



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.

Terminology:

“Most-significant” bit(s)
or “high-order” bit(s)

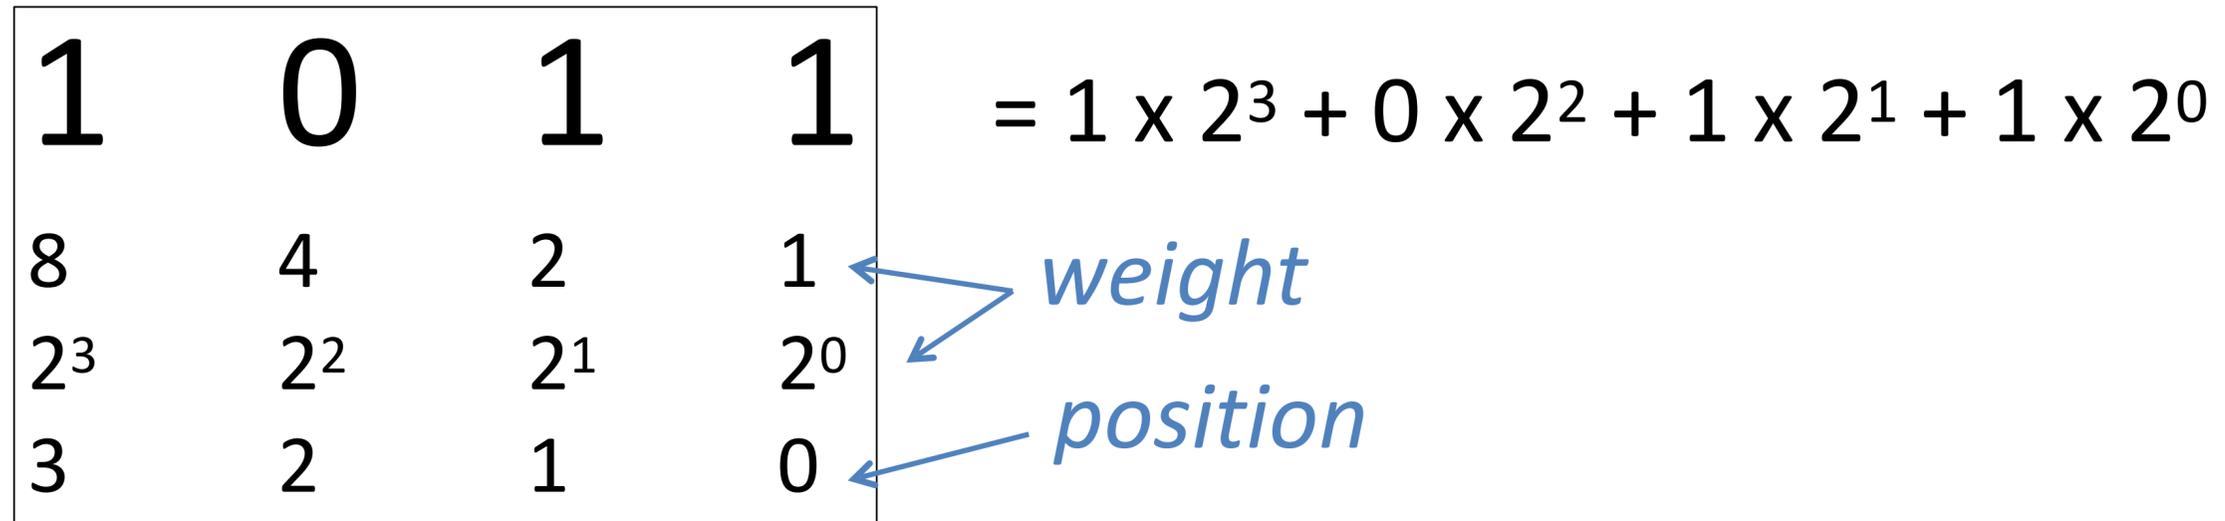
MSB

0110010110101001

“Least-significant” bit(s)
or “low-order” bit(s)

LSB

(4-bit) unsigned integer representation



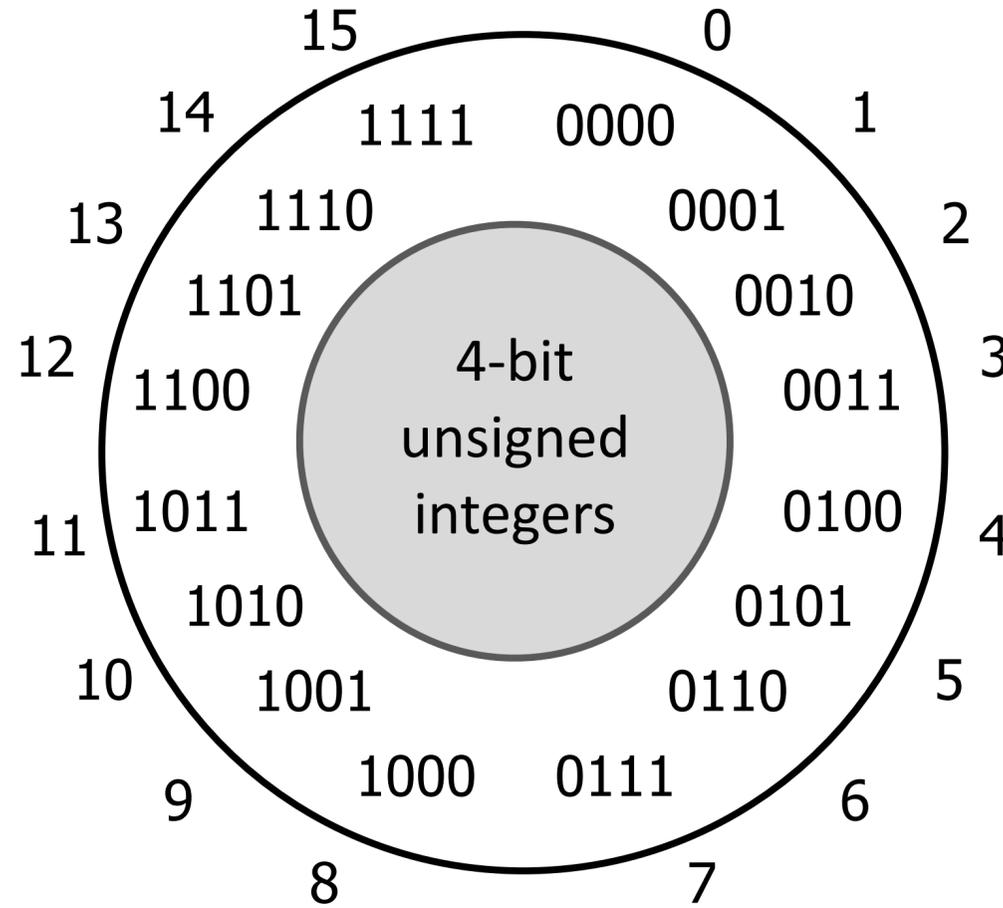
n -bit unsigned integers:

unsigned minimum = 0

unsigned maximum = $2^n - 1$

modular arithmetic, unsigned overflow

$$\begin{array}{r}
 11 \\
 + 2 \\
 \hline
 13
 \end{array}
 \qquad
 \begin{array}{r}
 1 \\
 1011 \\
 + 0010 \\
 \hline
 1101
 \end{array}$$



$$\begin{array}{r}
 13 \\
 + 5 \\
 \hline
 2
 \end{array}
 \qquad
 \begin{array}{r}
 111 \\
 1101 \\
 + 0101 \\
 \hline
 0010
 \end{array}$$

$x+y$ 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



| | | | |
|----------|----------|----------|----------|
| 1 | 0 | 1 | 1 |
| $-(2^3)$ | 2^2 | 2^1 | 2^0 |

$$= 1 \times -(2^3) + 0 \times 2^2 + 1 \times 2^1 + 1 \times 2^0$$

still only 2^n distinct values, half negative.

4-bit two's complement integers:

$$\text{signed minimum} = - (2^{(n-1)})$$

4-bit min: **1000**

$$\text{signed maximum} = 2^{(n-1)} - 1$$

4-bit max: **0111**

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:

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



$$\begin{array}{r} 0000100 \\ +1000011 \\ \hline \end{array}$$

Zero?

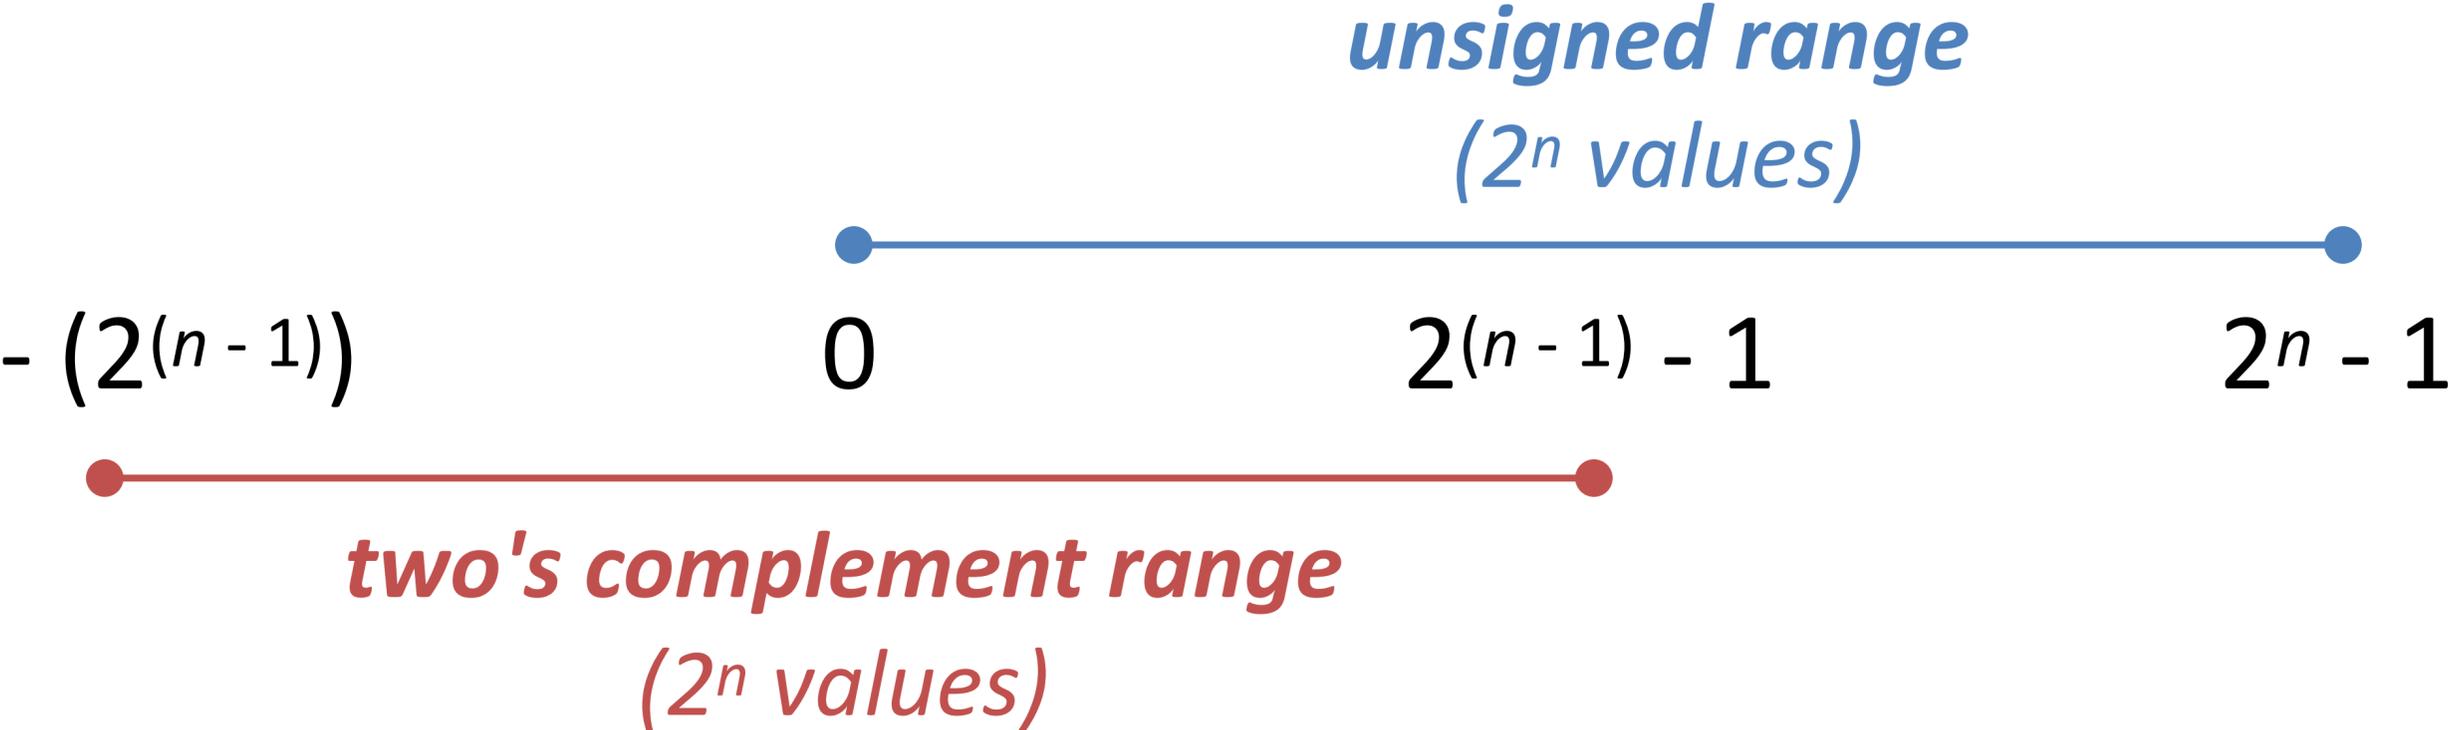


two's complement vs. unsigned

| | | | | | |
|--------------|-----------|-----|-------|-------|-------|
| — | — | ... | — | — | — |
| 2^{n-1} | 2^{n-2} | ... | 2^2 | 2^1 | 2^0 |
| $-(2^{n-1})$ | 2^{n-2} | ... | 2^2 | 2^1 | 2^0 |

unsigned places

two's complement places



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

1 0 1 1

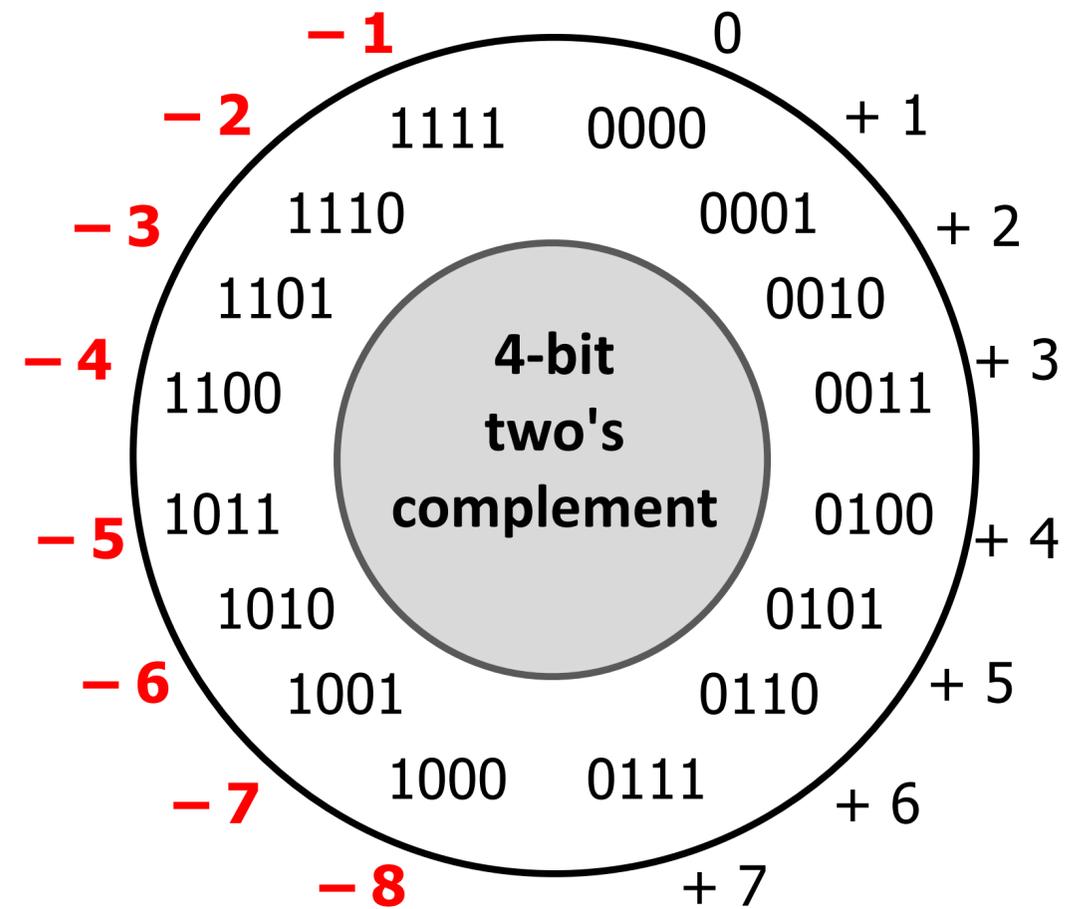
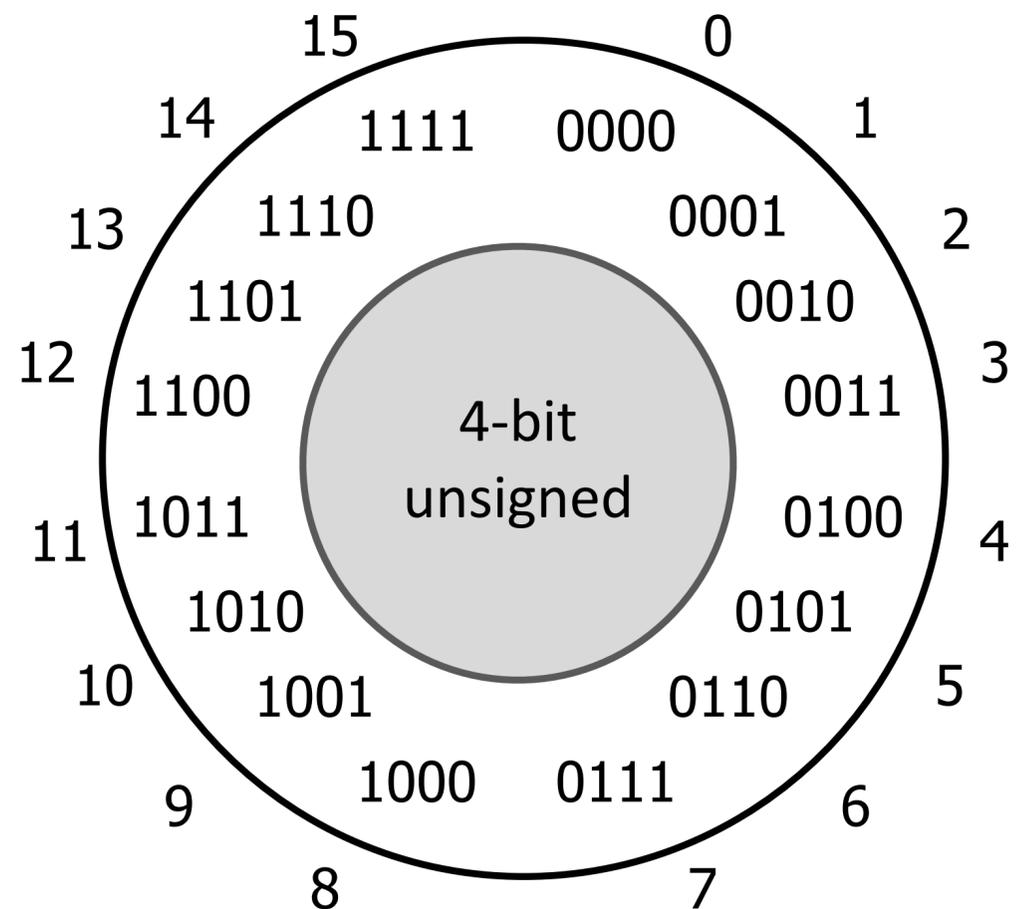
$$1 \times 2^3 + 0 \times 2^2 + 1 \times 2^1 + 1 \times 2^0$$

$$1 \times -2^3 + 0 \times 2^2 + 1 \times 2^1 + 1 \times 2^0$$

11

difference = ___ = 2___

-5



8-bit representations



0 0 0 0 1 0 0 1

1 0 0 0 0 0 0 1

1 1 1 1 1 1 1 1

0 0 1 0 0 1 1 1

n-bit two's complement numbers:

minimum =

maximum =

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

10110011

10110000

00101011

11000000

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

11101011

10110000

00101011

01010110

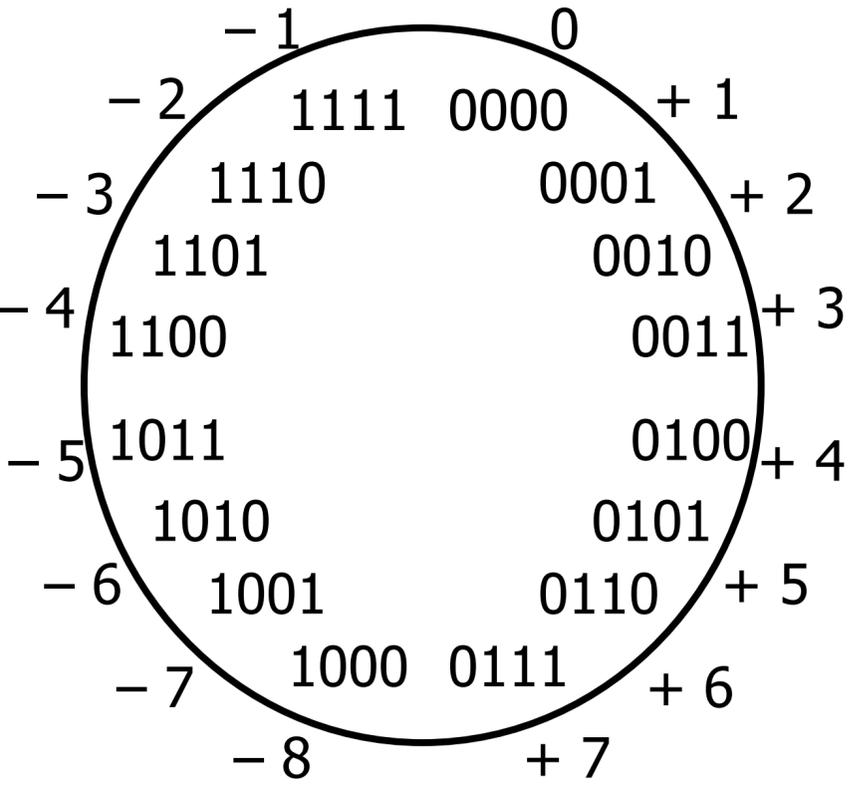
two's complement (signed) addition

$$\begin{array}{r}
 1 \\
 2 \quad 0010 \\
 + 3 \quad + 0011 \\
 \hline
 5 \quad 0101
 \end{array}$$

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

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

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



Modular Arithmetic

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** internal to the GCUs that will **overflow** 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

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`

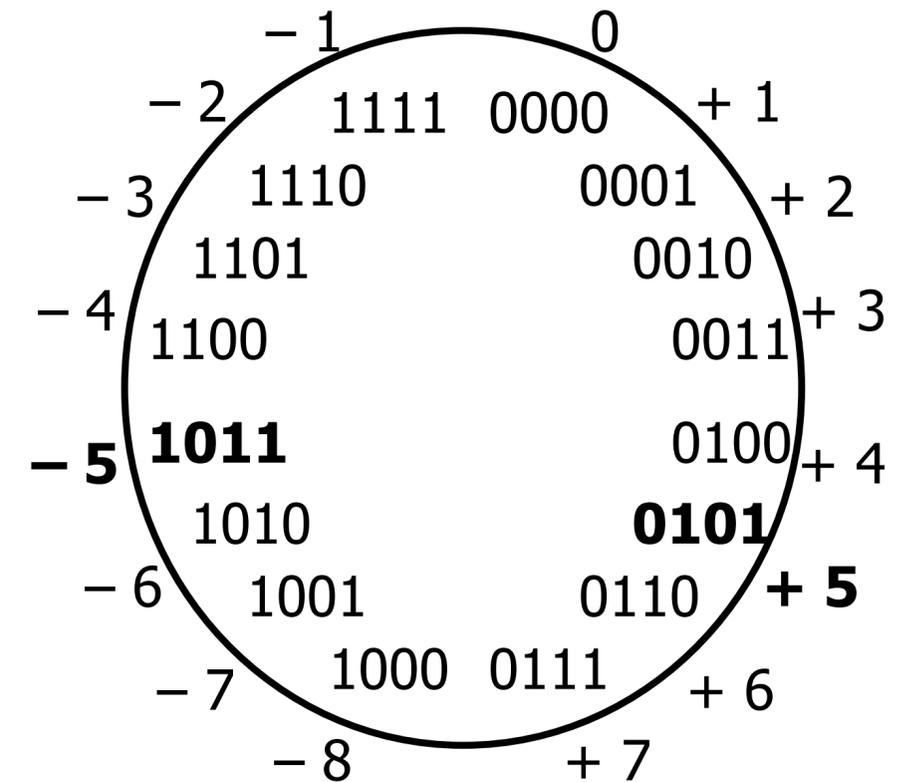
$\sim 0b0101$ is `0b1010`

$$\begin{array}{r} 0b1010 \\ + \quad 1 \\ \hline \end{array}$$

`0b1011` is -5

Even subtraction!

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



Another derivation



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.

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

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

Convert/cast signed number to larger type.

0 0 0 0 0 0 1 0

8-bit 2

0 0 0 0 0 0 1 0

16-bit 2

1 1 1 1 1 1 0 0

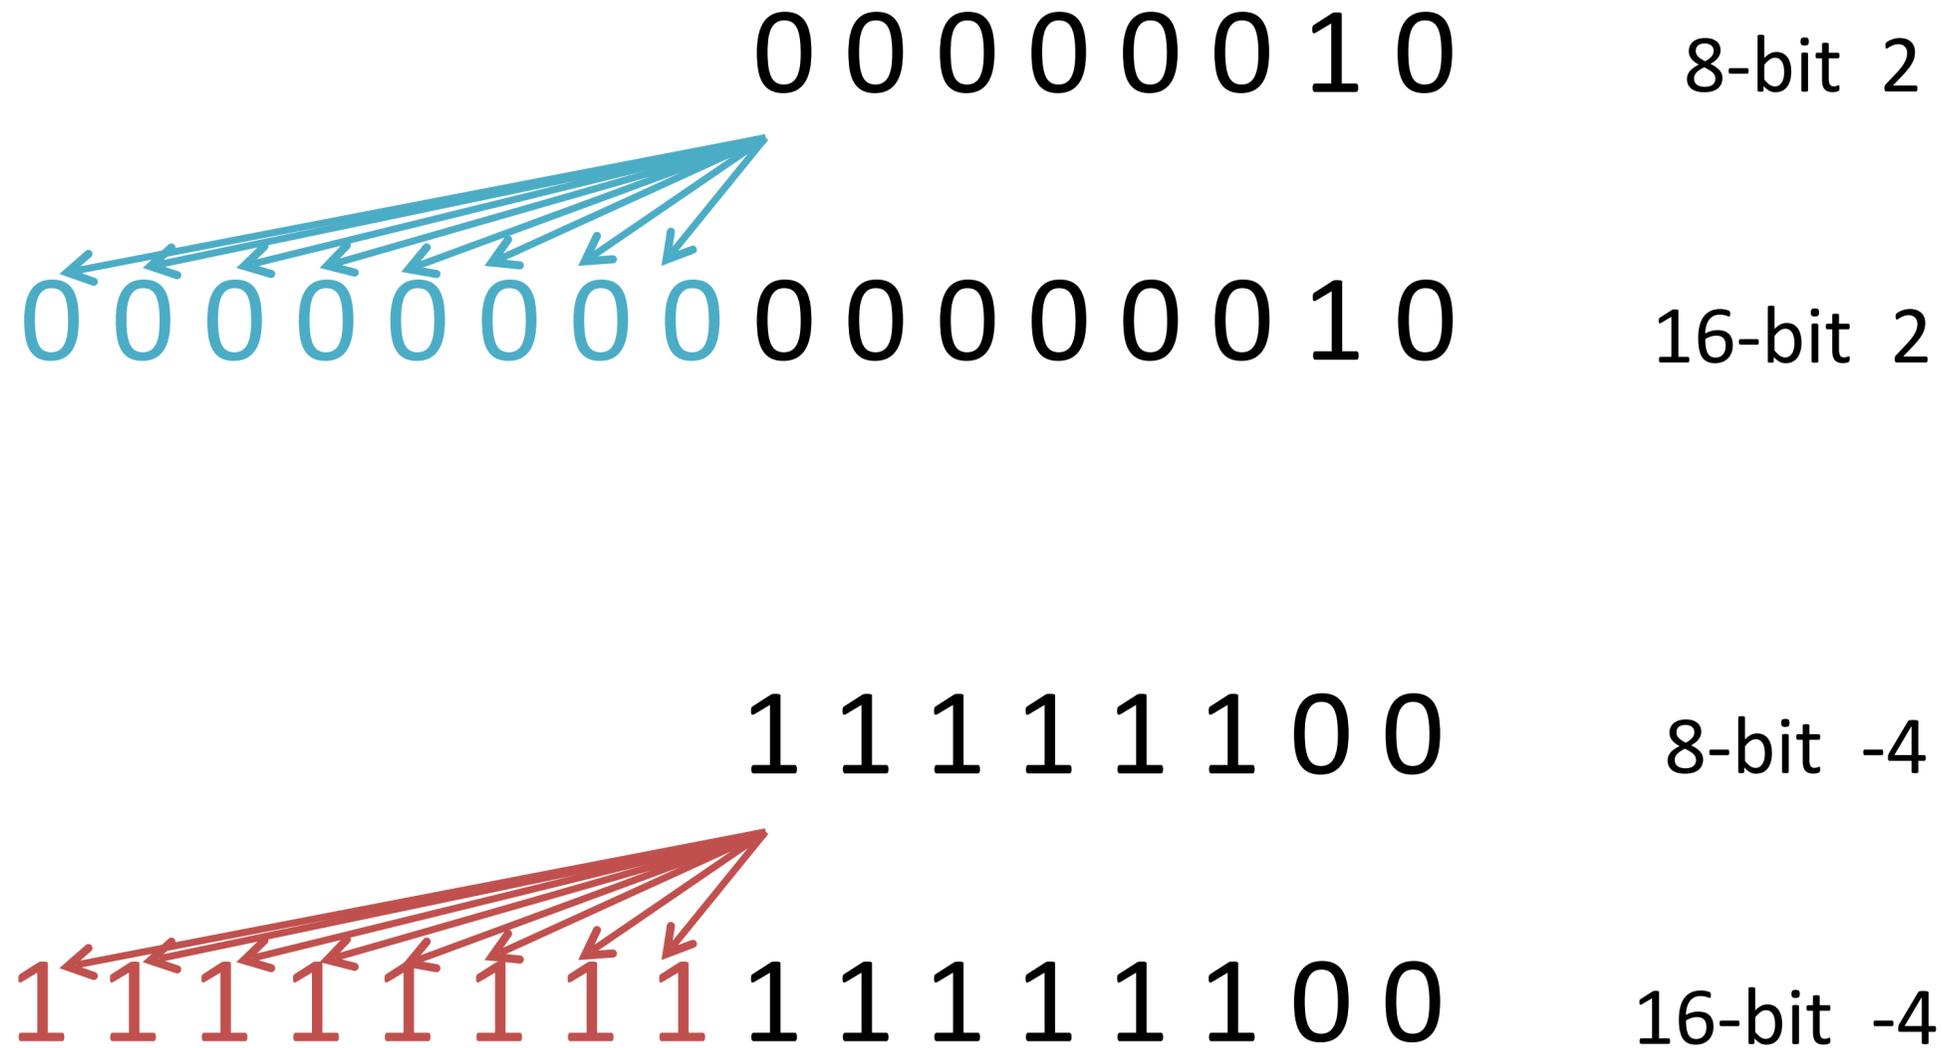
8-bit -4

1 1 1 1 1 1 0 0

16-bit -4

Rule/name?

Sign extension for two's complement



Casting from smaller to larger signed type does sign extension.

unsigned shifting and arithmetic

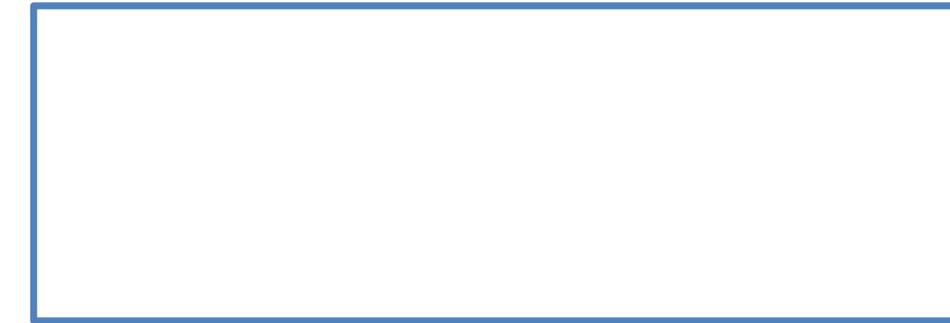
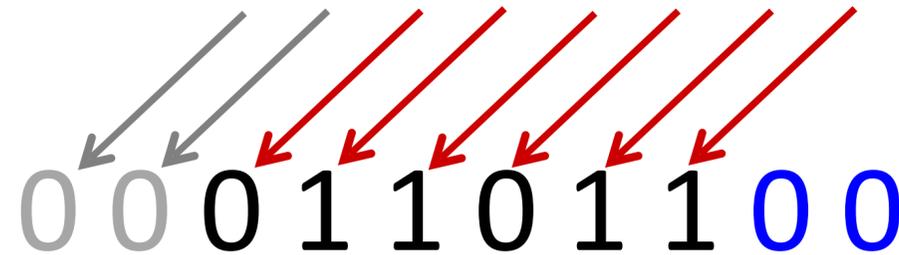
unsigned

$x = 27;$

$y = x \ll 2;$

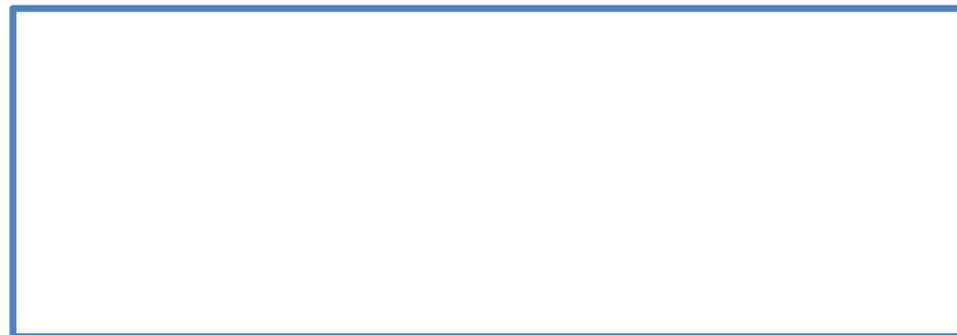
$y == 108$

0 0 0 1 1 0 1 1



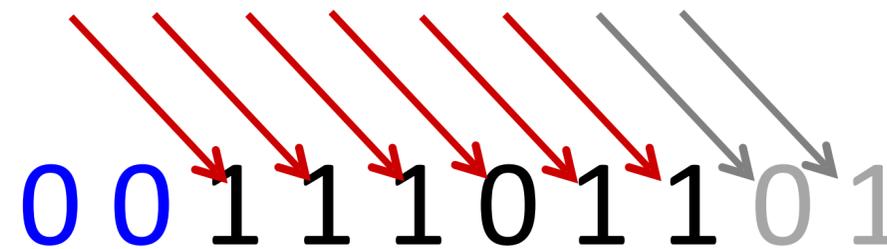
logical shift left

$n =$ shift distance in bits, $w =$ width of encoding in bits



logical shift right

1 1 1 0 1 1 0 1



unsigned

$x = 237;$

$y = x \gg 2;$

$y == 59$

two's complement shifting and arithmetic

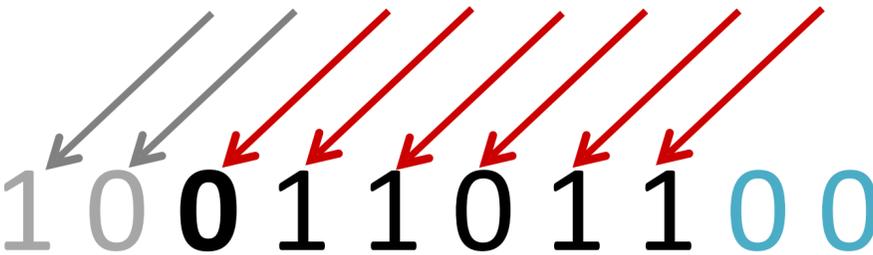
signed

x = -101;

y = x << 2;

y == 108

1 0 0 1 1 0 1 1



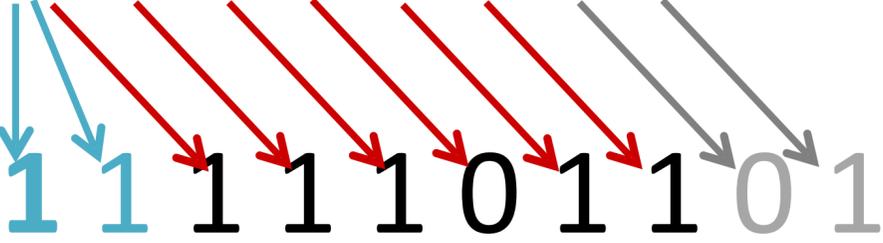
logical shift left

n = shift distance in bits, w = width of encoding in bits



arithmetic shift right

1 1 1 0 1 1 0 1



signed

x = -19;

y = x >> 2;

y == -5

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

10110011

00101011

11101100

11101011

shift-and-add

ex

Available operations

$x \ll k$ implements $x * 2^k$

$x + y$

Implement $y = x * 24$ using only \ll , $+$, and integer literals

$y = x * (16 + 8);$

$y = (x * 16) + (x * 8);$

$y = (x \ll 4) + (x \ll 3)$

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.

Aside: real-world connection to Alexa's research

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

Critical alexcrichton 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

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 + 0x7fffffff` = 36507222004 = ~34G 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.

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/C:H/I:H/A:H

CVE ID
CVE-2023-26489

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 alexcrichon

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!