

ECE/CS 250

Computer Architecture

Summer 2018

From C to Binary

Tyler Bletsch
Duke University

Slides are derived from work by
Daniel J. Sorin (Duke), Andrew Hilton (Duke), Alvy Lebeck (Duke),
Benjamin Lee (Duke), Amir Roth (Penn)

Also contains material adapted from CSC230: C and Software Tools developed by
the NC State Computer Science Faculty

Outline

- Previously:
 - Computer is machine that does what we tell it to do
- Next:
 - How do we tell computers what to do?
 - How do we represent data objects in binary?
 - How do we represent data locations in binary?

Representing High Level Things in Binary

- Computers represent **everything** in binary
- Instructions are specified in binary
- Instructions must be able to describe
 - Operation types (add, subtract, shift, etc.)
 - Data objects (integers, decimals, characters, etc.)
 - Memory locations
- Example:

```
int x, y;           // Where are x and y?  How to represent an int?
bool decision;     // How do we represent a bool?  Where is it?
y = x + 7;         // How do we specify "add"?  How to represent 7?
decision=(y>18);  // Etc.
```

Representing Operation Types

- How do we tell computer to add? Shift? Read from memory? Etc.
- Arbitrarily! 😊
- Each Instruction Set Architecture (ISA) has its own binary encodings for each operation type
- E.g., in MIPS:
 - Integer add is: 00000 010000
 - Read from memory (load) is: 010011
 - Etc.

Representing Data Types

- How do we specify an integer? A character? A floating point number? A bool? Etc.
- Same as before: binary!
- Key Idea: the same 32 bits might mean one thing if interpreted as an integer but another thing if interpreted as a floating point number

Basic Data Types

Bit (bool): 0, 1

Bit String: sequence of bits of a particular length

4 bits is a **nibble**

8 bits is a **byte**

16 bits is a **half-word** (for MIPS32)

32 bits is a **word** (for MIPS32)

64 bits is a **double-word** (for MIPS32)

128 bits is a **quad-word** (for MIPS32)

What is a **word**?

The standard unit of manipulation for a particular system. E.g.:

- **MIPS32: 32 bits**
- Original Nintendo: 8 bit
- Super Nintendo: 16 bit
- Intel x86 (classic): 32 bit
- Nintendo 64: 64 bit
- Intel x86_64 (modern): 64 bit

Integers (int, long):

"2's Complement" (32-bit or 64-bit representation)

Floating Point (float, double):

Single Precision (32-bit representation)

Double Precision (64-bit representation)

Extended (Quad) Precision (128-bit representation)

Character (char):


ASCII 7-bit code

Basic Binary

- Advice: memorize the following
 - $2^0 = 1$
 - $2^1 = 2$
 - $2^2 = 4$
 - $2^3 = 8$
 - $2^4 = 16$
 - $2^5 = 32$
 - $2^6 = 64$
 - $2^7 = 128$
 - $2^8 = 256$
 - $2^9 = 512$
 - $2^{10} = 1024$



Useful bit facts

- **If you have N bits, you can represent 2^N things.** 
- **The binary metric system:**
 - $2^{10} = 1024$.
 - This is *basically* 1000, so we can have an alternative form of metric units based on base 2.
 - 2^{10} bytes = 1024 bytes = 1kB.
 - Sometimes written as 1kiB
(pronounced "kibibyte" where the 'bi' means 'binary')
(but nobody says "kibibyte" out loud because it sounds stupid)
 - 2^{20} bytes = 1MB, 2^{30} bytes = 1GB, 2^{40} bytes = 1TB, etc.
 - Easy rule to convert between exponent and binary metric number:

$$2^{XY} \text{ bytes} = 2^Y \text{ <X_prefix> B}$$

$$2^{13} \text{ bytes} = 2^3 \text{ kB} = 8 \text{ kB}$$

$$2^{39} \text{ bytes} = 2^9 \text{ GB} = 512 \text{ GB}$$

$$2^{05} \text{ bytes} = 2^5 \text{ B} = 32 \text{ B}$$

Decimal to binary using remainders


?	Quotient	Remainder	
457 ÷ 2 =	228	1	_____
228 ÷ 2 =	114	0	_____
114 ÷ 2 =	57	0	_____
57 ÷ 2 =	28	1	_____
28 ÷ 2 =	14	0	_____
14 ÷ 2 =	7	0	_____
7 ÷ 2 =	3	1	_____
3 ÷ 2 =	1	1	_____
1 ÷ 2 =	0	1	_____

111001001

Decimal to binary using comparison

Num	Compare 2^n	$\geq ?$
457	256	1
201	128	1
73	64	1
9	32	0
9	16	0
9	8	1
1	4	0
1	2	0
1	1	1

111001001



Hexadecimal

Hex digit	Binary	Decimal
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
A	1010	10
B	1011	11
C	1100	12
D	1101	13
E	1110	14
F	1111	15

Indicates a hex number

0x**DEADBEEF**
1101 1110 1010 1101 1011 1110 1110 1111

0x**02468ACE**
0000 0010 0100 0110 1000 1010 1100 1110

0x**13579BDF**
0001 0011 0101 0111 1001 1011 1101 1111

Binary to/from hexadecimal

- $0101101100100011_2 \rightarrow$
- $0101\ 1011\ 0010\ 0011_2 \rightarrow$
- $5\ B\ 2\ 3_{16}$

$1\ F\ 4\ B_{16} \rightarrow$

$0001\ 1111\ 0100\ 1011_2 \rightarrow$

0001111101001011_2

Hex digit	Binary	Decimal
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
A	1010	10
B	1011	11
C	1100	12
D	1101	13
E	1110	14
F	1111	15

BitOps: Unary

- Bit-wise complement (\sim)
 - Flips every bit.

```
~0x0d    // (binary 00001101)
== 0xf2  // (binary 11110010)
```

Not the same as Logical NOT (!) or sign change (-)

```
char i, j1, j2, j3;
i = 0x0d;    // binary 00001101
j1 = ~i;     // binary 11110010
j2 = -i;     // binary 11110011
j3 = !i;     // binary 00000000
```

BitOps: Two Operands

- Operate **bit-by-bit** on operands to produce a result operand of the same length
- And (**&**): result 1 if both inputs 1, 0 otherwise
- Or (**|**): result 1 if either input 1, 0 otherwise
- Xor (**^**): result 1 if one input 1, but not both, 0 otherwise
- Operands **must** be of type integer

Two Operands... (cont'd)

- Examples


```
0011 1000
& 1101 1110
-----
0001 1000
```

```
0011 1000
| 1101 1110
-----
1111 1110
```

```
0011 1000
^ 1101 1110
-----
1110 0110
```


Shift Operations

- $x \ll y$ is left (**logical**) shift of x by y positions
 - x and y must both be integers
 - x **should** be unsigned or positive
 - y leftmost bits of x are discarded
 - zero fill y bits on the right



ShiftOps... (cont'd)

- $x \gg y$ is right (**logical**) shift of x by y positions
 - y rightmost bits of x are discarded
 - zero fill y bits on the left



Bitwise Recipes

- Set a certain bit to 1?

- Make a MASK with a *one* at every position you want to *set*:

```
m = 0x02; // 000000102
```

- OR the mask with the input:

```
v = 0x41; // 010000012
```

```
v |= m; // 010000112
```

- Clear a certain bit to 0?

- Make a MASK with a *zero* at every position you want to *clear*:

```
m = 0xFD; // 111111012 (could also write ~0x02)
```

- AND the mask with the input:

```
v = 0x27; // 001001112
```

```
v &= m; // 001001012
```

- Get a substring of bits (such as bits 2 through 5)?

Note: bits are numbered right-to-left starting with zero.

- Shift the bits you want all the way to the right then AND them with an appropriate mask:

```
v = 0x67; // 011001112
```

```
v >>= 2; // 000110012
```

```
v &= 0x0F; // 000010012
```

Binary Math : Addition

- Suppose we want to add two numbers:

$$\begin{array}{r} 00011101 \\ + 00101011 \\ \hline \end{array}$$

- How do we do this?

Binary Math : Addition

- Suppose we want to add two numbers:

$$\begin{array}{r} 00011101 \\ + 00101011 \\ \hline \end{array}$$

$$\begin{array}{r} 695 \\ + 232 \\ \hline \end{array}$$

- How do we do this?
 - Let's revisit decimal addition
 - Think about the process as we do it

Binary Math : Addition

- Suppose we want to add two numbers:

$$\begin{array}{r} 00011101 \\ + 00101011 \\ \hline \end{array}$$

$$\begin{array}{r} 695 \\ + 232 \\ \hline 7 \end{array}$$

- First add one's digit $5+2 = 7$

Binary Math : Addition

- Suppose we want to add two numbers:

$$\begin{array}{r} 00011101 \\ + 00101011 \\ \hline \end{array}$$

$$\begin{array}{r} 1 \\ 695 \\ + 232 \\ \hline 27 \end{array}$$

- First add one's digit $5+2 = 7$
- Next add ten's digit $9+3 = 12$ (2 carry a 1)

Binary Math : Addition

- Suppose we want to add two numbers:

$$\begin{array}{r} 00011101 \\ + 00101011 \\ \hline \end{array}$$

$$\begin{array}{r} 695 \\ + 232 \\ \hline 927 \end{array}$$

- First add one's digit $5+2 = 7$
- Next add ten's digit $9+3 = 12$ (2 carry a 1)
- Last add hundred's digit $1+6+2 = 9$

Binary Math : Addition

- Suppose we want to add two numbers:

$$\begin{array}{r} 00011101 \\ + 00101011 \\ \hline \end{array}$$

- Back to the binary:
- First add 1's digit $1+1 = \dots?$

Binary Math : Addition

- Suppose we want to add two numbers:

$$\begin{array}{r} 1 \\ 00011101 \\ + 00101011 \\ \hline 0 \end{array}$$

- Back to the binary:
- First add 1's digit $1+1 = 2$ (0 carry a 1)

Binary Math : Addition

- Suppose we want to add two numbers:

$$\begin{array}{r} 11 \\ 00011101 \\ + 00101011 \\ \hline 00 \end{array}$$

- Back to the binary:
 - First add 1's digit $1+1 = 2$ (0 carry a 1)
 - Then 2's digit: $1+0+1 = 2$ (0 carry a 1)
 - You all finish it out....

Binary Math : Addition

- Suppose we want to add two numbers:

111111

00011101 = 29

+ 00101011 = 43

01001000 = 72

- Can check our work in decimal

Issues for Binary Representation of Numbers

- **How to represent negative numbers?**
- There are many ways to represent numbers in binary
 - Binary representations are encodings → many encodings possible
 - What are the issues that we must address?
- Issue #1: Complexity of arithmetic operations
- Issue #2: Negative numbers
- Issue #3: Maximum representable number
- Choose representation that makes these issues easy for machine, even if it's not easy for humans (i.e., ECE/CS 250 students)
 - Why? Machine has to do all the work!

Sign Magnitude

- Use leftmost bit for + (0) or – (1):
- 6-bit example (1 sign bit + 5 magnitude bits):
- +17 = 010001
- -17 = 110001
- Pros:
 - Conceptually simple
 - Easy to convert
- Cons:
 - Harder to compute (add, subtract, etc) with
 - Positive and negative 0: 000000 and 100000

NOBODY DOES THIS

1's Complement Representation for Integers

- Use largest positive binary numbers to represent negative numbers
- To negate a number, invert ("not") each bit:
 - $0 \rightarrow 1$
 - $1 \rightarrow 0$
- Cons:
 - Still two 0s (yuck)
 - Still hard to compute with

0000	0
0001	1
0010	2
0011	3
0100	4
0101	5
0110	6
0111	7
1000	-7
1001	-6
1010	-5
1011	-4
1100	-3
1101	-2
1110	-1
1111	-0

NOBODY DOES THIS EITHER

2's Complement Integers

- Use large positives to represent negatives
- $(-x) = 2^n - x$
- This is 1's complement + 1
- $(-x) = 2^n - 1 - x + 1$
- So, just invert bits and add 1

0000	0
0001	1
0010	2
0011	3
0100	4
0101	5
0110	6
0111	7
1000	-8
1001	-7
1010	-6
1011	-5
1100	-4
1101	-3
1110	-2
1111	-1

6-bit examples:

$$010110_2 = 22_{10}; 101010_2 = -22_{10}$$

$$1_{10} = 000001_2; -1_{10} = 111111_2$$

$$0_{10} = 000000_2; -0_{10} = 000000_2 \rightarrow \text{good!}$$

EVERYBODY DOES THIS

Pros and Cons of 2's Complement

- Advantages:
 - Only one representation for 0 (unlike 1's comp): $0 = 000000$
 - Addition algorithm is much easier than with sign and magnitude
 - Independent of sign bits
- Disadvantage:
 - One more negative number than positive
 - Example: 6-bit 2's complement number
 $100000_2 = -32_{10}$; but 32_{10} could not be represented

All modern computers use 2's complement for integers

2's Complement Precision Extension

- Most computers today support 32-bit (int) or 64-bit integers
 - Specify 64-bit using gcc C compiler with `long long`
- To extend precision, use `sign bit extension`
 - Integer precision is number of bits used to represent a number

Examples

$14_{10} = 001110_2$ in 6-bit representation.

$14_{10} = 000000001110_2$ in 12-bit representation

$-14_{10} = 110010_2$ in 6-bit representation

$-14_{10} = 111111110010_2$ in 12-bit representation.

Binary Math : Addition

- Let's look at another binary addition:

$$\begin{array}{r} 01011101 \\ + 01101011 \\ \hline \end{array}$$

Binary Math : Addition

- What about this one:

1111111

01011101 = 93

+ 01101011 = 107

11001000 = -56

- But... that can't be right?
 - What do you expect for the answer?
 - What is it in 8-bit signed 2's complement?

Integer Overflow

- Answer should be 200
 - Not representable in 8-bit signed representation
 - No right answer
- This is called integer **Overflow**
- Real problem in programs

Subtraction

- 2's complement makes subtraction easy:
 - Remember: $A - B = A + (-B)$
 - And: $-B = \sim B + 1$
 - ↑ that means flip bits ("not")
 - So we just flip the bits and start with carry-in (CI) = 1
 - Later: No new circuits to subtract (re-use adder hardware!)

$$\begin{array}{r} 0110101 \\ - 1010010 \\ \hline \end{array} \quad \rightarrow \quad \begin{array}{r} 1 \\ 0110101 \\ + 0101101 \\ \hline \end{array}$$

What About Non-integer Numbers?

- There are infinitely many real numbers between two integers
- Many important numbers are real
 - Speed of light $\approx 3 \times 10^8$
 - Pi = 3.1415...
- Fixed number of bits limits range of integers
 - Can't represent some important numbers
- Humans use Scientific Notation
 - 1.3×10^4

Option 1: Fixed point

- Use normal integers, but $(X \cdot 2^k)$ instead of X
 - Example: 32 bit int, but use $X \cdot 65536$
 - $3.1415926 \cdot 65536 = 205887$
 - $0.5 \cdot 65536 = 32768$, etc..
- Pros:
 - Addition/subtraction just like integers (“free”)
- Cons:
 - Mul/div require renormalizing (divide by 64K)
 - Range limited (no good rep for large + small)
- Can be good in specific situations

Can we do better?

- Think about scientific notation for a second:
- For example:
 $6.02 * 10^{23}$
- Real number, but comprised of ints:
 - 6 generally only 1 digit here
 - 02 any number here
 - 10 always 10 (base we work in)
 - 23 can be positive or negative
- Can we do something like this in binary?

Option 2: Floating Point

- How about:
 $\pm X.YYYYYYY * 2^{\pm N}$
- Big numbers: large positive N
- Small numbers (<1): negative N
- Numbers near 0: small N
- This is “floating point” : most common way

IEEE single precision floating point


- Specific format called IEEE single precision:
+/- 1.YYYYYY * $2^{(N-127)}$
- “float” in Java, C, C++,....
- Assume first bit is always 1 (saves us a bit)
- 1 sign bit (+ = 0, 1 = -)
- 8 bit biased exponent (do N-127)
- Implicit 1 before *binary point*
- 23-bit *mantissa* (YYYYYY)

Binary fractions

- 1.YYYY has a binary point
 - Like a decimal point but in binary
 - After a decimal point, you have
 - tenths
 - hundredths
 - thousandths
 - ...
- So after a binary point you have...
 - Halves
 - Quarters
 - Eighths
 - ...

Floating point example

- Binary fraction example:
 $101.101 = 4 + 1 + \frac{1}{2} + \frac{1}{8} = 5.625$
- For floating point, needs normalization:
 $1.01101 * 2^2$
- Sign is +, which = 0
- Exponent = $127 + 2 = 129 = 1000\ 0001$
- Mantissa = $1.011\ 0100\ 0000\ 0000\ 0000\ 0000$



Floating Point Representation

Example:


What floating-point number is:

0xC1580000?

Answer

What floating-point number is
0xC1580000?

1100 0001 0101 1000 0000 0000 0000 0000



Sign = 1 which is negative

Exponent = $(128+2)-127 = 3$

Mantissa = 1.1011

$-1.1011 \times 2^3 = -1101.1 = -13.5$

Trick question

- How do you represent 0.0?
 - Why is this a trick question?
 - $0.0 = 000000000$
 - But need 1.XXXXX representation?
- Exponent of 0 is denormalized
 - Implicit 0. instead of 1. in mantissa
 - Allows 0000....0000 to be 0
 - Helps with very small numbers near 0
- Results in +/- 0 in FP (but they are “equal”)

Other Weird FP numbers


- Exponent = 1111 1111 also not standard
 - All 0 mantissa: +/- ∞
 - $1/0 = +\infty$
 - $-1/0 = -\infty$
 - Non zero mantissa: Not a Number (NaN)
 - $\text{sqrt}(-42) = \text{NaN}$

Floating Point Representation

- Double Precision Floating point:

64-bit representation:

- 1-bit **sign**
 - 11-bit (biased) **exponent**
 - 52-bit **fraction** (with implicit 1).
- “double” in Java, C, C++, ...



What About Strings?

- Many important things stored as strings...
 - E.g., your name
- How should we store strings?

Standardized ASCII (0-127)

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

One Interpretation of 128-255

128	Ç	144	É	161	í	177	☐	193	⊥	209	≠	225	β	241	±
129	ü	145	æ	162	ó	178	☐	194	⊥	210	π	226	Γ	242	≥
130	é	146	Æ	163	ú	179		195	⊥	211	⊥	227	π	243	≤
131	â	147	ô	164	ñ	180	⊥	196	—	212	⊥	228	Σ	244	∫
132	ä	148	ö	165	Ñ	181	⊥	197	⊥	213	∫	229	σ	245	∫
133	à	149	ò	166	ª	182		198	⊥	214	∫	230	μ	246	+
134	â	150	û	167	º	183	π	199		215	⊥	231	τ	247	≈
135	ç	151	ù	168	¿	184	∫	200	⊥	216	⊥	232	Φ	248	°
136	ê	152	—	169	—	185		201	∫	217	∫	233	⊙	249	·
137	ë	153	Ö	170	¬	186		202	⊥	218	∫	234	Ω	250	·
138	è	154	Û	171	½	187	∫	203	≠	219	■	235	δ	251	√
139	ï	156	£	172	¼	188	∫	204	⊥	220	■	236	∞	252	—
140	î	157	¥	173	¡	189	∫	205	=	221	■	237	φ	253	≈
141	ì	158	—	174	«	190	∫	206	⊥	222	■	238	ε	254	■
142	Ä	159	f	175	»	191	∫	207	⊥	223	■	239	∩	255	
143	Å	160	á	176	☐	192	L	208	⊥	224	α	240	≡		

Source: www.LookupTables.com

(This allowed totally sweet ASCII art in the 90s)



Sources:

- <http://roy-sac.deviantart.com/art/Cardinal-NFO-File-ASCII-35664604>
- <http://roy-sac.deviantart.com/art/Siege-ISO-nfo-ASCII-Logo-35940815>
- <http://roy-sac.deviantart.com/art/deviantART-ANSI-Logo-31556803>

CARDINALS
-^~ THE-TRAINER-MACHINE ^~-

```
RELEASE NFO
-----
TRAINED GAME      :
COMPANY           :
PIRACY GROUP      :

CODER             :
TRAINED ITEMS     :
STAMP             :
PACKAGER          :

GAME RATINGS:
-----
HARDWARE SUPPORT
-----
GFX: sVga [ ]   SOUND: GRAViS [ ]
    Vga [ ]     SB 16b [ ]
    Ega [ ]     SB PRO [ ]
    Cga [ ]     SB mono [ ]
                ROLAND [ ]
                PRO AUDIO [ ]
                ADLiB [ ]
                HONKER [ ]

    10 | 10 | 10 |
     9 | 9  | 9  |
     8 | 8  | 8  |
     7 | 7  | 7  |
     6 | 6  | 6  |
     5 | 5  | 5  |
     4 | 4  | 4  |
     3 | 3  | 3  |
     2 | 2  | 2  |
     1 | 1  | 1  |
      +---+ +---+ +---+
      GFX  SFX  FUN
```

ADDITIONAL NOTES:

GROUP GREETINGS:

PERSONAL GREETINGS:

MEMBERS

Outline

- Previously:
 - Computer is machine that does what we tell it to do
- Next:
 - How do we tell computers what to do?
 - How do we represent data objects in binary?
 - How do we represent data locations in binary?

Computer Memory

- Where do we put these numbers?
 - Registers [more on these later]
 - In the processor core
 - Compute directly on them
 - Few of them (~16 or 32 registers, each 32-bit or 64-bit)
 - Memory [Our focus now]
 - External to processor core
 - Load/store values to/from registers
 - Very large (multiple GB)

Memory Organization

- Memory: billions of locations...how to get the right one?
 - Each memory location has an [address](#)
 - Processor asks to read or write specific address
 - Memory, please load address 0x123400
 - Memory, please write 0xFE into address 0x8765000
 - Kind of like a giant array
 - Array of what?
 - Bytes?
 - 32-bit ints?
 - 64-bit ints?


Memory Organization

- Most systems: byte (8-bit) addressed
 - Memory is “array of bytes”
 - Each address specifies 1 byte
 - Support to load/store 8, 16, 32, 64 bit quantities
 - Byte ordering varies from system to system
- Some systems “word addressed”
 - Memory is “array of words”
 - Smaller operations “faked” in processor
 - Not very common

Word of the Day: Endianness


Byte Order

- **Big Endian:** byte 0 is 8 **most** significant bits IBM 360/370, Motorola 68k, MIPS, Sparc, HP PA
- **Little Endian:** byte 0 is 8 **least** significant bits Intel 80x86, DEC Vax, DEC Alpha



Memory Layout

- Memory is array of bytes, but there are conventions as to what goes where in this array
- Text: instructions (the program to execute)
- Data: global variables
- Stack: local variables and other per-function state; starts at top & grows down
- Heap: dynamically allocated variables; grows up
- What if stack and heap overlap????




Memory Layout: Example

```
int anumber = 3;

int factorial (int x) {
    if (x == 0) {
        return 1;
    }
    else {
        return x * factorial (x - 1);
    }
}

int main (void) {
    int z = factorial (anumber);
    printf("%d\n", z);
    return 0;
}
```



Summary: From C to Binary

- Everything must be represented in binary!
- Pointer is memory location that contains address of another memory location
- Computer memory is linear array of bytes
 - **Integers:**
 - **unsigned** $\{0..2^n-1\}$ vs **signed** $\{-2^{n-1} .. 2^{n-1}-1\}$ ("2's complement")
 - **char** (8-bit), **short** (16-bit), **int/long** (32-bit), **long long** (64-bit)
 - **Floats:** IEEE representation,
 - **float** (32-bit: 1 sign, 8 exponent, 23 mantissa)
 - **double** (64-bit: 1 sign, 11 exponent, 52 mantissa)
 - **Strings:** char array, ASCII representation
- Memory layout
 - **Stack** for local, **static** for globals, **heap** for malloc'd stuff (must free!)

POINTERS, ARRAYS, AND MEMORY ~AGAIN~

The following slides re-state a lot of what we've covered but in a different way. We'll likely skip it for time, but you can use the slides as an additional reference.

Let's do a little Java...

```
public class Example {  
    public static void swap (int x, int y) {  
        int temp = x;  
        x = y;  
        y = temp;  
    }  
    public static void main (String[] args) {  
        int a = 42;  
        int b = 100;  
        swap (a, b);  
        System.out.println("a =" + a + " b = " + b);  
    }  
}
```

- What does this print? Why?

Let's do a little Java...

```
public class Example {  
    public static void swap (int x, int y) {  
        int temp = x;  
        x = y;  
        y = temp;  
    }  
    public static void main (String[] args) {  
        int a = 42;  
        int b = 100;  
        → swap (a, b);  
        System.out.println("a =" + a + " b = " + b);  
    }  
}
```

Stack

main	
a	42
b	100

- What does this print? Why?

Let's do a little Java...

```
public class Example {  
    public static void swap (int x, int y) {  
→ int temp = x;  
    x = y;  
    y = temp;  
    }  
    public static void main (String[] args) {  
    int a = 42;  
    int b = 100;  
c0 → swap (a, b);  
    System.out.println("a =" + a + " b = " + b);  
    }  
}
```

Stack

main	
a	42
b	100

swap	
x	42
y	100
temp	???
RA	c0

- What does this print? Why?

Let's do a little Java...

```
public class Example {  
    public static void swap (int x, int y) {  
→ int temp = x;  
    x = y;  
    y = temp;  
    }  
    public static void main (String[] args) {  
        int a = 42;  
        int b = 100;  
c0 → swap (a, b);  
        System.out.println("a =" + a + " b = " + b);  
    }  
}
```



Stack

main	
a	42
b	100

swap	
x	42
y	100
temp	42
RA	c0

- What does this print? Why?

Let's do a little Java...

```
public class Example {  
    public static void swap (int x, int y) {  
        int temp = x;  
         x = y;  
        y = temp;  
    }  
    public static void main (String[] args) {  
        int a = 42;  
        int b = 100;  
         swap (a, b);  
        System.out.println("a =" + a + " b = " + b);  
    }  
}
```

Stack

main	
a	42
b	100

swap	
x	100
y	100
temp	42
RA	c0

- What does this print? Why?

Let's do a little Java...

```
public class Example {  
    public static void swap (int x, int y) {  
        int temp = x;  
        x = y;  
        → y = temp;  
    }  
    public static void main (String[] args) {  
        int a = 42;  
        int b = 100;  
        ← swap (a, b);  
        System.out.println("a =" + a + " b = " + b);  
    }  
}
```

Stack

main	
a	42
b	100

swap	
x	100
y	42
temp	42
RA	c0

- What does this print? Why?

Let's do a little Java...

```
public class Example {  
    public static void swap (int x, int y) {  
        int temp = x;  
        x = y;  
        y = temp;  
    }  
    public static void main (String[] args) {  
        int a = 42;  
        int b = 100;  
        swap (a, b);  
        System.out.println("a =" + a + " b = " + b);  
    }  
}
```

Stack

main	
a	42
b	100

- What does this print? Why?

Let's do some different Java...

```
public class Ex2 {
    int data;
    public Ex2 (int d) { data = d; }
    public static void swap (Ex2 x, Ex2 y) {
        int temp = x.data;
        x.data = y.data;
        y.data = temp;
    }
    public static void main (String[] args) {
        Example a = new Example (42);
        Example b = new Example (100);
        swap (a, b);
        System.out.println("a =" + a.data +
                           " b = " + b.data);
    }
}
```


Stack

main	
a	??
b	??

- What does this print? Why?

Let's do some different Java...


```
public class Ex2 {
    int data;
    public Ex2 (int d) { data = d; }
    public static void swap (Ex2 x, Ex2 y) {
        int temp = x.data;
        x.data = y.data;
        y.data = temp;
    }
    public static void main (String[] args) {
        Example a = new Example (42);
        → Example b = new Example (100);
        swap (a, b);
        System.out.println("a =" + a.data +
                           " b = " + b.data);
    }
}
```



- What does this print? Why?

Let's do some different Java...

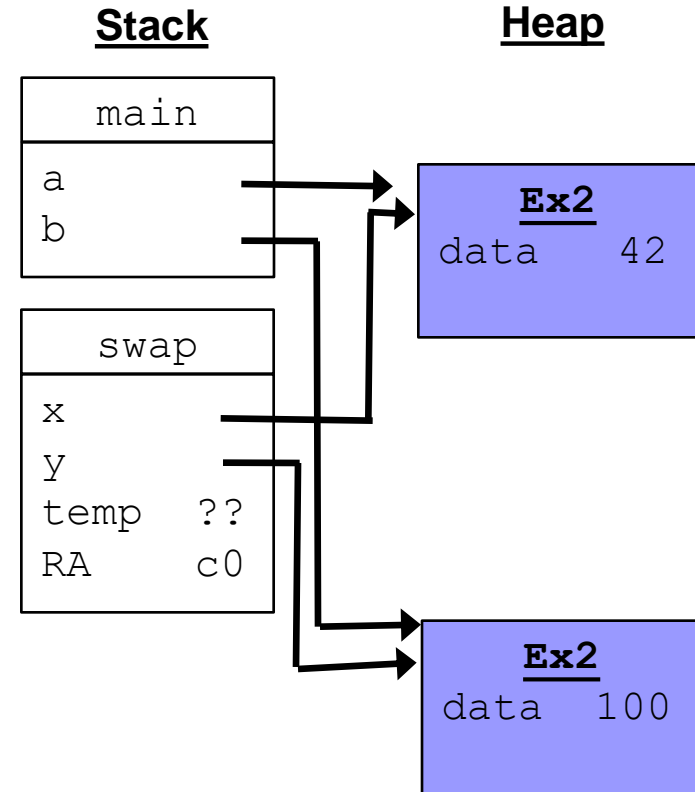
```
public class Ex2 {
    int data;
    public Ex2 (int d) { data = d; }
    public static void swap (Ex2 x, Ex2 y) {
        int temp = x.data;
        x.data = y.data;
        y.data = temp;
    }
    public static void main (String[] args) {
        Example a = new Example (42);
        Example b = new Example (100);
        → swap (a, b);
        System.out.println("a =" + a.data +
            " b = " + b.data);
    }
}
```



- What does this print? Why?

Let's do some different Java...

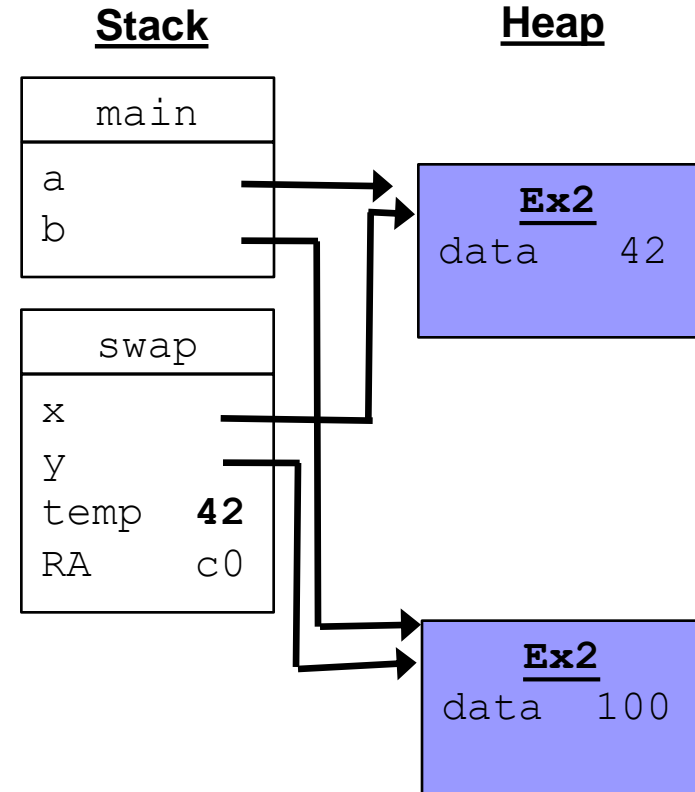
```
public class Ex2 {
    int data;
    public Ex2 (int d) { data = d; }
    public static void swap (Ex2 x, Ex2 y) {
        int temp = x.data;
        x.data = y.data;
        y.data = temp;
    }
    public static void main (String[] args) {
        Example a = new Example (42);
        Example b = new Example (100);
        swap (a, b);
        System.out.println("a =" + a.data +
            " b = " + b.data);
    }
}
```



- What does this print? Why?

Let's do some different Java...


```
public class Ex2 {
    int data;
    public Ex2 (int d) { data = d; }
    public static void swap (Ex2 x, Ex2 y) {
        int temp = x.data;
        x.data = y.data;
        y.data = temp;
    }
    public static void main (String[] args) {
        Example a = new Example (42);
        Example b = new Example (100);
        swap (a, b);
        System.out.println("a =" + a.data +
            " b = " + b.data);
    }
}
```



- What does this print? Why?

Let's do some different Java...

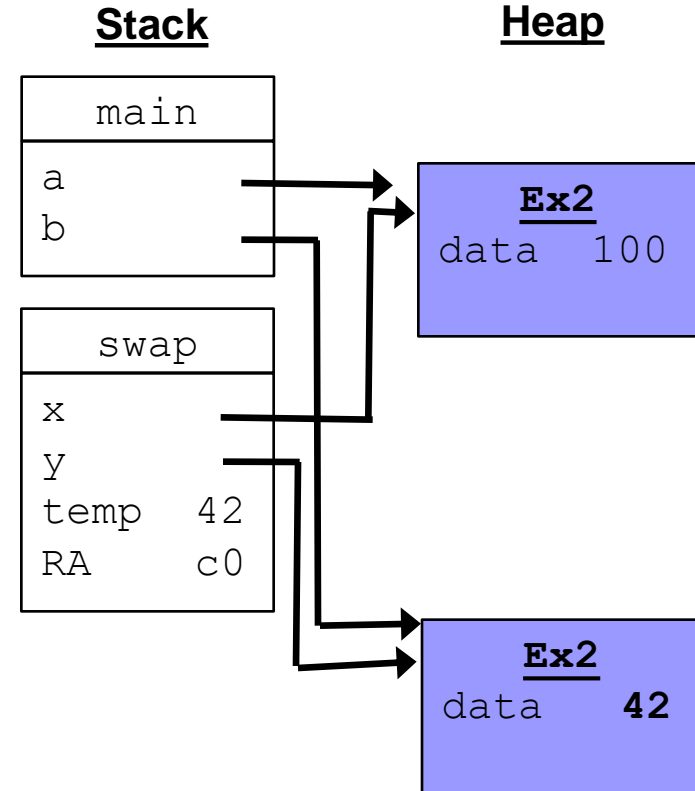
```
public class Ex2 {
    int data;
    public Ex2 (int d) { data = d; }
    public static void swap (Ex2 x, Ex2 y) {
        int temp = x.data;
        → x.data = y.data;
        y.data = temp;
    }
    public static void main (String[] args) {
        Example a = new Example (42);
        Example b = new Example (100);
        c0 → swap (a, b);
        System.out.println("a =" + a.data +
            " b = " + b.data);
    }
}
```



- What does this print? Why?

Let's do some different Java...


```
public class Ex2 {
    int data;
    public Ex2 (int d) { data = d; }
    public static void swap (Ex2 x, Ex2 y) {
        int temp = x.data;
        x.data = y.data;
        y.data = temp;
    }
    public static void main (String[] args) {
        Example a = new Example (42);
        Example b = new Example (100);
        swap (a, b);
        System.out.println("a =" + a.data +
            " b = " + b.data);
    }
}
```



- What does this print? Why?

Let's do some different Java...

```
public class Ex2 {  
    int data;  
    public Ex2 (int d) { data = d; }  
    public static void swap (Ex2 x, Ex2 y) {  
        int temp = x.data;  
        x.data = y.data;  
        y.data = temp;  
    }  
    public static void main (String[] args) {  
        Example a = new Example (42);  
        Example b = new Example (100);  
        swap (a, b);  
        System.out.println("a =" + a.data +  
            " b = " + b.data);  
    }  
}
```



- What does this print? Why?

References and Pointers (review)

- Java has **references**:
 - Any variable of object type is a reference
 - Point at objects (which are all in the heap)
 - Under the hood: is the memory address of the object
 - Cannot explicitly manipulate them (*e.g.*, add 4)
- Some languages (C,C++,assembly) have explicit **pointers**:
 - Hold the memory address of something
 - Can explicitly compute on them
 - Can **de-reference** the pointer (*ptr) to get thing-pointed-to
 - Can take the **address-of** (&x) to get something's address
 - Can do very **unsafe** things, shoot yourself in the foot

Pointers

- “address of” operator &
 - don't confuse with bitwise AND operator (&&)

Given

```
int x; int* p; // p points to an int
```

```
p = &x;
```

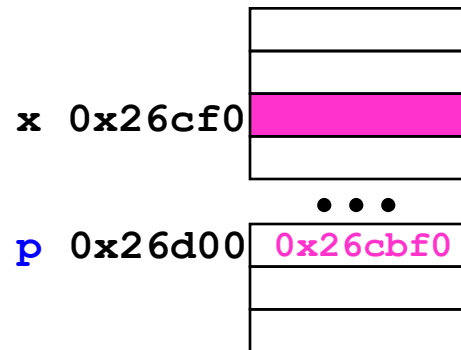
Then

```
*p = 2; and x = 2; produce the same result
```

Note: p is a pointer, *p is an int

- What happens for `p = 2;`;

On 32-bit machine, p is 32-bits



Back to Arrays

- Java:

```
int [] x = new int [nElems];
```


- C:

```
int data[42]; //if size is known constant
```


```
int* data = (int*)malloc (nElem * sizeof(int));
```

- `malloc` takes number of bytes
 - `sizeof` tells how many bytes something takes
- 

Arrays, Pointers, and Address Calculation


- x is a pointer, what is $x+33$?
- A pointer, but where?
 - what does calculation depend on?
- Result of adding an int to a pointer depends on size of object pointed to 
 - One reason why we tell compiler what type of pointer we have, even though all pointers are really the same thing (and same size)

```
int* a=malloc(100*sizeof(int));
```



$a[33]$ is the same as $*(a+33)$
if a is $0x00a0$, then $a+1$ is
 $0x00a4$, $a+2$ is $0x00a8$
(decimal 160, 164, 168)


```
double* d=malloc(200*sizeof(double));
```



$*(d+33)$ is the same as $d[33]$
if d is $0x00b0$, then $d+1$ is
 $0x00b8$, $d+2$ is $0x00c0$
(decimal 176, 184, 192)

More Pointer Arithmetic

- address one past the end of an array is ok for pointer comparison only
- what's at `*(begin+44)`?
- what does `begin++` mean?
- how are pointers compared using `<` and using `==` ?
- what is value of `end - begin`?




```
char* a = new char[44];  
char* begin = a;  
char* end = a + 44;
```

```
while (begin < end)  
{  
    *begin = 'z';  
    begin++;  
}
```

More Pointers & Arrays

```
int* a = new int[100];
```



a is a pointer

*a is an int

a[0] is an int (same as *a)

a[1] is an int

a+1 is a pointer

a+32 is a pointer

*(a+1) is an int (same as a[1])

*(a+99) is an int

***(a+100) is trouble**

Array Example


```
#include <stdio.h>

main()
{
    int* a = (int*)malloc (100 * sizeof(int));
    int* p = a;
    int k;


    for (k = 0; k < 100; k++)
    {
        *p = k;
        p++;
    }
    printf("entry 3 = %d\n", a[3])
}
```

Memory Manager (Heap Manager)

- malloc() and free()
- Library routines that handle memory **management for heap** (allocation / deallocation)
- Java has garbage collection (reclaim memory of unreferenced objects)
- C must use **free**, else memory leak



Strings as Arrays (review)



- A string is an array of characters with '\0' at the end
- Each element is one byte, ASCII code
- '\0' is null (ASCII code 0)

strlen() again

- `strlen()` returns the number of characters in a string
 - same as number elements in char array?

```
int strlen(char * s)
// pre: '\0' terminated
// post: returns # chars
{
    int count=0;
    while (*s++)
        count++;
    return count;
}
```

Vector Class vs. Arrays

- Vector Class
 - insulates programmers
 - array bounds checking
 - automagically growing/shrinking when more items are added/deleted
- How are Vectors implemented?
 - Arrays, re-allocated as needed
- Arrays can be more efficient