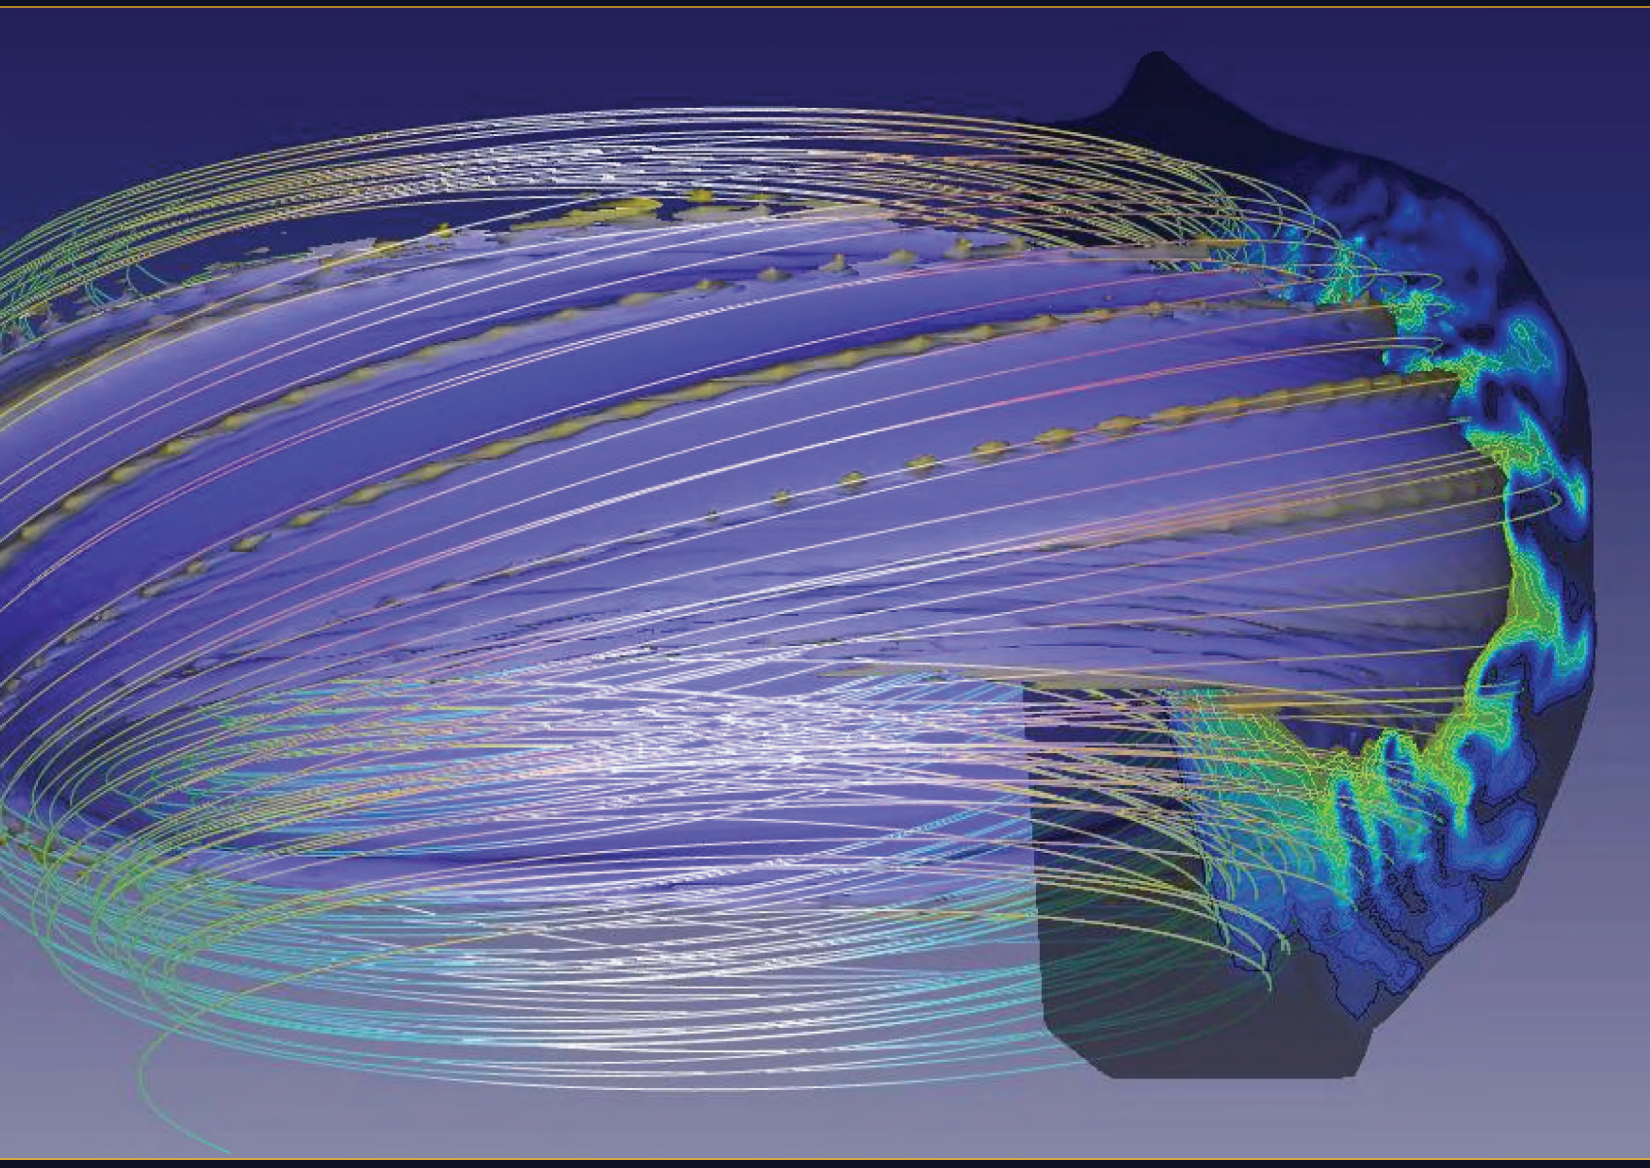

NERSC

2010 Annual Report

National Energy Research Scientific Computing Center

2010 Annual Report

Ernest Orlando Lawrence Berkeley National Laboratory
1 Cyclotron Road, Berkeley, CA 94720-8148

This work was supported by the Director, Office of Science, Office of Advanced Scientific Computing Research of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Lawrence Berkeley National Laboratory

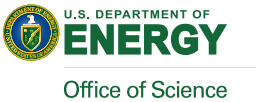

Table of Contents

The Year in Perspective

Even while NERSC is settling into the petascale era, we are preparing for the exascale era.

For both the NERSC user community and the center staff, 2010 was truly a transformational year, with an unprecedented number of new systems and upgrades, highlighted by NERSC's entry into petascale computing. As the primary computing center for the Department of Energy's Office of Science (DOE SC), our mission is to anticipate and support the computational science needs—in both systems and services—of DOE SC's six constituent offices. The science we support ranges from basic to mission-focused, from fundamental breakthroughs to practical applications. DOE's science mission currently includes a renewed emphasis on new energy sources, energy production and storage, and understanding the impacts of energy use on both regional and global climate systems, along with mitigation of those impacts through techniques such as carbon sequestration. A number of major improvements during the past year helped NERSC provide the critical support to meet these research needs.

The most notable improvement was the two-phase installation of Hopper, our newest supercomputer. Hopper Phase 1, a Cray XT5, was installed in October 2009 and entered production use in March 2010. Even with dedicated testing and maintenance times, utilization of Hopper-1 from December 15 to March 1 reached 90%, and significant scientific results are produced even during the testing period.

Hopper Phase 2, a petascale Cray XE6, was fully installed in September 2010 and opened for early use and testing late in the year. NERSC users responded in their usual fashion to this opportunity to run their codes at higher scale while at the same time helping us ensure the machine could meet the growing demands of our user community. Hopper's potential was evidenced by the 1.05 petaflop/s performance the system achieved running Linpack, making Hopper No. 5 on the November 2010 edition of the TOP500 list of the world's fastest computers. More importantly, the machine performs representative DOE applications even faster than anticipated due to the low-overhead Cray Gemini network and the fast shared memory communication within the AMD compute nodes.

Hopper represents an important step in a long-term trend of massive growth in on-chip parallelism and a shift away from a pure message passing (MPI) programming model. Hopper has 24-core nodes, in contrast to the dual-core nodes in the original Franklin system, a Cray XT4 system installed only three years earlier.

While the number of cores per node grew by an order of magnitude, the memory per node did not. As a result, it is important to share data structures as much a possible within nodes using some kind of shared memory programming. NERSC is taking a proactive role in evaluating new models and guiding the transition of the user community. In collaboration with Cray, Berkeley Lab's Computational Research Division, and representative users, NERSC has established a Center of Excellence for Multicore Programming to research programming models for Hopper and beyond. MPI alone can exhaust the available memory on Hopper, but initial research shows that a hybrid MPI + OpenMP model saves significant memory and achieves faster runtime for some applications.

Other new systems and upgrades in 2010 included installation of the Carver and Magellan clusters, the Euclid analytics server, and the Dirac GPU testbed; and upgrades of the NERSC Global Filesystem, the tape archive, and the power supply to the Oakland Scientific Facility. All of these systems and upgrades are described in the NERSC Center section of this report. But it's worth noting here that these installations required major efforts not only from the Systems Department staff, who planned and coordinated them, but also from the Services Department staff, who are responsible for software installation and user training. The challenge in deploying innovative, maximum-scale systems is always in making them available, usable, and failure tolerant. With their cross-group teamwork, resourcefulness, and commitment, NERSC's staff excels at meeting those challenges.

In the broader computing community, interest in and use of cloud computing grew significantly. NERSC used the Magellan system as a testbed to evaluate the viability of cloud computing to meet a portion of DOE's scientific computing demands for small-scale parallel jobs. The project generated a great deal of interest in the community, judging by the number of media inquiries, speaking invitations, and industry attention Magellan has generated. As part of a joint effort with Argonne National Laboratory and the Computational Research Division, NERSC is using Magellan to evaluate the costs, benefits, capabilities, and downsides of cloud computing. All of the cost advantages of centralization exist for scientific users as well as businesses using commercial clouds, but science applications require different types of systems and scheduling. The LBNL research showed that even modest scale scientific workloads slow down significantly—as much as a factor of 50—on standard cloud configurations, and the slowdowns are worse with more communication-intensive problems and as the parallelism within a job scales up. Even ignoring these performance concerns, the cost of commercial clouds is prohibitive—up to 20 cents per core hour for clouds configured with high speed networking and parallel job schedulers—compared to a total NERSC cost of around 2 cents per core hour, including the cost of NERSC's scientific consulting services. Still, the cloud computing analysis also revealed some important advantages of the commercial setting in the ability to run tailored software stacks and to obtain computing on demand with very low wait times.

Magellan also captured the attention of DOE's Joint Genome Institute (JGI), located 20 miles away in Walnut Creek, California. In March 2010, JGI became an early user of Magellan when the genomics center had a sudden need for increased computing resources. In less than three days, NERSC and JGI staff provisioned and configured a 120-node cluster in Magellan to match the computing environment available on JGI's local compute clusters, where the data resides. At the same time, staff at both centers collaborated with ESnet network engineers to deploy—in just 24 hours—a dedicated 9 Gbps virtual circuit between Magellan and JGI over ESnet's Science Data Network (SDN). In April, NERSC and JGI formalized their partnership, with NERSC managing JGI's six-person systems staff to integrate, operate, and

support systems at both NERSC and JGI. This strategy gives JGI researchers around the world increased computational capacity without any change to their software or workflow.

NERSC's focus on scientific productivity led us to expand the number of Science Gateways, which make massive scientific data sets and high performance computing resources available through the web. The NERSC Web Toolkit (NEWT) enables anyone familiar with HTML and Javascript to develop a Science Gateway. The current roster of active gateways includes the DeepSky astronomical image database, the Gauge Connection for lattice QCD data, the Coherent X-Ray Imaging Data Bank, the 20th Century Reanalysis Project climate database, the Daya Bay Neutrino Experiment, and the Earth System Grid climate data gateway. Science Gateways are extending the impact of computational science to new scientific domains, such as experimental communities who did not previously have easy access to HPC resources and data. This kind of access also helps make scientific conclusions more robust by making it easier to replicate analyses and reproduce results.

Even while NERSC is settling into the petascale era, we are preparing for the exascale era, with the understanding that exascale is all about energy-efficient computing. An exaflop computer built with today's standard technology would require 3 gigawatts of electricity; and even accounting for Moore's Law increases in chip density, an exaflop computer would still need over 100 megawatts of power, which is simply not affordable. Twenty megawatts is a more affordable target, but achieving that goal will require major innovations in computer architecture to design for low power, which will change the algorithms and programming techniques needed for exascale systems.

NERSC's research into improving performance on advanced computing architectures includes the Dirac GPU testbed—another collaboration with Berkeley Lab's Computational Research Division. Each of Dirac's 48 nodes is composed of two CPUs and one GPU. GPUs can offer energy-efficient and cost-effective performance boosts to traditional processors. The question that Dirac will help to answer is whether GPUs can do this for a broad scientific workload or for a more limited class of computations. About 100 users are running codes on Dirac, and it seems to be particularly popular with postdocs in the SciDAC-e program, which is improving and developing new scientific algorithms and software for energy-related science.

Several innovations to increase NERSC's energy efficiency were implemented in 2010. NERSC staff configured the installation of the IBM iDataplex system Carver (of which Dirac is a sub-cluster) in a manner so efficient that it reduces cooling costs by 50 percent, and in some places the cluster actually cools the air around it. NERSC has also instrumented its machine room with state-of-the-art wireless monitoring technology that gathers information on variables important to machine room operation, including air temperature, pressure, and humidity. Responding quickly to accurate machine room data not only enhances cooling efficiency, it helps assure the reliability of center systems.

All of these activities enhanced NERSC's focus of advancing scientific discovery. Over the course of the past year, our user community continued to grow, and NERSC supported over 3,200 users at universities, national labs, and other research organizations around the nation. With all the advanced technologies we deploy, it is still people who make the difference. I am grateful to our DOE SC sponsors for their continued endorsements, to our users for the science they produce using NERSC resources, and to the NERSC staff who make NERSC an excellent high performance computing center.

Katherine Yelick
NERSC Division Director

Research News

With 4,294 users in 2010 from universities, national laboratories, and industry, NERSC supports the largest and most diverse research community of any computing facility within the DOE complex. In their annual allocation renewals, users reported 1,864 refereed publications (accepted or submitted) enabled at least in part by NERSC resources in 2010. These scientific accomplishments are too numerous to cover in this report, but several representative examples are discussed in this section.

Research related to energy and climate included developing a rechargeable heat battery, turning waste heat into electricity, converting carbon dioxide into fuel, resolving uncertainties about carbon sequestration, explaining plasma instability in fusion experiments, and looking to earth's distant past to predict our future climate.

Other research accomplishments included new insight into the ignition of core-collapse supernovae, a surprising discovery about the composition of Earth's lower mantle, controlling the shape and conductivity of graphene nanostructures, exploring a new region of strange antimatter, and developing an important community resource for biological research—a database of protein folding pathways.

Early results from a few NERSC Initiative for Scientific Exploration (NISE) projects are also included below, along with a list of awards earned by NERSC users.

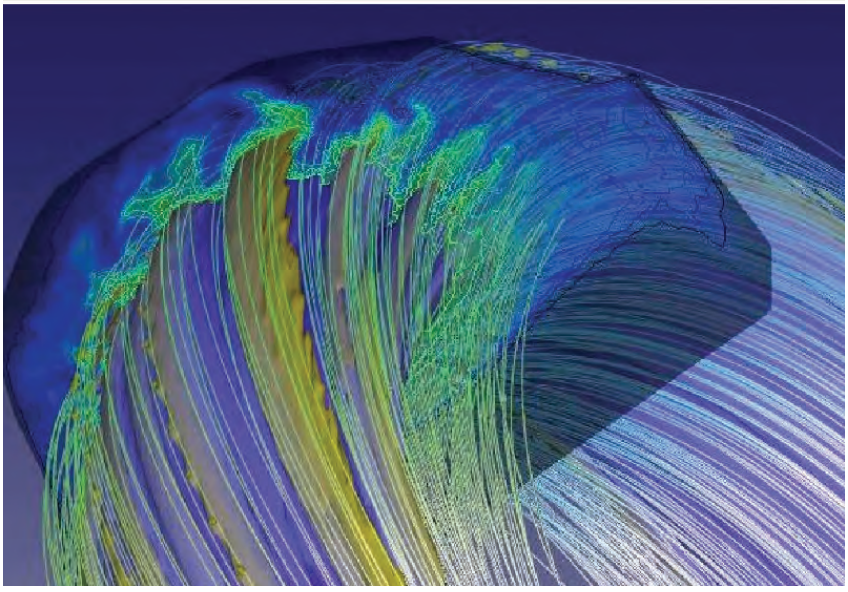

Overlooked Alternatives

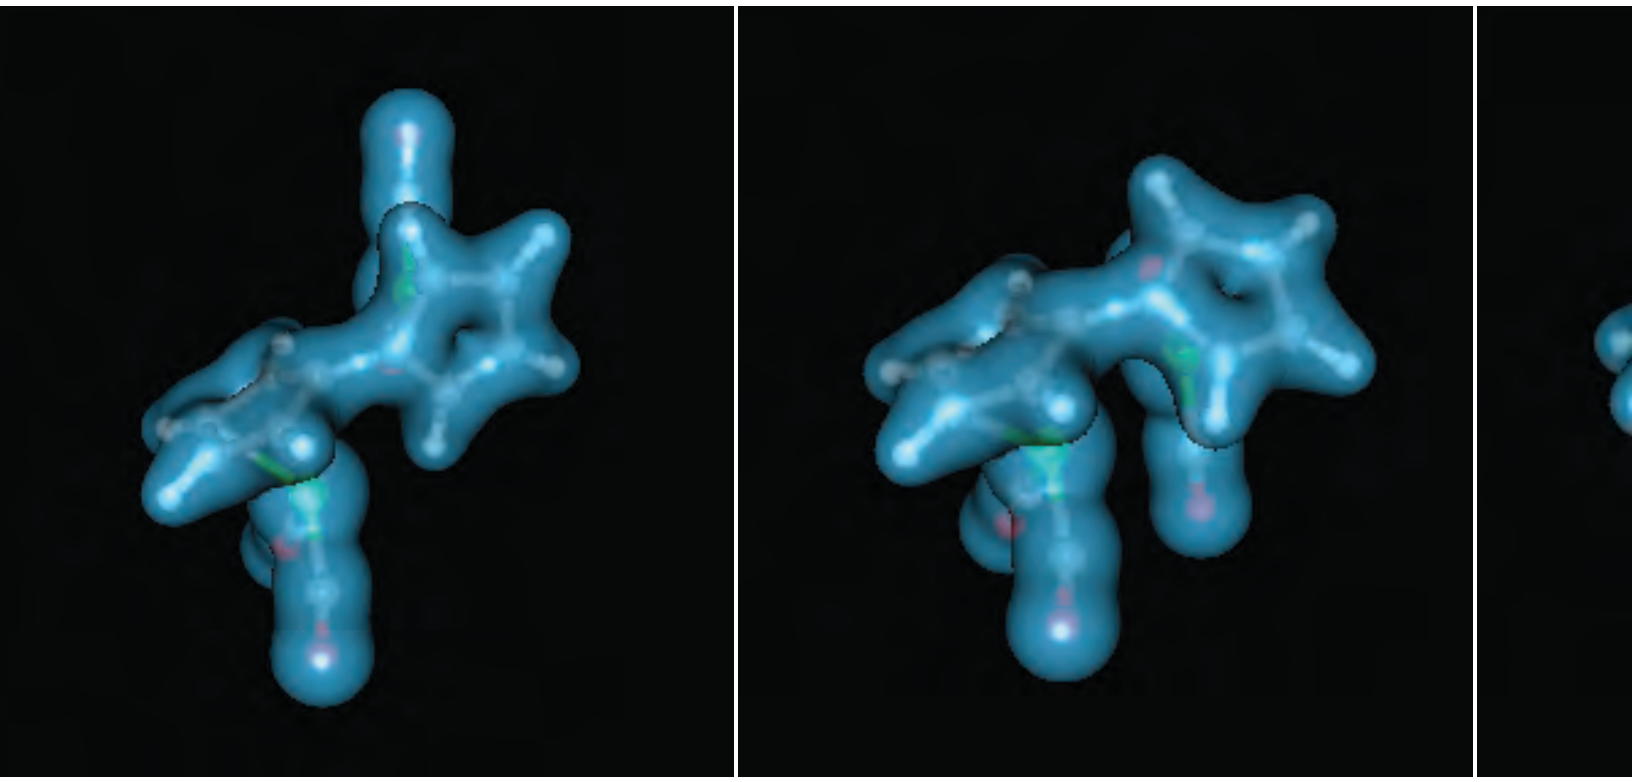

New research restores photoisomerization and thermoelectrics to the roster of promising alternative energy sources

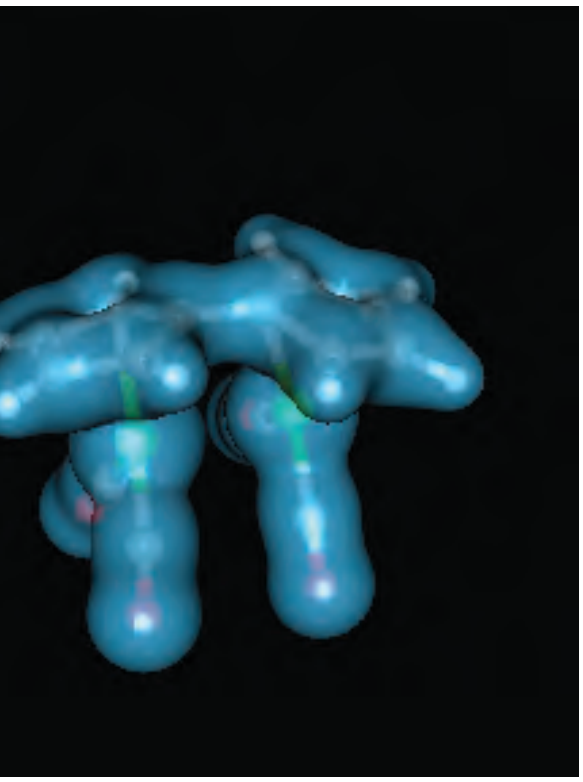

Project: Quantum Simulations of Nano-scale Solar Cells

PI: Jeffrey Grossman, Massachusetts Institute of Technology

Senior Investigators: Joo-Hyoung Lee and Vardha Srinivasan, MIT, and Yosuke Kanai, LLNL

Funding: BES, NSF, LBNL, MIT, UCB

Computing Resources: NERSC, Teragrid, LLNL

Alternative energy research might sound like a search for new ideas, but sometimes it involves putting a new spin on old ideas that have been abandoned or neglected. That's the case with two recent research projects of Jeffrey Grossman, an associate professor in the Department of Materials Science and Engineering at the Massachusetts Institute of Technology (MIT). Both projects involved collaborations with experimental groups at the University of California at Berkeley, where Grossman led a computational nanoscience research group before joining MIT in 2009, and Lawrence Berkeley National Laboratory.

"The focus of my research group at MIT is on the application and development of cutting-edge simulation tools to understand, predict, and design novel materials with applications in energy conversion, energy storage, thermal transport, surface phenomena, and synthesis," Grossman says.

His recent collaborations involved two methods of energy conversion and storage that most people have never heard of:

1. *Photoisomerization*, a thermochemical approach to capturing and storing the sun's energy which could be used to create a rechargeable heat battery.

2. *Thermoelectrics*, the conversion of temperature differences into an electric current, which could turn vast quantities of waste heat into usable energy.

A Rechargeable Heat Battery

Broadly speaking, there have been two approaches to capturing the sun's energy: photovoltaics, which turn the sunlight into electricity, or solar-thermal systems, which concentrate the sun's heat and use it to boil water to turn a turbine, or use the heat directly for hot water or home heating. But there is another approach whose potential was seen decades ago, but which was sidelined because nobody found a way to harness it in a practical and economical way.

This is the thermochemical approach, in which solar energy is captured in the configuration of certain molecules which can then release the energy on demand to produce usable heat. And unlike conventional solar-thermal systems, which require very effective insulation and even then gradually let the heat leak away,

the heat-storing chemicals could remain stable for years.

Researchers explored this type of solar thermal fuel in the 1970s, but there were big challenges: Nobody could find a chemical that could reliably and reversibly switch between two states, absorbing sunlight to go into one state and then releasing heat when it reverted to the first state. Such a compound was finally discovered in 1996 by researchers from the University of California at Berkeley and Lawrence Berkeley National Laboratory; but the compound, called fulvalene diruthenium (tetracarbonyl), included ruthenium, a rare and expensive element, so it was impractical for widespread energy storage. Moreover, no one understood how fulvalene diruthenium worked, which hindered efforts to find a cheaper variant.

Recently Peter Vollhardt of UC Berkeley and Berkeley Lab's Chemical Sciences division—the discoverer of the molecule—has begun construction of a working device as proof of principle. However, a basic understanding of what makes this particular molecule special was still missing. In order to tackle this question, Vollhardt teamed up with Grossman, who did quantum mechanical calculations at NERSC to find out exactly how the molecule accomplishes its energy storage and release.

Specifically, the theorists used the Quantum ESPRESSO code to perform density functional theory (DFT) calculations on 256 processors of the Franklin Cray XT4 supercomputer at NERSC, with additional computation at Lawrence Livermore National Laboratory.

The calculations at NERSC showed that the fulvalene diruthenium molecule, when it absorbs sunlight, undergoes a structural transformation called *photoisomerization*, putting the molecule into a higher-energy or charged state where it can remain stable indefinitely. Then, triggered by a small addition of heat or a catalyst, it snaps back to its original shape, releasing heat in the process (Figure 1). But the team found that the process is a bit more complicated than that.

"It turns out there's an intermediate step that plays a major role," said Grossman. In this intermediate step, the molecule forms a semistable configuration partway between the two previously known states (Figure 2). "That was unexpected," he said. The two-step process helps explain why the molecule is so stable, why the process is easily reversible, and also why substituting other elements for ruthenium has not worked so far.

In effect, explained Grossman, this makes it possible to produce a "rechargeable heat battery" that can repeatedly store and release heat gathered from sunlight or other sources. In principle, Grossman said, a fuel made from fulvalene diruthenium, when its stored heat is released, "can get as hot as 200

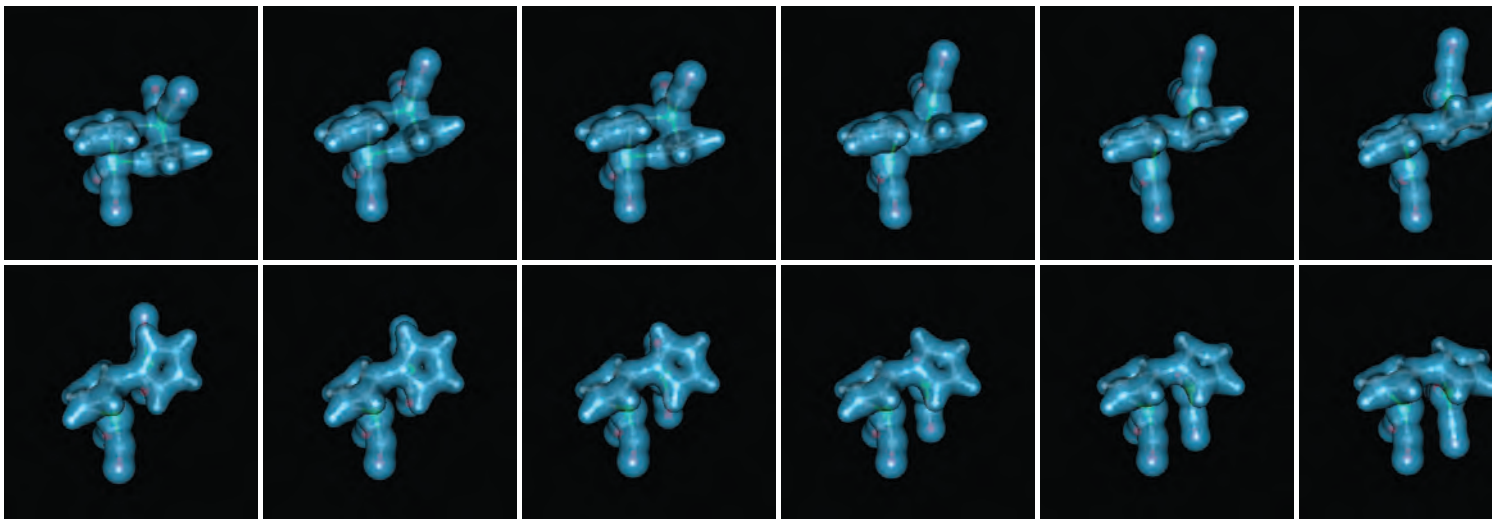

Figure 1. A molecule of fulvalene diruthenium, seen in diagram, changes its configuration when it absorbs heat, and later releases heat when it snaps back to its original shape. This series shows the reaction pathway from the high-energy state back to the original state. *Image: J. Grossman.*

degrees C, plenty hot enough to heat your home, or even to run an engine to produce electricity."

Compared to other approaches to solar energy, he said, "it takes many of the advantages of solar-thermal energy, but stores the heat in the form of a fuel. It's reversible, and it's stable over a long term. You can use it where you want, on demand. You could put the fuel in the sun, charge it up, then use the heat, and place the same fuel back in the sun to recharge."

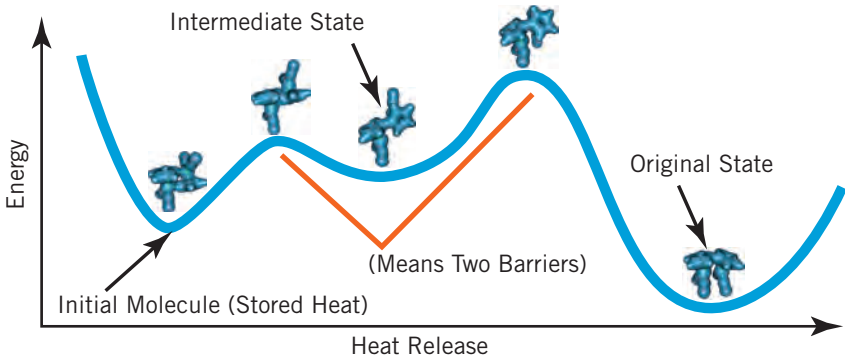

Figure 2. DFT calculations revealed an unexpected second barrier along the reaction pathway from the charged state to the original state. The two-step process helps explain the molecule's stability, and the relative barrier heights play a crucial role in its functionality.

In addition to Grossman, the theoretical work was carried out by Yosuke Kanai of Lawrence Livermore National Laboratory and Varadharajan Srinivasan of MIT, and supported by complementary experiments by Steven Meier and Vollhardt of UC Berkeley. Their report on the work, which was funded in part by the National Science Foundation, by the Sustainable Products and Solutions Program at UC Berkeley, and by an MIT Energy Initiative seed grant, was published on October 20, 2010, in the journal Angewandte Chemie International Edition.[1]

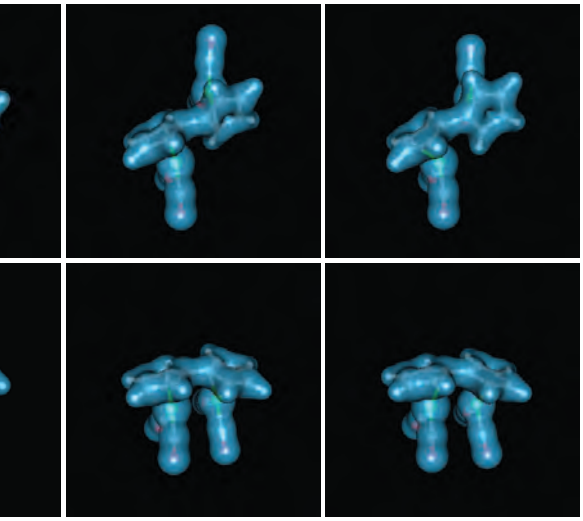

The problem of ruthenium's rarity and cost still remains as "a deal-breaker," Grossman said, but now that the fundamental mechanism of how the molecule works is understood, it should be easier to find other materials that exhibit the same behavior. This molecule "is the wrong material, but it shows it can be done," he said.

The next step, he said, is to use a combination of simulation, chemical intuition, and databases of tens of millions of known molecules to look for other candidates that have structural similarities and might exhibit the same behavior. "It's my firm belief that as we understand what makes this material tick, we'll find that there will be other materials" that will work the same way, Grossman said.

Grossman plans to collaborate with Daniel Nocera, a professor of energy and chemistry at MIT, to tackle such questions, applying the principles learned from this analysis in order to design new, inexpensive materials that exhibit this same reversible process. The tight coupling between computational materials design and experimental synthesis and validation, he said, should further accelerate the discovery of promising new candidate solar thermal fuels.

Soon after the Angewandte Chemie paper was published, Roman Boulatov, assistant professor of chemistry at the University of Illinois at Urbana-Champaign, said of this research that "its greatest accomplishment is to overcome significant challenges in quantum-chemical modeling of the reaction," thus enabling the design of new types of molecules that could be used for energy storage. But he added that other challenges remain: "Two other critical questions would have to be solved by other means, however. One, how easy is it to synthesize the best candidates? Second, what is a possible catalyst to trigger the release of the stored energy?"

[1] Y. Kanai, V. Srinivasan, S. K. Meier, K. P. C. Vollhardt, and J. C. Grossman, "Mechanism of thermal reversal of the (fulvalene)tetracarbonyldiruthenium photoisomerization: Toward molecular solar–thermal energy storage," Angewandte Chemie Int. Ed. **49**, 8926 (2010).

Recent work in Vollhardt's laboratory has provided answers to both questions. Vollhardt and other Berkeley Lab experimentalists, including Rachel Segalman and Arun Majumdar (now director of the U.S. Department of Energy's Advanced Research Projects Agency-Energy, or ARPA-E), have successfully constructed a microscale working device of fulvalene diruthenium and have applied for a U.S. patent. Using both theoretical and experimental results, they are working to maximize the device's efficiency and demonstrate that thermochemical energy storage is workable in a real system.

Turning Waste Heat to Electricity

Automobiles, industrial facilities, and power plants all produce waste heat, and a lot of it. In a coal-fired power plant, for example, as much as 60 percent of the coal's energy disappears into thin air. One estimate places the worldwide waste heat recovery market at one trillion dollars, with the potential to offset as much as 500 million metric tons of carbon per year.

One solution to waste heat has been an expensive boiler-turbine system; but for many companies and utilities, it does not recover enough waste heat to make economic sense. Bismuth telluride, a semiconductor material, has shown promise by employing a thermoelectric principle called the Seebeck effect to convert heat into an electric current, but it's also problematic: toxic, scarce, and expensive, with limited efficiency.

What's the solution? Junqiao Wu, a materials scientist at UC Berkeley and Berkeley Lab, believes the answer lies in finding a new material with spectacular thermoelectric properties that can efficiently and economically convert heat into electricity. (*Thermoelectrics* is the conversion of temperature differences in a solid—that is, differences in the amplitude of atomic vibration—into an electric current.)

Computations performed on NERSC's Cray XT4, Franklin, by Wu's collaborators Joo-Hyoung Lee and Grossman of MIT, showed that the introduction of oxygen impurities into a unique class of semiconductors known as highly mismatched alloys (HMAs) can substantially enhance the thermoelectric performance of these materials without the usual degradation in electric conductivity. The team published a paper on these results in the journal Physical Review Letters.[2]

"We are predicting a range of inexpensive, abundant, non-toxic materials in which the band structure can be widely tuned for maximal thermoelectric efficiency," says Wu. "Specifically, we've shown that the hybridization of electronic wave functions of alloy constituents in HMAs makes it possible to enhance thermopower without much reduction of electric conductivity, which is not the case for conventional thermoelectric materials."

In 1821, the physicist Thomas Johann Seebeck observed that a temperature difference between two ends of a metal bar created an electrical current in between, with the voltage being directly proportional to the temperature difference. This phenomenon became known as the Seebeck thermoelectric effect, and it holds great promise for capturing and converting into electricity some of the vast amounts of heat now being lost in the turbine-driven production

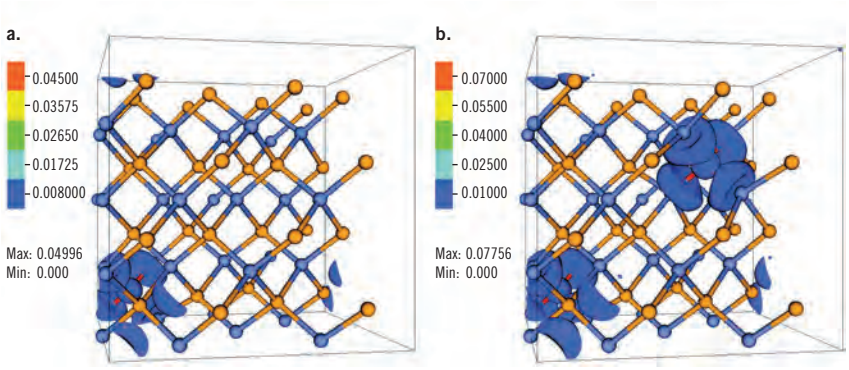

Figure 3. Contour plots showing electronic density of states in HMAs created from zinc selenide by the addition of (a) 3.125-percent oxygen atoms, and (b) 6.25 percent oxygen. The zinc and selenium atoms are shown in light blue and orange. Oxygen atoms (dark blue) are surrounded by high electronic density regions. *Image: J.-H. Lee, J. Wu, and J. Grossman.*

[2] Joo-Hyoung Lee, Junqiao Wu, and Jeffrey C. Grossman, "Enhancing the thermoelectric power factor with highly mismatched isoelectronic doping," Physical Review Letters **104**, 016602 (2010).

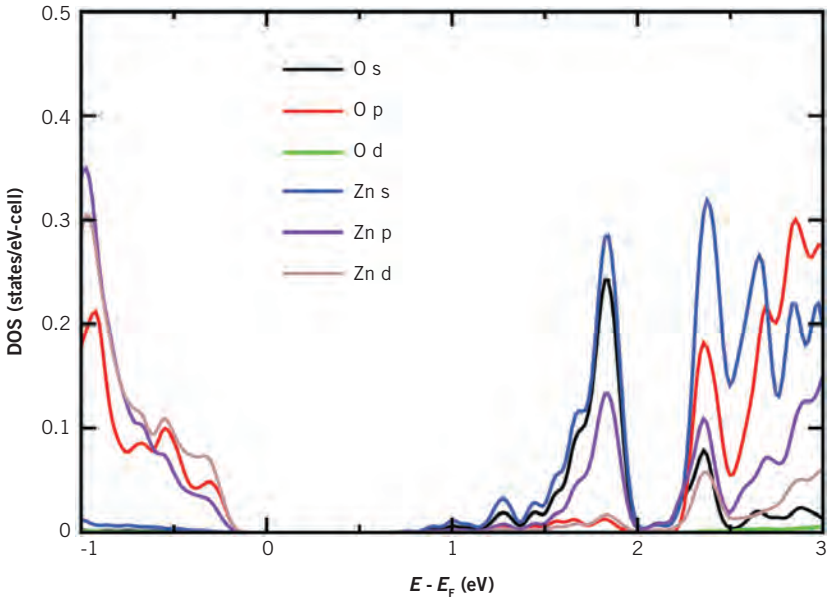

Figure 4. The projected density of states (DOS) of $ZnSe_{1-x}O_x$ with $x = 3{:}125\%$. *Image: J.-H. Lee, J. Wu, and J. Grossman.*

of electrical power. For this lost heat to be reclaimed, however, thermoelectric efficiency must be significantly improved.

"Good thermoelectric materials should have high thermopower, high electric conductivity, and low thermal conductivity," says Wu. "Enhancement in thermoelectric performance can be achieved by reducing thermal conductivity through nanostructuring. However, increasing performance by increasing thermopower has proven difficult because an increase in thermopower has typically come at the cost of a decrease in electric conductivity."

Wu knew from earlier research that a good thermoelectric material needed to have a certain type of density of states, which is a mathematical description of distribution of energy levels within a semiconductor. "If the density of states is flat or gradual, the material won't have very good thermoelectric properties," Wu says. "You need it to be spiky or peaky."

Wu also knew that a specialized type of semiconductor called highly mismatched alloys (HMAs) could be very peaky because of their hybridization, the result of forcing together two materials that don't want to mix atomically, akin to mixing water and oil. HMAs are formed from alloys that are highly mismatched in terms of electronegativity, which is a measurement of their ability to attract electrons. Their properties can be dramatically altered with only a small amount of doping. Wu hypothesized that mixing two semiconductors, zinc selenide and zinc oxide, into an HMA, would produce a peaky density of states.

In their theoretical work, Lee, Wu, and Grossman discovered that this type of electronic structure engineering can be greatly beneficial for thermoelectricity. Working with the semiconductor zinc selenide, they simulated the introduction of two dilute concentrations of oxygen atoms (3.125 and 6.25 percent respectively) to create model HMAs (Figure 3). In both cases, the oxygen impurities were shown to induce peaks in the electronic density of states above the conduction band minimum. It was also shown that charge densities near the density of state peaks were substantially attracted toward the highly electronegative oxygen atoms (Figure 4).

The researchers found that for each of the simulation scenarios, the impurity-induced peaks in the electronic density of states resulted in a sharp increase of both thermopower and electric conductivity compared to oxygen-free zinc selenide. The increases were by factors of 30 and 180 respectively.

"Furthermore, this effect is found to be absent when the impurity electronegativity matches the host that it substitutes," Wu says. "These results suggest that highly electronegativity-mismatched alloys can be designed for high performance thermoelectric applications." Wu and his research group are now working to actually synthesize HMAs for physical testing in the laboratory.

In addition to capturing energy that is now being wasted, Wu believes that HMA-based thermoelectrics can also be used for solid state cooling, in which a thermoelectric device is used to cool other devices or materials. "Thermoelectric coolers have advantages over conventional refrigeration technology in that they have no moving parts, need little maintenance, and work at a much smaller spatial scale," Wu says.

A Goldilocks Catalyst

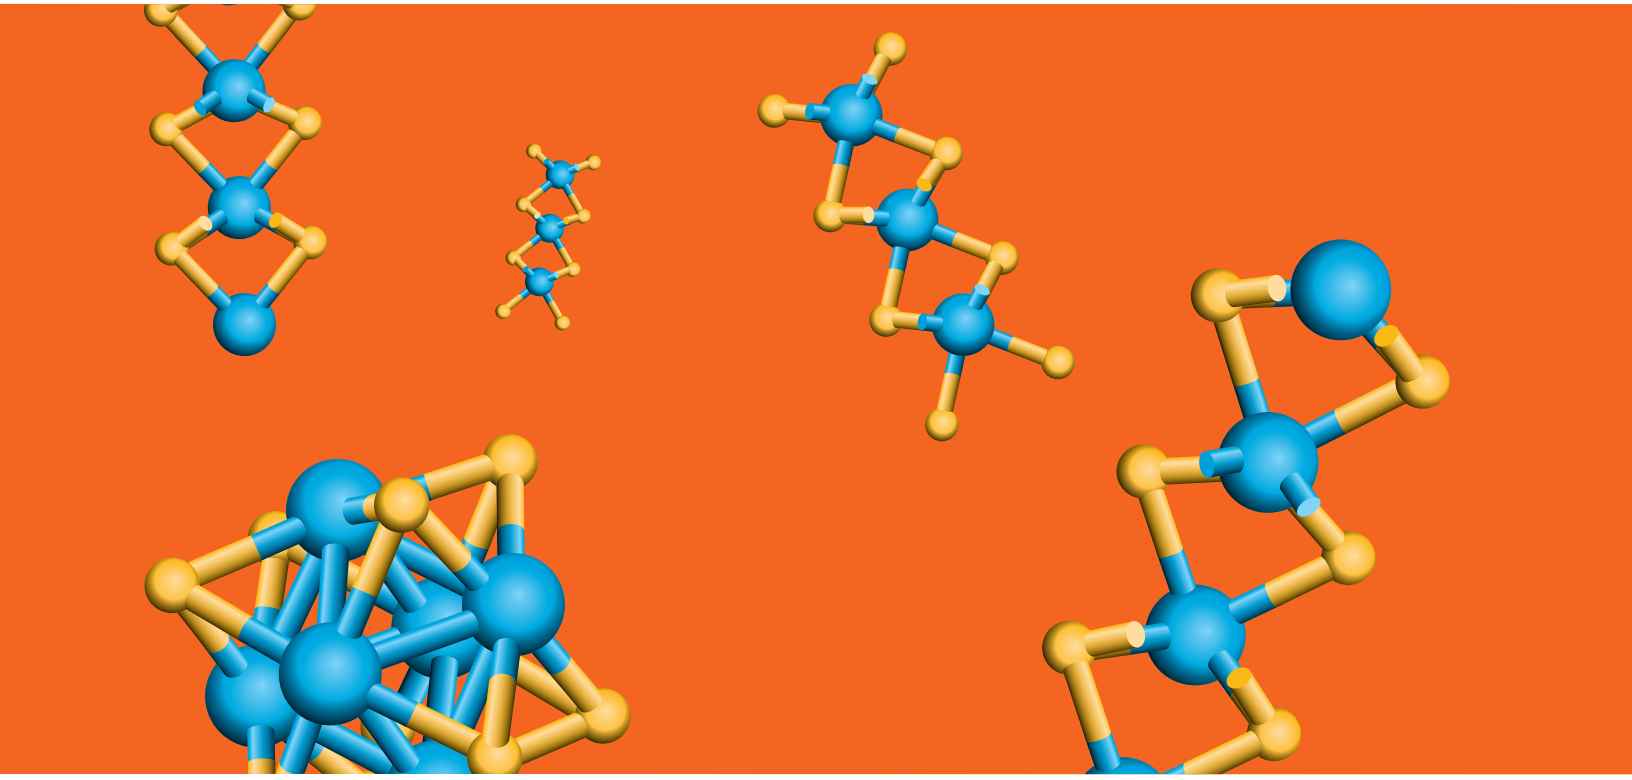

Calculations show why a nanocluster may be "just right" for recycling carbon dioxide by converting it to methanol

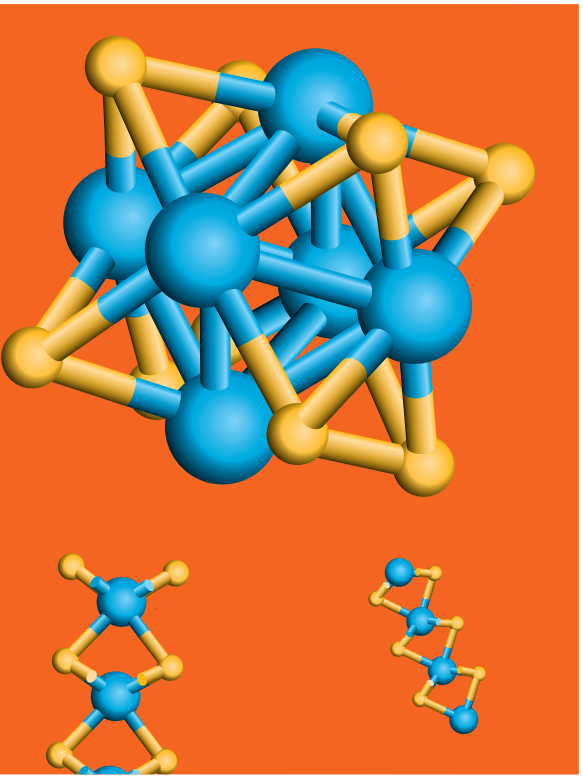

Project: Computational Studies of Supported Nanoparticles in Heterogeneous Catalysis

PI: Ping Liu, Brookhaven National Laboratory

Senior Investigators: YongMan Choi, BNL

Funding: BES, BNL

Computing Resources: NERSC, BNL

Carbon dioxide (CO_2) emissions from fossil fuel combustion are major contributors to global warming. Since CO_2 comes from fuel, why can't we recycle it back into fuel rather than releasing it into the atmosphere? An economical way to synthesize liquid fuel from CO_2 would help mitigate global warming while reducing our need to burn fossil fuels. Catalysts that induce chemical reactions could be the key to recycling carbon dioxide, but they need to have just the right activity level—too much, and the wrong chemicals are produced; too little, and they don't do their job at all.

Recycling CO_2 into alcohols such as methanol (CH_3OH) or ethanol (C_2H_5OH) by adding hydrogen would be sustainable if the hydrogen were produced by a renewable technology such as solar-thermal water splitting, which separates the hydrogen and oxygen in water. But synthesizing fuel by CO_2 hydrogenation requires a catalyst, and the most successful catalysts so far have been rhodium compounds. With a spot price averaging more than $4000 an ounce over the last five years, rhodium is too expensive to be practical on a large scale.

Molybdenum ("moly") sulfide is a more affordable catalyst, and experiments have identified it as a possible candidate for CO_2 hydrogenation; but bulk moly disulfide (MoS_2) only produces methane. Other metals can be added to the catalyst to produce alcohols, but even then, the output is too low to make the process commercially viable.

Nanoparticles, however, often have higher catalytic activity than bulk chemicals. Designing moly sulfide nanoparticles optimized for alcohol synthesis would depend on finding answers to basic questions about the chemical reaction mechanisms.

Many of those questions were answered in a paper featured on the cover of the March 25, 2010 special issue of the Journal of Physical Chemistry A devoted to "Green Chemistry in Energy Production" (Figure 5). In this paper, Brookhaven National Laboratory (BNL) chemist Ping Liu, with collaborators YongMan Choi, Yixiong Yang, and Michael G. White from BNL and the State University of New York (SUNY) Stony Brook, used density functional theory (DFT) calculations conducted at BNL and NERSC to investigate methanol synthesis using a moly sulfide nanocluster (Mo_6S_8) as the catalyst instead of the bulk moly sulfide powder used in previous experiments.[3]

"This study might seem like pure research, but it has significant implications for real catalysis," Liu says. "If you can recycle CO_2 and produce some useful fuels, that would be very significant. And that's what DOE has always asked people to do, to find some new alternatives for future fuel that will eliminate the environmental footprint."

The reason bulk moly disulfide only produces methane (CH_4) is, ironically, that it is too active a catalyst. "Catalysts are very tricky materials," Liu says. "You don't want it too active, which will break the CO bond and only form methane, an undesired product for liquid fuel production. If the catalyst is too inert, it can't dissociate CO_2 or H_2, and do the trick. So a good catalyst should have the right activity to do this catalysis."

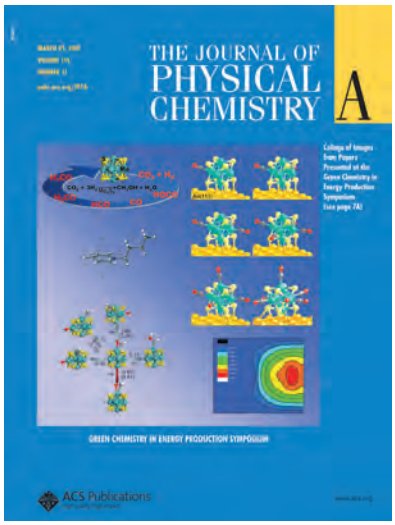

Figure 5. The top left, bottom left, and top right images illustrate aspects of methanol synthesis on a moly sulfide nanocluster. (The other two images are from different studies.) *Image: © 2010 American Chemical Society.*

Designing moly sulfide nanoparticles optimized for alcohol synthesis depends on finding answers to basic questions about the chemical reaction mechanisms.

From DFT calculations run using the VASP code on 16 to 32 cores of NERSC's Franklin system, Liu and colleagues discovered that the Mo_6S_8 nanocluster has catalytic properties very different from bulk MoS_2. This is not surprising, given the different geometries of the molecules: the cluster has a more complex, cagelike structure (Figure 6). The cluster is not as reactive as the bulk—it cannot break the carbon–oxygen bond in CO_2—and therefore it cannot produce methane. But, using a completely different reaction pathway (Figure 7), the cluster does have the ability to produce methanol.

"In this case, the nanocluster is not as active as the bulk in terms of bonding; but it has higher activity, or you could say higher selectivity, to synthesize methanol from CO_2 and H_2," Liu explains.

Figure 7 shows the details of the reaction pathway from carbon dioxide and hydrogen to methanol and water, as catalyzed by the Mo_6S_8 cluster, including the intermediate states, energetics, and geometries of the transition states—all of which were revealed for the first time by this research.

This detailed understanding is important for optimizing the catalyst's efficiency. For example, lowering the energy barrier in the CO to HCO reaction would allow increased production of methanol. Adding a C–C bond could result in synthesis of ethanol, which is usable in many existing engines. Both of these improvements might be achieved by adding another metal to the cluster, a possibility that Liu and her theorist colleagues are currently working on in collaboration with Mike White, an experimentalist.

The goal of catalyst research is to find a "Goldilocks" catalyst—one that does neither too much nor too little, but is "just right" for a particular application. By revealing why a specific catalyst enables a specific reaction, studies like this one provide the basic understanding needed to design the catalysts we need for sustainable energy production.

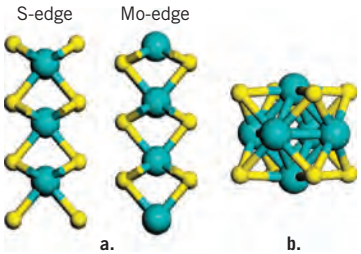

Figure 6. Optimized edge surface of (a) MoS_2 and (b) Mo_6S_8 cluster (Mo = cyan, S = yellow). *Image: Liu et al., 2010.*

[3] Ping Liu, YongMan Choi, Yixiong Yang, and Michael G. White, "Methanol synthesis from H_2 and CO_2 on a Mo_6S_8 cluster: A density functional study," J. Phys. Chem. A **114**, 3888 (2010).

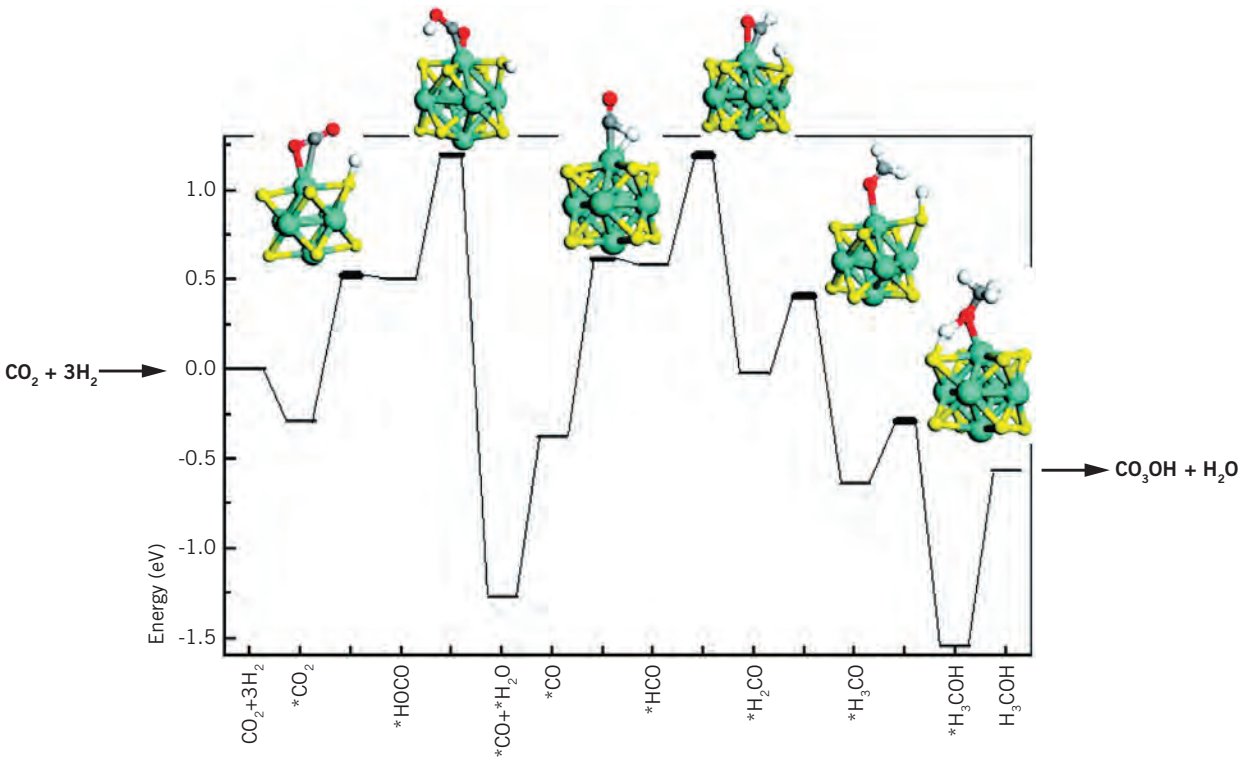

Figure 7. Optimized potential energy diagram for methanol synthesis from CO_2 and H_2 on a Mo_6S_8 cluster. Thin bars represent the reactants, products, and intermediates; thick bars stand for the transition states. The geometries of the transition states involved in the reaction are also included (Mo = cyan, S = yellow, C = gray, O = red, H = white). *Image: Liu et al., 2010.*

Down with Carbon Dioxide

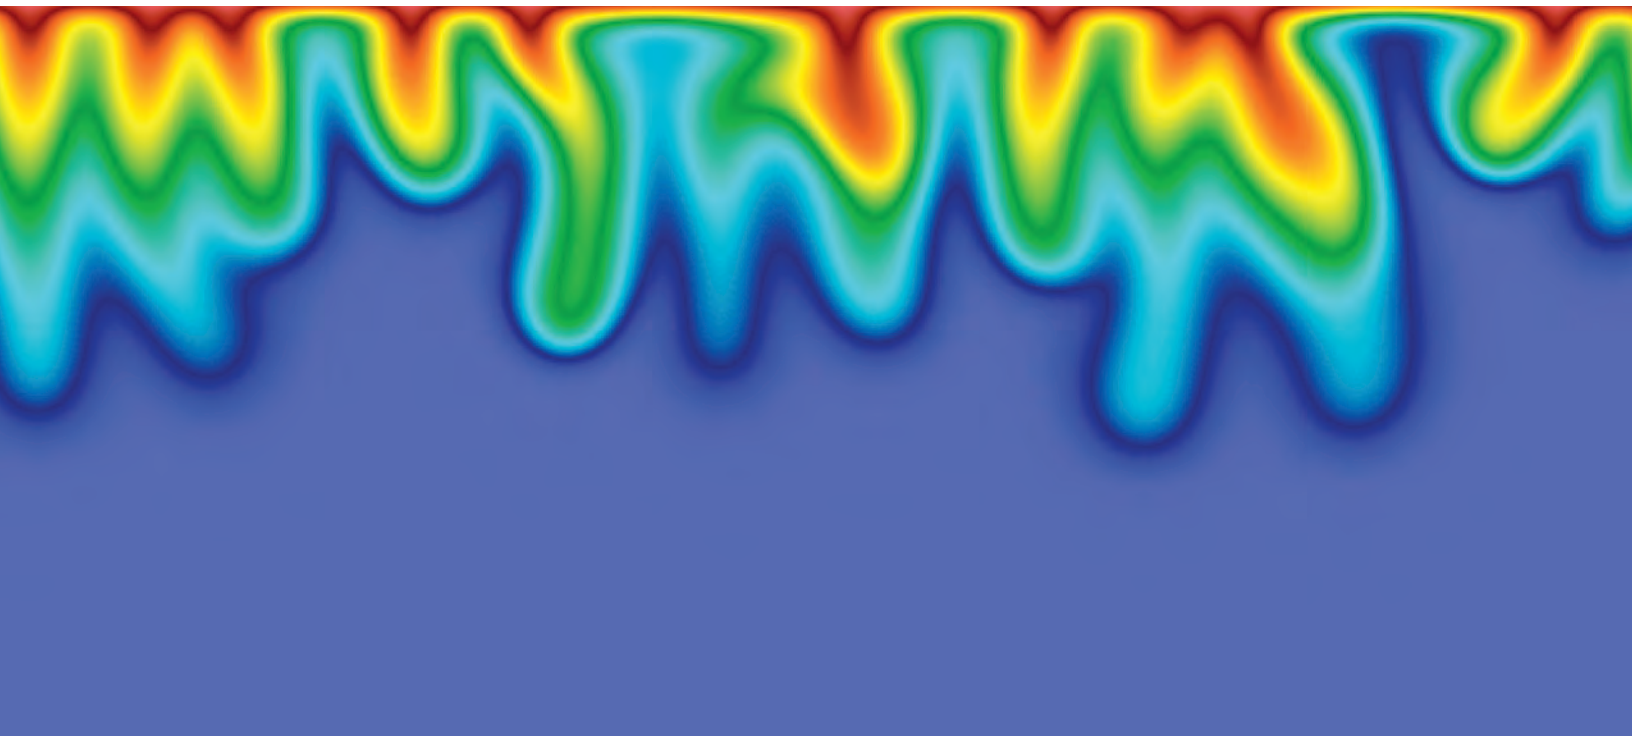

Detailed computational models help predict the performance of underground carbon sequestration projects

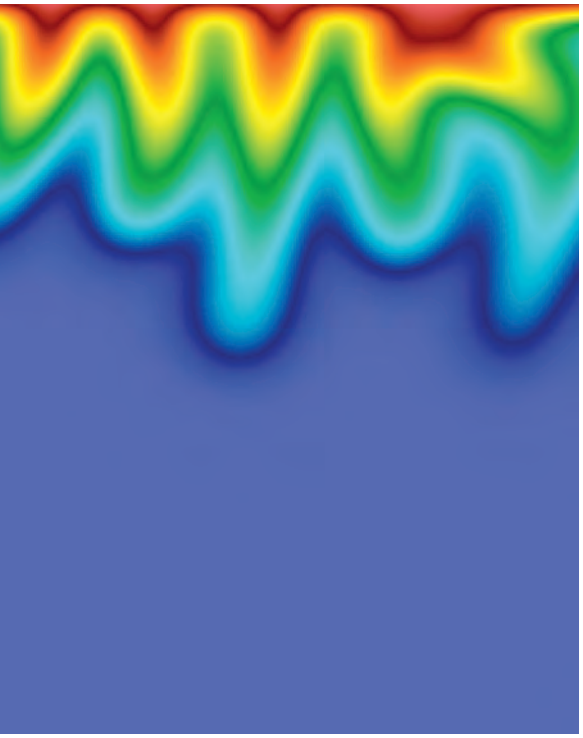

Despite progress in clean energy, Americans will continue to rely on fossil fuels for years to come. In fact, coal-, oil- and natural gas-fired power plants will generate 69 percent of U.S. electricity as late as 2035, according to the U.S. Energy Information Administration.

Such sobering projections have sparked a wide range of proposals for keeping the resulting carbon dioxide (CO_2) out of the atmosphere, where it traps heat and contributes to global warming. Berkeley Lab scientists are using computer simulations to evaluate one promising idea: Pump it into saltwater reservoirs deep underground.

Underground, or geologic, carbon sequestration "will be key tool in reducing atmospheric CO_2," says George Pau, a Luis W. Alvarez Postdoctoral Fellow with Berkeley Lab's Center for Computational Sciences and Engineering (CCSE). "By providing better characterizations of the processes involved, we can more accurately predict the performance of carbon sequestration projects, including the storage capacity and long-term security of a potential site."

Taking advantage of NERSC's massively parallel computing capacity, CCSE researchers led by John Bell have created the most detailed models yet of the mixing processes that occur at the interface of the sequestered carbon dioxide and brine.[4] These simulations—including the first-ever three-dimensional ones—will help scientists better predict the success of this kind of sequestration project.

Project: Simulation and Analysis of Reacting Flows

PI: John B. Bell, Lawrence Berkeley National Laboratory

Senior Investigators: Ann S. Almgren, Marcus S. Day, LBNL

Funding: ASCR, BES

Computing Resources: NERSC, LBNL

Selecting a Site

Carbon sequestration ideas run a wide gamut, from mixing CO_2 into cement to pumping it into the deepest ocean waters. However, geologic sequestration projects are promising enough that some are already under way (Figure 8). In this type of sequestration, CO_2 from natural gas-, oil-, or coal-fired power

[4] G. S. H. Pau, J. B. Bell, K. Pruess, A. S. Almgren, M. J. Lijewski, and K. Zhang, "High resolution simulation and characterization of density-driven flow in CO_2 storage in saline aquifers," Advances in Water Resources **33**, 443 (2010).

plants is pressure-injected deep into stable geologic formations. "The natural first picks [for these sites] are depleted oil and gas reservoirs," says Karsten Pruess, a hydrologist with Berkeley Lab's Earth Sciences Division and a co-author of the study.

However, the problem with this "put-it-back-where-it-comes-from" notion is that combustion generates about three times more CO_2, by volume, than the fossil fuels burned. "So, even if all these old reservoirs were good targets—and not all are—the capacity available isn't nearly enough," Pruess concludes.

Saline aquifers offer a promising alternative. The brine that permeates these porous rock formations is unusable for crops or drinking, but given enough time, it can lock away some CO_2 in the form of a weak acid. Meanwhile, the solid-rock lid that prevents saline from welling up to the surface (the caprock) also traps injected CO_2 underground.

According to Pruess, when CO_2 is first injected into an aquifer, it forms a separate layer, a bubble of sorts, above the saline. Over time, some CO_2 will diffuse into the brine, causing the brine density to increase slightly. As more of the CO_2 diffuses into the water, the resulting layer of carbonic acid (now less buoyant than the water below it) sinks. This churns up fresh brine, and a convective process kicks in

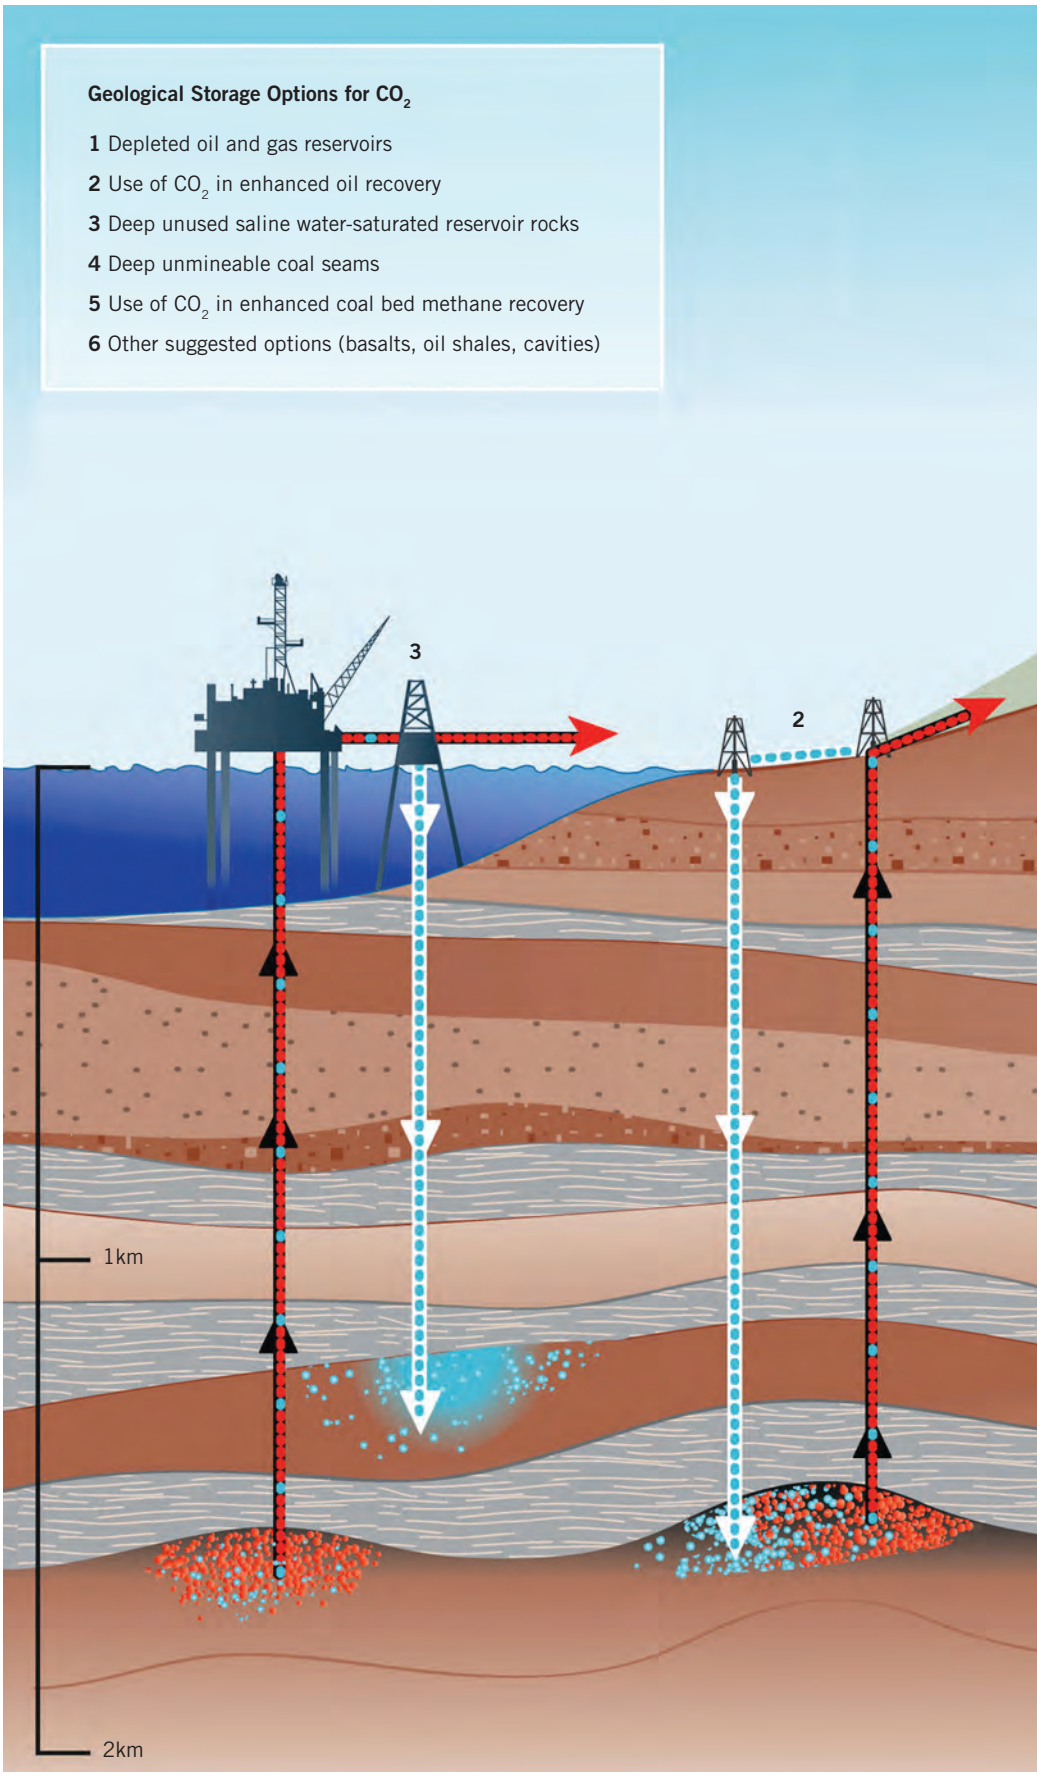

Figure 8. Geologic sequestration in saline aquifers (3) is shown in this illustration alongside other geologic sequestration ideas. *Image: Australian Cooperative Research Centre for Greenhouse Gas Technologies (CO2CRC).*

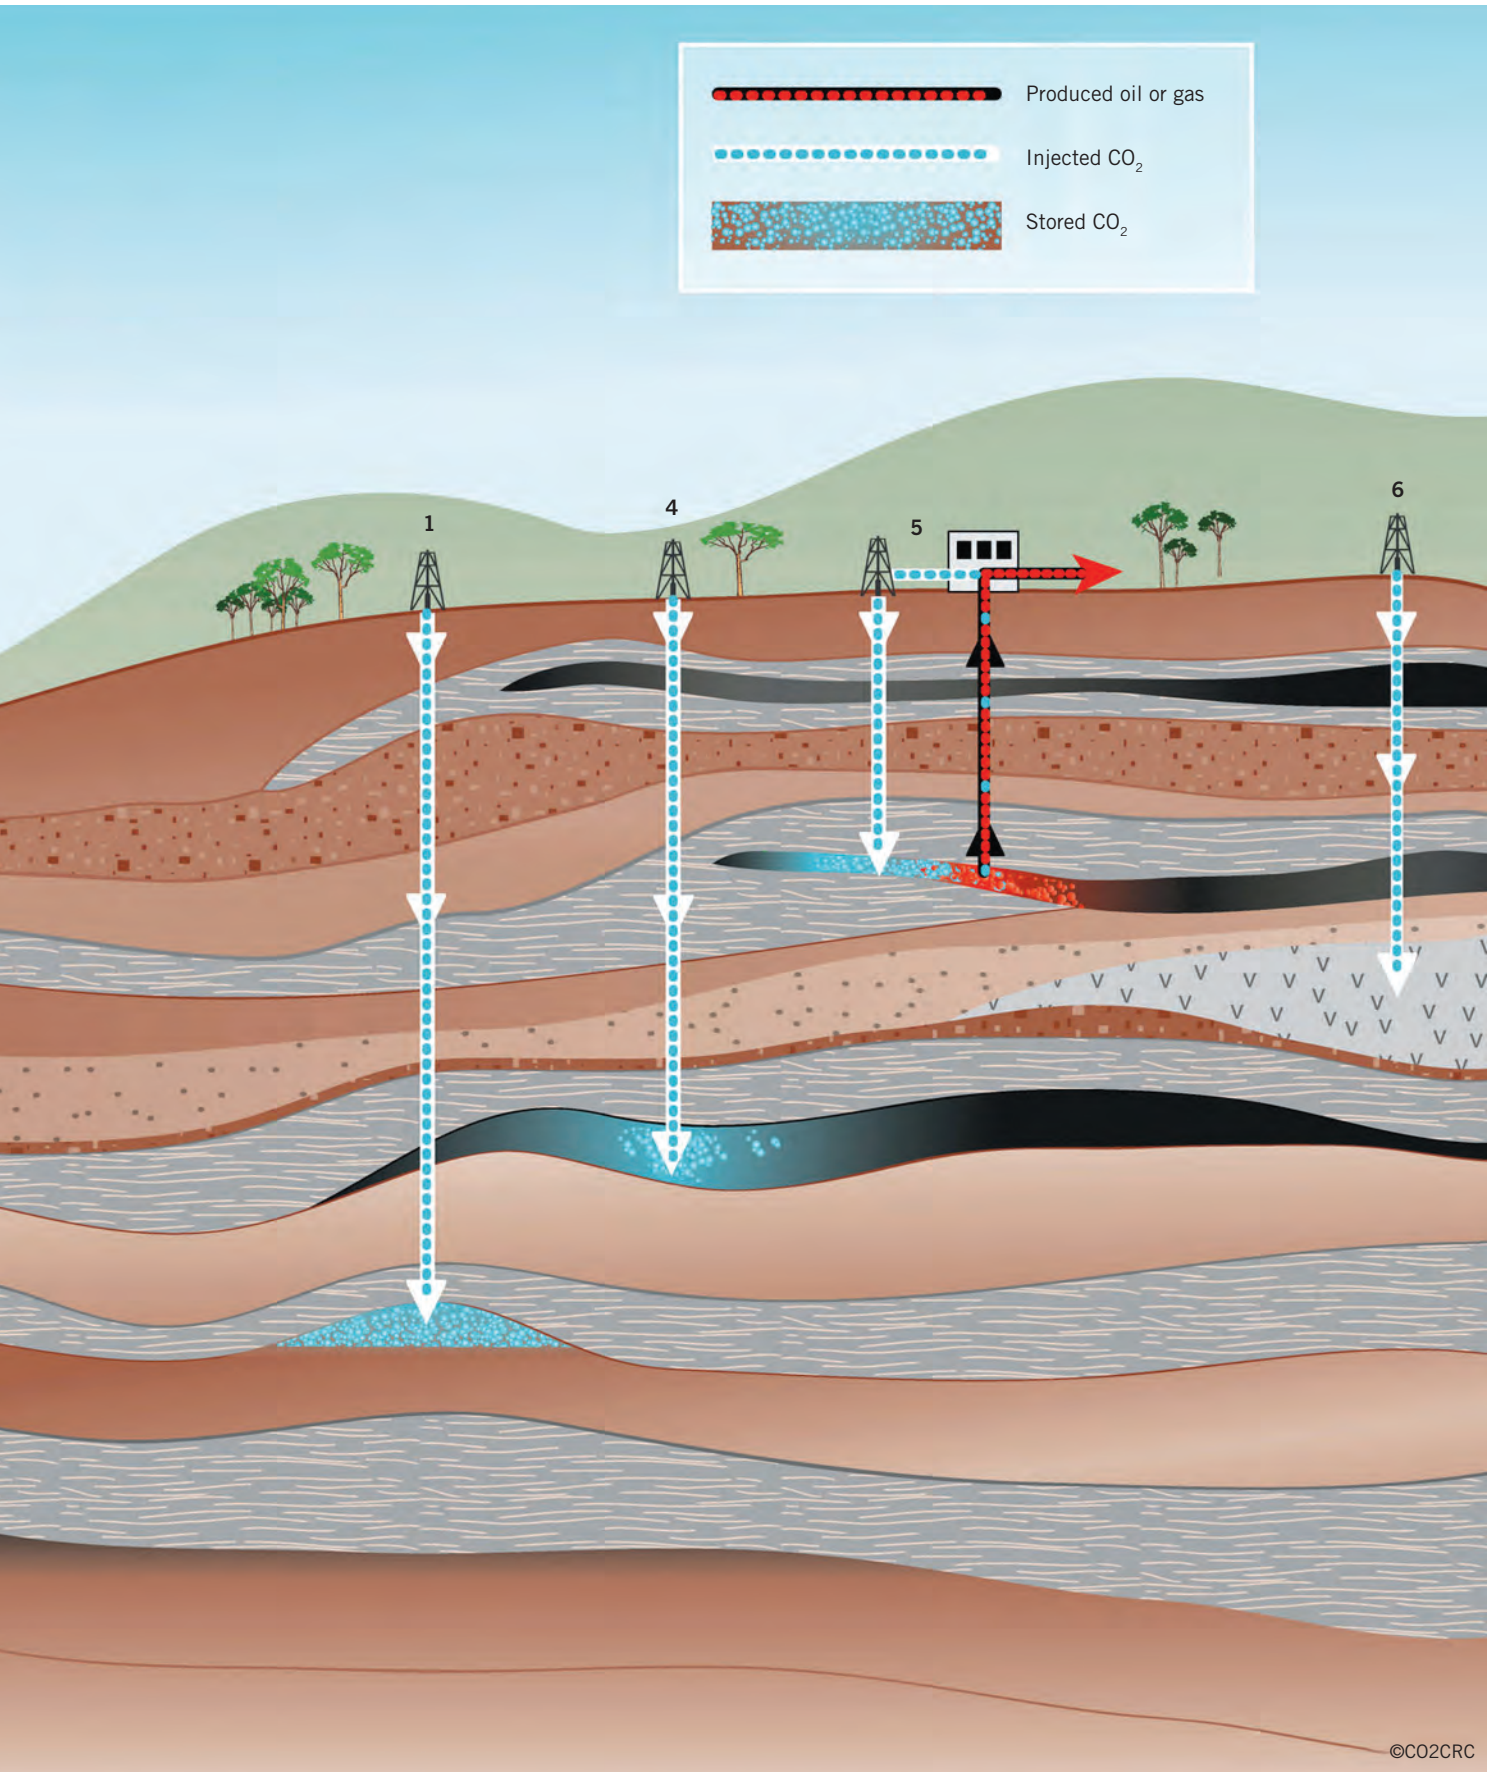
Produced oil or gas
Injected CO_2
Stored CO_2
1
4
5
6
©CO2CRC

which dramatically increases the rate of CO_2 uptake, compared to diffusion alone (Figures 9 and 10). This process can take hundreds, even thousands of years, but scientists haven't fully quantified the factors that influence and trigger convection.

"The processes by which CO_2 dissolves into the brine are of serious consequence," says Pau. "If you can determine the rate this will occur, then you can more accurately determine how much CO_2 can be stored in a given aquifer."

Because scientists can't wait hundreds of years for answers, they are using NERSC supercomputers to run mathematical models that simulate the processes involved.

Geological models that break the problem into fixed-size pieces cannot reveal phenomena needed to understand sequestration because they have either insufficient detail or require too much computation to be practical. "About 100 meters resolution is the best we can do," Pruess says. Because convection effects start at scales smaller than a millimeter, they are typically invisible to these geological models.

Cracking the Sequestration Code

That's why Berkeley Lab's CCSE and Pruess collaborated to develop a new simulation code. The code combines a computing technique called adaptive mesh refinement (AMR) with high performance parallel computing to create high-resolution simulations. The team's simulations were performed using 2 million processor-hours and running on up to 2,048 cores simultaneously on NERSC's Cray XT4 system, Franklin.

Called Porous Media AMR, or PMAMR, the code concentrates computing power on more active areas of a simulation by breaking them into finer parts (higher resolution) while calculating less active, less critical portions in coarser parts (lower resolution). The size of the chunks, and thus the resolution, automatically shifts as the model changes, ensuring that the most active, critical processes are also the most highly resolved. This ability to automatically shift focus makes AMR a powerful tool for modeling phenomena that start small and grow over time, says Pruess.

The first PMAMR models were of modest size, on the scale of meters. But the eventual goal is to be able to use the characteristics of an aquifer—such as the size of the pores in the sandstone, the amount of salts in the water, or the composition of the reservoir rocks—to model a given site and predict how much CO_2 it can accommodate.

3D simulations of carbon sequestration are a first step toward predicting how much CO_2 an aquifer can hold.

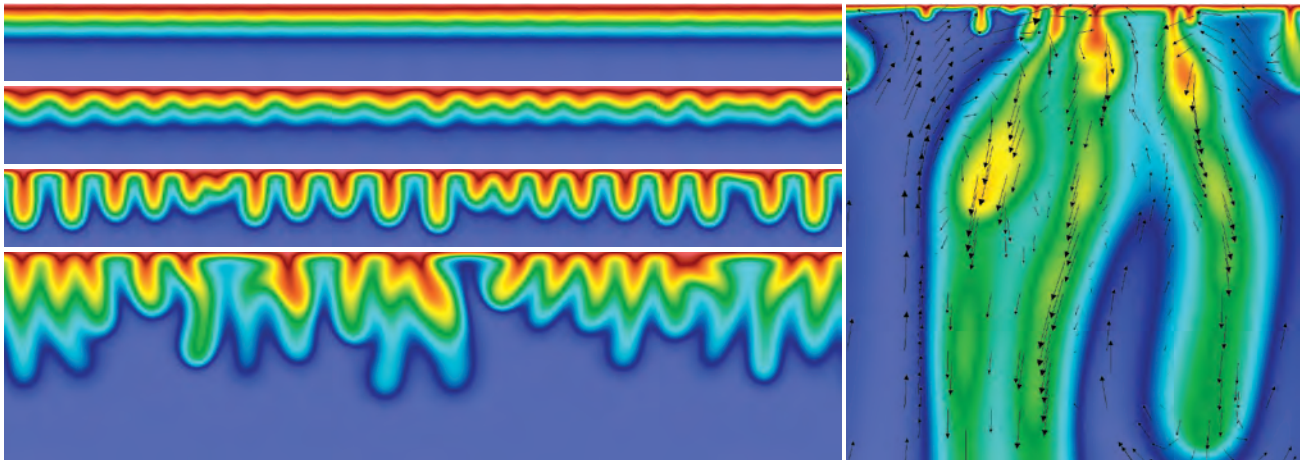

Figure 9. As CO_2 diffuses into the brine, the top, more concentrated layers (indicated in red) begin to sink, triggering a convective process that churns up fresh brine and speeds the salt water's CO_2 uptake. The arrows in the frame at the right indicate the direction of flow. Understanding convection onset is critical to understanding how much CO_2 a given aquifer can store. *Image: George Pau.*

PMAMR simulations might also be used to help geologists more accurately track and predict the migration of hazardous wastes underground or to recover more oil from existing wells. And the team's three-dimensional simulations could help scientists uncover details not apparent in two-dimensional models.

The next challenge the team faces involves extending and integrating their models with large-scale models. "If you look how much CO_2 is made by a large coal-fired power plant, and you want to inject all that underground for 30 years, you end up with a CO_2 plume on the order of 10 km in dimension," says Pruess. This free CO_2 will continue to migrate outward from the injection site and, if the aquifer has a slope, could eventually reach drinking water aquifers or escape at outcrops hundreds of kilometers away. "We are very interested in the long-term fate of this thing," says Pruess.

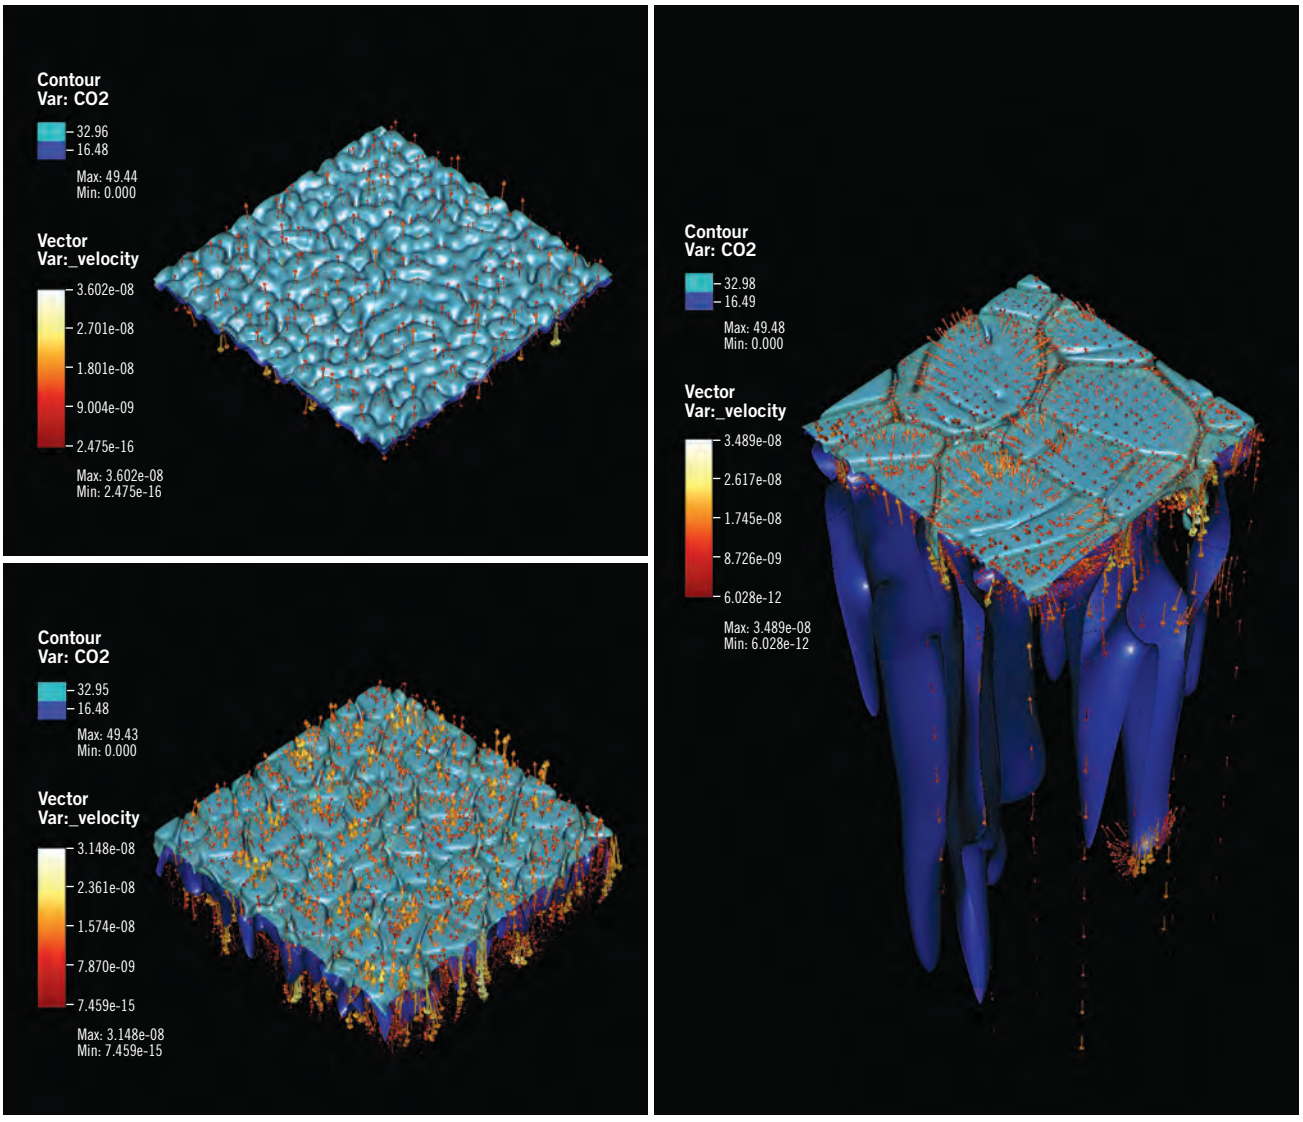

Figure 10. These three-dimensional simulation snapshots show the interface between sequestered CO_2 and brine over time. The sequence, starting at the upper left frame and proceeding counterclockwise, shows a dramatic evolution in the size of the convection cells with time. *Image: George Pau.*

Touching Chaos

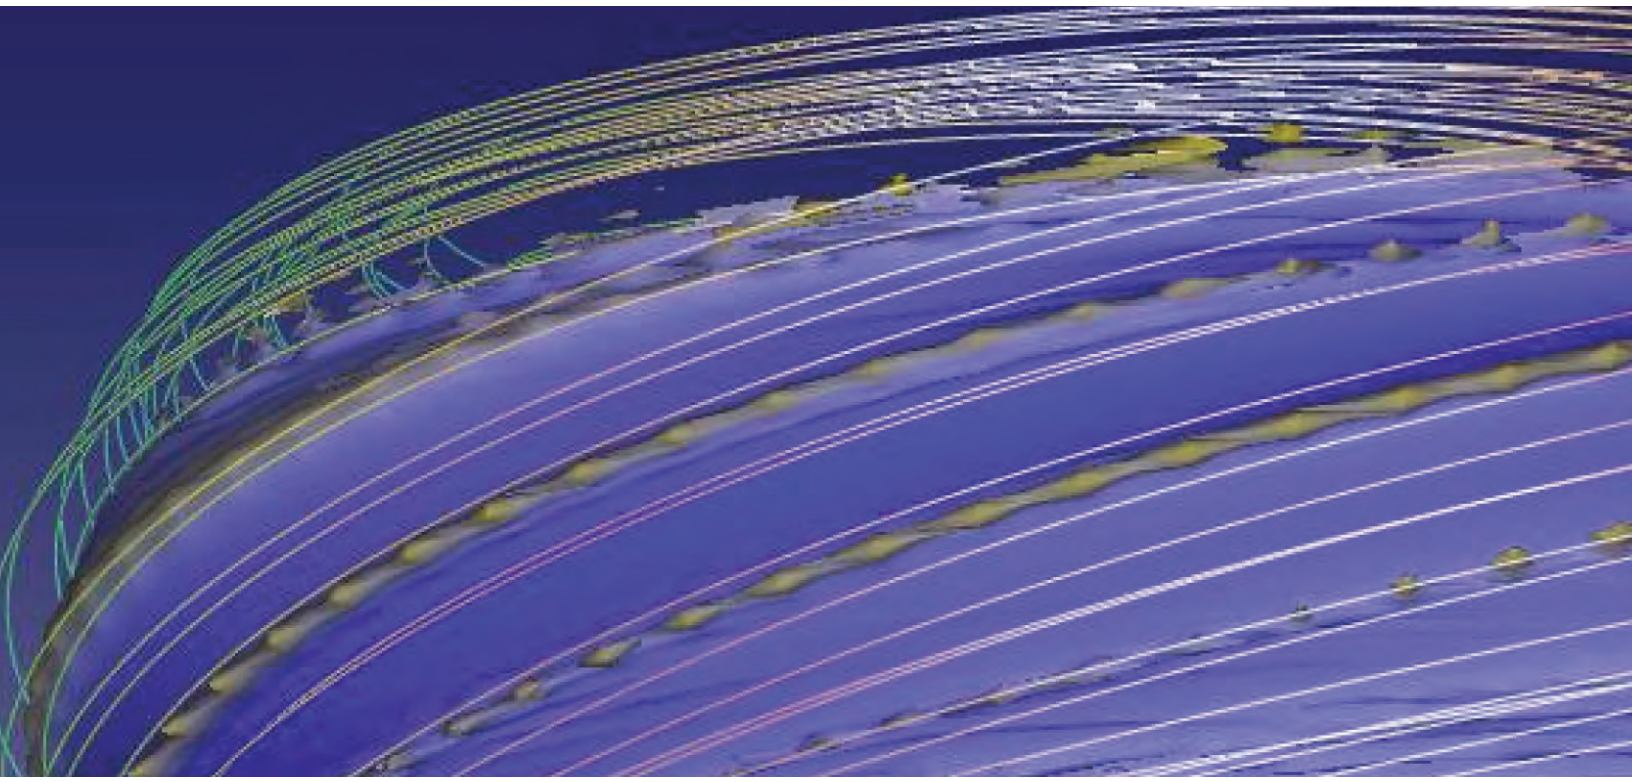

Simulations reveal mechanism behind a plasma instability that can damage fusion reactors

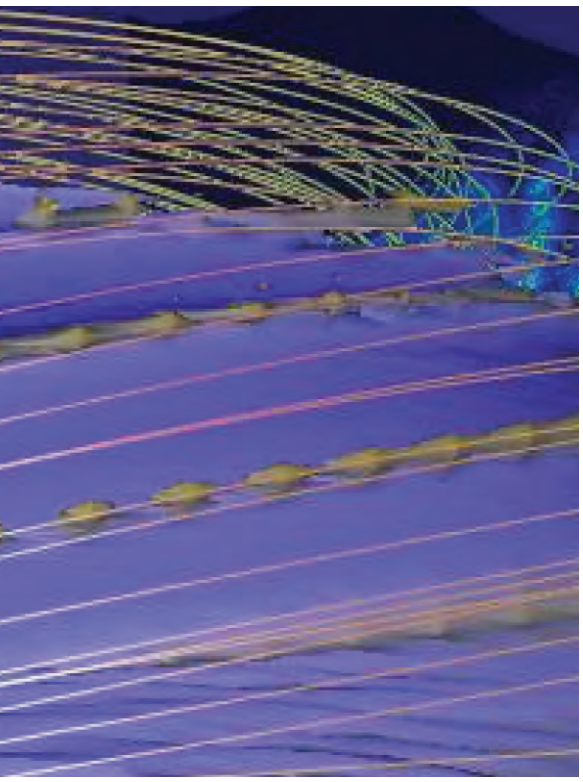

Project: 3D Extended MHD Simulation of Fusion Plasmas

PI: Stephen Jardin, Princeton Plasma Physics Laboratory

Senior Investigators: Joshua Breslau, Jin Chen, Guoyong Fu, Wonchull Park, PPPL; Henry R. Strauss, New York University; Linda Sugiyama, Massachusetts Institute of Technology

Funding: FES, SciDAC

Computing Resources: NERSC

If humans could harness nuclear fusion, the process that powers stars like our Sun, the world could have an inexhaustible energy source. In theory, scientists could produce a steady stream of fusion energy on Earth by heating up two types of hydrogen atoms—deuterium and tritium—to more than 100 million degrees centigrade until they become a gaseous stew of electrically charged particles, called plasma; then use powerful magnets to confine these particles until they fuse together, releasing energy in the process.

Although researchers have achieved magnetic fusion in experiments, they still do not understand the behavior of plasma well enough to generate a sustainable flow of energy. They hope the international ITER experiment, currently under construction in southern France, will provide the necessary insights to make fusion an economically viable alternative energy source. ITER, an industrial-scale reactor, will attempt to create 500 megawatts of fusion power for several-minute stretches. Much current research is focused on answering fundamental questions that will allow engineers to optimize ITER's design.

The DOE's Scientific Discovery through Advanced Computing (SciDAC) Center for Extended Magnetohydrodynamic (MHD) Modeling, led by Steve Jardin of the Princeton Plasma Physics Laboratory, has played a vital role in developing computer codes for fusion modeling. Recently, the collaboration created an extension to the Multilevel 3D (M3D) computer code that allows researchers to simulate what happens when charged particles are ejected from a hot plasma stew and splatter on the walls of a tokamak device, the doughnut-shaped "pot" used to magnetically contain plasmas. According to Linda Sugiyama of the Massachusetts Institute of Technology, who led the code extension development, these simulations are vital for ensuring the safety and success of upcoming magnetic confinement fusion experiments like ITER.

Modeling Fusion to Ensure Safety

Because hot plasma is extremely dangerous, it is imperative that researchers understand plasma instability so that they can properly confine it. For fusion to occur, charged particles must be heated to more than 100 million degrees centigrade. At these searing temperatures, the material is too hot to contain with most earthly materials. In tokamaks, the plasma is shaped into a torus,

or doughnut shape, and held in place within a vacuum by powerful magnetic fields, so that the plasma does not touch the surface of the containment vessel.

Tokamak plasmas typically have one or two magnetic "X-points"—so called because the magnetic field lines form an X shape—at the top or bottom of the plasma surface. X-points help promote efficient fusion: by channeling stray particles along the magnetic field lines to divertors designed to handle them, the X-points help control the plasma shape and prevent instabilities.

Until they don't.

For three decades, tokamak researchers have observed periodic disruptive instabilities, called edge localized modes (ELMs), on the edge of the plasma. ELMs allow blobs of plasma to break out of the magnetic confinement through the protective vacuum and splatter onto the tokamak wall before the plasma stabilizes and resumes its original shape. The number of unstable pulses affects how much plasma is thrown onto the walls, which in turn determines the extent of wall damage.

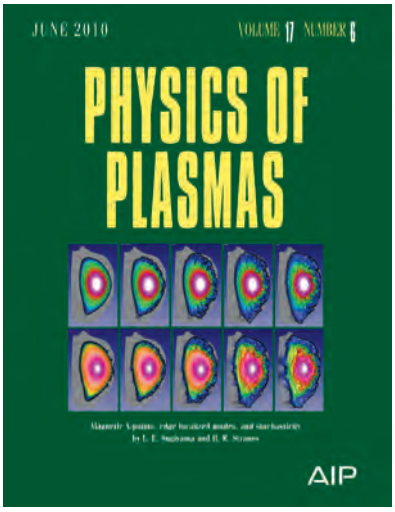

Figure 11. This ELM research was featured on the cover of the June 2010 issue of Physics of Plasmas. *Image: ©2010, American Institute of Physics.*

The origins and dynamics of ELMs—or of stable plasma edges, for that matter—have long resisted theoretical explanation. A breakthrough came when Sugiyama and a colleague, Henry Strauss of the Courant Institute of Mathematical Sciences at New York University, used the M3D extension on NERSC computers to simulate the development of the magnetic fields on a tokamak plasma surface. They reported their results in an article featured on the cover of the June 2010 issue of Physics of Plasmas (Figure 11).[5]

"Studies of nonlinear dynamics show that if you have X-points on the plasma boundary, which is standard for high temperature fusion plasmas, particles can fly out and hit the walls. But no one had ever seen what actually happens inside the plasma if you let it go," says Sugiyama. "This is the first time that we have seen a stochastic magnetic tangle, a structure well known from Hamiltonian chaos theory, generated by the plasma itself. Its existence also means that we will have to rethink some of our basic ideas about confined plasmas."

What Sugiyama and Strauss discovered is that ELMs constitute a new class of nonlinear plasma instability involving the mutual influence of growing plasma structures and the magnetic field. A minor instability anywhere on a plasma boundary that possesses an X-point can perturb the magnetic field, causing it to become chaotic. As the plasma instability grows, magnetic chaos increases, allowing the plasma bulges to penetrate deep into the plasma and also to break out of the boundary surface, resulting in a loss of plasma (Figures 12 and 13). This demonstration that the plasma motion can couple to a chaotic magnetic field may help explain the large range of ELM and ELM-free edge behaviors in fusion plasmas, as well as clarifying some properties of magnetic confinement.

Sugiyama notes that the primary causes of plasma instability vary in the different plasma regions of the tokamak, including the core, edge, and surrounding vacuum. To understand the most dangerous of these instabilities, computational physicists model each plasma region on short time and space scales. The plasma regions are strongly coupled through the magnetic and electric fields, and researchers must carry out an integrated simulation that shows how all the components interact in the big picture, over long time scales.

"Modeling plasma instabilities is computationally challenging because these processes occur on widely differing scales and have unique complexities," says Sugiyama. "This phenomenon couldn't be investigated without computers. Compared to other research areas, these simulations are not computationally

[5] L. E. Sugiyama and H. R. Strauss, "Magnetic X-points, edge localized modes, and stochasticity," Physics of Plasmas **17**, 062505 (2010).

large. The beauty of using NERSC is that I can run my medium-size jobs for a long time until I generate all the time steps I need to see the entire process accurately."

She notes that the longest job ran on 360 processors for 300 hours on NERSC's Cray XT4 Franklin system. However, she also ran numerous other jobs on the facility's Franklin and DaVinci systems using anywhere from 432 to 768 processors for about 200 CPU hours.

"I greatly appreciate the NERSC policy of supporting the work required to scale up from small jobs to large jobs, such as generating input files and visualizing the results," says Sugiyama. "The center's user consultants and analytics group were crucial to getting these results."

Her ongoing research has two major goals. One is to verify the simulation model against experimental observations. "This is proving surprisingly difficult," Sugiyama says, "primarily because the instability is extremely fast compared to most experimental measurements—changes occur over a few microseconds—and it is very difficult to measure most things occurring inside a hot dense plasma, since material probes are destroyed immediately. We have good contacts with all the major U.S. experiments and some international ones, but have been waiting for good data."

Her second goal is to characterize the ELM threshold in a predictive fashion for future fusion plasmas, including new experiments such as ITER. This involves studying small ELMs and ELM-stable plasmas, where there are other, smaller MHD instabilities. Demonstrating the predictive value of this ELM model, besides being a significant theoretical achievement, will help bring practical fusion energy one step closer to realization.

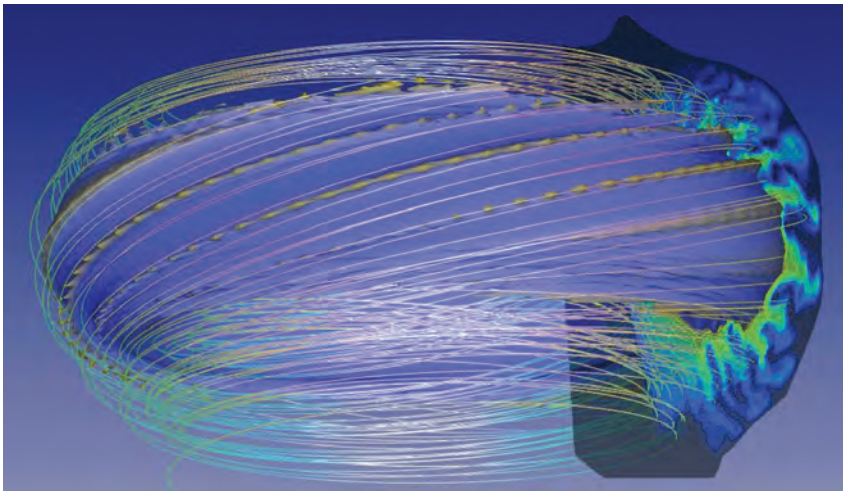

Figure 12. Temperature surface near the plasma edge shows helical perturbation aligned with the magnetic field. *Image: L. E. Sugiyama and the M3D Team.*

"This is the first time that we have seen a stochastic magnetic tangle … generated by the plasma itself."

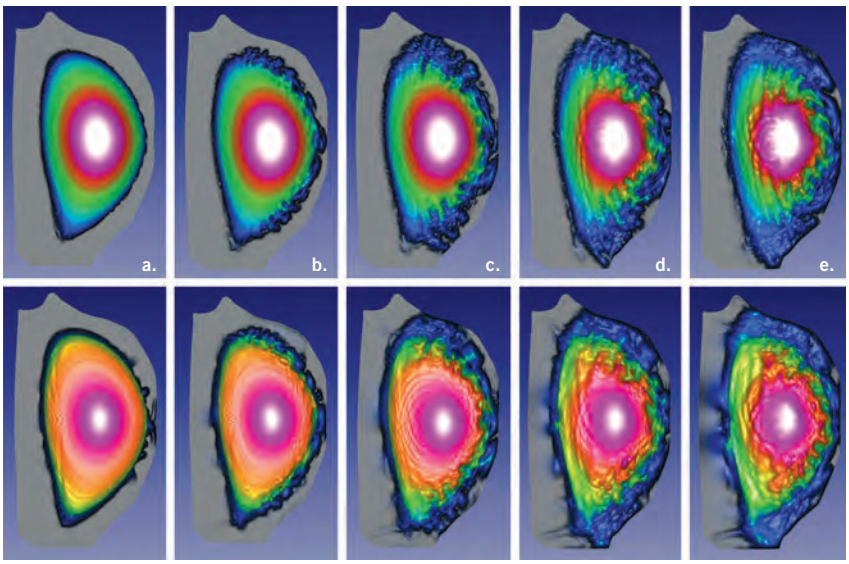

Figure 13. This image shows a simulation of the initial stage of a large edge instability in the DIII-D tokamak. The top row shows contours of the plasma temperature in a cross section of the torus with its central axis to the left. The bottom row shows corresponding density. The vacuum region between plasma and wall is grey; plasma expands rapidly into this region and hits the outer wall, then gradually subsides back to its original shape. *Image: L. E. Sugiyama and H. R. Strauss.*

Winds of Climate Change

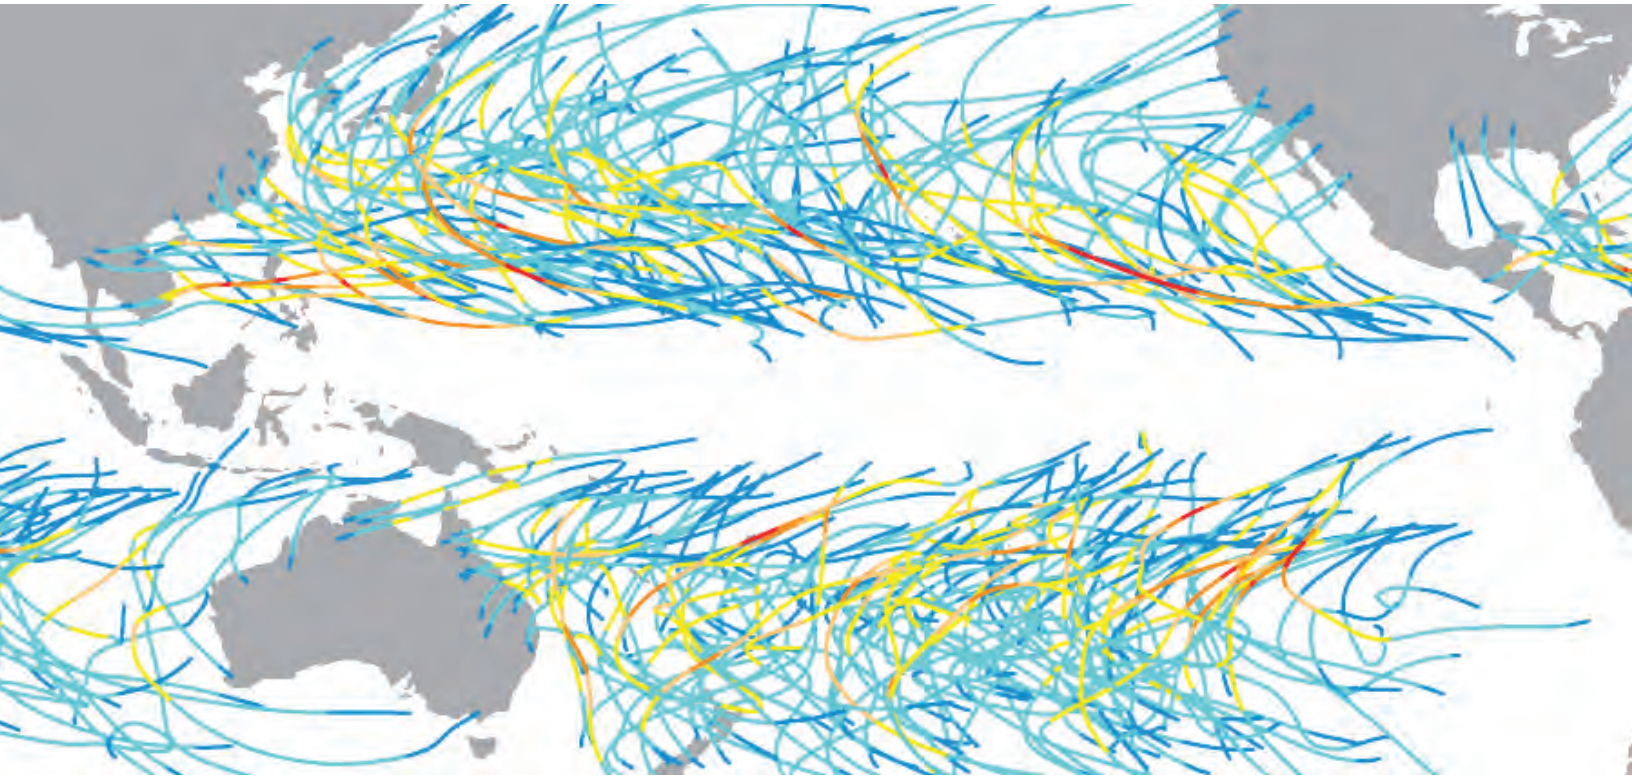

Cyclones in the warm Pliocene era may have implications for future climate

Scientists searching for clues to Earth's future climate are turning to its dim past, the Pliocene epoch. "We're interested in the Pliocene because it's the closest analog we have in the past for what could happen in our future," says Chris Brierley, a Yale climate researcher computing at NERSC.

Brierley and Yale colleague Alexey Fedorov are using simulations to help unravel a mystery that has bedeviled climatologists for years: Why was the Pliocene—under conditions similar to today—so much hotter? Their most recent work suggests that tropical cyclones (also called hurricanes or typhoons) may have played a crucial role.[6]

Working with Kerry Emanuel of MIT, the Yale researchers' computer simulations revealed twice as many tropical cyclones during the Pliocene; storms that lasted two to three days longer on average (compared to today); and, unlike today, cyclones occurred across the entire tropical Pacific Ocean (Figure 14). This wide-ranging distribution of storms disrupted a key ocean circulation pattern, warming seas and creating atmospheric conditions that encouraged more cyclones.

This positive feedback loop, they posit, is key to explaining the Pliocene's higher temperatures and persistent El Niño-like conditions. Their findings, featured on the cover of the journal Nature (Figure 15), could have implications for the planet's future as global temperatures continue to rise due to climate change.

What's the Pliocene Got to Do with It?

The Pliocene Epoch (the period between 5 and 3 million years ago) is especially interesting to climatologists because conditions then were similar to those of today's Earth, including the size and position of land masses, the levels of sunlight, and, crucially, atmospheric concentrations of the greenhouse gas carbon dioxide (CO_2). "In geological terms, it was basically the same," Brierley says.

Project: Tropical Cyclones and the Shallow Pacific Overturning Circulation

PI: Alexey Fedorov, Yale University

Senior Investigator: Christopher Brierley, Yale University

Funding: BER, NSF, David and Lucile Packard Foundation

Computing Resources: NERSC

[6] Alexey V. Fedorov, Christopher M. Brierley, and Kerry Emanuel, "Tropical cyclones and permanent El Niño in the early Pliocene epoch," Nature **463**, 1066 (2010).

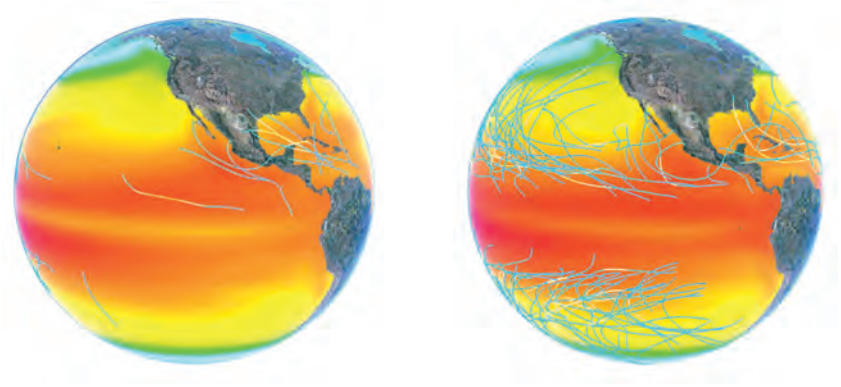

Figure 14. Simulated hurricane tracks under modern conditions (left) and Pliocene (right) overlay sea surface temperatures created by a fully coupled climate model. *Images: Chris M. Brierley.*

Yet the Pliocene was puzzlingly hotter. Global mean temperatures ranged up to 4° C (about 8° F) higher than today; sea levels were 25 meters higher; and sea ice was all but absent. The poles heated up, yet temperatures in the tropics did not exceed modern-day levels. And the planet was locked into a continuous El Niño-like state, as opposed to the on/off oscillation experienced today.

Fedorov and Brierley suspected cyclones played a role. "We knew from a previous study that additional vertical mixing [in the ocean] can cause warming in the Niño region" of the eastern tropical Pacific, says Brierley. "Hurricanes were one of the few sources of mixing that we could envisage changing," he explains.

Cyclones Stir the Pot

To test their idea, the team first simulated atmospheric flows for modern Earth and for the Pliocene (see sidebar). Because most other factors were basically the same, the main difference between the two was the Pliocene's higher sea surface temperatures.

They then randomly seeded vortices and tracked the results. "We found there were twice as many tropical cyclones during this period [versus modern times], and that they lasted longer," two to three days longer on average, says Brierley.

Hurricane paths shifted, too (Figure 16). "Rather than starting in the far western Pacific, hitting Japan, then heading back out into the northern Pacific, they were all over the place," striking all along two bands north and south of the equator, says Brierley.

The next step was to integrate the extra ocean-churning action of all those hurricanes into a more complete climate model. To do that, the team applied an idealized pattern of hurricane ocean-mixing to a fully coupled climate model, one that includes ocean, sea ice, land cover, and atmosphere.

Today, cold water originating off the coasts of California and Chile skirts the region of tropical cyclone activity on its way to the Equator, where it surfaces in a "cold tongue" that stretches west off the coast of South America. El Niño events, which affect precipitation and temperature patterns worldwide, are characterized by an uptick in the temperature of the cold tongue. It only took about 50 model-years for the cold tongue to begin warming in the researchers' Pliocene model. After 200 years, El Niño-like warming was persistent (Figure 17).

Brierley, Fedorov, and Emanuel suggest tropical cyclones were to blame, at least in part. More numerous and more widely distributed across the globe, the cyclones disrupted the cold tongue, raising ocean temperatures and leading to changes in the atmosphere that favored the formation of even more tropical storms. The cycle repeated, helping to sustain the Pliocene's higher overall temperatures.

"We're suggesting that one of the reasons for the differences [between the Pliocene and today] is this feedback loop that isn't operating in our present climate," says Brierley.

Figure 15. These simulations were featured on the cover of the 25 February 2010 issue of Nature. *Image ©2010 Macmillan Publishers Limited.*

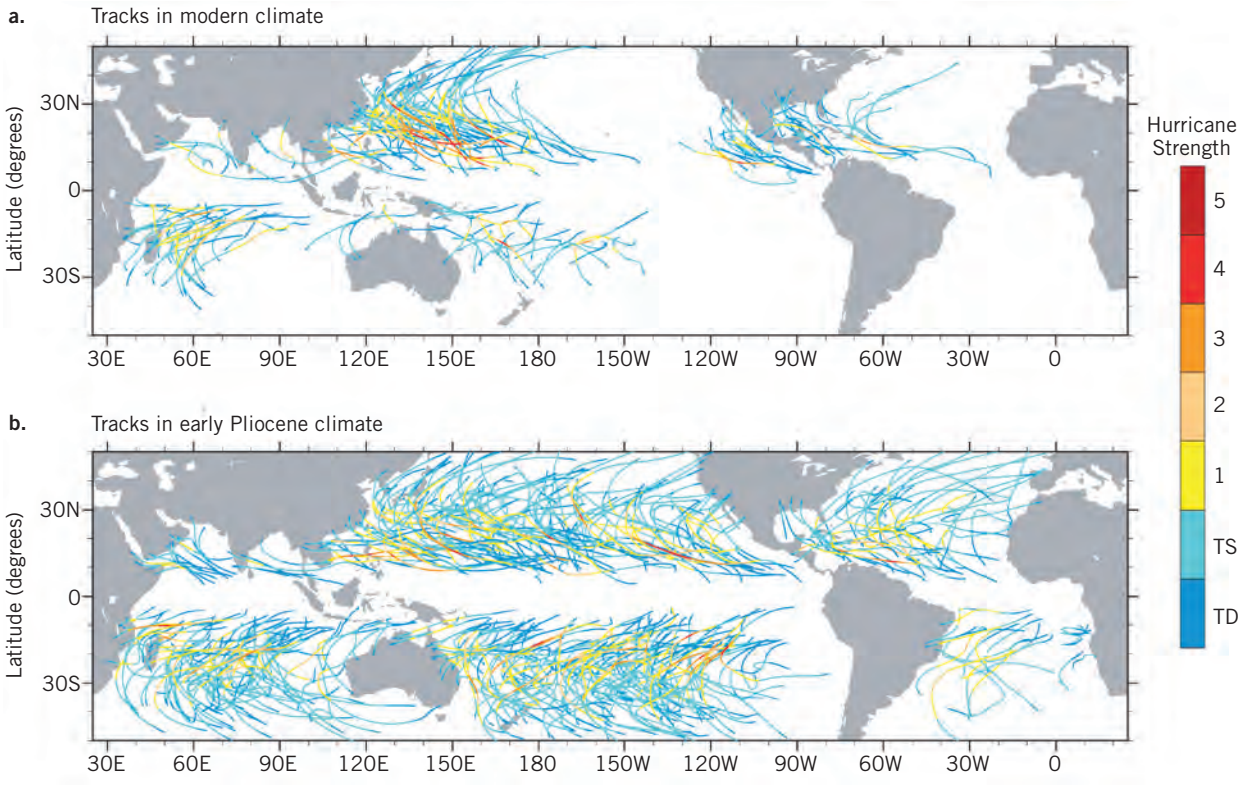

Figure 16. Simulated hurricane tracks in (a) modern climate and (b) Pliocene climate illustrate a two-year subsample from 10,000 simulated tropical cyclones. Colors indicate strength from tropical depression (blue) to category-5 hurricane (red). *Images: Chris M. Brierley.*

Powering Climate Models

Calculating the complex, multifaceted interactions of land, sea, atmosphere, and ice requires specialized and massive computing power and storage. NERSC's Cray XT4, Franklin, and the NERSC High Performance Storage System (HPSS) both played important roles in Fedorov and Brierley's enquiries.

For this investigation, the team used 125,000 processor hours on Franklin to run fully coupled climate models. "The current resources that we have locally achieve a rate of about three model years per day, whereas Franklin can achieve around 25 years per day," says Brierley. Models run on Franklin in just a month would have taken six months or more locally. "It's not just the number and speed of individual processors, but the speed at which they can pass information" that makes it feasible to run hundreds of years long simulations, Brierley says.

The team was also able to get their models up and running quickly because of Franklin's pre-optimized, pre-installed climate modeling code. "I only had to set up the problem to do the science, rather than dealing with all the trouble of getting the code running first," says Brierley.

Because of NERSC's mass storage facility, HPSS, "I was able to store the vast amounts of data generated and then only take the fields I need at the time to a local machine to run visualization software," Brierley says. Safely stored, the datasets can be used again in future rather than forcing scientists to repeat computationally expensive runs.

Will History Repeat?

The researchers don't know if modern-day climate will eventually develop the same feedback loop, because they aren't sure what set if off the first place. "There must have been a trigger, some sort of tipping point, but we don't yet know what that may have been," Brierley explains.

Fedorov cautions against drawing direct comparisons between our future climate and this past epoch. The current scientific consensus projects that the number of intense hurricanes will increase under global warming, but the overall number will actually fall, he says. "However, unless we understand the causes of these differences, we will not be sure whether our projections are correct," explains Fedorov. "Changes in the frequency and distribution of these storms could be a significant component of future climate conditions."

For their next step, the team is refining and rerunning their models at NERSC. "We've done all this assuming that you can average the behavior of hurricanes into one annual mean signal," says Brierley. The team is working on more realistic hurricane mixing, "where you have a very strong forcing for a short amount of time followed by a long quiescence," says Brierley. "We need to know if this makes a significant difference in our models."

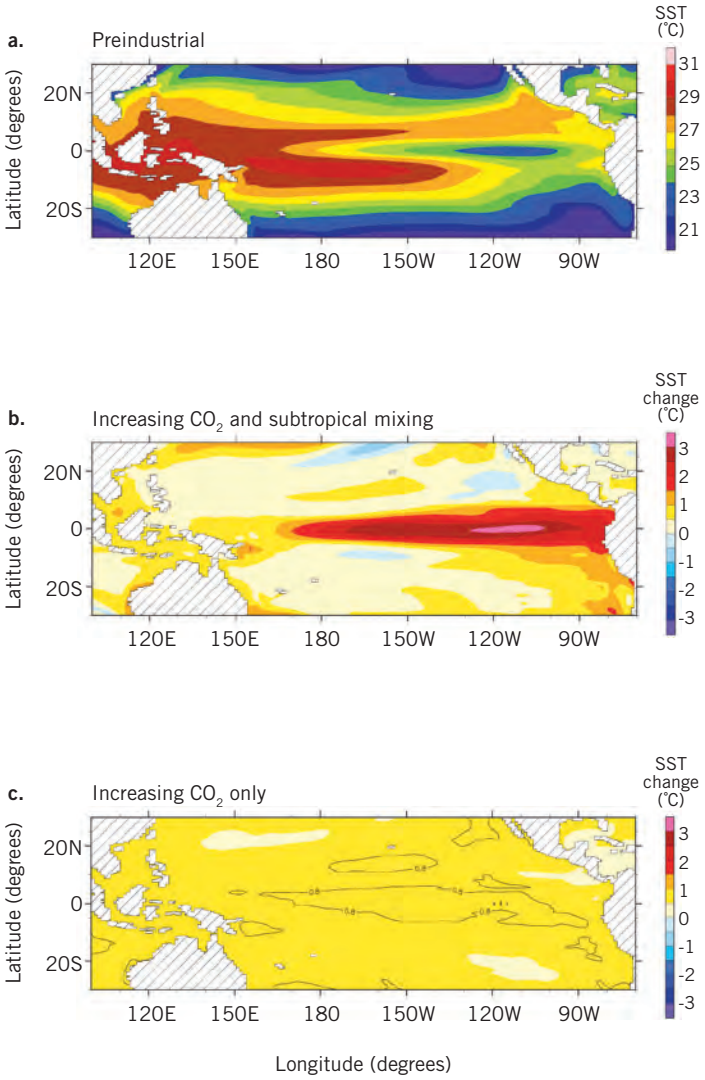

Figure 17. Sea surface temperature changes in the tropical Pacific as simulated by the coupled model are shown here. In the top frame is a baseline of temperatures using pre-industrial CO_2 levels. The center frame shows the amount of warming (over baseline) under Pliocene CO_2 levels with the water-churning effects of intense hurricane activity. The bottom frame shows the effects of higher CO_2 levels alone. After 200 years of calculations, the Pliocene scenario (center) shows warming in the eastern equatorial Pacific is high, the signature of El Niño-like conditions. *Images: Chris M. Brierley.*

Tropical cyclones disrupted the "cold tongue" in the Pacific Ocean, raising ocean temperatures and leading to changes in the atmosphere that favored the formation of even more tropical storms.

3D Is the Key

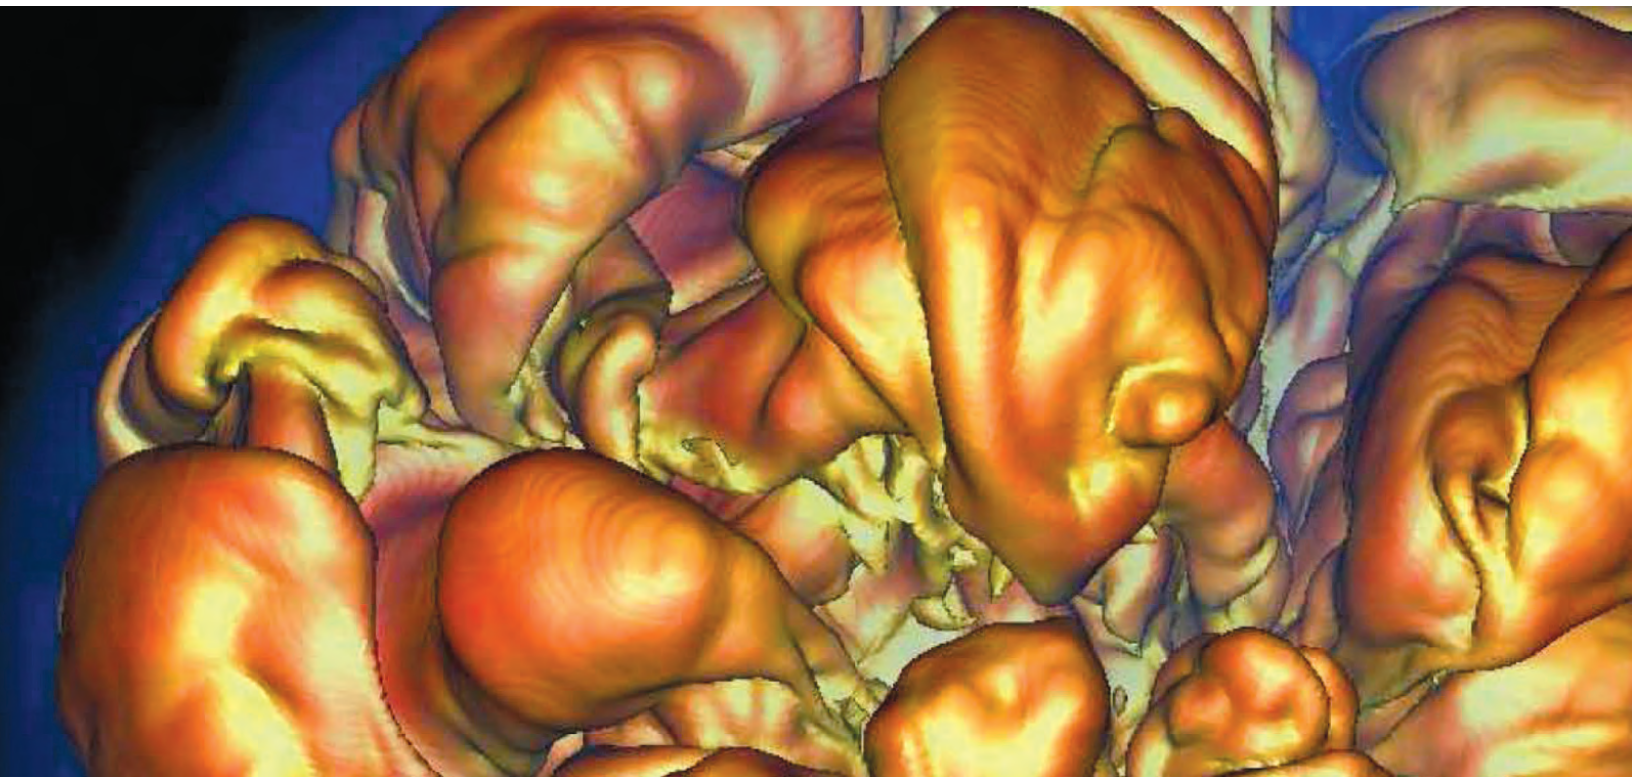

Simulated supernova explosions happen easier and faster when third dimension is added

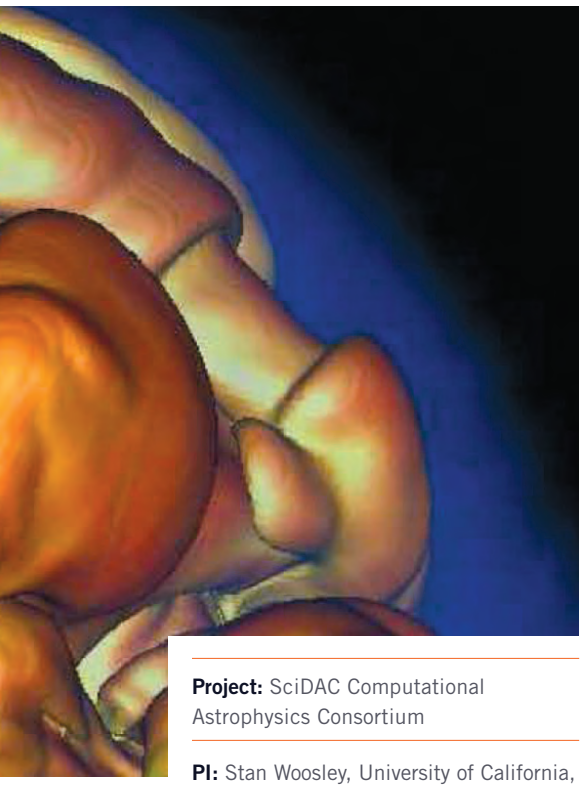

Project: SciDAC Computational Astrophysics Consortium

PI: Stan Woosley, University of California, Santa Cruz

Senior Investigators: Adam Burrows, Princeton University; Peter Nugent, John Bell, and Ann Almgren, Lawrence Berkeley National Laboratory; Michael Zingale, State University of New York, Stony Brook; Louis Howell and Robert Hoffman, Lawrence Livermore National Laboratory; Alex Heger, University of Minnesota; Daniel Kasen, University of California, Santa Cruz

Funding: HEP, SciDAC, NSF

Computing Resources: NERSC, TIGRESS, NICS, TACC, ALCF

For scientists, supernovae are true superstars—massive explosions of huge, dying stars that shine light on the shape and fate of the universe. Recorded observations of supernovae stretch back thousands of years, but only in the past 50 years have researchers been able to attempt to understand what's really happening inside a supernova via computer modeling. These simulations, even when done in only one or two dimensions, can lead to new information and help address longstanding problems in astrophysics.

Now researchers have found a new way to make three-dimensional computer simulations of "core collapse" supernova explosions (one of two major types). Writing in the Sept. 1, 2010 issue of the Astrophysical Journal, Jason Nordhaus and Adam Burrows of Princeton University, and Ann Almgren and John Bell of Lawrence Berkeley National Laboratory report that the new simulations are beginning to match the massive blowouts astronomers have witnessed when gigantic stars die.[7]

The new simulations are based on the idea that the collapsing star is not spherical but distinctly asymmetrical and affected by a host of instabilities in the volatile mix surrounding its core (Figure 18). The simulations were made possible by bigger and faster supercomputers and a new code that can take advantage of their power.

The researchers performed these 3D simulations with approximately 4 million computer processor hours on 8,000 to 16,000 cores of NERSC's Cray XT4 Franklin system, using a new multidimensional radiation/hydrodynamic code called CASTRO, the development of which was led by Almgren and Bell of Berkeley Lab's Center for Computational Sciences and Engineering.

In the past, simulated explosions represented in one and two dimensions often stalled, leading scientists to conclude that their understanding of the physics was incorrect or incomplete. Using the new CASTRO code on supercomputers many times more powerful than their predecessors, this team was able to create

[7] J. Nordhaus, A. Burrows, A. Almgren, and J. Bell, "Dimension as a key to the neutrino mechanism of core-collapse supernova explosions," Astrophysical Journal **720**, 694 (2010).

supernova simulations in 3D that were ~50% easier to explode than in 1D, and exploded significantly faster than in 2D. This dimensional effect is much larger than that of a variety of other physical effects such as nuclear burning, inelastic scattering, and general relativity.

"It may well prove to be the case that the fundamental impediment to progress in supernova theory over the last few decades has not been lack of physical detail, but lack of access to codes and computers with which to properly simulate the collapse phenomenon in 3D," the team wrote. "This could explain the agonizingly slow march since the 1960s toward demonstrating a robust mechanism of explosion."

Burrows, a professor of astrophysical sciences at Princeton, commented: "I think this is a big jump in our understanding of how these things can explode. In principle, if you could go inside the supernovae to their centers, this is what you might see."

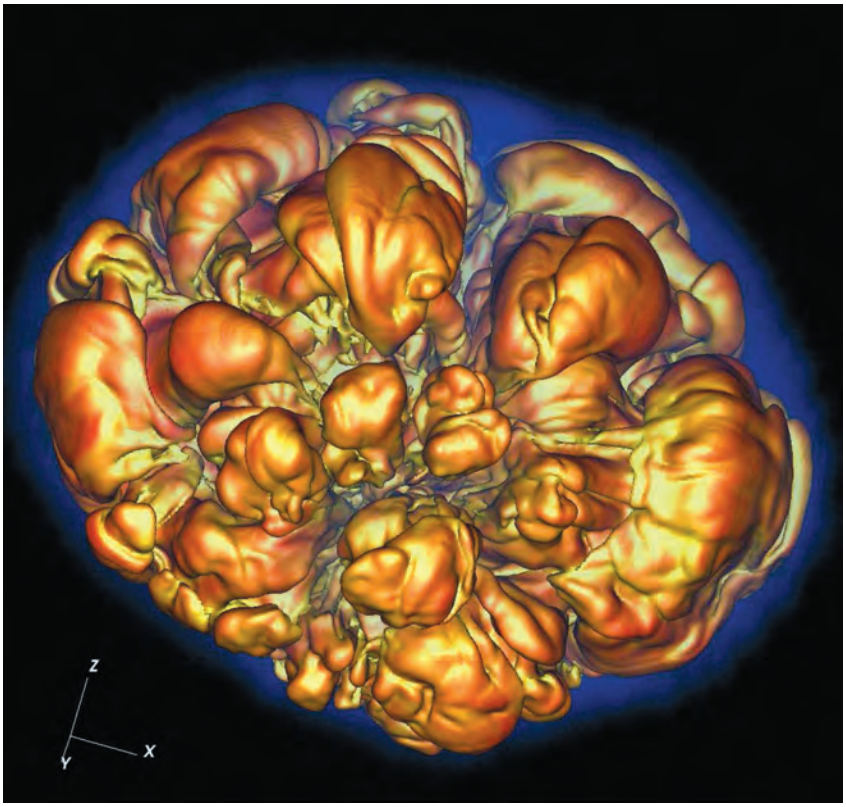

Figure 18. Isoentropy surface of the explosion of a massive star. The snapshot is taken ~50 milliseconds after the explosion. The multi-dimensional turbulent character of the blast is manifest. *Image: J. Nordhaus, A. Burrows, A. Almgren, and J. Bell.*

Birth of a Supernova

Supernovae are the primary source of heavy elements in the cosmos. Most result from the death of single stars much more massive than the sun.

As a star ages, it exhausts the hydrogen and helium fuel at its core. With still enough mass and pressure to fuse carbon and produce other heavier elements, it gradually becomes layered like an onion with the bulkiest tiers at its center. Once its core exceeds a certain mass, it begins to implode. In the squeeze, the core heats up and grows even denser.

"Imagine taking something as massive as the sun, then compacting it to something the size of the Earth," Burrows said. "Then imagine that collapsing to something the size of Princeton."

What comes next is even more mysterious.

At some point, the implosion reverses. Astrophysicists call it "the bounce." The core material stiffens up, acting like what Burrows calls a "spherical piston," emitting a shock wave of energy. Neutrinos, which are inert particles, are emitted too. The shock wave and the neutrinos are invisible.

Then there is a massive and very visible explosion, and the star's outer layers are ejected into space (Figure 19). This stage is what observers see as the supernova. What's left behind is an ultra-dense object called a neutron star. Sometimes, when an ultramassive star dies, a black hole is created instead.

Scientists have a sense of the steps leading to the explosion, but there is no consensus about the mechanism of the bounce. Part of the difficulty is that no one can see what is happening on the inside of a star. During this phase, the star looks undisturbed. Then, suddenly, a blast wave erupts on the surface. Scientists don't know what occurs to make the central region of the star instantly unstable. The emission of neutrinos is believed to be related, but no one is sure how or why.

"We don't know what the mechanism of explosion is," Burrows said. "As a theorist who wants to get to root causes, this is a natural problem to explore."

Multiple Scientific Approaches

The scientific visualization employed by the research team is an interdisciplinary effort combining astrophysics, applied mathematics, and computer science. The endeavor produces a representation, through computer-generated images, of three-dimensional phenomena. In general, researchers employ visualization techniques with the aim of making realistic renderings of quantitative information including surfaces, volumes, and light sources. Time is often an important component too, allowing researchers to create movies showing a simulated process in motion.

To do their work, Burrows and his colleagues came up with mathematical values representing the energetic behaviors of stars by using mathematical representations of fluids in motion—the same partial differential equations solved by geophysicists for climate modeling and weather forecasting. To solve these complex equations and simulate what happens inside a dying star, the team used the CASTRO code, which took into account factors that changed over time, including fluid density, temperature, pressure, gravitational acceleration, and velocity.

The calculations took months to process on supercomputers at NERSC, Princeton, and other centers.

The simulations are not ends unto themselves, Burrows noted. Part of the learning process is viewing the simulations and connecting them to real observations. In this case, the most recent simulations are uncannily similar to the explosive behavior of stars in their death throes witnessed by scientists. In addition, scientists often learn from simulations and see behaviors they had not expected.

"Visualization is crucial," Burrows said. "Otherwise, all you have is merely a jumble of numbers. Visualization via stills and movies conjures the entire phenomenon and brings home what has happened. It also allows one to diagnose the dynamics, so that the event is not only visualized, but understood."

To help them visualize these results, the team relied on Hank Childs of the NERSC Analytics Team. Nordhaus noted that Childs played an important role in helping the team create 3D renderings of these supernovae.

"The Franklin supercomputer at NERSC gives us the most bang for the buck," said Nordhaus. "This system not only has enough cores to run our 3D simulations, but our code scales extremely well on this platform. Having the code developers, John Bell and Ann Almgren, nearby to help us scale our simulations on the supercomputer is an added bonus."

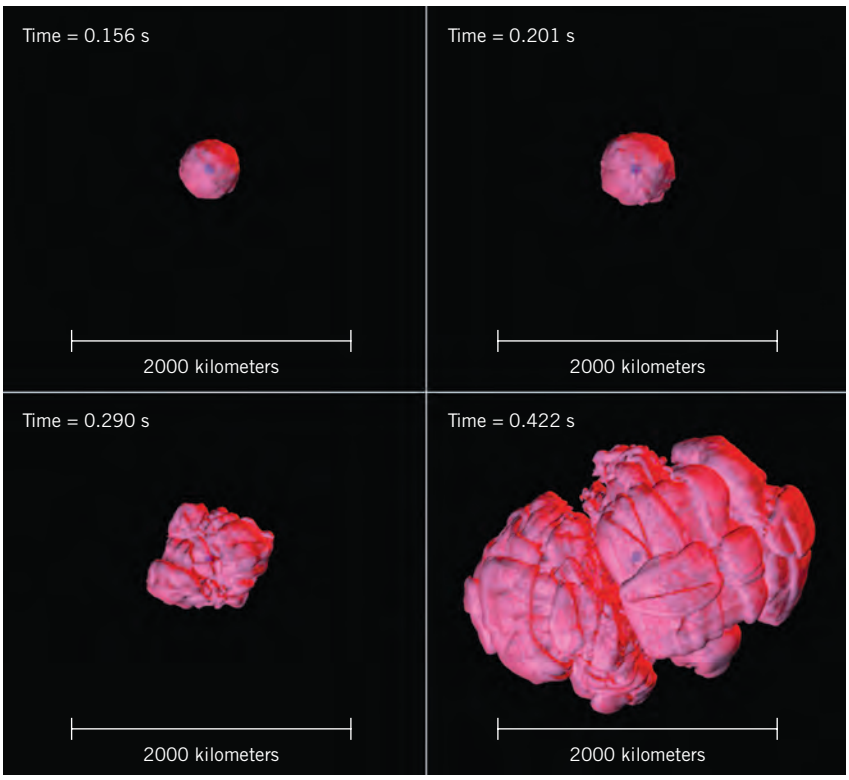

Figure 19. Evolutionary sequence of blast morphology for four different times after the bounce (0.156, 0.201, 0.289, and 0.422 seconds). The scale is more than 2000 km on a side. *Image: J. Nordhaus, A. Burrows, A. Almgren, and J. Bell.*

Surface Mineral

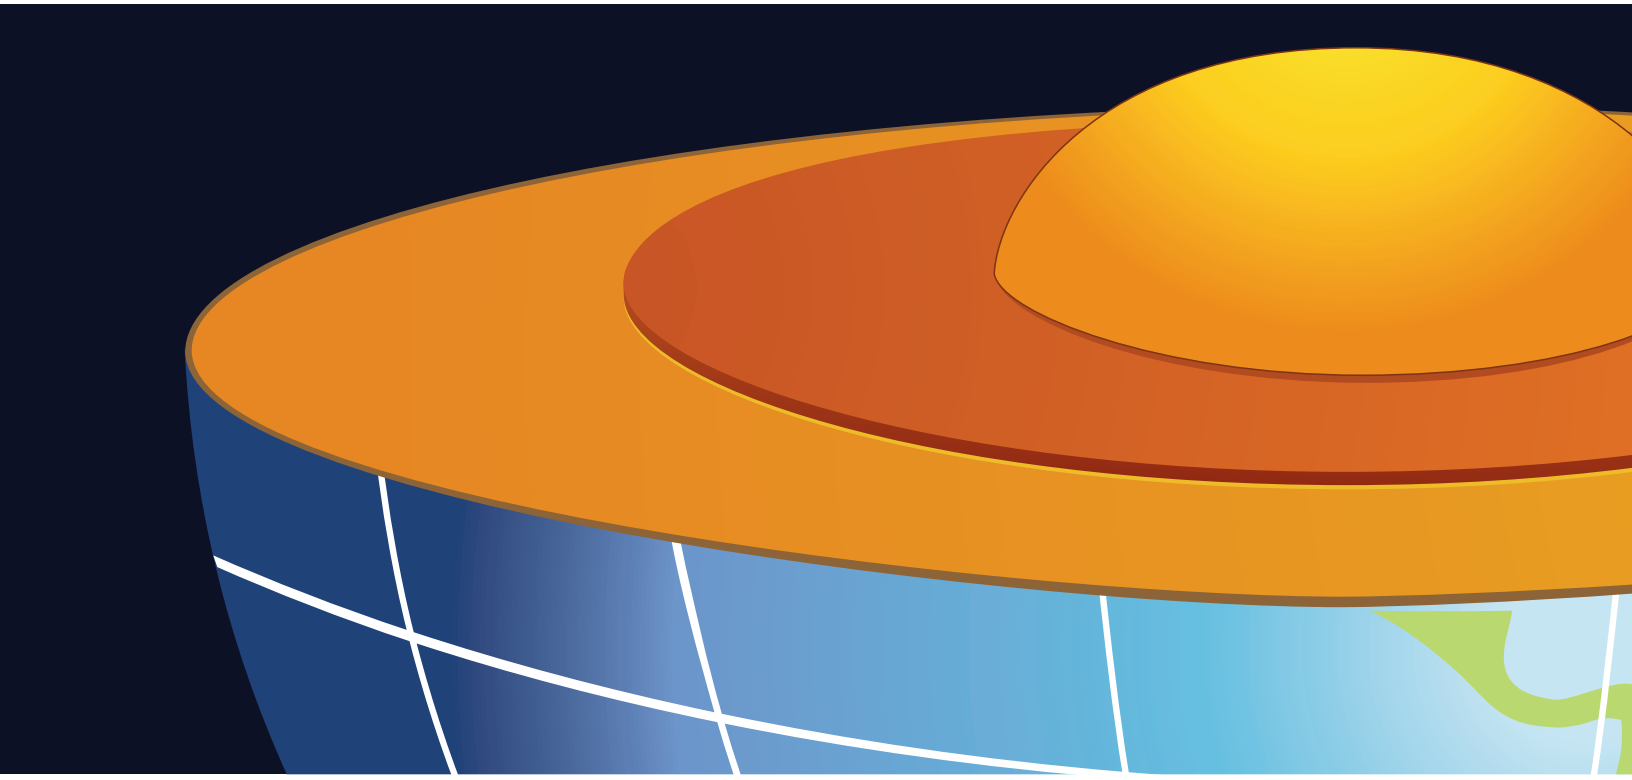

Innovative simulations reveal that Earth's silica is predominantly superficial

Silica is one of the most common minerals on Earth. Not only does it make up two-thirds of our planet's crust, it is also used to create a variety of materials from glass to ceramics, computer chips, and fiber optic cables. Yet new quantum mechanics results generated by a team of physicists from Ohio State University (OSU) and elsewhere show that this mineral only populates our planet superficially—in other words, silica is relatively uncommon deep within the Earth.

Using several of the largest supercomputers in the nation, including NERSC's Cray XT4 Franklin system, the team simulated the behavior of silica in high-temperature, high-pressure environments that are particularly difficult to study in a lab. These details may one day help scientists predict complex geological processes like earthquakes and volcanic eruptions. Their results were published in the May 25, 2010 edition of the Proceedings of the National Academy of Sciences (PNAS).[8]

"Silica is one of the simplest and most common minerals, but we still don't understand everything about it. A better understanding of silica on a quantum-mechanical level would be useful to earth science, and potentially to industry as well," says Kevin Driver, an OSU graduate student who was corresponding author of the paper. "Silica adopts many different structural forms at different temperatures and pressures, not all of which are easy to study in the lab."

Over the past century, seismology and high-pressure laboratory experiments have revealed a great deal about the general structure and composition of the Earth. For example, such work has shown that the planet's interior structure exists in three layers called the crust, mantle, and core (Figure 20). The outer two layers —the mantle and the crust—are largely made up of silicates, minerals containing silicon and oxygen. Still, the detailed structure and composition of the deepest parts of the mantle remain unclear. These details are important for complex geodynamical modeling that may one day predict large-scale events, such as earthquakes and volcanic eruptions.

Project: Modeling Dynamically and Spatially Complex Materials

PI: John Wilkins, Ohio State University

Funding: BES, NSF

Computing Resources: NERSC, NCAR, TeraGrid, NCSA, OSC, CCNI

[8] K. P. Driver, R. E. Cohen, Zhigang Wu, B. Militzer, P. López Ríos, M. D. Towler, R. J. Needs, and J. W. Wilkins, "Quantum Monte Carlo computations of phase stability, equations of state, and elasticity of high-pressure silica," PNAS **107**, 9519 (2010).

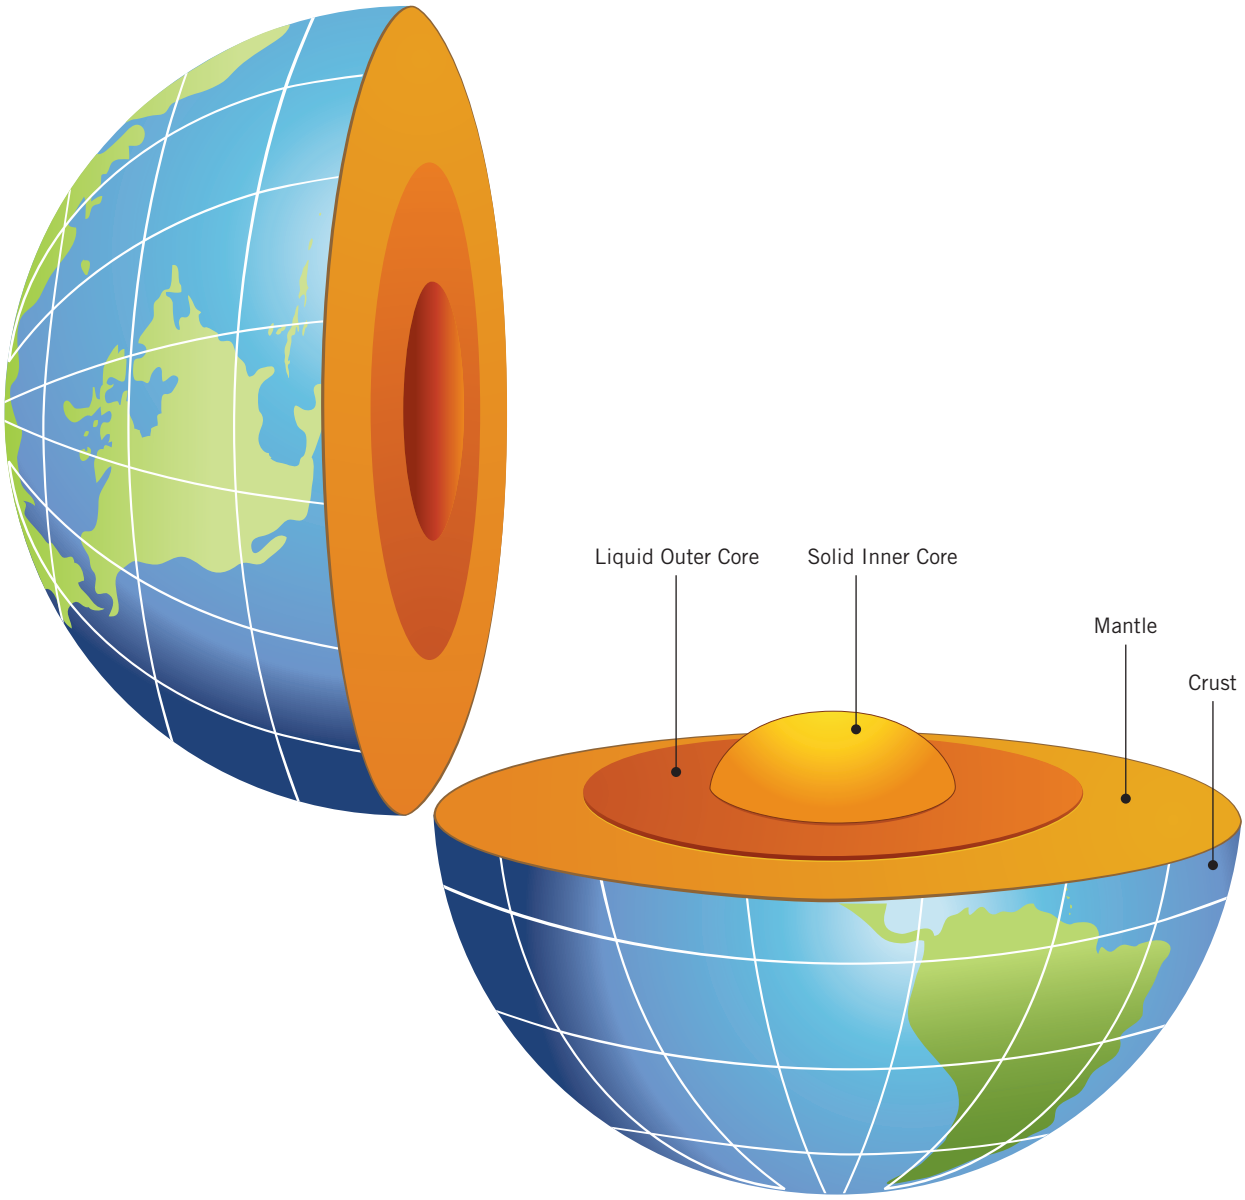

Figure 20. Cross-section of the Earth.

Driver notes that even the role that the simplest silicate—silica—plays in the Earth's mantle is not well understood. "Say you're standing on a beach, looking out over the ocean. The sand under your feet is made of quartz, a form of silica containing one silicon atom surrounded by four oxygen atoms. But in millions of years, as the oceanic plate below becomes subducted and sinks beneath the Earth's crust, the structure of the silica changes dramatically," he says.

As pressure increases with depth, the silica molecules crowd closer together, and the silicon atoms start coming into contact with more oxygen atoms from neighboring molecules. Several structural transitions occur, with low-pressure forms surrounded by four oxygen atoms and higher-pressure forms surrounded by six (Figure 21). With even more pressure, the structure collapses into a very dense form of the mineral, which scientists call the alpha-lead dioxide form.

Driver notes that it is this form of silica that likely resides deep within the earth, in the lower part

of the mantle, just above the planet's core. When scientists try to interpret seismic signals from that depth, they have no direct way of knowing what form of silica they are dealing with. They must rely on high-pressure experiments and computer simulations to constrain the possibilities. Driver and colleagues use a particularly high-accuracy, quantum mechanical simulation method to study the thermodynamic and elastic properties of different silica forms, and then compare the results to the seismic data.

In PNAS, Driver, his advisor John Wilkins, and their coauthors describe how they used a quantum mechanical method to design computer algorithms that would simulate the silica structures. When they did, they found that the behavior of the dense, alpha-lead dioxide form of silica did not match up with any global seismic signal detected in the lower mantle. This result indicates that the lower mantle is relatively devoid of silica, except perhaps in localized areas where oceanic plates have subducted.

"As you might imagine, experiments performed at pressures near those of Earth's core can be very challenging. By using highly accurate quantum mechanical simulations, we can offer reliable insight that goes beyond the scope of the laboratory," says Driver.

A New Application of Quantum Monte Carlo

The team's work was one of the first to show that quantum Monte Carlo (QMC) methods could be used to study complex minerals. QMC does not rely on density functional calculations, in which approximations can produce inaccurate results in these kinds of studies; instead, QMC explicitly treats the electrons and their interactions via a stochastic solution of Schrödinger's equation. Although QMC algorithms have been used for over half a century to simulate atoms, molecules, gases, and simple solids, Driver notes that applying them to a complex solid like silica was simply too labor- and computer-intensive until recently.

"In total, we used the equivalent of six million CPU hours to model four different states of silica. Three million of those CPU hours involved using NERSC's Franklin system to calculate a shear elastic constant for silica with the QMC method. This is the first time it had ever been done," said Driver. The Franklin results allowed him to measure how silica deforms at different temperatures and pressures.

"From computing hardware to the consulting staff, the resources at NERSC are really excellent," says Driver. "The size and speed of the center's machines is something that I don't normally have access to at other places."

Wilkins comments, "This work demonstrates both the superb contributions a single graduate student can make, and that the quantum Monte Carlo method can compute nearly every property of a mineral over a wide range of pressures and temperatures. The study will stimulate a broader use of quantum Monte Carlo worldwide to address vital problems."

"This work demonstrates … that the quantum Monte Carlo method can compute nearly every property of a mineral over a wide range of pressures and temperatures."

Wilkins and his colleagues expect that quantum Monte Carlo will be used more often in materials science in the future, as the next generation of computers goes online.

Coauthors on the paper included Ronald Cohen of the Carnegie Institution of Washington; Zhigang Wu of the Colorado School of Mines; Burkhard Militzer of the University of California, Berkeley; and Pablo López Ríos, Michael Towler, and Richard Needs of the University of Cambridge.

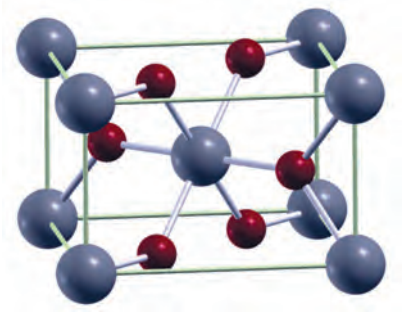

Figure 21. The structure of a six-fold coordinated silica phase called stishovite, a prototype for many more complex mantle minerals. Stishovite is commonly created in diamond anvil cells and often found naturally where meteorites have slammed into earth. Driver and co-authors computed the shear elastic constant softening of stishovite under pressure with quantum Monte Carlo using over 3 million CPU hours on Franklin. *Image: Kevin Driver.*

Soggy Origami

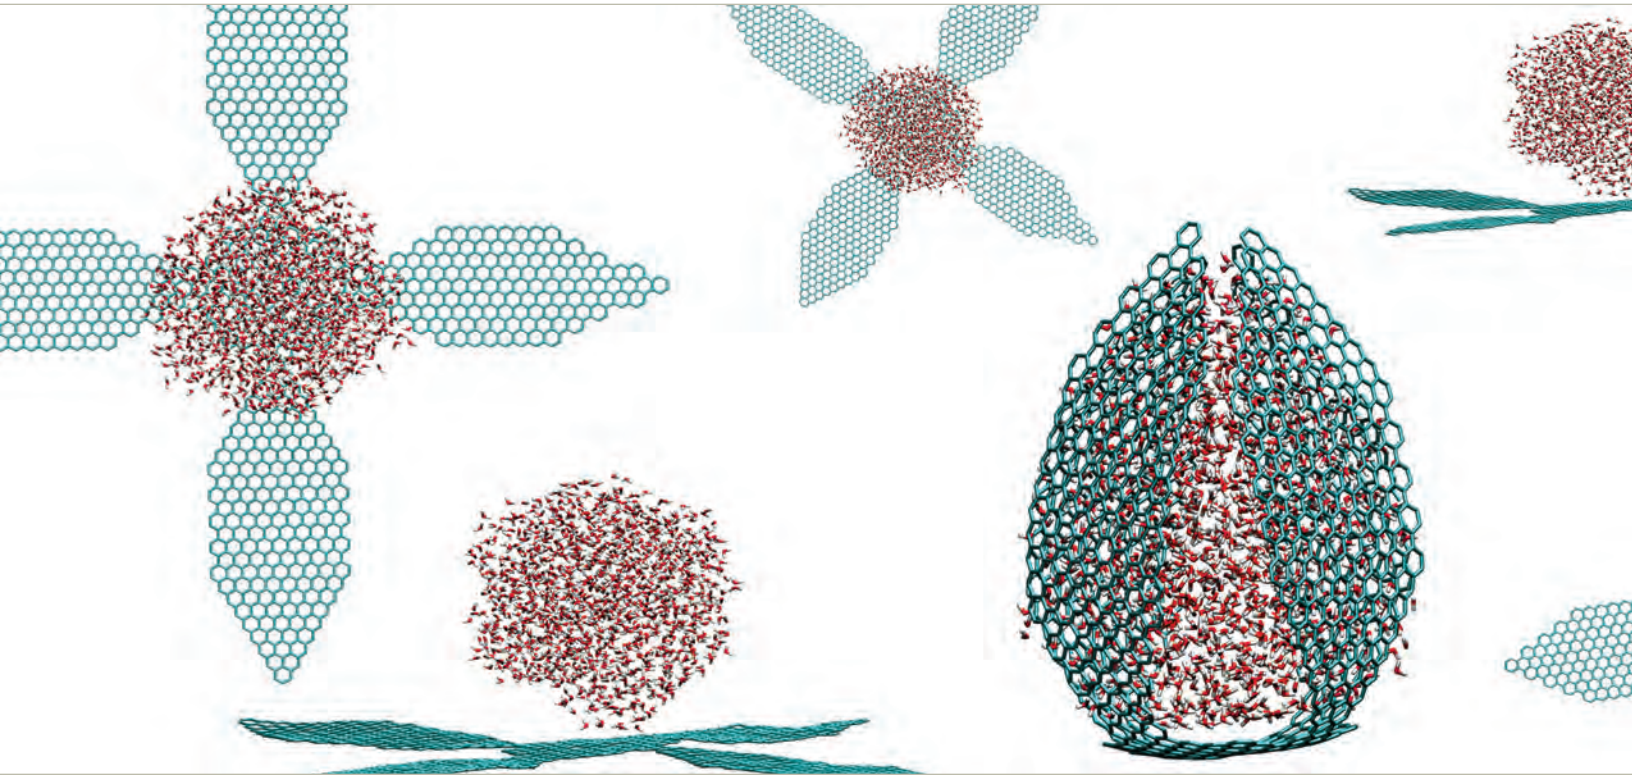

Nanotech research on graphene is blossoming at NERSC

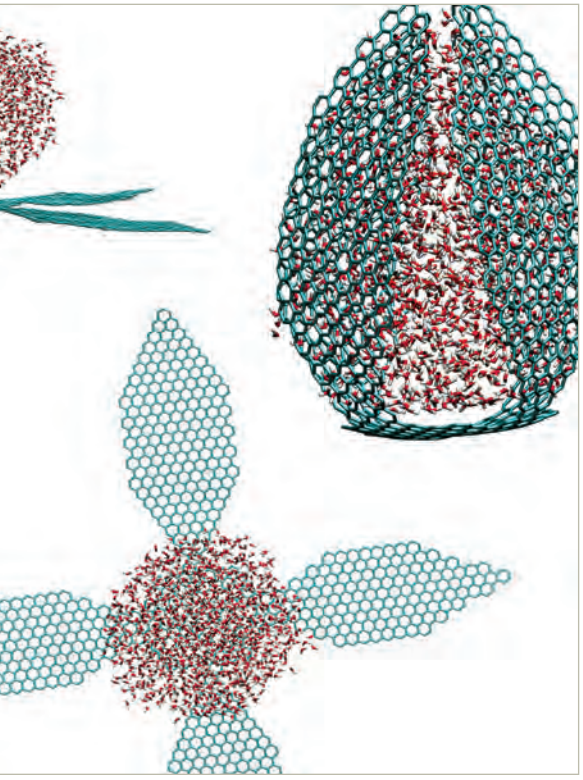

Project: Multiscale Modeling of Nanofluidic Systems

PI: Petr Král, University of Illinois at Chicago

Funding: BES, NSF

Computing Resources: NERSC

In a 1959 lecture titled "There's Plenty of Room at the Bottom," physicist and future Nobel laureate Richard Feynman predicted that humans would create increasingly smaller and more powerful machines by tapping into the quantum physics that drives atoms and molecules. Nearly 50 years later this prediction is realized in the field of materials nanoscience, where some researchers are designing materials atom by atom to build devices with unprecedented capabilities.

In fact, assistant professor Petr Král and his team of graduate students from the University of Illinois at Chicago (UIC) are making headway in this arena by using supercomputers at NERSC to investigate the molecular and atomic properties of graphene—a single-atom layer of carbon, densely packed in a honeycomb pattern. Discovered in 2004, graphene is the most conductive material known to humankind and could be used to make the next generation of transistors, computer memory chips, and high-capacity batteries.

Manipulating the "Miracle Material" with Water

Electrons blaze through graphene faster than in any other conductor, making it ideal for use in a variety of high-tech devices. With 200 times the breaking strength of steel, graphene is also stronger and stiffer than a diamond. The application potential for this material is so promising that the two scientists who discovered it, Andre Geim and Konstantin Novoselov, won the 2010 Nobel Prize in Physics.

But before this "miracle material" can be widely implemented in new technologies, researchers must figure out how to control its shape. That's exactly what Král's research team did. In an article published in Nano Letters[9] and highlighted in Nature's "News and Views" section,[10] Král and his former UIC graduate students Niladri Patra and Boyang Wang reported that nanodroplets of water could be used to shape sheets of graphene. Using the scalable NAMD

[9] Niladri Patra, Boyang Wang, and Petr Král, "Nanodroplet activated and guided folding of graphene nanostructures," Nano Letters **9**, 3766 (2009).
[10] Vincent H. Crespi, "Soggy origami," Nature **462**, 858 (2009).

code to study molecular dynamics and SIESTA to calculate electron structure on NERSC's Franklin system, they found that nanodroplets of water remove the energy barrier that prevents graphene from folding.

Because atoms exposed on the surface of materials tend to be less stable than those buried within, placing a water droplet on graphene —which is just a single layer of carbon atoms—will cause the material to deform so that the number of exposed atoms is minimized and the system's energy is reduced. Although spontaneous folding is not a characteristic unique to wet graphene, the computational models also showed that researchers could control how the material deforms. Being able to easily shape graphene is essential to its application in future technologies.

According to Král, a graphene sheet will naturally wrap around a water droplet. By positioning nanodroplets of water on graphene with guided microscopic needles, scientists can make the carbon sheet roll, bend, and slide; the material can also be manipulated into a variety of complex forms like capsules, sandwiches, knots, and rings, to serve as the foundation for nanodevices.

The end product depends on the graphene's initial shape and the water droplet's diameter. For example, placing a droplet of water in the center of the "petals" of a flower-shaped graphene ribbon causes the "flower" to close (Figure 22). A droplet placed on one end of a graphene ribbon causes

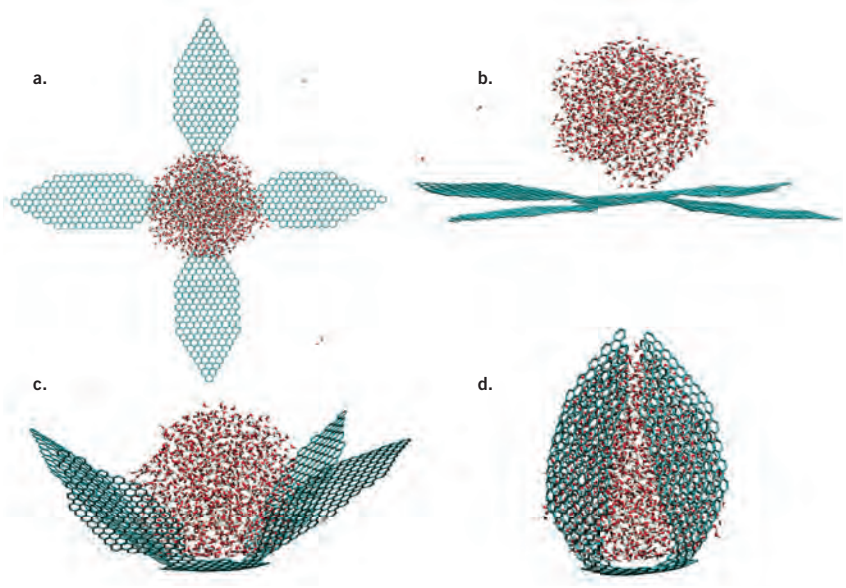

Figure 22. A droplet of water (pink) causes the "petals" of a flower-shaped graphene ribbon (green) about 13 nanometers across to fold up around it. *Image: N. Patra, B. Wang, and P. Král.*

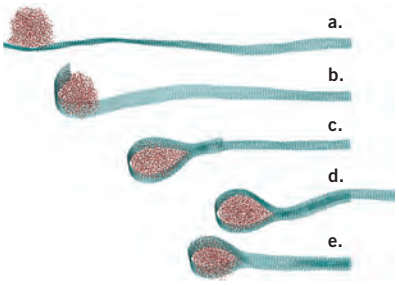

Figure 23. Placing a droplet of water on one end of a graphene ribbon causes the free end to wrap around it, thus folding the ribbon in half. *Image: N. Patra, B. Wang, and P. Král.*

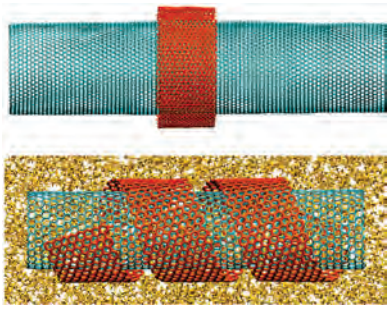

Figure 24. Placing a graphene nanoribbon on a nanotube causes it to fold on the nanotube surface, even in the presence of nonpolar solution. *Image: N. Patra, Y. Song and P. Král.*

Controlling graphene's shape and conductivity is essential to its application in future technologies.

[11] N. Patra, Y. Song, and P. Král, "Self-assembly of graphene nanostructures on nanotubes," ACS Nano (in press).

the free end to wrap around it, thus folding the ribbon in half (Figure 23). Another study, which will appear in a forthcoming issue of ACS Nano, shows that placing a graphene ribbon on a nanotube causes the ribbon to wrap itself around the tube, even in solvents (Figure 24).[11]

"Until these results it wasn't thought that we could controllably fold these structures," says Král. "Now, thanks to NERSC, we know how to shape graphene using weak forces between nanodroplets of water."

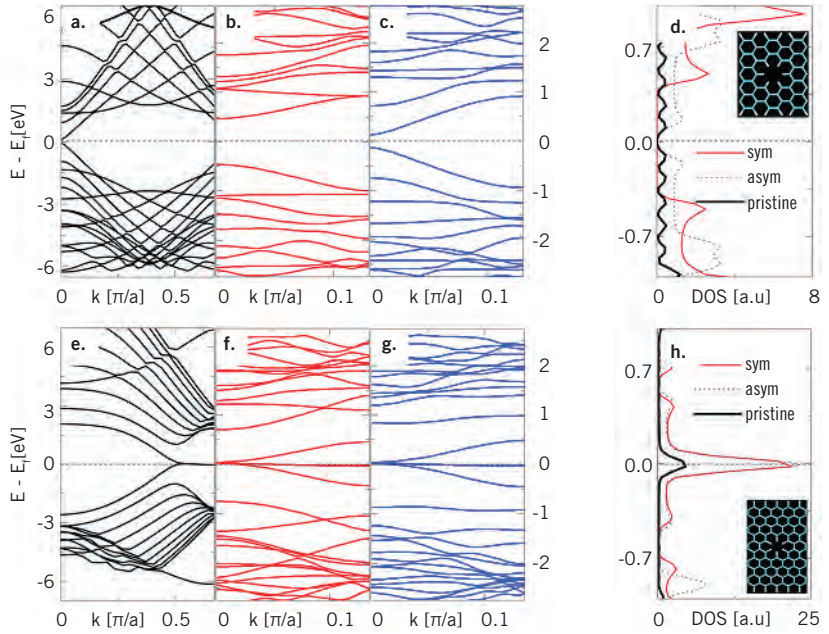

Figure 25. These graphs show the energies of electronic states as a function of their momentum (k). When the curves touch each other in the center, the graphene is a metal; otherwise it is a semiconductor. The graphs on the right side show densities of states. AGNR is armchair graphene nanoribbon and ZGNR is a zigzag. Band structure is shown for: (a) pristine 11-AGNR, (b) 11-AGNR with centered pore, (c) 11-AGNR with shifted pore, (e) pristine 10-ZGNR, (f) 10-ZGNR with centered pore, (g) 10-ZGNR with shifted pore. The notation 11- or 10- here means the width of the ribbons, namely, 11-AGNR and 10-ZGNR are nanoribbons of 5 honeycomb carbon rings width, as shown in the insets of (d) and (h). *Image: A. Baskin and P. Král.*

Manipulating the Conductivity of Graphene

In addition to shaping graphene, Král and another UIC graduate student, Artem Baskin, used large-scale density functional theory *ab initio* calculations at NERSC to confirm a belief long held by scientists—that they could manipulate the conductivity of graphene by strategically inserting holes into it.[12] The simulations revealed that the arrangement of these pores and their size could be used to control graphene's conductivity; and changing the distance between these pores can cause the material to switch from metal to semiconducting porous graphene (Figure 25).

In this project they explored the atomic characteristics of porous graphene-based materials, like nanoribbons, superlattices, and nanotubes. Based on their results, Baskin and Král propose a phenomenological model that can be used to predict the band structure features of porous nanocarbons without performing the time-consuming quantum mechanical calculations. The duo successfully tested this model on a number of systems and recently submitted a paper with this recommendation to Physical Review Letters.

The team used NERSC resources for algorithm testing as well as modeling. "NERSC resources have been indispensible for us," says Král. "Our projects require long-multiprocessor calculations, and access to the center's computation time was very important for progressing our research in a timely manner."

[12] A. Baskin and P. Král, "Electronic structures of porous nanocarbons," Physical Review Letters (submitted).

An Antimatter Hypernucleus

Discovery opens the door to new dimensions of antimatter

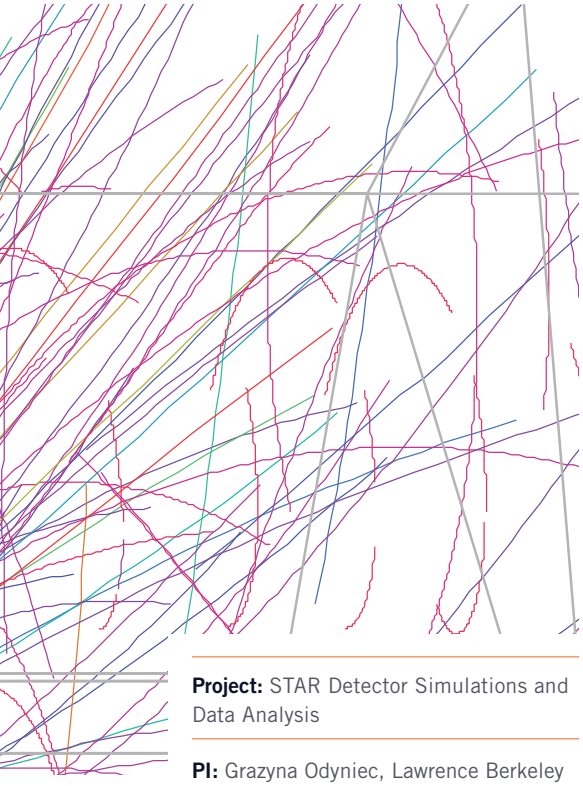

Project: STAR Detector Simulations and Data Analysis

PI: Grazyna Odyniec, Lawrence Berkeley National Laboratory

Senior Investigators: Nu Xu, LBNL; James Dunlop, Brookhaven National Laboratory; Olga Barannikova, University of Illinios, Chicago; Bernd Surrow, Massachusetts Institute of Technology

Funding: NP, HEP, NSF, Sloan Foundation, CNRS/IN2P3, STFC, EPSRC, FAPESP–CNPq, MESRF, NNSFC, CAS, MoST, MoE, GA, MSMT, FOM, NWO, DAE, DST, CSIR, PMSHE, KRF, MSES, RMST, ROSATOM

Computing Resources: NERSC, RCF, OSG

The strangest form of antimatter ever seen has been discovered by the STAR experiment at the Relativistic Heavy Ion Collider (RHIC) at DOE's Brookhaven National Laboratory in New York.[13] The new particle goes by the jawbreaking name "anti-hyper-triton."

Translated, that means a nucleus of antihydrogen containing one antiproton and one antineutron—plus one heavy relative of the antineutron, an antilambda hyperon. It's that massive antilambda that makes the newly discovered antinucleus "hyper."

"STAR is the only experiment that could have found an antimatter hypernucleus," says Nu Xu of Berkeley Lab's Nuclear Science Division, the spokesperson for the STAR experiment. "We've been looking for them ever since RHIC began operations. The discovery opens the door on new dimensions of antimatter, which will help astrophysicists trace back the story of matter to the very first millionths of a second after the big bang."

RHIC was designed to collide heavy nuclei like gold together at very high energies. The fireballs that blossom in the collisions are so hot and dense that the protons and neutrons in the nuclei are shattered and there is massive production of partons (Figure 26).

For a brief instant there exists a quark-gluon plasma that models the first moments of the early universe, in which quarks—and gluons, which strongly bind quarks together at lower temperatures—are suddenly free to move independently. But as the quark-gluon plasma cools, its constituents recombine in a variety of ways.

The commonest quarks, which are also the lightest, are the ups and downs; they make up the protons and neutrons in the nuclei of all the familiar ordinary atoms around us. Heavier particles also condense out of the quark-gluon plasma, including the lambda, made of an up quark, a down quark, and a massive strange quark.

[13] The STAR Collaboration, "Observation of an antimatter hypernucleus," Science **238**, 58 (2010).

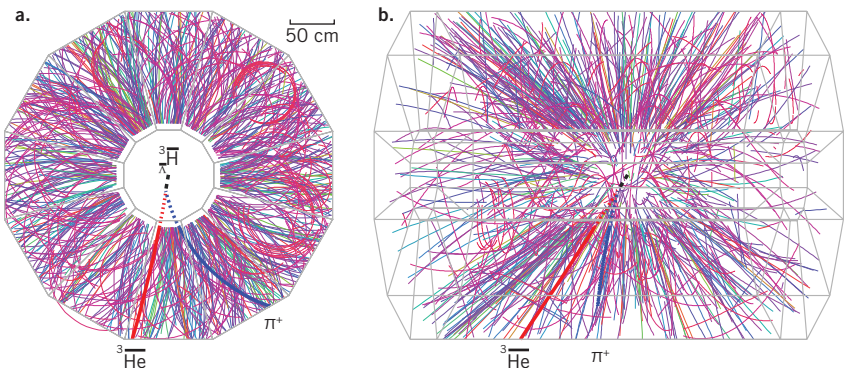

Figure 26. In a single collision of gold nuclei at RHIC, many hundreds of particles are emitted, most created from the quantum vacuum via the conversion of energy into mass in accordance with Einstein's famous equation $E = mc^2$. The particles leave telltale tracks in the STAR detector (shown here from the end and side). Scientists analyzed about a hundred million collisions to spot the new antinuclei, identified via their characteristic decay into a light isotope of antihelium and a positive pi-meson. Altogether, 70 examples of the new antinucleus were found. *Image: STAR Collaboration.*

The New Region of Strange Antimatter

The standard Periodic Table of Elements is arranged according to the number of protons, which determine each element's chemical properties. Physicists use a more complex, three-dimensional chart to also convey information on the number of neutrons, which may change in different isotopes of the same element, and a quantum number known as "strangeness," which depends on the presence of strange quarks (Figure 27). Nuclei containing one or more strange quarks are called hypernuclei.

For all ordinary matter, with no strange quarks, the strangeness value is zero and the chart is flat. Hypernuclei appear above the plane of the chart. The new discovery of strange antimatter with an antistrange quark (an antihyper-nucleus) marks the first entry below the plane.

"Ordinary" hypernuclei, in which a neutron or two are replaced by lambdas or other hyperons, have been made in particle accelerators since the 1950s—not long after the first lambda particles were identified in cosmic ray debris in 1947.

Quarks and antiquarks exist in equal numbers in the quark-gluon plasma, so the particles that condense from a cooling quark-gluon plasma come not only in ordinary-matter forms, but in their oppositely charged antimatter forms as well.

Antiprotons and antineutrons were discovered at Berkeley Lab's Bevatron in the 1950s. A few antihydrogen atoms with a single antiproton for a nucleus have even been made, but no heavier forms of antihydrogen. A hydrogen nucleus with one proton and two neutrons is called a triton. The STAR discovery is the first time anyone has observed an antihypertriton.

What made the discovery possible was the central element in the STAR experiment, its large time projection chamber (TPC). The TPC concept was invented at Berkeley Lab by David Nygren; STAR's TPC was built there and shipped to Brookhaven aboard a giant C-5A cargo plane.

"What a time projection chamber does is allow three-dimensional imaging of the tracks of every particle and its decay products in a single gold-gold collision," Nu Xu explains. "The tracks are curved by a uniform magnetic field, which allows their original mass and energy to be determined by following them back to their origin. We can reconstruct the 'topology of decay' of a chosen event and identify every particle that it produces."

So far the STAR collaboration has identified about 70 of the strange new antihypertritons, along with about 160 ordinary hypertritons, from 100 million collisions. As more data is collected, the STAR researchers expect to find even heavier hypernuclei and antimatter hypernuclei, including alpha particles, which are the nuclei of helium.

Data Simulation, Storage, and Analysis

To identify this antihypertriton, physicists used supercomputers at NERSC and other research centers to painstakingly sift through the debris of some 100 million collisions. The team also used NERSC's PDSF system to simulate detector response. These results allowed them to see that all of the charged particles within the collision debris left their mark by ionizing the gas inside RHIC's time projection chamber, while the antihypertritons revealed themselves through a unique decay signature—the two tracks left by a charged pion and an antihelium-3 nucleus, the latter

being heavy and so losing energy rapidly with distance in the gas.

"These simulations were vital to helping us optimize search conditions such as topology of the decay configuration," says Zhangbu Xu, a physicist at Brookhaven who is one of the STAR collaborators. "By embedding imaginary anti-matter in a real collision and optimizing the simulations for the best selection conditions, we were able to find a majority of those embedded particles."

As a STAR Tier 1 site, the PDSF stores a portion of the raw data from the experiment and all of the analysis-ready data files. The STAR databases at Brookhaven are mirrored to PDSF to provide direct access for PDSF users. STAR currently maintains around 800 TB of data in NERSC's HPSS archive.

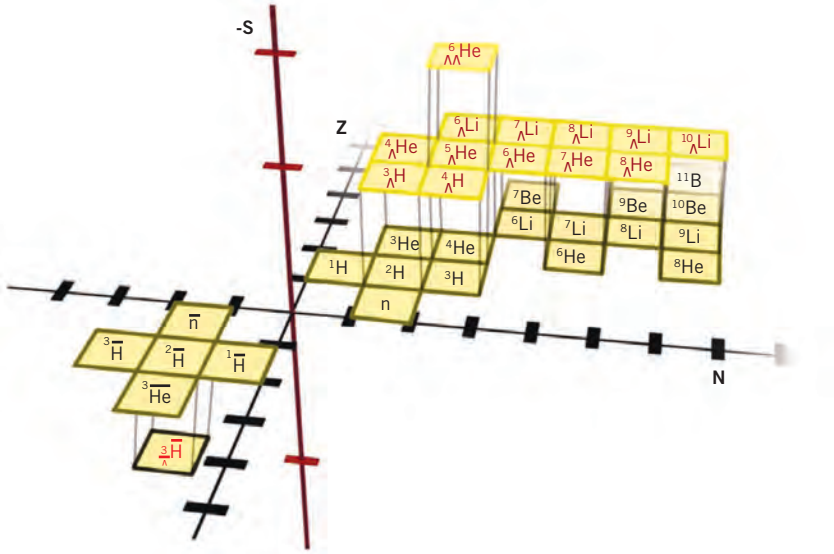

Figure 27. The diagram above is known as the 3D chart of the nuclides. The familiar Periodic Table arranges the elements according to their atomic number, *Z*, which determines the chemical properties of each element. Physicists are also concerned with the *N* axis, which gives the number of neutrons in the nucleus. The third axis represents strangeness, *S*, which is zero for all naturally occurring matter, but could be non-zero in the core of collapsed stars. Antinuclei lie at negative *Z* and *N* in the above chart, and the newly discovered antinucleus (red) now extends the 3D chart into the new region of strange antimatter. *Image: STAR Collaboration.*

Great Expectations

The STAR detector will soon be upgraded, and as the prospects for determining precise ratios for the production of various species of strange hypernuclei and their antimatter counterparts increase, so does the expectation that STAR will be able to shed new light on fundamental questions of physics, astrophysics, and cosmology, including the composition of collapsed stars and the ratio of matter and antimatter in the early universe.

"The strangeness value could be nonzero in the core of collapsed stars," says Jinhui Chen of the STAR collaboration, a postdoctoral researcher at Kent State University and staff scientist at the Shanghai Institute of Applied Physics, "so the present measurements at RHIC will help us distinguish between models that describe these exotic states of matter."

"Understanding precisely how and why there's a predominance of matter over antimatter remains a major unsolved problem of physics," says Zhangbu Xu. "A solution will require measurements of subtle deviations from perfect symmetry between matter and antimatter, and there are good prospects for future antimatter measurements at RHIC to address this key issue."

Why there is more matter than antimatter in the universe is one of the major unsolved problems in science. Solving it could tell us why human beings, indeed why anything at all, exists today.

Subtle deviations from perfect symmetry will help explain why there is more matter than antimatter.

Dynameomics

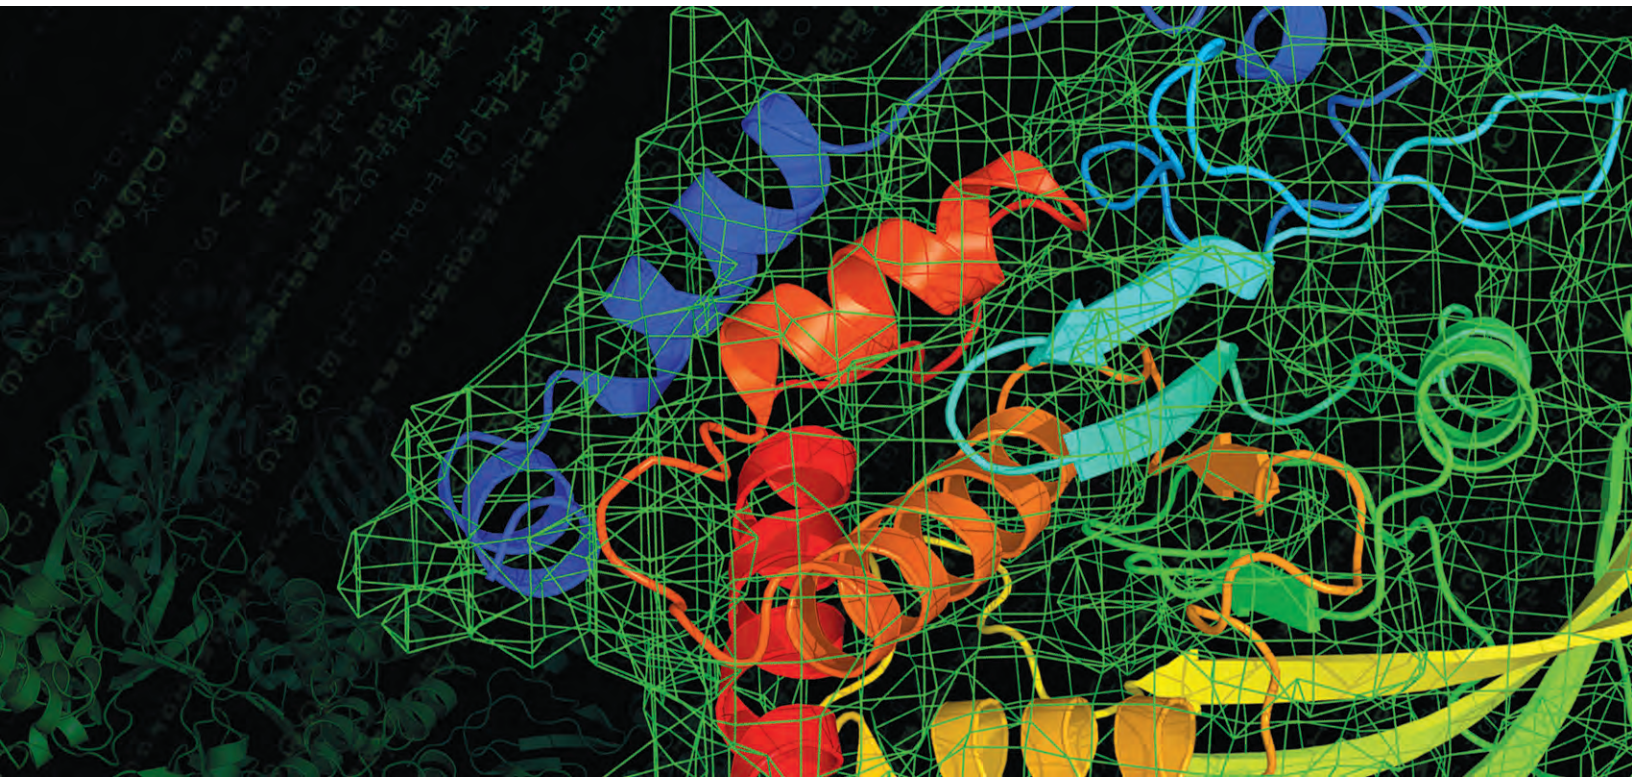

A new database helps scientists discover how proteins do the work of life

All organisms from the largest mammals to the smallest microbes are composed of proteins—long, non-branching chains of amino acids. In the human body, proteins play both structural and functional roles, making up muscle tissue and organs, facilitating bone growth, and regulating hormones, in addition to various other tasks.

Scientists know that the key to a protein's biological function depends on the unique, stable, and precisely ordered three-dimensional shape that it assumes just microseconds after the amino acid chain forms. Biologists call this the "native state," and suspect that incorrectly folded proteins are responsible for illnesses like Alzheimer's disease, cystic fibrosis, and numerous cancers. Over the years, NMR spectroscopy and X-ray crystallography experiments have provided a lot of information about the native state structure of most protein folds, but the scientific rules that determine their functions and cause them to fold remain a mystery. Researchers expect that a better understanding of these rules will pave the way for breakthroughs in medical applications, drug design, and biophysics.

This is where Valerie Daggett, professor of Bioengineering at the University of Washington, Seattle, believes supercomputers can help. Over the last six years, Daggett used a total of 12.4 million processor hours at NERSC to perform molecular dynameomics simulations. This technique combines molecular dynamics (the study of molecular and atomic motions) with proteomics (the study of protein structure and function) to characterize the folding and unfolding pathways of proteins. The culmination of this work has been incorporated into a public database, http://www.dynameomics.org/.

"We already have several examples of dynamic behavior linked to function that was discovered through molecular dynamic simulations and not at all evident in the static average experimental structures," Daggett says in a June 2010 Nature Methods article.[14]

Project: Molecular Dynameomics

PI: Valerie Daggett, University of Washington, Seattle

Funding: BER, NIH, Microsoft

Computing Resources: NERSC

[14] Allison Doerr, "A database of dynamics," Nature Methods **7**, 426 (2010).

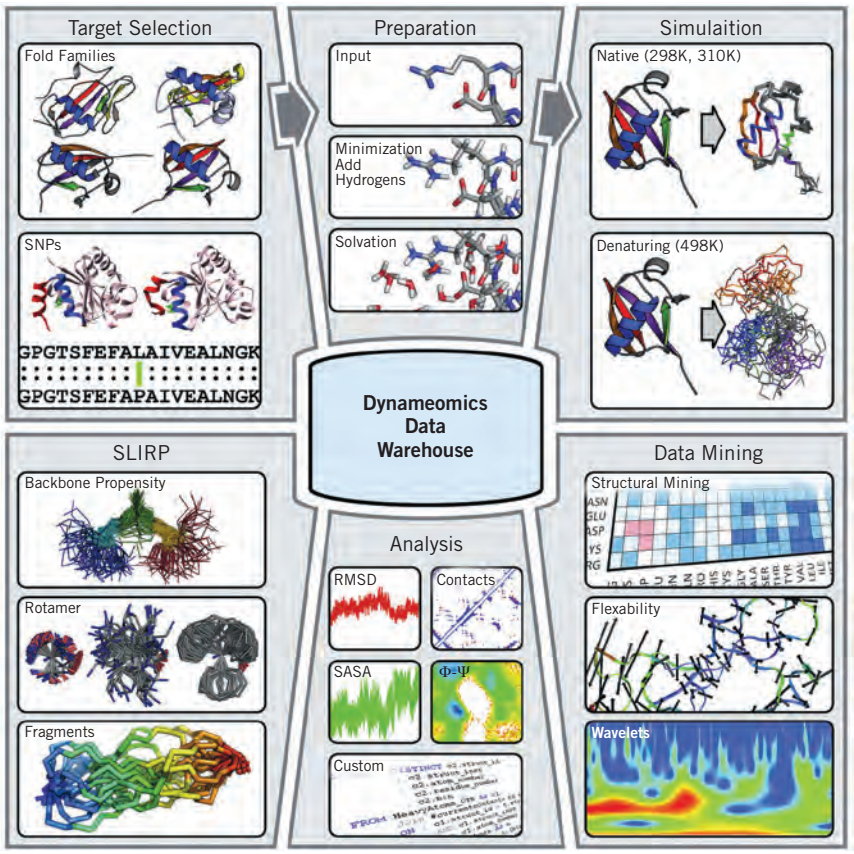

Figure 28. Overview of the generation and use of molecular dynameomics simulation data, organized through the Dynameomics Data Warehouse. The repository include a Structural Library of Intrinsic Residue Propensities (SLIRP). *Image: Daggett Group.*

Unraveling Proteins with Computer Simulations

Although the Protein Data Bank (PDB) of experimentally derived protein structures has been vital to many important scientific discoveries, these data do not tell the whole story of protein function. These structures are simply static snapshots of proteins in their native, or folded, state. In reality, proteins are constantly moving, and these motions are critical to their function. This is why Daggett's team wanted to create a complementary database of molecular dynamics structures for representatives of all protein folds, including their unfolding pathways. The Dynameomics Data Warehouse provides an organizing framework, a repository, and a variety of access interfaces for simulation and analysis data (Figure 28).

"With the simulation data stored in an easily queryable structured repository that can be linked to other sources of biological and experimental data, current scientific questions can be addressed in ways that were previously impossible or extremely cumbersome," writes Marc van der Kamp, a postdoctoral researcher in Daggett's group, along with a dozen co-authors in a paper featured on the cover of the April 14, 2010 issue of Structure (Figure 29).[2]

A protein will easily unfold, or denature, when enough heat is applied. This unfolding process follows the same transitions and intermediate structures as folding, except in reverse. To simulate this process computationally, Daggett's team uses the "in lucem Molecular Dynamics" (ilmm) code, which was developed by David Beck, Darwin Alonso, and Daggett, to solve Newton's equations of motion for every atom in the protein molecule and surrounding solvent. These simulations map the conformational changes of a protein with atomic-level detail, which can then be tested with experiments.

To determine the trajectories of protein unfolding, Daggett's team simulates each protein at least six times—once to explore the molecular dynamics of a native state for at least 51 nanoseconds at a temperature of 298 K (77° F), and five more times to study the dynamics of protein at 498 K (437° F), with two simulations of at least 51 ns and three short simulations of at least 2 ns each.

Daggett's group applies a careful strategy to select diverse protein folds to analyze. Using three existing protein fold classification systems—SCOP, CATH, and Dali—they identify domains for proteins available in the PDB and cluster them into "metafolds."[3] The team has simulated at least one member from

[2] M. W. van der Kamp, R. D. Schaeffer, A. L. Jonsson, A. D. Scouras, A. M. Simms, A. D. Toofanny, N. C. Benson, P. C. Anderson, E. D. Merkley, S. Rysavy, D. Bromley, D. A. C. Beck, and V. Daggett, "Dynameomics: A comprehensive database of protein dynamics," Structure **18**, 423 (2010).
[3] R. Dustin Schaeffer, Amanda L. Jonsson, Andrew M. Simms, and Valerie Daggett, "Generation of a consensus protein domain dictionary," Bioinformatics **27**, 46 (2011).

each protein metafold. Although experimental dynamic information is not available for most of their targets, Daggett notes that the simulations do compare well where experimental data are available.

Daggett began running these simulations at NERSC in 2005, when she received a Department of Energy INCITE allocation of 2 million processor hours on the facility's IBM SP "Seaborg" system. In that first year, she modeled the folding pathways of 151 proteins in approximately 906 simulations. Subsequently her team received additional allocations from DOE's Office of Biological and Environmental Research, as well as NERSC discretionary time, to continue their successful efforts. As of 2010, the group has completed over 11,000 individual molecular dynamics simulations on more than 2,000 proteins from 218 different organisms. Their complete database of simulations contains over 340 microseconds of native and unfolding simulations, stored as $>10^8$ individual structures, which is four orders of magnitude larger than the PDB.

"This is a data-mining endeavor to identify similarities and differences between native and unfolded states of proteins across all secondary and tertiary structure types and sequences," says Daggett. "It is necessary to successfully predict native states of proteins, in order to translate the current deluge of genomic information into a form appropriate for better functional identification of proteins and for drug design."

Traditional drug design strategies tend to focus on native-state structure alone. But the Dynameomics Data Warehouse allows researchers to systematically search for transient conformations that would allow the binding of drugs or chemical chaperones. These molecules could then stabilize the protein structure without interfering with function. This kind of information is readily available through simulation but not through standard experimental techniques.

Figure 29. The plane of structures on the bottom shows sample fold-representatives that were simulated. One example of a simulated protein, Protein-tyrosine phosphatase 1B (PDB code 2HNP), is highlighted above. *Cover image: Marc W. van der Kamp, Steven Rysavy, and Dennis Bromley. ©2010, Elsevier Ltd.*

Some of the proteins in the Dynameomics Data Warehouse include single-nucleotide polymorphism-containing mutant human proteins implicated in disease. In some cases, the computer simulations provided valuable insights into why certain mutations caused structural disruptions at the active site. The database also includes proteins from thermophilic organisms, which thrive in scorching environments. By comparing the native state and unfolding dynamics of thermophiles to their mesophilic counterparts, which thrive in moderate temperatures, researchers hope to discover the mechanisms that make thermophiles so stable at high temperatures.

"There are approximately 1,695 known non-redundant folds, of which 807 are self-contained autonomous folds, and we have simulated representatives for all 807," says Daggett. "We couldn't have come this far without access to DOE resources at NERSC, and with their continued suport we will simulate more targets for well-populated metafolds to explore the sequence determinants of dynamics and folding."

Dynameomics Data Suggests New Treatment Strategy for Amyloid Diseases

One of the most exciting applications of Dynameomics involves amyloids, the insoluble, fibrous protein deposits that play a role in Alzheimer's disease, type 2 diabetes, Parkinson's disease, heart disease, rheumatoid arthritis, and many other illnesses. The Daggett group has mined the simulations of amyloid producing proteins and found a common structural feature in the intermediate stages between normal and toxic forms. Compounds with physical properties complementary to these structural features were designed on computers and synthesized.

Experiments showed that these custom-designed compounds inhibit amyloid formation in three different protein systems, including those responsible for (1) Alzheimer's disease, (2) Creutzfeldt-Jakob disease and mad cow disease, and (3) systemic amyloid disease and heart disease. The experimental compounds preferentially bind the toxic forms of these proteins. These results suggest it may be possible to design drugs that can treat as many as 25 amyloid diseases, to design diagnostic tools for amyloid diseases in humans, and to design tests to screen our blood and food supplies. A patent application for this technology is in preparation, and further experiments are being planned.

NISE Program Encourages Innovative Research

The NERSC Initiative for Scientific Exploration (NISE) program provides special allocations of computer time for exploring new areas of research, such as:

- A new research area not covered by the existing ERCAP proposal: this could be a tangential research project or a tightly coupled supplemental research initiative.
- New programming techniques that take advantage of multicore compute nodes by using OpenMP, Threads, UPC or CAF: this could include modifying existing codes, creating new applications, or testing the performance and scalability of multicore programming techniques.
- Developing new algorithms that increase researchers' scientific productivity, for example by running simulations at a higher scale or by incorporating new physics.

Twenty-seven projects received NISE awards in 2010 (see sidebar). Although many NISE projects are exploratory, some have already achieved notable scientific results. A few of those are described briefly on the following pages.

Modeling the Energy Flow in the Ozone Recombination Reaction

The ozone-oxygen cycle is the process by which ozone (O_3) and oxygen (O_2) are continuously converted back and forth from one to the other in Earth's stratosphere. This chemical process converts solar ultraviolet radiation into heat and greatly reduces the amount of harmful UV radiation that reaches the lower atmosphere. The NISE project "Modeling the Energy Flow in the Ozone Recombination Reaction," led by Dmitri Babikov of Marquette University, aims to resolve the remaining mysteries of what determines the isotopic composition of atmospheric ozone and to investigate unusual isotope effects in other atmospheric species. These effects are poorly understood, mainly due to computational difficulties associated with the quantum dynamics treatment of polyatomic systems.

Using their NISE allocation, Babikov and his collaborator Mikhail V. Ivanov have developed and applied a mixed quantum-classical approach to the description of collisional energy transfer, in which the vibrational motion of an energized molecule is treated quantum mechanically using wave packets, while the collisional motion of the molecule plus quencher and the rotational motion of the molecule are treated using classical trajectories. For typical calculations, the number of classical trajectories needed is between 10,000 and 100,000, which can be propagated on different processors,

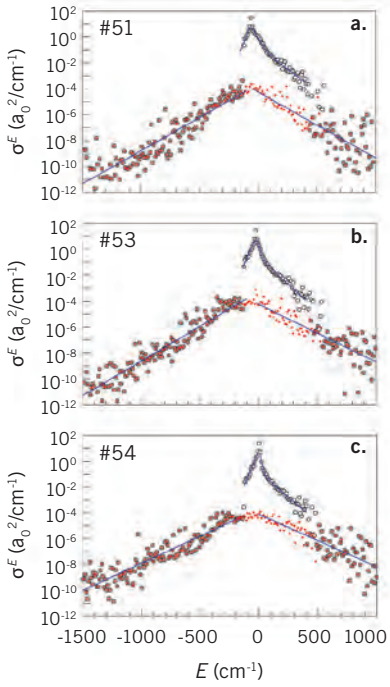

Figure 30. Energy transfer functions expressed as the differential (over energy) cross sections. Three frames correspond to three initial states. Transition and stabilization cross sections are shown as black circles and red dots, respectively. Presence of the van der Waals states in the range $-150 < E < +500$ cm^{-1} lifts circles with respect to the dots by several orders of magnitude. Outside of this range circles and dots coincide. The fit of each wing by the exponential model is shown as solid line. The elastic scattering peak is seen in each frame. *Image: M. V. Ivanov and D. Babikov.*

while the vibrational wave packet fits into the memory of one node.

Ivanov and Babikov applied this mixed quantum-classical approach to the kinetics of ozone formation, performing the most sophisticated calculations of the energy transfer step ever done.[17, 18] They identified the mechanism by which the van der Waals states of ozone contribute to the ozone-forming recombination reaction (Figure 30). This method of treating the dynamics of the energy transfer allows the incorporation of quantum zero-point energy, scattering resonances, and symmetry—the features essential for understanding the isotope effect.

Explaining the anomalous isotope effects in O_3, NO_2, and CO_2 will significantly improve our understanding of their production, chemistry, lifetime, and loss in the atmosphere. That knowledge will help to identify and remove pollution sources as well as monitor the ozone hole, with the possible impact on enhancing the security of all life on the planet. It will allow the isotopic composition of oxygen to be used as a reliable probe of its source and history and provide information for studying atmospheric chemistry, global climate change, atmospheres of other planets, and the history of the solar system.

Modeling Loss Mechanisms in Nitride Light Emitters

Nitride light-emitting devices already have a wide range of applications, such as the lasers in Blu-Ray players and the white LEDs of bicycle lights. In the future, they may also be used as general white light sources, replacing the existing incandescent and fluorescent light bulbs, or in tiny laser projectors that can fit inside a cell phone. At present, however, these devices are not as efficient

[17] M. V. Ivanov and D. Babikov, "Mixed quantum-classical theory for the collisional energy transfer and the rovibrational energy flow: Application to ozone stabilization," J. Chem. Phys. (in press).
[18] M. V. Ivanov and D. Babikov, "Collisional stabilization of van der Waals states of ozone," submitted to J. Chem. Phys.

at the high intensities required for these applications. The project "First Principles Modeling of Charged-Defect-Assisted Loss Mechanisms in Nitride Light Emitters," led by Emmanouil Kioupakis of the University of California, Santa Barbara, is investigating why nitride light-emitting devices lose their efficiency at high intensities and suggesting ways to fix this problem.

Several mechanisms have been blamed for this efficiency reduction,

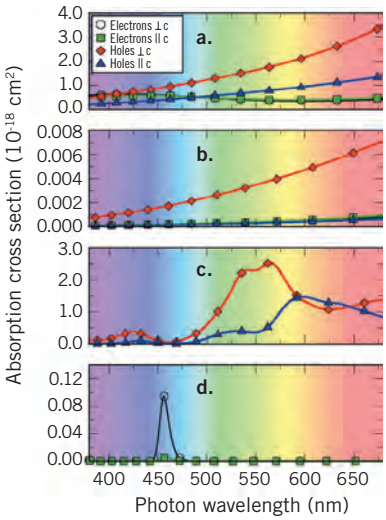

Figure 31. The phonon- and charged-defect-assisted absorption cross-section spectra by free electrons and holes of gallium nitride at room temperature are shown in (a) and (b), respectively. The direct absorption spectrum is plotted in (c) for holes and (d) for electrons. The four lines correspond to absorption by electrons or holes, for light polarized either parallel or perpendicular to the *c* axis. *Image: E. Kioupakis et al., ref. 19.*

including carrier leakage, recombination at dislocations, and Auger recombination, but the underlying microscopic mechanism remains an issue of intense debate. Using atomistic first-principles calculations, Kioupakis and his collaborators have shown that charge-defect-scattering is not the major mechanism that contributes to the efficiency degradation. Instead, the efficiency droop is caused by indirect Auger recombination, mediated by electron-phonon coupling and alloy scattering (Figure 31).[19, 20, 21] Due to the participation of phonons in the Auger process, increasing temperature contributes to the performance degradation. Moreover, the indirect Auger coefficients increase as a function of indium mole fraction and contribute to the efficiency reduction for devices operating at longer wavelengths (the "green gap" problem). By identifying the origin of the droop, these results provide a guide to addressing the efficiency issues in nitride LEDs and to engineering highly efficient solid-state lighting.

Electron Transport for Electronic, Spintronic, and Photovoltaic Devices

With the increasing ability to lithographically create atomic-scale devices, a full-scale quantum transport simulation is becoming necessary for efficiently designing nanoscale electronic, spintronic, and photovoltaic devices. At the nanoscale, material characteristics change significantly from the conventional due to pronounced effects of geometry and extreme quantum confinement in various dimensions. On the other hand, a full quantum mechanical simulation of the transport phenomena in nanostructured materials, with minimum approximations, is extremely difficult computationally.

The NISE project "Electron Transport in the Presence of Lattice Vibrations for Electronic, Spintronic, and Photovoltaic Devices," led by Sayeef Salahuddin of the University of California, Berkeley, aims to massively parallelize the NEST code for solving quantum transport problems, including spin, using the nonequilibrium Green's function formalism. NEST can be used to simulate transport in a wide variety of nanostructures where quantum effects are of significant importance, for example, nanoscale transistors and memory devices, solar cells, and molecular electronics.

Perfect scaling up to 8,192 processor cores has been achieved. This massive parallelization has enabled the researchers to solve problems that had been deemed intractable before. This resulted in multiple publications in leading journals, including a cover story in the July 19, 2010 issue of Applied Physics Letters (Figure 32).[22]

The starting point for that study was an experimental demonstration that a graphene nanoribbon (GNR)

[19] E. Kioupakis, P. Rinke, A. Schleiffe, F. Bechstedt, and C. G. Van de Walle, "Free-carrier absorption in nitrides from first principles," Phys. Rev. B **81**, 241201(R) (2010).
[20] E. Kioupakis, P. Rinke, and C. G. Van de Walle, "Determination of internal loss in nitride lasers from first principles," Appl. Phys. Express **3**, 082101 (2010).
[21] E. Kioupakis, P. Rinke, K. T. Delaney, and C. G. Van de Walle, "Indirect Auger recombination as a cause of efficiency droop in nitride LEDs," submitted (2011).
[22] Youngki Yoon and Sayeef Salahuddin, "Barrier-free tunneling in a carbon heterojunction transistor," Applied Physics Letters **97**, 033102 (2010).

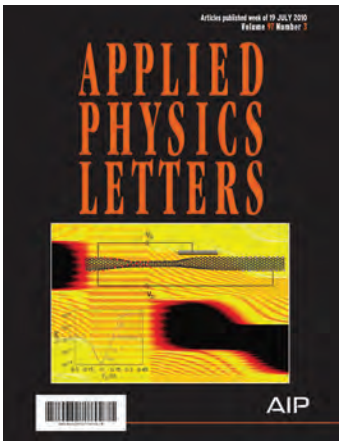

Figure 32. The illustration shows the "local density of states" (the quantum mechanically viable states) in the carbon heterojunction transistor under a gate bias when the transistor is turning ON. The black regions denote the band gap. The yellow regions show where a Klein tunneling is taking place, thereby creating the unique situation that increases energy efficiency in this transistor beyond classical limits. The inset shows the structure of the device and current-voltage characteristics. *Image: ©2011, American Institute of Physics.*

can be obtained by unzipping a carbon nanotube (CNT). This makes it possible to fabricate all-carbon heterostructures that have a unique interface between a CNT and a GNR. By performing a self-consistent nonequilibrium Green's function based calculation on an atomistically defined structure, Salahuddin and collaborator Youngki Yoon demonstrated that such a heterojunction may be utilized to obtain a unique transistor operation. They showed that such a transistor may reduce energy dissipation below the classical limit while not compromising speed—thus providing an alternative route toward ultralow-power, high-performance carbon-heterostructure electronics. This means significantly lower need for electrical power for plugged-in electronics (currently estimated to be 3% of total national electricity needs), as well as a significant increase in battery life for all portable devices.

NERSC Initiative for Scientific Exploration (NISE) 2010 Awards

Project: Modeling the Energy Flow in the Ozone Recombination Reaction
PI: Dmitri Babikov, Marquette University
NISE award: 260,000 hours
Research areas/applications: Atmospheric chemistry, chemical physics, climate science, geochemistry

Project: Bridging the Gaps Between Fluid and Kinetic Magnetic Reconnection Simulations in Large Systems
PI: Amitava Bhattacharjee, University of New Hampshire
NISE award: 1,000,000 hours
Research areas/applications: Plasma physics, astrophysics, fusion energy

Project: Decadal Predictability in CCSM4
PIs: Grant Branstator and Haiyan Teng, National Center for Atmospheric Research
NISE award: 1,600,000 hours
Research areas/applications: Climate science

Project: 3-D Radiation/Hydrodynamic Modeling on Parallel Architectures
PI: Adam Burrows, Princeton University
NISE award: 5,000,000 hours
Research areas/applications: Astrophysics

Project: Dependence of Secondary Islands in Magnetic Reconnection on Dissipation Model
PI: Paul Cassak, West Virginia University
NISE award: 200,000 hours
Research areas/applications: Plasma physics, astrophysics, fusion energy

Project: Simulation of Elastic Properties and Deformation Behavior of Lightweight Protection Materials under High Pressure
PI: Wai-Yim Ching, University of Missouri, Kansas City
NISE award: 1,725,000 hours
Research areas/applications: Materials science, protective materials

Project: Molecular Mechanisms of the Enzymatic Decomposition of Cellulose Microfibrils
PI: Jhih-Wei Chu, University of California, Berkeley
NISE award: 2,000,000 hours
Research areas/applications: Chemistry, biofuels

Project: Developing an Ocean-Atmosphere Reanalysis for Climate Applications (OARCA) for the Period 1850 to Present
PI: Gil Compo, University of Colorado, Boulder
NISE award: 2,000,000 hours
Research areas/applications: Climate science

Project: Surface Input Reanalysis for Climate Applications 1850–2011
PI: Gil Compo, University of Colorado at Boulder
NISE award: 1,000,000 hours
Research areas/applications: Climate science

NISE 2010 Awards (continued)

Project: Pushing the Limits of the GW/BSE Method for Excited-State Properties of Molecules and Nanostructures
PI: Jack Deslippe, University of California, Berkeley
NISE award: 1,000,000 hours
Research areas/applications: Materials science, protective materials

Project: Reaction Pathways in Methanol Steam Reformation
PI: Hua Guo, University of New Mexico
NISE award: 200,000 hours
Research areas/applications: Chemistry, hydrogen fuel

Project: Warm Dense Matter Simulations Using the ALE-AMR Code
PI: Enrique Henestroza, Lawrence Berkeley National Laboratory
NISE award: 450,000 hours
Research areas/applications: Plasma physics, fusion energy

Project: Modeling the Dynamics of Catalysts with the Adaptive Kinetic Monte Carlo Method
PI: Graeme Henkelman, University of Texas, Austin
NISE award: 1,700,000 hours
Research areas/applications: Chemistry, alternative energy

Project: First-Principles Study of the Role of Native Defects in the Kinetics of Hydrogen Storage Materials
PI: Khang Hoang, University of California, Santa Barbara
NISE award: 600,000 hours
Research areas/applications: Materials science, hydrogen fuel

Project: ITER Rapid Shutdown Simulation
PI: Valerie Izzo, General Atomics
NISE award: 1,200,000 hours
Research areas/applications: Fusion energy

Project: First-Principles Modeling of Charged-Defect-Assisted Loss Mechanisms in Nitride Light Emitters
PI: Emmanouil Kioupakis, University of California, Santa Barbara
NISE award: 900,000 hours
Research areas/applications: Materials science, energy-efficient lighting

Project: Thermodynamic, Transport, and Structural Analysis of Geosilicates using the SIESTA First Principles Molecular Dynamics Software
PI: Ben Martin, University of California, Santa Barbara
NISE award: 110,000 hours
Research areas/applications: Geosciences

Project: Computational Prediction of Protein-Protein Interactions
PI: Harley McAdams, Stanford University
NISE award: 1,880,000 hours
Research areas/applications: Biological systems science

Project: Light Propagation in Nanophotonics Structures Designed for High-Brightness Photocathodes
PIs: Karoly Nemeth and Katherine Harkay, Argonne National Laboratory
NISE award: 1,300,000 hours
Research areas/applications: Materials science, high-brightness electron sources

Project: International Linear Collider (ILC) Damping Ring Beam Instabilities
PI: Mauro Pivi, Stanford Linear Accelerator Center
NISE award: 250,000 hours
Research areas/applications: Accelerator physics

Project: Thermonuclear Explosions from White Dwarf Mergers
PI: Tomasz Plewa, Florida State University
NISE award: 500,000 hours
Research areas/applications: Astrophysics

Project: Beam Delivery System Optimization of Next-Generation X-Ray Light Sources
PIs: Ji Qiang, Robert Ryne, and Xiaoye Li, Lawrence Berkeley National Laboratory
NISE award: 1,000,000 hours
Research areas/applications: Accelerator physics

Project: Electron Transport in the Presence of Lattice Vibrations for Electronic, Spintronic, and Photovoltaic Devices
PI: Sayeef Salahuddin, University of California, Berkeley
NISE award: 600,000 hours
Research areas/applications: Materials science, ultra-low-power computing

Project: High Resolution Climate Simulations with CCSM
PIs: Warren Washington, Jerry Meehl, and Tom Bettge, National Center for Atmospheric Research
NISE award: 3,525,000 hours
Research areas/applications: Climate science

Project: Abrupt Climate Change and the Atlantic Meridional Overturning Circulation
PI: Peter Winsor, University of Alaska, Fairbanks
NISE award: 2,000,000 hours
Research areas/applications: Climate science

Project: Modeling Plasma Surface Interactions for Materials under Extreme Environments
PI: Brian Wirth, University of California, Berkeley
NISE award: 500,000 hours
Research areas/applications: Fusion energy

Project: Models for the Explosion of Type Ia Supernovae
PI: Stan Woosley, University of California, Santa Cruz
NISE award: 2,500,000 hours
Research areas/applications: Astrophysics

NERSC Users' Awards and Honors

National Medal of Science

Warren Washington
National Center for Atmospheric Research

Fellows of the American Academy of Arts and Sciences

Andrea Louise Bertozzi
University of California, Los Angeles

Adam Seth Burrows
Princeton University

Gary A. Glatzmaier
University of California, Santa Cruz

Member of the National Academy of Sciences

Gary A. Glatzmaier
University of California, Santa Cruz

Fellows of the American Association for the Advancement of Science

Liem Dang
Pacific Northwest National Laboratory

Alenka Luzar
Virginia Commonwealth University

Gregory K. Schenter
Pacific Northwest National Laboratory

Yousef Saad
University of Minnesota

Chris G. Van de Walle
University of California, Santa Barbara

Fellows of the American Physical Society (APS)

Mark Asta
University of California, Davis

David Bruhwiler
Tech-X Corporation

Liem Dang
Pacific Northwest National Laboratory

Francois Gygi
University of California, Davis

Julius Jellinek
Argonne National Laboratory

En Ma
Johns Hopkins University

Barrett Rogers
Dartmouth College

Philip Snyder
General Atomics

Ramona Vogt
Lawrence Livermore National Laboratory

Fuqiang Wang
Purdue University

Martin White
University of California, Berkeley and Lawrence Berkeley National Laboratory

Zhangbu Xu
Brookhaven National Laboratory

APS James Clerk Maxwell Prize for Plasma Physics

James Drake
University of Maryland

APS John Dawson Award for Excellence in Plasma Physics Research

Eric Esarey
Lawrence Berkeley National Laboratory

Cameron Geddes
Lawrence Berkeley National Laboratory

Carl Schroeder
Lawrence Berkeley National Laboratory

American Chemical Society (ACS) Ahmed Zewail Prize in Molecular Sciences

William H. Miller
University of California, Berkeley and Lawrence Berkeley National Laboratory

Peter Debye Award in Physical Chemistry

George C. Schatz
Northwestern University

Fellow of the Materials Research Society

Chris Van de Walle
University of California, Santa Barbara

IEEE Sidney Fernbach Award

James Demmel
University of California, Berkeley and Lawrence Berkeley National Laboratory

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 2010.

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

Andrea Louise Bertozzi
University of California, Los Angeles

Yousef Saad
University of Minnesota

SIAM Junior Scientist Prize

Kamesh Madduri
Lawrence Berkeley National Laboratory

International Council for Industrial and Applied Mathematics (ICIAM) Pioneer Prize

James Demmel
University of California, Berkeley and Lawrence Berkeley National Laboratory

Gordon Bell Prize

Aparna Chandramowlishwaran
Lawrence Berkeley National Laboratory

Logan Moon
Georgia Institute of Technology

Association for Computing Machinery (ACM) Distinguished Scientists

Wu-chun Feng
Virginia Polytechnic Institute and State University

Kesheng (John) Wu
Lawrence Berkeley National Laboratory

Academician of Academia Sinica

Inez Yau-Sheung Fung
University of California, Berkeley and Lawrence Berkeley National Laboratory

Presidential Early Career Award for Scientists and Engineers (PECASE)

Cecilia R. Aragon
Lawrence Berkeley National Laboratory

Joshua A. Breslau
Princeton Plasma Physics Laboratory

U.S. Air Force Young Investigator Research Program

Per-Olof Persson
University of California, Berkeley and Lawrence Berkeley National Laboratory

NASA Public Service Group Achievement Award

Julian Borrill
Lawrence Berkeley National Laboratory

Christopher Cantalupo
Lawrence Berkeley National Laboratory

Theodore Kisner
Lawrence Berkeley National Laboratory

The NERSC Center

NERSC is the primary computing center for the DOE Office of Science, serving approximately 4,000 users, hosting 400 projects, and running 700 codes for a wide variety of scientific disciplines. NERSC has a tradition of providing systems and services that maximize the scientific productivity of its user community. NERSC takes pride in its reputation for the expertise of its staff and the high quality of services delivered to its users. To maintain its effectiveness, NERSC proactively addresses new challenges in partnership with the larger high performance computing (HPC) community.

The year 2010 was a transformational year for NERSC, with an unprecedented number of new systems and upgrades, including the acceptance of Hopper Phase 1; the installation of Hopper Phase 2, the Carver and Magellan clusters, the Euclid analytics server, and the Dirac GPU testbed; upgrades of the NERSC Global Filesystem and the tape archive; and a major upgrade to the Oakland Scientific Facility power supply. We also made significant improvements in the Center's energy efficiency.

At the same time, NERSC staff were developing innovative ways to improve the usability of HPC systems, to increase user productivity, and to explore new computational models that can make the most effective use of emerging computer architectures.

The following pages describe some of the ways NERSC is meeting current challenges and preparing for the future.

Hopper Powers Petascale Science

The installation of NERSC's newest supercomputer, Hopper (Figure 1), marks a milestone—it is NERSC's first petascale computer, posting a performance of 1.05 petaflops (quadrillions of calculations per second) running the Linpack benchmark. That makes Hopper, a 153,216-processor-core Cray XE6 system, the fifth most powerful supercomputer in the world and the second most powerful in the United States, according to the November 2010 edition of the TOP500 list, the definitive ranking of the world's top computers.

NERSC selected the Cray system in an open competition, in which vendors ran NERSC application benchmarks from several scientific domains that use a variety of algorithmic methods. The Cray system showed the best application performance per dollar and per megawatt of power. It features external login nodes and an external filesystem for increased functionality and availability.

Hopper was installed in two phases. The first phase, a 5,344-core Cray XT5, was installed in October 2009 and entered production use in March 2010. Phase 1 helped NERSC staff optimize the external node architecture. "Working out the kinks in Phase 1 will ensure a more risk-free deployment when Phase 2 arrives," said Jonathan Carter, who led the Hopper system procurement as head of NERSC's User Services Group. "Before accepting the Phase 1 Hopper system, we encouraged all 300 science projects computing at NERSC to use the system during the pre-production period to see whether it could withstand the gamut of scientific demands that we typically see."

The new external login nodes on Hopper offer users a more responsive environment. Compared to NERSC's Cray XT4 system, called Franklin, the external login nodes have more memory, and in aggregate, have more computing power. This allows users to compile applications faster and run small post-processing or visualization jobs directly on the login nodes without interference from other users. Even with dedicated testing and maintenance times, utilization of Hopper-1 from December 15 to March 1 reached 90%, and significant scientific results were produced even during the testing period, including high-resolution simulations of magnetic reconnection, and simulations of peptide aggregation into amyloid fibrils.

Hopper Phase 2, the Cray XE6, was fully installed in September 2010 and opened for early use and testing late in the year. It has 6,384 nodes, with two 12-core AMD "Magny-Cours" chips per node, and features Cray's high-bandwidth, low-latency Gemini interconnect. With 168 GB/sec routing capacity and internode latency on the order of 1 microsecond, Gemini supports millions of MPI messages per second, as well as OpenMP, Shmem, UPC, and Co-Array Fortran. The innovative ECOphlex liquid cooling technology increases Hopper's energy efficiency.

Because Hopper has 2 PB of disk space and 70 GB/sec of bandwidth on the external filesystem, users with extreme data demands are seeing fewer bottlenecks when they move their data in and out of the machine. Additionally, the availability of dynamically loaded libraries enables even more applications to run on the system and adds support for popular frameworks like Python. This feature helps ensure that the system is optimized for scientific productivity.

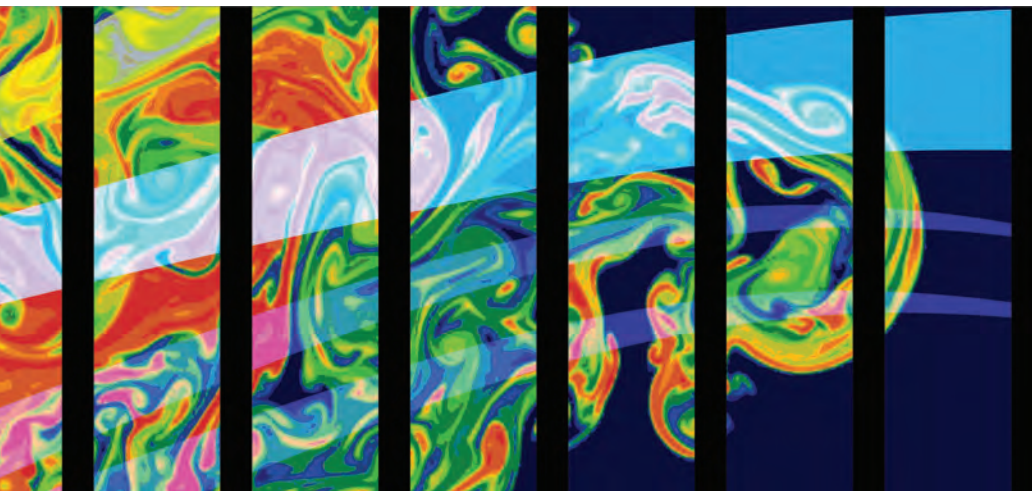

Figure 1. NERSC's newest supercomputer is named after American computer scientist Grace Murray Hopper, a pioneer in the field of software development and programming languages who created the first compiler. She was a champion for increasing the usability of computers, understanding that their power and reach would be limited unless they were made more user-friendly. *Cabinet façade design by Caitlin Youngquist, Berkeley Lab Creative Services Office. Grace Hopper photo courtesy of the Hagley Museum and Library, Wilmington, Delaware.*

Other New Systems and Upgrades

Carver and Magellan Clusters

The Carver and Magellan clusters (Figure 2), built on liquid-cooled IBM iDataPlex technology, were installed in March 2010. Carver, named in honor of American scientist and inventor George Washington Carver, entered production use in May, replacing NERSC's Opteron Jacquard cluster and IBM Power5 Bassi system, which were both decommissioned at the end of April.

Figure 2. Top view of the Carver and Magellan clusters showing trays of InfiniBand cables. *Photo: Roy Kaltschmidt, Berkeley Lab Public Affairs.*

Carver contains 800 Intel Nehalem quad-core processors, or 3,200 cores. With a theoretical peak performance of 34 teraflops, Carver has 3.5 times the capacity of Jacquard and Bassi combined, yet uses less space and power. The system's 400 compute nodes are interconnected by the latest 4X QDR InfiniBand technology, meaning that 32 GB/s of point-to-point bandwidth is available for high performance message passing and I/O. Carver is one of the first platforms of this scale to employ four Voltaire switches, allowing information to quickly transfer between the machine and its external filesystem.

NERSC staff have configured Carver's batch queue and scratch storage to handle a range of large and small jobs, making Carver NERSC's most popular system. A queue that allows use of up to 32 processors now has the option to run for 168 hours, or seven straight days.

Early scientific successes on Carver included calculations of the atomic interactions of titanium, other transition metals, and their alloys; and computational modeling of data from x-ray crystallography and small-angle x-ray scattering that reveals how protein atoms interact inside of a cell.

The Magellan cluster, with hardware identical to Carver's but different software, is a testbed system and will be discussed below on page 68.

Euclid Analytics Server and Dirac GPU Testbed

The Euclid analytics server and Dirac GPU testbed were installed in May 2010. Euclid, named in honor of the ancient Greek mathematician, is a Sun Microsystems Sunfire x4640 SMP. Its single node contains eight six-core Opteron 2.6 GHz processors, with all 48 cores sharing the same 512 GB of memory. The system's theoretical peak performance is 499.2 Gflop/s. Euclid supports a wide variety of data analysis and visualization applications.

Dirac, a 50-node hybrid CPU/GPU cluster with Intel Nehalem and NVIDIA Fermi chips, is an experimental system dedicated to exploring the performance of scientific applications on a GPU (graphics processing unit) architecture. Research conducted on Dirac is discussed on page 71.

NERSC Global Filesystem and HPSS Tape Archive Upgrades

"Historically, each new computing platform came with its own distinct storage hardware and software, creating isolated storage 'islands,'" wrote NERSC Storage Systems Group lead Jason Hick in HPC Source, a supplement to Scientific Computing magazine. "Users typically had different filesystem quotas on each machine and were responsible for keeping track of their data, from determining where the data was to transferring files between machines."[1]

In recent years, however, NERSC began moving from this "island" model toward global storage resources. The center was among the first to implement a shared filesystem, called the NERSC Global Filesystem (NGF), which can be accessed from any of NERSC's major computing platforms. This service facilitates file sharing between platforms, as well as file sharing among NERSC users working on a common project. This year, to allow users to access nearly all their data regardless of the system they log into, NGF was expanded from project directories to include scratch and home directories as well. (Project directories provide a large-capacity file storage resource for groups of users; home directories are used for permanent storage of source code and other relatively small files; and scratch directories are used for temporary storage of large files.) With these tools, users spend less time moving data between systems and more time using it.

Another aspect of NERSC's storage strategy is a multi-year effort to improve the efficiency of the High Performance Storage System (HPSS) tape archive by upgrading to higher-capacity media. In just the past year, the number of tapes has been reduced from 43,000 to 28,000 (Figure 3), while still providing capacity for the continuing exponential growth in data (Figure 4). Disk cache capacity and bandwidth have been doubled, and NERSC has invested in HPSS client software improvements to accommodate future increases.

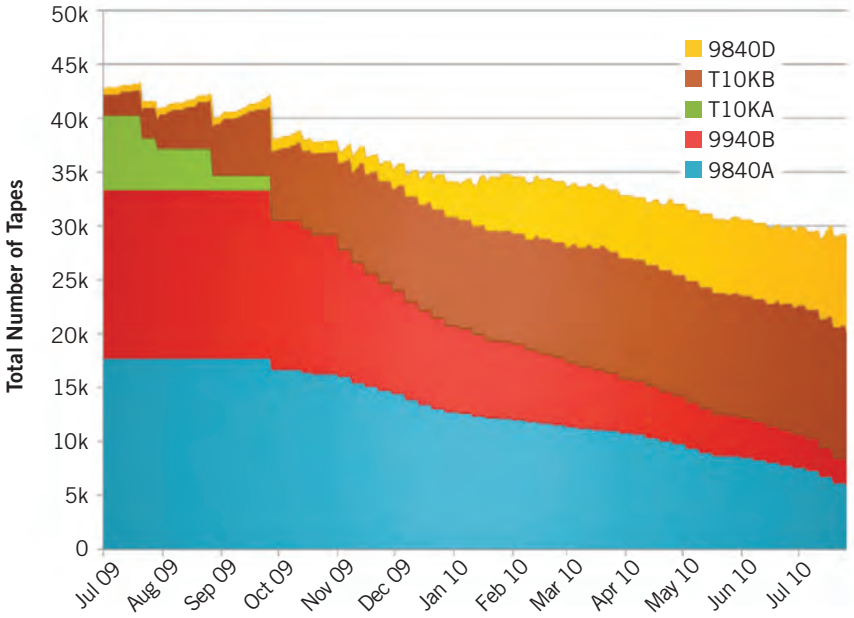

Figure 3. Number of HPSS tapes, July 2009 to August 2010.

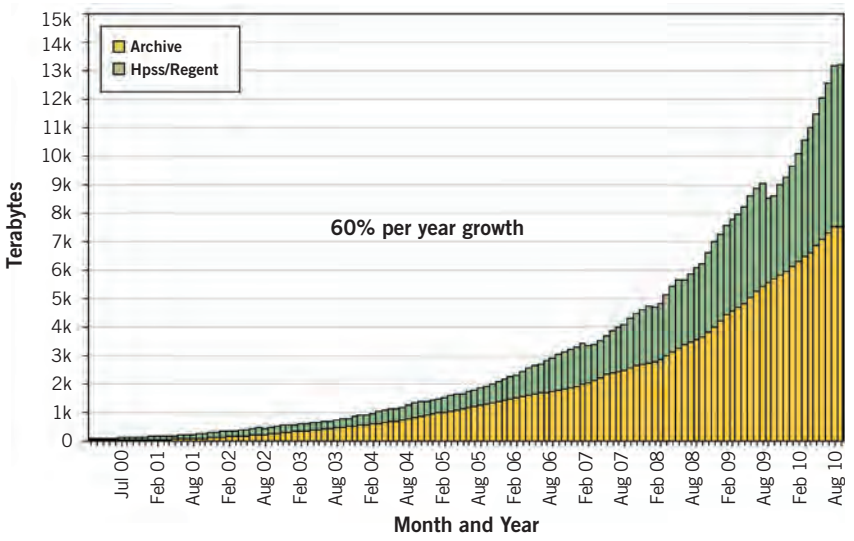

Figure 4. Cumulative storage by month and system, January 2000 to August 2010, in terabytes.

[1]Jason Hick, "Big-picture storage approach," HPC Source, Summer 2010, pp. 5–6.

Innovations to Increase NERSC's Energy Efficiency

In July 2010, NERSC scheduled a four-day, center-wide outage for installation of a 3 megawatt power supply upgrade to the computer room. The total 9 megawatt capacity was needed to accommodate Hopper Phase 2 along with existing and future systems.

But perhaps more important than the increased power supply were NERSC's innovative solutions to improve energy efficiency, which has become a major issue for high performance computing and data centers. The energy consumed by data center servers and related infrastructure equipment in the United States and worldwide doubled between 2000 and 2005 and is continuing to grow. A 2007 Environmental Protection Agency report to Congress estimated that in 2006, data centers consumed about 1.5 percent of total U.S. electricity consumption—equivalent to 5.8 million average households—and that federal servers and data centers accounted for 10 percent of that figure.

Cost and carbon footprint are critical concerns, but performance is suffering, too. In an interview in HPC Source, NERSC Director Kathy Yelick said, "The heat, even at the chip level, is limiting processor performance. So, even from the innards of a computer system, we are worried about energy efficiency and getting the most computing with the least energy."[2]

Two innovative projects to increase energy efficiency at NERSC are attracting attention throughout the HPC community: the innovative configuration and cooling system for the Carver cluster, and the extensive energy monitoring system now operating in the NERSC machine room.

Carver Installation Saves Energy, Space, Water

NERSC staff configured the IBM iDataplex system Carver in a manner so efficient that, in some places, the cluster actually cools the air around it. By setting row pitch at 5 ft, rather than the standard 6 ft, and reusing the 65° F water that comes out of the Franklin Cray XT4 to cool the system, the team was able to reduce cooling costs by 50 percent and to use 30 percent less floor space. NERSC collaborated with vendor Vette Corporation to custom-design a cooling distribution unit (CDU) for the installation.

By narrowing row pitch to 5 ft from cabinet-front to cabinet-front (a gap of only 30 inches between rows), the Carver team was able to use only one "cold row" to feed the entire installation. (A standard installation would require cooling vents along every other row.) At such a narrow pitch, the IBM iDataplex's transverse cooling architecture recirculates cool air from the previous row while the air is still at a low temperature. When the air passes through Carver's water-cooled back doors (the cool water coming from Franklin), it is further cooled and passed to the next row. When the air finally exits the last row of Carver, it can be several degrees cooler than when it entered.

Using a design concept from NERSC, Vette Corporation custom-designed a CDU that reroutes used water coming out of Franklin back into Carver. Water from the building's chillers enters Franklin's cooling system at 52° F. The water exits the system at about 65° F. Sending the water back to be chilled at such low temperatures is inefficient. Instead, the custom-designed CDU increases efficiency by rerouting the water through Carver to pick up more heat, sending water back to the chillers at 72° F and enabling one CDU to cool 250 kW of computing equipment.

A case study published by IBM, titled "NERSC creates an ultra-efficient supercomputer," observed:

> *The sheer size of the rear doors—which measure four feet by seven feet—means that the air can move more slowly through them and has more time to exchange heat with the water. This in turn enables slower fans, dramatically reducing the energy wasted as noise, and making the Carver cluster the quietest place in the machine room.*

NERSC has shared these innovations with the wider HPC community and with IBM, which now uses NERSC techniques in new installations; and Vette is working to turn this custom design into a commercial product.

Monitoring Energy Across the Center to Improve Efficiency

NERSC has instrumented its machine room with state-of-the-art wireless monitoring technology from SynapSense to study and optimize

[2]Mike May, "Sensing the future of greener data centers," HPC Source, Autumn 2010, pp. 6-8.

energy use. The center has installed 992 sensors, which gather information on variables important to machine room operation, including air temperature, pressure, and humidity. This data is collected by NERSC's SynapSense system software, which generates heat, pressure, and humidity maps (Figure 5).

"Just a glance at one of these maps tells you a lot," says Jim Craw, NERSC's newly appointed Risks and Energy Manager. "If you start developing hot or cold spots, you're getting an unambiguous visual cue that you need to investigate and/or take corrective action."

The system can also calculate the center's power usage effectiveness, an important metric that compares the total energy delivered to a data center with that consumed by its computational equipment.

Responding quickly to accurate machine room data not only enhances cooling efficiency, it helps assure the reliability of center systems. For example, after cabinets of the decommissioned Bassi system were shut down and removed, cold air pockets developed near Franklin's air handling units. The SynapSense system alerted administrators, who were able to quickly adjust fans and chillers in response to these changes in airflow and room temperatures before Franklin was negatively affected. The monitoring system also revealed changes in machine room airflow following the power upgrade. Repositioning some floor tiles and partitions corrected the airflow and rebalanced air temperatures.

In addition to the SynapSense sensors, NERSC's two Cray systems have their own internal sensors; and sensor readings from two different chiller plants and air handlers come in on their own separate systems as well. NERSC staff created several applications to bring the Cray and SynapSense data together, and work continues on integrating the others.

Future energy efficiency projects include controlling water flow and fan speeds within the air handling units based on heat load, and placing curtains or "top hats" around systems to optimize the air flows in their vicinity.

Improving HPC System Usability with External Services

Feedback from NERSC's Franklin users has led NERSC and Cray to jointly develop configuration

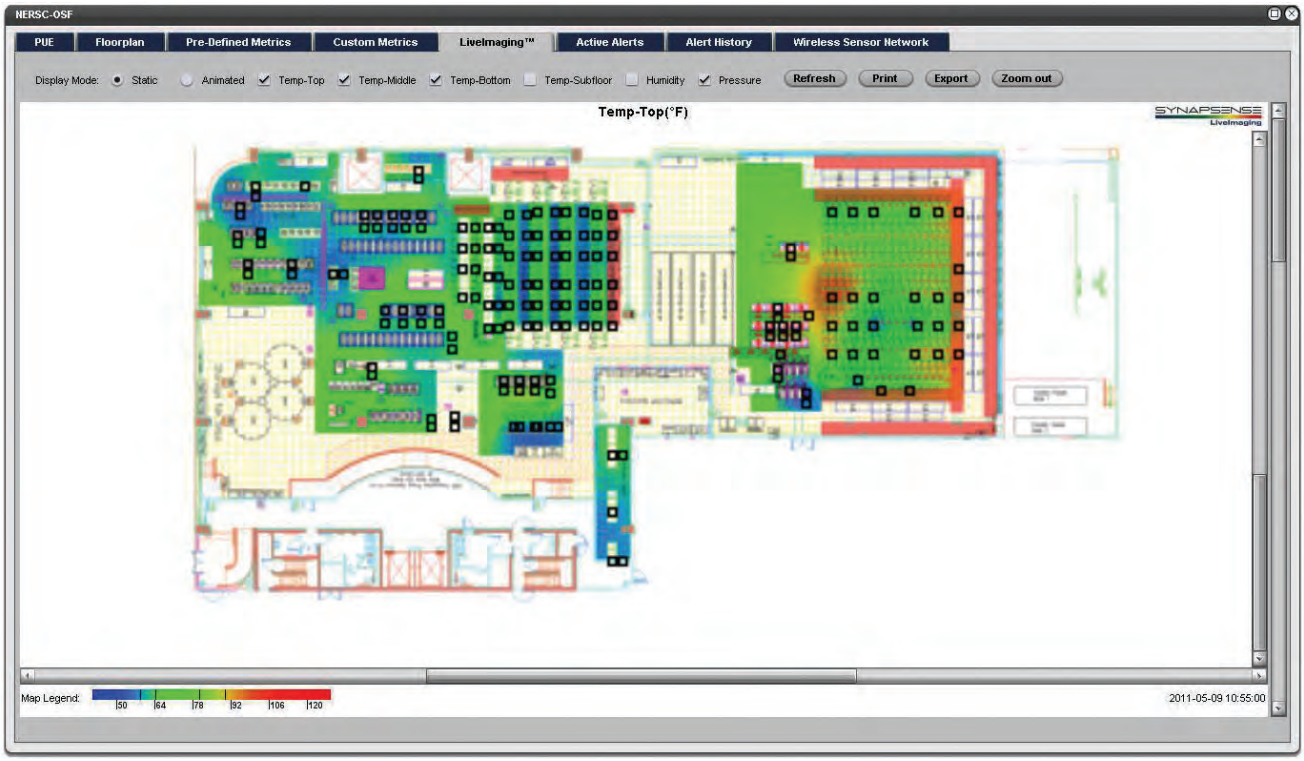

Figure 5. NERSC's energy monitoring system generates heat maps that can be used to look for areas where efficiency can be improved.

improvements for both Franklin and Hopper that improve their usability and availability. These improvements, called *external services*, are now available on all Cray XE6 systems, and include external login nodes, reconfigurable head nodes (also called machine-oriented miniserver, or MOM nodes), and an external global filesystem.

While NERSC users have made great use of the Franklin Cray XT4 system, the large number and variety of users, projects, applications, and workflows running at NERSC exposed several issues tied to the fact that almost all user-level serial processing (e.g., compiling, running analysis, pre- or post-processing, file transfer to or from other systems) is confined to a small number of shared service nodes. These nodes have limited memory (8 GB), no swap disk, and older processor technology. NERSC user applications had often run the Franklin service nodes out of memory, primarily due to long and involved application compilation, analytics applications, and file transfers. In the Franklin configuration, it was hard for NERSC to increase the number of login nodes as demand for this service increased.

During the contract negotiation for the Hopper system, NERSC staff worked with Cray to design, configure, and support eight external login nodes, located outside of the Hopper Gemini interconnect; they have more memory (128 GB), more cores (16), and more disk than the original Franklin login nodes. One external login node was also added to Franklin itself.

Head nodes are the nodes that launch parallel jobs and run the serial parts of these jobs. Initially on Franklin this function ran on the login nodes, and large interactive jobs could cause congestion and batch launch failures. On Franklin the solution was to divide the servers into a login pool and a head node pool. On Hopper, the hardware for the head nodes is the same as the hardware for the compute nodes, allowing compute nodes to be reconfigured into head nodes and vice versa at boot time. NERSC has initially deployed a higher ratio of head nodes to compute nodes on Hopper than on Franklin. This configuration provides better support for complex workflows. As a result, NERSC's "power login" users now have the capabilities they need, and they are no longer disruptive to other users.

In addition to external login nodes, login node balancers have improved interactive access to NERSC systems. Assigning login nodes to incoming interactive sessions in the traditional round-robin fashion can lead to congestion, or worse, the perception that an entire system is down when the user lands on a bad login node or one that has been taken offline. Load balancers now intercept incoming sessions and distribute them to the least-loaded active node. Also, staff can now take down nodes at will for maintenance, and the load balancer automatically removes out-of-service nodes from the rotation and redirects users attempting to connect to a cached (but offline) node to an active one. As a result, users no longer confuse a downed node with a downed system.

While load balancers are commonly used for web and database applications (where sessions tend to be extremely short and the amount of data transferred quite small), these devices have not typically been used in a scientific computing context (where connections are long-lived and data transfers large). To take advantage of this technology, NERSC staff worked closely with vendor Brocade. Together they conceived, adapted, and implemented three improvements to support HPC environments: (1) support for 10 GB/s connections, needed for HPC I/O; (2) support for jumbo frames, needed for large data sets; and (3) support for long-lasting login sessions. Brocade has integrated these innovations into their product line, making them available to other HPC centers.

Availability of files has also been improved. When Franklin was first acquired, it could not be integrated into the NERSC Global Filesystem (NGF), which uses GPFS software, because Franklin's I/O nodes only ran the Lustre filesystem software. As a result, files on Franklin were not accessible from other NERSC systems. So NERSC worked with Cray to test, scale, and tune the Data Virtualization Services (DVS) software that Cray had purchased and developed for the solution, and in 2010 the DVS software went into production. Now NGF hosts a global *project* filesystem that allow users to share files between systems and among diverse collaborators. In addition, NERSC introduced a GPFS-based global *home* filesystem that provides a convenient way for users to access source files, input files, configuration files, etc., regardless of the platform the user is logged onto; and a global *scratch* filesystem that can be used to temporarily store large amounts of data and access it from any system. When user applications load one of these dynamic shared libraries, the compute nodes use DVS to access and cache the libraries or data.

Taken together, the benefit of these external services (Figure 6) is improved usability and availability: users do not have to worry about whether Hopper is up or down, because filesystem and login nodes and software are outside the main system. So, users can log in, access files, compile jobs, submit jobs, and analyze data even if the main system is down.

Providing New Computational Models

Magellan Cloud Computing Testbed and Evaluation

Cloud computing can be attractive to scientists because it can give them on-demand access to compute resources and lower wait times than the batch-scheduled, high-utilization systems at HPC centers. Clouds can also offer custom-built environments such as the software stack used with their scientific code. And some clouds have implemented data programming models for data intensive computing such as MapReduce.

With American Recovery and Reinvestment Act (ARRA) funding, NERSC has deployed the Magellan science cloud testbed for the DOE Office of Science. It leverages NERSC experience in provisioning a specialized cluster (PDSF) for the high energy and nuclear physics community. NERSC is focused on understanding what applications map well to existing and emerging commercial solutions, and what aspects of cloud computing are best provided at HPC centers like NERSC.

To date we have demonstrated on-demand access to cycles for the DOE Joint Genome Institute (JGI) with a configuration using a 9 GB network provided by ESnet and 120 nodes from the Magellan cluster. This arrangement allows biological data to remain at JGI even though the compute servers are at NERSC. One early success on this cluster comes from a JGI team using Hadoop to remove errors from 5 billion reads of next-generation DNA sequence data.

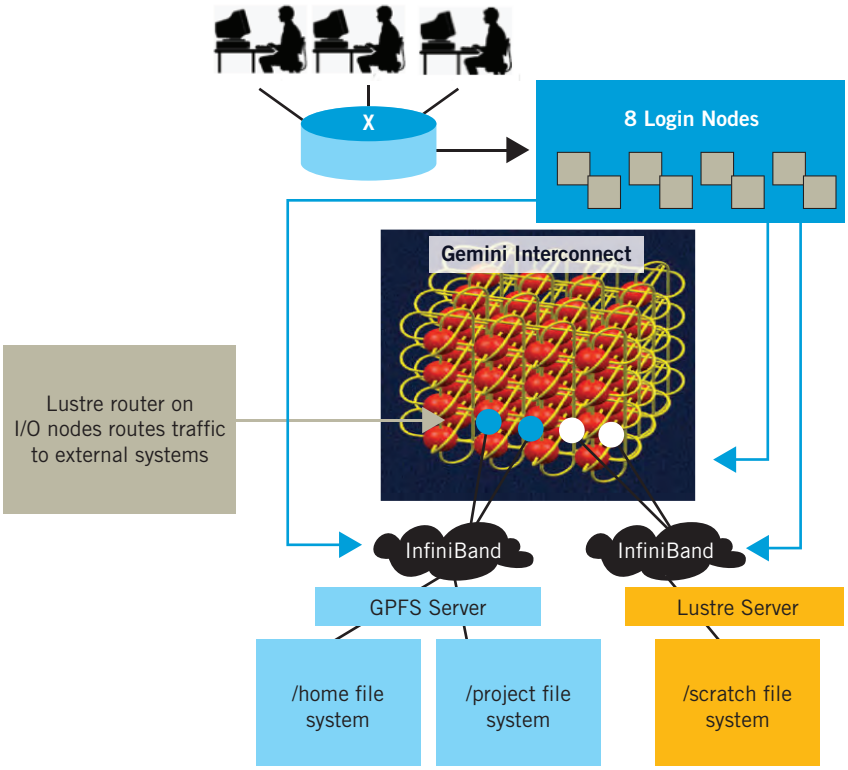

Figure 6. Schematic of external nodes on Hopper.

We have also deployed a MapReduce cluster running Hadoop, Apache's open-source implementation of MapReduce, which includes both a framework for running MapReduce jobs as well as the Hadoop Distributed File System (HDFS), which provides a locality aware, fault tolerant file system to support MapReduce jobs. NERSC is working with targeted communities such as bioinfomatics to investigate how applications like BLAST perform inside the MapReduce framework. Other users are using the Hadoop cluster to perform their own evaluations, including implementing algorithms from scratch in the programming model. These experiences will provide insight into how to best take advantage of this new approach to distributed computing.

The first published performance analysis research from Magellan collaborators shows that cloud computing will not work for science unless the cloud is optimized for it. The paper "Performance Analysis of High Performance Computing Applications on the Amazon Web Services Cloud,"[3] written by a team of researchers from Berkeley Lab's Computational Research Division, Information Technology Division,

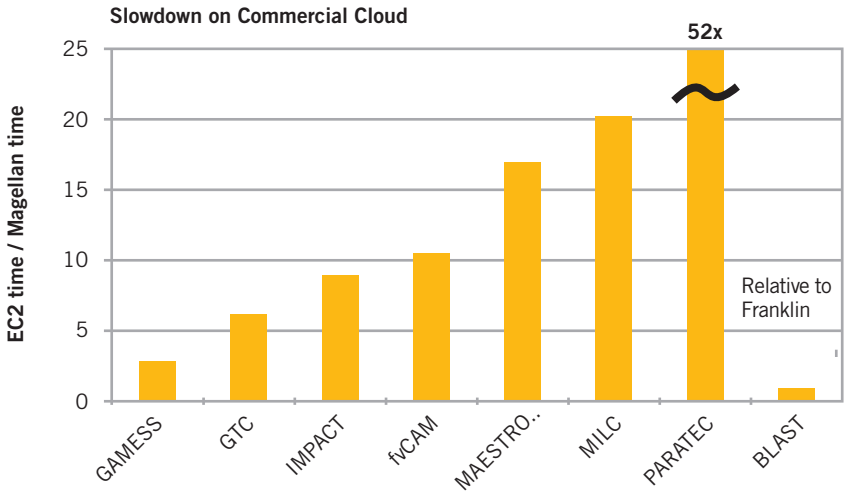

Figure 7. On traditional science workloads, standard cloud configurations see significant slowdown (up to 52x), but independent BLAST jobs run well.

and NERSC, was honored with the Best Paper Award at the IEEE's International Conference on Cloud Computing Technology and Science (CloudCom 2010), held November 30–December 1 in Bloomington, Indiana.

After running a series of benchmarks designed to represent a typical midrange scientific workload—applications that use less than 1,000 cores—on Amazon's EC2 system, the researchers found that the EC2's interconnect severely limits performance and causes significant variability. Overall, the cloud ran six times slower than a typical midrange Linux cluster, and 20 times slower than a modern high performance computing system.

"We saw that the communication pattern of the application can impact performance," said Keith Jackson, a computer scientist in Berkeley Lab's Computational Research Division (CRD) and lead author of the paper. "Applications like Paratec with significant global communication perform relatively worse than those with less global communication." He noted that the EC2 cloud performance varied significantly for scientific applications because of the shared nature of the virtualized environment, the network, and differences in the underlying non-virtualized hardware.

A similar comparison of performance on the standard EC2 environment versus an optimized environment on the Magellan testbed showed a slowdown of up to 52x on EC2 for most applications, although the BLAST bioinformatics code ran well (Figure 7). Another key issue is cost: EC2 costs 20 cents per CPU hour, while Magellan can run workloads at less than 2 cents per CPU hour, since it is run by a nonprofit government lab.

The Magellan Project was honored with the 2010 HPCwire Readers' Choice Award for "Best Use of HPC in the Cloud."

Developing Best Practices in Multicore Programming

The motivation for studying new multicore programming methods stems from concerns expressed by the DOE Office of Science users about how to rewrite their applications to keep ahead of the exponentially increasing HPC system parallelism. In particular, the increase in intra-node parallelism of the Hopper Phase 2 system presents challenges to existing application implementations—and this is just the first step in a long-term technology trend. Applications and algorithms will need to rely increasingly on fine-grained parallelism, strong scaling, and improved support for fault resilience.

NERSC has worked with Cray to establish a Cray Center of Excellence (COE) to investigate these issues, to identify the correct programming model to express fine-grained parallelism, and to guide the transition of the user community by preparing training materials. The long-term goals are to evaluate advanced programming models and identify a durable approach for programming on the path to exascale.

MPI (Message Passing Interface), which has been the dominant HPC programming model for nearly two decades, does not scale well beyond eight cores. The hybrid OpenMP/MPI model is the most mature approach to multicore programming, and it works well on Hopper; but getting the best performance depends on

[3] K. R. Jackson, L. Ramakrishnan, K. Muriki, S. Canon, S. Cholia, J. Shalf, H. J. Wasserman, and N. J. Wright, "Performance Analysis of High Performance Computing Applications on the Amazon Web Services Cloud," 2010 IEEE Second International Conference on Cloud Computing Technology and Science (CloudCom), pp.159–168, Nov. 30, 2010–Dec. 3, 2010, doi: 10.1109/CloudCom.2010.69.

customizing the implementation for each application.

While MPI is a collection of processes with no shared data, OpenMP is a collection of threads with both private and shared data; so OpenMP can save a significant amount of memory compared with MPI. The key question is how many OpenMP threads to put on a node vs MPI processes: One MPI process on each core? Or one MPI process on the whole node, and within the node use OpenMP thread parallelism? The running time for each option depends on the application: some run slower, some faster, some run faster for a while and then get slower. So the COE is developing guidelines for finding the "sweet spot" of how much OpenMP parallelism to add to various applications (Figure 8).

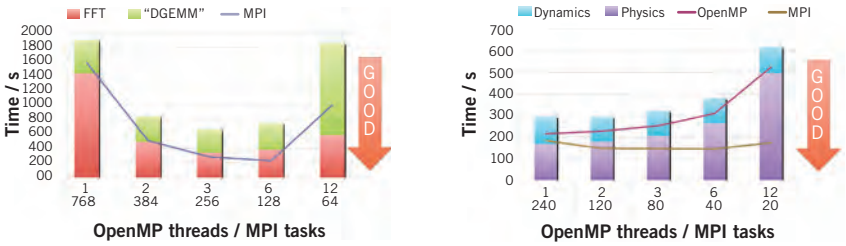

Figure 8. MPI + OpenMP performance on two of NERSC's benchmark applications: (a) Paratec and (b) fvCAM.

Computational Science and Engineering Petascale Initiative

To ensure that science effectively adapts to multicore computing, NERSC is receiving $3 million from the American Recovery and Reinvestment Act to develop the Computational Science and Engineering Petascale Initiative. This initiative identifies key application areas with specific needs for advanced programming models, algorithms, and other support. The applications are chosen to be consistent with the current mission of the Department of Energy, with a particular focus on applications that benefit energy research, those supported by other Recovery Act funding, and Energy Frontier Research Centers (EFRCs). The initiative pairs postdoctoral researchers (Figure 9) with these high-impact projects at NERSC. Here are the projects they are working on:

- Brian Austin is working with high-energy physicists to create computational models that will refine the design of Berkeley Lab's proposed soft x-ray, free electron laser user facility called the Next Generation Light Source.

- Kirsten Fagnan is working with researchers in Berkeley Lab's Center for Computational Sciences and Engineering (CCSE) to develop an adaptive mesh refinement capability for numerical simulation of carbon sequestration and porous media flow.

- Jihan Kim is helping chemists from the University of California, Berkeley accelerate their chemistry codes on NERSC supercomputers.

- Wangyi (Bobby) Liu (starting in 2011) is working with the Neutralized Drift Compression Experiment (NDCX) group at Berkeley Lab to develop a surface tension model within the current warm dense matter simulation framework—new tools for understanding inertial confinement fusion.

- Filipe Maia is working with researchers at Berkeley Lab's Advanced Light Source to accelerate biological imaging codes and with the Earth Sciences Division to accelerate geophysical imaging codes with graphics processing units (GPUs).

- Praveen Narayanan is working with CCSE researchers to run numerical experiments using high performance combustion solvers, and also working with tools to analyze performance of HPC codes slated for codesign at extremely high concurrency.

- Xuefei (Rebecca) Yuan (starting in 2011) is working with researchers in Berkeley Lab's Computational Research Division (CRD) to improve a hybrid linear software package and with physicists at Princeton Plasma Physics Laboratory (PPPL) to accelerate magnetohydrodynamic codes with GPUs.

- Robert Preissl works on advanced programming models for multicore systems and is applying this to physics codes from PPPL to optimize and analyze parallel magnetic fusion simulations.

Two of the "petascale postdocs," Robert Preissl and Jihan Kim, along with project leader Alice Koniges and collaborators from other institutions, were co-authors of the Best Paper winner at the Cray User Group meeting in Edinburgh, Scotland, in May 2010.[4] Titled "Application Acceleration on Current and Future

Cray Platforms," the paper described current bottlenecks and performance improvement areas for applications including plasma physics, chemistry related to carbon capture and sequestration, and material science. The paper discussed a variety of methods including advanced hybrid parallelization using multi-threaded MPI, GPU acceleration, and autoparallelization compilers.

GPU Testbed and Evaluation

Graphics processing units (GPUs) are becoming more widely used for computational science applications, and GPU testbeds have been requested at the requirements workshops NERSC has been conducting with its users. GPUs can offer energy-efficient performance boosts to traditional processors, since they contain massive numbers of simple processors, which are more energy-efficient than a smaller number of larger processors. They also are available at reasonable cost, since they are already being mass-produced for video gaming. The question is whether GPUs offer an effective solution for a broad scientific workload or for a more limited class of computations.

In April 2010 NERSC fielded a 50-node GPU cluster called Dirac in collaboration with the Computational Research Division at Berkeley Lab, with funding from the DOE ASCR Computer Science Research Testbeds program. Each node contains two InfiniBand-connected Intel Nehalem quad-core chips and one NVIDIA Fermi GPU (except for two cores that have four GPUs each). The system includes a complete set of development tools for hybrid computing, including CUDA, OpenCL, and PGI GPGPU-targeted compilers.

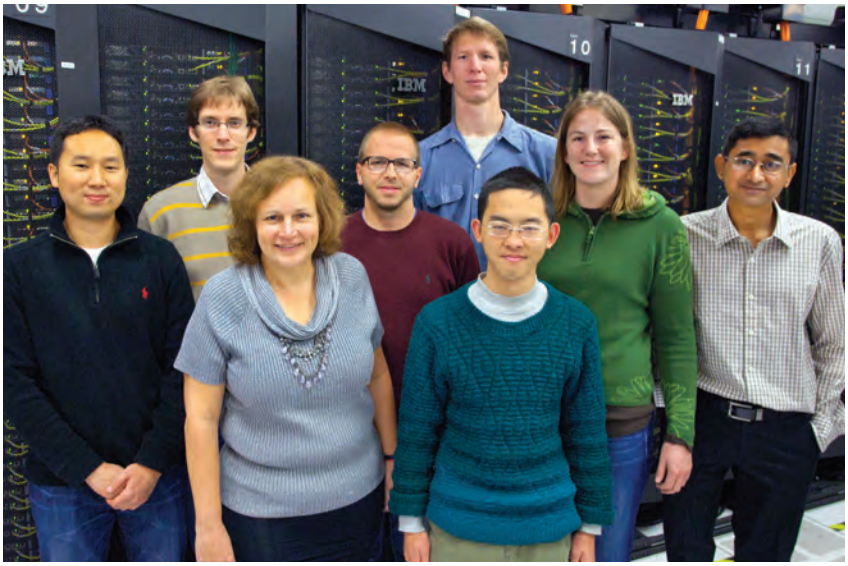

Figure 9. Alice Koniges (third from left) of NERSC's Advanced Technologies Group leads the Computational Science and Engineering Petascale Initiative, which pairs postdoctoral researchers with high-impact projects at NERSC. The postdocs are (from left) Jihan Kim, Filipe Maia, Robert Preissl, Brian Austin (back), Wangyi (Bobby) Liu (front), Kirsten Fagnan, and Praveen Narayanan. Not shown: Xuefei (Rebecca) Yuan. *Photo: Roy Kaltschmidt, Berkeley Lab Public Affairs.*

The Dirac testbed provides the opportunity to engage with the NERSC user community to answer the following questions:

- What parts of the NERSC workload will benefit from GPU acceleration?
- What portions of the workload see no benefit, or insufficient benefit to justify the investment?
- Do GPUs represent the future of HPC platforms, or a feature that will be beneficial to a subset of the community?

Porting an application to a GPU requires reworking the code, such as deciding which parts to run on a GPU and then figuring out how to best modify the code. For example, running an application efficiently on a GPU often requires keeping the data near the device to reduce the computing time taken up with moving data from the CPU or memory to the GPU. Also, a programmer must decide how to thread results from the GPU back into the CPU program. This is not a trivial process, so the Dirac web site provides users with some tips for porting applications to this system.

Over 100 NERSC users have requested and been given access to the Dirac cluster. With their help, we have collected a broad range of

[4] Alice Koniges, Robert Preissl, Jihan Kim, David Eder, Aaron Fisher, Nathan Masters, Velimir Mlaker, Stephane Ethier, Weixing Wang, Martin Head-Gordon, and Nathan Wichmann, "Application Acceleration on Current and Future Cray Platforms," CUG 2010, the Cray User Group meeting, Edinburgh, Scotland, May 2010.

preliminary performance data that support the idea that GPUs in their current configuration will be an important way to augment the effectiveness of NERSC systems for a subset of the user base.

Perhaps the largest collaboration using Dirac is the International Center for Computational Science (ICCS), a collaboration between Berkeley Lab, the University of California–Berkeley, NERSC, the University of Heidelberg, the National Astronomical Observatories of the Chinese Academy of Science, and Nagasaki University. ICCS explores emerging hardware devices, programming models, and algorithm techniques to develop effective and optimal tools for enabling scientific discovery and growth.

ICCS researchers have used Dirac and other GPU clusters for rigorous testing of an adaptive mesh refinement code called GAMER used for astrophysics and other science domains, and the *N*-body code phi-GPU, which is used to simulate dense star clusters with many binaries and galactic nuclei with supermassive black holes, in which correlations between distant particles cannot be neglected. The results played a significant role in developing metrics for workloads and respective speedups; and the phi-GPU team won the 2011 PRACE (Partnership for Advanced Computing in Europe) Award for the paper "Astrophysical Particle Simulations with Large Custom GPU Clusters on Three Continents."[5] In this paper, they present direct astrophysical *N*-body simulations with up to 6 million bodies using the parallel MPI-CUDA code phi-GPU on large GPUclusters in Beijing, Berkeley, and Heidelberg, with different kinds of GPU hardware. They reached about one-third of the peak performance for this code in a real application scenario with hierarchically blocked timesteps and a core-halo density structure of a stellar system.

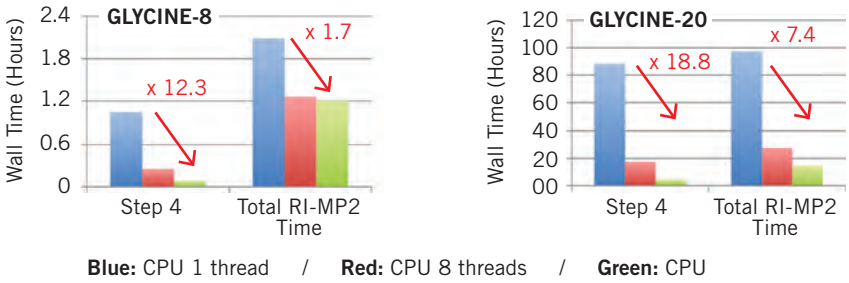

Figure 10. Q-Chem CPU/GPU performance comparisons for two proteins. CPU multi-thread performance is similar to GPU performance for smaller glycine molecules.

Other research results from Dirac include:

- NERSC postdoc Jihan Kim has found impressive GPU versus single-node CPU speedups for the Q-Chem application, but performance varies significantly with the input structure (Figure 10).

- A study by researchers from Berkeley Lab's Computational Research Division, NERSC, UC Berkeley, and Princeton Plasma Physics Laboratory demonstrated the challenges of achieving high GPU performance for key kernels from a gyrokinetic fusion application due to irregular data access, fine-grained data hazards, and high memory requirements.[6]

- Researchers from Argonne National Laboratory, testing coupled cluster equations for electronic correlation in small-to medium-sized molecules, found that the GPU-accelerated algorithm readily achieves a factor of 4 to 5 speedup relative to the multi-threaded CPU algorithm on same-generation hardware.[7]

Improving User Productivity

With over 400 projects and 700 applications, NERSC is constantly looking for ways to improve the productivity of all 4000 scientific users. NERSC closely collaborates with scientific users to improve their productivity and the performance of their applications. Additionally, because of NERSC's large number

[5] R. Spurzem, P. Berczik, T. Hamada, K. Nitadori, G. Marcus, A. Kugel, R. Manner, I. Berentzen, J. Fiestas, R. Banerjee, and R. Klessen, "Astrophysical Particle Simulations with Large Custom GPU Clusters on Three Continents," in press.
[6] Kamesh Madduri, Eun-Jin Im, Khaled Ibrahim, Samuel Williams, Stéphane Ethier, and Leonid Oliker, "Gyrokinetic particle-in-cell optimization on emerging multi- and manycore platforms," Parallel Computing Journal, in press.
[7] A. E. DePrince and J. R. Hammond, "Coupled Cluster Theory on Graphics Processing Units I. The Coupled Cluster Doubles Method," Journal of Chemical Theory and Computation, in press.

of users, we are always looking for collaborations with users where the results can be leveraged and applied to many applications. A performance optimization for one code is always useful, but extrapolating that optimization and applying it to a wider user base produces performance improvements on a much broader scale.

In addition to one-on-one deep collaborations, NERSC also provides rapid responses and advice to the hundreds of user requests received each month. Some simpler questions are answered immediately. However, users also call on the NERSC technical staff's deep knowledge of HPC to cut through complex issues that can otherwise hold up the user's research for months. This section describes a few of our recent collaborations with NERSC users.

Expanding the Science Gateway Infrastructure

NERSC is continuing to work with users to build Science Gateways that allow scientists to access data, perform computations, and interact with NERSC resources using web-based interfaces and technologies. The goal is to make it easier for scientists to use NERSC while creating collaborative tools for sharing data with the rest of the scientific community.

NERSC engages with science teams interested in using these new services, assists with deployment, accepts feedback, and tries to recycle successful approaches into methods that other teams can use. Science Gateways can be configured to provide public unauthenticated access to data sets and services as well as authenticated access if needed.

To help users create Science Gateways, we have developed the NERSC Web Toolkit (NEWT), a service that allows users to access computing resources at NERSC through a simple representational state transfer (REST)-based application programming interface (API). The NEWT API and web service let researchers interact with the NERSC center through simple HTTP URLs and commands. This makes it very easy to build powerful web applications using nothing but HTML and Javascript, providing services such as authentication, file management, job submission, and accounting interfaces. A number of other web access and development tools are also available.

New science gateways in 2010 include:

- **The 20th Century Reanalysis Project Ensemble Gateway.** This site provides data from the 20th Century Reanalysis Project, offering temperature, pressure, humidity, and wind predictions in 200 km sections all around the earth from 1871 to 2008, every six hours, based on historical data. The ensemble mean and standard deviation for each value were calculated over a set of 56 simulations. (This was the first gateway built with NEWT.)

- **The Coherent X-Ray Imaging Data Bank (CXIDB).** This website is dedicated to archiving and sharing data obtained in coherent x-ray imaging (CXI) experiments. The website also serves as the reference for the CXI file format, in which most of the experimental data on the database is stored.

- **The Daya Bay Neutrino Detector Gateway.** The Daya Bay Neutrino Experiment is an international neutrino-oscillation experiment designed to measure the mixing angle θ_{13} using anti-neutrinos produced by the reactors of the Daya Bay and Ling Ao nuclear power plants in China.

- **The Earth System Grid (ESG) Climate Gateway.** The SciDAC-funded Earth System Grid integrates supercomputers with large-scale data and analysis servers located at numerous national labs and research centers to create a powerful environment for next generation climate research. Access to ESG is provided through a system of federated data gateways that collectively allow access to massive data and services for global and regional climate models, IPCC research, and analysis and visualization software.

Speeding Up Turnaround Time for the Palomar Transient Factory

When astronomers search for supernovae, time is of the essence. These spectacular exploding stars don't flare for long, and in order to catch them at their brightest, when they can offer the most information, supernovae must be detected early.

The NERSC Analytics Group recently retooled the workflow of the Palomar Transient Factory (PTF) to speed up the turnaround time for identifying objects such as supernovae from 4 hours to 15 minutes. In addition, the group has helped speed identification of supernova candidates by crowd-sourcing some of the work through "citizen scientists" who log into the Galaxy Zoo Supernovae website (http://supernova.galaxyzoo.org/).

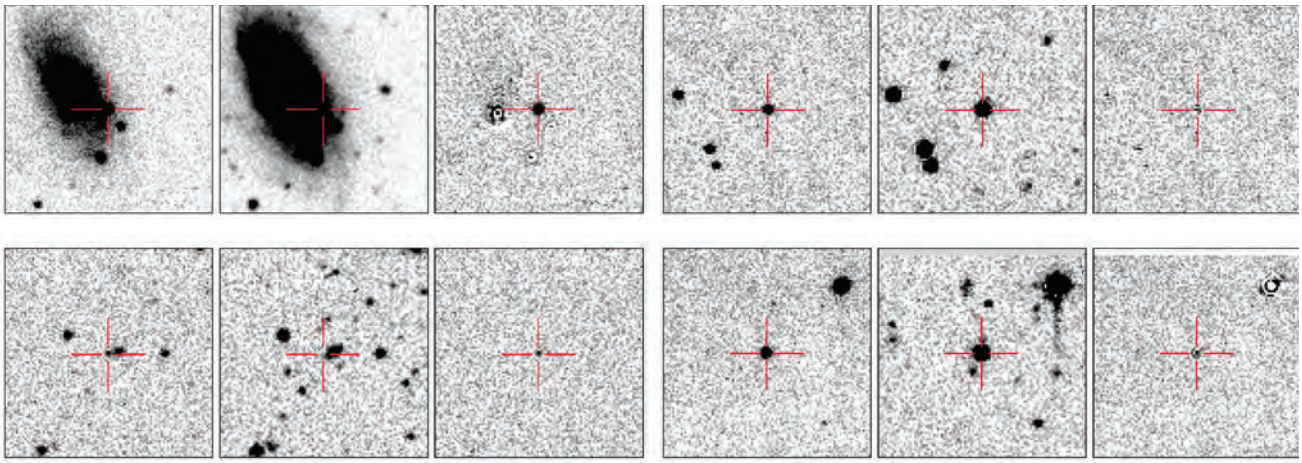

Figure 11. Four example detection triplets from PTF, similar to those uploaded to Galaxy Zoo Supernovae. Each image is 100 arcseconds on a side. In each triplet, the panels show (from left to right) the most recent image containing the candidate supernova light (the science image), the reference image generated from data from an earlier epoch with no supernova candidate light, and the subtraction or difference image—the science image minus the reference image—with the supernova candidate at the centre of the crosshairs. The two triplets on the left are real supernovae, and were highly scored by the Zoo; the triplets on the right are not supernovae and were given the lowest possible score. *Image: A. M. Smith et al., 2011.*

The PTF, which started operating in 2008, is a consortium of universities and telescope sites that image the sky every night. The resulting files are sent to NERSC, where software running on the Carver system compares each night's images to composites gathered over the last decade, a process called subtraction. Images in which stars appear to have moved or changed brightness are flagged for human follow-up. Some follow-up is done by scientists but recently the PTF also enlisted volunteers through Galaxyzoo.org to sort stray airplane lights from supernovae. Once a good candidate is identified, one telescope in the consortium is assigned to follow up.

For example, one group of astronomers has embarked on an ambitious Type 1a supernova search using the PTF to search for supernovae in nearby galaxies along with the Hubble Space Telescope. In order to get the Hubble to observe a "target of opportunity," in this case a nearby supernova burning at peak brightness, astronomers must find the exploding star, do follow-up observations to confirm that they are indeed looking at a supernova, then alert the Hubble—all within the first seven or eight days of onset. The PTF sent 30 such target of opportunity triggers to Hubble for that program, and it inspired a new program in which Hubble observes PTF supernova targets starting more than two weeks before peak brightness and follows them for a month.

In a proof of concept for the Galaxy Zoo Supernovae project, from April to July 2010, nearly 14,000 supernova candidates from PTF were classified by more than 2,500 volunteers within a few hours of data collection (Figure 11). The citizen scientists successfully identified as transients 93% of the ~130 spectroscopically confirmed supernovae that PTF located during the trial period, with no false positive identifications. This trial run shows that the Galaxy Zoo Supernovae model can in principle be applied to any future imaging survey, and the resulting data can also be used to improve the accuracy of automated machine-learning transient classifiers.[8]

Speeding the PTF workflow with a high-speed link and running the database on the NERSC Science Gateways in parallel with the Carver subtraction pipeline has enabled the PTF to become the first survey to target astrophysical transients in real time on a massive scale. In its first two and a half years, PTF has discovered over 1,100 supernovae.

NERSC and JGI Join Forces for Genomics Computing

A torrent of data has been flowing from the advanced sequencing

[8] A. M. Smith et al., "Galaxy Zoo Supernovae," Monthly Notices of the Royal Astronomical Society **412**, 1309 (2011).

platforms at the Department of Energy Joint Genome Institute (DOE JGI), among the world's leading generators of DNA sequence information for bioenergy and environmental applications. In 2010, JGI generated more than 5 terabases (5 trillion nucleotide letters of genetic code) from its plant, fungal, microbial, and metagenome (microbial community) user programs—five times more information than JGI generated in 2009 (Figure 12). On December 14, 2010, the Energy Sciences Network (ESnet) tweeted that JGI had just sent more than 50 terabits of genomics data in the past 10 hours at the rate of nearly 10 gigabits per second. JGI is now the largest single user of ESnet, surpassing even the Large Hadron Collider.

To ensure that there is a robust computational infrastructure for managing, storing, and gleaning scientific insights from this ever-growing flood of data, JGI has joined forces with NERSC. Genomics computing systems are now split between JGI's campus in Walnut Creek, California, and NERSC's Oakland Scientific Facility, which are 20 miles apart. NERSC also manages JGI's six-person systems staff to integrate, operate, and support these systems. NERSC performs this work on a cost-recovery basis similar to the way it manages and supports the Parallel Distributed Systems Facility (PDSF) for high energy and nuclear physics research.

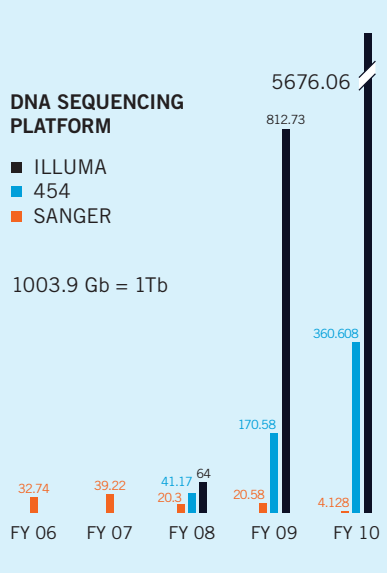

Figure 12. DOE JGI sequence output in billions of nucleotide bases (gigabases, Gb), 2006–2010.

"We evaluated a wide variety of options, and after a thorough review, it made perfect sense to partner with NERSC," said Vito Mangiardi, JGI's Deputy Director for Business Operations and Production. "With a successful track record in providing these kinds of services, they don't have the steep learning curve to climb."

"This is a great partnership, because data-centric computing is an important future direction for NERSC, and genomics is seeing exponential increases in computation and storage," said NERSC Division Director Kathy Yelick. "The computing requirements for the genomics community are quite different from NERSC's more traditional workload. The science is heavy in data analytics, and JGI runs web portals providing access to genomic information."

"These are critically important assets that we can now bring to bear on some of the most complex questions in biology," said JGI Director Eddy Rubin. "We really need the massive computation 'horsepower' and supporting infrastructure that NERSC offers to help us advance our understanding of the carbon cycle and many of the other biogeochemical processes in which microbes play a starring role and that the JGI is characterizing."

Data is currently flowing between JGI and NERSC over ESnet's Science Data Network (SDN). The SDN provides circuit-oriented services to enable the rapid movement of massive scientific data sets, and to support innovative computing architectures such as those being deployed at NERSC in support of JGI science.

In March 2010, JGI became an early user of NERSC's Magellan Cloud Computing System when the Institute had a sudden need for increased computing resources. In less than three days, NERSC and JGI staff provisioned and configured hundreds of processor cores on Magellan to match the computing environment available on JGI's local compute clusters. At the same time, staff at both centers collaborated with ESnet network engineers to deploy a dedicated 9 gigabit per second virtual circuit between Magellan and JGI over SDN within 24 hours.

This strategy gives JGI researchers around the world increased computational capacity without any change to their software or workflow. JGI users still log on to the Institute's network and submit scientific computing jobs to its batch queues, managed by hardware located in Walnut Creek. Once the jobs reach the front of the queue, the information travels 20 miles on reserved SDN bandwidth, directly to NERSC's Magellan system in Oakland. After a job has finished, the results are sent back to Walnut Creek on the SDN within milliseconds to be saved on filesystems at JGI.

"What makes this use of cloud computing so attractive is that JGI users do not notice a difference between computing on Magellan, which is 20 miles away at NERSC, or on JGI's computing environment in Walnut Creek," says Jeff Broughton, who heads NERSC's Systems Department.

Improving VisIt Performance for AMR Applications

Applications using adaptive mesh refinement (AMR) are an important component of the DOE scientific workload. VisIt, a 3D visualization tool, is widely used in data analysis and feature rendering of AMR data.

When Tomasz Plewa of Florida State University, a user of the FLASH AMR code, reported to Hank Childs of the NERSC Analytics Group that VisIt was performing very slowly, Childs found that VisIt's representations of AMR data structures were not well designed for very large numbers of patches. Childs spent several weeks adding new infrastructure to VisIt to represent AMR data more efficiently. He then modified VisIt's FLASH reader (in consultation with the VisIt-FLASH team at Argonne National Laboratory) to use this new infrastructure. As a result, VisIt brought up a plot of Plewa's data in only 13 seconds—40 times faster than the 8 minutes, 38 seconds previously required. The new performance time is reasonable considering the size of the data with respect to I/O performance.

The AMR tuning proved beneficial to other NERSC users as well. Haitao Ma and Stan Woosley of the University of California, Santa Cruz, and the SciDAC Computational Astrophysics Consortium called on Childs to create an animation for the annual SciDAC conference's Visualization Night ("Vis Night"). Their AMR data also had a large number of patches and thus was running slowly in VisIt. To address this, Childs modified the BoxLib reader to use the new AMR infrastructure, which allowed the data to be processed much more quickly. The movie was finished in time for the Vis Night contest, where it ultimately received one of the People's Choice awards (Figure 13). All of the changes (AMR infrastructure, FLASH, and BoxLib readers) have been committed to the VisIt source code repository and are now available to the entire VisIt user community.

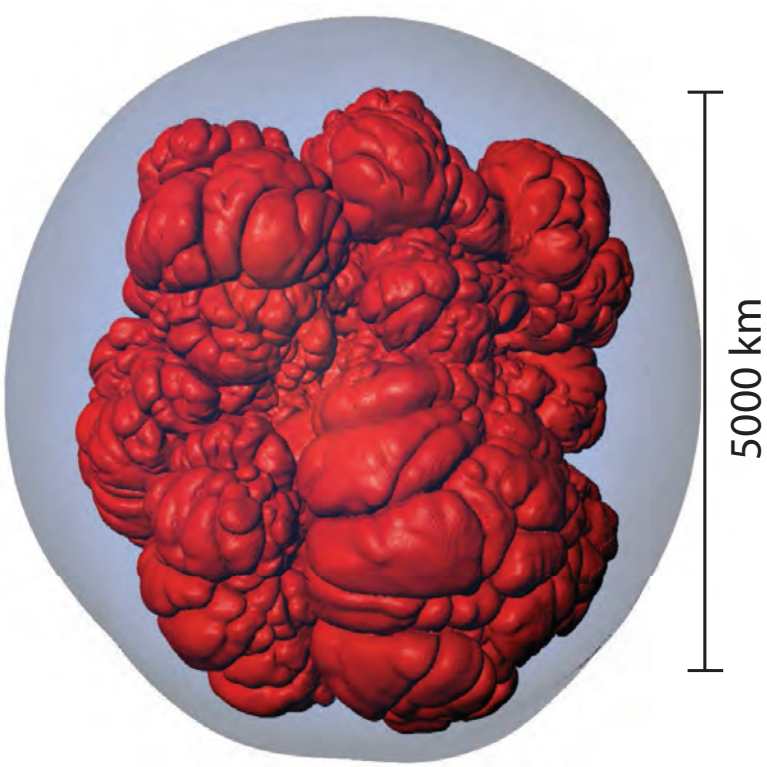

Figure 13. Flame front image from the video "Type 1a Supernova: Turbulent Combustion at the Grandest Scale" by Haitao Ma and Stan Woosley of UC Santa Cruz with John Bell, Ann Almgren, Andy Nonaka, and Hank Childs of Berkeley Lab.

Optimizing HPSS Transfers

Randall Hamper from Indiana University contacted NERSC because he needed to access data from the HPSS storage system and transfer it to his home university. Hamper collaborates with the Nearby Supernova Factory project sponsored by the Office of High Energy Physics, and he needed to access 1.6 million small files. A number of the files were stored as one group on the storage system, but many others were scattered around on various storage tapes.

Creating a simple list of files and attempting to access the files in the listed order resulted in the same tape media having to be mounted more than once. Because mounting the tape is the slowest part of retrieving a small file, possibly taking a minute or two, file access times crawled to a near halt. These numerous small requests also created contention on the system for other NERSC users attempting to access their own datasets.

Members of the HPSS storage team created a script that queries the HPSS metadata to determine which tape volume it belongs to. Then the files were grouped according to tape volume. This change improved data retrievals from HPSS by orders of magnitude. Additionally, this collaboration created a reusable technology, the script for grouping files. It has been distributed to a number of different users to help them retrieve data more optimally.

Uncovering Job Failures at Scale

Debugging applications at scale is essentially detective work. The issues only occur over a certain concurrency and are often tightly coupled to a system's architecture. Wexing Wang from the Princeton Plasma Physics Laboratory contacted NERSC consultants because his application was failing more than half the time when he ran on more than 8,000 processor cores. The cause of his problem was not obvious. The simulation did not fail in the same part of the code consistently. Sometimes the simulation would run for two hours before failing; other times it ran for 18 hours, though it rarely ran to completion successfully. Furthermore, the application code, GTS, had been known to scale to very large concurrencies on other systems.

NERSC put together a team of consultants, systems staff, and Cray staff to troubleshoot the issue. After much detective work—searching for unhealthy nodes in the system, examining the code's communication pattern and memory use—the I/O was analyzed, and it was discovered that the application was not doing I/O optimally. It was using a one-file-per-processor model to write out 20–30 MB of data per core on over 8,000 processors, roughly 240 GB per dump. A Lustre striping parameter was set to 32 disks, and with the user writing output frequently, it meant 7.6 million file segments were being created every 4 minutes. This large number of I/O transactions was flooding the system and causing the job to fail.

Once the problem was identified, the solution was straightforward. We recommended decreasing the striping parameter and reducing the restart frequency to every four hours. Smaller analysis output files continued to be produced every four minutes. Tests showed that making these two changes allowed the code to run successfully to completion.

Improving Network Transfer Rates

A user working with the Planck project, an international consortium studying the cosmic microwave background radiation, contacted the consulting group because he was experiencing slow data transfer rates between NERSC and the Leibniz Supercomputing Centre (LRZ) in Germany. Using the gridFTP tool, he only saw rates of about 5 MB/second, far too slow to transfer large observational astronomy images in a reasonable amount of time.

NERSC made recommendations to increase his TCP and gridFTP buffer settings, which increased performance from 5 to 30 MB/sec. Next, a NERSC networking engineer traced a packet from NERSC to LRZ and noticed that a router along the way was not enabled for jumbo frames, creating a bottleneck along the path between the two sites. Finding the administrator responsible for a specific router in Germany is not easy. However, with the help of staff from the LRZ, the administrator was found, and jumbo frames were enabled on the target router. This increased the transfer rate to 65 MB/sec for the Planck project user. Catching this misconfigured router will not only improve performance transfer rates between NERSC and LRZ for all users, but will improve performance for other DOE collaborators transferring data along this link.

Finally, NERSC gave advice to LRZ staff on reconfiguring their firewall. This last improvement increased transfer rate performance to 110 MB/sec, for an overall improvement of over 18 times the original speed, as shown in Figure 14.

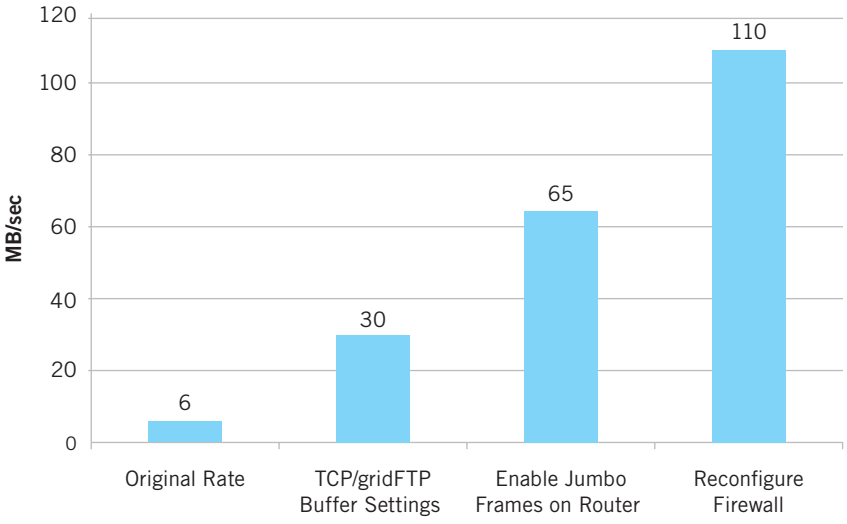

Figure 14. Transfer rates between NERSC and LRZ were boosted 18-fold.

Software Support

NERSC pre-compiles and supports dozens of software packages for our users, comprising over 13.5 million lines of source code. These software packages include debugging tools, analysis and post-processing software, and third-party applications that scientists use to run their primary simulations. Instead of each user installing his or her own version of a third-party software package, NERSC installs and supports optimized versions of many applications, allowing users to focus on their science research rather than on installing and supporting software.

In the past few years, NERSC has seen an increase in the number of scientists (currently 400) using these third-party applications. This is primarily due to the relative increase in the proportion of materials science and chemistry time allocated to NERSC projects and the popularity of third-party applications in these science areas (Tables 1 and 2).

Maintaining User Satisfaction

The annual NERSC user survey responses provide feedback about every aspect of NERSC's operation, help us judge the quality of our services, give DOE information on how well NERSC is doing, and point us to areas we can improve. Three hundred ninety-five users responded to the 2009/2010 survey, representing 11 percent of authorized users and 67 percent of total MPP usage. The respondents represent all six DOE science offices and a variety of home institutions. On a long list of questions, users rated us on a 7-point satisfaction scale (Table 3).

On the open-ended question "What does NERSC do well?" 132 users responded. Some representative comments are:

User support is fantastic—timely and knowledgeable, including follow-up service. New machines are installed often, and they are state-of-the-art. The queues are crowded but fairly managed.

Website is first class, especially the clear instructions for compiling and running jobs.

The account allocation process is very fast and efficient.

NERSC has proven extremely effective for running high resolution models of the earth's climate that require a large number of processors. Without NERSC I never would have been able to run these simulations at such a high resolution to predict future climate. Many thanks.

On the question "What can NERSC do to make you more productive?" 105 users responded. The top areas of concern were long queue turnaround times, the need for more computing resources, queue policies, and software support. Some of the comments from this section are:

There are a lot of users (it is good to be useful and popular), and the price of that success is long queues that lead to slower turn-around. A long-standing problem with no easy answer.

Allow longer jobs (such as one month or half year) on Carver and Hopper. Let science run to the course.

Provide an interface which can help the user determine which of the NERSC machines is more appropriate at a given time for running a job based on the number of processors and runtime that are requested.

I could possibly use some more web-based tutorials on various topics: MPI programming, data analysis with NERSC tools, a tutorial on getting VisIt (visualization tool) to work on my Linux machine.

Every year we institute changes based on the previous year's survey. In 2010

Table 1
Pre-Compiled Materials Science Codes at NERSC

Codes	Hopper	Franklin	Carver
ABINIT	•	•	•
CP2K	•	•	•
CPMD		•	
Quantum Espresso	•	•	•
LAMMPS	•	•	•
Qbox		•	
SIESTA	•	•	•
VASP	•	•	•
WEIN 2K			•

Table 2
Pre-Compiled Chemistry Codes at NERSC

Codes	Hopper	Franklin	Carver
AMBER		•	•
G09			•
GAMESS	•	•	•
GROMACS	•	•	•
MOLPRO	•	•	•
NAMD	•	•	•
NWChem	•	•	•
Q-Chem		•	•

Table 3
2009/2010 User Survey Satisfaction Scale

Satisfaction Score	Meaning	Number of Times Selected
7	Very Satisfied	8,053
6	Mostly Satisfied	6,219
5	Somewhat Satisfied	1,488
4	Neutral	1,032
3	Somewhat Dissatisfied	366
2	Mostly Dissatisfied	100
1	Very Dissatisfied	88

NERSC took a number of actions in response to suggestions from the 2008/2009 user survey, including improvements to the Franklin system and the NERSC web site.

On the 2008/2009 survey, Franklin uptime received the second lowest average score (4.91). In the first half of 2009, Franklin underwent an intensive stabilization period. Tiger teams were formed with close collaborations with Cray to address system instability. These efforts were continued in the second half of 2009 and throughout 2010, when NERSC engaged in a project to understand system-initiated causes of hung jobs, and to implement corrective actions to reduce their number. These investigations revealed bugs in the Seastar interconnect as well as in the Lustre file system. These bugs were reported to Cray and were fixed in March 2010, when Franklin was upgraded to Cray Linux Environment 2.2.i. As a result, Franklin's mean time between failures improved from a low of about 3 days in 2008 to 9 days in 2010.

On the 2010 survey, Franklin uptime received an average score of 5.99, a significant increase of 1.08 points over the previous year. Two other Franklin scores (overall satisfaction, and disk configuration and I/O performance) were significantly improved as well. Another indication of increased satisfaction with Franklin is that on the 2009 survey 40 users requested improvements in Franklin uptime or performance, whereas only 10 made such requests on the 2010 survey.

On the 2008/2009 survey, ten users requested improvements to the NERSC web site. In response, User Services staff removed older documentation and made sure that the remaining documentation was up to date. On the 2010 survey, the score for "ease of finding information on the NERSC web site" was significantly improved. Also, for the medium-scale MPP users, the scores for the web site overall and for the accuracy of information on the web showed significant improvement.

The complete survey results are available at https://www.nersc.gov/news-publications/publications-reports/user-surveys/2009-2010-user-survey-results/.

Leadership Changes

In September 2010, Associate Laboratory Director for Computing Sciences (and former NERSC Division Director) Horst Simon was named Deputy Director of Berkeley Lab. "Horst is a strong leader who has helped to lead a tremendously productive program in high performance computing that is world-class," said Berkeley Lab Director Paul Alivisatos. "As Deputy Director he'll help me lead major scientific initiatives, oversee strategic research investments, and maintain the intellectual vitality of Berkeley Lab."

A few days later, Berkeley Lab announced that Kathy Yelick was named Associate Lab Director for Computing Sciences. Yelick has been the director of the NERSC Division since 2008, a position she will continue to hold. "I am very pleased to see Kathy Yelick assume this critical role in senior leadership as Associate Lab Director for Computing Sciences," said Berkeley Lab Director Paul Alivisatos. "We will benefit immensely from her knowledge and vision. She will have tremendous opportunity to use Berkeley Lab's strengths in scientific computing and advanced networking to accelerate research that addresses some of our nation's most urgent scientific challenges."

In August 2010 Jonathan Carter, who had led NERSC's User Services Group since 2005, was named the new Computing Sciences Deputy, succeeding Michael Banda, who joined Berkeley Lab's Advanced Light Source. Carter was one of the first new employees hired when NERSC moved to Berkeley Lab in 1996, and had most recently led the NERSC-6 procurement team that selected the Cray XE6 Hopper system.

Succeeding Carter as leader of the User Services Group was Katie Antypas, whom Kathy Yelick describes as "a passionate advocate for the NERSC users." Antypas was a member of the NERSC-6 procurement team, co-lead ofthe NERSC-6 implementation team, and co-lead on a number of Franklin stabilization "tiger teams," as well as participating in various planning activities.

Energy efficiency has become a major issue for NERSC, with costs growing with each new system and the need for facility upgrades likely every few years. This requires increased monitoring, reporting, and optimization of energy use, as well as management of the associated risks. Rather than distributing the management of these activities across multiple individuals, NERSC Director Kathy Yelick created a new position of Energy Efficiency and Risk Manager and named Jim Craw to this new role. "Jim's extensive knowledge of NERSC's systems, operations and vendor relations will be invaluable to support the expanded emphasis on these matters," Yelick said. Craw had been leader of NERSC's Computational Systems Group; Systems Department Head Jeff Broughton is acting leader of the group until a permanent replacement is hired.

Outreach and Training

As science tackles increasingly complex problems, interdisciplinary research becomes more essential, because interdisciplinary teams bring a variety of knowledge, approaches, and skills to solve the problem. Computational science is no exception. That's why outreach and training, which promote diversity in the computational science community, are essential to NERSC's mission.

Collaborations between institutions are one fruitful avenue of outreach. Examples include NERSC's technical exchanges with KISTI (Korea), CSCS (Switzerland), and KAUST (Saudi Arabia); the data transfer working group activities with Oak Ridge and Argonne national labs; procurement and technical exchanges with Sandia; Cray systems administrator interactions with Oak Ridge; and center management interchanges between NERSC, Argonne, and Oak Ridge.

What began as an informal discussion two years ago between a physicist and a computer scientist has led to a new international collaboration aimed at creating computational tools to help scientists make more effective use of new computing technologies, including multicore processors. The International Center for Computational Science (ICCS) is located at Berkeley Lab/NERSC and the University of California at Berkeley, with partners at the University of Heidelberg in Germany, the National Astronomical Observatories of the Chinese Academy of Sciences, and Nagasaki University in Japan. ICCS explores the vast and rapidly changing space of emerging hardware devices, programming models, and algorithm techniques to develop effective and optimal tools for enabling scientific discovery and growth.

ICCS hosts several collaborative research programs, including ISAAC, a three-year NSF-funded project to focus on research and development of infrastructure to accelerate physics and astronomy applications using multicore architectures; and the GRACE Project and the Silk Road Project, which are developing programmable hardware for astrophysical simulations.

Top row, left to right: Horst Simon, Kathy Yelick, Jonathan Carter. **Bottom row, left to right:** Katie Antypas, Jim Craw, Jeff Broughton.

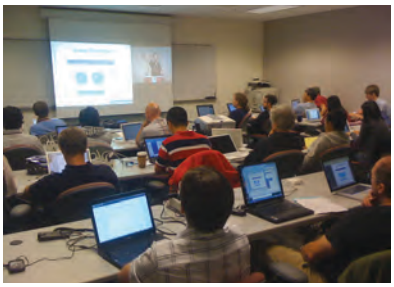

Figure 15. Researchers with scientific backgrounds ranging from astrophysics to biochemistry attended the first computational science summer school at NERSC. *Photo: Hemant Shukla.*

ICCS also offers training in programming multicore devices such as GPUs for advanced scientific computation. In August 2010, ICCS offered a five-day course at NERSC, "Proven Algorithmic Techniques for Many-Core Processors," in collaboration with the Virtual School of Computational Science and Engineering, a project of the University of Illinois and NVIDIA. Participants applied the techniques in hands-on lab sessions using NERSC computers (Figure 15). Seating was limited and many people had to be turned away, so ICCS offered a second five-day workshop at UC Berkeley in January 2011 on "Manycore and Accelerator-Based High-Performance Scientific Computing," which also included hands-on tutorials.

NERSC staff have been actively involved in developing and delivering tutorials at SciDAC meetings, at the SC conference, and for the TeraGrid. One hundred participants attended tutorials at the July 2010 SciDAC conference, which were organized by the SciDAC Outreach Center, which is hosted at NERSC. The SciDAC Outreach Center has provided HPC outreach to industry, identifying computing needs in industrial areas that reinforce the DOE mission with the Council on Competiveness. An example is working with United Technologies Corporation (UTC) on code profiling, HPC scalability, and code tuning, with the result of a 3.6x faster time to solution for one class of problems.

NERSC has also extended its outreach to local high schools. In June 2010, NERSC hosted 12 students from Oakland Technical High School's Computer Science and Technology Academy, a small academy within the larger high school, for an introduction to computational science and supercomputer architecture, and a tour of NERSC (Figure 16). In July, Berkeley Lab Computing Sciences hosted 14 local high school students for a four-day Careers in Computing Camp, with sessions at both the main Lab and NERSC's Oakland Scientific Facility (Figure 17). The program was developed with input from computer science teachers at Berkeley High, Albany High, Kennedy High in Richmond, and Oakland Tech, and included presentations, hands-on activities, and tours of various facilities. Staff from NERSC, CRD, ESnet, and the IT Division presented a wide range of topics including assembling a desktop computer, cyber security war stories, algorithms for combustion and

Figure 16. Dave Paul, a systems engineer, brought out computer nodes and parts from NERSC's older systems to show Oakland Tech students how the components have become both more dense and more power efficient as the technology has evolved over time. *Photo: Roy Kaltschmidt, Berkeley Lab Public Affairs.*

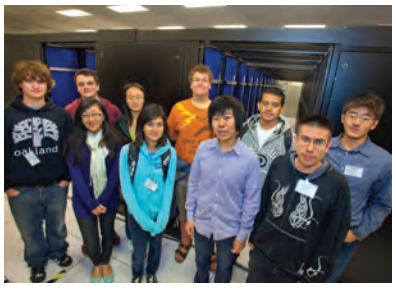

Figure 17. Students from the Berkeley Lab Careers In Computing Camp pose in front of NERSC's Cray XT5 supercomputer, Franklin. *Photo: Roy Kaltschmidt, Berkeley Lab Public Affairs.*

astrophysics, the role of applied math, networking, science portals, and hardware/software co-design.

To reach an even broader audience, NERSC, ESnet, and Computational Research Division staff are contributing to Berkeley Lab's Video Glossary, a popular website where scientific terms are defined in lay language (http://videoglossary.lbl.gov/). Computing-related videos to date include:

- bits and bytes (Inder Monga, ESnet)
- computational science (Associate Lab Director Horst Simon)
- green computing (John Shalf, NERSC)
- IPv6 (Michael Sinatra, ESnet)
- petabytes (Inder Monga, ESnet)
- petaflop computing (David Bailey, CRD)
- quantum computing (Thomas Schenkel, Accelerator and Fusion Research Division)
- science network (Eli Dart, ESnet)
- supercomputing (NERSC Director Kathy Yelick)

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.

In 2010, four conference papers co-authored by NERSC staff won Best Paper awards:

- **MPI-Hybrid Parallelism for Volume Rendering on Large, Multi-Core Systems.** Mark Howison, E. Wes Bethel, and Hank Childs. Proceedings of the Eurographics Symposium on Parallel Graphics and Visualization (EGPGV'10), Norrköping, Sweden, May 2–3, 2010.

- **Application Acceleration on Current and Future Cray Platforms.** Alice Koniges, Robert Preissl, Jihan Kim, David Eder, Aaron Fisher, Nathan Masters, Velimir Mlaker, Stephane Ethier, Weixing Wang, Martin Head-Gordon, and Nathan Wichmann. Proceedings of the 2010 Cray User Group Meeting, Edinburgh, Scotland, May 24–27, 2010.

- **Seeking Supernovae in the Clouds: A Performance Study.** Keith Jackson, Lavanya Ramakrishnan, Karl Runge, and Rollin Thomas. ScienceCloud 2010, the 1st Workshop on Scientific Cloud Computing, Chicago, Ill., June 2010.

- **Performance Analysis of High Performance Computing Applications on the Amazon Web Services Cloud.** Keith R. Jackson, Lavanya Ramakrishnan, Krishna Muriki, Shane Canon, Shreyas Cholia, John Shalf, Harvey J. Wasserman, and Nicholas J. Wright. Proceedings of the 2nd IEEE International Conference on Cloud Computing Technology and Science (CloudCom 2010), Bloomington, Ind., Nov. 30–Dec.1, 2010.

Alice Koniges of NERSC's Advanced Technologies Group and Gabriele Jost of the Texas Advanced Computing Center were guest editors of a special issue of the journal Scientific Programming titled "Exploring Languages for Expressing Medium to Massive On-Chip Parallelism" (Vol. 18, No. 3–4, 2010). Several NERSC researchers co-authored articles in this issue:

- **Guest Editorial: Special Issue: Exploring Languages for Expressing Medium to Massive On-Chip Parallelism.** Gabriele Jost and Alice Koniges.

- **Overlapping Communication with Computation using OpenMP Tasks on the GTS Magnetic Fusion Code.** Robert Preissl, Alice Koniges, Stephan Ethier, Weixing Wang, and Nathan Wichmann.

- **A Programming Model Performance Study using the NAS Parallel Benchmarks.** Hongzhang Shan, Filip Blagojević, Seung-Jai Min, Paul Hargrove, Haoqiang Jin, Karl Fuerlinger, Alice Koniges, and Nicholas J. Wright.

A paper co-authored by Daniela Ushizima of NERSC's Anaytics Group was on two successive lists of the "Top 25 Hottest Articles" based on number of downloads from the journal Digital Signal Processing, from April to September 2010:

- **SAR Imagery Segmentation by Statistical Region Growing and Hierarchical Merging.** E. A. Carvalho, D. M. Ushizima, F. N. S. Medeiros, C. I. O. Martins, R. C. P. Marques, and I. N. S. Oliveira. Digital Signal Processing, Volume 20, Issue 5, pp. 1365–1378, Sept 2010.

Other publications by NERSC staff in 2010 are listed below. (Not all co-authors are from NERSC.) Many of these papers are available online at https://www.nersc.gov/news-publications/publications-reports/nersc-staff-publications/.

ALE-AMR: A New 3D Multi-Physics Code for Modeling Laser/Target Effects. A. E. Koniges, N. D. Masters, A. C. Fisher, R. W. Anderson, D. C. Eder, T. B. Kaiser, D. S. Bailey, B. Gunney, P. Wang, B. Brown, K. Fisher, F. Hansen, B. R. Maddox, D. J. Benson, M. Meyers, and A. Geille. Proceedings of the Sixth International Conference on Inertial Fusion Sciences and Applications, Journal of Physics: Conference Series, Vol. 244 (3), 032019.

An Auto-Tuning Framework for Parallel Multicore Stencil Computations. S. Kamil, C. Chan, L. Oliker, J. Shalf, and S. Williams. 2010 IEEE International Symposium on Parallel & Distributed Processing (IPDPS), pp.1–12, April 19–23, 2010.

An Unusually Fast-Evolving Supernova. Poznanski, Dovi; Chornock, Ryan; Nugent, Peter E.; Bloom, Joshua S.; Filippenko, Alexei V.; Ganeshalingam, Mohan; Leonard, Douglas C.; Li, Weidong; Thomas, Rollin C. Science, Volume 327, Issue 5961, pp. 58– (2010).

Analyzing and Tracking Burning Structures in Lean Premixed Hydrogen Flames. P.-T. Bremer, G. H. Weber, V. Pascucci, M. S. Day, and John B. Bell. IEEE Transactions on Visualization and Computer Graphics, 16(2): 248–260, 2010.

Analyzing the Effect of Different Programming Models upon Performance and Memory Usage on Cray XT5 Platforms. H. Shan, H. Jin, K. Fuerlinger, A. Koniges, and N. Wright. Proceedings of the 2010 Cray User Group, May 24–27, 2010, Edinburgh, Scotland.

Assessment and Mitigation of Radiation, EMP, Debris and Shrapnel Impacts at Mega-Joule Class Laser Facilities. D. C. Eder, et al. J. Physics, Conf. Ser. 244: 0320018, 2010.

Auto-Tuning Stencil Computations on Diverse Multicore Architectures. K. Datta, S. Williams, V. Volkov, J. Carter, L. Oliker, J. Shalf, and K. Yelick. In *Scientific Computing with Multicore and Accelerators*, edited by: Jakub Kurzak, David A. Bader, and Jack Dongarra. CRC Press, 2010.

Automated Detection and Analysis of Particle Beams in Laser-Plasma Accelerator Simulations. Daniela Ushizima, Cameron Geddes, Estelle Cormier-Michel, E. Wes Bethel, Janet Jacobsen, Prabhat, Oliver Rübel, Gunther Weber, Bernard Hamann, Peter Messmer, and Hans Hagen. In-Tech, 367–389, 2010.

Building a Physical Cell Simulation and Comparing with Confocal Microscopy. C. Rycroft, D. M. Ushizima, R. Saye, C. M. Ghajar, J. A. Sethian. Bay Area Physical Sciences–Oncology Center (NCI) Meeting 2010, UCSF Medical Sciences, Sept. 2010.

Communication Requirements and Interconnect Optimization for High-End Scientific Applications. Shoaib Kamil, Leonid Oliker, Ali Pinar, and John Shalf. IEEE Trans. Parallel Distrib. Syst. 21(2): 188–202 (2010).

Compressive Auto-Indexing in Femtosecond Nanocrystallography. Filipe Maia, Chao Yang, and Stefano Marchesini. Lawrence Berkeley National Laboratory technical report LBNL-4008E, 2010.

Compressive Phase Contrast Tomography. Filipe Maia, Alastair MacDowell, Stefano Marchesini, Howard A. Padmore, Dula Y. Parkinson, Jack Pien, Andre Schirotzek, and Chao Yang. SPIE Optics and Photonics 2010, San Diego, California.

Core-Collapse Supernovae from the Palomar Transient Factory: Indications for a Different Population in Dwarf Galaxies. Arcavi, Iair; Gal-Yam, Avishay; Kasliwal, Mansi M.; Quimby, Robert M.; Ofek, Eran O.; Kulkarni, Shrinivas R.; Nugent, Peter E.; Cenko, S. Bradley; Bloom, Joshua S.; Sullivan, Mark; Howell, D. Andrew; Poznanski, Dovi; Filippenko, Alexei V.; Law, Nicholas; Hook, Isobel; Jönsson, Jakob; Blake, Sarah; Cooke, Jeff; Dekany, Richard; Rahmer, Gustavo; Hale, David; Smith, Roger; Zolkower, Jeff; Velur, Viswa; Walters, Richard; Henning, John; Bui, Kahnh; McKenna, Dan; Jacobsen, Janet. Astrophysical Journal, Volume 721, Issue 1, pp. 777–784 (2010).

Coupling Visualization and Data Analysis for Knowledge Discovery from Multi-Dimensional Scientific Data. Oliver Rübel, Sean Ahern, E. Wes Bethel, Mark D. Biggin, Hank Childs, Estelle Cormier-Michel, Angela DePace, Michael B. Eisen, Charless C. Fowlkes, Cameron G. R. Geddes, Hans Hagen, Bernd Hamann, Min-Yu Huang, Soile V. E. Keränen, David W. Knowles, Cris L. Luengo Hendriks, Jitendra Malik, Jeremy Meredith, Peter Messmer, Prabhat, Daniela Ushizima, Gunther H. Weber, and Kesheng Wu. Proceedings of International Conference on Computational Science, ICCS 2010, May 2010.

Defining Future Platform Requirements for E-Science Clouds. Lavanya Ramakrishnan, Keith R. Jackson, Shane Canon, Shreyas Cholia, and John Shalf. Proceedings of the 1st ACM Symposium on Cloud Computing (SoCC '10). ACM, New York, NY, USA, 101–106 (2010).

Discovery of Precursor Luminous Blue Variable Outbursts in Two Recent Optical Transients: The Fitfully Variable Missing Links UGC 2773-OT and SN 2009ip. Smith, Nathan; Miller, Adam; Li, Weidong; Filippenko, Alexei V.; Silverman, Jeffrey M.; Howard, Andrew W.; Nugent, Peter; Marcy, Geoffrey W.; Bloom, Joshua S.; Ghez, Andrea M.; Lu, Jessica; Yelda, Sylvana; Bernstein, Rebecca A.; Colucci, Janet E. Astronomical Journal, Volume 139, Issue 4, pp. 1451–1467 (2010).

Efficient Bulk Data Replication for the Earth System Grid. A. Sim, D. Gunter, V. Natarajan, A. Shoshani, D. Williams, J. Long, J. Hick, J. Lee, and E. Dart. International Symposium on Grid Computing, 2010.

Electromagnetic Counterparts of Compact Object Mergers Powered by the Radioactive Decay of r-Process Nuclei. Metzger, B. D.; Martínez-Pinedo, G.; Darbha, S.; Quataert, E.; Arcones, A.; Kasen, D.; Thomas, R.; Nugent, P.; Panov, I. V.; Zinner, N. T. Monthly Notices of the Royal Astronomical Society, Volume 406, Issue 4, pp. 2650–2662 (2010).

Exploitation of Dynamic Communication Patterns through Static Analysis. Robert Preissl, Bronis R. de Supinski, Martin Schulz, Daniel J. Quinlan, Dieter Kranzlmüller, and Thomas Panas. Proc. International Conference on Parallel Processing (ICPP), Sept. 13–16, 2010.

Extreme Scaling of Production Visualization Software on Diverse Architectures. H. Childs, D. Pugmire, S. Ahern, B. Whitlock, M. Howison, Prabhat, G. Weber, and E. W. Bethel. Computer Graphics and Applications, Vol. 30 (3), pp. 22–31, May/June 2010.

File System Monitoring as a Window into User I/O Requirements. A. Uselton, K. Antypas, D. M. Ushizima, and J. Sukharev. Proceedings of the 2010 Cray User Group Meeting, Edinburgh, Scotland, May 24–27, 2010.

H5hut: A High-Performance I/O Library for Particle-Based Simulations. Mark Howison, Andreas Adelmann, E. Wes Bethel, Achim Gsell, Benedikt Oswald, and Prabhat. Proceedings of IASDS10 — Workshop on Interfaces and Abstractions for Scientific Data Storage, Heraklion, Crete, Greece, September 2010.

High Performance Computing and Storage Requirements for HIgh Energy Physics. Richard Gerber and Harvey Wasserman, eds. Lawrence Berkeley National Laboratory technical report LBNL-4128E, 2010.

HPSS in the Extreme Scale Era: Report to DOE Office of Science on HPSS in 2018–2022. D. Cook, J. Hick, J. Minton, H. Newman, T. Preston, G. Rich, C. Scott, J. Shoopman, J. Noe, J. O'Connell, G. Shipman, D. Watson, and V. White. Lawrence Berkeley National Laboratory technical report LBNL-3877E, 2010.

Hybrid Parallelism for Volume Rendering on Large, Multi-Core Systems. M. Howison, E. W. Bethel, and H. Childs. Journal of Physics: Conference Series, Proceedings of SciDAC 2010, July 2010, Chattanooga TN.

Integrating Data Clustering and Visualization for the Analysis of 3D Gene Expression Data. O. Rübel, G. H. Weber, M-Y Huang, E. W. Bethel, M. D. Biggin, C. C. Fowlkes, C. Luengo Hendriks, S. V. E. Keränen, M. Eisen, D. Knowles, J. Malik, H. Hagen and B. Hamann. IEEE Transactions on Computational Biology and Bioinformatics, Vol. 7, No.1, pp. 64–79, Jan/Mar 2010.

International Exascale Software Project (IESP) Roadmap, version 1.1. Jack Dongarra et al. [incl. John Shalf, David Skinner, and Kathy Yelick]. October 18, 2010, www.exascale.org.

Large Data Visualization on Distributed Memory Multi-GPU Clusters. T. Fogal, H. Childs, S. Shankar, J. Kruger, R. D. Bergeron, and P. Hatcher. High Performance Graphics 2010, Saarbruken, Germany, June 2010.

Laser Ray Tracing in a Parallel Arbitrary Lagrangian Eulerian Adaptive Mesh Refinement Hydrocode. N. Masters, D. C. Eder, T. B. Kaiser, A. E. Koniges, and A. Fisher. J. Physics, Conf. Ser. 244: 032022, 2010.

Lessons Learned from Moving Earth System Grid Data Sets over a 20 Gbps Wide-Area Network. R. Kettimuthu, A. Sim, D. Gunter, W. Allcock, P.T. Bremer, J. Bresnahan, A. Cherry, L. Childers, E. Dart, I. Foster, K. Harms, J. Hick, J. Lee, M. Link, J. Long, K. Miller, V. Natarajan, V. Pascucci, K. Raffenetti, D. Ressman, D. Williams, L. Wilson, and L. Winkler. Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing (HPDC '10). ACM, New York, NY, USA, 316–319.

Magellan: Experiences from a Science Cloud. Lavanya Ramakrishnan, Piotr T. Zbiegel, Scott Campbell, Rick Bradshaw, Richard Shane Canon, Susan Coghlan, Iwona Sakrejda, Narayan Desai, Tina Declerck, and Anping Liu. Proceedings of the 2nd International Workshop on Scientific Cloud Computing (ScienceCloud '11). ACM, New York, NY, USA, 49–58.

Modeling Heat Conduction and Radiation Transport with the Diffusion Equation in NIF ALE-AMR. A. Fisher, D. S. Bailey, T. B. Kaiser, B. T. N. Gunney, N. D. Masters, A. E. Koniges, D. C. Eder, and R. W. Anderson. J. Physics, Conf. Ser. 244: 022075, 2010.

Multi-Level Bitmap Indexes for Flash Memory Storage. Kesheng Wu, Kamesh Madduri, and Shane Canon. IDEAS '10: Proceedings of the Fourteenth International Database Engineering and Applications Symposium, Montreal, QC, Canada, 2010.

Nearby Supernova Factory Observations of SN 2007if: First Total Mass Measurement of a Super-Chandrasekhar-Mass Progenitor. Scalzo, R. A.; Aldering, G.; Antilogus, P.; Aragon, C.; Bailey, S.; Baltay, C.; Bongard, S.; Buton, C.; Childress, M.; Chotard, N.; Copin, Y.; Fakhouri, H. K.; Gal-Yam, A.; Gangler, E.; Hoyer, S.; Kasliwal, M.; Loken, S.; Nugent, P.; Pain, R.; Pécontal, E.; Pereira, R.; Perlmutter, S.; Rabinowitz, D.; Rau, A.; Rigaudier, G.; Runge, K.; Smadja, G.; Tao, C.; Thomas, R. C.; Weaver, B.; Wu, C. Astrophysical Journal, Volume 713, Issue 2, pp. 1073–1094 (2010).

Nickel-Rich Outflows Produced by the Accretion-Induced Collapse of White Dwarfs: Light Curves and Spectra. Darbha, S.; Metzger, B. D.; Quataert, E.; Kasen, D.; Nugent, P.; Thomas, R. Monthly Notices of the Royal Astronomical Society, Volume 409, Issue 2, pp. 846–854 (2010).

Ocular Fundus for Retinopathy Characterization. D. M. Ushizima and J. Cuadros. Bay Area Vision Meeting, Feb. 5, Berkeley, CA, 2010.

Optimizing and Tuning the Fast Multipole Method for State-of-the-Art Multicore Architectures. A. Chandramowlishwaran, S. Williams, L. Oliker, I. Lashuk, G. Biros, and R. Vuduc. Proceedings of 2010 IEEE International Symposium on Parallel & Distributed Processing (IPDPS), April 2010.

Parallel Algorithms for Moving Lagrangian Data on Block Structured Eulerian Meshes. A. Dubey, K. B. Antypas, C. Daley. Parallel Computing, Vol. 37 (2): 101–113 (2011).

Parallel and Streaming Generation of Ghost Data for Structured Grids. M. Isenberg, P. Lindstrom, and H. Childs. Computer Graphics and Applications, Vol. 30 (3), pp. 50–62, May/June 2010.

Parallel I/O Performance: From Events to Ensembles. Andrew Uselton, Mark Howison, Nicholas J. Wright, David Skinner, Noel Keen, John Shalf, Karen L. Karavanic, and Leonid Oliker. Proceedings of 2010 IEEE International Symposium on Parallel & Distributed Processing (IPDPS), 19–23 April 2010; Atlanta, Georgia; pp.1-11.

Performance Analysis of Commodity and Enterprise Class Flash Devices. N. Master, M. Andrews, J. Hick, S. Canon, and N. Wright. Petascale Data Storage Workshop (PDSW) 2010, Nov. 2010.

Petascale Parallelization of the Gyrokinetic Toroidal Code. S. Ethier, M. Adams, J. Carter, and L. Oliker. VECPAR: High Performance Computing for Computational Science, June 2010.

Quantitative Visualization of ChIP-chip Data by Using Linked Views. Min-Yu Huang, Gunther H. Weber, Xiao-Yong Li, Mark D. Biggin, and Bernd Hamann. Proceedings IEEE International Conference on Bioinformatics and Biomedicine 2010 (IEEE BIBM 2010) Workshops, Workshop on Integrative Data Analysis in Systems Biology (IDASB), pp. 195–200.

Rapidly Decaying Supernova 2010X: A Candidate "Ia" Explosion. Kasliwal, Mansi M.; Kulkarni, S. R.; Gal-Yam, Avishay; Yaron, Ofer; Quimby, Robert M.; Ofek, Eran O.; Nugent, Peter; Poznanski, Dovi; Jacobsen, Janet; Sternberg, Assaf; Arcavi, Iair; Howell, D. Andrew; Sullivan, Mark; Rich, Douglas J.; Burke, Paul F.; Brimacombe, Joseph; Milisavljevic, Dan; Fesen, Robert; Bildsten, Lars; Shen, Ken; Cenko, S. Bradley; Bloom, Joshua S.; Hsiao, Eric; Law, Nicholas M.; Gehrels, Neil; Immler, Stefan; Dekany, Richard; Rahmer,

Gustavo; Hale, David; Smith, Roger; Zolkower, Jeff; Velur, Viswa; Walters, Richard; Henning, John; Bui, Kahnh; McKenna, Dan. Astrophysical Journal Letters, Volume 723, Issue 1, pp. L98–L102 (2010).

Recent Advances in VisIt: AMR Streamlines and Query-Driven Visualization. G. H. Weber, S. Ahern, E. W. Bethel, S. Borovikov, H. R. Childs, E. Deines, C. Garth, H. Hagen, B. Hamann, K. I. Joy, D. Martin, J. Meredith, Prabhat, D. Pugmire, O. Rübel, B. Van Straalen, and K. Wu. Numerical Modeling of Space Plasma Flows: Astronum-2009 (Astronomical Society of the Pacific Conference Series), Vol. 429, pp. 329–334, 2010.

Resource Management in the Tessellation Manycore OS. J. A. Colmenares, S. Bird, H. Cook, P. Pearce, D. Zhu, J. Shalf, K. Asanovic, and J. Kubiatowicz. Proceedings of the Second USENIX Workshop on Hot Topics in Parallelism (HotPar'10), Berkeley, California, June 2010.

Retinopathy Diagnosis from Ocular Fundus Image Analysis. D. M. Ushizima and F. Medeiros. Modeling and Analysis of Biomedical Image, SIAM Conference on Imaging Science (IS10), Chicago, Il, April 12–14, 2010.

SDWFS-MT-1: A Self-Obscured Luminous Supernova at z ~= 0.2. Kozłowski, Szymon; Kochanek, C. S.; Stern, D.; Prieto, J. L.; Stanek, K. Z.; Thompson, T. A.; Assef, R. J.; Drake, A. J.; Szczygieł, D. M.; Woźniak, P. R.; Nugent, P.; Ashby, M. L. N.; Beshore, E.; Brown, M. J. I.; Dey, Arjun; Griffith, R.; Harrison, F.; Jannuzi, B. T.; Larson, S.; Madsen, K.; Pilecki, B.; Pojmański, G.; Skowron, J.; Vestrand, W. T.; Wren, J. A. Astrophysical Journal, Volume 722, Issue 2, pp. 1624–1632 (2010).

Silicon Nanophotonic Network-On-Chip Using TDM Arbitration. Gilbert Hendry, Johnnie Chan, Shoaib Kamil, Lenny Oliker, John Shalf, Luca P. Carloni, and Keren Bergman. IEEE Symposium on High Performance Interconnects (HOTI) 5.1 (Aug. 2010).

SN 2008iy: An Unusual Type IIn Supernova with an Enduring 400-d Rise Time. Miller, A. A.; Silverman, J. M.; Butler, N. R.; Bloom, J. S.; Chornock, R.; Filippenko, A. V.; Ganeshalingam, M.; Klein, C. R.; Li, W.; Nugent, P. E.; Smith, N.; Steele, T. N. Monthly Notices of the Royal Astronomical Society, Volume 404, Issue 1, pp. 305–317 (2010).

Spectra and Hubble Space Telescope Light Curves of Six Type Ia Supernovae at 0.511 < z < 1.12 and the Union2 Compilation. Amanullah, R.; Lidman, C.; Rubin, D.; Aldering, G.; Astier, P.; Barbary, K.; Burns, M. S.; Conley, A.; Dawson, K. S.; Deustua, S. E.; Doi, M.; Fabbro, S.; Faccioli, L.; Fakhouri, H. K.; Folatelli, G.; Fruchter, A. S.; Furusawa, H.; Garavini, G.; Goldhaber, G.; Goobar, A.; Groom, D. E.; Hook, I.; Howell, D. A.; Kashikawa, N.; Kim, A. G.; Knop, R. A.; Kowalski, M.; Linder, E.; Meyers, J.; Morokuma, T.; Nobili, S.; Nordin, J.; Nugent, P. E.; Östman, L.; Pain, R.; Panagia, N.; Perlmutter, S.; Raux, J.; Ruiz-Lapuente, P.; Spadafora, A. L.; Strovink, M.; Suzuki, N.; Wang, L.; Wood-Vasey, W. M.; Yasuda, N.; Supernova Cosmology Project, The. Astrophysical Journal, Volume 716, Issue 1, pp. 712–738 (2010).

Streamline Integration using MPI-Hybrid Parallelism on Large Multi-Core Architecture. David Camp, Christoph Garth, Hank Childs, David Pugmire, and Kenneth Joy. IEEE Transactions on Visualization and Computer Graphics, 07 Dec. 2010.

Supernova PTF 09UJ: A Possible Shock Breakout from a Dense Circumstellar Wind. Ofek, E. O.; Rabinak, I.; Neill, J. D.; Arcavi, I.; Cenko, S. B.; Waxman, E.; Kulkarni, S. R.; Gal-Yam, A.; Nugent, P. E.; Bildsten, L.; Bloom, J. S.; Filippenko, A. V.; Forster, K.; Howell, D. A.; Jacobsen, J.; Kasliwal, M. M.; Law, N.; Martin, C.; Poznanski, D.; Quimby, R. M.; Shen, K. J.; Sullivan, M.; Dekany, R.; Rahmer, G.; Hale, D.; Smith, R.; Zolkower, J.; Velur, V.; Walters, R.; Henning, J.; Bui, K.; McKenna, D. Astrophysical Journal, Volume 724, Issue 2, pp. 1396–1401 (2010).

Transitioning Users from the Franklin XT4 System to the Hopper XE6 System. K. Antypas, Y. He. CUG Procceedings 2011, Fairbanks, Alaska.

Two-Stage Framework for a Topology-Based Projection and Visualization of Classified Document Collections. Patrick Oesterling, Gerik Scheuermann, Sven Teresniak, Gerhard Heyer, Steffen Koch, Thomas Ertl, and Gunther H. Weber. Proceedings IEEE Symposium on Visual Analytics Science and Technology (IEEE VAST), Salt Lake City, Utah, October 2010.

Type II-P Supernovae as Standard Candles: The SDSS-II Sample Revisited. Poznanski, Dovi; Nugent, Peter E.; Filippenko, Alexei V. Astrophysical Journal, Volume 721, Issue 2, pp. 956–959 (2010).

Unveiling the Origin of Grb 090709A: Lack of Periodicity in a Reddened Cosmological Long-Duration Gamma-Ray Burst. Cenko, S. B.; Butler, N. R.; Ofek, E. O.; Perley, D. A.; Morgan, A. N.; Frail, D. A.; Gorosabel, J.; Bloom, J.

S.; Castro-Tirado, A. J.; Cepa, J.; Chandra, P.; de Ugarte Postigo, A.; Filippenko, A. V.; Klein, C. R.; Kulkarni, S. R.; Miller, A. A.; Nugent, P. E.; Starr, D. L. Astronomical Journal, Volume 140, Issue 1, pp. 224–234 (2010).

Vessel Network Detection Using Contour Evolution and Color Components. D. M. Ushizima, F. N. S. Medeiros, J. Cuadros, C. I. O. Martins. Proceedings of the 32nd Annual International Conference of the IEEE Engineering in Medicine and Biology Society, Buenos Aires, Argentina. Sept 2010.

Visualization and Analysis-Oriented Reconstruction of Material Interfaces. J. Meredith and H. Childs. Eurographics/IEEE Symposium on Visualization (EuroVis), Bordeaux, France, June 2010.

What's Ahead for Fusion Computing? Alice Koniges, Robert Preissl, Stephan Ethier, and John Shalf. International Sherwood Fusion Theory Conference, April 2010, Seattle, Washington.

Appendix A

NERSC Policy Board

The NERSC Policy Board meets at least annually and provides scientific and executive-level advice to the Berkeley Lab Director regarding the overall NERSC program and, specifically, on such issues as resource utilization to maximize the present and future scientific impact of NERSC, and long-range planning for the program, including the research and development necessary for future capabilities. Policy Board members are widely respected leaders in science, computing technology, or the management of scientific research and/or facilities.

Daniel A. Reed, Policy Board Chair
Corporate Vice President, Extreme Computing Group
Microsoft Corporation

Achim Bachem
Chairman of the Board of Directors
Forschungszentrum Jülich GmH

Jack Dongarra
University Distinguished Professor
Director, Innovative Computing Laboratory
University of Tennessee

Stephane Ethier
Ex officio, NERSC Users Group Chair
Computational Plasma Physics Group
Princeton Plasma Physics Laboratory

Sidney Karin
Professor of Computer Science and Engineering
University of California, San Diego

Pier Oddone
Director
Fermi National Accelerator Laboratory

Edward Seidel
Assistant Director for Mathematics and Physical Sciences
National Science Foundation

Appendix B

NERSC Client Statistics

In support of the DOE Office of Science's mission, NERSC served 4,294 scientists throughout the United States in 2010. 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 2011 proposal forms, principal investigators reported 1,528 refereed publications (published or in press)—as well as 262 publications submitted to refereed journals—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 XT4 equivalent hours.

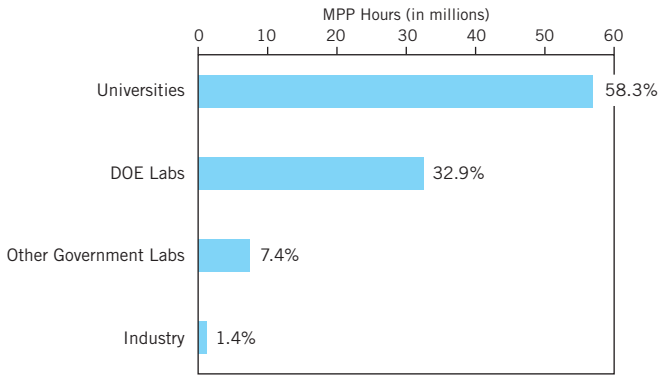

Figure 1. NERSC usage by institution type, 2010 (MPP hours).

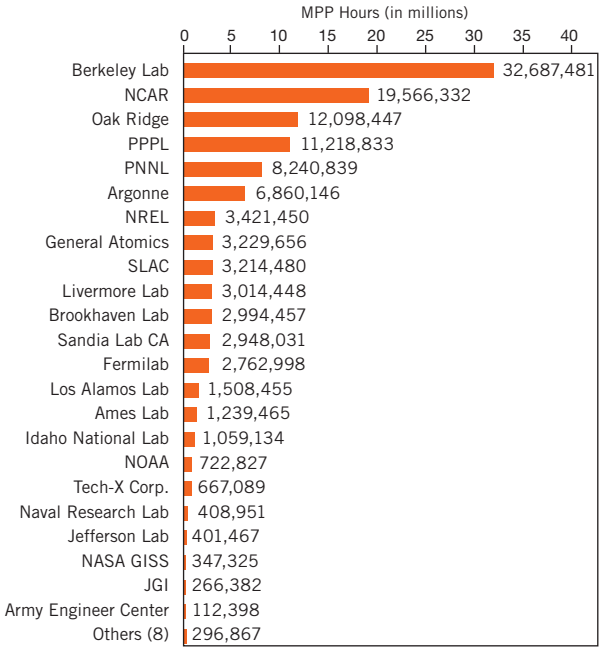

Figure 2. DOE and other laboratory usage at NERSC, 2010 (MPP hours).

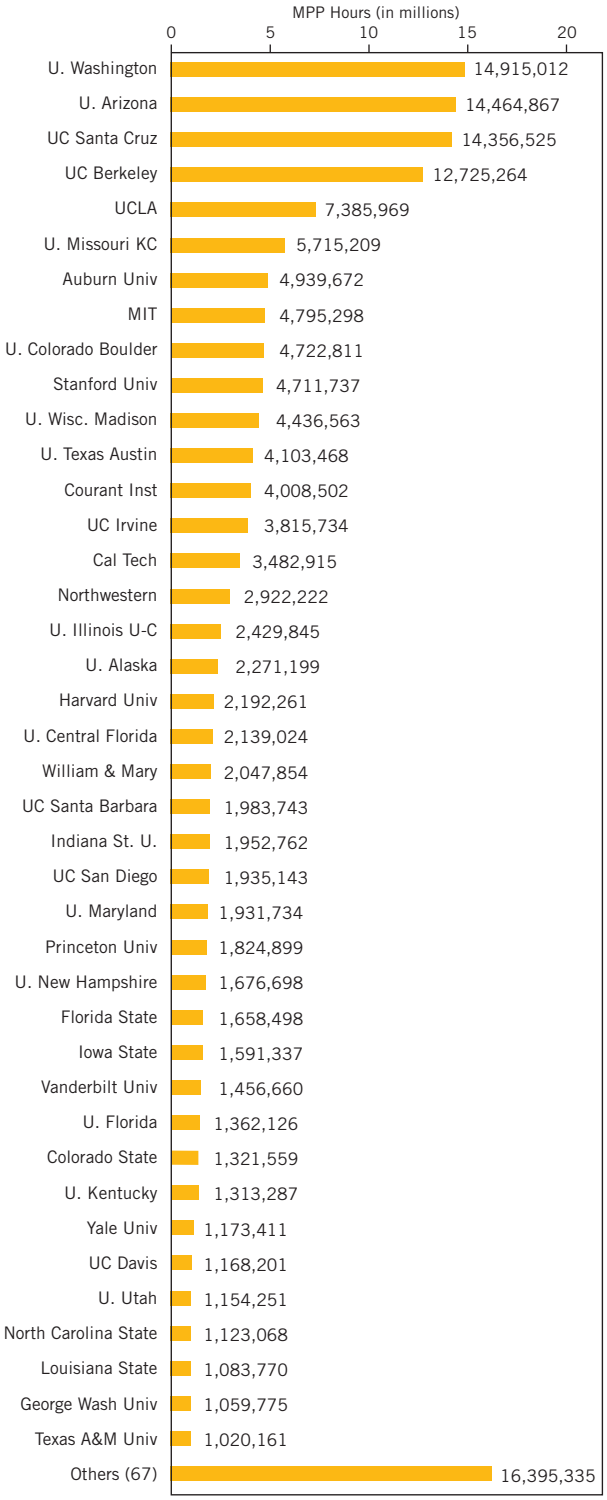

Figure 3. Academic usage at NERSC, 2010 (MPP hours).

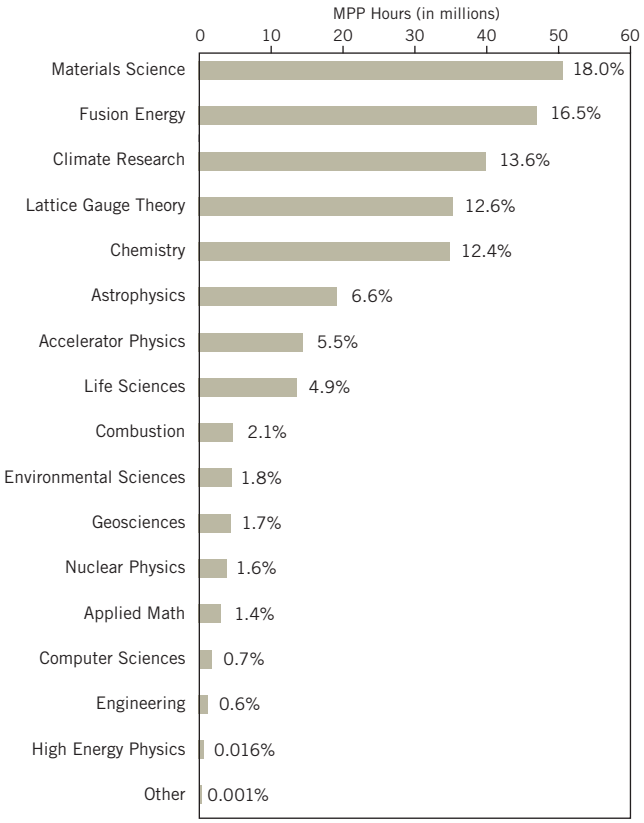

Figure 4. NERSC usage by scientific discipline, 2010 (MPP hours).

Appendix C

NERSC Users Group Executive Committee

Office of Advanced Scientific Computing Research
Mark Adams, Columbia University
Sergei Chumakov, Stanford University
Mike Lijewski, Lawrence Berkeley National Laboratory

Office of Basic Energy Sciences
Eric Bylaska, Pacific Northwest National Laboratory
Paul Kent, Oak Ridge National Laboratory
Brian Moritz, SLAC National Accelerator Laboratory

Office of Biological and Environmental Research
David Beck, University of Washington
Barbara Collignon, Oak Ridge National Laboratory**
Brian Hingerty, Oak Ridge National Laboratory*
Adrianne Middleton, National Center for Atmospheric Research

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
Cameron Geddes, Lawrence Berkeley National Laboratory*
Steven Gottlieb, Indiana University**
Frank Tsung (Vice Chair), University of California, Los Angeles

Office of Nuclear Physics
David Bruhwiler, Tech-X Corporation*
Daniel Kasen, Lawrence Berkeley National Laboratory**
Peter Messmer, Tech-X Corporation*
Tomasz Plewa, Florida State University
Martin Savage, University of Washington**

Members at Large
Angus Macnab, DyNuSim, LLC*
Gregory Newman, Lawrence Berkeley National Laboratory**
Jason Nordhaus, Princeton University
Ned Patton, National Center for Atmospheric Research*
Maxim Umansky, Lawrence Livermore National Laboratory

* Outgoing members
** Incoming members

Appendix D

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:

Daniel Hitchcock
Acting Associate Director, ASCR

Barbara Helland
Senior Advisor

Walter M. Polansky
Senior Advisor

Robert Lindsay
Special Projects

Julie Scott
Financial Management Specialist

Lori Jernigan
Program Support Specialist

Melea Baker
Administrative Specialist

Appendix D (continued)

Office of Advanced Scientific Computing Research

Facilities Division

Daniel Hitchcock
Division Director

Vincent Dattoria
Computer Scientist

Dave Goodwin
Physical Scientist

Barbara Helland
Computer Scientist

Sally McPherson
Program Assistant

Betsy Riley
Computer Scientist

Yukiko Sekine
Computer Scientist

Computational Science Research and Partnerships (SciDAC) Division

William Harrod
Division Director

Teresa Beachley
Program Assistant

Rich Carlson
Computer Scientist

Christine Chalk
Physical Scientist

Benjy Grover
Computer Scientist (Detailee)

Barbara Helland
Computer Scientist

Sandy Landsberg
Mathematician

Randall Laviolette
Physical Scientist

Steven Lee
Physical Scientist

Lenore Mullin
Computer 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

Appendix E

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. ASCAC's recommendations include advice on long-range plans, priorities, and strategies to address more effectively the scientific aspects of advanced scientific computing including the relationship of advanced scientific computing to other scientific disciplines, and maintaining appropriate balance among elements of the program. The Committee formally reports to the Director, Office of Science. The Committee primarily includes representatives of universities, national laboratories, and industries involved in advanced computing research. Particular attention is paid to obtaining a diverse membership with a balance among scientific disciplines, institutions, and geographic regions.

Roscoe C. Giles, Chair
Boston University

Marsha Berger
Courant Institute of Mathematical Sciences

Jackie Chen
Sandia National Laboratories

Jack J. Dongarra
University of Tennessee

Susan L. Graham
University of California, Berkeley

James J. Hack
Oak Ridge National Laboratory

Anthony Hey
Microsoft Research

Thomas A. Manteuffel
University of Colorado at Boulder

John Negele
Massachusetts Institute of Technology

Linda R. Petzold
University of California, Santa Barbara

Vivek Sarkar
Rice University

Larry Smarr
University of California, San Diego

William M. Tang
Princeton Plasma Physics Laboratory

Victoria White
Fermi National Accelerator Laboratory

Appendix F

Acronyms and Abbreviations

ACM	Association for Computing Machinery
ACS	American Chemical Society
AIP	American Institute of Physics
ALCF	Argonne Leadership Computing Facility
AMR	Adaptive mesh refinement
API	Application programming interface
APS	American Physical Society
ARPA-E	Advanced Research Projects Agency-Energy (DOE)
ARRA	American Recovery and Reinvestment Act
ASCR	Office of Advanced Scientific Computing Research (DOE)
BER	Office of Biological and Environmental Research (DOE)
BES	Office of Basic Energy Sciences (DOE)
BNL	Brookhaven National Laboratory
C_2H_5OH	Ethanol
CAF	Co-Array Fortran
CAS	Chinese Academy of Sciences
CATH	Class, Architecture, Topology and Homology
CCNI	Computational Center for Nanotechnology Innovations
CCSE	Center for Computational Sciences and Engineering (LBNL)
CDU	Cooling distribution unit
CH_3OH	Methanol
CH_4	Methane
CNRS	Centre National de la Recherche Scientifique, France
CNT	Carbon nanotube
CO_2	Carbon dioxide
COE	Center of Excellence
CPU	Central processing unit
CRD	Computational Research Division, Lawrence Berkeley National Laboratory
CSCS	Swiss National Supercomputing Centre
CSIR	Council of Scientific and Industrial Research, India
CUDA	Compute Unified Device Architecture
CXI	Coherent X-ray imaging
CXIDB	Coherent X-Ray Imaging Data Bank
DAE	Department of Atomic Energy, India
DFT	Density functional theory
DNA	Deoxyribonucleic acid
DOE	U.S. Department of Energy

DOS	Density of state
DST	Department of Science and Technology, India
DVS	Data Virtualization Services
EFRC	Energy Frontier Research Center
ELM	Edge localized mode
EPSRC	Engineering and Physical Sciences Research Council, UK
ESG	Earth System Grid
ESnet	Energy Sciences Network
FAPESP–CNPq	Fundação de Amparo à Pesquisa do Estado de São Paulo–Conselho Nacional de Desenvolvimento Científico e Tecnológico, Brazil
FES	Office of Fusion Energy Sciences (DOE)
FOM	Fundamenteel Onderzoek der Materie, Netherlands
GA	Grant Agency of the Czech Republic
GB	Gigabyte
Gflop/s	Gigaflop/s
GHz	Gigahertz
GNR	Graphene nanoribbon
GPFS	General Parallel File System
GPGPU	General-purpose graphics processing unit
GPU	Graphics processing unit
H_2	Hydrogen
HCO	Hydrogen carbonate (transition state)
HEP	Office of High Energy Physics (DOE)
HMA	Highly mismatched alloy
HPC	High performance computing
HPSS	High Performance Storage System
HTML	Hypertext Markup Language
HTTP	Hypertext Transfer Protocol
I/O	Input/output
ICCS	International Center for Computational Science
ICIAM	International Council for Industrial and Applied Mathematics
IEEE	Institute of Electrical and Electronics Engineers
IN2P3	Institut National de Physique Nucléaire et Physique des Particules, France
INCITE	Innovative and Novel Computational Impact on Theory and Experiment (DOE)
IPCC	Intergovernmental Program on Climate Change
IT	Information technology
ITER	Latin for “the way”; an international fusion energy experiment in southern France
JGI	Joint Genome Institute (DOE)
KAUST	King Abdullah University of Science and Technology
KISTI	Korea Institute of Science and Technology Information
KRF	Korea Research Foundation
LBNL	Lawrence Berkeley National Laboratory
LED	Light-emitting diode

LLNL	Lawrence Livermore National Laboratory
LRZ	Leibniz Supercomputing Centre
MB	Megabyte
MESRF	Ministry of Education and Science of the Russian Federation
MHD	Magnetohydrodynamic
MIT	Massachusetts Institute of Technology
Mo_6S_8	Molybdenum sulfide nanocluster
MoE	Ministry of Education of the People's Republic of China
MOM	Machine-oriented miniserver
MoS_2	Molybdenum disulfide
MoST	Ministry of Science and Technology of the People´s Republic of China
MPI	Message Passing Interface
MPP	Massively parallel processing
MSES	Ministry of Science, Education and Sports of the Republic of Croatia
MSMT	Ministerstvo Školství, Mládeže a Tělovýchovy, Czech Republic
NASA	National Aeronautics and Space Administration
NCAR	National Center for Atmospheric Research
NCSA	National Center for Supercomputing Applications
NERSC	National Energy Research Scientific Computing Center
NEWT	NERSC Web Toolkit
NGF	NERSC Global Filesystem
NICS	National Institute for Computational Science, University of Tennessee
NIH	National Institutes of Health
NISE	NERSC Initiative for Scientific Exploration
NMR	Nuclear magnetic resonance
NNSFC	National Natural Science Foundation of China
NO_2	Nitrogen dioxide
NOW	Nederlandse Organisatie voor Wetenschappelijk Onderzoek, Netherlands
NP	Office of Nuclear Physics (DOE)
NSF	National Science Foundation
NUG	NERSC Users Group
O_2	Oxygen
O_3	Ozone
OSC	Ohio Supercomputer Center
OSG	Open Science Grid
OSU	Ohio State University
PB	Petabyte
PDB	Protein Data Bank
PDSF	Parallel Distributed Systems Facility (NERSC)
PECASE	Presidential Early Career Award for Scientists and Engineers
Petaflops	Quadrillions of floating point operations per second
Pflops	Petaflops

PGI	Portland Group, Inc.
PMAMR	Porous Media AMR (computer code)
PMSHE	Polish Ministry of Science and Higher Education
PNAS	Proceedings of the National Academy of Sciences
PPPL	Princeton Plasma Physics Laboratory
PRACE	Partnership for Advanced Computing in Europe
PTF	Palomar Transient Factory
QMC	Quantum Monte Carlo
RCF	RHIC Computing Facility
REST	Representational state transfer
RHIC	Relativistic Heavy Ion Collider
RMST	Russian Ministry of Science and Technology
ROSATOM	State Atomic Energy Corporation, Russia
SC	Office of Science (DOE)
SciDAC	Scientific Discovery through Advanced Computing (DOE)
SCOP	Structural Classification of Proteins
SDN	Science Data Network (ESnet)
SIAM	Society for Industrial and Applied Mathematics
SLIRP	Structural Library of Intrinsic Residue Propensities
SMP	Symmetric multiprocessor
STFC	Science and Technology Facilities Council, UK
SUNY	State University of New York
TACC	Texas Advanced Computing Center
TB	Terabyte
TCP	Transmission Control Protocol
Teraflops	Trillions of floating point operations per second
TIGRESS	Terascale Infrastructure for Groundbreaking Research in Engineering and Science High Performance Computing Center, Princeton University
TPC	Time projection chamber
UC	University of California
UCB	University of California, Berkeley
UPC	Unified Parallel C (programming language)
URL	Uniform Resource Locator
UV	Ultraviolet

For more information about NERSC, contact:

Jon Bashor
NERSC Communications
Berkeley Lab, MS 50B4230
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
John Hules, Jon Bashor, Linda Vu, and Margie Wylie, Berkeley Lab Computing Sciences Communications Group; Paul Preuss and Lynn Yarris, Berkeley Lab Public Affairs Department; David L. Chandler, MIT News Office; Pam Frost Gorder, Ohio State University; Kitta MacPherson, Princeton University Communications Office; Rachel Shafer, UC Berkeley Engineering Marketing and Communications Office; Karen McNulty Walsh, Brookhaven National Laboratory.

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

CSO 20748

Matthew Andrews

Katie Antypas · Brian Austin · Clayton Bagwell

William Baird · Nicholas Balthaser · Jon Bashor · Elizabeth Bautista

Del Black · Jeremy Brand · Jeffrey Broughton · Gregory Butler · Tina Butler

David Camp · Scott Campbell · Shane Canon · Nicholas Cardo · Jonathan Carter · Steve Chan

Ravi Cheema · Hank Childs · Shreyas Cholia · Roxanne Clark · James Craw · Thomas Davis

Tina Declerck · David Donofrio · Brent Draney · Matt Dunford · Norma Early · Kirsten Fagnan · Aaron Garrett

Richard Gerber · Annette Greiner · Jeff Grounds · Leah Gutierrez · Patrick Hajek · Matthew Harvey

Damian Hazen · Helen He · Mark Heer · Jason Hick · Eric Hjort · John Hules · Wayne Hurlbert · John Hutchings

Christos Kavouklis · Randy Kersnick · Jihan Kim · Alice Koniges · Yu Lok Lam · Thomas Langley · Craig Lant

Stefan Lasiewski · Jason Lee · Rei Chi Lee · Bobby Liu · Fred Loebl · Burlen Loring · Steven Lowe · Filipe Maia

Ilya Malinov · Jim Mellander · Joerg Meyer · Praveen Narayanan · Robert Neylan · Trever Nightingale

Peter Nugent · Isaac Ovadia · R. K. Owen · Viraj Paropkari · David Paul · Lawrence Pezzaglia · Jeff Porter

Prabhat · Robert Preissl · Tony Quan · Lavanya Ramakrishnan · Lynn Rippe · Cody Rotermund

Oliver Ruebel · Iwona Sakrejda · John Shalf · Hongzhang Shan · Hemant Shukla · David Skinner

Jay Srinivasan · Suzanne Stevenson · David Stewart · Michael Stewart · David Turner · Alex Ubungen

Andrew Uselton · Daniela Ushizima · Francesca Verdier · Linda Vu · Howard Walter

Harvey Wasserman · Gunther Weber · Mike Welcome · Nancy Wenning · Cary Whitney

Nicholas Wright · Margie Wylie · Woo-Sun Yang · Yushu Yao

Katherine Yelick · Rebeccca Yuan · Brian Yumae

Zhengji Zhao

NERSC

BERKELEY LAB
Lawrence Berkeley National Laboratory

U.S. DEPARTMENT OF ENERGY
Office of Science