

Bits, Bytes, and Integers


CSci 2021: Machine Architecture and Organization
Lectures #2-4, January 23rd-28th, 2015

Your instructor: Stephen McCamant

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

1

Binary Representations



2

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

3

Byte-Oriented Memory Organization



Programs Refer to Virtual Addresses

- Conceptually very large array of bytes
- Actually implemented with hierarchy of different memory types
- System provides address space private to particular "process"
 - Program being executed
 - Program can clobber its own data, but not that of others

Compiler + Run-Time System Control Allocation

- Where different program objects should be stored
- All allocation within single virtual address space

4

Machine Words

Machine Has "Word Size"


- Nominal size of integer-valued data
 - Including addresses
- Most current machines use 32 bits (4 bytes) words
 - Limits addresses to 4GB
 - Becoming too small for memory-intensive applications
- High-end systems use 64 bits (8 bytes) words
 - Potential address space $\approx 1.8 \times 10^{19}$ bytes
 - x86-64 machines support 48-bit addresses: 256 Terabytes
- Machines support multiple data formats
 - Fractions or multiples of word size
 - Always integral number of bytes

5

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)



6

Data Representations

C Data Type	Typical 32-bit	Intel IA32	x86-64
char	1	1	1
short	2	2	2
int	4	4	4
long	4	4	8
long long	8	8	8
float	4	4	4
double	8	8	8
long double	8	10/12	10/16
pointer	4	4	8

7

Byte Ordering


- How should bytes within a multi-byte word be ordered in memory?
- Conventions
 - Big Endian: Sun, PPC Mac, Internet convention
 - Least significant byte has highest address
 - Little Endian: x86, VAX
 - Least significant byte has lowest address

8


Byte Ordering Example

- Big Endian**
 - Least significant byte has highest address
- Little Endian**
 - Least significant byte has lowest address
- Example**
 - Variable x has 4-byte representation 0x01234567
 - Address given by &x is 0x100

Big Endian



Little Endian



9

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
804836e:	83 bb 28 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

10

Examining Data Representations

- Code to Print Byte Representation of Data**
 - Casting pointer to unsigned char * creates byte array

```
void show_bytes(unsigned char *start, int len) {
    int 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

11

show_bytes Execution Example

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

Result (Linux):


```
int a = 15213;
0xbffffcb8 0x6d
0xbffffcb9 0x3b
0xbffffcba 0x00
0xbffffcbb 0x00
```

12


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
 (Covered later)

13

Representing Pointers

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



Different compilers & machines assign different locations to objects


14

Representing Strings

```
char S[6] = "18243";
```

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

15

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	lt	ln	^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

16

Today: Bits, Bytes, and Integers

- Representing information as bits
- (Logistics interlude)
- Bit-level manipulations
- Integers
 - Representation: unsigned and signed
 - Conversion, casting
 - Expanding, truncating
 - Addition, negation, multiplication, shifting
- Summary

17

Homework turn-in process

- For full credit: turn in at the beginning of class on the due date
 - On-time = 3:35pm, or when I start lecturing, whichever is later
 - Yes, this means you have to come to class on time (that day)
- We strongly recommend typing your assignments on a computer, not hand-writing
- Late submissions only will be online using the Moodle
- Do not turn in paper assignments at other times
 - This helps us stay organized

18

2021-dedicated VMs now available

- SSH into: `xA-B.cselabs.umn.edu`
 - Where A is 21, 22, or 23
 - And B is 01, 02, 03, 04, or 05
 - E.g., `x22-02.cselabs.umn.edu`
- 32-bit version of Ubuntu Linux version 14.04
- Do not run graphical programs (Firefox, etc.) on these machines (it would be slow anyway)
- If you prefer to use other CSE Labs Linux machines, give the `-m32` option to GCC to get 32-bit binaries

19

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

Not (math: \neg)

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

\sim	0	1
0	1	
1	0	

Or (math: \vee)

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

$ $	0	1
0	0	1
1	1	1

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


- $A^*B = 1$ when either $A=1$ or $B=1$, but not both

\wedge	0	1
0	0	1
1	1	0

20

Application of Boolean Algebra

- Applied to Digital Systems by Claude Shannon
 - 1937 MIT Master's Thesis
 - Reason about networks of relay switches
 - Encode closed switch as 1, open switch as 0



Connection when

$$A \& \sim B \mid \sim A \& B$$

$$= A^*B$$

21

General Boolean Algebras

Operate on Bit Vectors

- Operations applied bitwise

$$\begin{array}{cccc} 01101001 & 01101001 & 01101001 & 01101001 \\ \& 01010101 & | 01010101 & \wedge 01010101 & \sim 01010101 \\ \hline 01000001 & 01111101 & 00111100 & 10101010 & \end{array}$$

All of the Properties of Boolean Algebra Apply

22

Representing & Manipulating Sets

Representation

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

$$\begin{array}{ll} 01101001 & \{0, 3, 5, 6\} \\ 76543210 & \end{array}$$

$$\begin{array}{ll} 01010101 & \{0, 2, 4, 6\} \\ 76543210 & \end{array}$$

Operations

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

23

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)

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

24

Contrast: Logic Operations in C

- **Contrast to Logical Operators**
 - &&, ||, !
 - View 0 as "False"
 - Anything nonzero as "True"
 - Always return 0 or 1
 - Early termination (AKA "short-circuit evaluation")
- **Examples (char data type)**
 - !0x41 → 0x00
 - !0x00 → 0x01
 - !!0x41 → 0x01
 - 0x69 && 0x55 → 0x01
 - 0x69 || 0x55 → 0x01
 - p && *p (avoids null pointer access)

25

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
- **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 right
- **Undefined Behavior**
 - Shift amount < 0 or ≥ size

Argument x	01100010
<< 3	00010000
Log. >> 2	00011000
Arith. >> 2	00011000

Argument x	10100010
<< 3	00010000
Log. >> 2	00101000
Arith. >> 2	11101000

26

Exercise break: flip case

```

/* Convert lowercase to uppercase and vice-versa,
   return any other characters unchanged */
char flip_case(char c) {
    if (c >= 'A' && c <= 'Z') {
        /* 0x41 through 0x5A */
        return c | 0x20;
    } else if (c >= 'a' && c <= 'z') {
        /* 0x61 through 0x7A */
        return c & ~0x20;
    } else {
        return c;
    }
}

```

- Fill in the blanks, using bitwise operators

27

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

28

Encoding Integers

Unsigned

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

Two's Complement

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

```

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

29

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
Sum	15213	-15213

30

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

31

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 platform specific

32

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

33

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


34

Announcement interlude: Lab 1 out

- Lab 0 (Hello, world) is due tonight
- Lab 1 on data representation is out
- Basic idea: puzzles implementing operations with other operations
 - E.g., implement logical right shift using only arithmetic right shift
- Most problems relate to bitwise operations and two's complement rules
 - I.e., you can start working on them now
- Increasing difficulty, try the easier ones first
- Two questions relating to floating point

35

Mapping Between Signed & Unsigned



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

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



Large negative weight becomes Large positive weight


$$ux = \begin{cases} x & x \geq 0 \\ x + 2^w & x < 0 \end{cases}$$

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, 429496729U`

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 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

42

Code Security Example

```

/* 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;
}

```

- Similar to code found in FreeBSD's implementation of `getpeername`
- There are legions of smart people trying to find vulnerabilities in programs

43

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);
}

```

44

Malicious Usage

```

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

```

```

/* 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);
    ...
}

```

45

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!!

46


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

47

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



48

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

49

Summary: Expanding, Truncating: Basic Rules

- Expanding (e.g., short int to int)
 - Unsigned: zeros added ("zero extension")
 - 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 behaviour

50

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

51

Negation: Complement & Increment

- Claim: Following Holds for 2's Complement

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

- Complement

- Observation: $\sim x + x == 1111\dots111 == -1$

$$\begin{array}{r} x \quad 100111101 \\ + \sim x \quad 01100010 \\ \hline -1 \quad 11111111 \end{array}$$

- Where would we fill in gaps for a more complete proof?
- Note: operation can apply to unsigned as well
- Two values for which x and $-x$ have the same sign

52

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
~ 0	-1	FF FF	11111111 11111111
$\sim 0 + 1$	0	00 00	00000000 00000000

53

Unsigned Addition

Operands: w bits

$$\begin{array}{r} u \quad \square \square \square \dots \square \square \\ + v \quad \square \square \square \dots \square \square \\ \hline \text{True Sum: } w+1 \text{ bits} \\ u + v \quad \square \square \square \dots \square \square \square \\ \hline \text{Discard Carry: } w \text{ bits} \quad UAdd_w(u, v) \quad \square \square \square \dots \square \square \end{array}$$

- Standard Addition Function
 - Ignores carry output
- Implements Modular Arithmetic


$$s = UAdd_w(u, v) = u + v \text{ mod } 2^w$$

$$UAdd_w(u, v) = \begin{cases} u + v & u + v < 2^w \\ u + v - 2^w & u + v \geq 2^w \end{cases}$$

54

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



55

Visualizing Unsigned Addition

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



56

Mathematical Properties

- Modular Addition Forms an Abelian Group
 - Closed under addition
 - $0 \leq UAdd_w(u, v) \leq 2^w - 1$
 - Commutative
 - $UAdd_w(u, v) = UAdd_w(v, u)$
 - Associative
 - $UAdd_w(t, UAdd_w(u, v)) = UAdd_w(UAdd_w(t, u), v)$
 - 0 is additive identity
 - $UAdd_w(u, 0) = u$
 - Every element has additive inverse
 - Let $UComp_w(u) = 2^w - u$
 - $UAdd_w(u, UComp_w(u)) = 0$


57

Two's Complement Addition


Operands: w bits



True Sum: $w+1$ bits



Discard Carry: w bits




- 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 $s == t$

58

TAdd Overflow

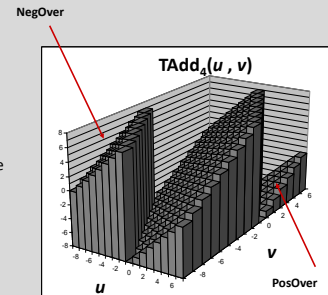
- Functionality
 - True sum requires $w+1$ bits
 - Drop off MSB
 - Treat remaining bits as 2's comp. integer



59

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




60

Characterizing TAdd

Functionality

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



$$TAdd_w(u, v) = \begin{cases} u + v + 2^w & u + v < TMin_w \text{ (NegOver)} \\ u + v & TMin_w \leq u + v \leq TMax_w \\ u + v - 2^w & TMax_w < u + v \text{ (PosOver)} \end{cases}$$

61

Mathematical Properties of TAdd

Isomorphic Group to unsigneds with UAdd

- $TAdd_w(u, v) = U2T(UAdd_w(T2U(u), T2U(v)))$
 - 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

$$TComp_w(u) = \begin{cases} -u & u \neq TMin_w \\ TMin_w & u = TMin_w \end{cases}$$

62

Exercise break: ten's complement

- Before digital computers, there were mechanical computers that used base 10
- There's an analog of two's complement called ten's complement that works in decimal
- Suppose we have an adding machine with 10 decimal digits, 10^4 instead of 2^{32} .
- What should be the ten's complement representation of -21?
- I.e., we want a number x so that adding x is the same as subtracting 21, when you only have 4 digits

63

Ten's complement answer

- We want $x \equiv -21 \pmod{10000}$, or $x + 21 + 10000k = 0$ for integer k
- $x = 10000 - 21 = 9979$
- (The equivalent of \sim is called nines' complement: $\sim 21 = 9978$)

64


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)



65

Multiplication

Computing Exact Product of w -bit numbers x, y

- Either signed or unsigned

Ranges


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

Maintaining Exact Results

- Would need to keep expanding word size with each product computed
- Done in software by "arbitrary precision" arithmetic packages

66

Unsigned Multiplication in C




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

67

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);
```



68

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;
}
```

69


XDR Vulnerability

`malloc(ele_cnt * ele_size)`

- What if:
 - $\text{ele_cnt} = 2^{20} + 1$
 - $\text{ele_size} = 4096 = 2^{12}$
 - Allocation = ??
- How can I make this function secure?

70

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

71

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

72

Compiled Multiplication Code

C Function

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

Compiled Arithmetic Operations

```
leal (%eax,%eax,2), %eax
sall $2, %eax
```

Explanation

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

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

73

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

74

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

75

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 >> 1	7606.5	7606	1D B6	00011101 10110110
x >> 4	950.8125	950	03 B6	00000011 10110110
x >> 8	59.4257813	59	00 3B	00000000 00111011

76

Compiled Unsigned Division Code

C Function

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

Compiled Arithmetic Operations

```
shrl $3, %eax
```

Explanation

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


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

77

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 >> 1	-7606.5	-7607	E2 49	11100010 01001001
y >> 4	-950.8125	-951	FC 49	11111100 01001001
y >> 8	-59.4257813	-60	FF C4	11111111 11000100


78

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

79

Correct Power-of-2 Divide (Cont.)

Case 2: Rounding



Biasing adds 1 to final result

80

Compiled Signed Division Code

C Function

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

Compiled Arithmetic Operations

```
testl %eax, %eax
js    L4
L3:  sarl $3, %eax
     ret
L4:  addl $7, %eax
     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

81

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 $/$'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)

82

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)

83

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

84

Today: Integers

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

85

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
 - $\text{UMult}_w(t, \text{UAdd}_w(u, v)) = \text{UAdd}_w(\text{UMult}_w(t, u), \text{UMult}_w(t, v))$

86

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
 - $\text{TMax} + 1 == \text{TMin}$
 - $15213 * 30426 == -10030$ (16-bit words)

87

Why Should I Use Unsigned?

- **Don't Use Just Because Number Nonnegative**
 - 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)
    . . .
```
- **Do Use When Performing Modular Arithmetic**
 - E.g., used in multiprecision arithmetic
- **Do Use When Using Bits to Represent Sets**
 - Logical right shift, no sign extension

88

Integer C Puzzles

- ```
int x = foo();
int y = bar();
unsigned ux = x;
unsigned uy = y;
```
- $x < 0 \Rightarrow ((x^2) < 0)$
  - $ux \geq 0$
  - $x \& 7 == 7 \Rightarrow (x < 30) < 0$
  - $ux > -1$
  - $x > y \Rightarrow -x < -y$
  - $x * x \geq 0$
  - $x > 0 \&\& y > 0 \Rightarrow x + y > 0$
  - $x \geq 0 \Rightarrow -x \leq 0$
  - $x \leq 0 \Rightarrow -x \geq 0$
  - $(x|-x) >> 31 == -1$
  - $ux >> 3 == ux/8$
  - $x >> 3 == x/8$
  - $x \& (x-1) != 0$

89