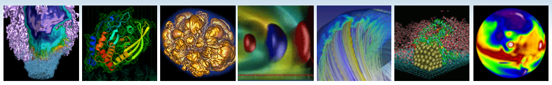

Saving the World with Computing

Kathy Yelick
EECS Professor, U.C. Berkeley

Associate Laboratory Director for Computing Sciences
Lawrence Berkeley National Laboratory

Using Computers for Science and Engineering

Computers are used to understand things that are:

- too big
- too small
- too fast
- too slow
- too expensive or
- too dangerous

for experiments

Understanding the universe

Proteins and diseases like Alzheimer's

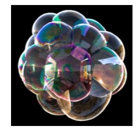

Industrial products and processes

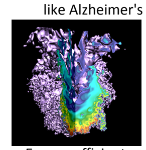

Energy-efficient combustion engines

Simulation: The "Third Pillar" of Science

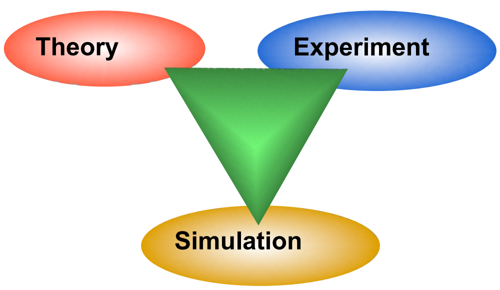

Addressing Challenges using Computing

- **Two of the most significant challenges**
  - **Our changing world:** understanding climate change, alternative energy sources, mitigation techniques, etc.
  - **Health and medicine:** understanding the human body, development of treatments, and disease prevention

Carbon Cycle 2.0 Initiative at Berkeley Lab

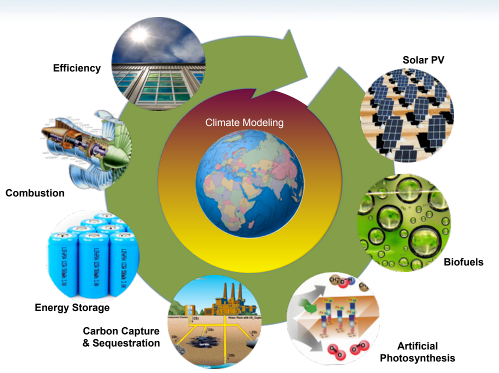

Carbon Cycle 2.0 Initiative at Berkeley Lab

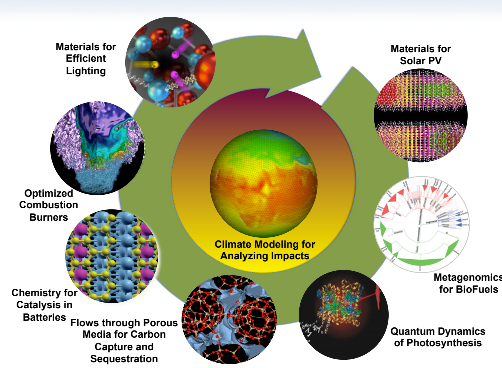

Simulations Show the Effects of Climate on Hurricanes

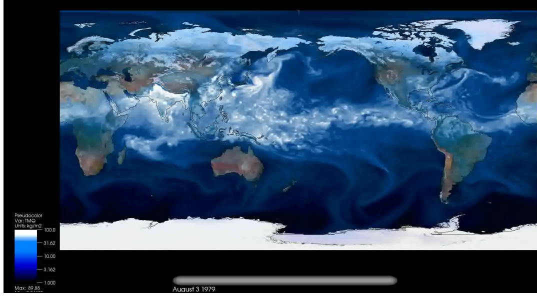

Mitigating Global Climate Change

Can global warming impacts be diminished if greenhouse gases are cut?

Average surface air temperatures rise by >3°C if emissions increase at current rate

Temperatures rise by <2°C across nearly all populated areas if emissions are cut by 70%

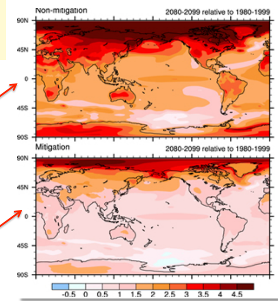

20th Century Climate Data Reconstructed

NERSC

Reconstructed weather in 6-hour intervals from 1871-2010

- **Data from military, meteorologists, ships' crews, military, and volunteers**
- **Over 10M hours of computing at NERSC using reverse Kalman filter algorithms**
- **Total science data storage at NERSC: 44 Petabytes (10^{15})**

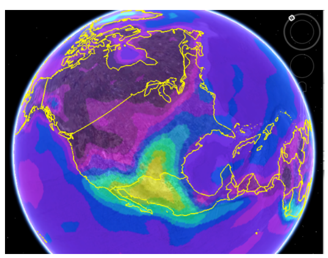

Relative Humidity for 1920-1929
Gil Compo, PI (U. Colorado)

Climate Change Requires Lots of Data

"validate" that the computer models are working as expected

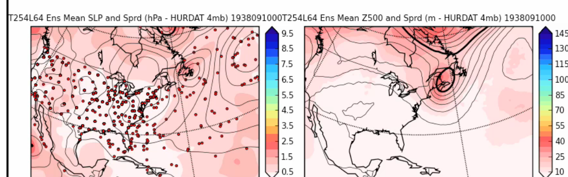

Data Structures for Simulations

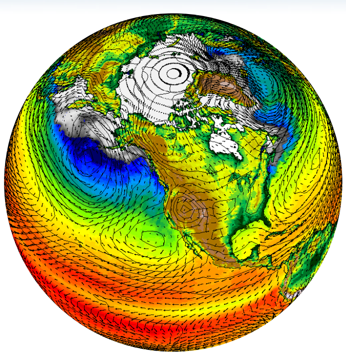

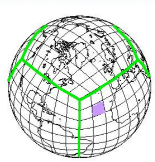

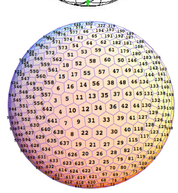

Faster Computers: More Detail

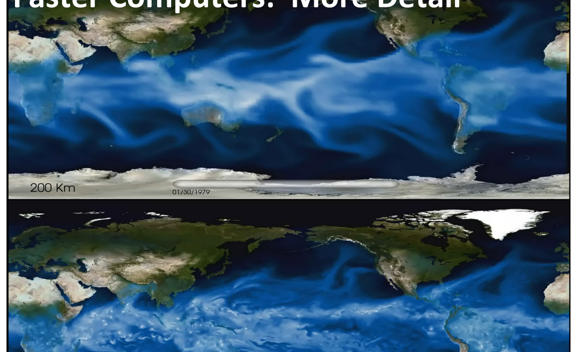

Michael Wehner, Prabhat, Chris Algieri, Fuyu Li, Bill Collins, Lawrence Berkeley National Laboratory; Kevin Reed, University of Michigan; Andrew Gettelman, Julio Bacmeister, Richard Neale, National Center for Atmospheric Research

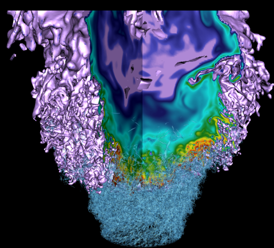

Simulations reveal features not visible in lab

Experiments demonstrate feasibility

Result: Low NOx burner technology licensed by industry

Basic Material Science Result Explains Practical Limits of LEDs

NERSC

- *LEDs are up to 3x more energy efficient than fluorescent lights and last 10x longer for low level light.*
- *At higher current levels LEDs lose of efficiency (called "LED droop") making them worse than fluorescents.*
- *Fundamental material science question: why does LED efficiency "droop"?*
  - *Answer: Auger recombination combined with carrier scattering.*
  - Researchers in universities and industry are now working on solutions.

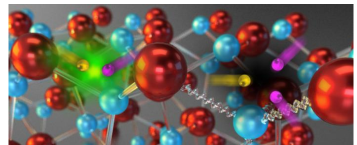

The illustration shows nitride-based LEDs. At left, an electron and electron hole recombine and release light. In Auger recombination (right) the electron and hole combine with a third carrier, releasing no photon. The energy loss is also assisted by indirect processes, vibrations in the crystal lattice shown as squiggles.

Simulations to Get Rid of CO_2

- Carbon sequestration: "The process of removing carbon from the atmosphere or from flue gasses and depositing it in a reservoir."
- $CaO + CO_2 \rightarrow CaCO_3$

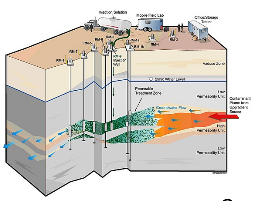

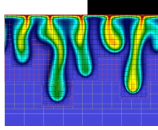

George Pau, LBNL

Towards a Digital Human: The 20+ Year Vision

- Imagine a "digital body double"
  - 3D image-based medical record
  - Includes diagnostic, pathologic, and other information
- Used for:
  - Diagnosis
  - Less invasive surgery-by-robot
  - Experimental treatments

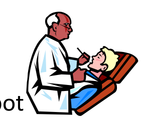

Digital Human Today: Imaging

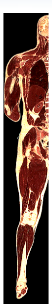

- The Visible Human Project
  - 18,000 digitized sections of the body
    - Male: 1mm sections, released in 1994
    - Female: .33mm sections, released in 1995
  - Goals
    - study of human anatomy
    - testing medical imaging algorithms
  - Current applications:
    - educational, diagnostic, treatment planning, virtual reality, artistic, mathematical and industrial
    - Used by > 1,400 licensees in 42 countries

Image Source: www.madsci.org

Experimental Data: Visible Human

The National Library of Medicine's

Visible Human Project (TM)

Human-Computer Interaction Lab
Univ. of Maryland at College Park

Heart Simulation

Movie from Charles Peskin and Dave McQueen at NYU

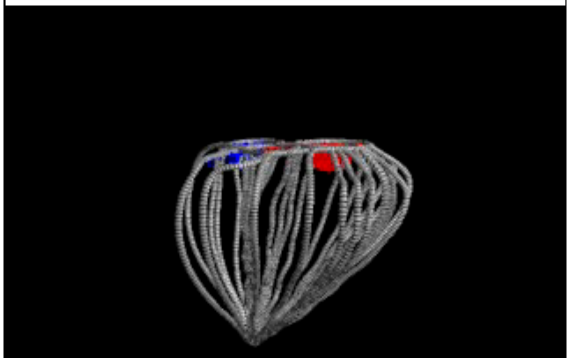

Heart Simulation

Movie from Boyce Griffith's PhD thesis, NYU

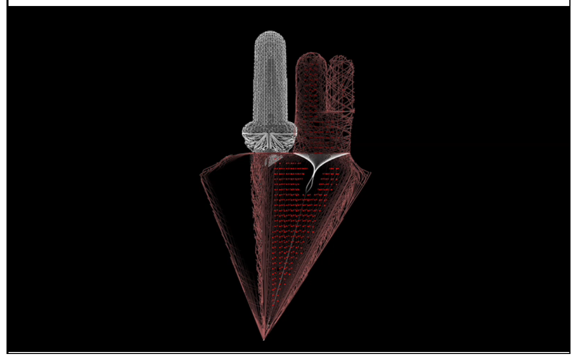

Organ Simulation

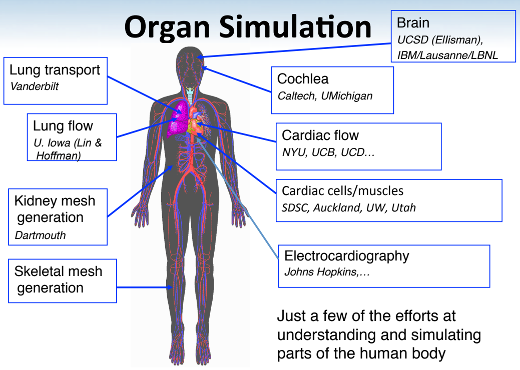

Just a few of the efforts at understanding and simulating parts of the human body

Screening Proteins

- Large number of simulations covering a variety of related proteins,…

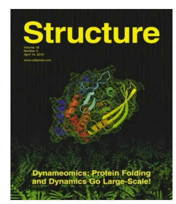

Dynameomics Database
Improve understanding of disease and drug design, e.g., 11,000 protein unfolding simulations stored in a public database. [V. Daggett, UW]

22

Real-Time Deformation and Fracture in a Game Environment

Eric Parker
Pixelux Entertainment

James O'Brien
U.C. Berkeley

Video Edited by Sebastian Burke

From the proceedings of SCA 2009, New Orleans

LBNL uses Big Data and Big Computing to Study Biology

IMG has over 16B Genes for metagenomics; Phylogeny Map shown

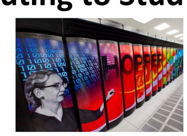

NERSC provides 1.3B computing hours to DOE researchers, connected by ESNet network

Advanced computing helps scientists define a healthy human microbiome

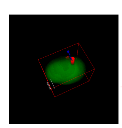

Algorithms distinguish normal vs. malignant cells

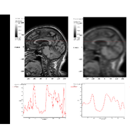

Imaging denoising uses of advanced architectures (GPUs)

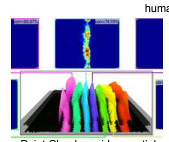

Point Cloud provides spatial & temporal gene expression in early Drosophila embryos

Advanced math provides foundation for virtual colonoscopy

Big D and Big C: Computing on Big Data to help Cure Cancer

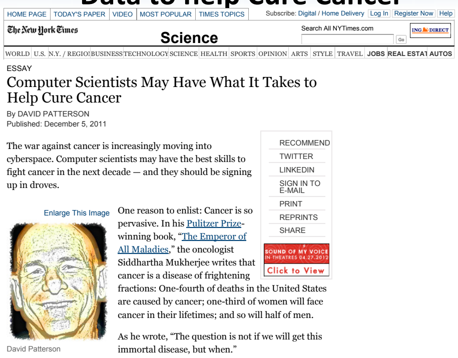
HOME PAGE | TODAY'S PAPER | VIDEO | MOST POPULAR | TIMES TOPICS
Subscribe: Digital / Home Delivery | Log In | Register Now | Help
The New York Times
Science
Search All NYTimes.com
WORLD | U.S. | N.Y. / REGION | BUSINESS | TECHNOLOGY | SCIENCE | HEALTH | SPORTS | OPINION | ARTS | STYLE | TRAVEL | JOBS | REAL ESTATE | AUTOS

ESSAY

Computer Scientists May Have What It Takes to Help Cure Cancer

By DAVID PATTERSON
Published: December 5, 2011

The war against cancer is increasingly moving into cyberspace. Computer scientists may have the best skills to fight cancer in the next decade — and they should be signing up in droves.

RECOMMEND
TWITTER
LINKEDIN
SIGN IN TO E-MAIL
PRINT
REPRINTS
SHARE

Enlarge This Image

One reason to enlist: Cancer is so pervasive. In his Pulitzer Prize-winning book, "The Emperor of All Maladies," the oncologist Siddhartha Mukherjee writes that cancer is a disease of frightening fractions: One-fourth of deaths in the United States are caused by cancer; one-third of women will face cancer in their lifetimes; and so will half of men.

As he wrote, "The question is not if we will get this immortal disease, but when."

David Patterson

Why Study Computer Science?

1) Because computers can help solve important problems
2) **Because programming is fun and there are plenty of new problems to solve**

Black Swans of Computing

with 1992 Technology

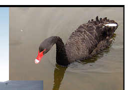

Units of Measure in High Performance Computing (HPC)

- **High Performance Computing (HPC) units are:**
  - **Flops: floating point operations**
  - **Flops/s: floating point operations per second**
  - **Bytes: size of data (a double precision floating point number is 8)**
- **Typical sizes are millions, billions, trillions…**

	Kilo	Kflop/s = 10^3 flop/sec	Kbyte = 2^{10} = 1024 ~ 1,000 bytes
	Mega	Mflop/s = 10^6 flop/sec	Mbyte = 2^{20} = 1048576 ~ 10^6 bytes
1988	Giga	Gflop/s = 10^9 flop/sec	Gbyte = 2^{30} ~ 10^9 bytes
1998	Tera	Tflop/s = 10^{12} flop/sec	Tbyte = 2^{40} ~ 10^{12} bytes
2008	Peta	Pflop/s = 10^{15} flop/sec	Pbyte = 2^{50} ~ 10^{15} bytes
	Exa	Eflop/s = 10^{18} flop/sec	Ebyte = 2^{60} ~ 10^{18} bytes
	Zetta	Zflop/s = 10^{21} flop/sec	Zbyte = 2^{70} ~ 10^{21} bytes
	Yotta	Yflop/s = 10^{24} flop/sec	Ybyte = 2^{80} ~ 10^{24} bytes

High End Computing Revolutions

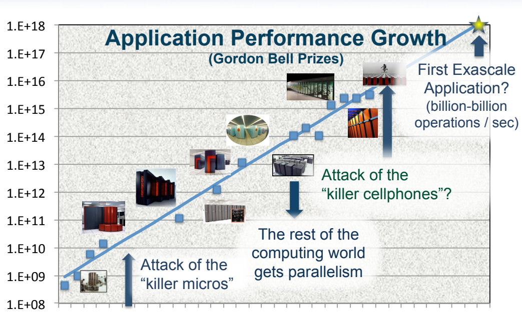

29

The Fastest Computers (for Science) Have Been Parallel for a Long Time

- Fastest Computers in the world: top500.org
- Our Hopper Computer has 150,000 cores and > 1 Petaflop (10^{15} math operations / second)
- Programming and "debugging" are challenging

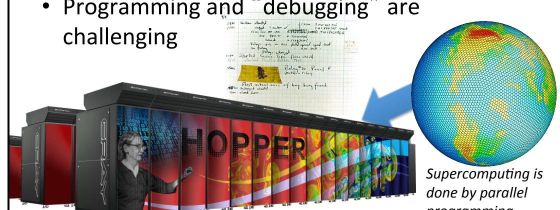

Supercomputing is done by parallel programming

30

The Programming Answer is Obvious…

More Regular

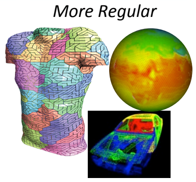

Message Passing Programming
Divide up domain in pieces
Compute one piece
Send/Receive data from others

MPI, and many libraries

More Irregular

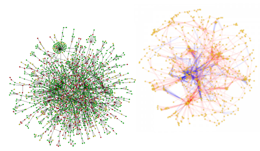

Global Address Space Programming
Each start computing
Grab whatever / whenever

UPC, CAF, X10, Chapel, GlobalArrays

31

Energy Challenge for Computing

At ~$1M per MW, energy costs are substantial

An exaflop in 2020 would use ~200 MW with "usual" scaling

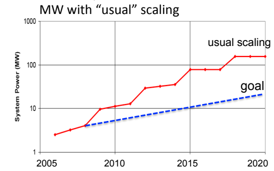

Google Details, and Defends, Its Use of Electricity

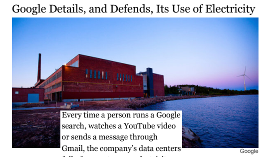

Every time a person runs a Google search, watches a YouTube video or sends a message through Gmail, the company's data centers full of computers use electricity. Those data centers around the world continuously draw almost 260 million watts — about a quarter of the output of a nuclear power plant.

NSA Maxes Out Baltimore Power Grid

August 6th, 2006 : Rich Miller

The National Security Agency's technology infrastructure at Fort Meade, Md. has maxed out the electric capacity of the Baltimore area power grid, creating a major challenge for the agency, sources told the Baltimore Sun. An excerpt:

The worldwide data center power in was about 26 gigawatts in 2010 (up from 17 in 2005)

32

New Processor Designs are Needed to Save Energy

Cell phone processor
(0.1 Watt, 4 Gflop/s)

Server processor
(100 Watts, 50 Gflop/s)

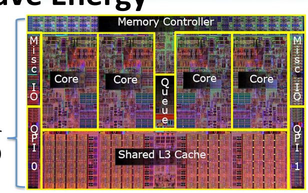

- The server is about 10x faster than the cell phone processor
- But uses 1000x more power → cell phone is 100x more efficient
- Why: Power is proportional to V^2f, and increasing frequency (f) also requires increase voltage V → cube
- Next computers built from graphics, games, cell phones,…

33

Power Limits Computing Performance Growth

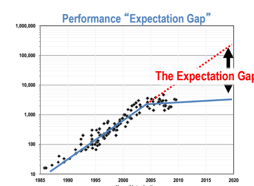

Processor industry was running at "maneuvering speed"
- David Liddle

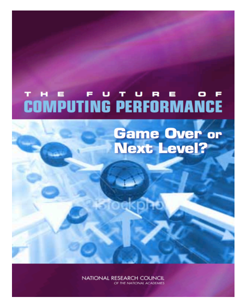

34

Computational Science is Necessarily Collaborative

*… as from the beginning the work has been a team effort involving many able and devoted co-workers in many laboratories. As I am sure you will appreciate, a **great many diverse talents** are involved in such developments and whatever measure of success is achieved is dependent on **close and effective collaboration**.*

Ernest O. Lawrence
UC Berkeley Professor of Physics
Founder of Lawrence Berkeley National Laboratory
In his Nobel Lecture, December 11, 1951

Internships Available: http://csee.lbl.gov/

Why Study Computer Science?

1) Because computers can help solve important problems
2) Because computers are fun to program
3) **Because computers make a good career**
4) **Because you will get to work with lots of great people**