



Integer Representation

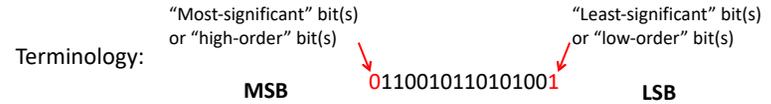
Representation of integers: unsigned and signed
Modular arithmetic and overflow
Sign extension
Shifting and arithmetic
Multiplication
Casting

Fixed-width integer encodings

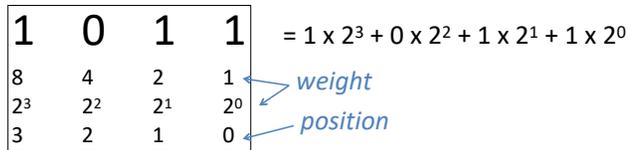
Unsigned $\subset \mathbb{N}$ non-negative integers only

Signed $\subset \mathbb{Z}$ both negative and non-negative integers

n bits offer only 2^n distinct values.



(4-bit) unsigned integer representation

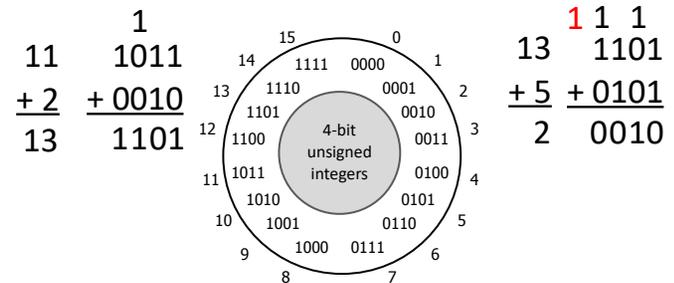


n -bit unsigned integers:

unsigned minimum = 0

unsigned maximum = $2^n - 1$

modular arithmetic, unsigned overflow



$x+y$ in n -bit unsigned arithmetic is $(x + y) \bmod 2^n$ in math

unsigned overflow = "wrong" answer = wrap-around = carry 1 out of MSB = math answer too big to fit

Unsigned addition overflows if and only if a carry bit is dropped.

(4-bit) two's complement signed integer representation

compare to unsigned

$$\begin{array}{cccc} \boxed{1} & \boxed{0} & \boxed{1} & \boxed{1} \\ \text{-(2}^3\text{)} & 2^2 & 2^1 & 2^0 \end{array} = 1 \times \text{-(2}^3\text{)} + 0 \times 2^2 + 1 \times 2^1 + 1 \times 2^0$$

still only 2^n distinct values, half negative.

4-bit two's complement integers:

signed minimum = $-(2^{n-1})$ 4-bit min: **1000**

signed maximum = $2^{n-1} - 1$ 4-bit max: **0111**

5

alternate signed attempt: sign-magnitude



Most-significant bit (MSB) is *sign bit*

0 means non-negative 1 means negative

Remaining bits are an unsigned magnitude

Note: this is *not* two's complement

8-bit sign-magnitude:

Anything weird here?

00000000 represents ____

01111111 represents ____

10000101 represents ____

10000000 represents ____

Arithmetic?

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

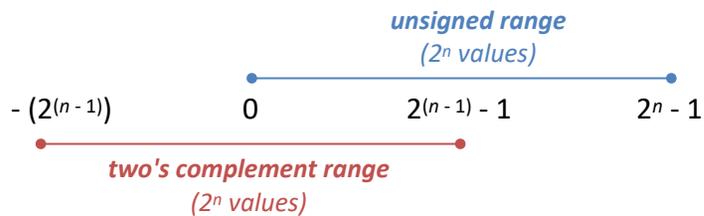
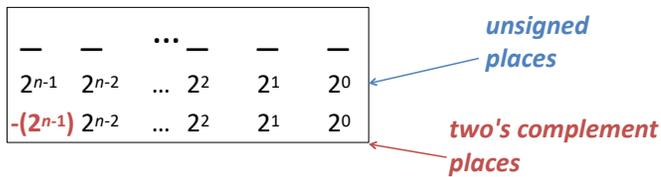
$$\begin{array}{r} 00000100 \\ + 10000011 \\ \hline \end{array}$$

Zero?

ex

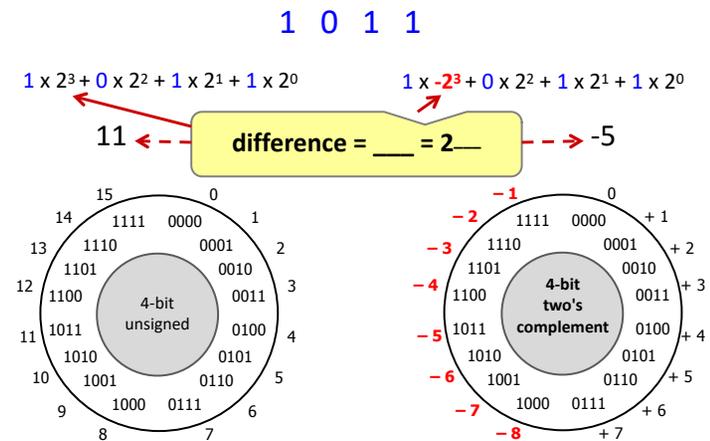
6

two's complement vs. unsigned



7

4-bit unsigned vs. 4-bit two's complement



8

8-bit representations

ex

00001001 10000001

11111111 00100111

n-bit two's complement numbers:

minimum =

maximum =

9

Consider a single byte: `\unsigned char x = 10101100;`
 What is the result of `\x << 2`?

10110011

10110000

00101011

11000000

Start the presentation to see live content. For screen share software, share the entire screen. Get help at polliv.com/app

Consider a single byte: `\unsigned char x = 10101100;`
 What is the result of `\x >> 2`?

11101011

10110000

00101011

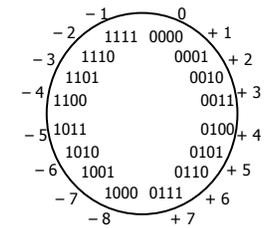
01010110

Start the presentation to see live content. For screen share software, share the entire screen. Get help at polliv.com/app

two's complement (signed) addition

$$\begin{array}{r} 1 \\ 2 \quad 0010 \\ + 3 \quad + 0011 \\ \hline 5 \quad 0101 \end{array} \qquad \begin{array}{r} 1 \quad 1 \\ -2 \quad 1110 \\ + -3 \quad + 1101 \\ \hline -5 \quad 1011 \end{array}$$

$$\begin{array}{r} 1 \quad 1 \\ -2 \quad 1110 \\ + 3 \quad + 0011 \\ \hline 1 \quad 0001 \end{array} \qquad \begin{array}{r} 2 \quad 0010 \\ + -3 \quad + 1101 \\ \hline -1 \quad 1111 \end{array}$$



Modular Arithmetic

12

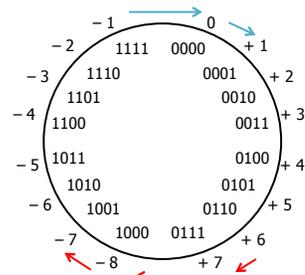
two's complement (signed) *overflow*

Addition overflows

if and only if the arguments have the same sign but the result does not.
if and only if the carry in and carry out of the sign bit differ.

$$\begin{array}{r} -1 \\ + 2 \\ \hline \end{array} \quad \begin{array}{r} 1\ 11 \\ 1111 \\ + 0010 \\ \hline 0001 \end{array}$$

$$\begin{array}{r} 6 \\ + 3 \\ \hline \end{array} \quad \begin{array}{r} 0\ 11 \\ 0110 \\ + 0011 \\ \hline 1001 \end{array}$$



Modular Arithmetic

Some CPUs/languages raise exceptions on overflow.
C and Java cruise along silently... Feature? Oops?

13

Recall: software correctness

Ariane 5 Rocket, 1996

Exploded due to cast of 64-bit floating-point number to 16-bit signed number.
Overflow.



Boeing 787, 2015



"... a **Model 787 airplane** ... can lose all alternating current (AC) electrical power ... caused by a **software counter** after **248 days** of continuous power. We are issuing this AD to prevent loss of all AC electrical power, which could result in **loss of control of the airplane.**"
--FAA, April 2015

14

A few reasons *two's complement* is awesome

Arithmetic hardware

The carry algorithm works for everything!

Sign

The MSB can be interpreted as a sign bit.

Negative one

-1_{10} is encoded as all ones: $0b11...1$

Complement rules

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

5 is $0b0101$

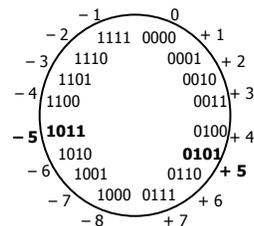
$\sim 0b0101$ is $0b1010$

$$+ \quad \underline{\quad 1}$$

$0b1011$ is -5

Even subtraction!

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



15

Another derivation

ex

How should we represent 8-bit negatives?

- For all positive integers x , we want the representations of x and $-x$ to sum to zero.
- We want to use the standard addition algorithm.

$$\begin{array}{r} 11111111 \\ 00000001 \\ \hline +11111111 \\ \hline 00000000 \end{array} \quad \begin{array}{r} 11111111 \\ 00000010 \\ \hline +11111110 \\ \hline 00000000 \end{array} \quad \begin{array}{r} 11111111 \\ 00000011 \\ \hline +11111101 \\ \hline 00000000 \end{array}$$

- Find a rule to represent $-x$ where that works...

16

Convert/cast signed number to larger type.

```

      00000010    8-bit 2
-----
00000010    16-bit 2

      11111100    8-bit -4
-----
11111100    16-bit -4
  
```

Rule/name?

Sign extension for two's complement

```

      00000010    8-bit 2
      / / / / / / / /
0000000000000010    16-bit 2

      11111100    8-bit -4
      / / / / / / / /
1111111111111100    16-bit -4
  
```

Casting from smaller to larger signed type does sign extension.

unsigned shifting and arithmetic

unsigned
 $x = 27;$
 $y = x \ll 2;$
 $y == 108$

```

00011011
/ / / / / / / /
0001101100
  
```

logical shift left

$n = \text{shift distance in bits}, w = \text{width of encoding in bits}$

```

11101101
/ / / / / / / /
0011101101
  
```

logical shift right

unsigned
 $x = 237;$
 $y = x \gg 2;$
 $y == 59$

two's complement shifting and arithmetic

signed
 $x = -101;$
 $y = x \ll 2;$
 $y == 108$

```

10011011
/ / / / / / / /
1001101100
  
```

logical shift left

$n = \text{shift distance in bits}, w = \text{width of encoding in bits}$

```

11101101
/ / / / / / / /
1111101101
  
```

arithmetic shift right

signed
 $x = -19;$
 $y = x \gg 2;$
 $y == -5$

Consider a single *signed* byte: `\`signed char x = 10101100;\``.
 What is the result of `\`x >> 2\``?

- 10110011
- 00101011
- 11101100
- 11101011

Start the presentation to see live content. For screen share software, share the entire screen. Get help at polllev.com/app

shift-and-add



Available operations

`x << k` implements `x * 2k`
`x + y`

Implement `y = x * 24` using only `<<`, `+`, and integer literals

```
y = x * (16 + 8);
y = (x * 16) + (x * 8);
y = (x << 4) + (x << 3)
```

Parenthesize shifts to be clear about precedence, which may not always be what you expect.

Casting Integers in C



Number literals: `37` is signed, `37U` is unsigned

Integer Casting: *bits unchanged, just reinterpreted.*

Explicit casting:

```
int tx = (int) 73U;    // still 73
unsigned uy = (unsigned) -4; // big positive #
```

Implicit casting: Actually does

```
tx = ux;    // tx = (int)ux;
uy = ty;    // uy = (unsigned)ty;
void foo(int z) { ... }
foo(ux);    // foo((int)ux);
if (tx < ux) ... // if ((unsigned)tx < ux) ...
```

More Implicit Casting in C



If you mix unsigned and signed in a single expression, then
signed values are implicitly cast to unsigned.

How are the argument bits interpreted?

Argument ₁	Op	Argument ₂	Type	Result
0	==	0U	unsigned	
-1	<	0	signed	
-1	<	0U	unsigned	
2147483647	<	-2147483647-1		
2147483647U	<	-2147483647-1		
-1	<	-2		
(unsigned)-1	<	-2		
2147483647	<	2147483648U		
2147483647	<	(int)2147483648U		

Note: $T_{min} = -2,147,483,648$ $T_{max} = 2,147,483,647$

T_{min} must be written as `-2147483647-1` (see pg. 77 of CSAPP for details)

Let x be an `int` (that is, a 32-bit signed two's complement integer). Write an expression in terms of `x` *without* using constants greater than `0xFF`.
Write an expression that sets all bytes other than the most significant byte to 0.

Nobody has responded yet.
Hang tight! Responses are coming in.

Start the presentation to see live content. For screen share software, share the entire screen. Get help at pollev.com/app

Aside: real-world connection to Alexa's research

Guest-controlled out-of-bounds read/write on x86_64

Critical alexrichton published GHSA-ff4p-7xrq-q5r8 on Mar 8

Package	Affected versions	Patched versions
@cranelift-codegen (Rust)	<= 0.93.0, >= 0.84.0	0.93.1, 0.92.1, 0.91.1
@wasmtime (Rust)	<= 6.0.0, >= 0.37.0	6.0.1, 5.0.1, 4.0.1

Severity
Critical 9.9 / 10

CVSS base metrics	
Attack vector	Network
Attack complexity	Low
Privileges required	Low
User interaction	None
Scope	Changed
Confidentiality	High
Integrity	High
Availability	High

CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:C/CH:H/A:H

CVE ID
CVE-2023-26489

Description

Impact

Wasmtime's code generator, Cranelift, has a bug on x86_64 targets where address-mode computation mistakenly would calculate a 35-bit effective address instead of WebAssembly's defined 33-bit effective address. This bug means that, with default codegen settings, a wasm-controlled load/store operation could read/write addresses up to 35 bits away from the base of linear memory. Wasmtime's default sandbox settings provide up to 6G of protection from the base of linear memory to guarantee that any memory access in that range will be semantically correct. Due to this bug, however, addresses up to $0xFFFFFFFF * 8 + 0x7FFFFFFC = 36587222884 = -346$ bytes away from the base of linear memory are possible from guest code. This means that the virtual memory 6G away from the base of linear memory up to -34G away can be read/written by a malicious module.

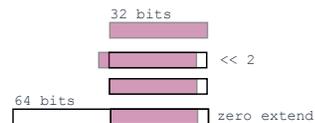
Security-critical bug in shift-and-extend code

Guest-controlled out-of-bounds read/write on x86_64

GHSA-ff4p-7xrq-q5r8 published on Mar 8 by alexrichton

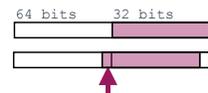
Conceptually, the compiler tried to convert this with a 32-bit x :

```
address + zero_extend_64(x << 2)
```



To this:

```
address + (zero_extend_64(x) << 2)
```



Incorrect address calculated!