

UNIVERSITY of WASHINGTON | L02: Memory & Data I | CSE391, Winter 2018

Memory, Data, & Addressing I

CSE 351 Winter 2018

Instructor:
Mark Wyse

Teaching Assistants:
Kevin Bi
Parker DeWilde
Emily Furst
Sarah House
Waylon Huang
Vinny Palaniappan

<http://xkcd.com/953/>

UNIVERSITY of WASHINGTON | L02: Memory & Data I | CSE391, Winter 2018

Administrivia

- ❖ Pre-Course Survey due tonight @ 11:59pm
- ❖ Lab 0 due Monday (1/8)
- ❖ Homework 1 due Wednesday (1/10)

- ❖ All course materials can be found on the website/schedule
- ❖ Course Overloads – fill out the Google Form linked in lecture 1!
- ❖ Make sure you're also enrolled in CSE391 (EEs and non-majors included!)

UNIVERSITY of WASHINGTON | L02: Memory & Data I | CSE391, Winter 2018

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
1000110100001000000010
1000100111000010
1100000111110100001111
```

OS: Windows 10, OS X, Linux

Computer system: CPU, RAM, Hard Drive

UNIVERSITY of WASHINGTON | L02: Memory & Data I | CSE391, Winter 2018

Hardware: Logical View

```

    graph TD
        CPU((CPU)) --- Bus[Bus]
        Memory[Memory] --- Bus
        Disks[(Disks)] --- Bus
        Net((Net)) --- Bus
        USB[USB] --- Bus
        Etc[Etc.] --- Bus
    
```

UNIVERSITY of WASHINGTON | L02: Memory & Data I | CSE391, Winter 2018

Hardware: Physical View

Labels: Bus connections, CPU (socket), I/O controller, Memory

Storage connections

UNIVERSITY of WASHINGTON | L02: Memory & Data I | CSE391, Winter 2018

Hardware: 351 View (version 0)

- ❖ CPU executes instructions; memory stores data
- ❖ To execute an instruction, the CPU must:
 - fetch an instruction;
 - fetch the data used by the instruction; and, finally,
 - execute the instruction on the data...
 - which may result in writing data back to memory

UNIVERSITY of WASHINGTON | L02: Memory & Data | CSE351, Winter 2018

Hardware: 351 View (version 1)

- ❖ The CPU holds instructions temporarily in the **instruction cache**
- ❖ The CPU holds data temporarily in a fixed number of **registers**
- ❖ **Instruction and operand fetching** is hardware-controlled
- ❖ **Data movement** is programmer-controlled (in assembly)
- ❖ We'll learn about the instructions the CPU executes – take CSE/EE 469 and 470 to find out how it actually executes them

7

UNIVERSITY of WASHINGTON | L02: Memory & Data | CSE351, Winter 2018

Hardware: 351 View (version 1)

- ❖ The CPU holds instructions temporarily in the **instruction cache**
- ❖ The CPU holds data temporarily in a fixed number of **registers**
- ❖ **Instruction and operand fetching** is hardware-controlled
- ❖ **Data movement** is programmer-controlled (in assembly)
- ❖ We'll learn about the instructions the CPU executes – take CSE/EE 469 and 470 to find out how it actually executes them

8

UNIVERSITY of WASHINGTON | L02: Memory & Data | CSE351, Winter 2018

Question 1:

- ❖ **Binary Encoding!**

9

UNIVERSITY of WASHINGTON | L02: Memory & Data | CSE351, Winter 2018

Question 1: Some Additional Details

- ❖ Because storage is finite in reality, everything is stored as “fixed” length
 - Data is moved and manipulated in fixed-length chunks
 - Multiple fixed lengths (e.g. 1 byte, 4 bytes, 8 bytes)
 - Leading zeros now *must* be included up to “fill out” the fixed length
- ❖ **Example:** the “eight-bit” representation of the number 4 is 0b00000100
 - Most Significant Bit (MSB)
 - Least Significant Bit (LSB)

10

UNIVERSITY of WASHINGTON | L02: Memory & Data | CSE351, Winter 2018

Question 2:

11

UNIVERSITY of WASHINGTON | L02: Memory & Data | CSE351, Winter 2018

Byte-Oriented Memory Organization

- ❖ Conceptually, memory is a single, large array of bytes, each with a unique **address** (index)
 - The value of each byte in memory can be read and written
- ❖ Programs refer to bytes in memory by their **addresses**
 - Domain of possible addresses = **address space**
- ❖ But not all values fit in a single byte... (e.g. 351)
 - Many operations actually use multi-byte values
- ❖ We can store addresses as data to “remember” where other data is in memory

12

UNIVERSITY of WASHINGTON | L02: Memory & Data | CSE351, Winter 2018

Peer Instruction Question

- ❖ If we choose to use 8-bit addresses, how big is our address space?
 - *i.e.* How much space can we “refer to” using our addresses?

A. 256 bits
 B. 256 bytes
 C. 8 bits
 D. 8 bytes
 E. We’re lost...

13

UNIVERSITY of WASHINGTON | L02: Memory & Data | CSE351, Winter 2018

Machine “Words”

- ❖ Instructions encoded into machine code (0’s and 1’s)
 - Historically (still true in some assembly languages), all instructions were exactly the size of a **word**
- ❖ Word size bounds the size of the *address space*
 - word size = address size = register size
 - word size = w bits $\rightarrow 2^w$ addresses
- ❖ Current x86 systems use **64-bit (8-byte) words**
 - Potential address space: 2^{64} addresses
 2^{64} bytes $\approx 1.8 \times 10^{19}$ bytes
 $= 18$ billion billion bytes = 18 EB (exabytes)
 - Actual physical address space: **48 bits**

14

UNIVERSITY of WASHINGTON | L02: Memory & Data | CSE351, Winter 2018

Word-Oriented Memory Organization

- ❖ Addresses still specify locations of *bytes* in memory
 - Addresses of successive words differ by word size (in bytes); *e.g.* 4 (32-bit) or 8 (64-bit)
 - Address of word 0, 1, ... 10?

15

UNIVERSITY of WASHINGTON | L02: Memory & Data | CSE351, Winter 2018

Word-Oriented Memory Organization

- ❖ Addresses still specify locations of *bytes* in memory
 - Addresses of successive words differ by word size (in bytes); *e.g.* 4 (32-bit) or 8 (64-bit)
 - Address of word 0, 1, ... 10?
- ❖ **Address of word** = address of *first* byte in word
 - The address of *any* chunk of memory is given by the address of the first byte
 - **Alignment**

16

UNIVERSITY of WASHINGTON | L02: Memory & Data | CSE351, Winter 2018

A Picture of Memory (64-bit view)

- ❖ A “64-bit (8-byte) word-aligned” view of memory:
 - In this type of picture, each row is composed of 8 bytes
 - Each cell is a byte
 - A 64-bit pointer will fit on one row

17

UNIVERSITY of WASHINGTON | L02: Memory & Data | CSE351, Winter 2018

A Picture of Memory (64-bit view)

- ❖ A “64-bit (8-byte) word-aligned” view of memory:
 - In this type of picture, each row is composed of 8 bytes
 - Each cell is a byte
 - A 64-bit pointer will fit on one row

18

UNIVERSITY of WASHINGTON | L02: Memory & Data | CSE351, Winter 2018

Addresses and Pointers

64-bit example (pointers are 64-bits wide) big-endian

- An *address* is a location in memory
- A *pointer* is a data object that holds an address
 - Address can point to *any* data
- Value 351 stored at address 0x08
 - $351_{10} = 15F_{16} = 0x00 \dots 00015F$
- Pointer stored at 0x38 points to address 0x08

Address: 0x00, 0x08, 0x10, 0x18, 0x20, 0x28, 0x30, 0x38, 0x40, 0x48

19

UNIVERSITY of WASHINGTON | L02: Memory & Data | CSE351, Winter 2018

Addresses and Pointers

64-bit example (pointers are 64-bits wide) big-endian

- An *address* is a location in memory
- A *pointer* is a data object that holds an address
 - Address can point to *any* data
- Pointer stored at 0x48 points to address 0x38
 - Pointer to a pointer!
- Is the data stored at 0x08 a pointer?
 - Could be, depending on how you use it

Address: 0x00, 0x08, 0x10, 0x18, 0x20, 0x28, 0x30, 0x38, 0x40, 0x48

20

UNIVERSITY of WASHINGTON | L02: Memory & Data | CSE351, Winter 2018

Data Representations

- Sizes of data types (in bytes)

Java Data Type	C Data Type	32-bit (old)	x86-64
boolean	bool	1	1
byte	char	1	1
char		2	2
short	short int	2	2
int	int	4	4
float	float	4	4
	long int	4	8
double	double	8	8
long	long	8	8
	long double	8	16
(reference)	pointer *	4	8

To use "bool" in C, you must #include <stdbool.h>

address size = word size

21

UNIVERSITY of WASHINGTON | L02: Memory & Data | CSE351, Winter 2018

More on Memory Alignment in x86-64

- For good memory system performance, Intel recommends data be aligned
 - However the x86-64 hardware will work correctly regardless of alignment of data
 - Design choice: x86-64 instructions are *variable* bytes long
- Aligned:** Primitive object of *K* bytes must have an address that is a multiple of *K*
 - More about alignment later in the course

K	Type
1	char
2	short
4	int, float
8	long, double, pointers

22

UNIVERSITY of WASHINGTON | L02: Memory & Data | CSE351, Winter 2018

Byte Ordering

- How should bytes within a word be ordered *in memory*?
 - Example:** store the 4-byte (32-bit) int: `0x a1 b2 c3 d4`
- By convention, ordering of bytes called *endianness*
 - The two options are **big-endian** and **little-endian**
 - In which address does the least significant *byte* go?
 - Based on *Gulliver's Travels*: tribes cut eggs on different sides (big, little)

23

UNIVERSITY of WASHINGTON | L02: Memory & Data | CSE351, Winter 2018

Byte Ordering

- Big-endian (SPARC, z/Architecture)
 - Least significant byte has highest address
- Little-endian (x86, x86-64)
 - Least significant byte has lowest address
- Bi-endian (ARM, PowerPC)
 - Endianness can be specified as big or little (default)
- Example:** 4-byte data 0xa1b2c3d4 at address 0x100

Big-Endian: 0x100 (a1), 0x101 (b2), 0x102 (c3), 0x103 (d4)

Little-Endian: 0x100 (d4), 0x101 (c3), 0x102 (b2), 0x103 (a1)

24

UNIVERSITY of WASHINGTON | L02: Memory & Data | CSE351, Winter 2018

Byte Ordering Examples

```
int x = 12345;
// or x = 0x3039;
```

IA32, x86-64 (little-endian) | SPARC (big-endian)

0x00	39	00	00
0x01	30	00	00
0x02	00	30	00
0x03	00	39	00

```
long int y = 12345;
// or y = 0x3039;
```

32-bit IA32 | 64-bit x86-64 | 32-bit SPARC | 64-bit SPARC

0x00	39	39	00	00	00
0x01	30	30	00	00	00
0x02	00	00	00	00	00
0x03	00	00	30	00	00
0x04		00	39	00	00
0x05		00	00	00	00
0x06		00	00	30	00
0x07		00	00	39	00

(A long int is the size of a word)

25

UNIVERSITY of WASHINGTON | L02: Memory & Data | CSE351, Winter 2018

Peer Instruction Question:

- ❖ We store the value $0 \times 01\ 02\ 03\ 04$ as a **word** at address 0×100 in a big-endian, 64-bit machine
- ❖ What is the **byte of data** stored at address 0×104 ?

A. 0×04
 B. 0×40
 C. 0×01
 D. 0×10
 E. We're lost...

64-bit SPARC

0x100	
0x101	
0x102	
0x103	
0x104	
0x105	
0x106	
0x107	

26

UNIVERSITY of WASHINGTON | L02: Memory & Data | CSE351, Winter 2018

Endianness

- ❖ **Endianness only applies to memory storage**
- ❖ Often programmer can ignore endianness because it is handled for you
 - Bytes wired into correct place when reading or storing from memory (hardware)
 - Compiler and assembler generate correct behavior (software)
- ❖ Endianness still shows up:
 - Logical issues: accessing different amount of data than how you stored it (e.g. store int, access byte as a char)
 - Need to know exact values to debug memory errors
 - Manual translation to and from machine code (in 351)

27

UNIVERSITY of WASHINGTON | L02: Memory & Data | CSE351, Winter 2018

Summary

- ❖ Memory is a long, *byte-addressed array*
 - Word size bounds the size of the *address space* and memory
 - Different data types use different number of bytes
 - Address of chunk of memory given by address of lowest byte in chunk
 - Object of K bytes is *aligned* if it has an address that is a multiple of K
- ❖ Pointers are data objects that hold addresses
- ❖ Endianness determines memory storage order for multi-byte data

28