

High-Performance Computing and NERSC



Presentation for CSSS Program

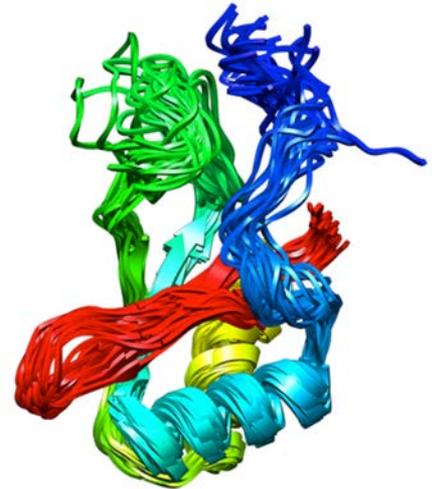
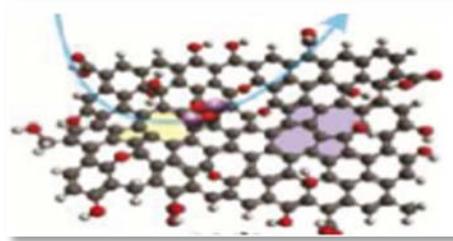
Rebecca Hartman-Baker, PhD
User Engagement Group Lead
June 3, 2021

High-Performance Computing Is...

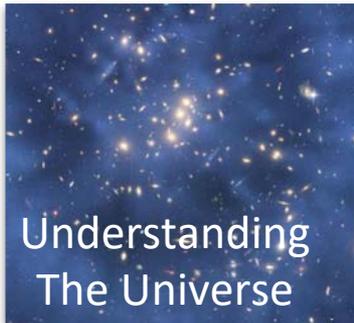
... the application of "supercomputers" to scientific computational problems that are either too large for standard computers or would take them too long.



Designing Better Batteries



Understanding How Proteins Work





What Is a Supercomputer?

A Supercomputer Is...



VS.

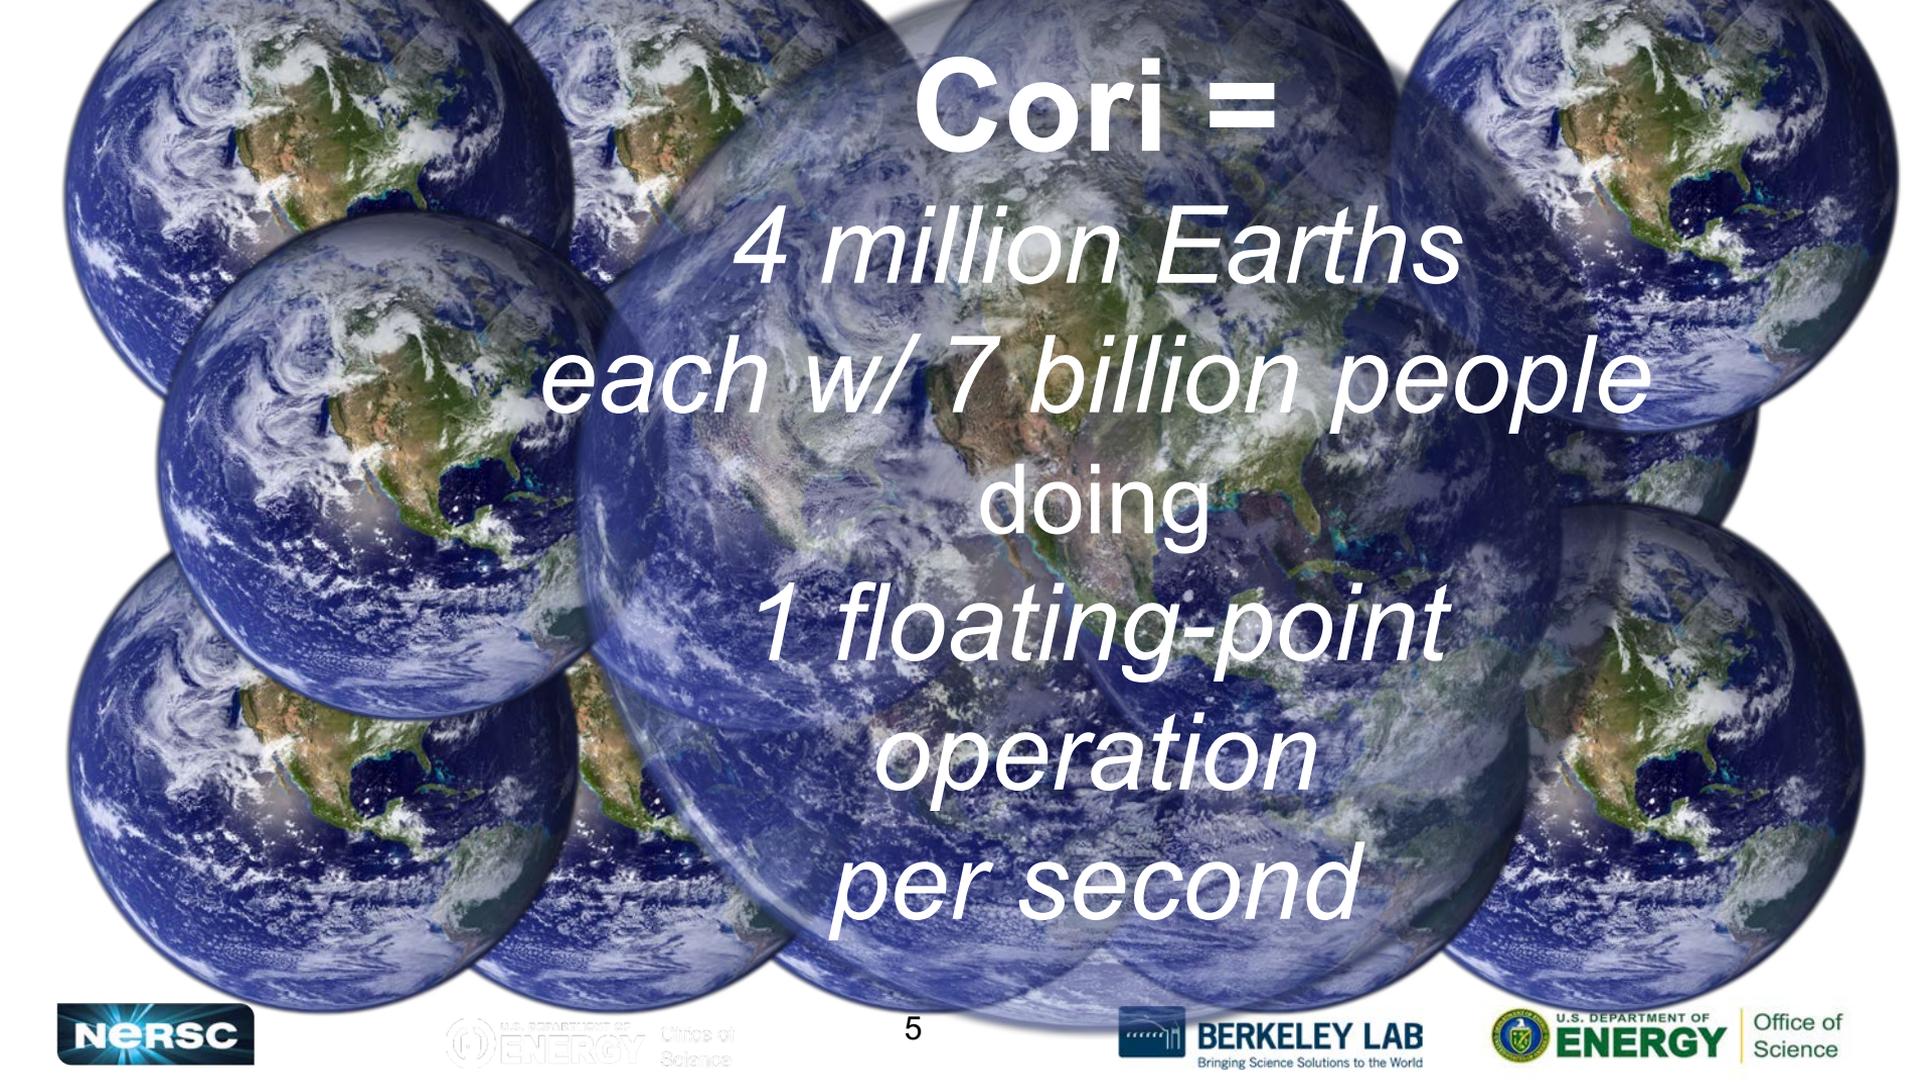


... not so different from a super high-end desktop computer.

Or rather, a lot of super high-end desktop computers.

Cori (left) has ~11,000 nodes (~ high-end desktop computers)

700,000 compute cores that can perform $\sim 3 \times 10^{16}$ calculations/second



Cori =
4 million Earths
each w/ 7 billion people
doing
1 floating-point
operation
per second

But Wait, There's More!

The nodes are all connected to each other with a high-speed, low-latency network.

This is what allows the nodes to “talk” to each other and work together to solve problems you could never solve on your laptop or even 150,000 laptops.

Typical point-to-point bandwidth

- Supercomputer: 10 GBytes/sec
 - Your home: 0.02* GBytes/sec
- 5,000 X

Latency

- Supercomputer: 1 μ s
 - Your home computer: 20,000* μ s
- 20,000 X

* If you're really lucky



Cloud systems have slower networks



...and Even More!

PBs of fast storage for files and data

- Cori: 30 PB
- Your laptop: 0.0005 PB
- Your iPhone: 0.00005 PB

Write data to permanent storage

- Cori: 700 GB/sec
- My iMac: 0.01 GB/sec



Cloud systems have slower I/O and less permanent storage



High-Performance Computing

High-Performance Computing...

- implies parallel computing
- In parallel computing, scientists divide a big task into smaller ones
- “Divide and conquer”

For example, to simulate the behavior of Earth’s atmosphere, you can divide it into zones and let each processor calculate what happens in each.

From time to time each processor has to send the results of its calculation to its neighbors.



Distributed-Memory Systems

This maps well to HPC “distributed memory” systems

- Many nodes, each with its own local memory and distinct memory space
- A node typically has multiple processors, each with multiple compute cores (Cori has 32 or 68 cores per node)
- Nodes communicate over a specialized high-speed, low-latency network
- SPMD (Single Program Multiple Data) is the most common model
 - Multiple copies of a single program (tasks) execute on different processors, but compute with different data
 - Explicit programming methods (MPI) are used to move data among different tasks



What is NERSC?

National Energy Research Scientific Computing Center

- NERSC is a national supercomputer center funded by the U.S. Department of Energy Office of Science (SC)
 - Supports SC research mission
 - Part of Berkeley Lab
- If you are a researcher with funding from SC, you can use NERSC
 - Other researchers can apply if research is in SC mission
- NERSC supports 8,000 users, 800 projects
 - From all 50 states + international; 65% from universities
 - Hundreds of users log on each day



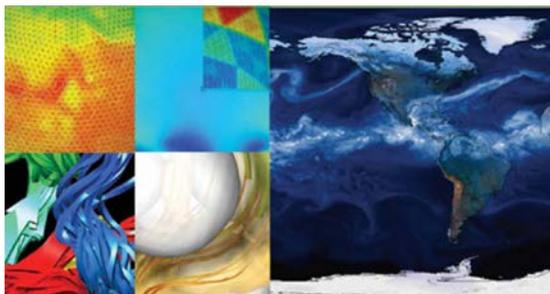
NERSC is the Production HPC & Data Facility for DOE Office of Science Research



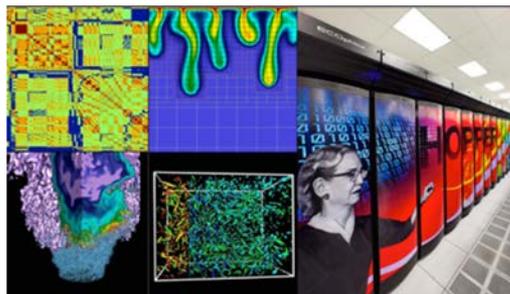
U.S. DEPARTMENT OF
ENERGY

Office of
Science

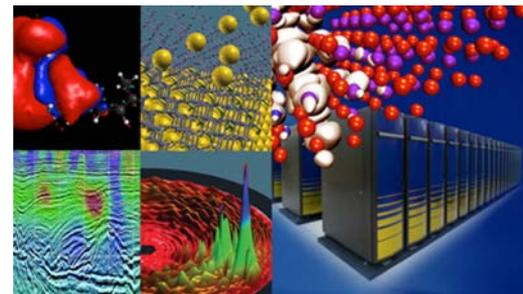
Largest funder of physical science
research in U.S.



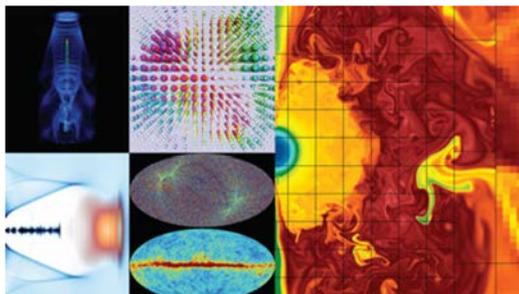
Bio Energy, Environment



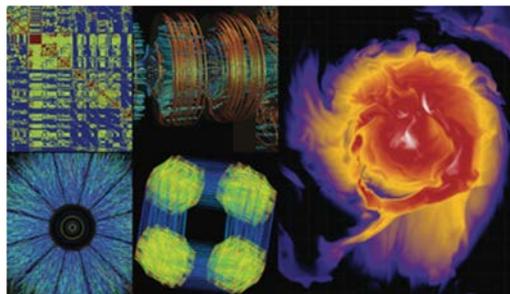
Computing



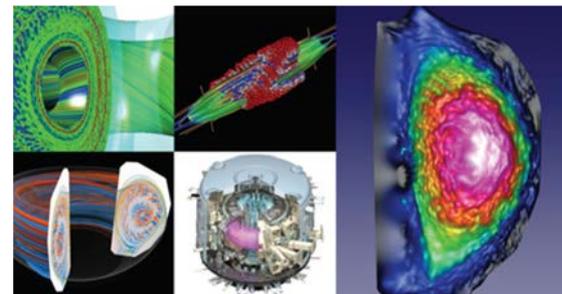
Materials, Chemistry, Geophysics



Particle Physics, Astrophysics



Nuclear Physics



Fusion Energy, Plasma Physics



NERSC: Science First!

***NERSC's mission is to accelerate
scientific discovery
at the DOE Office of Science
through high-performance computing
and data analysis.***



2018 Science Output

>2500 refereed publications

- Nature (14),
Nature Communications (31),
Other Nature journals (37)
- Science (11),
Science Advances (9)
- Proceedings of the National
Academy of Sciences (31)
- Physical Review Letters (67),
Physical Review B (85)



NERSC Nobel Prize Winners



2013
Chemistry

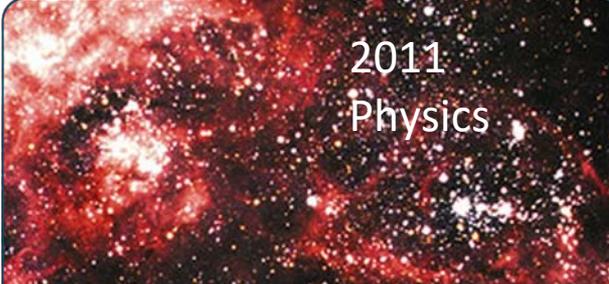


$R = 1.8 \text{ \AA}$

Martin Karplus



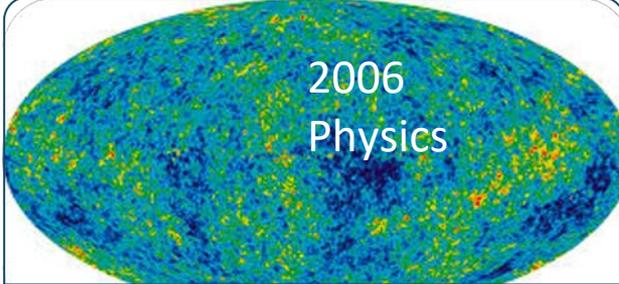
2011
Physics



Saul Perlmutter



2006
Physics



George Smoot



2007
Peace



Warren
Washington



2015 Nobel Prize in Physics



Scientific Achievement

The discovery that neutrinos have mass & oscillate between different types

Significance and Impact

The discrepancy between predicted & observed solar neutrinos was a mystery for decades. This discovery overturned the Standard Model interpretation of neutrinos as massless particles and resolved the “solar neutrino problem”

Research Details

The Sudbury Neutrino Observatory (SNO) detected all three types (flavors) of neutrinos & showed that when all three were considered, the total flux was in line with predictions. This, together with results from the Super Kamiokande experiment, was proof that neutrinos were oscillating between flavors & therefore had mass.

Nobel Recipients: Arthur B. McDonald, Queen's University (SNO)
Takaaki Kajita, Tokyo University (Super Kamiokande)



A SNO construction photo shows the spherical vessel that would later be filled with water.

Calculations performed on PDSF & data stored on HPSS played a significant role in the SNO analysis. The SNO team presented an autographed copy of the seminal *Physical Review Letters* article to NERSC staff.

Q. R. Ahmad et al. (SNO Collaboration). *Phys. Rev. Lett.* 87, 071301 (2001)



NERSC Science Stories

COVID-19 Text Mining & Knowledge Portal

Scientific Achievement

Working closely with NERSC, a team of materials scientists at Berkeley Lab – scientists who normally spend their time researching things like high-performance materials for thermoelectrics or battery cathodes – has created a text-mining tool to help the global scientific community synthesize the mountain of scientific literature on COVID-19 being generated every day.

Significance and Impact

COVID Scholar was developed in response to a March 16 call to action from the White House Office of Science and Technology Policy that asked artificial intelligence experts to develop new data and text-mining techniques to help find answers to key questions about COVID-19. COVID Scholar is also being used in MIT's Rapid Review Journal, the KG-COVID knowledge graph project, and in Pacific Northwest National Lab's literature analysis tool.

Research Details

COVID Scholar uses natural-language processing techniques to quickly scan and search tens of thousands of research papers and help draw insights and connections that may otherwise not be apparent. Assisted by quick access to NERSC's Cori supercomputer, the center's "Spin" internal cloud for edge services, and expert staff assistance from NERSC staff, a prototype was created in just 7 days. State-of-the-art natural-language processing models run daily on Cori, and the search portal is running on Spin.



Berkeley Lab researchers (clockwise from top left) Kristin Persson, John Dagdelen, Gerbrand Ceder, and Amalie Trewartha led development of COVID Scholar, a text-mining tool for COVID-19-related scientific literature. (Credit: Berkeley Lab)



covid scholar.org



NERSC Project PI: Kristin Persson, Berkeley Lab
NERSC Director's Reserve, Research Funded by Berkeley Lab LDRD



19



Regional Simulations of Building Response to Major Earthquakes

Scientific Achievement

Using the EQSIM fault-to-structure computational framework running on NERSC's Cori supercomputer, a research team from the University of Nevada-Reno, Berkeley Lab, UC Berkeley, and Livermore Lab are providing new insights into how buildings respond to major earthquakes over large regions.

Significance and Impact

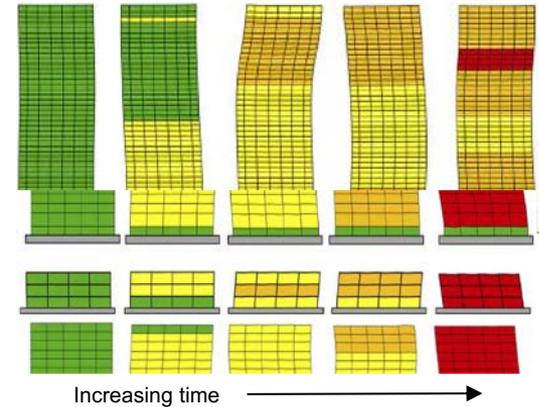
The risk to buildings and other structures over a geographical area due to major earthquakes is poorly understood. The effect on buildings depends on many complex factors, including where and how a fault ruptures, how the waves move through the Earth, and how the soil underneath the building interacts with the structure. This team's simulation framework, leveraging the computational capability of Cori, allows researchers to examine many scenarios of interest and inform civic planners so they can save lives and protect infrastructure.

Research Details

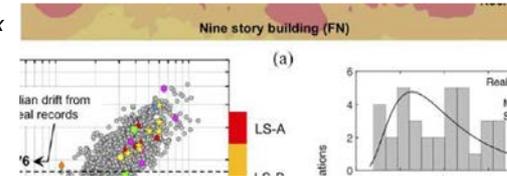
- The team, led by David McCallen from UNR and Berkeley Lab, simulated a magnitude 7.0 earthquake and looked in detail at the response of buildings of various sizes throughout a region.
- Using an allocation of computer time from the NERSC Director's Discretionary Reserve, the team ran the EQSIM framework, which is being developed with funding from the DOE Exascale Computing Project, on Cori using 2,048 nodes (139,000 processor cores). The project used 65 million hours of compute time in 2020.



New simulations provide insight into how earthquakes stress and distort buildings depending on their location, size, and shape. Red indicates the most distortion and green the least. Steel frame buildings were considered in this study.



Maps reveal the complex spatial distribution of distortions suffered by buildings following an earthquake.



McCallen, D.; Petrone, F.; Miah, M.; Pitarka, A.; Rodgers, A.; Abrahamson, N., "EQSIM—A multidisciplinary framework for fault-to-structure earthquake simulations on exascale computers, part II: Regional simulations of building response"; Earthquake Spectra 2020, 10.1177/8755293020970980

NERSC Project PI: David McCallen, University of Nevada-Reno & Berkeley Lab
NERSC Director's Reserve, Research Funded by the Exascale Computing Project

Urban Landscapes Enhance Destructive Storms

Scientific Achievement

Using NERSC supercomputers to model destructive thunderstorms, researchers from Pacific Northwest National Laboratory found that urban landscapes and human-made aerosols – particles suspended in the atmosphere – can make wind gusts stronger, rain heavier, and hail larger and even steer storms toward cities.

Significance and Impact

Urbanization has been a significant change in the earth's environment since industrialization and is expected to further expand during the coming decades. Many modeling and observational studies have shown that landscapes can impact weather and climate, and it is important to understand how this will affect the lives of millions of Americans.

Research Details

Researchers modeled two large historical storms: a supercell near Kansas City that produced hail, strong wind, and a tornado; and a sea-breeze-induced thunderstorm near Houston. Using a version of the Weather Research and Forecasting model that includes routines to model aerosols and detailed chemistry and physics, the team found that storms were stronger when urban landscapes and human-produced aerosols – for example, from auto exhausts or farming – were included in the simulations. In addition to harnessing the computational power of NERSC's Cori supercomputer, the team also relied on a special allocation of 50TB of data storage on Cori's high-speed file system.



*Fan, J.; Zhang, Y.; Li, Z.; Hu, J.; Rosenfeld, D.,
"Urbanization-induced land and aerosol impacts on sea-
breeze circulation and convective precipitation";
Atmospheric Chemistry and Physics, 20:14163-14182;
2020 , 10.5194/acp-20-14163-2020*

NERSC Project PI: Jiwen Fan, Pacific Northwest National Laboratory
DOE Mission Science, Research Funded by Office of Biological & Environmental Research

NERSC Usage Facts & Figures

In 2020, scientists used

8,044,000,000

>1,000,000 single-CPU-years

NERSC-hours
and currently store

220,000,000

4 million iPhones

GB of data
at NERSC



Homo erectus
~1,000,000 years ago

Data Storage

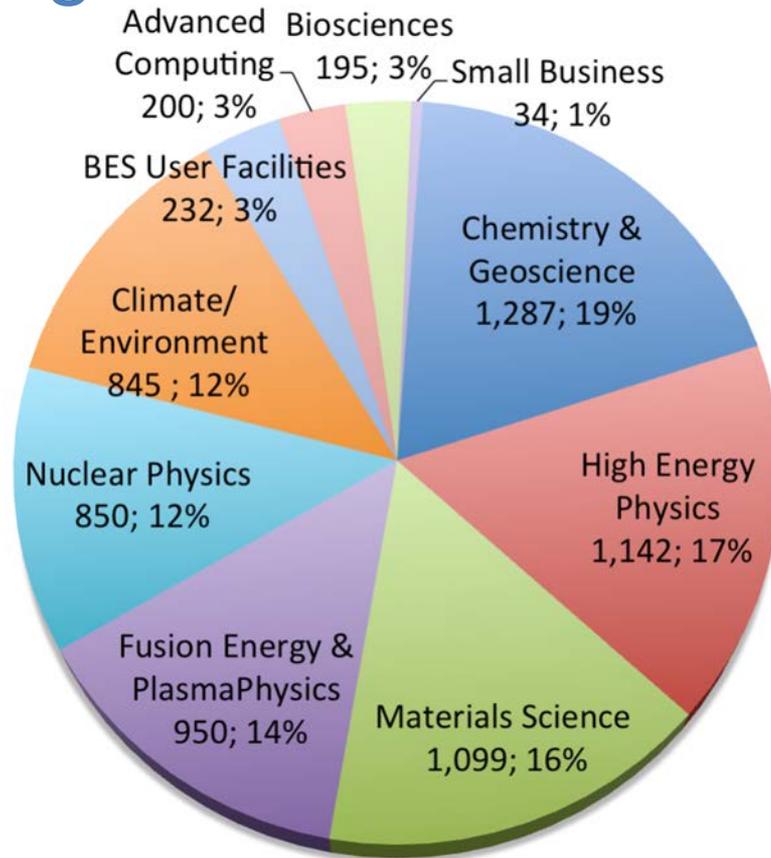


HPSS
200 Petabytes



Community
64 Petabytes

Compute Usage

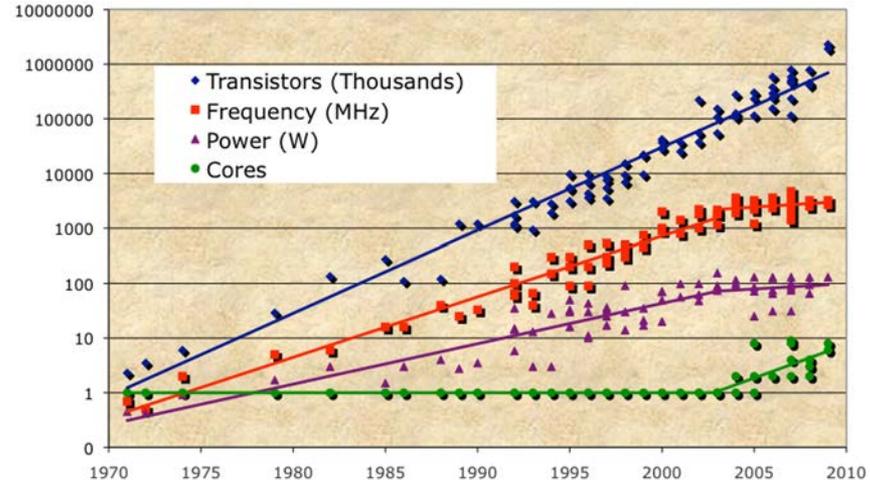
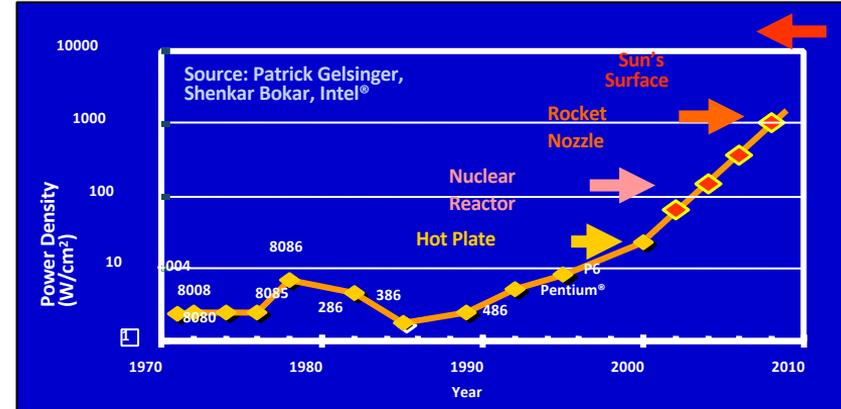




Challenges in HPC

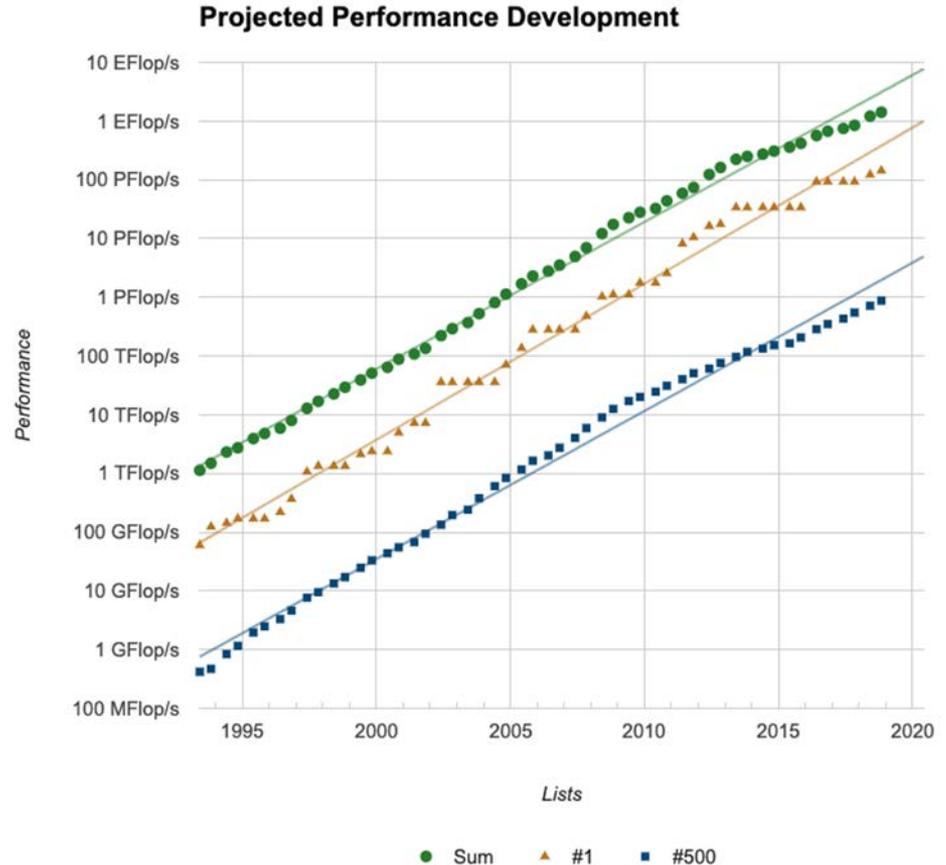
Power: the Biggest Architectural Challenge

- If we just kept making computer chips faster and more dense, they'd melt and we couldn't afford or deliver the power.
- Now compute cores are getting slower and simpler, but we're getting lots more on a chip.
 - GPUs and Intel Xeon Phi have 60+ "light-weight cores"



Revolution in Energy Efficiency Needed

- Energy efficiency is increasing, but today's top computers use 10s of Megawatts of power at ~\$1M/MW.
- With newest tech, get 4-5x performance with doubling of power



Energy-Efficient Tech for Facilities Needed

- PUE: Measure of energy efficiency in operation of data center facilities
 - Perfect operation, no overhead: PUE = 1
 - NERSC operations over past year: PUE = 1.08
 - Data-center industry average: PUE = 1.67
- Active research on facility energy efficiency at NERSC
 - Data centers can operate outside human temperature-comfort zone
 - Automation of intelligent fan/water settings
 - More ways to reduce PUE?

Programming for Advanced Architectures

- Advanced architectures (e.g., CPU+GPU offload) present challenges in programming and performance
 - Science expert must become expert on computer architectures and programming models
 - Performance on one architecture doesn't always translate to performance on another
 - Many codes not ported and many unsuitable for this type of architecture; complete overhaul required

Beyond Moore's Law

- Moore's law: doubling of performance every 18-24 months
 - There is an end, and it is soon
 - What do we do next?
- Pathfinding new architectures
 - Accelerators? FPGAs? Quantum?
 - How to program for these?

Data: Getting Bigger All the Time!

- Simulations producing more data
- Scientific instruments producing more data
 - SKA when comes fully online will produce more data in a day than currently exists!
- How do we
 - process this data?
 - manage it?
 - store it?
 - transfer it?
 - access it?
- Efficient workflows for data analysis and management needed

Your Challenges

- Figure out how to program the next generation of machines
- Find a way to make sense of all the data
- Build faster, more capable hardware that uses less energy
- Design energy-efficient facilities that reduce PUE
- Create effective data and job management workflows
- Bring new fields of science into HPC
- Tell the world about what you're doing!

Questions?

