

Computer Systems

CSE 410 Autumn 2013

3 - Integers

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**
- Machine code & C
- x86 assembly
- Procedures & stacks
- 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:



Encoding

Integer & Floating Point Numbers

- Representation of integers: unsigned and signed
- Unsigned and signed integers in C
- Arithmetic and shifting
- Sign extension

- Background: fractional binary numbers
- IEEE floating-point standard
- Floating-point operations and rounding
- Floating-point in C

- Reading: Bryant/O'Hallaron sec. 2.2-2.3

Unsigned Integers

- **Unsigned values are just what you expect**

- $b_7b_6b_5b_4b_3b_2b_1b_0 = b_72^7 + b_62^6 + b_52^5 + \dots + b_12^1 + b_02^0$
- Interesting aside: $1+2+4+8+\dots+2^{N-1} = 2^N - 1$

$$\begin{array}{r} 00111111 \\ +00000001 \\ \hline 01000000 \end{array}$$

$$\begin{array}{r} 63 \\ + \underline{1} \\ \hline 64 \end{array}$$

- **You add/subtract them using the normal “carry/borrow” rules, just in binary**

- **An important use of unsigned integers in C is pointers**

- There are no negative memory addresses


Signed Integers

- **Let's do the natural thing for the positives**
 - They correspond to the unsigned integers of the same value
 - Example (8 bits): $0x00 = 0$, $0x01 = 1$, ..., $0x7F = 127$
- **But, we need to let about half of them be negative**
 - Use the high order bit to indicate *negative*: call it the “sign bit”
 - Call this a “sign-and-magnitude” representation
 - Examples (8 bits):
 - $0x00 = 00000000_2$ is non-negative, because the sign bit is 0
 - $0x7F = 01111111_2$ is non-negative
 - $0x85 = 10000101_2$ is negative
 - $0x80 = 10000000_2$ is negative...

Sign-and-Magnitude Negatives

■ How should we represent -1 in binary?

- Sign-and-magnitude: 10000001_2
Use the MSB for + or -, and the other bits to give magnitude




Sign-and-Magnitude Negatives

■ How should we represent -1 in binary?

- Sign-and-magnitude: 10000001_2

Use the MSB for + or -, and the other bits to give magnitude

(Unfortunate side effect: there are two representations of 0!)




Sign-and-Magnitude Negatives

■ How should we represent -1 in binary?


- Sign-and-magnitude: 10000001_2
Use the MSB for + or -, and the other bits to give magnitude
(Unfortunate side effect: there are two representations of 0!)
- Another problem: math is cumbersome

- Example:

$$4 - 3 \neq 4 + (-3)$$



0100
+1011
1111




Two's Complement Negatives

■ How should we represent -1 in binary?

- Rather than a sign bit, let MSB have same value, but *negative* weight
 - W-bit word: Bits 0, 1, ..., W-2 add $2^0, 2^1, \dots, 2^{W-2}$ to value of integer when set, but bit W-1 adds -2^{W-1} when set
 - e.g. unsigned 1010_2 : $1*2^3 + 0*2^2 + 1*2^1 + 0*2^0 = 10_{10}$
 2's comp. 1010_2 : $-1*2^3 + 0*2^2 + 1*2^1 + 0*2^0 = -6_{10}$

- So -1 represented as 1111_2 ; all negative integers still have MSB = 1
- Advantages of two's complement: only one zero, simple arithmetic
- To get negative representation of any integer, take bitwise complement and then add one!

$$\sim x + 1 = -x$$



Two's Complement Arithmetic

- The same addition procedure works for both unsigned and two's complement integers
 - Simplifies hardware: only one adder needed
 - Algorithm: simple addition, discard the highest carry bit
 - Called “modular” addition: result is sum *modulo* 2^W
- Examples:

4 0100	4 0100	- 4 1100
+ 3 + 0011	- 3 + 1101	+ 3 + 0011
= 7 = 0111	= 1 1 0001	- 1 1111
	drop carry = 0001	

Two's Complement

■ Why does it work?

- Put another way: given the bit representation of a positive integer, we want the negative bit representation to always sum to 0 (ignoring the carry-out bit) when added to the positive representation

- This turns out to be the *bitwise complement plus one*

- What should the 8-bit representation of -1 be?

$$\begin{array}{r}
 00000001 \\
 + \underline{????????} \\
 \hline
 00000000
 \end{array}
 \quad \text{(we want whichever bit string gives the right result)}$$

$$\begin{array}{r}
 00000010 \\
 + \underline{????????} \\
 \hline
 00000000
 \end{array}
 \qquad
 \begin{array}{r}
 00000011 \\
 + \underline{????????} \\
 \hline
 00000000
 \end{array}$$

Two's Complement

■ Why does it work?

- Put another way: given the bit representation of a positive integer, we want the negative bit representation to always sum to 0 (ignoring the carry-out bit) when added to the positive representation
- This turns out to be the *bitwise complement plus one*
 - What should the 8-bit representation of -1 be?

$$\begin{array}{r}
 00000001 \\
 + \underline{11111111} \\
 \hline
 100000000
 \end{array}
 \quad \text{(we want whichever bit string gives the right result)}$$

$$\begin{array}{r}
 00000010 \\
 + \underline{????????} \\
 \hline
 00000000
 \end{array}
 \qquad
 \begin{array}{r}
 00000011 \\
 + \underline{????????} \\
 \hline
 00000000
 \end{array}$$

Two's Complement

■ Why does it work?

- Put another way: given the bit representation of a positive integer, we want the negative bit representation to always sum to 0 (ignoring the carry-out bit) when added to the positive representation
- This turns out to be the *bitwise complement plus one*
 - What should the 8-bit representation of -1 be?

$$\begin{array}{r}
 00000001 \\
 +11111111 \\
 \hline
 100000000
 \end{array}
 \quad \text{(we want whichever bit string gives the right result)}$$

$$\begin{array}{r}
 00000010 \\
 +11111110 \\
 \hline
 100000000
 \end{array}
 \qquad
 \begin{array}{r}
 00000011 \\
 +11111101 \\
 \hline
 100000000
 \end{array}$$


Unsigned & Signed Numeric Values

X	Unsigned	Signed
0000	0	0
0001	1	1
0010	2	2
0011	3	3
0100	4	4
0101	5	5
0110	6	6
0111	7	7
1000	8	-8
1001	9	-7
1010	10	-6
1011	11	-5
1100	12	-4
1101	13	-3
1110	14	-2
1111	15	-1


- Both signed and unsigned integers have limits
 - If you compute a number that is too big, you wrap: $6 + 4 = ?$ $15U + 2U = ?$
 - If you compute a number that is too small, you wrap: $-7 - 3 = ?$ $0U - 2U = ?$
 - Answers are only correct mod 2^b
- The CPU may be capable of “throwing an exception” for overflow on signed values
 - It won't for unsigned
- But C and Java just cruise along silently when overflow occurs...

Visualizations

- Same W bits interpreted as signed vs. unsigned:



- Two's complement (signed) addition: x and y are W bits wide



Numeric Ranges

■ Unsigned Values

- UMin = 0
 - 000...0
- UMax = $2^w - 1$
 - 111...1

■ Two's Complement Values

- TMin = -2^{w-1}
 - 100...0
- TMax = $2^{w-1} - 1$
 - 011...1

■ Other Values

- Negative 1
 - 111...1 0xFFFFFFFF (32 bits)

Values for $W = 16$

	Decimal	Hex	Binary
UMax	65535	FF FF	11111111 11111111
TMax	32767	7F FF	01111111 11111111
TMin	-32768	80 00	10000000 00000000
-1	-1	FF FF	11111111 11111111
0	0	00 00	00000000 00000000

Integer & Floating Point Numbers

- Representation of integers: unsigned and signed
- Unsigned and signed integers in C
- Arithmetic and shifting
- Sign extension

- Background: fractional binary numbers
- IEEE floating-point standard
- Floating-point operations and rounding
- Floating-point in C

Values for Different Word Sizes

	W			
	8	16	32	64
UMax	255	65,535	4,294,967,295	18,446,744,073,709,551,615
TMax	127	32,767	2,147,483,647	9,223,372,036,854,775,807
TMin	-128	-32,768	-2,147,483,648	-9,223,372,036,854,775,808

■ Observations

- $|TMin| = TMax + 1$
 - Asymmetric range
- $UMax = 2 * TMax + 1$

■ C Programming

- `#include <limits.h>`
- Declares constants, e.g.,
 - `ULONG_MAX`
 - `LONG_MAX`
 - `LONG_MIN`
- Values are platform specific
- See: `/usr/include/limits.h` on Linux

Signed vs. Unsigned in C

■ Constants

- By default are considered to be signed integers
- Use “U” suffix to force unsigned:
 - `0U, 4294967259U`

Signed vs. Unsigned in C

■ Casting

- `int tx, ty;`
- `unsigned ux, uy;`
- Explicit casting between signed & unsigned:
 - `tx = (int) ux;`
 - `uy = (unsigned) ty;`
- Implicit casting also occurs via assignments and function calls:
 - `tx = ux;`
 - `uy = ty;`
 - The gcc flag `-Wsign-conversion` produces warnings for implicit casts, but `-Wall` does not!
- How does casting between signed and unsigned work – what values are going to be produced?
 - *Bits are unchanged*, just interpreted differently!

Casting Surprises

■ Expression Evaluation

- If you mix unsigned and signed in a single expression, then *signed values implicitly cast to unsigned*
- Including comparison operations $<$, $>$, $==$, $<=$, $>=$
- Examples for $W = 32$: **TMIN = -2,147,483,648** **TMAX = 2,147,483,647**

■ Constant ₁	Constant ₂	Relation	Evaluation
0	0U	==	unsigned
-1	0	<	signed
-1	0U	>	unsigned
2147483647	-2147483647-1	>	signed
2147483647U	-2147483647-1	<	unsigned
-1	-2	>	signed
(unsigned)-1	-2	>	unsigned
2147483647	2147483648U	<	unsigned
2147483647	(int) 2147483648U	>	signed

Integer & Floating Point Numbers

- Representation of integers: unsigned and signed
- Unsigned and signed integers in C
- Arithmetic and shifting
- Sign extension

- Background: fractional binary numbers
- IEEE floating-point standard
- Floating-point operations and rounding
- Floating-point in C

Shift Operations for unsigned integers

- **Left shift:** $x \ll y$
 - Shift bit-vector x left by y positions
 - Throw away extra bits on left
 - Fill with 0s on right
- **Right shift:** $x \gg y$
 - Shift bit-vector x right by y positions
 - Throw away extra bits on right
 - Fill with 0s on left

x	00000110
$\ll 3$	00110000
$\gg 2$	00000001

x	11110010
$\ll 3$	10010000
$\gg 2$	00111100

Shift Operations for signed integers

- **Left shift:** $x \ll y$
 - Equivalent to multiplying by 2^y
 - (if resulting value fits, no 1s are lost)
- **Right shift:** $x \gg y$
 - Logical shift (for unsigned values)
 - Fill with 0s on left
 - Arithmetic shift (for signed values)
 - Replicate most significant bit on left
 - Maintains sign of x
 - Equivalent to dividing by 2^y
 - Correct rounding (towards 0) requires some care with signed numbers

x	01100010
$\ll 3$	00010000
Logical $\gg 2$	00011000
Arithmetic $\gg 2$	00011000

x	10100010
$\ll 3$	00010000
Logical $\gg 2$	00101000
Arithmetic $\gg 2$	11101000

**Undefined behavior when
 $y < 0$ or $y \geq \text{word_size}$**

Using Shifts and Masks

■ Extract 2nd most significant byte of an integer

- First shift: $x \gg (2 * 8)$
- Then mask: $(x \gg 16) \& 0xFF$

x	01100001 01100010 01100011 01100100
$x \gg 16$	00000000 00000000 01100001 01100010
$(x \gg 16) \& 0xFF$	00000000 00000000 00000000 11111111 00000000 00000000 00000000 01100010

■ Extracting the sign bit

- $(x \gg 31) \& 1$ - need the “& 1” to clear out all other bits except LSB

■ Conditionals as Boolean expressions (**assuming x is 0 or 1**)

- if (x) a=y else a=z; which is the same as $a = x ? y : z$;
- Can be re-written as: $a = ((x \ll 31) \gg 31) \& y + (!x \ll 31) \gg 31) \& z$;


Sign Extension

■ Task:

- Given w -bit signed integer x
- Convert it to $w+k$ -bit integer with same value

■ Rule:

- Make k copies of sign bit:
- $X' = \underbrace{x_{w-1}, \dots, x_{w-1}}_{k \text{ copies of MSB}}, x_{w-1}, x_{w-2}, \dots, x_0$



Sign Extension Example

- Converting from smaller to larger integer data type
- C automatically performs sign extension

```
short int x = 12345;
int      ix = (int) x;
short int y = -12345;
int      iy = (int) y;
```

	Decimal	Hex	Binary
x	12345	30 39	00110000 01101101
ix	12345	00 00 30 39	00000000 00000000 00110000 01101101
y	-12345	CF C7	11001111 11000111
iy	-12345	FF FF CF C7	11111111 11111111 11001111 11000111