

# Memory, Data, & Addressing I

CSE 351, Summer 2019

## Instructor:


Sam Wolfson

## Teaching Assistants:

Rehaan Bhimani

Daniel Hsu

Corbin Modica



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

# Administrivia

- ❖ Pre-Course Survey due tonight @ 11:59 pm
- ❖ Homework 1 due Friday (6/28)
- ❖ Lab 0 out today, due Monday (7/1)
  
- ❖ All course materials can be found on the website schedule (check it out!)
  
- ❖ **Get your machine set up for this class (VM or attu) as soon as possible!**
  - Bring your laptop to section tomorrow if you are having trouble.

# Converting to Base 10

- ❖ Can convert from any base *to* base 10
  - $0b110 = 110_2 = (1 \times 2^2) + (1 \times 2^1) + (0 \times 2^0) = 6_{10}$
  - $0xA5 = A5_{16} = (10 \times 16^1) + (5 \times 16^0) = 165_{10}$
- ❖ We learned to think in base 10, so this is fairly natural for us
- ❖ **Challenge:** Convert into other bases (e.g. 2, 16)

# Challenge Question

❖ Convert  $13_{10}$  into binary

❖ Hints:

■  $2^3 = 8$

■  $2^2 = 4$

■  $2^1 = 2$

■  $2^0 = 1$

**$13_{10} = ?$**

**$13 = 8 + 4 + 1$**

**Binary: 0b 1 1 0 1**

**Dec:            8 4   1**

❖ Think on your own for a minute, then discuss with your neighbor(s)

■ No voting for this question.

# Converting from Decimal to Binary

- ❖ Given a decimal number  $N$ :
  - List increasing powers of 2 from *right to left* until  $\geq N$
  - Then from *left to right*, ask is that (power of 2)  $\leq N$ ?
    - If **YES**, put a 1 below and subtract that power from  $N$
    - If **NO**, put a 0 below and keep going

❖ Example: 13 to binary

$2^4=16$	$2^3=8$	$2^2=4$	$2^1=2$	$2^0=1$
<b>0</b>	<b>1</b>	<b>1</b>	<b>0</b>	<b>1</b>

# Converting from Decimal to Base B

- ❖ Given a decimal number  $N$ :
  - List increasing powers of  $B$  from *right to left* until  $\geq N$
  - Then from *left to right*, ask is that (power of  $B$ )  $\leq N$ ?
    - If **YES**, put *how many of that power go into N* and subtract from  $N$
    - If **NO**, put a 0 below and keep going

- ❖ Example: 165 to hex

$16^2=256$	$16^1=16$	$16^0=1$
<b>0</b>	<b>A</b>	<b>5</b>

# Converting Binary $\leftrightarrow$ Hexadecimal

## ❖ Hex $\rightarrow$ Binary

- Substitute hex digits, then drop any **leading zeros**
- Example: 0x2D to binary
  - 0x2 is 0b0010, 0xD is 0b1101
  - Drop two leading zeros, answer is 0b101101

## ❖ Binary $\rightarrow$ Hex

- Pad with **leading zeros** until multiple of 4, then substitute each group of 4
- Example: 0b101101
  - Pad to 0b 0010 1101
  - Substitute to get 0x2D

Base 10	Base 2	Base 16
0	0000	0
1	0001	1
2	0010	2
3	0011	3
4	0100	4
5	0101	5
6	0110	6
7	0111	7
8	1000	8
9	1001	9
10	1010	A
11	1011	B
12	1100	C
13	1101	D
14	1110	E
15	1111	F

# Binary → Hex Practice

- ❖ Convert 0b100110110101101
  - How many digits? **15**
  - Pad: **0100 1101 1010 1101**
  - Substitute: **0x4DAD**

Base 10	Base 2	Base 16
0	0000	0
1	0001	1
2	0010	2
3	0011	3
4	0100	4
5	0101	5
6	0110	6
7	0111	7
8	1000	8
9	1001	9
10	1010	A
11	1011	B
12	1100	C
13	1101	D
14	1110	E
15	1111	F



# Base Comparison

- ❖ Why does all of this matter?
  - *Humans* think about numbers in **base 10**, but *computers* “think” about numbers in **base 2**
  - **Binary encoding** is what allows computers to do all of the amazing things that they do!
- ❖ You should have this table memorized by the end of the class
  - Might as well start now!

Base 10	Base 2	Base 16
0	0000	0
1	0001	1
2	0010	2
3	0011	3
4	0100	4
5	0101	5
6	0110	6
7	0111	7
8	1000	8
9	1001	9
10	1010	A
11	1011	B
12	1100	C
13	1101	D
14	1110	E
15	1111	F

# Numerical Encoding

- ❖ **AMAZING FACT: You can represent *anything* countable using numbers!**
  - Need to agree on an **encoding**
  - Kind of like learning a new language
- ❖ Examples:
  - Decimal Integers:  $0 \rightarrow 0b0$ ,  $1 \rightarrow 0b1$ ,  $2 \rightarrow 0b10$ , etc.
  - English Letters: CSE  $\rightarrow 0x435345$ , yay  $\rightarrow 0x796179$
  - Emoticons: 😊 0x0, 😞 0x1, 😎 0x2, 😇 0x3, 😈 0x4, 🙋 0x5

# Binary Encoding

- ❖ With  $N$  binary digits, how many “things” can you represent?
  - Need  $N$  binary digits to represent  $n$  things, where  $2^N \geq n$
  - Example: 5 binary digits for alphabet because  $2^5 = 32 > 26$
- ❖ A binary digit is known as a **bit**
- ❖ A group of 4 bits (1 hex digit) is called a **nybble**
- ❖ A group of 8 bits (2 hex digits) is called a **byte**
  - 1 bit  $\rightarrow$  2 things, 1 nybble  $\rightarrow$  16 things, 1 byte  $\rightarrow$  256 things

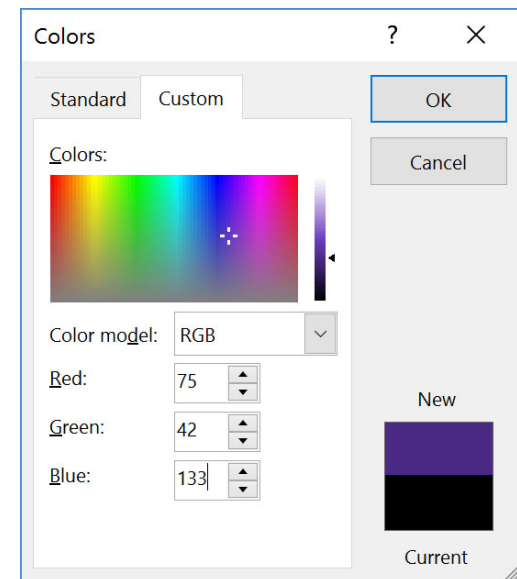

# So What's It Mean?

- ❖ *A sequence of bits can have many meanings!*
- ❖ Consider the hex sequence 0x4E6F21
  - Common interpretations include:
    - The decimal number 5140257
    - The characters “No!”
    - The background color of this slide
    - The real number  $7.203034 \times 10^{-39}$
- ❖ It is up to the program/programmer to decide how to **interpret** the sequence of bits

# Binary Encoding – Colors

## ❖ RGB – Red, Green, Blue

- Additive color model (light): byte (8 bits) for each color
- Commonly seen in hex (in HTML, photo editing, etc.)
- Examples: **Blue**→0x0000FF, **Gold**→0xFFD700,  
**White**→0xFFFFFF, **Deep Pink**→0xFF1493



# Binary Encoding – Characters/Text

## ❖ ASCII Encoding ([www.asciitable.com](http://www.asciitable.com))

### ■ American Standard Code for Information Interchange

Dec	Hx	Oct	Char	Dec	Hx	Oct	Html	Chr	Dec	Hx	Oct	Html	Chr	Dec	Hx	Oct	Html	Chr
0	0	000	<b>NUL</b> (null)	32	20	040	&#32;	<b>Space</b>	64	40	100	&#64;	<b>@</b>	96	60	140	&#96;	<b>`</b>
1	1	001	<b>SOH</b> (start of heading)	33	21	041	&#33;	<b>!</b>	65	41	101	&#65;	<b>A</b>	97	61	141	&#97;	<b>a</b>
2	2	002	<b>STX</b> (start of text)	34	22	042	&#34;	<b>"</b>	66	42	102	&#66;	<b>B</b>	98	62	142	&#98;	<b>b</b>
3	3	003	<b>ETX</b> (end of text)	35	23	043	&#35;	<b>#</b>	67	43	103	&#67;	<b>C</b>	99	63	143	&#99;	<b>c</b>
4	4	004	<b>EOT</b> (end of transmission)	36	24	044	&#36;	<b>\$</b>	68	44	104	&#68;	<b>D</b>	100	64	144	&#100;	<b>d</b>
5	5	005	<b>ENQ</b> (enquiry)	37	25	045	&#37;	<b>%</b>	69	45	105	&#69;	<b>E</b>	101	65	145	&#101;	<b>e</b>
6	6	006	<b>ACK</b> (acknowledge)	38	26	046	&#38;	<b>&amp;</b>	70	46	106	&#70;	<b>F</b>	102	66	146	&#102;	<b>f</b>
7	7	007	<b>BEL</b> (bell)	39	27	047	&#39;	<b>'</b>	71	47	107	&#71;	<b>G</b>	103	67	147	&#103;	<b>g</b>
8	8	010	<b>BS</b> (backspace)	40	28	050	&#40;	<b>(</b>	72	48	110	&#72;	<b>H</b>	104	68	150	&#104;	<b>h</b>
9	9	011	<b>TAB</b> (horizontal tab)	41	29	051	&#41;	<b>)</b>	73	49	111	&#73;	<b>I</b>	105	69	151	&#105;	<b>i</b>
10	A	012	<b>LF</b> (NL line feed, new line)	42	2A	052	&#42;	<b>*</b>	74	4A	112	&#74;	<b>J</b>	106	6A	152	&#106;	<b>j</b>
11	B	013	<b>VT</b> (vertical tab)	43	2B	053	&#43;	<b>+</b>	75	4B	113	&#75;	<b>K</b>	107	6B	153	&#107;	<b>k</b>
12	C	014	<b>FF</b> (NP form feed, new page)	44	2C	054	&#44;	<b>,</b>	76	4C	114	&#76;	<b>L</b>	108	6C	154	&#108;	<b>l</b>
13	D	015	<b>CR</b> (carriage return)	45	2D	055	&#45;	<b>-</b>	77	4D	115	&#77;	<b>M</b>	109	6D	155	&#109;	<b>m</b>
14	E	016	<b>SO</b> (shift out)	46	2E	056	&#46;	<b>.</b>	78	4E	116	&#78;	<b>N</b>	110	6E	156	&#110;	<b>n</b>
15	F	017	<b>SI</b> (shift in)	47	2F	057	&#47;	<b>/</b>	79	4F	117	&#79;	<b>O</b>	111	6F	157	&#111;	<b>o</b>
16	10	020	<b>DLE</b> (data link escape)	48	30	060	&#48;	<b>0</b>	80	50	120	&#80;	<b>P</b>	112	70	160	&#112;	<b>p</b>
17	11	021	<b>DC1</b> (device control 1)	49	31	061	&#49;	<b>1</b>	81	51	121	&#81;	<b>Q</b>	113	71	161	&#113;	<b>q</b>
18	12	022	<b>DC2</b> (device control 2)	50	32	062	&#50;	<b>2</b>	82	52	122	&#82;	<b>R</b>	114	72	162	&#114;	<b>r</b>
19	13	023	<b>DC3</b> (device control 3)	51	33	063	&#51;	<b>3</b>	83	53	123	&#83;	<b>S</b>	115	73	163	&#115;	<b>s</b>
20	14	024	<b>DC4</b> (device control 4)	52	34	064	&#52;	<b>4</b>	84	54	124	&#84;	<b>T</b>	116	74	164	&#116;	<b>t</b>
21	15	025	<b>NAK</b> (negative acknowledge)	53	35	065	&#53;	<b>5</b>	85	55	125	&#85;	<b>U</b>	117	75	165	&#117;	<b>u</b>
22	16	026	<b>SYN</b> (synchronous idle)	54	36	066	&#54;	<b>6</b>	86	56	126	&#86;	<b>V</b>	118	76	166	&#118;	<b>v</b>
23	17	027	<b>ETB</b> (end of trans. block)	55	37	067	&#55;	<b>7</b>	87	57	127	&#87;	<b>W</b>	119	77	167	&#119;	<b>w</b>
24	18	030	<b>CAN</b> (cancel)	56	38	070	&#56;	<b>8</b>	88	58	130	&#88;	<b>X</b>	120	78	170	&#120;	<b>x</b>
25	19	031	<b>EM</b> (end of medium)	57	39	071	&#57;	<b>9</b>	89	59	131	&#89;	<b>Y</b>	121	79	171	&#121;	<b>y</b>
26	1A	032	<b>SUB</b> (substitute)	58	3A	072	&#58;	<b>:</b>	90	5A	132	&#90;	<b>Z</b>	122	7A	172	&#122;	<b>z</b>
27	1B	033	<b>ESC</b> (escape)	59	3B	073	&#59;	<b>;</b>	91	5B	133	&#91;	<b>[</b>	123	7B	173	&#123;	<b>{</b>
28	1C	034	<b>FS</b> (file separator)	60	3C	074	&#60;	<b>&lt;</b>	92	5C	134	&#92;	<b>\</b>	124	7C	174	&#124;	<b> </b>
29	1D	035	<b>GS</b> (group separator)	61	3D	075	&#61;	<b>=</b>	93	5D	135	&#93;	<b>]</b>	125	7D	175	&#125;	<b>}</b>
30	1E	036	<b>RS</b> (record separator)	62	3E	076	&#62;	<b>&gt;</b>	94	5E	136	&#94;	<b>^</b>	126	7E	176	&#126;	<b>~</b>
31	1F	037	<b>US</b> (unit separator)	63	3F	077	&#63;	<b>?</b>	95	5F	137	&#95;	<b>_</b>	127	7F	177	&#127;	<b>DEL</b>

# Binary Encoding – Files and Programs

- ❖ At the lowest level, all digital data is stored as bits!
- ❖ Layers of abstraction keep everything comprehensible
  - Data/files are groups of bits interpreted by program
  - Program is actually groups of bits being interpreted by your CPU
- ❖ Computer Memory Demo (if time)
  - From vim: `%!xxd`
  - From emacs: `M-x hexl-mode`

# Summary

- ❖ Humans think about numbers in decimal; computers think about numbers in binary
  - Base conversion to go between them
  - Hexadecimal is more human-readable than binary
- ❖ All information on a computer is binary
- ❖ Binary encoding can represent *anything!*
  - Computer/program needs to know how to interpret the bits



# 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:



# Memory, Data, and Addressing


- ❖ Hardware - High Level Overview
- ❖ Representing information as bits and bytes
  - Memory is a byte-addressable array
  - Machine “word” size = address size = register size
- ❖ Organizing and addressing data in memory
  - Endianness – ordering bytes in memory
- ❖ Manipulating data in memory using C
- ❖ Boolean algebra and bit-level manipulations

# Hardware: Physical View





**Storage connections**

# Hardware: Logical View



# Hardware: 351 View (version 0)




- ❖ The CPU **executes** instructions
- ❖ Memory **stores** data
- ❖ Binary encoding!
  - Instructions *are* just data

How are data  
and instructions  
represented?

# Aside: Why Base 2?

- ❖ Electronic implementation
  - Easy to store with bi-stable elements
  - Reliably transmitted on noisy and inaccurate wires




- ❖ Other bases possible, but not yet viable:
  - DNA data storage (base 4: A, C, G, T) is a hot topic
  - Quantum computing

# Binary Encoding 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)


# Hardware: 351 View (version 0)



- ❖ To execute an instruction, the CPU must:
  - 1) Fetch the instruction
  - 2) (if applicable) Fetch data needed by the instruction
  - 3) Perform the specified computation
  - 4) (if applicable) Write the result back to memory




# Hardware: 351 View (version 1)



- ❖ More CPU details:
  - Instructions are held temporarily in the **instruction cache**
  - Other data are held temporarily in **registers**
- ❖ **Instruction fetching** is hardware-controlled
- ❖ **Data movement** is programmer-controlled (assembly)


# Hardware: 351 View (version 1)



- ❖ We will start by learning about Memory

How does a program find its data in memory?

# An Address Refers to a Byte of Memory




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

# 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**
- ❖ We have *chosen* to tie word size to address size/width
  - 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 x 10<sup>19</sup> bytes**  
= 18 billion billion bytes = 18 EB (exabytes)
  - Actual physical address space: **48 bits**

# Word-Oriented Memory Organization

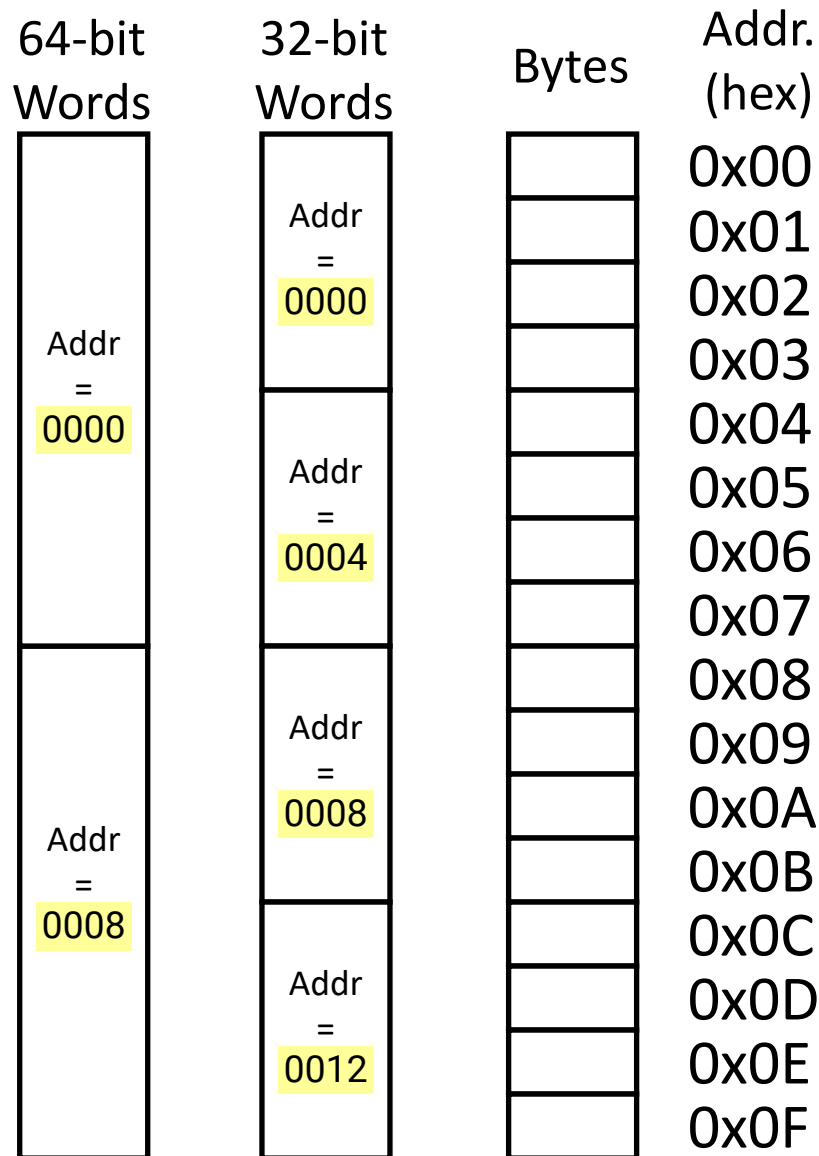
- ❖ Addresses still specify locations of *bytes* in memory
  - Addresses of successive words differ by word size (in bytes): *e.g.* 4 (32-bit) or 8 (64-bit)
  - Address of word 0, 1, ... 10?



# Address of a Word = Address of First Byte in the Word


- ❖ 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**




# A Picture of Memory (64-bit word 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



# A Picture of Memory (64-bit word 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






# Addresses and Pointers

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

big-endian

- ❖ An *address* refers to a location in memory
- ❖ A *pointer* is a data object that holds an address
  - Address can point to *any* data
- ❖ Value 504 stored at address **0x08**
  - $504_{10} = 1F8_{16}$   
= 0x 00 ... 00 01 F8
- ❖ Pointer stored at **0x38** points to address **0x08**




# Addresses and Pointers

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

big-endian

- ❖ An *address* refers to 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



# 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
<b>(reference)</b>	<b>pointer *</b>	<b>4</b>	<b>8</b>

address size = word size


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

# Memory Alignment

- ❖ **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

Bytes



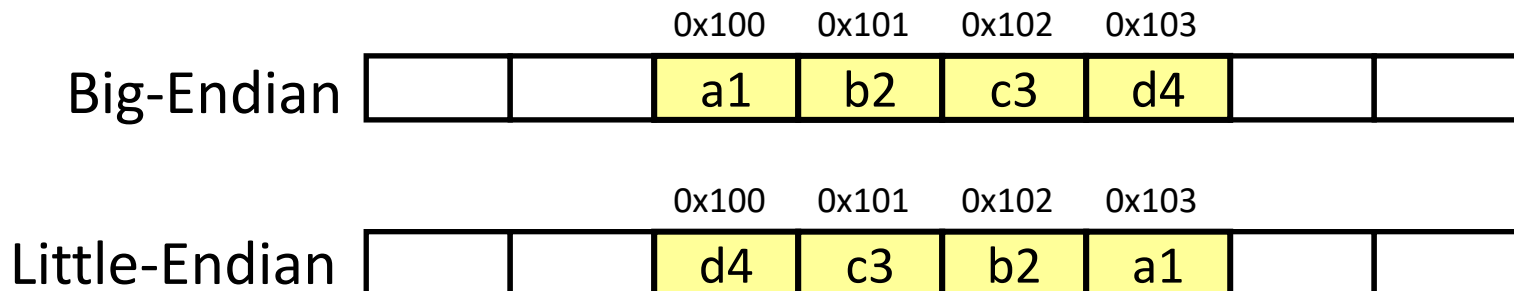
- ❖ For good memory system performance, Intel (x86) recommends data be aligned
  - However the x86-64 hardware will work correctly otherwise
    - Design choice: x86-64 instructions are *variable* bytes long

# 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  
(in decimal: 2712847316)
- ❖ 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)

# 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
- ❖ **Example:** 4-byte data 0xa1b2c3d4 at address 0x100



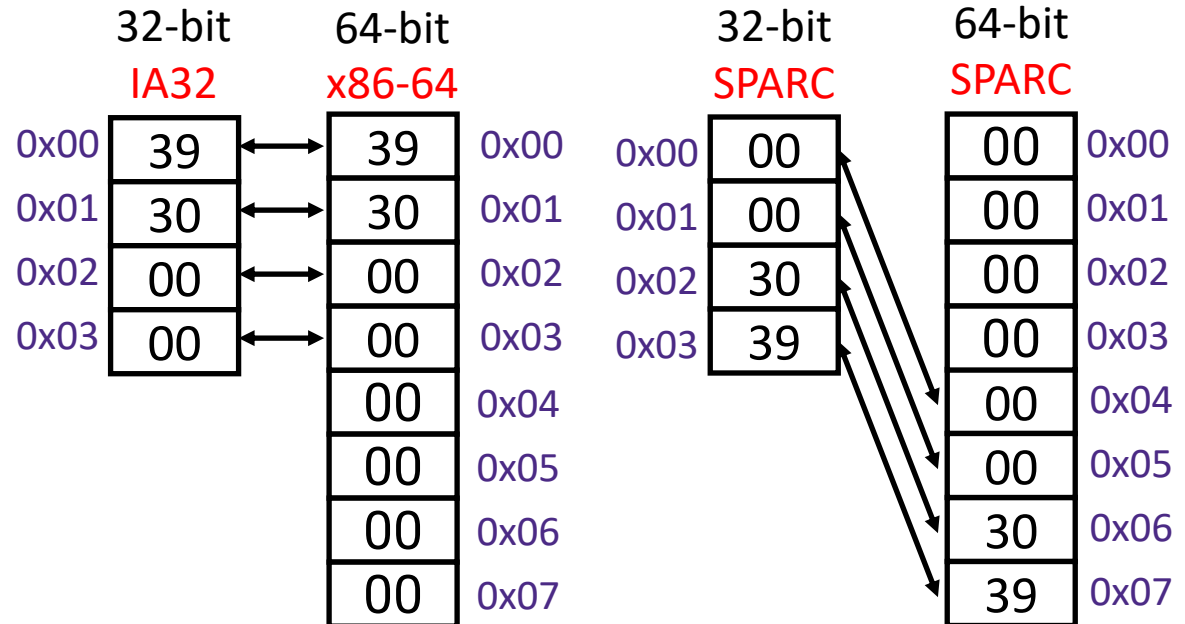
# Byte Ordering Examples

Decimal:	12345
Binary:	0011 0000 0011 1001
Hex:	3 0 3 9

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



```
long int y =
(long int) x;
```



(A long int is the size of a word)

# 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)



# 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