

```
29 #include <balloon.h>
30
31 int struct {
32     ScePspFVector3 mode;
33     ScePspFVector3 pos;
34     int sbuf[3];
35     float scnt;
36     t;
37 } BALLOONDAT;
38
39 static BALLOONDAT balloon;
40 static ScePspFVector3 sphere[28];
41 static ScePspFVector3 pole[28];
42
43 extern void DrawSphere(ScePspFVector3 *array, float r);
44 extern void DrawPole(ScePspFVector3 *array, float r);
45
46 void init_balloon(void)
47 {
48     int i;
49
50     balloon.mode=MODE_IN;
51     balloon.pos.x= 0.0f;
52     balloon.pos.y=-8.0f;
53     balloon.pos.z= 0.0f;
54     balloon.t=0.0f;
55     balloon.scnt=2;
56
57     for (i=0; i<3; i++)
58         balloon.sbuf[i]=24*RAM;
59
60 }
61
62 void draw_balloon(void)
63 {
64     ScePspFVector3 vec;
65     glEnable(SCEGU_TEXTURE);
66     glTranslatef(balloon.pos);
67 }
```

# Operating Systems and C Fall 2022 2. Representing Data

```
29 #include "balloon.h"
30
31 int struct {
32     ScePspFVector3 mode;
33     ScePspFVector3 pos;
34     int sbuf[3];
35     float scnt;
36     t;
37 } BALLOONDAT;
38
39 static BALLOONDAT balloon;
40 static ScePspFVector3 sphere[28];
41 static ScePspFVector3 pole[28];
42
43 extern void DrawSphere(ScePspFVector3 *array, float r);
44 extern void DrawPole(ScePspFVector3 *array, float r);
45
46 void init_balloon(void)
47 {
48     int i;
49
50     balloon.mode=MODE_IN;
51     balloon.pos.x= 0.0f;
52     balloon.pos.y=-8.0f;
53     balloon.pos.z= 0.0f;
54     balloon.t=0.0f;
55     balloon.scnt=2;
56
57     for (i=0; i<3; i++)
58         balloon.sbuf[i]=24+RAM;
59
60 }
61
62 void draw_balloon(void)
63 {
64     ScePspFVector3 vec;
65     glTranslatef(SCEGU_TEXTURE);
66     glTranslatef(balloon.pos);
67 }
```

# Operating Systems and C

## Fall 2022

### 2. Representing Data

why do we bother looking into this?

# Motivation: Performance

Data Representation is crucial for data science, e.g.,  
How to represent a binary matrix? (bits  $\Rightarrow$  ops bit-level)

```
#include <math.h>
#include <stdint.h>
#include <malloc.h>
#include <stdio.h>

#define N 8
#define M 8
#define n 0
#define m 1

int main(int argc, char *argv[])
{
    /* binary matrix allocation */
    uint64_t *matrix;
    matrix = (uint64_t *) malloc((size_t) ceil((N*M)/sizeof(uint64_t)));
    *matrix = 0x1001b;
    /* reading the [n][m]-th element of an NxM binary matrix */
    uint64_t mask = 1U << ((n*N+m)%64U);
    uint64_t result = mask & matrix[(size_t) ceil((n*N+m)/(64*1.0))];
    result = result >> ((n*N+m)%64U);
    printf("matrix[n][m]: %lu\n", result);
    return 0;
}
```

# Motivation: Security

l2ex1.c+

```
1 /* Kernel memory region holding user-accessible data */
2 #define KSIZE 1024
3 char kbuf[KSIZE];
4
5 /* Copy at most maxlen bytes from kernel region to user buffer */
6 int copy_from_kernel(void *user_dest, int maxlen) {
7     /* Byte count len is minimum of buffer size and maxlen */
8     int len = KSIZE < maxlen ? KSIZE : maxlen;
9     memcpy(user_dest, kbuf, len);
10    return len;
11 }
```

expects type `size_t` (unsigned),  
given an `int`. what happens? problem?

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

we need a deeper understand of how data is represented.

- 1. Two's complement**
2. Integer Arithmetic
3. Bit Manipulation

# Two issues

## 1. Finite representation

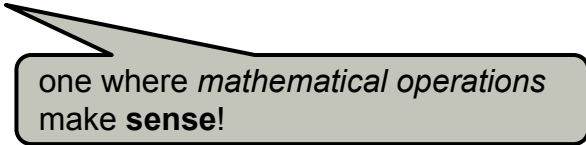
- There is a limit to the number of integers (put differently, a max value) that can be represented on a fixed number of bytes

## 2. Representing positive and negative integers

- Sign and values must be represented


More than one way to skin a cat.

What's a **sensible** way?



one where *mathematical operations* make **sense!**

# Spoiler: Algebra: Ring (Finite Field) - Unsigned

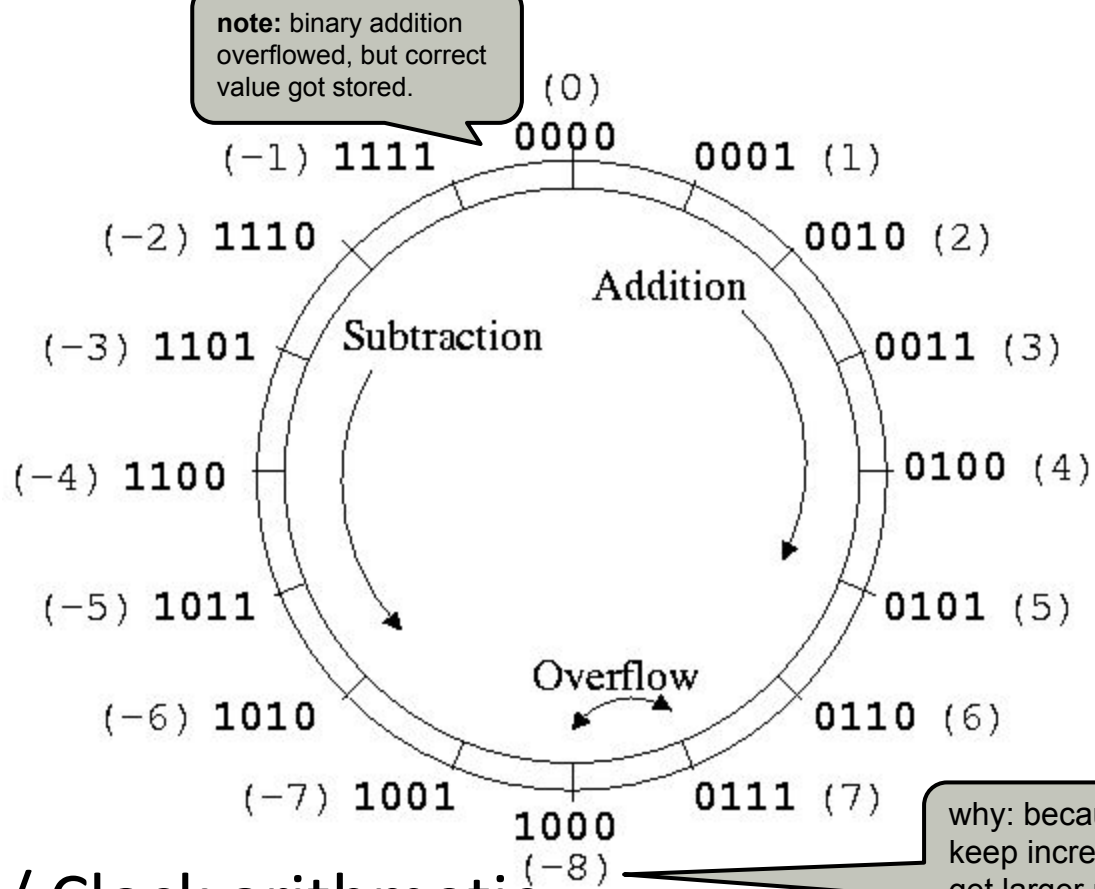


parentheses:  
how we *interpret*  
the bits.

Modular / Clock arithmetic.<sup>(8)</sup>

<https://stackoverflow.com/questions/7221409/is-unsigned-integer-subtraction-defined-behavior>

# Spoiler: Algebra: Ring (Finite Field) - Signed



note: binary addition overflowed, but correct value got stored.

parentheses: how we *interpret* the bits.

why: because then, if you keep incrementing, you get larger numbers.

note: binary addition of 3 data-bits overflowed, thus overwriting the sign bit

subtraction-defined-behavior

<https://dnscsite.wordpress.com/tag/twos-complement/>

## Modular / Clock arithmetic.

<https://stackoverflow.com/questions/722140>



# Integer Types in C

- Signed vs. unsigned
- Number of bytes
  - 1B: char
  - 2B: short
  - 4B: int
  - 4B (32 bits) or 8B (64 bits): long
  - 8B: long long

# Sequences of bits

hex: compact  
representation of bits.

Shorts in C are coded on 2B. Used for examples in these slides.

2B = 16 bits,  
or 4 Hexadecimals (prefixed with 0x)

see book on

dec-bin, bin-dec,  
dec-hex, hex-dec,  
bin-hex, hex-bin

or:

[onlinetoolz.net](http://onlinetoolz.net)

Operations on sequences of bits:

- **bitwise** operations:            and (&), or (|), not (~), xor (^), shift (<<, >>)
- (interpreted as Booleans) 0 is false; anything but 0 is true
- **logical** operations:            and (&&), or (||), not (!)
- (interpreted as integers)
- **arithmetic** operations:        +, -, \*, /

# Bitwise Operations: Examples

## Bitwise on cos

<https://github.com/mellowcandle/bitwise>

- $0x0000+1 = 0x0001$ ;  $0x000E+1 = 0x000F$ ;  
 $0x000F+1 = 0x0010$
- $1 \ll 15 = 0x8000$ ;  $1 \ll 8 = 0x0100$ ;  
 $1 \ll 3 = 0x0008$
- $0xEFFF+1 = 0xF000$ ;  $0x24BC+1 = 0x24BD$
- $\sim 0x0000 = 0xFFFF$ ;  $\sim 0x0001 = 0xFFFE$
- $X+\sim X = 0xFFFF$  ;  $X\&\sim X = 0$

# Signed vs. Unsigned Interpretation

A **sequence of bits** is interpreted differently depending on whether the integer type is unsigned or signed.

how much each bit contributes to value

Unsigned

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

40960

0xA000

-x is *not* just bits of x w/ a 1 in sign

Signed (**two's complement**)

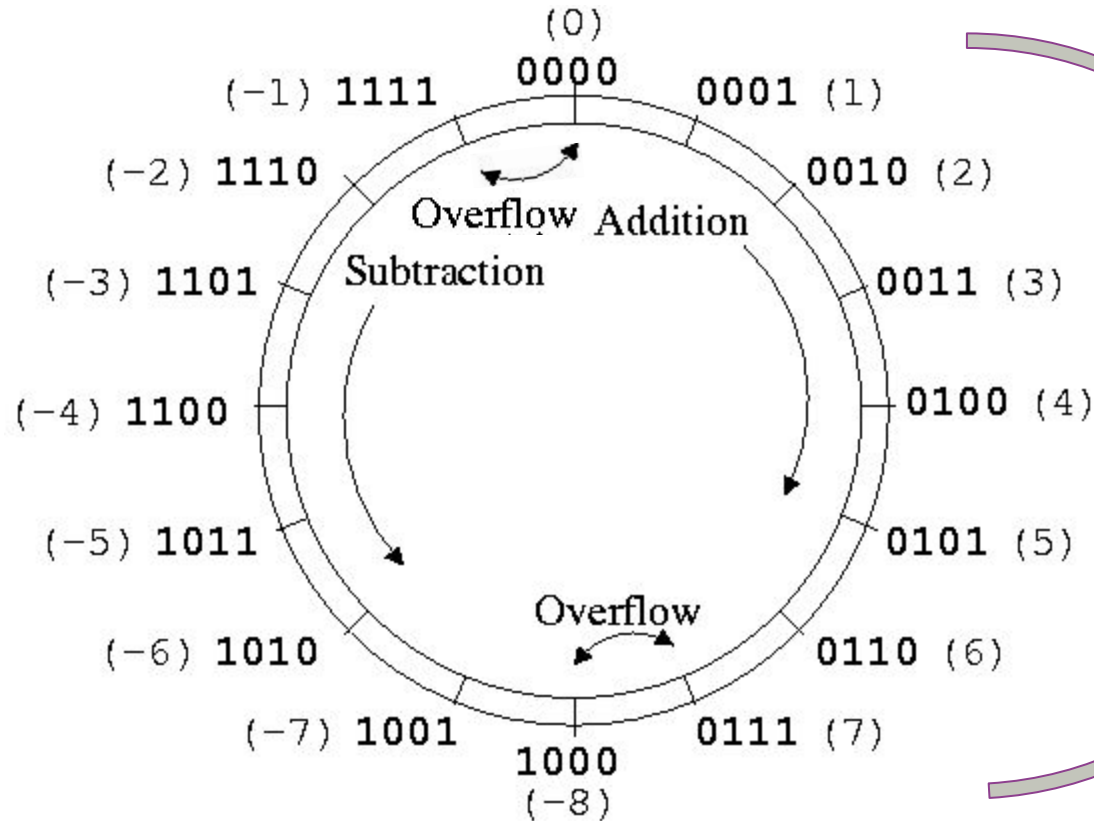
$$B2T(X) = -x_{w-1} \cdot 2^{w-1} + \sum_{i=0}^{w-2} x_i \cdot 2^i$$

-24576

Sign Bit

same bits; different interpretations as numbers.

# Identical Interpretation?



book has interesting results on when bits have identical unsigned & signed interpretation. (outside range, conversion either adds or subtracts  $2^w$ )

# Integer Types in C

\$ vi /usr/include/limits.h

```
/* Minimum and maximum values a 'signed short int' can hold. */  
# define SHRT_MIN      (-32768)  
# define SHRT_MAX      32767  
  
/* Maximum value an 'unsigned short int' can hold. (Minimum is 0.) */  
# define USHRT_MAX     65535
```

0x0000  $B2U_{min} = 0$

0xFFFF  $B2U_{max} = \sum_{i=0}^{15} 2^i$   
 $= 2^{16}-1$

0x8000  $B2T_{min} = -2^{15}$

0x7FFF  $B2T_{max} = \sum_{i=0}^{14} 2^i$   
 $= 2^{15}-1$

1 bit reserved  
for sign

there's 1 more  
neg than pos

# Signed Integers

given 0s and 1s,  
how do I interpret it?

Two's Complement. Representation on N bytes

- Positive numbers:
  - Sign bit is 0
  - Binary representation of value on  $(8*N)-1$  bits
- Negative numbers:
  - Binary representation of corresponding positive value on  $8*N$  bits
  - Invert all digits (0 becomes 1; 1 becomes 0)
  - Add one

o.O why?

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

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

yup, that checks out. but why?

# Why this works?

## EXAMPLE WITH 8 bits

Subtracting from 00000000 or subtracting from 100000000 is the same for all practical purpose (as the borrowing is carried forward beyond the 8th bit). So, in 2's complement  $-x$  is represented by

$$2^8 - x = 100000000 - x = 11111111 + 1 - x = (11111111 - x) + 1$$

As we have seen (now with 8 bits):

$$\sim x + x = 11111111 \Rightarrow (11111111 - x) = \sim x$$

recall:  $\sim x$  is  $x$  with bits flipped

In two's complement on  $w$  bits

$-x$  is represented as  $2^w - x$

$$\Rightarrow \sim x + 1$$

in other words:  $-x$  is  
“**what must I add to  $x$  to get 0?**”  
(answer:  $\sim x + 1$ )



# Two's complement

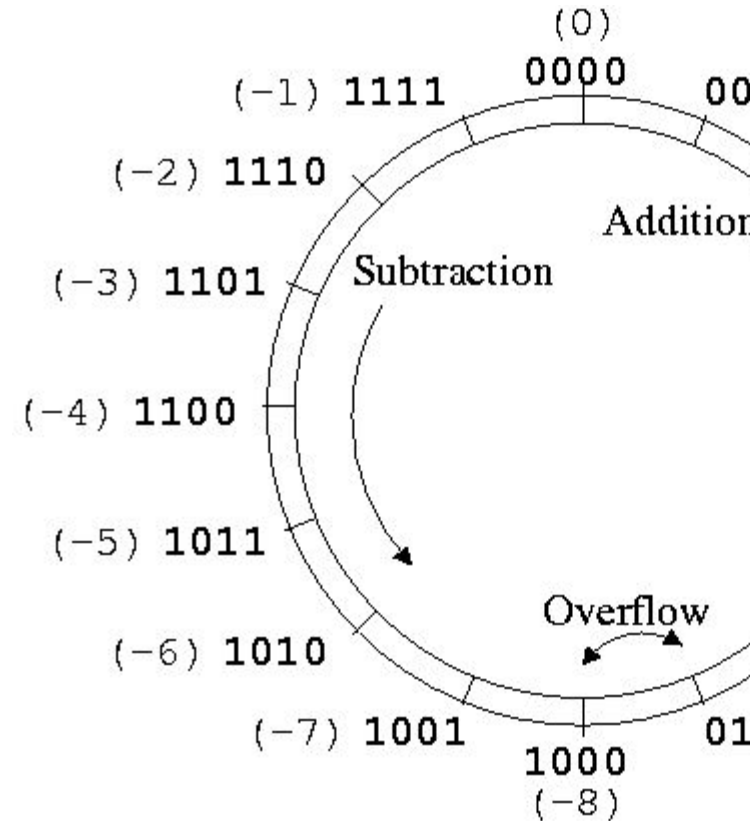
What does 0xFFFFFFFF represent?

How about 0x80000000?

# Two's complement



What does `0xFFFFFFFF` represent?

How about `0x80000000`?

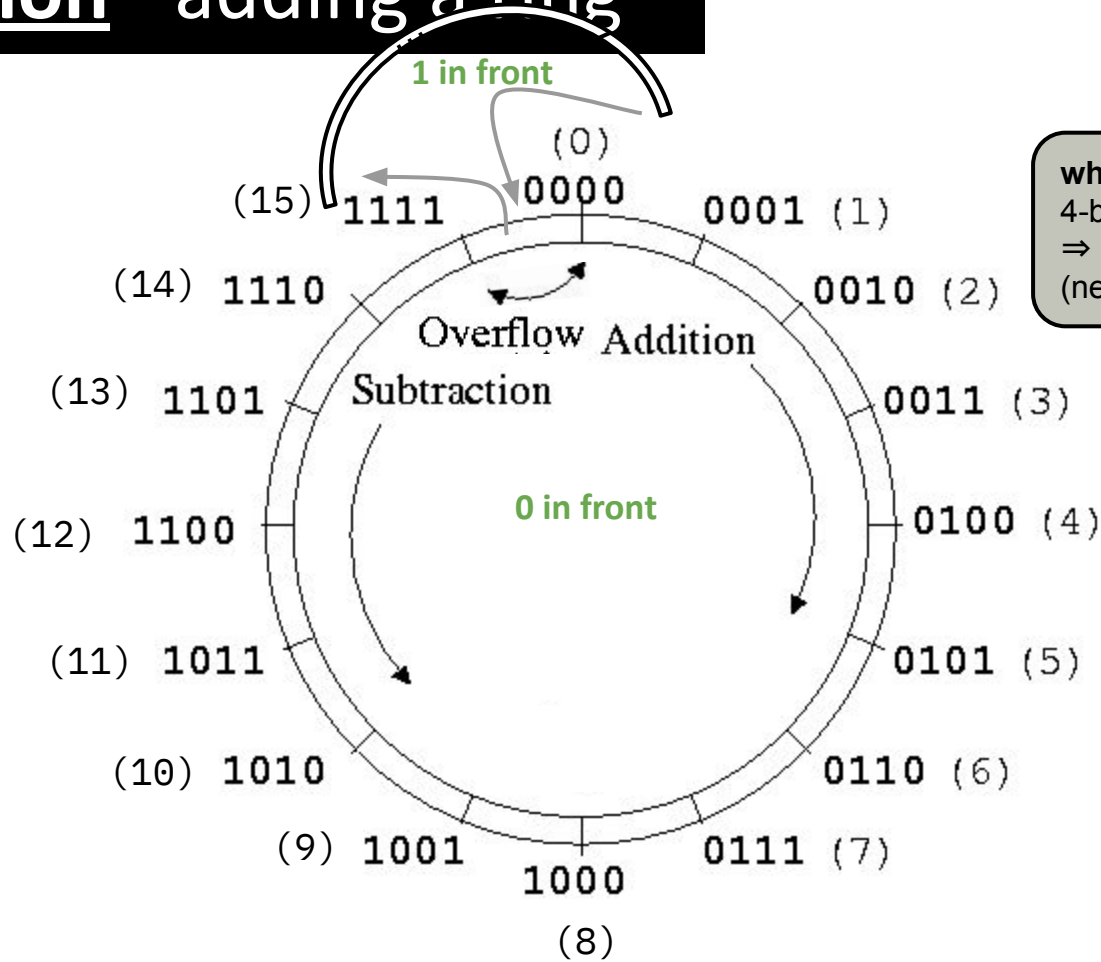


hint

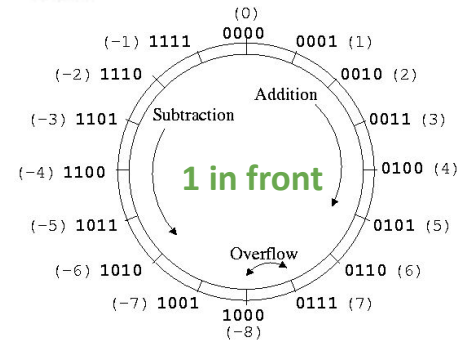

# Bit extension - adding a ring



# Bit extension - adding a ring

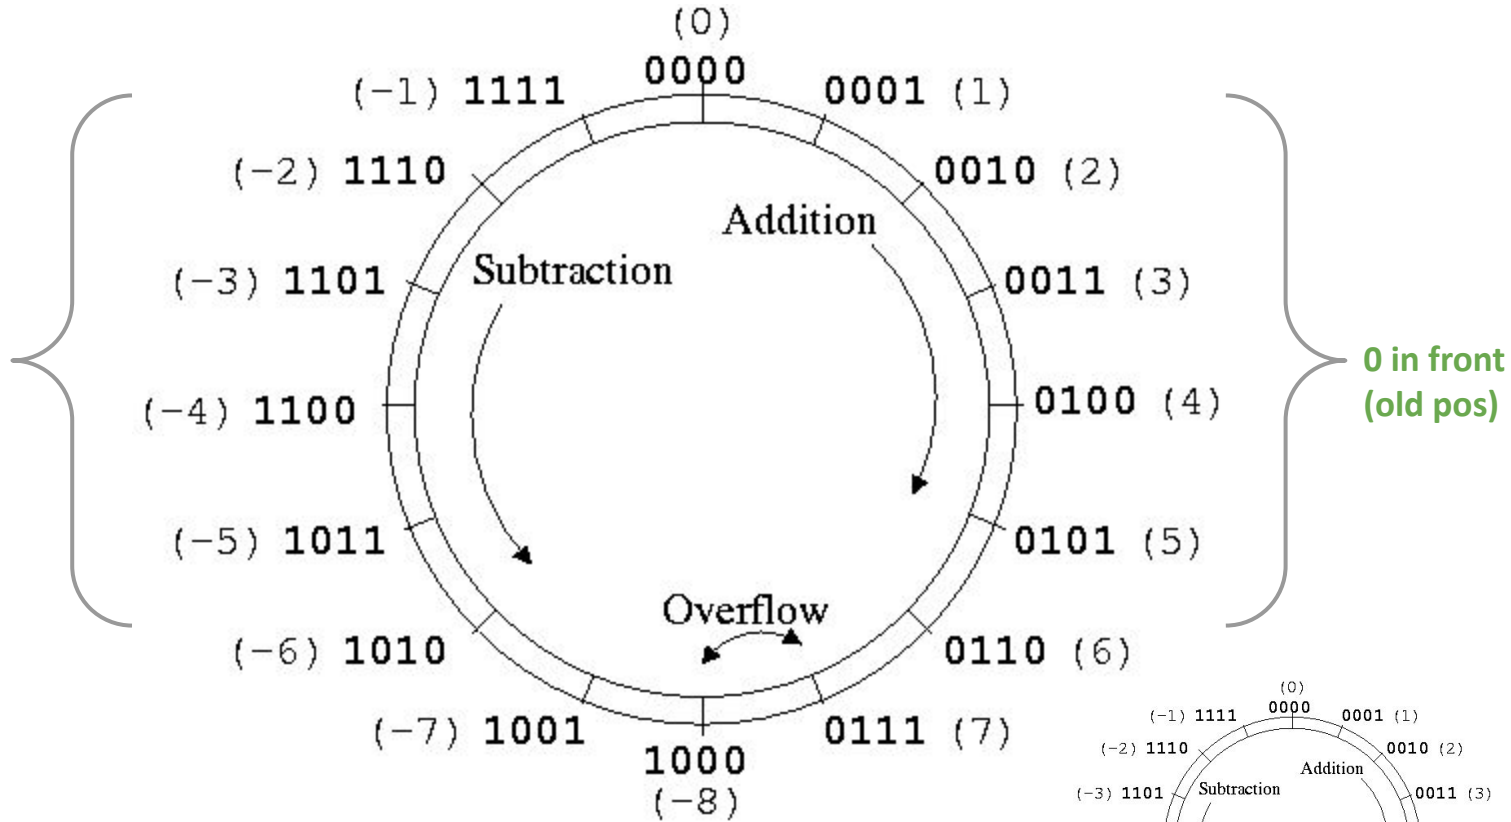


# Sign extension - adding a ring (bit extension, signed)




# Sign extension - adding a ring (bit extension, signed)

1 in front  
(old neg)



0 in front  
(old pos)


0 in front  
(new pos)



1 in front  
(new neg)

# Sign extension - adding a ring (bit extension, signed)

1 in front  
(old neg)



why this way: storing 4-bit value in 5-bit space  
⇒ old bits are unchanged!  
(5th bit is 0 if pos, 1 if neg)

0 in front  
(old pos)

this helps explain the following slide.  
(analogy: arithmetic right shift)

new pos

new neg

even more positive, even more negative

# Sign Extension


Expanding: Converting from smaller to larger integer data type

Make  $k$  copies of sign bit:

$$X' = \underbrace{x_{w-1}, \dots, x_{w-1}}_{k \text{ copies of Most Significant Bit}}, x_{w-2}, \dots, x_0$$

$k$  copies of Most Significant Bit

Most Significant Bit



the overall value of negative numbers does not change.



# Sign Extension

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

easy in C!  
(one of the nice properties of  
two's complement)

	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

C automatically performs sign extension

# Sign Extension

Expanding (e.g., short int to int)

Unsigned: zeros added (on left)

Signed: sign extension (as shown on previous slides)

Both yield expected result

Truncating (e.g., unsigned to unsigned short)

Unsigned/signed: bits are truncated

Result reinterpreted

Unsigned: mod operation

Signed: depends on bit pattern (large negative number might be truncated to positive number)

# Why do we care?

In C type system,

If there is a mix of signed and unsigned values in an expression, then signed values are **implicitly cast to unsigned**:

- The bit pattern is maintained
- But re-interpreted!!
- Can have unexpected effects => adding or subtracting  $2^N$

# Security, Revisited

```
12
13 #define MSIZE 528
14
15 void getstuff() {
16     char mybuf[MSIZE];
17     copy_from_kernel(mybuf, -MSIZE);
18     . . .
19 }
```

l2ex1.c+

```
1 /* Kernel memory region holding user-accessible data */
2 #define KSIZE 1024
3 char kbuf[KSIZE];
4
5 /* Copy at most maxlen bytes from kernel region to user buffer */
6 int copy_from_kernel(void *user_dest, int maxlen) {
7     /* Byte count len is minimum of buffer size and maxlen */
8     int len = KSIZE < maxlen ? KSIZE : maxlen;
9     memcpy(user_dest, kbuf, len);
10    return len;
11 }
```

MEMCPY(3)

Linux Programmer's Manual

MEMCPY(3)

## NAME

memcpy - copy memory area

## SYNOPSIS

```
#include <string.h>
```

```
void *memcpy(void *dest, const void *src, size_t n);
```

## DESCRIPTION

The `memcpy()` function copies `n` bytes from memory area `src` to memory area `dest`. The memory areas must not overlap. Use `memmove(3)` if the memory areas do overlap.

## RETURN VALUE

The `memcpy()` function returns a pointer to `dest`.

From [GNU glibc manual](#) /Appendix A – Language Features /Important Data Types:

Data Type: `size_t`

This is an unsigned integer type used to represent the sizes of objects. The result of the `sizeof` operator is of this type, and functions such as `malloc` (see [Unconstrained Allocation](#)) and `memcpy` (see [Copying Strings and Arrays](#)) accept arguments of this type to specify object sizes. On systems using the GNU C Library, this will be `unsigned int` or `unsigned long int`.

**Usage Note:** `size_t` is the preferred way to declare any arguments or variables that hold the size of an object.

# Security - Woops

```
12
13 #define MSIZE 528
14
15 void getstuff() {
16     char mybuf[MSIZE];
17     copy_from_kernel(mybuf, -MSIZE);
18     . . .
19 }
```

l2ex1.c+

```
1 /* Kernel memory region holding user-accessible data */
2 #define KSIZE 1024
3 char kbuf[KSIZE];
4
5 /* Copy at most maxlen bytes from kernel region to user buffer */
6 int copy_from_kernel(void *user_dest, int maxlen) {
7     /* Byte count len is minimum of buffer size and maxlen */
8     int len = KSIZE < maxlen ? KSIZE : maxlen;
9     memcpy(user_dest, kbuf, len);
10    return len;
11 }
```

-528 in two's complement:

**0xFFFFDF0**



Reinterpreted as unsigned within  
memcpy: 4294966768 (decimal)

# Why do we care?



Always be mindful/careful  
of unsigned integers!

1. Two's complement
2. **Integer Arithmetic**
3. Bit Manipulation



How to deal with finite representation?

Adding two integers encoded on  $w$  bytes

Should take  $w+1$  bytes

How to encode the sum on  $w$  bytes?

No magic!! The sum will **overflow**

# Unsigned Addition

Operands:  $w$  bits



True Sum:  $w+1$  bits



Discard Carry:  $w$  bits



overflow

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

# Signed Addition

Operands:  $w$  bits



+  $v$




overflow

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  $s == t$

# Signed Addition

**Beware:**  
This is undefined  
behavior in C!

```
l2ex2.c+
1 #include <stdio.h>
2 #include <limits.h>
3
4 int main() {
5     // max short +1
6     short int i = 0xFFFF + 1;
7     printf("signed: 0xFFFF + 1 = %d\n", i);
8
9     // adding to large negative numbers
10    short int j = 0x8000;
11    short int k = j + j;
12    printf("signed: 0x8000 + 0x8000 = %d\n", k);
13
14    return 0;
15 }
```

```
cc l2ex2.c -o l2ex2
l2ex2.c:6:24: warning: implicit conversion from 'int' to 'short' changes value from 65536 to 0 [-Wconstant-conversion]
  short int i = 0xFFFF + 1;
                   ~~~~~~
1 warning generated.
```

```
~/Documents/Tmp > ./l2ex2
signed: 0xFFFF + 1 = 0
signed: 0x8000 + 0x8000 = 0
```

# Unsigned Multiplication in C

Operands:  $w$  bits



True Product:  $2*w$  bits



Discard  $w$  bits:  $w$  bits

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




## Standard Multiplication Function

Ignores high order  $w$  bits

Implements Modular Arithmetic

$$\text{UMult}_w(u, v) = u \cdot v \bmod 2^w$$

# 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


# Power-of-2 Multiply with Shift

Operation

$$\mathbf{u} \ll \mathbf{k} \text{ gives } \mathbf{u} * 2^k$$

Both signed and unsigned


Operands:  $w$  bits



True Product:  $w+k$  bits



Discard  $k$  bits:  $w$  bits



Examples

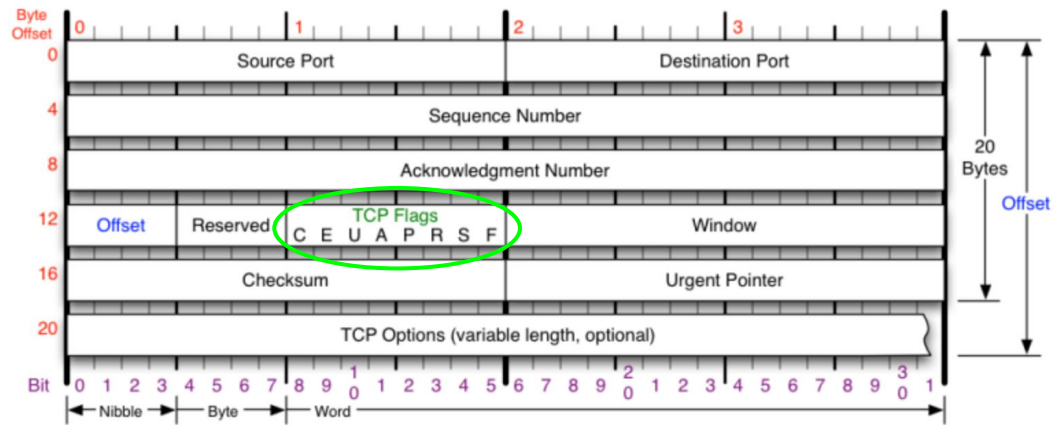
$$\mathbf{u} \ll 3 == \mathbf{u} * 8$$

$$\mathbf{u} \ll 5 - \mathbf{u} \ll 3 == \mathbf{u} * 24$$

1. Two's complement
2. Integer Arithmetic
3. **Bit Manipulation**



# Bit Representation



## TCP Header

TCP Flags	Congestion Notification	TCP Options	Offset																											
<b>C E U A P R S F</b>	ECN (Explicit Congestion Notification). See RFC 3168 for full details, valid states below.	0 End of Options List 1 No Operation (NOP, Pad) 2 Maximum segment size 3 Window Scale 4 Selective ACK ok 8 Timestamp	Number of 32-bit words in TCP header, minimum value of 5. Multiply by 4 to get byte count.																											
<p>Congestion Window</p> <ul style="list-style-type: none"> <li>C 0x80 Reduced (CWR)</li> <li>E 0x40 ECN Echo (ECE)</li> <li>U 0x20 Urgent</li> <li>A 0x10 Ack</li> <li>P 0x08 Push</li> <li>R 0x04 Reset</li> <li>S 0x02 Syn</li> <li>F 0x01 Fin</li> </ul>	<table border="1"> <thead> <tr> <th>Packet State</th> <th>DSB</th> <th>ECN bits</th> </tr> </thead> <tbody> <tr> <td>Syn</td> <td>00</td> <td>11</td> </tr> <tr> <td>Syn-Ack</td> <td>00</td> <td>01</td> </tr> <tr> <td>Ack</td> <td>01</td> <td>00</td> </tr> <tr> <td>No Congestion</td> <td>01</td> <td>00</td> </tr> <tr> <td>No Congestion</td> <td>10</td> <td>00</td> </tr> <tr> <td>Congestion</td> <td>11</td> <td>00</td> </tr> <tr> <td>Receiver Response</td> <td>11</td> <td>01</td> </tr> <tr> <td>Sender Response</td> <td>11</td> <td>11</td> </tr> </tbody> </table>	Packet State	DSB	ECN bits	Syn	00	11	Syn-Ack	00	01	Ack	01	00	No Congestion	01	00	No Congestion	10	00	Congestion	11	00	Receiver Response	11	01	Sender Response	11	11	<p>Checksum</p> <p>Checksum of entire TCP segment and pseudo header (parts of IP header)</p>	<p>RFC 793</p> <p>Please refer to RFC 793 for the complete Transmission Control Protocol (TCP) Specification.</p>
Packet State	DSB	ECN bits																												
Syn	00	11																												
Syn-Ack	00	01																												
Ack	01	00																												
No Congestion	01	00																												
No Congestion	10	00																												
Congestion	11	00																												
Receiver Response	11	01																												
Sender Response	11	11																												

Is Ack flag set?

Is any TCP flag set?

Is a single flag set?

How many TCP flags are set?

small puzzles;  
what the assignment is about

to answer these,  
you need to work with bitwise representations.

- 00000000 interpreted as **False**
- Byte containing at least a 1, e.g., 00100010 interpreted as **True**
- X is a sequence of bit
  - X interpreted as Boolean expression ( $X \neq 0$ )
  - !X interpreted as Boolean expression ( $X == 0$ )

Beware difference between:

- **Bitwise operations** applied on sequence of bits  
&, |, ~, ^, >>, <<
- **Logical operations** applied on Booleans  
&&, ||, !

# TCP Problem

Consider the TCP flags encoded on 8 bits as X.

How do you test whether the Ack flag (0x10) is set?


How do you test whether any flag is set?

How do you test if a single flag is set?

How do you count the number of flags set?

# Answers

- How do you test whether the Ack flag (0x10) is set?  
**(X & 0x10 )**  
=> interpreted as false if flag not set, true otherwise
- How do you test whether any flag is set?  
**(X)**  
=> interpreted as false if no flag is set , true otherwise
- How do you test if a single flag is set?  
**((X & (X-1)) == 0) && (X != 0)**  
=> (X & (X-1)) transforms the rightmost 1-bit into a 0-bit



small puzzles!  
some hate it; we like it (beauty)

How do you count the number of flags set?

Divide and conquer

Trivial for 2 bits

$0+0 \rightarrow 00$

$0+1 \rightarrow 01$

$1+0 \rightarrow 01$

$1+1 \rightarrow 10$

# Answers

x is a short: 2B



& 0x5555



& 0x3333

Masks



& 0x0F0F



& 0x00FF

Bitwise  
operations  
in parallel

$$x = (x \& 0x5555) + ((x \gg 1) \& 0x5555)$$

$$x = (x \& 0x3333) + ((x \gg 2) \& 0x3333)$$

$$x = (x \& 0x0F0F) + ((x \gg 4) \& 0x0F0F)$$

$$x = (x \& 0x00FF) + ((x \gg 8) \& 0x00FF)$$

# Motivation: Performance

## Fast binary matrix operations

```
#include <math.h>
#include <stdint.h>
#include <malloc.h>
#include <stdio.h>

#define N 8
#define M 8
#define n 0
#define m 1

int main(int argc, char *argv[])
{
    /* binary matrix allocation */
    uint64_t *matrix;
    matrix = (uint64_t *) malloc((size_t) ceil((N*M)/sizeof(uint64_t)));
    *matrix = 0x1001b;
    /* reading the [n][m]-th element of an NxM binary matrix */
    uint64_t mask = 1U << ((n*N+m)%64U);
    uint64_t result = mask & matrix[(size_t) ceil((n*N+m)/(64*1.0))];
    result = result >> ((n*N+m)%64U);
    printf("matrix[n][m]: %lu\n", result);
    return 0;
}
```

# Take Aways

Two complements used to represent signed integers

Two complements procedure is:

1. Write positive value in binary
2. Invert all digits
3. Add 1

Beware **implicit cast** to unsigned in C!!

security

Bit manipulation enables **efficient operations**.  
It is a rich algorithmic playground.

performance