

CS 61C:

Great Ideas in Computer Architecture
Performance and Floating Point Arithmetic

Instructors:

John Wawrzynek & Vladimir Stojanovic

<http://inst.eecs.berkeley.edu/~cs61c/>

New-School Machine Structures (It's a bit more complicated!)

Software

Hardware

- Parallel Requests
Assigned to computer
e.g., Search "Katz"
- Parallel Threads
Assigned to core
e.g., Lookup, Ads
- Parallel Instructions
>1 instruction @ one time
e.g., 5 pipelined instructions
- Parallel Data
>1 data item @ one time
e.g., Add of 4 pairs of words
- Hardware descriptions
All gates @ one time
- Programming Languages

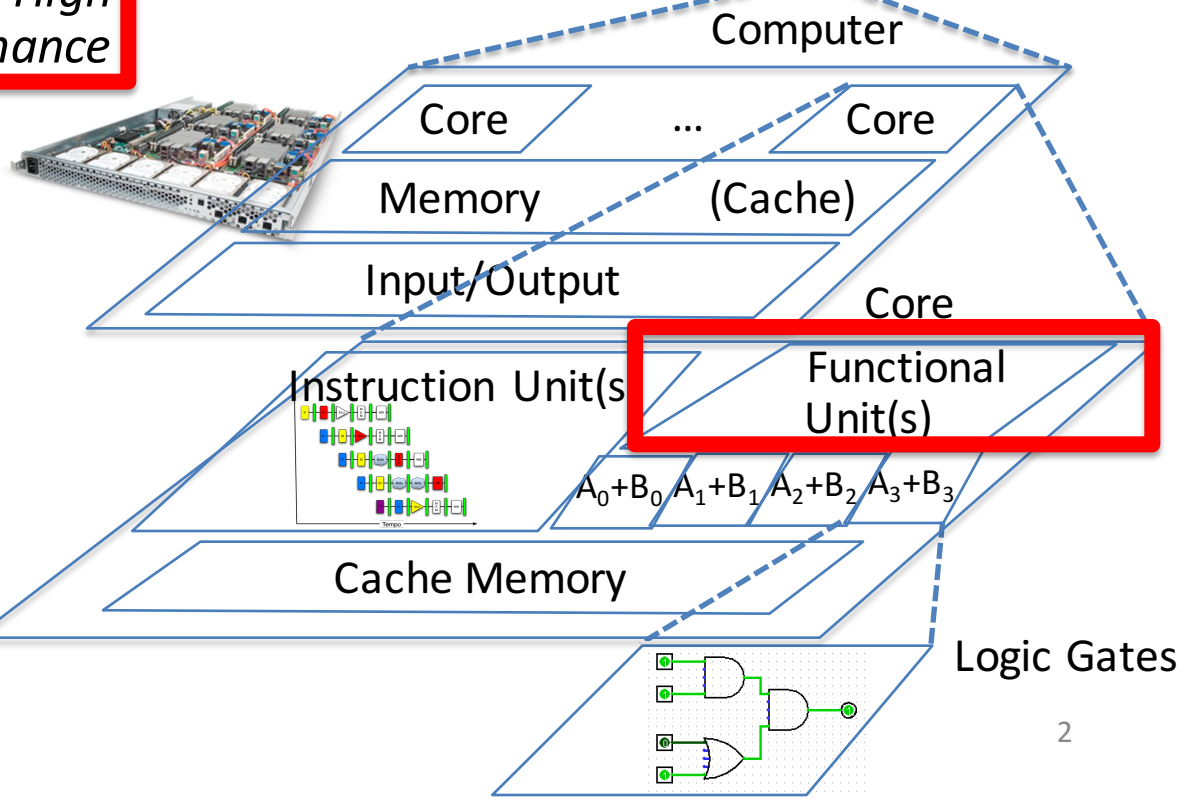
*Harness
Parallelism &
Achieve High
Performance*

Warehouse
Scale
Computer

How do
we know?



Smart
Phone



What is Performance?

- *Latency (or response time or execution time)*
 - Time to complete one task
- *Bandwidth (or throughput)*
 - Tasks completed per unit time

Cloud Performance: Why Application Latency Matters

Server Delay (ms)	Increased time to next click (ms)	Queries/user	Any clicks/user	User satisfaction	Revenue/User
50	--	--	--	--	--
200	500	--	-0.3%	-0.4%	--
500	1200	--	-1.0%	-0.9%	-1.2%
1000	1900	-0.7%	-1.9%	-1.6%	-2.8%
2000	3100	-1.8%	-4.4%	-3.8%	-4.3%

Figure 6.10 Negative impact of delays at Bing search server on user behavior [Brutlag and Schurman 2009].

- Key figure of merit: application responsiveness
 - Longer the delay, the fewer the user clicks, the less the user happiness, and the lower the revenue per user

Defining CPU Performance

- What does it mean to say X is faster than Y?

- Ferrari vs. School Bus?

- 2013 Ferrari 599 GTB

- 2 passengers, quarter mile in 10 secs

- 2013 Type D school bus

- 50 passengers, quarter mile in 20 secs

- *Response Time (Latency)*: e.g., time to travel $\frac{1}{4}$ mile

- *Throughput (Bandwidth)*: e.g., passenger-mi in 1 hour



Defining Relative CPU Performance

- $\text{Performance}_x = 1/\text{Program Execution Time}_x$
- $\text{Performance}_x > \text{Performance}_y \Rightarrow$
 $1/\text{Execution Time}_x > 1/\text{Execution Time}_y \Rightarrow$
 $\text{Execution Time}_y > \text{Execution Time}_x$
- Computer X is N times faster than Computer Y
 $\text{Performance}_x / \text{Performance}_y = N$ or
 $\text{Execution Time}_y / \text{Execution Time}_x = N$
- Bus to Ferrari performance:
 - Program: Transfer 1000 passengers for 1 mile
 - Bus: 3,200 sec, Ferrari: 40,000 sec

Measuring CPU Performance

- Computers use a clock to determine when events take place within hardware
- *Clock cycles*: discrete time intervals
 - aka clocks, cycles, clock periods, clock ticks
- *Clock rate* or *clock frequency*: clock cycles per second (inverse of clock cycle time)
- 3 GigaHertz clock rate
 - => clock cycle time = $1/(3 \times 10^9)$ seconds
 - clock cycle time = 333 picoseconds (ps)

CPU Performance Factors

- To distinguish between processor time and I/O, *CPU time* is time spent in processor
- CPU Time/Program
= Clock Cycles/Program
x Clock Cycle Time
- Or
CPU Time/Program
= Clock Cycles/Program ÷ Clock Rate

Iron Law of Performance

by Emer and Clark

- A program executes instructions
- CPU Time/Program
= Clock Cycles/Program x Clock Cycle Time
= Instructions/Program
x Average Clock Cycles/Instruction
x Clock Cycle Time
- 1st term called *Instruction Count*
- 2nd term abbreviated *CPI* for average *Clock Cycles Per Instruction*
- 3rd term is 1 / Clock rate

Restating Performance Equation

- $$\begin{aligned} \text{Time} &= \frac{\text{Seconds}}{\text{Program}} \\ &= \frac{\text{Instructions}}{\text{Program}} \times \frac{\text{Clock cycles}}{\text{Instruction}} \times \frac{\text{Seconds}}{\text{Clock Cycle}} \end{aligned}$$

What Affects Each Component?

A) Instruction Count, B) CPI, C) Clock Rate

	Affects What? (click in letter of component not affected)
Algorithm	
Programming Language	
Compiler	
Instruction Set Architecture	

What Affects Each Component?

Instruction Count, CPI, Clock Rate

	Affects What?
Algorithm	Instruction Count, CPI
Programming Language	Instruction Count, CPI
Compiler	Instruction Count, CPI
Instruction Set Architecture	Instruction Count, Clock Rate, CPI

Clickers

Computer	Clock frequency	Clock cycles per instruction	#instructions per program
A	1GHz	2	1000
B	2GHz	5	800
C	500MHz	1.25	400
D	5GHz	10	2000

- Which computer has the highest performance for a given program?

Workload and Benchmark

- *Workload*: Set of programs run on a computer
 - Actual collection of applications run or made from real programs to approximate such a mix
 - Specifies programs, inputs, and relative frequencies
- *Benchmark*: Program selected for use in comparing computer performance
 - Benchmarks form a workload
 - Usually standardized so that many use them

SPEC

(System Performance Evaluation Cooperative)

- Computer Vendor cooperative for benchmarks, started in 1989
- SPEC CPU2006
 - 12 Integer Programs
 - 17 Floating-Point Programs
- Often turn into number where bigger is faster
- *SPECratio*: reference execution time on old reference computer divide by execution time on new computer to get an effective speed-up

SPECINT2006 on AMD Barcelona

Description	Instruction Count (B)	CPI	Clock cycle time (ps)	Execution Time (s)	Reference Time (s)	SPEC-ratio
Interpreted string processing	2,118	0.75	400	637	9,770	15.3
Block-sorting compression	2,389	0.85	400	817	9,650	11.8
GNU C compiler	1,050	1.72	400	724	8,050	11.1
Combinatorial optimization	336	10.0	400	1,345	9,120	6.8
Go game	1,658	1.09	400	721	10,490	14.6
Search gene sequence	2,783	0.80	400	890	9,330	10.5
Chess game	2,176	0.96	400	837	12,100	14.5
Quantum computer simulation	1,623	1.61	400	1,047	20,720	19.8
Video compression	3,102	0.80	400	993	22,130	22.3
Discrete event simulation library	587	2.94	400	690	6,250	9.1
Games/path finding	1,082	1.79	400	773	7,020	9.1
XML parsing	1,058	2.70	400	1,143	6,900	6.0

Summarizing Performance ...

System	Rate (Task 1)	Rate (Task 2)
A	10	20
B	20	10

Clickers: Which system is faster?

A: System A

B: System B

C: Same performance

D: Unanswerable question!

... Depends Who's Selling

System	Rate (Task 1)	Rate (Task 2)	Average
A	10	20	15
B	20	10	15

Average throughput

System	Rate (Task 1)	Rate (Task 2)	Average
A	0.50	2.00	1.25
B	1.00	1.00	1.00

Throughput relative to B

System	Rate (Task 1)	Rate (Task 2)	Average
A	1.00	1.00	1.00
B	2.00	0.50	1.25

Throughput relative to A

Summarizing SPEC Performance

- Varies from 6x to 22x faster than reference computer

- *Geometric mean* of ratios:
N-th root of product
of N ratios

$$\sqrt[n]{\prod_{i=1}^n \text{Execution time ratio}_i}$$

- Geometric Mean gives same relative answer no matter what computer is used as reference
- Geometric Mean for Barcelona is 11.7

Administrivia

- HW3 released, due 11/01@23:59:59
- Project 3-2 released, due 11/01@23:59:59
- Midterm 2 on 11/10 in class
 - For midterm 2 re-scheduling (not dsp) contact Fred and William

CodeBears Announcement

- “CodeBears is a rapidly growing organization on campus aimed at providing members with the opportunity to compete with each other and against teams in other universities in algorithmic software competitions that require strong problem solving abilities. Winners of our intra-college programming competitions will be able to meet with industry leaders, network with founders of various organizations, and represent UCB at inter-college competitions. We hope to create a new hacking culture at Cal! Learn more about us at:
<https://callink.berkeley.edu/organization/codebears>
- We already have over 100 registered members and had a tremendous turnout rate at our first GM last Friday! Visit our Facebook page at:<https://www.facebook.com/CodeBears> to access last Friday’s slides and learn how to join CodeBears!

Review of Numbers

- **Computers are made to deal with numbers**
- **What can we represent in N bits?**
 - **2^N things, and no more!** They could be...

- **Unsigned integers:**

0 to $2^N - 1$

(for $N=32$, $2^N - 1 = 4,294,967,295$)

- **Signed Integers (Two's Complement)**

$-2^{(N-1)}$ to $2^{(N-1)} - 1$

(for $N=32$, $2^{(N-1)} = 2,147,483,648$)



What about other numbers?

1. Very large numbers? (seconds/millennium)
⇒ $31,556,926,000_{10}$ ($3.1556926_{10} \times 10^{10}$)
2. Very small numbers? (Bohr radius)
⇒ $0.0000000000529177_{10}\text{m}$ ($5.29177_{10} \times 10^{-11}$)
3. Numbers with both integer & fractional parts?
⇒ 1.5

First consider #3.

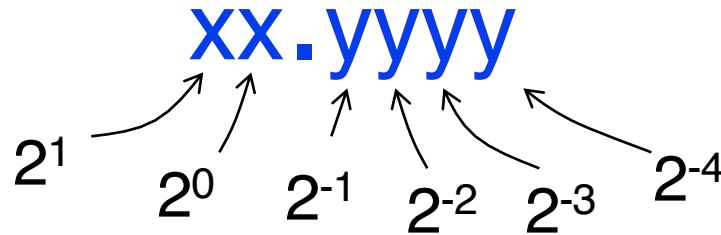
...our solution will also help with #1 and #2.



Representation of Fractions

“Binary Point” like decimal point signifies boundary between integer and fractional parts:

Example 6-bit representation:



$$10.1010_{\text{two}} = 1 \times 2^1 + 1 \times 2^{-1} + 1 \times 2^{-3} = 2.625_{\text{ten}}$$

If we assume “fixed binary point”, range of 6-bit representations with this format:

0 to 3.9375 (almost 4)



Fractional Powers of 2

i	2⁻ⁱ	
0	1.0	1
1	0.5	1/2
2	0.25	1/4
3	0.125	1/8
4	0.0625	1/16
5	0.03125	1/32
6	0.015625	
7	0.0078125	
8	0.00390625	
9	0.001953125	
10	0.0009765625	
11	0.00048828125	
12	0.000244140625	
13	0.0001220703125	
14	0.00006103515625	
15	0.000030517578125	



Representation of Fractions with Fixed Pt.

What about addition and multiplication?

Addition is straightforward:

	01.100	1.5 _{ten}		
	+ 00.100	0.5 _{ten}		
	<u>10.000</u>	2.0 _{ten}		
			01.100	1.5 _{ten}
			<u>00.100</u>	0.5 _{ten}

Multiplication a bit more complex:

		00	000
		000	00
		0110	0
		00000	
		00000	
		<u>00000</u>	
		0000110000	

Where's the answer, 0.11? (need to remember where point is)



Representation of Fractions

So far, in our examples we used a “fixed” binary point. What we really want is to “float” the binary point. Why?

Floating binary point most effective use of our limited bits (and thus more accuracy in our number representation):

example: put 0.1640625_{ten} into binary. Represent with 5-bits choosing where to put the binary point.

... 000000.001010100000...



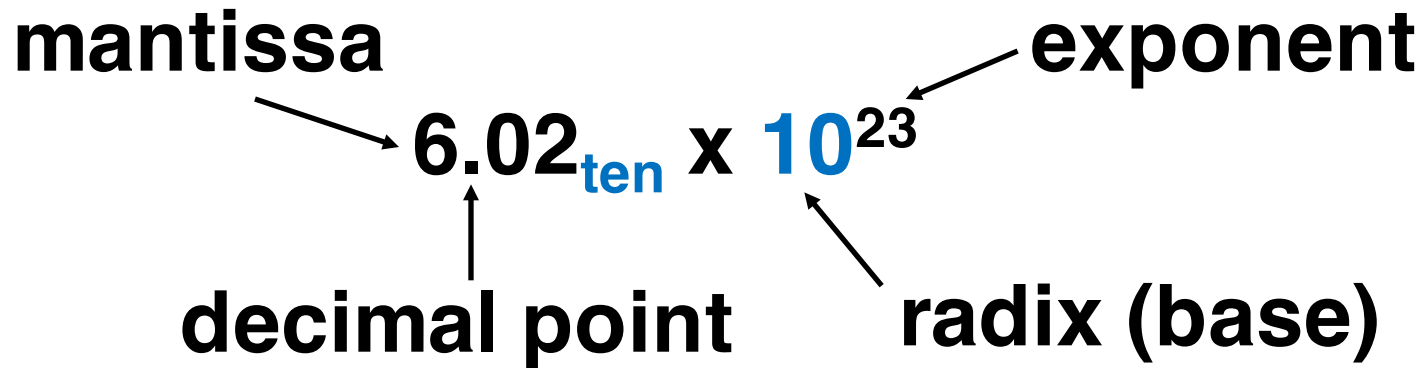
Store these bits and keep track of the binary point 2 places to the left of the MSB

Any other solution would lose accuracy!

With floating-point rep., each numeral carries an exponent field recording the whereabouts of its binary point.

The binary point **can be outside** the stored bits, so very large and small numbers can be represented.

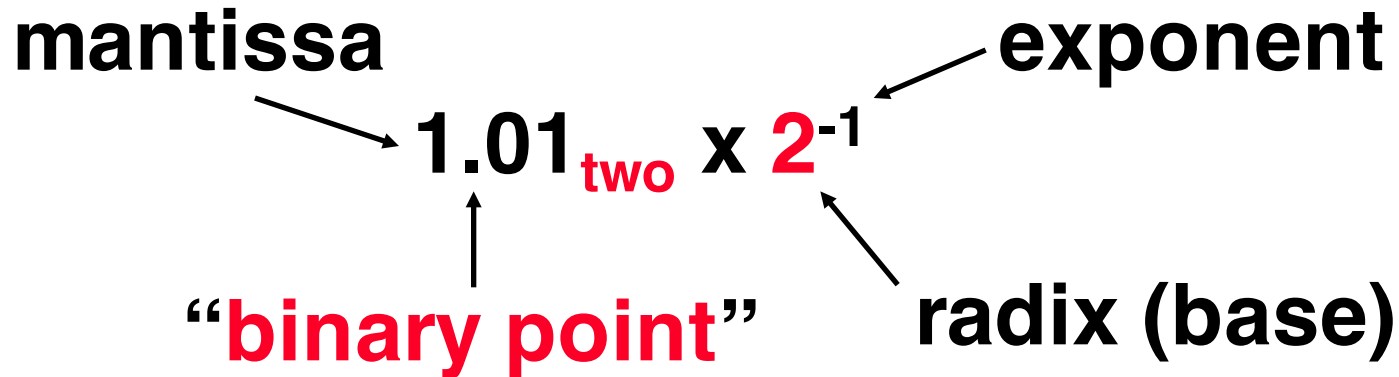
Scientific Notation (in Decimal)



- Normalized form: no leading 0s (exactly one digit to left of decimal point)
- Alternatives to representing 1/1,000,000,000
 - Normalized: 1.0×10^{-9}
 - Not normalized: $0.1 \times 10^{-8}, 10.0 \times 10^{-10}$



Scientific Notation (in Binary)

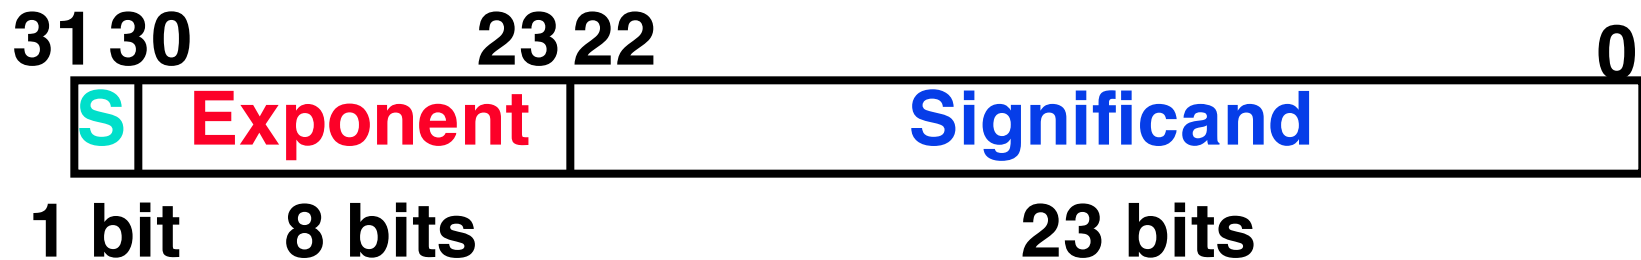


- Computer arithmetic that supports it called floating point, because it represents numbers where the binary point is not fixed, as it is for integers
 - Declare such variable in C as `float`
 - `double` for double precision.



Floating-Point Representation (1/2)

- Normal format: $+1.x_{\text{two}} \dots x_{\text{two}} * 2^{y_{\text{two}} \dots y_{\text{two}}}$
- Multiple of Word Size (32 bits)



- S represents Sign
Exponent represents y 's
Significand represents x 's
- Represent numbers as small as $2.0_{\text{ten}} \times 10^{-38}$ to as large as $2.0_{\text{ten}} \times 10^{38}$



Floating-Point Representation (2/2)

- What if result too large?

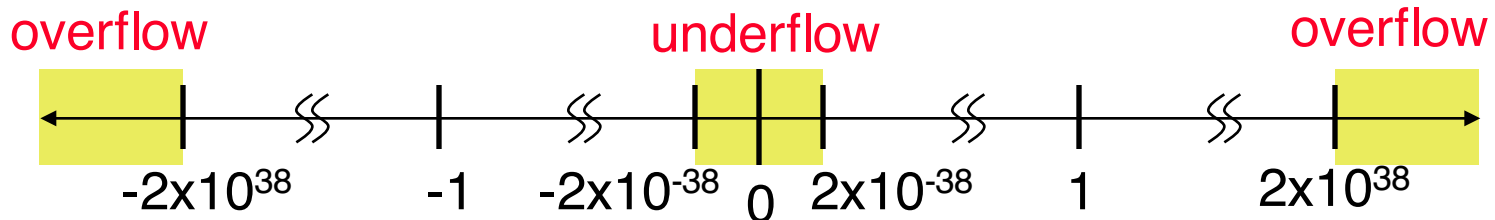
($> 2.0 \times 10^{38}$, $< -2.0 \times 10^{38}$)

- **Overflow!** \Rightarrow Exponent larger than represented in 8-bit Exponent field

- What if result too small?

(>0 & $< 2.0 \times 10^{-38}$, <0 & $> -2.0 \times 10^{-38}$)

- **Underflow!** \Rightarrow Negative exponent larger than represented in 8-bit Exponent field



- What would help reduce chances of overflow and/or underflow?



IEEE 754 Floating Point Standard (2/3)

- **IEEE 754 uses “biased exponent” representation**
 - Designers wanted FP numbers to be used even if no FP hardware; e.g., sort records with FP numbers using integer compares
 - Wanted bigger (integer) exponent field to represent bigger numbers
 - 2’s complement poses a problem (because negative numbers look bigger)
 - Use just magnitude and offset by half the range



IEEE 754 Floating Point Standard (3/3)

- Called **Biased Notation**, where bias is number subtracted to get final number
 - IEEE 754 uses bias of 127 for single prec.
 - Subtract 127 from Exponent field to get actual value for exponent

• Summary (single precision):



- $(-1)^S \times (1 + \text{Significand}) \times 2^{(\text{Exponent}-127)}$

- Double precision identical, except with exponent bias of 1023 (half, quad similar)



“Father” of the Floating point standard

IEEE Standard 754 for Binary Floating-Point Arithmetic.



Prof. Kahan

1989
ACM Turing
Award Winner!

www.cs.berkeley.edu/~wkahan/ieee754status/754story.html



Clickers

- Guess this Floating Point number:

1 1000 0000 1000 0000 0000 0000 0000 000

A: -1×2^{128}

B: $+1 \times 2^{-128}$

C: -1×2^1

D: $+1.5 \times 2^{-1}$

E: -1.5×2^1

Representation for $\pm \infty$

- In FP, divide by 0 should produce $\pm \infty$, not overflow.
- Why?
 - OK to do further computations with ∞
E.g., $X/0 > Y$ may be a valid comparison
- IEEE 754 represents $\pm \infty$
 - Most positive exponent reserved for ∞
 - Significands all zeroes



Representation for 0

- **Represent 0?**
 - **exponent all zeroes**
 - **significand all zeroes**
 - **What about sign? Both cases valid**

+0: 0 00000000 000000000000000000000000000000

-0: 1 00000000 000000000000000000000000000000



Special Numbers

- What have we defined so far?
(Single Precision)

Exponent	Significand	Object
0	0	0
0	<u>nonzero</u>	<u>???</u>
1-254	anything	+/- fl. pt. #
255	0	+/- ∞
255	<u>nonzero</u>	<u>???</u>

- Professor Kahan had clever ideas:
 - Wanted to use $\text{Exp}=0,255$ & $\text{Sig}\neq 0$



Representation for Not a Number

- What do I get if I calculate $\text{sqrt}(-4.0)$ or $0/0$?
 - If ∞ not an error, these shouldn't be either
 - Called Not a Number (NaN)
 - Exponent = 255, Significand nonzero
- Why is this useful?
 - Hope NaNs help with debugging?
 - They contaminate: $\text{op}(\text{NaN}, X) = \text{NaN}$
 - Can use the significand to identify which!



Representation for Denorms (1/2)

- **Problem: There's a gap among representable FP numbers around 0**

- **Smallest representable pos num:**

$$a = 1.0\dots_2 * 2^{-126} = 2^{-126}$$

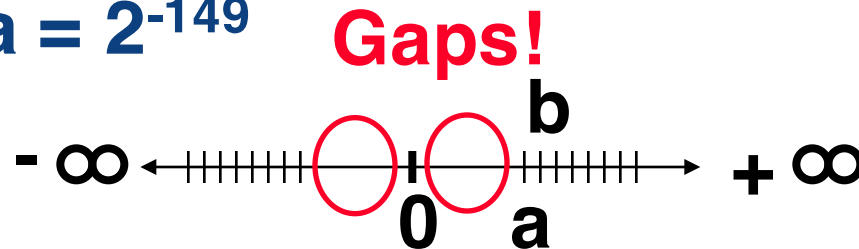
- **Second smallest representable pos num:**

$$\begin{aligned} b &= 1.000\dots1_2 * 2^{-126} \\ &= (1 + 0.00\dots1_2) * 2^{-126} \\ &= (1 + 2^{-23}) * 2^{-126} \\ &= 2^{-126} + 2^{-149} \end{aligned}$$

**Normalization
and implicit 1
is to blame!**

$$a - 0 = 2^{-126}$$

$$b - a = 2^{-149}$$



Representation for Denorms (2/2)

- **Solution:**

- We still haven't used Exponent = 0, Significand nonzero
- DENormalized number: no (implied) leading 1, **implicit exponent = -126.**
- Smallest representable pos num:

$$a = 2^{-149}$$

- Second smallest representable pos num:

$$b = 2^{-148}$$



Special Numbers Summary

- Reserve exponents, significands:

Exponent	Significand	Object
0	0	0
0	<u>nonzero</u>	<u>Denorm</u>
1-254	anything	+/- fl. pt. #
255	<u>0</u>	<u>+/- ∞</u>
255	<u>nonzero</u>	<u>NaN</u>



Conclusion

- **Floating Point lets us:**

- Represent numbers containing both integer and fractional parts; makes efficient use of available bits.
- Store **approximate** values for very large and very small #s.

- **IEEE 754 Floating-Point Standard** is most widely accepted attempt to standardize interpretation of such numbers (Every desktop or server computer sold since ~1997 follows these conventions)

Exponent tells Significand how much (2^i) to count by (... , 1/4, 1/2, 1, 2, ...)

Can store NaN, $\pm \infty$

- **Summary (single precision):**



- $(-1)^S \times (1 + \text{Significand}) \times 2^{(\text{Exponent}-127)}$

- **Double precision identical, except with exponent bias of 1023 (half, quad similar)**



And In Conclusion, ...

- Time (seconds/program) is measure of performance

$$= \frac{\text{Instructions}}{\text{Program}} \times \frac{\text{Clock cycles}}{\text{Instruction}} \times \frac{\text{Seconds}}{\text{Clock Cycle}}$$

- Floating-point representations hold approximations of real numbers in a finite number of bits