National Energy Research Scientific Computing Center

2023 ANNUAL REPORT

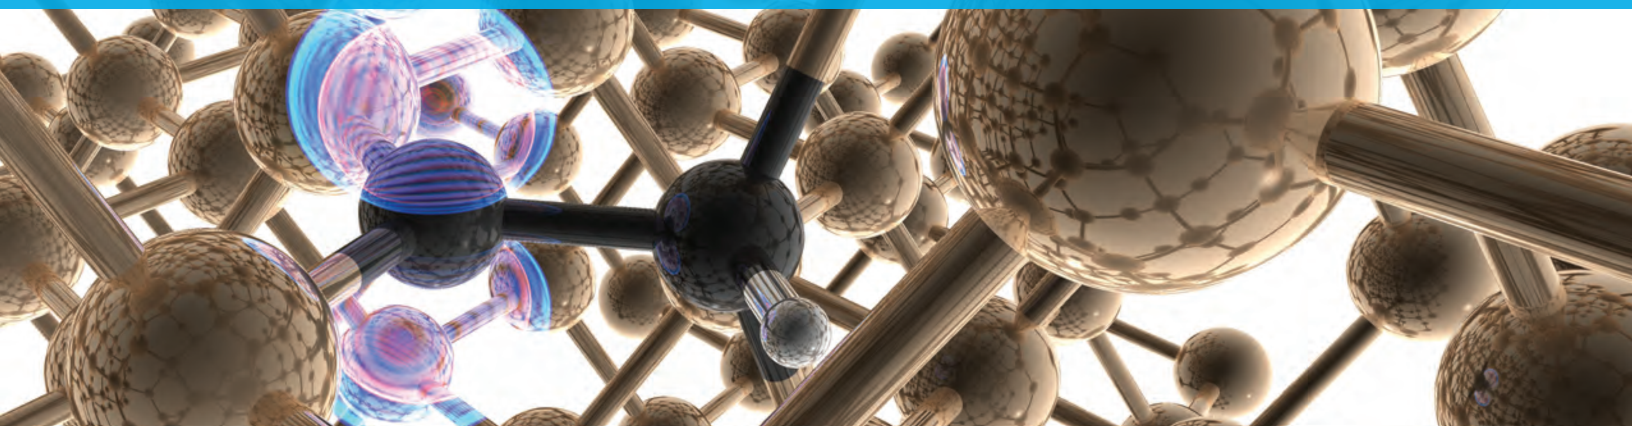

U.S. DEPARTMENT OF ENERGY | Office of Science

Table of Contents

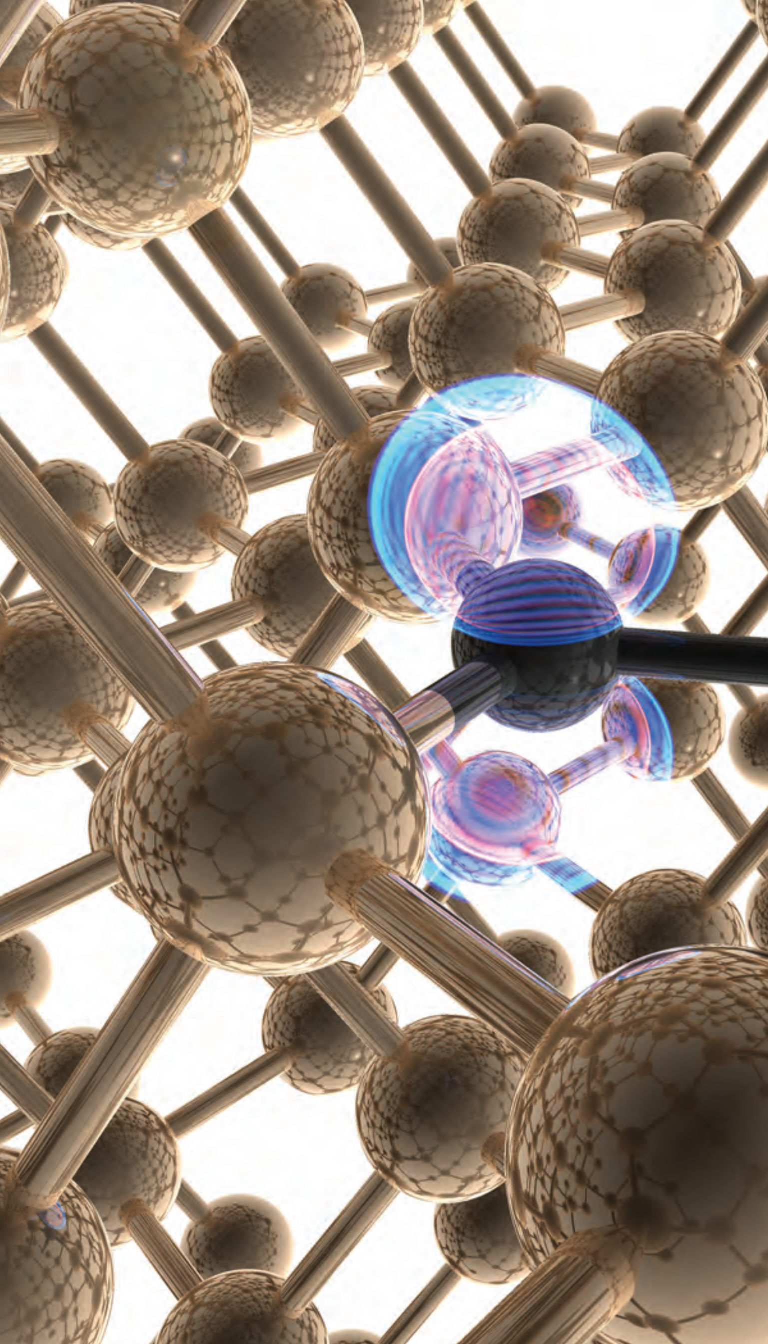

The National Energy Research Scientific Computing Center is the mission high-performance computing facility for the U.S. Department of Energy's Office of Science (SC). NERSC is managed by Lawrence Berkeley National Laboratory and funded by the U.S. Department of Energy's Office of Science Advanced Scientific Computing Research Office. NERSC's mission is to accelerate scientific discovery at the U.S. Department of Energy Office of Science through HPC and extreme data analysis.

Director’s Note

The year 2023 was an exciting one at NERSC—our flagship *Perlmutter* system came into full production mode, supporting thousands of world-changing research projects even as staff set their sights on new horizons in HPC for science. Emerging technologies came closer to becoming reality, offering glimpses of the future. And through intention and investment, the NERSC community grew both in numbers and in connection.

Here are some highlights from the year:

Preparing for Complex Workflows (Artificial Intelligence, Integrated Research Infrastructure, NESAP)

The future of HPC for science will go beyond the increasing demand for more computational cycles required by traditional simulation and modeling workloads. Continuing to fulfill the mission of DOE SC will require a paradigm shift in the way HPC systems are designed, configured, and operated. NERSC is a fundamental part of that shift and is leading the way through a variety of initiatives:

AI

We worked across areas of AI system deployment, deep engagement with science teams towards AI-enabled applications, and broad engagement and training of the wider scientific community. In 2023, NERSC offered eight workshops and training events to prepare NERSC users and the broader HPC community for the role of AI in future workflows.

IRI

Integrated workflows are the future of science research, and DOE is investing in integrated research infrastructure (IRI)—the seamless connection of world-class experimental research tools and user facilities with high performance computing to empower researchers and radically accelerate scientific discovery. In 2023, NERSC participated in this work by assisting with the DOE IRI Architecture Blueprint Activity, a road map for the vision, strategy, and implementation across the national laboratory system, and offering experienced staff to serve with the IRI Working Group. NERSC’s experience with world-class HPC systems and expertise surrounding seamless interoperability in the integrated research ecosystem make it a key stakeholder for the future of IRI.

NESAP

The NERSC Science Acceleration Program (NESAP) has been a vital tool for deep technical engagements with users and vendors and proven highly effective for previous NERSC systems. However, due to the expanded range of complex workflows expected to be supported by future systems, the program is changing to accommodate the paradigm shift from application performance to end-to-end workflow performance. In 2023, NERSC introduced the NESAP Pathfinding program to improve the ability to assimilate newer technologies in HPC space.

Perlmutter Hits its Stride

In 2023, *Perlmutter* reached its full potential, transforming from a pre-production system to a production system. *Perlmutter* provided significantly enhanced performance for a wide range of applications. Application performance continued to improve on *Perlmutter*, with notable improvements to applications at scale. We saw a 30% increase in measured app speedup across a diverse set of representative apps. Over the course of the year, roughly 50 experiments from across DOE SC used *Perlmutter* for simulation and data analysis.

Quantum

Quantum computers are expected to gain some advantage over classical supercomputers for applications that align with the NERSC workload in the coming years. With the potential value of quantum computing for the NERSC user base in mind, NERSC has developed staff expertise and engaged domain scientists, QIS researchers, and quantum computing companies to help NERSC move into the new era of quantum computing. In particular, 2023 saw the expansion of the popular QIS@*Perlmutter* program and a successful collaboration with the neutral-atom quantum computing company QuEra.

Investing in the NERSC Community

NERSC's most important assets are its staff and the vibrant community of over 10,000 users who rely on NERSC systems for their research. In 2023, NERSC invested deeply in both groups. Through the robust NERSC User Group, over 4,000 users attended NERSC training events, and NERSC staff began work on improving and revamping the already significant slate of training materials available, with an emphasis on new users. Additionally, the debut of the series of staff presentations called NERSC Talks led to an overall greater understanding of work across the organization, better visibility for new staff, and improved professional and personal connection.

Monterey Data Conference

Leaders in the data science community gathered on the California coast in August, putting their heads together at the annual Monterey Data Conference. Organized by a committee led by Berkeley Lab, the conference convened researchers from DOE labs and facilities, universities, and industry to exchange ideas about the latest advances and challenges in scientific data analysis and computing. The theme of the meeting was Connections, an exploration of increasing integration between science, data, and society: cross-facility workflows and IRI, AI for connected science, data management and processing at scale, and more.

Together, these accomplishments are steps in the path of NERSC's continually improving capacity to fulfill its mission of accelerating scientific discovery through HPC, today and in the future. As we look to new paradigms, systems, and possibilities, we can look back on the work we did this year with pride.

— Sudip Dosanjh
NERSC Division Director

NERSC by the Numbers

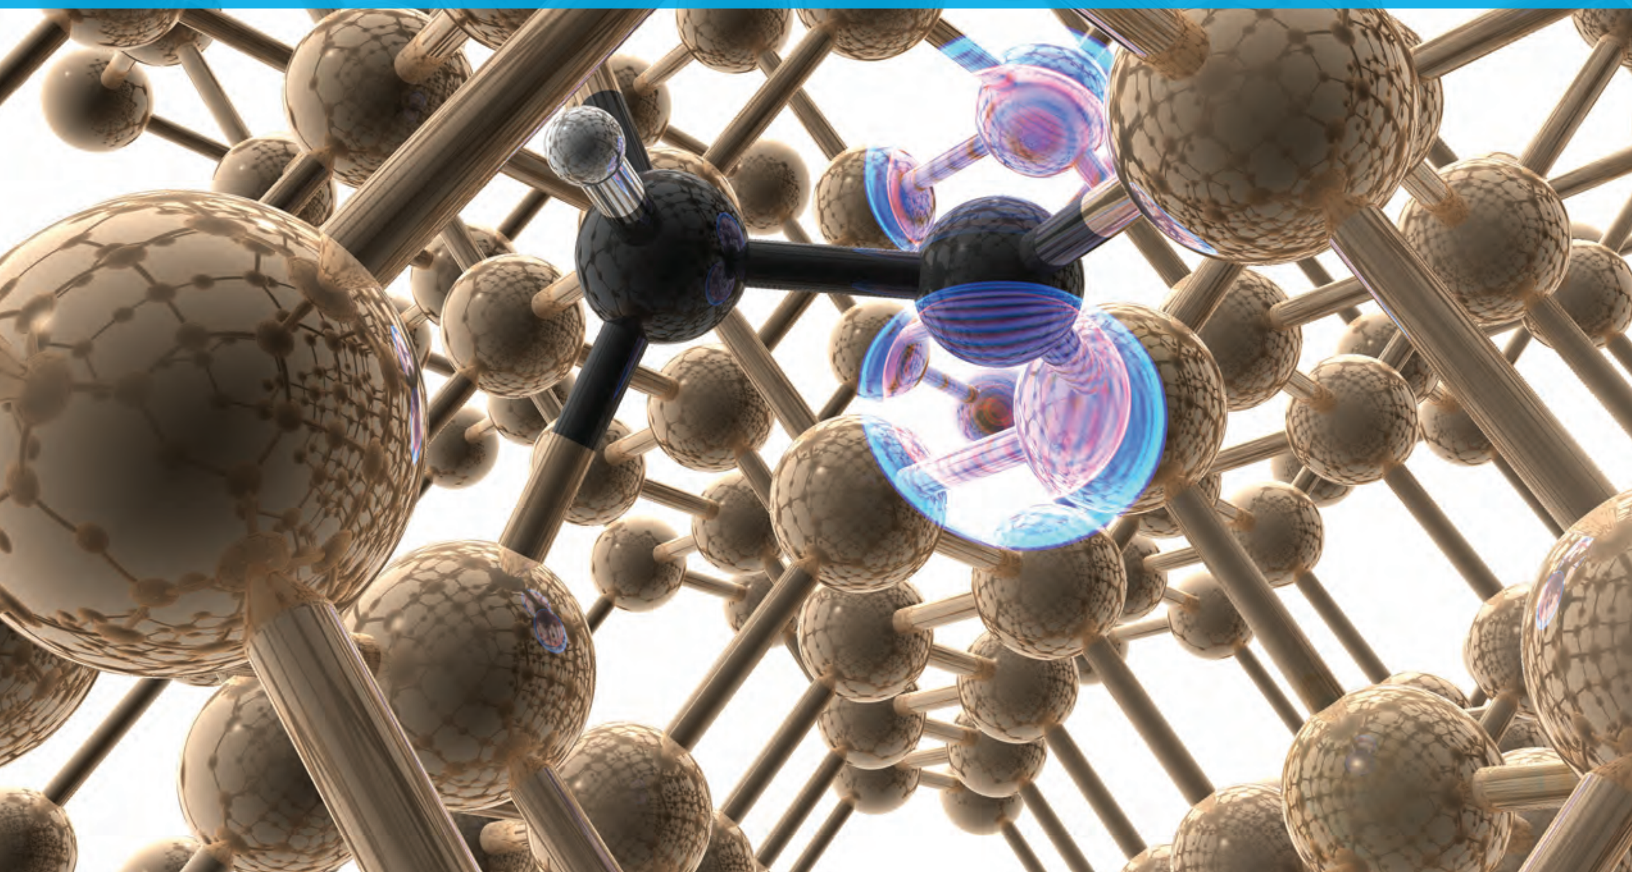

2023 NERSC USERS ACROSS US AND WORLD

50
States,
Washington D.C.
& Puerto Rico

46
Countries

~10,000 Annual Users from **~800** Institutions + National Labs

30% Graduate Students

18% Postdoctoral Fellows

16% Staff Scientists

14% University Faculty

9% Undergraduate Students

6% Professional Staff

60% Universities

28% DOE Labs

5% Other Government Labs

4% Industry

1% Small Businesses

<1% Private Labs

NERSC Systems 2023

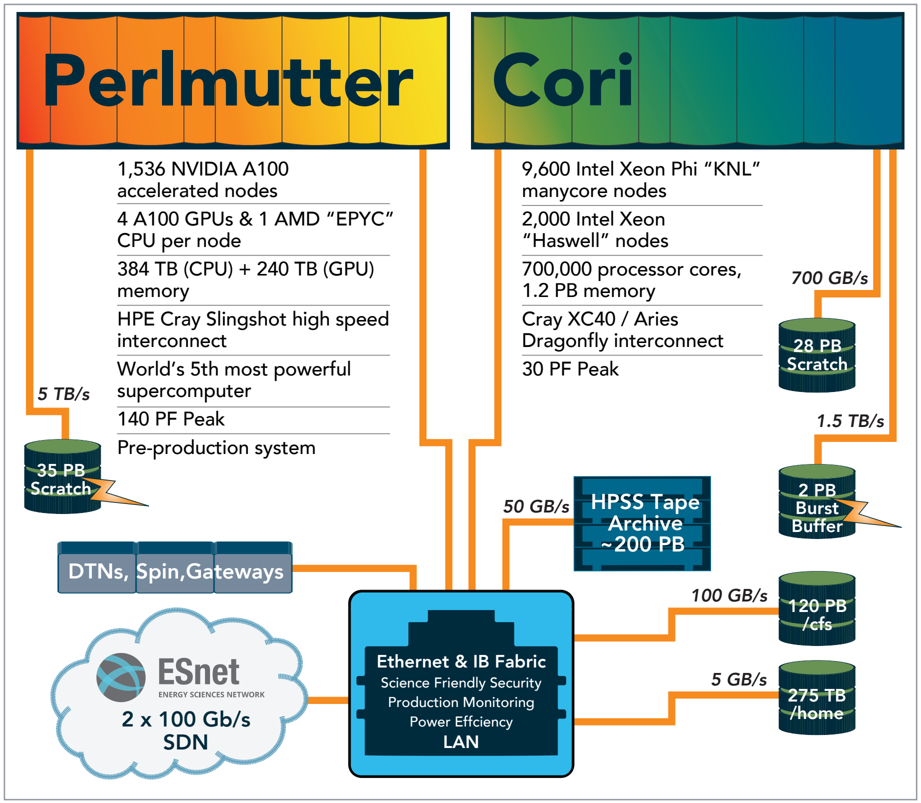

Figure 1. NERSC systems in 2023

TOP 5 DOE Labs by Users

- Berkeley Lab
- Oak Ridge National Laboratory
- Pacific Northwest National Laboratory
- Los Alamos National Laboratory
- Argonne National Laboratory

TOP 5 Universities by Users

- Berkeley, University of California
- Stanford University
- University of Washington
- Massachusetts Institute of Technology
- University of Michigan

Top Science Disciplines

(By computational hours used)

Chemical Sciences
Nuclear Physics
Fusion Energy
Scientific User Facilities
High Energy Physics
Biosciences
Geosciences
Small Business Innovation Research
Materials Sciences
Plasma Science
Climate and Environmental Science
Biological Systems Science

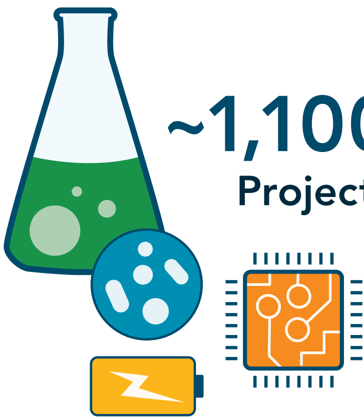

~1,100 Projects

Breakdown of Compute Used by DOE Program Office

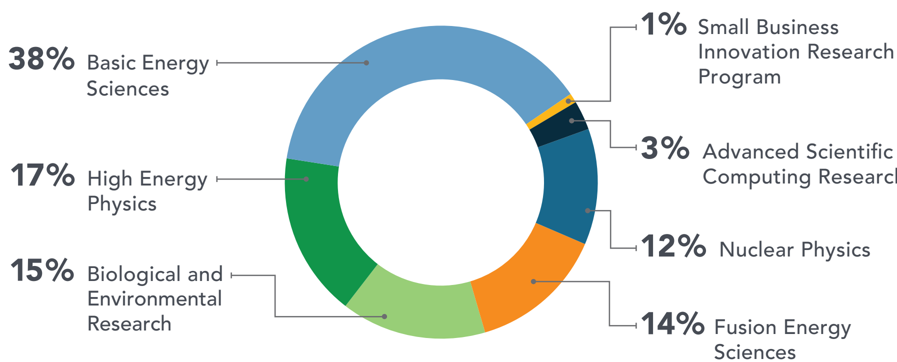

57 MILLION

COMPUTE HOURS USED IN 2023

>2,100

Refereed Publications cited NERSC

Data Stored

283

Petabytes

2023 Milestones

Preparing for Complex Workflows

For DOE SC to fulfill its mission to maintain and extend U.S. leadership in scientific discovery, a paradigm shift in the way HPC systems are designed, configured, and operated at ASCR facilities is required. NERSC is a fundamental part of that shift and is leading the way.

The future of HPC for science will go beyond the increasing demand for more computational cycles required by traditional simulation and modeling workloads. It will leverage new technologies and support emerging needs in AI and experimental and observational science to accelerate end-to-end workflows and enable new modes of scientific discovery through the integration of experiment, data analysis, and simulation. Future infrastructure will also address a growing need for more dynamic and programmable systems to accommodate complex workflows that may require running many interdependent simulations and/or analysis tasks. These workflows will move vast amounts of complex data among storage hierarchies, both for providing interactive real-time feedback to experiments and for integrating simulations with AI.

NERSC is preparing for the shifting future: planning new and fundamentally different systems, contributing thought leadership and infrastructure for IRI and AI, and providing resources to prepare NERSC users and the research community at large for the new paradigm.

NERSC-10

In 2023, NERSC made major strides toward its next flagship system, known as NERSC-10. Due in 2026, the NERSC-10 system will accelerate end-to-end DOE SC workflows and enable new modes of scientific discovery through the integration of experiment, data analysis, and simulation.

To help drive NERSC-10 system design, NERSC is engaging with both the vendor community and the user community to gain valuable input. In 2023, NERSC conducted an extensive market survey to better understand the changing vendor landscape and how new technologies, both hardware and software, could support DOE SC users' end-to-end workflows. NERSC has also gained insight around users' needs through technical deep dives into science workflows and outreach to the broader community.

Additionally, NERSC's successful NERSC Science Acceleration Program (NESAP) is growing and changing to accommodate the paradigm shift in HPC and continue to prepare the community for technologies to come.

IRI

Integrated workflows are the future of science research, and DOE is investing in integrated research infrastructure (IRI) — the seamless connection of world-class experimental research tools and user facilities with high performance computing facilities to empower researchers and radically accelerate scientific discovery. DOE made strong progress toward its IRI goals in 2023, publishing the IRI Architecture Blueprint Activity, a road map for the vision, strategy, and implementation across the national laboratory system.

NERSC has long been at the forefront of IRI, laying the groundwork with its Superfacility strategy beginning in 2015, continuing with the successful Superfacility Project, and contributing to the IRI Architecture Blueprint Activity through the active participation and leadership of many NERSC staff. NERSC's experience with world-class HPC systems and expertise surrounding seamless interoperability in the integrated research ecosystem make it a key stakeholder for the future of IRI.

In 2023, DOE began the implementation phase of IRI, beginning with the initiation of an IRI Working Group including NERSC staff. The charge of the working group is to coordinate closely between ASCR facilities, identifying areas of synergy where efforts can be more closely coordinated; engage with "pathfinder" science teams ready to explore IRI models; and identify potential initial areas of work. The team also began building the necessary management structures to develop and direct a complex program to support cross-facility science.

In addition to human infrastructure for IRI, the IRI Working Group also identified resources at every ASCR facility to be part of an IRI testbed, including the *Perlmutter*-on-Demand (POD) system at NERSC.

Perlmutter-on-Demand

In 2023, NERSC deployed POD as an IRI testbed to study the implementation of cross-facility and event-triggered workflow infrastructural support, though it will also play a role in investigating new capabilities for future NERSC systems. As an IRI testbed, POD will aid in the development and testing of

- co-scheduling capabilities across systems to have compute resources available when users need them,
- new workload management approaches to combine urgent compute with queued batch jobs in order to maximize system utilization,
- tools to identify bottlenecks in end-to-end workflows across multiple systems, which includes means to profile workflows across systems,
- advanced productivity capabilities (e.g., a fully containerized software deployment), and
- tools and infrastructure for edge computing that can surge to *Perlmutter*.

Additionally, POD is a tool for exploring approaches to increasing the resilience of the *Perlmutter* ecosystem, providing a place to experiment with a variety of software configurations and packages and to provide resources for urgent compute workloads if *Perlmutter* is undergoing maintenance.

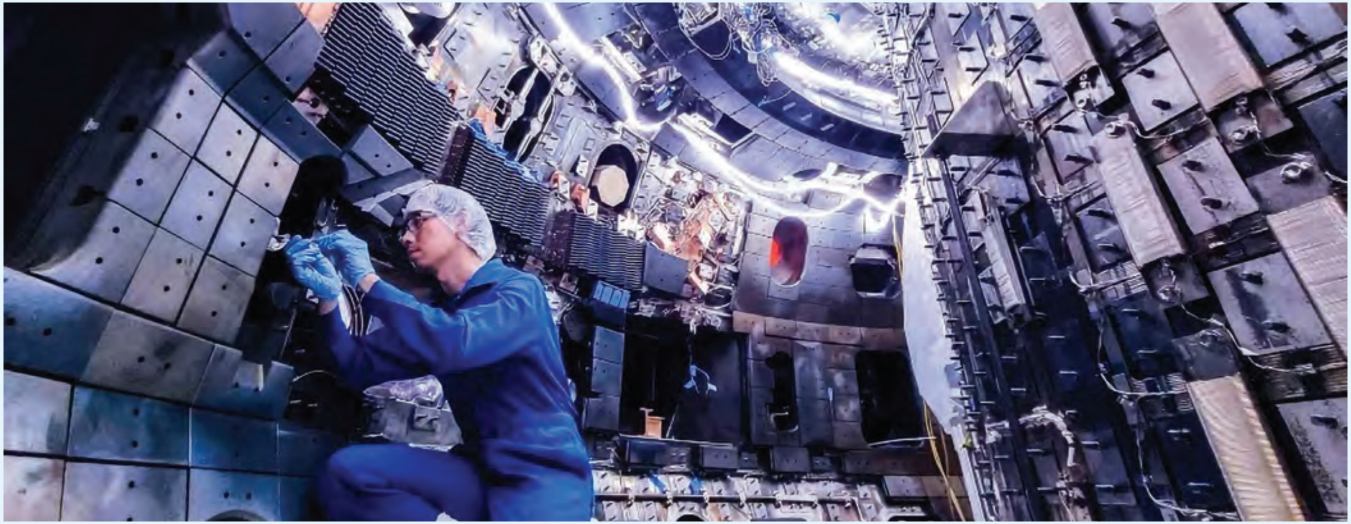

Figure 2. The DIII-D tokamak is a primary source of experimental information on fusion.

IRI in Action: DIII-D National Fusion Facility, NERSC, AMCR, and ESnet Collaboration Speeds Nuclear Fusion Research

As the global population continues to grow, demand for energy will also increase. One possible way of meeting that demand is through nuclear fusion, which does not emit carbon or consume scarce natural resources. Although hurdles remain before energy produced by nuclear fusion is widely available, DOE-funded researchers at the DIII-D National Fusion Facility, NERSC, Berkeley Lab's Applied Mathematics and Computational Research (AMCR) Division, and Energy Sciences Network (ESnet) are teaming up to bring that vision closer to reality.

At DIII-D, researchers study how to harness nuclear fusion energy as a viable energy source, a process that requires rapid processing of massive amounts of experimental data. To meet this challenge, a team from DIII-D, AMCR, NERSC, and ESnet collaborated to connect experimental resources at DIII-D and HPC resources at NERSC via ESnet — a collaboration to make data from fusion experiments more useful and available to the research community, speeding the realization of fusion energy production.

Data + High-Performance Networking + High-Speed Computing = Breakthroughs

“We have long recognized that all experimental science teams need a better way to connect with high-speed networks and HPC,” said NERSC Science Engagement & Workflows Department Head Debbie Bard, who leads the broader Superfacility work at Berkeley Lab. “We started the Superfacility work at Berkeley Lab as a broad initiative to develop the tools, infrastructure, and policies to enable these connections. DIII-D, AMCR, and ESnet have been key partners in this work.”

In the tokamak at the heart of the DIII-D National Fusion Facility, gas atoms are heated to temperatures hotter than the Sun, which causes them to disintegrate into their component electrons and nuclei. The free nuclei may crash into each other and fuse, releasing energy. Researchers study the behavior of short plasma discharges called shots, typically performed at 10- to 15-minute intervals with nearly 100 diagnostic and instrumentation systems capturing gigabytes of data during each shot. Between shots, scientists address any issues or evaluate how specific parameter settings affect plasma behavior.

Previously, these adjustments required extensive manual calculations by subject-matter experts, a time- and labor-intensive process that still offered limited information. Automating some of this work has always been a potential solution, but the computational needs of experiments at DIII-D were greater than could be addressed using standard computing systems, and gaining access to a separate supercomputing center has been prohibitively difficult and time-consuming.

Working together, staff from across the collaboration smoothed out the details of a data management workflow and decreased the time-to-solution by 80%, from 60 minutes to 11 minutes for a benchmark case. Between 2008 and 2022, before DIII-D was connected to NERSC, 4,000 hand-made reconstructions of the plasma’s behavior were produced. In the first six months of the DIII-D/NERSC Superfacility operation, DIII-D achieved more than 20,000 automated high-resolution magnetic field profile reconstructions for 555 shots.

“DIII-D (and fusion more generally) presents a use-case where returning results quickly really matters,” said Laurie Stephey, a member of the NERSC team. “Many types of fusion simulation and data analysis are too computationally demanding to be run on local resources between shots, so this often means that these analyses either never get done or are finished too late to be actionable. This Superfacility project combines DIII-D and HPC resources to produce something greater than the sum of their parts – just-in-time scientific results that would otherwise not be possible.”

AI

With the explosion of interest in and exploration of AI in the DOE science communities, HPC centers are preparing for a shift toward new AI-enhanced computational workflows. It is imperative that the scientific HPC community be ready to support this workload.

NERSC engages in research and development for scalable, powerful, and easy-to-use artificial intelligence and machine learning in scientific computing. This includes work across areas of AI system deployment, deep engagement with science teams toward AI-enabled applications and broad engagement and training of the wider scientific community. System deployment work has, for example, included benchmarking both existing and novel approaches on NERSC and other HPC systems as part of NERSC's leadership role in the industry-standard MLCommons organization. AI application engagement has involved the NESAP program, a slate of deep engagements with domain scientists applying AI to problems across the DOE Science portfolio. For broad training, NERSC offers robust education in topics related to AI for users and the HPC community as a whole, including a continuing tutorial at the Supercomputing Conference and a broad range of NERSC-user training events.

Organizing and Publishing the MLPerf HPC Benchmark Suite

MLPerf benchmarks from the MLCommons organization are the industry standard benchmark for AI performance—that is, current drivers of innovation in systems and software. NERSC co-founded the HPC working group within MLCommons, which developed the MLPerf HPC benchmark suite to address HPC and scientific AI workloads, and has continued to participate in planning and organization through the MLPerf HPC v3.0 submission round.

During this submission round, the group added the new OpenFold benchmark, a protein-folding application based on DeepMind's famous AlphaFold2 model, and laid the groundwork for a robust power measurement methodology in coordination with the MLCommons Training group. The MLPerf HPC v3.0 results were published in November 2023 and featured 30 performance results (a 50% increase over the previous year), new submitters, and impressive speed-ups over previous submission rounds—the latest DeepCAM results are 14X faster than when the benchmark debuted.

NERSC partnered with HPE to submit results to MLPerf HPC v3.0 using *Perlmutter* in its final configuration with upgraded Slingshot 11 network. The *Perlmutter* system achieved excellent results across all workloads and factors of node scaling with impressive speedups over the previous v1.0 results, including 5.3X faster results on OpenCatalyst, 2X faster results on CosmoFlow, 1.4X faster results on DeepCAM, and highly competitive results on the new OpenFold benchmark.

NESAP

The NERSC Science Acceleration Program (NESAP) is a collaborative effort in which NERSC partners with code teams, vendors, and library and tool developers to prepare for advanced architectures and new systems. NESAP began in late 2014 to help users prepare for the Cori manycore Knights Landing/Xeon Phi architecture. The program was crucial in improving code offloading to GPUs and improving computational efficiencies of prominent codes running on *Perlmutter*. Recently, the demands of the NERSC-10 system have led NERSC to reframe NESAP in order to support the integration of HPC applications into complex and cross-facility workflows.

NERSC-10 and NESAP Pathfinding

NESAP has been a vital tool for deep technical engagements with users and vendors and proven itself highly effective for previous systems. However, due to the expanded range of complex workflows expected to be supported by NERSC-10 and future systems, the NESAP program is changing to accommodate the paradigm shift from application performance to end-to-end workflow performance.

To improve the ability to assimilate newer technologies in HPC space, NERSC has introduced NESAP Pathfinding, a one-year limited project engagement. This strategic change is driven by the observation that the workflows community is evolving quickly, necessitating a more agile approach to better anticipate future user needs. In October 2023, NERSC released the NESAP Pathfinding call for proposals and received 53 applications, of which 24 were selected for staff and *Perlmutter* time.

Sudip Dosanjh presents at the SC24 conference.
(Credit: Bonnie Powell, Berkeley Lab)

One of the driving forces for the workflow capability requirements in NERSC-10 is the IRI program; the IRI Blueprint Activity Final Report identified several categories of workflow patterns crucial to IRI. As a result, the NESAP for Workflows Pathfinding Activity devotes significant resources to exploring cross-facility workflows.

As part of this exploration of advanced workflow capabilities, NERSC has also begun incorporating technical groups, which facilitate collaborations between teams sharing similar technical requirements. In 2024, the program will pilot the Superfacility technical group to evaluate its effectiveness in fostering synergy and knowledge exchange among participating teams. By leveraging collective expertise and resources, NESAP aims to optimize technical support and accelerate the pace of scientific discovery within the NERSC user community.

Perlmutter supercomputer at the National Energy Research Scientific Computing Center (NERSC). *(Credit: Thor Swift, Berkeley Lab)*

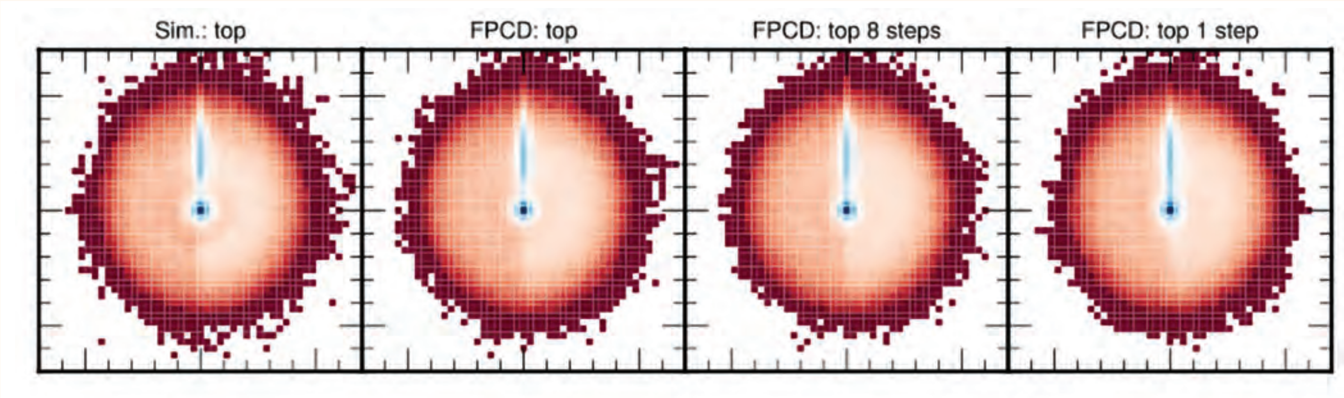

Figure 3. Comparison of true simulated particles (further left) with generated particles using different generation speeds (second to fourth from the left). Generation quality is retained even at the fastest setting, resulting in speedups of a factor of 500 to 2000 compared to the traditional simulation.

NESAP in Action: Machine Learning for High-Energy Physics

Simulations are used in high-energy physics (HEP) for a wide array of tasks, including data selection, statistical inference, and experimental design. As detectors become more sophisticated and data volumes grow, it's expected that the computational cost of efficiently producing simulations of current and next- generation HEP experiments will soon surpass the computing budget available. For example, detector simulation for the ATLAS and CMS experiments consumed 40% of the grid central processing unit (CPU) during Run 2 of the Large Hadron Collider (LHC), and the expected CPU time needed to simulate a single LHC event is expected to increase by a factor of three or more after the High-Luminosity LHC upgrade in the coming years. While running the full simulation ensures high data fidelity, the computational becomes prohibitive as many billions of simulated events are required to describe different Standard Model and Beyond the Standard Model processes.

In recent works supported by NESAP, NERSC and Berkeley Lab physics division researchers developed new algorithms for fast simulation using the *Perlmutter* supercomputer. These algorithms are capable of reproducing the underlying physics process with precision 500 to 2000 times faster than traditional algorithms.

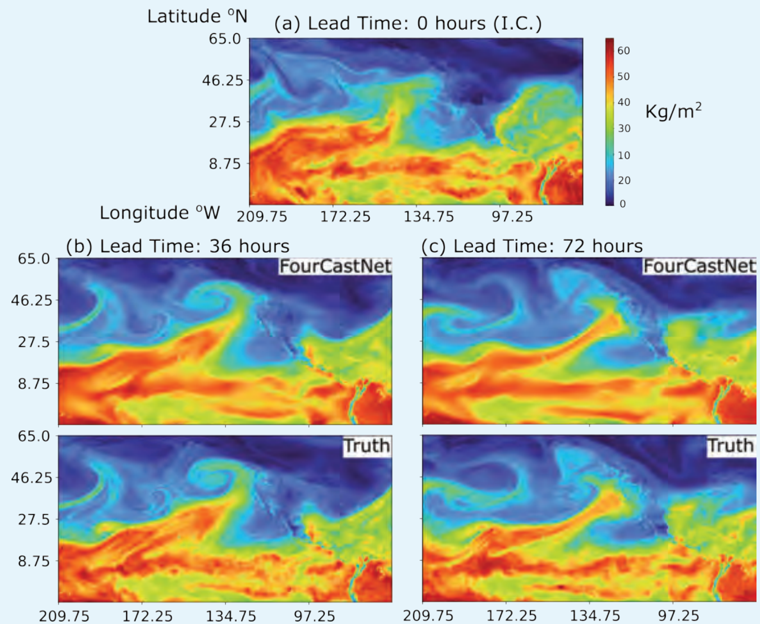

Figure 4. The Fourier Forecasting Neural Network (FourCastNet) is a data-driven weather forecasting model integrating HPC and AI.

NESAP in Action: FourCastNet

One successful project supported by the NESAP program is the Fourier Forecasting Neural Network (FourCastNet), a data-driven weather forecasting model integrating HPC and AI to provide accurate fine-grained (0.25°, or about 30 km x 30 km near the equator) forecasting across the globe. FourCastNet offers improved model complexity, computational cost, and scalability compared with previous models.

The rapid rise of deep learning in numerical weather prediction (NWP) has led to a proliferation of models that forecast atmospheric variables with equal or greater success than traditional physics-based NWP. These leading models encompass a wide variety of model architectures and training techniques. Further, the parts of AI models that contribute most to an effective model have not been thoroughly studied, making it hard to discern which components are most critical.

Researchers at NERSC and Berkeley Lab have shown that it is possible to attain high forecast skill even with relatively off-the-shelf architectures, simple training procedures, and moderate compute budgets. For example, they trained a minimally modified SwinV2 transformer on ERA5 reanalysis data and found that it performed better than the existing operational state-of-the-art Integrated Forecast System from the European Centre for Medium-Range Weather Forecasts. The models were built using Pytorch Shifter containers on *Perlmutter* and trained on 64 GPUs using data parallelism, and completed training in less than a day.

The remarkable forecast speed of models like FourCastNet enables new possibilities in weather and climate analysis, particularly in the ability to generate large ensemble forecasts at a fraction of the computational cost required by traditional NWP models. In collaboration with climate scientists from Berkeley Lab Earth Sciences and researchers from NVIDIA, the FourCastNeXT NESAP project has demonstrated for the first time the use of a deep learning forecast model to generate ensemble forecasts at full global scale with over 7000 ensemble members. Generating an ensemble of this size required developing new methods for initializing the ensemble in order to achieve ensemble spread and forecast skill comparable to NWP models, and the production forecasting runs (using data from the summer of 2023, the hottest on record) ran on 256 *Perlmutter* GPUs for three days, generating almost 3 PB of data to be analyzed. The ensemble results demonstrated that it is possible to reliably sample extreme events as far as 4 sigma from the ensemble mean and improve uncertainty quantification for extreme weather by over an order of magnitude.

Perlmutter Hits Its Stride

In 2023, NERSC transformed *Perlmutter* from a pre-production system to a hardened, resilient production system. The system reached its fully operational form in July and achieved nearly 95% utilization in the June-to-December timeframe after its predecessor, Cori, retired.

Since completing the installation process and testing phases, *Perlmutter* has proven a nimble and reliable system for science, running 6.7 million jobs in the second half of the year with minimal interruptions for researchers. At the SC23 conference in November, two science teams using *Perlmutter* were awarded the prestigious Gordon Bell Prize for outstanding achievement in HPC.

Perlmutter Supports Two Gordon Bell-Winning Projects in 2023

One, the project “The Simple Cloud-Resolving E3SM Atmosphere Model (SCREAM) Running on the Frontier Exascale System,” part of the Energy Exascale Earth System Model (E3SM), won first prize for the new Gordon Bell Climate Modeling award.

SCREAM is a full-featured atmospheric general-circulation model developed for very fine-resolution simulations on exascale machines, incorporating state-of-the-art parameterizations for fluid dynamics, microphysics, moist turbulence, and radiation. The project is a multi-lab effort led by Lawrence Livermore National Laboratory and includes NERSC staff member Noel Keen as an integral team member. SCREAM demonstrated a record-setting simulation of 1.26 simulated years per day for a realistic cloud-resolution model on the Frontier system at ORNL, the first time a new algorithm for efficient and accurate simulations of deep convective clouds could run at scale on nearly the entire exascale supercomputer.

While the focus of this project was its use on the new Frontier exascale supercomputer, much of its developmental work was performed on *Perlmutter*. In particular, *Perlmutter*’s CPU nodes were instrumental in debugging and improving the performance of SCREAM before its scaled runs on Frontier. The first simulations with all the integral pieces of E3SM and SCREAM ran on *Perlmutter*, making SCREAM the first such model to run on both AMD GPUs and NVIDIA GPUs, as well as the first to run at scale on nearly an entire exascale system.

Figure 5. A multi-institutional team won the first-ever Gordon Bell Prize for Climate Modelling; they used *Perlmutter* to improve the performance of their algorithm.

Figure 6. A multi-institutional team won the Gordon Bell Prize for their work on large-scale materials modeling at quantum accuracy, using *Perlmutter* to simulate the stability of quasicrystals in a ytterbium-cadmium alloy.

Additionally, the team behind the Gordon Bell Prize-winning project "Large-scale Materials Modeling at Quantum Accuracy: Ab Initio Simulations of Quasicrystals and Interacting Extended Defects in Metallic Alloys" also used *Perlmutter* in their work. Although most simulations for this work were performed at ORNL, the team separately ran simulations on *Perlmutter* to study the stability of quasicrystals—an ordered but not periodic structure—in a ytterbium-cadmium alloy.

Retiring Cori

After nearly seven years of service, thousands of user projects, and tens of billions of compute hours, the Cori supercomputer retired on May 31, 2023. With its first cabinets installed in 2015 and the system fully deployed by 2016, Cori had been in service longer than any supercomputer in NERSC's 49-year history and enabled more than 10,000 scientific publications. Its technological innovations reflect the dynamic evolution of HPC over the past decade, paving the way for the next generation of scientific computing.

Cori was designed through a partnership with Intel, Cray (now HPE), and Los Alamos and Sandia National Laboratories. The Cray XC40 system was named in honor of biochemist Gerty Cori, the first American woman to win a Nobel Prize in science and the first woman to win a Nobel Prize for Physiology or Medicine. It comprised 2,388 Intel Xeon Haswell processor nodes, 9,688 Intel Xeon Phi Knight's Landing (KNL) nodes, and a 1.8 PB Cray Data Warp Burst Buffer and had a peak performance of ~30 petaflops; when it debuted in 2017, it ranked fifth on the TOP500 list.

Cori was also the first supercomputer to be installed from the ground up in Shyh Wang Hall, influencing the building's infrastructure and prompting the implementation of numerous energy-efficiency innovations on the system and throughout the facility. In addition, the introduction of Cori's manycore KNL architecture changed the way NERSC interacts with users, leading to the implementation of the NESAP program, which is only growing and accelerating in the *Perlmutter* GPU era.

Figure 7. The Cori supercomputer was installed in 2016. When it was retired in May 2023, it was NERSC's longest-serving system.

Other innovative features introduced on Cori that have influenced current and next-generation architectures include the Burst Buffer (which laid the foundation for *Perlmutter*'s all-flash file system), high-bandwidth memory, increased vector capability, real-time queues, deep learning library support, Globus sharing connections, and workflow service nodes.

Quantum

Quantum information science (QIS) and quantum computing (QC) are up-and-coming technologies that hold immense promise and also require intense investment in order to become reality. These technologies may one day provide a powerful approach to solving complex computational problems relevant to NERSC's mission.

Quantum computers are expected to gain some advantage over classical supercomputers for applications that align with the NERSC workload in the coming years. Specifically, applications in chemistry, materials science, and high-energy physics will likely benefit from QC, and QC technology may open new research directions impossible to pursue with classical supercomputers.

With the potential value of QC for the NERSC user base in mind, NERSC has developed staff expertise and engaged domain scientists, QIS researchers, and quantum computing companies to help NERSC move into the new era of quantum computing.

Supporting Users

The QIS@*Perlmutter* program, established in Allocation Year 2022 to advance research into QIS for HPC workloads, was extended into Allocation Year 2023. In 2023, 17 new projects were awarded over 150,000 GPU node hours. The program resulted in many new scientific publications and has been continued for a third year.

NERSC's second Quantum for Science Day event celebrating the role and position of quantum computing in science research took place in November, welcoming over 50 people in person and at least 50 more virtually. The program featured a panel discussion on quantum computing efforts at HPC centers as well as talks and tutorials delivered by quantum computing vendors and QIS@*Perlmutter* users.

Collaborating with Industry

In March 2023, NERSC announced its first R&D partnership with QuEra computing, a full-stack quantum computing company based in Boston that is building hardware based on neutral atom technology. In 2023, NERSC staff worked with QuEra scientists on a broad range of activities, including:

- an introductory two-day virtual user training event that attracted over 70 participants
- a series of low-level training sessions on QuEra computing quantum technologies for NERSC staff delivered by QuEra scientists
- testing and verification of Aquilla, QuEra computing's analog neutral atom quantum computer
- joint application co-design research for which both teams worked toward multiple scientific publications

In addition, NERSC staff continued developing strategic relationships with vendors in the US and the international quantum computing industry, leveraging virtual briefings and onsite visits with vendors developing a range of qubit technologies. These relationships place NERSC in a position to assess the evolution of quantum technology and engage with vendors across the technological landscape.

At the 2023 New User Group meeting, Xanadu Technologies provided NERSC users with an introduction to using their software, Pennylane, to simulate quantum circuits. Xanadu was also selected as a participant in the 2024 NESAP program. NERSC staff will work with Xanadu scientists to run large-scale simulations of quantum machine learning applications on *Perlmutter*. The goal is to simulate well beyond the state of the art to improve NERSC's understanding of the potential of quantum machine learning algorithms.

Benchmarking the Future

NERSC staff has continued developing and studying new quantum computer applications in a series of publications released in 2023. With collaborators from a range of institutions, the team developed HamLib, a curated dataset of quantum Hamiltonians designed for benchmarking quantum algorithms and hardware. Additionally, NERSC researchers have been awarded a DOE Pathfinder grant to develop HamPerf, a suite of quantum benchmarks built on top of HamLib. NERSC researchers were also part of a large, multi-institutional team awarded a research grant from the Wellcome Leap to study the potential of quantum computing algorithms for applications in biology and life sciences.

NERSC researchers are also analyzing the resources required to support future large-scale quantum computations. A paper published as part of the Performance Modeling, Benchmarking, and Simulation of HPC Systems (PMBS) workshop at SC23 showed that quantum advantage-scale computations could start at $1 million under current cloud pricing models. In a second work, the NERSC quantum team modeled the classical compute requirements for supporting large-scale quantum computations during run time and found that about one petaflop of compute suffices for real-time quantum error decoding at scale.

Investing in the NERSC User Community

NERSC is a valued resource for 10,000 researchers working in a range of fields and hailing from around the globe. NERSC engages deeply with its user community, prioritizing collaboration and taking a multifaceted approach to user success.

The vibrant NERSC User Group (NUG) and its executive committee (NUGEx) connect with users through monthly webinars and an annual user meeting. The NESAP and IRI engagements constitute deep partnerships with science teams and their application and workflow development efforts. An annual user survey gathers key data: how users interact with NERSC, what they need to succeed, and how their experience can be improved. And to contribute to the growth and health of the user community, NERSC delivers a varied set of outreach events, including hackathons, workshops, conferences, and training events.

In 2023, training events spanned a broad range of topics appealing to a wide range of audiences, from novice to experienced. Topics included getting started at NERSC, coding for GPUs, machine learning/deep learning, using performance tools, running applications, programming models, and services. Attendance ranged from 15–20 to more than 300 people per event, for a total of more than 4,000 users.

Figure 8. Top NERSC User Training Needs and Key Training Events Addressing Them

TOPIC	EVENTS
Leveraging the power of *Perlmutter's* GPUs: developing and using GPU capable applications	• N-Ways to GPU Programming Bootcamp • July GPU Hackathon • GPUs for Science Day • Optimizing Applications for GPUs using OpenMP Offload and OpenACC
Moving codes and workflows from Cori, which retired in May 2023, onto *Perlmutter*	• Migrating from Cori to *Perlmutter* Training
Leveraging the power of AI techniques	• AI for Scientific Computing Bootcamp
Developing complex workflows and integration of HPC with data from experimental facilities	• DOE Cross-Facility Workflows Training • NERSC SuperFacility API Training • 5 Spin Training Sessions
Learning how to use NERSC for the approximately 33% of users who are new to the facility	• New User Training 2023

In particular, in response to user feedback, NERSC is investing in improving the experiences of new users, aiming to support HPC novices in current work and build their skills for the future. In 2023, NUG offered two live two-day new-user training events and began developing a set of asynchronous resources for onboarding new users, including a welcome brochure providing important, basic information regarding NERSC operations and user opportunities and a more comprehensive and detailed onboarding guide available in the technical documentation and as an online course.

Overall, users have an overwhelmingly positive view of NERSC; in the 2023 user survey, users rated NERSC at 5.21 on a six-point scale and offered the following feedback:

"[T]he experts who answer tickets are amazing and extremely knowledgeable... [W]hen I ask them to explain things for my own benefit so I can [understand] instead of just resolving the problem, many give detailed and insightful responses."

"[T]he support group for tickets are phenomenal. I get quick responses and my problems are always solved quickly."

"Training offerings [are] always valuable & excellent."

"NERSC is a friendly institution, with users in mind, approachable, unparalleled by other labs. I am extremely satisfied with their support and computing power."

The "documentation [is] exceptional. I can almost always get the help that I need there."

"NERSC provides a good selection of optimized, up-to-date software that performs well on their CPU- and GPU-based resources."

"NERSC has been very helpful in my research journey, my queries were resolved in swift time. Super fast and efficient and communication is effective."

"NERSC does an excellent job of providing world-class computing resources and responding to users' requests. NERSC offers an excellent user experience."

Figure 9. NERSC staff gather quarterly to connect and learn from one another.

Investing in the NERSC Staff Community

NERSC TALKS

To encourage staff to innovate as a team and help provide an understanding of work across the organization, NERSC has implemented a communication platform called NERSC Talks, in which each staff member presents an onsite, fifteen-minute technical talk with a five-minute question-and-answer session.

NERSC Talks have resulted in an overall greater understanding of work across the organization, better visibility for new staff, and improved professional and personal connection. In addition to the opportunity to gather onsite for these talks, many staff took the opportunity to gather socially during the following lunch periods to discuss recent talks, create potential collaborations, and learn more about the research or projects at hand.

NERSC Talks have been so successful that they have been continued for another year.

Science Highlights

NERSC'S strategic planning and engagement activities helped enable yet another highly productive year of science for thousands of users — the most in the DOE national lab complex. Here are just a few of the projects and scientific discoveries NERSC supported in 2023:

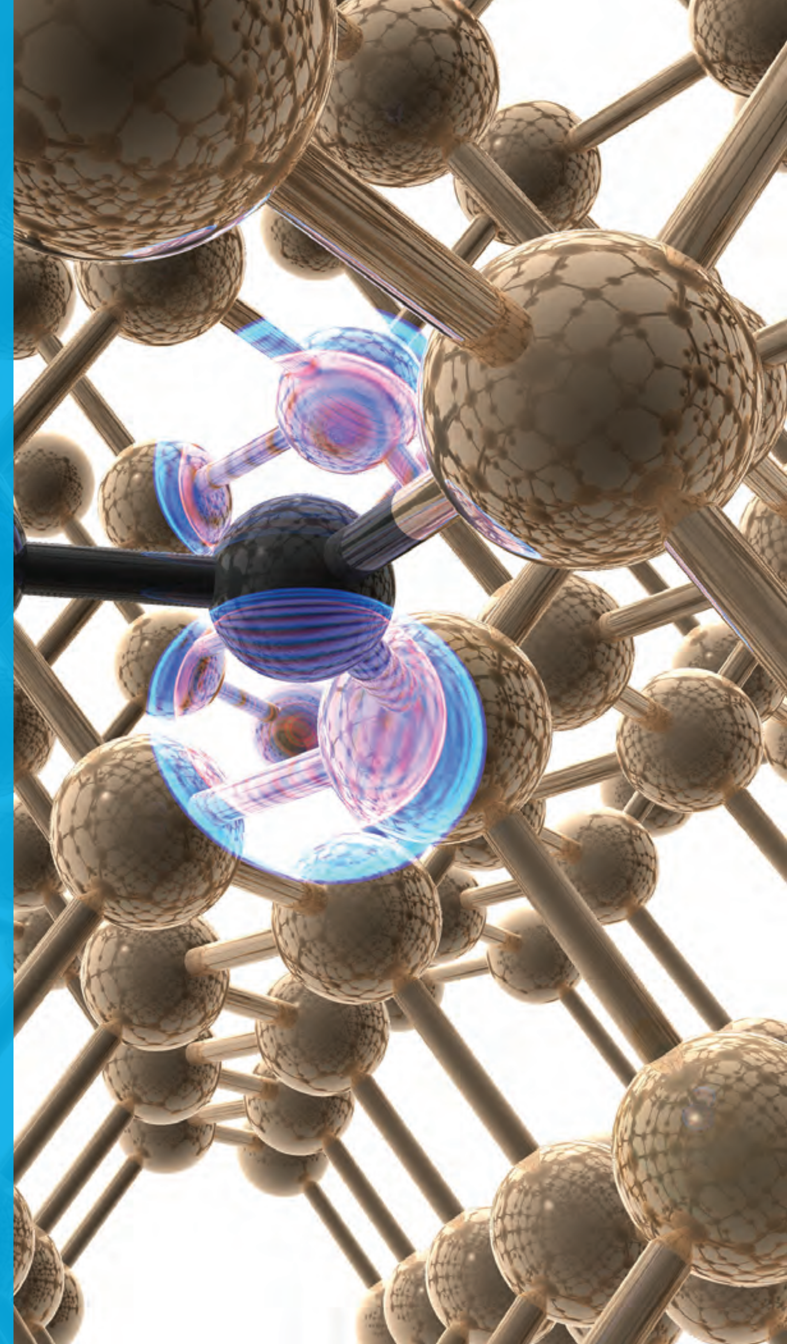

NERSC Supports First All-GPU Full-Scale Physics Simulation

THE SCIENCE. Using supercomputers at NERSC, researchers have completed a simulation of a detector of neutrino interactions that's designed to run exclusively on GPUs – the first simulation of its kind and an example of using GPUs' highly parallel structure to process large amounts of physical data. The simulation is part of the preparation for the Deep Underground Neutrino Experiment (DUNE), an international collaboration studying the neutrino that includes DOE resources.

Neutrinos are the most abundant particle of matter in the universe, and efforts to measure the neutrino's mass and understand its relationship to matter in the universe have been underway for decades. Researchers typically study neutrinos using a device called a liquid argon time projection chamber (LArTPC), in which neutrinos flow through liquid argon and leave ions and electrons in their wake, which are received by arrays of sensing wires. With the wires' position and the charges' time of arrival for input, the wires produce 2D images of the neutrino interactions. However, a new method pioneered at Berkeley Lab replaces the wires with sensors known as pixels, which add a third sensing dimension and yield 3D images instead – an increase in information and much more data to analyze and store. (In this case, a "pixel" is a small sensor that is unrelated to the pixels found on consumer electronics screens.) Before construction of the detectors begins, the team uses digital simulations to ensure that both the physical detectors and the workflows around them will work as planned.

THE IMPACT. Because GPUs are uniquely suited to executing many calculations in parallel, they represent a much faster way of dealing with large quantities of data. At NERSC, *Perlmutter's* thousands of GPU nodes allowed the researchers to simulate the detector over many nodes simultaneously, greatly increasing the compute capability relative to using only CPUs. This workflow yielded an associated increase in speed – the simulation of the signal from each pixel took about one millisecond on the GPU compared with ten seconds on a CPU, a major step forward for the DUNE experiment as well as for the field of high-energy physics, particularly the question of why the universe is made of matter and not antimatter.

ADDITIONAL DETAILS. The research was published in the Journal of Instrumentation in April.

NERSC PI: Dan Dwyer, Lawrence Berkeley National Laboratory

PROJECT FUNDING AND ALLOCATION AWARD: DOE Office of Science, Office of High Energy Physics

PUBLICATION: Abed Abud, A., et al., "Highly-parallelized simulation of a pixelated LArTPC on a GPU." JINST 18, P04034 (2023). https://doi.org/10.1088/1748-0221/18/04/P04034

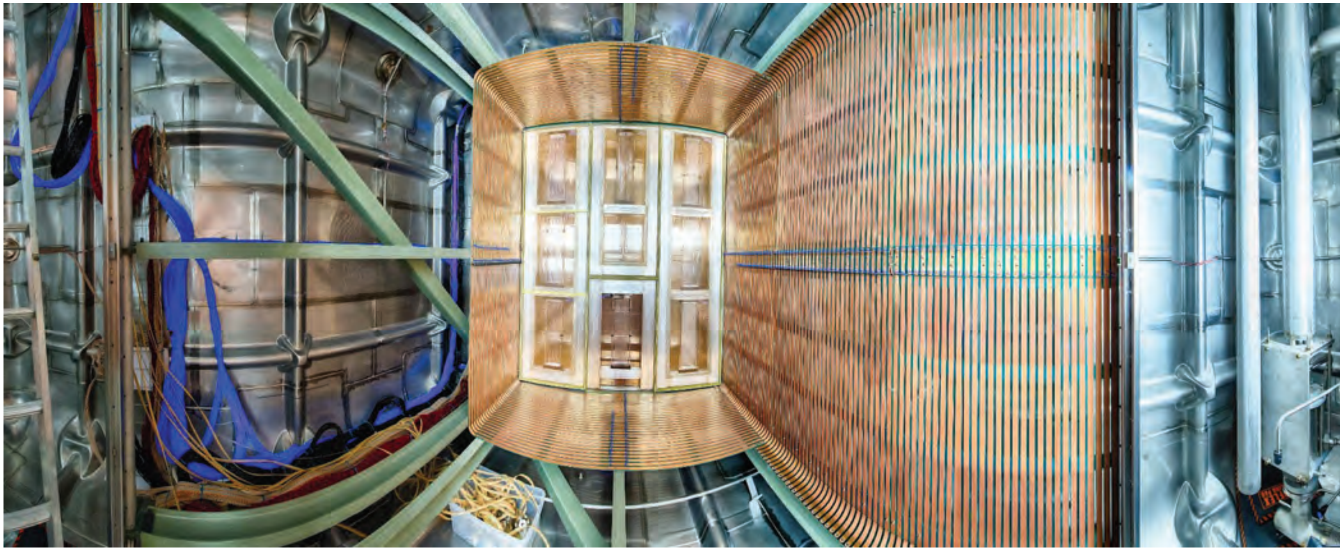

Figure 10. A prototype cryostat with 35-ton capacity is part of the infrastructure needed for the DUNE project.

NERSC Collaboration Shines a Light on Microbial Dark Matter

THE SCIENCE. Researchers in Berkeley Lab's Applied Mathematics and Computational Research Division (AMCR) and at the Joint Genome Institute (JGI), with the help of supercomputers at NERSC and support from the DOE's Exabiome Exascale Computing Project, have developed new tools to solve the problem of "microbial functional dark matter", grouping and identifying billions of protein sequences and vastly increasing the number of protein sequences recognized by science.

THE IMPACT. Metagenomics is a method of identifying all the genetic entities in a particular environment—researchers take comprehensive environmental samples and analyze the genetic profile of the entire sample at once, and reference databases of known genomes help them identify and sort what they find. However, only a small portion of the data matches up with those known protein families; as much as 93% has no match and is considered microbial functional dark matter, and simply discarded from the data.

The illumination of so much microbial dark matter represents a significant advancement in biologists' use of HPC.

ADDITIONAL DETAILS. The project analyzed unknown proteins from 26,931 metagenomics databases, having already subtracted the genes with matches to known proteins in a database of over 100,000 reference genomes. They identified 1.17 billion protein sequences of longer than 35 amino acids with no previously known match and organized them into 106,198 sequence clusters of more than 100 members each. Then, they annotated the families according to their taxonomy, habitat, geography, and gene neighborhood distributions and, in some cases, predicted their structures in 3D. By organizing them in this way, researchers can begin to understand what each protein's function might be across species.

Aydin Buluç, AMCR research scientist Oguz Selvitopi, and Berkeley Lab affiliate research scientist Ariful Azad used HipMCL, a massively parallel implementation of the Markov Clustering algorithm the researchers previously developed for the Exascale Computing Project's Exabiome program. Using 2,500 compute nodes on the Cori system – about a quarter of the system – they took advantage of HipMCL's ability to run on distributed-memory computers and crunch these previously uncrunchable numbers.

NERSC PI: Aydin Buluç, Lawrence Berkeley National Laboratory

PROJECT FUNDING AND ALLOCATION AWARD: DOE Office of Science, Office of Biological and Environmental Research

PUBLICATIONS: Pavlopoulos, G.A., Baltoumas, F.A., Liu, S. et al. Unraveling the functional dark matter through global metagenomics. Nature 622, 594–602 (2023). https://doi.org/10.1038/s41586-023-06583-7

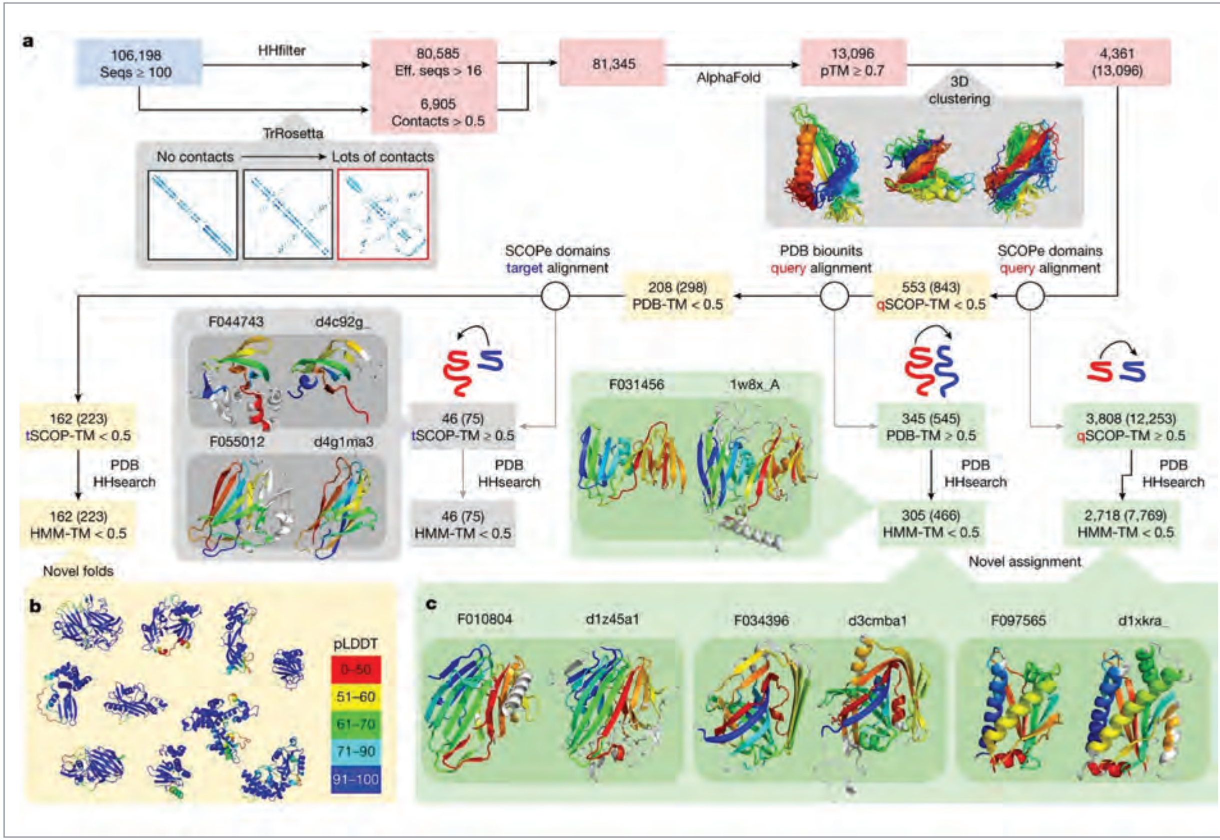

Figure 11. The researchers organized 1.17 billion protein sequences into over 106,000 families using computational means.

Perlmutter Provides Peek into the Interior of Ice Giant Planets

THE SCIENCE. The ice giant planets Uranus, Neptune, and sub-Neptune exoplanets are believed to be composed primarily of the elements hydrogen, carbon, nitrogen, and oxygen; however, the details of their interior structures are largely unknown due to unclarity around the behavior of compounds H2O, NH3, and CH4 at high pressures and temperatures. Using *Perlmutter*, UC Berkeley Earth and Planetary Science Ph.D. student Kyla de Villa and colleagues produced ab initio simulations of 13 H-C-N-O compounds and demonstrated that at high temperatures, they assume a superionic state – a phase of matter between solid and liquid in which hydrogen ions diffuse like a liquid through a solid sublattice of larger nuclei. At yet higher temperatures, four of the 13 compounds demonstrated double superionicity: in addition to hydrogen, either carbon or oxygen also diffused like a liquid through the sublattice of the remaining large nuclei.

THE IMPACT. Ultimately, the researchers could perform AI-accelerated molecular dynamics simulations of a sample of 1925 atoms on a single *Perlmutter* node in 1/25 the time required to simulate 96 atoms on two nodes using VASP. Additionally, *Perlmutter* provided a much faster simulation experience than Cori: the researchers saw a 6x speedup on *Perlmutter* for VASP simulations, a 30x speedup for training the neural network, and more than a 100x speedup for molecular dynamics simulations using LAMMPS. These performance improvements represent major progress for H-C-N-O chemistry and researchers' understanding of ice giant planets.

Applications for superionic properties aren't limited to planetary sciences, however. Superionic materials are ionically conductive due to the flow of charged particles, even at pressures and temperatures that are too low to allow them to be electronically conductive. Many solid-state batteries incorporate superionic materials, and a better understanding of superionicity could help scientists create better batteries for various applications.

ADDITIONAL DETAILS. The researchers first performed DFT calculations using VASP to simulate each compound. However, to work around the computational cost and limitations to size and timescale of DFT simulations and allow for scaled-up explorations, they went on to generate machine-learning potentials. They did this by running many DFT calculations and feeding the results into a neural network, generating a potential energy surface that could be used to run much faster and less expensive molecular dynamics simulations using LAMMPS. The team's work was published in Nature Communications.

NERSC PI: Burkhard Militzer, University of California, Berkeley

PROJECT FUNDING AND ALLOCATION AWARD: DOE Office of Science, Office of Fusion Energy Science

PUBLICATION: de Villa, K., González-Cataldo, F. & Militzer, B. Double superionicity in icy compounds at planetary interior conditions. Nat Commun 14, 7580 (2023). https://doi.org/10.1038/s41467-023-42958-0

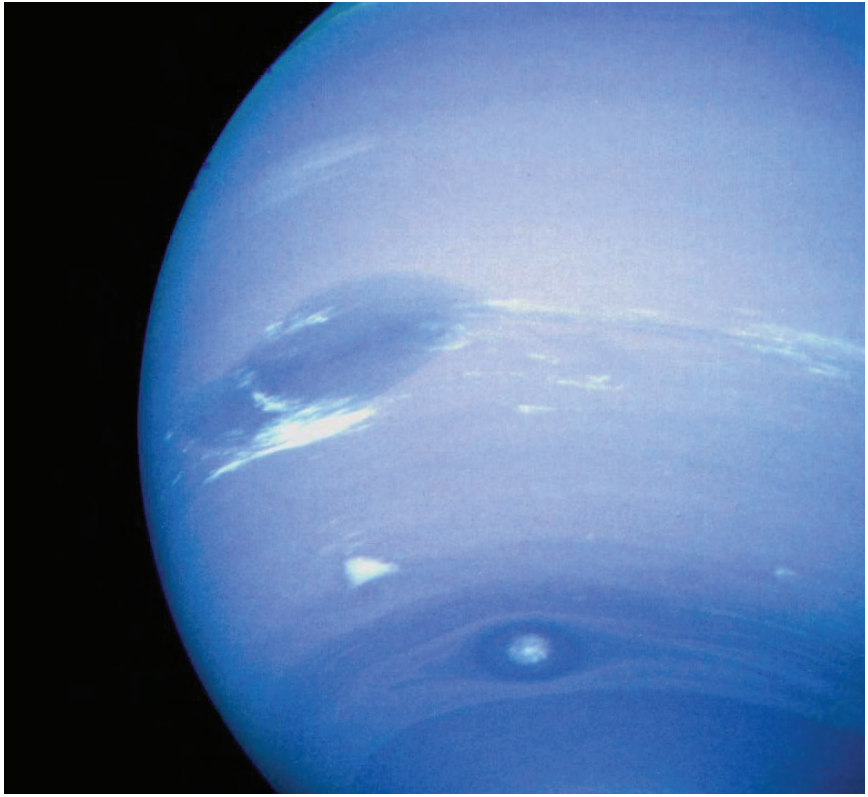

Figure 12. Using AI-accelerated molecular dynamics simulations, the researchers glimpsed possible conditions inside ice giant planets like Neptune.

Cori and *Perlmutter* Support New Understanding of Salt-Based Nuclear Reactors

THE SCIENCE. Using computing resources at NERSC, researchers at the University of Iowa and Oak Ridge National Laboratory (ORNL) have shown how electrons interact with ions of molten salts, providing insights into the processes that could occur inside salt-based nuclear reactors known as molten-salt reactors (MSRs). According to a new paper in the Journal of Physical Chemistry B, an excess electron can react with the salts to form any of three different states, or species, each with distinct properties.

The researchers simulated the introduction of an extra electron into molten zinc chloride salt. They found that the electron reacted with the molten salts to form any of three different states (part of a molecular radical including two zinc ions, localized on a single zinc ion, or diffused over multiple salt ions), each with distinct properties.

THE IMPACT. MSRs are a promising high-efficiency, inherently safer technology for electric power production. In this type of reactor, the nuclear fuel is dissolved in molten salt that also operates as the circulating heat transfer fluid. MSRs may be safer than other reactors because they operate at lower pressure and are designed to shut down into a passively safe state.

While early research indicated that the salts in MSRs are unreactive to radiation, more recent studies suggest that the ionic melts in these reactors can react with radiation-generated electrons under certain conditions. Understanding those interactions can help predict the impact of radiation on the performance of MSRs.

ADDITIONAL DETAILS. To gain insight into the interactions inside MSRs, the researchers simulated the introduction of an extra electron into molten zinc chloride salt and observed the results. They found that the electron reacted with the molten salts to form any of three different species (part of a molecular radical including two zinc ions, localized on a single zinc ion, or diffused over multiple salt ions), each with different properties.

To perform this work, the researchers used both Cori and *Perlmutter* to run ab initio molecular dynamics trajectories including a solvated electron using the Generalized Gradient Approximation with the PBE-D3 density functional and single-point energy calculations for selected configurations using the PBE0-D3 hybrid density function. The calculations were performed via a code called CP2K; the input files produced at NERSC were later used as starting points for the periodic time-dependent density functional theory calculations performed on the Cades Condo research cluster at ORNL.

Additionally, after performing simulations on the NERSC systems, the researchers used the NERSC community file system to share and transfer trajectories composed of several terabytes of data between their teams at ORNL and the University of Iowa.

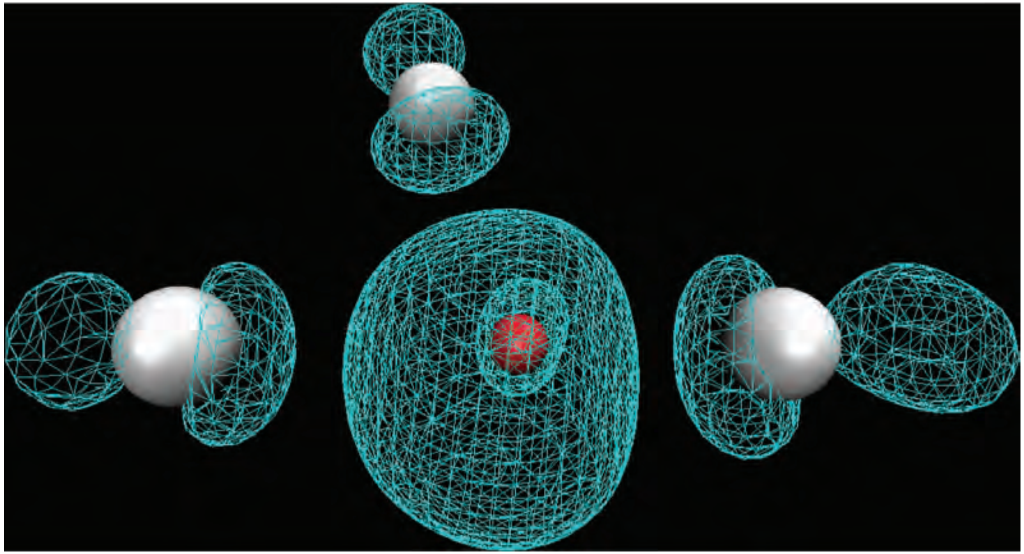

Figure 13. A simulation snapshot of a typical structure of ZnCl3•2 shows the metal ion in red and the Cl– ions in white. (Credit: Hung H. Nguyen, et al)

NERSC PI: Hung H. Nguyen, The University of Iowa

PROJECT FUNDING AND ALLOCATION AWARD: The Molten Salts in Extreme Environments (MSEE) Energy Frontier Research Center, which supported this work, is funded by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences.

PUBLICATION: Nguyen, H. H., Bryantsev, V. S., & Margulis, C. J., "Are High-Temperature Molten Salts Reactive with Excess Electrons? Case of ZnCl2," J. Phys. Chem. B 2023, 127 (42), 9155-9164.

Computational Methods Zoom in on Mechanisms Raising Hurricane Risk

THE SCIENCE. Researchers from Pacific Northwest National Lab (PNNL) used NERSC supercomputers to understand more about how climate change causes hurricanes to become stronger and strike more often on the U.S. Gulf and lower East coasts. In a paper published in Science Advances in April, the scientists describe the newly understood mechanism that fuels these stronger and more frequent hurricanes: a combination of warming sea surface temperatures and corresponding atmospheric circulation changes.

THE IMPACT. Hurricanes are among the leading cause of economic damage, and associated displacement of people, in the United States today. As the planet continues to warm, the risk of hurricanes and hurricane damage is expected to rise worldwide—but studying hurricanes at the global scale can miss important regional differences in storm behavior, and previous methods of modeling hurricanes accurately for many regions have been very computationally expensive. By developing a less resource-intensive simulation method, the researchers enable more research that will offer insight into how hurricanes will behave in a warming climate.

ADDITIONAL DETAILS. Researchers used the Risk Analysis Framework for Tropical Cyclones (RAFT), which combines physics, statistics, and machine learning to produce large ensembles of storms. RAFT is essentially a "synthetic hurricane generator" that runs in combination with climate models and makes powerful statistical analyses possible. According to the researchers in this study, RAFT allowed them to predict risk into the future more robustly with the ability to simulate many storms at a very low computational cost. Given the limited amount of observational data on hurricanes and the high computational cost of running high resolution climate models of the storms, modeling frameworks such as RAFT allow scientists to overcome those challenges.

NERSC PI: L. Ruby Leung, Pacific Northwest National Laboratory

PROJECT FUNDING AND ALLOCATION AWARD: DOE Office of Science, Office of Biological and Environmental Research

PUBLICATION: K. Balaguru, W. Xu, C. Chang, et al, "Increased U.S. coastal hurricane risk under climate change." Science Advances, 7 April 2023. Vol. 9, Issue 14. DOI: 10.1126/sciadv.adf0259

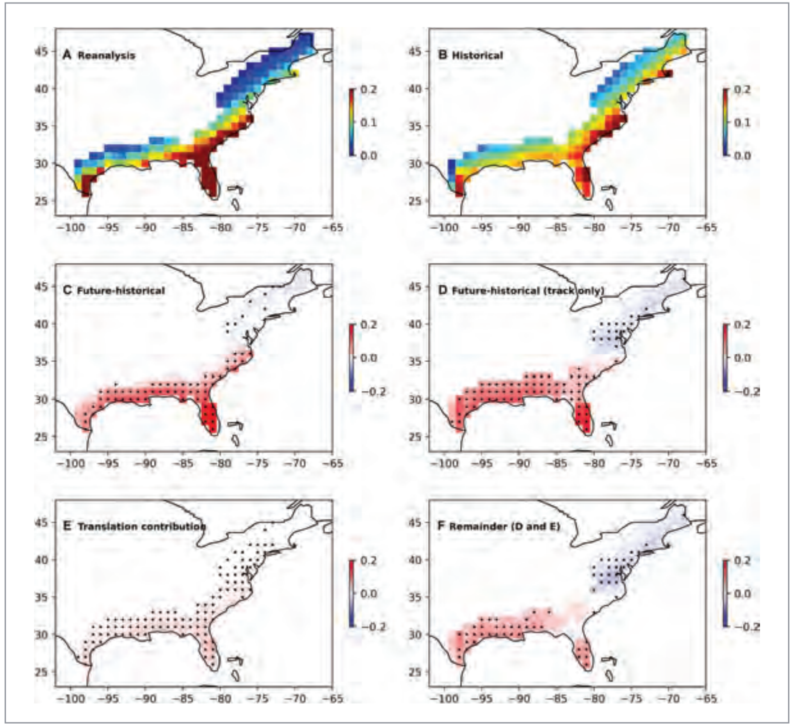

Figure 14. The researchers used the RAFT framework to model the projected change in hurricane frequency in the southeast United States under a warming climate.

The Gluonic Structure, Pressure, and Shear Distributions of the Nucleon

THE SCIENCE. Combining results achieved using NERSC resources over the last several years, a team of researchers from MIT, Fermilab, and UC Berkeley has achieved the first determination of the separate quark and gluon contributions to the gravitational form factors of two particles, the pion and the proton. These results are an important advancement in helping nuclear physicists understand the structure of particles that make up 99.9% of the mass in the visible universe.

THE IMPACT. One of the fundamental goals of nuclear physics is to understand the proton's quark and gluon structure. Protons, along with neutrons, make up the nucleus of all matter in the universe, but a complete understanding of the internal nature of their constituents—quarks and gluons—has not been achieved.

Over the last 60 years, significant progress has been made theoretically and experimentally toward understanding how the proton responds to electromagnetic and weak nuclear forces. The details are encoded in quantities known as form factors, now known with high precision. However, there are additional form factors – gravitational form factors – that are not well determined, yet are critical to understanding the nature of quarks and gluons. Calculating these form factors is exceptionally computationally demanding, and the results of these studies provide the first results for separate quark and gluon contributions to the form factors. The gravitational form factors that the team computed also determine the energy, pressure, and shear distributions within the proton. These results are timely given the first recent experimental determination of the quark and gluon form factor most relevant to pressure. The results set important benchmarks for the Electron-Ion Collider program, whose primary mission is to experimentally study gluons in the proton and other states.

ADDITIONAL DETAILS. This calculation was undertaken within the lattice Quantum Chromodynamics (LQCD) framework, the only known first-principles approach to studying QCD in the nonperturbative regime. The team used the software known as Chroma, a program for performing LQCD calculations. The QUDA library provides GPU acceleration, which is highly optimized for performing lattice QCD calculations on GPUs. The Chroma developers collaborated with the NERSC NESAP program and the DOE Exascale Computing Project to develop and optimize the code for the latest generation of supercomputers. The teams used hundreds of thousands of GPU hours on *Perlmutter* computing to obtain their results.

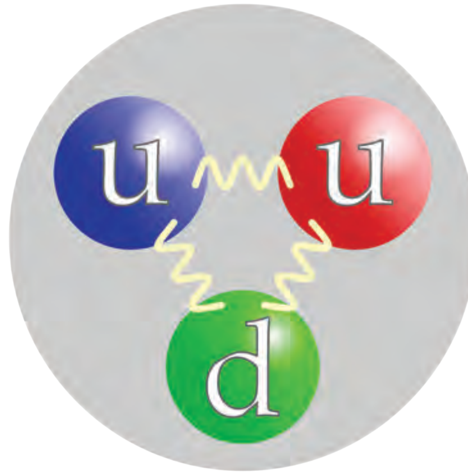

Figure 15. Protons comprise quarks (u,d) and gluons that bind them together. The way quarks and gluons interact and how those interactions lead to the observable features of the proton are not completely understood.

NERSC PI: Phiala Shanahan, Massachusetts Institute of Technology

PROJECT FUNDING AND ALLOCATION AWARD: DOE Office of Science, Office of Nuclear Physics

PUBLICATION: Hackett, Daniel C. and Oare, Patrick R. and Pefkou, Dimitra A. and Shanahan, Phiala E, "Gravitational form factors of the pion from lattice QCD." Physical Review D, 14 December 2023, 108, 114504. DOI: 10.1103/PhysRevD.108.114504; Daniel C. Hackett and Dimitra A. Pefkou and Phiala E. Shanahan, "Gravitational form factors of the proton from lattice QCD," 2003, https://arxiv.org/pdf/2310.08484.pdf

Acronyms and Abbreviations

ALS
Advanced Light Source, Lawrence Berkeley National Laboratory

AMR
Adaptive Mesh Refinement

ANL
Argonne National Laboratory

API
Application Programming Interface

ASCR
Office of Advanced Scientific Computing Research

BER
Office of Biological and Environmental Research

BES
Office of Basic Energy Sciences

BNL
Brookhaven National Laboratory

CERN
European Organization for Nuclear Research

CMB
Cosmic Microwave Background

CPU
Central Processing Unit

CSCS
Swiss National Supercomputing Centre

DESI
Dark Energy Spectroscopic Instrument

DFT
Density Functional Theory

DTN
Data Transfer Node

ECP
Exascale Computing Project

FES
Office of Fusion Energy Sciences

GB
Gigabytes

Gbps
Gigabits Per Second

GPCNet
Global Performance and Congestion Network

GPU
Graphics Processing Unit

HDF5
Hierarchical Data Format 5

HEP
Office of High Energy Physics

HPC4Mfg
High Performance Computing for Manufacturing

JGI
Joint Genome Institute

KNL
Knights Landing Processors

LANL
Los Alamos National Laboratory

LCLS
Linac Coherent Light Source

LLNL
Lawrence Livermore National Laboratory

LZ Dark Matter Experiment
LUX-Zeplin Dark Matter Experiment

MFA
Multi-Factor Authentication

MHD
Magnetohydrodynamic

NCEM
National Center for Electron Microscopy

NESAP
NERSC Exascale Scientific Application Program

NIM
NERSC Information Management

NOAA
National Oceanic and Atmospheric Administration

NP
Office of Nuclear Physics

OLCF
Oak Ridge Leadership Computing Facility

OpenMP
Open Multi-Processing

OpenMSI
Open Mass Spectrometry Imaging

PB
Petabytes

PNNL
Pacific Northwest National Laboratory

PPPL
Princeton Plasma Physics Laboratory

PUE
Power Usage Effectiveness

SENSE
Software-defined Network for End-to-End Networked Science at Exascale

SciDAC
Scientific Discovery Through Advanced Computing

SDN
Software-defined Networking

SLURM
Simple Linux Utility for Resource Management

TAP
Trusted Access Platform

TB
Terabytes

FOR MORE INFORMATION ABOUT NERSC, CONTACT:

NERSC Communications
Lawrence Berkeley National Laboratory
1 Cyclotron Road
Berkeley, CA 94720-8148
Email: CSComms@lbl.gov

NERSC's Web Site
www.nersc.gov

NERSC Annual Report Editor
Elizabeth Ball

Design
Design, layout, illustration, photography, and printing coordination:
Berkeley Lab Creative Services

Cover Image
This rendering visualizes the electronic structure of the T center, a promising quantum defect in silicon that functions as a spin-photon interface at telecom wavelengths, a key component for quantum networks. Quantum simulations performed on NERSC's *Perlmutter* supercomputer revealed the T center's distinctive defect wavefunction, an insight that can guide the design of novel quantum defects.
Credit: Weiru Chen and Yihuang Xiong, Dartmouth College

BERKELEY LAB

NERSC

24-CSAO-9079