CS 190C
Science Education in Computational Thinking

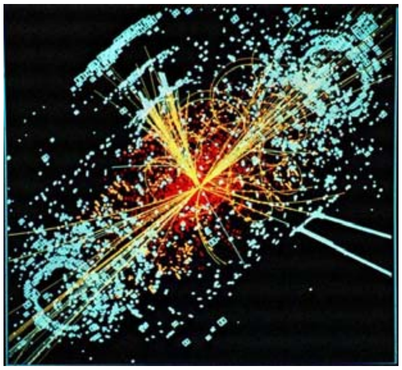

SCIENTIFIC COMPUTING:

- Applications
- High performance computing
- Integration of functions
- Sampling and statistics
- Problem & Algorithm Stability
- Geometric Algebra

Hoffmann, 2008

Applications of Scientific Computing

Playing Cosmic Billiards (or, How to Change a Spacecraft's Speed and Direction)

One of the first things taught in geometry is that "the shortest distance between two points is a straight line." But in the case of Project Galileo, an interplanetary spacecraft bound for Jupiter, the most efficient path is a six-year journey that initially heads towards Venus instead of Jupiter.

Galileo is currently on the last leg of a flight path that has included one close flyby of Venus and two close flybys of Earth. Each of these flybys allowed the spacecraft to use the planet's gravity to accelerate Galileo to greater speeds, rather like a slingshot. The exact flight path depends on how close Galileo flies to the planet, so navigation engineers work like expert billiard players, carefully lining up each maneuver and encounter so that Galileo arrives right on target. Without the boost provided by these flybys, Galileo would need an extra 10,900 kilograms of propellant—about twelve times more than was on board at launch.

Having accomplished the three planetary flybys, Galileo is right on schedule to reach Jupiter on December 7, 1995. Along the way, there have been opportunities to observe two asteroids; these pictures and other scientific information are the first close-up views we have had of these ancient remnants of the primordial solar system.

Close flybys and gravity assists will also be used to enable Galileo to make a complex tour of Jupiter's system. An additional 3,600 kilograms of propellant (about four times the total amount of propellant on the spacecraft at launch) would be needed to fly the tour without the billiards-like gravity assist technique.

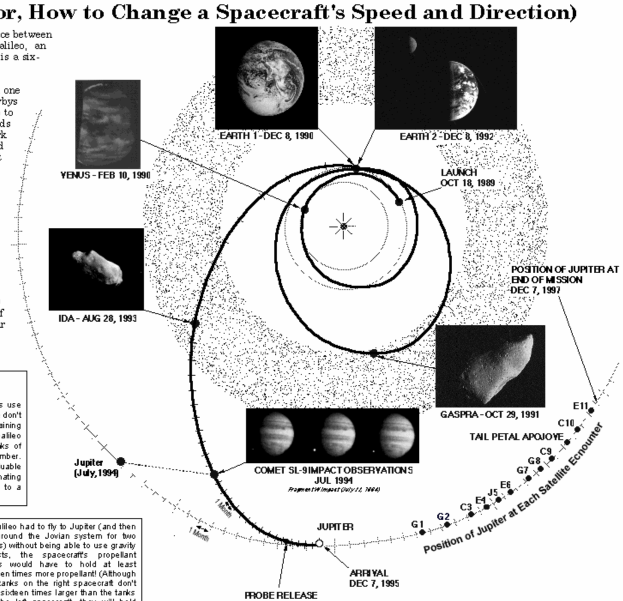

What Type of Rocket Propellant Does Galileo Use?

Galileo's fuel tanks can't just be filled up at the local gas station, as the engines use monomethyl hydrazine for fuel. Fuel has to be oxidized (burned) in order to ignite; we don't have to worry about this on Earth, where we are surrounded by an atmosphere containing oxygen, but an interplanetary spacecraft like Galileo has to carry its own oxidizer (Galileo uses nitrogen tetroxide). Instead of pumps, the spacecraft uses two separate tanks of helium pressurant that force the fuel and oxidizer together into the combustion chamber. These two liquids are stored at about room temperature on the spacecraft. One valuable property of these liquids is that they burn on contact with each other, eliminating complicated ignition systems (such as spark plugs; Galileo doesn't have access to a mechanic for regularly scheduled tune-ups!).

How "Fuel-Efficient" are Gravity Assists?

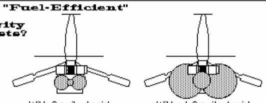

If Galileo had to fly to Jupiter (and then fly around the Jovian system for two years) without being able to use gravity assists, the spacecraft's propellant tanks would have to hold at least sixteen times more propellant! (Although the tanks on the right spacecraft don't look sixteen times larger than the tanks on the left spacecraft, they will hold sixteen times more propellant, since volume is proportional to radius cubed)

Structure of Protein

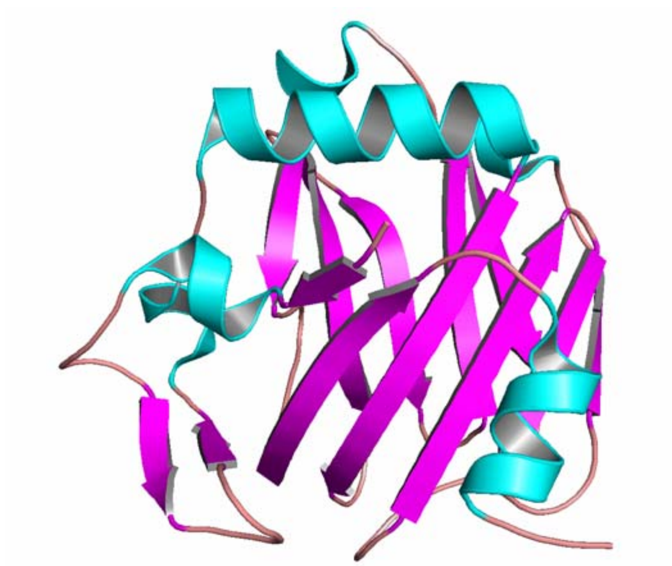

Stellar Evolution

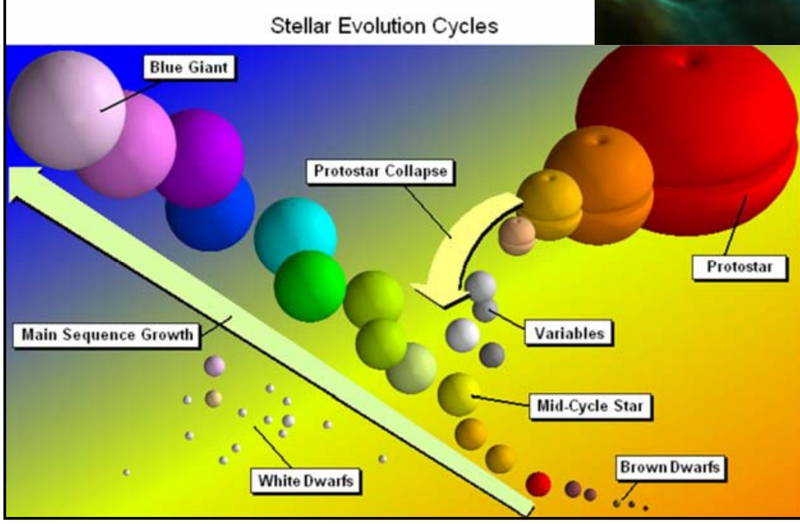

5

- Will it pop?
- Will it stand?
- Will it run?
- Will it fly?

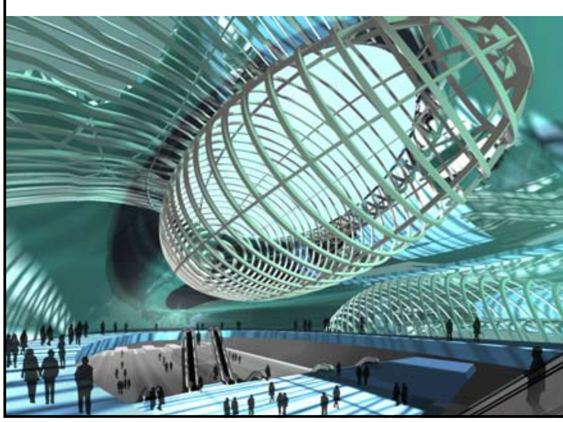

6

Scientific Computing

- Grand-challenge problems in physics, biology, chemistry:
  - CMS and the Higgs boson
  - Mechanism for dialysis in kidneys
  - Molecular interaction of enzymes
- Huge computational challenges
  - Blue waters peta-scale machine

High Performance Computing
High Throughput Computing

Moore’s Law: CPU doubles every 2 years

- Transistor count a proxy for uniprocessor speed
- Memory barrier
- Heat barrier

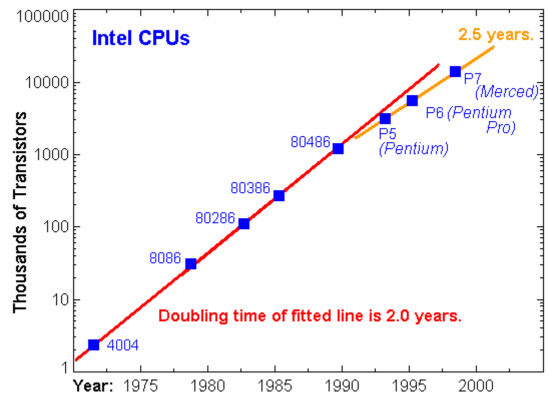

Typical Cluster

- Node has 2 processors plus on-board memory
- Processor 4 cores, each with a cache
- Node-to-node interconnect fabric
- Racks have disk drives

- Total cost of ownership:
  - staff to run the machine
  - electricity to power the machine
  - A/C to dissipate the heat
- Lear has
  - 512 nodes, 1024 cores.
  - Uses 180kW; at $0.042 electricity costs $5,400 a month
  - A/C cost is extra

Top 500 as of Nov. 2007

Rank	Site	System	Cores	R_{max} (GFlop/s)	R_{peak} (GFlop/s
1	DOE/NNSA/LLNL United States	eServer Blue Gene Solution IBM	212992	478200	596378
2	Forschungszentrum Juelich (FZJ) Germany	Blue Gene/P Solution IBM	65536	167300	222822
3	SGI/New Mexico Computing Applications Center (NMCAC) United States	SGI Altix ICE 8200, Xeon quad core 3.0 GHz SGI	14336	126900	172032
4	Computational Research Laboratories, TATA SONS India	Cluster Platform 3000 BL460c, Xeon 53xx 3GHz, Infiniband Hewlett-Packard	14240	117900	170880
5	Government Agency Sweden	Cluster Platform 3000 BL460c, Xeon 53xx 2.66GHz, Infiniband Hewlett-Packard	13728	102800	146430
6	NNSA/Sandia National Laboratories United States	Sandia/ Cray Red Storm, Opteron 2.4 GHz dual core Cray Inc.	26569	102200	127531
7	Oak Ridge National Laboratory United States	Cray XT4/XT3 Cray Inc.	23016	101700	119350
8	IBM Thomas J. Watson Research Center United States	eServer Blue Gene Solution IBM	40960	91290	114688
9	NERSC/LBNL United States	Cray XT4, 2.6 GHz Cray Inc.	19320	85368	100464
10	Stony Brook/BNL, New York Center for Computational Sciences United States	eServer Blue Gene Solution IBM	36864	82161	103219

Clicker Question

- How soon does the fastest supercomputer drop off the top 500 list?

A. Two years?

B. Four years?

C. Six years?

D. Eight years?

Answer:

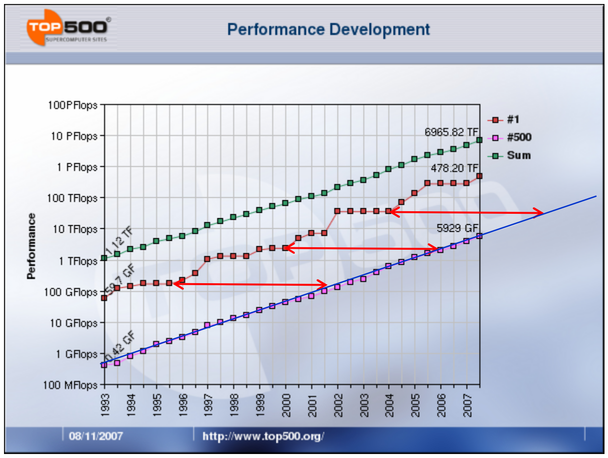

Some Purdue HPC/HTC Projects

- CMS (particle physics)
- STEREO and other NASA projects (AAE)
- PRISM (ME and others)
- nCn (ECE)
- Virus structure (Bio)
- Climate modeling (EAS)

…

CMS Data Challenge

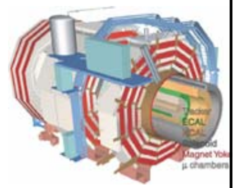

- Large Hadron collider (CERN)
- CMS generating events at the rate of 500 GB/sec (http://cms.cern.ch/)
- Looking for interesting events ($1:10^{16}$)

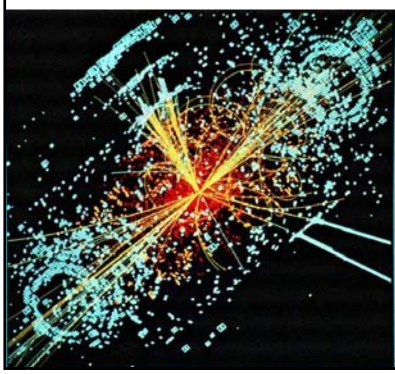

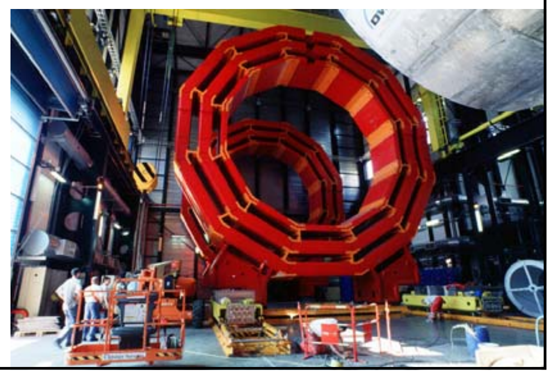

STEREO:

Solar TErrestrial RElations Observatory

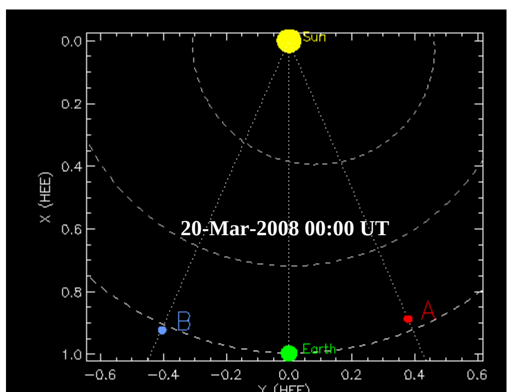

stereo.gsfc.nasa.gov/where.shtml

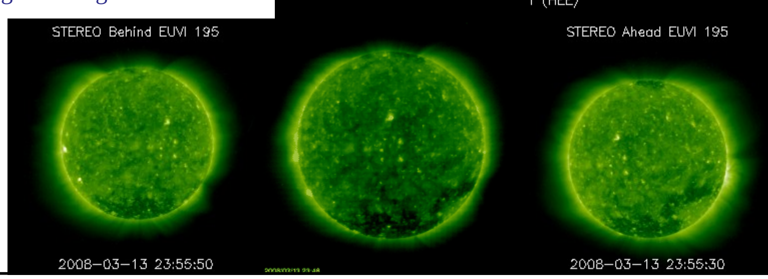

Function Integration

$$\int_a^b f(x)dx$$

$$f(x) = e^{-x}\sin(8\sqrt[3]{x}) + 1$$

$$a = 0, \quad b = 2$$

Basic algorithm in scientific computing

Midpoint

$$\int_a^b f(x)dx \approx (b-a)f(c), \qquad c = \frac{a+b}{2}$$

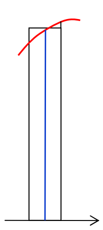

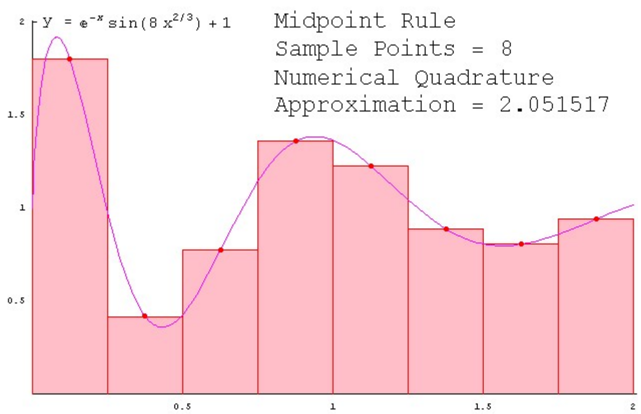

Trapezoidal

$$\int_a^b f(x)dx \approx h(\frac{1}{2}f(x_0) + f(x_1) + \ldots + f(x_{n-1}) + \frac{1}{2}f(x_n))$$

- Piecewise linear approximation
- Subdivision into *n* intervals

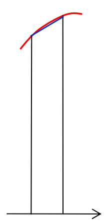

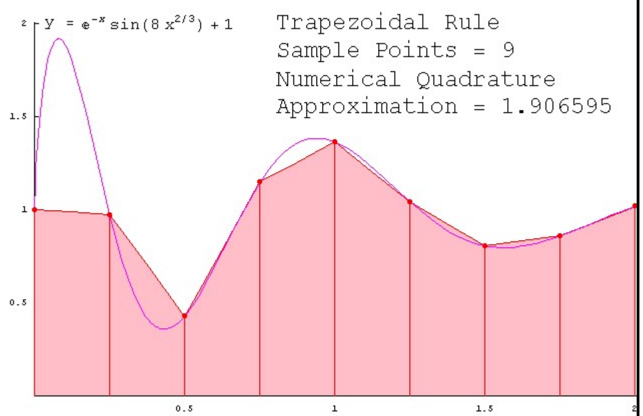

Simpson

$$\int_a^b f(x)dx \approx h(\frac{1}{3}f(x_0) + \frac{4}{3}f(x_1) + \frac{2}{3}f(x_2) + \frac{4}{3}f(x_3) + \ldots + \frac{2}{3}f(x_{n-2}) + \frac{4}{3}f(x_{n-1}) + \frac{1}{3}f(x_n))$$

- Approximation by parabolic arcs
- Subdivision into *n* intervals, *n* even