

## Bits, Bytes, and Integers

CSci 2021: Machine Architecture and Organization  
September 14th-19th, 2018

Your instructor: Stephen McCamant


Based on slides originally by:  
Randy Bryant, Dave O'Hallaron

## Today: Bits, Bytes, and Integers

- Representing information as bits
- Bit-level manipulations
- Integers
  - Representation: unsigned and signed
  - Conversion, casting
  - Expanding, truncating
  - Addition, negation, multiplication, shifting
  - Summary
- Representations in memory, pointers, strings

## Everything is bits

- Each bit is 0 or 1
- By encoding/interpreting sets of bits in various ways
  - Computers determine what to do (instructions)
  - ... and represent and manipulate numbers, sets, strings, etc...
- Why bits? Electronic Implementation
  - Easy to store with bistable elements
  - Reliably transmitted on noisy and inaccurate wires



## For example, can count in binary

- Base 2 Number Representation
  - Represent  $15213_{10}$  as  $11101101101101_2$
  - Represent  $1.20_{10}$  as  $1.0011001100110011[0011]..._2$
  - Represent  $1.5213 \times 10^4$  as  $1.1101101101101_2 \times 2^{13}$

## Encoding Byte Values

- Byte = 8 bits
  - Binary  $0000000_2$  to  $11111111_2$
  - Decimal:  $0_{10}$  to  $255_{10}$
  - Hexadecimal  $00_{16}$  to  $FF_{16}$ 
    - Base 16 number representation
    - Use characters '0' to '9' and 'A' to 'F'
  - Write  $FA1D37B_{16}$  in C as
    - $0xFA1D37B$
    - $0xfa1d37b$

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

## Aside: ASCII table

	0	1	2	3	4	5	6	7	8	9	a	b	c	d	e	f
0x0_	\0	^A	^B	^C	^D	^E	^F	^G	^H	^I	^J	^K	^L	^M	^N	^O
0x1_	^P	^Q	^R	^S	^T	^U	^V	^W	^X	^Y	^Z	ESC	FS	GS	RS	US
0x2_	SPC	!	"	#	\$	%	&	'	(	)	*	+	,	-	.	/
0x3_	0	1	2	3	4	5	6	7	8	9	:	;	<	=	>	?
0x4_	@	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
0x5_	P	Q	R	S	T	U	V	W	X	Y	Z	[	\	]	^	_
0x6_	`	a	b	c	d	e	f	g	h	i	j	k	l	m	n	o
0x7_	p	q	r	s	t	u	v	w	x	y	z	{		}	~	DEL

## Example Data Representations

C Data Type	Typical 32-bit	Typical 64-bit	x86-64
char	1	1	1
short	2	2	2
int	4	4	4
long	4	8	8
float	4	4	4
double	8	8	8
long double	-	-	10/16
pointer	4	8	8

## Today: Bits, Bytes, and Integers

- Representing information as bits
- Bit-level manipulations
- Integers
  - Representation: unsigned and signed
  - Conversion, casting
  - Expanding, truncating
  - Addition, negation, multiplication, shifting
  - Summary
- Representations in memory, pointers, strings

## Boolean Algebra

- Developed by George Boole in 19th Century
  - Algebraic representation of logic
  - Encode "True" as 1 and "False" as 0

### And (math: $\wedge$ )

- $A \& B = 1$  when both  $A=1$  and  $B=1$

$\&$	0	1
0	0	0
1	0	1

### Or (math: $\vee$ )

- $A | B = 1$  when either  $A=1$  or  $B=1$

	0	1
0	0	1
1	1	1

### Not (math: $\neg$ )

- $\sim A = 1$  when  $A=0$

$\sim$	0	1
0	1	
1	0	

### Exclusive-Or "xor" (math: $\oplus$ )

- $A \oplus B = 1$  when either  $A=1$  or  $B=1$ , but not both

$\oplus$	0	1
0	0	1
1	1	0

## General Boolean Algebras

- Operate on bit vectors
  - Operations applied bitwise

01101001	01101001	01101001	01101001
$\&$ 01010101	01010101	$\wedge$ 01010101	$\sim$ 01010101
01000001	01111101	00111100	10101010

- All of the properties of Boolean algebra apply

## Example: Representing & Manipulating Sets

### Representation

- Width  $w$  bit vector represents subsets of  $\{0, \dots, w-1\}$
- $a_j = 1$  if  $j \in A$

- 01101001  $\{0, 3, 5, 6\}$
- **76543210**

- 01010101  $\{0, 2, 4, 6\}$
- **76543210**

### Operations

- $\&$  Intersection 01000001  $\{0, 6\}$
- $|$  Union 01111101  $\{0, 2, 3, 4, 5, 6\}$
- $\wedge$  Symmetric difference 00111100  $\{2, 3, 4, 5\}$
- $\sim$  Complement 10101010  $\{1, 3, 5, 7\}$

## Bit-Level Operations in C

### Operations $\&$ , $|$ , $\sim$ , $\wedge$ Available in C

- Apply to any "integral" data type
  - long, int, short, char, unsigned
- View arguments as bit vectors
- Arguments applied bit-wise

### Examples (Char data type)

- $\sim 0x41 \rightarrow 0xBE$ 
  - $\sim 01000001_2 \rightarrow 10111110_2$
- $\sim 0x00 \rightarrow 0xFF$ 
  - $\sim 00000000_2 \rightarrow 11111111_2$
- $0x69 \& 0x55 \rightarrow 0x41$ 
  - $01101001_2 \& 01010101_2 \rightarrow 01000001_2$
- $0x69 | 0x55 \rightarrow 0x7D$ 
  - $01101001_2 | 01010101_2 \rightarrow 01111101_2$

## Contrast: Logic Operations in C

### Contrast to Logical Operators

- `&&`, `||`, `!`
  - View 0 as "False"
  - Anything non-zero is "True"
  - Always evaluates both operands
  - Early exit
- Example
  - `!0x41` → `0`
  - `!0x00` → `1`
  - `!!0x41` → `1`
  - `0x69 && 0x01` → `0`
  - `0x69 || 0x55` → `0x01`
- `p && *p` (avoids null pointer access)

Watch out for `&&` vs. `&` (and `||` vs. `|`)... one of the more common oopsies in C programming

## Shift Operations

### Left Shift: `x << y`

- Shift bit-vector `x` left `y` positions
  - Throw away extra bits on left
  - Fill with 0's on right

Argument <code>x</code>	01100010
<code>&lt;&lt; 3</code>	00010000
Log. <code>&gt;&gt; 2</code>	00011000
Arith. <code>&gt;&gt; 2</code>	00011000

### Right Shift: `x >> y`

- Shift bit-vector `x` right `y` positions
  - Throw away extra bits on right
  - Logical shift: fill with 0's on left
  - Arithmetic shift: replicate most significant bit on left

Argument <code>x</code>	10100010
<code>&lt;&lt; 3</code>	00010000
Log. <code>&gt;&gt; 2</code>	00101000
Arith. <code>&gt;&gt; 2</code>	11101000

### Undefined Behavior

- Shift amount `< 0` or `≥` word size
- Signed shift into or out of sign bit (i.e., arith. behavior not assured)

## Interlude: ChimeIn

<https://chimein.cla.umn.edu/course/view/2021>

### Which of the following numbers represented in hex is not a multiple of 4?

- `0x2248730c`
- `0xf4d56f9e`
- `0x1c841a94`
- `0x0970bd90`
- `0xacc3f6978`

### Idea: it is enough to look at the last digit

- Like divisibility by 10, 2 and 5 for decimal
- `0x0`, `0x4`, `0x8`, and `0xc` = decimal 12 are multiples of 4
- `0xe` = 14 is even but not a multiple of 4

## Today: Bits, Bytes, and Integers

- Representing information as bits
- Bit-level manipulations
- Integers
  - Representation: unsigned and signed
  - Conversion, casting
  - Expanding, truncating
  - Addition, negation, multiplication, shifting
  - Summary
- Representations in memory, pointers, strings
- Summary

## Binary Number Property

### Claim

$$1 + 1 + 2 + 4 + 8 + \dots + 2^{w-1} = 2^w$$

$$1 + \sum_{i=0}^{w-1} 2^i = 2^w$$

- `w = 0`:
  - $1 = 2^0$
- Assume true for `w-1`:
  - $1 + 1 + 2 + 4 + 8 + \dots + 2^{w-1} + 2^w = 2^w + 2^w = 2^{w+1}$
  - $= 2^w$

## Encoding Integers

### Unsigned

$$B2U(X) = \sum_{i=0}^{w-1} x_i \cdot 2^i$$

### Two's Complement

$$B2T(X) = -x_{w-1} \cdot 2^{w-1} + \sum_{i=0}^{w-2} x_i \cdot 2^i$$

```
short int x = 15213;
short int y = -15213;
```

Sign Bit

### C short 2 bytes long

	Decimal	Hex	Binary
<code>x</code>	15213	3B 6D	00111011 01101101
<code>y</code>	-15213	C4 93	11000100 10010011

### Sign Bit

- For 2's complement, most significant bit indicates sign
  - 0 for nonnegative
  - 1 for negative

### Four-bit Example, unsigned

Unsigned:  $\begin{matrix} 8 & 4 & 2 & 1 \\ \downarrow & \downarrow & \downarrow & \downarrow \\ 1 & 1 & 1 & 1 \end{matrix} = 15_{10}$

This approach can represent 0 through 15

23

### Four-bit Example, sign + magnitude

Sign + magnitude:  $\begin{matrix} 1 & 1 & 1 & 1 \\ \uparrow & \uparrow & \uparrow & \uparrow \\ \text{Negate} & & & \\ \text{if 1} & & & \\ 4 & 2 & 1 & \end{matrix} = 7_{10}$

This approach can represent -7 through 7  
Disadvantages: special cases, -0

24

### Four-bit Example, two's complement

Two's complement:  $\begin{matrix} 1 & 1 & 1 & 1 \\ \uparrow & \uparrow & \uparrow & \uparrow \\ -8 & 4 & 2 & 1 \end{matrix} = -1_{10}$

This approach can represent -8 through 7

25

### Four-bit Example

Unsigned:  $\begin{matrix} 8 & 4 & 2 & 1 \\ \downarrow & \downarrow & \downarrow & \downarrow \\ 1 & 1 & 1 & 1 \end{matrix} = 15_{10}$

Two's complement:  $\begin{matrix} 1 & 1 & 1 & 1 \\ \uparrow & \uparrow & \uparrow & \uparrow \\ -8 & 4 & 2 & 1 \end{matrix} = -1_{10}$

26

### Encoding Integers

#### Unsigned

$$B2U(X) = \sum_{i=0}^{w-1} x_i \cdot 2^i$$

#### Two's Complement

$$B2T(X) = -x_{w-1} \cdot 2^{w-1} + \sum_{i=0}^{w-2} x_i \cdot 2^i$$

```
short int x = 15213;
short int y = -15213;
```

#### ■ C short 2 bytes long

	Decimal	Hex	Binary
x	15213	3B 6D	00111011 01101101
y	-15213	C4 93	11000100 10010011

#### ■ Sign Bit

- For 2's complement, most significant bit indicates sign
  - 0 for nonnegative
  - 1 for negative

Sign Bit

27

### Two-complement Encoding Example (Cont.)

```
x = 15213: 00111011 01101101
y = -15213: 11000100 10010011
```

Weight	15213	-15213
1	1	1
2	0	0
4	1	4
8	1	8
16	0	0
32	1	32
64	1	64
128	0	0
256	1	256
512	1	512
1024	0	0
2048	1	2048
4096	1	4096
8192	1	8192
16384	0	0
-32768	0	0
<b>Sum</b>	<b>15213</b>	<b>-15213</b>

28

## 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

- Minus 1  
111...1

Values for  $W = 16$

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

## Values for Different Word Sizes

	W			
	8	16	32	64
<b>UMax</b>	255	65,535	4,294,967,295	18,446,744,073,709,551,615
<b>TMax</b>	127	32,767	2,147,483,647	9,223,372,036,854,775,807
<b>TMin</b>	-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 platform specific

## Unsigned & Signed Numeric Values

X	B2U(X)	B2T(X)
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

### Equivalence

- Same encodings for nonnegative values

### Uniqueness

- Every bit pattern represents unique integer value
- Each representable integer has unique bit encoding

### ⇒ Can Invert Mappings

- $U2B(x) = B2U^{-1}(x)$ 
  - Bit pattern for unsigned integer
- $T2B(x) = B2T^{-1}(x)$ 
  - Bit pattern for two's comp integer

## Today: Bits, Bytes, and Integers

### Representing information as bits


### Bit-level manipulations

### Integers

- Representation: unsigned and signed
- Conversion, casting**
  - Expanding, truncating
  - Addition, negation, multiplication, shifting
  - Summary

### Representations in memory, pointers, strings

## Mapping Between Signed & Unsigned



- Mappings between unsigned and two's complement numbers: Keep bit representations and reinterpret**

## Mapping Signed ↔ Unsigned


Bits	Signed	Unsigned
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	-7	9
1010	-6	10
1011	-5	11
1100	-4	12
1101	-3	13
1110	-2	14
1111	-1	15

Arrows indicate the T2U and U2T mappings between the Signed and Unsigned columns.

## Mapping Signed ↔ Unsigned

Bits	Signed	Unsigned
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	-7	9
1010	-6	10
1011	-5	11
1100	-4	12
1101	-3	13
1110	-2	14
1111	-1	15


## Relation between Signed & Unsigned



## Conversion Visualized

### 2's Comp. → Unsigned

- Ordering Inversion
- Negative → Big Positive



## Signed vs. Unsigned in C

### Constants

- By default are considered to be signed integers
- Unsigned if have "U" as suffix  
0U, 4294967259U

### Casting

- Explicit casting between signed & unsigned same as U2T and T2U  

```
int tx, ty;
unsigned ux, uy;
tx = (int) ux;
uy = (unsigned) ty;
```
- Implicit casting also occurs via assignments and procedure calls  

```
tx = ux;
uy = ty;
```

## Casting and Comparison Surprises

### Expression Evaluation

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

Constant <sub>1</sub>	Constant <sub>2</sub>	Relation	Evaluation
0	0U	==	unsigned
-1	0	<	signed
-1	0U	>	unsigned

<https://chimein.cla.umn.edu/course/view/2021>

## Casting and Comparison Surprises

### Expression Evaluation

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

Constant <sub>1</sub>	Constant <sub>2</sub>	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

## Brief announcements

- About 1 week left to go on HA1
  - Keep your questions coming, don't put it off
- My office hours tomorrow will move to 3-4pm
  - Still in 4-225E Keller

## Summary

### Casting Signed ↔ Unsigned: Basic Rules

- Bit pattern is maintained
- But reinterpreted
- Can have unexpected effects: adding or subtracting  $2^w$
  
- Expression containing signed and unsigned int
  - int is cast to unsigned!!

## Typical Usage

```
/* Kernel memory region holding user-accessible data */
#define KSIZE 1024
char kbuf[KSIZE];

/* Copy at most maxlen bytes from kernel region to user buffer */
int copy_from_kernel(void *user_dest, int maxlen) {
    /* Byte count len is minimum of buffer size and maxlen */
    int len = KSIZE < maxlen ? KSIZE : maxlen;
    memcpy(user_dest, kbuf, len);
    return len;
}
```

```
#define MSIZE 528

void getstuff() {
    char mybuf[MSIZE];
    copy_from_kernel(mybuf, MSIZE);
    printf("%s\n", mybuf);
}
```

## Malicious Usage

```
/* Declaration of library function memcpy */
void *memcpy(void *dest, void *src, size_t n);
typedef unsigned long size_t;
```

```
/* Kernel memory region holding user-accessible data */
#define KSIZE 1024
char kbuf[KSIZE];

/* Copy at most maxlen bytes from kernel region to user buffer */
int copy_from_kernel(void *user_dest, int maxlen) {
    /* Byte count len is minimum of buffer size and maxlen */
    int len = KSIZE < maxlen ? KSIZE : maxlen;
    memcpy(user_dest, kbuf, len);
    return len;
}
```

```
#define MSIZE 528


void getstuff() {
    char mybuf[MSIZE];
    copy_from_kernel(mybuf, -MSIZE);
    . . .
}
```

## Today: Bits, Bytes, and Integers

- Representing information as bits
- Bit-level manipulations
- Integers
  - Representation: unsigned and signed
  - Conversion, casting
  - Expanding, truncating
  - Addition, negation, multiplication, shifting
  - Summary
- Representations in memory, pointers, strings

## 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:
  - $X' = X_{w-1} \dots X_{w-1} X_{w-1} X_{w-1} X_{w-2} \dots X_0$



## Sign Extension Example

```
short int x = 15213;
int ix = (int) x;
short int y = -15213;
int iy = (int) y;
```

	Decimal	Hex	Binary
x	15213	3B 6D	00111011 01101101
ix	15213	00 00 3B 6D	00000000 00000000 00111011 01101101
y	-15213	C4 93	11000100 10010011
iy	-15213	FF FF C4 93	11111111 11111111 11000100 10010011

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

## Summary: Expanding, Truncating: Basic Rules


- Expanding (e.g., short int to int)
  - Unsigned: zeros added
  - Signed: sign extension
  - Both yield expected result
- Truncating (e.g., unsigned to unsigned short)
  - Unsigned/signed: bits are truncated
  - Result reinterpreted
  - Unsigned: mod operation
  - Signed: similar to mod
  - For small numbers yields expected behavior

## Today: Bits, Bytes, and Integers


- Representing information as bits
- Bit-level manipulations
- Integers
  - Representation: unsigned and signed
  - Conversion, casting
  - Expanding, truncating
  - Addition, negation, multiplication, shifting
- Representations in memory, pointers, strings
- Summary

## Unsigned Addition


Operands:  $w$  bits



True Sum:  $w+1$  bits




Discard Carry:  $w$  bits



- Standard Addition Function
  - Ignores carry output
- Implements Modular Arithmetic
 
$$s = UAdd_w(u, v) = u + v \bmod 2^w$$



## Visualizing (Mathematical) Integer Addition

- Integer Addition
  - 4-bit integers  $u, v$
  - Compute true sum  $Add_4(u, v)$
  - Values increase linearly with  $u$  and  $v$
  - Forms planar surface



## Visualizing Unsigned Addition

- Wraps Around
  - If true sum  $\geq 2^w$
  - At most once






## Mathematical Properties

### Modular Addition Forms an Abelian Group

- Closed** under addition  
 $0 \leq \text{UAdd}_w(u, v) \leq 2^w - 1$
- Commutative**  
 $\text{UAdd}_w(u, v) = \text{UAdd}_w(v, u)$
- Associative**  
 $\text{UAdd}_w(t, \text{UAdd}_w(u, v)) = \text{UAdd}_w(\text{UAdd}_w(t, u), v)$
- 0** is additive identity  
 $\text{UAdd}_w(u, 0) = u$
- Every element has additive **inverse**
  - Let  $\text{UComp}_w(u) = 2^w - u$   
 $\text{UAdd}_w(u, \text{UComp}_w(u)) = 0$

## Two's Complement Addition




### TAdd and UAdd have Identical Bit-Level Behavior

- Signed vs. unsigned addition in C:  
`int s, t, u, v;`  
`s = (int) ((unsigned) u + (unsigned) v);`  
`t = u + v;`
- Will give `s == t`

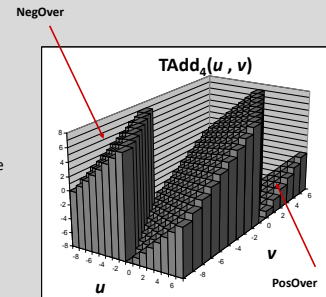
## TAdd Overflow

- Functionality**
  - True sum requires  $w+1$  bits
  - Drop off MSB
  - Treat remaining bits as 2's comp. integer



## Visualizing 2's Complement Addition

- Values**
  - 4-bit two's comp.
  - Range from -8 to +7
- Wraps Around**
  - If  $\text{sum} \geq 2^{w-1}$ 
    - Becomes negative
    - At most once
  - If  $\text{sum} < -2^{w-1}$ 
    - Becomes positive
    - At most once



## Mathematical Properties of TAdd

### Isomorphic Group to unsigneds with UAdd

- $\text{TAdd}_w(u, v) = \text{U2T}(\text{UAdd}_w(\text{T2U}(u), \text{T2U}(v)))$ 
  - Since both have identical bit patterns

### Two's Complement Under TAdd Forms a Group


- Closed, Commutative, Associative, 0 is additive identity
- Every element has additive inverse

$$\text{TComp}_w(u) = \begin{cases} -u & u \neq \text{TMin}_w \\ \text{TMin}_w & u = \text{TMin}_w \end{cases}$$

## Characterizing TAdd

### Functionality


- True sum requires  $w+1$  bits
- Drop off MSB
- Treat remaining bits as 2's comp. integer



$$\text{TAdd}_w(u, v) = \begin{cases} u + v + 2^w & u + v < \text{TMin}_w \text{ (NegOver)} \\ u + v & \text{TMin}_w \leq u + v \leq \text{TMax}_w \\ u + v - 2^w & \text{TMax}_w < u + v \text{ (PosOver)} \end{cases}$$

## Signed/Unsigned Overflow Differences

- **Unsigned:**
  - Overflow if carry out of last position
  - Also just called "carry" (C)
- **Signed:**
  - Result wrong if input signs are the same but output sign is different
  - In CPUs, unqualified "overflow" usually means signed (O or V)



61

## Sign bit table, signed ordering

	-4	-3	-2	-1	0	1	2	3
3	-1	0	1	2	3	-4	-3	-2
2	-2	-1	0	1	2	3	-4	-3
1	-3	-2	-1	0	1	2	3	-4
0	-4	-3	-2	-1	0	1	2	3
-1	3	-4	-3	-2	-1	0	1	2
-2	2	3	-4	-3	-2	-1	0	1
-3	1	2	3	-4	-3	-2	-1	0
-4	0	1	2	3	-4	-3	-2	-1

62

## Sign bit table, unsigned ordering

	0	1	2	3	-4	-3	-2	-1
-1	-1	0	1	2	3	-4	-3	-2
-2	-2	-1	0	1	2	3	-4	-3
-3	-3	-2	-1	0	1	2	3	-4
-4	-4	-3	-2	-1	0	1	2	3
3	3	-4	-3	-2	-1	0	1	2
2	2	3	-4	-3	-2	-1	0	1
1	1	2	3	-4	-3	-2	-1	0
0	0	1	2	3	-4	-3	-2	-1

63

## Negation: Complement & Increment

- **Claim: Following Holds for 2's Complement**

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

- **Complement**

- Observation:  $\sim x + x = 1111..111 = -1$

$$\begin{array}{r} x \quad 10011101 \\ + \sim x \quad 01100010 \\ \hline -1 \quad 11111111 \end{array}$$

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

64

## Complement & Increment Examples

x = 15213

	Decimal	Hex	Binary
x	15213	3B 6D	00111011 01101101
$\sim x$	-15214	C4 92	11000100 10010010
$\sim x + 1$	-15213	C4 93	11000100 10010011
y	-15213	C4 93	11000100 10010011

x = 0

	Decimal	Hex	Binary
0	0	00 00	00000000 00000000
$\sim 0$	-1	FF FF	11111111 11111111
$\sim 0 + 1$	0	00 00	00000000 00000000

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

65

## Multiplication

- **Goal: Computing Product of w-bit numbers x, y**

- Either signed or unsigned

- **But, exact results can be bigger than w bits**

- Unsigned: up to 2w bits
  - Result range:  $0 \leq x * y \leq (2^w - 1)^2 = 2^{2w} - 2^{w+1} + 1$
- Two's complement min (negative): Up to 2w-1 bits
  - Result range:  $x * y \geq (-2^{w-1}) * (2^{w-1} - 1) = -2^{2w-2} + 2^{w-1}$
- Two's complement max (positive): Up to 2w bits, but only for  $(TMin_w)^2$ 
  - Result range:  $x * y \leq (-2^{w-1})^2 = 2^{2w-2}$


- **So, maintaining exact results...**

- would need to keep expanding word size with each product computed
- is done in software, if needed
  - e.g., by "arbitrary precision" arithmetic packages

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition


66

## Unsigned Multiplication in C



- Standard Multiplication Function
  - Ignores high order  $w$  bits
- Implements Modular Arithmetic
  - $UMult_w(u, v) = u \cdot v \text{ mod } 2^w$

## Signed Multiplication in C




- Standard Multiplication Function
  - Ignores high order  $w$  bits
  - Some of which are different for signed vs. unsigned multiplication
  - Lower bits are the same

## Code Security Example #2

- SUN XDR library
  - Widely used library for transferring data between machines

```
void* copy_elements(void *ele_src[], int ele_cnt, size_t ele_size);
```



## XDR Code

```
void* copy_elements(void *ele_src[], int ele_cnt, size_t ele_size) {
    /*
     * Allocate buffer for ele_cnt objects, each of ele_size bytes
     * and copy from locations designated by ele_src
     */
    void *result = malloc(ele_cnt * ele_size);
    if (result == NULL)
        /* malloc failed */
        return NULL;
    void *next = result;
    int i;
    for (i = 0; i < ele_cnt; i++) {
        /* Copy object i to destination */
        memcpy(next, ele_src[i], ele_size);
        /* Move pointer to next memory region */
        next += ele_size;
    }
    return result;
}
```


## XDR Vulnerability

```
malloc(ele_cnt * ele_size)
```

- What if:
  - $ele\_cnt = 2^{20} + 1$
  - $ele\_size = 4096 = 2^{12}$
  - Allocation = ??
- Chime in: <https://chimein.cla.umn.edu/course/view/2021> (Question 16257)
- $(2^{20} + 1) \cdot 2^{12} = 2^{20} \cdot 2^{12} + 2^{12} = 2^{32} + 2^{12} \cong 2^{12}$
- How can I make this function secure?

## Power-of-2 Multiply with Shift

- Operation
  - $u \ll k$  gives  $u * 2^k$
  - Both signed and unsigned



- Examples
  - $u \ll 3 == u * 8$
  - $(u \ll 5) - (u \ll 3) == u * 24$
  - Most machines shift and add faster than multiply

## Compiled Multiplication Code

### C Function

```
long mul12(long x)
{
    return x*12;
}
```

### Compiled Arithmetic Operations

```
leaq(%rax,%rax,2),%rax
salq$2,%rax
```

### Explanation

```
t <- x+x*2
return t << 2;
```

- C compiler automatically generates shift/add code when multiplying by constant

## Background: Rounding in Math

- How to round to the nearest integer?
- Cannot have both:
  - $\text{round}(x + k) = \text{round}(x) + k$  (k integer), "translation invariance"
  - $\text{round}(-x) = -\text{round}(x)$  "negation invariance"
- $\lfloor x \rfloor$ , read "floor": always round down (to  $-\infty$ ):
  - $\lfloor 2.0 \rfloor = 2, \lfloor 1.7 \rfloor = 1, \lfloor -2.2 \rfloor = -3$
- $\lceil x \rceil$ , read "ceiling": always round up (to  $+\infty$ ):
  - $\lceil 2.0 \rceil = 2, \lceil 1.7 \rceil = 2, \lceil -2.2 \rceil = -2$
- C integer operators mostly use round to zero, which is like floor for positive and ceiling for negative

## Division in C

- Integer division  $/$ : rounds towards 0
  - Choice (settled in C99) is historical, via FORTRAN and most CPUs
- Division by zero: undefined, usually fatal
- Unsigned division: no overflow possible
- Signed division: overflow *almost* impossible
  - Exception: TMin/-1 is un-representable, and so undefined
  - On x86 this too is a default-fatal exception


## Undefined behavior

- Many things you should not do are officially called "undefined" by the C language standard
  - Meaning: compiler can do anything it wants
- Examples:
  - Accessing beyond the ends of an array
  - Dividing by zero
  - Overflow in signed operations
  - Shifts of negative values
- Bad interaction with improving compiler optimizers
- Gap between standard and lenient practical compilers not yet resolved

Things you do in this section of the course!

## Unsigned Power-of-2 Divide with Shift

- Quotient of Unsigned by Power of 2
  - $u \gg k$  gives  $\lfloor u / 2^k \rfloor$
  - Uses logical shift



	Division	Computed	Hex	Binary
$x$	15213	15213	3B 6D	00111011 01101101
$x \gg 1$	7606.5	7606	1D B6	00011101 10110110
$x \gg 4$	950.8125	950	03 B6	00000011 10110110
$x \gg 8$	59.4257813	59	00 3B	00000000 00111011

## Compiled Unsigned Division Code

### C Function

```
unsigned long udiv8
(unsigned long x)
{
    return x/8;
}
```

### Compiled Arithmetic Operations

```
shrq$3,%rax
```

### Explanation


```
# Logical shift
return x >> 3;
```

- Uses logical shift for unsigned
- For Java Users
  - Logical shift written as  $\gg$

## Signed Power-of-2 Divide with Shift

- Quotient of Signed by Power of 2

- $x \gg k$  gives  $\lfloor x / 2^k \rfloor$
- Uses arithmetic shift
- Rounds wrong direction when  $u < 0$




	Division	Computed	Hex	Binary
$y$	-15213	-15213	C4 93	11000100 10010011
$y \gg 1$	-7606.5	-7607	E2 49	11100010 01001001
$y \gg 4$	-950.8125	-951	FC 49	11111100 01001001
$y \gg 8$	-59.4257813	-60	FF C4	11111111 11000100

## Correct Power-of-2 Divide

- Quotient of Negative Number by Power of 2

- Want  $\lceil x / 2^k \rceil$  (Round Toward 0)
- Compute as  $\lfloor (x+2^k-1) / 2^k \rfloor$ 
  - In C:  $(x + (1 \ll k) - 1) \gg k$
  - Biases dividend toward 0


Case 1: No rounding



*Biasing has no effect*

## Correct Power-of-2 Divide (Cont.)

Case 2: Rounding



*Biasing adds 1 to final result*

## Compiled Signed Division Code

C Function

```
long idiv8(long x)
{
    return x/8;
}
```

Compiled Arithmetic Operations

```
testq %rax, %rax
js L4
L3:
    sarq $3, %rax
    ret
L4:
    addq $7, %rax
    jmp L3
```

Explanation

```
if x < 0
    x += 7;
# Arithmetic shift
return x >> 3;
```

- Uses arithmetic shift for int
- For Java Users
  - Arith. shift written as  $\gg$

## Remainder operator

- Written as % in C
- $x \% y$  is the remainder after division  $x / y$
- E.g.,  $x \% 10$  is the lowest digit of non-negative  $x$
- Behavior for negative values matches  $\lceil \cdot \rceil$ 's rounding toward zero
  - $b * (a / b) + (a \% b) = a$
- I.e. sign of remainder matches sign of dividend
- (Some other languages have other conventions: sign of result equals sign of divisor, sometimes distinguished as "modulo", or always positive)

## Today: Bits, Bytes, and Integers

- Representing information as bits
- Bit-level manipulations
- Integers
  - Representation: unsigned and signed
  - Conversion, casting
  - Expanding, truncating
  - Addition, negation, multiplication, shifting
  - Summary
- Representations in memory, pointers, strings

## Arithmetic: Basic Rules

- **Addition:**
  - Unsigned/signed: Normal addition followed by truncate, same operation on bit level
  - Unsigned: addition mod  $2^w$ 
    - Mathematical addition + possible subtraction of  $2^w$
  - Signed: modified addition mod  $2^w$  (result in proper range)
    - Mathematical addition + possible addition or subtraction of  $2^w$
- **Multiplication:**
  - Unsigned/signed: Normal multiplication followed by truncate, same operation on bit level
  - Unsigned: multiplication mod  $2^w$
  - Signed: modified multiplication mod  $2^w$  (result in proper range)

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

85

## Arithmetic: Basic Rules

- **Unsigned ints, 2's complement ints are isomorphic rings: isomorphism = casting**
- **Left shift**
  - Unsigned/signed: multiplication by  $2^k$
  - Always logical shift
- **Right shift**
  - Unsigned: logical shift, div (division + round to zero) by  $2^k$
  - Signed: arithmetic shift
    - Positive numbers: div (division + round to zero) by  $2^k$
    - Negative numbers: div (division + round away from zero) by  $2^k$Use biasing to fix

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

86

## Properties of Unsigned Arithmetic

- **Unsigned Multiplication with Addition Forms Commutative Ring**
  - Addition is commutative group
  - Closed under multiplication
$$0 \leq \text{UMult}_w(u, v) \leq 2^w - 1$$
  - Multiplication Commutative
$$\text{UMult}_w(u, v) = \text{UMult}_w(v, u)$$
  - Multiplication is Associative
$$\text{UMult}_w(t, \text{UMult}_w(u, v)) = \text{UMult}_w(\text{UMult}_w(t, u), v)$$
  - 1 is multiplicative identity
$$\text{UMult}_w(u, 1) = u$$
  - Multiplication distributes over addition
$$\text{UMult}_w(t, \text{UAdd}_w(u, v)) = \text{UAdd}_w(\text{UMult}_w(t, u), \text{UMult}_w(t, v))$$

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

87

## Properties of Two's Comp. Arithmetic

- **Isomorphic Algebras**
  - Unsigned multiplication and addition
    - Truncating to  $w$  bits
  - Two's complement multiplication and addition
    - Truncating to  $w$  bits
- **Both Form Rings**
  - Isomorphic to ring of integers mod  $2^w$
- **Comparison to (Mathematical) Integer Arithmetic**
  - Both are rings
  - Integers obey ordering properties, e.g.,
$$u > 0 \Rightarrow u + v > v$$
$$u > 0, v > 0 \Rightarrow u \cdot v > 0$$
  - These properties are not obeyed by two's comp. arithmetic
$$\begin{aligned} TMax + 1 &== TMin \\ 15213 * 30426 &== -10030 \end{aligned} \quad (16\text{-bit words})$$

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

88

## Why Should I Use Unsigned?

- **Don't use without understanding implications**
  - Easy to make mistakes

```
unsigned i;
for (i = cnt-2; i >= 0; i--)
    a[i] += a[i+1];
```
  - Can be very subtle

```
#define DELTA sizeof(int)
int i;
for (i = CNT; i-DELTA >= 0; i-= DELTA)
    . . .
```

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

89

## Counting Down with Unsigned

- **Proper way to use unsigned as loop index**

```
unsigned i;
for (i = cnt-2; i < cnt; i--)
    a[i] += a[i+1];
```
- **See Robert Seacord, *Secure Coding in C and C++***
  - C Standard guarantees that unsigned addition will behave like modular arithmetic
    - $0 - 1 \rightarrow UMax$
- **Even better**

```
size_t i;
for (i = cnt-2; i < cnt; i--)
    a[i] += a[i+1];
```

  - Data type `size_t` defined as unsigned value with length = word size
  - Code will work even if `cnt = UMax`
  - What if `cnt` is signed and `< 0`?

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

90

## Why Should I Use Unsigned? (cont.)

- **Do Use When Performing Modular Arithmetic**
  - Multiprecision arithmetic
- **Do Use When Using Bits to Represent Sets**
  - Logical right shift, no sign extension

## Today: Bits, Bytes, and Integers

- Representing information as bits
- Bit-level manipulations
- Integers
  - Representation: unsigned and signed
  - Conversion, casting
  - Expanding, truncating
  - Addition, negation, multiplication, shifting
  - Summary
- Representations in memory, pointers, strings

## Byte-Oriented Memory Organization




- **Programs refer to data by address**
  - Conceptually, envision it as a very large array of bytes
    - In reality, it's not, but can think of it that way
  - An address is like an index into that array
    - and, a pointer variable stores an address
- **Note: system provides private address spaces to each "process"**
  - Think of a process as a program being executed
  - So, a program can clobber its own data, but not that of others

## Machine Words

- **Any given computer has a "Word Size"**
  - Nominal size of integer-valued data
    - and of addresses
  - Until recently, most machines used 32 bits (4 bytes) as word size
    - Limits addresses to 4GB ( $2^{32}$  bytes)
  - Increasingly, machines have 64-bit word size
    - Potentially, could have 18 EB (exabytes) of addressable memory
    - That's  $18.4 \times 10^{18}$
  - Machines still support multiple data formats
    - Fractions or multiples of word size
    - Always integral number of bytes

## Word-Oriented Memory Organization

- **Addresses Specify Byte Locations**
  - Address of first byte in word
  - Addresses of successive words differ by 4 (32-bit) or 8 (64-bit)



## Example Data Representations


C Data Type	Typical 32-bit	Typical 64-bit	x86-64
char	1	1	1
short	2	2	2
int	4	4	4
long	4	8	8
float	4	4	4
double	8	8	8
long double	-	-	10/16
pointer	4	8	8

## Byte Ordering

- So, how are the bytes within a multi-byte word ordered in memory?
- Conventions
  - Big Endian: Sun, PPC Mac, Internet
    - Least significant byte has highest address
  - Little Endian: x86, ARM processors running Android, iOS, and Windows
    - Least significant byte has lowest address

## Byte Ordering Example


- Example
  - Variable x has 4-byte value of 0x01234567
  - Address given by &x is 0x100




## Representing Integers

Decimal: 15213  
 Binary: 0011 1011 0110 1101  
 Hex: 3 B 6 D


int A = 15213;



long int C = 15213;



int B = -15213;



Two's complement representation

## Examining Data Representations

- Code to Print Byte Representation of Data
  - Casting pointer to unsigned char \* allows treatment as a byte array

```
typedef unsigned char *pointer;

void show_bytes(pointer start, size_t len) {
    size_t i;
    for (i = 0; i < len; i++)
        printf("%p\t0x%.2x\n", start+i, start[i]);
    printf("\n");
}
```

Printf directives:  
 %p: Print pointer  
 %x: Print hexadecimal

## show\_bytes Execution Example


```
int a = 15213;
printf("int a = 15213;\n");
show_bytes((pointer) &a, sizeof(int));
```

Result (Linux x86-64):

```
int a = 15213;
0x7fffb7f71dbc 6d
0x7fffb7f71dbd 3b
0x7fffb7f71dbe 00
0x7fffb7f71dbf 00
```

## Representing Pointers

```
int B = -15213;
int *P = &B;
```



Different compilers & machines assign different locations to objects


Even get different results each time run program



## Representing Strings

```
char S[6] = "18213";
```

- **Strings in C**
  - Represented by array of characters
  - Each character encoded in ASCII format
    - Standard 7-bit encoding of character set
    - Character "0" has code 0x30
      - Digit  $i$  has code  $0x30+i$
  - String should be null-terminated
    - Final character = 0
- **Compatibility**
  - Byte ordering not an issue



## Reading Byte-Reversed Listings

- **Disassembly**
  - Text representation of binary machine code
  - Generated by program that reads the machine code
- **Example Fragment**

Address	Instruction Code	Assembly Rendition
8048365:	5b	pop %ebx
8048366:	81 c3 ab 12 00 00	add \$0x12ab,%ebx
804836c:	83 bb 28 00 00 00 00	cmpl \$0x0,0x28(%ebx)

- **Deciphering Numbers**
  - Value: 0x12ab
  - Pad to 32 bits: 0x000012ab
  - Split into bytes: 00 00 12 ab
  - Reverse: ab 12 00 00

## Integer C Puzzles

1.  $x < 0 \Rightarrow ((x*2) < 0)$
2.  $ux > -1$
3.  $x > 0 \ \&\& \ y > 0 \Rightarrow x + y > 0$
4.  $(x|-x)>>31 == -1$

### Initialization

```
int x = foo();
int y = bar();
unsigned ux = x;
unsigned uy = y;
```

<https://chimein.cla.umn.edu/course/view/2021>

## Bonus: More Integer C Puzzles

- $x < 0 \Rightarrow ((x*2) < 0)$
- $ux \geq 0$
- $x \ \& \ 7 == 7 \Rightarrow (x \ll 30) < 0$
- $ux > -1$
- $x > y \Rightarrow -x < -y$
- $x * x \geq 0$
- $x > 0 \ \&\& \ y > 0 \Rightarrow x + y > 0$
- $x \geq 0 \Rightarrow -x \leq 0$
- $x \leq 0 \Rightarrow -x \geq 0$
- $(x|-x)>>31 == -1$
- $ux \gg 3 == ux/8$
- $x \gg 3 == x/8$
- $x \ \& \ (x-1) != 0$

### Initialization

```
int x = foo();
int y = bar();
unsigned ux = x;
unsigned uy = y;
```