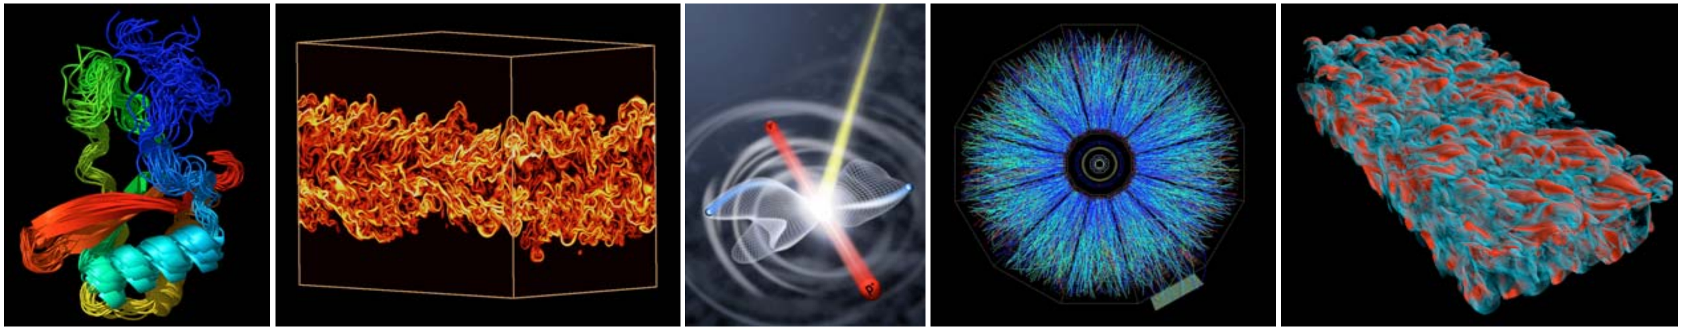

Saving the World with Computing

and other reasons to study Computer Science

Kathy Yelick

EECS Professor, U.C. Berkeley

Associate Laboratory Director for

Computing Sciences and NERSC Director,

Lawrence Berkeley National Laboratory

NERSC Overview

NERSC represents science needs

- Over 3000 users, 400 projects, 500 code instances
- Over 1,600 publications in 2009
- Time is used by university researchers (65%), DOE Labs (25%) and others

1 Petaflop Hopper system, late 2010

- High application performance
- Nodes: 2 12-core AMD processors
- Low latency Gemini interconnect

NERSC Systems

Large-Scale Computing Systems

Franklin (NERSC-5): Cray XT4
- 9,532 compute nodes; 38,128 cores
- ~25 Tflop/s on applications; 356 Tflop/s peak

Hopper (NERSC-6): Cray XE6
- Phase 1: Cray XT5, 668 nodes, 5344 cores
- Phase 2: > 1 Pflop/s peak (late 2010 delivery)

Clusters

140 Tflops total

Carver
- IBM iDataplex cluster

PDSF (HEP/NP)
- ~1K core throughput cluster

Magellan Cloud testbed
- IBM iDataplex cluster

GenePool (JGI)
- ~5K core throughput cluster

NERSC Global Filesystem (NGF)

Uses IBM's GPFS
- 1.5 PB capacity
- 5.5 GB/s of bandwidth

HPSS Archival Storage

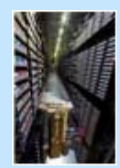

- 40 PB capacity
- 4 Tape libraries
- 150 TB disk cache

Analytics

Euclid

(512 GB shared memory)

Dirac GPU testbed (48 nodes)

Why Study Computer Science?

- **Because computers can help solve important problems**

Using Computers for Science

Computers are used to understand things that are:

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

for experiments

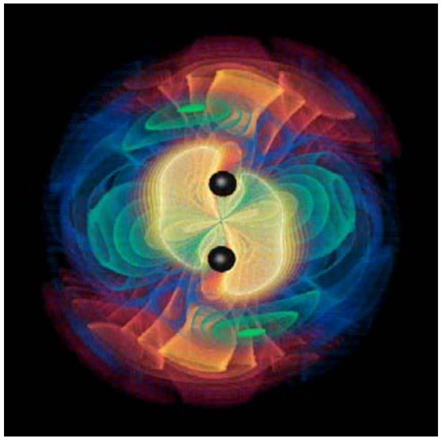

Black holes

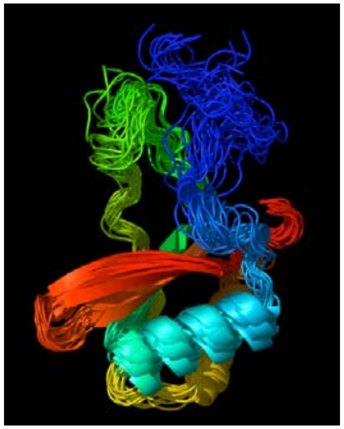

Proteins and diseases like Alzheimer's

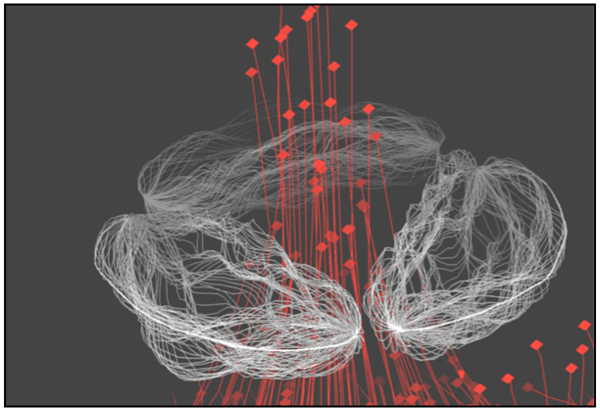

Blood flow in the heart

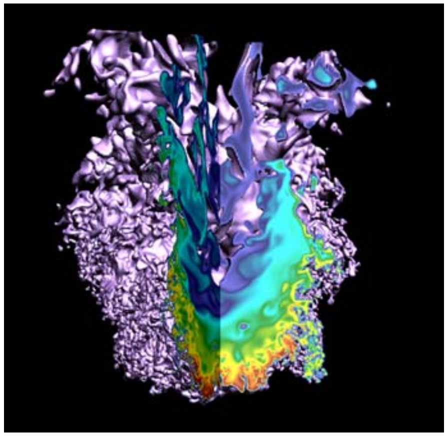

Energy-efficient combustion engines

Simulation:
The “Third Pillar” of Science

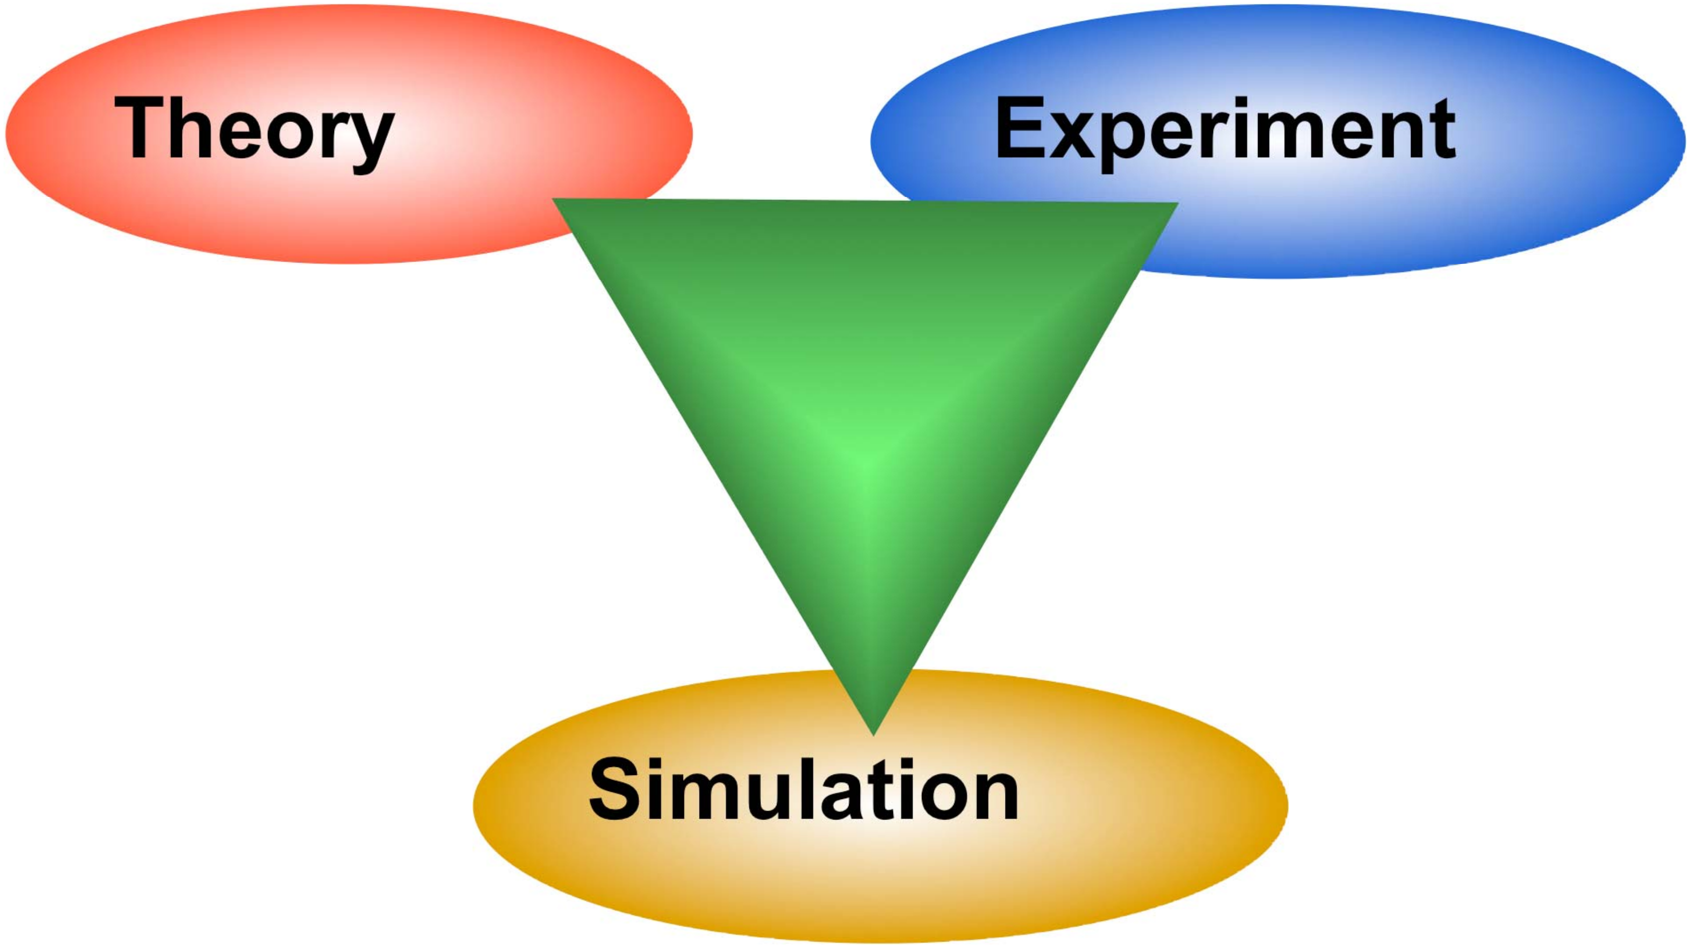

Global Challenges using Computing

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

Energy Science at NERSC

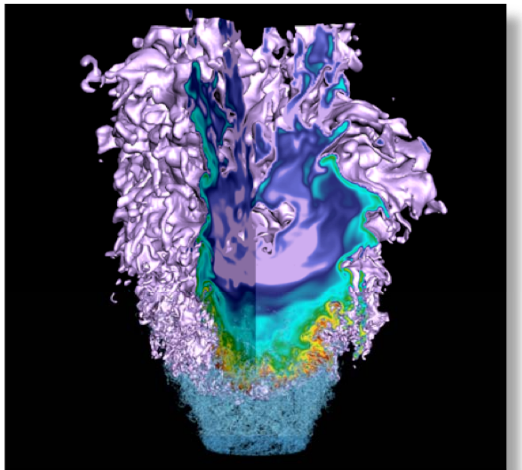

Combustion: New algorithms (AMR) coupled to experiments

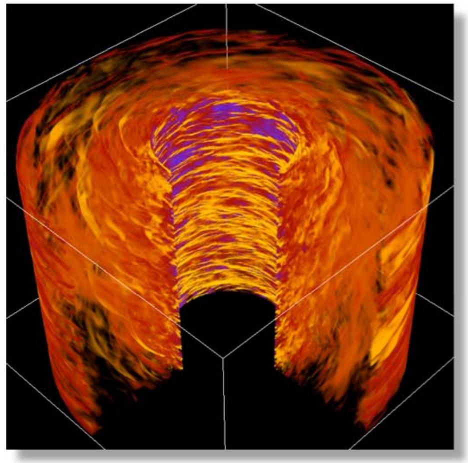

Fusion: Simulations of Fusion devices at ITER scale

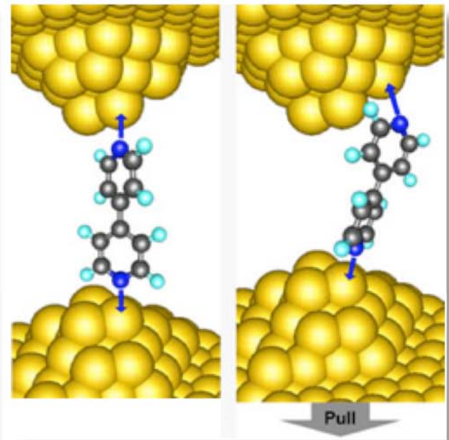

Nano devices: New single molecule switching element

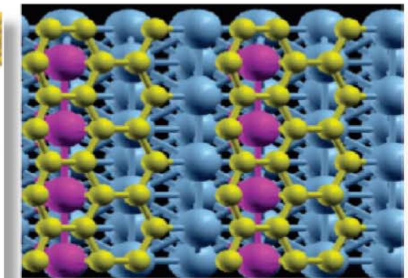

Energy storage: Catalysis for improved batteries and fuel cells

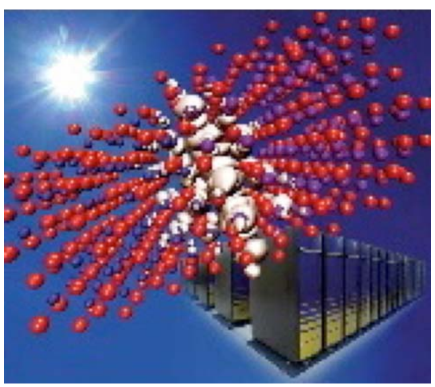

Materials: For solar panels and other applications.

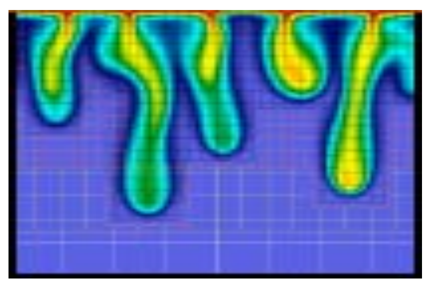

Capture & Sequestration: EFRCs

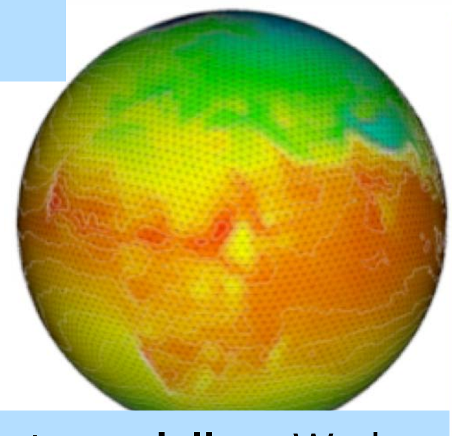

Climate modeling: Work with users on scalability of cloud-resolving models

Supercomputers "Discovered" Climate Change

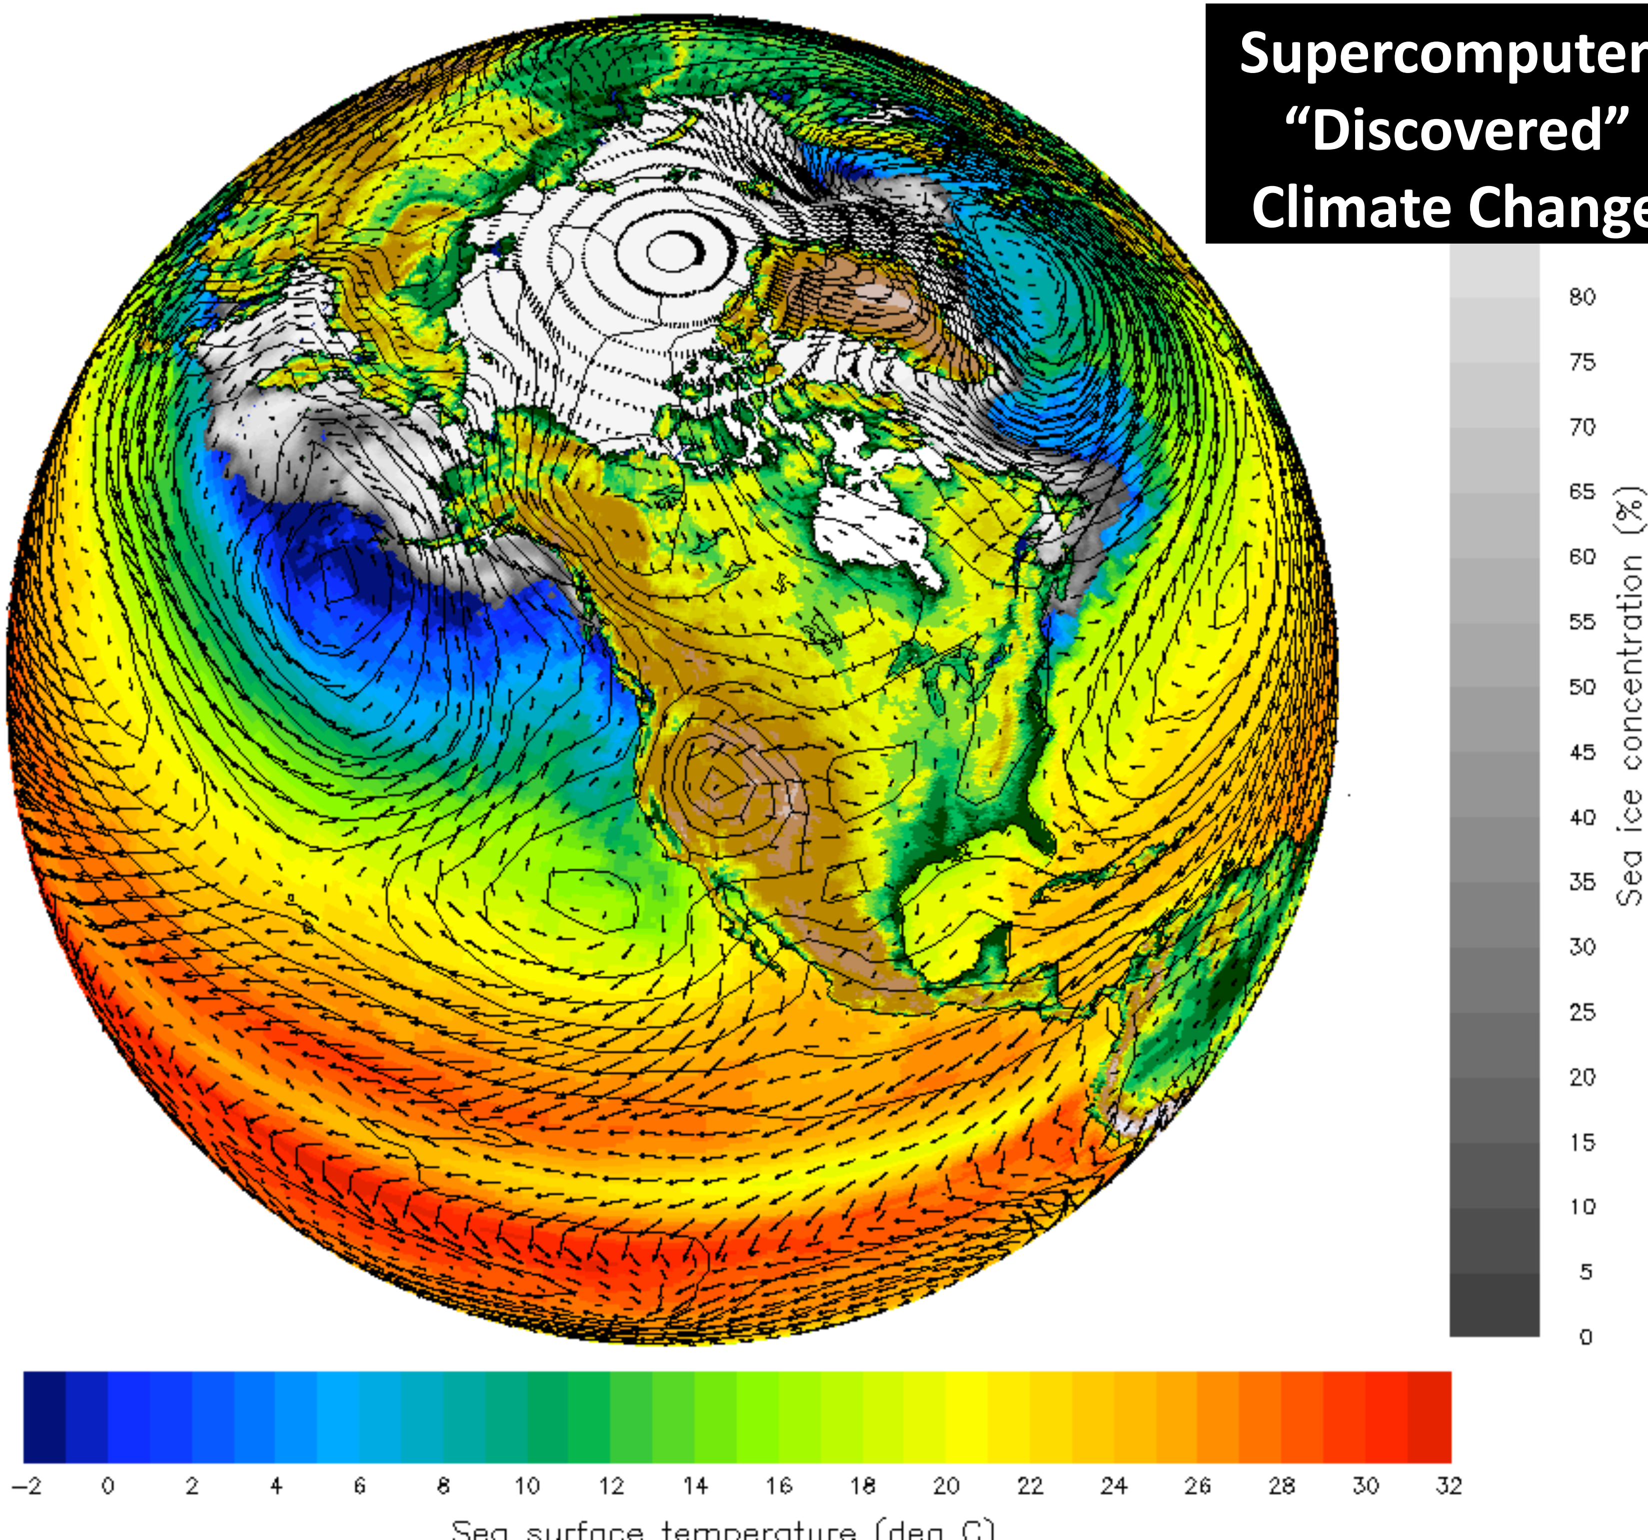

To study future climate change, we start by simulating the past

Simulation of 1938 hurricane hitting New York

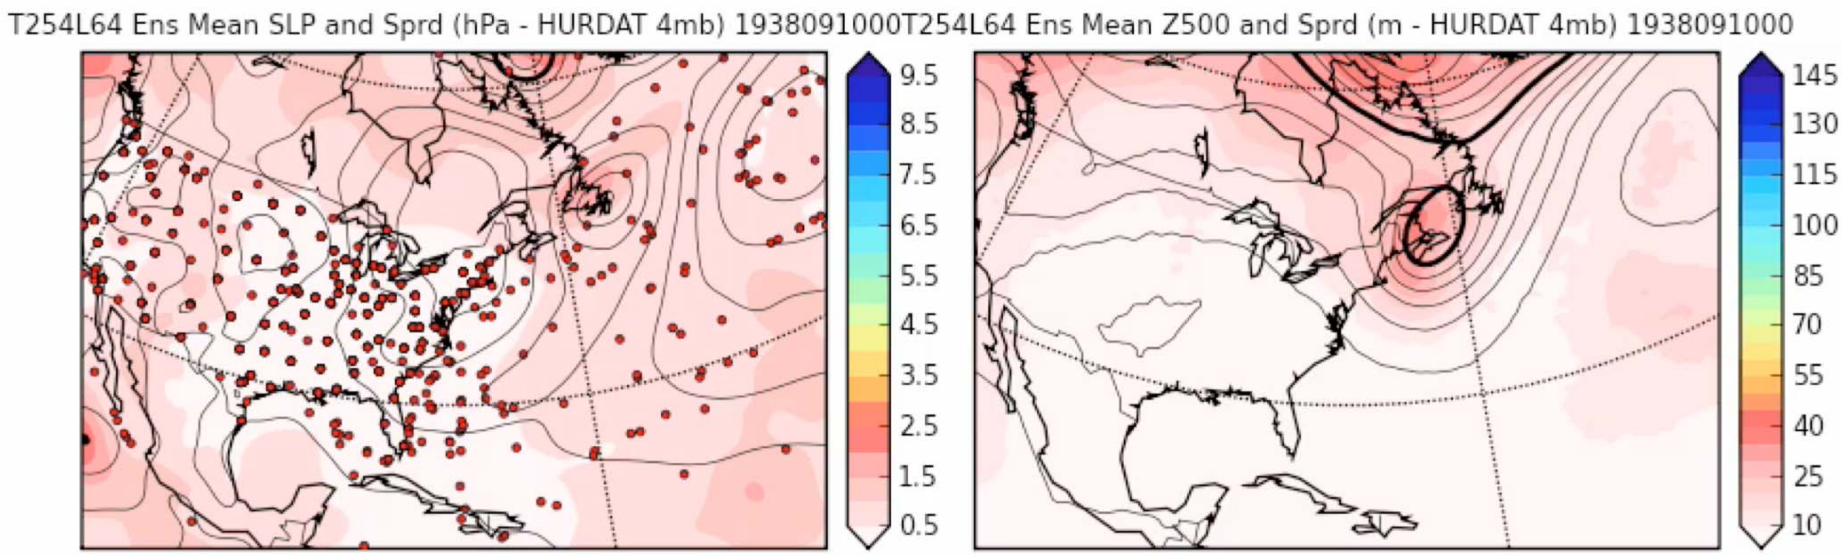

Mitigating Global Climate Change

Objective: **Determine if global warming can still be diminished if society cuts emissions of greenhouse gases.**

Implications: **Provide policymakers with appropriate research so they can make informed decisions to avoid the worst impacts of climate change.**

Accomplishments: **CCSM used at NERSC, ORNL, ANL, & NCAR to study a century of climate conditions, two CO_2 scenarios.**

- **70% cut in emissions would save arctic ice, reduce sea level rise**
- **Current students on catastrophic chance in Atlantic Meridional overturning circulation.**

BER

W. Washington (NCAR)

Simulations show how average surface air temperatures could rise if greenhouse gas emissions continue to climb at current rates (top), or if emissions are cut by 70% (bottom).

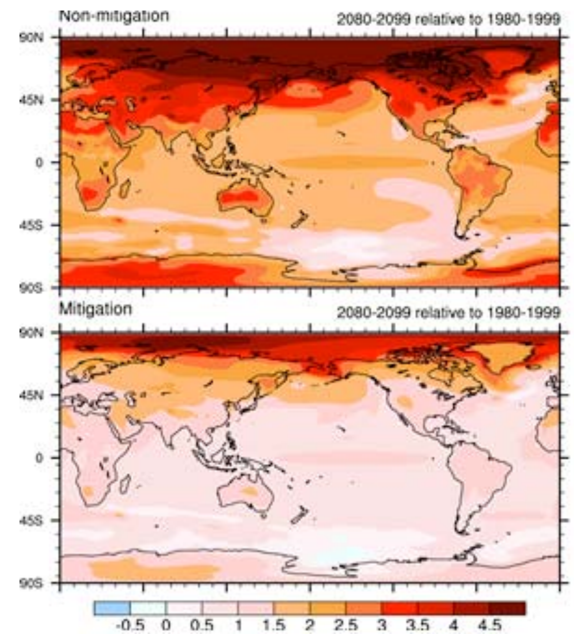

Temperatures rise by <2°C across nearly all populated areas if emissions are cut but unchecked emissions could lead to warming of >3°C in those areas.

Geophys. Res. Lett. 36, 08703 (2009)

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 Effort
  - Lead by the Federation of American Scientists

Digital Human Today: Imaging

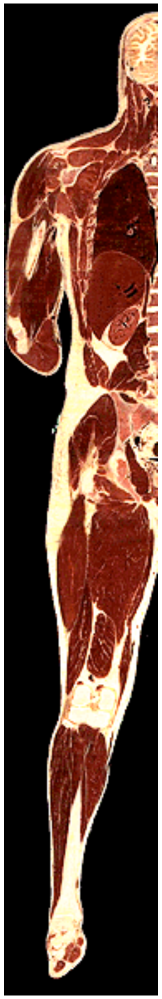

- The Visible Human Project
  - 18,000 digitized sections of the body
    - Male: 1mm sections, released in 1994
    - Female: .33mm sections, released in 1995
  - Goals
    - study of human anatomy
    - testing medical imaging algorithms
  - Current applications:
    - educational, diagnostic, treatment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

THE VISIBLE HUMAN PROJECT®

Image Source: www.madsci.org

Organ Simulation

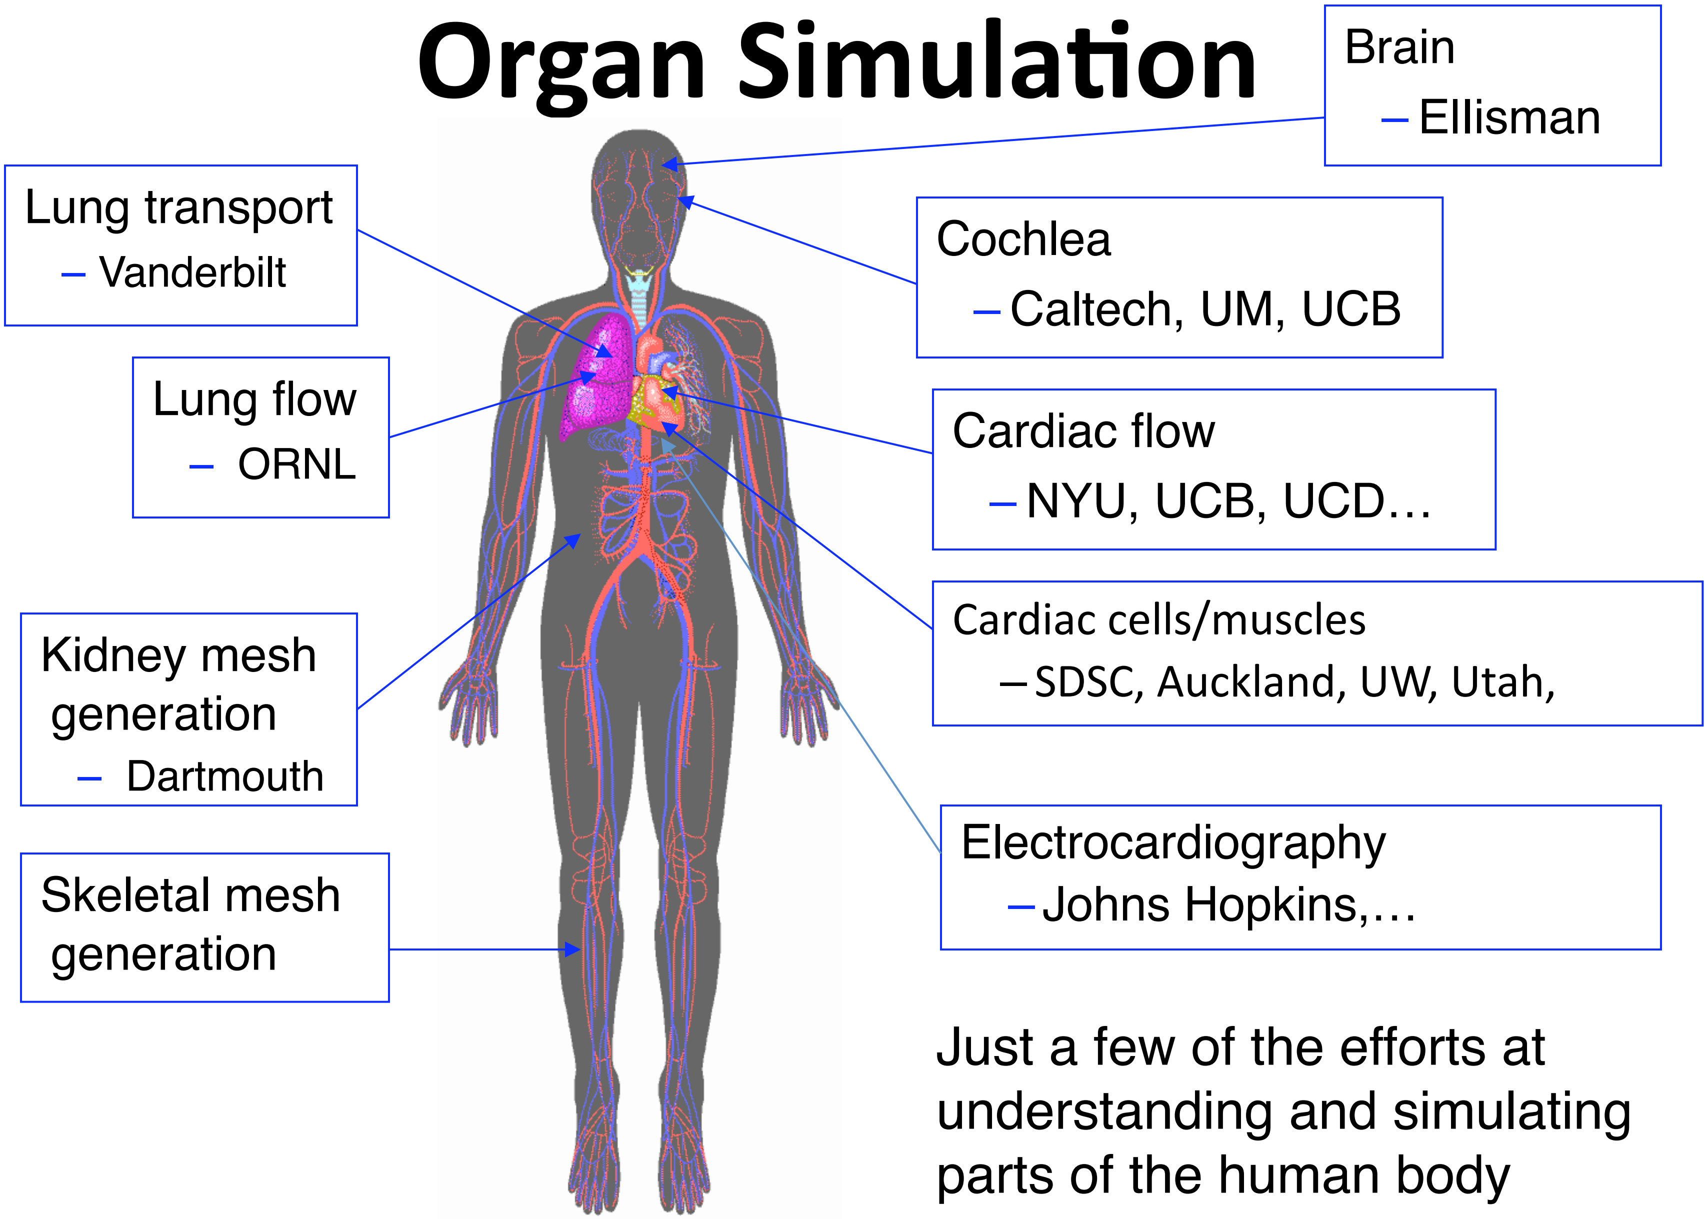

Just a few of the efforts at understanding and simulating parts of the human body

Heart Model

- Full structure shows cone shape
- Includes atria, ventricles, valves, and some arteries
- The rest of the circulatory system is modeled by
  - sources: inflow
  - sinks: outflow

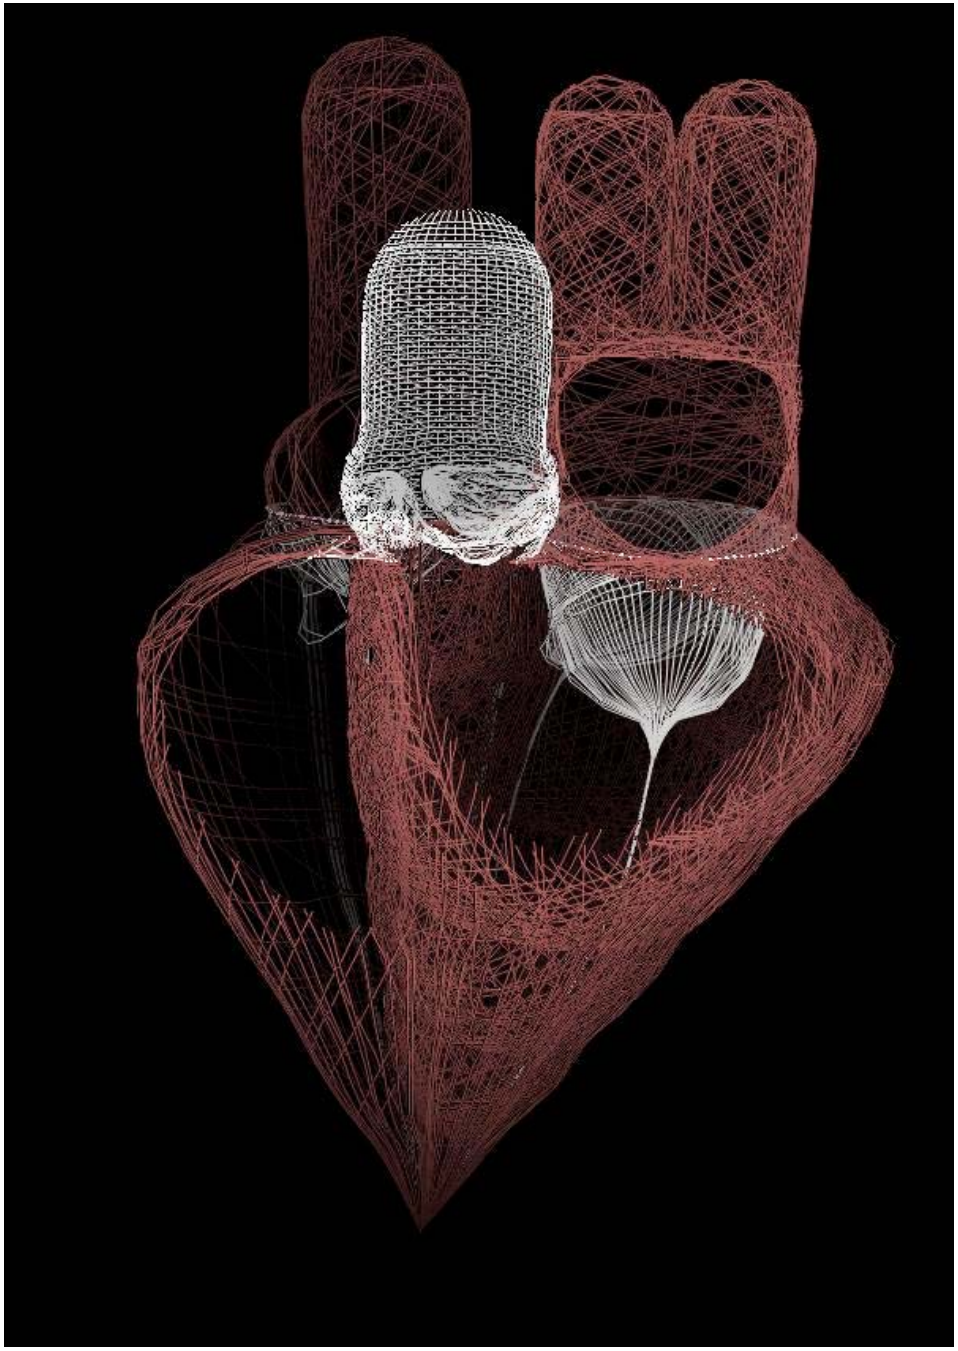

Heart Simulation

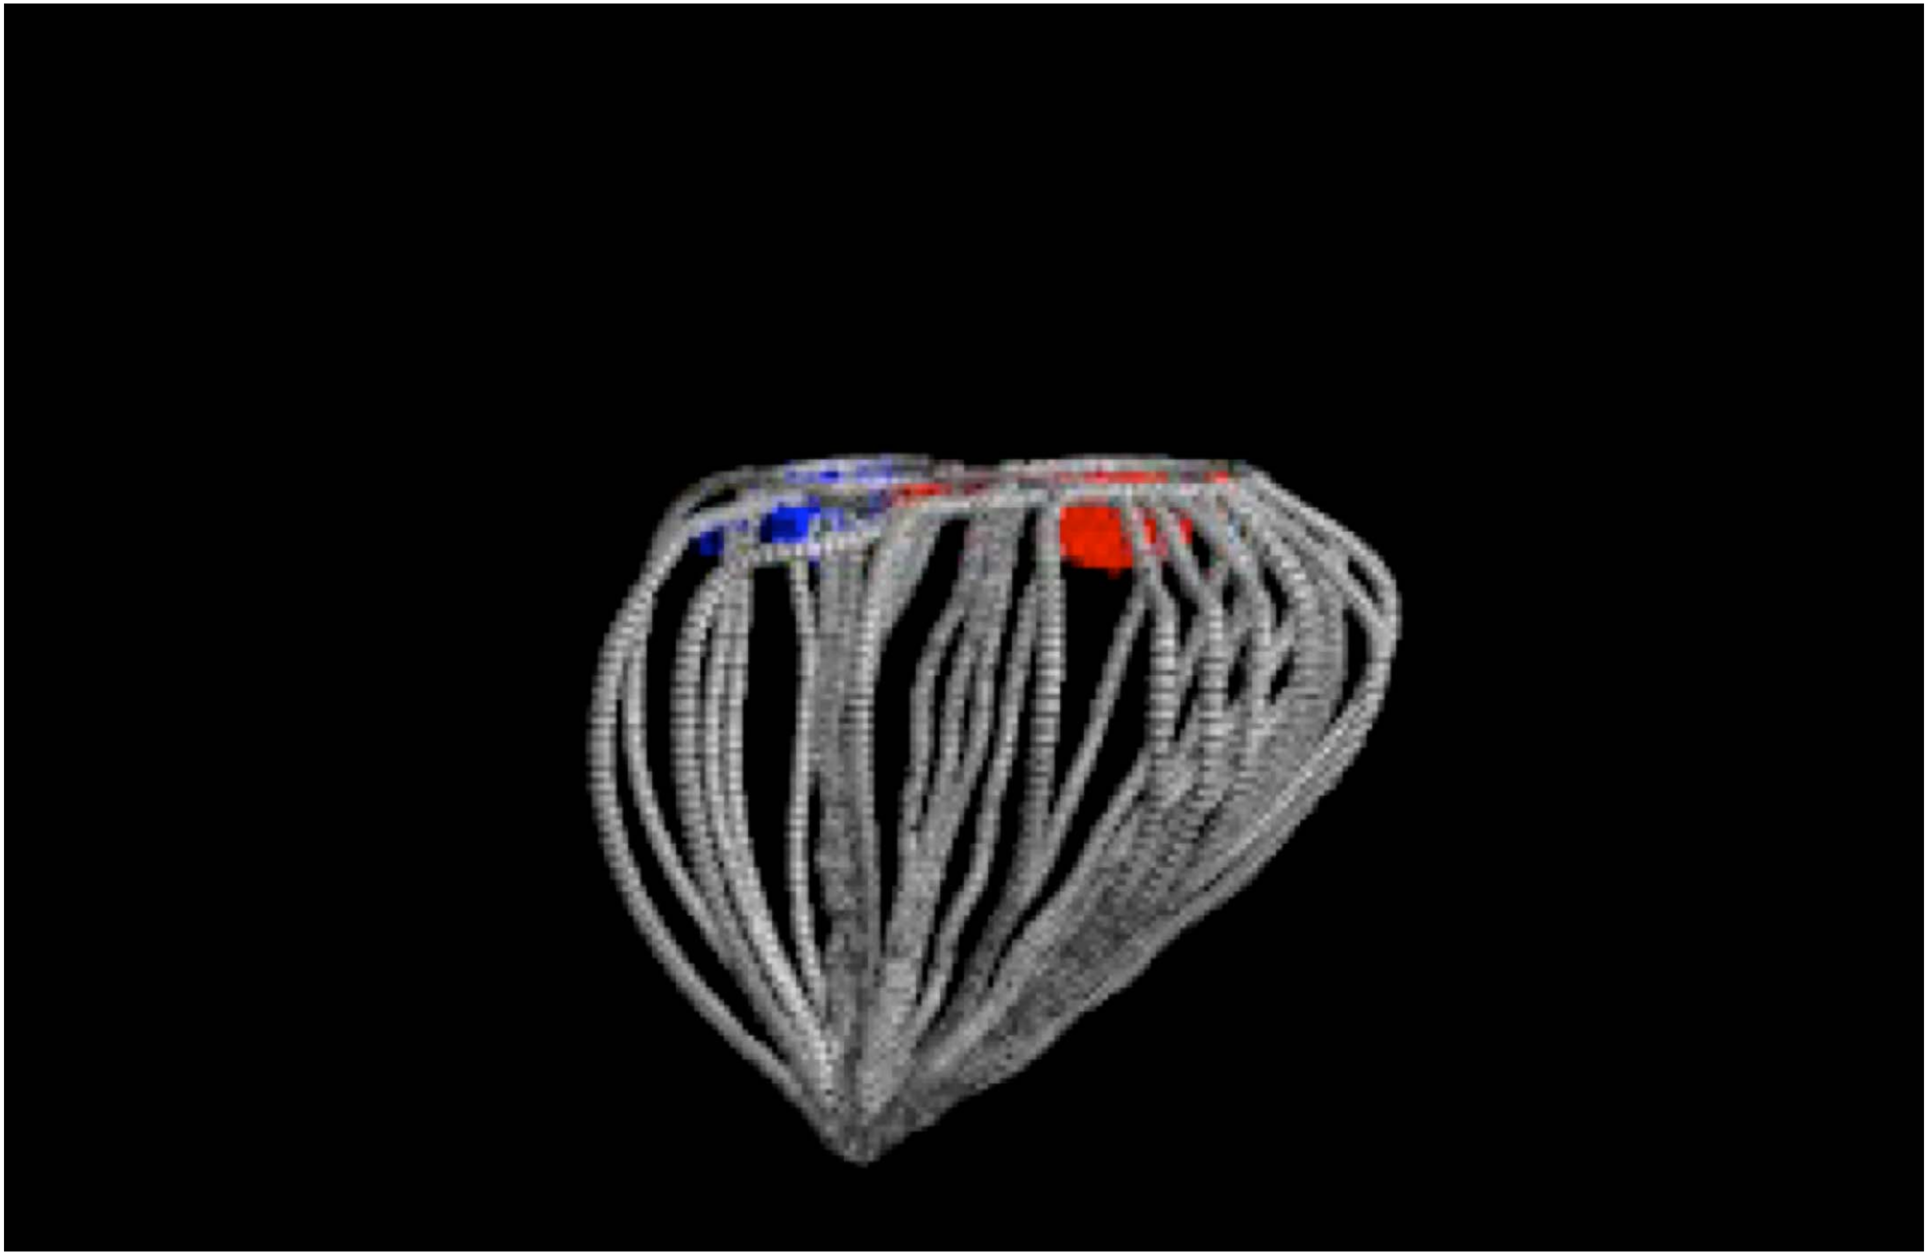

Immersed Boundary Method Structure

4 steps in each timestep

2D Dirac Delta Function

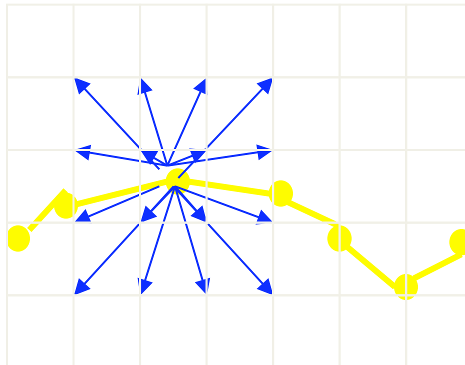

Material Points

1.Material activation & force calculation

2. Spread Force

Fluid Lattice

3. Navier-Stokes Solver

Interaction

4. Interpolate & move material

Why Study Computer Science?

- Because computers can help solve important problems
- **Because computers are fun to program**

Supercomputers are Parallel Computers

LBL's Franklin machine

- 38,000-cores (processors)
- Next one will have over 150,000

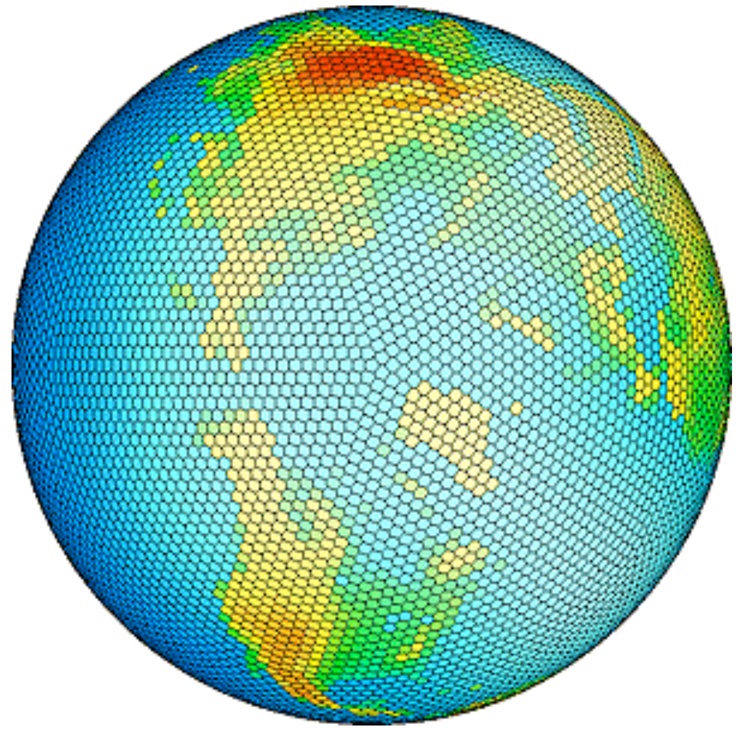

Supercomputing is done by parallel programming

Too Hot to Handle

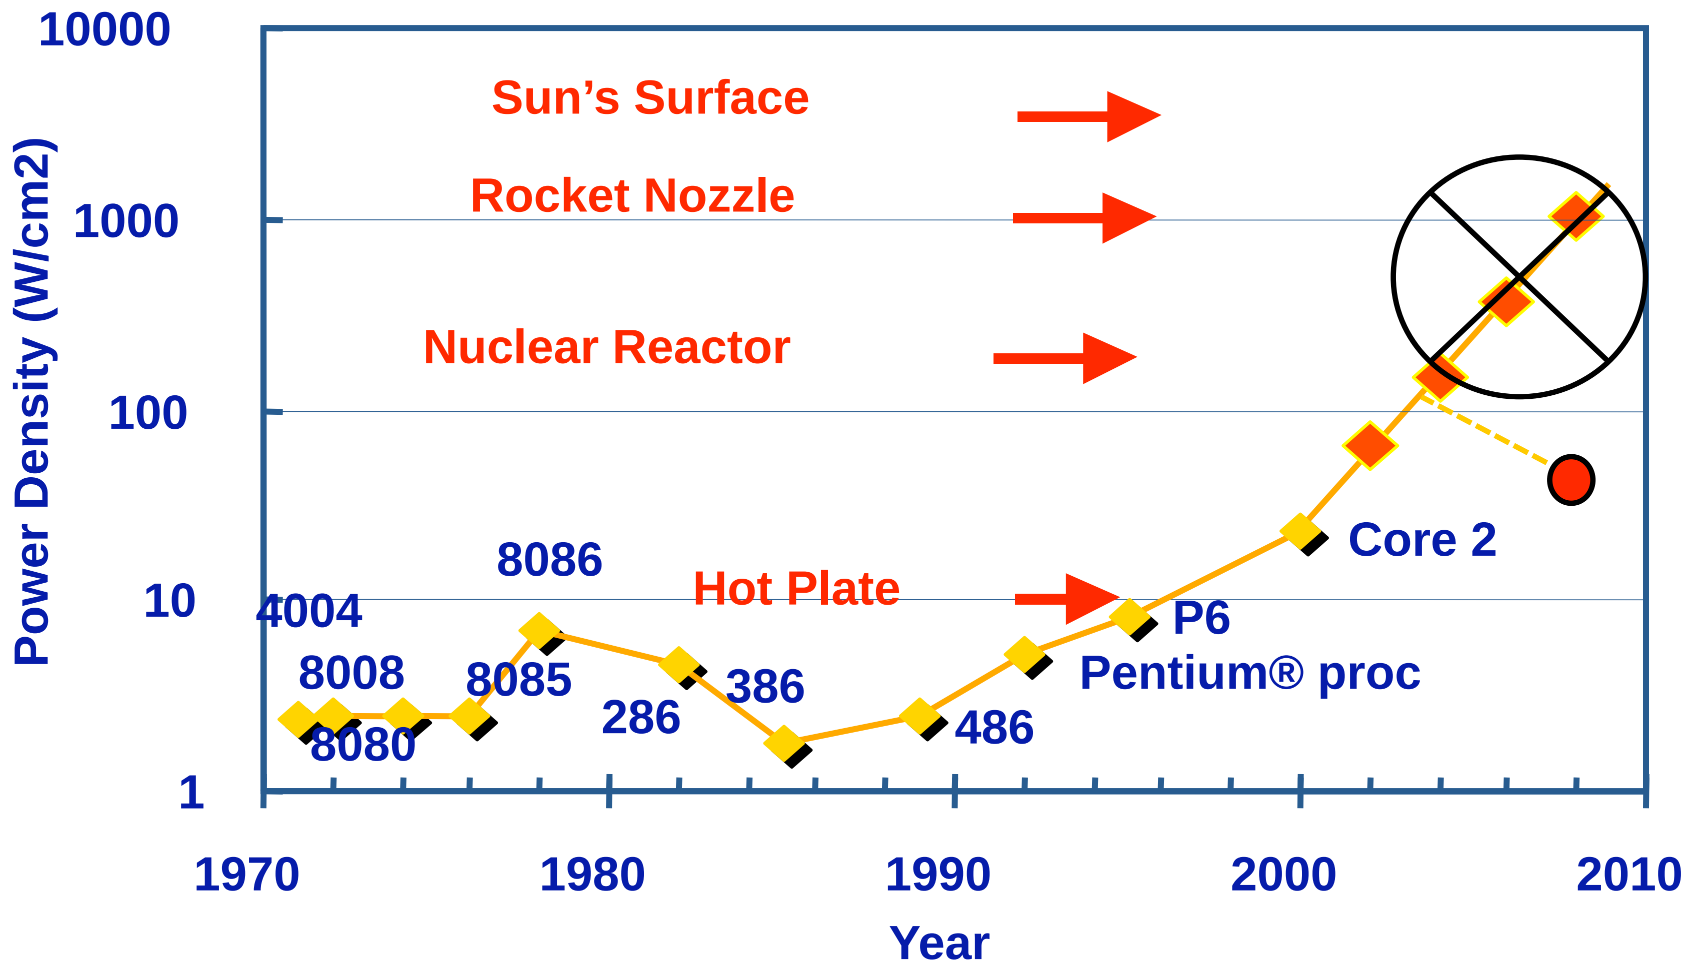

Source: S. Borkar (Intel)

More heat means more wasted energy

All Computers are Parallel

- Power density limit single processor clock speeds
- Cores per chip is growing
- How to program them?
  - Parallel "loops"
  - Parallel map
  - Parallel divide-and-conquer
  - (Message passing)

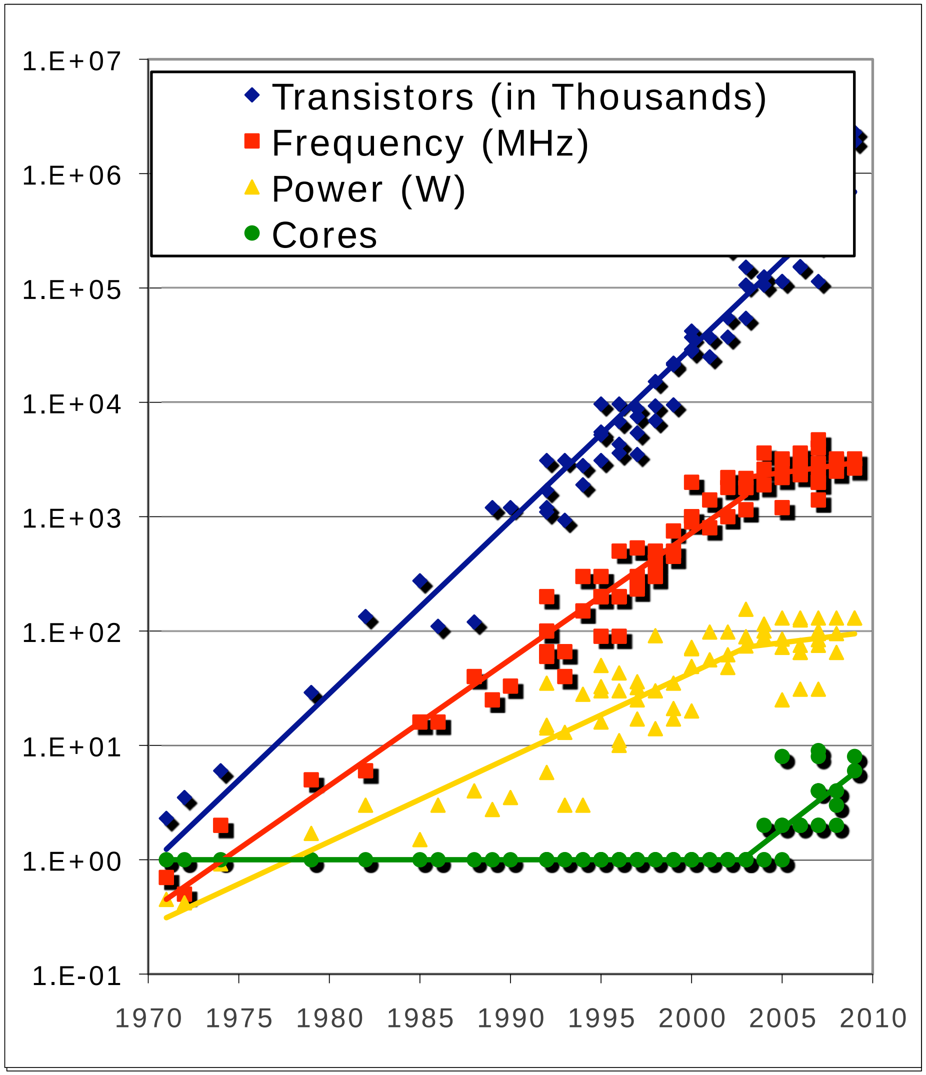

Parallelism is Green

The server processor is about 3x faster than the simple cell phone processor

But it uses 1300x more power so the cell phone is 400x more efficient

Why: Power is proportional to $V^2 f$, increasing frequency (f) also requires increase voltage V → cube

Can we build computers from lots of simple processors and save 100x in power?

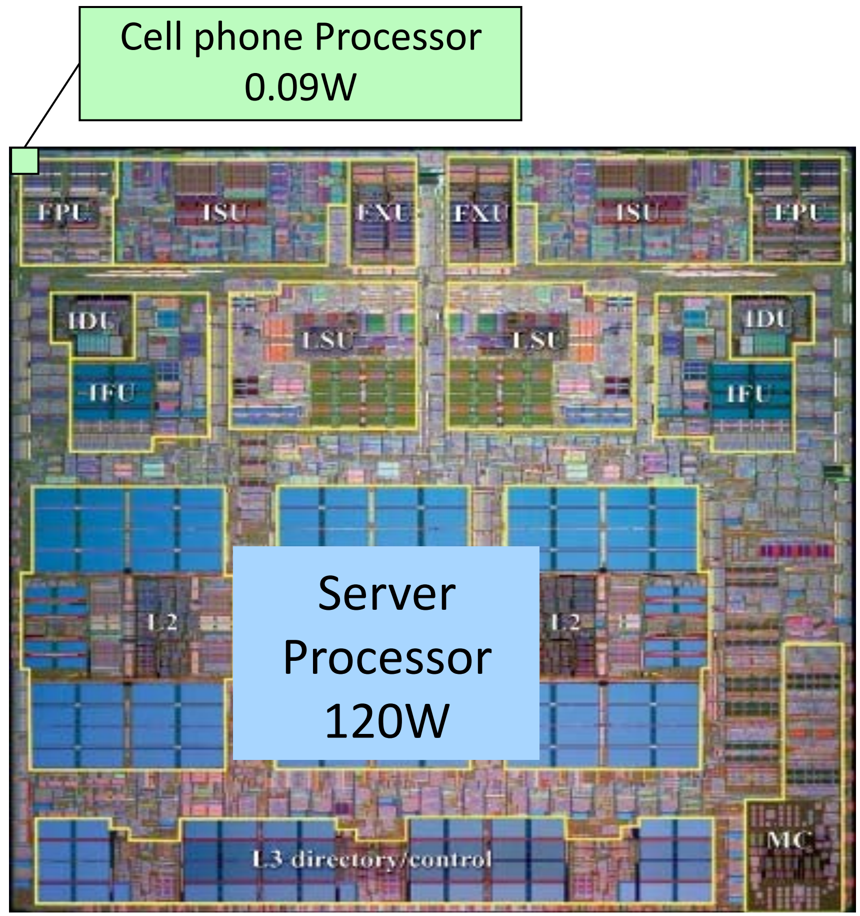

Why Study Computer Science?

- Because computers can help solve important problems
- Because computers are fun to program
- **Because computers make a good career**

Applications: What are the problems?

"Who needs 100 cores to run M/S Word?"

- Need compelling apps that use 100s of cores

How did we pick applications?

1. Enthusiastic expert application partner, leader in field, promise to help design, use, evaluate our technology
2. Compelling in terms of likely market or social impact, with short term feasibility and longer term potential
3. Requires significant speed-up, or a smaller, more efficient platform to work as intended
4. As a whole, applications cover the most important
   - Platforms (handheld, laptop, games)
   - Markets (consumer, business, health)

Computation in Music

(David Wessel)

- Musicians have an insatiable appetite for computation
  - More channels, instruments, more processing, more interaction!
  - Latency must be low (5 ms)
  - Must be reliable (No clicks)
- Music Enhancer
  - Enhanced sound delivery systems for home sound systems using large microphone and speaker arrays
  - Laptop/Handheld recreate 3D sound over ear buds
- Hearing Augmenter
  - Handheld as accelerator for hearing aid
- Novel Instrument User Interface
  - New composition and performance systems beyond keyboards
  - Input device for Laptop/Handheld

Berkeley Center for New Music and Audio Technology (CNMAT) created a compact loudspeaker array: 10-inch-diameter icosahedron incorporating 120 tweeters.

Computation in Photo Managment

(Kurt Keutzer)

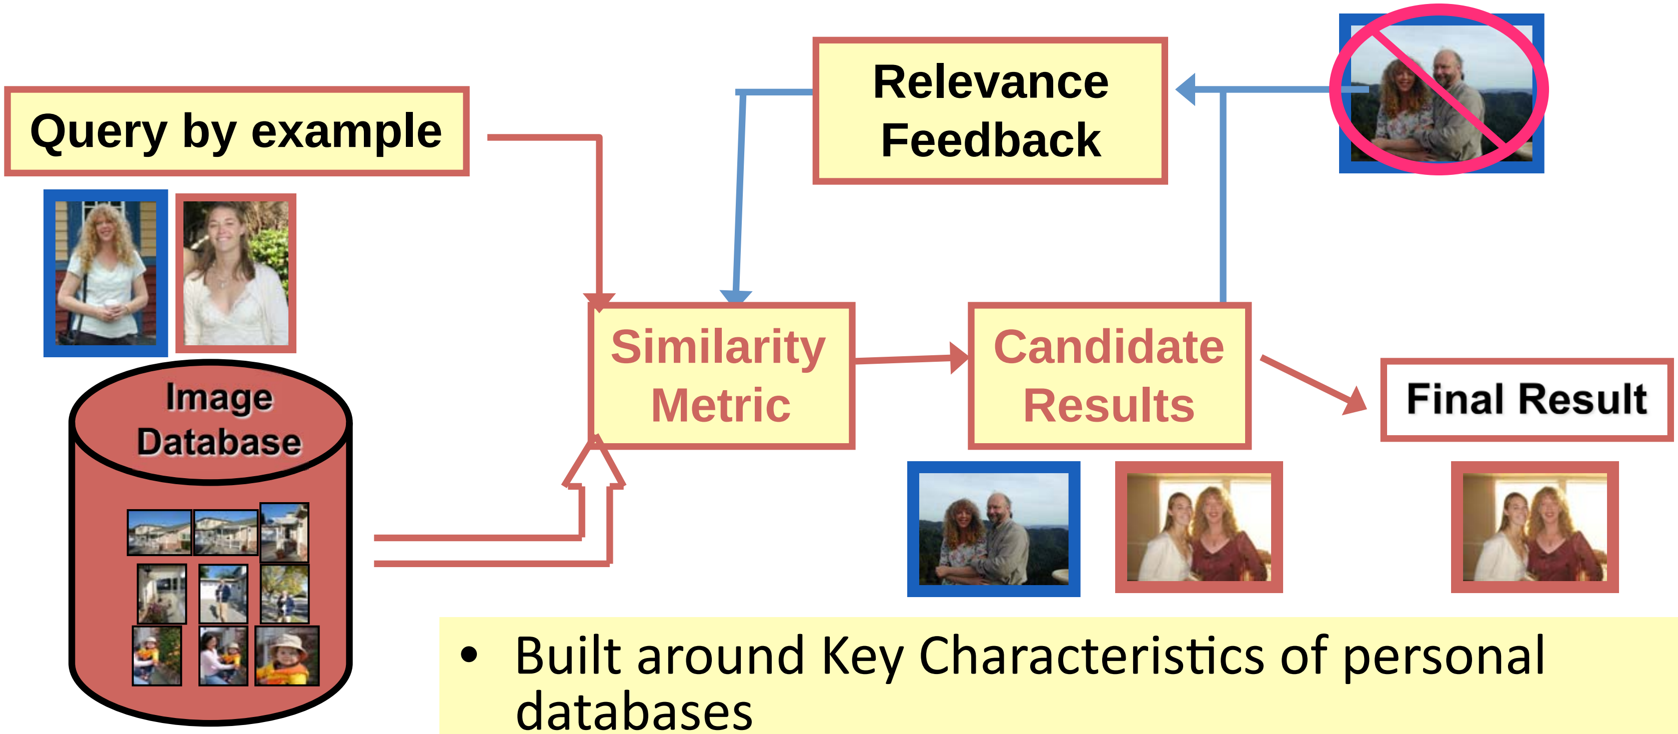

1000's of images

- Built around Key Characteristics of personal databases
  - Very large number of pictures (>5K)
  - Non-labeled images
  - Many pictures of few people
  - Complex pictures including people, events, places, and objects

Meeting Diarist and Teleconference Aid

(Nelson Morgan)

- Meeting Diarist
  - Laptops/ Handhelds at meeting coordinate to create speaker identified, partially transcribed text diary of meeting

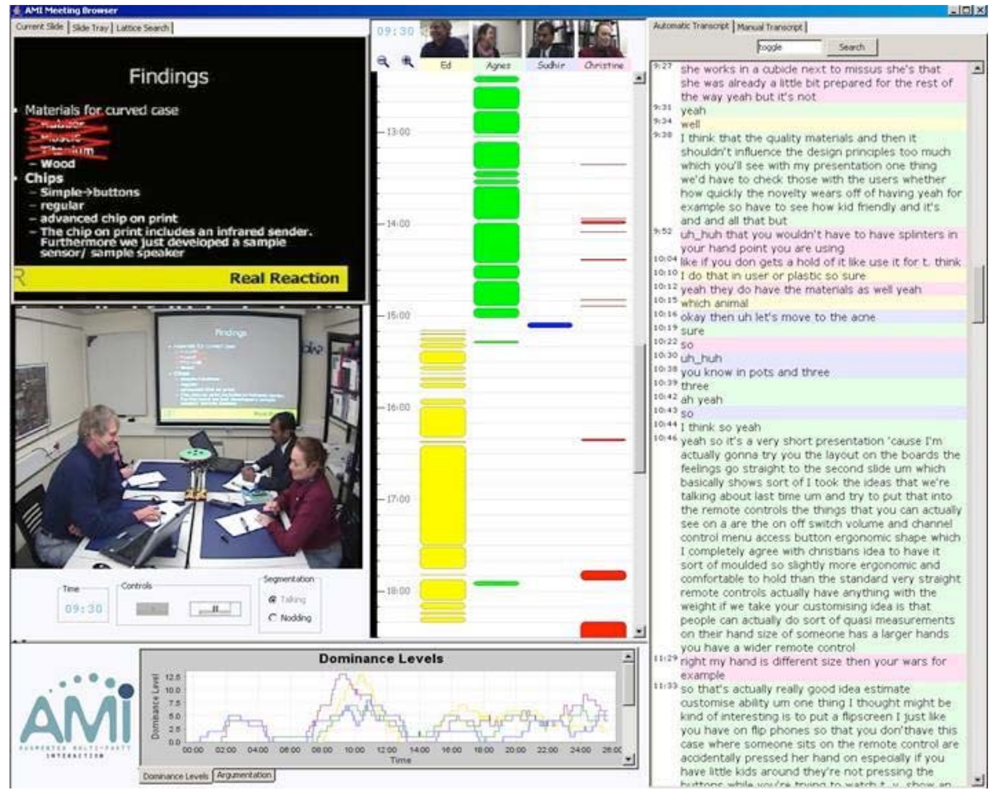

- **Teleconference speaker identifier**
  - **L/Hs used for teleconference, identifies who is speaking, "closed caption" hint of what being said**

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

Computational Science is Necessarily Collaborative

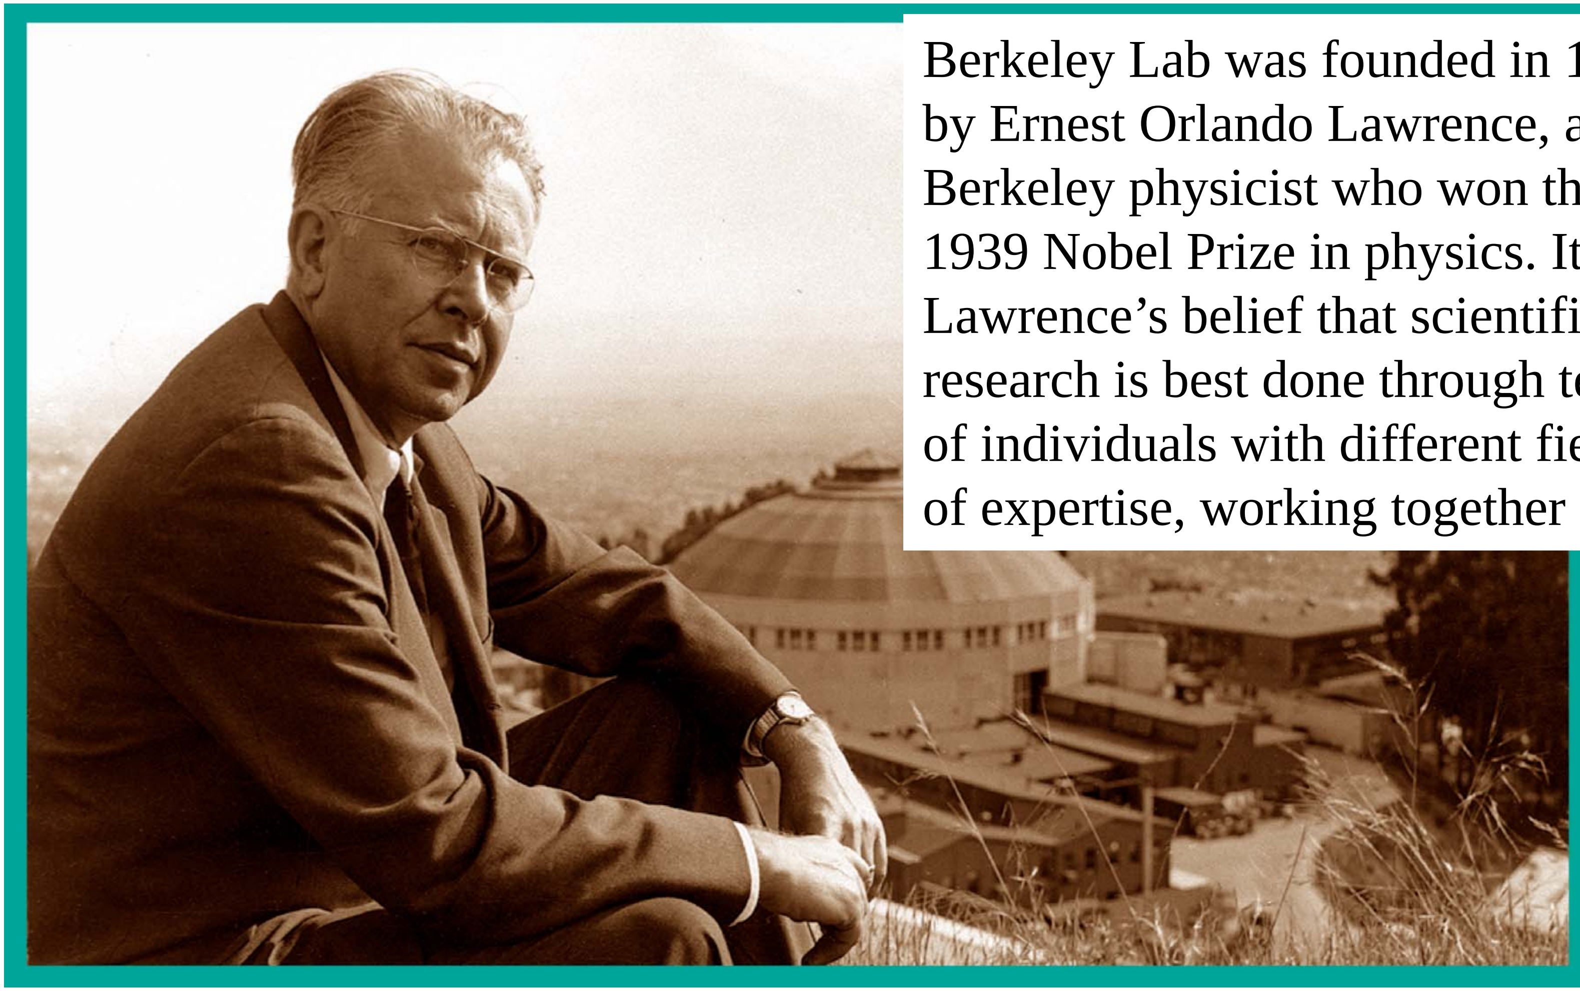

Berkeley Lab was founded in 1931 by Ernest Orlando Lawrence, a UC Berkeley physicist who won the 1939 Nobel Prize in physics. It was Lawrence's belief that scientific research is best done through teams of individuals with different fields of expertise, working together

… 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
The evolution of the cyclotron
Nobel Lecture, December 11, 1951

A National Laboratory Next to a University Campus

Berkeley Lab's Major Scientific Facilities Serving Universities, Industry, and Government

Advanced Light Source

88-Inch Cyclotron

Joint Genome Institute

National Center for Electron Microscopy

National Energy Research Scientific Computing Center

Molecular Foundry

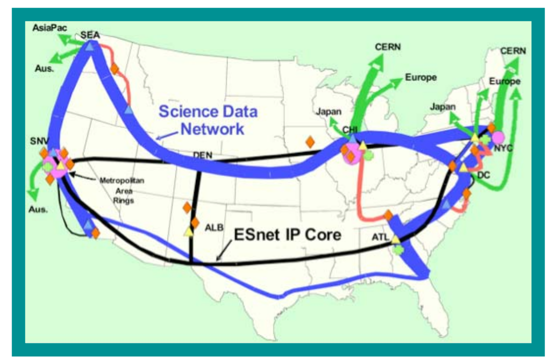

Energy Sciences Network (ESnet)