






National Energy Research Scientific Computing Center

2012 Annual Report



National Energy Research Scientific Computing Center

2012 Annual Report

Table of Contents

4 The Year in Perspective

6 Research News

8	Automated Search Leads to Discovery of Unique Supernova
10	Large-Scale Studies Help Define the Healthy Human Microbiome
12	Modeling Feat Sheds Light on How Protein Gatekeepers Work
14	Why Onion-Like Carbons Make High-Energy Supercapacitors
16	New Computer Model Pinpoints Prime Materials for Carbon Capture
18	Protons Take Unexpected Path Through a Fuel Cell Membrane
20	Researchers Design Polymer That Filters Out Carbon Dioxide
22	Bioinorganic Compound Can Convert Sunlight to Chemical Energy
24	Rising Sea Levels Due to Global Warming Are Unstoppable
26	Thriving Tundra Bushes Add Fuel to Northern Thaw
28	Physics of Intrinsic Plasma Rotation Explained for the First Time
30	First-Ever Analysis and Visualization of a Trillion Particle Dataset
32	New Model Will Help Predict Stability of CO ₂ Reservoirs
34	Researchers Discover a New Kind of Neutrino Transformation
36	Looking for No Neutrinos
38	New Method Enables Fine Tuning of Electronic and Solar Materials
40	A New, More Efficient Approach to Water Desalination
42	Intrinsic and Engineered Piezoelectricity in 2D Materials
44	NISE Program Encourages Innovative Research
46	NERSC Users' Awards and Honors

52 The NERSC Center

54 Computational Systems and Facilities

- 54 Cray Delivers Phase 1 of “Edison” XC30 System
- 55 After Five Years, NERSC’s Franklin Retires
- 55 Breaking Ground on Computational Research and Theory Facility

57 Innovations

- 57 Adoption and Integration of Externally Developed Tools
- 58 Enabling Web-Based Access to Computation and Data
- 60 Supporting Data-Intensive Computing and New Workloads
- 64 Operational Innovations to Enhance Efficiency and Security
- 66 Understanding the Impact of Job Placement on Application Runtime Variability

66 User Support and Outreach

- 67 End-to-End Network Tuning Sends Data Screaming from NERSC to NOAA
- 69 Code and Runtime Performance Improvements
- 70 Functionality Improvements
- 71 Data Analytics Support
- 72 NERSC Participates in “Materials for Energy Applications” Workshop
- 73 Inspiring Careers in Science Research
- 73 State Department’s TechWomen Visit NERSC

75 Research and Development by NERSC Staff


80 Appendix A: NERSC User Statistic

82 Appendix B: NERSC Users Group Executive Committee

83 Appendix C: Office of Advanced Scientific Computing Research


84 Appendix D: Advanced Scientific Computing Advisory Committee

85 Appendix E: Acronyms and Abbreviations



The Year in Perspective

The National Energy Research Scientific Computing Center (NERSC) is the high-end scientific computing facility for the Department of Energy's Office of Science (DOE SC). With more than 4,500 users from universities, national laboratories and industry, we support the largest and most diverse research community of any computing facility within the DOE complex.



Our primary mission is to accelerate the pace of scientific discovery by providing high-performance computing, information, data, and communications services to the DOE Office of Science community. In 2012, our users co-authored more than 1,900 refereed papers based on the computations performed at NERSC. Computing the properties of neutrinos from the Daya Bay Neutrino Experiment led to the discovery of a new type of neutrino oscillation which may help solve the riddle of matter-antimatter asymmetry in the universe (one of Science Magazine's Top 10 Breakthroughs of 2012). NERSC's integrated resources and services enabled the earliest-ever discovery of a supernova—within hours of its explosion—providing new information about supernova explosion dynamics. Addressing critical energy and climate issues, our users achieved breakthroughs in seawater desalination methods, carbon dioxide sequestration, solar materials, lithium batteries, and the impact of climate change mitigations on sea level.

In 2012, we provided over 1.25 billion computer hours to our users. Our petascale system, Hopper, was available over 97 percent of the time and 37 percent of the simulations used more than 16,000 cores. In November we took delivery of phase 1 of the new NERSC-7 system (called Edison), a 10,624 core Cray XC30 system. When phase 2 is delivered during the second half of 2013, Edison's peak performance will be more than 2 petaflops. Perhaps even more importantly, Edison will have very high memory and bisection bandwidth. We expect our applications to perform very well on Edison because the performance of our applications is usually limited by data movement and not by peak flops.

Our users' computing and data needs will continue to increase at a very dramatic pace during the next decade. In order to even approach these needs, we must transition our users to more energy efficient architectures when we deploy NERSC-8 in late 2015. The mission need for NERSC-8 was approved in November of 2012. NERSC-8 will provide a significant increase in computational capabilities over Hopper, at least 10x on a set of representative DOE benchmarks.

In order to effectively site future NERSC systems, we began construction of the Computational Research and Theory (CRT) facility at Berkeley Lab. CRT will be a highly energy-efficient, state-of-the-art computing facility that can provide over 40 MW of power and 30,000 square feet of space for computing and storage. NERSC will be co-located with ESnet and Berkeley Lab's Computational Research Division in CRT, collaborating and sharing ideas. NERSC-8 will be deployed in CRT and Edison will move with us in early 2015.

My first year at NERSC has been very exciting and rewarding. I have enjoyed getting to know our staff, our users and our program managers. I have been very impressed by the incredible talent and dedication of the entire NERSC team, and I look forward to working with them during the next decade to enable a wide range of scientific breakthroughs by our users.

Sudip Dosanjh
NERSC Division Director





Research News

Since it was established in 1974, NERSC has been one of the nation's leading centers for computational science. In supporting the largest and most diverse research community of any computing facility within the DOE complex, NERSC provides large-scale, state-of-the-art computing for DOE'S unclassified research programs in alternative energy sources, climate change, energy efficiency, environmental science, and other fundamental science areas within the DOE mission.

Because NERSC is the DOE Office of Science's mission computing facility, DOE program managers across SC allocate most of the center's computing resources. In 2012, the six programmatic offices allocated 80 percent of NERSC computing time, while DOE's ASCR Leadership Computing Challenge allocated 10 percent. The remaining 10 percent was allocated through the NERSC Initiative for Scientific Exploration (NISE) program. In Allocation Year 2012 (January 10, 2012 to January 7, 2013), 1.25 billion NERSC computer hours were allocated in total.

NERSC's primary mission is to accelerate scientific discovery at the DOE Office of Science through high performance computing and extreme data analysis. Data resources at NERSC played a major role in one of Science Magazine's Top Ten Breakthroughs of 2012—the measurement of the Θ_{13} neutrino mixing angle. In 2012, our users co-authored more than 1,500 refereed papers based on the computations performed at NERSC, including 17 journal cover stories. Several representative examples are highlighted in this section, followed by the 2012 NISE awards and a listing of NERSC users' awards and honors.

Automated Search Leads to Discovery of Unique Supernova

First direct evidence of a Type Ia progenitor system containing a red giant star

Exploding stars called Type Ia supernova are ideal for measuring cosmic distance because they are bright enough to spot across the Universe and have relatively the same luminosity everywhere. But just because Type Ia supernovae look the same, that doesn't mean they are all born the same way.

In the August 24, 2012 issue of *Science*, the Palomar Transient Factory (PTF) collaboration published the first direct observations of a Type Ia supernova progenitor system, and the first evidence of a progenitor system containing a red giant star. They also showed that the system previously underwent at least one much smaller nova eruption before it ended its life in a destructive supernova. The system, called PTF 11kx, is located 600 million light years away in the constellation Lynx.


The PTF survey uses a robotic telescope at Palomar Observatory in southern California to scan the sky nightly. As the observations are taken, the data travels more than 400 miles via high-speed networks to NERSC. There, the Real-Time Transient Detection Pipeline uses supercomputers, a high-speed parallel file system, and sophisticated machine learning algorithms to automatically sift through 600 images per night and identify events for scientists to follow up on.

The pipeline detected PTF 11kx on January 16, 2011, and researchers immediately followed up with spectroscopy observations from the University of California's Lick Observatory, which revealed unusually strong calcium signals in the gas and dust surrounding the supernova. The signals were so peculiar that the team triggered a Target of Opportunity observation using the Keck Telescope in Hawaii.

From the Keck observations, astronomers noticed that the clouds of gas and dust surrounding PTF 11kx were moving too slowly to be coming from the recent supernova, but too quickly to be stellar wind. They suspected that the star had previously erupted as a nova, propelling a shell of material outwards. The material, they surmised, must be slowing down as it collided with wind from a nearby red giant star. But for this theory to be true, the material from the recent supernova should eventually catch up and collide with gas and dust from the previous nova. And that's exactly what they observed 58 days after the supernova explosion.

It's not unusual for a white dwarf star orbiting a red giant to undergo nova eruptions more than once. But astronomers previously believed that Type Ia supernovae would not occur in recurring nova systems, because the white dwarf would lose too much mass in the nova eruptions to ever become a supernova. PTF 11 kx gave the first observational evidence that a star can explode first as a nova, then as a Type Ia supernova.

Indirect observations in 2011 of another Type Ia supernova progenitor system, called SN 2011fe, showed no evidence of a red giant star. Taken together, these observations show that just because Type Ia supernovae look the same, that doesn't mean they are all born the same way.





// In the left photo, the supernova PTF 11kx can be seen as the blue dot on the galaxy. The image was taken when the supernova was near maximum brightness by the Faulkes Telescope North. (*B. J. Fulton, Las Cumbres Observatory Global Telescope Network*) On the right is an artist's rendition of the PTF 11kx by Romano Corradi and the Instituto de Astrofísica de Canarias.

PROJECT

Palomar Transient Factory

PROJECT LEADER

Peter Nugent, Lawrence Berkeley National Laboratory

NERSC RESOURCES

Caryer, DeepSky Science Gateway (154 TB of data imported per year)

DOE OFFICE

High Energy Physics

Full story: <https://www.nersc.gov/news-publications/news/science-news/2012/ptf11kx/>

Publication: B. Dilday et al., "PTF 11kx: A Type Ia Supernova with a Symbiotic Nova Progenitor," *Science* **337**, 942–945 (2012), doi:10.1126/science.1219164.

Large-Scale Studies Help Define the Healthy Human Microbiome

Two million computer hours to analyze 3.5 terabases of genome sequence data

You're outnumbered. There are ten times as many microbial cells in your body as there are human cells.

The human microbiome—as scientists call the communities of microorganisms that inhabit your skin, mouth, gut, and other parts of your body by the trillions—plays a fundamental role in keeping you healthy. These communities are also thought to cause disease when they're perturbed. But our microbiome's exact function, good and bad, is poorly understood. That could all change now that the normal microbial makeup of healthy humans has been mapped for the first time.

The research will help scientists understand how our microbiome carries out vital tasks such as supporting our immune system and helping us digest food. It will also shed light on our microbiome's role in diseases such as ulcerative colitis, Crohn's disease, and psoriasis.

Five years of research by the Human Microbiome Project (HMP) Consortium, organized by the National Institutes of Health (NIH), resulted in several major reports published June 14, 2012 in *Nature* and in journals of the Public Library of Science. The HMP Consortium includes about 200 researchers from nearly 80 research institutions, including the DOE Joint Genome Institute (JGI), Lawrence Berkeley National Laboratory, and NERSC.

HMP researchers sampled 242 healthy volunteers up to three times over 22 months, collecting specimens from 15 body sites in men and 18 body sites in women—sites such as the mouth, nose, skin, and lower intestine. The microbial communities in each body site can be as different as the microbes in the Amazon rainforest versus the Sahara desert.

Researchers then purified all human and microbial DNA in more than 5,000 samples and ran them through DNA sequencing machines. The result is about 3.5 terabases of genome sequence data. A terabase is one trillion subunits of DNA.

Two million computer hours were allocated at NERSC to carry out HMP data integration as well as sift through HMP data for 16S ribosomal RNA genes, which can be used to identify individual species. Focusing on this microbial signature allowed HMP researchers to subtract the human genome sequences and analyze only bacterial DNA.

The analysis helped scientists determine the wide diversity of microbial species within a person, including within different body sites in a person. It also revealed the extent to which microbial communities vary widely between people. While scientists had previously isolated only a few hundred bacterial species from the body, HMP researchers now calculate that more than 10,000 species occupy the human ecosystem. We owe much of our biology and our individuality to the microbes that live on and in our bodies.

PROJECT

Microbial Genome and
Metagenome Data
Processing and Analysis

PROJECT LEADER


Victor Markowitz,
Lawrence Berkeley
National Laboratory

NERSC RESOURCES

2 million hours on Carver,
Hopper, Genepool

DOE OFFICE

Biological and
Environmental Research



// Relative abundances of 11 distinct *Streptococcus* species in tongue samples from 90 subjects. Individual species within a single habitat demonstrate a wide range of compositional variation. Inset illustrates average tongue sample composition. (HMP Consortium)

Full story: <https://www.nersc.gov/news-publications/news/science-news/2012/scientists-help-define-the-healthy-human-microbiome/>

Publication: The Human Microbiome Project Consortium, "Structure, function and diversity of the healthy human microbiome," *Nature* **486**, 207–214 (2012), doi:10.1038/nature11234.

Modeling Feat Sheds Light on How Protein Gatekeepers Work

Long-timescale simulation reconciles conflicting experimental results

Using NERSC supercomputers, chemists have managed for the first time to simulate the biological function of a channel called the Sec translocon, which allows specific proteins to pass through membranes. The feat required bridging timescales from the realm of nanoseconds all the way up to full minutes, exceeding the scope of earlier simulation efforts by more than six orders of magnitude. The result is a detailed molecular understanding of how the translocon works.


The Sec translocon is a channel in cellular membranes involved in the targeting and delivery of newly made proteins. Such channels are needed because the proteins that are synthesized at ribosomes must travel to other regions of the cell or outside the cell in order to perform their functions; however, the cellular membranes prevent even the smallest of molecules, including water, from passing through them willy-nilly. In many ways, channels such as the Sec translocon serve as gatekeepers—once the Sec translocon determines that a given protein should be allowed to pass through, it opens up and allows the protein to do one of two things: to be integrated into the membrane, or to be secreted completely out of the cell.


Scientists have disagreed about how the fate of a given protein entering the translocon is determined. Based on experimental evidence, some have argued that a protein's amino-acid sequence is what matters—that is, how many of its amino acids interact favorably with water and how many clash. This argument treats the process as one in equilibrium, where the extremely slow rate at which a ribosome adds proteins to the channel can be considered infinitely slow. Other researchers have shown that slowing down the rate of protein insertion into the channel actually changes the outcome, suggesting that kinetic effects can also play a role.

The long-timescale simulations developed by Thomas Miller and Bin Zhang of the California Institute of Technology revealed that, in fact, both equilibrium and kinetically controlled processes are happening—but in a way that was not obvious until researchers could see everything working together.

Beyond elucidating how the translocon works and reconciling seemingly disparate experimental results, the new simulation also lets the researchers perform experiments computationally that have yet to be tried in the lab. For example, they have run simulations with longer proteins and observed that at such lengths—unlike what has been seen with shorter proteins—the equilibrium picture begins to be affected by kinetic effects.

"This could bring the two experimental camps together, and to have led that would be kind of exciting," Miller says.





// The ribosome (red-blue) in complex with the translocon channel (green), which is embedded in the cell membrane (yellow, white). Proteins that are inserted via the ribosome into the channel can either be laterally integrated into the cell membrane or secreted across the cell membrane (inset). (Bin Zhang and Thomas Miller)

PROJECT

Sampling Diffusive Dynamics on Long Timescales, and Simulating the Coupled Dynamics of Electrons and Nuclei

PROJECT LEADER

Thomas Miller, California Institute of Technology

NERSC RESOURCES

28 million hours on Franklin, Hopper, and Carver

OTHER COMPUTING RESOURCES

PNSF, Pittsburgh Supercomputing Center

DOE OFFICE

Basic Energy Sciences

Full story: <http://www.nersc.gov/news-publications/news/science-news/2012/modeling-feat-sheds-light-on-protein-channel-s-function/>

Publication: Bin Zhang and Thomas F. Miller, "Long-Timescale Dynamics and Regulation of Sec-Facilitated Protein Translocation," *Cell Reports* **2**, 927–937 (2012), doi:10.1016/j.celrep.2012.08.039.

Why Onion-Like Carbons Make High-Energy Supercapacitors

Simulations explain experimental results for electrical storage devices

The two most important electrical storage technologies are batteries and capacitors. Batteries can store a lot of energy, but have slow charge and discharge rates. Capacitors generally store less energy but have very fast (nearly instant) charge and discharge rates, and last longer than rechargeable batteries. Developing technologies that combine the optimal characteristics of both will require a detailed understanding of how these devices work at the molecular level.


Electrical double-layer (EDL) capacitors, also called supercapacitors or ultracapacitors, offer the usual fast charging and discharging rates of conventional capacitors but also have higher power density, high capacitance, and excellent durability, thus bridging the gap between batteries and conventional capacitors. Commercial supercapacitors are currently used to power electric vehicles, portable electronic equipment, and other devices.

Carbon-based materials are the most widely used electrodes for supercapacitors, and researchers looking for ways to improve supercapacitor technology are especially interested in graphene because of its unique electrical, thermal, mechanical, and chemical properties. But given that graphene can form stripes, tubes, balls, ribbons, and other useful shapes, an important question is how shape affects EDL electrical properties.

Experiments have shown that onion-like carbons (OLCs), which consist of concentric graphene spheres, offer ultrahigh energy density and charging/discharging rates in supercapacitors, but the experiments did not reveal the physical origin of this phenomenon. Researchers in the DOE's Fluid Interface Reactions, Structures and Transport (FIRST) Energy Frontier Research Center designed molecular dynamics simulations, using a variety of software made available by NERSC staff, to explain the relationship between capacitance and electrode potential in supercapacitors consisting of an OLC electrode suspended in a room-temperature ionic liquid.

By varying the radius of the OLC spheres, they were able to understand the influence of electrode curvature and size. Simulations showed that the surface charge density in OLCs increases almost linearly with the potential applied at the OLC's electric double layer. This leads to a nearly flat differential capacitance-versus-potential curve—unlike the bell or camel shaped curves observed for planar electrodes—and could potentially inspire the design of supercapacitors with much more stable capacitive performance. The simulations also explained why the capacitance of the OLC increases as the size, and thus curvature, of the OLC decreases. More compact or more uniform onions and better electrolytes could boost performance, especially for applications that require large bursts of power, long lifetimes, and high storage capacities.





// Capacitance and geometry effects revealed by molecular dynamics simulations. The OLC and the ionic liquid that were the basis of the simulation are shown in the lower left. (Guang Feng, De-en Jiang, Peter T. Cummings, © ACS Publications)

PROJECT

Computational Resources for the FIRST (Fluid Interface Reactions, Structures and Transport) EFR Center

PROJECT LEADER

Peter Cummings, Vanderbilt University

NERSC RESOURCES

1.7 million hours on Franklin and Hopper

OTHER COMPUTING RESOURCES

Clemson University

DOE OFFICE

Basic Energy Sciences

Publication: Guang Feng, De-en Jiang, and Peter T. Cummings, "Curvature Effect on the Capacitance of Electric Double Layers at Ionic Liquid/Onion-Like Carbon Interfaces," *Journal of Chemical Theory and Computation* **8**, 1058–1063 (2012), doi:10.1021/ct200914j.

New Computer Model Pinpoints Prime Materials for Carbon Capture

Millions of solid, porous materials can be screened for optimum performance

Approximately 45 percent of electricity used in the United States is produced by coal-burning power plants that spew carbon dioxide (CO_2) into the atmosphere and contribute to global warming. This effect could be mitigated by capturing CO_2 from power plant flue gas before it reaches the atmosphere, but doing this with current technologies would substantially drive up the cost of electricity. Dubbed “energy parasites,” these methods use about one-third of the total energy generated by the plant.

But less energy expensive technologies are on the horizon, say researchers at the University of California, Berkeley. Using resources at NERSC, the team developed a new computer model that suggests solid porous materials like zeolites and metal oxide frameworks (MOFs) could be more efficient for capturing CO_2 from plant emissions.


“The current on-the-shelf process of carbon capture has problems, including environmental ones, if you do it on a large scale,” said Berend Smit, a professor at UC Berkeley, senior scientist in Berkeley Lab’s Materials Sciences Division, and director of the DOE Energy Frontier Research Center for Gas Separations Relevant to Clean Air Technologies. “Our calculations show that we can reduce the parasitic energy costs of carbon capture by 30 percent with porous materials like MOFs and zeolites, which should encourage the industry and academics to look at them.”


Although the number of possible zeolite structures has been estimated to be more than 2.5 million, only about 190 structures have been synthesized to date. The main goal of this project is to identify structures with optimal performance for CO_2 separation. This is being done by performing efficient screening of material databases with a combination of molecular simulation and chemical informatics techniques.

Smit and his team worked with electric power industry scientists to establish the best criteria for a good carbon capture material, focusing on the energy costs of capture, release, and compression, and then developed a computer model to calculate this energy consumption for any material. Smit then obtained a database of 4 million zeolite structures compiled by Rice University scientists and ran the structures through his model. The team also computed the energy efficiency of 10,000 MOF structures, which are composites of metals with organic compounds that together form a porous structure.

“The surprise was that we found many materials, some already known but others hypothetical, that could be synthesized and work more energy efficiently than [current technology],” said Smit. “This best part about this model is that it will work for structures other than zeolites and MOFs.”

By incorporating this model into a Carbon Capture Materials Database, the team is enabling any researcher in the world to upload the structure of a proposed material to this site and calculate whether it offers improved performance over the energy consumption figures of today’s best technology for removing carbon.






PROJECT
Computational
Characterization
of Porous Materials


PROJECT LEADER
Berend Smit, Lawrence
Berkeley National
Laboratory and University
of California, Berkeley

NERSC RESOURCES
2.5 million hours on Dirac,
Franklin, and Hopper

DOE OFFICE
Basic Energy Sciences



// One of the 50 best zeolite structures for capturing carbon dioxide. Zeolite is a porous solid made of silicon dioxide, or quartz. The red balls are oxygen, tan balls are silicon. The blue-green area is where carbon dioxide prefers to nestle when it adsorbs.
(Berend Smit laboratory)



Full story: <https://www.nerisc.gov/news-publications/news/science-news/2012/new-computer-model-pinpoints-prime-materials-for-carbon-capture/>

Publication: Li-Chiang Lin et al., "In silico screening of carbon-capture materials," *Nature Materials* **11**, 633–641 (2012), doi:10.1038/nmat3336.

Protons Take Unexpected Path Through a Fuel Cell Membrane

Shunning liquid highway, protons prefer scenic route to more direct path

Meandering along the water-filled highways laced through an alternative membrane for fuel cells, protons visit sulfur and oxygen atom clusters along the way, according to scientists at Pacific Northwest National Laboratory (PNNL). The highways are lined with identical hook-shaped molecules. Each channel contains a shell that keeps the water in, formed by the molecules' backbones. The small clusters, known as sulfonate groups, dangle inside the channel. In the middle is a thin "wire" of water.

"Conventional wisdom was that the protons would zip along the center of the water channel, where there were no sulfonate groups to slow them down," said Ram Devanathan, a materials scientist at PNNL who led the study. "What we found was that the proton doesn't go through the highway, but rather goes on sightseeing trips mainly along the sulfonate groups."

Understanding how protons are transported through membranes could answer one of the key problems of adopting fuel cells. With the potential to revolutionize transportation and portable power, as well as bring electricity to remote locations, fuel cells are promising energy conversion devices. But, the cells are costly. If a membrane could be designed that withstands higher temperatures and effectively allows protons to pass, fuel cell designers could replace the platinum with a less costly material, which would drive down the costs.

"Whether you are using membranes for separation or water purification, you are interested in selective transport through the membrane," said Michel Dupuis, a theoretical chemist who worked on the study. "Specifically, we are interested in how changing the design would affect proton transport through the membrane."

Using supercomputers at NERSC and PNNL's Environmental Molecular Sciences Laboratory (EMSL), the team ran simulations of thousands of molecules of the short side chain perfluorosulfonic acid membrane. They wanted to determine what would happen if they changed different aspects of the molecules used to construct the membrane. The team ran simulations where they changed different aspects of the molecule, including the length of the side chain that dangles the sulfonate group into the water wire.

Conventional wisdom said that making the side chains shorter would make the central water channel larger and give the protons more room to move, allowing them to go faster. Surprisingly, the team found that the chain length did not matter. When the chain is long, it folds back and doesn't block the channel. When the chain is short, it just stays out of the way.

The team's results compared well with experimental data collected on the perfluorosulfonic acid membrane. Now, Devanathan and Dupuis are conducting detailed studies to determine what influences protons bouncing between the sulfonate groups. They are also examining how to transfer the protons effectively in the absence of water, perhaps with an ionic liquid.





// The cover represents the environment around the side chain. The right side is the water network that exists between the sulfonate groups shown in yellow. The left side is the short chain with the sulfonate group. (Ram Devanathan and Michel Dupuis, © the Owner Societies)

PROJECT

Molecular Transport in Ionic Liquids, Electrolytes, and Polymers

PROJECT LEADER

Ram Devanathan,
Pacific Northwest
National Laboratory

NERSC RESOURCES

Franklin and Hopper

OTHER COMPUTING RESOURCES

EMSL

DOE OFFICE

Basic Energy Sciences

Full story: <https://www.nerisc.gov/news-publications/news/science-news/2012/the-path-a-proton-takes-through-a-fuel-cell-membrane/>

Publication: R. Devanathan and M. Dupuis, "Insight from molecular modelling: does the polymer side chain length matter for transport properties of perfluorosulfonic acid membranes?" *Physical Chemistry Chemical Physics* **14**, 11281–11295 (2012), doi:10.1039/C2CP24132C.

Researchers Design Polymer That Filters Out Carbon Dioxide

May be 100 times more effective than the best existing materials

Carbon dioxide gas separation is important for many environmental and energy applications, especially for removing carbon dioxide from power plant exhausts, a process which is currently prohibitively expensive. But now, based on simulations using supercomputers at NERSC, researchers at Haverford College have come up with a new type of two-dimensional polymer, PG-ES1, which allows, in theory, for highly efficient and economical separation of carbon dioxide.

PG-ES1 is predicted to be more than 100 times more permeable to carbon dioxide than the best existing materials, while maintaining a rejection of nitrogen and methane gases that meets or exceeds the best existing materials. This allows it to act as a molecular filter that lets the carbon dioxide pass through easily, while preventing other gases from escaping.

The key to the new process is to utilize both the preferential adsorption of carbon dioxide gas molecules on the surface and the ability to create small, nanometer-sized pores in the surface.


"Nitrogen and carbon dioxide are linear molecules, and the holes are too small to allow them to enter in any way other than along their skinniest dimensions," says Joshua Schrier, Assistant Professor of Chemistry at Haverford College.

"As it turns out, carbon dioxide is a little skinnier than nitrogen, which allows it to pass through the hole more readily. Although it is unlikely that a random molecule would have the correct orientation, the surface adsorption helps increase the local concentration of carbon dioxide and allows each carbon dioxide molecule to try several attempts at different orientations until it finds the correct one, which stacks the deck in favor of carbon dioxide passage. Nobody has previously considered the role of surface adsorption on the barrier crossing process, but it is absolutely crucial for performing this type of separation."

Using NERSC's Hopper system, Schrier performed detailed molecular simulations which followed the paths of the molecules under various conditions and tested their interactions with different types of materials. He was able to do this thanks to the support of a NERSC Initiative for Scientific Exploration (NISE) award of 500,000 processor hours, which doubled his regular allocation from the DOE's Chemical Sciences program in the Office of Basic Energy Sciences.

NERSC staff members Zhengji Zhao and Brian Austin also played critical roles in this work. Zhao made some crucial modifications to the LAMMPS molecular dynamics code, enabling it to treat the quadrupolar interactions of nitrogen and oxygen. Austin improved the performance of the code by a factor of three by finding ways to distribute the computational work more efficiently across processors.

Schrier is now working with Zhao and Austin to build on this work with PG-ES1 and similar polymers to simulate helium separation from natural gas, as well as paraffin and olefin separation. Other Haverford researchers are synthesizing the building blocks needed to construct PG-ES1, and several of them have filed for a provisional patent on the polymer.



// Molecular dynamics simulations were used to characterize a two-dimensional hydrocarbon polymer, PG-ES1, that uses a combination of surface adsorption and narrow pores to separate carbon dioxide from nitrogen, oxygen, and methane gases. (Joshua Schrier)

PROJECT

Two-Dimensional Membranes for Energy Efficient Gas Separation: Computational Studies of Molecular Ion Beam and Polymeric Synthesis Approaches

PROJECT LEADER

Joshua Schrier, Haverford College

NERSC RESOURCES

1 million hours on Hopper

DOE OFFICE

Basic Energy Sciences

Photo credit: Ting XU group at UC Berkeley and MSD-LBNL

Full story: <https://www.nerisc.gov/news-publications/news/science-news/2012/haverford-college-researchers-create-carbon-dioxide-separating-polymer/>

Publication: Joshua Schrier, "Carbon Dioxide Separation with a Two-Dimensional Polymer Membrane," ACS Applied Materials & Interfaces **4**, 3745–3752 (2012), doi:10.1021/am300867d.

Bioinorganic Compound Can Convert Sunlight to Chemical Energy

Commercialization depends on getting the molecules to stop wobbling

Rare-earth elements in traditional photovoltaic cells raise their cost and keep solar power from being competitive with fossil fuels. But if we learn from plants, which use only common elements—hydrogen, nitrogen, carbon, oxygen, and some others—to convert sunlight into energy, then we'll be able to bring down the cost of solar power. This is why researchers are looking at bio-inspired materials as possible resources for solar energy.

In the 1990s, an Arizona State University research group led by Devens Gust made a huge advance in this field by creating the carotenoid-porphyrin-C60 molecular triad, a novel material that converts sunlight into chemical energy by mimicking photosynthesis. However, the material has been difficult to commercialize because it can only be controlled in experimental labs.


Using 2 million computer hours at NERSC and 2.5 million computer hours at the Texas Advanced Computing Center (TACC), Margaret Cheung, Assistant Professor of Physics at the University of Houston, and her research team explored the role that confinement, temperature, and solvents play in the stability and energy efficiency of the light-harvesting triad. Their results provide a way to test, tailor, and engineer nano-capsules with embedded triads that, when combined in large numbers, could greatly increase the ability to produce clean energy.

Unlike solid-state photovoltaic cells, the carotenoid-porphyrin-C60 triad is a bioinorganic compound, combining biological and inorganic components: a carotenoid (an organic pigment, similar to the chromophore in plants); a fullerene or buckyball (a carbon-based molecule that forms a hollow sphere); and a porphyrin (an organic compound that can bind ligands to metals, as in hemoglobin).

When photons hit the triad, the molecule becomes excited, driving the electrons into a polarized distribution, like a dipole. This separation of positive and negative charges in the system becomes the stored chemical potential from which energy can be produced. The problem is that bioinorganic compounds are flexible and don't want to stay in a fixed configuration. "If the vehicle that carries the charge-separated state is wobbling all the time, then it's not very reliable," Cheung said.

The wobbliness of the triad also challenged efforts to simulate its dynamics in the past. Cheung had to pioneer new methods that combine quantum chemistry approaches, molecular dynamics simulations, and statistical physics to take into account the microscopic landscape of the molecules and the many configurations that the triad might be in when photons hit the material.

Cheung and her team simulated the triad in solution at many different temperatures and confinement conditions to map the impact of these changes on the behavior of the molecule. They discovered that the triad conformation distribution could be manipulated by temperature fluctuations in the solvent and by confinement. Their goal is to use information from the computer simulations to design a scalable system that maximizes the generation of chemical energy while maintaining the triad's stability.



// Snapshot of an all-atomistic representation of the triad solvated in a spherocylindrical confinement. Hydrogen atoms are white, carbon atoms are cyan, nitrogen are blue, and oxygen are red. The carbon atoms of porphyrin are colored purple for visual guidance.
(Margaret Cheung)

PROJECT

Solar Energy to Chemical Fuel: Design of Efficient Artificial Photosynthetic Materials

PROJECT LEADER

Margaret Cheung,
University of Houston

NERSC RESOURCES

1 MILLION HOURS ON
Hopper, Carver

OTHER COMPUTING RESOURCES

Texas Advanced
Computing Center,
University of Houston,
Rice University

DOE OFFICE

Basic Energy Sciences

Full story: <https://www.nersc.gov/news-publications/news/science-news/2012/learning-from-photosynthesis-to-create-electricity/>

Publication: Guoxiong Su, Arkadiusz Czader, Dirar Homouz, Gabriela Bernardes, Sana Mateen, and Margaret S. Cheung, "Multiscale Simulation on a Light-Harvesting Molecular Triad," *Journal of Physical Chemistry B* **116**, 8460–8473 (2012), doi:10.1021/jp212273n.

Rising Sea Levels Due to Global Warming Are Unstoppable

Mitigation can slow down but not prevent sea level rise for centuries to come

A reduction in greenhouse gas emissions could greatly lessen the impacts of climate change. However, the gases already added to the atmosphere ensure a certain amount of sea level rise to come, even if future emissions are reduced. A study by National Center for Atmospheric Research (NCAR) scientist Gerald Meehl and colleagues quantifies the impact on oceans of the “climate commitment” being made now by human activity.

As the atmosphere heats up, the world's oceans absorb heat more slowly, due to the physical properties of sea water, the oceans' vast depth, and other characteristics. So there is lag time as the heat makes its way deeper and deeper into the oceans. This ever-increasing volume of warm water slowly expands—a process that occurs over centuries.


Thermal expansion accounts for about 0.8 millimeters per year of sea level rise since the 1970s, almost half of the estimated total sea level rise per year of 1.8 millimeters (0.075 inch). The other contributions come from melting glaciers and ice sheets, as well as other sources such as piping groundwater to the surface. If the most stringent measures to slow warming were put in place now, thermal expansion alone would still cause substantial sea level rise well into the future, the study finds.


Meehl's team analyzed future warming and potential sea level rise under four mitigation scenarios developed for the next report of the Intergovernmental Panel on Climate Change. In the most aggressive scenario, the world's energy needs are met largely through renewable energy, nuclear power, and carbon sequestration—a combination that is technologically feasible, though difficult to enact. Under this scenario, there could be “negative emissions” of carbon dioxide by 2070, meaning that more CO₂ would be taken out of the atmosphere than emitted.

Even under such aggressive measures, which could result in global cooling by 2100, sea level still continues to rise for centuries. In this scenario, thermal expansion alone would push sea level beyond its 1986–2005 average by 14.2 centimeters (5.5 inches) in the year 2100 and by 24.2 centimeters (9.4 inches) in 2300. These values do not take into account additional sea level rise from melting of glaciers and ice sheets. In comparison, the scenario with least mitigation produces sea level rise of 32.3 cm (12.6 in) by 2100 to 139.4 cm (54.4 in) by 2300 from thermal expansion alone.

The wild card for future sea level rise is uncertainty about the contributions from the Greenland and Antarctic ice sheets, says Meehl. How fast the ice sheets could disintegrate has proven difficult to model, making it challenging for researchers to precisely estimate the total rise in sea level.

According to Meehl, the findings should not dampen efforts to rein in greenhouse gases but should revitalize those efforts. Under the most aggressive scenario, he points out, “Even though we're already committed to a certain amount of sea level rise, it's not going to go up nearly as much as if we don't do anything.”





// Sea level rise from thermal expansion alone with moderate mitigation (dashed green line), including already committed expansion (red line). Shading indicates uncertainty in the projections, with lighter shading becoming less certain. (NCAR)

PROJECT

Climate Change
Simulations with CESM:
Moderate and High
Resolution Studies

PROJECT LEADER

Warren Washington,
National Center for
Atmospheric Research

NERSC RESOURCES

Hopper, 34 million hours

OTHER COMPUTING RESOURCES

NCAR, Oak Ridge
Leadership Computing
Facility

DOE OFFICE

Biological and
Environmental Research

Publication: G. A. Meehl, A. Hu, C. Tebaldi, J. M. Arblaster, W. M. Washington, H. Teng, B. Sanderson, T. Ault, W. G. Strand, and J. B. White III, "Relative Outcomes of Climate Change Mitigation Related to Global Temperature versus Sea Level Rise," *Nature Climate Change* **2**, 576–580 (2012), doi:10.1038/nclimate1529.

Thriving Tundra Bushes Add Fuel to Northern Thaw

Spread of taller vegetation could exacerbate warming in northern latitudes

Carbon-gobbling plants are normally allies in the fight to slow climate change, but in the frozen north, the effects of thriving vegetation may actually push temperatures higher. In a series of climate simulations performed at NERSC, a group of researchers found that the spread of bushes, taller ones especially, could exacerbate warming in northern latitudes by anywhere from 0.6°C to 1.8°C per year.

What's more, taller species have the potential to warm tundra soil more deeply, threatening to thaw permafrost in some areas. That means more of the greenhouse gases now locked out of the atmosphere by a year-round freeze could be released into the atmosphere, increasing warming even more.

"Until now, most climate model studies have only focused on the climate effects induced by a complete tundra-to-forest conversion," said Celine Bonfils, a scientist at Lawrence Livermore National Laboratory (LLNL) and principal investigator in this study. While warming could eventually lead to northward forest expansion, "we don't expect a full-scale conversion of tundra to forest anytime in the next century," said Tom Phillips, one of Bonfils's collaborators on the study, also of LLNL. "More likely you'll see shorter shrubs or new species that are taller moving in gradually."

That is, in fact, what has been happening to some tundra regions. The authors point out that one study found a 1.2 percent increase per year in coverage in Alaska.


To better understand the more likely short-term fate of expanding and invading shrubs, Bonfils and colleagues simulated two, more realistic, although still idealized, scenarios: the expansion of the short shrubs that already grow on the tundra, and the invasion of taller shrubs. In both cases the team found that temperatures rose, as expected based on earlier simulations involving forest expansion, but the rise caused by taller shrubs was striking, says Phillips, a coauthor of the study.

"The physics of what's happening here is pretty straightforward," says Phillips. "Shrubs tend to darken the land surface, so it absorbs more sunlight, warming the ground and increasing evaporation." Plants also suck water from the soil and transpire it into the air, a process called evapotranspiration. That water vapor, in turn, acts as a greenhouse gas, trapping heat, raising temperatures, and creating more favorable growing conditions.

"These runs, performed at NERSC, show for the first time that the strength and timing of these two mechanisms greatly depends on the height of the shrubs, and the time at which branches and leaves protrude above the snow," explained Bonfils. Taller shrubs darken the ground earlier in the spring and transpire more efficiently than shorter shrubs, thereby increasing soil warming and making the permafrost less stable.

"These simulations mainly focus on the biophysical effects of shrub height and expansion, but they did not include the biogeochemical effects. In other words, the plants were not allowed to gobble carbon as they grow," added Bonfils. "In expanding and growing taller, the plants will remove some greenhouse gases from the atmosphere, but probably not enough to negate the warming and permafrost destabilization induced through reduced reflection of sunlight and increased transpiration. New climate simulations would be needed to verify this hypothesis."





// Left: Current deciduous shrub distribution. Right: Simulation of bare ground converted to deciduous shrubs. (Celine Bonfils)

PROJECT

Investigation of the Magnitudes and Probabilities of Abrupt Climate Transitions (IMPACTS)

PROJECT LEADER

William Collins, Lawrence Berkeley National Laboratory

NERSC RESOURCES

Franklin

DOE OFFICE

Biological and Environmental Research

Publication: C.J.W. Bonfils et al., "On the influence of shrub height and expansion on northern high latitude climate," *Environmental Research Letters* **7**, 015503 (2012), doi:10.1088/1748-9326/7/1/015503

Physics of Intrinsic Plasma Rotation Explained for the First Time

Key understanding for modeling future fusion reactors such as ITER

The quality of a fusion reaction is determined by plasma confinement at the edge, which is not yet completely understood due to the complicated interactions between multiscale physics. The SciDAC Center for Edge Physics Simulation (EPSI) uses large-scale simulation to understand the edge physics from first-principles equations, and ultimately to provide predictions of fusion performance.

This work is relevant to existing magnetic fusion experiments and essential for next-generation burning plasma experiments such as ITER.

The plasma edge presents a set of multi-physics, multi-scale problems involving complex 3D magnetic geometry. Perhaps the greatest computational challenge is the overlap of multiple temporal and spatial scales, which prevents the problem from being broken up into smaller units; a full kinetic simulation requires the use of at least 50,000 computing cores, and sometimes more.

One key factor in plasma performance is the stabilizing influence of toroidal rotation. In today's tokamaks, rotation is driven mainly by external beams, but beam drive will be less effective in larger future devices such as ITER. On the other hand, self-acceleration provides an intrinsic rotation—a spontaneous rotation without external momentum input. This phenomenon has been observed in many experiments, but its origin has not been well understood until now.

The SciDAC-developed XGC1 code is the world's first and only gyrokinetic code able to simulate the multiscale turbulence and background physics in realistic edge geometries including the magnetic separatrix and the material wall.

The EPSI team used XGC1 to model relevant multi-scale physics over the entire plasma volume of the DIII-D fusion reactor and explain the intrinsic rotation phenomenon. The simulation demonstrated for the first time that the toroidal momentum is generated at the edge, propagates inward as the edge pedestal is formed, and is redistributed by turbulence in the plasma core.

A useful analogy for this phenomenon is a heat engine, which converts some of the energy in a temperature difference into mechanical work. In fusion reactors that generate intrinsic rotation, the energy from hot spots at the plasma edge is converted to kinetic energy, which causes the plasma to rotate.

This is a key finding for modeling of future reactors such as ITER, and it is currently being validated by experiments in current reactors such as the NSTX. Multi-physics simulations such as this one are possible only on petascale computers, but modeling of larger devices like ITER will require exascale systems.



PROJECT

SciDAC Center for Edge
Physics Simulation

PROJECT LEADER


C.S. Chang, Princeton
Plasma Physics Laboratory

NERSC RESOURCES

10-million-hour simulation
using one-half of Hopper.
Total project usage: 78
million hours

DOE OFFICE

Fusion Energy Sciences



// Flamelets or hot spots along the plasma edge (a) drive turbulence intensity (b), temperature intensity (c), and intrinsic torque (d) inward, converting heat into toroidal rotation. (S. Ku et al.)

Publication: S. Ku, J. Abiteboul, P.H. Diamond, G. Dif-Pradalier, J.M. Kwon, Y. Sarazin, T.S. Hahm, X. Garbet, C.S. Chang, G. Latu, E.S. Yoon, Ph. Ghendrih, S. Yi, A. Strugarek, W. Solomon, and V. Grandgirard, "Physics of intrinsic rotation in flux-driven ITG turbulence," *Nuclear Fusion* **52**, 063013 (2012), doi:10.1088/0029-5515/52/6/063013.

First-Ever Analysis and Visualization of a Trillion Particle Dataset

New insights from 3D model of energetic particles in magnetic reconnection

Modern research tools like supercomputers, particle colliders, and telescopes are generating so much data, so quickly, that many scientists fear they will not be able to keep up with the deluge. That's why Surendra Byna and several of his colleagues from Berkeley Lab's Computational Research Division teamed up with researchers from the University of California, San Diego (UCSD), Los Alamos National Laboratory, Tsinghua University, and Brown University to develop novel software strategies for storing, mining, and analyzing massive datasets—more specifically, for data generated by a state-of-the-art plasma physics code called VPIC.

When the team ran VPIC on NERSC's Hopper Cray XE6 supercomputer, they generated a three-dimensional magnetic reconnection dataset of a trillion particles. VPIC simulated the process in thousands of time steps, periodically writing a massive 32 terabyte (TB) file to disk.


Using H5Part, a particle data extension of parallel HDF5 that enables high performance parallel I/O, the researchers wrote each 32 TB file to disk in about 20 minutes, at a sustained rate of 27 gigabytes per second (GB/s). By applying an enhanced version of the FastQuery tool, the team indexed this massive dataset in about 10 minutes, then queried the dataset in three seconds for interesting features to visualize.

"This is the first time anyone has ever queried and visualized 3D particle datasets of this size," says Homa Karimabadi, who leads the space physics group at UCSD.

Magnetic reconnection is a process in which the magnetic topology in a plasma (a gas made up of charged particles) is rearranged, leading to an explosive release of energy in form of plasma jets, heated plasma, and energetic particles. Reconnection is the mechanism behind the aurora borealis and solar flares, as well as fractures in Earth's protective magnetic field—fractures that allow energetic solar particles to enter our planet's magnetosphere and wreak havoc on electronics, power grids, and space satellites.

Identifying mechanisms leading to particle energization remains an important unsolved problem in plasma physics. But until recently, the closest that any researcher has come to studying this is by looking at 2D simulations. Although these datasets are much more manageable, containing at most only billions of particles, they leave out many important details, such as flux rope interactions and resulting turbulence, which only occur in 3D simulations with trillions of particles.

To address the challenges of analyzing 3D particle data, Karimabadi and a team of astrophysicists joined forces with the ExaHDF5 team, a DOE-funded collaboration to develop high performance I/O and analysis strategies for future exascale computers. With the new query-based visualization techniques, the researchers were finally able to verify the localization behavior of energetic particles, gain insights into the relationship between magnetic field structure and energetic particles, and discover the agyrotropic distribution of particles near the reconnection hotspot.



// After querying a dataset of 114,875,956,837 particles for those with energy values less than 1.5, FastQuery identifies 57,740,614 particles, which are mapped on this plot. (Oliver Rübel)

PROJECT

Petascale Kinetic Simulations in Laboratory and Space Plasmas; High Performance Visualization, Analytics, and I/O

PROJECT LEADER

Homa Karimabadi, UCSD;
E. Wes Bethel, LBNL

NERSC RESOURCES

24 million hours on Hopper

OTHER COMPUTING RESOURCES

HECC, NICS, OLCF

DOE OFFICES

Fusion Energy Sciences,
Advanced Scientific
Computing Research

Full story: <https://www.nersc.gov/news-publications/news/science-news/2012/sifting-through-a-trillion-electrons/>

Publication: S. Byna et al., "Parallel I/O, Analysis, and Visualization of a Trillion Particle Simulation," Proceedings of SC12, Salt Lake City, UT, November 2012.

New Model Will Help Predict Stability of CO₂ Reservoirs

First large-scale simulation of reactive transport at the pore scale

Geologic carbon sequestration—trapping carbon dioxide (CO₂) from fossil-fuel-burning power plants in underground aquifers—is a key strategy for mitigating global warming. Developing tools and techniques to predict and control the long-term stability of these CO₂ reservoirs is the mission of the DOE Energy Frontier Research Center (EFRC) for Nanoscale Control of Geologic CO₂ (NCGC), a collaborative effort led by Lawrence Berkeley National Laboratory, and including three other national laboratories and four universities.

As CO₂ is pumped underground, it can react chemically with underground minerals and brine in various ways, sometimes resulting in mineral dissolution and precipitation, which can change the porous structure of the aquifer. But predicting these changes is difficult because these processes take place at the pore scale and cannot be calculated using macroscopic models. The dissolution rates of many minerals have been found to be slower in the field than those measured in the laboratory; understanding this discrepancy requires modeling the pore-scale interactions between reaction and transport processes, then scaling them up to reservoir dimensions.


To meet this challenge, the NCGC team combined the SciDAC-developed Chombo platform for direct numerical simulation of flow and transport with the complex geochemistry code CrunchFlow. With the integrated Chombo-Crunch code, and using a NERSC NISE allocation on Hopper, they performed the first-ever large-scale simulation of reactive transport pore-scale processes on real-pore-space geometry as obtained from experimental data.

Textures affecting flow, transport, and reaction—as might be found in porous rock—were explicitly introduced into the pore-scale model of a capillary tube packed with crushed calcite, with full resolution of all interfaces at the micron scale. This was the largest pore-scale reactive flow simulation ever attempted, involving 1.6 billion grid points, and requiring 8.7 million processor hours on Hopper and up to 138,000 cores. The grid resolution of the model is now very close to that of imaging experiments, enabling a key EFRC goal of combining simulation and experiment.

The new high-resolution model demonstrated that the mineral dissolution rate depends on the pore structure of the aquifer: the average reaction rate depends on the flow pattern within the porous rock and the concentration gradients that develop at the pore scale. Lack of uniformity in the flow field at the pore scale (due to varying rock structures) has the effect of decreasing the overall reactivity of the system, and this effect becomes more pronounced as the variety of rock structures increases and the fluid flow slows down. This result is the first step in explaining the discrepancy between laboratory experiments and field measurements of mineral dissolution rates.

The integration of microscopic imaging, pore-scale modeling, and carefully controlled experimentation will provide an unprecedented new understanding of how fundamental molecular-scale processes of flow, transport, and reaction play out at the pore scale. New insights into the pore-scale dynamics that merge under the far-from-equilibrium conditions associated with the injection of CO₂ will greatly improve researchers' ability to develop scientifically defensible predictive models of CO₂ reservoirs.





// Reactive transport in a capillary tube packed with crushed calcite from experimental image data. Resolution of the simulation grid (top) is 1 micron, as shown in the second slice of calculated calcite saturation (bottom). (David Trebotich)

PROJECT

Advanced Simulation
of Pore Scale Reactive
Transport Processes
Associated with Carbon
Sequestration

PROJECT LEADER

David Trebotich, Lawrence
Berkeley National
Laboratory

NERSC RESOURCES

8.7 million hours on
Hopper

DOE OFFICE

Advanced Scientific
Computing Research

Publication: S. Molins, D. Trebotich, C. I. Steefel, and C. Shen,
"An Investigation of the Effect of Pore Scale Flow on Average
Geochemical Reaction Rates using Direct Numerical Simulation,"
Water Resources Research **48**, W03527 (2012),
doi:10.1029/2011WR011404.

Researchers Discover a New Kind of Neutrino Transformation

One of Science Magazine's Top 10 Breakthroughs of 2012

Neutrinos, the wispy particles that flooded the universe in the earliest moments after the Big Bang, are continually produced in the hearts of stars and other nuclear reactions. Untouched by electromagnetism, they respond only to the weak nuclear force and even weaker gravity, passing mostly unhindered through everything from planets to people.

Years ago scientists discovered that although neutrinos come in three basic "flavors"—electron, muon, and tau—neutrinos and their corresponding antineutrinos can transform from one flavor to another while they are traveling close to the speed of light. How they do this has been a longstanding mystery.

But some new, and unprecedentedly precise, measurements from the multinational Daya Bay Neutrino Experiment are revealing how electron antineutrinos "oscillate" into different flavors as they travel. This finding opens a gateway to a new understanding of fundamental physics and may eventually help solve the riddle of why there is far more matter than antimatter in the universe today.

Nuclear reactors of the China Guangdong Nuclear Power Group at Daya Bay and nearby Ling Ao produce millions of quadrillions of electron antineutrinos every second. Six massive detectors buried in the mountains adjacent to the reactors make up the Daya Bay Experiment. Researchers in the collaboration count the number of electron antineutrinos detected in the halls nearest the reactors and calculate how many would reach the detectors in the Far Hall if there were no oscillation. The number that apparently vanishes on the way (oscillating into other flavors, in fact) gives the value of theta one-three (written θ_{13}), the last neutrino "mixing angle" to be precisely measured.

Shortly after experimental data is collected, it travels across the Pacific Ocean to the NERSC Global Filesystem, from which the data is processed automatically on the PDSF and Carver systems, stored on the PDSF file system, and shared with collaborators around the world via the Daya Bay Offline Data Monitor, a web-based "science gateway" hosted by NERSC. NERSC is the only U.S. site where all of the raw, simulated, and derived Daya Bay data are analyzed and archived (see story on page 62).

The first Daya Bay results show that θ_{13} , once feared to be near zero, instead is "comparatively huge," remarks Luk Kam-Biu Luk of Berkeley Lab and the University of California, Berkeley. Luk is co-spokesperson of the Daya Bay Experiment and heads U.S. participation. "What we didn't expect was the sizable disappearance, equal to about six percent."

"This experiment could not have been done without NERSC and ESnet," says Craig Tull of Berkeley Lab's Computational Research Division. "The software and computing components of the Daya Bay Neutrino Experiment have been extremely crucial to science. Thanks to the computing expertise at Berkeley Lab, we were able to see antineutrinos ... in the first filled detectors within 24 hours. We were able to see an anti-neutrino deficit in the far hall within days. And finally, we have been able to extract a high-quality θ_{13} value within only 75 days of start of far-hall running."


PROJECT
 Daya Bay Reactor Neutrino Experiment

PROJECT LEADER
 Craig Tull, Lawrence
 Berkeley National
 Laboratory

NERSC RESOURCES
 PDSF, Carver, NGF, HPSS

**OTHER COMPUTING
 RESOURCES**
 Brookhaven National
 Laboratory

DOE OFFICE
 High Energy Physics



// Detectors count antineutrinos near the Daya Bay nuclear reactor in China. By calculating how many would be seen if there were no oscillation and comparing to measurements (inset), a 6.0% rate deficit provides clear evidence of the new transformation. (Photo: Roy Kaltschmidt, LBNL; Graph: Daya Bay Collaboration)

Publication: F. P. An et al., "Observation of Electron-Antineutrino Disappearance at Daya Bay," *Physical Review Letters* 108, 171803 (2012), doi:10.1103/PhysRevLett.108.171803.

Looking for No Neutrinos

Physicists set stringent limits on neutrinoless double beta decay

With help from supercomputers at NERSC, the Enriched Xenon Observatory experiment (EXO-200) has placed the most stringent constraints yet on the nature of a process called neutrinoless double beta decay. In doing so, the physicists have narrowed down the range of possible masses for the neutrino—a tiny uncharged particle that rarely interacts with anything, passing right through people and planets at nearly the speed of light.

In normal double-beta decay, two neutrons in an unstable atomic nucleus turn into two protons. And in the process, two electrons and two antineutrinos—the antimatter counterparts of neutrinos—are emitted. But physicists have suggested that two neutrons could also decay into two protons by emitting two electrons without producing any antineutrinos. This implies that the two neutrinos produced in this neutrinoless double beta decay somehow cancel each other out.


For this to happen, a neutrino must be its own antiparticle, allowing one of the two neutrinos to act as an antineutrino and annihilate the other. However, the widely accepted scientific theory that describes how all elementary particles behave and interact—called the Standard Model—does not predict that a neutrino can be its own antiparticle. So if this neutrinoless process does indeed exist, physicists would be forced to revise the Standard Model.

The international EXO-200 Collaboration has a detailed computer model that predicts the physics processes occurring inside the observatory. After each experiment, researchers compare or “fit” the experimental results to find the parameters of the computational model that describes the results best. In normal double-beta decay, half of a given sample would decay after 10^{21} years—a half-life roughly 100 billion times longer than the time that has elapsed since the Big Bang.

In this case, the first results in almost seven months of data showed no signal for neutrinoless double beta decay. When the team fit these results with their computational parameter that describes the rate of neutrinoless double-beta decay of Xe-136, they were able to rule out possible values for the half-life of the neutrinoless process. In their paper, the team proclaims at 90 percent confidence level that the neutrinoless double-beta decay of Xe-136 half-life cannot be shorter than 1.6×10^{25} years, or a quadrillion times older than the age of the universe.

“This confidence level statement essentially means that if we were to run this experiment 1,000 times, we would expect to calculate limits on the neutrinoless double-beta decay rate of Xe-136 that enclose the true rate 900 times,” says Michael Marino, a postdoctoral researcher at Germany’s Technical University of Munich and member of the EXO-200 team. “To check this in practice, we cannot run this experiment 1,000 times, but we did run simulations of this experiment 1,000 times on Hopper.”

With the value of the half-life pinned down, physicists can calculate the mass of a neutrino—another longstanding mystery. The new data suggest that a neutrino cannot be more massive than about 0.140 to 0.380 electron volts (eV, a unit of mass commonly used in particle physics); an electron, by contrast, is about 500,000 eV, or about 9×10^{-31} kilograms.



PROJECT

Data Analysis for EXO,
the Enriched Xenon
Observatory

PROJECT LEADER


Giorgio Gratta, Stanford
University

NERSC RESOURCES

Hopper

DOE OFFICE

High Energy Physics



// Located at the Waste Isolation Pilot Plant near Carlsbad, New Mexico, EXO-200 is a 200-kg detector using liquid xenon, enriched in the isotope 136, to detect neutrinoless double beta decay. (*Enriched Xenon Observatory*)

Publication: M. Auger, et al. (EXO Collaboration), "Search for Neutrinoless Double-Beta Decay in ^{136}Xe with EXO-200," *Physical Review Letters* 109, 032505 (2012), doi:10.1103/PhysRevLett.109.032505.

New Method Enables Fine Tuning of Electronic and Solar Materials

Optical absorption spectrum of silicon calculated from first principles

Solar electricity is a clean and abundant energy source that can satisfy the world's long-term energy needs, but bridging the gap between the present utilization and its full potential requires significant advances in solar-cell engineering. First-principles computational studies, which require only knowledge of the fundamental properties of the material, can provide theoretical insights that will guide experimental research efforts in developing the next generation of efficient and inexpensive photovoltaic devices.

Using a NERSC Initiative for Scientific Exploration (NISE) computer allocation, and supported in part by the Center for Energy Efficient Materials, a DOE Energy Frontier Research Center, Emmanouil Kioupakis of the University of Michigan and collaborators at the University of California in Berkeley and Santa Barbara were able to compute the phonon-assisted interband optical absorption spectrum of silicon entirely from first principles—an important process for optoelectronic and photovoltaic applications that cannot be addressed with simple models.


A phonon is a quantum mechanical description of a special type of vibrational motion, in which a lattice uniformly oscillates at the same frequency. Lattice vibrations or defects provide additional momentum that allows intraband light absorption, and silicon is a commercially successful photovoltaic material because of such optical transitions. In fact, nearly all commercially available photovoltaic cells currently depend on this absorption process. Despite its importance, however, only a very limited number of first-principles studies of phonon-assisted optical absorption spectra exist, primarily because of the high computational cost of modeling the phonon interactions.

A key ingredient in the new method was the BerkeleyGW software written by NERSC User Services Group consultant Jack Deslippe. BerkeleyGW is a massively parallel code that takes as input the mean-field results from various electronic structure codes and calculates the quasiparticle properties and the optical responses of a large variety of materials, from bulk periodic crystals to nanostructures such as slabs, wires, and molecules. It can be used in conjunction with many density-functional theory codes for ground-state properties. Kioupakis and colleagues used BerkeleyGW with a custom Wannier interpolation formalism to determine quasiparticle energies and electron-phonon coupling matrix elements associated with the optical transition.

The resulting high-resolution calculated spectra near the absorption onset are in very good agreement with experimental results, especially in the energy range between the indirect and direct band gaps. This is a spectral region that cannot be accessed by standard model calculations and thus had never before been calculated entirely from first principles. This region covers the entire visible spectrum and is important for optoelectronic applications.

Nanomaterials and other complex materials have unique and highly tunable optical-electronic properties. The new method is general enough to predict and analyze the phonon-assisted optical absorption spectrum of any material and can address questions that are not accessible by experiment. It will enable better evaluation and fine-tuning of material properties with applications in efficient lighting and solar energy conversion.





// Absorption spectrum of silicon for two temperatures showing excellent agreement between computed and experimental values. (*J. Noffsinger et al.*)

PROJECT
Electronic and Optical
Properties of Novel
Photovoltaic and
Thermoelectric Materials
from First Principles

PROJECT LEADER
Emmanouil Kioupakis,
University of Michigan

NERSC RESOURCES
4 million hours on Hopper

**OTHER COMPUTING
RESOURCES**
University of California,
Santa Barbara

DOE OFFICE
Basic Energy Sciences

Publication: J. Noffsinger, E. Kioupakis, C. G. Van de Walle, S. G. Louie, and M. L. Cohen, "Phonon-assisted optical absorption in silicon from first principles," *Physical Review Letters* **108**, 167402 (2012), doi:10.1103/PhysRevLett.108.167402.

A New, More Efficient Approach to Water Desalination

In simulation, graphene sheets purify water faster or at lower pressure

The availability of fresh water is dwindling in many parts of the world, a problem that is expected to grow with populations. One promising source of potable water is the world's virtually limitless supply of seawater, but so far desalination technology has been too expensive for widespread use.

Now, using supercomputers at NERSC, researchers from the Massachusetts Institute of Technology (MIT) have come up with a new approach using a different kind of filtration material: sheets of graphene, a one-atom-thick form of the element carbon, which they say can be far more efficient and possibly less expensive than existing desalination systems.

"We were very pleasantly surprised by how well graphene performed compared to existing systems in computer simulations," says Jeffrey Grossman, Professor of Power Engineering at MIT, who is the senior author of a paper describing the new process in the journal *Nano Letters*.

"NERSC proved to be an invaluable resource for our research," says David Cohen-Tanugi, an MIT graduate student who is the lead author of the paper.


Using NERSC's Hopper and Carver systems, Grossman and Cohen-Tanugi aimed to control the properties of the material down to the atomic level, producing a graphene sheet perforated with precisely sized holes. They also added other elements to the material, causing the edges of these minuscule openings to interact chemically with water molecules—either repelling or attracting them.

One common method of desalination, called reverse osmosis, uses membranes to filter the salt from the water. But these systems require extremely high pressure—and hence, high energy use—to force water through the thick membranes, which are about a thousand times thicker than graphene. The new graphene system operates at much lower pressure—with permeability rates 100 to 1,000 times higher than conventional polymeric reverse osmosis membranes—and thus could purify water at far lower cost.

While reverse osmosis has been used for decades, "really basic mechanisms of separating salt from water are not well understood," Cohen-Tanugi says, adding that it's very difficult to do experiments at the scale of individual molecules and ions. But the graphene-based system, he says, works "hundreds of times faster than current techniques, with the same pressure"—or, alternatively, the system could run at similar rates to present systems, but with lower pressure.

The key to the new process is very precise control over the size of the holes in the graphene sheet. "There's a sweet spot, but it's very small," Grossman says—between pores so large that salt could pass through and ones so small that water molecules would be blocked. The ideal size is just about one nanometer, or one billionth of a meter, he says. If the holes are just a bit smaller—0.7 nanometers—the water won't flow through at all.

And because graphene is such a strong material—pound for pound, it's the strongest material known—the membranes should be more durable than those presently used for reverse osmosis, Grossman says.



// When water molecules (red and white) and sodium and chlorine ions (green and purple) in saltwater, on the right, encounter a sheet of graphene (pale blue, center) perforated by holes of the right size, the water passes through (left side), but the sodium and chlorine of the salt are blocked. (David Cohen-Tanugi)

PROJECT

Quantum Simulations
of Nanoscale Energy
Conversion

PROJECT LEADER

Jeffrey Grossman,
Massachusetts Institute of
Technology

NERSC RESOURCES

2 million hours on Hopper
and Carver

DOE OFFICE

Basic Energy Sciences

Full story: <https://www.nersc.gov/news-publications/news/science-news/2012/a-new-approach-to-water-desalination/>

Publication: David Cohen-Tanugi and Jeffrey C. Grossman, "Water Desalination across Nanoporous Graphene," *Nano Letters* **12**, 3602–3608 (2012), doi:10.1021/nl3012853

Intrinsic and Engineered Piezoelectricity in 2D Materials

Surprising discoveries may lead to a new array of nanoscale devices

Nanoelectromechanical systems (NEMS) and nanoscale electronics are the final frontier in the push for miniaturization that has been one of the dominant themes of technological progress for over 50 years. Low-dimensional materials including nanoparticles, nanotubes, and atomically thin sheets have emerged as the key components for the next generation of nano devices, including radios, switches, tweezers, a wide variety of sensors, light-emitting diodes (LEDs), generators, relays, transistors, motors, and robots.

Many of today's larger-scale technologies, from ultrasound scanners to microphones and inkjet printers, rely on the piezoelectric effect, also known as "pressure electricity." Under normal circumstances, piezoelectric materials are electrically homogeneous: a positive charge in one place cancels out a negative charge nearby. But under pressure, the crystals that make up these materials deform, upsetting the spatial balance of positive and negative and producing a voltage. Part of what makes piezoelectric materials so useful is that the effect can also be reversed—applying an electric field to a crystal will cause it to change shape or deform.


Sometimes graphene is referred to as the "miracle material" of the 21st century. Composed of a single sheet of carbon atoms, graphene is tougher than diamond, more conductive than copper, and has potential applications in a variety of nanoscale technologies. Unfortunately, graphene is not intrinsically piezoelectric and therefore is not amenable to control with moderate electric fields.

But researchers led by Evan Reed in the Department of Materials Science and Engineering at Stanford University, conducting density functional theory (DFT) calculations on NERSC supercomputers, have discovered that with a little chemical doping on one side of the sheet, graphene can be transformed into a controllable piezoelectric material, with piezoelectric magnitudes comparable to those in 3D materials.

"The strain or deformation generated in piezoelectric materials is proportional to the magnitude of the electric field applied, so this phenomenon would allow us to dynamically control the electronic properties of graphene," says Mitchell Ong, lead author of the paper that reported this discovery.

Ong notes that previous experiments have shown that absorbates can be spatially patterned on the surface of graphene, which has led the researchers to the idea of engineering devices by inducing piezoelectricity selectively. "This experimental work, combined with our results, hints that devices could be created where strains are spatially concentrated to control local electrical and chemical properties," says Ong.

Encouraged by this discovery, the research team turned to a group of widely studied atomically thin sheets composed of transition metal compounds. As 3D crystals, none of these metals have piezoelectric properties. But surprisingly, DFT calculations show that as 2D sheets, some of these materials exhibit stronger piezoelectric coupling than bulk materials commonly used in today's piezoelectric devices. This discovery opens up an entirely new arsenal of "out of the box" active components for NEMS and piezo-electronics. Armed with their NERSC results, Reed and his collaborators are now testing these simulations in laboratory experiments.



// Schematic illustration of mechanical strain on a monolayer transition metal sheet producing electrical current. (Duerloo, Ong, and Reed)

PROJECT

Dynamics of Materials at Extreme Conditions

PROJECT LEADER

Evan Reed, Stanford University

NERSC RESOURCES

Franklin, Hopper

OTHER COMPUTING RESOURCES

Stanford National Nanotechnology Infrastructure Network Computing Facility

DOE OFFICE

Basic Energy Sciences

Full story: <https://www.nersc.gov/news-publications/news/science-news/2012/a-roadmap-for-engineering-piezoelectricity-in-graphene/>

Publication: Mitchell T. Ong and Evan J. Reed, "Engineered Piezoelectricity in Graphene," *ACS Nano* 6, 1387–1394 (2012), doi:10.1021/nn204198g. Karel-Alexander N. Duerloo, Mitchell T. Ong, and Evan J. Reed, "Intrinsic Piezoelectricity in Two-Dimensional Materials," *Journal of Physical Chemistry Letters* 3, 2871–2876 (2012), doi:10.1021/jz3012436.

NISE Program Encourages Innovative Research

The NERSC Initiative for Scientific Exploration (NISE) is a mechanism used for allocating the NERSC reserve (10 percent of the total allocation). It is a competitive allocation administered by NERSC staff and management. In 2012 we were particularly interested in large-scale or data-intensive proposals. We received 61 NISE requests and awarded 108 million hours to 23 projects, as listed below.

Project Title	Investigator	Hours Awarded	DOE Office & Program
Ocean-Atmosphere Reanalysis for Climate Applications (OARCA) 1850-2013	Gilbert Compo University of Colorado at Boulder	10,000,000	BER Climate Research
WeFold: A Collaborative Effort for Protein Structure Prediction	Silvia Crivelli Lawrence Berkeley National Laboratory	1,550,000	BER Biological Systems Science
A Multi-Decadal Reforecast Data Set to Improve Weather Forecasts for Renewable Energy Applications	Thomas Hamill National Oceanic & Atmospheric Administration	1,675,000	BER Climate Research
Excited-State Dynamics and Properties of Complex Interfaces for Energy Conversion	Yosuke Kanai University of North Carolina at Chapel Hill	2,000,000	BES Materials Science
Kinetic Simulations in Laboratory and Space Plasmas	Homa Karimabadi University of California, San Diego	7,000,000	FES Fusion Energy
Turbulent Reacting Flows for Multi-Physics Model Development	Colleen Kaul Stanford University	4,500,000	BES Chemistry
Systems Biology Knowledge Base	Keith Keller Lawrence Berkeley National Laboratory	100,000	BER Biosciences
Electronic Properties of Novel Nitride Nanostructures	Emmanouil Kioupakis	3,000,000	BES Materials Science
Integrated Carbon Cycle Data Assimilation with NCAR Carbon-Climate Model	Junjie Liu Jet Propulsion Laboratory	5,100,000	BER Climate Research
Environmental Fluctuations and Gating in Bio-Inorganic Proton-Coupled Electron Transfer	Thomas Miller California Institute of Technology	7,000,000	BES Chemistry

Project Title	Investigator	Hours Awarded	DOE Office & Program
Turbulence over Complex Terrain: A Wind Energy Perspective	Edward Patton National Center for Atmospheric Research	3,200,000	BER Climate Research
The Materials Genome	Kristin Persson Lawrence Berkeley National Laboratory	11,500,000	BES Materials Science
Global Full-Waveform Seismic Tomography	Barbara Romanowicz University of California, Berkeley	1,500,000	BES Geosciences
Quantum Transport Simulation of Nano Scale Electronic Devices for Ultra Low Power Computing	Sayeeff Salahuddin University of California, Berkeley	300,000	BES Materials Science
Carbon Dioxide Gas Separation Using Nanoporous Graphene	Joshua Schrier Haverford University	700,000	BES Chemistry
Attribute-Based Unified Data Access Service	Arie Shoshani Lawrence Berkeley National Laboratory	1,525,000	ASCR Computer Science
Mean-Field Solutions for Heavy Nuclei: Structure and Dynamics	Ionel Stetcu Los Alamos National Laboratory	1,100,000	NP Nuclear Theory
Attribution of Extreme Weather Risk to Anthropogenic Emissions	Daithi Stone Lawrence Berkeley National Laboratory	8,650,000	BER Climate Research
Advanced Simulation of Pore Scale Reactive Transport Processes Associated with Carbon Sequestration	David Trebotich Lawrence Berkeley National Laboratory	7,500,000	ASCR Applied Math & BES Energy Frontier Research Centers
Computational Prediction and Discovery of Magnet Materials	Cai-Zhuang Wang Ames Laboratory Iowa State University	13,100,000	BES Materials Science
Spin-Lattice Coupling in Magnetic Phase Transition	Yi Wang Pennsylvania State University	6,000,000	BES Materials Science
Understanding Multiple Exciton Generation and Charge Extraction in Nanoparticle-Based Solar Cells	Stefan Wippermann University of California, Davis	5,500,000	BES Chemistry
Guest-Host Interactions in Hydrate Lattices: Implications for Hydrogen Storage and Carbon Dioxide Sequestration	Sotiris Xantheas Pacific Northwest National Laboratory	7,500,000	BES Chemistry

NERSC Users' Awards and Honors

NERSC HPC Achievement Awards

NERSC announced the winners of the inaugural High Performance Computing (HPC) Achievement Awards in February 2013 at the annual NERSC Users Group meeting. The awardees are all NERSC users who have either demonstrated an innovative use of HPC resources to solve a scientific problem, or whose work has had an exceptional impact on scientific understanding or society. In an effort to encourage young scientists who are using HPC in their research, NERSC also presented two early career awards.

Nominations were made by NERSC principal investigators, project managers, PI proxies, and NERSC staff. Winners were chosen by representatives from the NERSC Users Group Executive Committee and NERSC staff.

"High performance computing is changing how science is being done and facilitating breakthroughs that would have been impossible a decade ago," says NERSC Director Sudip Dosanjh. "The 2013 NERSC Achievement Award winners highlight some of the ways this trend is expanding our fundamental understanding of science, and how we can use this knowledge to benefit humanity."

"The winning projects represent just a small sample of the groundbreaking research being done with NERSC resources. We received so many great nominations, and I look forward to seeing these projects and researchers spotlighted in future HPC Achievement Awards," says Richard Gerber, Deputy Lead of NERSC's User Services Group.

"I am especially impressed by our Early Career award winners," says Dave Goodwin, who manages NERSC for Department of Energy's Office of Science. "These young researchers are the future of science, and it is really gratifying to see that they recognize the important role of HPC in science and achieve such incredible scientific results so early in their career."

This year's winners include:

NERSC Award for High Impact Scientific Achievement

Jeff Grossman and David Cohen-Tanugi, *Massachusetts Institute of Technology*

Using supercomputers at NERSC, Grossman and Cohen-Tanugi came up with a new approach for desalinating seawater using sheets of graphene, a one-atom-thick form of the element carbon. This method holds the promise of being far more efficient and possibly less expensive than existing desalination systems. With world populations expected to keep growing, and potable water projected to grow scarcer over the coming century, a practical and cheap means of desalinating seawater is one of materials science's holy grails.

The key to this new approach is very precise control over the size of the holes in the graphene sheet. Using NERSC's Hopper and Carver systems, the researchers aimed to control the properties of the material down to the atomic level, producing a graphene sheet perforated with precisely sized holes. They also added other elements to the material,

causing the edges of these minuscule openings to interact chemically with water molecules—either repelling or attracting them. They found that the ideal pore size is just about one nanometer, or one-billionth of a meter. If the holes are just a bit smaller—0.7 nanometers—the water won't flow through at all.

One common method of desalination, called reverse osmosis, uses membranes to filter the salt from the water. But these systems require extremely high pressure—and hence, energy use—to force water through the thick membranes, which are about a thousand times thicker than graphene. The new graphene system operates at much lower pressure, and thus could purify water at far lower cost. In December 2012, Smithsonian Magazine named this result the fifth “Surprising Scientific Milestone of 2012.”

More information:

<http://www.nerisc.gov/news-publications/news/science-news/2012/a-new-approach-to-water-desalination/>


NERSC Award for High Impact Scientific Achievement—Early Career

Tanmoy Das, Postdoctoral Researcher, Los Alamos National Laboratory

Das was nominated for his computational work to understand fundamental materials aspects in three different areas: (1) the role of Fermi surface anisotropy on the superconducting gap structure in multiband iron-based superconductors in the presence of rotating magnetic fields; (2) spin-orbit ordering effects in two-dimensional electron gases and in the hidden order state of URu₂Si₂; and (3) his seminal contributions to the self-consistent spin-fluctuation theory applied to real materials—the intermetallic actinides.

Using his own MPI algorithm, Das used approximately 256 cores primarily on NERSC's Hopper to compute the first-ever field-temperature phase diagram of the four-fold oscillations in the specific heat and the thermal conductivity of the new iron-based superconductor using material-specific Fermi surface parameterization, which showed a large effect of the Fermi surface on oscillations in thermal observables in a rotating magnetic field, even for fully isotropic superconducting gaps.

He also demonstrated for the first-time that the peak-dip-hump structure in the spectral function could be explained in terms of spin fluctuations in the particle-hole spectrum. He did this by running his MPI algorithm on 2,000 to 3,000 cores to compute renormalized spectral functions in a self-consistent many-body self-energy approximation for the intermetallic Pu-115 and UCoGa₅ compounds. This work could have important implications for future technologies.



// Jeff Grossman (top) and David Cohen-Tanugi (bottom), Massachusetts Institute of Technology (MIT).



// Tanmoy Das, Los Alamos National Laboratory.

NERSC Award for Innovative Use of High Performance Computing

Peter Nugent and the PTF Team, *Lawrence Berkeley National Laboratory*

The detection and analysis of transient events could lead to a greater understanding of astrophysical objects like supernovae, active galaxies, and gamma-ray bursts, among a variety of other known and unknown cosmic phenomena. But a major challenge has been identifying transient objects in real time among a scene of normal cosmic variations. The Palomar Transient Factory (PTF) was the first project dedicated solely to finding and following up transient events in real time, and they worked with NERSC to develop an automated system to sift through terabytes of astronomical data every night to find interesting events.



// Peter Nugent, Lawrence Berkeley National Laboratory

Every night for about four years, the PTF camera—a 100-megapixel machine mounted on the 48-inch Samuel Oschin Telescope at Palomar Observatory in Southern California—automatically snapped pictures of the sky, then sent those observations to NERSC, where computers running machine learning algorithms in the Real-Time Transient Detection Pipeline scoured this data for transients, or cosmic objects that change in brightness or position, by comparing the new observations with all of the data collected from previous nights. Once an interesting event was discovered, an automated system sent its coordinates to ground-based telescopes around the world for follow-up observations. NERSC also archived this data and allowed collaborators to access it over the Internet through a web-based science gateway called DeepSky.

In 2011, the PTF Real-Time Transient Detection Pipeline discovered the closest Type Ia supernova—approximately 21 million light-years away from Earth—in a generation, just hours after it exploded. Within hours of identifying the event, this automated system sent the coordinates to telescopes around the world for follow-up observations. By catching the supernova early, astronomers caught a rare glimpse of the outer layers of the explosion, which contained hints about the star as it once was. In 2010 the PTF discovered a new class of superluminous supernova. While rare, these are among the brightest explosions in the universe. The PTF pipeline also found the first-ever direct observations of a Type Ia supernova progenitor system. They were able to determine that the system contained a red giant star and that it previously underwent at least one smaller nova eruption before it ended its life in a destructive supernova. And recently, while digging through the PTF data archive at NERSC, astronomers found the first causal evidence that massive stars, which can explode as Type II supernovae, shed huge amounts of material in a “penultimate outburst” before their final destructive detonation.

More information:

<http://www.nersc.gov/news-publications/news/science-news/2011/supernova-caught-in-the-act/>

<http://www.nersc.gov/news-publications/news/science-news/2012/ptf11kx/>

<http://www.nersc.gov/news-publications/news/science-news/2013/a-massive-stellar-burst-before-the-supernova/>

<http://www.nersc.gov/news-publications/news/science-news/2009/nersc-helps-discover-cosmic-transients/>

NERSC Award for Innovative Use of High Performance Computing — Early Career

Edgar Solomonik, *University of California, Berkeley*

Solomonik was nominated for developing novel algorithms for massively parallel tensor contractions and applying them to quantum chemistry problems, specifically coupled-cluster theory, which is the de facto standard for important scientific applications in the thermochemistry of combustion and excited states of systems where density-functional theory (DFT) breaks down.

His algorithms represent a major development in the area of tensor computations. Rather than using one-sided communication and dynamic load-balancing as is done in NWChem and other codes like it, Solomonik has transformed the irregular nature of symmetric tensor contractions into highly regular computations that are solved using topology-aware and communication-avoiding dense linear algebra methods developed in Jim Demmel's group at UC Berkeley.

Solomonik's algorithmic developments are instantiated in the Cyclops Tensor Framework (CTF), which has been used on some of the largest supercomputers in the world, including the NERSC Hopper system and the IBM Blue Gene/Q systems at the Lawrence Livermore National Laboratory and Argonne Leadership Computing Facility. To date the largest application has achieved ~0.5 petaflop/s at ~30 percent of theoretical peak on Blue Gene/Q. On the Cray XE6 system at NERSC, CTF is faster than NWChem, the previous state-of-the-art coupled-cluster code.

Other Awards and Honors

Every year a significant number of NERSC users are honored for their scientific discoveries and achievements. Listed below are some of the most prominent awards given to NERSC users in 2012.

Fellow of the American Academy of Arts and Sciences

Eli Yablonovitch, *University of California, Berkeley*

Fellow of the American Association for the Advancement of Science

Wim Leemans, *Lawrence Berkeley National Laboratory*

Member of the National Academy of Sciences

John B. Bell, *Lawrence Berkeley National Laboratory*

Fellows of the American Physical Society (APS)

Alan Aspuru-Guzik, *Harvard University*

James P. Colgan, *Los Alamos National Laboratory*

Amalie Frischknecht, *Sandia National Laboratories*

Lloyd E. Knox, *University of California, Davis*

Zoltan Ligeti, *Lawrence Berkeley National Laboratory*

Dirk Klaus Morr, *University of Illinois, Chicago*

Doreen Wackeroth, *State University of New York, Buffalo*

APS George E. Valley Jr. Prize

Jinhui Chen, *Chinese Academy of Science*

APS John Dawson Award for Excellence in Plasma Physics Research

Laurent Divol, *Lawrence Livermore National Laboratory*

APS Outstanding Doctoral Thesis Research in Beam Physics Award

Daniel Ratner, *Stanford University*

APS Dissertation Award in Nuclear Physics

Phillip S. Barbeau, *Stanford University*

APS Nicholas Metropolis Award for Outstanding Doctoral Thesis Work in Computational Physics

Justin Weber, *University of California, Santa Barbara*

Advanced Accelerator Concepts Prize

Wim Leemans, *Lawrence Berkeley National Laboratory*

Royal Society of Chemistry Faraday Lectureship Prize

Richard Saykally, *University of California, Berkeley*

R&D 100 Award

Qing Ji, *Lawrence Berkeley National Laboratory*

Fellow of the Society for Industrial and Applied Mathematics (SIAM)

Barry F. Smith, *Argonne National Laboratory*

Institute of Electrical and Electronics Engineers (IEEE) Photonics Award

Eli Yablonovitch, *University of California, Berkeley*

Association for Computing Machinery (ACM) Fellow

Katherine Yelick, *Lawrence Berkeley National Laboratory*

ACM Distinguished Scientist

E. Wes Bethel, *Lawrence Berkeley National Laboratory*

ACM IEEE-CS George Michael Memorial HPC Fellowship

Yanhua Sun, *University of Illinois at Urbana-Champaign*

Department of Energy Ernest Orlando Lawrence Award

Barry Smith, *Argonne National Laboratory*

Department of Energy Office of Science Early Career Research Awards

Pavan Balaji, *Argonne National Laboratory*
Céline Bonfils, *Lawrence Livermore National Laboratory*
Hank Childs, *Lawrence Berkeley National Laboratory*
Aleksandar Donev, *New York University*
Gary E. Doublerly, *University of Georgia*
Huaiyu Duan, *University of New Mexico*
Daniel Kasen, *University of California, Berkeley*
Andreas J. Kemp, *Lawrence Livermore National Laboratory*
Svetlana Kilina, *North Dakota State University*
Jaime Marian, *Lawrence Livermore National Laboratory*
Nikhil Padmanabhan, *Yale University*
Felix Parra Diaz, *Massachusetts Institute of Technology*
Pavel Snopok, *Illinois Institute of Technology*
Lisa Whitehead, *University of Houston*

The NERSC Center

Sudip Dosanjh Is Named New NERSC Director

Sudip Dosanjh, a leader in extreme-scale computing at Sandia National Laboratories in Albuquerque, was named director of the NERSC Division at Lawrence Berkeley National Laboratory in August 2012 and began work in November. Dosanjh is the sixth director of NERSC, which was established in 1974.

"As one of the world's best-known centers for computational science, NERSC helps the nation's science community advance discovery by tapping into the capabilities of petascale supercomputers," Dosanjh said. "NERSC will need to work closely with the research community to ensure that scientific progress continues through the disruption in computer architectures that is underway."

Dosanjh was previously the group lead for extreme-scale computing at Sandia, and co-led the development of the exascale technology roadmap that was presented at the DOE Architectures and Technology workshop in December 2009. He has given numerous presentations on exascale computing and played a key role in establishing co-design as a DOE strategy for achieving exascale computing. Earlier in his career, Dosanjh worked extensively in developing large-scale parallel scientific applications in areas such as materials modeling, nuclear reactor safety, combustion, and heat transfer.

"In addition to his leadership in the field of exascale computing, Sudip also brings experience in developing productive partnerships in areas such as computer architecture and algorithms, and managing research programs—all of which make him well suited to lead NERSC," said Berkeley Lab Associate Laboratory Director Kathy Yelick. "As an experienced computational scientist and recognized leader, Sudip will also help us continue to provide innovative services to users and to attract the best people to the NERSC organization."

Yelick had served as NERSC Director since January 2008. She announced her intention to step down earlier in 2012 to focus on her duties as Associate Laboratory Director for Computing Sciences, which comprises NERSC, the Computational Research Division, and DOE's Energy Sciences Network (ESnet).



Dosanjh joined Sandia National Labs in 1986 and has held a number of positions, all involving high performance computing and computational science. Most recently he managed Sandia's Computer Systems and Software Environments (CSSE) Program, which is funded by DOE's Advanced Simulation and Computing (ASC) Program. He was co-director of the Alliance for Computing at Extreme Scale (ACES), a Los Alamos/Sandia center that designs, deploys, and operates capability supercomputers for the ASC program. He also served on DOE's Exascale Initiative Steering Committee for several years.

He earned his bachelor's degree in engineering physics in 1982, his master's degree (1984) and Ph.D. (1986) in mechanical engineering. All three of his degrees were earned at the University of California, Berkeley.

John Shalf Is Named Chief Technology Officer for NERSC

On December 5, 2012, John Shalf was named NERSC's Chief Technology Officer (CTO) by NERSC Director Sudip Dosanjh. Shalf also continues to serve in his role as head of the Computer and Data Sciences Department in Berkeley Lab's Computational Research Division (CRD).

As Chief Technology Officer, Shalf is helping NERSC develop a plan to achieve exascale performance—1 quintillion or 10^{18} computer operations per second—a thousand-fold increase over current petascale systems. Exascale is expected to require computers containing 10 million to 100 million processing elements or cores, as well as new programming techniques and software. Shalf will also help coordinate exascale research activities with future NERSC procurements.

"NERSC is responsible for moving hundreds of science codes that serve the diverse needs of the Office of Science to next generation architectures," Dosanjh said. "Because of the breadth of our applications, we cannot build application-specific architectures—rather, we must broadly support DOE's mission. This poses special challenges for data movement (memory and interconnect being examples). Our goal is to ensure that scientific progress is uninterrupted on the road to exascale. John's demonstrated intellectual leadership in advanced architecture development will help NERSC choose the most productive path forward."

Shalf added, "A key goal of DOE's exascale program is to develop high performance scientific computers that deliver a thousand times the performance of today's most powerful computers at all scales, while using less than twice the power, by the end of the next decade. The demands of energy efficiency are driving deep changes that will change the way we do computing at all scales, not just exascale. NERSC will take an active role to work with industry as a public/private partnership to guide HPC designs and bring the DOE user community along in this time of great transition."

Shalf has served on the DOE Exascale Steering Committee and is co-principal investigator of the Computer Architecture Laboratory (CAL), a joint activity of Sandia National Laboratories and Lawrence Berkeley National Laboratory for the DOE Office of Advanced Scientific Computing Research. CAL is coordinating hardware architecture research and development activities across the DOE. He also leads the CRD/NERSC Green Flash project, which is developing a novel HPC system design (hardware and software) for kilometer-scale global climate modeling that is hundreds of times more energy efficient than conventional approaches.

Kathy Yelick, Associate Laboratory Director for Computing Sciences, views Shalf's appointment as an important step in combining the organization's research and facilities expertise

to address the challenges to continue growing computing performance for science. Computing Sciences, which comprises NERSC, CRD, and ESnet, provides the high performance computing and networking resources and the applied math and computer science tools and technologies needed to advance computational science.

CRD Director David Brown said Shalf's dual role will promote closer ties between NERSC and CRD to support Computing Sciences' strategic efforts for exascale computing and science engagement at the lab. "As computation plays an increasingly important role in research from accelerators and astrophysics to materials and genomics, John's expertise and experience in both applications and software should lead to greater strategic collaboration between NERSC, CRD, and the other scientific divisions at Berkeley Lab," Brown said.

Shalf is a co-author of over 60 publications in the field of software frameworks and HPC technology, including three best papers and the widely cited report "The Landscape of Parallel Computing Research: A View from Berkeley" (with David Patterson and others), as well as "ExaScale Software Study: Software Challenges in Extreme Scale Systems," which sets the Defense Advanced Research Project Agency's (DARPA's) information technology research investment strategy for the next decade. He was a member of the Berkeley Lab/NERSC team that won a 2002 R&D 100 Award for the RAGE robot.

Shalf received a bachelor's degree in electrical engineering and a master's degree in computer engineering from Virginia Polytechnic

Institute and State University. Before joining Berkeley Lab in 2000, he was a research programmer at the National Center for Supercomputing Applications at the University of Illinois and a visiting scientist at the Max-Planck-Institut für Gravitationsphysik/Albert Einstein Institute in Potsdam, Germany, where he co-developed the Cactus code framework for computational astrophysics.



// Chief Technology Officer John Shalf

Computational Systems and Facilities

Cray Delivers Phase 1 of “Edison” XC30 System

On June 27, 2012, NERSC announced that Cray would install its next-generation supercomputer system. Consisting of products and services, the multi-year, multi-phase project is valued at more than \$40 million. It includes delivery of a future-generation Cray supercomputer code-named “Cascade” and a next-generation Cray Sonexion storage system.

Phase 1 of NERSC’s newest supercomputer, named Edison after U.S. inventor and businessman Thomas Alva Edison, was delivered on November 27, 2012. When completed, the new system will deliver a peak performance of more than 2 petaflops, equivalent to more than two-quadrillion calculations per second. The full system is expected to go into production in 2013. NERSC is also home to a Cray XE6 supercomputer, named “Hopper.”

“From energy efficient batteries to climate change, NERSC’s 4,500 users are tackling problems that are of vital importance

to our nation’s competitiveness and sustainability, so it is critical that our next system, NERSC-7, deliver readily accessible performance on real-world applications,” said Kathy Yelick, Associate Laboratory Director of Computing Sciences at Berkeley Lab.

According to Yelick, it is also important that NERSC provide supercomputing resources to users in an energy efficient manner, and she said the new Cray system will enable many pioneering features on this front, including the ability to run year-round using “free cooling” at the NERSC site.

“This approach utilizes water from cooling towers only, not mechanical chillers, to provide exceptional energy efficiency,” said Jeff Broughton, head of NERSC’s System’s Department. “The moderate Bay Area climate combined with Cray’s new design will allow us to keep power for cooling to less than 10 percent of the power used for computing.”



// Phase 1 of NERSC’s Cray XC30 system.



// Over its five-year lifetime, Franklin has delivered 1.18 billion processor hours to scientific research.

The Cray XC30 is the next step in Cray's Adaptive Supercomputing vision. The system features major advancements of the Cray Linux Environment, Cray's high performance computing (HPC) optimized programming environment, and the next-generation Aries interconnect chipset. It also features support for Intel Xeon processors—a first for Cray's high-end systems. The Cascade supercomputer is in part made possible by Cray's participation in the Defense Advanced Research Project Agency's (DARPA) High Productivity Computing Systems program.

The Cray Sonexion storage system delivered to NERSC will scale to more than 6 petabytes of usable storage and more 140 gigabytes per second of sustained aggregate I/O (input/output) performance. Sonexion brings together an integrated file system, software, and storage offering that has been designed specifically for a wide range of HPC workloads, providing users with an integrated, scalable Lustre solution that is easy to install and maintain. Cray's Sonexion storage system combines powerful servers, the latest Lustre parallel file system, and efficient management

// From left, Berkeley Mayor Tom Bates, Office of Science Director Bill Brinkman, UC Berkeley Chancellor Robert Birgeneau, Energy Secretary Steven Chu, Lab Director Paul Alivisatos, UC President Mark Yudof, and Associate Lab Director for Computing Sciences Kathy Yelick.

software into a modular and scalable storage product that is tested at scale, and supported as a complete solution by Cray.

After Five Years, NERSC's Franklin Retires

In May 2012, NERSC retired one of its most scientifically prolific supercomputers to date—a Cray XT4 named Franklin, in honor of the United States' pioneering scientist Benjamin Franklin.

From climate to energy sciences, Franklin has tackled some of the most challenging scientific problems facing humanity. In November 2008, it was the world's seventh most powerful supercomputer, according to the TOP500 list. By 2009, Franklin underwent a significant upgrade that doubled its number of processor cores (to 38,288) and increased its theoretical peak performance to 352 teraflops—three and a half times that of the original system. Over its five-year lifetime, Franklin has delivered 1.18 billion processor hours to scientific research. Its last listing on the TOP500 list placed it as the 38th most powerful supercomputer in the world.

In addition to its scientific accomplishments, Franklin was also a trailblazer in supercomputing software. Franklin was the first Cray system to fully utilize the Cray Linux Environment (CLE) operating system. Today, all Cray systems run CLE.

"Franklin has proven to be a tremendous benefit to the DOE computational science community," said Kathy Yelick. "From the day it went online, Franklin has been a scientific workhorse for many important computational science projects. Now that the system is ready to retire, we are all grateful for its contributions."

Breaking Ground on Computational Research and Theory Facility

Department of Energy Secretary Steven Chu, along with Berkeley Lab and University of California leaders, broke ground on the Lab's Computational Research and Theory (CRT) facility



on February 1, 2012. The CRT will be at the forefront of high-performance supercomputing research and will be DOE's most efficient facility of its kind.


Joining Secretary Chu as speakers were Berkeley Lab Director Paul Alivisatos, University of California President Mark Yudof, the Energy Department's Office of Science Director Bill Brinkman, and UC Berkeley Chancellor Robert Birgeneau. The festivities were emceed by Associate Lab Director for Computing Sciences Kathy Yelick. The speeches were followed by the symbolic "shovel and hard hat" photo opportunity.

The CRT will include a 140,000 gross square foot, \$145 million computer facility and office structure with associated infrastructure. The facility will accommodate up to 300 staff and will bring together under one roof for the first time these world-class programs within the Computing Sciences organization of Berkeley Lab:

- The National Energy Research Scientific Computing Center (NERSC) is the Department of Energy's most scientifically productive supercomputing center, serving more than 4,500 researchers who use NERSC to generate around 1,900 papers a year.

- The Computational Research Division (CRD) at Berkeley Lab fosters scientific breakthroughs in applied computer science R&D, data management and analytics, networking and distributed computing, and applied mathematics.
- The Scientific Networking Division promotes global and scientific collaborations through its operation of the Energy Department's Energy Sciences Network (ESnet), which provides high-bandwidth connections for researchers worldwide to work together on such problems as global climate change, developing fusion energy, and expanding our understanding of the universe.
- Berkeley Lab Computing Sciences has joined with the University of California at Berkeley to create a program called Computational Science and Engineering at Berkeley (CSE). This program trains Ph.D. students in the increasingly important fields of large-scale simulation and analysis of large data sets. CSE also serves students at UC campuses in Davis, Merced, and Santa Cruz. The training lectures are freely available online.

Secretary Chu also used the occasion to highlight President Obama's State of the Union address and the administration's commitment to America's leadership in scientific innovation, particularly in the area of energy



// Figure 1. Library usage on Hopper, June 21, 2012 to January 17, 2013.

Innovations

Adoption and Integration of Externally Developed Tools

Four of NERSC's innovations in 2012 came from implementing and augmenting tools that were developed at sister HPC computing centers:


- the Automated Library Tracking Database developed by staff at the Oak Ridge Leadership Computing Facility (OLCF), the National Institute for Computational Sciences (NICS), and the Swiss National Supercomputing Centre (CSCS)
- the Darshan I/O profiler developed at Argonne National Laboratory (ANL)
- the Lustre Monitoring Tool (LMT) that was developed at Lawrence Livermore National Laboratory
- Globus Online, developed by the international Globus Alliance.

The Automated Library Tracking Database (ALTD) is a tool that transparently tracks library and third-party application usage on the Cray systems. Understanding which libraries are heavily used and which are not is important for determining how NERSC spends a limited software budget and allocates support. NERSC staff have worked closely with staff from Cray, OLCF, and NICS to implement this tool on Hopper and have benefitted from the

collaborative effort. Figure 1 was created using library data captured by the ALTD tool over six months and shows the top libraries used at NERSC ordered by number of unique users. The y-axis is on a log scale and shows the importance of key tools such as MPICH, parallel-I/O, and math libraries to the Department of Energy Office of Science workload.


The Darshan tool is a scalable HPC I/O characterization tool developed by ANL for the BlueGene/P Intrepid system. Through collaborations with ANL and ALCF staff, NERSC has ported and deployed Darshan on the Hopper Cray XE6 system and the Edison Cray XC30 system. The Darshan tool captures application I/O calls and records them in a central log file. To use Darshan, users must recompile applications. After initial testing by a small number of NERSC users with no performance degradations observed, on November 15, 2012 the Darshan library was loaded onto Hopper as a default software module, meaning that any user who recompiled his or her application on Hopper would have the I/O in their application profiled.

NERSC went one step farther by parsing the Darshan logfile output and putting it into a database that could be easily queried to show Darshan I/O statistics on the NERSC website. Now, every time a user runs an application that has been profiled by Darshan, the results are readily available from the popular "Completed Jobs" web page (Figure 2). The web page reports information such as the I/O read and write rate in MB/sec, the amount of I/O read and written and the size of I/O transactions, and the percentage of time spent in I/O. Understanding these metrics is key to improving an application's I/O performance.



// Figure 2. An example of the Darshan I/O profiling information available to NERSC users on the Completed Jobs web page.

The Lustre Monitoring Tool (LMT) monitors Lustre File System servers (MDT, OST, and LNET routers) and collects server-side data that characterizes the current load on the file system. The original implementation of LMT allowed capture of the total amount of data read and written and the current bandwidth utilization. NERSC developed an add-on to LMT that allows recording of the size of each individual I/O transaction. This additional feature was key to enabling the determination of the load on the file system, a quantity that was not previously measurable (Figure 3). This innovation will allow NERSC to understand in much greater detail the characteristics of the I/O workload, allowing us to make more informed



// **Figure 3.** The load (black), read (red), and write (blue) bandwidths for two different days on the Hopper scratch file system. In (a) the bandwidth measurements and the load are high. In (b) note that although the bandwidths are low, the load on the file system is in fact high due to many small I/O transactions. Before this enhancement to LMT developed by NERSC, we would have wrongly assumed that (b) was a day where the file system was unloaded.

future purchasing decisions as well as to improve the quality of service delivered to NERSC users in an operational setting.

Globus Online and HPSS. Archiving data has traditionally involved running command-line tools, but several NERSC users have requested a graphical user interface (GUI) to move data between HPSS and compute systems. A GUI makes it easier to back up or restore relevant data. Using Globus Online, along with a special HPSS-enabled GridFTP server, NERSC has deployed a drag-and-drop archiving capability. The NERSC HPSS Archive system is available on the Globus Online software-as-a-service web offering. With the service, NERSC HPSS users can use the HPSS endpoint to drag and drop files between the HPSS system and another endpoint in the service directly from their web browsers.

At GlobusWORLD 2011, held at Argonne National Laboratory, NERSC was awarded “Most Innovative Use of Globus Online” for using it to archive data onto HPSS.

Enabling Web-Based Access to Computation and Data

NEWT API

Increasingly we all expect to have web-based interfaces to the resources and information on which we depend. HPC is no different, and NERSC is working to improve accessibility and broaden the impact of its resources by embracing web-based HPC. Recognizing the need to make web-based application programming interfaces (APIs) for science, NERSC developed the NEWT (NERSC Web Toolkit) API, which brings high performance computing to the web through easy-to-write web applications. NEWT makes HPC resources accessible and useful to the next generation of scientists, who are more comfortable with the web than they are with back-end command-line tools.


NEWT is a web service that allows researchers to access resources at NERSC through a simple ReSTful API. The NEWT API and web service lets them interact with NERSC through simple HTTP URLs and commands. NEWT responds to client requests using JSON (JavaScript Object Notation), a lightweight data-interchange format that is easy for humans to read and write and easy for machines to parse and generate. This makes it very easy to build powerful web applications that can interact with NERSC HPC resources using HTML and JavaScript. NEWT enables access to all major compute and data systems, and supports features like authentication, job submission, file transfer, system status, NIM (NERSC Information Management) accounting information, interactive commands, and persistent object storage.

NEWT has spurred the growth of several web-based tools and Science Gateways to access NERSC in recent years, including the Daya Bay Offline Data Monitor, the Materials Project, the 20th Century Climate Reanalysis Portal, the Gauge Connection, the NERSC mobile app, and the MyNERSC tool. In 2012, the ALS Portal and OpenMSI, both described below, joined the growing list of web-based tools for scientific computing.

NERSC Mobile Apps

In an effort to make NERSC resources more accessible to its users, the staff are rolling out a number of applications that allow researchers to access scientific data on their web browsers, tablets, and smart phones. In April 2012, NERSC announced two new applications now available to its users:

The NERSC mobile user portal (<http://m.nersc.gov>) allows researchers to check the current status of NERSC systems and user Message of the Day (MOTD), as well as log into their



// Figure 4. Growing ALS data rates.


account on their mobile phones to view recently completed, queued, and running jobs.

The NOVA portal (<https://portal-auth.nerisc.gov/nova>) is an experimental web application that allows licensed VASP users to submit jobs to NERSC systems. VASP (Vienna Ab initio Simulation Package) is a computer program for atomic-scale materials modeling. Based on user feedback with this early system, the team hopes to provide additional features in the future.

NERSC is continuing to improve these applications and create more products for the web.

ALS Portal


The Advanced Light Source (ALS) at Berkeley Lab is a third-generation synchrotron and DOE Office of Basic Energy Sciences (BES) user facility that serves 10,000 scientists and engineers every year from a variety of different fields using a variety of different x-ray probes and detectors. Several factors are contributing to higher data rates at light sources like the ALS: increased light source luminosity, increased detector resolution, and sample automation (allowing for more samples to be studied in a given time). The projected ALS data rates are growing exponentially and expected to be above 1 PB in 2013 (Figure 4).




// Figure 5. An example of 3D tomographic reconstruction.

The amount of data being collected creates a need that exceeds the ability of the computational resources available on site to archive and analyze the data. In particular, for the case of tomographic data, analysis and reconstruction of raw images presents a serial bottleneck when taking data, requiring expert intervention and limiting the use of the facility (Figure 5).

For this reason, NERSC has engaged ALS scientists and set up working data and analysis pipelines from the ALS through ESnet to NERSC's compute and archival systems (Figure 6). The data pipeline consists of data transfer nodes (DTNs) at both the ALS and NERSC, a data-suitcasing model, new APIs at NERSC, and improvements to ALS analysis and simulation software enabling work to be performed on NERSC compute machines. The data suitcasing model consists of organizing data sets into single HDF5 files with associated metadata (also populated into a MongoDB database) for high performing I/O during analysis.



// Figure 6. ALS data pipeline to NERSC.



// Figure 7. ALS Data and Simulation Portal sample screens. (Data and images, Hannah Wood, ALS)

The web APIs constructed at NERSC use Python Django, HDF5, and MongoDB libraries and custom HPSS tools to manage the ALS data at NERSC and expose it to ALS users via a web user interface.

The web user interface, built with the custom ALS API in conjunction with NEWT, currently allows users to search for datasets, retrieve the sets (from disk or tape), and perform simulations (Figure 7). In the near future, real-time analysis will be enabled on multiple beamlines.

OpenMSI

OpenMSI (Open Mass Spectrometry Imaging) is a startup effort at NERSC that aims to provide a web-based gateway for management and storage of MSI data, the visualization of the hyper-dimensional contents of the data, and the statistical analysis of the data. As part of this project, the scientists also plan to investigate the use of HPC to speed up a select set of compute-intensive analysis tasks, such as peak detection and filtering. In 2012 they created the first prototype web-based viewer for the mass spectrometry imaging data running at NERSC on the Science Gateway nodes.


Supporting Data-Intensive Computing and New Workloads

NERSC Global Filesystem Played a Key Role in Discovery of the Last Neutrino Mixing Angle

Discovery of the last neutrino mixing angle—one of Science magazine's top ten breakthroughs of the year 2012—was announced in March 2012, just a few months after the Daya Bay Neutrino Experiment's first detectors went online in southeast China. Collaborating scientists from China, the United States, the Czech Republic, and Russia were thrilled that their experiment was producing more data than expected, and that a positive result was available so quickly.

But that result might not have been available so quickly without the NERSC Global Filesystem (NGF) infrastructure, which allowed NERSC staff to rapidly scale up disk and node resources to accommodate the surprisingly large influx of data.

NGF, which NERSC started developing in the early 2000s and deployed in 2006, was designed to make computational science more productive, especially for data-intensive projects like Daya Bay. NGF provides shared access to large-capacity data storage for researchers working on the same project, and it enables access to this data from any NERSC computing system. The end result is that scientists don't waste time moving large data sets back and forth from one system to another—which used to be the case when each computer had its own file system. NERSC was one of the first supercomputer centers to provide a center-wide file system.



// Daya Bay Neutrino Facility in China.

The Daya Bay experiment had been expected to produce large data sets—after all, the nuclear reactors of the China Guangdong Nuclear Power Group at Daya Bay and nearby Ling Ao produce millions of quadrillions of electron antineutrinos every second. However, neutrinos and antineutrinos are very difficult to detect—they pass through most materials, even our bodies, without leaving a trace. Nevertheless, the Daya Bay physicists expected the eight massive detectors buried in the nearby mountains to produce enough high-quality data to measure, for the first time, θ_{13} (pronounced “theta-one-three”)—the last unobserved parameter of neutrino oscillation, and a key to understanding the difference between matter and antimatter.

Craig Tull, head of Berkeley Lab’s Science Software Systems Group and U.S. manager of the Daya Bay experiment’s overall software and computing effort, had been working with NERSC and ESnet staff to ensure that data coming from Daya Bay could be processed, analyzed, and shared in real time, providing scientists immediate insight into the quality of the physics data and the performance of the detectors.

The plan was this: Shortly after the experimental data was collected in China, it would travel via the National Science Foundation’s GLORIAD network and DOE’s ESnet to NERSC in Oakland, California. At NERSC the data would be processed

automatically on the PDSF and Carver systems, stored on the PDSF file system (a dedicated resource for high energy and nuclear physics), and shared with collaborators around the world via the Daya Bay Offline Data Monitor, a web-based Science Gateway hosted by NERSC. NERSC is the only U.S. site where all of the raw, simulated, and derived Daya Bay data were to be analyzed and archived.

The first six detectors to go online performed beyond expectations. Between December 24, 2011, and February 17, 2012, they recorded tens of thousands of interactions of electron antineutrinos, about 50 to 60 percent more than anticipated, according to Tull—about 350 to 400 gigabytes (GB) of data per day. “There were three things that caught us by surprise,” he says. “One was that we were taking more data than projected. Two was that the initial processing of the data did not reduce the volume as expected. And three was that the ratio of I/O operations to CPU operations was much higher for Daya Bay analysis than, say, a simulation on a NERSC computer. The combination of those three things created larger than anticipated I/O needs.”

In short, Tull was worried that the PDSF file system did not have the space and bandwidth to handle all the incoming data. So in January 2012 he met with Jason Hick, leader of NERSC’s Storage

Systems Group, to find a solution. “NGF was already running at a petabyte scale,” Hick says, “and Craig’s requirements were largely satisfied by simply scaling up the current system. It’s exciting to see NERSC’s vision for NGF realized when we can directly help scientists like Craig achieve their aims.”

In less than a week, the Storage Systems team, including Greg Butler and Rei Lee, added more disk space to NGF and updated its configuration, then worked with Tull to reroute the daily Daya Bay data runs to NGF. As a result, processed data was made available to Daya Bay collaborators within two hours after it was collected at the detectors. Only two months later, in March 2012, the Daya Bay Collaboration announced that θ_{13} equals 8.8° , outpacing other neutrino experiments in France, South Korea, Japan, and the U.S. (see page 36 for science highlight).

In May 2012, under NERSC’s new Data Intensive Computing Pilot Program (see below), Daya Bay was one of eight projects to receive an award of resources, including disk space, archival storage, and processor hours, to improve data processing performance in anticipation of the experiment going into full production, which occurred in October 2012 when the last two detectors went online. The experiment will continue taking data for another three years to improve the accuracy of the results and to do a variety of analyses.

By December 2012, the project had accumulated half a petabyte of data, and it was time to reprocess all of the data using the latest refinements in software and calibrations, in preparation for a January 2013 collaboration meeting in Beijing. Tull consulted with Iwona Sakrejda of NERSC’s Computational Systems group and received a temporary loan of 1,600 cores for processing the data.

“We were predicting that it would take three to four weeks to do the whole analysis,” Tull says, “but with Jason and Iwona’s help, we were able to do it in eight days instead of three to four weeks. It was the combination of NGF and the extra cores which shortened the time to finish production.”

The Daya Bay story is typical of what NERSC expects to see more of in the future: data-intensive science that requires high-bandwidth, high-volume, and low-response-time data systems.

“For the last four years, NERSC has been a net importer of data,” says NERSC Division Director Sudip Dosanjh. “About a petabyte of data is typically transferred to NERSC every month for storage, analysis, and sharing, and monthly I/O for the entire center is in the 2 to 3 petabyte range. In the future, we hope to acquire more resources and personnel so that science teams with the nation’s largest data-intensive challenges can rely on NERSC to the same degree they already do for modeling and simulation, and the entire scientific community can easily access, search, and collaborate on the data stored at NERSC.”

NERSC Launches Data Intensive Pilot Program

In April 2012, NERSC launched a new initiative to support DOE-relevant, data-intensive science pilot projects for up to 18 months.

“NERSC has long understood the importance of data intensive science and has supported the analysis of data streams from telescopes, detectors, and sequencers in addition to data coming from simulations run at NERSC,” said Kathy Yelick, Associate Laboratory Director for Computing Sciences at Berkeley Lab. “DOE has unique data challenges arising from their large experimental facilities.”

While all applications were considered, NERSC was particularly interested in supporting experiments that are generating data at rates beyond their current analysis capabilities. “Many of these fields are generating data at increasing rates and struggling to marshal sufficient resources to realize the full potential from new instruments and detectors,” said Shane Canon, who heads NERSC’s Technology Integration Group. “This effort is aimed at addressing that gap.”

“Many of the big data challenges that have long existed in the particle and high energy physics world are now percolating other areas of science. At NERSC we’ve seen an increase in user requests for more computing resources and bigger storage allocations to deal with bigger datasets,” said David Skinner, who leads NERSC’s Outreach, Software and Programming Group. “So the goal of this pilot is to see how new technologies and software environments will help these scientists better manage, analyze, store and share their growing datasets.”

Those selected for the pilot program received access to large data stores (up to 1 petabyte of disk and tape storage), priority access to a 6 terabyte flash-based file system (with 15 gigabits per second transfer speeds), and priority access to Hadoop-style computing resources on NERSC’s Carver Infiniband cluster (with access to a 1 TB memory node). They may also use NERSC’s Science Gateways for web access.

“We’ve seen overall science data traffic growing at a rate of 70 percent per year since 1990. We expect the trend to continue, and even accelerate, in coming years,” said Greg Bell, acting director of ESnet. “The challenge of getting that data from instruments to analysis—or even from scientist to scientist—shouldn’t be underestimated.”

Data transfers to and from NERSC can take advantage of the center’s high-speed connectivity to ESnet. NERSC to ESnet connectivity was upgraded from 20 gigabits per second to 100 Gbps by the end of 2012. In addition to an array of storage, computing and data-transfer resources, awardees were also assigned a staff member to support and advocate for each project within NERSC.

Eight projects were selected for the Data Intensive Computing Program:

1. **High Throughput Computational Screening of Energy Materials**
Gerbrand Ceder, Massachusetts Institute of Technology
2. **Analysis and Serving of Data from Large-Scale Cosmological Simulations**
Salman Habib, Argonne National Laboratory
3. **Interactive Real-Time Analysis of Hybrid Kinetic-MHD Simulations with NIMROD**
Charlson C. Kim, University of Washington
4. **Next-Generation Genome-Scale In Silico Modeling: The Unification of Metabolism, Macromolecular Synthesis, and Gene Expression Regulation**
Bernhard Palsson and Joshua Lerman, University of California, San Diego
5. **Transforming X-Ray Science toward Data-Centrism**
Amedeo Perazzo, Stanford University
6. **Data Processing for the Daya Bay Reactor Neutrino Experiment's Search for θ_{13}**
Craig Tull, Lawrence Berkeley National Laboratory
7. **Data Globe**
John Wu, Lawrence Berkeley National Laboratory
8. **Integrating Compression with Parallel I/O for Ultra-Large Climate Data Sets**
Jian Yin, Pacific Northwest National Laboratory

More information on these projects is available at <https://www.nersc.gov/users/accounts/awarded-projects/data-2012-awards/>.

Parallel Databases Testbed

This project deployed two emerging scalable parallel database systems, SciDB and Aster Data, on hardware called Jessup that is a part of the Carver testbed nodes. SciDB is an open-source analytical database oriented toward the data management needs of scientists. Aster Data's nCluster is a massively parallel processing analytic database management system.

In order to evaluate how scientists might take advantage of these parallel databases, we created subprojects by pairing up with science teams to help port and analyze their data. These science projects are expected to produce and analyze terabytes of data in a short term. Three such collaborations are described here:

1. **Analysis of 1D spectra.** The data contains millions of simulation-generated 1D spectra (wavelength vs.

magnitude); 1.5 million spectra have a size of around 200 GB. A set of metadata (simulation parameters) is associated with each spectrum. When a user uploads an observed spectrum from the web portal, the system is required to return a simulation spectrum from the database that best fits the given one, with certain restrictions on the metadata. SciDB can calculate all 1.5 million chi-square statistics, sort the result, and return the best match within 250 seconds. This enables web users to run this operation on the dataset at near real time without downloading the data. Traditional SQL databases are not able to handle this type of workload at scale.

2. **Palomar Transient Factory.** In the PTF catalogs of observed objects, such as stars and galaxies, each entry contains the position of the object and a set of properties such as the magnitude. There can be up to 1 billion objects in each catalog (300–400 GB in size). The goal is to cross-match two catalogs (e.g., objects originating from two telescopes), each with one billion objects, and determine the subset of all objects that appear in both catalogs at the same part of the sky. Two objects are considered potentially identical if they appear at approximately the same position in the sky given some uncertainty in their measured positions. The spatial query (for example, return all the stars within a certain region of the sky) is an expensive operation that greatly benefits from special indexing mechanisms integrated into the database. With SciDB, cross-matching two catalogs of 1 billion objects each can be done within five minutes, which is 50 times faster than PostgreSQL.

3. **Extreme weather simulation.** The CAM5 climate simulation output is a 3D array of variables such as surface temperature at a given longitude, latitude, and time. For each variable, one batch of simulation data amounts to ~200 GB. The goal of this work is to calculate the 97th percentile of surface temperature on each grid location (Long-Lat) over time, and save only the top 3 percent of the values to a file for plotting. With SciDB, the sort/percentile and filtering operation can be achieved with a single query. For the test dataset of 200 GB, ~60,000 sort/percentile operations take 20 minutes in parallel, and filtering the data takes only seconds.

By working with science teams, we identified two types of uses for such a parallel database system: (1) as a cluster for interactive analysis of terabytes of data, and (2) as a powerful back end for web portals, giving real-time response to complicated queries that go through terabytes of data.

The results to date highlight the potential value of these technologies to data-intensive science and data exploration.

This project also illustrates the value of operating testbed platforms to evaluate emerging technologies. The next steps in this project are to continue working with the selected science teams to learn more, and then to decide on a solution and configuration for general production use at NERSC.

Supporting Serial and High-Throughput Workloads on the Cray Systems

The XE6 architecture and runtime environment have been optimized for tightly coupled applications with MPI use, and we typically have policies in place to allow users to efficiently schedule and run jobs that request a large number of processors. However, these policies can penalize users whose science requires running throughput-oriented workloads. We have investigated a number of ways to allow throughput workloads to run efficiently on the Cray systems. These include (1) providing alternative frameworks (such as a TaskFarmer approach, Hadoop on Demand, or even a personal scheduler instance) that run within the constraints of the workload management system and runtime environment of the XE6, but allow users to efficiently execute throughput workloads; and (2) providing a throughput-oriented scheduling environment in parallel with the more traditional scheduling mechanism.

The TaskFarmer approach to supporting a throughput workload involves running a single large parallel job within the constraints and policies of the scheduler on the Cray. Within this job, we use an internal mechanism to launch tasks on the allocated resources, using a client-server method.

Hadoop on Demand is implemented as a tool called “MyHadoop,” which allows users to instantiate a private Hadoop cluster, requested as a single parallel job through the standard batch system on the Cray. To ensure this works, we had to utilize the Cray’s scratch Lustre file system to replace HDFS, allowing for a pseudo-local file system for each compute node. We also had to ensure that the compute nodes had a complete Linux runtime environment, which was achieved by utilizing the Dynamic Shared Library (DSL) environment that Cray provides.

MySGE allows users to instantiate a virtual private cluster (VPC) on the Cray by submitting a request for a large parallel job but starting up a private instance of the GridEngine scheduling software within that job. The user can then submit a traditional throughput workload to this VPC scheduler instance.

One feature of the above methods is that they are all restricted to a single user’s workload. In order to support a more general throughput workload that is also transparent to multiple users, we can utilize the Cluster Compatibility Mode (CCM) to provision a separate set of resources for throughput computing. CCM builds on the DSL facility that Cray provides and essentially transforms individual compute nodes into cluster-like nodes familiar to users of Linux clusters. A designated special user

then requests a large parallel job, and starts up a batch client instance on each of the CCM compute nodes. Along with this, a batch server and scheduler is started up. Users are then able to submit jobs to this secondary batch system, without the constraints that the primary batch system has. Specifically, this allows us to share compute nodes among multiple users for serial jobs and to tune scheduling policies on the secondary scheduler for throughput-oriented jobs.

Operational Innovations to Enhance Efficiency and Security



Cluster Consolidation across Domains

NERSC strives for continuous improvement in operational efficiency. One area of focus is in the management of cluster resources. NERSC operates cluster systems that serve both the general NERSC users as well as specific communities. The latter include the users of PDSF, the Joint Genome Institute, and the Materials Project.

In 2012 NERSC needed to deploy new cluster resources for all of these communities. Rather than procure, deploy, and operate these resources individually, NERSC procured them under a common request for proposals and deployed a consolidated cluster named Mendel. Using a number of innovative practices, NERSC is able to efficiently provide systems that are tailored to the specific needs of these communities. This includes the use of the open source xCAT system for scalable cluster management, and the NERSC-developed CHOS system to provide tailored environments. CHOS is a software package that provides a mechanism for simultaneously supporting multiple Linux environments on a single Linux system.

While tools for cluster management have improved dramatically over the past decade, this continues to be an area of rapid evolution. NERSC has evaluated and used a number of cluster management tools through the years, including ones from Linux Networks, IBM, Bright, and homegrown tools. NERSC has recently adopted and customized xCAT for the majority of its self-supported clusters. Using xCAT enables NERSC to manage a large collection of resources in a scalable manner and easily re-provision resources.

NERSC first began to use xCAT on the PDSF system, continued its use on the Carver/Magellan systems, and is now expanding its use throughout the center. Specifically, all the cluster nodes run the same boot image, while domain-specific customization (such as the file systems that need to be mounted and the CHOS environments that are available) is done at boot time. We have also developed a system of add-ons to the images (which are extracted at boot time) that provide additional functionality (such as KVM, InfiniBand stack, OSG stack) without increasing the size of the base system image. In addition, the boot image is built on a dedicated-build virtual machine and



// Figure 8. Job placement within the Hopper torus network.

is automatically version-controlled, which allows for easy reversion to older images.

NERSC has also expanded the use of CHOS to allow us to provide customized environments to meet the needs of specific communities. CHOS was first developed and deployed in 2004 on PDSF to meet the growing need for specific Linux operating system versions for some projects. A growing number of projects were certifying their analysis tools on specific Linux distributions and versions. These tools were often dependent on specific compilers and library versions. Rather than partition resources for these specific projects and risk underutilization and balkanization, NERSC developed CHOS to enable users to select the required Linux environment. CHOS even allows applications running on the same host to see different environments. This was accomplished without the use of virtualization, which can have performance penalties and increases complexity.

While CHOS has been in use on PDSF for almost a decade, NERSC is now expanding its use to provide customized

environments for the JGI users and Materials Project. This has required testing with new OS environments based on Debian (PDSF primarily uses RedHat based distributions) and integrating with parallel MPI jobs. NERSC is evaluating other ways to extend CHOS to enable users to have even more flexibility over their environment. This could achieve many of the perceived benefits of cloud-based models while still providing easy access to shared resources like parallel global file systems.

100 Gbps Security Monitoring

NERSC has upgraded its border router capability to 100 Gbps. Accompanying this upgrade comes the need to monitor and protect NERSC from attacks at these greater bandwidths. While several well-understood technologies exist for addressing security issues at 10 Gbps, a number of significant hurdles exist at 100 Gbps that need to be addressed before effective monitoring actions can be provisioned. By studying traffic patterns in live traffic, we can take advantage of these inherent characteristics and develop a useful and effective framework for security monitoring.

NERSC has inserted an optical “tap” into its 100 Gbps uplink to capture both incoming and outgoing traffic from the center. The optical feeds from the tap are then fed into two 100 Gbps ports on an enterprise-class router. The router then reflects the traffic on these feeds out to a group of ports in parallel. These ports connect to a clustered IDS (intrusion detection system).

The implementation at NERSC leverages a combination of the heavy tail effect and bandwidth distribution behavior to reduce the traffic that an IDS system is required to inspect, while maintaining a high level of security monitoring effectiveness. Our key insight is that the largest flows tend to exist for significant periods of time, consume a significant portion of available bandwidth, and after an initial survey period have little or no security value. Having identified and characterized these large flows, the IDS can have the router remove these large flows from the forwarded data stream, thereby freeing up analysis resources.

The collection of optical tapping, splitting traffic into multiple individual feeds, the shunting of data-heavy flows and use of a clustered IDS system allow NERSC to monitor traffic at unprecedented levels. NERSC has what we believe to be one of the first 100 Gbps IDS systems in production.

Understanding the Impact of Job Placement on Application Runtime Variability

Because NERSC systems run applications of all sizes for varying lengths of time, an application’s placement within the Hopper

torus network can vary from run to run, depending on the nodes that are available. Different node allocations can lead to variations in an application’s runtime.

When we began exploring this issue in 2011, the Cray scheduler allocated nodes in a specified list order, starting from the beginning of the list and allocating a job to any subsequent free nodes in the list. Because of Hopper’s stability, running at times over a month between outages or maintenance, the node list could become very fragmented, and an application could be allocated to nodes spread across the torus network.

To help visually understand this behavior, a NERSC summer intern, using the open-source visualization package Blender, created a visualization tool to display how nodes were allocated in the network. Using one of NERSC’s Hopper benchmarks, the QCD application MILC, he ran the application a number of times, studying the way application placement on nodes affected runtime variability (Figure 8). This analysis led Cray to improve their algorithm for allocating an application across nodes, from using a simple first-free-node algorithm to one that attempts to group larger numbers of nodes together.

When NERSC users saw this analysis, they were very interested in the node placement for their own application. Our summer intern generalized his visualization tool for any application running on Hopper, and NERSC staff made the node placement visualization available to every user and every job that runs on Hopper. The output is available on the NERSC completed jobs page.

User Support and Outreach

NERSC offers a wide range of user support services. The User Services Group provides the direct interface between NERSC and its user community. The services and engagements through Users Services are many-fold, from resetting passwords to one-on-one consulting, from organizing Office of Science-wide requirements reviews to coordinating its active user group, NUG. NERSC’s website features a prominent “For Users” top-level section that makes it easy for users to find information; and our training activities provide additional opportunities for communicating with users.

NERSC’s consultants and account support staff are available to users via email, an online web interface, and on the phone during business hours (8 a.m. – 5 p.m. Pacific Time). Basic account support (password resets, resetting login failures) is available via the NERSC operations staff 24 x 7, 365 days per year.

NERSC’s 12 consultants include nine experts in high performance computing, six with Ph.D. degrees. NERSC’s two account support personnel each have 10+ years of experience. When users contact NERSC, they immediately are in contact with highly trained HPC specialists, who can usually solve issues directly, or immediately route the request to the appropriate systems engineers. NERSC’s policy is to respond to all inquiries within four business hours, and either solve the problem or communicate a work plan to the user within three business days for all reported incidents. The latter is a metric we report to DOE.

When users require more extensive assistance—such as programming or optimization help—they are routed to a consultant with the most appropriate skill set, who will engage in a longer-term collaboration with the user. Of course, with thousands of users and only a handful of consultants, extensive

one-on-one assistance is not always possible. NERSC's strategy is to deliver assistance at scale by heavily relying on extensive web documentation and an active training program. Likewise, many account support services that were once labor intensive (e.g., password resets, updating contact information) have been converted to self-service operations via the NERSC Information Management (NIM) web interface (nim.nersc.gov). These efficiency improvements result in fewer tickets per user, as shown in Figure 9.


NERSC also provides special assistance to strategic projects by making special staff assignments for limited periods. These deep engagements often have high payoffs. For example, NERSC was awarded a 2011 Scientific Computing Award for HPC Innovations Excellence for enabling the 20th Century Reanalysis Project, a reconstruction of historic weather and climate conditions. NERSC staff provided significant support that was needed to successfully transfer all of the dataset to NERSC, and collaborated with the science team on Science Gateway software to construct a portal that provides access to the data, which currently hosts the master copy of the entire dataset.

Other deep engagements leverage NERSC's extensive HPC technology expertise. For example, we have developed tools to help users create Science Gateways (web portals). These tools have contributed to NERSC hosting a number of prominent gateways for science teams, including those for the Materials Project, Deep Sky, the Coherent X-Ray Imaging Data Bank, and the Daya Bay Neutrino Detector. Work is under way to set up workflows to connect NERSC with the Advanced Light Source at Berkeley Lab and with users of the LINAC Coherent Light Source at SLAC.

The Data Analytics Team often spearheads some of our deep engagements as we endeavor to help our users get the most science out of their simulations or data analyses. A recent example involved the analysis of a leading plasma physics simulation with over a trillion particles on 120,000 Hopper cores (see page 32). This unprecedented simulation produced 30 TB of data per time step. Here the team worked with both NERSC and Cray staff to optimize the performance of HDF5 and MPI-IO to obtain peak I/O rates on 144 Lustre object storage targets (OSTs). In addition, leveraging our ties to Berkeley Lab's Computational Research Division, we were able to apply FastBit-based indexing/querying to the dataset. This allowed the scientists to query the indexed one trillion particle dataset in approximately three seconds.

End-to-End Network Tuning Sends Data Screaming from NERSC to NOAA

When it comes to moving large datasets between NERSC and his home institution in Boulder, Colo., Gary Bates is no slouch.



// Figure 9. Number of NERSC users and user tickets created per year.

As an associate scientist in the Earth System Research Lab of the National Oceanic and Atmospheric Administration (NOAA), Bates has transferred hundreds of thousands of files to and from NERSC, as part of a “reforecasting” weather forecasting project.


The “reforecasting” project, led by NOAA's Tom Hamill, involves running several decades of historical weather forecasts with the same (2012) version of NOAA's Global Ensemble Forecast System (GEFS). Among the advantages associated with a long reforecast dataset is that model forecast errors can be diagnosed from the past forecasts and corrected, thereby dramatically increasing the forecast skill, especially in forecasts of relatively rare events and longer-lead forecasts.

The GEFS weather forecast model used in this project is the same as that which is currently run in real time by the National Weather Service. In the reforecast project, GEFS forecasts were made on a daily basis from 1984 through early 2012, out to a forecast lead of 16 days. To further improve forecasting skills, an ensemble of 11 realizations was run each day, differing only slightly in their initial conditions.

In 2010, the NOAA team received an ASCR Leadership Computing Challenge (ALCC) allocation of 14.5 million processor hours on NERSC supercomputers to perform this work. In all, the 1984–2012 historical GEFS dataset now totals over 800 terabytes, stored on the NERSC HPSS archival system. Of the 800 terabytes at NERSC, the NOAA team sought to bring about 170 terabytes back to NOAA Boulder for further processing and to make it more readily available to other researchers. Because of the large quantity of data involved, having the data move as quickly and easily as possible is important, both on the network and at both end points in Oakland and Boulder.

Bates was able to bring over the bulk of the 170 terabytes earlier in 2012, using a machine that NOAA Boulder's Network Operations Center staff (BNOC) had temporarily set up as a Globus Online endpoint. Globus Online is a cloud-based tool

(a) Analyzed 24-h Accumulated Precipitation 1995 01 09
 (b) Forecast Precipitation 1-day fcst, 1995 01 09
 (c) Pr (precip > 50mm), 1-day forecast, 1995 01 09



// Figure 10. Reforecasting with historical data can improve the accuracy of forecasts: (a) 24-hour observed precipitation amounts for January 9, 1995; (b) average one-day precipitation forecasts; (c) a new forecast for the same day calibrated with old reforecasts and precipitation analyses. (NOAA Earth Systems Research Laboratory)

for high-performance data transfers. However, when the remainder of the data was ready to be moved in the summer of 2012, that machine was being used for other tasks and was no longer available.

When he tried to use an FTP server located behind NOAA's firewall for the remaining transfers, Bates discovered that data trickled in at about 1–2 megabytes per second. So Keith Holub of the BNOF staff set up a new, dedicated server with a data path unencumbered by legacy firewalls. The BNOF staff configured the new Globus Online endpoint node using information from ESnet's Fasterdata website (<http://fasterdata.es.net/>), a knowledge base of tips and tricks for speeding up end-to-end data transfers. This kind of configuration is an example of ESnet's Science DMZ model for high-performance systems supporting data intensive science. The change was instantly noticeable—much faster transfer rates than ever achieved previously were realized.

"Whoa! Transfer from NERSC to the BNOF data transfer node using Globus is screaming!" Bates wrote to his team and Eli Dart of ESnet. "I transferred 273 files with a total size of 239.5 gigabytes in just over 10 minutes. I calculate that's a rate of 395 megabytes per second. I've never gotten anything close to that before. Transferring the same 239.5 gigabytes from BNOF data node down to my local data storage is slower but still very good: it took about 81 minutes, or 49 MB/s."

Dart learned about Bates's data transfer challenges when Bates was trying to move data to NERSC, uploading the files from tape onto an FTP server. Dart points out that Bates's transfer rate adds up to more than 1 terabyte per hour.

"Now the system really rocks," Dart said. "NERSC has a well-configured data transfer infrastructure that operates very well. When all the right things are done at the other end, everything runs well—this is the way it's supposed to work."

Damian Hazen of NERSC's Mass Storage Group regularly helps users, including Bates, who are looking to move large datasets to and from NERSC's data archive as quickly and easily as possible. Because the center deals with a large number of users, NERSC staff are diligent about tuning local transfer nodes for the best performance, Hazen said. The staff also helps users try to track down the problem, which in Bates's case was the firewall at his institution. By understanding the characteristics of the file systems and the network, NERSC and ESnet staff have a pretty good idea of what the optimal performance should be, then work with users and other staff to get as close to that level as they can.

Another way to improve a user's overall performance is to help streamline the workflow. For example, it can take 90 seconds or more to locate and mount the archive tape holding the requested data. If several of the requested files are on the same tape, grouping the requests means the files are read consecutively. It seems like a small thing, Hazen said, but it can make a big difference if a lot of files are involved.

Hazen credits Bates with being eager to implement new tools, such as Globus Online. In fact, Globus Online named Bates the "user of the month" in October 2011 for being one of the biggest users in terms of the amount of data transferred. "In fact, with his work over the past two weeks, Gary now sits as the Top 2 user ever in terms of data moved," Globus Online noted in announcing the honor.

NERSC staff also helped Bates port and tune his application to speed up its performance on Franklin and Hopper, the center's Cray supercomputers, and Carver, an IBM system. Helen He in NERSC's User Services Group specializes in helping users who run climate and weather codes. In Bates's case, she rewrote some of the code so it could be ported from NOAA's IBM computer to the Cray systems at NERSC. In the process, she helped speed up the runtime for part of the code from 24 minutes to just a few seconds. She assisted in redesigning post-processing workflow for better throughput by working with the queue structures on Carver. The project also benefited significantly with the queue boost and dedicated compute nodes.


"By helping him improve the workflow of his application, we've gotten better throughput and faster runtimes," He said.

Bates agreed, saying "It's a very good system to use—the support is there and the people are very helpful."

Code and Runtime Performance Improvements

Improvements to BerkeleyGW Code

The BerkeleyGW package is a materials science application that calculates electronic and optical properties with quantitative accuracy, a critical need in materials design for more efficient and cost-effective solar light harvesting and energy conversion (Figure 11). Its usage at NERSC has grown such that it is now in the top 20 codes used by NERSC users based on wall hours.



// **Figure 11.** An example of systems of interest for study with BerkeleyGW: (A) the P3HT/PCBM interface, and (B) a Zn-phthalocyanine derivative dye molecule. Both can now be studied from first principles at NERSC due to BerkeleyGW improvements.

An optimized BerkeleyGW has been provided to NERSC users through a software module installed on the NERSC systems Hopper and Edison. In order to improve the scalability of the code and prepare the application for future architectures, OpenMP was added to the application by a NERSC consultant. In addition, the I/O for BerkeleyGW was transformed to use the parallel HDF5 library, which resulted in a performance improvement of 10x or more in common cases.

BerkeleyGW is one of many materials science codes whose development is directly or indirectly guided by NERSC's large user base and expert consultants.

Improving Shared Library Performance

The performance of dynamic shared libraries is crucial to some of the NERSC workload, especially for those large-scale applications that use Python as the front-end interface. Unfortunately, the shared library performance was so poor at large scale on Hopper that it took a long time to just start up the code. A user reported that the startup time of a Python application code, WARP (a code modeling high-intensity ion beams), increases linearly with the concurrency, so that if running at 40K concurrency, it would take five hours just to launch the code. To work around this, the WARP code developer had to build Python and all its dependent packages statically, which requires tremendous effort to maintain the code.

In collaboration with Cray, NERSC has provided a solution to this problem, which enabled the WARP code to start up in seven minutes at 40K concurrency on Hopper (Figure 1-10). The results can be applied to any projects using Python or other shared libraries on Hopper and Edison at large scale. This work was shared with the community at the Cray User Group CUG 2012 meeting.¹

Improving MPI-IO Performance from Hopper to GPFS-Based File Systems

In January 2012, NERSC formed the I/O Center of Excellence (COE) to collectively work to improve MPI-IO performance to GPFS (IBM) based parallel file systems. The GPFS client, unlike the Lustre file system client, is not natively supported on the Cray compute nodes, and instead goes through an I/O forwarding layer called the Data Virtualization Service (DVS). The DVS clients run on the compute nodes and forward I/O to a set of I/O nodes where the GPFS client is installed.

It was through this DVS layer that shared-file MPI-IO was losing performance. File-per-processor I/O on Hopper achieved a high percentage of peak file system performance, but MPI-IO shared file performance achieved only 7 percent of file-per-processor

¹ Zhengji Zhao, Mike Davis, Katie Antypas, Yushu Yao, Rei Lee, and Tina Butler, "Shared Library Performance on Hopper," CUG 2012, Stuttgart, Germany, April 29–May 3, 2012.

Show 25 ▼ entires
 Search:

#	Package	Platform	Category	Version	Module	Install Date	Date Made Default
1353	Quantum ESPRESSO	carver	applications/ materials sciences	4.1.2	espresso/4.1.2	2010-03-13	2010-03-14
1354	Quantum ESPRESSO	carver	applications/ materials sciences	4.2.1	espresso/4.2.1	2010-12-01	2010-03-22
1355	Quantum ESPRESSO	carver	applications/ materials sciences	4.3.2	espresso/4.3.2	2012-01-10	
1356	Quantum ESPRESSO	carver	applications/ materials sciences	5.0.0	espresso/5.0.0	2012-05-29	
1357	Quantum ESPRESSO	edison	applications/ materials sciences	5.0.2	espresso/5.0.2	2013-01-22	2013-01-22
1358	Quantum ESPRESSO	hopper	applications/ materials sciences	4.2.1	espresso/4.2.1	2010-11-01	2010-11-01
1359	Quantum ESPRESSO	hopper	applications/ materials sciences	4.3.2	espresso/4.3.2	2012-01-10	
1360	Quantum ESPRESSO	hopper	applications/ materials sciences	4.3.2	espresso/4.3.2	2012-04-04	2012-04-04
1361	Quantum ESPRESSO	hopper	applications/ materials sciences	4.3.2-2	espresso/4.3.2-2	2012-01-10	

// Figure 12. Sample web display of some currently installed software modules.

read performance and 16 percent of write performance. The Carver IBM iDataPlex system natively runs the GPFS client on compute nodes and demonstrated 92 percent of file-per-processor read performance and 65 percent write performance, showing that achieving a high rate is possible.

NERSC worked with Cray DVS and MPI-IO experts in an iterative process to tune the file system and application behavior to achieve higher I/O performance. The tuning parameters include number of DVS nodes used, the DVS block size, collective buffering, and collective behavior of the MPI-IO library itself.

Functionality Improvements

Implementing Cluster Compatibility Mode on Hopper

Cluster Compatibility Mode (CCM) is a new product developed by Cray that allows independent software vendor (ISV) Linux applications to run “out of the box” on Cray systems. NERSC staff were immediately interested in CCM because it can enable TCP/IP applications, which is one of our important workloads, on Hopper and Edison. Previously some applications, such as Gaussian, NAMD Replica Exchange, and WIEN2k, could not run on Hopper, as they require SSH or socket operations between

compute nodes, which were not supported. As a result, our TCP/IP users had to rely on our smaller Linux cluster, Carver, and suffered a long queue wait time, while other users could get a better queue turnaround on the larger system, Hopper.

To help these TCP/IP users, NERSC deployed CCM. After ironing out some initial issues, we did functionality and performance tests; built libraries, tools, and applications for CCM; and established a complete environment where users not only could run jobs “out of the box,” but also could compile, debug, profile, analyze, and visualize data. In April 2012, CCM was put into production on Hopper.

During the course of the year, 10 million MPP hours were used by more than 130 CCM users. Our effort allowed more of the NERSC workload to have access to NERSC’s largest resources and balanced out the long queue wait times on the Carver cluster. In addition, CCM is analogous to generic Linux clusters, which offered a chance to test a new support model, i.e., having one large machine support all our diverse workloads, which is potentially more cost-effective than maintaining several midrange clusters and one large system. This work was presented at the CUG 2012 meeting.²

² Zhengji Zhao, Yun (Helen) He, and Katie Antypas, “Cray Cluster Compatibility Mode on Hopper,” and Richard S. Canon, Jay Srinivasan, and Lavanya Ramakrishnan, “My Cray Can Do That? Supporting Diverse Workloads on the Cray XE-6,” CUG 2012, Stuttgart, Germany, April 29–May 3, 2012.

Infrastructure Improvements

In addition to providing direct support for NERSC user applications, NERSC staff continuously work to improve the infrastructure and efficiency of the NERSC center so more time can be devoted to helping users. In 2012, NERSC completed two improvement projects—software modules and Message Of The Day (MOTD)—that have improved our communication with users and the timeliness of information displayed on the NERSC website.

Software modules database modernization. NERSC has improved our mechanism for installing, documenting, and displaying supported software to the NERSC website. Previously, when a software package was installed by a NERSC staff member, the online web documentation needed manual updating, and the NERSC website often became out of date. NERSC improved the software modules infrastructure by versioning modules with SVN and implementing a commit hook to automatically reflect any changes in the repository into a database whose information is accessible via the web (Figure 12). This effort substantially improves the modules management process because it reduces the burden on staff to make cumbersome updates in multiple places, assures that the software list on the web is up to date for users, and creates a backup of module data in case recovery is necessary.

Improved MOTD database. Some of the most valued information NERSC provides to users via the web is the current status of NERSC resources and a list of past, current, and future planned outages. NERSC's main way of communicating this

information is via the MOTD system. Until recently, this system consisted of a manually edited block of ASCII text that was displayed on users' terminals when they logged in to a system and was also displayed in a preformatted text field on the NERSC website.

Over the past year, partially due to the proliferation of science gateways built with NEWT, the NERSC Web Toolkit, there has been a growing need among both users and staff to query a NERSC resource's status and outage log. For example, a common use case is the design of workflow pipelines that are aware of various NERSC resources' status and future outage times and that can automatically and dynamically route, plan, and, in the case of a machine suffering an unplanned outage, reschedule work appropriately.

The major limitation of the previous MOTD approach was that the valuable status and outage information was difficult to extract from the MOTD text; attempts to use string-parsing techniques to programmatically get a system's status or future outage information were prone to failure. For this reason, the MOTD system was retooled from the ground up, based on a more flexible database/web-form approach. The status and outage information is now stored in a queryable SQL database whose contents are produced by the Operations Group via a web form that is structured to ensure consistency and accuracy of the entered content.

The new approach allows us to display the MOTD information (or relevant subsets) in appropriately formatted ways for the web (Figure 13), terminal, and local applications, as well as to incorporate the information into dynamic science gateway and workflow software.

LIVE STATUS				
Current MOTD				
System	Status	Jobs Running	Cores in Use	Description/Notes
Carver:	Up	689	7456	
Dirac:	Up	16	128	
Edison:	Down			02/20/13 6:00-18:00 PST Scheduled maintenance. (02/08/13 16:02 PST) Edison is currently only available to users in the early testing program.
Genepool:	Up			
Hopper:	Up	274	90120	
HPSS backup:	Up			
HPSS User:	Down			02/20/13 8:45-13:00 PST Scheduled maintenance.
Jesup:	Up			
NGF:	Up			
PDSF:	Up			
Service Status				
All services are available.				
Planned Outages				
No Planned Outages				
Past 24 Hour Outages				
No Outages in Past 24 Hours				

Data Analytics Support

Analytics is the intersection of visualization, analysis, scientific data management, human-computer interfaces, cognitive science, statistical analysis, and reasoning. The primary focus of the Data Analytics Team is to provide visualization and scientific data management solutions to the NERSC user community to better understand complex phenomena “hidden” in scientific data. The responsibilities of the team range from applying off-the-shelf commercial software to developing advanced tools to realizing new solutions where none previously existed. The following are examples of collaborations and support provided to users in 2012.

// Figure 13. Sample dynamically generated web display of MOTD featuring system status and outages.

R and ImageJ/Fiji. R is a language and environment for statistical computing and graphics; ImageJ/Fiji is an image processing and analysis package. The Data Analytics Team provided maintenance, updates, and customization of these software packages for users on several NERSC machines. Some examples of science problems that were impacted include:

- stochastic population dynamics: this is a CSGF (Computational Science Graduate Fellowship) project to detect early warning signals of ecosystem collapse
- image analysis of experimental data that comes from imaging facilities such as the Advanced Light Source and the National Center for Electron Microscopy
- climate model simulations.

Parallel I/O for Trillion Particle Simulation. The Data Analytics Team provided support to a project studying the kinetic effects in plasmas of the planetary magnetospheres in running a trillion particle plasma physics simulation (see page 36). The simulation ran on 120,000 Hopper cores and produced 350 TB of data in the scratch file system. The Analytics Team members worked closely with other staff to schedule these runs and troubleshoot I/O performance. They were able to successfully demonstrate about 35 GB/s peak and about 25 GB/s average bandwidth utilization using parallel HDF5 and H5Part libraries. The Team also provided analysis and visualization of the project's data by creating subsets of energetic particles using FastBit, as well as a number of plots to examine the evolution of energetic particles and their properties.

The analysis led to a number of first-time scientific discoveries:

- This is the first self-consistent 3D simulation to demonstrate a power law in the spectrum of energetic particles.
- Aggyrotropy was discovered near the magnetic reconnection hot spot.
- The project discovered preferential acceleration of particles in a direction parallel to the magnetic field and confirmed their correlation with flux ropes.

3D Visualization of Sloan Digital Sky Survey (SDSS). The Data Analytics Team worked with the Sloan Digital Sky Survey (SDSS) collaboration to produce a stereoscopic 3D movie, "Fly Through the Known Universe." The movie used location data and galaxy photos to reproduce the real universe. The Carver cluster at NERSC was used to render this movie, which had its world-premiere screening at the 3D Film Festival in Los Angeles in September 2012. We have been invited to present the details of this work at the 2013 HPC User Forum.

Parallel Gaussian Processes. The Data Analytics Team worked with a climate research project to develop a "bigGP" package for R. The package can perform Gaussian processes (statistical estimators that infer the value of a random field at an unobserved location) on a large number of observational points and predict values with associated uncertainty bounds at unknown locations. The package uses ScaLAPACK-like operations to create covariance matrices and perform Cholesky factorizations and back solves in a completely distributed fashion. The package is also able to utilize ACML and MKL to speed up BLAS operations.

The package was used to fit generalized extreme value distributions to precipitation and temperature data and was run on concurrencies ranging from 1,000 to 10,000 cores on Hopper. It was used to check the behavior of CAM5 models in reproducing extreme rainfall and temperature patterns. This package was also applied to a supernovae spectrum dataset consisting of ~100,000 data points. We were able to fill gaps in the supernovae data and are working on tracking specific features in the spectra.

ParaView. The Data Analytics Team installed and provided user support for ParaView on NERSC's HPC systems through the 3.14 and 3.98 releases. We answered a number of user tickets and provided specialized user support to advanced users. In one case, working with an advanced user of Hopper, we developed extensions to ParaView to compute FTLE and map magnetic field topology to investigate the role that turbulence plays in driving magnetic reconnection. Using ParaView with their new extensions, we analyzed a 3072 x 3072 x 1024 plasma simulation dataset on Hopper using 4,096 cores. The scientists are currently successfully using the new analysis codes and ParaView on Hopper in their ongoing research.

VisIt. The Data Analytics Team installed VisIt on Carver, adding support to VisIt for launching jobs on Carver's large memory nodes. We also maintained the VisIt install on Hopper, adding improvements such as using the CRAY_ROOTFS environment variable to enable loading system libraries without copying them into the VisIt directory.

NERSC Participates in "Materials for Energy Applications" Workshop

From January 30 to February 1, 2012, Berkeley Lab hosted an invitation-only workshop on Materials for Energy Applications, which was jointly sponsored by all 17 DOE national laboratories. This three-day conference—the first of a planned series—was held to increase industry awareness of relevant research capabilities within the DOE national laboratory system, to deepen the national laboratories' understanding of the technical challenges facing industry, and to identify and improve paths forward for collaboration.

David Skinner of NERSC, who also heads the SciDAC Outreach Center, represented the Center with a poster on “Software Opportunities: Industry, ISVs [independent software vendors] and SciDAC.” The poster pointed out the breadth of the SciDAC software portfolio and how it can be leveraged in the manufacturing sector and the private sector more generally, increasing return on investment and decreasing time to solution.

Through a catalog of software that the center maintains and a list of “worked examples” from the five years of the SciDAC2 program, the poster is part of a larger effort to diversify the stakeholders in scientific software. “The impact of our investments in HPC software is measured in their end-to-end impacts over the long haul; they are not just about proof of principle,” said Skinner. “Finding ways for research-funded software to bring increased efficiency and innovation to a broad set of stakeholders beyond the SciDAC research community benefits all involved.”

“The past decades have seen a vigorous effort in bringing parallel computing to extreme problem sizes and extreme performance,” Skinner continued. “The time is ripe for the expertise developed within DOE/ASCR to be applied to the many parties who are newly interested in parallelism, whether that is motivated by the transition to multicore or the need to study systems of larger size. Expanding interest in parallel computing provides a way for the research community to see their work improved and sustained by that larger customer base.”

The last keynote address at the workshop was given by Michael McQuade, Senior Vice President for Science and Technology at United Technologies. The SciDAC Outreach Center and UTC won an IDC Innovation Award in 2011 for their application tuning work optimizing fuel spray nozzles to improve fuel efficiency. That outreach engagement yielded a 3.6x application speedup and a 66 percent reduction in nozzle design cycles. McQuade mentioned that he looks forward to the output of two similar joint projects that are currently under way.



// David Turner shows Lowell High School students around NERSC’s computer room.

Inspiring Careers in Science Research

In an effort to expose high school students to careers in research, Berkeley Lab’s Computing Sciences Diversity Outreach Program partnered with San Francisco’s Lowell High School Science Research Program, an after school program that aims to give highly motivated juniors and seniors a chance

to develop research projects with professional guidance with the intent to have the students enter the Intel Science Talent Search, a competition sponsored by Intel that offers college scholarships for outstanding scientific work.

As part of this collaboration, 32 Lowell students got a tour of NERSC in January 2012. Here, the students got to see Hopper—the world’s eighth most powerful supercomputer—and talk to the center’s system administrators, user consultants, and supercomputer analysts about their day-to-day work. Throughout the month of January, a number of Berkeley Lab Computing Sciences staff also trekked out to Lowell to talk about their research and career paths.

Elizabeth Bautista, who heads NERSC’s Computer Operations and Network Support Group, coordinated the Lowell High School partnership with Berkeley Lab Computing Sciences. Over time the plan is to expand the speakers and tours to provide the students with exposure to science areas across the Laboratory.

State Department’s TechWomen Visit NERSC

Thirty women from the United States Department of State’s 2012 TechWomen cultural exchange toured NERSC on September 27, 2012. Berkeley Lab researchers in the Computational Research, Physics, and Physical Biosciences divisions also hosted three TechWomen—from Algeria, Lebanon, and Tunisia—during the month of September.

“One of the goals of this program is to expose these women to organizations like Berkeley Lab, which have a wide range of expertise—in computer hardware and software, computational science, applied mathematics, as well as networking and middleware—and specialize in applying these skills to solve scientific problems, as well as support a national research community,” says Taghrid Samak, a postdoctoral researcher in Berkeley Lab’s Computational Research Division (CRD) and TechWomen mentor.

Launched by Secretary of State Hillary Rodham Clinton in 2011, TechWomen is an international exchange that uses technology as a means to empower women and girls in the Middle East and North Africa. The exchange builds on Clinton’s vision of “smart power” embracing the full range of diplomatic tools, in this case technology, to bring people together for greater understanding and to empower women and girls worldwide.

“TechWomen gave us a valuable opportunity to learn and collaborate with our American companions,” says Faïda Mansouri Joumâde, a TechWoman from Tunisia. “Although we come from different countries and cultures, as scientists we all share the same interests and concerns. For instance, issues like climate change and drought cross national

boundaries and affect us all.” Joumade holds an engineering degree in agronomy and specializes in soil science and geometrics. During her U.S. visit, she was mentored by Samak and Deb Agarwal in Berkeley Lab’s CRD.

“I had never seen a synchrotron light source before my visit to Berkeley Lab,” says Zeina Hobaika, a TechWoman from Lebanon. An Assistant Professor in the Faculty of Sciences at Saint Joseph University of Beirut, Hobaika holds a PhD in Structure, Function, and Proteins Engineering and specializes in anti-HIV drug design. Corie Ralston of Berkeley Lab’s Physical Biosciences Division mentored Hobaika during her stay in the U.S.



// Thirty women from the U.S. Department of State’s TechWomen cultural exchange toured NERSC.

“Professionally, I learned a lot from Corie. It was interesting to see how she interacted with her team, and we planned together some experiments that we will be applying to my research in anti-HIV therapies,” says Hobaika. “This was my first visit to the U.S., and I am leaving here with a new collaborator and friend. It was truly a once in a lifetime experience. I would like to thank all of the Berkeley Lab mentors for their generosity.”

“During my time here at Berkeley Lab, I attended seminars on campus, got to give a lecture about my research, and made some wonderful professional contacts,” says Essma Redouan-Salah, an Algerian high-energy physicist and Assistant Professor at M’sila University’s Faculty of Sciences. Marjorie Shapiro in Berkeley Lab’s Physics Division hosted Salah during her U.S. visit.

TechWomen brought a total of 41 women working in the technology sector from Algeria, Egypt, Jordan, Lebanon, Morocco, the Palestinian Territories, Tunisia, and Yemen to the U.S. in 2012 for a five-week mentoring program with their American counterparts. Later in the year, U.S. mentors traveled to the Middle East and North Africa to conduct workshops for women in the technology sector and young girls who have expressed an interest in pursuing a tech-based career.

TechWomen is an initiative of the U.S. department of State’s Bureau of Educational and Cultural Affairs (ECA). TechWomen is managed by the Institute of International Education (IIE) and implemented in partnership with the Anita Borg Institute for Women in Technology. For more information on TechWomen, see <http://www.techwomen.org/>.



// NERSC’s Elizabeth Bautista (center) takes the 2012 TechWomen on a tour of the facility’s computer room.

Research and Development by NERSC Staff

Staying ahead of the technological curve, anticipating problems, and developing proactive solutions are part of NERSC's culture. Many staff members collaborate on computer science research projects, scientific code development, and domain-specific research, as well as participating in professional organizations and conferences and contributing to journals and proceedings. The NERSC user community benefits from the results of these activities as they are applied to systems, software, and services at NERSC and throughout the HPC community.

Publications and presentations by NERSC staff in 2012 are listed below. (Not all co-authors are from NERSC.) Online links to many of these reports and presentations can be found at <https://www.nersc.gov/news-publications/publications-reports/nersc-staff-publications-and-presentations/>.

Dan Gunter, Shreyas Cholia, Anubhav Jain, Michael Kocher, Kristin Persson, Lavanya Ramakrishnan, Shyue Ping Ong, Gerbrand Ceder, **"Community Accessible Datastore of High-Throughput Calculations: Experiences from the Materials Project,"** 5th IEEE Workshop on Many-Task Computing on Grids and Supercomputers (MTAGS) 2012, November 12, 2012.

Alice Koniges, Katherine Yelick, Rolf Rabenseifner, Reinhold Bader, David Eder, Filip Blagojevic, Robert Preissl, Paul Hargrove, **"Introduction to PGAS (UPC and CAF) and Hybrid for Multicore Programming,"** SC12 Full Day Tutorial, November 2012.

Z. Liu, M. Veeraraghavan, Z. Yan, C. Tracy, J. Tie, I. Foster, J. Dennis, J. Hick, Y. Lik and W. Yang, **"On using virtual circuits for GridFTP transfers,"** 2012 International Conference for High Performance Computing, Networking, Storage and Analysis, November 12, 2012.

Richard A. Gerber, **"Job Analytics,"** NUG Teleconference, November 8, 2012.

R. Barrett, S. Dosanjh, et al., **"Towards Codesign in High Performance Computing Systems,"** IEEE/ACM International Conference on Computer-Aided Design (ICCAD), San Jose, CA, November 5, 2012.

Michael A. Heroux, Alice E. Koniges, David F. Richards, Richard F. Barrett, Thomas Brunner, **"Using Application Proxies for Co-design of Future HPC Computer Systems and**

Applications," SC12 Full Day Tutorial, November 2012.

Hank Childs, Eric Brugger, Brad Whitlock, Jeremy Meredith, Sean Ahern, David Pugmire, Kathleen Biagas, Mark Miller, Cyrus Harrison, Gunther H. Weber, Hari Krishnan, Thomas Fogal, Allen Sanderson, Christoph Garth, E. Wes Bethel, David Camp, Oliver Rübel, Marc Durant, Jean M. Favre, Paul Navratil, **"VisIt: An End-User Tool For Visualizing and Analyzing Very Large Data,"** in *High Performance Visualization—Enabling Extreme-Scale Scientific Insight* (CRC Press, October 31, 2012), pages 357–371.

E. Wes Bethel, David Camp, Hank Childs, Christoph Garth, Mark Howison, Kenneth I. Joy, David Pugmire, **"Hybrid Parallelism,"** in *High Performance Visualization—Enabling Extreme-Scale Scientific Insight* (CRC Press, October 31, 2012).

Wangyi Liu, John Barnard, Alice Koniges, David Eder, Nathan Masters, Aaron Fisher, Alex Friedman, **"A numerical scheme for including surface tension effects in hydrodynamic simulation: a full Korteweg type model without parasitic flows,"** APS DPP 2012.

Daya Bay Collaboration, F. P. An, Q. An, J. Z. Bai et al., **"Improved Measurement of Electron Antineutrino Disappearance at Daya Bay,"** *Chinese Phys. C* 37 011001.

Hongzhang Shan, Brian Austin, Nicholas Wright, Erich Strohmaier, John Shalf, Katherine Yelick, **"Accelerating Applications at Scale Using One-Sided Communication,"** The 6th Conference on Partitioned Global Address Programming Models, Santa Barbara, CA, October 10, 2012.

Richard A. Gerber, **"Batch Strategies of Maximizing Throughput and Allocation,"** NERSC Users Group Monthly Webinar, October 4, 2012.

Thomas Ludwig, Costas Bekas, Alice Koniges, Kengo Nakajima, **"Topic 15: High Performance and Scientific Applications,"** Euro-Par 2012 Parallel Processing, Springer, 2012, Lecture: 779–780.

David Camp, Hank Childs, Christoph Garth, David Pugmire, Kenneth I. Joy, **"Parallel Stream Surface Computation for Large Data Sets,"** Proceedings of IEEE Symposium on Large Data Analysis and Visualization (LDAV), October 1, 2012.

Alice Koniges, **"Opportunities and Challenges for Domain-Specific Languages in Fusion Applications,"** ASCR Exascale PI Meeting, October 2012.

Clayton Bagwell, **"How to Submit an ERCAP Request,"** September 6, 2012.

Allison Dzubak, Li-Chiang Lin, Jihan Kim, Joseph Swisher, Roberta Poloni, Sergei Maximoff, Berend Smit, Laura Gagliardi, **"Ab initio Carbon Capture in Open-Site Metal Organic Frameworks,"** Nature Chemistry 4, 810–816 (2012), doi:10.1038/nchem.1432.

Jihan Kim, Richard Martin, Oliver Ruebel, Maciej Haranczyk, Berend Smit, **"High-throughput Characterization of Porous Materials Using Graphics Processing Units,"** Journal of Chemical Theory and Computation, 2012, 8 (5), pp 1684–1693, DOI: 10.1021/ct200787v.

Jihan Kim, Li-Chiang Lin, Richard Martin, Joseph Swisher, Maciej Haranczyk, Berend Smit, **"Large Scale Computational Screening of Zeolites for Ethene/Ethane Separation,"** Langmuir, 2012, 28 (32), 11914.

Richard Martin, Thomas Willems, Li-Chiang Lin, Jihan Kim, Joseph Swisher, Berend Smit, Maciej Haranczyk, **"Similarity-driven Discovery of Porous Materials for Adsorption-based Separations,"** ChemPhysChem 13 (2012) 13, 3595–3597.

Li-Chiang Lin, Adam Berger, Richard Martin, Jihan Kim, Joseph Swisher, Kuldeep Jariwala, Chris Rycroft, Abhoyjit Brown, Michael Deem, Maciej Haranczyk, Berend Smit, **"In Silico Screening of Carbon Capture Materials,"** Nature Materials 11, 633–641 (2012), doi:10.1038/nmat3336.

Jihan Kim, Berend Smit, **"Efficient Monte Carlo Simulations of Gas Molecules Inside Porous Materials,"** Journal of Chemical Theory and Computation, 2012, 8 (7), pp 2336–2343, doi:10.1021/ct3003699.

Richard A. Gerber, **"Uses for High Performance Computing,"** June 12, 2012.

Richard A. Gerber, **"Introduction to High Performance Computers,"** June 12, 2012.

Richard A. Gerber, **"Challenges in HPC,"** June 12, 2012.

Damian Hazen, Jason Hick, **"MIR Performance Analysis,"** June 12, 2012, LBNL-5896E.

Alice Koniges, Mike Heroux, **"Using Application Proxies for Co-design of Future HPC Computer Systems and Applications,"** Dagstuhl Co-design Workshop, May 22, 2012.

M. Reinsch, B. Austin, J. Corlett, L. Doolittle, P. Emma, G. Penn, D. Prosnitz, J. Qiang, A. Sessler, M. Venturini, J. Wurtele, **"Machine Parameter Studies for and FEL Facility using STAFF,"** Proceedings of IPAC2012, New Orleans, Louisiana, USA, May 20, 2012, 1768.

J. Kim, A. Koniges, R.L. Martin, J. Swisher, M. Haranczyk, B. Smit, **"Computational Screening of Novel Carbon Capture Materials,"** 2012 GTC GPU Conference, 2012.

Megan Bowling, Zhengji Zhao and Jack Deslippe, **"The Effects of Compiler Optimizations on Materials Science and Chemistry Applications at NERSC,"** a paper presented in the Cray User Group meeting, April 29–May 3, 2012, Stuttgart, Germany.

Zhengji Zhao, Mike Davis, Katie Antypas, Yushu Yao, Rei Lee and Tina Butler, **"Shared Library Performance on Hopper,"** a paper presented in the Cray User Group meeting, April 29–May 3, 2012, Stuttgart, Germany.

Megan Bowling, Zhengji Zhao and Jack Deslippe, **"The Effects of Compiler Optimizations on Materials Science and Chemistry Applications at NERSC,"** a talk in the Cray User Group meeting, April 29–May 3, 2012, Stuttgart, Germany.

Zhengji Zhao, Yun (Helen) He and Katie Antypas, **"Cray Cluster Compatibility Mode on Hopper,"** a paper presented in the Cray User Group meeting, April 29–May 3, 2012, Stuttgart, Germany.

Yun (Helen) He and Katie Antypas, **"Running Large Jobs on a Cray XE6 System,"** Cray User Group Meeting, April 29–May 3, 2012, Stuttgart, Germany.

Yun (Helen) He, **"Programming Environments, Applications, and Documentation SIG,"** Cray User Group meeting, April 29–May 3, 2012, Stuttgart, Germany.

N. Balthaser, J. Hick, W. Hurlbert, **"StorageTek Tape Analytics: Pre-Release Evaluation at LBNL,"** LTUG 2012, April 25, 2012.

Larry Pezzaglia, **"CHOS in Production: Supporting Multiple Linux Environments on PDSF at NERSC,"** a talk at the HEPiX Spring 2012 Workshop, Prague, Czech Republic, April 25, 2012.

Eric Hjort, Larry Pezzaglia, Iwona Sakrejda, **"PDSF at NERSC: Site Report,"** a talk at the HEPiX Spring 2012 Workshop, Prague, Czech Republic, April 24, 2012.

E. Wes Bethel, David Camp, Hank Childs, Mark Howison, Hari Krishnan, Burlen Loring, Joerg Meyer, Prabhat, Oliver Ruebel, Daniela Ushizima, Gunther Weber, **"Towards Exascale: High Performance Visualization and Analytics — Project Status Report,"** DOE Exascale Research Conference, April 1, 2012, LBNL 5767E.

Richard A. Gerber, Harvey J. Wasserman, **"Large Scale Computing and Storage Requirements for Nuclear Physics,"** Workshop, March 26, 2012, LBNL LBNL-5355E.

F. P. An, J. Z. Bai, A. B. Balantekin, et al., **"Observation of electron-antineutrino disappearance at Daya Bay,"** Phys. Rev. Lett. 108, 171803 (2012).

A.C. Usselton, K.B. Antypas, D. Ushizima, J. Sukharev, **"File System Monitoring as a Window into User I/O Requirements,"** CUG Proceedings, Edinburgh, Scotland, March 1, 2012.

Richard Shane Canon, **"Magellan Project: Clouds for Science?"** Coalition for Academic Scientific Computation, February 29, 2012.

Richard Gerber, David Skinner, **"Debugging and Optimization Tools,"** Presented to UC Berkeley CS 267 Class, Applications of Parallel Computer, February 16, 2012.

David Skinner, **"Tools for Performance Debugging HPC Applications,"** Presented to UC Berkeley CS 267 Class, Applications of Parallel Computer, February 16, 2012.

Scott Campbell, Jason Lee, **"Prototyping a 100G Monitoring System,"** 20th Euromicro International Conference on Parallel, Distributed, and Network-Based Processing (PDP 2012), February 12, 2012.

"NERSC Exceeds Reliability Standards With Tape-Based Active Archive," Active Archive Alliance Case Study, February 10, 2012.

Kathy Yelick, **"NERSC Accomplishments and Plans,"** NERSC User Group meeting, February 3, 2012.

Richard Gerber, **"User Requirements Gathered for the NERSC-7 Procurement,"** NERSC User Group meeting, February 3, 2012.

Zhengji Zhao and Helen He, **"Using Cray Cluster Compatibility Mode on Hopper,"** NERSC User Group meeting, Feb. 2, 2012.

Yun (Helen) He and Woo-Sun Yang, **"Using Hybrid MPI/OpenMP, UPC, and CAF at NERSC,"** NERSC User Group Meeting, Oakland, CA, February 2, 2012.

K. Antypas, **"Best Practices for Reading and Writing Data on HPC Systems,"** NERSC User Group Meeting, February 1, 2012.

Joshua S. Bloom, Daniel Kasen, Ken J. Shen, Peter E. Nugent, Nathaniel R. Butler, Melissa L. Graham, D. Andrew Howell,

Ulrich Kolb, Stefan Holmes, Carole A. Haswell, Vadim Burwitz, Juan Rodriguez, and Mark Sullivan, **"SN 2010jp (PTF10aaxi): a jet in a Type II supernova,"** Monthly Notices of the Royal Astronomical Society, Volume 420, Issue 2, pages 1135–1144, February 2012.

Richard A. Gerber, Harvey J. Wasserman, **"Large Scale Computing and Storage Requirements for Advanced Computational Science Research,"** Workshop, January 2012, LBNL-5249E.

J. Hick, **"NERSC Site Update (NGF),"** SPXXL Winter 2012, January 10, 2012.

Jack Deslippe, Georgy Samsonidze, David Strubbe, Manish Jain, Marvin L. Cohen, Steven G. Louie, **"BerkeleyGW: A Massively Parallel Computer Package for the Calculation of the Quasiparticle and Optical Properties of Materials,"** Computer Physics Communications, Volume 183, Issue 6, June 2012, Pages 1269–1289.

Gunther H. Weber, Peer-Timo Bremer, Valerio Pascucci, **"Topological Cacti: Visualizing Contour-based Statistics,"** in *Topological Methods in Data Analysis and Visualization II: Theory, Algorithms, and Applications* (Springer-Verlag: January 1, 2012), pages: 63–76.

D.Y. Zubarev, B.M. Austin, W.A. Lester Jr., **"Practical Aspects of Quantum Monte Carlo for the Electronic Structure of Molecules,"** in *Practical Aspects of Computational Chemistry I: An Overview of the Last Two Decades and Current Trends*, (Springer-Verlag: January 1, 2012), p. 255.

D.Y. Zubarev, B.M. Austin, W.A. Lester Jr., **"Quantum Monte Carlo for the x-ray absorption spectrum of pyrrole at the nitrogen K-edge,"** Journal of Chemical Physics, January 1, 2012, 136:144301.

Erin LeDell, Prabhat, Dmitry Yu Zubarev, Brian Austin, Jr., William A. Lester, **"Classification of Nodal Pockets in Many-Electron Wave Functions via Machine Learning,"** Journal of Mathematical Chemistry, January 1, 2012, 50:2043.

Joshua S. Bloom, Daniel Kasen, Ken J. Shen, Peter E. Nugent, Nathaniel R. Butler, Melissa L. Graham, D. Andrew Howell, Ulrich Kolb, Stefan Holmes, Carole A. Haswell, Vadim Burwitz, Juan Rodriguez, and Mark Sullivan, **"A Compact Degenerate Primary-star Progenitor of SN 2011fe,"** Astrophysical Journal, January 2012, 744:L17.

Wangyi Liu, Andrea Bertozzi, and Theodore Kolokolnikov, **"Diffuse interface surface tension models in an expanding flow,"** Comm. Math. Sci., 2012, 10(1): 387–418.

Johannes Lischner, Jack Deslippe, Manish Jain, Steven G Louie, **"First-Principles Calculations of Quasiparticle Excitations of Open-Shell Condensed Matter Systems,"** Physical Review Letters, 2012, 109:36406.

Kaihui Liu, Jack Deslippe, Fajun Xiao, Rodrigo B Capaz, Xiaoping Hong, Shaul Aloni, Alex Zettl, Wenlong Wang, Xuedong Bai, Steven G Louie, others, **"An atlas of carbon nanotube optical transitions,"** Nature Nanotechnology, 2012, 7:325–329.

Daniela M. Ushizima, Dmitriy Morozov, Gunther H. Weber, Andrea G.C. Bianchi, James A. Sethian, E. Wes Bethel, **"Augmented Topological Descriptors of Pore Networks for Material Science,"** IEEE Transactions on Visualization and Computer Graphics (Proceedings IEEE Vis 2012), 2012, 18:2041–2050.

Hank Childs, David Pugmire, Sean Ahern, Brad Whitlock, Mark Howison, Prabhat, Gunther H. Weber, E. Wes Bethel, **"Visualization at Extreme-Scale Concurrency,"** in *High Performance Visualization: Enabling Extreme-Scale Scientific Insight*, (CRC Press: 2012) pages: 291–306.

Oliver Rübel, Soile V. E. Keränen, Mark D. Biggin, David W. Knowles, H. Weber, Hans Hagen, Bernd Hamann, E. Wes Bethel, **"Linking Advanced Visualization and MATLAB for the Analysis of 3D Gene Expression Data,"** in *Mathematical Methods for Visualization in Medicine and Life Sciences II*, (Springer-Verlag: 2012) pages: 267–286, LBNL 4891E.

Kenes Beketayev, Gunther H. Weber, Dmitriy Morozov, Aidos Abzhanov, Bernd Hamann, **"Geometry-Preserving Topological Landscapes,"** Proceedings of the Workshop at SIGGRAPH Asia 2012, New York, NY, USA, ACM, 2012, 155–160, LBNL 6082E.

Dogan Demir, Kenes Beketayev, Gunther H. Weber, Peer-Timo Bremer, Valerio Pascucci, Bernd Hamann, **"Topology Exploration with Hierarchical Landscapes,"** Proceedings of the Workshop at SIGGRAPH Asia 2012, New York, NY, USA, ACM, 2012, 147–154, LBNL 6083E.

Allen R. Sanderson, Brad Whitlock, Oliver Rübel, Hank Childs, Gunther H. Weber, Prabhat, Kesheng Wu, **"A System for Query Based Analysis and Visualization,"** Third International EuroVis Workshop on Visual Analytics EuroVA 2012, Vienna, Austria, 2012, 25–29, LBNL 5507E.

Gunther H. Weber, Peer-Timo Bremer, **"In-situ Analysis: Challenges and Opportunities,"** 2012, LBNL 5692E.

Gunther H. Weber, Kenes Beketayev, Peer-Timo Bremer, Bernd Hamann, Maciej Haranczyk, Mario Hlawitschka, Valerio Pascucci, **"Comprehensible Presentation of Topological Information,"** 2012, LBNL 5693E.

Surendra Byna, Jerry Chou, Oliver Rübel, Prabhat, Homa Karimabadi, William S. Daughton, Vadim Roytershteyn, E. Wes Bethel, Mark Howison, Ke-Jou Hsu, Kuan-Wu Lin, Arie Shoshani, Andrew Uselton, Kesheng Wu, **"Parallel I/O, analysis, and visualization of a trillion particle simulation,"** SC 12, Los Alamitos, CA, USA, IEEE Computer Society Press, 2012, 59:1–59:1.

Oliver Rübel, Cameron G.R. Geddes, Min Chen, Estelle Cormier-Michel, E. Wes Bethel, **"Feature-based Analysis of Plasma-based Particle Acceleration Data,"** IEEE Transactions on Visualization and Computer Graphics (TVCG), 2012.

Michael Wehner, Surendra Byna, Prabhat, Thomas Yopes, John Wu, **"Atmospheric Rivers in the CMIP3/5 Historical and Projection Simulations,"** World Climate Research Programme (WCRP) Workshop on CMIP5 Model Analysis, 2012.

E. Wes Bethel, Surendra Byna, Jerry Chou, Estelle Cormier-Michel, Cameron G. R. Geddes, Mark Howison, Fuyu Li, Prabhat, Ji Qiang, Oliver Rübel, Robert D. Ryne, Michael Wehner, Kesheng Wu, **"Big Data Analysis and Visualization: What Do LINACS and Tropical Storms Have In Common?"** 11th International Computational Accelerator Physics Conference, ICAP 2012, Rostock-Warnemünde, Germany, 2012.

Babak Behzad, Joseph Huchette, Huong Luu, Suren Byna, Yushu Yao, Prabhat, **"Auto-Tuning of Parallel I/O Parameters for HDF5 Applications,"** SC 2012.

Surendra Byna, Jerry Chou, Oliver Rübel, Prabhat, Homa Karimabadi, William S. Daughton, Vadim Roytershteyn, E. Wes Bethel, Mark Howison, Ke-Jou Hsu, Kuan-Wu Lin, Arie Shoshani, Andrew Uselton, Kesheng Wu, **"Parallel Data, Analysis, and Visualization of a Trillion Particles,"** XLDB 2012.

E. Wes Bethel, Rob Ross, Wei-Ken Liao, Prabhat, Karen Schuchardt, Peer-Timo Bremer, Oliver Rübel, Surendra Byna, Kesheng Wu, Fuyu Li, Michael Wehner, John Patchett, Han-Wei Shen, David Pugmire, Dean Williams, **"Recent Advances in Visual Data Exploration and Analysis of Climate Data,"** SciDAC 3 Principal Investigator Meeting, 2012.

Oliver Ruebel, Cameron Geddes, Min Chen, Estelle Cormier, Ji Qiang, Rob Ryne, Jean-Luc Vey, David Grote, Jerry Chou, Kesheng Wu, Mark Howison, Prabhat, Brian Austin, Arie Shoshani, E. Wes Bethel, **"Scalable Data Management, Analysis and Visualization of Particle Accelerator Simulation Data,"** SciDAC 3 Principal Investigator Meeting, 2012.

Prabhat, Suren Byna, Kesheng Wu, Jerry Chou, Mark Howison, Joey Huchette, Wes Bethel, Quincey Koziol, Mohammad Chaarawi, Ruth Aydt, Babak Behzad, Huong Luu, Karen Schuchardt, Bruce Palmer, **"Updates from the ExaHDF5 project: Trillion particle run, Auto-Tuning and the Virtual Object Layer,"** DOE Exascale Research Conference, 2012.

E. Wes Bethel, David Camp, Hank Childs, Mark Howison, Hari Krishnan, Burlen Loring, Jörg Meyer, Prabhat, Oliver Rübel, Daniela Ushizima, Gunther Weber, **"Towards Exascale: High Performance Visualization and Analytics — Project Status Report,"** 2012.

Oliver Rübel, Wes Bethel, Prabhat, Kesheng Wu, **"Query Driven Visualization and Analysis,"** in *High Performance Visualization: Enabling Extreme-Scale Scientific Insight*, (CRC Press: 2012).

Prabhat, **"13 TB, 80,000 cores and TECA: The search for extreme events in climate datasets,"** American Geophysical Union Meeting, 2012.

Mehmet Balman, Eric Pouyoul, Yushu Yao, Burlen Loring, E. Wes Bethel, Prabhat, John Shalf, Alex Sim, Brian L. Tierney, **"Experiences with 100G Network Applications,"** Proceedings of the 5th International Workshop on Data Intensive and Distributed Computing (DIDC 2012), Delft, Netherlands, 2012.

Prabhat, Oliver Rübel, Surendra Byna, Kesheng Wu, Fuyu Li, Michael Wehner, E. Wes Bethel, **"TECA: A Parallel Toolkit for Extreme Climate Analysis,"** Third Workshop on Data Mining in Earth System Science (DMESS 2012) at the International Conference on Computational Science (ICCS 2012), Omaha, Nebraska, 2012.

Michael F. Wehner, Prabhat, Surendra Byna, Fuyu Li, Erin LeDell, Thomas Yopes, Gunther Weber, Wes Bethel, and William D. Collins, **"TECA, 13 TB, 80,000 Processors: Or: Characterizing Extreme Weather in a Changing Climate,"** Second Workshop on Understanding Climate Change from Data, 2012.

Michael Wehner, Kevin Reed, Prabhat, Surendra Byna, William D. Collins, Fuyu Liand, Travis O'Brien, Julio Bacmeister, Andrew Gettelman, **"High Resolution CAM5.1 Simulations,"** 14th International Specialist Meeting on the Next Generation Models of Climate and Sustainability for Advanced High Performance Computing Facilities, 2012.

Kevin Reed, Michael Wehner, Christiane Jablonowski, **"Towards Projecting Changes in Tropical Cyclones with High Resolution CAM5.1,"** CESM Workshop, 2012.

Wei-Chen Chen, George Ostrouchov, David Pugmire, Prabhat, Michael Wehner, **"Model Based Clustering Analysis of Large Climate Simulation Datasets,"** Joint Statistical Meeting, 2012.

Hongzhang Shan, Nicholas J. Wright, John Shalf, Katherine Yelick, Marcus Wagner, Nathan Wichmann, **"A preliminary evaluation of the hardware acceleration of Cray Gemini interconnect for PGAS languages and comparison with MPI,"** SIGMETRICS Performance Evaluation Review, 2012, 40(2):92-98.

Lavanya Ramakrishnan, R. Shane Canon, Krishna Muriki, Iwona Sakrejda, Nicholas J. Wright, **"Evaluating Interconnect and Virtualization Performance for High Performance Computing,"** SIGMETRICS Performance Evaluation Review, 2012, 40(2):55-60.

Xuefei Yuan, Jin Chen, Xiaoye Li, Ichitaro Yamazaki, Stephen Jardin, Alice Koniges, David Keyes, **"Scalability improvement of M3D-C1 via the block Jacobi preconditioner using PDSLin,"** Bulletin of the American Physical Society, 2012.

Alice Koniges, Xuefei Yuan, Wangyi Liu, Praveen Narayanan, Robert Preissl, Stephane Ethier, Weixang Wang, Stephen Jardin, Jeff Candy, **"Enabling Fusion Codes for Upcoming Exascale Platforms,"** American Physical Society, 2012, <http://meetings.aps.org/link/BAPS.2012.DPP.BP8.173>.

Gabriele Jost, Alice Koniges, **"Hardware Trends and Implications for Programming Models,"** in *Handbook of Research on Computational Science and Engineering: Theory and Practice*, (IGI Global 2012) pages: 1–21.


Appendix A:

NERSC User Statistics


In support of the DOE Office of Science's mission, NERSC served 5,778 scientists throughout the United States in 2012. These researchers work in DOE laboratories, universities, industry, and other Federal agencies. Figure 1 shows the proportion of NERSC usage by each type of institution, while Figures 2 and 3 show laboratory, university, and other organizations that used large allocations of computer time. Computational science conducted at NERSC covers the entire range of scientific disciplines, but is focused on research that supports the DOE's mission and scientific goals, as shown in Figure 4.

On their Allocation Year 2013 proposal forms, principal investigators reported 1,925 refereed publications (published or in press) for the preceding 12 months, based, at least in part, on using NERSC resources. Lists of publications resulting from use of NERSC resources are available at <https://www.nersc.gov/news-publications/publications-reports/nersc-user-publications/>.


The MPP hours reported here are Cray XE6 equivalent hours.




// Figure 1. NERSC usage by institution type, 2011 (MPP hours in millions).



// Figure 2. NERSC usage by scientific discipline, 2012 (MPP hours in millions).



// Figure 3. Academic usage at NERSC, 2012 (MPP hours in millions).



// Figure 4. DOE, other laboratory, and industry usage at NERSC, 2011 (MPP hours in millions).

Appendix B:

NERSC Users Group Executive Committee

Office of Advanced Scientific Computing Research

Mark Adams, Columbia University

Anubhav Jain, Lawrence Berkeley National Laboratory

Vadim Roytershteyn, SciberQuest, Inc.

Office of Basic Energy Sciences

Eric Bylaska, Pacific Northwest National Laboratory

Gary Grest, Sandia National Laboratories

Paul Kent, Oak Ridge National Laboratory

Office of Biological and Environmental Research

Thomas Bettge, National Center for Atmospheric Research

Rudesh Toofanny, University of Washington

Guishan Zheng, Harvard University

Office of Fusion Energy Sciences

Stephane Ethier (Chair), Princeton Plasma Physics Laboratory

Linda Sugiyama, Massachusetts Institute of Technology

Jean-Luc Vay, Lawrence Berkeley National Laboratory

Office of High Energy Physics

Julian Borrill, Lawrence Berkeley National Laboratory

Steven Gottlieb, Indiana University

Frank Tsung (Vice Chair), University of California, Los Angeles

Office of Nuclear Physics

Daniel Kasen, University of California, Berkeley

Tomasz Plewa, Florida State University

Martin Savage, University of Washington

Members at Large

Cameron Geddes, Lawrence Berkeley National Laboratory

Konstantin Kemenov, CRAFT Tech

Gregory Newman, Lawrence Berkeley National Laboratory

Appendix C:

Office of Advanced Scientific Computing Research

The mission of the Advanced Scientific Computing Research (ASCR) program is to discover, develop, and deploy computational and networking capabilities to analyze, model, simulate, and predict complex phenomena important to the Department of Energy (DOE). A particular challenge of this program is fulfilling the science potential of emerging computing systems and other novel computing architectures, which will require numerous significant modifications to today's tools and techniques to deliver on the promise of exascale science.

To accomplish its mission and address those challenges, the ASCR program is organized into two subprograms: Mathematical, Computational, and Computer Sciences Research; and High Performance Computing and Network Facilities.

- The Mathematical, Computational, and Computer Sciences Research subprogram develops mathematical descriptions, models, methods, and algorithms to describe and understand complex systems, often involving processes that span a wide range of time and/or length scales. The subprogram also develops the software to make effective use of advanced networks and computers, many of which contain thousands of multi-core processors with complicated interconnections, and to transform enormous data sets from experiments and simulations into scientific insight.
- The High Performance Computing and Network Facilities subprogram delivers forefront computational and networking capabilities and contributes to the development of next-generation capabilities through support of prototypes and testbeds.

Berkeley Lab thanks the program managers with direct responsibility for the NERSC program and the research projects described in this report:

Barbara Helland,
Associate Director (Acting), ASCR

Robert Lindsay,
Special Projects

Julie Stambaugh, Financial
Management Specialist

Lori Jernigan, Program
Support Specialist

Melea Baker, Administrative Specialist

Research Division

William Harrod, Director

Teresa Beachley, Program Assistant

Rich Carlson, Computer Scientist

Christine Chalk, Physical Scientist

Sandy Landsberg, Mathematician

Carolyn Lauzon, AAAS S&T Fellow

Randall Laviolette, Physical Scientist

Steven Lee, Physical Scientist

Thomas Ndousse-Fetter,
Computer Scientist

Lucy Nowell, Computer Scientist

Karen Pao, Mathematician

Sonia Sachs, Computer Scientist

Ceren Susut, Physical Scientist

Angie Thevenot, Program Assistant

Facilities Division

Barbara Helland, Director

Vincent Dattoria, Computer Scientist

Dave Goodwin, Physical Scientist

Sally McPherson, Program Assistant

Betsy Riley, Computer Scientist

Appendix D:

Advanced Scientific Computing Advisory Committee

The Advanced Scientific Computing Advisory Committee (ASCAC) provides valuable, independent advice to the Department of Energy on a variety of complex scientific and technical issues related to its Advanced Scientific Computing Research program.

Roscoe C. Giles

Chair, Boston University

Marsha Berger

Courant Institute of Mathematical Sciences

Marjory S. Blumenthal

Georgetown University

Vinton G. Cerf

Google, Inc.

Vincent Chan

General Atomics

Barbara M.P. Chapman

University of Houston

Jackie Chen

Sandia National Laboratories

Jack J. Dongarra

University of Tennessee

Sharon C. Glotzer

University of Michigan

Susan L. Graham

University of California, Berkeley

Anthony Hey

Microsoft Research

Gwendolyn L. Huntoon

Pittsburgh Supercomputing Center

Juan Meza

University of California, Merced

John Negele

Massachusetts Institute of Technology

Linda R. Petzold

University of California, Santa Barbara

Vivek Sarkar

Rice University

Victoria White

Fermi National Accelerator Laboratory

Dean N. Williams

Lawrence Livermore National Laboratory

Appendix E:

Acronyms and Abbreviations

ACM - Association for Computing Machinery

ACS - American Chemical Society

ALCC - ASCR Leadership Computing Challenge

ALS - Advanced Light Source, Lawrence Berkeley National Laboratory

ALTD - Automated Library Tracking Database

ANL - Argonne National Laboratory

API - Application programming interface

APS - American Physical Society

ASCII - American Standard Code for Information Interchange

ASCR - Office of Advanced Scientific Computing Research (DOE)

BER - Office of Biological and Environmental Research (DOE)

BES - Office of Basic Energy Sciences (DOE)

BNOC - Boulder Network Operations Center (NOAA)

CAL - Computer Architecture Laboratory

CCM - Cluster Compatibility Mode

CLE - Cray Linux Environment

CRT - Computational Research and Theory Facility (under construction), Lawrence Berkeley National Laboratory

CO₂ - Carbon dioxide

CRD - Computational Research Division, Lawrence Berkeley National Laboratory

CSCS - Swiss National Supercomputing Centre

CSE - Computational Science and Engineering

CSGF - Computational Science Graduate Fellowship

CTF - Cyclops Tensor Framework

CTO - Chief Technology Officer

DARPA - Defense Advanced Research Project Agency

DFT - Density functional theory

DNA - Deoxyribonucleic acid

DOE - U.S. Department of Energy

DOI - Digital object identifier

DSL - Dynamic Shared Library

DTN - Data transfer node

DVS - Data Virtualization Service

EDL - Electrical double layer

EFRC - DOE Energy Frontier Research Center

EMSL - Environmental Molecular Science Laboratory at Pacific Northwest National Laboratory

EPSI - SciDAC Center for Edge Physics Simulation

ESnet - Energy Sciences Network (DOE)

eV - Electron volts

EXO - Enriched Xenon Observatory

FES - Office of Fusion Energy Sciences (DOE)

FIRST - Fluid Interface Reactions, Structures and Transport: a DOE Energy Frontier Research Center

FTP - File Transfer Protocol

GB - Gigabytes

GB/s - Gigabytes per second

GEFS - Global Ensemble Forecast System

GPFS - Global Parallel File System (IBM)

GUI - Graphical user interface

HECC - NASA's High-End Computing Capability Project

HEP - Office of High Energy Physics (DOE)

HMP - Human Microbiome Project

HPC - High performance computing

HPSS - High Performance Storage System

HTML - Hypertext Markup Language

HTTP - Hypertext Transfer Protocol

IMPACTS - Investigation of the Magnitudes and Probabilities of Abrupt Climate Transitions

IDS - Intrusion detection system

IEEE - Institute of Electrical and Electronics Engineers

I/O - Input/output

ITER - Latin for "the way"; an international fusion energy experiment in southern France

JGI - Joint Genome Institute (DOE)

JSON - JavaScript Object Notation

LBNL - Lawrence
Berkeley National
Laboratory

LED - Light-emitting
diode

LLNL - Lawrence
Livermore National
Laboratory

LMT - Lustre
Monitoring Tool

MIT - Massachusetts
Institute of Technology

MOF - Metal oxide
framework

MOTD - Message of
the Day

MPI - Message Passing
Interface

MPP - Massively parallel
processing

NCAR - National Center
for Atmospheric Research

NCGC - Center for
Nanoscale Control of
Geologic CO₂: a DOE
Energy Frontier Research
Center

NEMS -
Nanoelectromechanical
systems

NERSC - National Energy
Research Scientific
Computing Center

NEWT - NERSC Web
Toolkit

NGF - NERSC Global
Filesystem

NICS - National Institute
for Computational
Sciences at the University
of Tennessee/Oak Ridge
National Laboratory

NIH - National Institutes
of Health

NIM - NERSC Information
Management

NISE - NERSC Initiative
for Scientific Exploration

NOAA - National
Oceanic and Atmospheric
Administration

NP - Office of Nuclear
Physics (DOE)

NSF - National Science
Foundation

NUG - NERSC Users
Group

OLC - Onion-like carbon

OLCF - Oak Ridge
Leadership Computing
Facility

OpenMSI - Open Mass
Spectrometry Imaging

PDSF - Parallel
Distributed Systems
Facility (NERSC)

PI - Principal investigator

PNNL - Pacific Northwest
National Laboratory

PPPL - Princeton Plasma
Physics Laboratory

PTF - Palomar Transient
Factory

QCD - Quantum
chromodynamics

RNA - Ribonucleic acid

SC - DOE Office
of Science

SciDAC - Scientific
Discovery through
Advanced Computing
(DOE)

SDSS - Sloan Digital
Sky Survey

SIAM - Society for
Industrial and Applied
Mathematics

SQL - Structured Query
Language

TACC - Texas Advanced
Computing Center

TB - Terabytes

TCP/IP - Transmission
Control Protocol/Internet
Protocol

UCSD - University of
California, San Diego

URL - Universal
Resource Locator

UTC - United
Technologies Corporation

VASP - Vienna Ab initio
Simulation Package

VPC - Virtual private
cluster

Xe - Xenon

For more information about NERSC, contact:

Jon Bashor

NERSC Communications
Berkeley Lab, MS 50B2215
1 Cyclotron Road
Berkeley, CA 94720-8148

email jbashor@lbl.gov
phone (510) 486-5849
fax (510) 486-4300

NERSC's web site

<http://www.nersc.gov/>

NERSC Annual Report Editor

John Hules

Contributing writers and editors

John Hules, Jon Bashor, Linda Vu, and Margie Wylie, *Berkeley Lab Computing Sciences Communications Group*

Katie Antypas, Shane Canon, Richard Gerber, Jason Hick, Peter Nugent, David Skinner, Jay Srinivasan, Francesca Verdier, and Harvey Wasserman, *NERSC*

Dan Krotz, *Berkeley Lab Communications Department*

Kimm Fesenmaier and Marcus Woo, *California Institute of Technology*

Robert Sanders, *University of California, Berkeley*

Communications Office, *Pacific Northwest National Laboratory*

Rebecca Raber, *Haverford College*

Aaron Dubrow, *Texas Advanced Computing Center*

Communications Office, *National Center for Atmospheric Research*

David L. Chandler, *MIT News Office*

Design, layout, illustrations, photography, and printing coordination provided by the Creative Services Office of Berkeley Lab's Public Affairs Department.

CSO26585

DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, or The Regents of the University of California.

Ernest Orlando Lawrence Berkeley National Laboratory
is an equal opportunity employer.

