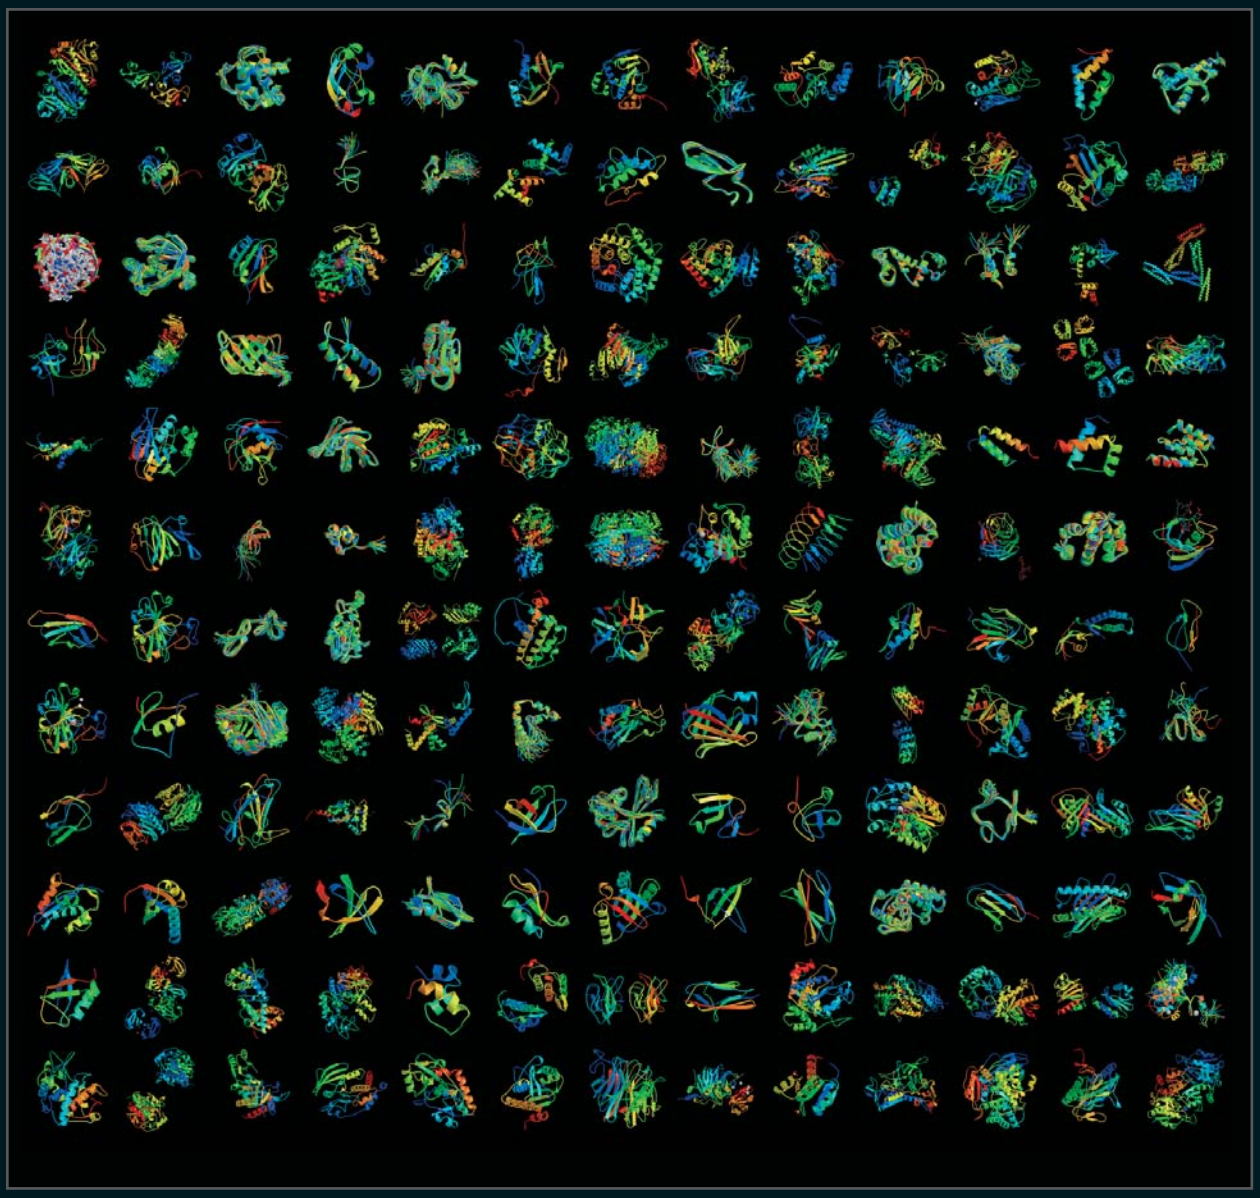

05

NERSC

National Energy Research
Scientific Computing Center

2005 Annual Report

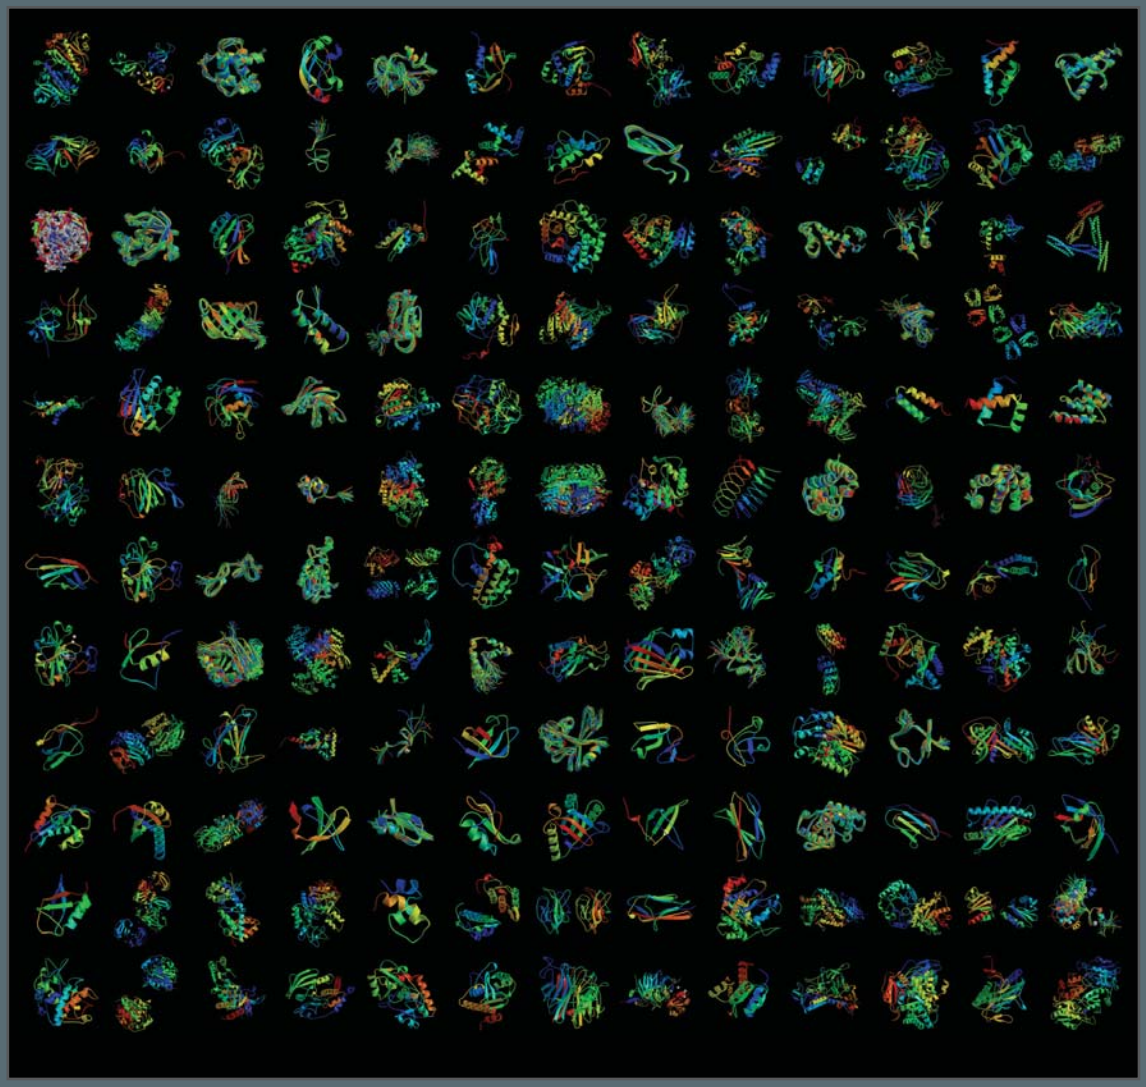

NERSC

National Energy Research
Scientific Computing Center

2005 Annual Report

Lawrence Berkeley National Laboratory / 1 Cyclotron Road / Berkeley CA / 94720

LBNL-60296

TABLE OF CONTENTS

THE YEAR IN PERSPECTIVE

During the spring of 2006, we find ourselves reflecting on the "early days" when NERSC moved to Berkeley a decade ago. We and many of the NERSC staff can now look back over a full ten years of our professional lives connected to NERSC. Memories of 1996, from the first system installation in a not yet completed computer room to the dedication of the Cray T3E, come back and remind us how far NERSC has progressed since then. So many exciting events have happened during that time, it would be impossible to list them all here.

Instead we will proudly repeat what Secretary of Energy Samuel Bodman recently said about NERSC: "NERSC has a well-earned reputation for providing highly reliable systems, fast turnaround on critical projects, and dedicated support for users." He made these comments when announcing a renewed allocation of NERSC computer resources to the U.S. Army Corps of Engineers to simulate Gulf Coast hurricanes. In the aftermath of hurricane Katrina, the Corps plans to use the NERSC cycles for simulations of the Texas, Louisiana, and Mississippi coastlines. These simulations will allow researchers to produce more accurate models for calculating the effects of future hurricanes. "Because these simulations could literally affect the lives of millions of Americans, we want to ensure that our colleagues in the Corps of Engineers have access to supercomputers which are up to the task," the Secretary stated, giving NERSC credit for its proven record of delivering highly reliable production supercomputing services.

We intentionally chose not to commemorate ten years in Berkeley with a big celebration this year. With rapid new developments expected in the next couple of years, our focus must be on the future, not on reminiscing. In 2005 NERSC set the stage for the next five years of its development through a series of important planning activities. Through the well-established Greenbook process, our active user community provided their input to the planning process. NERSC management then developed a new five-year plan for 2006 to 2010, which is discussed in

detail in this annual report. This plan was then thoroughly reviewed in a programmatic review by DOE. We can be proud about the outcome of this review of our plans, which were fully endorsed: "NERSC is a strong, productive, and responsive science-driven center that possesses the potential to significantly and positively impact scientific progress.... NERSC is extremely well run with a lean and knowledgeable staff." This strong endorsement, together with the continued support of our program office at DOE and the good budget news for 2007 and beyond, gives us the most positive outlook we have had in the last several years.

One of the reasons for excitement in the near future is the upcoming introduction of the NERSC-5 system in 2006 and 2007. As we eagerly await the completion of our contract negotiations, we can confidently state that NERSC-5 will be a major step forward in providing a new capability for our users, unlike any of the upgrades in the past half decade. The new system will provide 16 teraflop/s sustained performance on our benchmarks (more than 100 teraflop/s peak), which is about an order of magnitude more sustained performance than NERSC offers today in early 2006.

In addition to our increase of capability computing in the near future, NERSC added significant capacity in 2005 by introducing two new clusters, named "Jacquard" and "Bassi." This increase in capacity was highly welcomed by the NERSC community. Our users make the most progress in computational science with exactly this type of system that provides high performance in a reliable, predictable environment. One of our users, Robert Duke of the University of North Carolina, commented on the two systems as follows: "I have to say that both of these machines are really nothing short of fabulous. While Jacquard is perhaps the best-performing commodity cluster I have seen, Bassi is the best machine I have seen, period." By delivering this quality, NERSC makes it possible for its users to concentrate on their science. Not surprisingly, our user community enjoys continued and unparalleled scientific productivity. In 2005 we were able to list more than 1100 scientific publications that were written on the basis of simulations carried out at NERSC.

We also saw a long-term goal accomplished that will significantly set apart the NERSC production environment from other centers. In early 2006, NERSC deployed the NERSC Global Filesystem (NGF) into production, providing seamless data access from all of the Center's computational and analysis resources. NGF is intended to facilitate sharing of data between users and between machines. NGF's single unified namespace makes it easier for users to manage their data across multiple systems. Users no longer need to keep track of multiple copies of programs and data, and they no longer need to copy data between NERSC systems for pre- and post-processing. NGF provides several other benefits as well: storage utilization is more efficient because of decreased fragmentation, and computational resource utilization is more efficient because users can more easily run jobs on an appropriate resource.

Thus today NERSC users have the benefit of scalable high-end capability computing in Seaborg (and the soon-to-be-added NERSC-5), along with reliable capacity machines, in an integrated environment that offers a global filesystem, analytics support, and a seemingly infinite mass storage system (now close to 40 petabytes). Our five-year plan will move this integrated system environment to the near petaflop/s performance level, which we will reach in 2010 with the planned introduction of NERSC-6. Scalability to tens of thousands of processors, both for applications and systems software, managing petabytes of data, and at the same time continuing the expected level of support, reliability, and quality of service, will be the big challenges ahead. Thanks to the ongoing support from our program management at the Office of Advanced Scientific Computing Research at DOE, the NERSC budget has been set at a level that makes these ambitious plans feasible. Thus we are confidently looking forward to another year of both scientific and computing accomplishments at NERSC. As always, this progress would not be possible without the NERSC staff, who continue to tirelessly dedicate their time, skill, and effort to make NERSC the best scientific computing resource in the world. Our special thanks to all of you.

Horst D. Simon, *NERSC Center Division Director*

William T. C. Kramer, *Division Deputy and NERSC Center General Manager*

RESEARCH NEWS

Research at NERSC spans a wide range of topics, from subatomic particles to the formation of the cosmos, and from alternative energy sources to the behavior of proteins. This year's Research News presents a sampling of important investigations and discoveries made using NERSC's computational resources.

Global warming and energy issues were frequently in the news this past year. The article "Delayed reactions" describes three investigations into the dynamics and results of global warming; they conclude that climate change is no longer avoidable. "Burning questions" and "Combustion up close" report on breakthrough studies of the basic processes of combustion; studies such as these contribute to the scientific understanding needed if we are to improve the efficiency of our energy use. "Hailstones in hell" discusses an important issue in the design of fusion reactors, a long-anticipated alternative source of energy that is moving closer to realization.

Computational studies of matter and its behavior cover the widest possible range of scales, from subatomic to molecular to nanoscale to cosmic. "A perfect liquid" reports on the remarkable newly created state of quark-gluon matter. "Whispers from underground" relates how ghostly geoneutrinos hint at Earth's inner secrets. "Breaking up is hard to calculate" tells of the first complete numerical solution of the fragmentation of a system with four charged particles. At the nanoscale, "Talent scouting" and "Surface charge" discuss important discoveries about catalysts. And at the cosmic scale, "Magnetic disks in space" explores the origins of instability in accretion disks.

Finally, computing is becoming an indispensable tool for biology and medicine. The article "Proteins in motion" tells how biochemists are systematically simulating the unfolding pathways of all known protein folds, hoping to discover the general principles of protein folding, one of the basic mechanisms of life.

The heat is on

New studies reveal the inevitable and possibly self-perpetuating consequences of global warming

A significant rise in sea level is inevitable by the end of this century; ocean temperatures over the last 40 years have risen far above naturally occurring climate cycles; and continued greenhouse gas emissions may overwhelm the capacity of land and ocean to absorb excess carbon dioxide, thus speeding up global warming. Those are the conclusions of three recent climate modeling studies, all pointing to the inevitability of global warming and the need to develop mitigation strategies.

Even if the concentrations of all greenhouse gases had leveled off in the year 2000, we would still be committed to a warmer Earth and rising sea levels in this century, according to a study by a team of climate modelers at the National Center for Atmospheric Research (NCAR).[1] The models were run on supercomputers at NCAR and several DOE labs, including NERSC, and on the Earth Simulator in Japan.

The modeling study quantifies the relative rates of sea level rise and global temperature increase that we are already committed to in the 21st century. Even if no more greenhouse gases were added to the atmosphere, globally averaged surface air temperatures would rise about 0.5° C (1° F) and global sea levels would rise another 11 centimeters (4.3 inches) from thermal expansion alone by 2100.

"Many people don't realize we are committed right now to a significant amount of global warming and sea level rise because of the greenhouse gases we have already put into the atmosphere," said lead author Gerald Meehl. "Even if we stabilize greenhouse gas concentrations, the climate will continue to warm, and there will be proportionately even more sea level rise. The longer we wait, the more climate change we are committed to in the future."

The half-degree temperature rise is similar to that observed during the second half of the 20th century, but the projected sea level rise is more than twice the 5-centimeter (2-inch) rise that occurred during that period. These numbers do not take into account fresh water from melting ice sheets and glaciers, which could at least double the sea level rise caused by thermal expansion alone.

Though temperature rise shows signs of leveling off 100 years after stabi-

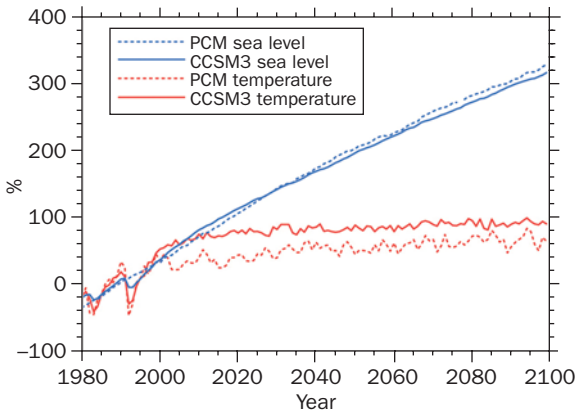

FIGURE 1. Percent increase of globally averaged surface air temperature and sea level rise from the two models computed relative to values for the base period 1980–1999, for the experiment in which greenhouse gas concentrations and all other atmospheric constituents were stabilized at the end of the 20th century.

[1] G. A. Meehl, W. M. Washington, W. D. Collins, J. M. Arblaster, A. Hu, L. E. Buja, W. G. Strand, and H. Teng, "How much more global warming and sea level rise?" Science **307**, 1769 (2005). Research funding: NSF, BER, NCAR, MEXT. Computational resources: NCAR, NERSC, LANL, NCCS/ORNL, ES.

lization in the study, ocean waters continue to warm and expand, causing global sea level to rise unabated (Figure 1).

"With the ongoing increase in concentrations of greenhouse gases, every day we commit to more climate change in the future," Meehl explained. "When and how we stabilize concentrations will dictate, on the time scale of a century or so, how much more warming we will experience. But we are already committed to ongoing large sea level rise, even if concentrations of greenhouse gases could be stabilized."

The inevitability of the climate changes described in the study is the result of thermal inertia, mainly from the oceans, and the long lifetime of carbon dioxide and other greenhouse gases in the atmosphere. Thermal inertia refers to the process by which water heats and cools more slowly than air because it is denser than air.

This study quantifies future committed climate change using coupled global three-dimensional climate models. Coupled models link major components of Earth's climate in ways that allow them to interact with each other. Meehl and his NCAR colleagues ran the same scenario a number of times and averaged the results to create ensemble simulations from each of two global climate models. Then they compared the results from each model.

The scientists also used the two models to compare possible 21st century climate scenarios in which greenhouse gases continue to build in the atmosphere at low, moderate, or high rates. The worst-case scenario projects an average temperature rise of 3.5° C (6.3° F) and sea level rise from thermal expansion of 30 centimeters (12 inches) by 2100 (Figure 2). All scenarios analyzed in the study will be assessed by international teams of scientists for the next report by the Intergovernmental Panel on Climate Change, due out in 2007.

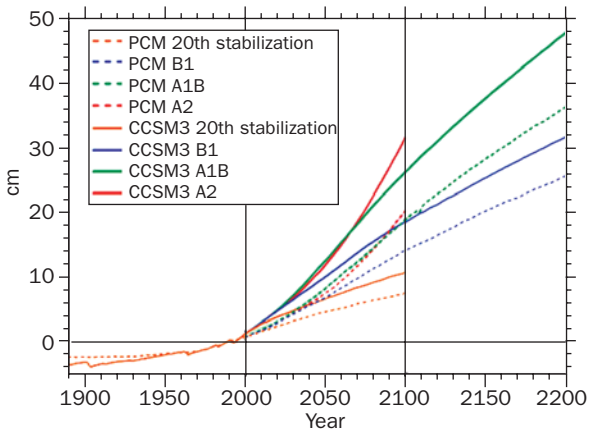

FIGURE 2. Globally averaged sea level rise from thermal expansion for various scenarios of greenhouse gas buildup.

The NCAR team used the Parallel Climate Model (PCM), developed by NCAR and the Department of Energy, and the new Community Climate System Model (Version 3). The CCSM3 was developed at NCAR with input from university and federal climate scientists around the country and principal funding from the National Science Foundation (NCAR's primary sponsor) and the Department of Energy. The CCSM3 shows slightly higher temperature rise and sea level rise from thermal expansion and greater weakening of the thermohaline circulation in the North Atlantic. Otherwise, the results from the two models are similar. Three of the researchers later performed the same analysis using 12 different general climate models, and all showed similar results.

The warming of the oceans over the past 40 years has not been uniform, and another recent study[2] investigated warming at different depths on an ocean-by-ocean basis, using observational data as well as simulations from PCM and the HadCM3 model from Britain's Hadley Centre. The researchers noted that the vertical structure of the temperature data varies widely by ocean, and they explored three possible causes for the warming trend: (1) natural variability internal to the coupled ocean– atmosphere system; (2) external natural variability, such as solar or volcanic forcing; and (3) forcing arising from human activity (emission of greenhouse gases and sulfate aerosols).

The study showed that the warming trend is much stronger than would be expected from natural variability. In Figure 3, the hatched blue region shows simulated natural internal variability; the green triangles represent combined solar and volcanic forcing; and the red dots represent observed temperature trends. On the time and space scales used in this study, the solar plus volcanic signals are generally indistinguishable from the natural internal variability, but the observed temperatures bear little resemblance to the natural variations. Figure 4 shows that the actual temperature trends do match the simulated results of greenhouse gases and sulfate aerosols, indicating that ocean warming is caused by human activities.

"This is perhaps the most compelling evidence yet that global warming is happening right now and it shows that we can successfully simulate its past and likely future evolution," said Tim

[2] T. P. Barnett, D. W. Pierce, K. M. AchutaRao, P. J. Gleckler, B. D. Santer, J. M. Gregory, and W. M. Washington, "Penetration of human-induced warming into the world's oceans," Science **309**, 284 (2005). Research funding: NOAA, BER, DEFRA, GMRDP. Computational resources: NCAR, NERSC, ORNL, LANL, Hadley Center.

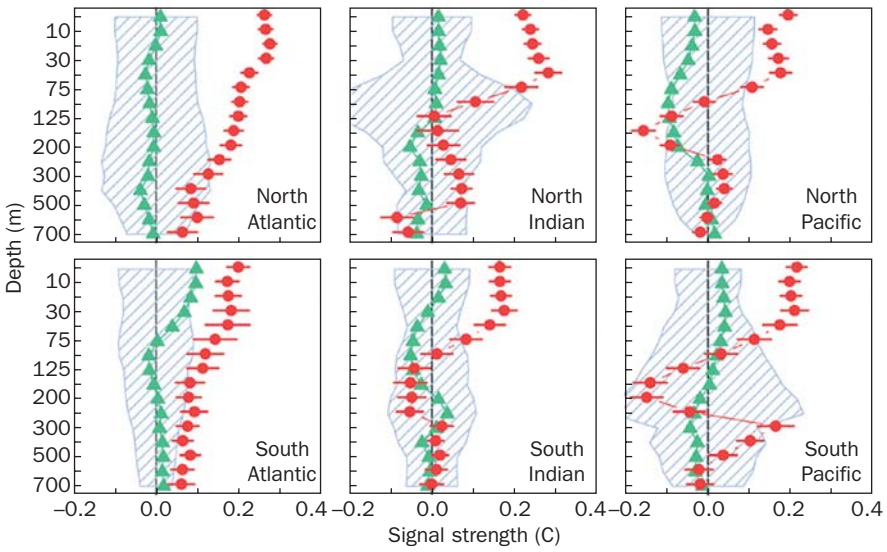

FIGURE 3. Observed ocean temperature trends from the last 40 years (red dots), projected onto a PCM simulation of natural variability internal to the coupled ocean–atmosphere system (blue hatched area) and combined solar and volcanic forcing (green triangles). Actual temperatures are much higher than would be expected from natural variability.

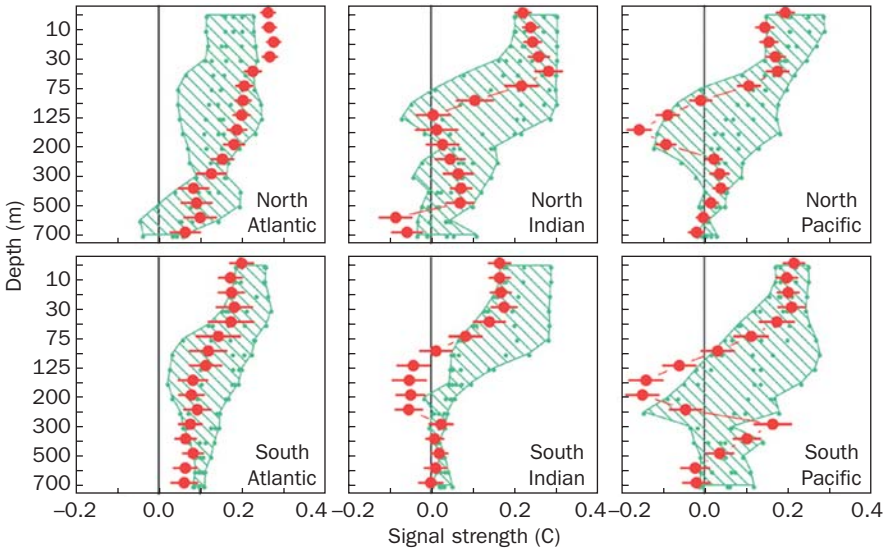

FIGURE 4. Observed ocean temperature trends (red dots) projected onto a PCM simulation of the results of greenhouse gases and sulfate aerosols (green hatched area). There is excellent agreement at most depths in all oceans, indicating that ocean warming is the result of human activities. Similar results were obtained with the HadCM3 model.

Barnett, lead author of the study and a research marine physicist at the Scripps Institution of Oceanography. Barnett said he was "stunned" by the results because the computer models reproduced the penetration of the warming signal in all the oceans. "The statistical significance of these results is far too strong to be merely dismissed and should wipe out much of the uncertainty about the reality of global warming."

A third study, based on a climate model that includes the effects of Earth's carbon cycle, indicates there are limits to our planet's ability to absorb increased emissions of carbon dioxide.[3] If the current release of carbon from fossil fuels continues unabated, by the end of the century the land and oceans will be less able to take up carbon than they are today, the model indicates.

"If we maintain our current course of fossil fuel emissions or accelerate our emissions, the land and oceans will not be able to slow the rise of carbon dioxide in the atmosphere the way they're doing now," said Inez Y. Fung at the University of California, Berkeley, who is director of the Berkeley Atmospheric Sciences Center, co-director of the new Berkeley Institute of the Environment, and professor of earth and planetary science and of environmental science, policy and management. "It's all about rates. If the rate of fossil fuel emissions is too high, the carbon storage capacity of the land and oceans decreases and climate warming accelerates."

Currently, the land and oceans absorb about half of the carbon dioxide produced by human activity, most of it resulting from the burning of fossil fuels, Fung said. Some scientists have suggested that the land and oceans will continue to absorb more

[3] I. Y. Fung, S. C. Doney, K. Lindsay, and J. John, "Evolution of carbon sinks in a changing climate," Proceedings of the National Academy of Sciences **102**, 11201 (2005). Research funding: NSF, NASA, BER, LBNL, WHOI. Computational resources: NCAR, NERSC.

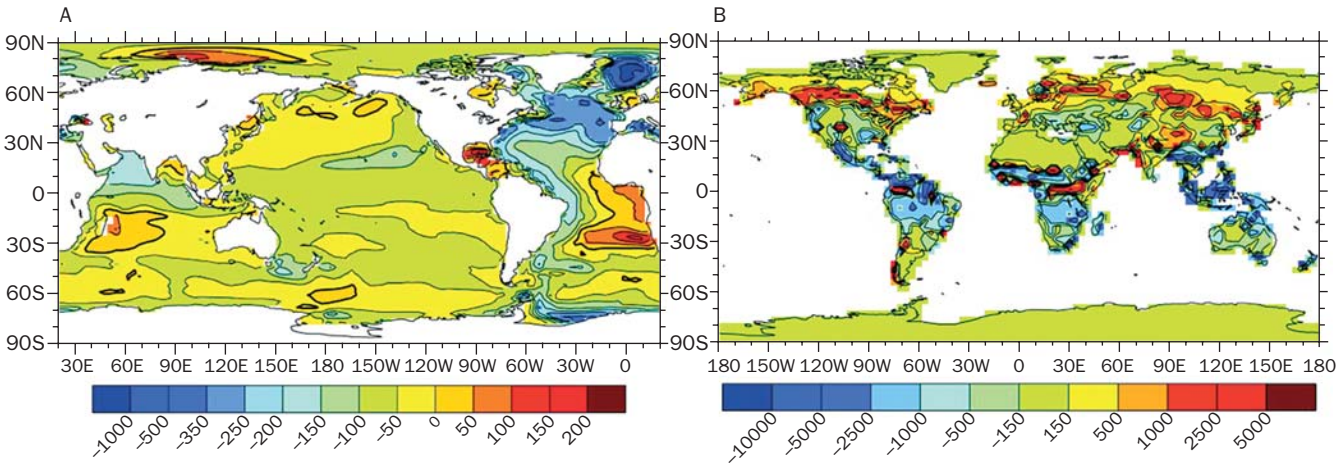

FIGURE 5. The impact of carbon–climate feedback on carbon storage for the years 2001–2100 (a) in the oceans and (b) on land. The yellow, green, and blue areas indicates less carbon storage, while the orange and red areas indicate more carbon storage.

and more CO_2 as fossil fuel emissions increase, making plants flourish and the oceans bloom. Fung's computer model, however, indicates that the "breathing biosphere" can absorb carbon only so fast. Beyond a certain point, the planet will not be able to keep up with carbon dioxide emissions (Figure 5).

"The reason is very simple," Fung said. "Plants are happy growing at a certain rate, and though they can accelerate to a certain extent with more CO_2, the rate is limited by metabolic reactions in the plant, by water and nutrient availability, et cetera."

Fung and her colleagues found that increasing temperatures and drought frequencies lower plant uptake of CO_2 as plants breathe in less to conserve water. At some point, the rate of fossil fuel CO_2 emissions will outstrip the ability of the vegetation to keep up, leading to a rise in atmospheric CO_2, increased greenhouse temperatures and increased frequency of droughts. An amplifying loop leads to ever higher temperatures, more droughts, and higher CO_2 levels.

The oceans exhibit a similar trend, Fung said, though less pronounced. There, mixing by turbulence in the ocean is essential for moving CO_2 down into the deep ocean, away from the top 100 meters of the ocean, where carbon absorption from the atmosphere takes place. With increased temperatures, the ocean stratifies more, mixing becomes harder, and CO_2 accumulates in the surface ocean instead of in the deep ocean. This accumulation creates a back pressure, lowering CO_2 absorption.

In all, business as usual would lead to a 1.4° C (2.5° F) rise in global temperatures by the year 2050. This estimate is at the low range of projected increases for the 21st century, Fung said, though overall, the model is in line with others predicting large ecosystem changes, especially in the tropics.

With voluntary controls that flatten fossil fuel CO_2 emission rates by the end of the century, the land and oceans could keep up with CO_2 levels and continue to absorb at their current rate, the model indicates.

"This is not a prediction, but a guideline or indication of what could happen," Fung said. "Climate prediction is a work in progress, but this model tells us that, given the increases in greenhouse gases, the Earth will warm up; and given warming, hot places are likely to be drier, and the land and oceans are going to take in carbon at a slower rate; and therefore, we will see an amplification or acceleration of global warming."

"The Earth is entering a climate space we've never seen before, so we can't predict exactly what will happen," she added. "We don't know where the threshold is. A two-degree increase in global temperatures may not sound like much, but if we're on the threshold, it could make a big difference."

Fung and colleagues have worked for several decades to produce a model of the Earth's carbon cycle that includes not only details of how vegetation takes up and releases carbon, but also details of decomposition by microbes in the soil, the carbon chemistry of oceans and lakes, the influence of rain and clouds, and many other sources

and sinks for carbon. The model takes into account thousands of details, ranging from carbon uptake by leaves, stems, and roots to the different ways that forest litter decomposes, day-night shifts in plant respiration, the salinity of oceans and seas, and effects of temperature, rainfall, cloud cover and wind speed on all these interactions.

All of today's climate models are able to incorporate the climate effects of carbon dioxide in the atmosphere, but only with concentrations of CO_2 specified by the modelers. Fung's model does not specify atmospheric CO_2 levels, but rather predicts the levels, given fossil fuel emissions. The researchers used observations of the past two centuries to make sure that their model is reasonable, and then used the model to project what will happen in the next 100 years.

The climate model coupled with the carbon cycle has been her goal for decades, as she tried to convince climate modelers that "whether plants are happy or not happy has an influence on climate projections. To include interactive biogeochemistry in climate models, which up to now embrace primarily physics and dynamics, is new."

She admits, however, that much work remains to be done to improve modeling. Methane and sulfate cycles must be included, plus effects like changes in plant distribution with rising temperatures, the possible increase in fires, disease, or insect pests, and even the effects of dust in the oceans.

"We have created a blueprint, in terms of a climate modeling framework, that will allow us to go beyond the physical climate models to more sophisticated models," she said. "Then, hopefully, we can understand what is going on now and what could happen. This understanding could guide our choices for the future."

This article written by: NCAR Media Relations Office, Robert Sanders (UC Berkeley), Mario Aguilera and Cindy Clark (Scripps Institution of Oceanography), John Hules (Berkeley Lab).

Burning questions

New 3D simulations are closing in on the holy grail of combustion science: turbulence–chemistry interactions

Controlling fire to provide heat and light was one of humankind's first great achievements, and the basic chemistry of combustion—what goes in and what comes out—was established long ago. But a complete quantitative understanding of what happens during the combustion process has remained as elusive as the ever-changing shape of a flame.

Even a simple fuel like methane (CH_4), the principal component of natural gas, burns in a complex sequence of steps. Oxygen atoms gradually replace other atoms in the hydrocarbon molecules, ultimately leaving carbon dioxide and water.

Methane oxidation involves about 20 chemical species for releasing energy, as well as many minor species that can become pollutants. Turbulence can distort the distribution of species and the redistribution of thermal energy which are required to keep the flame burning. These turbulence–chemistry interactions can cause the flame to burn faster or slower and to create more or less pollution.

The holy grail of combustion science has been to observe these turbulence–chemistry effects. Over the past few decades amazing progress has been made in observational techniques, including the use of lasers to excite molecules of a given species and produce a picture of the chemical distribution. But laser imaging is limited in the species and concentrations that can be reliably observed, and it is difficult to obtain simultaneous images to correlate different chemical species.

Because observing the details of combustion is so difficult, progress in combustion science has largely coincided with advances in scientific computing. For example, while basic concepts for solving one-dimensional flat flames originated in the 1950s, it only became possible to solve the 1D flame equations some 30 years later using Cray-1 supercomputers. Those calculations,

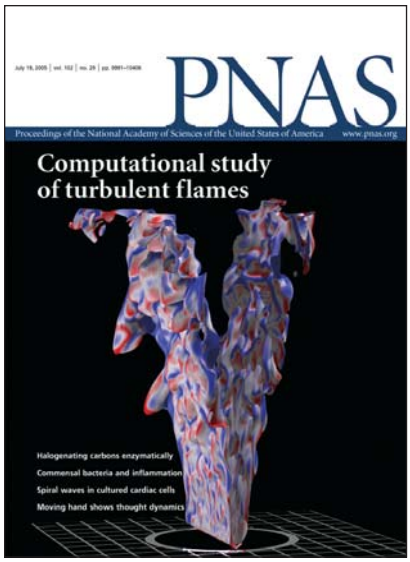

FIGURE 6. The calculated surface of a turbulent premixed laboratory methane flame.

which are routine on personal computers today, enabled a renaissance in combustion science by allowing chemists to observe the interrelationships among the many hypothesized reaction processes in the flame.

Simulating three-dimensional turbulent flames took 20 more years of advances in applied mathematics, computer science, and computer hardware, particularly massively parallel systems. But the effort has been worth it. New 3D simulations are beginning to provide the kind of detailed information about the structure and dynamics of turbulent flames that will be needed to design new low-emission, fuel-efficient combustion systems.

The first 3D simulation of a laboratory-scale turbulent flame from first principles—the result of a SciDAC-funded collaboration between computational and experimental scientists at Berkeley Lab—was featured on the cover of the July 19, 2005 Proceedings of the National Academy of Sciences (Figure 6).[4] The article, written by John Bell, Marc Day, Ian Shepherd, Matthew Johnson, Robert Cheng, Joseph Grcar, Vincent Beckner, and Michael Lijewski, describes the simulation of "a laboratory-scale turbulent rod-stabilized premixed methane V-flame." This simulation was unprecedented in several aspects—the number of chemical species included, the number of chemical reactions modeled, and the overall size of the flame.

This simulation employed a different mathematical approach than has typically been used for combustion. Most combustion simulations designed for basic research use compressible flow equations that include sound waves, and are calculated with small time steps on very fine, uniform spatial grids—all of which makes them very computationally expensive. Because of limited computer time, such simulations often have been restricted to only two dimensions, to scales less than a centimeter, or to just a few carbon species and reactions.

In contrast, the Center for Computational Sciences and Engineering (CCSE), under Bell's leadership, has developed an algorithmic approach that combines *low Mach-number equations,* which remove sound waves from the computation, with *adaptive mesh refinement,* which bridges the wide range of spatial scales relevant to a laboratory experiment. This combined methodology strips away relatively unimportant aspects of the simulation and focuses computing resources on the most important processes, thus slashing the computational cost of combustion simulations by a factor of 10,000.

Using this approach, the CCSE team has modeled turbulence and turbulence–chemistry interactions for a three-dimensional flame about 12 cm (4.7 in.) high, including 20 chemical species and 84 fundamental chemical reactions. The simulation was realistic enough to be compared directly with experimental diagnostics.

The simulation captured with remarkable fidelity some major features of the experimental data, such as flame-generated outward deflection in the unburned gases, inward flow convergence, and a centerline flow acceleration in the burned gases (Figure 7).

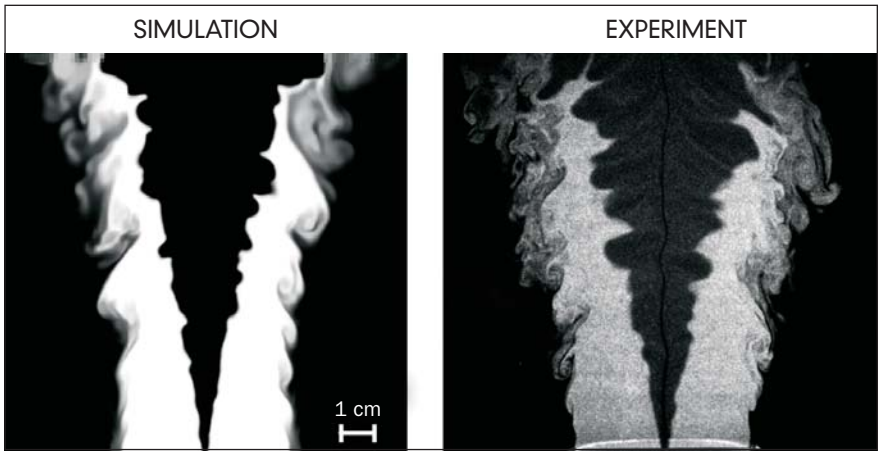

FIGURE 7. Left: A typical centerline slice of the methane concentration obtained from the simulation. Right: Experimentally, the instantaneous flame location is determined by using the large differences in Mie scattering intensities from the reactants and products to clearly outline the flame. The wrinkling of the flame in the computation and the experiment is of similar size and structure.

[4] J. B. Bell, M. S. Day, I. G. Shepherd, M. Johnson, R. K. Cheng, J. F. Grcar, V. E. Beckner, and M. J. Lijewski, "Numerical simulation of a laboratory-scale turbulent V-flame," Proceedings of the National Academy of Sciences **102**, 10006 (2005).

The simulation results were found to match the experimental results within a few percent. This agreement directly validated both the computational method and the chemical model of hydrocarbon reaction and transport kinetics in a turbulent flame.

The results demonstrate that it is possible to simulate a laboratory-scale flame in three dimensions without having to sacrifice a realistic representation of chemical and transport processes. This advance has the potential to greatly increase our understanding of how fuels behave in the complicated environments inside turbulent flames.

Research funding: BES, ASCR, SciDAC
Computational resources: NERSC
This article written by: John Hules, Jon Bashor

Combustion up close

Direct numerical simulations reveal small-scale structures and processes that experiments cannot measure

Efficient mixing is one of the basic challenges in turbulent combustion, especially in situations where the fuel and air streams are initially separated (so-called *nonpremixed combustion*), as in jet aircraft engines and direct-injection diesel and gasoline engines. The mixing of fuel and air molecules is necessary for the chemical reactions of combustion to take place.

Turbulence increases the mixing rate, so the right amount of turbulence improves the efficiency of combustion. But too much turbulent mixing can interrupt the chemical reaction process and lead to partial extinction of the flame. If extinguished pockets of unburned fuel–air mixture fail to reignite promptly, they may be emitted in the exhaust. Thus extinction may lead to reduced fuel efficiency, increased harmful emissions, and if pervasive enough, destabilization or blowout of the flame—a potentially catastrophic effect for aircraft.

Small-scale turbulence–chemistry effects such as extinction and reignition, flow and flame unsteadiness, and differential diffusion of chemical

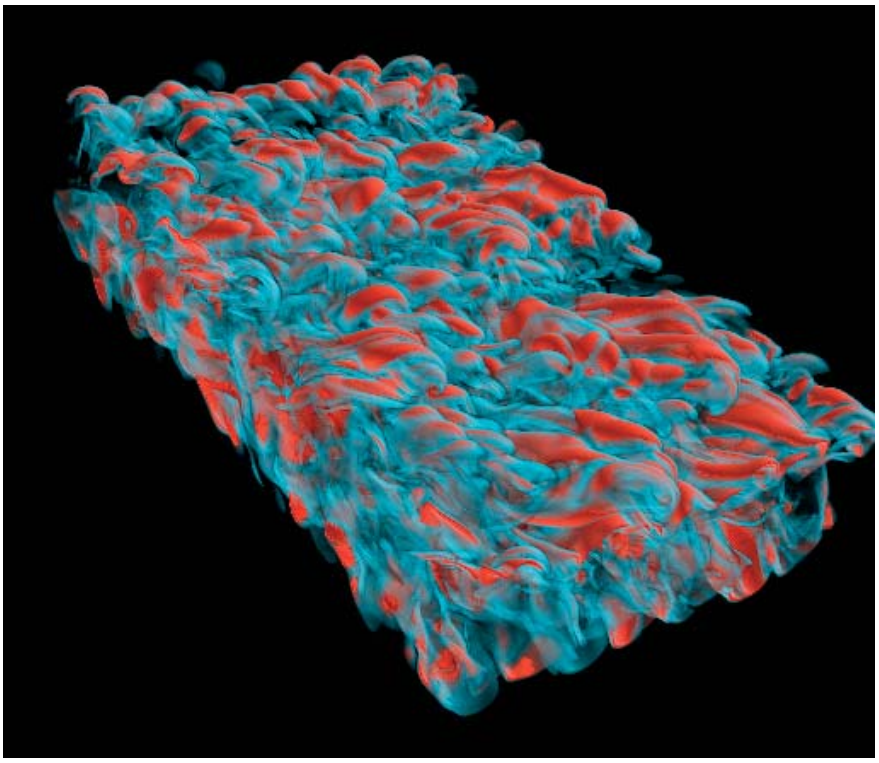

FIGURE 8. A simulated planar jet flame, colored by the rate of molecular mixing (scalar dissipation rate), which is critical for determining the interaction between reaction and diffusion in a flame. The image shows that high scalar dissipation regions exist in thin, highly intermittent structures aligned with principal strain directions. (Visualization created using an application written by Hongfeng Yu and Kwan-Liu Ma, UC Davis.)

species are difficult, if not impossible, to measure experimentally. Engineering-scale combustion simulations, in order to be computationally feasible, must use a spatial grid that is too coarse to capture these small-scale phenomena; to approximate their effects, mathematical models extrapolated from experimental results must be used. But the rapid growth of computational capabilities in recent years has presented new opportunities for high-resolution direct numerical simulations (DNS) of turbulent combustion that address unanswered questions about turbulence–chemistry interactions.

In the 2005 INCITE project "Direct Numerical Simulation of Turbulent Nonpremixed Combustion—Fundamental Insights towards Predictive Modeling," Jacqueline Chen, Evatt Hawkes, and Ramanan Sankaran of Sandia National Laboratories have performed the first 3D DNS simulations of a turbulent nonpremixed H_2/CO–air flame with detailed chemistry (Figure 8). The simulations, including 11 chemical species and 33 reactions, were performed with up to 500 million grid points and 100,000 time steps, and ran for 2.5 million processor hours on Seaborg and 1.5 million hours on Bassi.

An initial analysis of the data shows for the first time how detailed transport and chemistry effects can influence the mixing of chemical molecules.[5] (The mathematical representations of these molecules are called *scalars*.) This is an example of how DNS can be used to assess the validity of assumptions used in combustion models. Mixing models typically assume that the mixing rate is the same for each different scalar and corresponds to the turbulence time scale. In fact, the INCITE simulation reveals that the mixing rates of active (reacting) scalars and passive scalars may vary by a factor of 3. Models may need to incorporate these different mixing rates, because a poor mixing model could incorrectly predict a stable flame when actually extinction occurs. This finding and others yet to come from further analysis may lead to more accurate models for use in lower-resolution simulations.

To ramp up to the large INCITE calculations, Chen's team started by running many test calculations with their highly scalable S3D code, using different numbers of grid points on a variety of large-scale computing systems. After analyzing and optimizing the code's performance with the help of NERSC staff, they improved the code's efficiency by 45%.

The code improvements helped the researchers achieve the highest-ever Reynolds number (ratio of inertial to viscous forces) in a 3D fully resolved DNS of a nonpremixed flame with detailed chemistry. Simulations with different Reynolds numbers shed light on a phenomenon called intermittency—intense, localized fluctuations of any quantity in a turbulent flow—which can result in localized extinction and reignition of the combustion process. The results show a relationship between scalar intermittency and the Reynolds number (Figure 9).

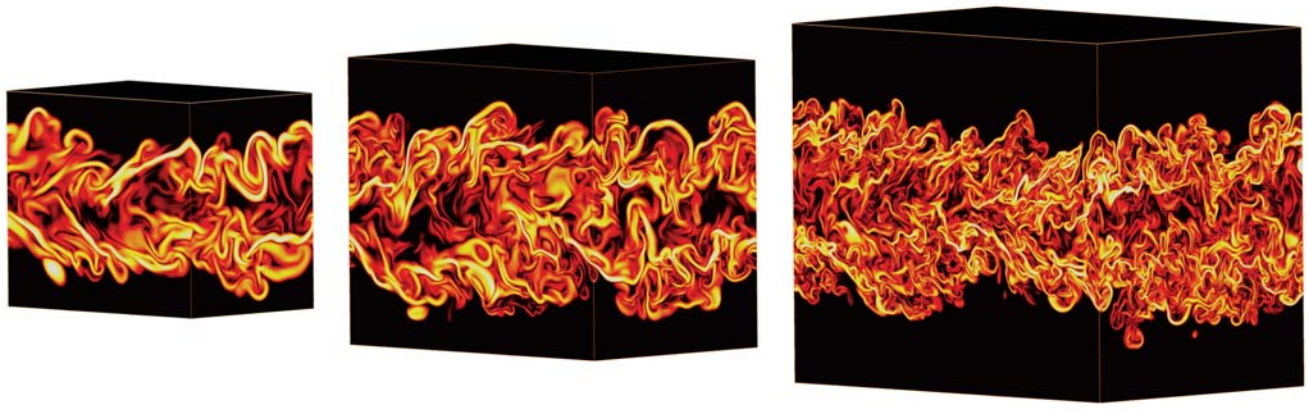

FIGURE 9. Instantaneous isocontours of the total scalar dissipation rate field for successively higher Reynolds numbers at a time when reignition following extinction in the domain is significant. The dissipation fields are organized into thin sheet-like lamellar structures, with lengths far exceeding their thickness, consistent with experimental observations in nonreactive flows. Increasingly fine-scaled structures are observed at higher Reynolds numbers. (From E. R. Hawkes, R. Sankaran, J. C. Sutherland, and J. H. Chen, "Direct Numerical Simulation of Temporally-Evolving Plane Jet Flames with Detailed CO/H_2 Kinetics," submitted to the 31st International Symposium on Combustion, 2006.)

[5] E. R. Hawkes, R. Sankaran, J. C. Sutherland, and J. H. Chen, "Direct numerical simulation of turbulent combustion: Fundamental insights towards predictive models," Journal of Physics Conference Series **16** (SciDAC 2005), 65 (2005).

This project has generated 10 terabytes of raw DNS data, which they intend to share with an ongoing international collaboration of experimentalists and modelers called the Turbulent Nonpremixed Flame Workshops.6 The DNS data will serve as a numerical benchmark, complementary to experimental data, for model validation and the advancement of basic scientific understanding of turbulent combustion.

Extracting and analyzing useful information from this massive dataset is a daunting task. Chen's group is working with Kwan-Liu Ma of the University of California at Davis to develop automated feature extraction tools and multi-variable visualization capabilities that will help researchers understand how turbulent mixing interacts with chemical reactions. Phenomena they are working to visualize include flame edges and cusps, extinction pockets, and ignition kernels.

Research funding: BES, INCITE, SciDAC
Computational resources: NERSC, MSCF, NCCS
This article written by: John Hules

Hailstones in hell

Simulations examine the behavior of frozen fuel pellets in fusion reactors

What happens when you shoot one of the coldest materials into the hottest environment on earth? The answer may help solve the world's energy crisis.

Imagine that there is a large chamber in hell that's shaped like a doughnut, and that you'd like to shoot a series of hailstones into that chamber so that as they melt, the water vapor penetrates as deeply as possible and disperses as evenly as possible throughout the chamber. Setting aside for a moment the question of *why* you would want to do that, consider the multitude of *how* questions: Should you shoot in the hailstones from outside the ring of the doughnut or from inside the hole? What size hailstones should you use? At what speed and angle should they enter the chamber?

This strange scenario is actually an analogy for one of the questions facing fusion energy researchers: how to refuel a tokamak. A tokamak is a machine that produces a toroidal (doughnut-shaped) magnetic field. In that field, two isotopes of hydrogen—deuterium and tritium—are heated to about 100 million degrees Celsius (more than six times hotter than the interior of the sun), stripping the electrons from the nuclei. The magnetic

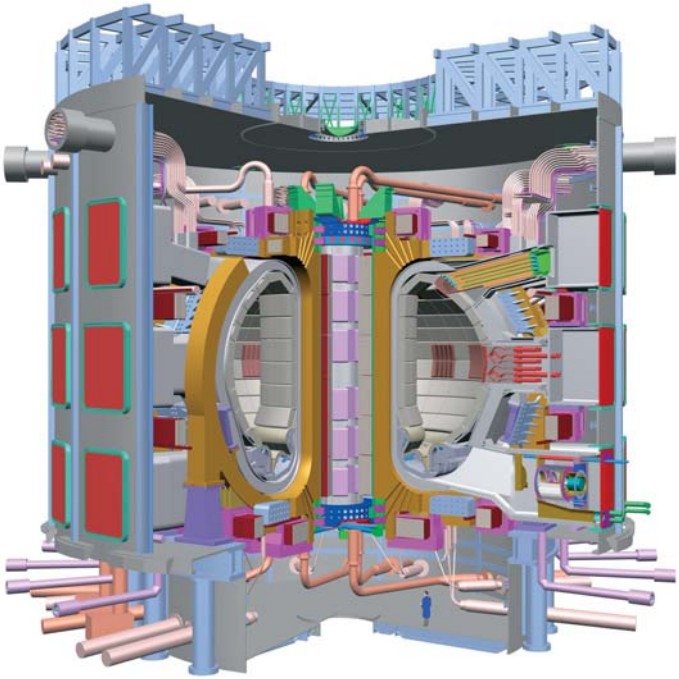

FIGURE 10. Cutaway illustration of the ITER tokamak.

[6] http://www.ca.sandia.gov/TNF/abstract.html

field makes the electrically charged particles follow spiral paths around the magnetic field lines, so that they spin around the torus in a fairly uniform flow and interact with each other, not with the walls of the tokamak. When the hydrogen ions (nuclei) collide at high speeds, they fuse, releasing energy. If the fusion reaction can be sustained long enough that the amount of energy released exceeds the amount needed to heat the plasma, researchers will have reached their goal: a viable energy source from abundant fuel that produces no greenhouse gases and no long-lived radioactive byproducts.

High-speed injection of frozen hydrogen pellets is an experimentally proven method of refueling a tokamak. These pellets are about the size of small hailstones (3–6 mm) and have a temperature of about 10 degrees Celsius above absolute zero. The goal is to have these pellets penetrate as deeply as possible into the plasma so that the fuel disperses evenly.

Pellet injection will be the primary fueling method used in ITER (Latin for "the way"), a multinational tokamak experiment to be built at Cadarache in southern France (Figure 10). ITER, one of the highest strategic priorities of the DOE Office of Science, is expected to produce 500 million thermal watts of fusion power—10 times more power than is needed to heat the plasma—when it reaches full operation around the year 2016. As the world's first production-scale fusion reactor, ITER will help answer questions about the most efficient ways to configure and operate future commercial reactors.

However, designing a pellet injection system that can effectively deliver fuel to the interior of ITER represents a special challenge because of its unprecedented large size and high temperatures. Experiments have shown that a pellet's penetration distance into the plasma depends strongly on how the injector is oriented in relation to the torus. For example, an "inside launch" (from inside the torus ring) results in better fuel distribution than an "outside launch" (from outside the ring).

In the past, progress in developing an efficient refueling strategy for ITER has required lengthy and expensive experiments. But thanks to a three-year, SciDAC-funded collaboration between the Computational Plasma Physics Theory Group at Princeton Plasma Physics Laboratory and the Advanced Numerical Algorithms Group at Lawrence Berkeley National Laboratory, computer codes have now reproduced some key experimental findings, resulting in significant progress toward the scientific goal of using simulations to predict the results of pellet injection in tokamaks.

"To understand refueling by pellet injection, we need to understand two phases of the physical process," said Ravi Samtaney, the Princeton researcher who is leading the code development effort. "The first phase is the transition from a frozen pellet to gaseous hydrogen, and the second phase is the distribution of that gas in the existing plasma."

The first phase is fairly well understood from experiments and theoretical studies. In this phase, called *ablation,* the outer layer of the frozen hydrogen pellet is quickly heated, transforming it from a solid into an expanding cloud of dense hydrogen gas surrounding the pellet. This gas quickly heats up, is ionized, and merges into the plasma. As ablation continues, the pellet shrinks until all of it has been gasified and ionized.

The second phase—the distribution of the hydrogen gas in the plasma—is less well understood. Ideally, the injected fuel would simply follow the magnetic field lines and the "flux surfaces" that they define, maintaining a stable and uniform plasma pressure.

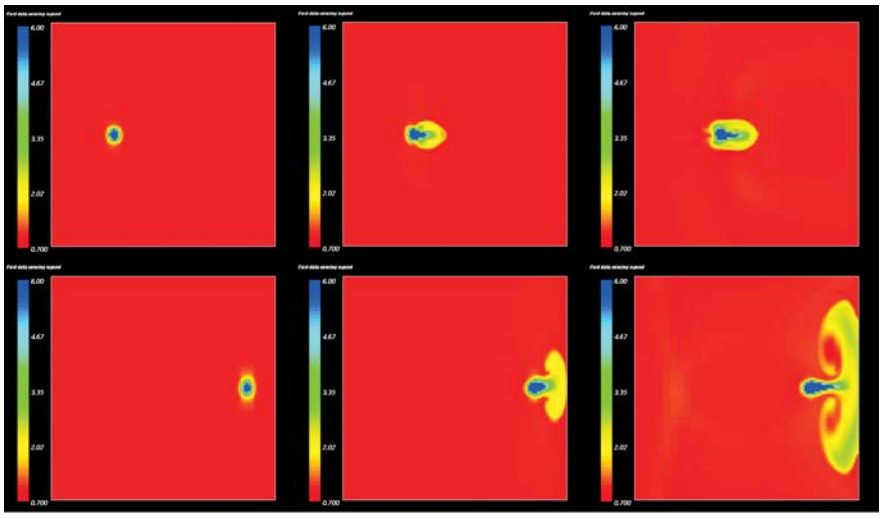

FIGURE 11. Time sequence of 2D slices from a 3D simulation of the injection of a fuel pellet into a tokamak plasma. Injection from outside the torus (bottom row, injection from right) results in the pellet stalling and fuel being dispersed near the plasma boundary. Injection from inside the torus (top row, injection from left) achieves fuel distribution in the hot central region as desired.

But experiments have shown that the high-density region around the pellet quickly heats up to form a local region of high pressure, higher than can be stably confined by the local magnetic field. A form of "local instability" (like a mini-tornado) then develops, causing the high-density region to rapidly move across, rather than along, the field lines and flux surfaces—a motion referred to as "anomalous" because it deviates from the large-scale motion of the plasma.

Fortunately, researchers have discovered that they can use this instability to their advantage by injecting the pellet from inside the torus ring, because from this starting point, the anomalous motion brings the fuel pellet closer to the center of the plasma, where it does the most good. This anomalous motion is one of the phenomena that Samtaney and his colleagues want to quantify and examine in detail.

Figure 11 shows the fuel distribution phase as simulated by Samtaney and his colleagues in the first detailed 3D calculations of pellet injection.[7] The inside launch (top row) distributes the fuel in the central region of the plasma, as desired, while the outside launch (bottom row) disperses the fuel near the plasma boundary, as shown in experiments.

Simulating pellet injection in 3D is difficult because the physical processes span several decades of time and space scales. The large disparity between pellet size and tokamak size, the large density differences between the pellet ablation cloud and the ambient plasma, and the long-distance effects of electron heat transport all pose severe numerical challenges. To overcome these difficulties, Samtaney and his collaborators used an algorithmic method called adaptive mesh refinement (AMR), which incorporates a range of scales that change dynamically as the calculation progresses. AMR allowed this simulation to run more than a hundred times faster than a uniform-mesh simulation.

While these first calculations represent an important advance in methodology, Samtaney's work on pellet injection is only beginning. "The results presented in this paper did not include all the detailed physical processes which we're starting to incorporate, along with more realistic physical parameters," he said. "For example, we plan to develop models that incorporate the perpendicular transport of the ablated mass. We also want to investigate other launch locations. And, of course, we'll have to validate all those results against existing experiments."

This pellet injection model will eventually become part of a comprehensive predictive capability for ITER, which its supporters hope will bring fusion energy within reach as a commercial source of electrical power.

Research funding: FES, SciDAC
Computational resources: NERSC
This article written by: John Hules

A perfect liquid

The newly created state of quark-gluon matter is more remarkable than predicted and raises many new questions

The four experimental groups conducting research at the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory have created a new state of hot, dense matter out of the quarks and gluons that are the basic particles of atomic nuclei, but it is a state quite different and even more remarkable than had been predicted. In an announcement at the April 2005 meeting of the American Physical Society in Tampa, Florida, and in peer-reviewed papers summarizing the first three years of RHIC findings, the scientists said that instead of behaving like a gas of free quarks and gluons, as was expected, the matter created in RHIC's heavy ion collisions appears to be more like a liquid.

The papers, which the four RHIC collaborations (BRAHMS, PHENIX, PHOBOS, and STAR) had been working on for nearly a year, were published simultaneously by the journal Nuclear Physics A[8] and were also compiled in a special Brookhaven report.[9] These summaries indicate that some of the observations at RHIC fit with the theoretical predictions for a quark-gluon plasma (QGP), the type of matter postulated to have existed just microseconds after the Big Bang. Indeed, many

[7] R. Samtaney, S. C. Jardin, P. Colella, and D. F. Martin, "3D adaptive mesh refinement simulations of pellet injection in tokamaks," Computer Physics Communications **164**, 220 (2004).
[8] Nuclear Physics A, Volume 757, Issues 1–2, August 8, 2005.
[9] BRAHMS, PHENIX, PHOBOS, and STAR Collaborations, *Hunting the Quark Gluon Plasma: Results from the First 3 Years at RHIC* (Upton, NY: Brookhaven National Laboratory report No. BNL-73847-2005).

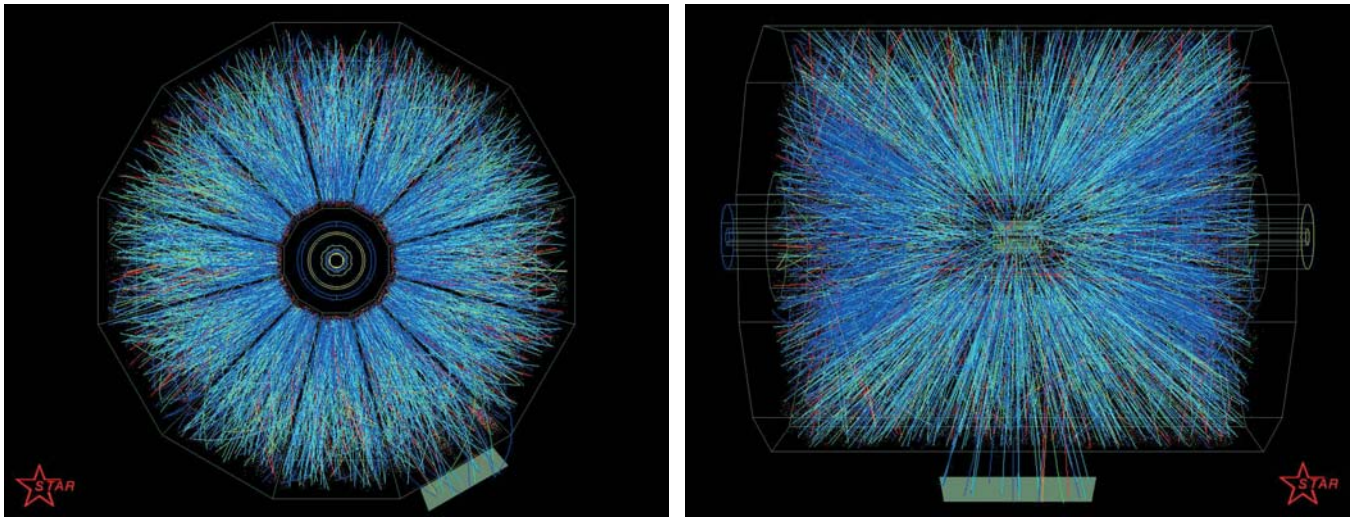

FIGURE 12. Two views of one of the first full-energy collisions between gold ions at Brookhaven Lab's Relativistic Heavy Ion Collider, as captured by the Solenoidal Tracker At RHIC (STAR) detector. The tracks indicate the paths taken by thousands of subatomic particles produced in the collisions as they pass through the STAR Time Projection Chamber, a large, 3D digital camera.

theorists have concluded that RHIC has already demonstrated the creation of quark-gluon plasma. However, all four collaborations note that there are discrepancies between the experimental data and early theoretical predictions based on simple models of quark-gluon plasma formation.

"We know that we've reached the temperature [up to 150,000 times hotter than the center of the sun] and energy density [energy per unit volume] predicted to be necessary for forming such a plasma," said Sam Aronson, Brookhaven's Associate Laboratory Director for High Energy and Nuclear Physics. But analysis of RHIC data from the start of operations in June 2000 through the 2003 physics run reveals that the matter formed in RHIC's head-on collisions of gold ions (Figure 12) is more like a liquid than a gas.

That evidence comes from measurements of unexpected patterns in the trajectories taken by the thousands of particles produced in individual collisions. These measurements indicate that the primordial particles produced in the collisions tend to move collectively in response to variations of pressure across the volume formed by the colliding nuclei (Figure 13).

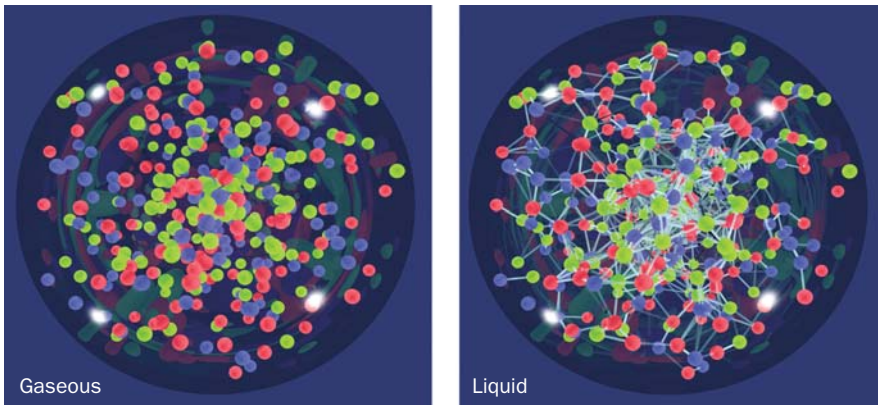

FIGURE 13. These images contrast the degree of interaction and collective motion, or "flow," among quarks in the predicted gaseous quark-gluon plasma state (left) vs. the liquid state that has been observed in gold-gold collisions at RHIC (right). The "force lines" and collective motion in the observed collisions show the much higher degree of interaction and flow among the quarks in what is now being described as a nearly "perfect" liquid. An animation demonstrating the differences between the expected gas and the observed liquid can be viewed at http://real.bnl.gov/ramgen/bnl/RHIC_animation.rm.

Scientists refer to this phenomenon as "flow," since it is analogous to the properties of fluid motion.

However, unlike ordinary liquids, in which individual molecules move

about randomly, the hot matter formed at RHIC seems to move in a pattern that exhibits a high degree of coordination among the particles—somewhat like a school of fish that responds as one entity while moving through a changing environment.

"This is fluid motion that is nearly 'perfect,'" Aronson said, meaning it can be explained by equations of hydrodynamics. These equations were developed to describe theoretically "perfect" fluids—those with extremely low viscosity and the ability to reach thermal equilibrium very rapidly due to the high degree of interaction among the particles. While RHIC scientists don't have a direct measure of viscosity, they can infer from the flow pattern that, qualitatively, the viscosity is very low, approaching the quantum mechanical limit.

Together, these facts present a compelling case. "In fact, the degree of collective interaction, rapid thermalization, and extremely low viscosity of the matter being formed at RHIC make this the most nearly perfect liquid ever observed," Aronson said.

In results reported earlier, other measurements at RHIC have shown "jets" of high-energy quarks and gluons being dramatically slowed down as they traverse the hot fireball produced in the collisions. This "jet quenching" demonstrates that the energy density in this new form of matter is extraordinarily high—much higher than can be explained by a medium consisting of ordinary nuclear matter.

"The current findings don't rule out the possibility that this new state of matter is in fact a form of the quark-gluon plasma, just different from what had been theorized," Aronson said. Many scientists believe this to be the case, and detailed measurements are now under way at RHIC to resolve this question.

Theoretical physicists, whose standard calculations cannot incorporate the strong coupling observed between the quarks and gluons at RHIC, are also revisiting some of their early models and predictions. To try to address these issues, they are running massive numerical simulations on some of the world's most powerful computers. Others are attempting to incorporate quantitative measures of viscosity into the equations of motion for fluid moving at nearly the speed of light. One subset of calculations uses the methods of string theory to predict the viscosity of the liquid being created at RHIC and to explain some of the other surprising findings. Such studies will provide a more quantitative understanding of how "nearly perfect" the liquid is.

The unexpected findings also introduce a wide range of opportunity for new scientific discovery regarding the properties of matter at extremes of temperature and density previously inaccessible in a laboratory.

The STAR Collaboration, consisting of 614 researchers from 52 institutions in 12 countries, performs its data analysis on NERSC's PDSF system. The computations carried out at NERSC are focused around analysis of the processed data, comparison of the data with experimental models, and studies of detector performance and acceptance. A storage resource allocation at NERSC supports an important collaboration between STAR, the Berkeley Lab Scientific Data Management Group, and the Particle Physics Data Grid (PPDG) project, involving development and deployment of data grid technology to move terabytes of data efficiently and securely.

One of STAR's next priority goals at NERSC is to study the spin structure of the proton, focusing first on measurements of the gluon contribution.

Research funding (STAR only): NP, NSF, BMBF, IN2P3, RA, RPL, EMN, EPSRC, FAPESP, RMST, MEC, NNSFC, GACR, FOM, DAE, DST, CSIR, SNSF, PSCSR, STAA
Computational resources: NERSC, BNL
This article written by: Karen McNulty Walsh (Brookhaven National Laboratory), John Hules

Whispers from underground

Ghostly geoneutrinos hint at Earth's inner secrets

Neutrinos and antineutrinos are elusive particles. With no charge and almost no mass, they rarely react with other matter, even as they zip through the Earth (and us) at nearly the speed of light. Recording the traces of rare neutrino interactions requires specialized, large-volume detectors, and finding those few traces among the mass of collected data is a little like listening for a whisper at a rock concert.

But the same qualities that make neutrinos elusive also make them valuable tools for probing environments, like the core of the Earth, that resist penetration by more common observational techniques.

So when an international group of researchers reported the first observation of geologically produced antineutrinos in a paper[10] that was featured on the cover of *Nature* (Figure 14), the geological community responded with enthusiasm. This discovery is expected to lead to better estimates of the abundance and distribution of radioactive elements in the Earth, and of the Earth's overall heat budget.

The geoneutrinos were detected at the Kamioka Liquid Scintillator Anti-Neutrino Detector in Japan, better known as KamLAND. Located in a mine cavern beneath the mountains of Japan's main island of Honshu, near the city of Toyama, KamLAND is the largest low-energy antineutrino detector ever built. It consists of a weather balloon 13 meters (43 feet) in diameter, filled with about a kiloton

FIGURE 14. The July 28, 2005 issue of *Nature* reported that electron antineutrinos emanating from the earth (geoneutrinos) may serve as a unique window into the interior of our planet, revealing information that is hidden from other probes. The left half of the cover image shows the production distribution for the geoneutrinos detected at KamLAND, and the right half shows the geologic structure.

of liquid scintillator, a chemical soup that emits flashes of light when an incoming antineutrino collides with a proton. These light flashes are detected by a surrounding array of 1,879 photomultiplier light sensors which convert the flashes into electronic signals that computers can analyze.

Most of the geoneutrino data produced at KamLAND was stored on the High Performance Storage System (HPSS) at NERSC and analyzed using NERSC's PDSF cluster. Together, these systems allowed scientists to find the scientific equivalent of a whisper at a rock concert.

KamLAND records data 24 hours a day, seven days a week. This data is shipped on tapes from the experimental site to Berkeley Lab, where it is read off the tapes and stored in the HPSS. KamLAND records about 200 GB of data each day, and HPSS currently has more than 250 TB of KamLAND data stored, making KamLAND the second-largest user of NERSC's HPSS system.

During dedicated production periods at NERSC, the KamLAND data are read out of HPSS and run through reconstruction software to convert the waveforms (essentially oscilloscope traces) from the photomultiplier tubes into physically meaningful quantities such as energy and position of the event inside the detector. This reduces the data volume by a factor of 60–100, and the reconstructed events are stored on disk for further analysis.

"The event reconstruction requires a lot of computing power, and with over 600 CPUs, PDSF is a great facility to run this kind of analysis," said Patrick Decowski, a Berkeley Lab physicist who works with NERSC staff on the project. "PDSF has been essential for our measurements."

With the data on disk, specialized analysis programs run over the reconstructed events to extract the geoneu-

[10] T. Araki et al., "Experimental investigation of geologically produced antineutrinos with KamLAND," Nature **436**, 499 (2005).

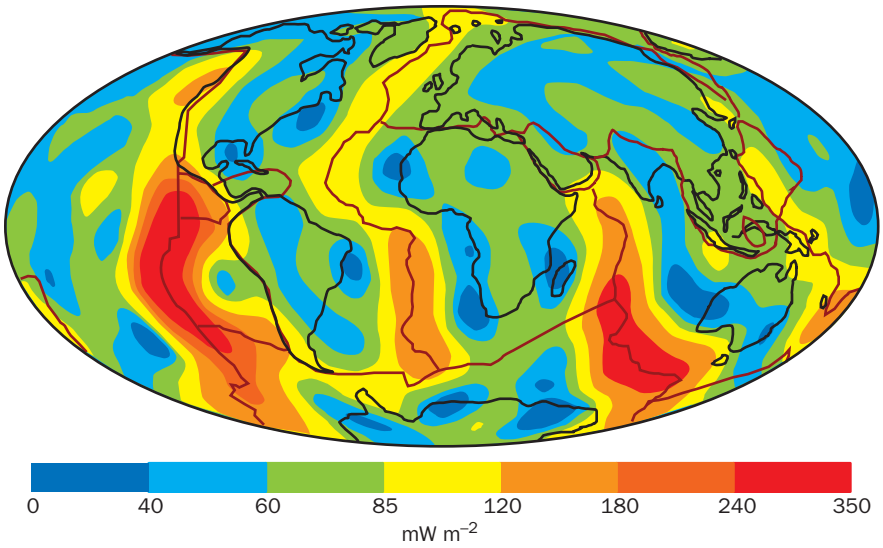

FIGURE 15. Earth's conductive heat flow is estimated to be about 30–45 terawatts. Radioactivity is known to account for perhaps half of this heat, but there has been no accurate way to measure radiogenic heat production. (Image from H. N. Pollack, S. J. Hurter, and J. R. Johnson, Reviews of Geophysics **31,** 267 [1993])

trino events and perform the final analysis. PDSF is also used for various simulation tasks in order to better understand the background signals in the detector.

"The whole analysis is like looking for a needle in a haystack—out of more than 2 billion events, only 152 candidates were found," Decowski said. "And of these, about 128 are background events."

Forty years ago, the late John Bahcall proposed the study of neutrinos coming from the sun to understand the fusion processes inside the sun. The measurement of a persistent deficit of the observed neutrino flux relative to Bahcall's calculations led to the 2002 Nobel Prize for Ray Davis and the discovery of neutrino oscillation, a quantum mechanical phenomenon whereby a neutrino created with a specific lepton "flavor" (electron, muon, or tau) can later be measured to have a different flavor. KamLAND provided the most direct evidence of neutrino oscillation in 2004.

Today, antineutrinos are being used to study the interior of the Earth, which is still little known. The deepest borehole ever drilled is less than 20 km in depth, while the radius of the Earth is more than 6000 km. While seismic events have been used to deduce the interior composition of the Earth's three basic regions—the core, the mantle and the crust—there are no direct measurements of the chemical makeup of the deeper regions.

An important observation for understanding the Earth is the measurement of the heat coming from within (Figure 15). These measurements show that the Earth produces somewhere between 30 and 45 terawatts of heat (1 TW is 10^{12} watts; the range depends on certain model assumptions). Two important sources of heat generation are the primordial energy released from planetary accretion and latent heat from core solidification. However, it is believed that heat produced by radioactivity also plays an important role in the Earth's heat balance, contributing perhaps half of the total heat.

Neutrinos can help researchers understand the Earth's internal structure and heat generation. Three important isotopes that are part of current Earth models—potassium, uranium, and thorium—produce electron antineutrinos (the so-called geoneutrinos) in their radioactive decay.

"KamLAND is the first detector sensitive enough to measure geoneutrinos produced in the earth from the decay of uranium-238 and thorium-232," said Stuart Freedman, a nuclear physicist with a joint appointment at Berkeley Lab and the University of California at Berkeley, who is a co-spokesperson for the U.S. team at KamLAND. "Since the geoneutrinos produced from the decay chains of these isotopes have exceedingly small interaction cross sections, they propagate undisturbed in the Earth's interior, and their measurement near the Earth's surface can be used to gain information on their sources."

In measuring geoneutrinos generated in the decay of natural radioactive elements in the earth's interior, scientists believe it should be possible to get a three-dimensional picture of the earth's composition and shell structure. This could provide answers to such questions as how much terrestrial heat comes from radioactive decays and how much is a primordial remnant from the birth of our planet. It might also help identify the source of Earth's magnetic field and what drives the geodynamo.

The research is a multinational effort, as shown by the fact that the Nature article represented the work of 87 authors from 14 institutions spread across four nations.

Research funding: NP, JAMSTEC
Computational resources: NERSC
This article written by: John Hules, Jon Bashor, Lynn Yarris

Breaking up is hard to calculate

But researchers have achieved the first complete numerical solution of the fragmentation of a system with four charged particles

Need to understand the details of how a molecule is put together? Want to see the effects of the intricate dance that its electrons do to make a chemical bond? Try blowing a molecule to bits and calculating what happens to all the pieces. That's the approach taken by an international group of collaborators from the University of California at Davis, universities in Spain and Belgium, and Lawrence Berkeley National Laboratory.

When a hydrogen molecule, H_2, is hit by a photon with enough energy to send both its electrons flying, the two protons left behind—the hydrogen nuclei—repel each other in a so-called Coulomb explosion. In this event, called the double photoionization of H_2, the paths taken by the fleeing electrons have much to say about how close together the two nuclei were at the moment the photon struck, and just how the electrons were correlated in the molecule (Figure 16).

Correlation means that properties of the particles like position and momentum cannot be calculated independently. When three or more particles are involved, calculations are notoriously intractable, both in classical physics and quantum mechanics. But in the December 16, 2005 issue of Science, the researchers reported on the first-ever complete quantum mechanical solution of a system with four charged particles.[11]

The groundbreaking calculations were inspired by earlier experiments on the photofragmentation of deuterium (heavy hydrogen) molecules, performed at Berkeley Lab's Advanced Light Source (ALS) in 2003 by a group of scientists from Germany, Spain, and several institutions in the United States. The experimenters were led by Thorsten Weber, then with the ALS and now at the University of Frankfurt.

"If you were trying to do this experiment and you didn't have access to the Advanced Light Source and a COLTRIMS experimental device"—a sophisticated, position-sensitive detector for collecting electrons and ions—"you'd just fire photons at a random sample of hydrogen molecules and measure the electrons that came out," said Thomas Rescigno of Berkeley Lab's Chemical Sciences Division, one of the authors of the *Science* paper. "What made this experiment special was that they could measure what happened to all four particles. From their precise positions and energy they could reconstruct the state of the molecule when it was hit."

Weber presented early experimental data at a seminar attended by Rescigno, William McCurdy of the Lab's Chemical Sciences Division, who is also a professor of chemistry at the University of California at Davis, and Wim Vanroose, a postdoctoral fellow at Berkeley Lab who is now at the Department of Computer Science at the Katholieke Universiteit Leuven in Belgium.

Said Rescigno, "Thorsten teased us with his results, some of which were extremely nonintuitive. What was remarkable was that very small differences in the internuclear distance"—the distance between the two protons at the moment the photon was absorbed—"made for radical differences in the ways the electrons were ejected."

"When I saw the results of the molecular experiments, in which small changes in the internuclear distance produced large and unexpected changes in the electron ejection patterns, it immediately occurred to me that the differences were because of the molecule's effects on electron correlations," McCurdy said.

McCurdy had recently been working with Fernando Martín, a professor of chemistry at the Universidad Autónoma de Madrid, merging computational techniques developed by Martín with a method McCurdy, Rescigno, and others had developed for calculating systems of three charged particles. Martín and McCurdy extended these methods to the helium atom, a system that, technically speaking, has four charged particles. But because the helium atom's two protons are bound together in the nucleus, the calculated distribution of electrons ejected by the absorption of an energetic photon tend to be quite symmetrical around the nucleus, with most pairs flying off in opposite directions.

The picture can look quite different for a hydrogen or deuterium molecule, in which a plot of the likelihood that

[11] Wim Vanroose, Fernando Martín, Thomas N. Rescigno, and C. William McCurdy, "Complete photo-induced breakup of the H_2 molecule as a probe of molecular electron correlation," Science **310**, 1787 (2005).

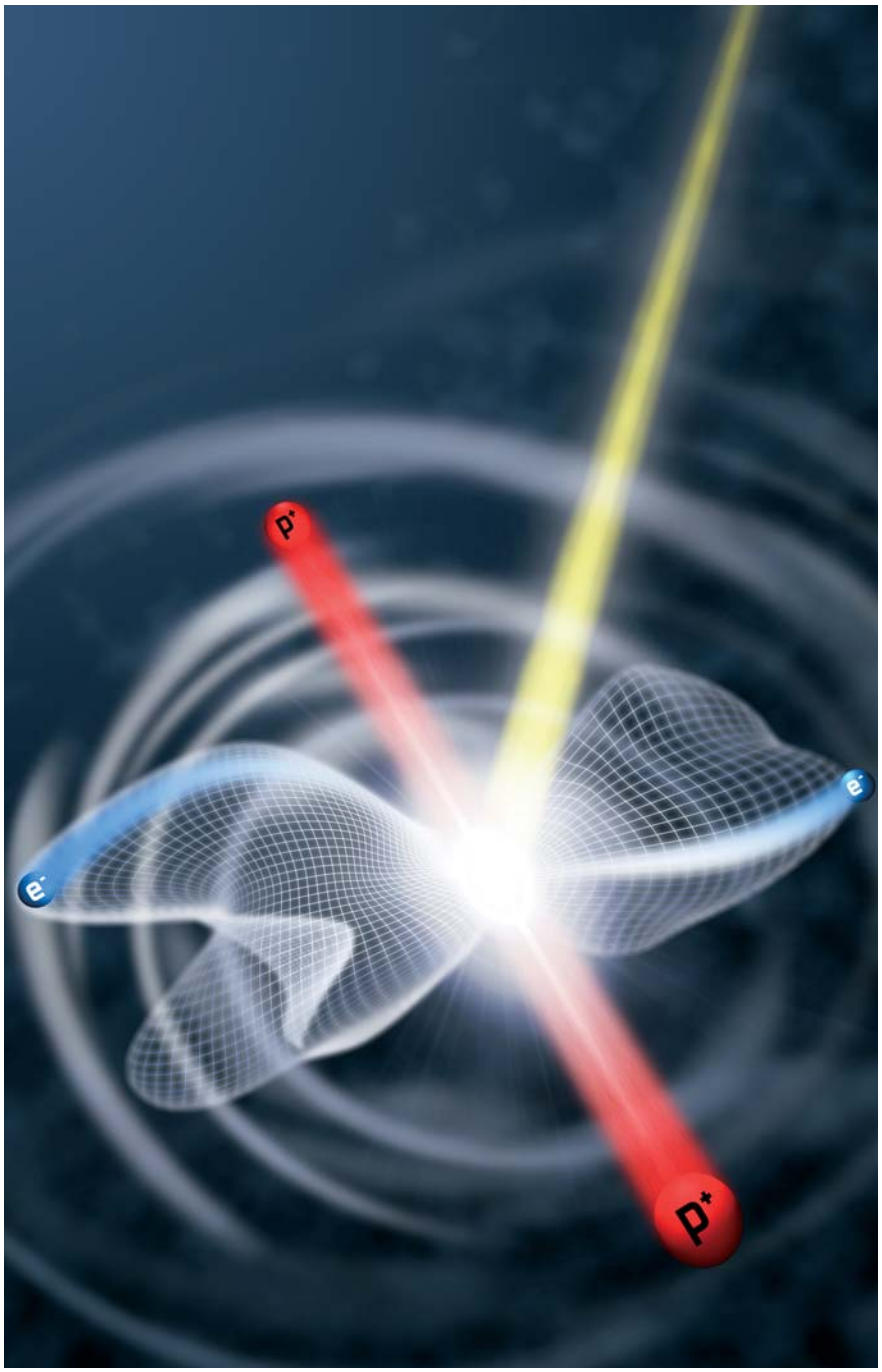

FIGURE 16. A hydrogen molecule hit by an energetic photon breaks apart. The ejected electrons, blue, take paths whose possible trajectories (here represented by net-like lobes) depend on how far apart the hydrogen nuclei, red, are at the moment the photon strikes. The bond length at that instant reflects how the molecule's electrons are correlated. (Courtesy Wim Vanroose)

electrons will be ejected at certain angles groups into lobes that grow increasingly asymmetric as the bond length between the two hydrogen atoms grows longer. McCurdy read this as the effect of the bond length on the correlation of the shared electrons. Indeed, this is what Weber and his colleagues speculated when they published the results of their deuterium photofragmentation studies in Nature in 2004.

Rescigno pointed out a fly in the ointment, however—namely that instead of being caused by electron correlations, large differences in ejection patterns caused by small differences in internuclear distance "could just be kinematics."

In other words, the scattered electrons might be sharing some of the potential energy stored by the Coulomb repulsion between the two like-charged protons. The closer together these two nuclei are at the moment the photon breaks up the molecule, the more energy goes into the Coulomb explosion, some of which could be transmitted to the outgoing electrons and affect their flight paths.

How to decide between kinematic effects or electron correlations? The experimental results could not address the question, since all the data were collected at the same photon energy; whether the electrons were acquiring additional kinetic energy was unknown.

But, said Vanroose, "because we were doing computations, we could do experiments the experimenters couldn't do. We had much more flexibility to fix the conditions."

Using supercomputers at NERSC, at UC Berkeley, and in Belgium, Vanroose was able to rerun the hydrogen molecule experiments "in silico," this time with different photon energies, distributed so that the outgoing elec-

trons always shared exactly the same kinetic energy no matter what the distance between the protons at the moment of photon absorption.

The results turned out to be remarkably similar in all cases. Even when kinetic energy made no contribution, the electrons flew off in patterns determined by the length of the bond between the nuclei (Figure 17). Therefore the differences were due almost entirely to the way the electrons were correlated in their orbital paths around the molecule's two nuclei.

Martín of UA Madrid sees the new calculations, which are a complete numerical solution for the Schrödinger equation of the photoionization of H_2, as "just the beginning. Probing the complicated physics of electron correlations will lead the way to more comprehensive methods combining theory and experiment to address some of the most pressing problems in chemistry."

Vanroose credits their success to day-in, day-out collaboration between top-notch theorists and experimenters "who are talking to each other all the time. The ability of experimentalists to call on the latest computational techniques is good for both; it's why we're two years ahead of other theorists in this field."

To Rescigno, the latest results show that "what began as blue-sky physics theory is now connecting with the nuts and bolts of practical experiment."

Said McCurdy, "These large-scale theoretical calculations, stimulated by the need to interpret novel experiments at the ALS, are already stimulating new experiments and establishing a new line of inquiry at Berkeley Lab."

Research funding: BER, MECS, BSPO
Computational resources: NERSC, UCB, KUL
This article written by: Paul Preuss

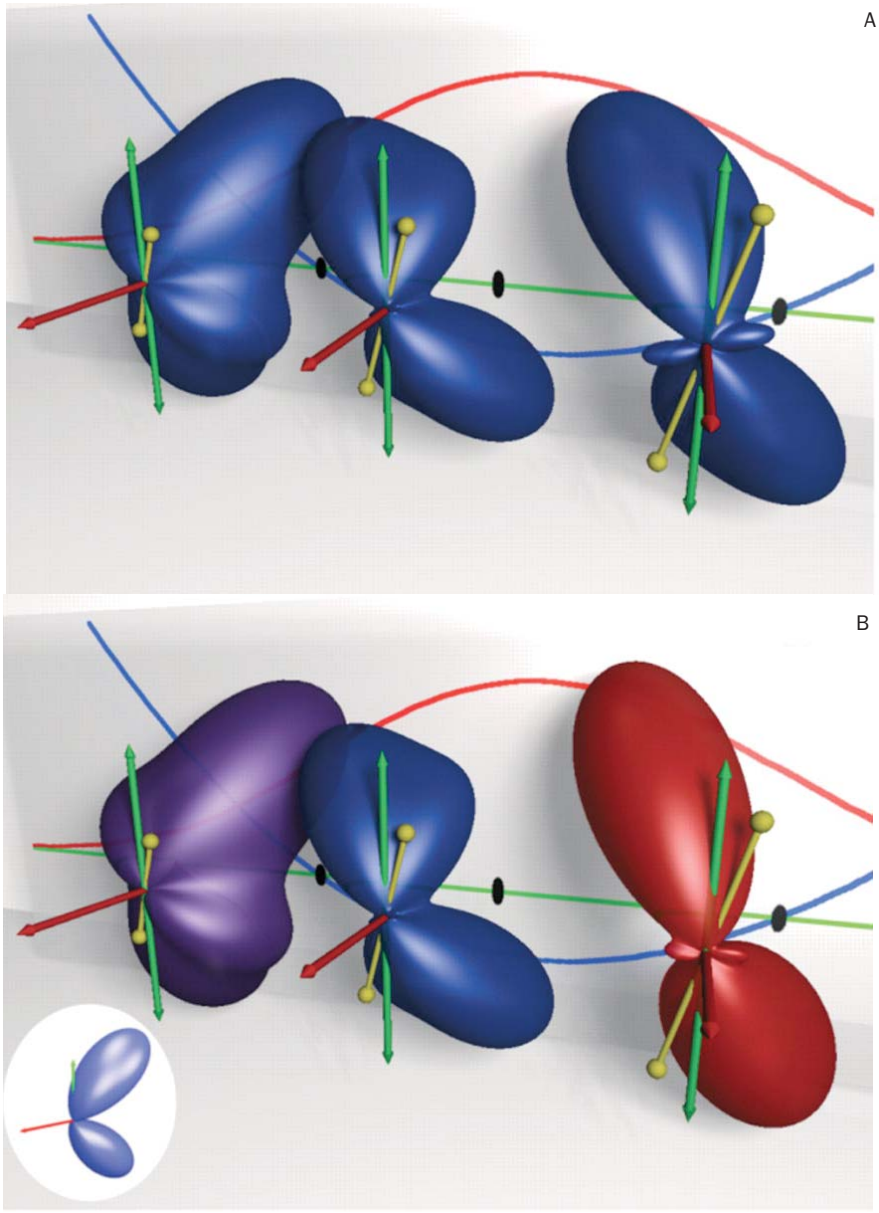

FIGURE 17. Effect of the internuclear distance on the possible electron trajectories (angular distributions, shown as lobes). (A) Angular distributions shown with molecular axis (yellow) 15° from polarization vector (green arrow) and fixed electron (red arrow) leaving perpendicular to that plane with 50% of the available ejection energy. (Left) Inner vibrational turning point; (middle) equilibrium internuclear distance; (right) outer turning point. Background shows ground state potential curve (blue), ground state vibrational wave function (red), and energy (green line, with black dots indicating equilibrium and inner and outer turning points). (B) The effect on the angular distribution of varying internuclear distance is primarily due to changes in molecular electronic correlation. The angular distributions closely resemble those in (A) when the photon energy is varied, so as to produce the same amount of final kinetic energy to be shared by the outgoing electrons, regardless of the internuclear distance at which ionization occurs. Inset shows the corresponding atomic case (helium).

Talent scouting

Virtual "auditions" of many alloys find a few with the potential to become "superstar" catalysts

A reality show called "Density Functional Theory" or "DFT" is not going to be challenging "American Idol" in the television ratings anytime soon. After all, DFT does not involve snarky judges and viewer voting; it is a quantum mechanical computational method used to calculate quantities such as the binding energy of molecules. But in the right hands, DFT has something in common with the popular talent contest—the ability to sort through a mob of "wannabes" and find a small number of candidates with unusual talents.

In a series of three recent papers,[12] chemical engineers Jeff Greeley and Manos Mavrikakis[13] of the University of Wisconsin-Madison reported the promising results of their talent search using DFT calculations—a search for better catalysts.

Catalysts reduce the amount of energy needed to start a chemical reaction, enabling the reaction to happen faster or at a lower temperature. They play a major role in the chemical industry for the production and processing of materials such as plastics, fuel, and pesticides, and in the pharmaceutical industry for drug synthesis.

Hydrogen catalysts are particularly important to the DOE Hydrogen Program, whose mission is to research and develop fuel cell and hydrogen production, delivery, and storage technologies, thus ensuring an abundant and affordable supply of clean energy. One of the Hydrogen Program's priority research areas is design of catalysts at the nanoscale, where reduced size often maximizes catalytic properties. Catalysis is vital to the success of the program because of its role in producing hydrogen from water or from carbon-containing fuels such as coal and biomass, and its role in producing electricity from hydrogen in fuel cells.

"Anything that would reduce the cost of materials in low-temperature fuel cells would certainly make the technology more economically competitive and thus bring fuel cells closer to being implemented in our everyday lives," Mavrikakis said. "A major fraction of that cost is the cost of noble metals, such as platinum, that are used as catalysts." Combining platinum with less expensive metals is one possible way of reducing the price of fuel cells, but only if the alloys are efficient catalysts—and even today's expensive catalysts are not efficient enough to be economically competitive.

So the research goal is clear: find less expensive, more efficient catalysts for hydrogen-related reactions. However, the discovery of catalysts by experimental methods is an expensive process of trial and error, often involving hundreds or even thousands of candidate materials. But a new theoretical approach—rational design of catalysts from first principles, based on the reactions we want them to facilitate—can speed up the discovery process by eliminating most of the candidates that are unlikely to succeed.

Using this approach, Greeley and Mavrikakis found a new class of near-surface alloys (NSAs) that exhibit superior catalytic behavior for hydrogen-related reactions (Figure 18). NSAs are alloys in which the composition near the surface differs from the bulk composition. Two idealized NSA structures (Figure 19) were used in the study: an *overlayer* of the solute metal on the host metal, and a *subsurface alloy,* in which a layer of solute lies just below the surface of the host.

Earlier studies had discovered a few NSAs that had unusual catalytic properties, so Greeley and Mavrikakis developed a novel and systematic DFT framework for screening a large number of NSA structures and compositions to determine their stability in hydrogen-rich environments. The stable NSAs were then examined for their hydrogen binding and dissociation energies. (In fuel cells, dissociation is the separation of the hydrogen atoms' protons and electrons, which allows a flow of electrical energy to be created.)

After screening 46 different NSAs, the researchers found a group of five that bind atomic hydrogen (H) as weakly as noble metals, such as gold, copper, and silver, while at the same time dissociating molecular hydrogen (H_2) much more easily. This is an unusual combination of properties: most materials with weak H binding are inactive for H_2 dissociation. "This unique set of properties may permit these alloys to serve as low-temperature, highly selective catalysts for

[12] J. Greeley and M. Mavrikakis, "Alloy catalysts designed from first principles," Nature Materials **3**, 810 (2004); "Surface and subsurface hydrogen: Adsorption properties on transition metals and near-surface alloys," Journal of Physical Chemistry B **109**, 3460 (2005); "Near-surface alloys for hydrogen fuel cell applications," Catalysis Today **111**, 52 (2006).

[13] Mavrikakis is Associate Professor and head of the Computational Surface Science and Catalysis Group in the Department of Chemical and Biological Engineering at UWM. Since receiving his Ph.D. from UWM, Greeley has accepted a position in the Center of Atomic-Scale Materials Physics at the Technical University of Denmark.

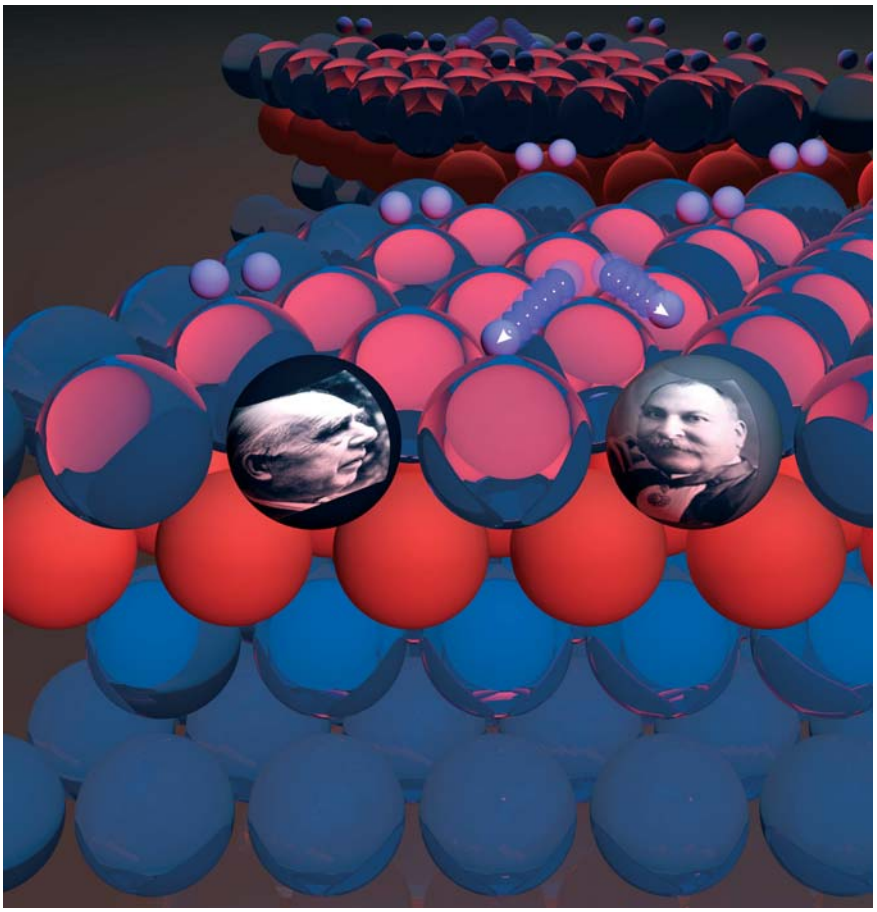

FIGURE 18. Illustration of bimetallic near surface alloys (NSAs) that bind H as weakly as noble metals but activate H_2 much more easily. Niels Bohr (1885–1962, left) and Paul Sabatier (1854–1941, right) represent the founders of quantum mechanics and catalysis theory.

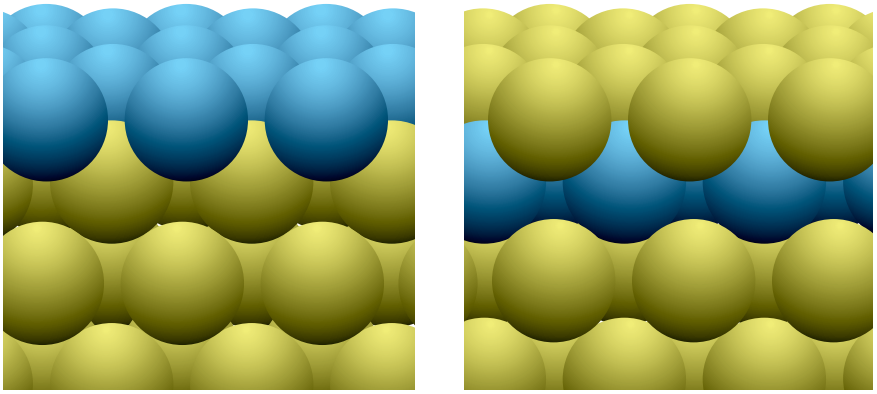

FIGURE 19. Two types of near-surface alloys were studied: overlayers (left) and subsurface alloys (right).

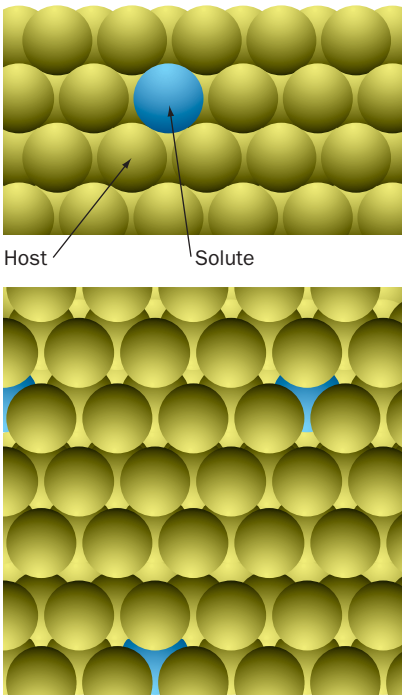

FIGURE 20. Schematic representation of an NSA defect site in cross-section and top view. The gold circles denote host metal atoms, while the blue circles indicate the subsurface solute impurity.

pharmaceuticals production and as robust fuel-cell anodes," Greeley and Mavrikakis wrote in their Nature Materials article.

Another important consideration for fuel cell anodes is resistance to carbon monoxide (CO) poisoning. Small amounts of CO are always present in hydrogen fuel streams, and those traces can bind to a platinum anode, ruining its ability to act as a catalyst. However, the NSAs with weak H binding also exhibit weak CO binding, making them ideal candidates for fuel cell anodes. The DFT results identified the following subsurface alloys as having all of these desirable properties: vanadium in platinum, tantalum in palladium, tungsten in platinum, molybdenum in platinum, and tantalum in platinum.

[14] http://www.nano.gov/

How relevant are these computational results, based on idealized structures, to NSAs in the real world? In their Catalysis Today paper, Greeley and Mavrikakis admitted, "In reality, it is fairly uncommon for NSAs to have a pure solute monolayer in either the first or second metal layers." But, they argued, "Calculations have shown … that modest deviations … do not *qualitatively* change the unusual catalytic properties of NSAs." In fact, they suggest that minor impurities, or "defect sites" (Figure 20), may have unusual catalytic properties in their own right, combining fast kinetics for bond-breaking or bond-making events with resistance to poisoning by the reactants or products of those reactions.

In any case, deviations from the ideal structure would produce *quantitative* changes in NSA properties. "Thus, to fully benefit from the unique catalytic properties of NSAs, it will be necessary to achieve more precise control of metal catalyst surface structure, perhaps by implementing improved catalyst nanofabrication techniques," the researchers wrote. Several promising nanofabrication techniques are already being developed, and perfecting such techniques is part of the mission of DOE's new Nanoscale Science Research Centers, established under the National Nanotechnology Initiative.[14]

One additional advantage of the "talented" NSAs is that they allow easy diffusion of atomic hydrogen into the bulk of the metal, raising the possibility of using them as light hydrogen-storage media. This characteristic will be studied in more detail in future research.

Research funding: BES, NSF
Computational resources: NERSC, NPACI, MSCF
This article written by: John Hules

Surface charge

Electrical charging is found to increase the reactivity of a gold nanocluster catalyst

Researchers at the Georgia Institute of Technology and Technical University Munich have discovered evidence of a phenomenon that may lead to drastically lowering the cost of manufacturing of materials from plastics to fertilizers. Studying nano-sized clusters of gold on a magnesium oxide surface, scientists found direct evidence for electrical charging of a nano-sized catalyst.[15] This is an important factor in increasing the rate of chemical reactions.

"The fabrication of most synthetic materials that we use involves using catalysts to promote reaction rates," said Uzi Landman, director of the Center for Computational Materials Science at Georgia Tech. "Designing catalysts that are more efficient, more selective and more specific to a certain type of reaction can lead to significant savings in manufacturing expenses. Understanding the principles that govern nanocatalysis is key to developing more effective catalysts."

The current study builds on joint research done since 1999 by the two groups that found gold, which is non-reactive in its bulk form, is a very effective catalyst when it is in nanoclusters of eight to about two dozen atoms in size. Those specific sizes allow the gold clusters to take on a three-dimensional structure, which is important for its reactivity.

"It is possible to tune the catalytic process not only by changing the composition of the materials, but also by changing the cluster's size atom by atom," explained Ueli Heiz, professor of chemistry at Technical University Munich. In these earlier studies of the reaction where carbon monoxide and molecular oxygen combine to form carbon dioxide, Landman's group used computer simulations to predict that when gold nanoclusters of eight atoms are used as the catalyst and magnesium oxide is used as the catalytic bed, reactions would occur when the bed had defects in the form of missing oxygen atoms, but would not occur when the magnesium oxide was defect-free.

Heiz's experiments confirmed this prediction and the teams concluded that the gold must be anchoring itself to the defect where it picks up an electron, giving the gold a slight negative charge. Theoretical simulations showed that the electronic structure of the gold clusters matches up with the oxygen and carbon monoxide (Figure 21). The charged gold transfers an electron to the reacting molecules, weakening the chemical bonds that

[14] http://www.nano.gov/

[15] B. Yoon, H. Häkkinen, U. Landman, A. S. Wörz, J.-M. Antonietti, S. Abbet, K. Judai, and U. Heiz, "Charging effects on bonding and catalyzed oxidation of CO on Au_8 clusters on MgO," Science **307**, 403 (2005)

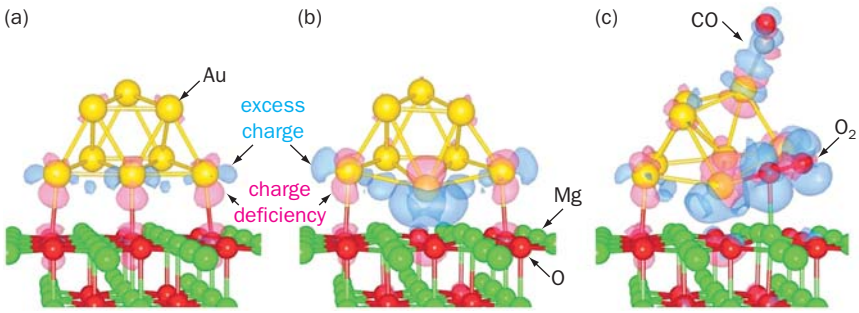

FIGURE 21. Charge differences for three configurations of an eight-atom gold nanocluster on a magnesium oxide (MgO) surface: (a) gold cluster on a perfect MgO surface; (b) gold cluster anchored on an oxygen vacancy defect, showing enhanced charging of the cluster; (c) same as (b) but with adsorbed reactants, showing charging of the carbon monoxide and oxygen molecules.

keep them together. Once the bond is weak enough it breaks, allowing reactions to occur.

Now, in this latest study, the group has found direct evidence that this is indeed what is happening. Using eight-atom gold clusters as the catalyst and magnesium oxide as the catalytic bed, the team measured and calculated the strength of the bonds in the carbon monoxide by recording the frequency of the molecule's vibrations.

"If carbon monoxide is a strong bond, then there is a certain frequency to this vibration," explained Landman. "If the bond of the carbon monoxide becomes weaker, then the frequency becomes lower. That's exactly what we saw—when we had defects in the magnesium oxide, we had larger shifts than when we had magnesium oxide without defects."

Lead author of the paper Bokwon Yoon, a senior research scientist in Landman's group, commented, "The agreement between the predicted and the measured values of the vibrational frequency shifts is very gratifying, confirming the charging and bonding mechanisms."

"And all this happens at low temperatures," said Heiz. Typically, reactions requiring catalysts need heat or pressure to get the reaction going, and that adds to the cost of manufacturing, but that isn't the case here. Since the properties of the catalytic beds can increase the rate of reactions for nanocatalysts, new and better low-temperature catalysts may be found.

"We knew the specific number of atoms in the catalyst and that defects in the catalytic beds are important. Now we know why those defects are so essential—because they allow the catalyst to be electrically charged. We hope these guidelines will lead to more research in search of nano-sized catalysts. It's possible that at the nanoscale you may find catalysts that can do things under more gentle and cheaper conditions," said Landman.

Research funding: BES, USAF-OSR, DFG, LBW, AVHF, JSPS, SNSF, GMODGO
Computational resources: NERSC, HPCMP
This article written by: David Terraso (Georgia Tech)

Magnetic disks in space

A new generation of simulations and experiments aim to pinpoint the origins of instability in accretion disks

Accretion disks—structures formed by material falling toward and then orbiting around a gravitational source—are a key part of the process by which newly born stars and black holes accumulate mass and grow (Figure 22). The gravitational energy released by accretion is believed to power many of the energetic phenomena observed in the universe, such as the jets streaming from galactic nuclei (Figure 23).

Accretion disks form because most of the matter being attracted toward the gravitational well has some angular momentum, which results in the matter orbiting the proto-star or black hole rather than falling right in. And that creates a problem for theorists, according to Fausto Cattaneo, a researcher at Argonne National

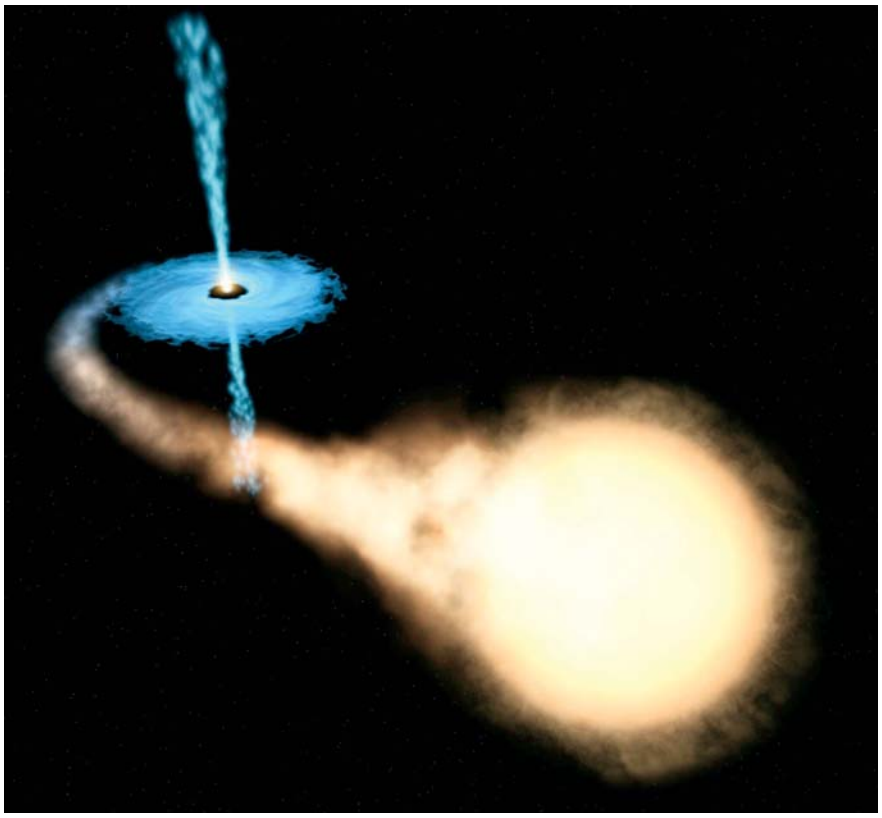

FIGURE 22. An artist's impression of a black hole with a closely orbiting companion star. Infalling matter forms an accretion disk, which feeds matter into the black hole. Hot gas rushes from the vicinity of the black hole, creating jets perpendicular to the disk. (Image courtesy of NASA.)

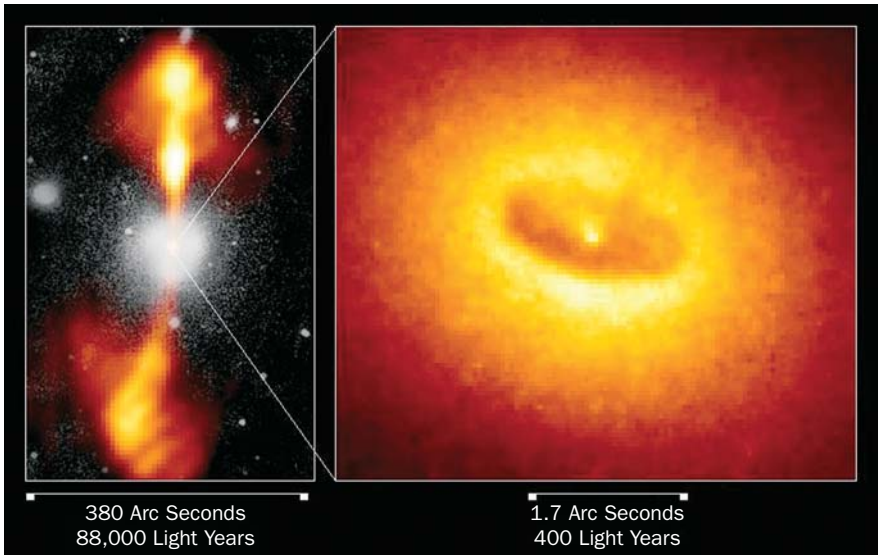

Laboratory and the Center for Magnetic Self-Organization at the University of Chicago.

"If the angular momentum is conserved," Cattaneo said, "the material will happily orbit the central object and never fall in. In order for material to fall, it must lose its angular momentum."

In the early days of accretion disk theory, the loss of angular momentum was not considered a problem, because two obvious explanations were easily available: collisional processes like friction or viscosity, and shear-driven turbulence.

If two particles in the disk collide, they will exchange angular momentum, and the faster-rotating one will slow down and fall into a lower orbit. If the collision rate is high enough, there will be a spiral of matter flowing down into the central gravitational well. This hypothesis makes sense until you do the math. "In most astrophysical situations, collisional processes are many orders of magnitude too small to account for the observed energy release rates," Cattaneo explained. Accretion disks are simply too big to be destabilized by such processes.

Turbulence is a more promising hypothesis, because turbulent eddies

FIGURE 23. The giant elliptical galaxy NGC 4261, one of the 12 brightest galaxies in the Virgo cluster. (Left) Composite ground-based optical and radio view of the galaxy. Photographed in visible light (white), the galaxy appears as a fuzzy disk of hundreds of billions of stars. A radio image (orange) shows a pair of opposed jets emanating from the nucleus. (Right) Hubble Space Telescope image of the core of NGC 4261: The disk is tipped enough to provide a clear view of the bright hub, which presumably harbors a black hole. The dark, dusty disk represents a cold outer region which leads to an ultra-hot accretion disk within a few hundred million miles from the black hole. (Image courtesy of NASA.)

provide a much more efficient mechanism for transporting angular momentum. And models of accretion disks that assume a reasonable amount of turbulence have produced credible accretion rates. "The challenge has been to provide an underlying physical mechanism to generate the turbulence in the first place," Cattaneo said. Shear-driven instabilities, arising from the varying rotational speeds within the disk, were another attractive possibility that had to be ruled out. "The problem is that the rotation profiles—the average angular velocity as a function of radius—of typical accretion disks are close to following Kepler's laws of planetary motion, and so they are very stable to infinitesimal perturbations."

"Two decades of accretion disk research failed to provide a single local instability mechanism that could operate in disks and lead to the desired state of turbulence," Cattaneo continued. "This lamentable state of affairs improved significantly in the early 1990s when it was realized that the stability properties of near-Keplerian disks could be dramatically affected by magnetic fields. Disks that were hydrodynamically stable could become hydromagnetically unstable, and remarkably, all it took was a weak magnetic field. Since most astrophysical accretion disks are ionized and hence good electrical conductors, this *magneto-rotational instability,* as it came to be called, could provide the desired mechanism for turbulence production."

The physical origin of magneto-rotational instability can be illustrated by a simple analogy: Imagine two spacecraft orbiting near each other at slightly different altitudes. If the orbits are nearly circular, the inner spacecraft will have higher angular velocity than the outer one. Now assume that a light elastic cord joins the two spacecraft, providing a weak tension. The effect of the tension on the inner spacecraft is to slow it down, i.e., to reduce its angular velocity, and therefore to move it into a lower orbit. In contrast, the effect of the tension on the outer spacecraft is to accelerate it, i.e., to increase its angular velocity, thereby moving it to a higher orbit. With one spacecraft going lower and the other going higher, it is hard to predict what will happen when they stretch their cord to its full length—will they start bouncing wildly, or will the cord break and propel them further apart?—but clearly, the system will become unstable.

A magnetic field in an electrically conducting disk can produce magnetic tension effects that are analogous to the elastic tension in the spacecraft example. By connecting portions of the disk that are rotating at different speeds, the magnetic field turns angular velocity into a source of instability. And because the central gravitational field does not play a direct role in magneto-rotational instability, laboratory experiments have been developed that can test many aspects of this mechanism on a small scale. In these experiments, the space between two coaxial rotating cylinders is filled with an electrically conducting fluid, typically liquid sodium or gallium. The rotation rates of the inner and outer cylinders are chosen so that the rotation profile is hydrodynamically stable. Then a weak external magnetic field is applied so that the origins and development of magneto-rotational instability can be studied.

"Although the physical input provided by these experiments is invaluable," Cattaneo said, "they suffer from two limitations: First, in a flow of liquid metal, it's impossible to see, and difficult even to detect, what is happening." The second limitation concerns the Reynolds number, a mathematical expression that measures the strength of advective effects relative to diffusion. "In any laboratory setup, the magnetic Reynolds number of liquid metals is a few tens, and with extreme efforts it can be raised slightly to exceed one hundred," he said. "This should be contrasted with the typical astrophysical situation in which the magnetic Reynolds number is huge—millions to billions."

Fortunately, this is a case in which numerical simulations can elucidate and reach even further than the experiments. In an INCITE-funded project titled "Magneto-Rotational Instability and Turbulent Angular Momentum Transport," Cattaneo and his colleagues Paul Fischer and Aleksandr Obabko have created three-dimensional numerical simulations that reproduce the geometry of the laboratory experiments conducted by Hantao Ji of the Princeton Plasma Physics Laboratory and Jeremy Goodman of Princeton University. But unlike the experiments, the simulations have magnetic Reynolds numbers exceeding 50,000—not quite as high as in actual accretion disks, but a valuable extension of the liquid metal results. The purpose of these simulations is to clarify the mechanisms of angular momentum transport in magnetized, rotationally constrained turbulence, and to apply them to the problem of astrophysical accretion flows.

"If you can do the research both computationally and experimentally, you are much better off than just using one approach," Cattaneo said. "Seeing is knowing, and in numerical simulations, you can visualize anything, so it is much easier to identify the beginnings of the instability and the structures that transport angular momentum."

In these simulations, as in most experiments to date, the magnetic field that catalyzes the magneto-rota-

tional instability is aligned with the rotation axis. Visualizations show that in this case, the coherent structures that transport the angular momentum outward are toroidal vortex rings emanating from the inner cylinder (Figure 24). Knowing what the structures look like and how they behave advances the theoretical understanding of this instability and also makes possible the future creation of phenomenological models to describe this phenomenon. These models could then be plugged into simulations that can run on workstations instead of supercomputers, or that can use parameters closer to those found in real astrophysical accretion disks.

Although Cattaneo and his colleagues will be spending the next year analyzing all the data from their simulations, some intriguing possibilities are already emerging. "If you turn the cylinders on their side," he said, "the transport of angular momentum looks similar to the transport of heat in a convecting fluid, like a pot of water on a hotplate. We're going to examine this in more detail to see whether some of the ideas of convection theory can be applied to fluid flow in accretion disks."

The close collaboration between the theorists in Chicago and the experimentalists in Princeton facilitates the exchange of results from both approaches. "We used the experimental results to validate our computational approach," Cattaneo said. "And our simulations are giving the experimental researchers a better idea of what to look for and where to look for it. They are already using our results to improve the design of their experiments."

Research funding: INCITE, HEP, NSF
Computational resources: NERSC, CMSO, BGL
This article written by: John Hules

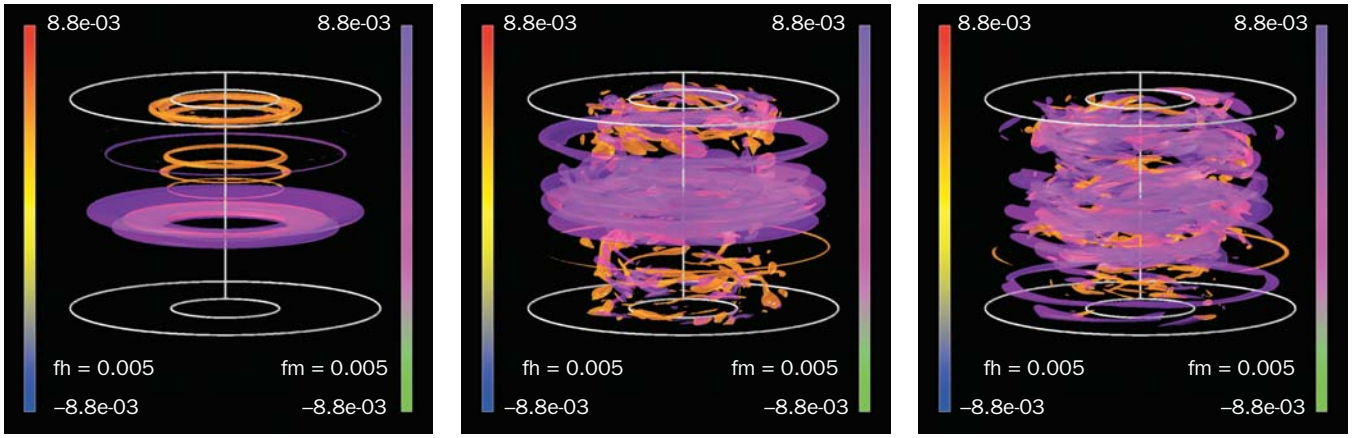

FIGURE 24. Visualization of the time evolution of the outward transport of angular momentum in a magnetic fluid bounded by rotating cylinders. The two colors correspond to the transport by hydrodynamic (orange) and hydromagnetic (purple) fluctuations. (Simulations by F. Cattaneo, P. Fischer, and A. Obabko, with visualization assistance from Cristina Siegerist of the Berkeley Lab/NERSC Visualization Group. The animation can be viewed at http://flash.uchicago.edu/~cattaneo/Animations/fh-fm_1024x768_1.mpg.)

Proteins in motion

Biochemists are systematically simulating the unfolding pathways of all known protein folds, hoping to discover the general principles of protein folding

Proteins do the work of life; they are essential to the structures and functions of all living cells. Some proteins play structural or mechanical roles, others catalyze chemical reactions, and still others are involved in storage or transportation or immune response.

A protein's biological function depends on its shape. Proteins are synthesized as long, non-branching chains of amino acids; but in order to perform their proper functions, each type of protein must assume a unique, stable, and precisely ordered three-dimensional folded structure, which biologists call its *native state.* Protein folding might be compared to a shoelace tying itself, except the shapes of proteins are much more complicated than knots, and the folding process is much faster, being completed only microseconds to seconds after the protein is created.

Incorrectly folded proteins are responsible for many illnesses, including Alzheimer's disease, cystic fibrosis, Huntington's disease, Parkinson's disease, Type II diabetes, Creutzfeldt-Jakob disease (the human counterpart of mad cow disease), and many cancers. The degradation of protein folding may even play a key role in aging. But despite the importance of protein folding, the phenomenon remains largely a mystery.

"Protein folding is one of the fundamental unsolved problems in molecular biology," said Valerie Daggett, Professor of Medicinal Chemistry and Adjunct Professor of Biochemistry, Biomedical and Health Informatics, and Bioengineering at the University of Washington. "Although we know a lot about of the structural details of the native folded conformation of proteins, very little is known about the actual folding process. Experimental approaches only provide limited amounts of information on the structural transitions and interactions occurring during protein folding. Given that protein folding is of such widespread importance to human health, we are using computer simulation methods in an attempt to delineate the important forces acting during this process."

Daggett is principal investigator of an INCITE project called "Molecular Dynameomics," which was awarded 2 million processor-hours at NERSC. Just as genomics studies the complete genetic makeup of an organism, and *proteomics* studies the complete protein structure and functioning of an organism, *dynameomics* is an ambitious attempt to combine molecular dynamics and proteomics, using molecular dynamics simulations to characterize and catalog the folding/unfolding pathways of representative proteins from all known protein folds. There are approximately 1,130 non-redundant protein folds, and Daggett's group has used their INCITE grant to simulate proteins from the 151 most common folds, which represent about 75% of all known protein structures (Figure 25).

"Structure prediction remains one of the elusive goals of protein chemistry," Daggett said. "In order to translate the current deluge of genomic information into knowledge about protein functions, which we can then use for drug design, we have to be able to successfully predict the native states of proteins starting with just the amino acid sequence."

The novelty of her group's computational approach, however, is that they work in the opposite direction, starting not with a protein's amino acid sequence but with its folded structure, as determined by NMR spectroscopy and X-ray crystallography experiments. A protein can easily be unfolded or denatured, in both experiments and simulations, by simply raising the temperature. (Anyone who has ever cooked an egg has seen the results of denaturation: the heat causes the bonds within the egg white proteins to break, unfolding the individual protein molecules and allowing them to form bonds with other molecules, thus solidifying the egg white.) Daggett's research over the last several years suggests that protein unfolding follows the same pattern as folding, moving in reverse through the same transition states and intermediate structures. The high temperature of the simulations, 498 Kelvin (437° F), just speeds up the process without modifying the overall pathway, she says.

To perform these simulations, Daggett's group uses a code called "*in lucem* Molecular Mechanics" or *il*mm, which was developed by David Beck, Darwin Alonso, and Daggett. The code solves

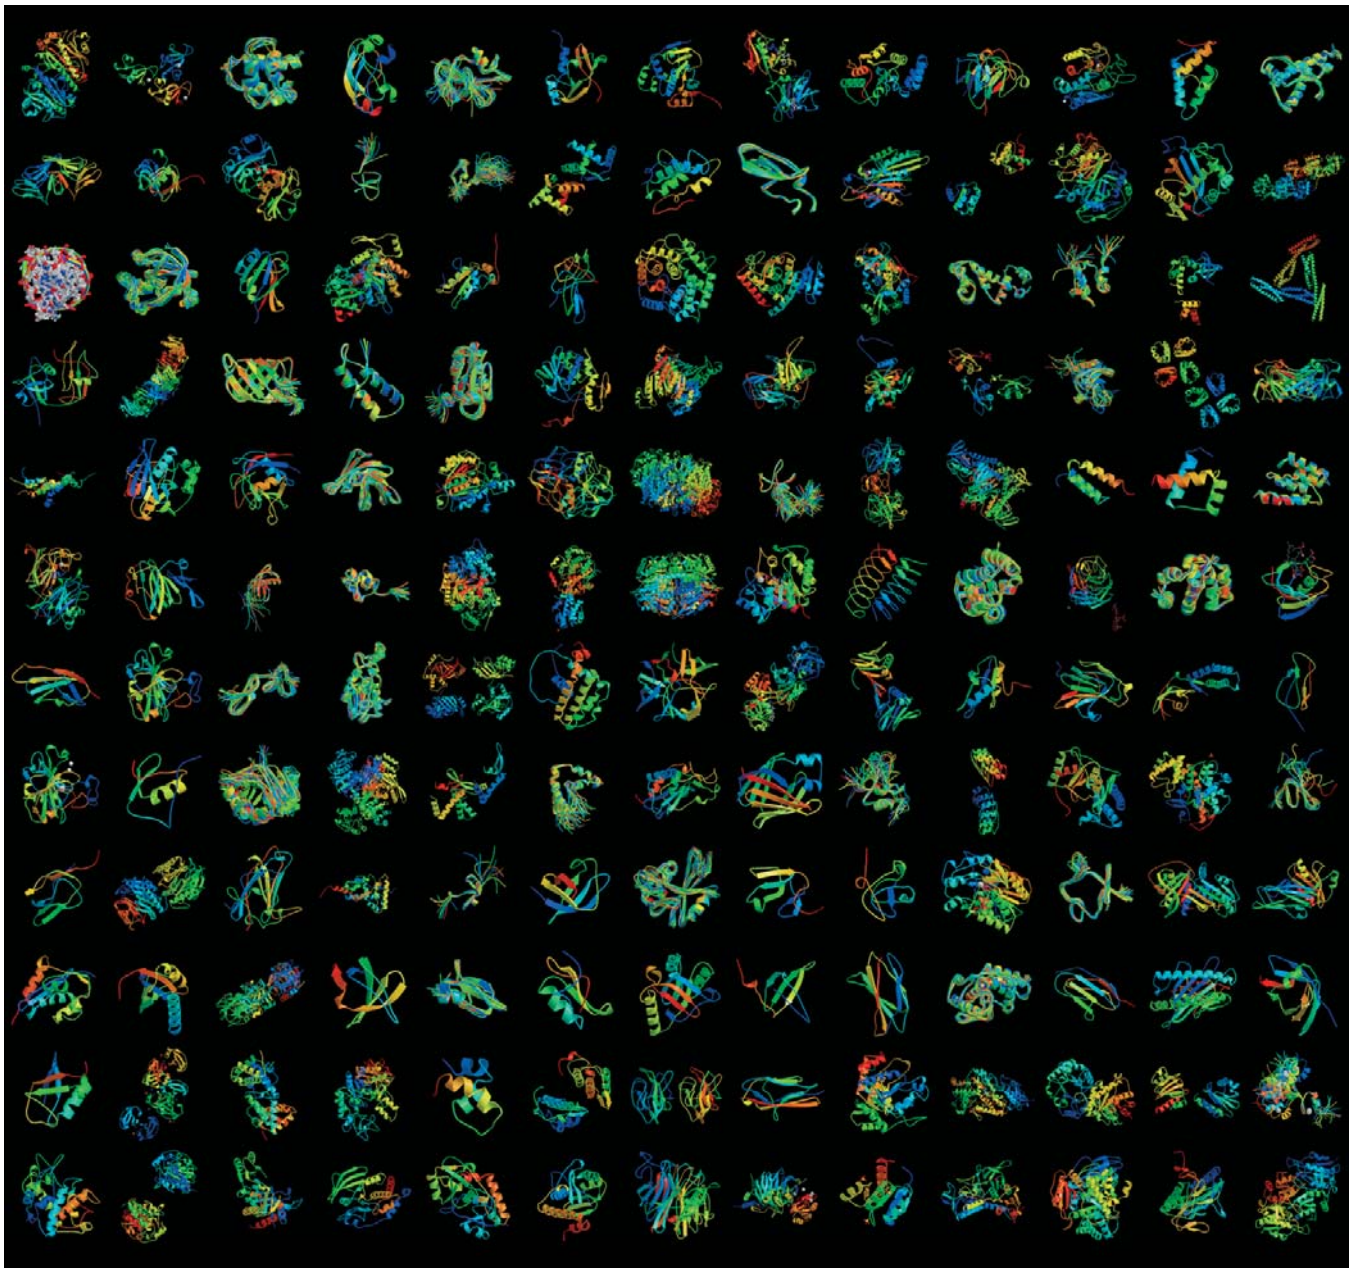

FIGURE 25. The first 156 dynameomics simulation targets.

Newton's equations of motion for every atom in the protein molecule and surrounding solvent environment, determining the trajectories of the atoms over time. Obtaining a reasonably good picture of protein unfolding pathways requires six simulations for each protein: one 20 nanosecond (ns) simulation of the native (folded) molecular dynamics at a temperature of 298 K or 77° F (Figure 26); and five simulations at 498 K, three at 2 ns duration and two at 20 ns (Figure 27).

Thus the 151 proteins studied in the INCITE project required a total of 906 simulations and produced 2 terabytes of compressed data. Even with the allocation of 2 million processor-hours, the researchers had to limit the size of the chosen proteins to no more than 300 amino acids. Fortunately, the larger, more complex proteins are generally structured as if they were assembled from a number of smaller shapes, so understanding the smaller shapes should reveal principles that are applicable to all proteins.

In fact, Daggett hopes that cataloging the folding strategies for every type of protein fold will enable the discovery of general rules of protein folding, which so far have eluded researchers, and that those rules will be used to improve algorithms for predicting protein structure. The results of the protein unfolding simulations will be made publicly available on the Dynameomics website (www.dynameomics.org). Just as the Protein Data Bank has become an important repository of experimental results that has enabled many further discoveries, Daggett envisions the Dynameomics simulation database as a resource for the whole research community, in which data mining may produce unexpected findings and a dynamic description of proteins will aid in understanding their biological functions.

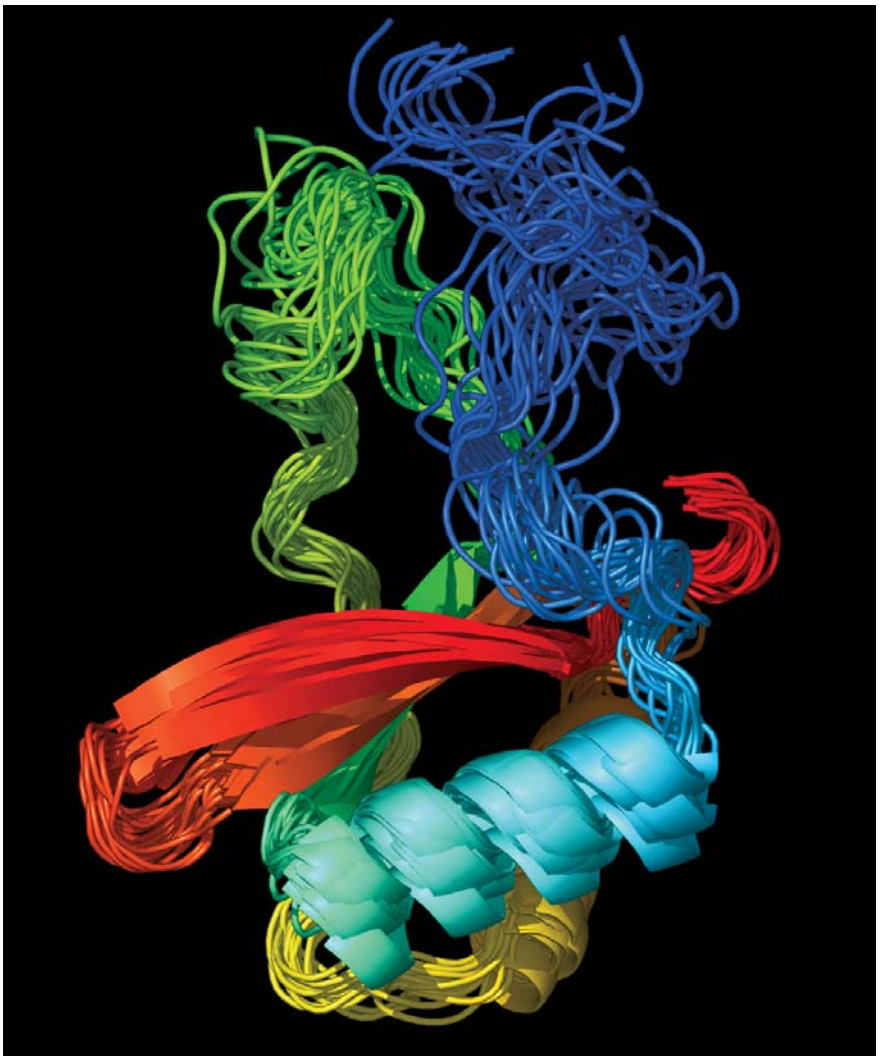

FIGURE 26. Protein target 3rub: A schematic representation of the secondary structure taken from a native state simulation of the enzyme RuBisCO, the most abundant protein in leaves and possibly the most abundant protein on Earth. Water, hydrogen, and protein atoms included in the simulation are not shown. The highly dynamic region (image top) is responsible for finding the protein's binding partner (not included in the simulation). RuBisCO catalyzes the first major step of carbon fixation, a process by which atmospheric carbon dioxide is made available to organisms in the form of energy-rich molecules such as sucrose.

David Beck, a graduate student in the Daggett lab, worked with NERSC consultants to optimize *il*mm's performance on Seaborg. "The INCITE award gave us a unique opportunity to improve the software, as well as do good science," Beck said. Improvements included load balancing, which sped up the code by 20%, and parallel efficiency, which reached 85% on 16-processor nodes. The INCITE award enabled the team to do five times as many simulations as they had previously completed using other computing resources.

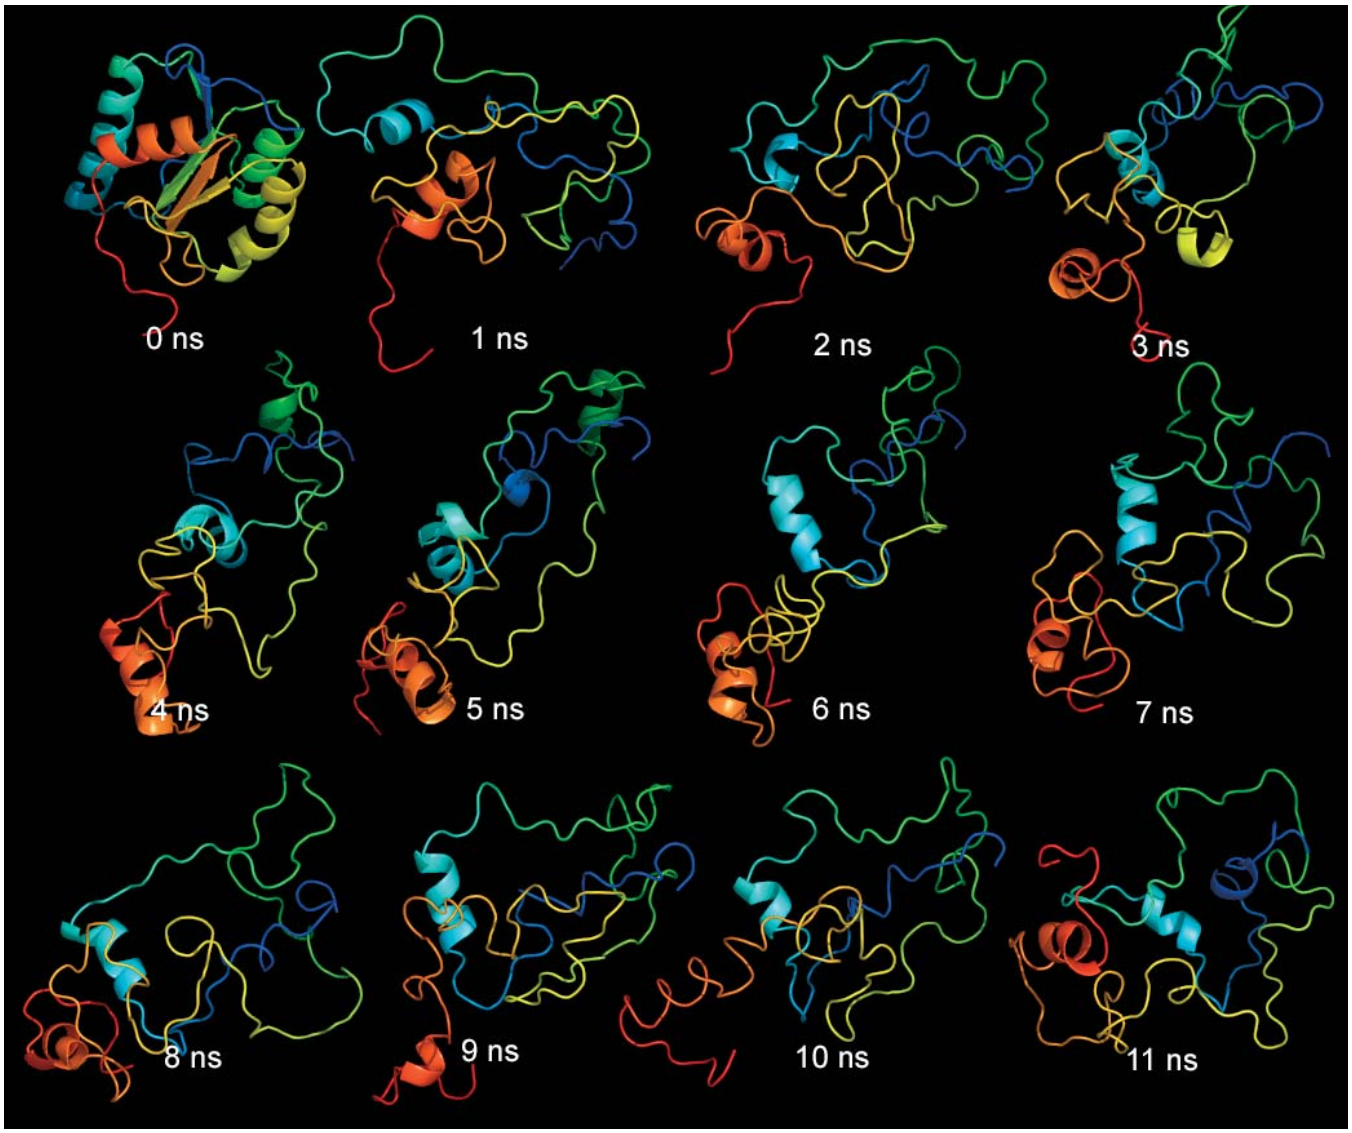

FIGURE 27. Protein target 1imf: Schematic representation of secondary structures taken at 1 ns intervals from a thermal unfolding simulation of inositol monophosphatase, an enzyme that may be the target for lithium therapy in the treatment of bipolar disorder.

The researchers are now analyzing the simulation results and comparing them to experimental results whenever possible. "Our preliminary analysis of this data set focused on the native control simulations and the endpoint of the thermal unfolding simulations—the denatured state," Daggett said. "In particular, we are evaluating the interactions between amino acids. We have already found a shift in the prevalence and types of interactions as the proteins unfold. Characterizing the interactions in the denatured state provides much-needed information about the starting point of folding. This represents the first step towards providing crucial information for prediction of the folding process and ultimately the folded, native structure."

And, of course, there are still hundreds more protein folds to simulate.

Research funding: INCITE, BER
Computational resources: NERSC
This article written by: John Hules

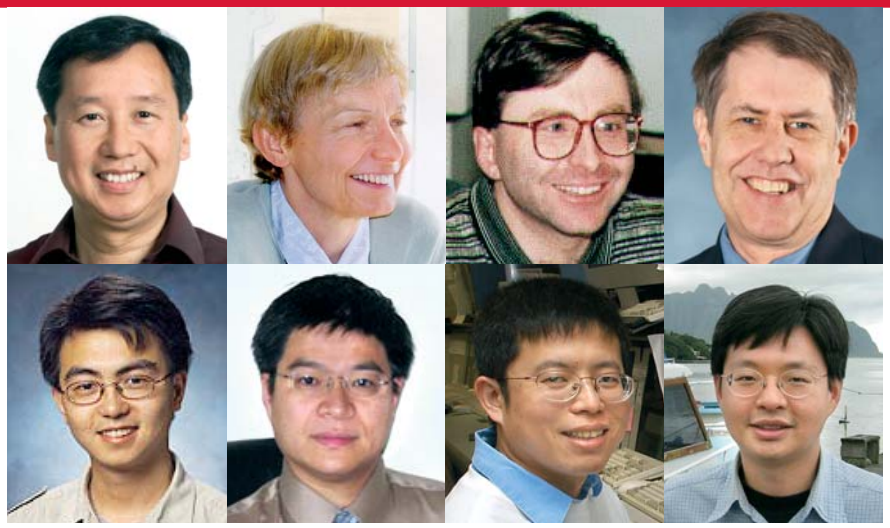

Top row, left to right: Louie, Romanowicz, Norman, and Chelikowsky; bottom row, left to right: Cai, Qin, Xu, and Tseng

NERSC Users Honored

Two NERSC users with joint appointments at the University of California, Berkeley and Lawrence Berkeley National Laboratory were elected to the National Academy of Sciences: Steven G. Louie, Professor of Physics and an expert on nanotubes and fullerenes, is one of the 25 most highly cited authors in nanoscience and one of the 100 most cited researchers in all of physics. Barbara A. Romanowicz, Director of the Berkeley Seismological Laboratory, studies deep earth structure and dynamics; she discovered that the mysterious waves that constantly reverberate through the Earth, even on days without earthquakes, originate from storm-driven sea waves.

Michael Norman, Professor of Physics at the University of California, San Diego, was elected a fellow of the American Academy of Arts and Sciences. Norman is PI of the 2006 INCITE project "Precision Cosmology Using the Lyman Alpha Forest," which aims to measure the cosmological parameters that describe the shape, matter-energy contents, and expansion history of the Universe.

The American Physical Society presented the David Adler Lectureship Award to James Chelikowsky, Professor of Physics, Chemical Engineering, and Chemistry and Biochemistry at the University of Texas, "for his creative and outstanding research in computational materials physics and for his effectiveness in communicating research results through lectures and publications."

Three NERSC users were recipients of the Presidential Early Career Awards for Scientists and Engineers: Wei Cai, an Assistant Professor of Mechanical Engineering at Stanford University, was honored for work he did at Lawrence Livermore National Laboratory in developing a computer model that simulates the dynamics of crystals as they deform. Hong Qin, a physicist at Princeton Plasma Physics Laboratory, was recognized for his contributions to the physics of high-intensity particle beams and for his work on electromagnetic effects in magnetically confined plasmas. Zhangbu Xu, a physicist at Brookhaven National Laboratory, was honored for his novel research techniques which led to the detection of subatomic particles termed "short-lived resonances" and "open charm" at the Relativistic Heavy Ion Collider.

Yu-Heng Tseng, a computational climate researcher at Lawrence Berkeley National Laboratory and a lecturer at the University of California, Berkeley, was honored as an Outstanding Overseas Young Scientist by the Foundation for the Advancement of Outstanding Scholarship in Taiwan.

THE NERSC CENTER

NERSC traditionally provides systems and services that maximize the scientific productivity of its user community, and takes pride in its reputation for the expertise of its staff and the high quality of services delivered to its users. In support of the DOE Office of Science's mission, the NERSC Center served 2,677 scientists throughout the United States in 2005. Appendix B provides statistical information about NERSC's clients, their primary affiliations, and their scientific disciplines.

Scientific productivity is the most important measure of NERSC's success, as evidenced by more than 1,400 scientific publications in 2005 that were based entirely or in part on calculations done at NERSC.[1] To maintain its effectiveness, NERSC proactively addresses new challenges. How the NERSC Center is working to improve that productivity even more in the face of growing scientific and technological challenges is discussed in the following pages.

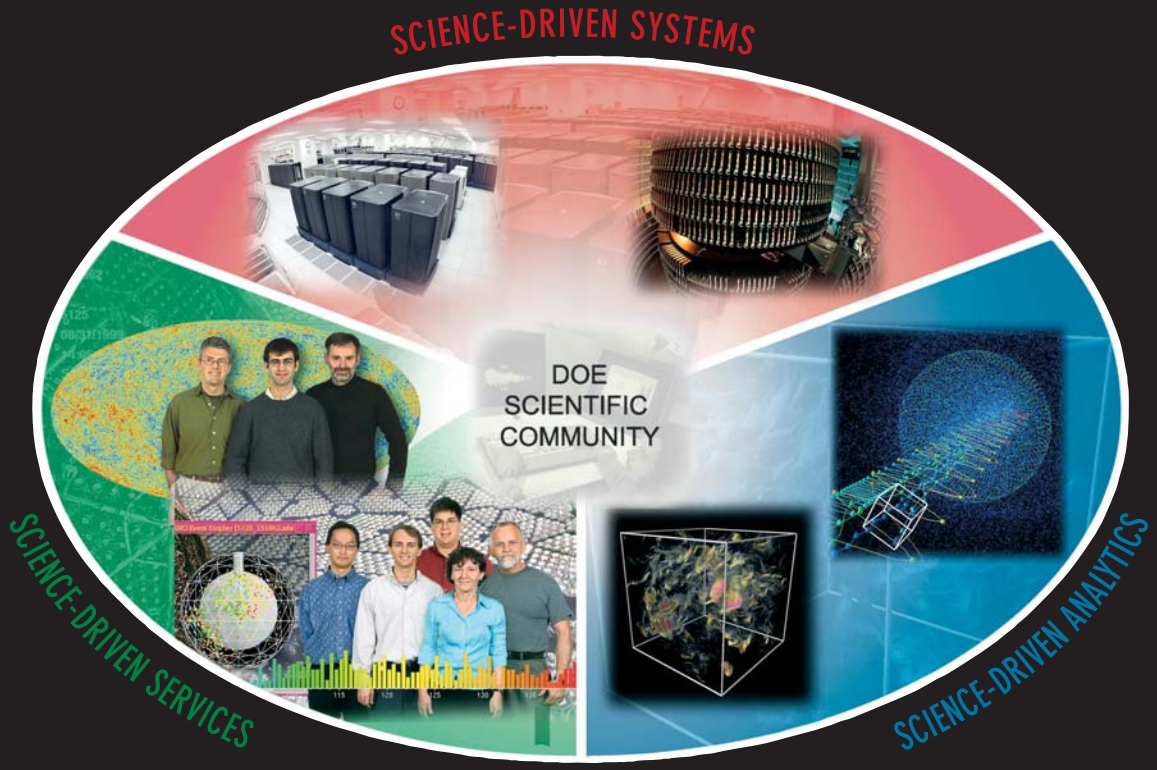

Science-Driven Computing

NERSC is continuously reassessing its approach to supporting high-end scientific computing, and from time to time undertakes a major reevaluation and realignment. In 2005 this involved several activities: the NERSC Users' Group writing and publishing the latest DOE Greenbook, the development of NERSC's five-year plan for 2006– 2010, and a programmatic peer review conducted for the DOE. The overall theme of this evaluation and planning effort was *Science-Driven Computing.*

DOE Greenbook Published

The *DOE Greenbook: Needs and Directions in High-Performance Computing for the Office of Science* was compiled by Steve Jardin for the NERSC Users Group and published in June 2005 (Figure 1). With contributions from 37 scientists from a variety of disciplines and organizations, this report documents the computational science being done at NERSC and other DOE computing centers and provides examples of computational challenges and opportunities that will guide the evolution of these centers over the next few years.

According to the Greenbook, researchers in all of the disciplines supported by the Office of Science are finding that large-scale computational capabilities are now essential for the advancement of their research. Today's most powerful computers and scientific application codes are being used to produce new and more precise scientific results at the cutting edge of each discipline, and this trend is destined to continue for years to come. The Greenbook presents many examples of the impact of large-scale computations on the sciences.

However, the Greenbook points out that

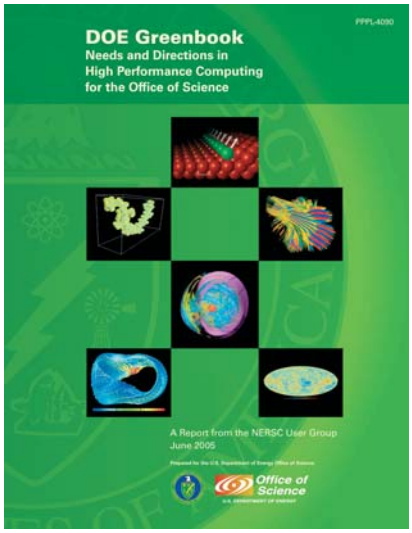

FIGURE 1. The DOE Greenbook, prepared by the NERSC Users Group, is available online at http://www.nersc.gov/news/greenbook/2005greenbook.pdf.

the current computational resources available through NERSC are saturated, and the lack of additional computing resources is becoming "a major bottleneck in the scientific research and discovery process." The report advises, "A large increase in computer power is needed in the near future to take the understanding of the science to the next level and to help secure the U.S. DOE SC leadership role in these fundamental research areas."

The Greenbook's specific recommendations include:

- Expand the high performance computing resources available at NERSC, maintaining an appropriate system balance to support the wide range of large-scale applications involving production computing and development activities in the DOE Office of Science.
- Configure the computing hardware and queuing systems to minimize the time-to-completion of large jobs, as well as to maximize the overall efficiency of the hardware.
- Actively support the continued improvement of algorithms, software, and database technology for improved performance on parallel platforms.
- Significantly strengthen the computational science infrastructure at NERSC that will enable the optimal use of current and future NERSC supercomputers.
- Carefully evaluate the requirements of data- or I/O-intensive scientific applications in order to support as wide a range of science as possible.

NERSC's Five-Year Plan

With guidance from the DOE Greenbook and other interactions with the NERSC user community, as well as monitoring of technology trends, NERSC has developed a five-year plan focusing on three components: Science-Driven Systems, Science-Driven Services, and Science-Driven Analytics (Figure 2). NERSC management and staff have observed three trends that need to be addressed over the next several years:

- the widening gap between application performance and peak performance of high-end computing systems
- the recent emergence of large, multidisciplinary computational science teams in the DOE research community
- the flood of scientific data from both simulations and experiments, and the convergence of computational simulation with experimental data collection and analysis in complex workflows.

NERSC's responses to these trends are the three components of the science-

[1] A list is available at http://www.nersc.gov/news/reports/ERCAPpubs05.php.

driven strategy that NERSC will implement and realize in the next five years:

- *Science-Driven Systems:* Balanced introduction of the best new technologies for complete computational systems—computing, storage, networking, visualization and analysis—coupled with the activities necessary to engage vendors in addressing the DOE computational science requirements in their future roadmaps.

- *Science-Driven Services:* The entire range of support activities, from high-quality operations and user services to direct scientific support, that enable a broad range of scientists to effectively use NERSC systems in their research. NERSC will concentrate on resources needed to realize the promise of the new highly scalable architectures for scientific discovery in multidisciplinary computational science projects.

- *Science-Driven Analytics:* The architectural and systems enhancements and services required to integrate NERSC's powerful computational and storage resources to provide scientists with new tools to effectively manipulate, visualize, and analyze the huge data sets derived from simulations and experiments.

This balanced set of objectives will be critical for the future of the NERSC Center and its ability to serve the DOE scientific community. Elements of this strategy that are currently being implemented are discussed in the following pages. The full five-year plan can be read at http://www.nersc.gov/news/reports/LBNL-57582.pdf.

DOE Review of NERSC

On May 17–19, 2005, a Programmatic Review of NERSC was conducted for the Department of Energy. The peer review committee was chaired by Frank

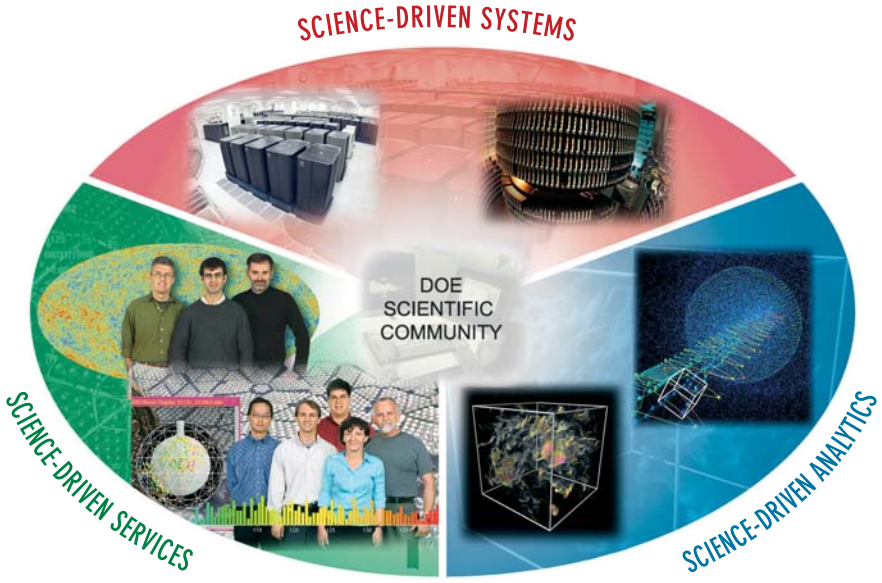

FIGURE 2. Conceptual diagram of NERSC's plan for 2006–2010.

Williams of the Arctic Region Supercomputing Center at the University of Alaska, Fairbanks. Other members were Walter F. Brooks of the Advanced Supercomputing Division at NASA Ames Research Center, Lawrence Buja of the National Center for Atmospheric Research, Cray Henry of the Defense Department's High Performance Computing Modernization Program, Robert Meisner of the National Nuclear Security Administration, José L. Muñoz of the National Science Foundation, and Tomasz Plewa of the Center for Astrophysical Thermonuclear Flashes at the University of Chicago.

In addition to reviewing the DOE Greenbook and NERSC's five-year plan, the review panel heard presentations and engaged in conversations covering all aspects of NERSC's operations. The DOE had requested that the panel address a number of specific topics, but they were also given the freedom to look into any aspect of NERSC and to comment accordingly. The panel responded by presenting a detailed list of findings and recommendations to DOE and NERSC managers, who are now using those findings to improve NERSC's operations.

The overall conclusions of the review committee included a strong endorsement of NERSC's approach to enabling computational science:

> NERSC is a strong, productive, and responsive science-driven center that possesses the potential to significantly and positively impact scientific progress by providing users with access to high performance computing systems, services, and analytics beneficial to the support and advancement of their science....
>
> Members of the review panel each report that NERSC is extremely well run with a lean and knowledgeable staff. The panel members saw evidence of strong and committed leadership, and staff who

are capable and responsive to users' needs and requirements. Widespread, high regard for the center's performance, reflected in such metrics as the high number of publications supported by NERSC, and its potential to positively impact future advancement of computational science, warrants continued support.

Organizational Changes

In order to implement the new initiatives and the changes in emphasis derived from the planning and review process, in November 2005 NERSC announced several organizational changes, including two new associate general managers, two new teams, and a new group.

"In order to efficiently carry out our plan and meet the expectations of our users and sponsors, we are modifying the NERSC Center organization," General Manager Bill Kramer wrote in announcing the changes. "In addition to the Division, Department and Group components of the organization, we will have two other components: Functional Areas and Teams."

NERSC has created two functional areas—*Science-Driven Systems* and *Science-Driven Services.* The majority of the NERSC staff will work in these two areas (Figure 3). The functional areas are responsible for carrying out the responsibilities and tasks discussed in the respective sections of NERSC's five-year plan. Functional areas will be led by Associate General Managers (AGMs), who are responsible for coordinating activities across the groups and teams in their areas. Francesca Verdier is associate general manager for Science-Driven Services, and Howard Walter is associate general manager for Science-Driven Systems (Figure 4).

The Accounts and Allocations Team, the Analytics Team, the Open Software and Programming Group, and the User Services Group will report to the Science-Driven Services AGM. The Computational Systems Group, the Computer, Operations and ESnet Support Group, the Mass Storage Group, and the Networking, Security and Servers Group will report to the Science-Driven Systems AGM.

The reorganization includes the creation of one new group and two new teams. They are:

- *Analytics Team:* Analytics is the intersection of visualization, analysis, scientific data management, human-computer interfaces, cognitive science, statistical analysis, and reasoning. The primary focus of the Analytics Team is to provide visualization and scientific data management solutions to the NERSC user community to better understand complex phenomena hidden in scientific data. The responsibilities of the team span the range from applying off-the-shelf commercial software to advanced development

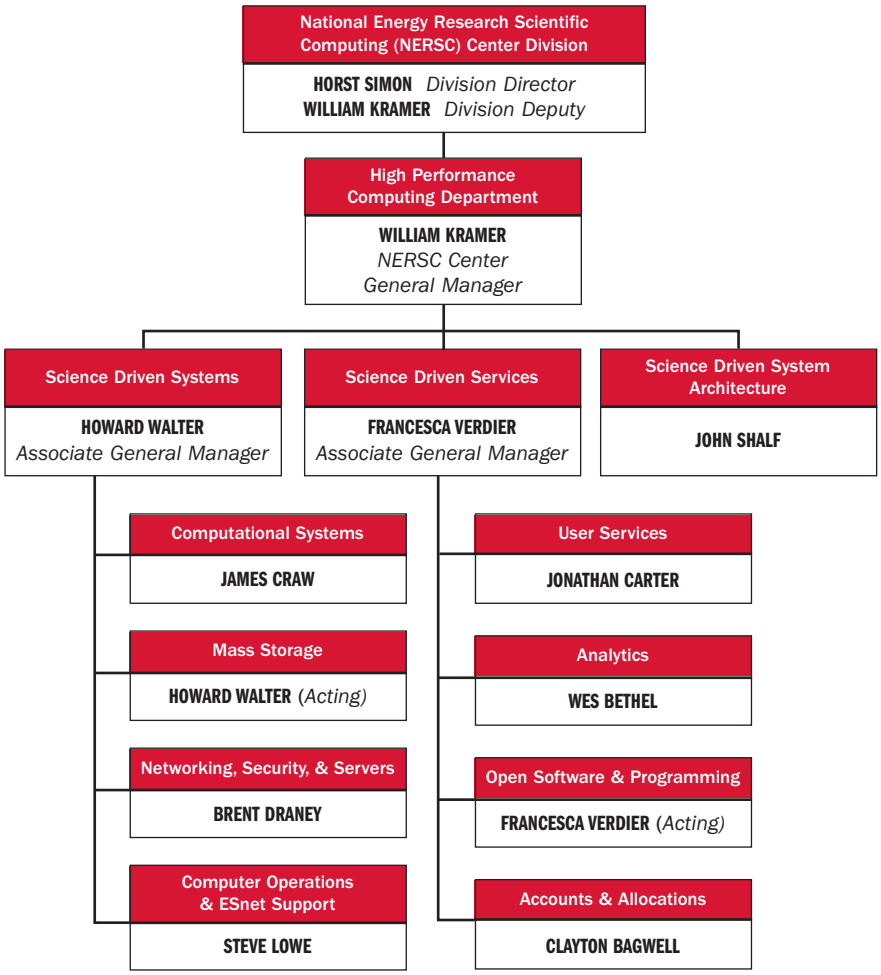

FIGURE 3. NERSC's new organization reflects new priorities and promotes coordination across groups and teams.

FIGURE 4. NERSC's new associate general managers, Francesca Verdier and Howard Walter, are responsible for Science-Driven Services and Science-Driven Systems, respectively.

FIGURE 5. Wes Bethel heads NERSC's new Analytics Team, while John Shalf heads the new Science-Driven Systems Architecture Team.

to realizing new solutions where none previously existed. The Analytics Team is a natural expansion of the visualization efforts that have been part of NERSC since it moved to Berkeley Lab. Wes Bethel is the team leader (Figure 5). NERSC's analytics strategy is discussed in more detail beginning on page 50.

- *Open Software and Programming (OSP) Group:* The growing use of open-source software and partially supported software requires a change of approach to NERSC's needs for the future. These areas now are a key component of NERSC's ability to provide high-quality systems and services. This group is responsible for the support and improvement of open-source and other partially supported software, particularly the software that NERSC uses for infrastructure, operations, and delivery of services. Key efforts include open-source engineering, development and support of middleware (Grid and Web tools), and NERSC's software infrastructure. Francesca Verdier is the acting group leader until a permanent one is recruited.

- *Science-Driven System Architecture (SDSA) Team:* This team performs ongoing evaluation and assessment of technology for scientific computing. The SDSA Team has expertise in benchmarking, system performance evaluations, workload monitoring, use of application modeling tools, and future algorithm scaling and technology assessment. Using scientific methods, the team will develop methods for analyzing possible technical alternatives and will create a clear understanding of current and future NERSC workloads. The SDSA Team will engage with vendors and the general research community to advocate technological features that will enhance the effectiveness of systems for NERSC scientists. The team is responsible for ongoing management of a suite of benchmarks that NERSC and Berkeley Lab use for architectural evaluation and procurement. This includes composite benchmarks and metrics such as SSP, ESP, variation, reliability, and usability. The team will matrix staff from both NERSC and Berkeley Lab's Computational Research Division for specific areas such as algorithm tracking and scaling, which are designed to develop and document future algorithmic requirements. The scientific focus for this effort will change periodically and will start with applied mathematics and astrophysics. The SDSA Team leader is John Shalf (Figure 5). Their current activities are discussed beginning on page 46.

Completing the reorganization, Jonathan Carter succeeds Francesca Verdier as leader of the User Services Group, and Brent Draney succeeds Howard Walter as leader of the Networking, Security and Servers Group (Figure 6).

FIGURE 6. Promoted to group leaders were Brent Draney of the Networking, Security and Servers Group and Jonathan Carter of the User Services Group.

Science-Driven Systems

NERSC deploys five major types of hardware systems—computational systems, filesystems, storage systems, network, and analytics and visualization systems—all of which must be carefully designed and integrated to maximize productivity for the entire computational community supported by the DOE Office of Science. In 2005 NERSC implemented major upgrades or improvements in all five categories, as described below.

System specifications must meet requirements that are constantly evolving because of technological progress and the changing needs of scientists. Therefore, NERSC is constantly planning for the future, not just by tracking trends in science and technology and planning new system procurements, but also by actively influencing the direction of technological development through efforts such as the Science-Driven System Architecture collaborations.

Two New Clusters: Jacquard and Bassi

In August 2005 NERSC accepted a 722-processor Linux Networx Evolocity cluster system named "Jacquard" for full production use (Figure 7). The acceptance test included a 14-day availability test, during which a select group of NERSC users were given full access to the Jacquard cluster to thoroughly test the entire system in production operation. Jacquard had a 99 percent availability uptime during the testing while users and scientists ran a variety of codes and jobs on the system.

The Jacquard system is one of the largest production InfiniBand-based Linux cluster systems and met rigorous acceptance criteria for performance, reliability, and functionality that are unprecedented for an InfiniBand cluster. Jacquard is the first system to deploy Mellanox 12x InfiniBand uplinks in its fat-tree interconnect, reducing network hot spots and improving reliability by dramatically reducing the number of cables required.

The system has 640 AMD 2.2 GHz Opteron processors devoted to computation, with the rest used for I/O, interactive work, testing, and interconnect management. Jacquard has a peak performance of 2.8 teraflop/s. Storage from DataDirect Networks provides 30 TB of globally available formatted storage.

Following the tradition at NERSC, the system was named for someone who has had an impact on science or computing. In 1801, Joseph-Marie Jacquard invented the Jacquard loom, which was the first programmable machine. The Jacquard loom used punched cards and a control unit that allowed a skilled user to program detailed patterns on the loom.

In January 2006, NERSC launched an 888-processor IBM cluster named "Bassi" into production use (Figure 8). Earlier, during the acceptance testing, users reported that codes ran from 3 to 10 times faster on Bassi than on NERSC's other IBM supercomputer, Seaborg, leading one tester to call the system the "best machine I have seen."

Bassi is an IBM p575 POWER5 system, and each processor has a theoretical peak performance of 7.6 gigaflop/s. The processors are distributed among 111 compute nodes with eight processors per node. Processors on each node have a shared memory pool of 32 GB. A Bassi node is an example of a shared memory processor, or SMP.

The compute nodes are connected to each other with a high-bandwidth, low-latency switching network. Each node runs its own full instance of the standard AIX operating system. The disk storage system is a distributed, paral-

FIGURE 7. Jacquard is a 722-processor Linux Networx Evolocity cluster system with a theoretical peak performance of 2.8 teraflop/s.

lel I/O system called GPFS (IBM's General Parallel File System). Additional nodes serve exclusively as GPFS servers. Bassi's network switch is the IBM "Federation" HPS switch, which is connected to a two-link network adapter on each node.

One of the test users for NERSC's two new clusters was Robert Duke of the University of North Carolina, Chapel Hill, the author of the PMEMD code, which is the parallel workhorse in modern versions of the popular chemistry code AMBER. PMEMD is widely used for molecular dynamics simulations and is also part of NERSC's benchmark applications suite. Duke has worked with NERSC's David Skinner to port and improve the performance of PMEMD on NERSC systems.

"I have to say that both of these machines are really nothing short of fabulous," Duke wrote to Skinner. "While Jacquard is perhaps the best-performing commodity cluster I have seen, Bassi is the best machine I have seen, period."

Other early users during the acceptance testing included the INCITE project team "Direct Numerical Simulation of Turbulent Nonpremixed Combustion" (see page 11). "Our project required a very long stretch of using a large fraction of Bassi processors—512 processors for essentially an entire month," recounted Evatt Hawkes. "During this period we experienced only a few minor problems, which is exceptional for a pre-production machine, and enabled us to complete our project against a tight deadline. We were very impressed with the reliability of the machine."

Hawkes noted that their code also ported quickly to Bassi, starting with

FIGURE 8. Bassi is an 888-processor IBM p575 POWER5 system with a theoretical peak performance of 6.7 teraflop/s.

a code already ported to Seaborg's architecture. "Bassi performs very well for our code. With Bassi's faster processors we were able to run on far fewer processors (512 on Bassi as opposed to 4,096 on Seaborg) and still complete the simulations more rapidly," Hawkes added. "Based on scalar tests, it is approximately 7 times faster than Seaborg and $1^{1}/_{2}$ times faster than a 2.0 GHz Opteron processor. Also, the parallel efficiency is very good. In a weak scaling test, we obtained approximately 78 percent parallel efficiency using 768 processors, compared with about 70 percent on Seaborg."

The machine is named in honor of Laura Bassi, a noted Newtonian physicist of the eighteenth century. Appointed a professor at the University of Bologna in 1731, Bassi was the first woman to officially teach at a European university.

New Visual Analytics Server: DaVinci

In mid-August, NERSC put into production a new server specifically tailored to data-intensive visualization and analysis. The 32-processor SGI Altix, called DaVinci (Figure 9), offers interactive access to large amounts of large memory and high performance I/O capabilities well suited for analyzing large-scale data produced by the NERSC high performance computing systems (Bassi, Jacquard, and Seaborg).

With its 192 gigabytes (GB) of RAM and 25 terabytes (TB) of disk, DaVinci's system balance is biased toward memory and I/O, which is different from the other systems at NERSC. This balance favors data-intensive analysis and interactive visualization. DaVinci has 6 GB of memory per processor, compared to 4 GB per processor on Jacquard and Bassi and 1 GB on Seaborg.

Users can obtain interactive access to 80 GB of memory from a single application (or all 192 GB of memory by prior arrangement), whereas the interactive limits on production NERSC supercomputing systems restrict interactive tasks to a smaller amount of memory (256 MB on login nodes). While DaVinci is available primarily for interactive use, the system is also configured to run batch jobs, especially those jobs that are data intensive.

The new server runs a number of visualization, statistics, and mathematics applications, including IDL, AVS/Express, CEI Ensight, VisIT (a parallel visualization application from Lawrence Livermore National Laboratory), Maple, Mathematica, and MatLab. Many users depend on IDL and MatLab to process or reorganize data in preparation for visualization. The large memory is particularly beneficial for these types of jobs.

FIGURE 9. DaVinci is a 32-processor SGI Altix with 6 GB of memory per processor and 25 TB of disk memory, a configuration designed for data-intensive analysis and interactive visualization.

DaVinci is connected to the NERSC Global Filesystem (see below), High Performance Storage System (HPSS), and ESnet networks by two independent 10 gigabit Ethernet connections.

With DaVinci now in production, NERSC has retired the previous visualization server, Escher, and the math server, Newton.

NERSC Global Filesystem

In early 2006, NERSC deployed the NERSC Global Filesystem (NGF) into production, providing seamless data access from all of the Center's computational and analysis resources. NGF is intended to facilitate sharing of data between users and/or machines. For example, if a project has multiple users who must all access a common set of data files, NGF provides a common area for those files. Alternatively, when sharing data between machines, NGF eliminates the need to copy large datasets from one machine to another. For example, because NGF has a single unified namespace, a user can run a highly parallel simulation on Seaborg, followed by a serial or modestly parallel post-processing step on Jacquard, and then perform a data analysis or visualization step on move a single data file.

NGF's single unified namespace makes it easier for users to manage their data across multiple systems (Figure 10). Users no longer need to keep track of multiple copies of programs and data, and they no longer need to copy data between NERSC systems for pre- and post-processing. NGF provides

several other benefits as well: storage utilization is more efficient because of decreased fragmentation; computational resource utilization is more efficient because users can more easily run jobs on an appropriate resource; NGF provides improved methods of backing up user data; and NGF improves system security by eliminating the need for collaborators to use "group" or "world" permissions.

"NGF stitches all of our systems together," said Greg Butler, leader of the NGF project. "When you go from system to system, your data is just there. Users don't have to manually move their data or keep track of it. They can now see their data simultaneously and access the data simultaneously."

NERSC staff began adding NGF to computing systems in October 2005, starting with the DaVinci visualization cluster and finishing with the Seaborg system in December. To help test the system before it entered production, a number of NERSC users were given preproduction access to NGF. Early users helped identify problems with NGF so they could be addressed before the filesystem was made available to the general user community.

"I have been using the NGF for some time now, and it's made my work a lot easier on the NERSC systems," said Martin White, a physicist at Berkeley Lab. "I have at times accessed files on NGF from all three compute platforms (Seaborg, Jacquard, and Bassi) semi-simultaneously."

NGF also makes it easier for members of collaborative groups to access data, as well as ensure data consistency by eliminating multiple copies of critical data. Christian Ott, a Ph.D. student and member of a team studying core-collapse supernovae, wrote that "the project directories make our collaboration much more efficient. We can now easily look at the output of the runs managed by other team members and monitor their progress. We are also sharing standard input data for our simulations."

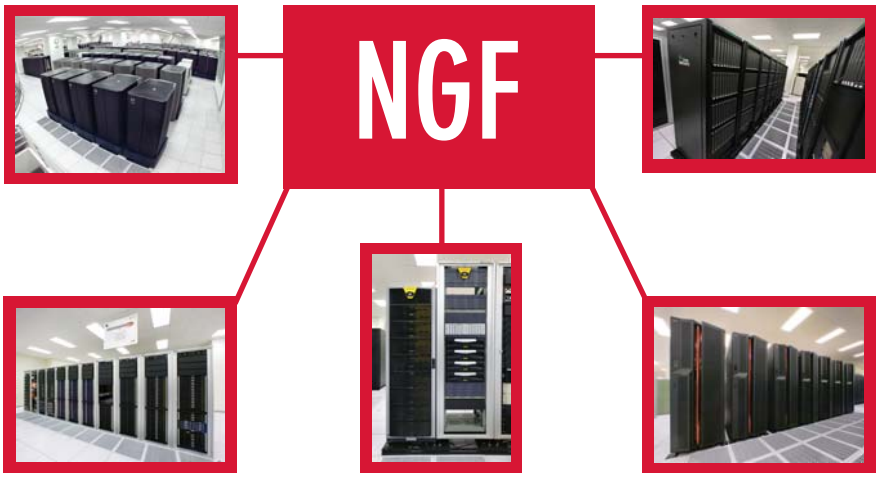

FIOURE 10. NGF is the first production global filesystem spanning five platforms (Seaborg, Jacquard, PDSF, DaVinci and Bassi), three architectures, and four different vendors.

NERSC General Manager Bill Kramer said that as far as he knows, NGF is the first production global filesystem spanning five platforms (Seaborg, Bassi, Jacquard, DaVinci, and PDSF), three architectures, and four different vendors. While other centers and distributed computing projects such as the National Science Foundation's TeraGrid may also have shared filesystems, Butler said he thinks NGF is unique in its heterogeneity.

The heterogeneous approach of NGF is a key component of NERSC's five-year plan. This approach is important because NERSC typically procures a major new computational system every three years, then operates it for five years to support DOE research. Consequently, NERSC operates in a heterogeneous environment with systems from multiple vendors, multiple platforms, different system architectures, and multiple operating systems. The deployed filesystem must operate in the same heterogeneous client environment throughout its lifetime.

Butler noted that the project, which is based on IBM's proven GPFS technology (in which NERSC was a research partner), started about five years ago. While the computing systems, storage, and interconnects were mostly in place, deploying a shared filesystem among all the resources was a major step beyond a parallel filesystem. In addition to the different system architectures, there were also different operating systems to contend with. However, the last servers and storage have now been deployed. To keep everything running and ensure a graceful shutdown in the event of a power outage, a large uninterruptible power supply has been installed in the basement of the Oakland Scientific Facility.

While NGF is a significant change for NERSC users, it also "fundamentally changes the Center in terms of our perspective," Butler said. For example, when the staff needs to do maintenance on the filesystem, the various groups need to coordinate their efforts and take all the systems down at once.

Storage servers, accessing the consolidated storage using the shared-disk filesystems, provide hierarchical storage management, backup, and archival services. The first phase of NGF is focused on function and not raw performance, but in order to be effective, NGF has to have performance comparable to native cluster filesystems. The current capacity of NGF is approximately 70 TB of user-accessible storage and 50 million inodes (the data structures for individual files). Default project quotas are 1 TB and 250,000 inodes. The system has a sustainable bandwidth of 3 GB/sec bandwidth for streaming I/O, although actual performance for user applications will depend on a variety of factors. Because NGF is a distributed network filesystem, performance will be only slightly less than that of filesystems that are local to NERSC compute platforms. This should only be an issue for applications whose performance is I/O bound.

NGF will grow in both capacity and bandwidth over the next several years, eventually replacing or dwarfing the amount of local storage on systems. NERSC is also working to seamlessly integrate NGF with the HPSS data archive to create much larger "virtual" data storage for projects. Once NGF is completely operational within the NERSC facility, Butler said, users at other centers, such as the National Center for Atmospheric Research and NASA Ames Research Center, could be allowed to remotely access the NERSC filesystem, allowing users to read and visualize data without having to execute file transfers. Eventually, the same capability could be extended to experimental research sites, such as accelerator labs.

NGF was made possible by IBM's decision to make its GPFS software available across mixed-vendor supercomputing systems. This strategy was a direct result of IBM's collaboration with NERSC. "Thank you for driving us in this direction," wrote IBM Federal Client Executive Mike Henesy to NERSC General Manager Bill Kramer when IBM announced the project in December 2005. "It's quite clear we would never have reached this point without your leadership!"

NERSC's Mass Storage Group collaborated with IBM and the San Diego Supercomputer Center to develop a Hierarchical Storage Manager (HSM) that can be used with IBM's GPFS. The HSM capability with GPFS provides a recoverable GPFS filesystem that is transparent to users and fully backed up and recoverable from NERSC's multi-petabyte archive on HPSS. GPFS and HPSS are both cluster storage software: GPFS is a shared disk filesystem, while HPSS supports both disk and tape, moving less-used data to tape while keeping current data on disk.

One of the key capabilities of the GPFS/HPSS HSM is that users' files are automatically backed up on HPSS as they are created. Additionally, files on the GPFS which have not been accessed for a specified period of time are automatically migrated from online resources as space is needed by users for files currently in use. Since the purged files are already backed up on HPSS, they can easily be automatically retrieved by users when needed, and the users do not need to know where the files are stored to access them. "This gives the user the appearance of almost unlimited disk storage space without the cost," said NERSC's Mass Storage Group Leader Nancy Meyer.

This capability was demonstrated in the Berkeley Lab and IBM booths at the SC05 conference. Bob Coyne of IBM, the industry co-chair of the HPSS executive committee, said, "There are at least ten institutions at SC05 who are both HPSS and GPFS users, many with over a petabyte of data, who have expressed interest in this capability. HPSS/GPFS will not only serve these existing users but will be an important step in simplifying the storage tools of the largest supercomputer centers and making them available to research institutions, universities, and commercial users."

"Globally accessible data is becoming the most important part of Grid computing," said Phil Andrews of the San Diego Supercomputer Center. "The immense quantity of information demands full vertical integration from a transparent user interface via a high performance filesystem to an enormously capable archival manager. The integration of HPSS and GPFS closes the gap between the long-term archival storage and the ultra high performance user access mechanisms."

The GPFS/HPSS HSM will be included in the release of HPSS 6.2 in spring 2006.

Integrating HPSS into Grids

NERSC's Mass Storage Group is currently involved in another development collaboration, this one with Argonne National Laboratory and IBM, to integrate HPSS accessibility into the Globus Toolkit for Grid applications.

At Argonne, researchers are adding functionality to the Grid file transfer daemon[2] so that the appropriate class of service can be requested from HPSS. IBM is contributing the development of an easy-to-call library of parallel I/O routines that work with HPSS structures and are also easy to integrate into the file transfer deamon. This library will ensure that Grid file transfer requests to HPSS movers are handled correctly.

[2] A daemon is a program that runs continuously in the background until it is activated by a particular event. The word *daemon* is Greek for "an attendant power or spirit."

NERSC is providing the HPSS platform and testbed system for IBM and Argonne to do their respective development projects. As pieces are completed, NERSC tests the components and works with the developers to help identify and resolve problems.

The public release of this capability is scheduled with HPSS 6.2, as well as future releases of the Globus Toolkit.

Bay Area MAN Inaugurated

On August 23, 2005, the NERSC Center became the first of six DOE research sites to go into full production on the Energy Science Network's (ESnet's) new San Francisco Bay Area Metropolitan Area Network (MAN). The new MAN provides dual connectivity at 20 to 30 gigabits per second (10 to 50 times the previous site bandwidths, depending on the site using the ring) while significantly reducing the overall cost.

The connection to NERSC consists of two 10-gigabit Ethernet links. One link is used for production scientific computing traffic, while the second is dedicated to special networking needs, such as moving terabyte-scale datasets between research sites or transferring large datasets which are not TCP-friendly.

"What this means is that NERSC is now connected to ESnet at the same speed as ESnet's backbone network," said ESnet engineer Eli Dart.

The new architecture is designed to meet the increasing demand for network bandwidth and advanced network services as next-generation scientific instruments and supercomputers come on line. Through a contract with Qwest Communications, the San Francisco Bay Area MAN provides dual connectivity to six DOE sites—the Stanford Linear Accelerator Center, Lawrence Berkeley National Laboratory, the Joint Genome Institute, NERSC, Lawrence Livermore National Laboratory, and Sandia National Laboratories/California (Figure 11). The MAN also provides high-speed access to Califor-nia's higher education network (CENIC), NASA's Ames Research Center, and DOE's R&D network, Ultra Science Net. The Bay Area MAN connects to both the existing ESnet production backbone and the first segments of the new Science Data Network backbone.

The connection between the MAN and NERSC was formally inaugurated on

FIGURE 12. DOE Office of Science Director Raymond Orbach (left) and Berkeley Lab Director Steven Chu made the ceremonial connection between NERSC and ESnet in June. After testing, the full production connection was launched in August.

FIGURE 11. ESnet's new San Francisco Bay Area Metropolitan Area Network provides dual connectivity at 20 to 30 gigabits per second to six DOE sites and NASA Ames Research Center.

June 24 by DOE Office of Science Director Raymond Orbach and Berkeley Lab Director Steven Chu (Figure 12).

Another Checkpoint/Restart Milestone

On the weekend of June 11 and 12, 2005, IBM personnel used NERSC's Seaborg supercomputer for dedicated testing of IBM's latest HPC Software Stack, a set of tools for high performance computing. To maximize system utilization for NERSC users, instead of "draining" the system (letting running jobs continue to completion) before starting this dedicated testing, NERSC staff checkpointed all running jobs at the start of the testing period. "Checkpointing" means stopping a program in progress and saving the current state of the program and its data—in effect, "bookmarking" where the program left off so it can start up later in exactly the same place.

This is believed to be the first full-scale use of the checkpoint/restart software with an actual production workload on an IBM SP, as well as the first checkpoint/restart on a system with more than 2,000 processors. It is the culmination of a collaborative effort between NERSC and IBM that began in 1999. Of the 44 jobs that were checkpointed, approximately 65% checkpointed successfully. Of the 15 jobs that did not checkpoint successfully, only 7 jobs were deleted from the queuing system, while the rest were requeued to run again at a later time. This test enabled NERSC and IBM staff to identify some previously undetected problems with the checkpoint/restart software, and they are now working to fix those problems.

In 1997 NERSC made history by being the first computing center to achieve successful checkpoint/restart on a massively parallel system, the Cray T3E.

Science-Driven System Architecture

The creation of NERSC's Science-Driven System Architecture (SDSA) Team formalizes an ongoing effort to monitor and influence the direction of technology development for the benefit of computational science. NERSC staff are collaborating with scientists and computer vendors to refine computer systems under current or future development so that they will provide excellent sustained performance per dollar for the broadest possible range of large-scale scientific applications.

While the goal of SDSA may seem ambitious, the actual work that promotes that goal deals with the nitty-gritty of scientific computing—for example, why does a particular algorithm perform well on one system but poorly on another—at a level of detail that some people might find tedious or overwhelming, but which the SDSA team finds fascinating and challenging.

"All of our architectural problems would be solvable if money were no object," said SDSA Team Leader John Shalf, "but that's never the case, so we have to collaborate with the vendors in a continuous, iterative fashion to work towards more efficient and cost-effective solutions. We're not improving performance for its own sake, but we are improving system effectiveness."

Much of the SDSA work involves performance analysis: how fast do various scientific codes run on different systems, how well do they scale to hundreds or thousands of processors, what kinds of bottlenecks can slow them down, and how can performance be improved through hardware development. A solid base of performance data is particularly useful when combined with workload analysis, which considers what codes and algorithms are common to NERSC's diverse scientific workload. These two sets of data lay a foundation for assessing how that workload would perform on alternative system architectures. Current architectures may be directly analyzed, while future architectures may be tested through simulations or predictive models.

The SDSA Team is investigating a number of different performance modeling frameworks, such as the San Diego Supercomputer Center's Memory Access Pattern Signature (MAPS), in order to assess their accuracy in predicting performance for the NERSC workload. SDSA team members are working closely with San Diego's Performance Modeling and Characterization Laboratory to model the performance of the NERSC-5 SSP benchmarks and compare the performance predictions to the benchmark results collected on existing and proposed HPC systems.

Another important part of the SDSA team's work is sharing performance and workload data, along with benchmarking and performance monitoring codes, with others in the HPC community. Benchmarking suites, containing application codes or their algorithmic kernels, are widely used for system assessment and procurement. NERSC has recently shared its SSP benchmarking suite with National Science Foundation (NSF) computer centers. With the Defense Department's HPC Modernization Program, NERSC has shared benchmarks and jointly developed a new one.

Seemingly mundane activities like these can have an important cumulative impact: as more research institutions set specific goals for application performance in their system procurement specifications, HPC vendors have to respond by offering systems that are specifically designed and tuned to meet the needs of scientists and engineers, rather than proposing strictly

off-the-shelf systems. By working together and sharing performance data with NERSC and other computer centers, the vendors can improve their competitive position in future HPC procurements, refining their system designs to redress any architectural bottlenecks discovered through the iterative process of benchmarking and performance modeling. The end result is systems better suited for scientific applications and a better-defined niche market for scientific computing that is distinct from the business and commercial HPC market.

The SDSA Team also collaborates on research projects in HPC architecture. One key project, in which NERSC is collaborating with Berkeley Lab's Computational Research Division and computer vendors, is ViVA, or Virtual Vector Architecture. The ViVA concept involves hardware and software enhancements that would coordinate a set of commodity scalar processors to function like a single, more powerful vector processor. ViVA would enable much faster performance for certain types of widely used scientific algorithms, but without the high cost of specialized processors. The research is proceeding in phases. ViVA-1 is focused on a fast synchronization register to coordinate processors on a node or multicore chip. ViVA-2 is investigating a vector register set that hides latency to memory using vector-like semantics. Benchmark scientific kernels are being run on an architectural simulator with ViVA enhancements to assess the effectiveness of those enhancements.

Perhaps the most ambitious HPC research project currently under way is the Defense Advanced Research Projects Agency's (DARPA's) High Productivity Computer Systems (HPCS) program. HPCS aims to develop a new generation of hardware and software technologies that will take supercomputing to the petascale level and increase overall system productivity ten-fold by the end of this decade. NERSC is one of several "mission partners" participating in the review of proposals and milestones for this project.

Proposals for New System Evaluated

As part of NERSC's regular computational system acquisition cycle, the NERSC-5 procurement team was formed in October 2004 to develop an acquisition plan, select and test benchmarks, and prepare a request for proposals (RFP). The RFP was released in September 2005; proposals were submitted in November and are currently being evaluated. The RFP set the following general goals for the NERSC-5 system:

- Support the entire NERSC workload, specifically addressing the DOE Greenbook recommendations (see page 36).
- Integrate with the NERSC environment, including the NERSC Global Filesystem, HPSS, Grid software, security and networking systems, and the user environment (software tools).
- Provide the optimal balance of the following system components:
  - computational: CPU speed, memory bandwidth, and latency
  - memory: aggregate and per parallel task
  - global disk storage: capacity and bandwidth
  - interconnect: bandwidth, latency, and scaling
  - external network bandwidth.

The RFP also stated specific goals for performance, disk storage, space and power requirements, software, etc. NERSC-5 is expected to significantly increase computational time for NERSC users beginning in the 2007 allocation year.

Two recent reports[3,4] on high-end computing recommended interagency collaboration on system procurements. The National Research Council report stated, "Joint planning and coordination of acquisitions will increase the efficiency of the procurement processes from the government viewpoint and will decrease variability and uncertainty from the vendor viewpoint."[5] NERSC-5 is possibly the first procurement involving collaboration with other government agencies. This collaboration includes the sharing of benchmarks with DOD and NSF as mentioned above. In addition, four organizations—the DOD HPC Modernization Program, the National Center for Supercomputing Applications, the Pittsburgh Supercomputing Center, and Louisiana State University—sent representatives to observe NERSC's Best Value Source Selection process. The NSF centers have adopted several of NERSC's procurement practices.

SC05 Conference Honor

In November 2005, NERSC demonstrated a production version of NGF at the SC05 conference, using it for the StorCloud Challenge. StorCloud was a special initiative for building a high performance computing storage capability showcasing HPC storage technologies (topologies, devices, interconnects) and applications. The StorCloud Challenge invited applicants from science and engineering communities to use the unique StorCloud infrastructure to demonstrate emerging techniques or applications, many of which consume enormous amounts of network and storage resources.

A NERSC/Berkeley Lab team, led by Will Baird of NERSC's Computational Systems Group, won an award for

[3] *Federal Plan for High-End Computing: Report of the High-End Computing Revitalization Task Force (HECRTF).* Washington, D.C.: National Coordination Office for Information Technology Research and Development, May 10, 2004.
[4] Susan L. Graham, Marc Snir, and Cynthia A. Patterson, eds., *Getting Up to Speed: The Future of Supercomputing,* Committee on the Future of Supercomputing, National Research Council. Washington, D.C.: The National Academies Press, 2005.
[5] Ibid., p. 171.

"Best Deployment of a Prototype for a Scientific Application." The team, which also included Jonathan Carter and Tavia Stone of NERSC and Michael Wehner, Cristina Siegerist and Wes Bethel of Berkeley Lab's Computational Research Division, used the StorCloud infrastructure "to test the wide-area deployment of an unprecedented system in support of a groundbreaking climate modeling application," according to the award. The application was fvCAM—or Finite Volume Community Atmospheric Model—which is being used to predict hurricane formation.

Baird designed the TRI Data Storm prototype around the concept of using an integrated, multisystem filesystem to improve the analysis of results produced by the demanding HPC application—the Community Atmospheric Model (CAM).

"Our aim was to take the output of CAM from a high-resolution grid, filter out the data of interest, and visualize the formation of storms in the North Atlantic basin," according to Baird. "This tool will be used in a study comparing real hurricane data with simulations. While this is a fairly generic workflow that could hold true for virtually any HPC application, the unique aspect to our approach is that there is a single high-performance parallel filesystem serving all of the different systems and applications."

For TRI Data Storm, the team used an externalized GPFS system, shared out by a dedicated cluster and mounted on all of the different computational resources used by their tool: the IBM POWER5 cluster Bassi and the PDSF Linux cluster at NERSC, the GPFS servers and storage at the conference, and an SGI Altix cluster in the Berkeley Lab booth at SC05.

"All the communication between the systems was simply through the filesystem, not through anything else," Baird said.

Science-Driven Services

NERSC's Science-Driven Services are designed to strike a balance between meeting the special needs of leading-edge computational research projects and providing responsive, comprehensive services for routine operations. Major issues addressed in NERSC's five-year plan include working with users to lower the gap between peak performance of terascale computing systems and the performance realized by scientific applications; providing special support for large-scale projects while at the same time maintaining the high level of support for all users; and helping users make productive use of the flood of scientific data from simulations and experiments.

Aspects of this strategy discussed below include support for large-scale projects, development of new data storage and retrieval strategies, and ongoing improvements in NERSC's everyday operations and services based on feedback from clients.

Support for Large-Scale Projects

NERSC works directly with scientists on major projects that require extensive scientific computing capabilities, such as the SciDAC and INCITE collaborations. These projects are often characterized by large collaborations, the development of community codes, and the involvement of computer scientists and applied mathematicians. In addition to high-end computing, these large projects handle issues in data management, data analysis, and data visualization, as well as automation features for resource management.

NERSC provides its highest level of support to these researchers, including special service coordination for queues, throughput, increased limits, etc.; and specialized consulting support, which may include algorithmic code restructuring to increase performance, I/O optimization, visualization support—whatever it takes to make the computation scientifically productive. The three INCITE projects for 2005 are good examples of this kind of support.

The INCITE project "Magneto-Rotational Instability and Turbulent Angular Momentum Transport," led by Fausto Cattaneo of the University of Chicago, is attempting to understand the forces that help newly born stars and black holes increase in size by simulating laboratory experiments that study magnetically caused instability (see page 26). "With the help of NERSC staff, we were able to tune our software for Seaborg's hardware and realize performance improvements that made additional simulations possible," Catteneo said. NERSC also provided crucial help in creating animated visualizations of the simulation results, which involve the formation of complex, three-dimensional structures that need to be seen to be understood.

For the INCITE project "Direct Numerical Simulation of Turbulent Nonpremixed Combustion," Jacqueline Chen, Evatt Hawkes, and Ramanan Sankaran of Sandia National Laboratories have performed the first 3D direct numerical simulations of a turbulent nonpremixed flame with detailed chemistry (see page 11). After analyzing and optimizing the code's performance with the help of NERSC staff, the researchers improved the code's efficiency by 45%. The simulations generated 10 TB of raw data, and NERSC consultants helped the researchers figure out the best strategy for efficiently transferring all that data from NERSC systems to the researchers' local cluster. "The assistance we

received from the NERSC computing staff in optimizing our code and with terascale data movement has been invaluable," Chen said. "The INCITE award has enabled us to extend our computations to three dimensions so that we may investigate interactions between turbulence, mixing, and finite-rate detailed chemistry in combustion."

The "Molecular Dynameomics" INCITE project, led by Valerie Daggett of the University of Washington, is an ambitious attempt to use molecular dynamics simulations to characterize and catalog the folding/unfolding pathways of representative proteins from all known protein folds (see page 30). David Beck, a graduate student in the Daggett lab, worked with NERSC consultants to optimize the performance of the group's code on Seaborg. "The INCITE award gave us a unique opportunity to improve the software, as well as do good science," Beck said. Improvements included load balancing, which sped up the code by 20%, and parallel efficiency, which reached 85% on 16-processor nodes. The INCITE award enabled the team to do five times as many simulations as they had previously completed using other computing resources. "We are quite satisfied with our experience at NERSC," Daggett commented.

Archiving Strategies for Genome Researchers

When researchers at the Production Genome Facility of DOE's Joint Genome Institute (JGI) found they were generating data faster than they could find somewhere to store the files, a collaboration with NERSC's Mass Storage Group developed strategies for improving the reliability of data storage while also making retrieval easier.

JGI is one of the world's leading facilities in the scientific quest to unravel the genetic data that make up living things. With advances in automatic sequencing of genomic information, scientists at the JGI's Production Genome Facility (PGF) found themselves overrun with sequence data, as their production capacity had grown so rapidly that data had overflowed the existing storage capacity (Figure 13). Since the resulting data are used by researchers around the world, PGF has to ensure that the data are reliably archived as well as easily retrievable.

As one of the world's largest public DNA sequencing facilities, the PGF produces 2 million files per month of trace data (25 to 100 KB each), 100 assembled projects per month (50 MB to 250 MB), and several very large assembled projects per year (~50 GB). In aggregate, this averages about 2,000 GB per month.

In addition to the amount of data, a major challenge is the way the data are produced. Data from the sequencing of many different organisms are produced in parallel each day, resulting in a daily archive that spreads the data for a particular organism over many tapes.

DNA sequences are considered the fundamental building blocks for the rapidly expanding field of genomics. Constructing a genomic sequence is an iterative process. The trace fragments are assembled, and then the sequence is refined by comparing it with other sequences to confirm the assembly. Once the sequence is assembled, information about its function is gleaned by comparing and contrasting the sequence with other sequences from both the same organism and other organisms. Current sequencing methods generate a large

FIGURE 13. JGI's automated sequencing facilities threatened to produce more genomic data than they could store or manage.

FIGURE 14. NERSC's High-Performance Storage System (HPSS) is capable of storing and retrieving all of JGI's genomic data efficiently.

volume of trace files that have to be managed—typically 100,000 files or more. And to check for errors in the sequence or make detailed comparisons with other sequences, researchers often need to refer back to these traces. Unfortun-ately, these traces are usually provided as a group of files with no information as to where the traces occur in the sequence, making the researchers' job more difficult.

This problem was compounded by the PGF's lack of sufficient online storage, which made organization (and subsequent retrieval) of the data difficult and led to unnecessary replication of files. This situation required significant staff time to move files and reorganize filesystems to find sufficient space for ongoing production needs; and it required auxiliary tape storage that was not particularly reliable.

Staff from NERSC's Mass Storage Group and the PGF agreed to work together to address two key issues facing the genome researchers. The most immediate goal was for NERSC's HPSS to become the archive for the JGI data, replacing the less-reliable local tape operation and freeing up disk space at the PGF for more immediate production needs (Figure 14). The second goal was to collaborate with JGI to improve the data handling capabilities of the genome sequencing and data distribution processes.

NERSC storage systems are robust and available 24 hours a day, seven days a week, as well as highly scalable and configurable. NERSC has high-quality, high-bandwidth connectivity to the other DOE laboratories and major universities provided by ESnet.

Most of the low-level data produced by the PGF are now routinely archived at NERSC, with ~50 GB of raw trace data being transferred from JGI to NERSC each night.

The techniques used in developing the archiving system allow it to be scaled up over time as the amount of data continues to increase—up to billions of files can be handled with these techniques. The data have been aggregated into larger collections which hold tens of thousands of files in a single file in the NERSC storage system. This data can now be accessed as one large file, or each individual file can be accessed without retrieving the whole aggregate.

Not only will the new techniques be able to handle future data, they also helped when the PGF staff discovered raw data that had been previously processed by software that had an undetected bug. The staff were able to retrieve the raw data from NERSC and reprocess it in about a month and a half, rather than go back to the sequencing machines and produce the data all over again—which would have taken about six months. In addition to saving time, this also saved money—a rough estimate is that the original data collection comprised up to 100,000 files per day at a cost of $1 per file, which added up to $1.2 million for processing six months' worth of data. Comparing this figure to the cost of a month and a half of staff time, the estimated savings are about $1 million—and the end result is a more reliable archive.

User Survey Provides Valuable Feedback

The results from the 2005 user survey show generally high satisfaction with NERSC's systems and support. Areas with the highest user satisfaction include account support services, the reliability and uptime of the HPSS mass storage system, and HPC consulting. The largest increases in satisfaction over last year's survey include the NERSC CVS server, the Seaborg batch queue structure, PDSF compilers, Seaborg uptime, available computing hardware, and network connectivity.

Areas with the lowest user satisfaction include batch wait times on both Seaborg and Jacquard, Seaborg's queue structure, PDSF disk stability, and Jacquard's performance and debugging tools. Only three areas were rated significantly lower this year: PDSF overall satisfaction and uptime, and the amount of time taken to resolve consulting issues. The introduction of three major systems in the last year combined with a reduction in consulting staff explain the latter.

Eighty-two users answered the question "What does NERSC do well?" Forty-seven respondents stated that NERSC gives them access to powerful computing resources without which they could not do their science; 32 mentioned excellent support services and NERSC's responsive staff; 30 pointed to very reliable and well managed hardware; and 11 said "Everything."

Sixty-five users responded to "What should NERSC do differently?" The areas of greatest concern are the interrelated issues of queue turnaround times (24 comments), job scheduling and resource allocation policies (22 comments), and the need for more or different computational resources (17 comments). Users also voiced concerns about data management, software, group accounts, staffing, and allocations.

As in the past, comments from the previous survey led to changes in 2005, including a restructuring of Seaborg's queuing polices, the addition of the new Jacquard and Bassi clusters, the upgrade of ESnet's connectivity to NERSC to 10 gigabits per second, and the installation of additional visualization software.

The complete survey results can be found at https://www.nersc.gov/news/survey/2005/.

Science-Driven Analytics

Simulations and experiments are generating data faster than it can be analyzed and understood. Addressing this bottleneck in the scientific discovery process is the emerging discipline of *analytics,* which has the simple goal of understanding data.

The term *analytics* refers to a set of interrelated technologies and intellectual disciplines that combine to produce insight and understanding from large, complex, disparate, and sometimes conflicting datasets. These technologies and disciplines include data management, visualization, analysis, and discourse aimed at producing specific types of understanding. These in turn rely on the computational infrastructure, expertise in using that infrastructure, and close cooperation between domain scientists, computational scientists, and computer scientists.

More specifically, the term *visual analytics* is the science of analytic reasoning facilitated by interactive visual interfaces. Its objective is to enable analysis of overwhelming amounts of information, and it requires human judgment to make the best possible evaluation of incomplete, inconsistent, and potentially erroneous information.

NERSC's analytics strategy builds on two of the Center's existing strengths: (1) proven expertise in effectively managing large, complex computing, infrastructure, and data storage systems to solve scientific problems of scale; and (2) exemplary user services, consulting, and domain scientific knowledge that help the NERSC user community effectively employ the Center's resources to solve challenging scientific problems. On this foundation, NERSC's analytics strategy adds an increased emphasis on facilities, infrastructure, expertise, and alliances that can be used to realize analytics solutions.

With the establishment of its new Analytics Team, NERSC is realigning its resources to support analytics activities. The NERSC Center's infrastructure is being broadened to include elements such as database deployment and support, with an increased focus on data analysis and scientific data management to support analytics. The existing visualization program is being expanded to include information visualization and integrated data management, analysis, and distributed computing. The goal is a well-rounded service and technology portfolio that is responsive to the analytics needs of NERSC's user community.

NERSC's analytics strategy includes five elements:

1. *Taking a proactive role in deploying emerging technologies.* NERSC will increasingly become a conduit for prototype technologies that emerge from the DOE computer science research community. Analytics will require adapting and deploying technologies from several different areas—data management, analysis, visualization, dissemination—into a unified workflow that functions effectively in a time-critical production environment. The role of NERSC staff will include deploying new system and support software, helping applications software engineers effectively use NERSC resources, and playing a proactive role in providing feedback to the original computer science researchers and developers to address security or performance concerns.

2. *Enhancing NERSC's data management infrastructure.* The NERSC Global Filesystem offers increased performance for all applications, including data-intensive analytics

tasks. It also helps streamline distributed workflows and provides high I/O rates, which are important for large datasets. NERSC also plans to increase its archival storage to nearly 40 PB over the next five years. In the near term, NERSC will evaluate and deploy software that provides distributed, file-level data management.

3. *Expanding NERSC's visualization and analysis capabilities.* One of the most significant activities performed by the NERSC visualization staff is in-depth, one-on-one consulting services, such as those provided to INCITE and other large projects. These activities typically involve finding or engineering solutions where none exist off the shelf. In addition, visualization staff will evaluate new visualization hardware and software technologies to determine which are beneficial to the user community. These technologies may include information visualization, which differs from the better-known scientific visualization in that the underlying data does not readily lend itself to spatial mapping—for example, comparing the results of genome alignment across multiple species. As data size and complexity grow, it will become increasingly crucial to use analysis technologies to reduce the processing load through the computational and visualization pipelines, as well as to reduce the "scientific processing load" on the humans who must interpret and understand the results. A portfolio of commercial, production, open-source, and research-grade technologies is expected to be most effective in meeting users' scientific needs.

4. *Enhancing NERSC's distributed computing infrastructure.* NERSC's strategy for supporting distributed computing will be tailored to provide services that have the broadest possible benefit and that conform to security requirements. In addition to providing low-level infrastructure such as the Open Grid Services Architecture (OGSA) and similar technologies that provide authentication and secure data movement across the network, NERSC will investigate and deploy higher-level applications and services that emerge from research and applications communities like the Particle Physics Data Grid and the Earth System Grid. Both of those projects rely on standard services for brokering access to data and tools that serve large, distributed user communities. NERSC will work closely with the user community to provide the documentation and assistance they need to construct analytics workflows.

5. *Understanding the analytics needs of the user community.* To be effective, NERSC's new program focus on Science-Driven Analytics will require additional information from the user community. To that end, the entire user community was surveyed in early 2006 to identify their most pressing analytics needs. The findings from this survey have been instrumental in shaping and prioritizing the emerging analytics effort. NERSC will continue soliciting input from users as well as tracking analytics trends in the larger scientific community.

APPENDIX A

NERSC Policy Board

Daniel A. Reed (Chair)
University of North Carolina, Chapel Hill

David Dean
(ex officio, NERSC Users Group Chair)
Oak Ridge National Laboratory

Robert J. Goldston
Princeton Plasma Physics Laboratory

Tony Hey
Microsoft Corporation

Sidney Karin
University of California, San Diego

Pier Oddone
Fermi National Accelerator Laboratory

Tetsuya Sato
Earth Simulator Center/Japan Marine Science and Technology Center

Stephen L. Squires
Hewlett-Packard Laboratories

APPENDIX B

NERSC Client Statistics

NERSC served 2,677 scientists throughout the United States in 2005. 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.

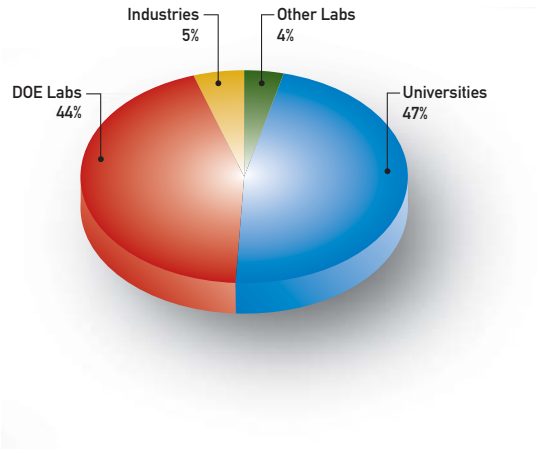

FIGURE 1. NERSC MPP usage by institution type, 2005.

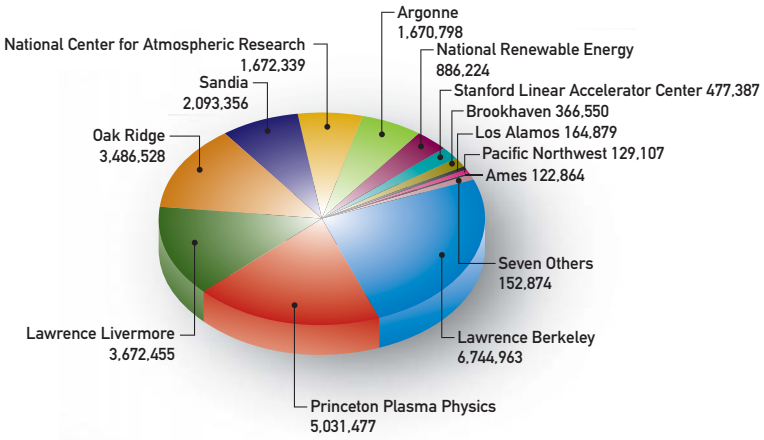

FIGURE 2. DOE and other Federal laboratory usage at NERSC, 2005 (MPP hours).

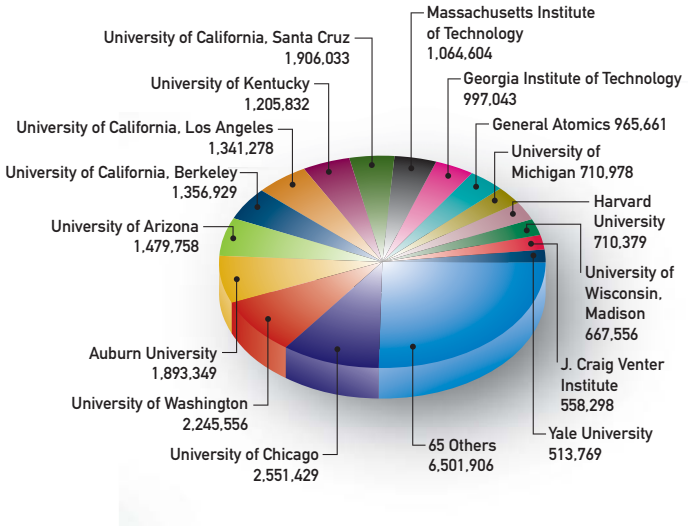

FIGURE 3. Academic and private laboratory usage at NERSC, 2005 (MPP hours).

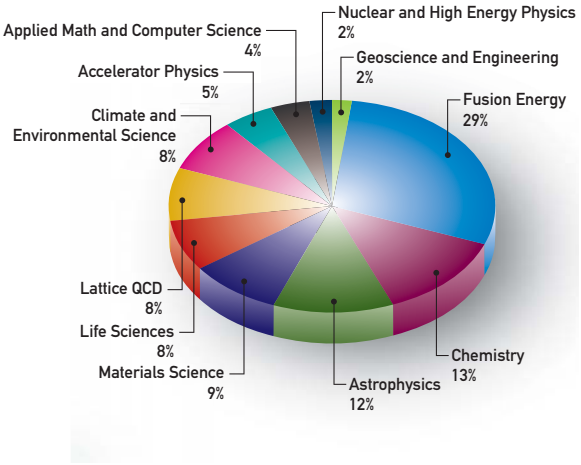

FIGURE 4. NERSC usage by scientific discipline, 2005.

APPENDIX C

NERSC Users Group Executive Committee

Members at Large

Yuen-Dat Chan, Lawrence Berkeley National Laboratory

Ricky A. Kendall, Ames Laboratory, Iowa State University

Donald A. Spong, Oak Ridge National Laboratory

Douglas Swesty, State University of New York at Stony Brook

Office of Advanced Scientific Computing Research

David Bernholdt, Oak Ridge National Laboratory

David Keyes, Old Dominion University

David Turner, Ames Laboratory

Office of Basic Energy Sciences

Feng Liu, University of Utah

G. Malcolm Stocks, Oak Ridge National Laboratory

Theresa L. Windus, Pacific Northwest National Laboratory

Office of Biological and Environmental Research

Michael Colvin, University of California, Merced, and Lawrence Livermore National Laboratory

Brian Hingerty, Oak Ridge National Laboratory

Vincent Wayland, National Center for Atmospheric Research

Office of Fusion Energy Sciences

Stephane Ethier (Vice Chair), Princeton Plasma Physics Laboratory

Alex Friedman, Lawrence Livermore and Lawrence Berkeley National Laboratories

Carl Sovinec, University of Wisconsin, Madison

Office of High Energy Physics

Olga Barranikova, Purdue University

Warren Mori, University of California, Los Angeles

Donald Sinclair, Argonne National Laboratory

Office of Nuclear Physics

David Dean (Chair), Oak Ridge National Laboratory

Patrick Decowski, Lawrence Berkeley National Laboratory

Douglas L. Olson, Lawrence Berkeley National Laboratory

APPENDIX D

Supercomputing Allocation Commitee

The Supercomputing Allocation Committee (SAC) is responsible for setting the policies associated with the utilization of computing, data storage, and other auxiliary services available to DOE Office of Science (SC) researchers and otherwise coordinating SC's computational projects. The Committee sets the distribution of NERSC and other available Office of Advanced Scientific Computing Research resources for scientific computing applications every year.

Committee Chair

Barbara Helland

Office of Advanced Scientific Computing Research

Barbara Helland

Office of Basic Energy Sciences

Richard Hilderbrandt

Office of Biological and Environmental Research

Michael Riches

Office of Fusion Energy Sciences

John Mandrekas

Office of High Energy Physics

Glen Crawford

Office of Nuclear Physics

Sidney Coon

APPENDIX E

Office of Advanced Scientific Computing Research

Mathematical, Information, and Computational Sciences Division

The primary mission of the Advanced Scientific Computing Research (ASCR) program, which is carried out by the Mathematical, Information, and Computational Sciences (MICS) subprogram, is to discover, develop, and deploy the computational and networking tools that enable researchers in the scientific disciplines to analyze, model, simulate, and predict complex phenomena important to the Department of Energy. To accomplish this mission, the program fosters and supports fundamental research in advanced scientific computing—applied mathematics, computer science, and networking—and operates supercomputer, networking, and related facilities. In fulfilling this primary mission, the ASCR program supports the Office of Science Strategic Plan's goal of providing extraordinary tools for extraordinary science as well as building the foundation for the research in support of the other goals of the strategic plan. In the course of accomplishing this mission, the research programs of ASCR have played a critical role in the evolution of high performance computing and networks. Berkeley Lab thanks the program managers with direct responsibility for the NERSC program and the MICS research projects described in this report:

Michael R. Strayer
Associate Director, ASCR, and Acting Director, MICS

Walter M. Polansky
ASCR Senior Technical Advisor for Project Management

Anil Deane
Program Manager for Applied Mathematics

Barbara Helland
Program Manager for NERSC and Leadership Computing Facility

Daniel Hitchcock
Program Manager for ESnet

Frederick C. Johnson
Program Manager for Computer Research

Thomas D. Ndousse-Fetter
Program Manager for Networking

Mary Anne Scott
Program Manager for Collaboratory Research

Yukiko Sekine
Program Manager for Scientific Data Management and Visualization

George R. Seweryniak
Program Manager

APPENDIX F

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.

Jill P. Dahlburg, Chair
Naval Research Laboratory

John W. Connolly, Vice Chair
University of Kentucky

F. Ronald Bailey
NASA Ames Research Center (retired)

Gordon Bell
Microsoft Bay Area Research Center

David J. Galas
Battelle Memorial Institute

Roscoe C. Giles
Boston University

James J. Hack
National Center for Atmospheric Research

Thomas A. Manteuffel
University of Colorado at Boulder

Ellen B. Stechel
Sandia National Laboratories

Virginia Torczon
College of William and Mary

Stephen S. Wolff
Cisco Systems

APPENDIX G

Acronyms and Abbreviations

AGM . . . . . Associate General Manager

AIX . . . . . Advanced Interactive eXecutive (an IBM operating system)

ALS. . . . . Advanced Light Source at Berkeley Lab

AMBER . . . Assisted Model Building with Energy Refinement

AMR . . . . . Adaptive mesh refinement

ASC . . . . . Advanced Simulation and Computing Program, National Nuclear Security Administration (DOE)

ASCR-MICS Office of Advanced Scientific Computing Research, Mathematical, Information, and Computational Sciences Division (DOE)

AVHF. . . . . Alexander von Humboldt Foundation (Germany)

BER . . . . . Office of Biological and Environmental Research (DOE)

BES . . . . . Office of Basic Energy Sciences (DOE)

BGL . . . . . Blue Gene/L, Argonne National Laboratory

BMBF . . . . Bundesministerium für Bildung und Forschung (Germany)

BNL . . . . . Brookhaven National Laboratory

BRAHMS . . Broad Range Hadron Magnetic Spectrometer at Brookhaven National Laboratory

BSPO . . . . Belgian Science Policy Office

CAM . . . . . Community Atmospheric Model

CCSE . . . . Center for Computational Sciences and Engineering (Berkeley Lab)

CCSM . . . . Community Climate System Model

CENIC . . . . Corporation for Education Network Initiatives in California

CMSO . . . . Center for Magnetic Self-Organization at the University of Chicago

COLTRIMS . Cold target recoil ion momentum spectroscopy

CPU . . . . . Central processing unit

CSIR . . . . . Council of Scientific and Industrial Research (India)

CVS . . . . . Concurrent Versions System

DAE . . . . . Department of Atomic Energy (India)

DARPA . . . Defense Advanced Research Projects Agency

DEFRA. . . . Department for Environment, Food and Rural Affairs (UK)

DFG . . . . . Deutsche Forschungsgemeinschaft (Germany)

DFT . . . . . Density functional theory

DNS . . . . . Direct numerical simulation

DOE . . . . . U.S. Department of Energy

DST . . . . . Department of Science and Technology (India)

EPSRC . . . Engineering and Physical Sciences Research Council (UK)

ES . . . . . Earth Simulator (Japan)

ESnet . . . . Energy Sciences Network (DOE)

FAPESP . . . Fundação de Amparo à Pesquisa do Estado de São Paulo (Brazil)

FES. . . . . Office of Fusion Energy Sciences (DOE)

FOM . . . . . Fundamenteel Onderzoek der Materie (Netherlands)

GACR . . . . Grant Agency of the Czech Republic

GB . . . . . Gigabyte

GMODGO. . Graduiertenkolleg Molekulare Organisation und Dynamik an Grenz- und Oberflächen (Germany)

GMRDP . . . Government Meteorological Research and Development Programme (UK)

GPFS . . . . General Parallel File System (IBM)

HadCM3 . . Hadley Centre Coupled Model, version 3

HEP . . . . . Office of High Energy Physics (DOE)

HPC . . . . . High performance computing

HPCMP . . . High Performance Computing Modernization Program (U.S. Department of Defense)

HPCS . . . . High Productivity Computer Systems

HPSS . . . . High Performance Storage System

HSM . . . . Hierarchical Storage Manager

*il*mm . . . . . *in lucem* Molecular Mechanics

IN2P3 . . . . Institut National de la Physique Nucléaire et de la Physique des Particules (France)

INCITE. . . . Innovative and Novel Computational Impact on Theory and Experiment (DOE)

I/O . . . . . Input/output

ITER . . . . . A multinational tokamak experiment to be built in France (Latin for "the way")

JGI . . . . . Joint Genome Institute (DOE)

JSPS . . . . . Japan Society for the Promotion of Science

KamLAND . Kamioka Liquid Scintillator Anti-Neutrino Detector (Japan)

KUL . . . . . Katholieke Universiteit Leuven (Belgium)

LANL . . . . . Los Alamos National Laboratory

LBNL . . . . . Lawrence Berkeley National Laboratory

LBW . . . . . Landesstiftung Baden-Württemberg (Germany)

MAN . . . . . Metropolitan area network

MAPS . . . . Memory Access Pattern Signature

MB . . . . . Megabyte

MEC . . . . . Ministry of Education of China

MECS . . . . Ministerio de Educación y Ciencia (Spain)

MEXT . . . . Ministry of Education, Culture, Sports, Science and Technology (Japan)

MPP . . . . . Massively parallel processing

MSCF . . . . Molecular Science Computing Facility at Pacific Northwest National Laboratory

NASA . . . . National Aeronautics and Space Administration

NCAR . . . . National Center for Atmospheric Research

NCCS . . . . National Center for Computational Sciences at Oak Ridge National Laboratory

NERSC . . . National Energy Research Scientific Computing Center

NGF . . . . . NERSC Global Filesystem

NMR . . . . . Nuclear magnetic resonance

NNSFC . . . National Natural Science Foundation of China

NOAA . . . . National Oceanographic and Atmospheric Administration

NP . . . . . Office of Nuclear Physics (DOE)

NPACI . . . . National Partnership for Advanced Computational Infrastructure

NSA . . . . . Near-surface alloy

NSF . . . . . National Science Foundation

OGSA . . . . Open Grid Services Architecture

ORNL . . . . Oak Ridge National Laboratory

OSP . . . . . Open Software and Programming Group (NERSC)

PB . . . . . Petabyte

PCM . . . . . Parallel Climate Model

PDSF . . . . Parallel Distributed Systems Facility (NERSC)

PGF . . . . . Production Genome Facility (JGI)

PHENIX . . . Pioneering High Energy Nuclear Interaction Experiment at Brookhaven National Laboratory

PHOBOS . . A detector experiment at Brookhaven National Laboratory

PMEMD . . . Particle mesh Ewald molecular dynamics

PPDG . . . . Particle Physics Data Grid

PSCSR . . . Polish State Committee for Scientific Research

QGP . . . . . Quark-gluon plasma

RFP . . . . . Request for proposals

RHIC . . . . Relativistic Heavy Ion Collider at Brookhaven National Laboratory

RMST . . . . Russian Ministry of Science and Technology

SC . . . . . Office of Science (DOE)

SciDAC . . . Scientific Discovery through Advanced Computing (DOE)

SDSA . . . . Science-Driven System Architecture (NERSC)

SNSF . . . . Swiss National Science Foundation

SSP . . . . . Sustained System Performance benchmark

STAA . . . . . Science and Technology Assistance Agency (Slovakia)

STAR . . . . Solenoidal Tracker at RHIC

TB . . . . . Terabyte

TCP . . . . . Transmission Control Protocol

UC . . . . . University of California

UCB . . . . . University of California, Berkeley

USAF-OSR . . U.S. Air Force Office of Scientific Research

UWM . . . . University of Wisconsin-Madison

ViVA . . . . . Virtual Vector Architecture

WHOI . . . . Woods Hole Oceanographic Institution

FOR MORE INFORMATION

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/

Published by the Berkeley Lab Creative Services Office in collaboration with NERSC researchers and staff. JO#11554

Editor and principal writer: John Hules.

Contributing writers: Jon Bashor, Lynn Yarris, and Paul Preuss (Berkeley Lab); Robert Sanders (UC Berkeley); Mario Aguilera and Cindy Clark (Scripps Institution of Oceanography); Karen McNulty Walsh (Brookhaven National Laboratory); David Terraso (Georgia Institute of Technology); NCAR Media Relations Office.

Lawrence Berkeley National Laboratory / 1 Cyclotron Road / Berkeley CA / 94720

LBNL-60296