Floating point numbers, arithmetic realities, and random numbers

Coming events …

- **Project 3** has been posted
  - You don't have to be a physicist to complete the ideal gas simulation
  - Everything is handed in together – one deadline!
- Solution to Project 2 has been posted – understand it
- **Exam 2** on Thursday, April 2
  - Writing short programs in the lab
  - Choose 3 out of 5
- Bioinformatics lectures - week of April 6

Floating Point Numbers

- Similar to numbers in scientific notation, but in binary (base 2) instead base 10
- IEEE 754 binary floating point standard is the most widely used standard
  - since 1985, updated in 2008
- Not covered in Zelle; some mention in Learning Python, see http://docs.python.org/tutorial/floatingpoint.html for a Python related discussion

•3 •Monday, March 23,2009

Scientific Numbers

- 1.23664×10^4:
  - 10: base
  - 4: exponent
  - 1.23664: mantissa (example has 6 decimal precision)
  - sign
- $1.23664 \times 10^4 + 1.56333 \times 10^7$
  - shift-align the mantissa 1.23664×10^4
  - do the operation $\underline{+1563.33000 \times 10^4}$
  - round to 6-digit precision 1564.56664×10^4
  - result 1.56457×10^4

•4 •Monday, March 23,2009

Binary representation

Every decimal integer can be represented in binary.

181 = $\mathbf{1x2^7} + 0x2^6 + \mathbf{1x2^5} + \mathbf{1x2^4} + 0x2^3 + \mathbf{1x2^2} + 0x2^1 + \mathbf{1x2^0}$

= 128 + 32 + 16 + 4 + 1

This gives 10110101 in binary.

Remainders of a repeated division by 2 also give the binary number.

Conversion tool

http://acc6.its.brooklyn.cuny.edu/~gurwitz/core5/nav2tool.html

Binary representation

Not every decimal number can be represented in binary.

0.825 = ½ + ¼ + 1/16 = 13/16

This gives 0.1101 in binary.

1/10 = 1/16 + 1/32 + 1/256 + 1/512 + 1/4096 + …

1/10 cannot be represented as a binary fraction

Floating Point Numbers (IEEE 754)

- Float (32 bits)
  1bit sign **s** + 8 bit exponent **e** + 23 bit mantissa **m**

- Double (64 bits)
  1 bit sign + 11 bit exponent + 52 bit mantissa

```
>>> 1.0/10
0.10000000000000001
>>> x = 0.085
>>> x
0.085000000000000006
```


Arithmetic reality

- Not every number can be represented accurately as a floating point number
- 32-bit floating point number may not be enough to represent a number: round-off errors, overflows
- Arithmetic operations can magnify small errors: digit cancellation errors
- Other sources of errors: measurement errors, discretization errors, statistical errors

Results in potentially dangerous situations ….

Ariane 5 Rocket (June 4, 1996)

- 10 year, $7 billion ESA rocket self-destructed after launch
- 64-bit float converted to 16-bit signed integer resulted in unanticipated overflow
- Considered the most expensive computer bug in history

Patriot Missile accident (February 25, 1991)

- Failed to track Iraqi scud; scud hit Army barrack killed 28
- Inaccuracy in representing 1/10; after 100 hours the inaccuracy corresponded to 34 seconds

Vancouver stock exchange (November 1983)

- Index undervalued by 44%
- 22 months of accumulated truncation error (chopping versus rounding)

Clicker question

The Y2K problem was a computer bug resulting from the practice of representing the year with two digits.

Which statement is **false**?

A. The Y2K caused no significant computer problems.
B. Y2K helped launch the globalization of the software industry in India.
C. Microsoft Excel had no Y2K problem
D. Countries that spent little on preparing for the Y2K problem performed as well as others.

Precision = number of digits

Accuracy = number of correct digits

p = 3.133333 is precise to 7 decimals, but has only two digits of accuracy

Roundoff errors

- Arithmetic with integers is exact, except when the number is outside the range that can be represented (overflow)
- Floating point arithmetic is not exact
  - roundoff errors can propagate through the calculation

Arithmetic Issues

- Decimal with 3-digit mantissa:
  - $351 = 0.351 \times 10^3$
  - $12.43 = 0\ .124 \times 10^2$
- Round-off error:

  $0.124 \times 0.351 = 0.043524$

  using 3-digits, we get 0.435×10^{-1}

  difference $= 0.24 \times 10^{-4}$
- Digit cancellation error:

  $0.127 - 0.124 = 0.003$

  0.300×10^{-2}
- Large summations can be problematic, as are iterated computations in general…

Plotting f(x) = (1–cos(x))/(x^2)

- Plot f(x) from x = $-4*10^{-8}$ to $4*10^{-8}$
- The value of f(x) in this region is about 0.5
- Alternative computation

  $\cos(x) \approx 1 - x^2/2! + x^4/4! - x^6/6! + \ldots$

  $1 - \cos(x) \approx x^2/2 - x^4/24 +$ lower order terms

  $(1 - \cos(x))/x^2 \approx 1/2 - x^2/24$
- See program CancelGraph.py
- Reason for f(x)'s wrong results: digit cancellation

Correct graph for f(x)?

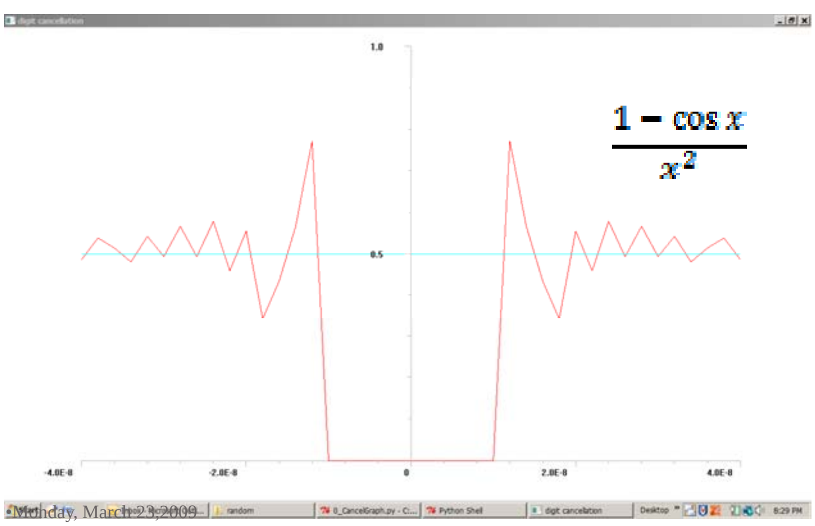

Another example with a cancellation problem

$e^x = 1 + x + x^2/2! + x^3/3! + x^4/4! + \ldots$

- Series converges, but for many negative values a program will have few correct digits
- Reason: consider x=-25
  terms $25^{24}/24!$ and $-25^{25}/25!$ should cancel out, but don't

Clicker Question

Which is not a Random Event?

A. Color of first car crossing Stadium Ave after 12 noon
B. 101st digit in expansion of Pi is even
C. Fair coin toss
D. June 1, 2009, is a cloudy day

Random numbers

Who needs them

- Gambling, statistical analysis, computer simulations, cryptographic systems

Random number generator

- A method or device generating a sequence of numbers lack any pattern

Pseudo random number generator

- An algorithm that produces long sequences of numbers with good random properties; eventually the sequence will repeat

Pseudo-Random Numbers

- Look randomly distributed
- Pass various tests of acceptability
  - e.g. uniform distribution and bins

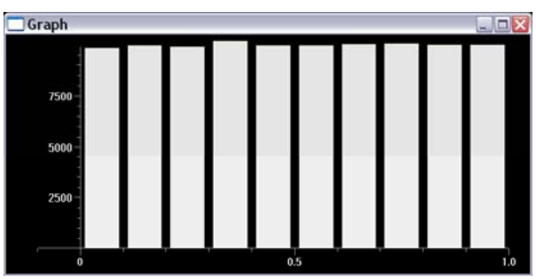

100,000 trials
First bin 9,964 vs. 10,000, less than 0.4%

Beyond Uniform Distributions

- Random number in a fixed range, each value equally likely
  - Example: toss of a coin
  - Probability of having a run of 10 consecutive tosses being *heads* is 1:1024
- Binomial distribution
  - Probability of getting k successes out of n trials
  - Drop balls into Francis Galton's board (1889)
    - a device for statistical experiments
    - the path of each ball is a Bernoulli trial

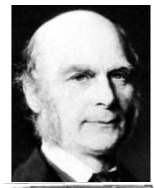

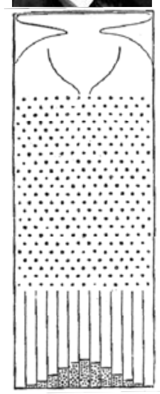

•Monday, March 23,2009

Normal (Gaussian) Distribution

- Bell curve:
  - standard deviation σ, mean μ
  - area 1 under curve

$$\frac{1}{\sigma\sqrt{2\pi}}\exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$$

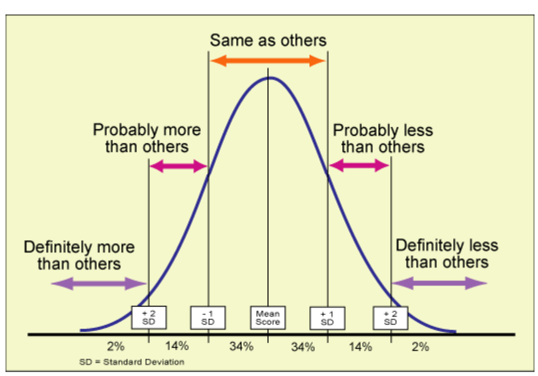

•Monday, March 23,2009

Visual Randomness Test

- Histogram is the basic tool

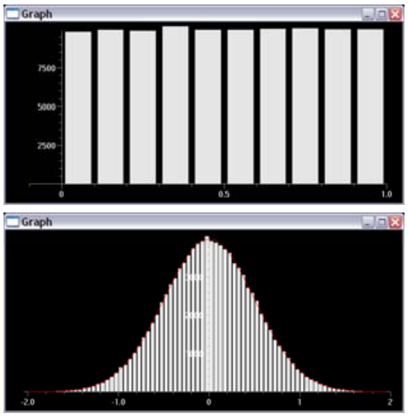

100,000 trials each

•Monday, March 23,2009