NERSC: Scientific Discovery through Computation with High-Performance Computing

Presentation for CSSS Program

Rebecca Hartman-Baker, PhD
User Engagement Group Lead
Charles Lively III, PhD
Science Engagement Engineer

June 8, 2023

Introductions - Rebecca

- User Engagement Group (UEG) Lead, NERSC
- World-famous violinist*
- Enthusiastic picker of fruits
- Mom to Vinny (16) & Elena (8)
- Kentucky native, honorary Aussie
- Algorithm enthusiast
- PhD, Computer Science, University of Illinois at Urbana-Champaign

Rebecca Hartman-Baker

*Slight exaggeration; I have played publicly in 3 countries on 2 continents

Introductions - Charles

- Science Engagement Engineer in UEG @NERSC
- Husband, son, brother, uncle, godfather
- Fur Daddy to Bella and Monte
- PhD, Computer Engineering, Texas A&M University
- Co-founded 2 Start-ups and served as Technical Advisor/Mentor for over 20 start-ups
- Theoretical Physicist in another life
- Avid Peloton rider

Charles Lively

The Plot

- What is NERSC?
- Science and NERSC’s mission
- What is High-Performance Computing?
- What is a Supercomputer?
- The User Engagement Group (UEG)
- Future Challenges in HPC
- Career Paths at NERSC/LBL

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 and you need resources at the scale and complexity NERSC provides, then you are eligible to apply to use NERSC
  - Other researchers can apply if their research is relevant the SC mission
- NERSC supports 9,000 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.

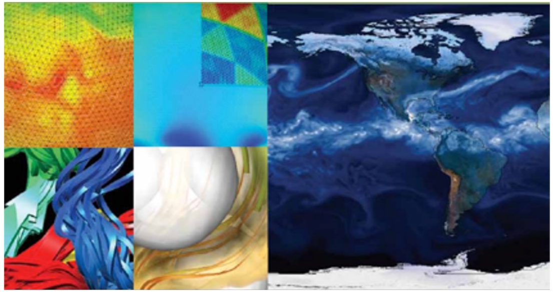
Bio Energy, Environment

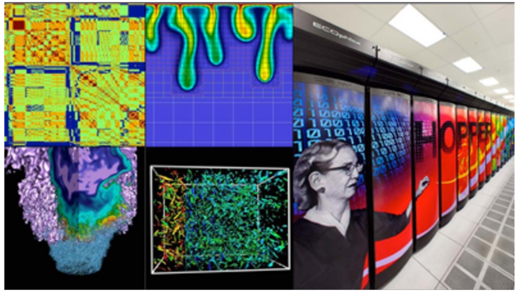

Computing

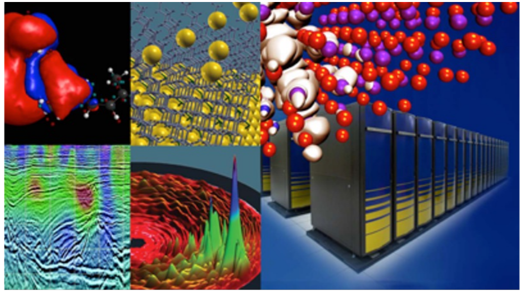
Materials, Chemistry, Geophysics

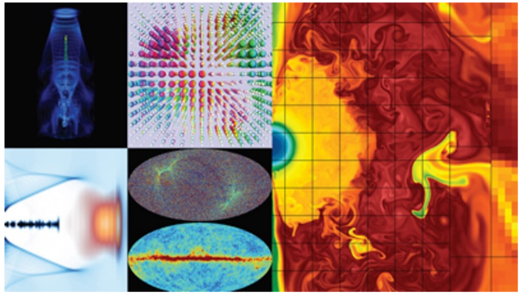
Particle Physics, Astrophysics

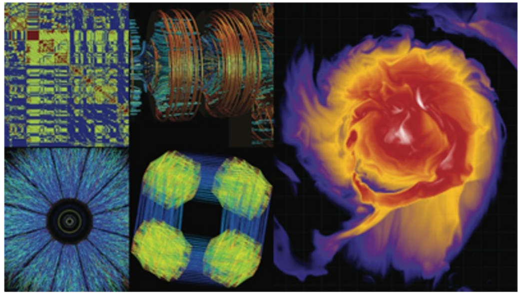
Nuclear Physics

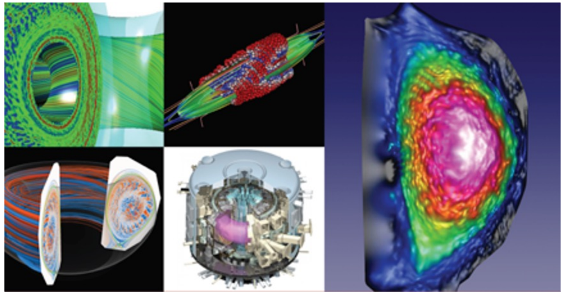
Fusion Energy, Plasma Physics

NERSC: Science First!

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

What is Science?

What is Science?

- Science is a systematic and organized approach to acquiring knowledge and understanding the natural world.
- It involves formulating questions, developing hypotheses, conducting experiments or observations, and analyzing data to draw conclusions.
- Science relies on evidence-based reasoning and follows established methods and principles.
- It aims to explain phenomena, predict outcomes, and improve our understanding of the universe.

In Your Mind?

- What does science mean to you?
- What is important about science?

We use High-Performance Computing

… to solve scientific computational problems that are either too large for standard computers or would take them too long.

Cheap & Efficient Solar

Designing Better Batteries

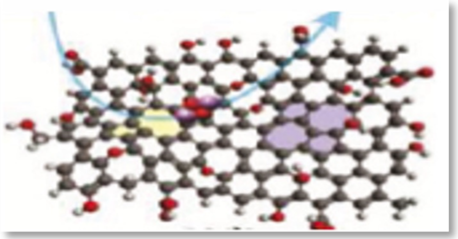

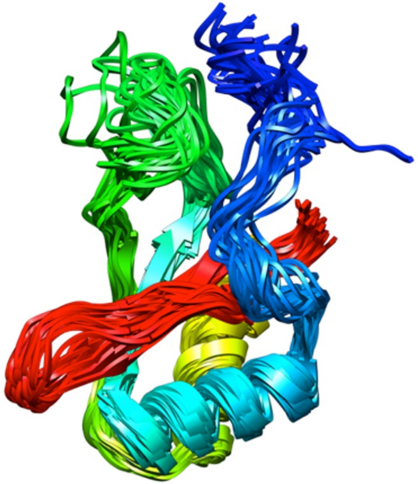

Understanding How Proteins Work

Understanding The Universe

Biofuels

Extreme Climate Events

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 (Perlmutter has 64 CPU cores and 256 GPU cores per node) or 128 cores per node for CPU-Only)
- 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

History of HPC

1970s
The Cray-1 supercomputer
Used Vector processing technique revolutionizes supercomputing, enabling the processing of multiple data elements simultaneously.

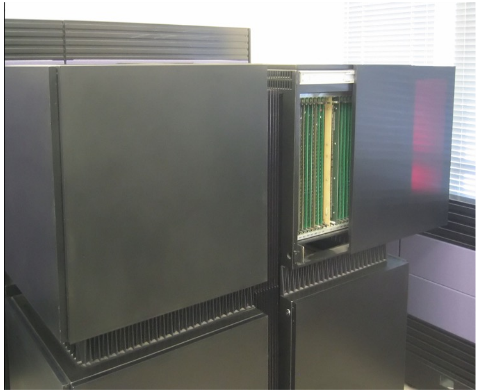

1980s
The Connection Machine, a visually striking representation of massively parallel processing.
Emergence of parallel processing and the Connection Machine drives high-performance computing to new heights.

1990s (Cray SuperServer CS6400)
A cluster of interconnected computers, symbolizing the rise of cluster computing.
Cluster computing and distributed systems gain popularity, enabling collaborative and accessible high-performance computing.

History of HPC

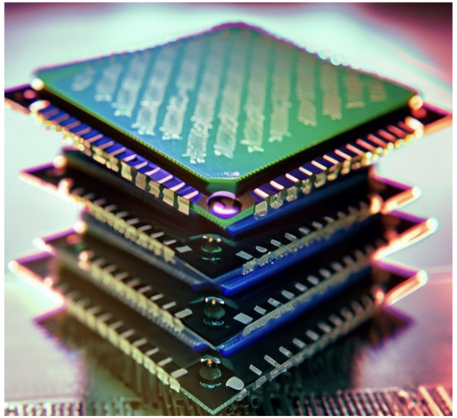

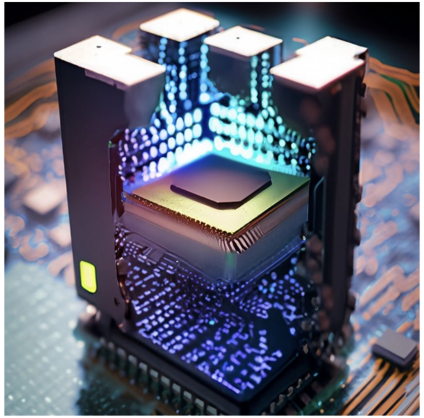

2000s
Multicore Processors
Multi-core processors become mainstream, unleashing significant computational power within a single chip.

2010s
A CPU-GPU hybrid system, representing the rise of heterogeneous computing.
Heterogeneous computing and accelerators like GPUs reshape supercomputing, delivering specialized processing power.

2020s (Present):
Towards exascale supercomputer
Exascale computing and quantum computing research drive the exploration of new frontiers in computational capabilities.

What Is a Supercomputer?

A Supercomputer Is...

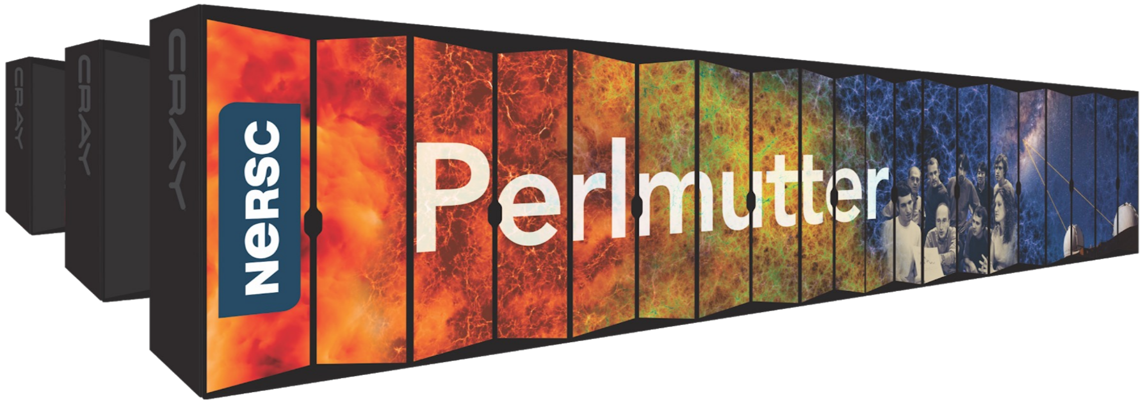

vs.

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

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

Perlmutter (left) has ~4800 nodes (~ high-end desktop computers)

Over 760,000 compute cores

NERSC-9 is named after Saul Perlmutter

- Shared 2011 Nobel Prize in Physics for discovery of the accelerating expansion of the universe.
- Supernova Cosmology Project, led by Perlmutter, was a pioneer in using NERSC supercomputers combine large scale simulations with experimental data analysis
- Login “saul.nersc.gov”

Perlmutter =

20 million Earth-like Planets each w/ 7 billion people

doing

1 floating-point operation per second

NERSC Systems Roadmap

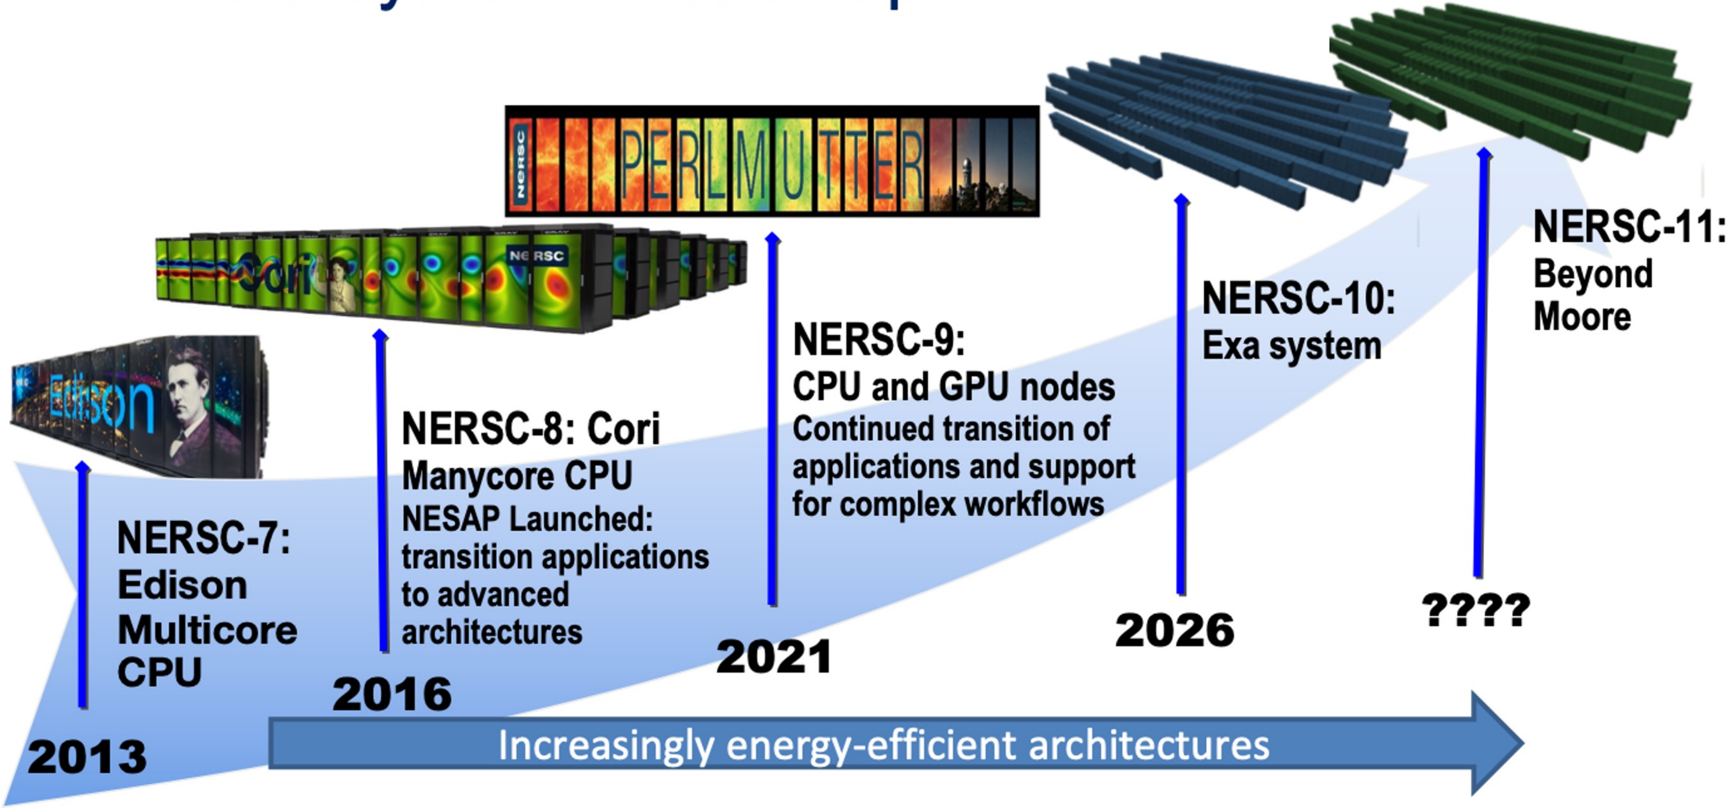

But Wait, There's More!

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 μs
- Your home computer: 20,000* μs

20,000 X

* If you're really lucky

Cloud systems have slower networks

...and Even More!

PBs of fast storage for files and data

- Perlmutter: 35 PB
- Your laptop: 0.0005 PB
- Your iPhone: 0.00005 PB

Write data to permanent storage

- Perlmutter: 5 TB/sec
- My iMac: 0.01 GB/sec

Cloud systems have slower I/O and less permanent storage

NERSC's Users Produce Groundbreaking Science

Materials Science
Revealing Reclusive Mechanisms for Solar Cells

NERSC PI: C. Van de Walle, UC Santa Barbara, *ACS Energy Letters*

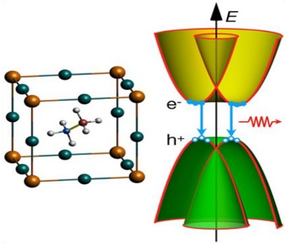

High Energy Physics
Shedding Light on Luminous Blue Variables

NERSC PI: Yan-Fei Jiang, UC Santa Barbara. *Nature*

2,500 Refereed Publications per Year

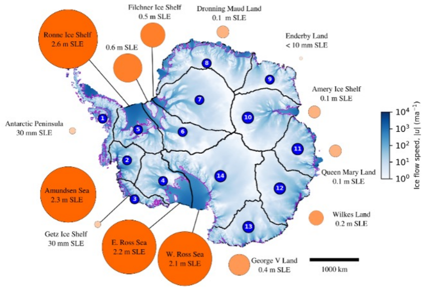

Earth Sciences
Simulations Probe Antarctic Ice Vulnerability

NERSC PIs: D. Martin, Berkeley Lab; E. Ng, Berkeley Lab; S. Price, LANL. *Geophysical Research Letters*

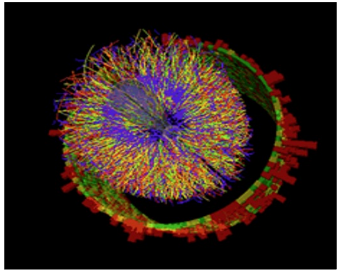

Nuclear Physics
Enabling Science Discovery for STAR

NERSC PI: J. Porter, Berkeley Lab. *J. Phys.: Conference Series*

Advanced Computing
Scalable Machine Learning in HPC

NERSC PI: L. Oliker, Berkeley Lab, *21st International Conference on AI and Statistics*

Plasma Physics
Plasma Propulsion Systems for Satellites

NERSC PI: I. Kaganovich, Princeton Plasma Physics Lab, *Physics of Plasmas*

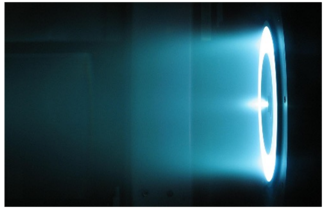

Nobel-Prize Winning Users

for the development of multiscale models for complex chemical systems

2013 Chemistry

Martin Karplus

for the discovery of the accelerating expansion of the Universe through observations of distant supernovae

2011 Physics

Saul Perlmutter

for the discovery of the blackbody form and anisotropy of the cosmic microwave background radiation

2006 Physics

George Smoot

for their efforts to build up and disseminate greater knowledge about man-made climate change

2007 Peace

Warren Washington

for developing cryo-electron microscopy for the high-resolution structure determination of biomolecules in solution

2017 Chemistry

Joachim Frank

for the discovery of neutrino oscillations, which shows that neutrinos have mass

2015 Physics

SNO Collaboration

HPC is Already Amongst You too!

- Large Language Model Training
  - ChatGPT
  - Generative AI
- Self-Driving Technologies
  - Sensor Fusion
  - Trajectory Planning
  - Supervised and Unsupervised Learning
- Video Game Technologies
  - Graphics Rendering
  - Game Testing and Quality Assurance
  - Procedural and Contextual Generation

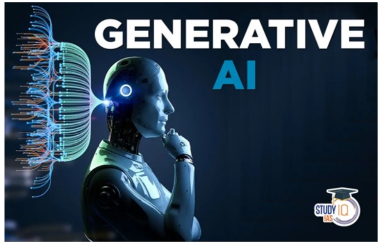

Supporting NERSC Researchers and Users: The User Engagement Group (UEG)

Our People


Justin Cook


Kevin Gott


Lipi Gupta


Rebecca Hartman-Baker


Helen He


Kadidia Konate


Charles Lively


Erik Palmer


Kelly Rowland


Shahzeb Siddiqui

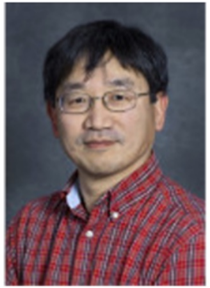
Woo-Sun Yang

Alumni:

Tiffany Connors
Zhengji Zhao
Steve Leak

Bringing Science Solutions to the World

UEG Mission

The User Engagement Group engages with the NERSC user community to increase user productivity via **advocacy, support, training,** and the **provisioning of usable computing environments.**

Bringing Science Solutions to the World

UEG Mission: Advocacy

- Determine user needs via
  - Directly working with users
  - User surveys
  - Discovering their habits, behaviors, etc. through analysis of user data
- Advocate for those needs in future systems, training offerings, etc.
- Build NERSC community through initiatives such as the NERSC User Group (NUG), NUG Executive Committee (NUGEX), NERSC User Community of Practice, etc.

UEG Mission: Support



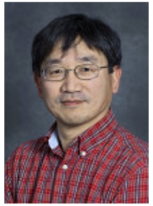

- Support NERSC users via
  - Tickets in ServiceNow
  - User appointments
  - Office Hours on special topics
  - Documentation
  - Communications (e.g., weekly email)
  - Automation of user processes
  - Special interest groups
  - and more!

UEG Mission: Training

- Oversee the NERSC user training program
  - Set direction for user training, taking user needs into account
- Coordinate across groups to provide NERSC user training
- Each year, we provide 20+ user training opportunities

The Future Of High-Performance Computing

Perlmutter: Optimized for Science

- HPE Cray System with 3-4x capability of Cori
- GPU-accelerated and CPU-only nodes
- HPE Cray Slingshot high-performance network
- All-Flash filesystem
- Application readiness program (NESAP)

Phase I: Arrived in 2021

- 1,536 GPU-accelerated nodes
- 1 AMD “Milan” CPU + 4 NVIDIA A100 GPUs per node
- 256 GB CPU memory and 40 GB GPU high BW memory
- 35 PB FLASH scratch file system
- User access and system management nodes

Phase II Addition: Arrived in 2022

- 3,072 CPU only nodes
- 2 AMD “Milan” CPUs per node
- 512 GB memory per node
- Upgraded high speed network
- CPU partition exceeded performance of entire Cori system

Future NERSC Systems

NERSC-7:
Edison
Multicore CPU

2013

NERSC-8: Cori
Manycore CPU
NESAP Launched: transition applications to advanced architectures

2016

NERSC-9:
CPU and GPU nodes
Continued transition of applications and support for complex workflows

2021

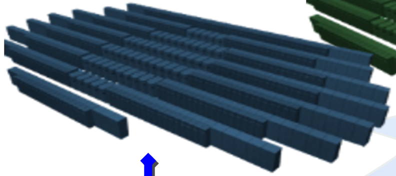

NERSC-10:
Exa system

2026

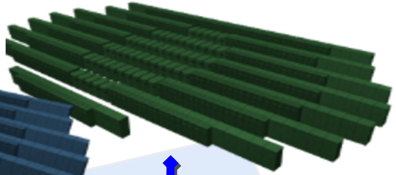

NERSC-11:
Beyond Moore

~2030

Increasingly energy-efficient architectures

Future NERSC Systems

- Not completely clear what NERSC-10 (~2026) will look like
    - Likely heterogeneous, Exaflop-level
    - Could include ASICs or other novel architectures
- NERSC-11(~2030) is even less predictable

NERSC-10: Heterogeneous within Nodes?

Edge Services

Interconnec
t

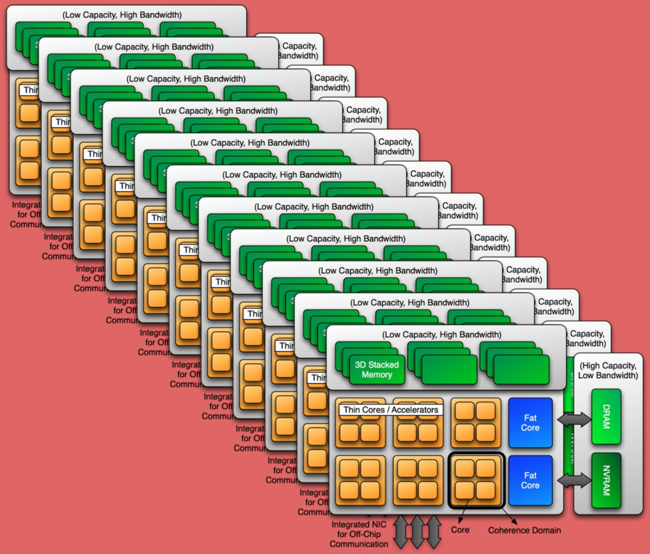

Heterogeneous nodes

Storage

...Or Heterogeneous Node Types?

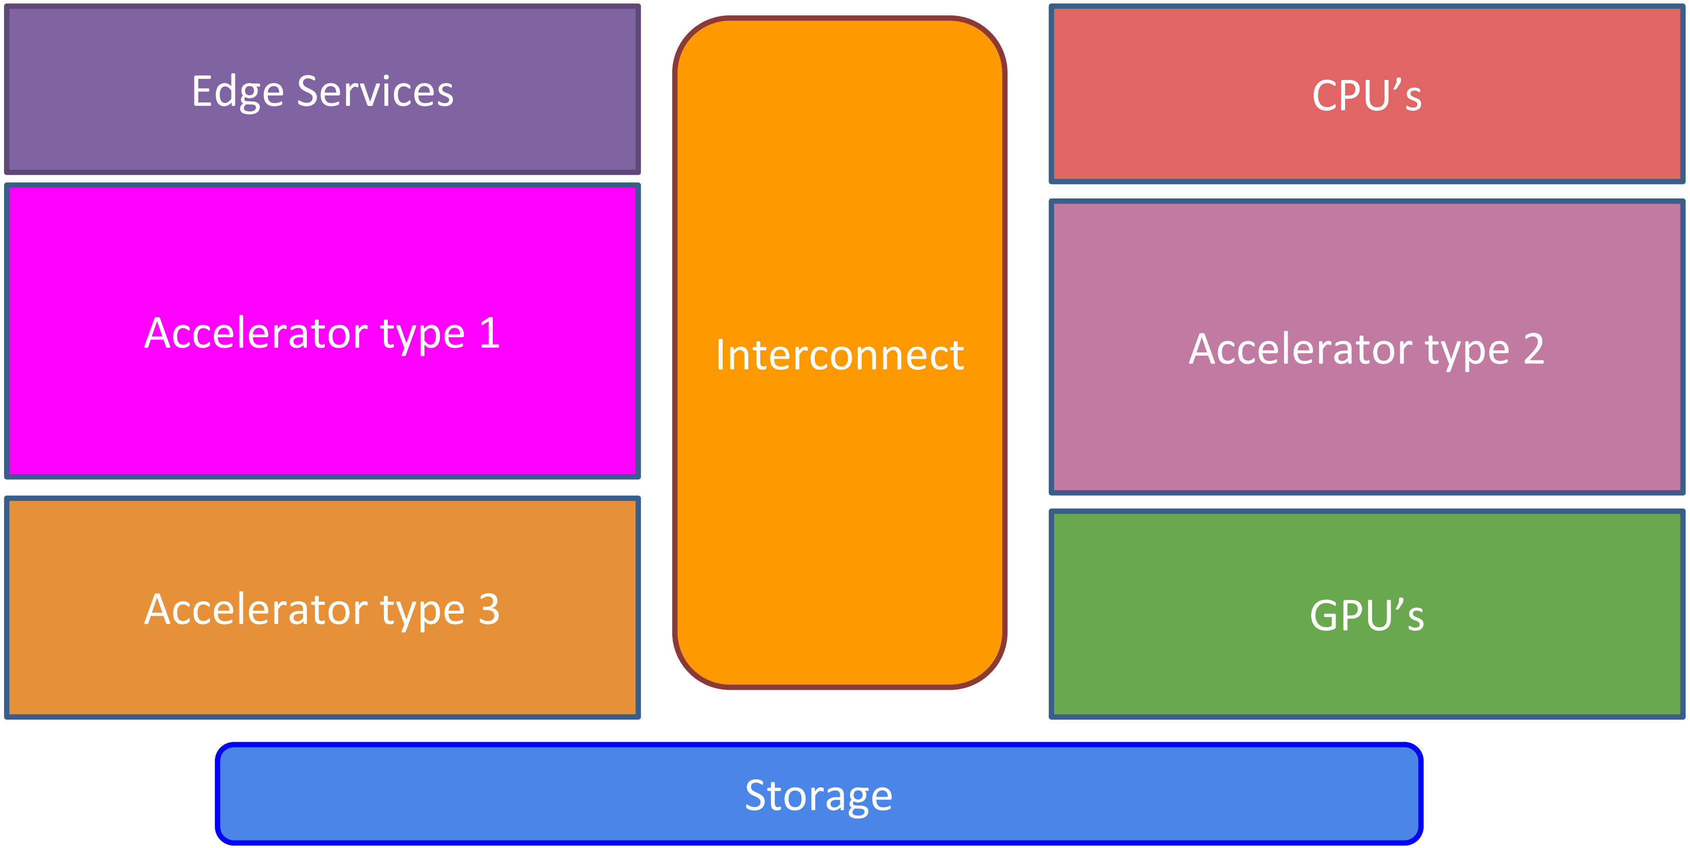

Quantum Computing Could Apply to >50% of NERSC Workload

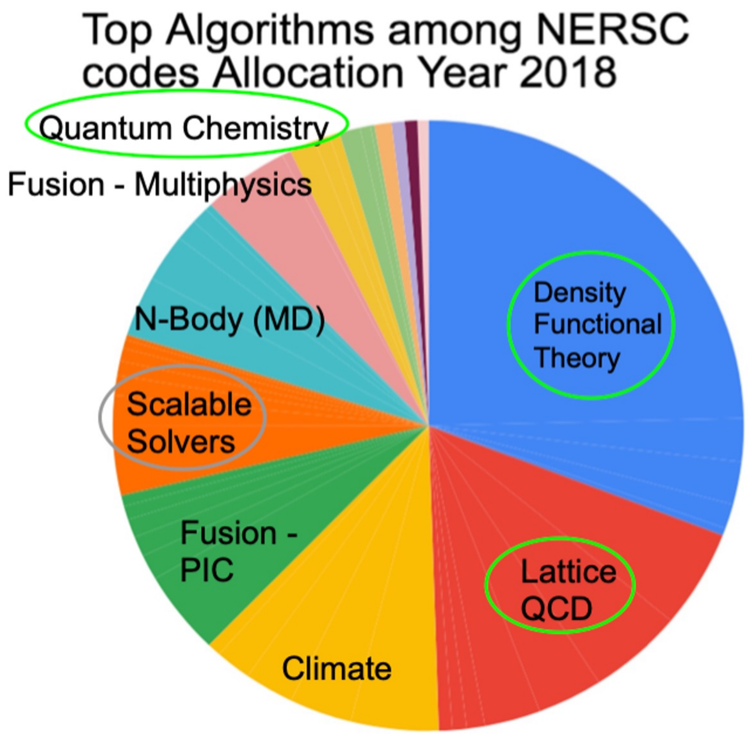

	Logical Qubits	Note
Quantum Chemistry	∝ active orbitals 10^1-10^2	Possible NISQ "killer app" - NAS
Density Functional Theory	∝ bands 10^3 - 10^5	Algorithm published. Like ab initio, but larger systems.
Lattice QCD	∝ lattice sites. 10^6-10^9	Algorithm published.
Machine Learning	???	Frameworks published. TensorFlow Quantum, TorchQuantum
Scalable Solvers	???	Kernels published. (Ax=b, FFT)

NERSC Quantum Computing Roadmap

2022	2022-2024	2024-2028	2028-203?
• **Ramp up** engagement with QIS community • Director's Discretionary Reserve Call for quantum information science (QIS) on **Perlmutter**	• **Engage** with quantum hardware companies and gov labs • Enable **user access** to quantum hardware • **Development** of hybrid algorithms • Identify **opportunities** for quantum accelerated HPC codes • **Benchmarking** quantum hardware	• **Integration** of near-term (NISQ) quantum hardware becoming standard • **Users** requesting both classical and quantum resources	• **High-performing quantum hardware** becoming available • **Full integration** with traditional HPC • Users **routinely** solve problems using quantum hardware !

Optimal integration of classical and quantum processors is an open area of research

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 have 100s of "light-weight cores"

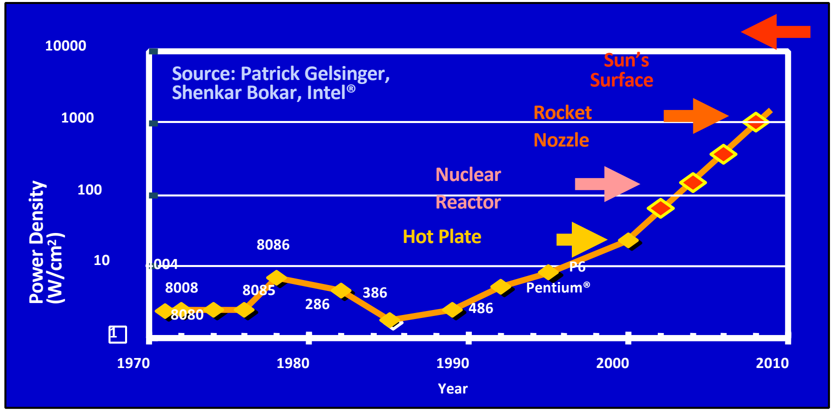

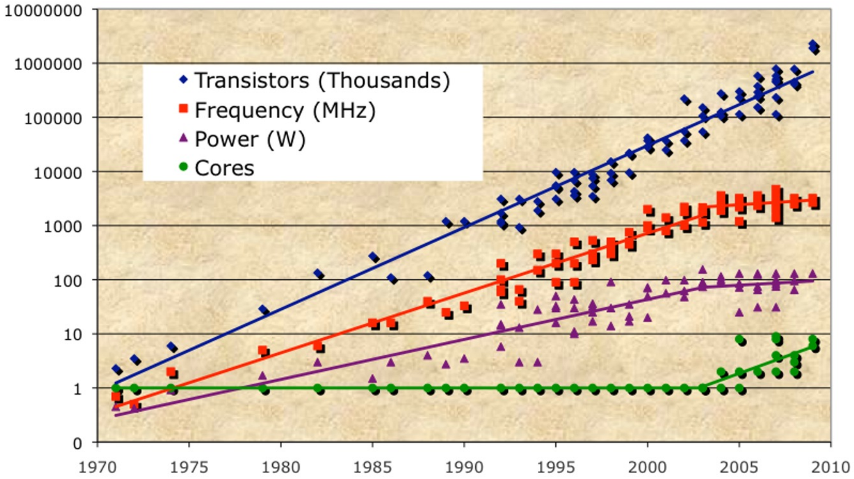

Programming for Advanced Architectures

- Advanced architectures (e.g., CPU+GPU offload) present challenges in programming and performance
  - Science expert must become computer architectures & programming models expert
  - 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
- Data: Getting Bigger All the Time!
  - Simulations producing more data
  - Scientific instruments producing more data
    - Square Kilometre Array, when comes fully online, will produce more data in a day than currently exists!
  - 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!

Career Paths in HPC

The Awesome Groups @ NERSC

TINA DECLERCK
Division Deputy for Operations

RICHARD GERBER
Senior Science Advisor

SUDIP DOSANJH
Division Director

HAI AH NAM
NERSC-10 Project Director

BECCI TOTZKE
NERSC-10 Project Manager

NICHOLAS WRIGHT
NERSC-10 Chief Architect

HIGH PERFORMANCE COMPUTING DEPARTMENT
RICHARD GERBER
Department Head

- **ADVANCED TECHNOLOGIES** — NICHOLAS WRIGHT, Group Leader
- **APPLICATION PERFORMANCE** — JACK DESLIPPE, Group Leader
- **BUSINESS OPERATIONS & SERVICES** — BECCI TOTZKE, Group Leader
- **PROGRAMMING ENVIRONMENTS & MODELS** — BRANDON COOK, Group Leader (Acting)
- **USER ENGAGEMENT** — REBECCA HARTMAN-BAKER, Group Leader

DATA DEPARTMENT
WAHID BHIMJI
Department Head (Acting)

- **DATA & ANALYTICS SERVICES** — WAHID BHIMJI, Group Leader
- **DATA SCIENCE ENGAGEMENT** — DEBORAH BARD, Group Leader
- **STORAGE SYSTEMS** — KRISTY KALLBACK-ROSE, Group Leader

SYSTEMS DEPARTMENT
TINA DECLERCK
Department Head

- **BUILDING INFRASTRUCTURE** — BEN MAXWELL, Group Leader
- **COMPUTATIONAL SYSTEMS** — DOUG JACOBSEN, Group Leader
- **INFRASTRUCTURE SERVICES** — CORY SNAVELY, Group Leader
- **OPERATIONS TECHNOLOGY** — ELIZABETH BAUTISTA, Group Leader
- **SECURITY & NETWORKING** — JAY SRINIVASAN, Group Leader (Acting)

HPC and You - Career Paths

- HPC Consultant
- HPC Research Scientist
- HPC Performance Engineer
- HPC Architect
- HPC Data Scientist
- HPC System Administrator
- HPC Application Developer
- HPC Cloud Architect
- HPC Educator/Trainer

Questions?