

## Bits, Bytes, and Integers

with contributions from Dr. Bin Ren, College of William & Mary

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

**1**

## 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**

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

2

## 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 bi-stable elements
  - Reliably transmitted on noisy and inaccurate wires



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

3

## The Decimal System and Bases

- **base 10 (decimal): digits 0-9**
  - e.g.,  $316_{10} = 3 \times 10^2 + 1 \times 10^1 + 6 \times 10^0 = 300 + 10 + 6$
- **in the decimal system, 10 is the base, or radix**
- **any integer > 1 can be a base**
- **base 2 has two digits: 0 and 1**
  - bit = binary digit
  - $1011_2 = 1 \times 2^3 + 0 \times 2^2 + 1 \times 2^1 + 1 \times 2^0 =$

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

4

### Converting to Decimal

- **Base eight (octal): digits 0-7**
  - $474_8 =$
- **Base 16 (hexadecimal): digits 0-9 and A-F**
  - $13C_{16} =$
- **Base 2 (binary): digits 0, 1**
  - $100110_2 =$
- **In general, radix  $r$  representations use the first  $r$  chars in {0...9, A...Z} and have the form  $d_{n-1}d_{n-2}...d_1d_0$ .**
  - Summing  $d_{n-1} \times r^{n-1} + d_{n-2} \times r^{n-2} + \dots + d_0 \times r^0$  converts to base 10.

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

5

## Converting from Decimal to Binary

- convert 1693 to binary
- use a divisor of 2 to obtain the following sequence of quotients and remainders

dividend	quotient	remainder
1693	846	1
846	423	0
423	211	1
211	105	1
105	52	1
52	26	0
26	13	0
13	6	1
6	3	0
3	1	1
1	0	1

- read remainders in reverse order  $1693_{10} = 11010011101_2$

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

6

## More Base Conversion Practice

- convert to base 10 by multiplication of powers
  - $10012_3 = ( \quad )_{10}$
- convert from base 10 by repeated division
  - $632_{10} = ( \quad )_8$
- converting base x to base y: convert base x to base 10 then convert base 10 to base y

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

7

## More Base Conversion Practice

- Convert from base 10
  - $123_{10} = ( \quad )_3$  and check
  - $1234_{10} = ( \quad )_{16}$  and check
- Another way to convert from decimal to base n

	$n^8$	$n^7$	$n^6$	$n^5$	$n^4$	$n^3$	$n^2$	$n^1$	$n^0$
for $n = 2$	256	128	64	32	16	8	4	2	1

- From LEFT TO RIGHT, ask "how many" and subtract
- $(219)_{10} = ( \quad )_2 = ( \quad )_{16}$

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

8

## Converting Between Hex and Binary

### chart of values

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

- to convert from binary to hex
  - start at right of binary number
  - convert each group of 4 digits into a hex value
  - e.g., convert  $11011101100_2$  to hex
 

binary:	0110	1110	1100
hex:	6	E	C

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

9

## Converting Between Hex and Binary

### chart of values

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

- to convert from hex to binary
  - replace each hex digit with its binary equivalent
  - e.g., convert  $8A5_{16}$  to binary
 

hex:	8	A	5
binary:	1000	1010	0101

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

10

## Octal


- $2^4 = 16$  and  $2^3 = 8$ 
  - power = # of bits per hex/octal digit
- Binary to Hex
  - every 4 bits = 1 hex digit
- Octal – base 8
  - digits 0-7
- Binary to Octal
  - Every 3 bits = 1 octal digit

DEC	OCT	HEX	BIN	Notes
	0	0	0	-
1	1	1	1	$2^0$
2	2	2	10	$2^1$
3	3	3	11	
4	4	4	100	$2^2$
5	5	5	101	
6	6	6	110	
7	7	7	111	
8	10	8	1000	$2^3$
9	11	9	1001	
10	12	A	1010	
11	13	B	1011	
12	14	C	1100	
13	15	D	1101	
14	16	E	1110	
15	17	F	1111	

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

11

## Every Base is Base 10



<http://cowbirdsinlove.com/43>

In general,  $10_x = X_{10}$

$10_2 = 2$
$10_3 = 3$
$10_4 = 4$
$10_5 = 5$
$10_6 = 6$
$10_7 = 7$
$10_8 = 8$
$10_9 = 9$
$10_{10} = 10$

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

12

## Other Binary Numbers

### ■ 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  $1.1101101101101_2 \times 2^{13}$

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

13

## Encoding Byte Values

### ■ Byte = 8 bits

- Binary  $00000000_2$  to  $11111111_2$
- Decimal:  $0_{10}$  to  $255_{10}$
- Hexadecimal  $00_{16}$  to  $FF_{16}$ 
  - useful for writing binary values concisely
  - 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

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

14

## 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

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

15

## 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

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

16

## Boolean Algebra

### ■ Developed by George Boole in 19th Century

- Algebraic representation of logic
  - Encode "True" as 1 and "False" as 0

#### And

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

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

#### Or

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

	0	1
0	0	1
1	1	1

#### Not

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

$\sim$	0	1
0	1	
1	0	

#### Exclusive-Or (Xor)

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

$\wedge$	0	1
0	0	1
1	1	0

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

17

## General Boolean Algebras

### ■ Operate on Bit Vectors

- Operations applied bitwise

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

- All of the Properties of Boolean Algebra Apply

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

18

## 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     $\{0, 2, 3, 4, 5, 6\}$
- ^ Symmetric difference    00111100     $\{2, 3, 4, 5\}$
- ~ Complement    10101010     $\{1, 3, 5, 7\}$

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

19

19

## Bit-Level Operations in C

### Operations &, |, ~, ^ 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)

- ~0x41 → 0xBE
- ~01000001<sub>2</sub> → 10111110<sub>2</sub>
- ~0x00 → 0xFF
- ~00000000<sub>2</sub> → 11111111<sub>2</sub>
- 0x69 & 0x55 → 0x41
- 01101001<sub>2</sub> & 01010101<sub>2</sub> → 01000001<sub>2</sub>
- 0x69 | 0x55 → 0x7D
- 01101001<sub>2</sub> | 01010101<sub>2</sub> → 01111101<sub>2</sub>

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

20

20

## Contrast: Logic Operations in C

### Contrast to Bit-level Operators

- &&, ||, !
  - View 0 as "false"
  - Anything nonzero as "true"
  - Always return 0 or 1
  - Early termination

### Examples (char data type)

- !0x41 → 0x00
- !0x00 → 0x01
- !!0x41 → 0x01
- 0x69 && 0x55 → 0x01
- 0x69 || 0x55 → 0x01
- p && \*p (avoids null pointer access)

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

21

21

## Contrast: Logic Operations in C

### Contrast to Bit-level Operators

- &&, ||, !
  - View 0 as "False"
  - Anything nonzero as "True"
  - Always return 0 or 1
  - Early termination

### Example

- !0x41 → 0x00
- !0x00 → 0x01
- !!0x41 → 0x01
- 0x69 && 0x55 → 0x01
- 0x69 || 0x55 → 0x01
- p && \*p (avoids null pointer access)

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

22

22

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

## Shift Operations

### Left Shift: $x \ll y$

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

### Right Shift: $x \gg y$

- 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

### Undefined Behavior

- Shift amount  $< 0$  or  $\geq$  word size

Argument $x$	01100010
$\ll 3$	00010000
Log. $\gg 2$	00011000
Arith. $\gg 2$	00011000

Argument $x$	10100010
$\ll 3$	00010000
Log. $\gg 2$	00101000
Arith. $\gg 2$	11101000

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

23

23

## 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

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

24

24

## Finite Precision

- the space to store integers in a computer is limited
  - forced to deal with finite precision
- different machines use a varying number of bits for its word size, from 4 to 256 bits
  - nominal size of integer and pointer data
  - 32 and 64 bits are the current preferred sizes
- in general, we can store  $2^n$  different values with  $n$  bits
  - 1 bit: 2 values (0 and 1)
  - 2 bits: 4 values (00, 01, 10, 11)
  - 4 bits: 16 values
    - we've seen 0..15, but no negative values

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

25

## Number of Values

- Address space depends on word size  $\rightarrow 2^{\text{word-size-in-bits}}$

- Is it big enough?
  - 64-bit high-end machines becoming more prevalent
  - Portability issues – insensitive to sizes of different data types

# bytes	# bits	# of values ( $2^{\text{#bits}}$ )	low	high
1	8	256		
2	16	65536		
3	24	16777216		
4	32	4294967296		
5	40	1.09951E+12		
6	48	2.81475E+14		
7	56	7.20576E+16		
8	64	1.84467E+19		
9	72	4.72237E+21		
10	80	1.20893E+24		
11	88	3.09485E+26		
12	96	7.92282E+28		
13	104	2.02824E+31		
14	112	5.1923E+33		
15	120	1.32923E+36		
16	128	3.40282E+38		

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

26

## Negative Values

- so far, we've seen the number of positive integers possible, but no negative values
- common sense tells us to split the number of bit patterns into two groups of roughly the same size: one for positive values and one for negative values
  - don't forget 0
- many ways to split these values have been developed over the years
  - two's complement is the most popular
  - unsigned represents only non-negative values (positive values and 0)

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

27

## 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
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

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

28

## Two-complement Encoding Example (Cont.)

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

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

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

29

## Two's Complement – Simple Conversion

- conversion from positive to negative using a two-step process
  - reverse the bits of the positive representation
  - add 1 to the result
- e.g.,
 

00001001	9
11110110	reverse all bits
11110111	add 1 = -9
- only one representation for 0
 

00000000	
11111111 + 1 =	00000000
- one more negative number than positive number

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

30

## Two's Complement – Alternate Conversion

- alternate conversion using a two-step process
  - reading from right to left, copy all values up to and including the first 1
  - reverse the remainder of the bits
  - e.g.,
 

00011100	28
11100100	-28
- positive numbers do not need conversion

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

31

31

## 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
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

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

32

32

## 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

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

33

33

## 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 complement integer

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

34

34

## 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

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

35

35

## Mapping Between Signed & Unsigned



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


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

36

36

## 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



37

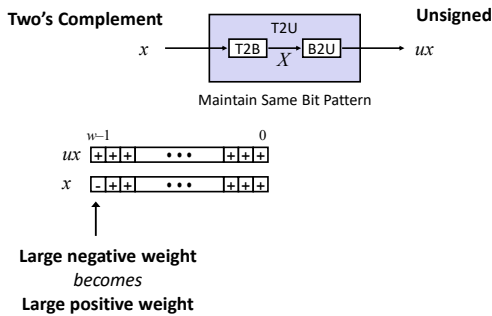
## Mapping Signed ↔ Unsigned

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



38

## Relation between Signed & Unsigned




39

## Conversion Visualized

### 2's Comp. → Unsigned

- Ordering Inversion
- Negative → Big Positive



40

## Signed vs. Unsigned in C

### Constants

- By default are considered to be signed integers
- Unsigned if have "U" as suffix  
0U, 4294967295U

### 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;
```

41

## Casting 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

42

## Summary Casting Signed $\leftrightarrow$ 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!!

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

43

43

## 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


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

44

44

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



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

45

45

## 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

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

46

46

## 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

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

47

47

## 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

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

48

48



## Unsigned Addition

Operands:  $w$  bits

$$\begin{array}{r} u \\ + v \\ \hline \end{array}$$

True Sum:  $w+1$  bits

$$u + v$$

Discard Carry:  $w$  bits

$$\text{UAdd}_w(u, v)$$

### Standard Addition Function

- Ignores carry output

### Implements Modular Arithmetic

$$s = \text{UAdd}_w(u, v) = u + v \bmod 2^w$$


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

49

## Visualizing (Mathematical) Integer Addition

### Integer Addition

- 4-bit integers  $u, v$
- Compute true sum  $\text{Add}_4(u, v)$
- Values increase linearly with  $u$  and  $v$
- Forms planar surface





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

50

## Visualizing Unsigned Addition

### Wraps Around

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



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

51

## Two's Complement Addition

Operands:  $w$  bits

$$\begin{array}{r} u \\ + v \\ \hline \end{array}$$

True Sum:  $w+1$  bits

$$u + v$$

Discard Carry:  $w$  bits

$$\text{TAdd}_w(u, v)$$

### 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  $s == t$


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

52

## TAdd Overflow

### Functionality

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



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

53

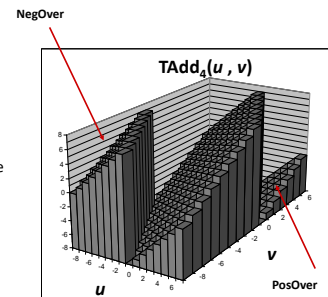
## Visualizing 2's Complement Addition

### Values

- 4-bit two's comp.
- Range from -8 to +7

### Wraps Around

- If sum  $\geq 2^{w-1}$ 
  - Becomes negative
  - At most once
- If sum  $< -2^{w-1}$ 
  - Becomes positive
  - At most once



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

54

## Binary Addition

- Goal: compute sum of  $w$ -bit numbers  $x, y$ 
  - Either signed or unsigned
- Binary Addition Basics

$$\begin{array}{r} 0 \\ + 0 \\ \hline 0 \end{array} \quad \begin{array}{r} 1 \\ + 0 \\ \hline 1 \end{array} \quad \begin{array}{r} 0 \\ + 1 \\ \hline 1 \end{array} \quad \begin{array}{r} 1 \\ + 1 \\ \hline 10 \end{array}$$

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

55

## Binary Addition

- Examples

Problem in base ten		Problem in two's complement		Answer in base ten
$\begin{array}{r} 3 \\ + 2 \\ \hline \end{array}$	→	$\begin{array}{r} 0011 \\ + 0010 \\ \hline 0101 \end{array}$	→	5
$\begin{array}{r} -3 \\ + -2 \\ \hline \end{array}$	→	$\begin{array}{r} 1101 \\ + 1110 \\ \hline 1011 \end{array}$	→	-5
$\begin{array}{r} 7 \\ + -5 \\ \hline \end{array}$	→	$\begin{array}{r} 0111 \\ + 1011 \\ \hline 0010 \end{array}$	→	2

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

56

## Binary Addition

- Using 6 bits we can represent values from -32 to 31, so what happens when we try to add 19 plus 14 or -19 and -14

$$\begin{array}{r} 19 \\ + 14 \\ \hline 33 \end{array}$$

$$\begin{array}{r} 010011 \\ + 001110 \\ \hline 100001 \end{array}$$

we have added two positive numbers and gotten a negative result – this is positive overflow

$$\begin{array}{r} -19 \\ - 14 \\ \hline -33 \end{array}$$

$$\begin{array}{r} 101101 \\ + 110010 \\ \hline 011111 \end{array}$$

we have added two negative numbers and gotten a positive result – this is negative overflow

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

57

## Binary Addition

- 8-bit binary addition

$$\begin{array}{r} 0101\ 0101 \\ + 0001\ 1001 \\ \hline 0110\ 1110 \end{array}$$

no left bit discarded  
no overflow

$$\begin{array}{r} 1111\ 0101 \\ + 1101\ 1011 \\ \hline 1\ 1101\ 0000 \end{array}$$

left bit discarded  
no overflow

$$\begin{array}{r} 0111\ 0101 \\ + 0101\ 1011 \\ \hline 1101\ 0000 \end{array}$$

no left bit discarded  
positive overflow

$$\begin{array}{r} 1111\ 0101 \\ + 1000\ 0011 \\ \hline 1\ 0111\ 1000 \end{array}$$

left bit discarded  
negative overflow

$$\begin{array}{r} 0111\ 0101 \\ + 1101\ 1011 \\ \hline 1\ 0101\ 0000 \end{array}$$

left bit discarded  
no overflow

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

58


## 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

59

## Unsigned Multiplication in C




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

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

60

## 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

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

61

## Binary Multiplication

### 8-bit binary multiplication

```

01010101
x 00011001
01010101
01010101
100001001101

```

truncated: 0100 1101


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

62

## 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
  - Compiler generates this code automatically


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

63

## 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

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

64

## Shifting and Overflow

### since an arithmetic left shift is the same as multiplying by 2, we may run out of space, resulting in overflow

- ex., 8-bit unsigned:  $0010\ 1110 \ll 3 = 0111\ 0000$  ( $46 * 8 = 368$ , not 112)
- ex., 8-bit signed:  $0010\ 1110 \ll 2 = 1011\ 1000$  ( $46 * 4 = 184$ , not -72)
- ex., 8-bit signed:  $1110\ 1110 \ll 3 = 0111\ 0000$  ( $-18 * 8 = -144$ , not 112)
- ex., 8-bit signed:  $1110\ 1110 \ll 2 = 1011\ 1000$  ( $-18 * 4 = -72$  OK)

### overflow limitations

- not valid with logical shifts
- not possible using right shifts
- determined by 1 bits shifting off left
  - if 1 bits used for sign extension, no overflow unless sign change
- can also occur by 0 bits shifting off left
  - change sign of result

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

65

## 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

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

66

## 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

67

## 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

68

## 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

69

## 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

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

70

## 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

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

71

## 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

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

72

## 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

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

73

## 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)



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

74

## 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

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

75

## 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

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

76

## Byte Ordering Example

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




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

77


## 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

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

78

## 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

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

79

## 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
```

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

80

## Representing Pointers

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

Sun	IA32	x86-64
EF	AC	3C
FF	28	1B
FB	F5	FE
2C	FF	82
		FD
		7F
		00
		00

Different compilers & machines assign different locations to objects

Even get different results each time run program

81

## 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

IA32	Sun
31	31
38	38
32	32
31	31
33	33
00	00

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

82