

# **Lecture 03**

# **Data Operations**

CS213 – Intro to Computer Systems  
Branden Ghen a – Winter 2026

Slides adapted from:

St-Amour, Hardavellas, Bustamente (Northwestern), Bryant, O'Hallaron (CMU), Garcia, Weaver (UC Berkeley)

# Course schedule change

- I decided to change things so you'd have more study time.
  - Details posted to Piazza ("Important: Schedule Change")
- Biggest impacts:
  - Midterm exam 1 – moved to Tuesday, January 27
    - Two weeks from today
  - Homework 1 – due Thursday, January 22
- Future changes are quite unlikely.

# Office hours schedule

- Office hours schedule is up
  - See Canvas homepage for office hours times. Starts tomorrow!
  - Mostly in-person with a few online hours
    - Online uses gather.town (Room D)
  - Office hours queue on the Canvas homepage when things get busy
  - **Note:** no office hours next Monday for MLK Day
- Today only: online hours from 4-6pm

# Pack Lab details

- Pack Lab should be out later today!
  - Sometime this evening
- Pack Lab partnership survey on Piazza
  - If you want a partner but don't know who you want to work with
  - I'll pair people from the list right after I post the lab
- You'll do Pack Lab on one of the EECS servers
  - Usually we use Moore, but any EECS server should be fine for this lab
  - SSH + Command Line interface
  - See Piazza post with some details on accessing the servers

# Today's Goals

- Finish encodings thoughts from last time
- Explore operations we can perform on integers and more generally on binary numbers
- Understand the edge cases of those operations

# Outline

- **Integer Operations**
  - **Addition**
  - Negation and Subtraction
  - Multiplication and Division
- Binary Operations
  - Boolean Algebra
  - Shifting
  - Bit Masks

# C versus the hardware

- Operations you can perform on types have edge conditions
  - Usually going above or below the bounds for the given size/encoding
- If we say what happens in that scenario, it'll be what "the hardware" (i.e., a computer) does
  - In today's examples, pretty much every computer does the same thing
- That is not the same as what C does
  - Unclear choices are left as: **UNDEFINED BEHAVIOR** 🤪
  - Which is to say, the compiler can make any choice it wants

# Unsigned Addition

- Like grade-school addition, but in base 2, and ignores final carry
  - If you want, can do addition in base 10 and convert to base 2. Same result! But here we're going to understand what the hardware is doing.
- **Example: Adding two 4-bit numbers**

$$\begin{array}{r} \overset{1}{0}\overset{1}{1}\overset{1}{0}1 \\ + \quad 0011 \\ \hline 1000 \end{array}$$

- **$5_{10} + 3_{10} = 8_{10} \checkmark$**


# Unsigned Addition and Overflow

- What happens if the numbers get too big?
- **Example: Adding two 4-bit numbers**

$$\begin{array}{r} \phantom{+} \overset{1}{1} \overset{1}{1} \overset{1}{1} \overset{1}{1} \\ \phantom{+} 1101 \\ + 0011 \\ \hline \textcolor{red}{1}0000 \end{array}$$


- **$13_{10} + 3_{10} = 16_{10}$** 
  - Too large for 4 bits!
  - Truncate most-significant bits that do not fit
    - Gives us modular (= modulo) behavior:  $16 \text{ modulo } 2^4 = 0$
- Final result is the 4 least significant bits (all we can fit): so  $0_{10}$

# Modulo behavior in binary numbers



# Unsigned addition is modular

- Implements modular arithmetic
  - $\text{UAdd}_w(u, v) = (u + v) \bmod 2^w$
- Need to drop carry-out bit, otherwise results will keep getting bigger
  - Example in base 10:  $80_{10} + 40_{10} = 120_{10}$  (2-digit inputs become a 3-digit output!)



- Warning: C does not tell you that the result was truncated!
  - **Unsigned** addition in C silently truncates most-significant bits beyond the limit

# Signed (2's Complement) Addition

- Works exactly the same as unsigned addition!
  - Just add the numbers in binary, and the result will work out
- Signed and unsigned sum have the exact same bit-level representation
  - Computers use the same machine instruction and the same hardware!
  - That's a big reason 2's complement is so nice! Shares operations with unsigned

# Signed addition example

- Same addition method as unsigned
- **Example: Adding two 4-bit signed numbers**

$$\begin{array}{rcl} & \overset{1}{1} & \overset{1}{} \\ & 1011 & (-8 + 3 = -5) \\ + & 0011 & (+3) \\ \hline & 1110 & (-8 + 6 = -2) \end{array}$$

- $-5_{10} + 3_{10} = -2_{10} \checkmark$

# Combining negative and positive numbers


- Truncation makes signed addition work!
- **Example: Adding two 4-bit signed numbers**

$$\begin{array}{r} \overset{1}{1} \overset{1}{1} \overset{1}{1} \overset{1}{0} 1 \quad (-8 + 5 = -3) \\ + \quad 0011 \quad (+3) \\ \hline \textcolor{red}{1}0000 \end{array}$$

- **$-3_{10} + 3_{10} = 0_{10}$** 
  - Too large for 4 bits! Drop the carry bit
  - Result is what we expect as long as we truncate

# Overflow of signed values

- Addition can lead to **overflow**
  - Unsigned: Result wraps around from big to small or vice versa
  - Signed: Both operands are positive and result is negative
  - Signed: Both operands are negative and result is positive



# Overflow: hardware vs C standard

- Hardware implementations for unsigned and signed addition are the same
  - Both implement truncation of overflowing bits, leads to modular arithmetic
- Unsigned overflow in C is defined as modular arithmetic
  - Safe to use. Big number wraps around to a small number
- Signed overflow in C is **UNDEFINED BEHAVIOR**
  - Compiler *probably* does modular result
  - But there are no promises about this and it can make *assumptions*
  - So don't rely on it
  - Generally: overflow is bad and is to be avoided!

# Overflow in the real world

- In Switzerland, there's a rule that trains are not allowed to have exactly 256 axles

## 3.7.4 Zugbildung

Um das ungewollte Freimelden von Streckenabschnitten durch das Rückstellen der Achszähler auf Null und dadurch Zuggefährdungen zu vermeiden, darf die effektive Gesamtachszahl eines Zuges nicht 256 Achsen betragen.

"To avoid falsely signalling a section of track as clear by resetting the axle counter to zero, and thus to avoid [collisions], the total number of axles in a train must not equal 256."

# Video games bugs can be exploited

- Dream Devourer
  - Special boss in the Nintendo DS edition
- Wanted to make it even more challenging
  - ~32000 hit points
  - Takes *forever* to defeat
- Hit points stored as a 16-bit signed integer
  - Range: -32768 to +32767
- **How do speedrunners defeat the boss?**



# Chrono Trigger signed overflow bug

- Solution: heal it
- Hit points go negative and it dies



# Outline

- **Integer Operations**

- Addition
- **Negation and Subtraction**
- Multiplication and Division

- **Binary Operations**

- Boolean Algebra
- Shifting
- Bit Masks

# Negating a number

- In C:
  - $x = -y;$
- Operation
  - Determine the negative, signed version of the number (two's complement)
  - Hardware method: flip bits and add one
- Complement operator ( $\sim$ )
  - Flips all bits: zeros become a one and ones become a zero
  - $\sim 0b1011 \rightarrow 0b0100$

# Negating via Complement & Increment

- Claim: The following is true for 2's complement

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

- Complement

- Observation:  $\sim x + x == 1111...11_2 == -1$

$$\begin{array}{r}
 x \quad \boxed{1} \boxed{0} \boxed{0} \boxed{1} \boxed{1} \boxed{1} \boxed{0} \boxed{1} \\
 + \quad \sim x \quad \boxed{0} \boxed{1} \boxed{1} \boxed{0} \boxed{0} \boxed{0} \boxed{1} \boxed{0} \\
 \hline
 -1 \quad \boxed{1} \boxed{1} \boxed{1} \boxed{1} \boxed{1} \boxed{1} \boxed{1} \boxed{1}
 \end{array}$$

- Increment

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

- Example, 4 bits:  $6_{10} = 0110_2$

- Complement:  $1001_2 \rightarrow$  Increment =  $1010_2 = -8 + 2 = -6_{10}$

# Subtraction in two's complement

- Subtraction becomes addition of the negative number
  - $5 - 3 = 5 + -3 = 2$
- Both unsigned and signed subtraction
  - Convert subtrahend to its two's complement negative form
    - i.e., negate it
  - Then do addition
  - Treat result as original encoding

$$\begin{array}{r} \overset{1}{0} \overset{1}{1} \overset{1}{0} 1 \quad (+5) \\ + \quad 1 1 0 1 \quad (-3) \\ \hline \textcolor{red}{1} 0 0 1 0 \end{array}$$

# C rules vs hardware rules

- Exact same overflow rules apply
- Unsigned subtraction overflow can wrap below zero to make a large number
  - Modular arithmetic
- Signed subtraction overflow is **UNDEFINED BEHAVIOR**
  - And therefore should not be trusted

# Break + practice

- Adding two 8-bit binary numbers:
  - Also determine the decimal version of the result
  - I'm intentionally not telling you the encoding

$$\begin{array}{r} 00010101 \\ + \underline{10110001} \end{array}$$

# Break + practice

- Adding two 8-bit binary numbers:
  - Also determine the decimal version of the result
  - I'm intentionally not telling you the encoding

|                                                                                                                                                                                                                                                      | Unsigned encoding                                                                   |    | Signed encoding                                                                       |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|----|---------------------------------------------------------------------------------------|
| $\begin{array}{r} \phantom{+} \overset{1}{0}\overset{1}{0}\phantom{0}\overset{1}{1}\overset{1}{0}\overset{1}{1} \\ + \phantom{0}\overset{1}{1}\overset{0}{0}\overset{1}{1}\overset{0}{0}\overset{0}{0}\overset{1}{1} \\ \hline 11000110 \end{array}$ | $\begin{array}{l} 16+4+1 = 21 \\ 128+32+16+1 = 177 \\ 128+64+4+2 = 198 \end{array}$ | OR | $\begin{array}{l} 16+4+1 = 21 \\ -128+32+16+1 = -79 \\ -128+64+4+2 = -58 \end{array}$ |

# Break + practice

- Adding two 8-bit binary numbers:
  - Also determine the decimal version of the result
  - I'm intentionally not telling you the encoding

|                                                                                                                  | Unsigned encoding                                          |    | Signed encoding                                              |
|------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------|----|--------------------------------------------------------------|
| $\begin{array}{r} \phantom{0}1\phantom{0}1\phantom{0}1 \\ 00010101 \\ + 10110001 \\ \hline 11000110 \end{array}$ | $16+4+1 = 21$<br>$128+32+16+1 = 177$<br>$128+64+4+2 = 198$ | OR | $16+4+1 = 21$<br>$-128+32+16+1 = -79$<br>$-128+64+4+2 = -58$ |

What about unsigned subtraction 21-79?

That would treat the result as unsigned, with the value 198  
Modular arithmetic in action


# Outline

- **Integer Operations**
  - Addition
  - Negation and Subtraction
  - **Multiplication and Division**
- Binary Operations
  - Boolean Algebra
  - Shifting
  - Bit Masks

# Multiplication

- Goal: Compute the Product of two  $w$ -bit numbers  $x, y$ 
  - Either signed or unsigned
- But, exact results can be bigger than  $w$  bits
  - Double the size ( $2w$ ), in fact!
  - Example in base 10:  $50_{10} * 20_{10} = 1000_{10}$ 
    - (2-digit inputs become a 4-digit output!)
- As with addition, result is truncated to fit in  $w$  bits
  - Because computers are finite, results can't grow indefinitely

# Unsigned Multiplication



- **Standard Multiplication Function**

- Equivalent to grade-school multiplication
- But ignores most significant  $w$  bits of the result
- As a person, we can do base 10 multiplication, convert to base 2, then truncate

- Implements modular arithmetic like addition does

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


# Unsigned multiplication

- **Example: Multiplying two 4-bit numbers**

$$\begin{array}{r} 0010 \\ \times 0101 \\ \hline 0010 \\ 00000 \\ 001000 \\ + 0000000 \\ \hline \cancel{000}1010 \end{array}$$

$$2_{10} * 5_{10} = 10_{10} \quad \checkmark$$

# Signed (2's Complement) Multiplication



- **Standard Multiplication Function** (same as unsigned)
  - Ignores most significant  $w$  bits
  - Lower bits still give the correct result
    - So we can use same machine instruction for both!
    - Again, that's one reason why 2's complement is so nice
- **In C, signed overflow is undefined**
  - ...but probably you'll see the two's complement behavior

# What about divide?

- Annoying operation, not going to discuss in this class
  - Similar to long division process
  - Tedious and complicated to get right
- I've worked on computers that don't have hardware support for division at all!!
- Important thing to remember is that integers don't have fractional parts
  - In C:  $1 / 2 == 0$
  - We'll need a different encoding for fractional numbers: floating point

# Outline

- Integer Operations
  - Addition
  - Negation and Subtraction
  - Multiplication and Division
- **Binary Operations**
  - **Boolean Algebra**
  - Shifting
  - Bit Masks


# Boolean algebra

- You've programmed with **and** and **or** in earlier classes
  - C and C++ logical operations: **&&** (and), **||** (or)
- **Boolean algebra is a generalization of that**
  - A mathematical system to represent logic (propositional logic)
  - 2 truth values: true = **1**, false = **0**
  - Operations: and **&**, or **|**, not (or complement) **~**
    - And others

# Truth tables for Boolean algebra

- For each possible value of each input, what is the output
  - Column for each input
  - Column for the output operation

Complement (NOT)



**$\sim A$**

| A | $\sim A$ |
|---|----------|
| 0 | 1        |
| 1 | 0        |

OR



**A | B**

| A | B | A   B |
|---|---|-------|
| 0 | 0 | 0     |
| 0 | 1 | 1     |
| 1 | 0 | 1     |
| 1 | 1 | 1     |

AND



**A & B**

| A | B | A & B |
|---|---|-------|
| 0 | 0 | 0     |
| 0 | 1 | 0     |
| 1 | 0 | 0     |
| 1 | 1 | 1     |

# Performing Boolean algebra

- **Follow the rules for each operation to compute results**
  - Rules are like those you know from programming

• OR: |    AND: &    NOT: ~    1: True    0: False




$$(1 \mid 0) \& 0 \longrightarrow 1 \& 0 \longrightarrow 0$$

$$(1 \& 1) \& \sim(0 \mid 0) \longrightarrow 1 \& \sim(0) \longrightarrow 1 \& 1 \longrightarrow 1$$

# Exclusive Or (xor)

| <b>A ^ B</b> |          |              |
|--------------|----------|--------------|
| <b>A</b>     | <b>B</b> | <b>A ^ B</b> |
| 0            | 0        | 0            |
| 0            | 1        | 1            |
| 1            | 0        | 1            |
| 1            | 1        | 0            |



- An operation you likely haven't used before:
  - Xor - either A or B, but not both
  - ^ symbol in C
- We can build Xor out of &, |, and ~
  - $A \wedge B = (\sim A \ \& \ B) \mid (A \ \& \ \sim B)$ 
    - (exactly one of A and B is true)
  - $A \wedge B = (A \mid B) \ \& \ \sim(A \ \& \ B)$ 
    - (either is true but not both are true)
- The two definitions are equivalent
  - Produce the same Truth Table

# Practice problem

|   |   | <b>(A &amp; B)   B</b> |  |
|---|---|------------------------|--|
| A | B | (A&B) B                |  |
| 0 | 0 |                        |  |
| 0 | 1 |                        |  |
| 1 | 0 |                        |  |
| 1 | 1 |                        |  |

# Practice problem

| <b>(A &amp; B)   B</b> |   |         |
|------------------------|---|---------|
| A                      | B | (A&B) B |
| 0                      | 0 | 0       |
| 0                      | 1 | 1       |
| 1                      | 0 | 0       |
| 1                      | 1 | 1       |

# Practice problem

| <b>(A &amp; B)   B</b> |   |         |
|------------------------|---|---------|
| A                      | B | (A&B) B |
| 0                      | 0 | 0       |
| 0                      | 1 | 1       |
| 1                      | 0 | 0       |
| 1                      | 1 | 1       |

This is equivalent to B  
(A has no influence on the solution)

# Generalized Boolean algebra

- Boolean operations can be extended to work on collections of bits (i.e., bytes)
- Operations are applied one bit at a time: ***bitwise***

|            |          |            |            |
|------------|----------|------------|------------|
| 01101001   | 01101001 | 01101001   |            |
| & 01010101 | 01010101 | ^ 01010101 | ~ 01010101 |
| <hr/>      | <hr/>    | <hr/>      | <hr/>      |
| 01000001   | 01111101 | 00111100   | 10101010   |

- All of the properties of Boolean algebra still apply
  - Relationships between operations, etc.
- Bitwise operations are usable in C: **&, |, ~, ^**
  - Can operate on any integer type (long, int, short, char, signed or unsigned)

# Warning: bitwise operations are NOT logical operations

- Logical operations in C: **|**, **&**, **!** (logical Or, And, and Not)
  - Only operate on a single bit
    - View 0 as "False"
    - View *anything nonzero* as "True"
    - Always return 0 or 1
  - Short-circuit evaluation: only checks the first operand if that is sufficient
- Examples
  - `!0x41 -> 0x00`                      `!0x00 -> 0x01`                      `!!0x41 -> 0x01`
  - `0x59 & 0x35 -> 0x01`
  - `(p != NULL) & *p` (short circuit evaluation avoids null pointer access)
- Don't confuse the two!! It's a common C mistake

# Break + Practice: C example of bitwise operators

```
unsigned char x = 13;  
unsigned char y = 11;  
unsigned char z = x & y;
```

- What decimal value is in `z` now?
  - Remember: `unsigned char` is an 8-bit value

# Break + Practice: C example of bitwise operators

```
unsigned char x = 13;  
unsigned char y = 11;  
unsigned char z = x & y;
```

- What decimal value is in `z` now?
  - Remember: `unsigned char` is an 8-bit value
  - `x`: 0b00001101
  - `y`: 0b00001011
  - `z`: 0b00001001 -> 9

# Outline

- Integer Operations
  - Addition
  - Negation and Subtraction
  - Multiplication and Division
- **Binary Operations**
  - Boolean Algebra
  - **Shifting**
  - Bit Masks

# Left Shift: $x \ll y$

- Shift bit-vector  $x$  left by  $y$  positions
  - Truncate back to maximum size
    - Throws away bits on the left
  - Fill empty bits with 0
    - Same behavior for signed or unsigned

|              |                                 |
|--------------|---------------------------------|
| Argument $x$ | 00000010                        |
| $\ll 3$      | <del>000</del> 00010 <u>000</u> |

|              |                                 |
|--------------|---------------------------------|
| Argument $x$ | 10100010                        |
| $\ll 3$      | <del>101</del> 00010 <u>000</u> |

- Equivalent to multiplying by  $2^y$ 
  - And then taking modulo (i.e. truncating overflow bits)
- Undefined behavior in C when:
  - $y < 0$ , or  $y \geq \text{bit\_width}(x)$
  - Also when some non-0 bits get shifted off (*probably* they get truncated)

# Right Shift: $x \gg y$

- Shift bit-vector  $x$  right  $y$  positions
  - Throw away extra bits on right
- But how to fill the new bits that open up?
  - Will depend on signed vs unsigned
- Unsigned: Logical shift
  - Always fill with 0's on left
- Signed: Arithmetic shift
  - Replicate most significant bit on left
  - Necessary for two's complement integer representation (sign extension!)
- Undefined behavior in C when:
  - $y < 0$ , or  $y \geq \text{bit\_width}(x)$

|                |                  |
|----------------|------------------|
| Argument $x$   | <u>0</u> 1100010 |
| Logi. $\gg 2$  | <u>00</u> 011000 |
| Arith. $\gg 2$ | <u>00</u> 011000 |

|                |                  |
|----------------|------------------|
| Argument $x$   | <u>1</u> 0100010 |
| Logi. $\gg 2$  | <u>00</u> 101000 |
| Arith. $\gg 2$ | <u>11</u> 101000 |

# Practice shifting in C

```
unsigned char x = 0b10100010;
```

```
x << 3 = ? 0b00010000
```

Steps:

0b10100010**000**

0b**101**00010**000**

```
unsigned char x = 0b10100010;
```

```
x >> 2 = ? 0b00101000
```

Steps:

0b**00**10100010

0b**00**101000**10**

```
signed char x = 0b10100010;
```

```
x >> 2 = ? 0b11101000
```

Steps:

0b**11**10100010

0b**11**101000**10**

## Note:

GCC supports the prefix **0b** for binary literals (like **0x...** for hex) directly in C. This is not part of the C standard! It may not work on other compilers.

# Concept: Not all operations are equally expensive!

- Some operations are pretty simple to perform in hardware
  - E.g., addition, shifting, bitwise operations
  - Also true of doing the same by hand on paper
- Others are much more involved
  - E.g., multiplication, or even more so division
  - Consider long multiplication / long division; quite tedious!
  - Hardware is not doing the exact same thing, but similar principle
- For best performance: swap expensive operations with simple ones!
  - Doesn't work in all cases, but often does when mult/div by constants
  - Don't do this yourself: the compiler will do it for you automatically

# Outline

- Integer Operations
  - Addition
  - Negation and Subtraction
  - Multiplication and Division
- **Binary Operations**
  - Boolean Algebra
  - Shifting
  - **Bit Masks**

# Bit Masking

- How do you manipulate certain bits within a number?
- Combines some of the ideas we've already learned
  - $\sim$ ,  $\&$ ,  $|$ ,  $\ll$ ,  $\gg$
- Steps
  1. Create a "bit mask" which is a pattern to choose certain bits
  2. Use  $\&$  or  $|$  to combine it with your number
  3. Optional: Use  $\gg$  to move the bits to the least significant position

# How to operate on bits

- Selecting bits, use the AND operation

- 1 means to select that bit
- 0 means to not select that bit

Select bottom four bits:

```
num & 0x0F
```

- Writing bits

- Writing a one, use the OR operation

- 1 means to write a one to that position
- 0 is unchanged

Set 6<sup>th</sup> bit to one:

```
num | (1 << 6)
num | (0b01000000)
```

- Writing a zero, use the AND operation

- 0 means to write a zero to that position
- 1 is unchanged

Clear 6<sup>th</sup> bit to zero:

```
num & (~ (1 << 6))
num & (~ (0b01000000))
num & (0b10111111)
```

## Example: selecting bits

- Select bits 2 and 3 from a number

**Input:** 0b01100100

**Mask:** 0b00001100

```
0b01100100
& 0b00001100
-----
0b00000100
```

Finally, shift right by two to get the values in the least significant position:

```
0b00000001
```

In C:

```
result = (input & 0x0C) >> 2;
```

# Practice: determine if bit 7 is 1 or 0

- Need to select bit 7, and shift it to the least-significant position
  - Multiple ways to achieve this. Try it yourself!

# Practice: determine if bit 7 is 1 or 0

- Need to select bit 7, and shift it to the least-significant position
  - Multiple ways to achieve this. Try it yourself!

```
bool result = (value & (1<<7)) >> 7;
```

- Grab the seventh bit, then shift it to the zero position

```
bool result = (value & 0b10000000) >> 7;
```

```
bool result = (value & 0x80) >> 7;
```

- Same as above, but hardcoded (and easier to make mistakes)

```
bool result = (value >> 7) & 1;
```

- Shift desired bit to the zero position, then select it

```
bool result = (value & (1<<7)) != 0;
```

- Select desired bit, then use logical operators to reduce to a single bit

# Outline

- Integer Operations
  - Addition
  - Negation and Subtraction
  - Multiplication and Division
- Binary Operations
  - Boolean Algebra
  - Shifting
  - Bit Masks

# Outline


- Dividing with bit shift
- Bonus material isn't required and won't be on an exam
  - Unless it becomes main lecture material in a different lecture
- Usually the material is just for students who want more depth
  - As is the case here

# Unsigned Power-of-2 Divide with Right Shift

- **Quotient of unsigned by power of 2**

- $u \gg k$  gives  $\lfloor u / 2^k \rfloor$
- Uses logical shift
- Pink part would be remainder / fractional part (right of the point)
  - Shift just drops it: equivalent to rounding **down**

$\lfloor x \rfloor$  : round x down  
 $\lceil x \rceil$  : round x up




|        | Division   | Computed | Hex   | Binary                    |
|--------|------------|----------|-------|---------------------------|
| x      | 15213      | 15213    | 3B 6D | 00111011 01101101         |
| x >> 1 | 7606.5     | 7606     | 1D B6 | <b>0</b> 0011101 10110110 |
| x >> 4 | 950.8125   | 950      | 03 B6 | <b>0000</b> 0011 10110110 |
| x >> 8 | 59.4257813 | 59       | 00 3B | <b>00000000</b> 00111011  |

# Signed Power-of-2 Divide with Shift (Almost)

- **Quotient of signed by power of 2**

- $x \gg k$  gives  $\lfloor x / 2^k \rfloor$
- Uses arithmetic shift
- Also rounds down, again by dropping bits
  - But signed division should round **towards 0!** (that's its math definition)
  - That means rounding **up** for negative numbers!




- **Example, 4 bits:  $-6 / 4 = -1.5$  (should round towards 0, to -1)**

- $1010_2 \gg 2 = \textcolor{red}{11}10_2 = -2_{10}$
- Rounds the wrong way!

# Correct Signed Power-of-2 Divide

- Want  $\lceil x / 2^k \rceil$  (round towards 0)
  - Math identity:  $\lceil x / y \rceil = \lfloor (x + y - 1) / y \rfloor$
  - Compute negative case as  $\lfloor (x + 2^k - 1) / 2^k \rfloor \rightarrow$  gets us correct rounding!
  - Computing both cases in C:  $(x < 0 ? (x + (1 << k) - 1) : x) >> k$ 
    - Biases dividend toward 0

## • Case 1: No rounding




***Biasing has no effect; all affected bits are dropped***

- **Example, 4 bits:  $-8 / 2^2 = -2$       bias =  $(1 << 2) - 1 = 3$** 
  - $(1000 + 0011) >> 2 = 1011 >> 2 = 1110 = -2_{10}$  (correct, no rounding)

# Correct Signed Power-of-2 Divide (Cont.)

## Case 2: Rounding



*Biasing adds 1 to final result; just what we wanted*

- **Example, 4 bits:  $-6 / 2^2 = -1$**       **bias =  $(1 \ll 2) - 1 = 3$** 
  - $(1010 + 0011) \gg 2 = 1101 \gg 2 = \textcolor{red}{11}11 = -1_{10}$  (correct, rounds towards 0)
- **Compiler does that for you (but you need to be able to read it!)**