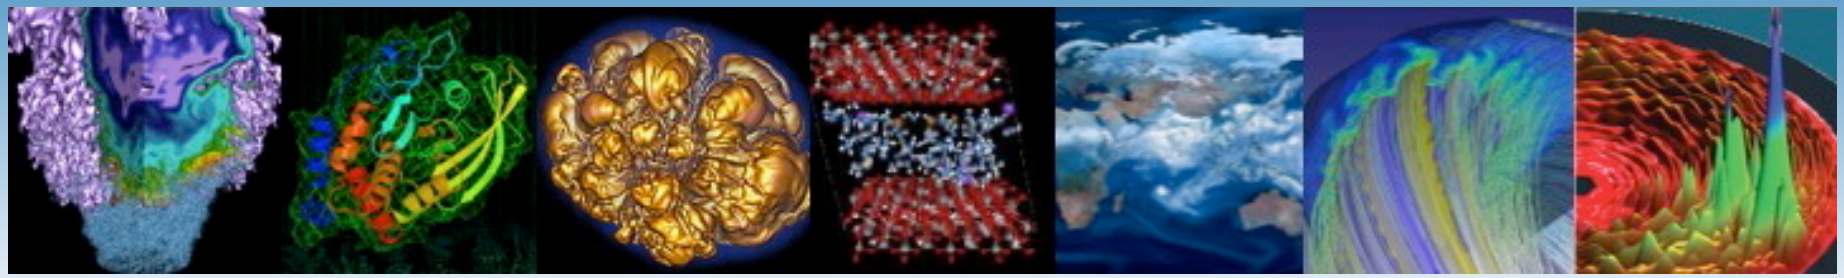

Introduction to HPC

Richard Gerber

National Energy Research Scientific Computing Center

Lawrence Berkeley National Laboratory

June 12, 2012

BERKELEY LAB

NERSC

Outline

- What is HPC?
- Who uses HPC and what are they doing with it?
- What is a "supercomputer?"
- Challenges in HPC
- How do you use a supercomputer?

Productivity?

What is High Performance Computing?

The application of "supercomputers" to computational problems that are either too large for standard computers or would take too long.

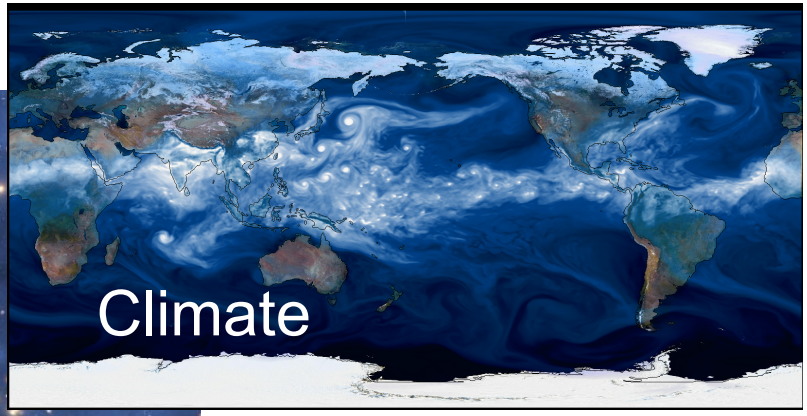

Protein Dynamics

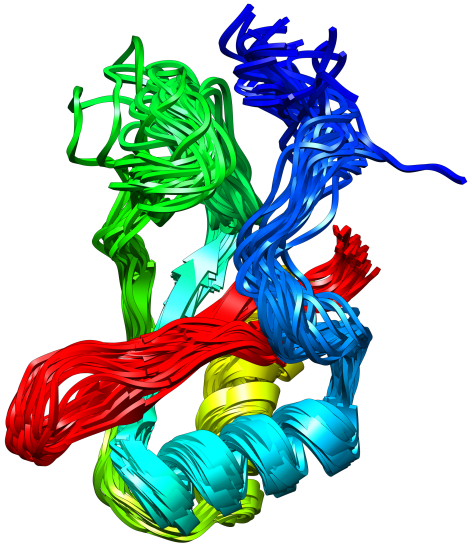

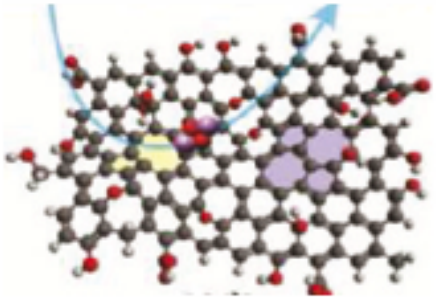

Better Batteries

What's So "Super"?

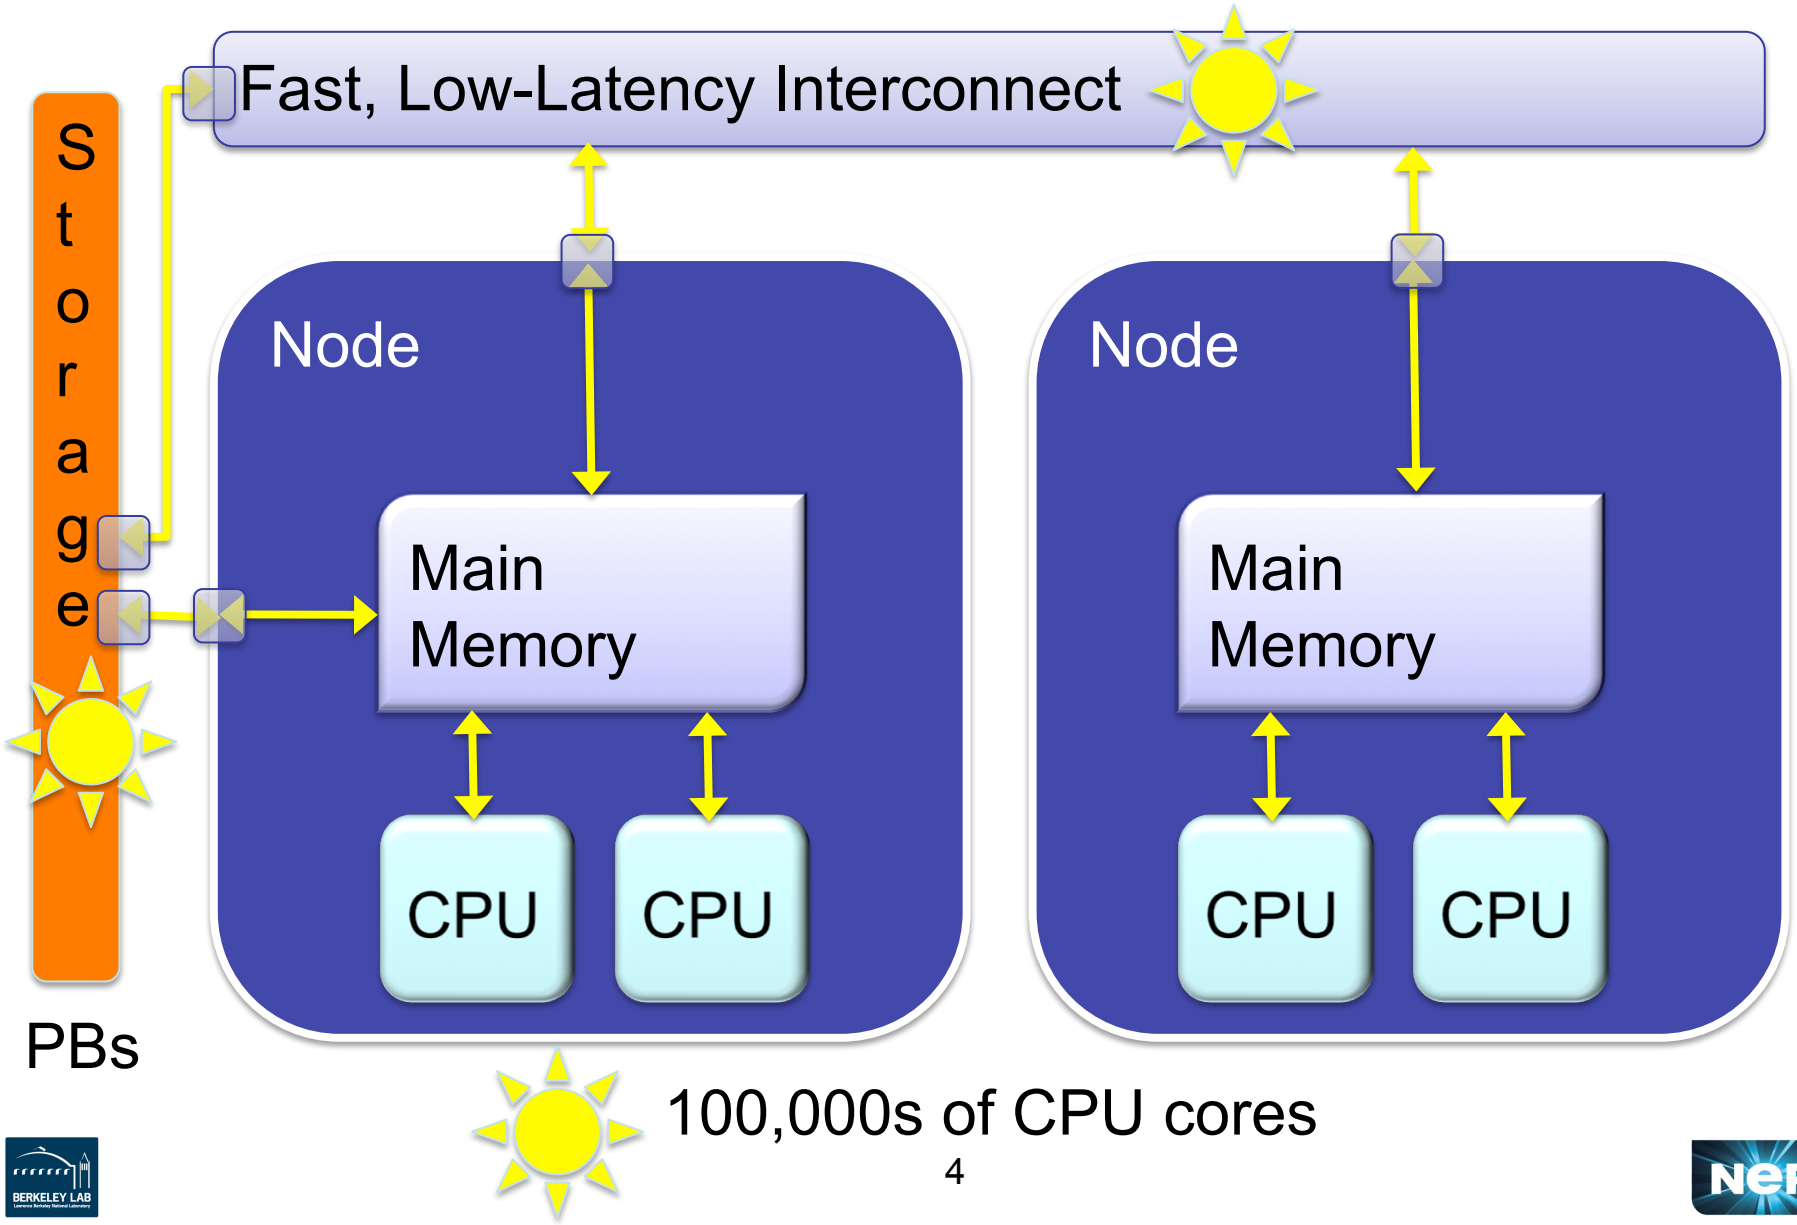

What is a Supercomputer?

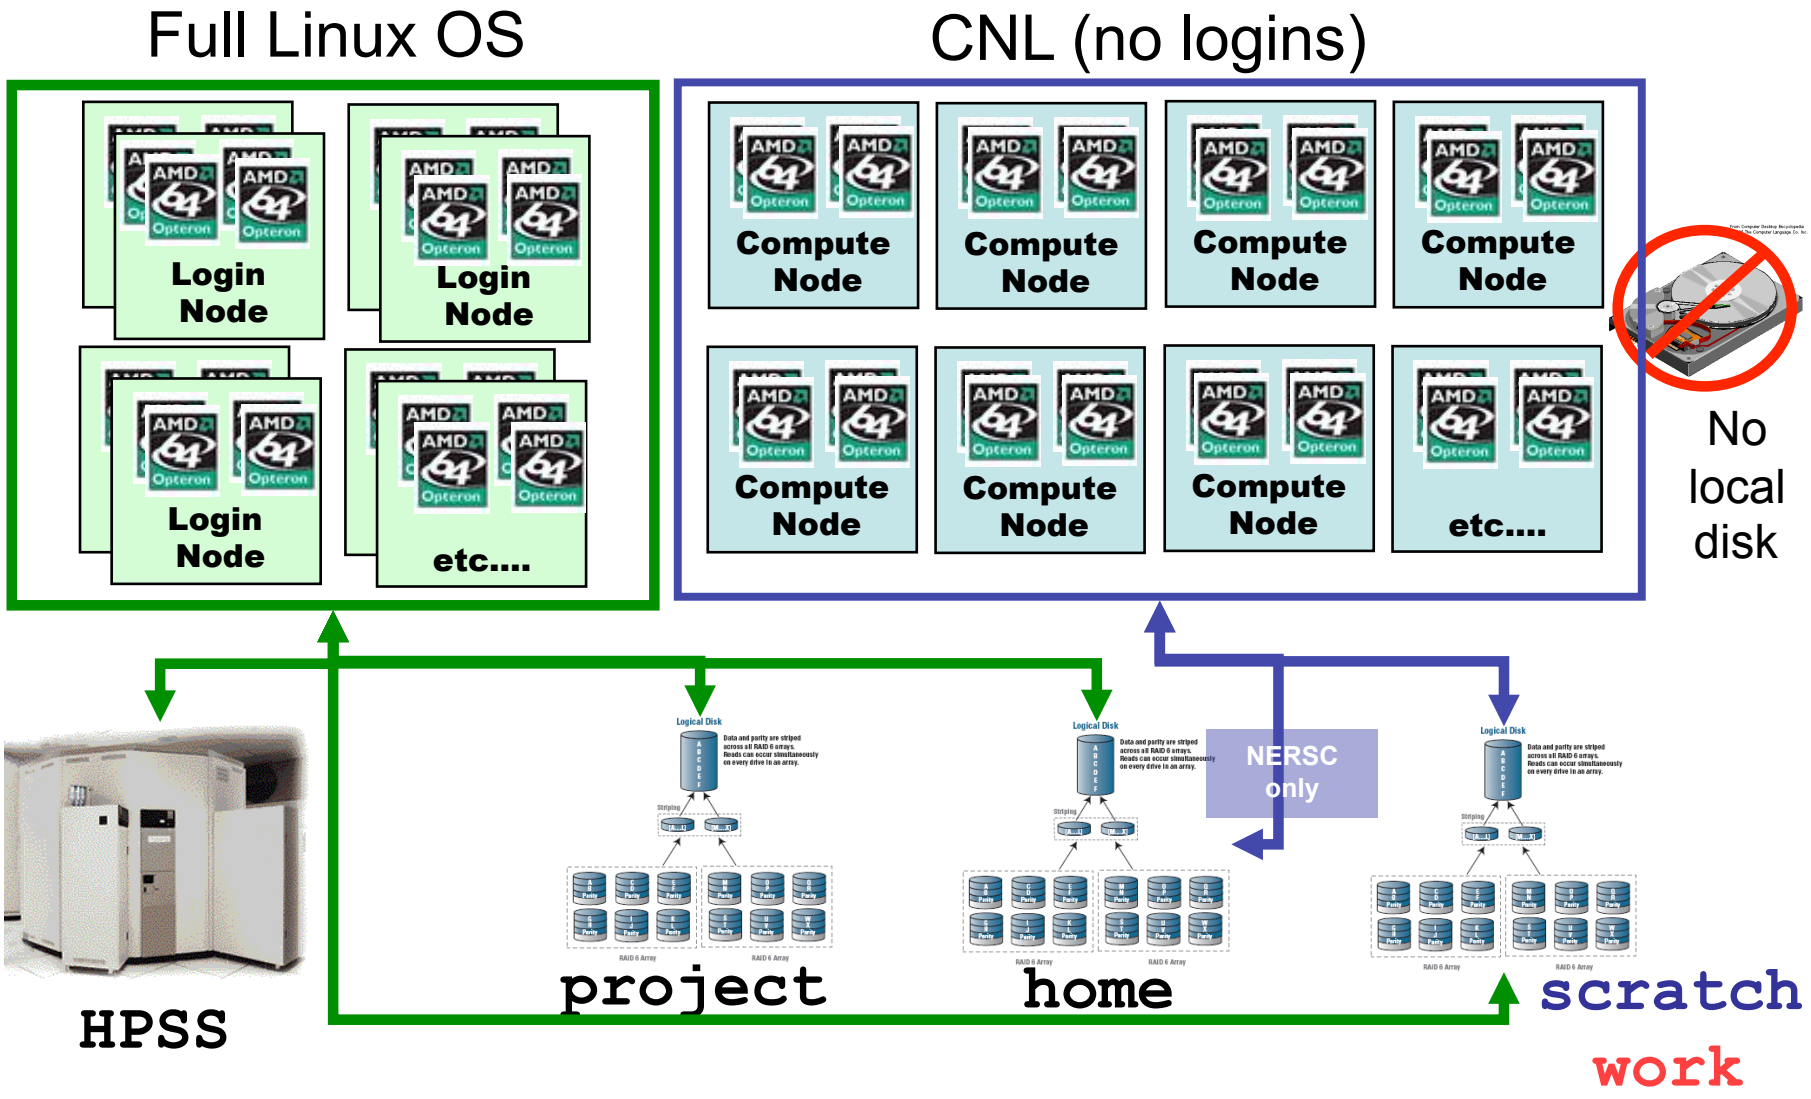

Who uses HPC?

- Scientists and engineers
  - climate prediction
  - protein folding simulations
  - oil and gas discovery
  - defense and aerospace work
  - automotive design
  - financial forecasting, etc
- Corporations
  - customer records
  - inventory management
  - employee details

Why Use Supercomputing?

Length (m)	Phenomena
10^{-18}-10^{-15}	quarks, strings
10^{-15}-10^{-12}	proton, neutron
10^{-12}-10^{-9}	gamma rays, X rays, hydrogen atom
10^{-9}-10^{-6}	DNA, virus, optical light
10^{-6}-10^{-3}	bacteria, fog, human hair
10^{-3}-10^{0}	mosquito, golf ball, football
10^{0}-10^{3}	people, football field, Eiffel tower
10^{3}-10^{6}	Mt. Everest, Panama Canal, asteroid
10^{6}-10^{9}	Moon, Earth, light-second
10^{9}-10^{12}	Sun, light-minute, Earth's orbit
10^{12}-10^{15}	Solar System
10^{15}-10^{18}	light-year, nearest star
10^{18}-10^{21}	galactic arm
10^{21}-10^{24}	Milky Way, distance to Andromeda galaxy
10^{24}-10^{26}	visible universe

Direct Human Experience

Why Use Supercomputing

Time Scale (s)	Phenomena
10^{-44}	Planck time
10^{-24}	light crosses nucleus
10^{-15}	atomic vibration, visible light
10^{-12}	IBM SiGe transistor
10^{-9}	1 Gz CPU
10^{-6}	protein folding, lightning bolt
10^{-3}	hard disk seek time, blink of an eye
10^{0}	earthquakes
10^{2}	tornadoes
10^{5}	hurricanes
10^{7}	year
10^{9}	human life span
10^{10}	deep ocean mixing time
10^{12}	first homo sapiens
10^{15}	Milky Way rotation period
10^{17}	age of universe

Direct Human Experience

Why Use Supercomputing?

- Design
- Prediction
- Explore dangerous or inaccessible domains

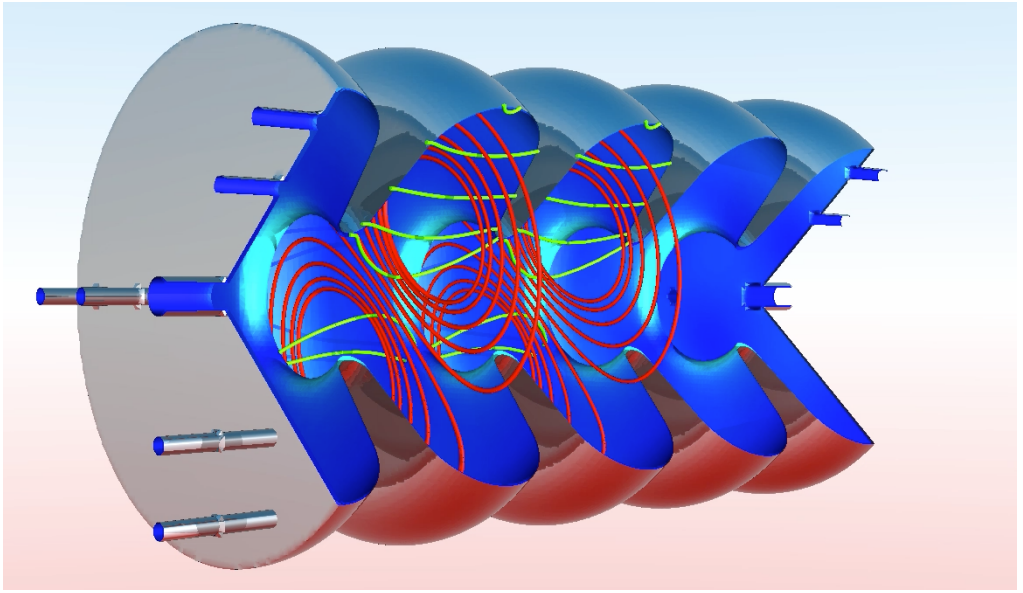

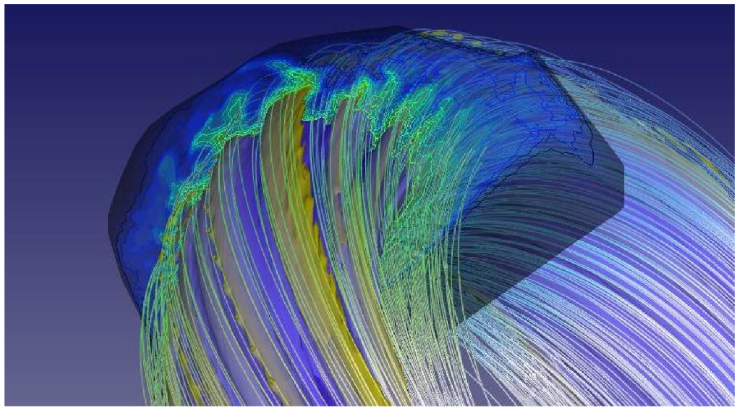

Supercomputers: Access to the Universe
Past, Present, & Future

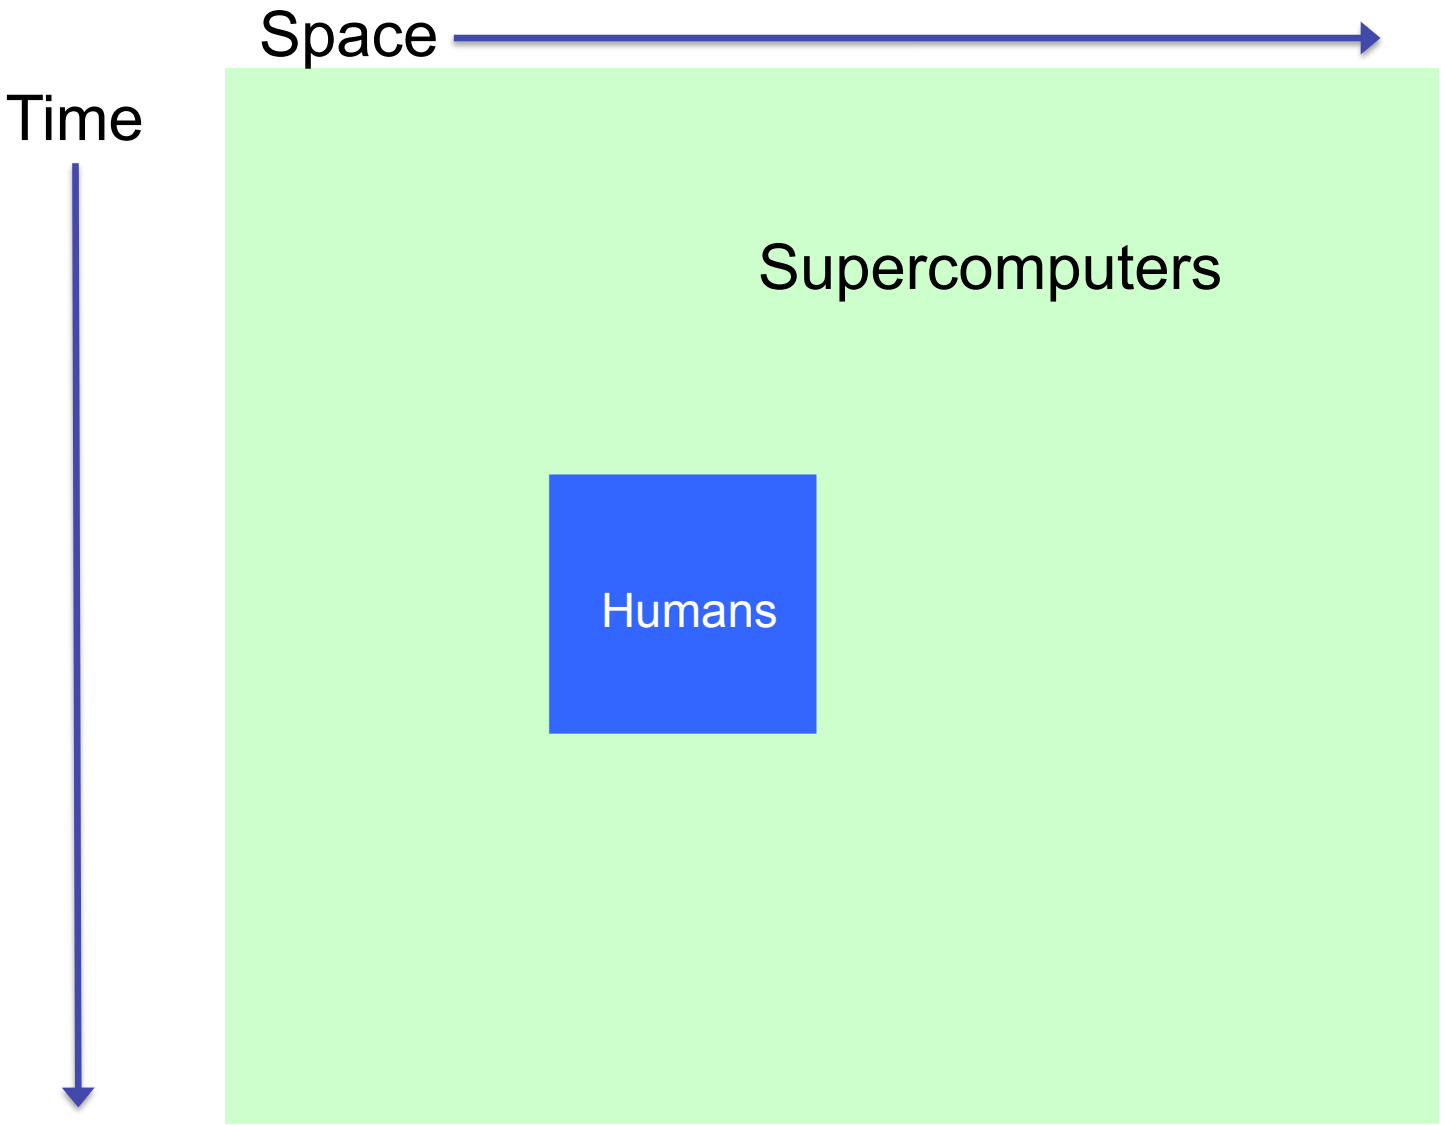

NERSC www.nersc.gov

NERSC computing for science

- 4000 users, 500 projects
- From 48 states; 65% from universities
- Hundreds of users each day
- ***1500 publications per year***

Systems designed for science

- 1.3PF Petaflop Cray system, Hopper
  - 8th Fastest computer in US
  - Fastest open Cray XE6 system

National Energy Research Scientific Computing Center (NERSC)

- NERSC is a national supercomputer center funded by the U.S. Department of Energy Office of Science (SC)
  - Located at Berkeley Lab
- SC supports a broad spectrum of energy-related research
  - Largest funding source for basic physical science research in the U.S.
- If you have SC funding, you can use NERSC
  - Other researchers can apply if research is in SC mission
- Breadth of SC (and therefore NERSC) research is represented by 6 OS Program Offices

Current NERSC Systems

Large-Scale Computing Systems

Hopper (NERSC-6): Cray XE6

- 6,384 compute nodes, 153,216 cores
- 144 Tflop/s on applications
- 1.3 Pflop/s peak

Midrange

140 Tflops total

Carver

- IBM iDataplex cluster
- 9884 cores; 106TF

PDSF (HEP/NP)

- ~1K core cluster

GenePool (JGI)

- ~5K core cluster
- 2.1 PB Isilon File System

NERSC Global Filesystem (NGF)

Uses IBM's GPFS

- 8.5 PB capacity
- 15GB/s of bandwidth

HPSS Archival Storage

- 240 PB capacity
- 5 Tape libraries
- 200 TB disk cache

Analytics & Testbeds

Euclid
(512 GB shared memory)

Dirac 48 Fermi GPU nodes

Magellan Hadoop

Basic Energy Sciences (BES)

Batteries

Breakthroughs in battery technologies may extend the range of electric cars.
(D. Mei, PNNL, L-W. Wang, LBNL)

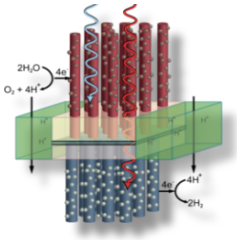

Artificial Photosynthesis

Simulation is playing a key role in highly visible quest to develop artificial photosynthesis.
(L. Wang, LBNL)

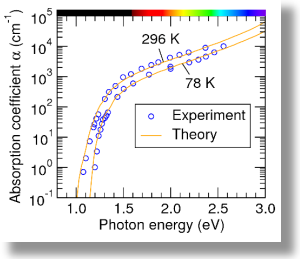

Solar Energy

A NERSC "NISE" award and software by a NERSC consultant yield an important new method for characterizing solar energy materials.
(E. Kioupakis, U. Michigan)

Coal Gasification

NERSC resources were used to model a real coal gasifier with a Large Eddy Simulation code.
(P. Smith, U. Utah)

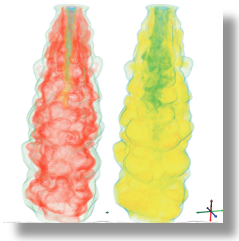

Nanotechnology

Computation explains the size-sensitive melting behavior of metal nanoclusters.
(S. Wei, NREL)

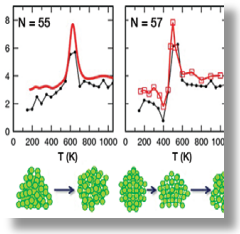

Biological and Environmental Research (BER)

Extreme Climate

New techniques help detect extreme events buried in immense data sets.
(Prabhat, M. Wehner, LBNL)

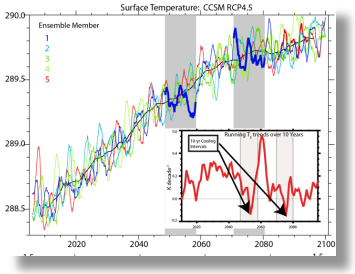

Oceanic Heat Reservoirs

Key finding that deep oceans can mask global warming for decade-long periods
(G. Meehl, A. Hu, NCAR)

Genomes

Genomes pipeline at NERSC can process 100 million genes in few days, a task that used to require weeks at the JGI.
(V. Markowitz, LBNL)

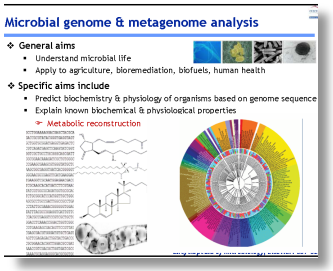

Aerosol Effects

Atmospheric scientists have shown how small-scale effects of aerosols contribute to errors in climate models.
(W. Gustafson, PNNL)

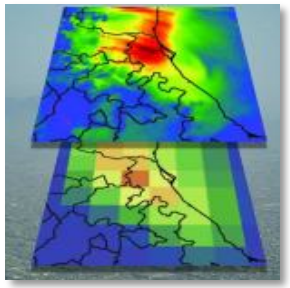

High Energy Physics (HEP)

Acceleration of the Universe

NERSC played a key role in the discovery that led to the 2011 Nobel Prize in Physics.
(S. Perlmutter, UC Berkeley/LBNL)

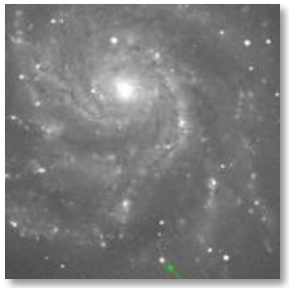

Supernova

The earliest-ever detection of a supernova was made possible by NERSC and Esnet.
(P. Nugent, LBNL)

Neutrino Decay

An important piece of the neutrino puzzle has fallen into place, thanks to data transfer, storage, analysis, archive, and gateway capabilities at NERSC and ESnet.
(K.-B. Luk, LBNL)

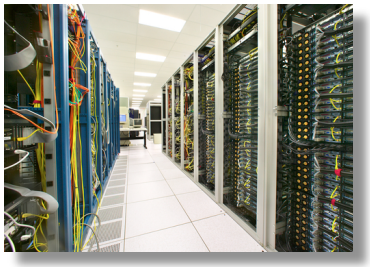

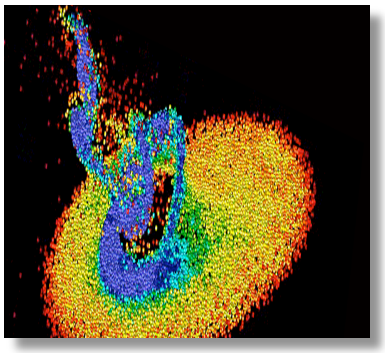

Dark Matter

Simulations done at NERSC helped validate a key new method that reveals a dark companion to the Milky Way.
(S. Chakrabarti, UC Berkeley)

Nuclear Physics (NP)

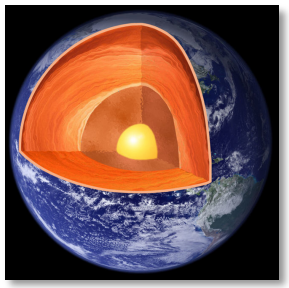

Nuclear Decay Heating

The KamLAND neutrino experiment showed that radioactivity cannot be Earth's only heat source; it accounts for only ½ of it.
(S. Freedman, LBNL)

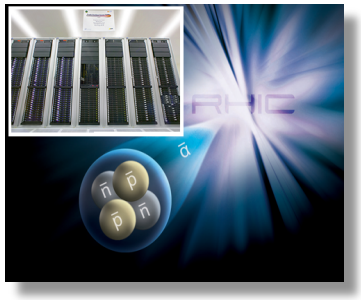

Antimatter

The heaviest antimatter particle has been discovered with NERSC help. Antihelium-4 is likely to hold the title for decades.
(STAR Collaboration at RHIC, BNL/ LBNL)

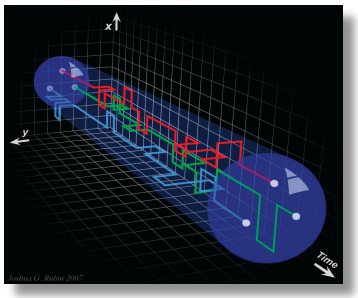

6-Quark Nucleons

Computations done at NERSC suggest the possible existence of a so-called H-dibaryon bound state, an exotic nucleus first envisaged in 1977.
(M. Savage, U. Washington)

Fusion Energy Sciences (FES)

Magnetic Reconnection

Magnetic reconnection simulations done at NERSC along with NASA Voyager probe data help shake up prevailing views of the solar system's outer reaches.
(J. Drake, U. Maryland)

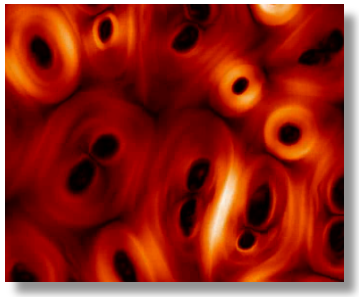

Voyager Surprise Explained

Study by award-winning researcher explained magnetic reconnection phenomenon observed by *Voyager* spacecrafts.
(J. Drake, U. Maryland)

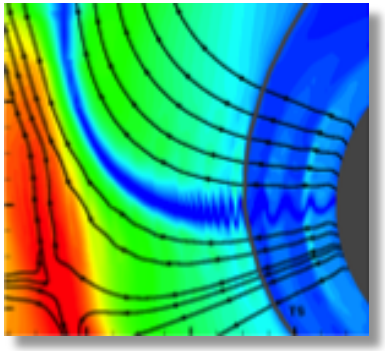

ITER Design

Study supported by NERSC NISE award suggests ITER might require an alternate mitigation strategy for "runaway electron" current.
(V. Izzo, General Atomics)

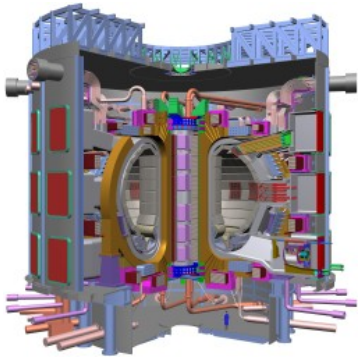

Advanced Scientific Computing Research (ASCR)

Visualization Technology

Demonstrated that visualization R&D has produced technology that can ingest and process tomorrow's "datasets" today.
(W. Bethel, LBNL)

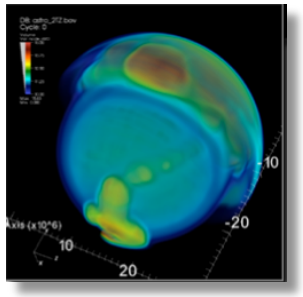

I/O Optimization

NERSC staff profiled & optimized HDF5 for Lustre and helped Cray optimize their MPI-IO, achieving a 10X improvement.

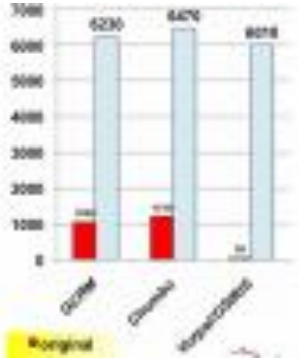

Programming Languages

Unified Parallel C (**UPC**) is an extension of the C programming language designed for high performance computing on large-scale parallel machines.
(K. Yelick, UC Berkeley/ LBNL)

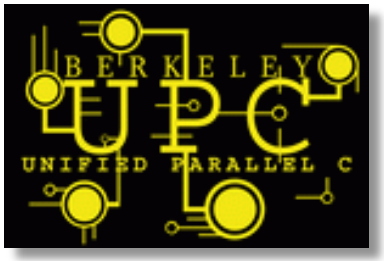

NERSC's Mission is to Enable Science

NERSC Mission: To *accelerate the pace of scientific discovery* by providing high-performance computing, data systems and services to the DOE Office of Science community.

NERSC has over 4500 users in 650 projects that produce about **1500 publications per year!**

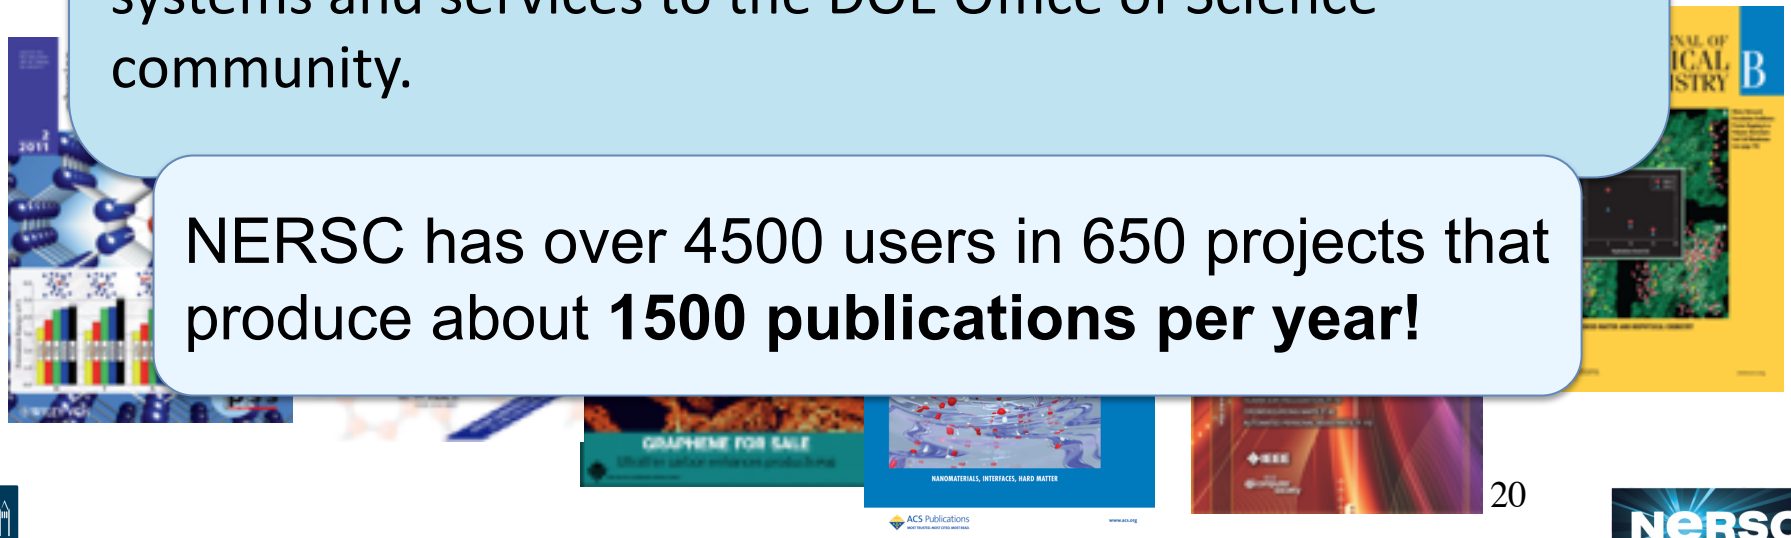

Video