## Rounding errors

## Example

Show demo: "Waiting for 1 ".
Determine the double-precision machine representation for 0.1

| $\# \times \mathbf{2}$ | Integer <br> part | Fractional <br> part |
| :--- | :--- | :--- |
| 0.2 | 0 | 0.2 |
| 0.4 | 0 | 0.4 |
| 0.8 | 0 | 0.8 |
| 1.6 | 1 | 0.6 |
| 1.2 | 1 | 0.2 |
| 0.4 | 0 | 0.4 |
| 0.8 | 0 | 0.8 |
| 1.6 | 1 | 0.6 |
| 1.2 | 1 | 0.2 |

$$
\begin{aligned}
& 0.1=(0.000110011 \overline{0011} \ldots)_{2} \\
& =(1.100110011 \ldots)_{2} \times 2^{-4} \\
& \quad s=0 \\
& \quad f=100110011 \ldots 00110011010 \\
& m=-4 \\
& c=m+1023=1019=(01111111011)_{2}
\end{aligned}
$$

$00111111101110011 \ldots 0011 \ldots 0011010$
(52-bit)
Roundoff error in its basic form!

## Machine floating point number

- Not all real numbers can be exactly represented as a machine floating-point number.
- Consider a real number in the normalized floating-point form:

$$
x= \pm 1 . b_{1} b_{2} b_{3} \ldots b_{n} \ldots \times 2^{m}
$$

- The real number $x$ will be approximated by either $x_{-}$or $x_{+}$, the nearest two machine floating point numbers.


Without loss of generality, let's see what happens when trying to represent a positive machine floating point number:

Exact number: $x=1 . b_{1} b_{2} b_{3} \ldots b_{n} \ldots \times 2^{m}$
$x_{-}=1 . b_{1} b_{2} b_{3} \ldots b_{n} \times 2^{m}$ (rounding by chopping)
$x_{+}=1 . b_{1} b_{2} b_{3} \ldots b_{n} \times 2^{m}+\underbrace{0.000 \ldots 01}_{\epsilon_{m}} \times 2^{m}$


Exact number: $x=1 . b_{1} b_{2} b_{3} \ldots b_{n} \ldots \times 2^{m}$

$$
\begin{aligned}
& x_{-}=1 . b_{1} b_{2} b_{3} \ldots b_{n} \times 2^{m} \\
& x_{+}=1 . b_{1} b_{2} b_{3} \ldots b_{n} \times 2^{m}+\underbrace{0.000 \ldots 01}_{\epsilon_{m}} \times 2^{m}
\end{aligned}
$$

Gap between $x_{+}$and $x_{-}:\left|x_{+}-x_{-}\right|=\epsilon_{m} \times 2^{m}$

Examples for single precision:
$x_{+}$and $x_{-}$of the form $q \times 2^{-10}:\left|x_{+}-x_{-}\right|=2^{-33} \approx 10^{-10}$
$x_{+}$and $x_{-}$of the form $q \times 2^{4}:\left|x_{+}-x_{-}\right|=2^{-19} \approx 2 \times 10^{-6}$
$x_{+}$and $x_{-}$of the form $q \times 2^{20}:\left|x_{+}-x_{-}\right|=2^{-3} \approx 0.125$
$x_{+}$and $x_{-}$of the form $q \times 2^{60}:\left|x_{+}-x_{-}\right|=2^{37} \approx 10^{11}$
The interval between successive floating point numbers is not uniform: the interval is smaller as the magnitude of the numbers themselves is smaller, and it is bigger as the numbers get bigger.

## Gap between two successive machine floating point numbers

A "toy" number system can be represented as $x= \pm 1 . b_{1} b_{2} \times 2^{m}$ for $m \in[-4,4]$ and $b_{i} \in\{0,1\}$.
$(1.00)_{2} \times 2^{0}=1$
$(1.00)_{2} \times 2^{1}=2$
$(1.00)_{2} \times 2^{2}=4.0$
$(1.01)_{2} \times 2^{0}=1.25$
$(1.01)_{2} \times 2^{1}=2.5$
$(1.01)_{2} \times 2^{2}=5.0$
$(1.10)_{2} \times 2^{0}=1.5$
$(1.10)_{2} \times 2^{1}=3.0$
$(1.10)_{2} \times 2^{2}=6.0$
$(1.11)_{2} \times 2^{0}=1.75$
$(1.11)_{2} \times 2^{1}=3.5$
$(1.11)_{2} \times 2^{2}=7.0$
$(1.00)_{2} \times 2^{3}=8.0$
$(1.00)_{2} \times 2^{4}=16.0$
$(1.00)_{2} \times 2^{-1}=0.5$
$(1.01)_{2} \times 2^{3}=10.0$
$(1.01)_{2} \times 2^{4}=20.0$
$(1.01)_{2} \times 2^{-1}$
$(1.10)_{2} \times 2^{3}=12.0$
$(1.10)_{2} \times 2^{4}=24.0$
$=0.625$
$(1.11)_{2} \times 2^{3}=14.0$
$(1.11)_{2} \times 2^{4}=28.0$
$(1.10)_{2} \times 2^{-1}=0.75$
$(1.11)_{2} \times 2^{-1}$
$\left.(1.00)_{2} \times 2^{-2}=0.25 \quad(1.00)_{2} \times 2^{-3}=0.125 \quad=0.875\right)_{2} \times 2^{-4}=0.0625$
$(1.01)_{2} \times 2^{-2}=0.3125 \quad(1.01)_{2} \times 2^{-3}=0.15625$
$(1.01)_{2} \times 2^{-4}=0.078125$
$(1.10)_{2} \times 2^{-2}=0.375$
$(1.10)_{2} \times 2^{-3}=0.1875$
$(1.10)_{2} \times 2^{-4}=0.09375$
$(1.11)_{2} \times 2^{-2}=0.4375 \quad(1.11)_{2} \times 2^{-3}=0.21875 \quad(1.11)_{2} \times 2^{-4}=0.109375$

## Rounding

The process of replacing $x$ by a nearby machine number is called rounding, and the error involved is called roundoff error.


Round by chopping: $f l(x)=x_{-}$

|  | $x$ is positive number | $x$ is negative number |
| :--- | :---: | :---: |
| Round up (ceil) | $f l(x)=x_{+}$ <br> Rounding towards $+\infty$ | $f l(x)=x_{-}$ <br> Rounding towards zero |
| Round down (floor) | $f l(x)=x_{-}$ <br> Rounding towards zero | $f l(x)=x_{+}$ <br> Rounding towards $-\infty$ |

Round to nearest: either round up or round down, whichever is closer

## Rounding (roundoff) errors

Consider rounding by chopping:

- Absolute error:

$$
\begin{gathered}
|\mathrm{fl}(x)-x| \leq\left|x_{+}-x_{-}\right|=\epsilon_{m} \times 2^{m} \\
|\mathrm{fl}(x)-x| \leq \epsilon_{m} \times 2^{m}
\end{gathered}
$$

- Relative error:

$$
\begin{gathered}
\frac{|\mathrm{fl}(x)-x|}{|x|} \leq \frac{\epsilon_{m} \times 2^{m}}{1 . b_{1} b_{2} b_{3} \ldots b_{n} \ldots \times 2^{m}} \\
\frac{|\mathrm{fl}(x)-x|}{|x|} \leq \epsilon_{m}
\end{gathered}
$$

## Rounding (roundoff) errors

$$
\frac{|\tilde{x}-x|}{|x|} \leq 2^{-23} \approx 1.2 \times 10^{-7} \quad \frac{|\tilde{x}-x|}{|x|} \leq 2^{-52} \approx 2.2 \times 10^{-16}
$$

Single precision: Floating-point math consistently introduces relative errors of about $10^{-7}$. Hence, single precision gives you about 7 (decimal) accurate digits.

Double precision: Floating-point math consistently introduces relative errors of about $10^{-16}$. Hence, double precision gives you about 16 (decimal) accurate digits.

## Iclicker question

Assume you are working with IEEE single-precision numbers. Find the smallest number $a$ that satisfies

$$
2^{8}+a \neq 2^{8}
$$

A) $2^{-1074}$
B) $2^{-1022}$
C) $2^{-52}$
D) $2^{-15}$
E) $2^{-8}$

Demo

Arithmetic with machine numbers

## Mathematical properties of FP operations

Not necessarily associative:
For some $x, y, z$ the result below is possible:

$$
(x+y)+z \neq x+(y+z)
$$

## Not necessarily distributive:

For some $x, y, z$ the result below is possible:

```
In [5]:(np.pi+1e100)-1e100
Out[5]: 0.0
```

In [6]: (np.pi)+(1e100-1e100)
Out [6]: 3.141592653589793
In [7]: $\begin{aligned} & \mathrm{b} \\ & \mathrm{a}=1 \mathrm{e}=1 \mathrm{e}= \\ & \mathrm{a}=1\end{aligned}$
print $(a+(b-b))$
print ( $(\mathrm{a}+\mathrm{b})-\mathrm{b})$
100.0
0.0
In [3]: $\begin{aligned} & \operatorname{print}(100 *(0.1+0.2)) \\ & \operatorname{print}(100 * 0.1+100 * 0.2)\end{aligned}$
30.000000000000004 30.0
$\operatorname{In}[4]: 100 *(0.1+0.2)==100 * 0.1+100 * 0.2$
Out[4]: False

Not necessarily cumulative:
Repeatedly adding a very small number to a large number may do nothing

## Floating point arithmetic (basic idea)

$$
x=(-1)^{s} 1 . f \times 2^{m}=\begin{array}{|l|l|l|}
\hline s & c & f \\
\hline
\end{array}
$$

- First compute the exact result
- Then round the result to make it fit into the desired precision
- $x+y=f l(x+y)$
- $x \times y=f l(x \times y)$


## Floating point arithmetic

Consider a number system such that $x= \pm 1 . b_{1} b_{2} b_{3} \times 2^{m}$ for $m \in[-4,4]$ and $b_{i} \in\{0,1\}$.

Rough algorithm for addition and subtraction:

1. Bring both numbers onto a common exponent
2. Do "grade-school" operation
3. Round result

- Example 1: No rounding needed

$$
\begin{aligned}
& a=(1.101)_{2} \times 2^{1} \\
& b=(1.001)_{2} \times 2^{1} \\
& c=a+b=(10.110)_{2} \times 2^{1}=(1.011)_{2} \times 2^{2}
\end{aligned}
$$

## Floating point arithmetic

Consider a number system such that $x= \pm 1 . b_{1} b_{2} b_{3} \times 2^{m}$ for $m \in[-4,4]$ and $b_{i} \in\{0,1\}$.

- Example 2: Require rounding

$$
\begin{aligned}
& a=(1.101)_{2} \times 2^{0} \\
& b=(1.000)_{2} \times 2^{0} \\
& c=a+b=(10.101)_{2} \times 2^{0} \approx(1.010)_{2} \times 2^{1}
\end{aligned}
$$

- Example 3:

$$
\begin{aligned}
a & =(1.100)_{2} \times 2^{1} \\
b & =(1.100)_{2} \times 2^{-1} \\
c & =a+b=(1.100)_{2} \times 2^{1}+(0.011)_{2} \times 2^{1}=(1.111)_{2} \times 2^{1}
\end{aligned}
$$

## Floating point arithmetic

Consider a number system such that $x= \pm 1 . b_{1} b_{2} b_{3} b_{4} \times 2^{m}$ for $m \in[-4,4]$ and $b_{i} \in\{0,1\}$.

- Example 4:

$$
\begin{aligned}
& a=(1.1011)_{2} \times 2^{1} \\
& b=(1.1010)_{2} \times 2^{1} \\
& c=a-b=(0.0001)_{2} \times 2^{1}
\end{aligned}
$$

Or after normalization: $\quad c=(1 . ? ? ? ?)_{2} \times 2^{-3}$
Unfortunately there is not data to indicate what the missing digits should be. The effect is that the number of significant digits in the result is reduced. Machine fills them with its best guess, which is often not good (usually what is called spurious zeros). This phenomenon is called Catastrophic Cancellation.

## Loss of significance

Assume $a \gg b$. For example
$a=1 . a_{1} a_{2} a_{3} a_{4} a_{5} a_{6} \ldots a_{n} \ldots \times 2^{0}$
$b=1 . b_{1} b_{2} b_{3} b_{4} b_{5} b_{6} \ldots b_{n} \ldots \times 2^{-8}$

In Single Precision (without loss of generality):
$f l(a)=1 . a_{1} a_{2} a_{3} a_{4} a_{5} a_{6} \ldots a_{22} a_{23} \times 2^{0}$
$f l(b)=1 . b_{1} b_{2} b_{3} b_{4} b_{5} b_{6} \ldots b_{22} b_{23} \times 2^{-8}$

1. $a_{1} a_{2} a_{3} a_{4} a_{5} a_{6} a_{7} a_{8} a_{9} \ldots a_{22} a_{23} \times 2^{0}$
$+0.00000001 b_{1} b_{2} b_{3} b_{4} b_{5} \ldots b_{14} b_{15} \times 2^{0}$

In this example, the result $f l(a+b)$ includes 15 bits of precision from $f l(b)$. Lost precision!

## Cancellation

$$
\begin{aligned}
& a=1 . a_{1} a_{2} a_{3} a_{4} a_{5} a_{6} \ldots a_{n} \ldots \times 2^{m 1} \\
& b=1 . b_{1} b_{2} b_{3} b_{4} b_{5} b_{6} \ldots b_{n} \ldots \times 2^{m 2}
\end{aligned}
$$

Suppose $a \approx b$ and single precision (without loss of generality)

$$
\begin{aligned}
& a=1 . a_{1} a_{2} a_{3} a_{4} a_{5} a_{6} \ldots a_{20} a_{21} 10 a_{24} a_{25} a_{26} a_{27} \ldots \times 2^{m} \\
& b=1 . a_{1} a_{2} a_{3} a_{4} a_{5} a_{6} \ldots a_{20} a_{21} 11 b_{24} b_{25} b_{26} b_{27} \ldots \times 2^{m} \\
& \substack{\text { Lost due to } \\
\text { rounding }}
\end{aligned}
$$

$$
f l(b-a)=0.0000 \ldots 0001 \times 2^{m}=1 . ? ? ? ? ? ? \ldots ? ? \times 2^{-n+m}
$$

$$
f l(b-a)=1.000 \ldots 00 \times 2^{-n+m}
$$

Not significant bits (precision lost, not due to $f l(b-a)$ but due to rounding of a, $b$ from the beginning

Example of cancellation:
Suppose

$$
\begin{aligned}
& a=1.1011 a_{5} a_{6} a_{7} \ldots \times 2^{1} \\
& b=1.1010 b_{5} b_{6} b_{7} \ldots \times 2^{1}
\end{aligned}
$$

using machine where $n=4 \Rightarrow a=1.1011 \times 2^{\prime}$

$$
b=1.1010 \times 2^{1}
$$

$$
\begin{aligned}
& \text { machine } 0.0001 \times 2^{\prime} \text { when done by "hand } \\
& \text { resulting } \\
& \text { with } \\
& \text { cancellation } \\
& \begin{array}{l}
1.0000 \times 2^{-3} \\
\text { not significant digits }
\end{array} \begin{array}{l}
l_{1} c_{2} c_{3} c_{4} \times 2^{-3} \\
\text { significant digits } \\
\text { from } a_{5} a_{6} a_{7} a_{8} \\
\Theta b_{5} b_{6} b_{7} b_{8}
\end{array}
\end{aligned}
$$

## Loss of Significance

How can we avoid this loss of significance? For example, consider the function $f(x)=\sqrt{x^{2}+1}-1$

If we want to evaluate the function for values $x$ near zero, there is a potential loss of significance in the subtraction.

For example, if $x=10^{-3}$ and we use five-decimal-digit arithmetic $f\left(10^{-3}\right)=\sqrt{\left(10^{-3}\right)^{2}+1}-1=0$

How can we fix this issue?

## Loss of Significance

Re-write the function as $f(x)=\frac{\left(\sqrt{x^{2}+1}-1\right) \times\left(\sqrt{x^{2}+1}+1\right)}{\sqrt{x^{2}+1}+1}=\frac{x^{2}}{\sqrt{x^{2}+1}+1}$ (no subtraction!)

Evaluate now the function for $x=10^{-3}$ using five-decimal-digit arithmetic

$$
f\left(10^{-3}\right)=\frac{\left(10^{-3}\right)^{2}}{\sqrt{\left(10^{-3}\right)^{2}+1}+1}=\frac{10^{-6}}{2}
$$

## Example:

If $x=0.3721448693$ and $y=0.3720214371$ what is the relative error in the computation of $(x-y)$ in a computer with five decimal digits of accuracy?

Using five decimal digits of accuracy, the numbers are rounded as:

$$
\mathrm{fl}(\mathrm{x})=0.37214 \text { and } \mathrm{fl}(\mathrm{y})=0.37202
$$

Then the subtraction is computed:

$$
\mathrm{fl}(\mathrm{x})-\mathrm{fl}(\mathrm{y})=0.37214-0.37202=0.00012
$$

The result of the operation is: $\mathrm{fl}(\mathrm{x}-\mathrm{y})=1.20000 \times 10^{-2}$ (the last digits are filled with spurious zeros)
The relative error between the exact and computer solutions is given by

$$
\frac{|(x-y)-f l(x-y)|}{|(x-y)|}=\frac{0.0001234322-0.00012}{0.000123432}=\frac{0.0000034322}{0.000123432} \approx 3 \times 10^{-2}
$$

Note that the magnitude of the error due to the subtraction is large when compared with the relative error due to the rounding

$$
\frac{|\mathrm{x}-\mathrm{fl}(\mathrm{x})|}{|\mathrm{x}|} \approx 1.3 \times 10^{-5}
$$

