September 10, 2024

Announcements

- Exam 1
- HW2

Goals

- floating P
- Lin. Systems Th.

Review

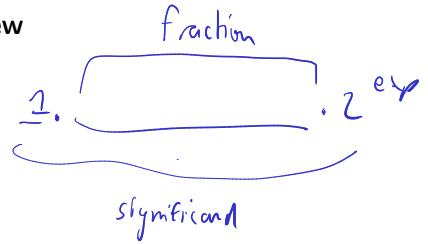

subnormals: based on sp. exp.
- no implicit leading one

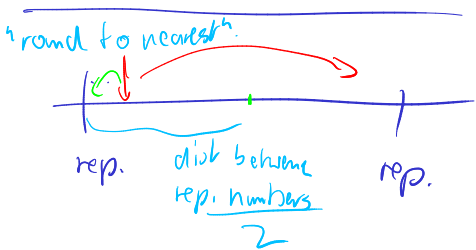

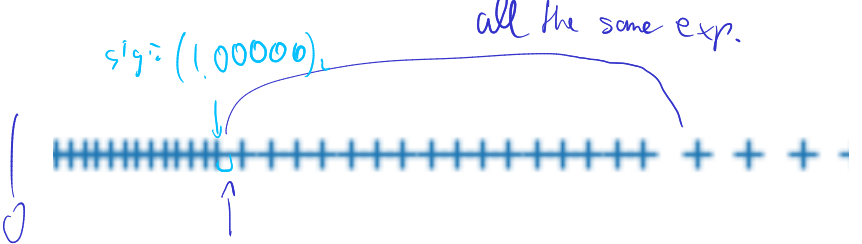
all the same exp.
sig: $(1.00000)_2$

Rounding Modes

How is rounding performed? (Imagine trying to represent π.)

$$(\underbrace{1.1101010}_{\text{representable}}11)_2$$

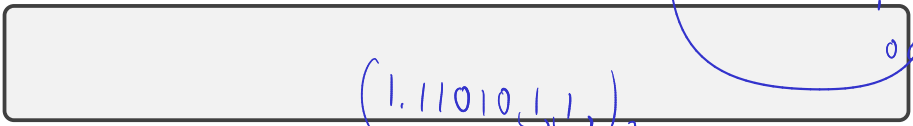

What is done in case of a tie? $0.5 = (0.1)_2$ ("Nearest"?)

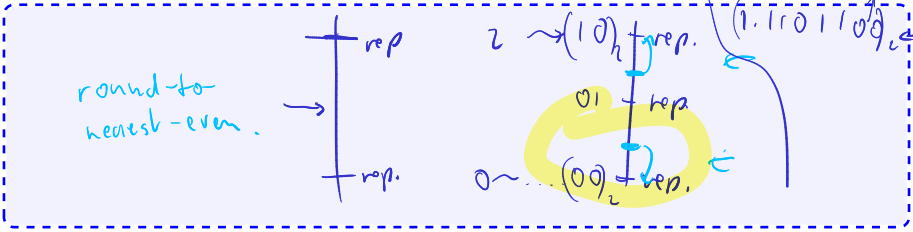

Demo: Density of Floating Point Numbers [cleared]

Rounding Modes

How is rounding performed? (Imagine trying to represent π.)

$$\big(\underbrace{1.1101010}_{\text{representable}}11\big)_2$$

What is done in case of a tie? $0.5 = (0.1)_2$ ("Nearest"?)

Up or down? It turns out that picking the same direction every time introduces *bias*. Trick: *round-to-even*.

$$0.5 \to 0, \qquad 1.5 \to 2$$

Demo: Density of Floating Point Numbers [cleared]

Smallest Numbers Above...

1.0000

- What is smallest FP number > 1? Assume 4 stored bits (5 total) in the significand.

1.0001

What's the smallest FP number > 1024 in that same system?

$(1.000)_2 \cdot 2^{10}$

Can we give that number a name?

Unit Roundoff

Unit roundoff or *machine precision* or *machine epsilon* or $\varepsilon_{\text{mach}}$ is...

smallest number $\varepsilon > 0$ such that

$$fl(1+\varepsilon) \neq 1$$

For round-to-even:

$$\varepsilon_{\text{mach}} = \frac{ULP}{2}$$

FP: Relative Rounding Error

What does this say about the relative error incurred in floating point calculations?

$$\hat{x} = x + x\varepsilon_{mach} = x(1+\varepsilon_{mach})$$

$\hat{x}$: smallest number $> x$

$$\frac{|\hat{x}-x|}{|x|} = \frac{|x(1+\varepsilon_{mach})-x|}{|x|} = \varepsilon_{mach}$$

$$\frac{x_{sig}\cdot 2^{x_{exp}} + \hat{x}_{sig}\cdot 2^{x_{exp}}}{x_{sig}\cdot 2^{x_{exp}}}$$

FP: Machine Epsilon

What's machine epsilon for double-precision floating point with round-to-nearest? (52 stored bits in the significand, 53 total)

Demo: Floating point and the Harmonic Series [cleared]

Problems with FP Addition

What happens if you subtract two numbers of very similar magnitude?
As an example, consider $a = (1.1011)_2 \cdot 2^0$ and $b = (1.1010)_2 \cdot 2^0$.

$\rightarrow a = (1.1011)_2 \cdots$

$\rightarrow b = (1.1010)_2$

$a - b = (0.0001)_2$

$(1.???? \; 0000)_2 \cdot 2^{-4}$

Demo: Catastrophic Cancellation [cleared]

Supplementary Material

- Josh Haberman, Floating Point Demystified, Part 1
- David Goldberg, What every computer programmer should know about floating point
- Evan Wallace, Float Toy
- Julia Evans, Examples of Floating Point Problems, 2022

Outline