combination in x86-64

↳ Sounds true, but cannot move anything to an Immediate either 😊

Summary

- ❖ Converting between integral and floating point data types *does* change the bits
 - Floating point rounding is a HUGE issue!
 - Limited mantissa bits cause inaccurate representations
 - Floating point arithmetic is NOT associative or distributive
- ❖ x86-64 is a complex instruction set computing (CISC) architecture
- ❖ **Registers** are named locations in the CPU for holding and manipulating data
 - x86-64 uses 16 64-bit wide registers
- ❖ Assembly operands include immediates, registers, and data at specified memory locations