NATIONAL ENERGY RESEARCH
SCIENTIFIC COMPUTING CENTER

ANNUAL REPORT 2006

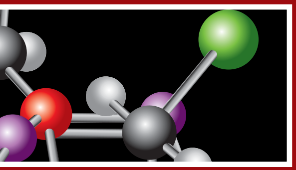

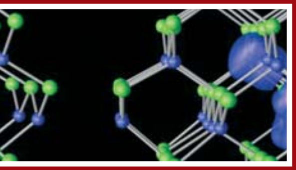

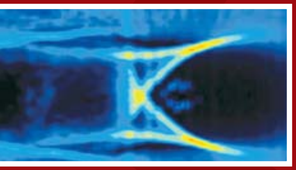

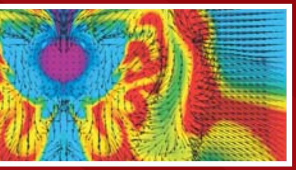

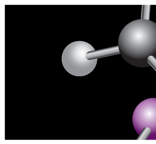

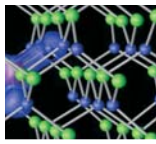

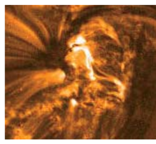

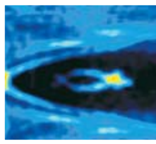

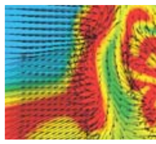

NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

ANNUAL REPORT 2006

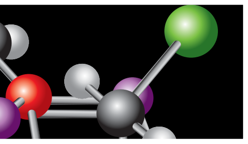

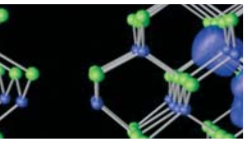

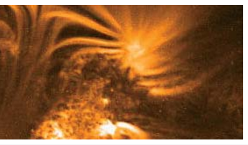

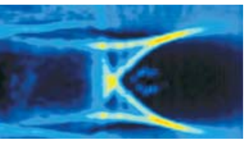

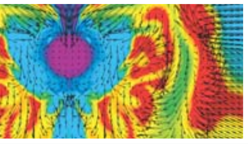

Ernest Orlando Lawrence
Berkeley National Laboratory

1 Cyclotron Road, Berkeley, CA
94720-8148

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

LBNL-63342, August 2007

Table of Contents

The Year in Perspective

This will be the last time that I will write this introduction to the NERSC annual report. When you read this, it will be almost 12 years since I returned to Berkeley and took on the challenge together with many dedicated colleagues to rebuild NERSC as a new center. One of the purposes of moving NERSC to Berkeley Lab in 1996 was to bring the benefits of supercomputing to a wider array of applications than was previously supported. One of the first new additions to NERSC's scientific roster was astrophysics, with Saul Perlmutter's Supernova Cosmology Project and with studies of cosmic microwave background radiation (CMB) by George Smoot's research group. Both groups had already been using smaller computers for data analysis but were eager to tackle larger datasets using NERSC's resources.

Those new collaborations bore fruit almost immediately, with Perlmutter's co-discovery in 1998 of the dark energy that is accelerating the expansion of the Universe, and with the BOOMERANG Consortium's discovery in 2000, based on analysis of CMB data, that the geometry of the Universe is flat. The latter finding was based on Smoot's earlier discovery, with John Mather and colleagues, of the blackbody form and anisotropy of the CMB, for which Smoot and Mather have been honored with the 2006 Nobel Prize for physics.

Over the past ten years, George Smoot and his colleagues have used nearly 5 million processor-hours and tens of terabytes of disk space at NERSC, and around 100 analysts from a dozen CMB experiments are now NERSC users. In fact, NERSC has become the computational center around which this community has coalesced as both its computational demands and results increase dramatically. The CMB sky map that Smoot and Mather produced in 1992 used only 6,144 pixels, but the newest CMB model can map 75 billion observations to 150 million pixels, making cosmology a highly precise science.

Research at NERSC over the past year by other astrophysicists has led to new insights into the formation of high-mass stars as well as the forces that make supernovas explode. Getting back down to earth, climate researchers have shown that the global warming that has already happened will produce bigger hurricanes, longer heat waves, and more extreme weather by the end of this century.

Materials scientists and chemists have also had a very productive year at NERSC. One research group achieved a breakthrough in surface plasmon resonance, a common but, up till now, expensive technique for measuring binding interactions, such as those between DNA and proteins. They developed a low-cost crystal array to make a portable and highly sensitive sensor that can be used in diagnostic bioassays and in research ranging from drug discovery to immunology, virology, and other fields. Another group's calculations show that zigzag graphene nanoribbons could serve as the basis for nanosized spintronic devices. Other researchers have shown why doping strengthens grain boundaries and how to overcome nanocrystals' resistance to doping.

Plasma physics saw the development of the first self-consistent model for the spontaneous onset of fast magnetic reconnection. This new finding may help scientists to better predict which solar storms pose the greatest threat to communications and other satellites, and it may also lead to a better understanding of how to control plasmas in fusion reactors.

An INCITE award of 2.5 million processor hours at NERSC was used to create full-scale, three-dimensional, explicit particle models that revealed important physical details of laser wakefield accelerator experiments that accelerated electron beams to energies exceeding a billion electron volts (GeV) in a distance of just 3.3 centimeters. These simulations coupled with experiments are developing the detailed understanding of laser acceleration needed to apply this technology to future higher energy particle physics experiments and to compact machines for medicine and laboratory science.

On the technical side, NERSC has taken a major step forward with the acquisition of what will be the largest Cray XT4 system in the world. When completed later in 2007, the system will have more than 19,000 processors and will deliver sustained performance of at least 16 teraflop/s when running a suite of diverse scientific applications at scale. This system will increase NERSC's sustained computational capability by almost a factor of 10.

NERSC has also laid the foundation for an analytics infrastructure that combines hardware, software, and the development and application of analytics technologies such as data management, data analysis and data mining, visual data exploration, and workflow management. These technologies will help NERSC users spend more time doing research and less time managing data and struggling with analytics software.

When the DOE's Advanced Scientific Computing Research Advisory Committee (ASCAC) formed a subcommittee to develop performance metrics for petascale facilities, it was natural for NERSC to take a leadership role, since we have long used goals and metrics to ensure that what we do is meeting the needs of DOE and its scientists. Among the subcommittee's recommendations were project-specific services, like those NERSC provides to SciDAC and INCITE projects, and the use of a standard user survey based on the one NERSC has used for several years to measure and improve service.

Looking to the future, NERSC is collaborating with computer scientists to meet the software challenges of petascale computing and manycore architectures, which will have hundreds to thousands of cores per processor. The explosion in hardware parallelism will require a complete redesign of applications, libraries, and algorithms—and possibly new programming languages—to fully utilize petascale systems. We are heading into a time of great innovation in high performance computing, and as always, NERSC will be influencing and implementing those innovations.

As I look back on my decade as director of the NERSC Center Division, I feel tremendous gratitude for the dedication, skills, and accomplishments of the NERSC staff, who are never satisfied with the status quo and are always looking for ways to make NERSC an even more productive resource for scientific researchers. When I hand over leadership of the division to my successor, I will do so with the confidence that he or she will have the support of the best scientific computing organization in the world.

Horst D. Simon
NERSC Center Division Director

Research News

a WARMER, STORMIER WO

RLD

By the end of this century, bigger hurricanes, longer heat waves, and more extreme weather will be evident

The verdict is in from the Intergovernmental Panel on Climate Change (IPCC). The summary of its Fourth Assessment Report says the world is already committed to centuries of warming, shifting weather patterns, and rising seas from the human production of greenhouse gases, but warming can be substantially blunted with prompt action.

Those conclusions were based on the research of thousands of scientists worldwide, including the climate simulations created by Warren Washington and his colleagues at the National Center for Atmospheric Research (NCAR) and elsewhere using CCSM3, a climate code whose development was funded primarily by the National Science Foundation (NSF) and the Department of Energy (DOE). These simulations investigate the response of the Earth's climate to future emissions scenarios that represent different policy choices for energy use and global development.

Data produced by these simulations are freely available to the research and education community via the DOE Earth System Grid. Among the recent studies based on these and other simulations are two that forecast more severe storms and more extreme weather in general.

GOING TO THE EXTREMES

Many previous studies have looked at how average temperature or rainfall might change in the next century as greenhouse gases increase. However, a new study titled "Going to the extremes: An intercomparison of model-simulated historical and future changes in extreme events"[1] looks more specifically at how weather extremes could change.

"It's the extremes, not the averages, that cause the most damage to society and to many ecosystems," said NCAR scientist Claudia Tebaldi, lead author for the report. "We now have the first model-based consensus on how the risk of dangerous heat waves, intense rains, and other kinds of extreme weather will change in the next century."

Tebaldi and colleagues based their work on simulations from nine different climate models, including CCSM3, for the periods 1980–1999 and 2080–2099. The simulations were created on supercomputers at NERSC and other research centers in France, Japan, Russia, and the United States. Each model simulated the 2080–2099 interval three times, varying the extent to which greenhouse gases accumulate in the atmosphere. These three scenarios were used to account for uncertainty over how fast society may act to reduce emissions of carbon dioxide and other greenhouse gases over coming decades.

From the model output, the scientists computed ten different indices of climate extremes, with five related to temperature and five to moisture. For instance, a frost days index measures how many days per year temperatures dip below 32 degrees Fahrenheit, while a dry days index measures the length of each year's longest consecutive string of days without rain or snow. Because the impact of a given index can be stronger in one climatic zone than another, the authors expressed the results in terms of statistical significance at each location.

For all three greenhouse-gas scenarios, the models agree that by 2080–2099:

- The number of extremely warm nights and the length of heat waves will increase significantly over nearly all land areas across the globe. During heat waves, very warm nights are often associated with fatalities because people and buildings have less chance to cool down overnight.
- Most areas above about 40 degrees latitude north will see a significant jump in the number of days with heavy precipitation (days with more than 0.4 inches). This includes the northern tier of U.S. states, Canada, and most of Europe.

PROJECT

Climate Change Simulations with CCSM3: Moderate and High Resolution Studies

PRINCIPAL INVESTIGATOR

Warren Washington, National Center for Atmospheric Research

SENIOR INVESTIGATORS

Jerry Meehl, Lawrence Buja, NCAR

FUNDING

BER, NSF

---

[1] Claudia Tebaldi, Katharine Hayhoe, Julie M. Arblaster, and Gerald A. Meehl, "Going to the extremes: An intercomparison of model-simulated historical and future changes in extreme events," Climatic Change **79,** 185 (2006). Funding: BER, NSF, EPA.

FIGURE 1. A thunderstorm cloud passes over the plains east of Denver. The number of days with heavy precipitation is expected to increase in the northern tier of U.S. states. (Photo by Carlye Calvin, ©UCAR, used with permission)

- Dry spells could lengthen significantly across the western United States, southern Europe, eastern Brazil, and several other areas. Dry spells are one of several factors in producing and intensifying droughts.
- The average growing season could increase significantly across most of North America and Eurasia.

Most of these trends are significantly weaker for the lowest-emission scenario than for the moderate and high-emission scenarios. Thus, the authors add, lowering the output of greenhouse gases over the next century should reduce the risk that the most severe changes will occur.

BREEDING BIGGER HURRICANES

Rising ocean temperatures in key hurricane breeding grounds of the Atlantic and Pacific oceans are due primarily to human-caused increases in greenhouse gas concentrations, according to a study published in the September 11, 2006 issue of the Proceedings of the National Academy of Sciences (PNAS).[2]

Using 22 different computer models of the climate system, including CCSM3, Benjamin Santer and six other atmospheric scientists from Lawrence Livermore National Laboratory, together with Tom Wigley, Gerald Meehl, and Warren Washington from NCAR, and collaborators from eight other research centers, have shown that the warming sea surface temperatures (SSTs) of the tropical Atlantic and Pacific oceans over the last century are linked to human activities.

"We've used virtually all the world's climate models to study the causes of SST changes in hurricane formation regions," Santer said.

Research published during the past year has uncovered evidence of a link between rising ocean temperatures and increases in hurricane intensity. This has raised concerns about the causes of the rising temperatures, particularly in parts of the Atlantic and Pacific where hurricanes and other tropical cyclones form.

---

[2] B. D. Santer, T. M. L. Wigley, P. J. Gleckler, C. Bonfils, M. F. Wehner, K. AchutaRao, T. P. Barnett, J. S. Boyle, W. Brüggemann, M. Fiorino, N. Gillett, J. E. Hansen, P. D. Jones, S. A. Klein, G. A. Meehl, S. C. B. Raper, R. W. Reynolds, K. E. Taylor, and W. M. Washington, "Forced and unforced ocean temperature changes in Atlantic and Pacific tropical cyclogenesis regions," PNAS **103,** 13905 (2006).

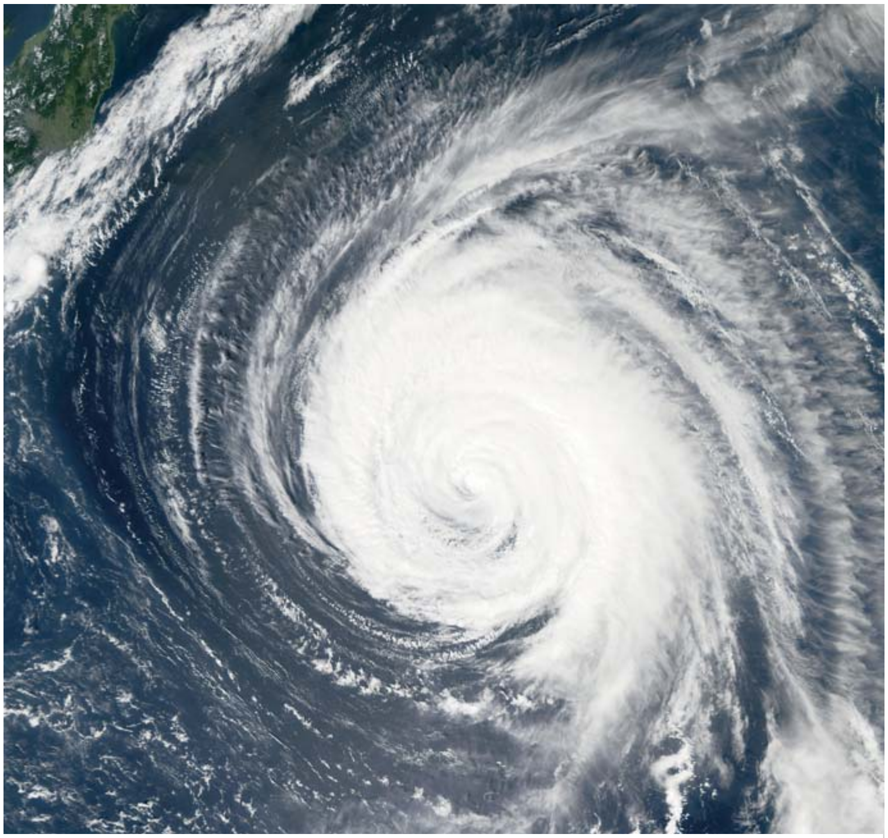

FIGURE 2. Hurricane Ioke passes by the Hawaiian Islands on August 21, 2006, with 132-mile-per-hour winds in this satellite image. The storm, renamed Typhoon Ioke as it moved west across the International Date Line, later intensified to become the most powerful central Pacific storm on record. (Image: Hal Pierce, SSAI/NASA GSFC)

Previous efforts to understand the causes of changes in SSTs have focused on temperature changes averaged over very large ocean areas, such as the entire Atlantic or Pacific basins. The new research specifically targets SST changes in much smaller hurricane formation regions.

"The important conclusion is that the observed SST increases in these hurricane breeding grounds cannot be explained by natural processes alone," said Wigley. "The best explanation for these changes has to include a large human influence."

Hurricanes are complex phenomena that are influenced by a variety of physical factors, such as SSTs, wind shear, water vapor, and atmospheric stability. The increasing SSTs in the Atlantic and Pacific hurricane formation regions are not the sole determinant of hurricane intensity, but they are likely to be one of the most significant influences.

"It is important to note that we expect global temperatures and SSTs to increase even more rapidly over the next century," Wigley said. According to Santer, "In a post-Katrina world, we need to do the best job we possibly can to understand the complex influences on hurricane intensity, and how our actions are changing those influences."

Other institutions contributing to this study include the University of California, Merced; Lawrence Berkeley National Laboratory; Scripps Institution of Oceanography; the University of Hamburg; the University of East Anglia; Manchester Metropolitan University; NASA's Goddard Institute for Space Studies; and NOAA's National Climatic Data Center.

IMPROVING HURRICANE DEFENSES

After the devastation caused by hurricanes Katrina and Rita, the Federal Emergency Management Agency (FEMA) asked the U.S. Army Corps of Engineers to run a series of simulations estimating hurricane-induced storm surge elevations to help improve hurricane defenses along the Gulf Coast. To assist in this effort, the DOE Office of Science allocated 800,000 processor hours of supercomputing time at NERSC to this project.

"NERSC … has a well-earned reputation for providing highly reliable systems, fast turnaround on critical projects, and dedicated support for users," said Secretary of Energy Samuel Bodman when announcing the allocation. "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."

As hurricanes move from the ocean toward land, the force of the storm causes the seawater to rise as it surges inland. The Corps of Engineers used its DOE supercomputer allocations to create revised models for predicting the effects of 100-year storm-surges—the worst-case scenario based on 100 years of hurricane data—along the Louisiana, Mississippi, and Texas coast lines (Figures 3 and 4). In particular, simulations were generated for the critical five-parish area of Louisiana surrounding New Orleans and the Lower Mississippi River.

PROJECT

Coastal Storm Surge Analyses

PRINCIPAL INVESTIGATOR

Jay Ratcliff, US Army Corps of Engineers

FUNDING

ASCR, USACOE

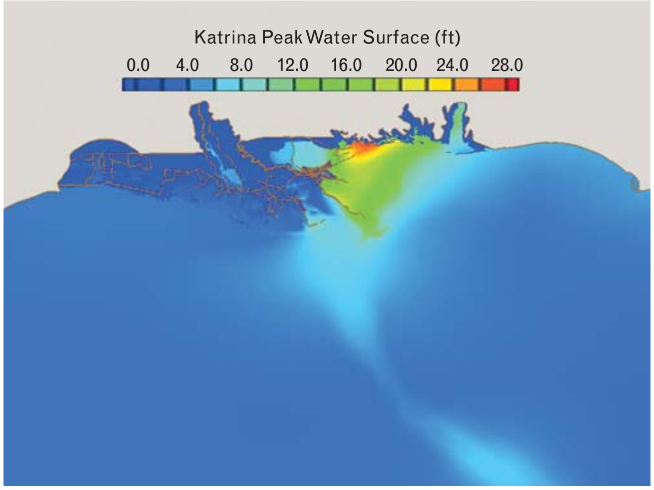

FIGURE 3. Overview simulation showing elevated storm surges along the Gulf Coast.

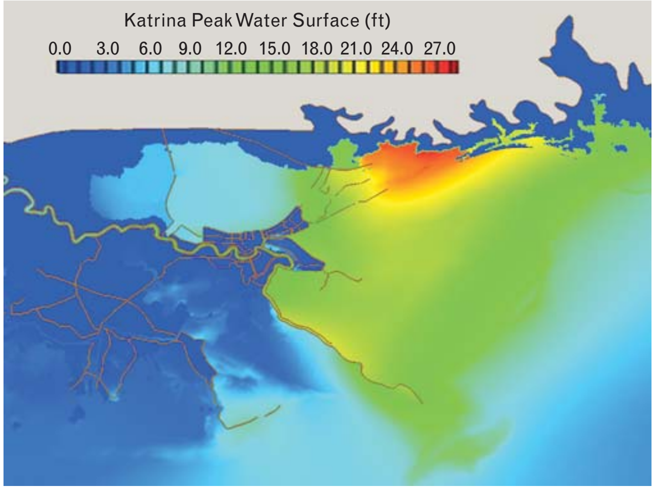

FIGURE 4. Simulation detail showing highest surge elevation (in red) striking Biloxi, Miss. New Orleans is the dark blue crescent to the lower left of Biloxi.

These revised models of the effects known as "storm-surge elevations" are serving as the basis of design for levee repairs and improvements currently being designed and constructed by the Corps of Engineers in the wake of Hurricane Katrina's destruction in the New Orleans Metro Area.

Additionally, Gulf Coast Recovery Maps were generated for Southern Louisiana based on FEMA's revised analysis of the frequency of hurricanes and estimates of the resulting waves. These maps are being used on an advisory basis by communities currently rebuilding from the 2005 storms.

Having access to the NERSC supercomputers allowed the Corps of Engineers to create more detailed models of the effects of Hurricane Rita and other storms along the Gulf Coast. Increased detail gave the Corps of Engineers and FEMA more information about the local effects of such storms.

For example, storm surge elevations are greatly influenced by local features such as roads and elevated railroads. Representing these details in the model greatly improves the degree to which computed elevations match observed storm surge high-water marks and allows the Corps to make better recommendations to protect against such surges.

The Corps of Engineers team also ran hurricane simulations on the DoD Major Shared Resource computers at the Engineering Research and Development Center (ERDC). Due to the tremendous computational requirements of these hurricane protection projects and urgent timelines, only by working together and using both DOE and DoD resources was the Corps able to provide high-quality engineering solutions.

As a result of the computer runs, the Corps determined that the applications produced incorrect results at topographic boundaries in some instances, and codes were modified to improve the accuracy of the results. For example, the runs at NERSC have improved the Corps' ability to model the effects of vegetation and land use on storm surges which propagate far inland, as Hurricane Rita did on Sept. 24, 2005.

This article written by: David Hosansky, NCAR; Jon Bashor and John Hules, Berkeley Lab.

UNDER the COMPUTATION MICRO

Materials scientists and chemists are using computational tools to understand, improve, and create new kinds of materials

AL

SCOPE

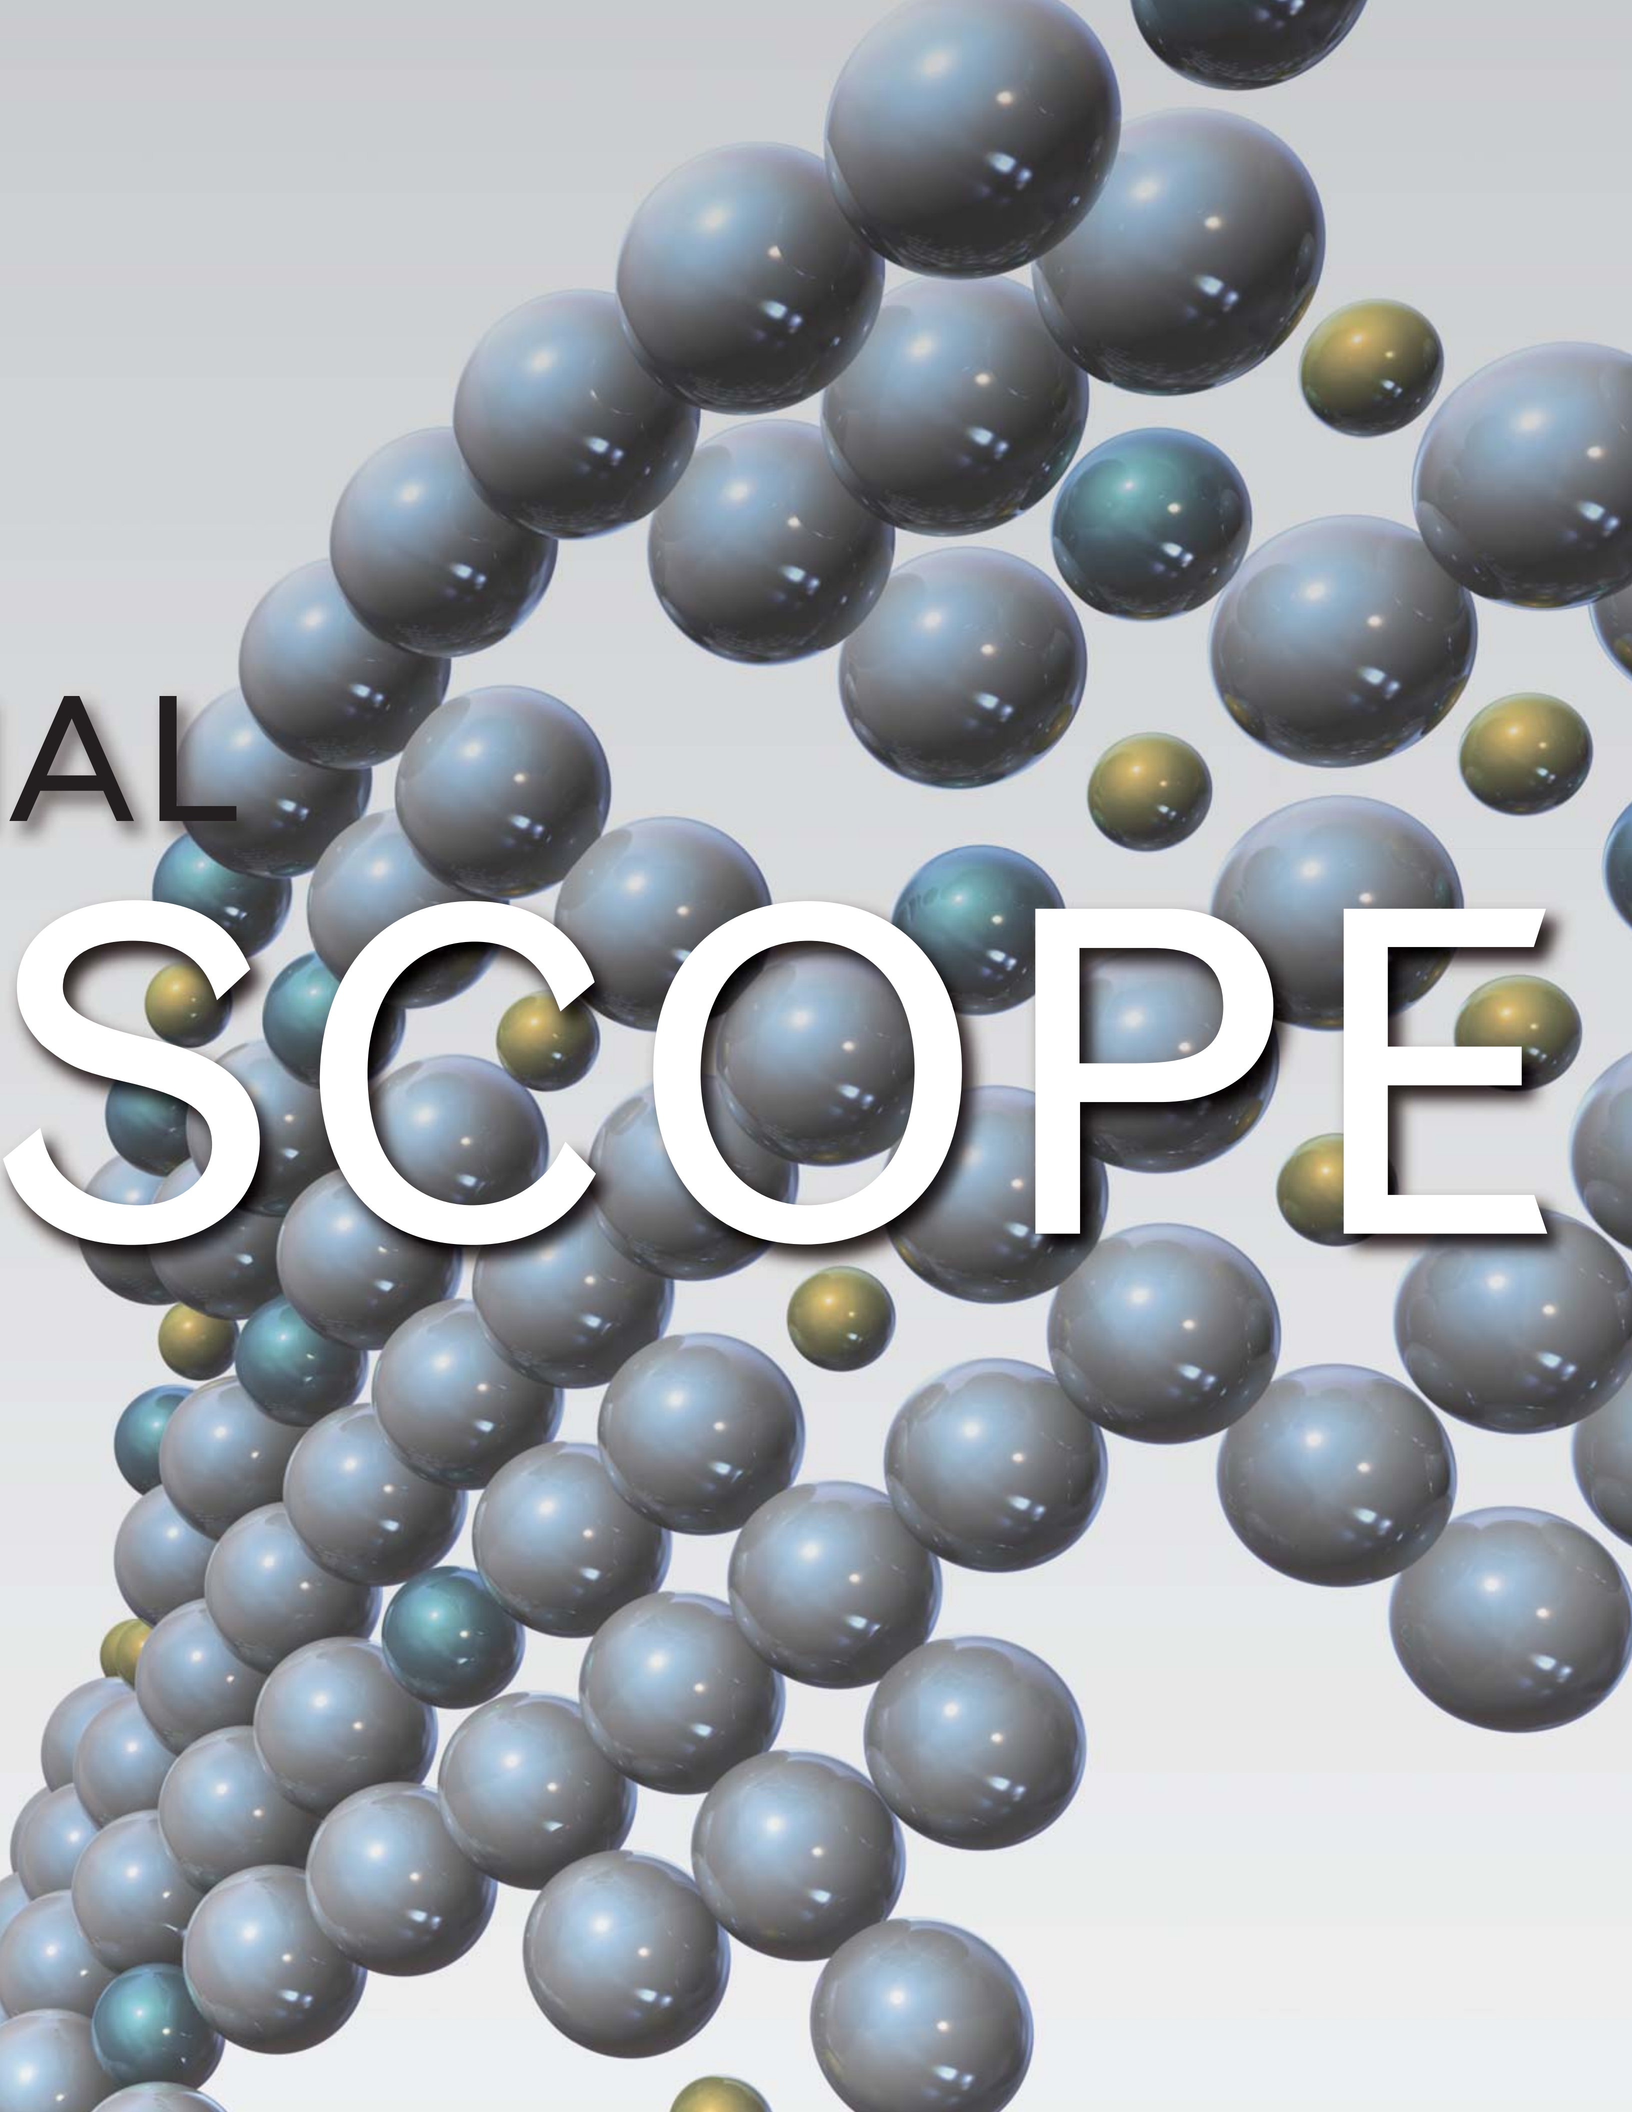

Optical microscopy and, later, electron microscopy may have revolutionized the study of nature, but computational modeling is filling in the blanks in the study of the very small. Computational simulations can reveal the structures and reaction mechanisms of unstable or reactive chemicals that are difficult to study experimentally; and they can also help scientists interpret the results of their experiments by elucidating electrical, magnetic, optical, and structural properties at the smallest scales. This section highlights some of this year's achievements in chemistry and materials science.

GRAPHENE NANORIBBONS: A NEW PATH TO SPINTRONICS

Spintronics—the principle behind electronic devices based on the spin of an electron, in addition to its charge—is the gleam in the collective eyes of the computer industry. With the discovery of a new pathway towards realizing the spintronics dream, that gleam should light up even brighter.

Marvin Cohen and Steven Louie, theorists who hold joint appointments with Berkeley Lab's Materials Sciences Division and the University of California at Berkeley, together with postdoctoral researcher Young-Woo Son, have calculated that nanoribbons of graphene—single-layered sheets of hexagonally-arranged carbon atoms—with zigzag edges can be made to carry a spin current.[1] Zigzag graphene nanoribbons could therefore serve as the basis for nanosized spintronic devices.

"Our calculations show that zigzag graphene nanoribbons are magnetic and can carry a spin current in the presence of a sufficiently large electric field," said Cohen. "An applied transverse electric field transforms the ribbon from a semiconductor with a small gap to a metal with carriers that are 100 percent spin-polarized. By carefully controlling the electric field, it should be possible to generate, manipulate, and detect electron spins and spin currents in spintronics applications."

Louie added, "There are, of course, many challenges to confront before this concept can be used for applications. However, if electric fields can be made to produce and manipulate a 100 percent spin-polarized carrier system through a chosen geometric structure, it will revolutionize spintronics technology."

Spintronic devices promise to be smaller, faster, and far more versatile than today's electronic devices. Spin is a quantum mechanical property that arises when the intrinsic rotational momentum of a particle, in this case an electron, creates a tiny magnetic field. For the sake of simplicity, spin is given the direction of either "up" or "down." The up or down values of spin can be used to encode data in the 0s and 1s of the binary system, just like the positive or negative values of an electrical charge. However, unlike charge-based data storage, spin-based data storage does not disappear when the electric current stops.

PROJECT

Ab Initio Prediction of Structural and Electronic Properties of Materials

PRINCIPAL INVESTIGATOR

Marvin Cohen, University of California, Berkeley, and Lawrence Berkeley National Laboratory

SENIOR INVESTIGATORS

Steven Louie, UC Berkeley and LBNL

FUNDING

BES, NSF

COMPUTING RESOURCES

NERSC, SDSC

One of the keys to the future development of spintronic technology is the curious, extremely rare class of materials known as "half metals." These materials are unique because their conducting electrons are all spin-polarized in either the up or down orientation. Conduction takes place by charge carriers exclusively oriented in a single spin direction; the spin polarization of the carriers in half metals is theoretically 100 percent, making them ideal for spintronic device structures.

The search for half metals among semiconductors has been intense, but to date there have been few investigations into organic materials, even though carbon-based nanostructures hold significant promise for future electronic devices.

"Although there are organic magnets in molecular and polymeric forms, carbon in the crystalline form is not magnetic," Louie said. "This is probably the reason why people haven't been looking for half metals in carbon-based nanostructures such as graphene."

Using first-principles calculations, which can predict a material's electrical and mag-

---

[1] Young-Woo Son, Marvin L. Cohen, and Steven G. Louie, "Half-metallic graphene nanoribbons," Nature **444**, 347, (2006).

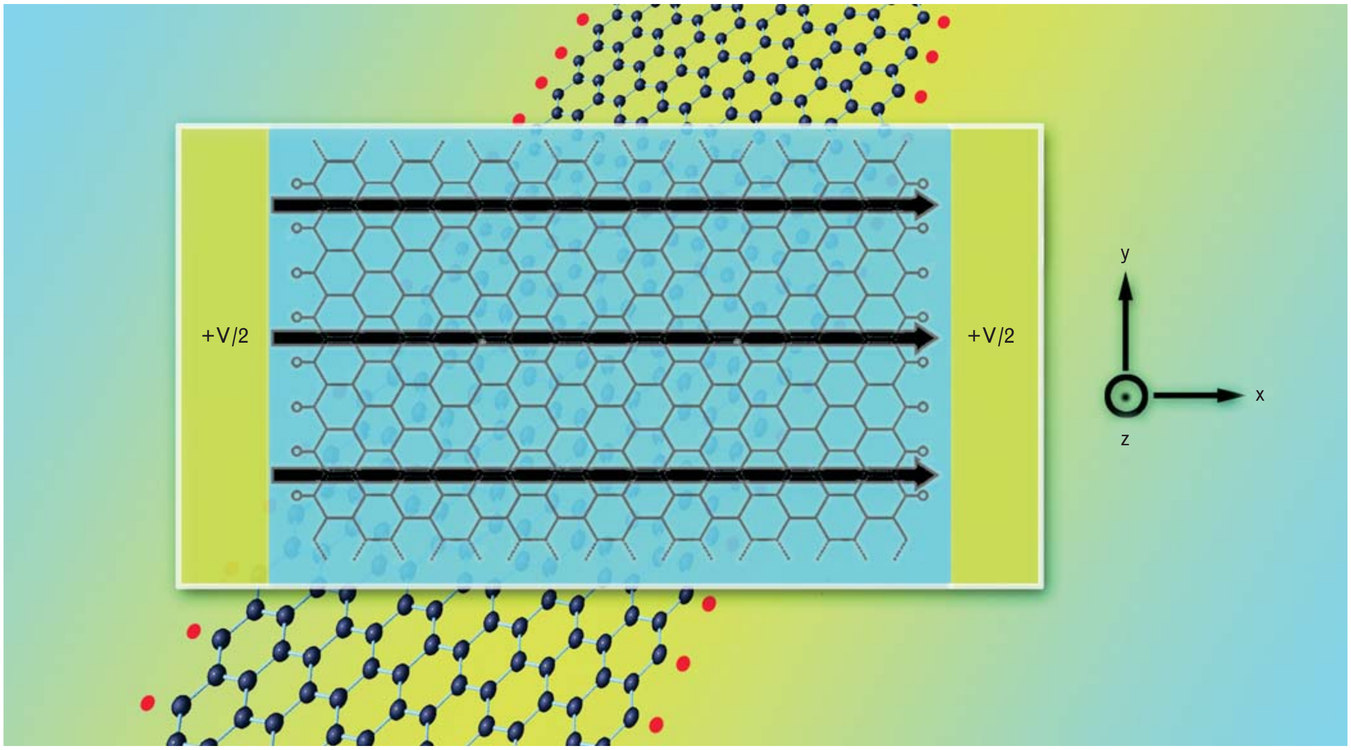

FIGURE 1. An external transverse electric field has been applied in the x direction across this zigzag graphene nanoribbon, which is between 1.5 to 6.7 nanometers wide (billionths of a meter). The field makes the nanoribbon magnetically positive towards the right side, so that the application of a small longitudinal field would generate spin-polarized currents along the y direction. Hydrogen atoms on the edges are denoted by circles or red dots.

netic properties from the atomic number and mass of its constituent atoms, plus the brute computational power of NERSC, Cohen, Louie, and Son were able to demonstrate the half-metallic phenomenon in graphene nanoribbons. They showed that the half-metallic property emerges when homogeneous electric fields are applied across graphene nanoribbons whose zigzag-shaped edges are attached to voltage contacts (Figure 1). Zigzag edges form when the carbon chemical bonds of a honeycombed sheet of graphene are uniformly cut.

"The electric fields can be used to directly manipulate the spin orientation of carriers in zigzag graphene nanoribbons by shifting the energy of different magnetic states via the electric field," Louie said. "We believe this is the first demonstration of such an effect."

Cohen commented, "It's very hard to polarize spins and even harder to manipulate them. Usually one needs magnetic fields that involve large pieces of equipment, which makes the production of small devices difficult. Here we get to have our cake and eat it too, because it is the electric field rather than the magnetic field that gives us the highest degree of polarization."

Basing spin polarization on electric fields makes it much easier to work with small devices. "Also, the zigzag graphene nanoribbons are resistant to heat buildup, relatively insensitive to temperature, and are non-toxic," said Cohen. "All we need now is for clever nanoscientists to use our first-principles calculations to design and make these or similar systems."

PRACTICAL PLASMONIC CRYSTAL BIOSENSORS

In the realm of myths, legends, and the occult, crystals have always been believed to have extraordinary powers, from protection and healing with crystal talismans to foretelling the future with crystal balls. Scientific discoveries about crystals may be less dramatic, but they are no less amazing. In a recent example, researchers at the University of Illinois, Urbana-Champaign (UIUC) and Argonne National Laboratory developed a small, low-cost crystal array that makes a highly sensitive biosensor, and used computational modeling to explain how it works.

What the researchers achieved was a breakthrough in a common but, until now, expensive technique for measuring binding interactions, such as those between DNA and proteins, based on changes in the refractive index near a metal surface. The technique, called surface plasmon resonance or SPR, is used in diagnostic bioassays and in research ranging from drug discovery to immunology, virology, and other fields.

Just as a pebble tossed into a pond produces waves on the surface of the water, a beam of light shining on a plasmonic crystal produces electron waves on the crystal's surface. “SPR is simply light causing a collective excitation of electrons near the surface of a metal,” explained Stephen Gray, a chemist at Argonne National Laboratory, who created the simulations that analyzed the experimental results.

“Those collective excitations of electrons are like waves on the metal's surface,” Gray continued, “and light shining around different objects above the metal creates different waves. When we use plasmonic crystals as sensors, small changes in the material specimen produce changes in the refraction index which can be measured with a spectrophotometer. From those responses, you can infer what the material is and how it has changed.”

PROJECT

Nanoscale Electrodynamics

PRINCIPAL INVESTIGATOR

Stephen Gray, Argonne National Laboratory

FUNDING

BES

COMPUTING RESOURCES

NERSC, LCRC

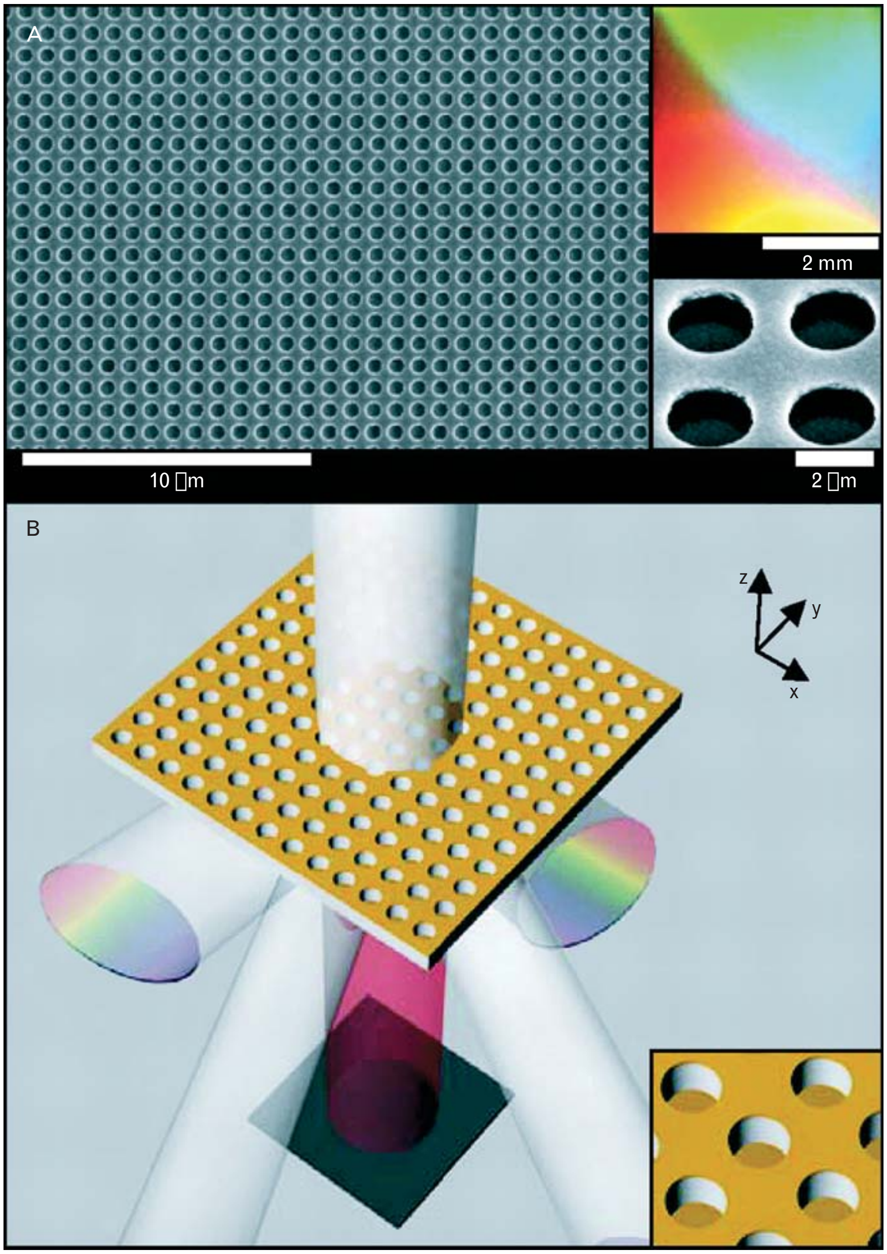

FIGURE 2. Images and schematic illustrations of a quasi-3D plasmonic crystal. (A) Scanning electron micrograph (SEM) of a crystal. (Upper Inset) A low-resolution optical image illustrating the diffraction colors produced by these structures. (Lower Inset) A high-magnification SEM that shows the upper and lower levels of gold. (B) Schematic illustration of the normal incidence transmission mode geometry used to probe these devices. The intensity of the undiffracted, transmitted light is monitored across the UV, visible, and near-infrared regions of the spectrum. (Inset) A close-up schematic illustration of the crystal.

While the scientific understanding of SPR may be modern, its application has a long history. “The ruby red color in some medieval stained-glass windows is the result of surface plasmon resonance,” Gray pointed out. “Gold nanoparticles that were added to the glass scatter and absorb light in a way that produces a pleasing color.”

As a sensing technology, SPR has the advantage of not requiring that fluorescent labels be added to samples, as in fluorescence mi-

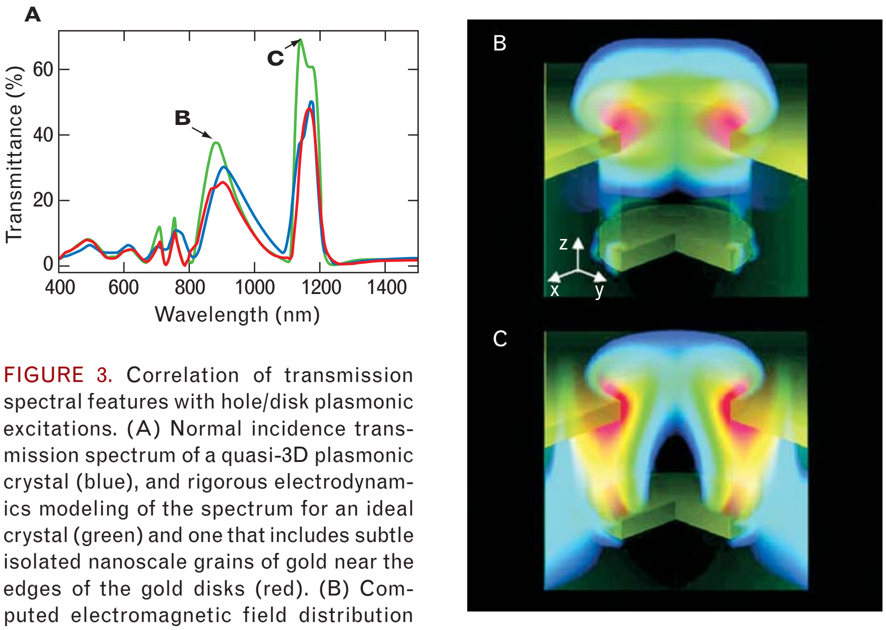

FIGURE 3. Correlation of transmission spectral features with hole/disk plasmonic excitations. (A) Normal incidence transmission spectrum of a quasi-3D plasmonic crystal (blue), and rigorous electrodynamics modeling of the spectrum for an ideal crystal (green) and one that includes subtle isolated nanoscale grains of gold near the edges of the gold disks (red). (B) Computed electromagnetic field distribution associated with the resonance at 883 nm (labeled B in A). The intensity is concentrated at the edges of the nanoholes in the upper level of the crystal. (C) Field distribution associated with the resonance at 1,138 nm (labeled C in A), showing strong coupling between the upper and lower levels of the crystal.

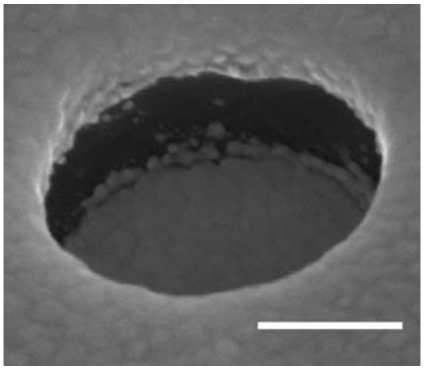

FIGURE 4. Tilted SEM image of an individual nanohole showing grains of gold at the edges of the recessed gold disk.

croscopy. But SPR to date has had a variety of limitations. The first SPR systems used prisms, but those systems were too bulky to be portable. More recent systems employ plasmonic crystals in the form of nanostructured films or nanoparticles; these systems are more portable but less sensitive, and fabricating large and uniform arrays of plasmonic crystals has been prohibitively expensive.

But all that may be changing. In an experiment reported in the Proceedings of the National Academy of Sciences,[2] Ralph Nuzzo, John Rogers and co-workers at UIUC's Frederick Seitz Materials Research Laboratory developed a low-cost crystal array to make a highly sensitive sensor. Using soft nanoimprint lithography, a technique that uses a soft polymeric mold to stamp and create structures on a substrate, the researchers created a plasmonic crystal consisting of a regular array of cylindrical wells in gold film on a polyurethane substrate (Figure 2). The SPR effects were produced on the nanoscale holes in the gold film and on the separate gold disks at the bottoms of the wells. SPR waves can be modeled mathematically using Maxwell's equations, so Gray was able to do a detailed computational analysis of the optical properties of the new crystals and the complex electromagnetic field distributions around the multilevel nanostructured features (Figure 3).

Interestingly, Gray's initial idealized crystal model produced spectral features (Figure 3A, green line) that did not quite match the experimental results (blue line); but when he added small defects in the form of isolated grains of gold on the sides of the wells near the bottom, the match was close to perfect (red line). Scanning electron micrographs confirmed that there were indeed grains of gold at the edges of the recessed gold disks (Figure 4).

"This showed how, at the nanoscale, very small defects can have important effects," Gray said. It also shows how computational modeling could be used to figure out how to fine-tune the performance of the system.

The unusual geometry and uniformity of these crystals gives them high sensitivity to multiple wavelengths over large sample areas with micrometer spatial resolution. The research team used a well studied ligand–receptor pair, biotin and avidin, as a model system to illustrate the functionality of these crystals in a quantitative analytical bioassay, and they were able to detect molecular binding in a single layer.

Because these plasmonic crystal arrays are smaller than typical sensing or imaging systems, and in view of their high sensitivity, low-cost fabrication, and simple readout apparatus, this technology could be used in developing the next generation of portable diagnostic sensors. They could easily be integrated into microfluidic lab-on-a-chip instrumentation.

OVERCOMING NANOCRYSTALS' RESISTANCE TO DOPING

The word *doping* has different meanings in sports and in materials science, but both meanings have one thing in common: enhanced performance through intentionally introduced impurities. In sports, of course, the

---

[2] Matthew E. Stewart, Nathan H. Mack, Viktor Malyarchuk, Julio A. N. T. Soares, Tae-Woo Lee, Stephen K. Gray, Ralph G. Nuzzo, and John A. Rogers, "Quantitative multispectral biosensing and 1D imaging using quasi-3D plasmonic crystals," PNAS **103,** 17143 (2006).

"impurities" are drugs. But in semiconductor production, doping pure silicon with small quantities of impurities such as boron or phosphorus enables engineers to tailor the electrical conductivity and other properties of the material for specific electronic applications.

A small amount of dopant can make a big difference: one atom of dopant per 10,000 silicon atoms is considered heavy doping in today's semiconductor manufacturing; light doping could mean a ratio of 1:100,000,000. But materials scientists would like to dope semiconductor nanocrystals that may have fewer than 10,000 atoms to begin with, because nanoelectronics holds the promise of more efficient solar cells, electroluminescent devices, computers, and much more.

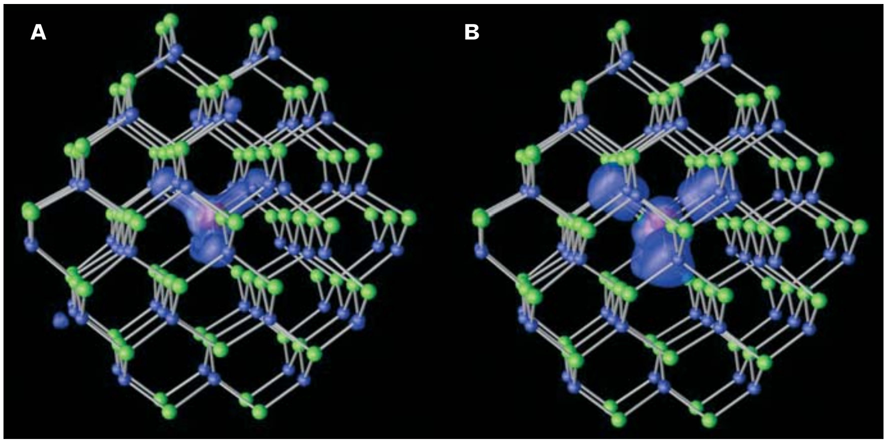

FIGURE 5. Charge density plot showing magnesium impurities in cadmium-selenium nanocrystals at two different gap levels: (a) resonant and (b) hybrid.

Nanocrystals tend to have fewer defects or impurities than bulk materials, and experiments have shown that it is often more difficult to dope a nanocrystal than a bulk material. But the reason for this difficulty has been a matter of debate. Since nanocrystals are so small—they can have as much surface as inner volume—one obvious explanation is that it is easy for a dopant atom to move to the surface and escape during the process of crystal formation, a process called *self-purification.*

That "kinetic argument" was not good enough for Gustavo Dalpian and James Chelikowsky of the Institute for Computational Engineering and Sciences at the University of Texas at Austin. (Dalpian is now at the Universidade Federal do ABC, Santo Andre, Brazil.) They thought the kinetic argument was too vague, leading to assumptions and speculations, and that an energetic or thermodynamic study might provide a more precise explanation.

PROJECT

Scalable Methods for Electronic Excitations and Optical Responses of Nanostructures: Mathematics to Algorithms to Observables

PRINCIPAL INVESTIGATOR

James Chelikowsky, University of Texas at Austin

SENIOR INVESTIGATOR

Gustavo Dalpian, Universidade Federal do ABC

FUNDING

BES, NSF

COMPUTING RESOURCES

NERSC, TACC

Using electronic structure calculations, Dalpian and Chelikowsky examined the stability of magnetic impurities in spherical cadmium selenide nanocrystals ranging in size from 1.4 to 2.6 nm in diameter—a maximum of 293 cadmium and selenium atoms.[3] At the center of the simulated nanocrystals they placed a single manganese impurity atom (Figure 5).

Their calculations showed that this impurity changes the energy level in the "band gap" between the occupied and unoccupied electron energy bands. As the size of the nanocrystal decreases, it takes more "formation energy" to insert the impurity, making doping more difficult, as shown in experiments.

"Our conclusion," said Chelikowsky, "is that the high energy of formation makes it difficult to stabilize a dopant within a nanocrystal. Since a dopant in a nanocrystal is intrinsically unstable, self-purification may occur for thermodynamic as well as kinetic reasons."

But their results also suggest a thermodynamic solution to the doping problem: if the dopant is a cation, then increasing the anion concentration in the crystal-growing solution would reduce the impurity formation energy and make successful doping more likely.

The significance of this research was highlighted by a feature article in the Nanozone News section of the Nature web site.[4]

WHY DOPING STRENGTHENS GRAIN BOUNDARIES

Ceramic engines have been a topic of research for decades because they are lighter and more fuel-efficient than metal engines. So why aren't we seeing cars and trucks with ceramic engines on the streets yet? The biggest problem with ceramics is durability. Their microscopic structure is granular, and under heat and pressure, the boundaries between the grains can act like tiny seismic faults, allowing small-scale slippage that may grow into cracks and fractures—a dangerous possibility in an engine.

Aluminum oxide or alumina (Al_2O_3) is one of the most promising ceramics for engines because of its hardness—it is widely used for abrasives, like sandpaper, and cutting tools. One drawback is that at high temperatures, alumina is prone to microscopic creep at

---

[3] Gustavo M. Dalpian and James R. Chelikowsky, "Self-purification in semiconductor nanocrystals," Physical Review Letters **96,** 226802 (2006).

[4] Philip Ball, "Why nanotech fails the dope test," Nanozone News, 22 June 2006, http://www.nature.com/materials/nanozone/news/060622/portal/m060622-2.html.

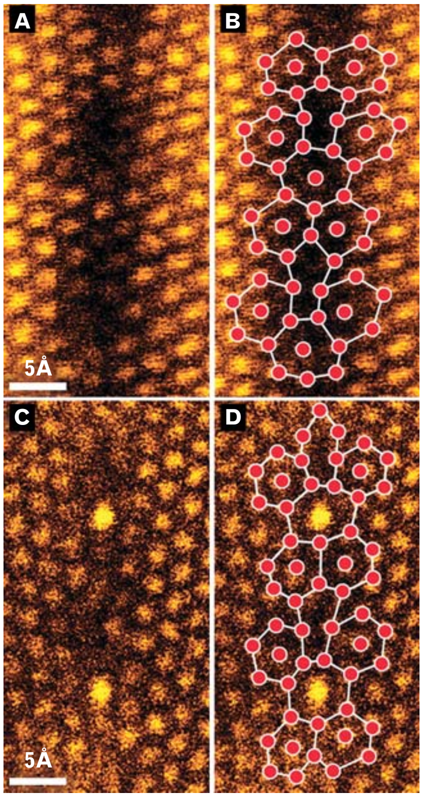

FIGURE 6. STEM images of undoped and yttrium-doped alumina grain boundaries. (A) Undoped alumina; (B) same image with overlay to illustrate the aluminum atomic column arrangement; (C) yttrium-doped alumina; (C) same image with structural overlay.

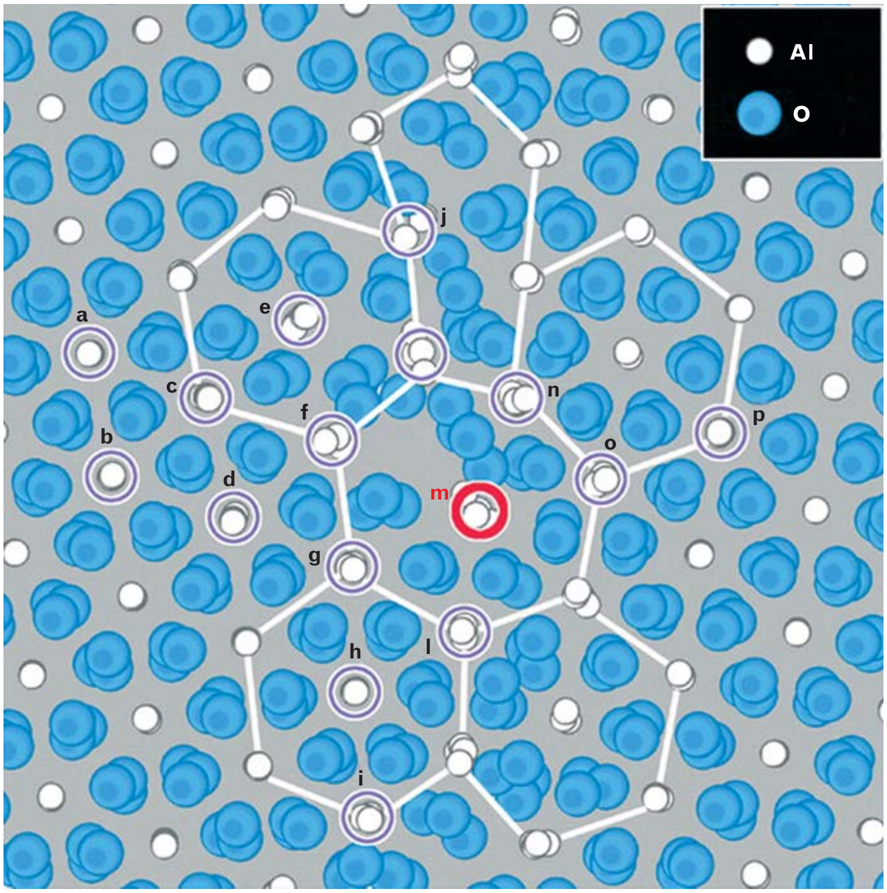

FIGURE 7. Theoretical grain boundary structure obtained by static lattice calculations. Aluminum atoms are white, oxygen blue. Bold lines mark the grain boundary structure as observed in the STEM images. Yttrium segregation energies were investigated for columns a through p, and column m showed the lowest segregation energy.

grain boundaries. However, researchers have found that doping alumina with rare earth elements, such as yttrium (Y), improves its resistance to creep. The dopant has been shown to settle in the grain boundaries, but how it prevents creep at the atomic scale has been controversial.

PROJECT

Theoretical Studies of the Electronic Structures and Properties of Complex Ceramic Crystals and Novel Materials

PRINCIPAL INVESTIGATOR

Wai-Yim Ching, University of Missouri–Kansas City

FUNDING

BES, JSPS

Now a collaboration of researchers from the universities of Tokyo and Missouri–Kansas City may have settled the issue. They examined both undoped and doped grain boundaries with scanning transmission electron microscopy (STEM), then analyzed the grain boundary structure and bonding using a combination of static lattice and first principles calculations.[5]

Figure 6 shows STEM images of undoped (A, B) and yttrium-doped (C, D) alumina grain boundaries. The orange spots correspond to atomic columns of aluminum (the oxygen is not visible), and the yellow spots in C and D are yttrium columns. The schematic overlay in B and D highlights the periodic structural units along the boundary plane, with seven-member rings of Al ions forming a large open structure. These images reveal that Y doping does not alter the basic grain boundary structure; instead, Y simply replaces Al at the center of some seven-member rings in the grain boundary.

The theoretical part of the study was conducted by Wai-Yim Ching, Curators' Professor of Physics at the University of Missouri–Kansas City, along with Japanese researchers and post-doctoral fellow Jun Chen. They first used static lattice calculations to determine the lowest energy structure of the undoped grain boundary. The calculated structure reproduced the experimentally observed seven-member ring structure at the grain boundary (Figure 7).

To locate the most energetically stable site for

---

[5] J. P. Buban, K. Matsunaga, J. Chen, N. Shibata, W. Y. Ching, T. Yamamoto, and Y. Ikuhara, "Grain boundary strengthening in alumina by rare earth impurities," Science **311,** 212 (2006).

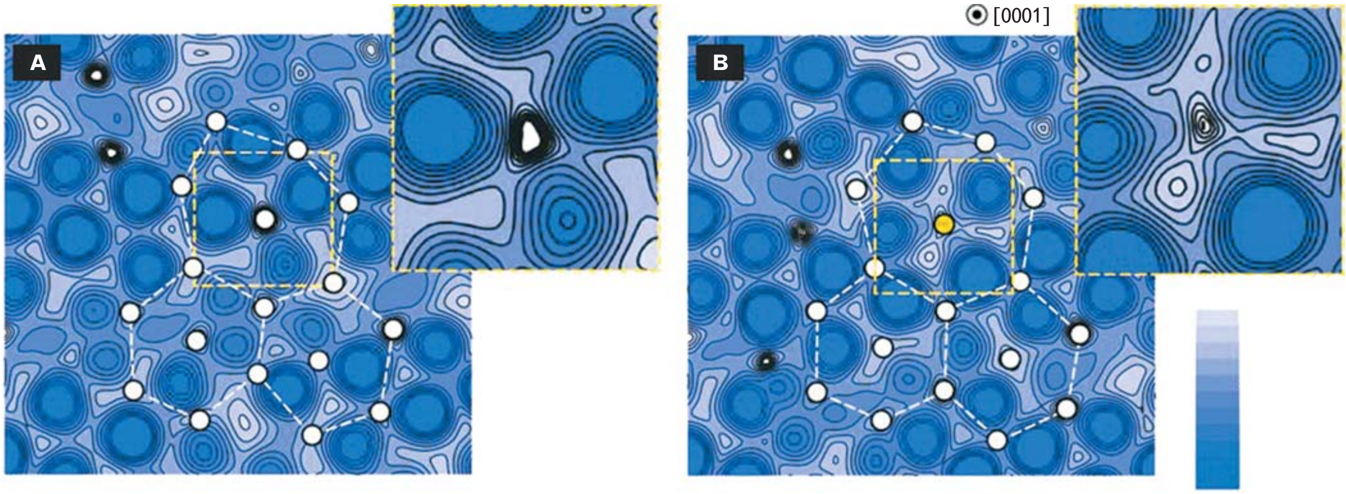

FIGURE 8. Charge density map for undoped (A) and Y-doped (B) grain boundaries.

Y segregation, the theoretical researchers substituted a single Y ion in various columns at or near the grain boundary (Figure 7, *a* through *p*). Site *m*, in the middle of a seven-member ring, had the lowest segregation energy, just as the experiment showed. The main difference between Y and Al ions is their ionic radius: 67.5 picometers for Al, and 104 picometers for Y. The researchers believe the larger area within the seven-member ring can accommodate the larger ion better than a six-member ring.

To investigate local atomic bonding and charge distributions, the researchers then used *ab initio* calculations to construct a large periodic supercell with 700 atoms. Figure 8 shows charge density maps for the undoped (A) and Y-doped (B) grain boundaries. White circles show the location of Al ions; graduated blue spots show the charge density from O ions; and the yellow circle in B is a Y ion. Figure 8A shows sharp nodes between the O charge densities and the charge density from the Al ion in the center of the seven-member ring. In contrast, Figure 8B shows that the O electron densities are elongated toward the Y ion, indicating a stronger covalency (sharing of electrons) between the Y–O bonds. Further calculations showed that more bonds formed between Y and O ions than between Al and O ions; the larger number of bonds in the Y-doped case contributed to lowering the grain boundary energy.

Although the actual mechanism for grain boundary creep is still not well understood, this study advances our understanding of creep resistance. Creep requires the continuous breaking and reforming of atomic bonds as two grains move in opposite directions. Grain boundaries with more bonds and higher bond strength, will, therefore, be more resistant to creep. The undoped seven-member rings in this study have fewer bonds than the interior of the grain, which is why the grain boundaries are mechanical weak points. But the Y-doped rings have more and stronger bonds between Y and O ions, which explains why Y doping increases creep resistance in alumina.

Ching sees the significance of these results in a larger context. "This work demonstrates the importance of combining theoretical computation and experimental observation, the effectiveness of international collaboration, and the need of top-line supercomputers for modern materials research," he said.

A RANDOM WALK ALONG AN INTERFACE

Interfaces are an important class of defects whose distribution affects the properties of the otherwise pristine material, both in nature and in technology. This is especially the case in polycrystals, thin films, multiphase materials, and composites, where the mechanical, chemical, and transport properties are sensitive to the underlying interfacial microstructure.

"In fact, tailoring this microstructure is an emerging paradigm for engineering high performance, multifunctional materials," said Zachary Trautt, a graduate research assistant and the first author of the study "Interface Mobility from Interface Random Walk," which appeared in the October 27, 2006 issue of Science.[6] In that paper, researchers at the Colorado School of Mines and Northeastern University reported a novel computational methodology aimed at quantifying the kinetics of interfaces in diverse material systems.

The interfacial microstructure is subject to several driving forces during material synthesis and function. More often than not, these driving forces are large enough to cause the interfaces to move and the microstructure (or its precursor) to evolve (Figure 9). Naturally, controlling the final microstructure requires accurate models (Figure 10) that relate the interface motion to the driving forces in effect.

A quantitative measure of interface kinetics is the *interface mobility*, the ratio of the interface velocity to the driving force. Past studies on individual homophase crystalline

---

[6] Zachary T. Trautt, Moneesh Upmanyu, and Alain Karma, "Interface mobility from interface random walk," Science **314,** 632 (2006).

PROJECT

Microstructural Evolution Based on Fundamental Interfacial Properties

PRINCIPAL INVESTIGATOR

Anthony Rollett, Carnegie Mellon University

SENIOR INVESTIGATORS

James Morris, Oak Ridge National Laboratory; Mark Asta, University of California, Davis; Alain Karma, Northeastern University

FUNDING

BES, ONR, ATC

COMPUTING RESOURCES

NERSC, NOOMSRC

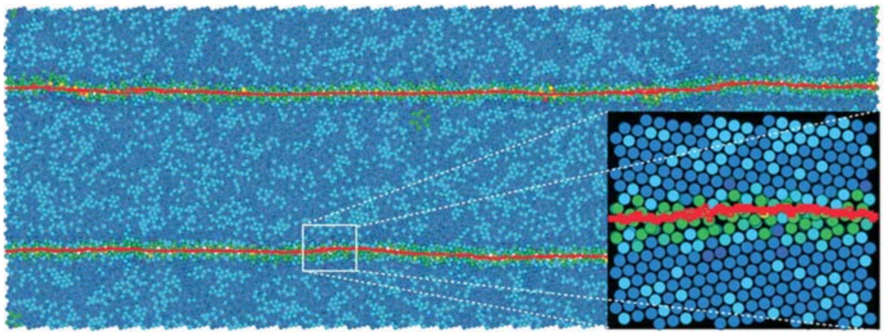

FIGURE 9. A slice through a computational cell consisting of two grain boundaries which separate two perfect crystals of different orientations. The average position of the grain boundaries $\bar{h}$ is calculated from their fluctuating profile. The color reflects the interaction energy—green indicates high-energy states. The superimposed red dotted curve is the result of the algorithm used to identify the interface position. (Images courtesy of Science)

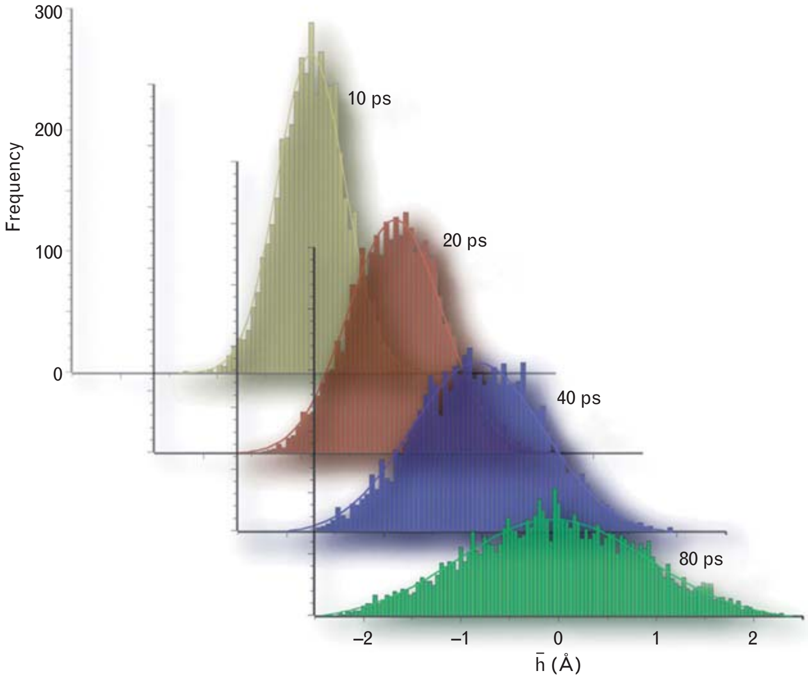

FIGURE 10. The time evolution in picoseconds of the distribution of the average grain boundary position $\bar{h}$. The distributions are normal, as predicted by the theory.

interfaces (or grain boundaries) in several high-purity metals show an interesting trend: the experimental mobilities are orders of magnitude smaller than those extracted via computations. The discrepancy is often attributed to the presence of impurities, fueling speculation that even minute quantities of impurities significantly retard interface motion.

“An often overlooked fact is that computations are limited to tens of nanoseconds,” said Moneesh Upmanyu, co-author and the lead researcher in the study. “As a result, they are performed at driving forces orders of magnitude greater than those commonly observed in experiments,” he explained. This further weakens the comparison, and there is a need to extend the computational studies to more realistic driving forces and include the effect of impurities.

“Our computational methodology offers a way to address both these challenges, efficiently and with setups that are relatively simple,” said Trautt. The basis for the methodology is the pioneering theoretical work by Einstein, Smulochowski, and Langevin on Brownian motion in the early 1900s. “Just as their study related the dance of macroscopic particles to their diffusivity, the microscopic thermal fluctuations result in interface forces that conspire towards a one-dimensional dance [random walk] of the average interface position, which in turn yields its mobility in the zero driving force limit,” said Alain Karma, also a co-author in the study. “The technique is remarkably efficient,” noted Upmanyu. “The computations on pure aluminum yielded mobilities within a nanosecond, a significant savings in computational resources.”

Comparisons with previous experiments and computations reveal that the retarding effect of impurities is much more severe than previously thought. The authors are now working on extending the theory and the computations to directly quantify the impurity drag effect.

Trautt is a graduate research assistant in the Group for Simulation and Theory of Atomic-scale Material Phenomena (stAMP), Engineering Division, Colorado School of Mines; Upmanyu is group leader of stAMP and Assistant Professor in the Engineering Division and Materials Science Program. Karma is Distinguished Professor in the Department of Physics and Center for Interdisciplinary Research on Complex Systems, Northeastern University, and a senior investigator in the NERSC project "Microstructural Evolution Based on Fundamental Interfacial Properties," led by Anthony Rollett, Department Head and Professor of Materials Science and Engineering at Carnegie Mellon University.

BUILDING A CHEMICAL TOOL BOX

To coax the next blockbuster drug out of chemical compounds, scientists must explore how these compounds work with each other. This process of assembling the right ingredients takes years and millions of dollars, and the result can lead to life-saving medicines.

How do chemists figure out the right mix? It takes ample knowledge about the characteristics of various compounds and a dose of serendipitous discovery. This process of synthesizing compounds is what has driven an INCITE project investigator to examine how lithium enolates, lithium carbenoids and the blend of the two interact with other lithium compounds. "Pharmaceutical chemists always need to synthesize new compounds," said Larry Pratt, principal investigator of the project and assistant professor at Fisk University in Nashville, Tennessee. "I am helping to build the tool box."

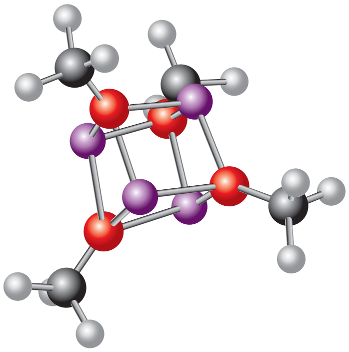

PROJECT

Reactions of Lithium Carbenoids, Lithium Enolates, and Mixed Aggregates

PRINCIPAL INVESTIGATOR

Lawrence Pratt, Fisk University

FUNDING

BES, NSF

Lithium, the lightest of the solid elements, is a soft white metal that oxidizes quickly in air and water. It is commonly found in portable batteries that power laptops and cell phones. In pharmacology, lithium compounds such as lithium carbonate and lithium citrate are commonly used to control mood swings.

The lithium compounds that have captivated Pratt's interest are lithium enolates and lithium carbenoids. They are two important classes of reagents in organic chemistry, which involves the study of the properties, synthesis and reactions of carbon-containing compounds. Pratt has drawn from experiments conducted by him and other researchers to calculate the structures, energy, and other properties that resulted from mixing these lithium compounds.

In cooking, when you throw in a bunch of spices, you can expect the mixture to give off an altogether different smell and taste than what you get from individual ingredients. The same thing happens when you blend different compounds into an *aggregate*—a cluster of molecules that chemically reacts as a single molecule. A mixed aggregate may have properties different than those of either pure component, and the aggregate may react very differently from either of its components. For the creative chemist, the result may not be what was desired. But the more that is known about the structure and reactivity of the aggregate, the better the chance of success.

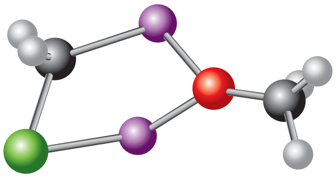

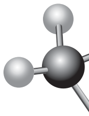

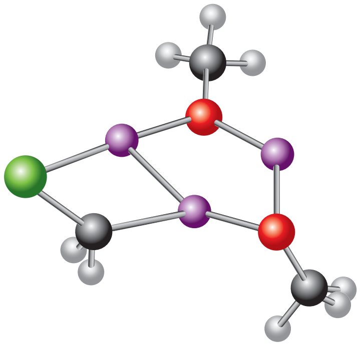

In solution, almost all lithium compounds form aggregates. Typically, the larger the molecule, the less likely it will form a complex bond. For example, the lithium enolate of acetaldehyde may exist as a four- or six-molecule compound (tetramer or hexamer) in tetrahydrofuran, a common ether-type solvent, while lithium cyclohexanone enolate may show up as a single- or dual-molecule compound (monomer or dimer) in the same solvent. As mixed aggregates, lithium enolates and lithium carbenoids can exhibit different characteristics, depending on the molecular structure, temperature, and the type of solvent.

Using mathematical models to determine the structures and potential reactions of lithium aggregates is a good way to advance the understanding of these substances, particularly because observing their step-by-step reaction in an experiment can be

FIGURE 11. Optimized geometries of lithium carbenoid mixed aggregates with lithium methoxide. Gray represents carbon; white, hydrogen; violet, lithium; green, chlorine; and red, oxygen.

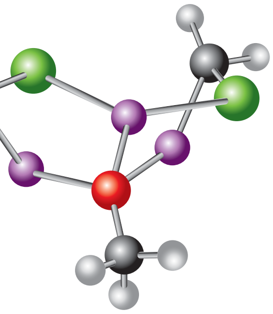

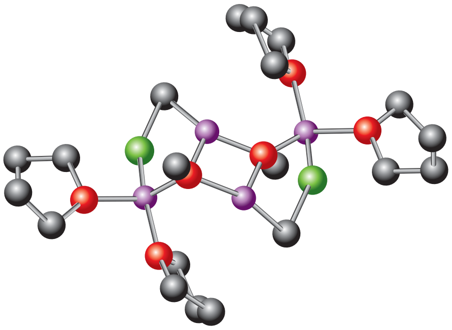

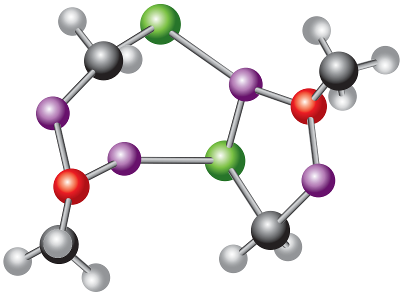

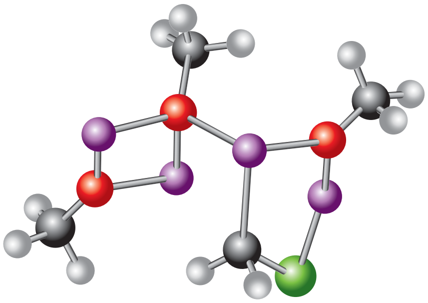

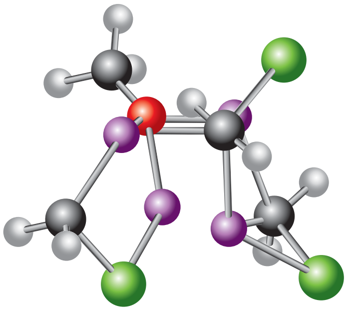

difficult. Many lithium carbenoids are reactive and unstable; they tend to decompose at room temperature, although some can be observed at low temperatures.

Pratt uses *ab initio* and density functional theory methods to investigate the structures and interactions between the compounds. After modeling all reasonable geometric structures of the compounds and aggregates (see Figure 11 for an example), he uses the results to figure out which structures are likely to exist in significant concentrations in the gas phase and in solution. Then he examines the reactions via each potential reactive species, such as monomer, dimer, tetramer, or mixed aggregate. Transition structures are then located and the activation energies calculated. Characterizing these reactions will help other researchers figure out new uses for these compounds.

This article written by: John Hules, Lynn Yarris, and Ucilia Wang, Berkeley Lab; and Laura Shea, Northeastern University.

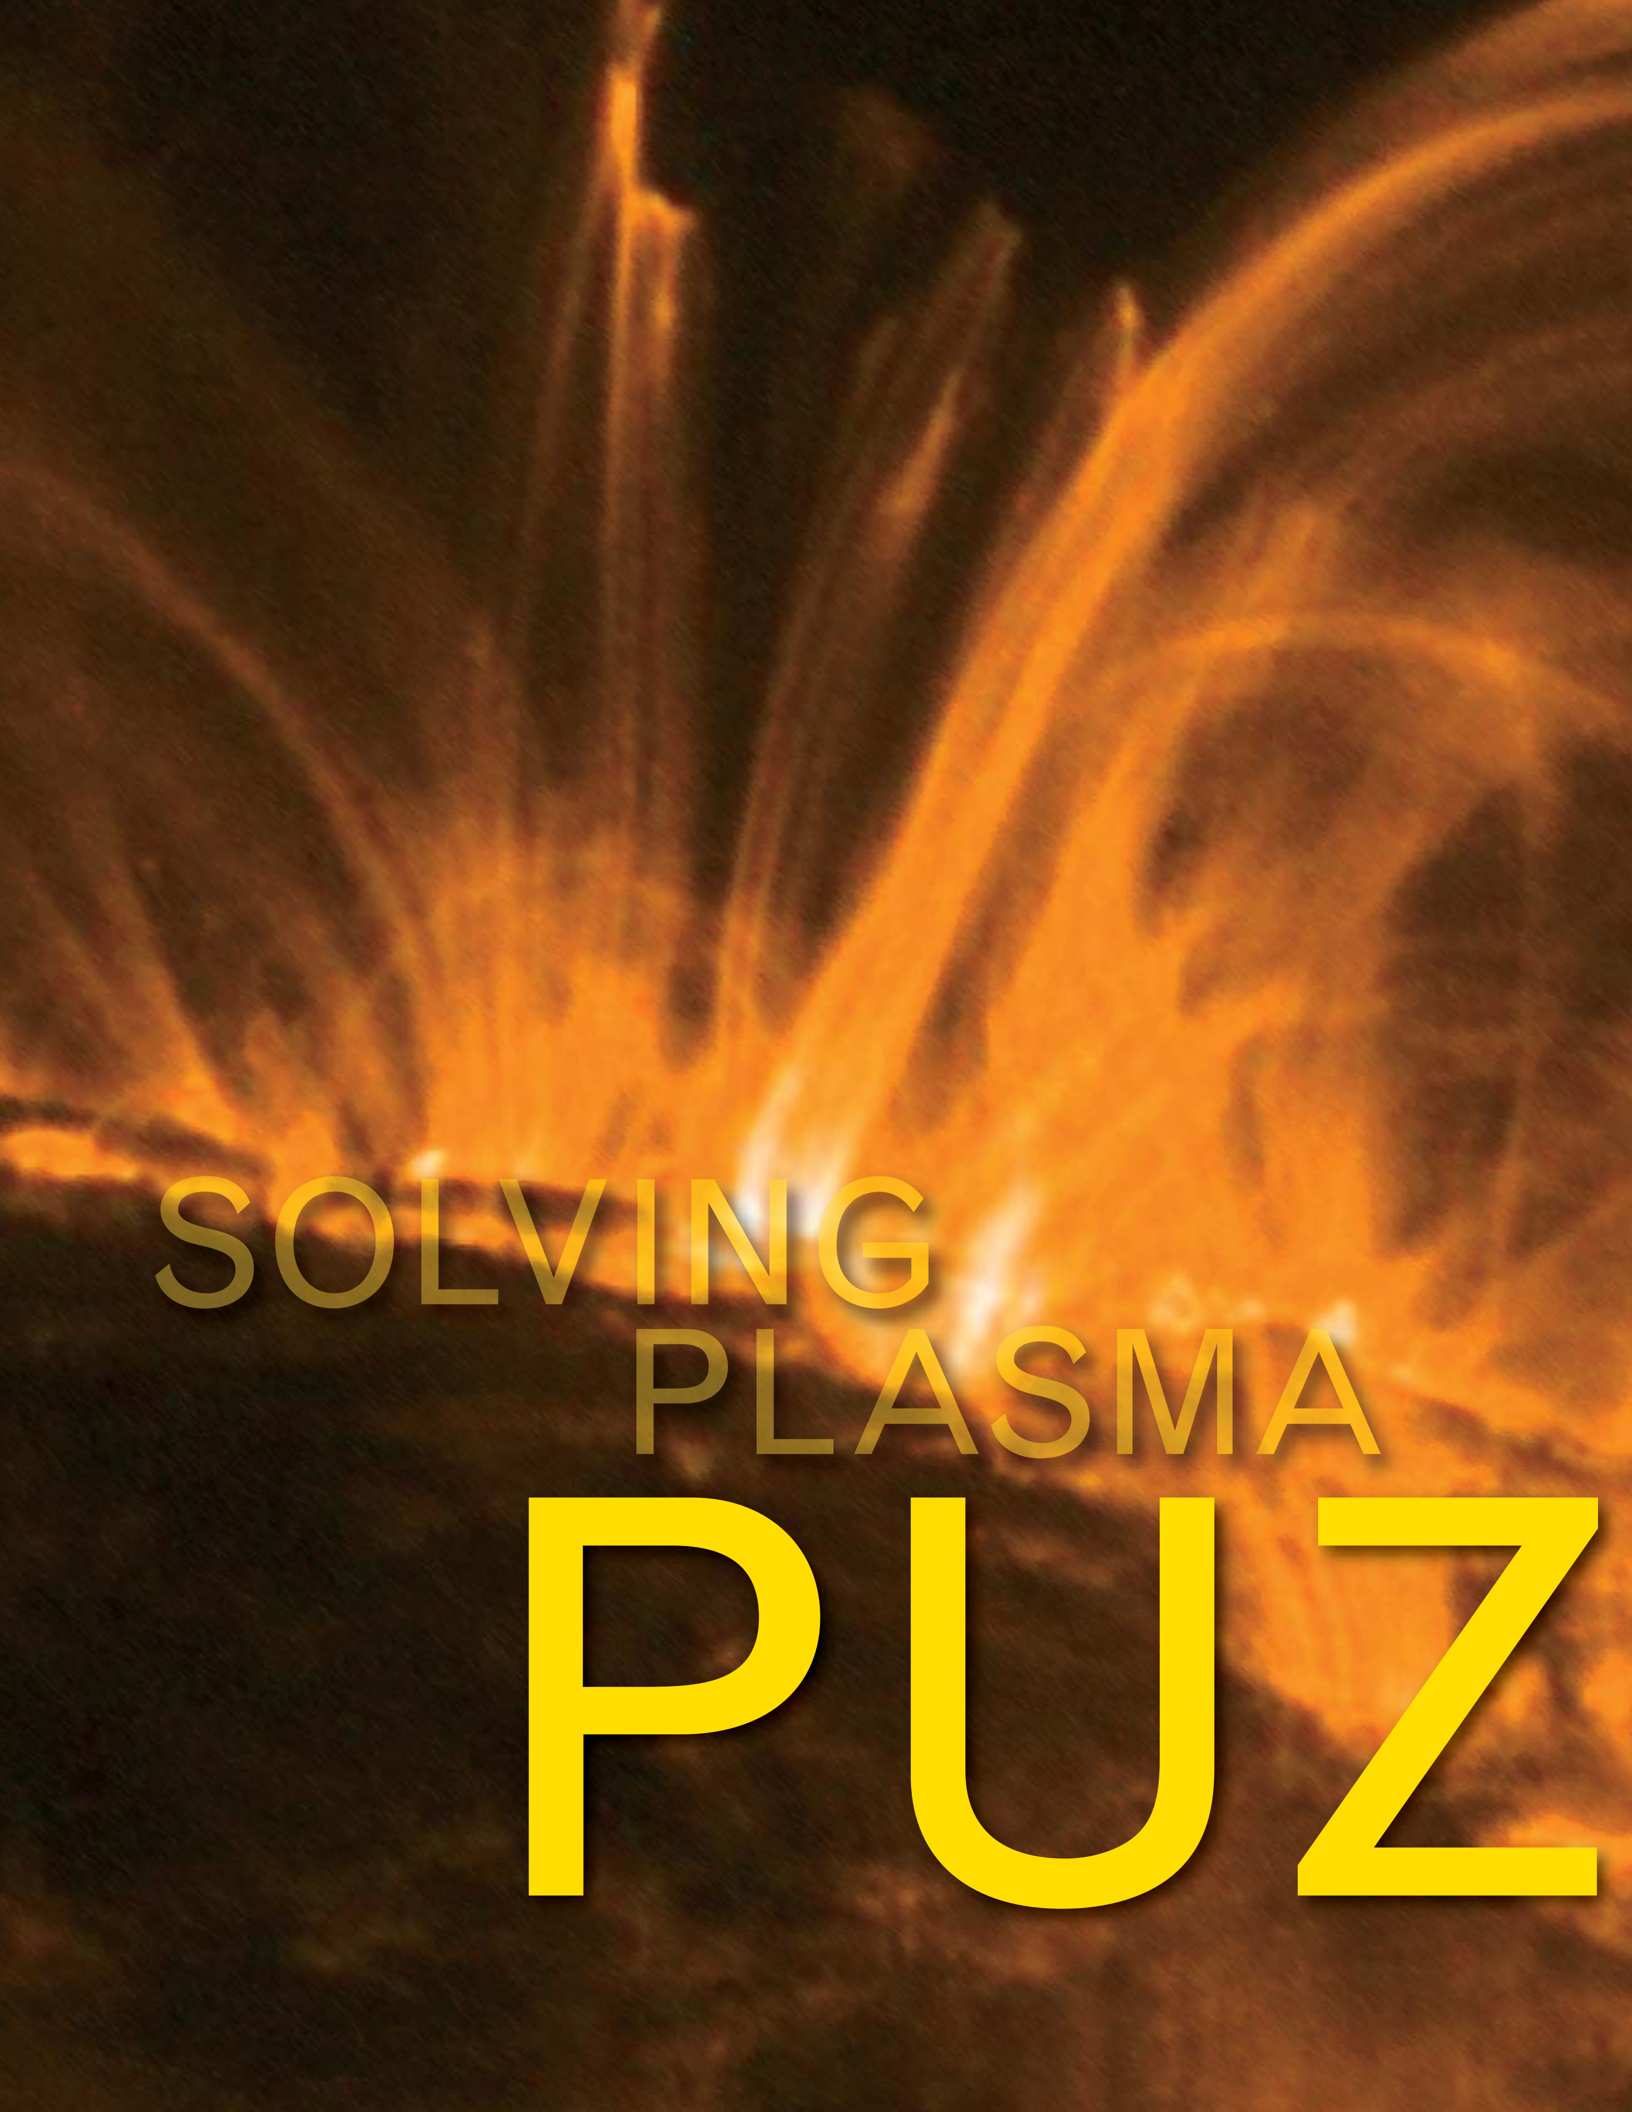

SOLVING PLASMA PUZ

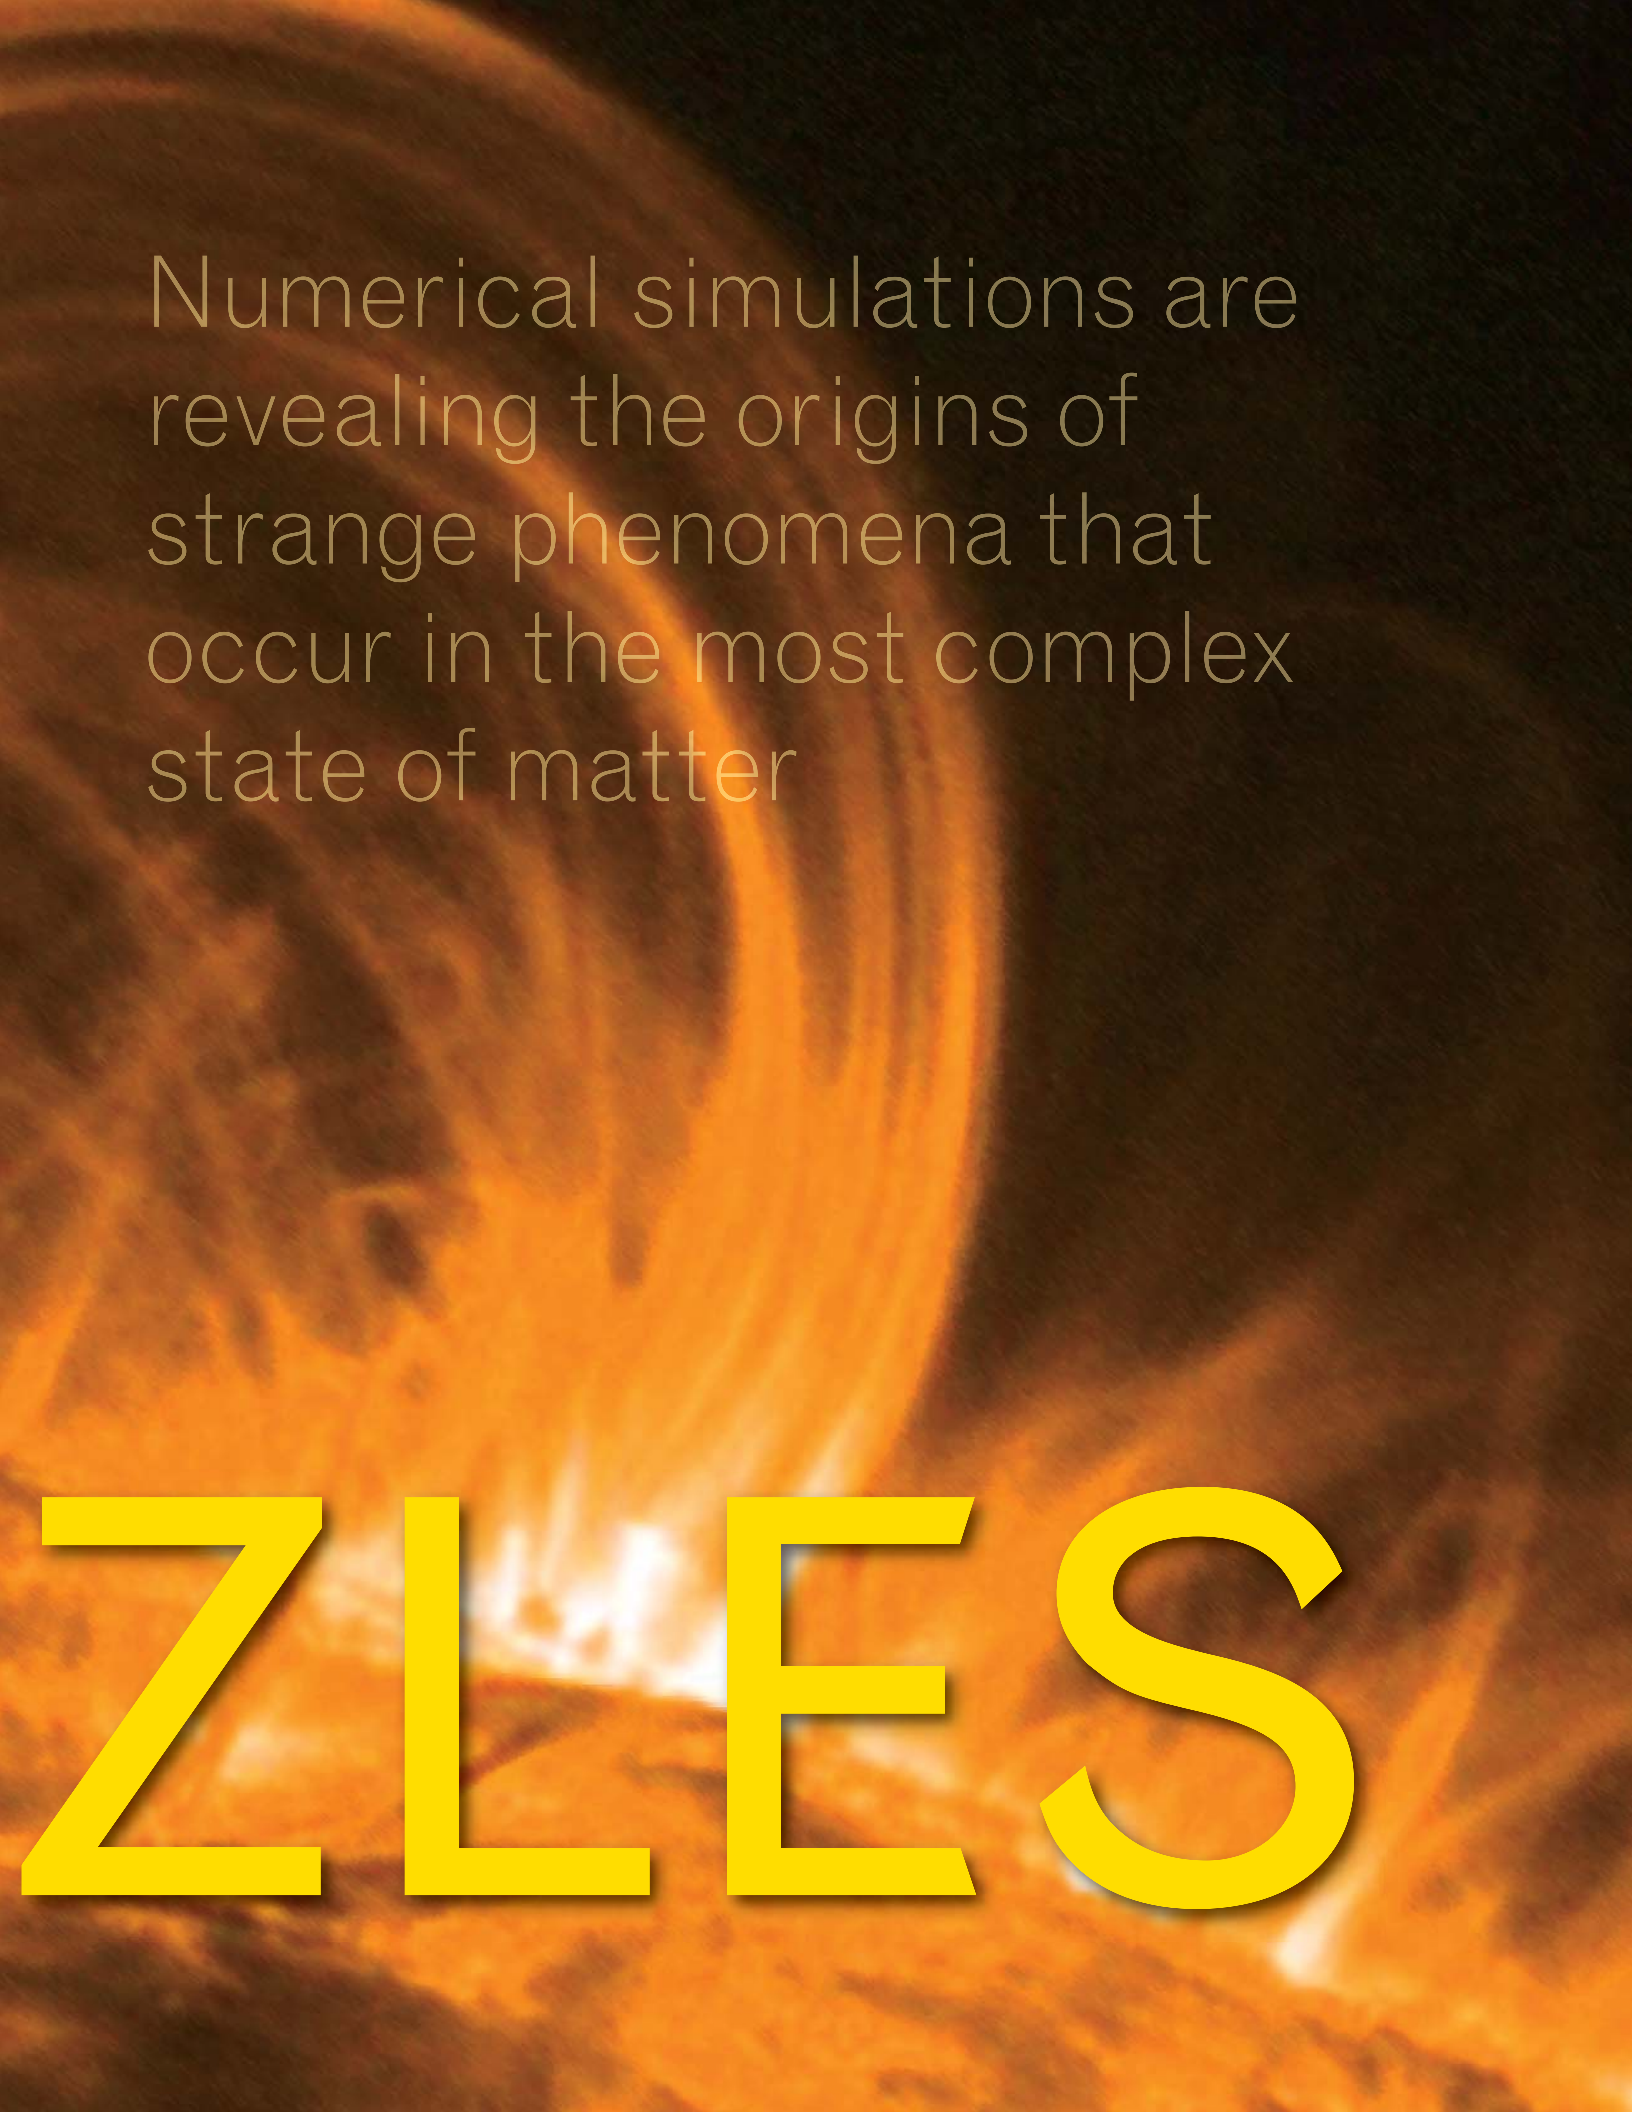

ZLES

Numerical simulations are revealing the origins of strange phenomena that occur in the most complex state of matter

The light and warmth of the sun, the romance and mystery of the stars, the drama and power of lightning, the breathtaking beauty of auroras—over the course of human history, these wonders have inspired a variety of thoughts and feelings, stories and art works. From the viewpoint of physics, all of the these phenomena are plasmas, but that fact does not detract from the mystery; for plasma (ionized gas), with its ability to carry electrical currents and generate magnetic fields, is perhaps the most complex and difficult to understand state of matter.

Plasmas are the most common form of baryonic matter, which is currently believed to constitute about 4% of the universe (the rest being dark matter and dark energy). Plasmas surround the Earth (ionosphere and magnetosphere), permeate the solar system, and pervade interstellar and intergalactic space.

With a growing understanding of plasmas over the past century, researchers have found dozens of practical applications for them. The most familiar applications are fluorescent lights, neon signs, and plasma display screens for televisions and computers; the most promising application is nuclear fusion as a renewable energy source. Recent discoveries from numerical simulations are advancing both the basic science of plasmas and fusion energy experiments.

UNDERSTANDING MAGNETIC EXPLOSIONS

On December 6, 2006, Global Positioning System (GPS) devices suddenly started malfunctioning all over the Earth. The culprit: a solar flare. Solar flares can eject a billion tons of charged particles into space at a speed of 1 million km per hour, disrupting navigation and communications satellites, and sometimes even electrical grids on Earth, while producing bright auroras in the polar regions.

How so much energy can be released so quickly has perplexed scientists for decades. In 1946 Ronald Giovanelli conceived the idea of *magnetic reconnection* to explain solar flares. The basic idea is that the churning of ionized gas amplifies the magnetic fields in a plasma by twisting and folding them—kinetic energy being converted into magnetic energy. When the magnetic field lines touch or cross, they break, reconnect, and reverse direction (Figure 1). The process may take months or, in the case of a solar flare, as little as 30 minutes, in which case vast amounts of magnetic energy are converted back to kinetic energy with explosive force.

"Magnetic reconnection differs from a conventional explosion in that the energy is not released equally in all directions," explained James F. Drake, Professor of Physics at the University of Maryland, whose recent research has focused on this subject. "Instead, the plasma flows in from one direction and flows out in another.

"Magnetic reconnection has broad importance for almost all areas of plasma physics, including solar flares, storms in the Earth's magnetosphere, and disruptions in laboratory fusion experiments," Drake added. "It's a fascinating topic and a challenging research area."

One of the puzzles in this sudden release of massive energy is how that much energy could have built up in the first place. If reconnection were always fast and occurred frequently, the magnetic fields would never be strong enough to reach explosive force. A long period of slow reconnections might allow the magnetic energy to accumulate. But why would reconnection happen at two different speeds?

The first mathematical model for magnetic reconnection, known as the Sweet-Parker model, was developed in the late 1950s. This model generates a slow and steady release of energy, but not the explosive events that the theory is supposed to explain. In this model the electrons and ions move together, and the heavier ions slow down the plasma flow. The more recent Hall reconnection model suggests that the movements of ions become

PROJECT

Turbulence, Transport and Magnetic Reconnection in High Temperature Plasma

PRINCIPAL INVESTIGATOR

William Dorland, University of Maryland

SENIOR INVESTIGATORS

James Drake, Parvez Guzdar, Adil Hassam, and Robert Kleva, University of Maryland

FUNDING

FES, NSF, NASA, CMPD, CISM

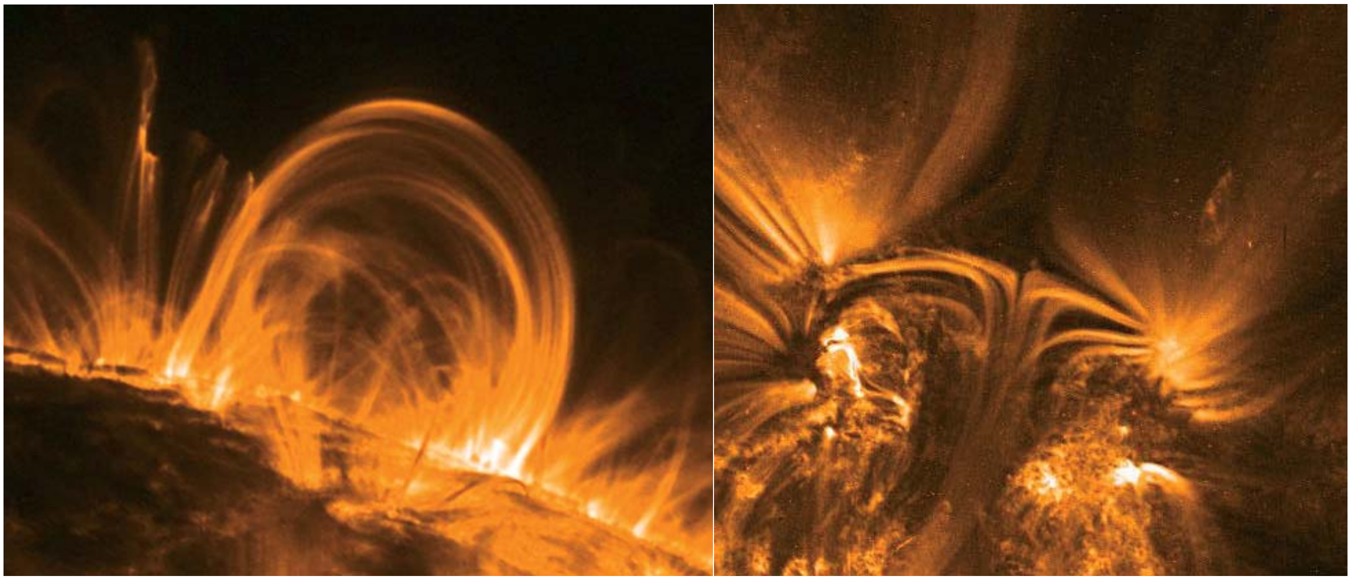

FIGURE 1. (A) Glowing loops of plasma illuminate the magnetic field structure around a sunspot. The planet Earth would easily fit under one of these loops. (B) Constantly in motion, the field lines sometimes touch or cross and reverse direction in a process called *magnetic reconnection.* Open field lines instead of loops show that plasma is being ejected outward as a solar flare. (Images courtesy of NASA)

decoupled from electrons and magnetic fields in the boundary layers where magnetic field lines reconnect. The result is much faster plasma flow. The signatures of the Hall model have been confirmed by satellite measurements in the magnetosphere and by laboratory experiments, but this model still does not explain the origin of the magnetic explosion.

In two recent papers by Paul Cassak, Drake, and Michael Shay, the two models have converged in a self-consistent model for the spontaneous onset of fast reconnection.[1] The researchers' calculations showed that slow reconnection can continue for a long time, during which magnetic stresses continue to build up. As progressively stronger magnetic fields are drawn into the reconnection region, when the available free energy crosses a critical threshold, the system abruptly transitions to fast reconnection, manifested as a magnetic explosion.

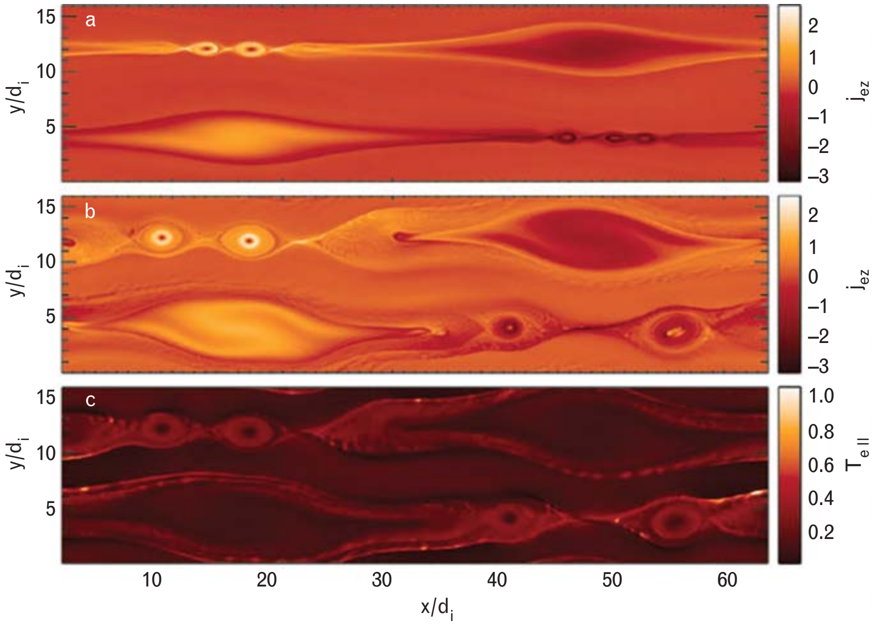

FIGURE 2. Computer simulations of island formation and electron acceleration during magnetic reconnection. The electron current is shown at two time steps in (a) and (b); (c) shows the electron temperature, with intense heating caused by electron acceleration along the rims of the islands.

This new model is consistent with solar flare observations. For example, extreme-ultraviolet observations of the sun's corona have shown one instance of slow reconnection lasting for 24 hours, followed by fast reconnection lasting for 3 hours. The change was sudden, with no visible trigger mechanism, and the energy released during fast reconnection was comparable to the energy accumulated during slow reconnection.

Solar observations have suggested that at least 50% of the energy released during flares is in the form of energetic electrons, and energetic electrons have also been measured during disruptions in laboratory nuclear fusion experiments. The source of these

---

[1] P. A. Cassak, M. A. Shay, and J. F. Drake, "Catastrophe model for fast magnetic reconnection onset," Physical Review Letters **95,** 235002 (2005); P. A. Cassak, J. F. Drake, and M. A. Shay, "A model for spontaneous onset of fast magnetic reconnection," Astrophysical Journal **644,** L145 (2006).

energetic electrons has been a puzzle. Large numbers of these low-mass particles travel at speeds far higher than can be explained by the energy of the moving magnetic field lines that propel them. Drake and Shay, along with Michael Swisdak and Haihong Che, proposed an answer to this question in a paper published in Nature.[2]

In their simulations of magnetic reconnection, the process is more turbulent than it was once thought to be—magnetic islands form, grow, contract, and merge as the field lines converge (Figure 2). The electrons gain speed by reflecting off the ends of contracting islands, just as a ball would gain speed if it were bouncing off two walls that were moving toward one another. But as the temperature in an island goes up, back pressure slows down the shrinking, thus slowing down reconnection and converting more of the magnetic energy into electron acceleration. The repeated interactions of electrons with many islands allow them to be accelerated to high speeds.

"Ours is the first mechanism that explains why electrons gain so much energy during magnetic reconnection," said Drake. "From a practical standpoint, these new findings can help scientists to better predict which solar storms pose the greatest threat to communications and other satellites. And they may give us a better understanding of how to control plasmas in fusion reactors."

Drake explained that the strongest confirming evidence for the new theory was the surprising agreement between the model and data from NASA's WIND satellite. "We were as surprised as the WIND scientists when the distribution of energetic electrons seen by their spacecraft popped right out of our model. Such a match isn't something you see very often," he said.

Drake computes at NERSC under the project "Turbulence, Transport and Magnetic Reconnection in High Temperature Plasma," led by William Dorland. In addition to magnetic reconnection, this project also studies the mechanisms by which plasma particles, energy, and momentum are transported across, rather than along, magnetic field lines—the so-called "anomalous transport" problem.

MODELING MICROTURBULENCE IN FUSION PLASMAS

Most people do not think about turbulence very often, except when they are flying and the captain turns on the "Fasten Seat Belts" sign. The kind of turbulence that may cause problems for airplane passengers involves swirls and eddies that are a great deal larger than the aircraft. But in fusion plasmas, much smaller-scale turbulence, called *microturbulence,* can cause serious problems—specifically, instabilities and heat loss that could stop the fusion reaction.

In fusion research, all of the conditions necessary to keep a plasma dense and hot long enough to undergo fusion are referred to as *confinement.* The retention of heat, called *energy confinement,* can be threatened by microturbulence, which can make particles drift across, rather than along with, the plasma flow. At the core of a fusion reactor such as a tokamak, the temperatures and densities are higher than at the outside edges. As with weather, when there are two regions with different temperatures and densities, the area between is subject to turbulence. In a tokamak, turbulence can allow charged particles in the plasma to move toward the outer edges of the reactor rather than fusing with other particles in the core. If enough particles drift away, the plasma loses temperature and the fusion reaction cannot be sustained.

The growth of the microinstabilities that lead to turbulent transport has been extensively studied over the years. Understanding this process is an important practical problem, and it is also a true scientific grand challenge which is particularly well suited to be addressed by modern terascale computational resources. One of the leading research groups exploring this issue is the SciDAC-funded Center for Gyrokinetic Particle Simulations of Turbulent Transport in Burning Plasmas and Multiscale Gyrokinetics (GPSC), headed by Wei-li Lee of the Princeton Plasma Physics Laboratory. This team developed the Gyrokinetic Toroidal Code (GTC) to simulate instabilities in tokamak plasmas using the particle-in-cell (PIC) method.

"Particle-in-cell simulation, which began in the late sixties, uses finite size particles on a grid to dramatically reduce the numerical noise associated with close encounters between the particles, while leaving intact their long range interactions outside the grid," Lee explained. "This approximation reduced the number of calculations for particle interactions and greatly reduced the computational time.

"For simulations of magnetic fusion plasmas," he continued, "further improvements came in the eighties and nineties with the development of the gyrokinetic particle simulation and perturbative particle simulation methods. Briefly, under the gyrokinetic approximation, the spiral motion of a charged particle is represented as a charged ring centered around its gyro-center; and perturbative methods are used to greatly reduce the discrete particle noise."

Over the past dozen years, these simulation methods have produced some impressive discoveries, including the identification of ion temperature gradient (ITG) drift turbulence as the most plausible process responsible for the thermal transport observed in tokamak experiments; the reduction of such transport by self-generated zonal flows; and the confinement scaling trends associated with the size of the plasma and also with the ionic isotope species.[3]

PROJECT

Center for Gyrokinetic Particle Simulations of Turbulent Transport in Burning Plasmas and Multiscale Gyrokinetics

PRINCIPAL INVESTIGATOR

Wei-li Lee, Princeton Plasma Physics Laboratory

SENIOR INVESTIGATORS

Scott Parker, University of Colorado; Zhihong Lin, University of California, Irvine; Viktor Decyk, University of California, Los Angeles

FUNDING

FES, SciDAC

COMPUTING RESOURCES

NERSC, LCF

---

[2] J. F. Drake, M. Swisdak, H. Che, and M. A. Shay, "Electron acceleration from contracting magnetic islands during reconnection," Nature **443,** 553 (2006).

[3] W. W. Lee, S. Ethier, W. X. Wang, and W. M. Tang, "Gyrokinetic particle simulation of fusion plasmas: Path to petascale computing," Journal of Physics: Conference Series (SciDAC 2006) **46,** 73 (2006).

With the availability of terascale computers in recent years, the GPSC team has been able to carry out simulations of experiment-sized plasmas with improved physics fidelity. Typical global PIC simulations of this type have used one billion particles with 125 million grid points over 7000 time steps to produce significant physics results. Simulations of this size would not be feasible on smaller computers.

With nearly two orders of magnitude increase in particle numbers, the GPSC project has been able to resolve longstanding uncertainty about the effect of discrete particle noise on the long-term transport predictions of turbulent gyrokinetic PIC simulations. The "noise" referred to here involves not just particle interactions that are not relevant to energy transport, but primarily numerical sampling noise, because PIC simulations involve Monte-Carlo sampling of a collection of "marker" particles. Recent work shows that this numerical noise has little effect on the resulting energy transport when a reasonable number of particles is used (Figure 1).[3,4]

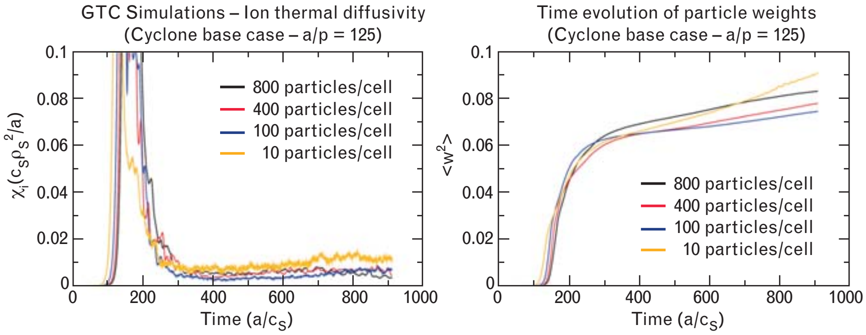

FIGURE 1. Particle number convergence studies for the ITG simulation: thermal diffusivity and the time rate of change of entropy for 10, 100, 400, and 800 particles per cell. The only numerical noise comes from the 10 particle run.

When the GTC code was applied to a geometry similar to the ITER experiment, an interesting new phenomenon was discovered: the turbulence spreads radially from a localized region to eventually cover most of the poloidal plane (Figure 2).[3,5] This discovery was made possible by a simulation volume that is large enough to allow a clear scale separation between the turbulence eddy size and the device size.

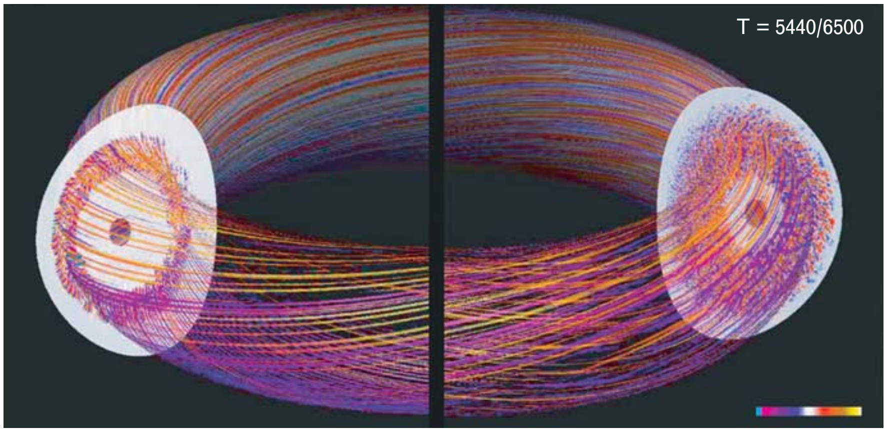

FIGURE 2. Turbulence spreading (left to right) as depicted by the perturbed potentials of ITG turbulence on the poloidal plane as they follow the magnetic field lines around the torus.

"The simulation clearly shows that small-scale turbulence eddies are typically generated in the unstable region and flow along the streamers to the stable region," Lee said. "In addition, the streamers are found to break and reconnect, resulting in a very complex dynamical evolution. These new results have raised intense interest in the fusion theory community on the fundamental physics of turbulence spreading."

Clearly, there is a lot more work to be done in modeling tokamak plasmas, and with petascale computers coming online and the addition of more detailed physics to the GTC code, the GPSC team is eager to continue. With trillion particle simulations, they hope to find detailed solutions to problems such as electron thermal transport, the scaling of confinement with plasma size, and the effects of different ionic isotope species such as tritium on plasma burning.

"Our long-range goal is to carry out integrated simulations for ITER plasmas for a wide range of temporal and spatial scales, including high-frequency short-wavelength wave heating, low-frequency meso-scale transport, and low-frequency large-scale magnetohydrodynamic physics," Lee said.

The success of these efforts will depend on close collaboration with other SciDAC centers, including the Terascale Optimal PDE Simulations (TOPS) Center, the Scientific Data Management (SDM) Center, the Ultrascale Visualization Center, and the Visualization and Analytics Center for Enabling Technologies (VACET).

This article written by: John Hules, Berkeley Lab.

---

[4] Thomas G. Jenkins and W. W. Lee, "Fluctuations and discrete particle noise in gyrokinetic simulation of drift waves," Physics of Plasmas **14,** 032307 (2007).

[5] W. X. Wang, Z. Lin, W. M. Tang, W. W. Lee, S. Ethier, J. L. V. Lewandowski, G. Rewoldt, T. S. Hahm, and J. Manickam, "Gyro-kinetic simulation of global turbulent transport properties in tokamak experiments," Physics of Plasmas **13,** 092505 (2006).

SURFING the PLASMA W

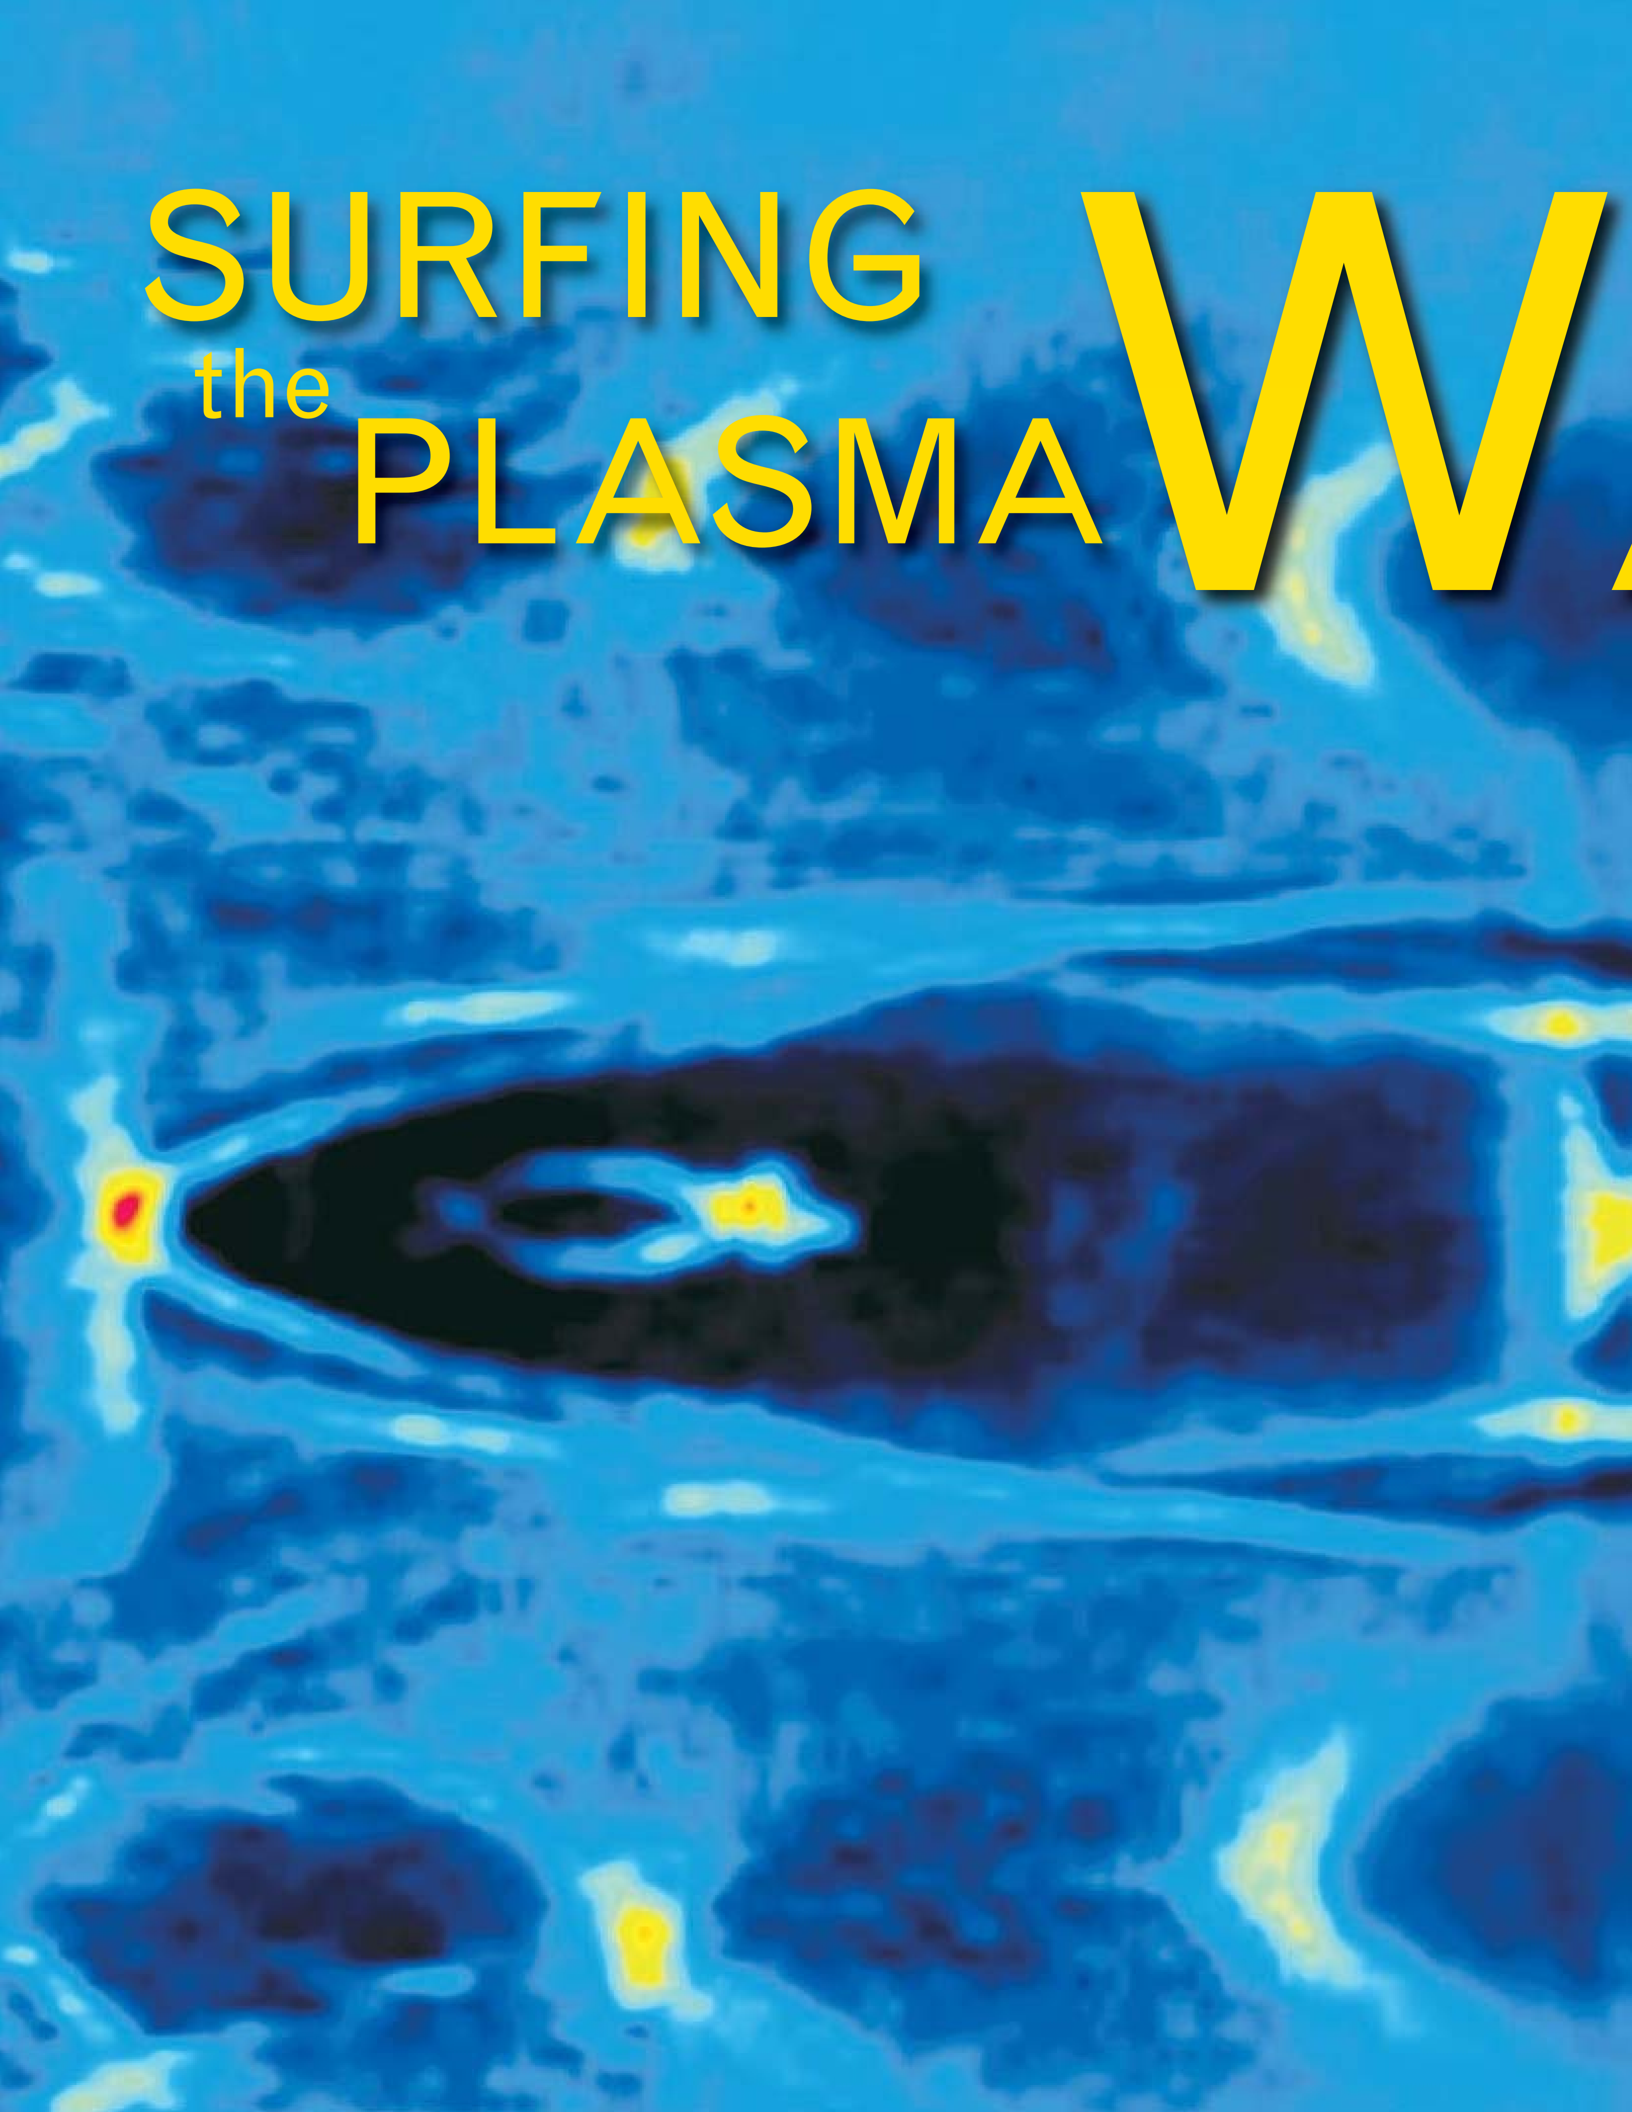

AVES

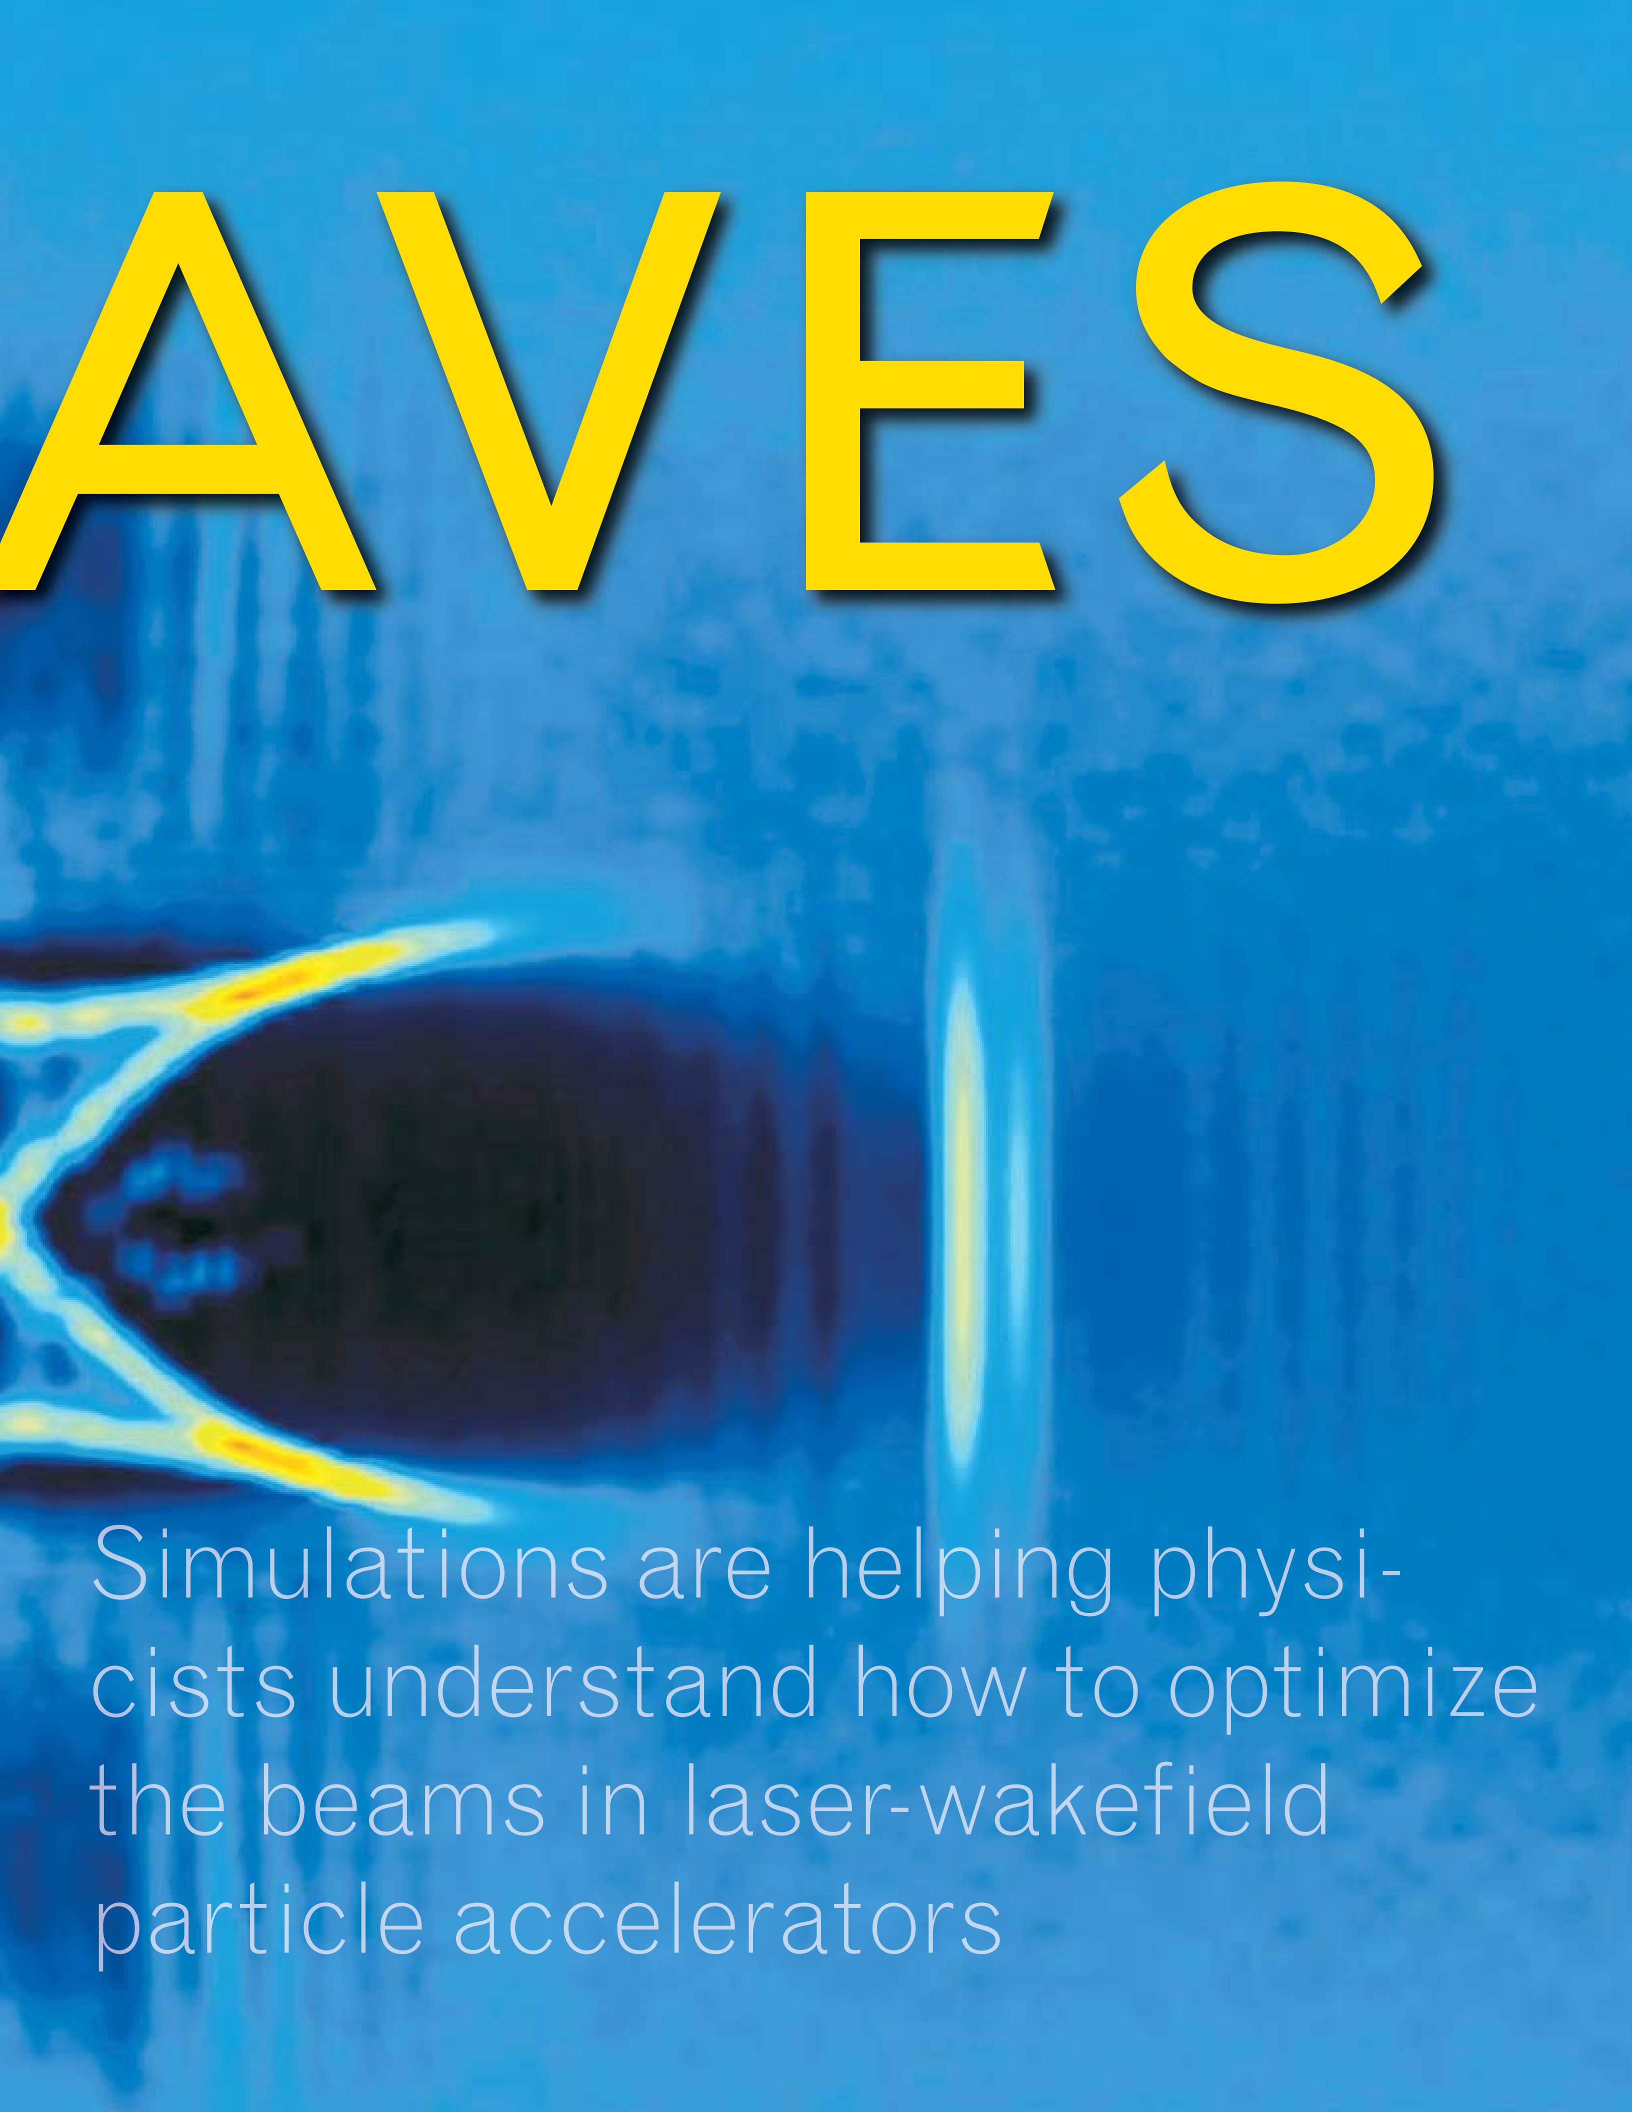

Simulations are helping physicists understand how to optimize the beams in laser-wakefield particle accelerators

The science and technology of laser-driven wakefield particle accelerators took two leaps forward in 2006. Scientists at Lawrence Berkeley National Laboratory, working with colleagues at the University of Oxford, accelerated electron beams to energies exceeding a billion electron volts (GeV) in a distance of just 3.3 centimeters. And the Berkeley team used an INCITE award of 2.5 million processor hours at NERSC to create full-scale, three-dimensional, explicit particle models that revealed important physical details of the experiments.

The experimental breakthrough, reported in the October 2006 issue of Nature Physics,[1] was spearheaded by the LOASIS (Laser Optical Accelerator Systems Integrated Studies) program at Berkeley Lab, headed by Wim Leemans. Defying the prediction that petawatt-class lasers would be needed to reach GeV energies, Leemans and his collaborators channeled a 40 terawatt peak-power laser pulse in a gas-filled capillary waveguide to produce a high-quality electron beam with 1 GeV energy.

The simulation project, led by Cameron Geddes of the LOASIS team, created the first high-resolution 3D models of these laser wakefield experiments, providing crucial understanding of the nonlinear laser–plasma interactions and particle distribution effects. The simulation activities are tightly tied to the experimental and theoretical efforts of the LOASIS program and collaborators at the University of Colorado and Tech-X Corporation.

Laser wakefield accelerators use laser-driven plasmas to accelerate particles in as little as a thousandth of the length required by conventional radiofrequency accelerators. The plasma-based accelerators are not subject to electrical breakdown that limits conventional accelerators and have demonstrated accelerating gradients thousands of times those obtained in conventional machines. Thus plasma-based accelerators offer a path to more compact, ultrafast particle and radiation sources for probing the subatomic world, for studying new materials and new technologies, and for medical applications.

"INCITE advanced the understanding of particle beam formation and evolution in these devices through a very large simulation in three dimensions as well as a large series of two-dimensional cases to evaluate accelerator optimization," said Geddes. "These simulations are computationally intensive because the laser wavelength (micron) must be resolved over the acceleration length of centimeters. Coupled with experiments, these simulations are developing the detailed understanding of laser acceleration needed to apply this technology to future higher energy particle physics experiments and to compact machines for medicine and laboratory science."

In laser wakefield accelerators, plasma is formed by heating hydrogen gas enough to disintegrate its atoms into their constituent protons and electrons. A laser pulse traveling through this plasma creates a wake in which bunches of free electrons are trapped and ride along, much like surfers riding the wake of a big ship (Figures 1 and 2). After propagating for a distance known as the "dephasing length," the electrons outrun the wake. This limits how far they can be accelerated and thus limits their energy. The LOASIS team's method for increasing the acceleration length is to provide a guide channel for the drive-laser pulse that creates the plasma wakefield.

PROJECT

Particle-in-Cell Simulation of Laser Wakefield Particle Acceleration

PRINCIPAL INVESTIGATOR

Cameron Geddes

SENIOR INVESTIGATORS

Wim Leemans, Eric Esarey, Carl Schroeder, William Isaacs, Berkeley Lab; John Cary, University of Colorado and Tech-X Corporation; David Bruhwiler, Tech-X Corporation

FUNDING

INCITE, HEP, NSF, AFOSR

Particle-in-cell simulations are a crucial tool in interpreting these experiments and planning the next generation because they can resolve kinetics and particle trapping. These simulations have revealed why recent experiments succeeded in producing a narrow energy spread: the trapping of an initial bunch of electrons loads the wake, suppressing further injection and forming a bunch of electrons isolated in phase space; at the dephasing point, as the bunch begins to outrun the wake, the particles are then concentrated near a single energy, and a high quality bunch is obtained (Figure 3). Only a single wake period contributes to the high energy bunch, and hence the electron bunch length is near 10 femtoseconds, indicating that a compact ultrafast electron source with high beam quality has been developed.

---

[1] W. P. Leemans, B. Nagler, A. J. Gonsalves, Cs. Tóth, K. Nakamura, C. G. R. Geddes, E. Esarey, C. B. Schroeder, and S. M. Hooker, "GeV electron beams from a centimetre-scale accelerator," Nature Physics **2,** 696 (2006). Funding: HEP, EPSRC.

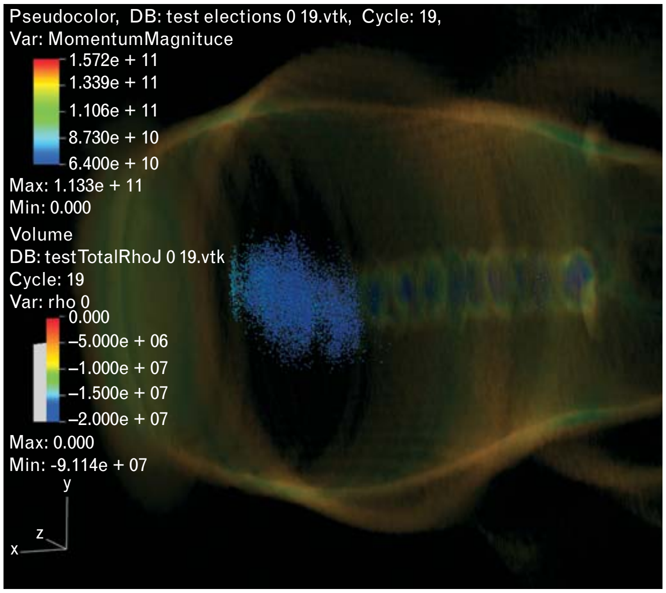

FIGURE 1. A 3D visualization showing the density of the plasma wave driven by the laser (volume shading), and positions of particles accelerated by that wave (blue spheres). (Simulation by John Cary and Cameron Geddes. Visualization by Cristina Siegerist.)

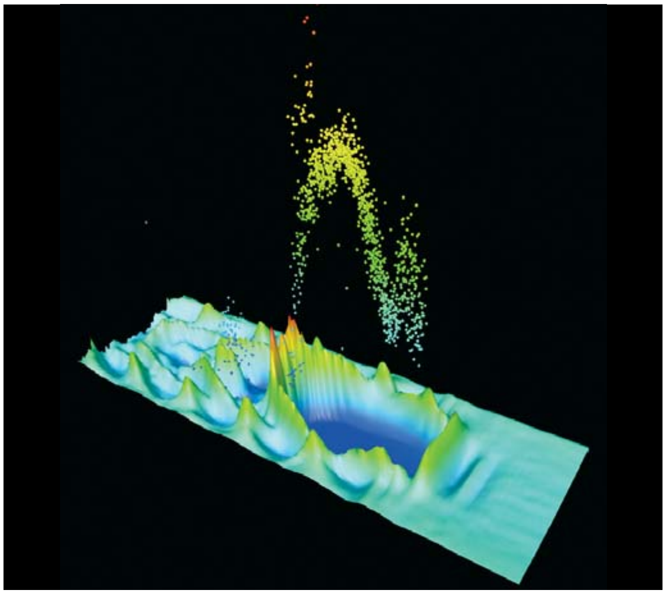

FIGURE 2. A two-dimensional cut through a 3D simulation shows the plasma wave density (surface height) and reveals the particle momentum distribution versus position (spheres, height and color = momentum). (Simulation by Cameron Geddes. Visualization by Cameron Geddes and Peter Messmer.)

Impressive as this is, "It's the tip of the iceberg," says Leemans. "We are already working on injection"—inserting an already energetic beam into an accelerating cavity—"and staging," the handoff of an energetic beam from one capillary to the next and subsequently to others, until very high energy beams are achieved. "Brookhaven physicist Bill Weng has remarked that achieving staging in a laser wakefield accelerator would validate 25 years of DOE investment in this field."

Leemans' group and their collaborators look forward to the challenge with confidence. "In DOE's Office of Science, the High Energy Physics office has asked us to look into what it would take to go to 10 GeV. We believe we can do that with an accelerator less than a meter long—although we'll probably need 30 meters' worth of laser path."

While it has often been said that laser wakefield acceleration promises high-energy accelerators on a tabletop, the real thing may not be quite that small. But laser wakefield acceleration does indeed promise electron accelerators potentially far more powerful than any existing machine—neatly tucked inside a small building.

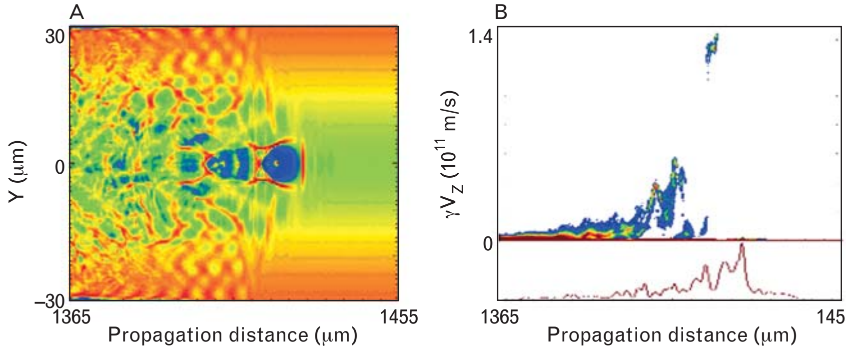

FIGURE 3. Particle-in-cell simulations show that a high-quality bunch is formed when beam loading and pulse evolution turn off injection after the loading of an initial bunch, resulting in an isolated bunch (A). These particles are then concentrated in energy at the dephasing length (B), forming a low-energy-spread bunch. The laser pulse envelope is shown at the bottom of pane (B), showing self modulation into sub-pulses at the plasma period.

This article written by: John Hules, Jon Bashor, Paul Preuss, and Ucilia Wang (Berkeley Lab)

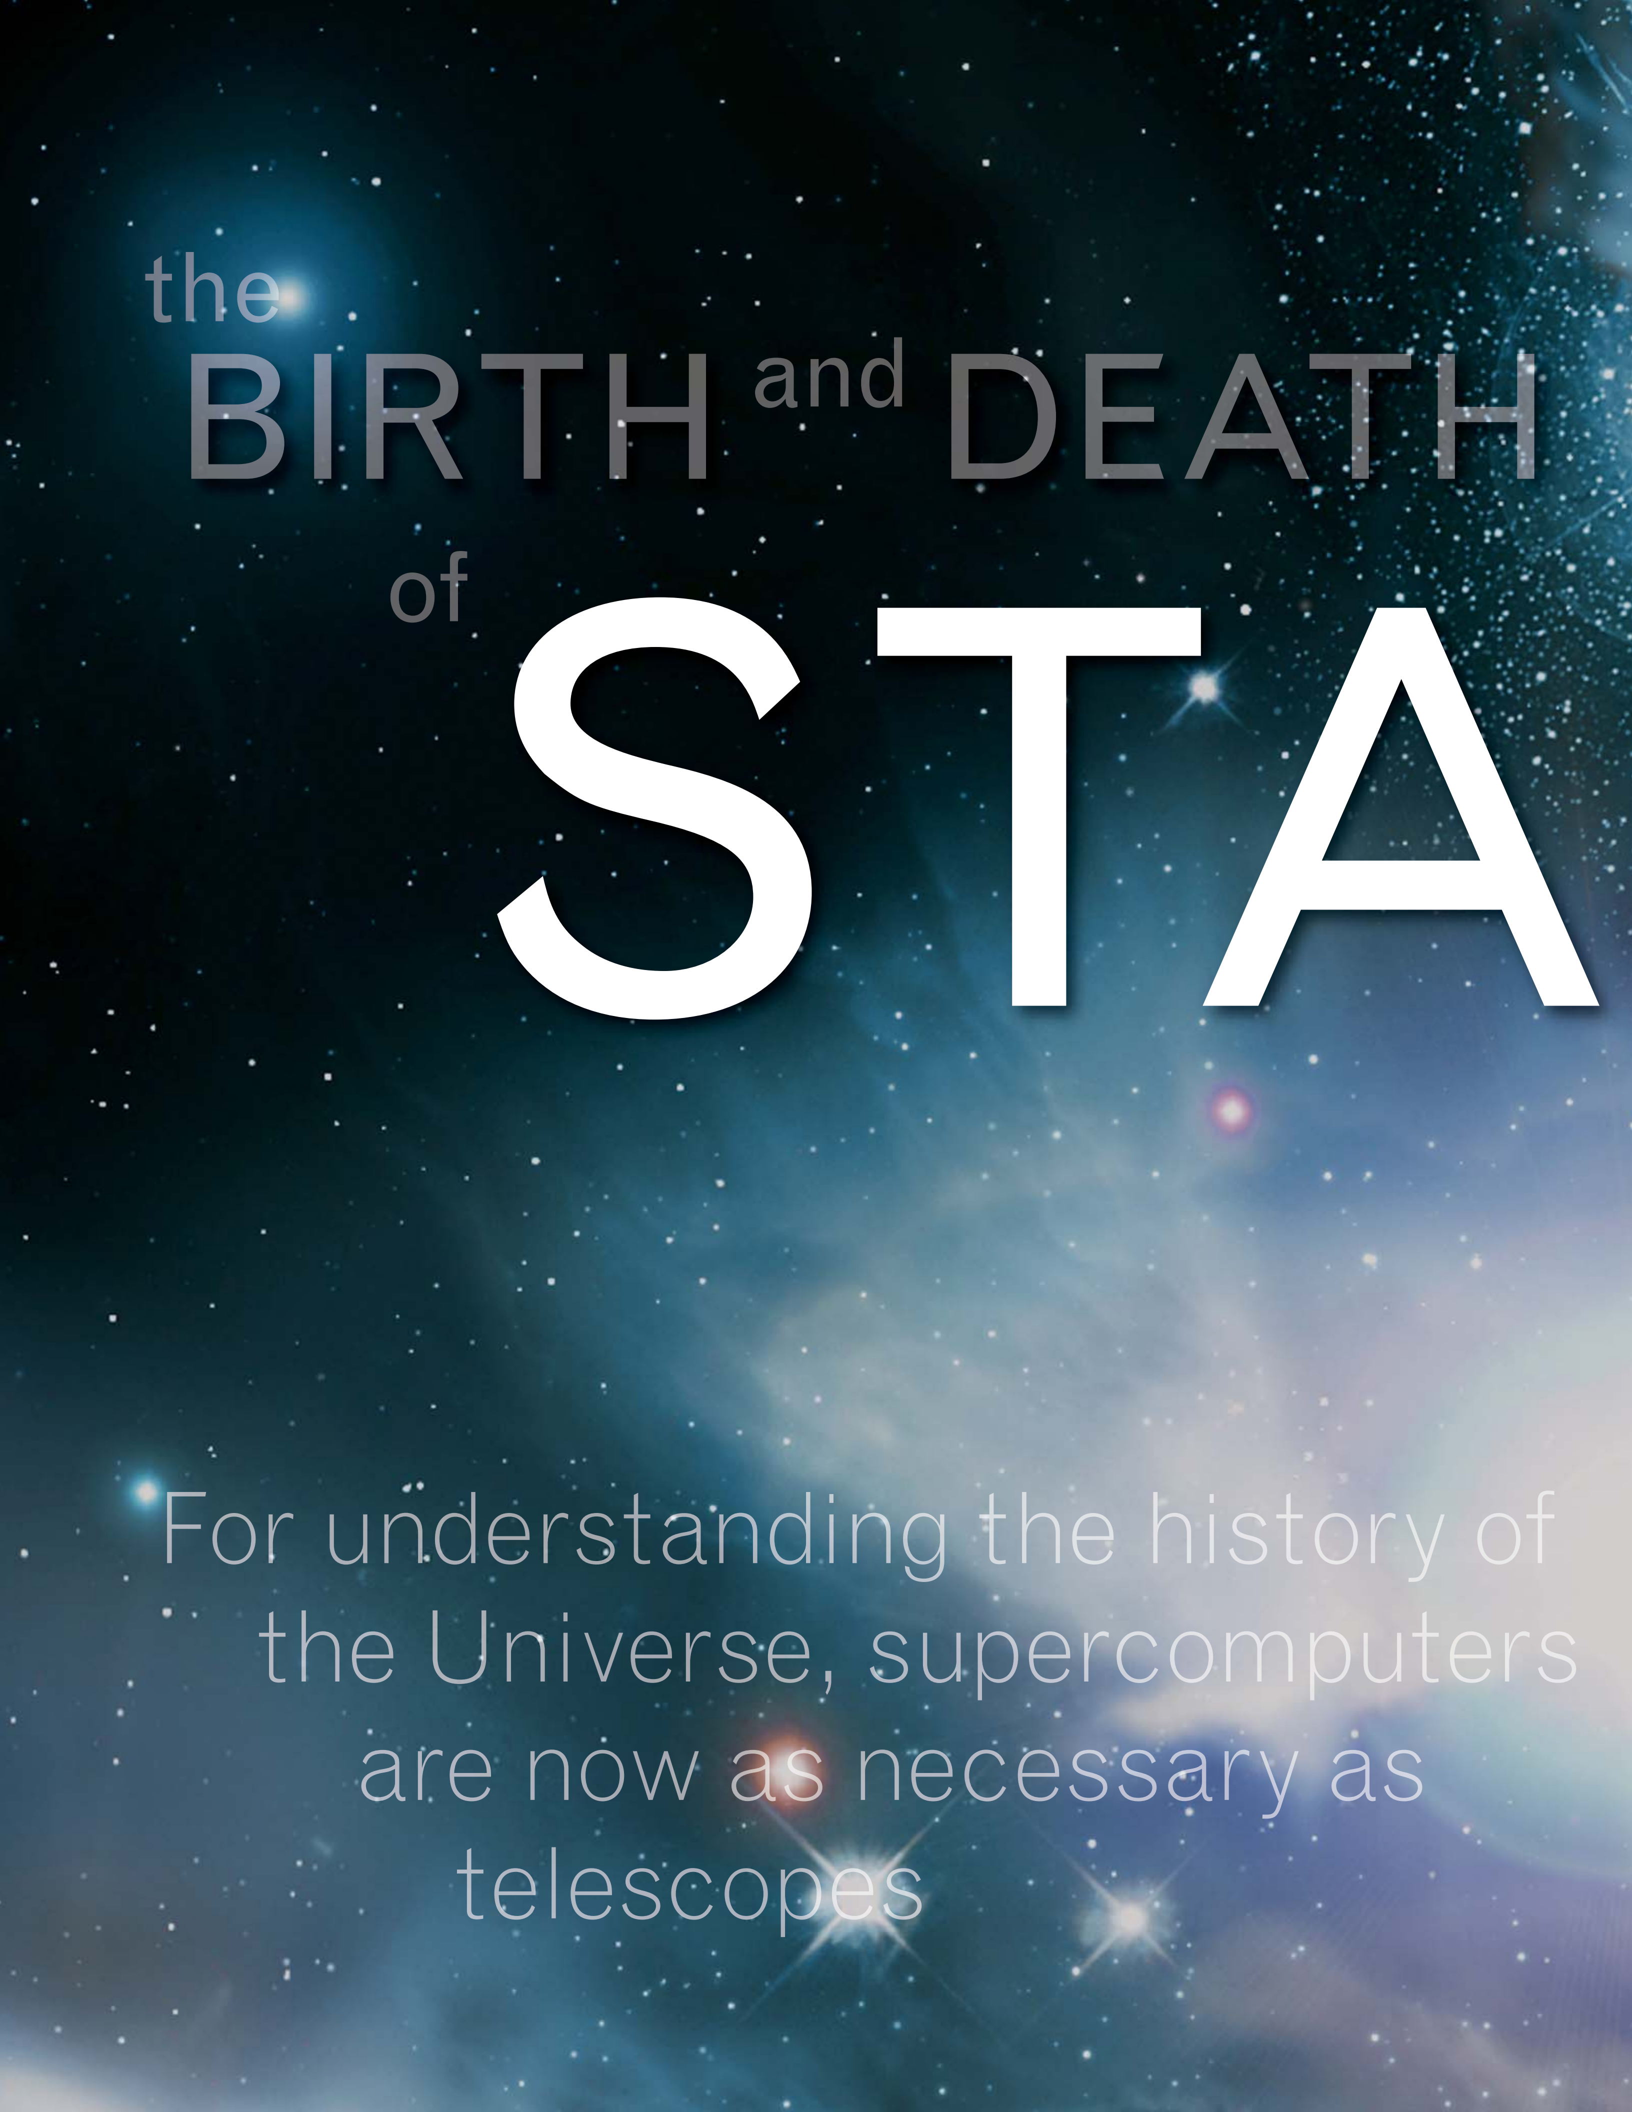

the BIRTH and DEATH of STA

For understanding the history of the Universe, supercomputers are now as necessary as telescopes

RS

Stars may seem distant, but we would not be here without them. The Big Bang produced hydrogen, helium, and lithium, but all of the heavier elements that make up our planet and our bodies were synthesized in stars and supernovae. To understand our Universe, we need to understand the formation and life cycles of stars. And several teams of researchers are using NERSC's supercomputers to do just that.

HIGH-MASS STAR FORMATION

The formation of high-mass stars remains one of the most significant unsolved problems in astrophysics. These stars, with masses from 10 to 100 times the mass of our sun, eventually explode as supernovae and produce most of the heavy elements in the Universe. They also have a major influence on the structure and evolution of galaxies. But observing the formation of massive stars is difficult, because they are born in distant, dense, and dusty regions of space, and they swallow up much of their birth environment as they are born.

Massive star formation also poses major theoretical challenges. Massive stars begin burning their nuclear fuel and radiating prodigious amounts of energy while still accreting mass from the dense clouds of mostly hydrogen gas surrounding them. But this radiation exerts a repellent effect on molecules in the accreting material that could theoretically exceed the attractive force of gravity. This paradox poses a question: How can a massive protostellar core sustain a high-mass accretion rate despite its repellent radiation pressure on the surrounding matter?

That is only one of the questions that Richard Klein and his collaborators, Christopher McKee and Mark Krumholz, are determined to answer. Klein is an adjunct professor of astronomy at UC Berkeley and a researcher at the Lawrence Livermore National Laboratory; McKee is a physics and astronomy professor at UC Berkeley; and Krumholz is now a post-doc at Princeton University. Together they are working toward a comprehensive theory of star formation, and they have already made major progress.

For a long time, scientists have understood that stars form when interstellar matter inside giant clouds of molecular hydrogen undergoes gravitational collapse, but the puzzle remained of how the protostars could grow to become high-mass stars in spite of the strong radiation and stellar winds that they generate. That question led to two competing theories on how massive stars come into being.

In the *competitive accretion* theory, the cloud gravitationally collapses to produce clumps containing small protostellar cores. These cores are the seeds which undergo growth by gravitationally pulling in matter from around them, competing with other cores in the process, and sometimes colliding and merging with other cores, eventually accreting many times their original mass.

The rival *direct gravitational collapse* theory, which Klein and his collaborators subscribe to, contends that the protostellar cores are already large soon after the star-forming clouds have fragmented into clumps. These cores subsequently collapse to make individual high-mass stars or small multiple systems, in either case continuing to accrete some matter from the parent clump, but not enough to change their mass substantially.

PROJECT:

Toward a Comprehensive Theory of Star Formation

PRINCIPAL INVESTIGATOR:

Richard Klein, University of California, Berkeley, and Lawrence Livermore National Laboratory

SENIOR INVESTIGATORS:

Christopher McKee, UC Berkeley; Mark Krumholz, Princeton University

FUNDING:

NP, NASA, NSF

COMPUTING RESOURCES:

NERSC, SDSC, LLNL

Krumholz, McKee, and Klein gave a major boost to the direct gravitational collapse theory in the November 17, 2005 issue of Nature,[1] where they reported the results of star formation simulations carried out at NERSC and the San Diego Supercomputer Center. "Our work was the first attempt with fully three-dimensional simulations to show how high-mass stars are formed, and it dealt a serious blow to the competitive accretion theory," said Klein.

---

[1] Mark R. Krumholz, Christopher F. McKee, and Richard I. Klein, "The formation of stars by gravitational collapse rather than competitive accretion," *Nature* **438**, 332 (2005).

A STAR IS BORN

Figure 1 shows a time sequence of the evolution of one simulation run, starting from the initial state shown in the top row. Turbulence delays the gravitational collapse for a while, but as the turbulence decays, gas starts to collapse. The primary star appears 5,300 years after the start of the simulation. It forms in a shocked filament, which continues to accrete mass and by 6,000 years is beginning to form a flattened protostellar disk (second row).

As the evolution continues, several more dense condensations appear, but most of these are sheared apart in the primary protostellar disk before they can collapse and form a protostar. After 12,200 years, a second protostar forms, but it falls into the primary star and merges with it at 12,700 years, before it has accreted one-tenth of a solar mass of gas. The primary star is already 2.1 solar masses, so the mass it gains in the merger is negligible. The third row shows the state of the simulation at 12,500 years, about halfway between when the second protostar appears and when it merges with the primary.

Only after 14,400 years does one of the condensations collapse to form a second protostar that is not immediately accreted, as shown in the fourth row of Figure 1. At this point the primary star is 3.2 solar masses and has a well-defined massive disk. The condensation from which the new protostar forms is already visible in the third row. Unlike several others, it is able to collapse and form a protostar because it is fairly distant from the primary protostar, which reduces the amount of radiative heating to which it is subjected.

The next significant change in the system occurs when one of the arms of the disk becomes unstable and fragments to form a third protostar at 17,400 years, as shown in the fifth row. At this point the central star mass is 4.3 solar masses; the fragment is very small in comparison.

The configuration after 20,000 years of evolution, shown in the sixth row, is substantially similar. Two more small disk fragments form, but they both collide with the primary star almost immediately after formation. At the end of 20,000 years, the primary star is 5.4 solar masses, the second star is 0.34 solar masses, and the third star, which formed in the disk of the first, is 0.2 solar masses. The disk itself is 3.4 solar masses. The system is well on its way to forming a massive star, and thus far the vast majority of the collapsed mass has concentrated into a single object. A larger plot of the full core at this point is shown in Figure 2.

Thus far the researchers have carried out 3D simulations, not yet complete, that show stars greater then 35 solar masses forming and still accreting gas from the surrounding turbulent core.

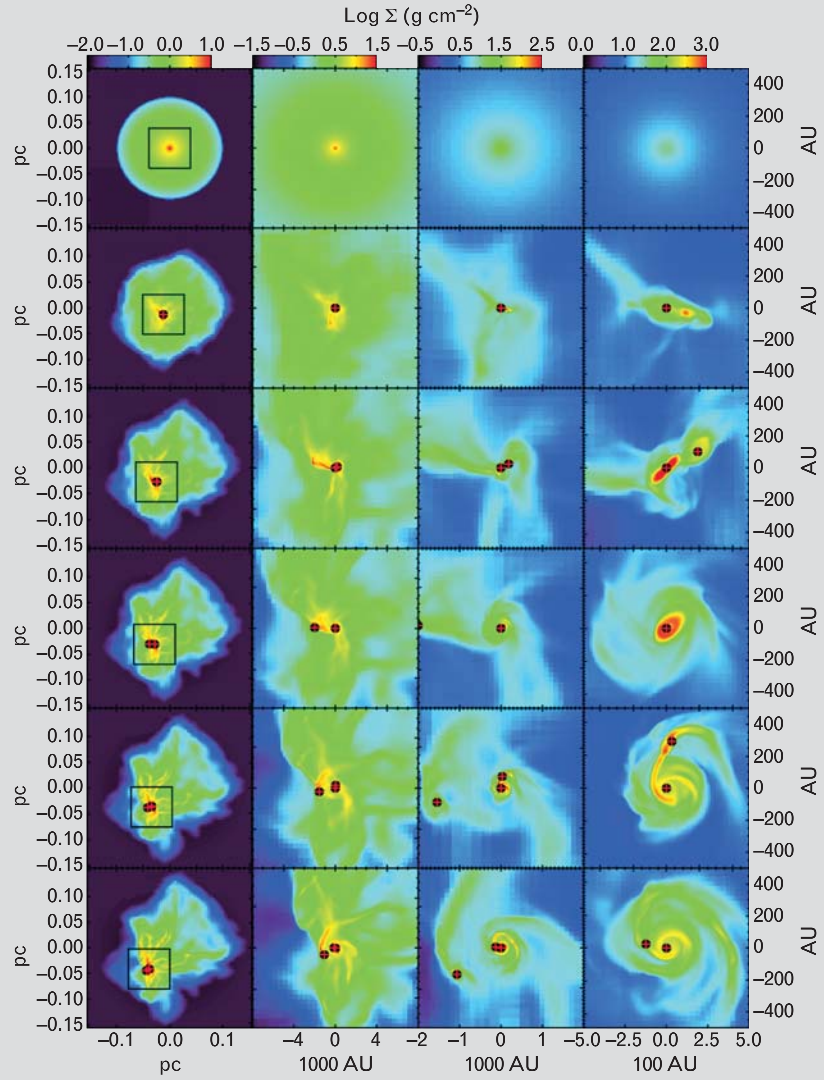

FIGURE 1. Column density as a function of time in one of a series of star formation simulations. From top to bottom, the rows show the cloud state over time. From left to right, each column "zooms out" to show 4 times more area. In the left column, the image is always centered on the point of origin, and the region shown in the second column is indicated by the black box. In the other columns, the image is centered on the location of the primary star at that time. Stars are indicated by red plus signs.

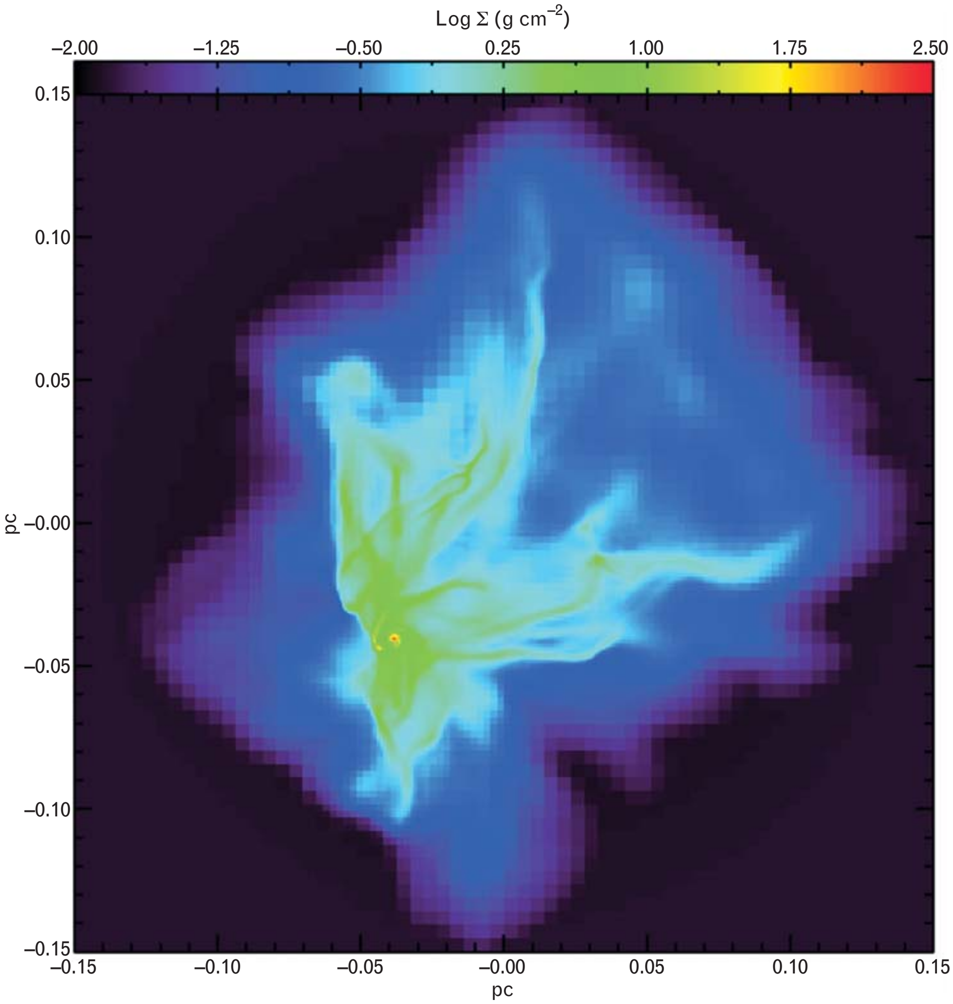

FIGURE 2. Column density of a simulated protostellar core 20,000 years after the beginning of gravitational collapse.

The 3D simulations determined the conditions in which competitive accretion can occur: low turbulence (the ratio of turbulent kinetic energy to gravitational potential energy) in the gas clumps in which the cores are formed, and low-mass clumps (a few solar masses). Earlier three-dimensional simulations with particle-based codes that appeared to support competitive accretion were based on these assumptions, and also had the turbulence taper out quickly in the star-forming process.

But do these two conditions necessary for competitive accretion actually exist? After reviewing observations of a broad sample of star-forming regions, Klein and his team found no evidence to support these two assumptions. On the contrary, star-forming regions show significant turbulence, and the clumps tend to have several thousand solar masses. "Every observation of these large clouds indicates that a mechanism, perhaps protostellar winds, must be present that keeps stirring the clouds to keep the turbulence around," Klein said.

The researchers have also demonstrated that radiation pressure is a much less significant barrier to massive star formation than has previously been thought. In proof-of-principal calculations of the recently observed gas outflows from massive protostars, they found that an outflow can substantially change the radiation field and radiation pressure around the protostar. The outflow cavity in the surrounding gaseous envelope provides a thin channel through which radiation can escape, significantly reducing the radiation pressure and allowing accretion to continue. "Surprisingly," they concluded, "outflows that drive gas out of a collapsing envelope may increase rather than decrease the size of the final massive star."[2]

Another issue for the direct gravitational collapse theory to resolve is fragmentation: why wouldn't a massive core collapse into many fragmented, low-mass protostars rather than one or a few high-mass stars? In three-dimensional simulations with a wide range of initial conditions, the researchers found that radiation feedback from accreting protostars inhibits the formation of fragments, so that the vast majority of the gas collapses into a handful of objects, with the majority of the mass accreting onto one primary object (see sidebar on previous page).[3] The emerging picture, then, is that massive cores are the direct progenitors of massive stars, without an intermediate phase of competitive accretion or stellar collisions.

Klein's team created these simulations using a code called Orion, which employs adaptive mesh refinement (AMR) to create three-dimensional simulations over an enormous range of spatial scales. "AMR enabled us for the first time to cover the full dynamic range with numerical simulations on a large scale, not just in star formation but in cosmology," Klein said. "We want to solve the entire problem of the formation of high-mass stars."

FROM SOUNDWAVES TO SUPERNOVAE

Once every 30 to 50 years in our galaxy—and then just for a few milliseconds—an exploding star known as a core-collapse (Type II) supernova emits as much energy in neutrinos as is emitted in photons by all the other stars in the Universe combined.

Supernovae have been documented for 1,000 years, and astrophysicists know a lot about how they form, what happens during the explosion and what's left afterward. But for the past 40 years, one problem has

---

[2] Mark R. Krumholz, Christopher F. McKee, and Richard I. Klein, "How protostellar outflows help massive stars form," *Astrophysical Journal* **618,** L33 (2005).

[3] Mark R. Krumholz, Richard I. Klein, and Christopher F. McKee, "Radiation-hydrodynamic simulations of collapse and fragmentation in massive protostellar cores," *Astrophysical Journal* **656,** 959 (2007).

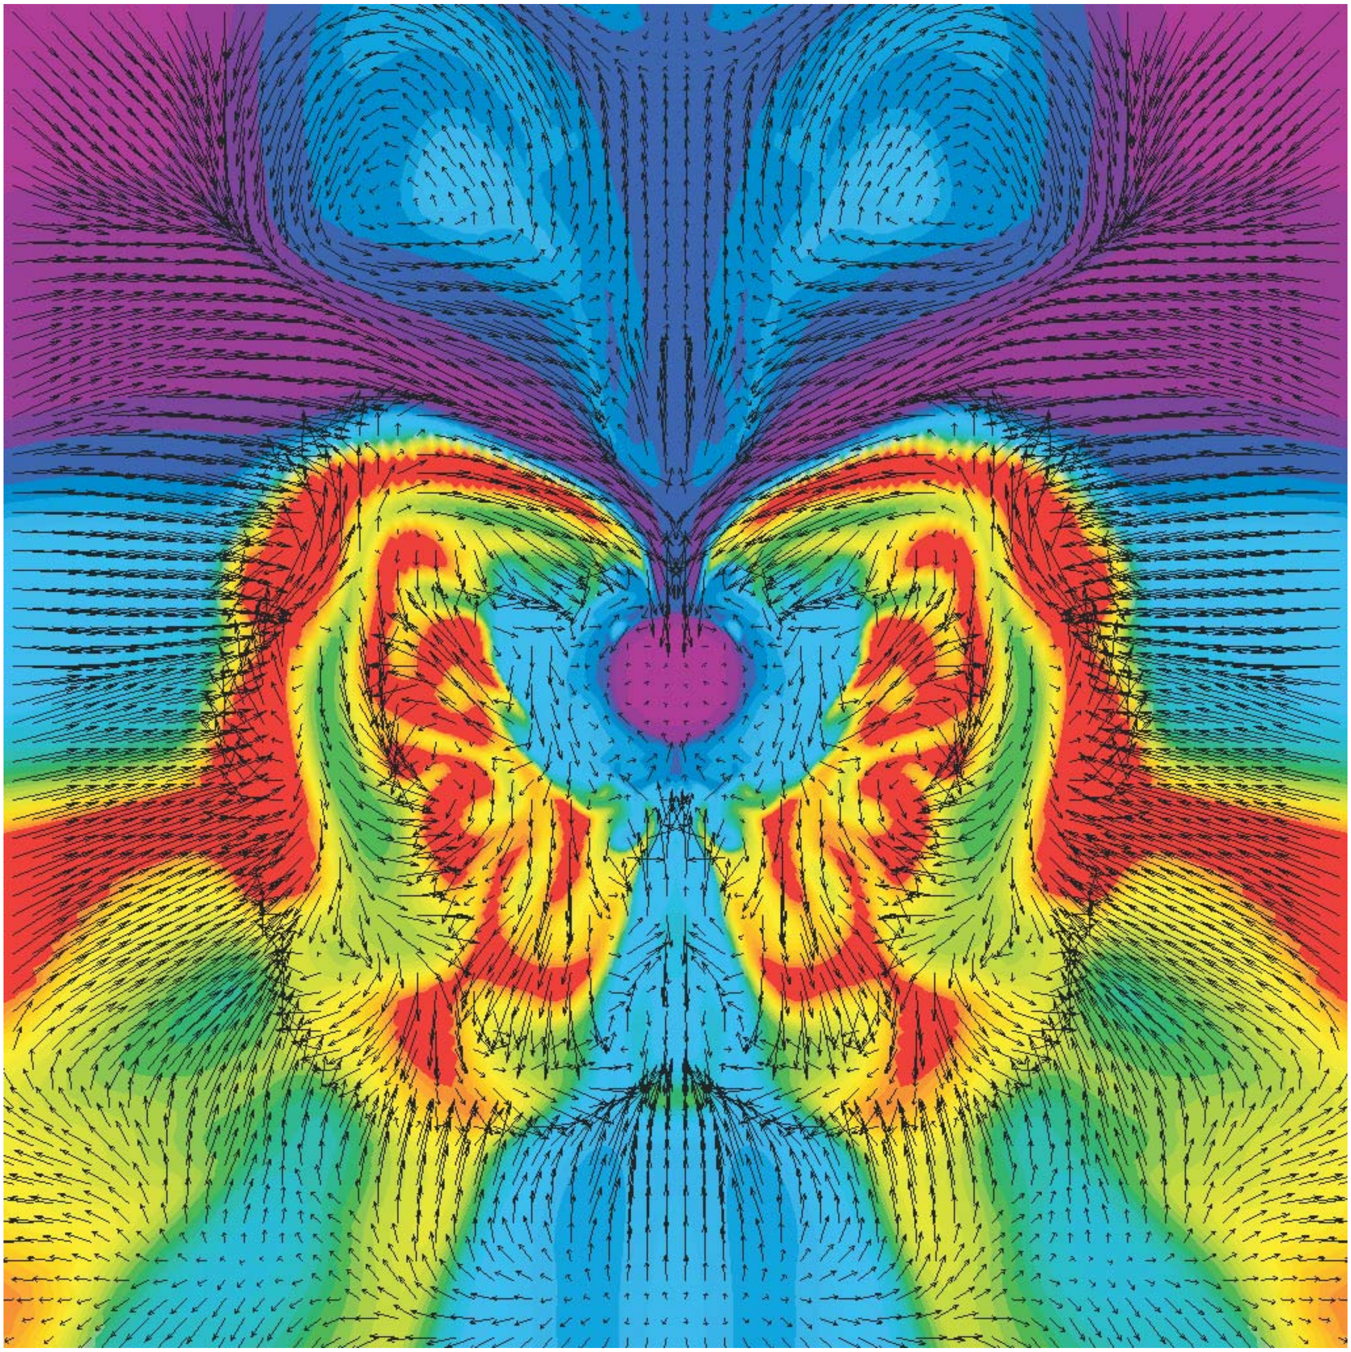

FIGURE 3. A 2D rendition of the entropy field of the early blast in the inner 500 km of an exploding supernova. Velocity vectors depict the direction and magnitude of the local flow. The bunching of the arrows indicates the crests of the sound waves that are escalating into shock waves. These waves are propagating outward, carrying energy from the core to the mantle and helping it to explode. The purple dot is the protoneutron star, and the purple streams crashing in on it are the accretion funnels.

dogged astrophysicists—what is the mechanism that actually triggers the massive explosion? Hydrodynamic, neutrino, convective, viscous, and magnetic mechanisms for driving core-collapse supernova explosions have all been proposed and investigated. One thing that is known is that Type II supernovae produce neutrinos, particles with very little mass which travel through space and everything in their path. Neutrinos carry energy from the deep interior of the star, which is being shaken around like a jar of supersonic salad dressing, and deposit the energy on the outer region. One theory holds that if the neutrinos deposit enough energy throughout the star, this may trigger the explosion.

To study this, a group led by Adam Burrows, professor of astronomy at the University of Arizona and a member of the SciDAC Computational Astrophysics Consortium, developed codes for simulating the behavior of a supernova core in two dimensions. While a 3D version of the code would be optimum, it would take at least five more years to develop and would require up to 300 times as much computing time. As it was, the group ran 1.5 million hours of calculations at NERSC.

But the two-dimensional model is suitable for Burrows' work, and the instabilities his group is interested in studying can be seen in 2D. What they found was that there is a big overturning motion in the core, which leads to wobbling, which in turn creates sound waves. These waves then carry energy away from the core, depositing it farther out near the mantle.

According to Burrows, these oscillations could provide the power that puts the star over the edge and causes it to explode. To imagine what such a scenario would look like, think of a pond into which rocks are thrown, causing waves to ripple out. Now think of the pond as a sphere, with the waves moving throughout the sphere. As the waves move from the denser core to the less dense mantle, they speed up. According to the model, they begin to crack like a bullwhip, which creates shockwaves. It is these shockwaves, Burrows believes, which could trigger the explosion (Figures 3 and 4).

So, what led the team to this new model of an acoustic triggering mechanism? They came up with the idea by following the pulsar—the neutron star which is the remains of a supernova. They wanted to explore the origin of the high speed which pulsars seem to be born with, and this led them to create a code that allowed the core to move. However, when they implemented this code, the core not only recoiled, but oscillated and generated sound waves.

The possible explosive effect of the oscillations had not been considered before because previous simulations of the condi-

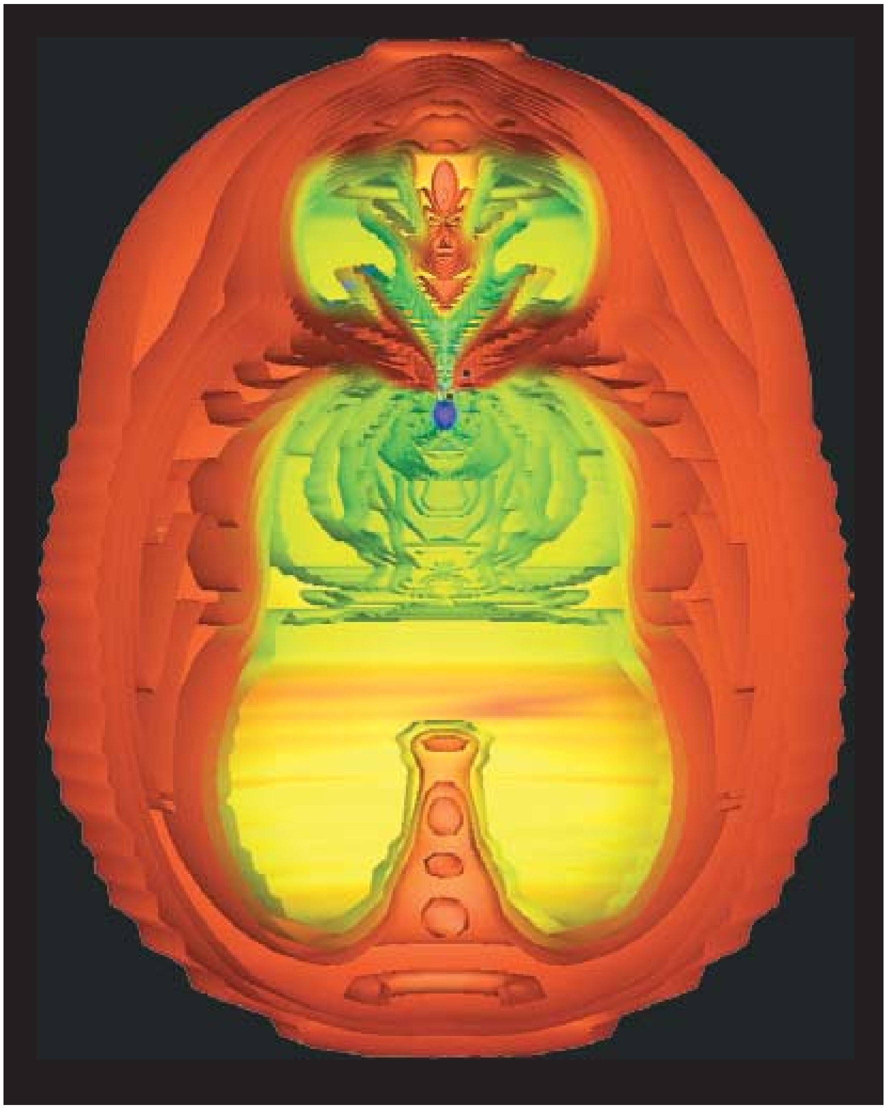

FIGURE 4. This shell of isodensity contours, colored according to entropy values, shows simultaneous accretion on the top and explosion on the bottom. The inner green region is the blast, and the outer orange region is the unshocked material that is falling in. The purple dot is the newly formed neutron star, which is accumulating mass through the accretion funnels (in orange).

PROJECT:

Computational Astrophysics Consortium

PRINCIPAL INVESTIGATOR:

Stan Woosley, University of California, Santa Cruz

SENIOR INVESTIGATORS:

Adam Burrows, University of Arizona; Alex Heger, Los Alamos National Laboratories; Rob Hoffman, Lawrence Livermore National Laboratory; Jon Arons, Richard Klein, and Christopher McKee, University of California, Berkeley; Roger Blandford, Stanford University; Gary Glatzmaier, University of California, Santa Cruz; Peter Nugent, John Bell, and Saul Perlmutter, Lawrence Berkeley National Laboratory; Mike Zingale, State University of New York, Stony Brook

FUNDING:

HEP, SciDAC, NSF, ISF, JINA

COMPUTING RESOURCES:

NERSC, AEI

tions inside the core used smaller time steps, which consumed more computing resources. With this limitation, the simulations ran their course before the onset of oscillations. With SciDAC support, however, Burrows' team was able to develop new codes with larger time steps, allowing them to model the oscillations for the first time.

In the paper resulting from this research,[4] neutrino transfer is included as a central theme in a 2D multi-group, multi-neutrino, flux-limited transport scheme—the first truly 2D neutrino code with results published in the archival literature. The results are approximate but include all the important components.

Calling the simulation a "real numerical challenge," Burrows said the resulting approach "liberated the inner core to allow it to execute its natural multidimensional motion." This motion led to the excitation of the core, causing the oscillations at a distinct frequency.

The results look promising, but as is often the case, more research is needed before a definitive mechanism for triggering a Type II supernova is determined. For example, if a simulation with better numerics or three dimensions produces a neutrino triggering mechanism that explodes the star earlier, then the acoustic mechanism would be aborted. Whether this happens remains to be seen and is the subject of intense research at NERSC and elsewhere.

"The problem isn't solved," Burrows said. "In fact, it's just beginning."

CALIBRATING COSMOLOGY

Type Ia supernovae are extraordinarily bright, remarkably uniform exploding stars which make excellent "standard candles" for measuring the expansion rate of the Universe at different times in its history. Researchers use supernovae's distance and the redshift of their home galaxies to calculate the speed at which they are moving away from us as the Universe expands. In 1998, by comparing the redshifts of dozens of supernovae, scientists discovered that, contrary to expectations, the expansion of the Universe is speeding up, and they coined the term *dark energy* to designate the unknown force behind this acceleration. Subsequent observations and calculations have determined that dark energy makes up about 70 percent of the Universe.

The nature of dark energy has been called the deepest mystery in physics, and its resolution may revolutionize our understanding of matter, space, and time. "This discovery has revolutionized cosmology," said Greg Aldering, head of the Nearby Supernova Factory (SNfactory) at Lawrence Berkeley National Laboratory. "Now astrophysicists want to understand the physical cause for the dark energy. This requires more precise measurements, and with large numbers of accurately measured Type Ia supernovae this should be possible."

The SNfactory is an international collaboration between astrophysicists at Berkeley Lab, at several institutions in France, and at Yale University. The aim of the collaboration is to discover as many nearby (low-redshift) supernovae as possible and to study them in detail in order to reduce the statistical uncertainties in previous data. Supernovae can then be used more effectively as cosmological distance indicators to measure the expansion history of the Universe and explore the nature of dark energy.

"The ingredients which go into a supernova explosion are fairly well known," Aldering continued, "and although computer modelers are not yet able to accurately predict the properties of supernovae in great detail, they do know something about how supernova properties change when the input ingredients are changed. Since measuring the change in the expansion rate of the Universe requires only relative distances, astrophysicists simply need to understand how supernovae will change in brightness when their input ingredients are changed by small amounts. This question can be explored empirically using nearby supernovae, which have a wide range of values for these input ingredients. Such exploration—and comparison with computational studies—is the basic goal of the Nearby Supernova Factory."

PROJECT:

The Nearby Supernova Factory

PRINCIPAL INVESTIGATOR:

Greg Aldering, Lawrence Berkeley National Laboratory

SENIOR INVESTIGATORS:

Saul Perlmutter, Peter Nugent, Cecilia Aragon, and Stewart Loken, Berkeley Lab

FUNDING:

HEP, SciDAC, ASCR, NSF, GBMF, CNRS/IN2P3, CNRS/INSU, PNC

Of course, the first step in studying nearby supernovae is to find them. "For the studies needed, we would like to discover the supernovae as soon as possible after they explode," Aldering explained. "This requires imaging the night sky repeatedly, returning to the same fields every few nights, and then quickly processing the data."

To that end, the SNfactory collaboration has built an automated system consisting of custom-built hardware and software that systematically searches the sky for new supernovae, screens potential candidates, then performs multiple spectral and photometric observations on each supernova. The imaging is done by a powerful CCD (charge-coupled device) camera built by the QUEST group that delivers 300 MB of imaging data every 100 seconds, amounting to 50–100 GB per night. These data are transferred via the High Performance Research and Education Network and ESnet to NERSC, where digital image subtraction software, running on the PDSF cluster, compares the new images to images of the same field archived on NERSC's HPSS system to find the light of any new supernovae.

"The processing and storage of these images requires resources on a scale only available at NERSC," Aldering said. "Near-line storage is critical for this project since archival images must be rapidly recalled so they can be compared with new images."

---

[4] A. Burrows, E. Livne, L. Dessart, C. D. Ott, and J. Murphy, "A new mechanism for core-collapse supernova explosions," *Astrophysical Journal* **640,** 878 (2006).

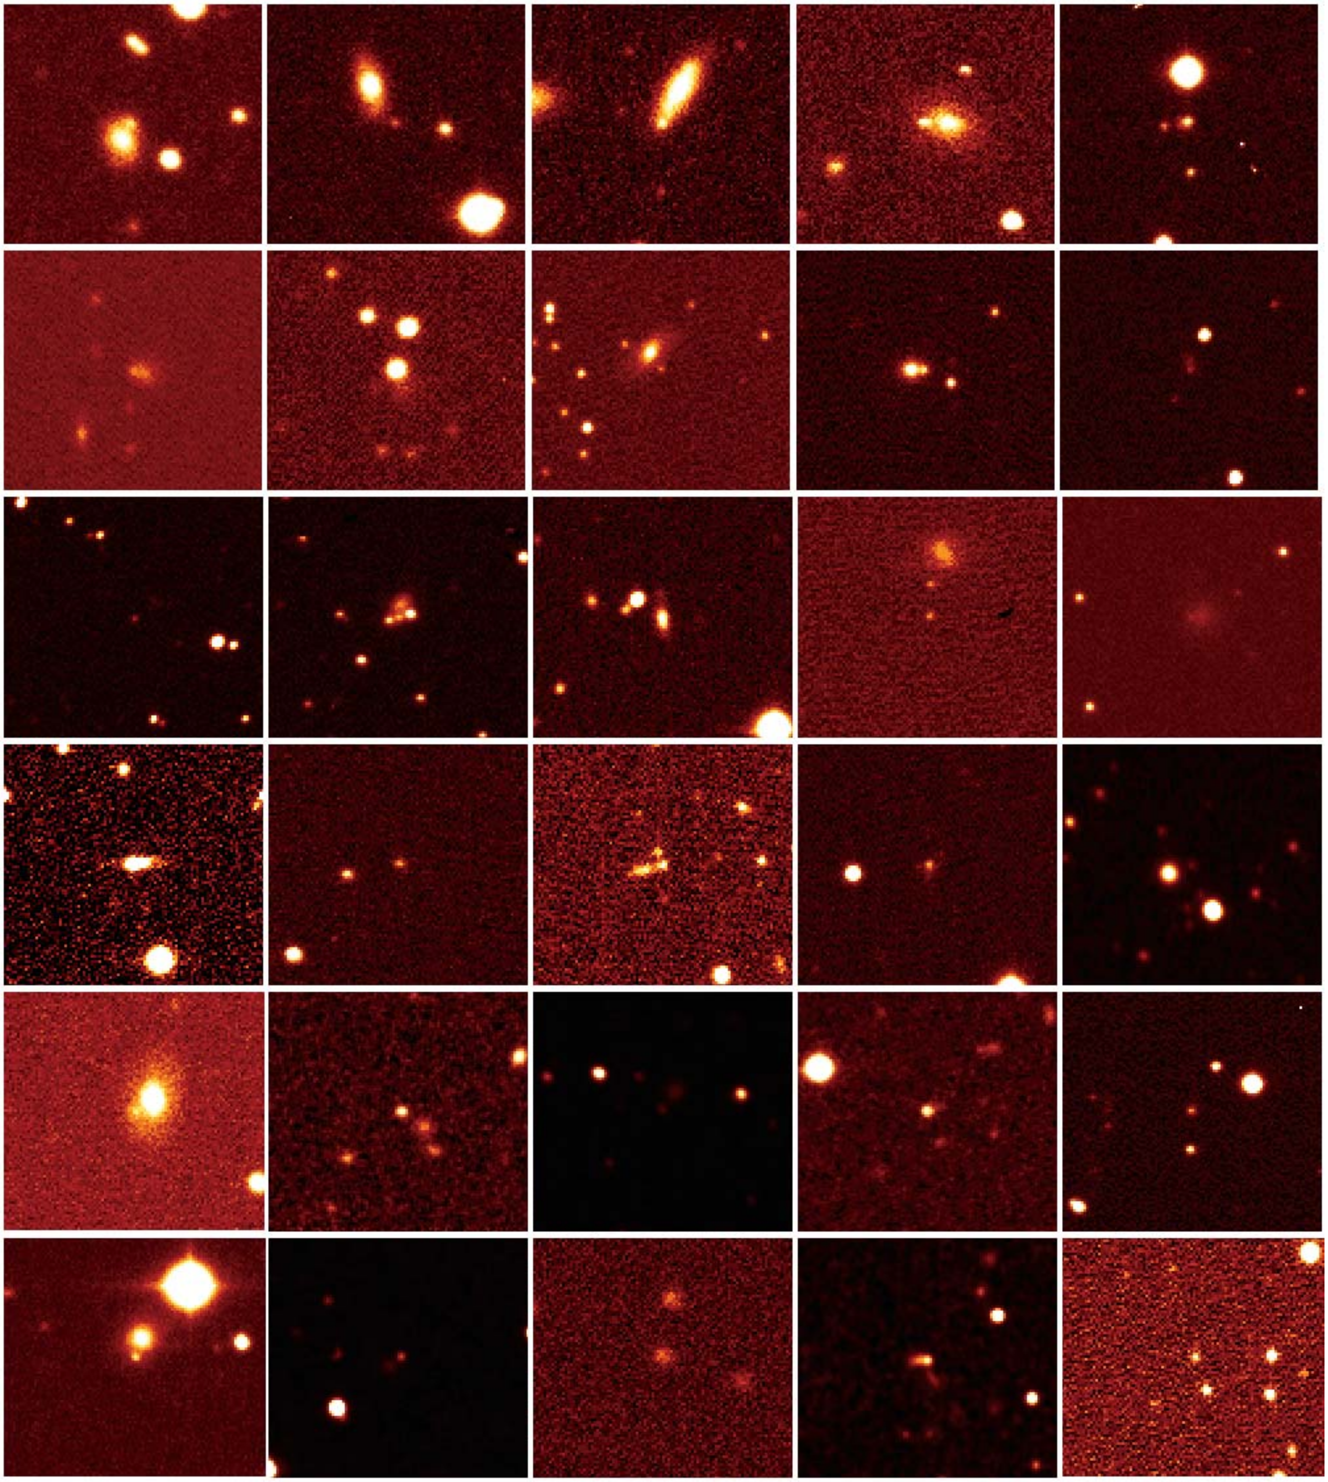

FIGURE 5. A small sample of the hundreds of supernovae discovered by the Nearby Supernova Factory.

In 2006 the SNfactory processed 4,172,340 images, which corresponds to 11 TB of raw data and 356,185 square degrees worth of images (Figure 5). These images covered approximately half the sky, 20 times over, at a resolution of 0.7 microradians (or one-fourth of a millidegree). This data resulted in 603,518 subtractions processed (one PDSF job per subtraction), more than 16 million output files, and a total of 17 TB of output data. The database of supernova spectra obtained by the SNfactory is now the most extensive in the world, and it soon will be available to researchers worldwide.

From this data, 249 supernovae were discovered and confirmed photometrically in 2006. In addition, using the SNfactory's SuperNova Integral Field Spectrograph (SNIFS), remotely operated on the University of Hawaii's 2.2-meter telescope on Mauna Kea, Aldering's team spectroscopically confirmed 136 supernovae: 89 Type Ia, 41 Type II, and 6 Type Ib/c. They also have used their trove of images to eliminate some objects as variable stars or quasars—as opposed to supernovae—based on their brightness behavior over the past six years in which the SNfactory has been archiving images.

During the past year, the SNfactory has implemented machine learning techniques that decreased the time required for human verification by a factor of 10, allowing the researchers to focus their attention on the best candidates immediately. This improved efficiency has resulted in a one-day turnaround—a supernova imaged at Palomar Observatory on one night is discovered by the SNfactory the next day and confirmed as a supernova the following night by SNIFS in Hawaii.

One example of the importance of early spectroscopy is a supernova designated SN 2006D. The SNfactory obtained SNIFS spectra of SN 2006D only three days after it was discovered by the Brazilian Supernova Search team and one week before it reached maximum brightness. The SN 2006D spectra provided the most definitive evidence to date of unburned carbon in a Type Ia supernova.[5] The white dwarf stars that explode as Type Ia supernovae are composed primarily of carbon and oxygen, most of which is burned into heavier elements by nuclear fusion during the explosion; and as expected, SN 2006D's carbon signature dissipated as the supernova approached peak brightness. But the presence of detectable unburned carbon in the early light of SN 2006D provides valuable data for researchers seeking to understand variations in supernova progenitors as well as explosion mechanisms.

To make their archived data even more useful, the SNfactory is now in the process of coadding (calculating average measurements for) the millions of distinct images of the night sky they have obtained over the past six years. Typically each sky location has been visited 40 times in the past six years, so the resulting coadded sky images should be very robust. Each set of 40 images covering one square degree of sky will be combined to construct a 4000-by-4000-pixel image, for a total of roughly 80,000 overlapping images. These images cover two-thirds of the entire sky and will constitute the greatest combination of depth and sky coverage of any sky atlas ever generated.

This article written by: John Hules, Jon Bashor, and Ucilia Wang, Berkeley Lab.

---

[5] R. C. Thomas et al. (The Nearby Supernova Factory), "Nearby Supernova Factory Observations of SN 2006D: On Sporadic Carbon Signatures in Early Type Ia Supernova Spectra," *Astrophysical Journal* **654,** L53 (2007).

Nobel Prize for Physics Leads List of Honors for NERSC Users in 2006

George Smoot

George F. Smoot, leader of a research team that was able to image the infant universe, revealing a pattern of minuscule temperature variations which evolved into the universe we see today, was awarded the 2006 Nobel Prize for physics. Smoot, an astrophysicist at Lawrence Berkeley National Laboratory and a University of California at Berkeley physics professor, shares the award with John C. Mather of NASA Goddard Space Flight Center. The citation reads "for their discovery of the blackbody form and anisotropy of the cosmic microwave background radiation."

On May 1, 1992, at a meeting of the American Physical Society, Smoot made an announcement that essentially silenced all the scientific critics of the Big Bang theory and helped change the course of future investigations into the origin and evolution of the universe. Smoot and his research team, after analyzing hundreds of millions of precision measurements in the data they had gathered from an experiment aboard NASA's Cosmic Background Explorer (COBE) satellite, had produced maps of the entire sky which showed "hot" and "cold" regions with temperature differences of a hundred-thousandth of a degree. These temperature fluctuations, produced when the universe was smaller than a single proton, were consistent with Big Bang predictions and are believed to be the primordial seeds from which our present universe grew.

The COBE data analysis was conducted using computers at Berkeley Lab and NASA's Goddard Space Flight Center, but as subsequent experiments produced ever larger data sets, faster computers and bigger data archives were required. So, over the past ten years, Smoot and his colleagues studying cosmic microwave background (CMB) data have used nearly 5 million processor-hours and tens of terabytes of disk space at NERSC. In fact, around 100 analysts from a dozen CMB experiments are now NERSC users.

Two years ago, to simulate processing an entire year's worth of data from Planck, the third-generation CMB space mission, Smoot's team used 6,000 processors on NERSC's Seaborg supercomputer for nearly two hours—the first time virtually all of the processors were used on a single code—mapping 75 billion observations to 150 million pixels. For comparison, the COBE sky map had used only 6,144 pixels.

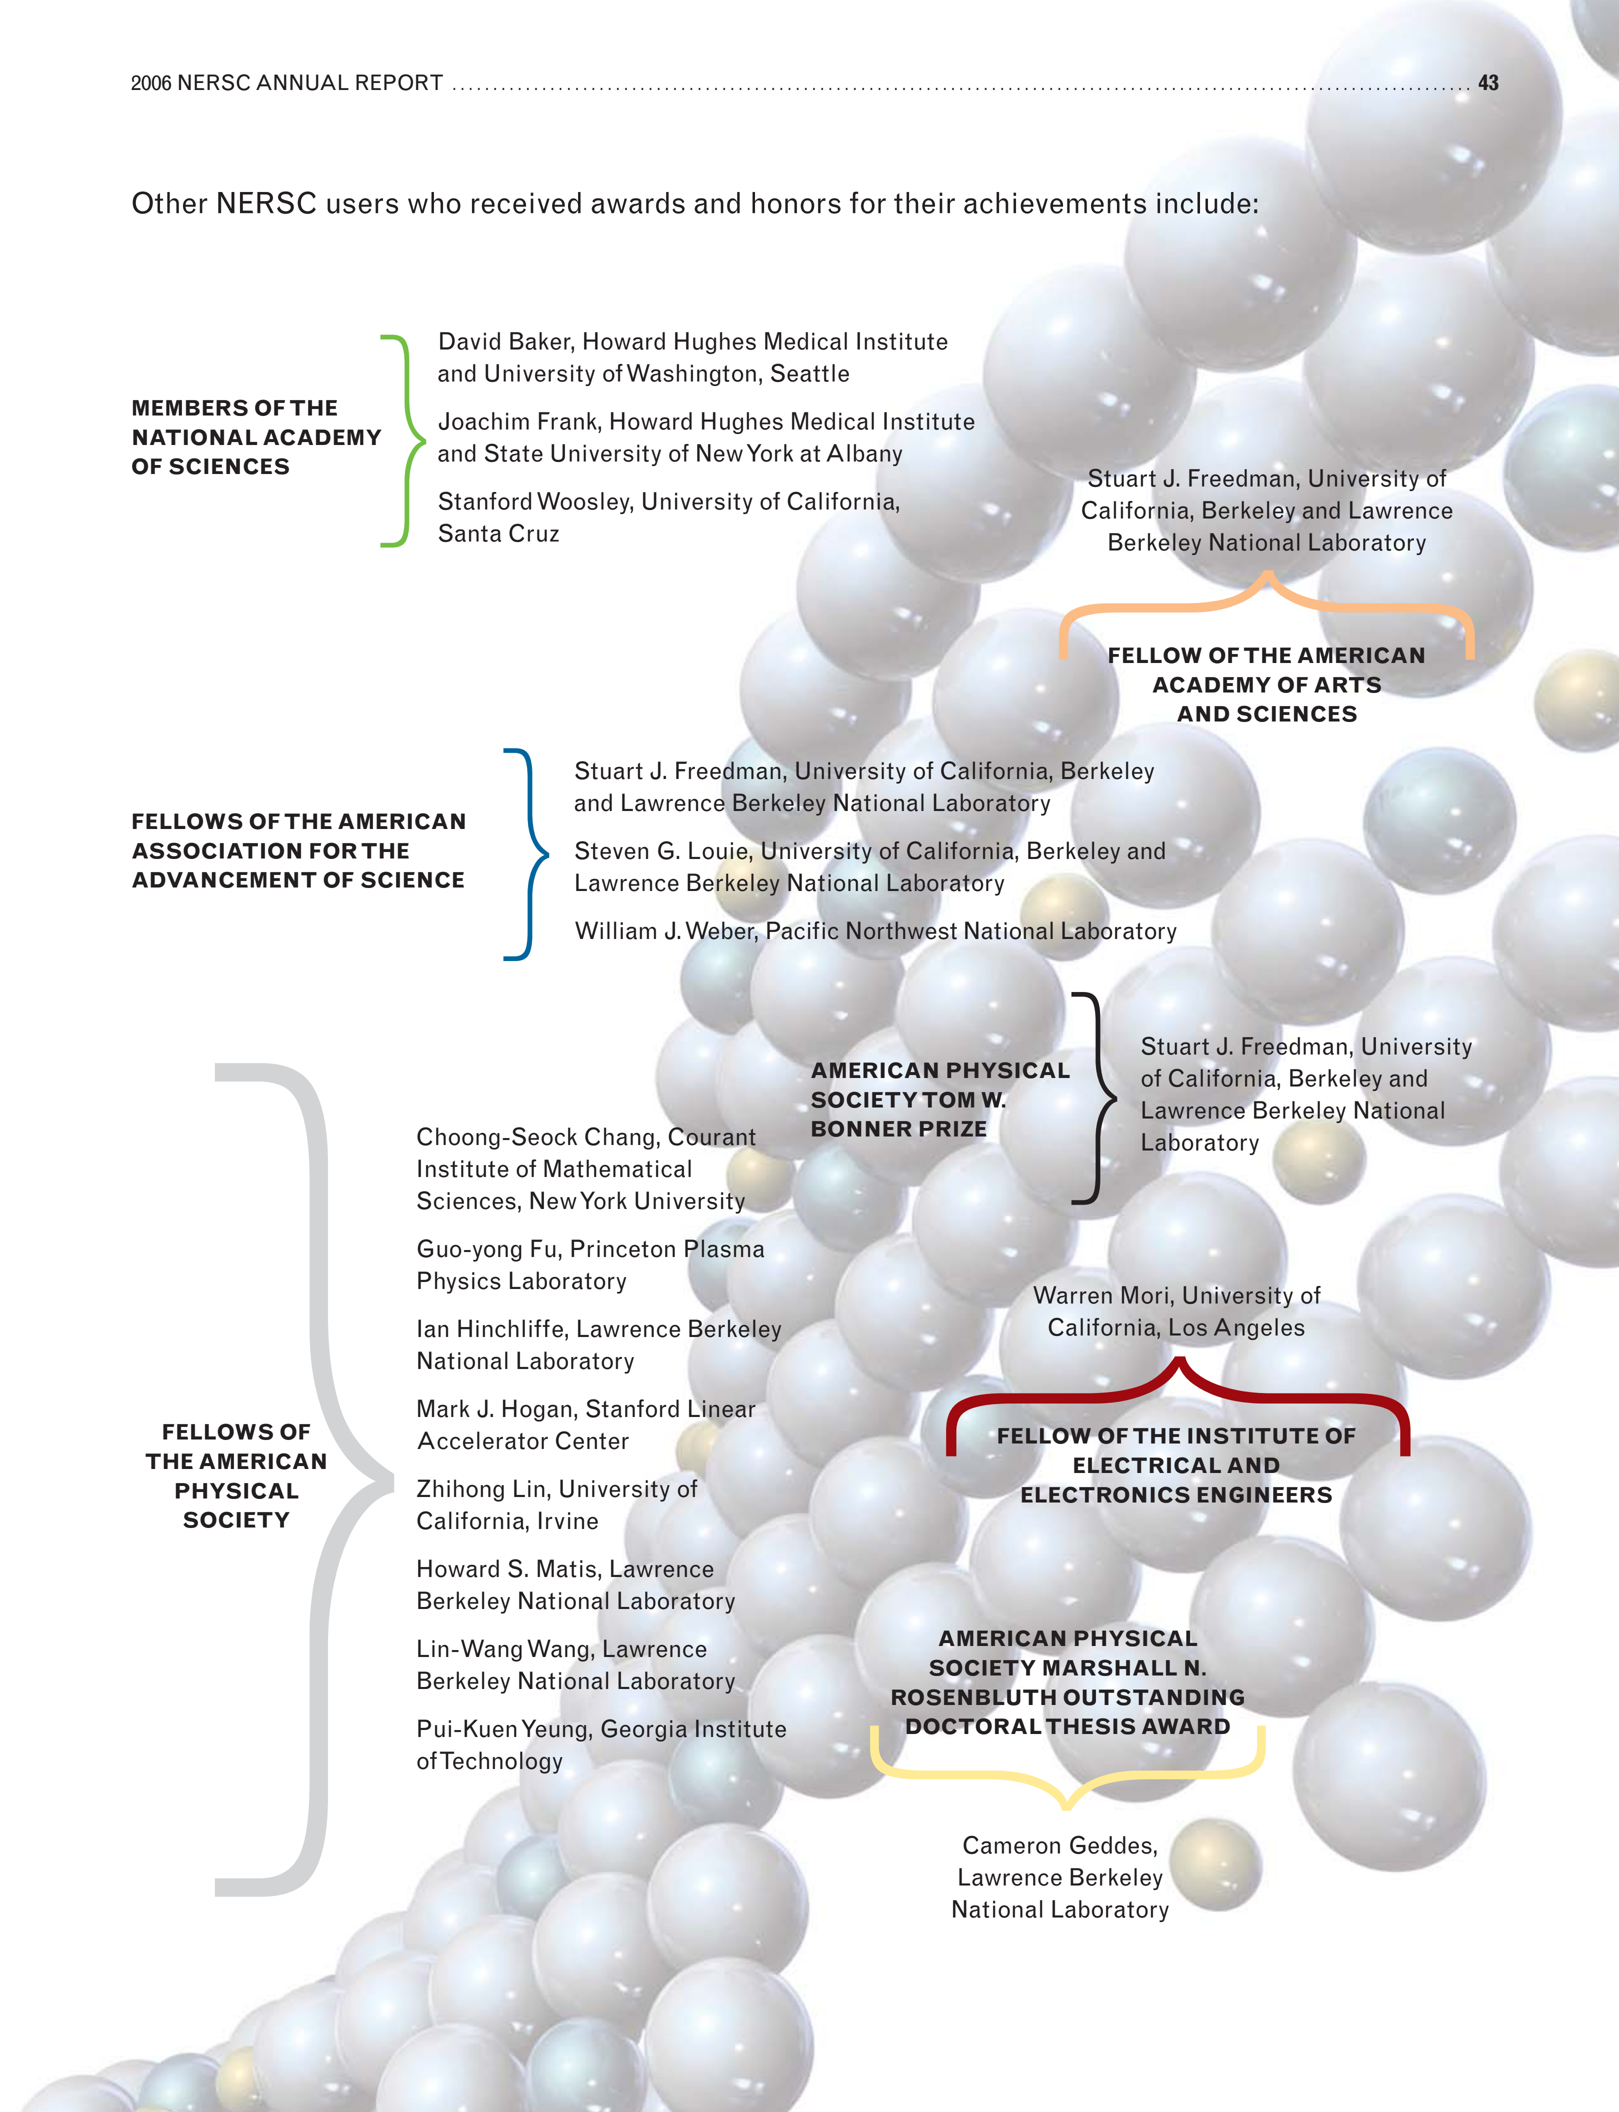

Other NERSC users who received awards and honors for their achievements include:

MEMBERS OF THE NATIONAL ACADEMY OF SCIENCES

David Baker, Howard Hughes Medical Institute and University of Washington, Seattle

Joachim Frank, Howard Hughes Medical Institute and State University of New York at Albany

Stanford Woosley, University of California, Santa Cruz

FELLOW OF THE AMERICAN ACADEMY OF ARTS AND SCIENCES

Stuart J. Freedman, University of California, Berkeley and Lawrence Berkeley National Laboratory

FELLOWS OF THE AMERICAN ASSOCIATION FOR THE ADVANCEMENT OF SCIENCE

Stuart J. Freedman, University of California, Berkeley and Lawrence Berkeley National Laboratory

Steven G. Louie, University of California, Berkeley and Lawrence Berkeley National Laboratory

William J. Weber, Pacific Northwest National Laboratory

AMERICAN PHYSICAL SOCIETY TOM W. BONNER PRIZE

Stuart J. Freedman, University of California, Berkeley and Lawrence Berkeley National Laboratory

FELLOWS OF THE AMERICAN PHYSICAL SOCIETY

Choong-Seock Chang, Courant Institute of Mathematical Sciences, New York University

Guo-yong Fu, Princeton Plasma Physics Laboratory

Ian Hinchliffe, Lawrence Berkeley National Laboratory

Mark J. Hogan, Stanford Linear Accelerator Center

Zhihong Lin, University of California, Irvine

Howard S. Matis, Lawrence Berkeley National Laboratory

Lin-Wang Wang, Lawrence Berkeley National Laboratory

Pui-Kuen Yeung, Georgia Institute of Technology

FELLOW OF THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS

Warren Mori, University of California, Los Angeles

AMERICAN PHYSICAL SOCIETY MARSHALL N. ROSENBLUTH OUTSTANDING DOCTORAL THESIS AWARD

Cameron Geddes, Lawrence Berkeley National Laboratory

The NERSC Center

HORST SIMON STEPS DOWN AS NERSC DIRECTOR

(Reprinted with permission from the December 15, 2006 issue of HPCwire)

Horst Simon, who has been director of NERSC since early 1996, announced last month that he was stepping down in order to focus his energy on the two other positions he holds at Lawrence Berkeley National Laboratory. Once a new director for NERSC is hired, Simon will concentrate on his duties as Associate Laboratory Director for Computing Sciences and Computational Research Division (CRD) Director. With the search for a new NERSC leader officially under way, Simon took some time to talk about his decision and how he sees his future.

Question: For the past 10 years, you've served as the director of the NERSC Division at Berkeley Lab. Just over a month ago, you announced that you intend to step down from this position. Can you elaborate on your reasons?

Simon: For the past three years, I've officially held three jobs—director of the NERSC and Computational Research divisions and Associate Laboratory Director (ALD) for Computing Sciences. During the past two years, the ALD position has become much more demanding. We are also initiating a new Computational Science and Engineering partnership with the UC Berkeley campus, while also planning to construct a new building to house NERSC and CRD.

I simply had to reduce my workload and get additional help to give Computing Sciences at LBNL adequate management strength. Relinquishing leadership of NERSC was the obvious choice as I've held the position for 10 years. The center is very well established, has secure funding, is being guided by a five-year strategic plan and we have very supportive program management from the Office of Advanced Scientific Computing Research in DOE's Office of Science.

The new director will have one or two years to fully learn the position, then can begin to look for new directions. It's always better to recruit in a situation of stability versus a crisis.

In short, I've been overcommitted and this is the right time to make a change.

Question: It seems that much of your identity is as the head of NERSC. What's been the reaction to your decision?

Simon: A number of the comments have expressed surprise. After that, people ask me what I'm going to do next. I'm going to stay at Berkeley Lab and build the new Computational Science and Engineering program at the Lab and on campus. And I'll be around to work with the next NERSC director. After I explain this, people say, "I see."

Question: You're only the fourth director of NERSC since the center was established 32 years ago. Looking at other large centers, there is a similar low turnover of directors. Does this surprise you?

Simon: At first, it did seem surprising. But thinking about it further, it was not really surprising. The job is a fascinating, all-engaging activity. What could be more exciting than leading a center with large national and international recognition? The technology evolves at a rapid pace and there are always interesting political ramifications. This is a field that clearly doesn't stay still. For a person who needs a daily dose of intellectual stimulation, this is one of the greatest jobs you can have.

Question: Given that such positions don't often become available, what is Berkeley Lab looking for in the new NERSC Director?

Simon: I think we've spelled it out pretty well in the position posting.... I think the challenge for the new director will be to bring both operational and scientific excellence to the job. The person should have experience managing large-scale scientific projects or facilities, and have recognition and standing in the HPC community. We also want someone with a background in computer science or computational science—a well-established track record.

Question: What attracted you to NERSC?

Simon: In the early 1990s, my dream job was to have an office in Berkeley, on the hill, running a computing center. I was working at SGI in 1995 when DOE made the decision to move NERSC from Lawrence Livermore to Lawrence Berkeley. When this happened, I could see the potential of NERSC in a more academic setting and was captivated. My reaction was "I gotta go there!" And I did.

Question: Of course, the NERSC Division Director position is only one of the three roles you currently have at LBNL. To pose a common job-interview question, where do you see yourself in five years?

Simon: I plan to still be sitting here. By then, we should have completed our new Computational Research and Theory building here at the Lab and have both NERSC and CRD moved into the facility. And we should have a large number of computational science programs built around our strengths in areas ranging from nanotechnology to astrophysics—from the smallest scale to the largest. I also hope we will have a strong partnership program with campus in five years.

In the field of HPC, I hope we will have learned how to make effective use of highly parallel systems with hundreds of thousands of processors. In the 1990s, we first began to talk about "massively parallel" systems, but those systems typically had hundreds or a few thousand processors. The new machines coming on line now are truly massively parallel. We will have to make the transition to those systems, just as we will be transitioning to running applications on multicore chips.

Of course maintaining and growing our base research programs in mathematics and computer science, as well as our SciDAC projects will continue to require my attention. With a strong partner who will lead NERSC, I expect to be in a much better position to put more attention to nurture and promote computational research in the future.

Question: Anything else you'd like to add?

Simon: When heads of other centers have stepped down, they have remained close to HPC, which I also intend to do. There is a lot of work to be done in Berkeley, with computational researchers both on campus and at the Lab. With our new partnership, we hope to create a new environment to foster computational science, and NERSC will obviously be included.

One other area I intend to focus on is energy efficiency for HPC. I believe that increasing efficiency and lowering power consumption will be one of our big challenges during the next decade. I intend to devote a lot of my future efforts to finding solutions.

SCIENCE-DRIV

Introducing and deploying the best new technologies for complete computational systems

VEN SYSTEMS

NERSC's Science-Driven Systems strategy includes 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.

CRAY PROVIDES THE NEXT MAJOR NERSC SYSTEM

On August 10, 2006, Cray Inc. and the DOE Office of Science announced that Cray had won the contract to install a next-generation supercomputer at NERSC. The systems and multi-year services contract includes delivery of a Cray XT4 supercomputer, with options for future upgrades that would quadruple the size of the system and eventually boost performance to one petaflop/s (1,000 trillion floating point operations per second) and beyond.

A successor to the massively parallel Cray XT3 supercomputer, the XT4 system installed at NERSC will be among the world's fastest general-purpose systems and will be the largest XT4 system in the world. It will deliver sustained performance of more than 16 trillion calculations per second when running a suite of diverse scientific applications at scale. The system uses thousands of AMD Opteron processors running a tuned, lightweight operating system and interfaced to Cray's unique SeaStar network.

Cray began building the new supercomputer at the manufacturing facility in late 2006 and delivered it in early 2007 (Figure 1), with completion of the installation and acceptance scheduled for the fall.

FIGURE 1. NERSC's Cray XT4 supercomputer, when complete, will deliver sustained performance of at least 16 teraflop/s.

As part of a competitive procurement process (see detailed discussion on next page), The NERSC procurement team evaluated systems from a number of vendors using the Sustained System Performance (SSP) metric. The SSP metric, developed by NERSC, measures sustained performance on a set of codes designed to accurately represent the challenging computing environment at the Center.

"While the theoretical peak speed of supercomputers may be good for bragging rights, it's not an accurate indicator of how the machine will perform when running actual research codes," said NERSC Director Horst Simon. "To better gauge how well a system will meet the needs of our 2,500 users, we developed SSP. According to this test, the new system will deliver over 16 teraflop/s on a sustained basis."

"The Cray proposal was selected because its price/performance was substantially better than other proposals we received, as determined by NERSC's comprehensive evaluation criteria of more than 40 measures," said Bill Kramer, General Manager of the NERSC Center.

The XT4 supercomputer at NERSC will consist of almost 20,000 AMD Opteron 2.6-gigahertz processor cores (19,344 compute CPUs), with two cores per socket making up one node. Each node has 4 gigabytes (4 billion bytes) of memory and a dedicated SeaStar connection to the internal network. The full system will consist of over 100 cabinets with 39 terabytes (39 trillion bytes) of aggregate memory capacity. When completely installed, the system will increase NERSC's sustained computational capability by almost a factor of 10, with an SSP of at least 16.01 teraflop/s (as a reference, Seaborg's SSP is 0.89 Tflop/s, and Bassi's SSP is 0.8 Tflop/s). The system will have a bisection bandwidth of 6.3 terabytes per second and 402 terabytes of usable disk.

In keeping with NERSC's tradition of naming supercomputers after world-class scientists, the new system will be called "Franklin" in honor of Benjamin Franklin, America's first scientist. The year 2006 was the 300th anniversary of Franklin's birth.

"Ben Franklin's scientific achievements included fundamental advances in electricity, thermodynamics, energy efficiency, material science, geophysics, climate, ocean currents, weather, population growth, medicine and health, and many other areas," said Kramer. "In the tradition of Franklin, we expect this system to make contributions to science of the same high order."

BVSS AND PERCU: A COMPREHENSIVE APPROACH TO HPC PROCUREMENT

NERSC's consistency in deploying reliable and robust high-end computing systems is due in large part to flexible procurement practices based on a process that can be summed up with the acronyms BVSS and PERCU—*Best Value Source Selection and Performance, Effectiveness, Reliability, Consistency, and Usability.*

Originally developed at Lawrence Livermore National Laboratory (LLNL) for the procurement of ASCI systems, BVSS has been used to procure all the major HPC systems installed at NERSC since the center moved to Berkeley Lab in 1996. The intent of BVSS is to reduce procurement time, reduce costs for technical evaluations, and provide an efficient and cost-effective way of conducting complex procurements to select the most advantageous offer. The flexibility of BVSS allows vendors to propose (and buyers to consider) different solutions than may have been envisioned at the outset, and allows buyers to evaluate and compare features in addition to price, focusing on the strengths and weaknesses of proposals. The end result at NERSC is usually a firm, fixed-price contract with hundreds of criteria that both NERSC and the vendor agree on.

Based on its success at NERSC, BVSS has since been adopted by Pacific Northwest National Laboratory and other organizations. And an offer to other supercomputing centers to get a firsthand look at the process by observing the procurement of NERSC-5 (resulting in the choice of the XT4) drew representatives from several National Science Foundation and Department of Defense facilities.

Within the BVSS framework, NERSC translates scientific requirements into about 50 high-level factors that reflect the attributes computational scientists want in a large system:

- Performance: How fast will a system process their work if everything is perfect?
- Effectiveness: What is the likelihood they can get the system to do their work?
- Reliability: The system is available to do work and operates correctly all the time.
- Consistency/variability: How often will the system process their work as fast as it can?
- Usability: How easy is it for them to get the system to go as fast as possible?

NERSC uses this PERCU methodology (developed by Bill Kramer as part of his Ph.D. research) to assess systems not just before purchase but throughout their life. PERCU includes the Sustained System Performance (SSP) and Effective System Performance (ESP) metrics, which NERSC uses to assure its client community and stakeholders that the systems will be highly productive and cost effective. SSP provides a quantitative assessment of sustained computer performance over time with a complex workload, while ESP is used to monitor the impact of configuration changes and software upgrades in existing systems. NERSC now has a web site for all the SSP benchmarks, to which other computer centers can download tests and report their own results (see http://www.nersc.gov/projects/SDSA/software/?benchmark=ssp).

NERSC General Manager Bill Kramer has

Bill Kramer

shared this procurement expertise by organizing sessions at the two largest HPC conferences, the 2006 International Supercomputer Conference (ISC) held in Dresden, Germany, in June, and SC2006, held in Tampa, Fla., in November. At ISC2006, Kramer and Michael Resch of the Stuttgart Supercomputing Center co-chaired a panel discussion on "Acquisition and Operation of an HPC System." The session, which drew approximately 50 attendees, also included presentations by representatives of NASA, LLNL, and the National Center for High-Performance Computing (NCHC) in Taiwan. That panel discussion led to a workshop organized by Kramer at the SC2006 conference in Tampa. The goal of this workshop was to serve as a starting point for accumulating and disseminating the shared expertise of the HPC community in assessing and acquiring HPC systems, with the expectation of creating a document of best practices for HPC system procurements.

NERSC GLOBAL FILESYSTEM MARKS FIRST FULL YEAR IN PRODUCTION

NERSC has historically been a leader in providing new systems and services to help users make the best use of the Center's computing resources. The NERSC Global Filesystem (NGF), which allows users to create and access a single file from any HPC system on the machine room floor, marked its first full year in production in 2006. NGF, which currently provides 70 terabytes (TB) of usable storage for users, ended the year at 88 percent of capacity and is scheduled to be upgraded to 140 TB in 2007. The goal of NGF is to increase scientific productivity by simplifying data management and access.

NGF does this by creating a single data file, using a single unified namespace, which can be used on any of NERSC's computing architectures. NGF's single unified namespace makes it easier for users to manage their data across multiple systems. Advantages include:

- Users no longer need to keep track of multiple copies of programs and data or copy data between NERSC systems.
- 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.
- NGF improves system security by eliminating the need for collaborators to use "group" or "world" permissions.

While the single unified namespace feature is important due to the heterogeneous computing environment of NERSC, NGF also proved its worth when Seaborg was temporarily taken out of service in mid-2006 due to security concerns. Those users who had asked to have project directories set up for their research were able to easily move their jobs to the other systems. In all, 77 projects have requested to use NGF, and these typically represent large users. For example, just 14 projects account for 70 percent of the storage used in NGF and of these, four projects account for 50 percent of the storage. The largest users are in the field of astrophysics and include the Planck satellite data project, the cosmic microwave background radiation project and the Nearby Supernova Factory. The groups have found that NGF helps them use the best computing system depending on the nature of the task and also provides flexibility for mapping out workflows to different systems.

One of the challenges facing both large and small research groups is consistency of the data being analyzed. NGF helps ensure that all members are using the same data file, rather than inadvertently working on different versions of a file. This also helps users keep track of the location of each file and means that files do not have to be moved from system to system as a job progresses.

NGF provides users with immediate access to data as soon as it is created. This instant availability of data, such as files generated on the visualization server DaVinci, enables users to computationally "steer" jobs running on NERSC systems. Future plans call for possible use of NGF over wide area networks from other computing centers and to have home and scratch file systems globally accessible.

To protect user data, NGF files are backed up to NERSC's HPSS biweekly, and NGF is scheduled to become fully integrated with HPSS in the future.

As deployed at NERSC, NGF is expected to have a long life of 10 to 15 years or more. It is expected that during this time the file system will change and evolve, just as the systems in the Center that are connected to it will. It is also expected that the user data will have long-term persistence in the file system, ranging from months and years up to the deployed life of the file system, at the discretion of the users.

INTEGRATING NERSC'S STORAGE AND FILE SYSTEMS

Over the past year, NERSC's Mass Storage Group, with Jason Hick as its new group leader, has improved the storage system's reliability, performance, and availability while managing an unprecedented amount of data. The storage systems, in fact, reached 3 petabytes of stored data in November 2006.

The Mass Storage Group participates in the GPFS-HPSS Integration (GHI) collaboration with IBM. The collaboration has focused on designing a new interface that allows HPSS to be more fully integrated

Jason Hick

with new GPFS software features. In the new version of GHI, HPSS will transparently serve as an additional hierarchy of storage so data will move between GPFS and HPSS as it is needed. The new version also explicitly performs highly parallel backups of GPFS.

Damian Hazen of the Mass Storage Group developed and demonstrated a proof of concept for the new GHI design at SC06 in November. The demonstration was hailed as a great success and served as the foundation for the new software, which is expected to be available in 2007–08.

Once in production, the software will provide a unified global name space to users by allowing users to access their files and data through any GPFS file system, including NGF. As time goes on, the data will move automatically to an HPSS system while leaving a file stub in GPFS. The user can still access the data through GPFS—the software will handle automated data movement between HPSS and GPFS when necessary.

"As we prepare to enter the petascale computing era, data is sure to increase, and integration of the storage and file system at NERSC is one approach to easing the data management challenges that users are sure to face," said Hick.

In 2007, the group is focusing on several projects, such as upgrading to a new version of HPSS, improving storage system bandwidth and capacity, and providing a Globus-based GridFTP server capable of accessing HPSS directly.

The group continues to prepare for the upgrade to HPSS version 6.2, which is expected to occur toward the end of 2007. The upgrade will remove HPSS's dependence on the legacy Distributed Computing Environment (DCE) software.

Deploying a new tape technology that will expand capacity and handle increased demand is also under way. The new technology will more than double the previous tape capacity and bandwidth capabilities, holding 500 gigabytes of uncompressed data in one cartridge, or nearly 1 terabyte in a compressed format.

The Globus-based GridFTP server is expected to be available after the new version of HPSS is in place. This GridFTP server will provide HPSS access to GridFTP clients. GridFTP is gaining support in scientific community as the data movement interface of choice. Providing a GridFTP server that accesses HPSS directly will provide a reliable and high performance wide area network transfer option for HPSS data.

OSF POWER SUPPLY IS UPGRADED

While NERSC systems staff members implement new approaches to improve the usability and reliability of computing resources, the Center also upgraded its electrical infrastructure in 2006. To accomplish this, a planned power outage was carried out during the week of October 30 at the Oakland Scientific Facility (OSF). The outage allowed the NERSC computer room to be safely upgraded to accommodate a new uninterruptible power supply (UPS) and future computing systems, including Franklin, NERSC's new Cray supercomputer. Several carefully timed email notices during the previous month had informed all NERSC users about the outage.

The electrical substations in the OSF basement could deliver up to 6 megawatts (MW) of power, but only 2 MW were actually used in the machine room. However, NERSC needs 4 MW to power the increased computing capability and cooling requirements of Franklin and future machines.

To meet these needs, Pacific Gas and Electric Company (PG&E) upgraded its connection to the building, and connected new 480V feeds between the basement and the machine room to deliver the increased power. The chilled water piping under the machine room floor was also rearranged to improve the air flow, since each of Franklin's 102 racks will need 2300 cubic feet of cooled air per minute.

In February 2007, NERSC completed the power upgrade by installing its first uninterruptible power supply (UPS) to protect critical data in the NERSC Global Filesystem (NGF) and HPSS. With the UPS in operation, if an unscheduled power outage does happen, the UPS will allow a graceful shutdown of NERSC's critical storage disks and databases. That added margin of safety will benefit NERSC staff and users with increasing reliability and decreasing the amount of time required to recover from power failures.

MEETING SECURITY CHALLENGES AND SHARING EXPERTISE

NERSC computer security efforts aim to protect NERSC systems and users' intellectual property from unauthorized access or modification. This a formidable challenge, considering that NERSC resources are accessed remotely by more than 2500 scientists via a wide range of non-dedicated networks, and that government agencies and supercomputer centers are attractive targets for hackers. Security measures in an open science environment must strike a delicate balance between permitting legitimate researchers unencumbered access to resources and preventing illegitimate adversaries from compromising those same resources.

The biggest challenge of the year came on October 5, 2006, when NERSC's security team discovered that a security incident had the potential of compromising some users' passwords, SSH keys, and Grid certificates. The security response team quickly disabled all the affected accounts, notified the users by email, and took Seaborg offline for several days while they performed a comprehensive security check and took remedial steps to assure the system was not at risk. NERSC staff worked around the clock and from two continents to restore the system to service in one-third the expected time.

Apparently no actual damage was done to any NERSC systems or user files, but as an extra precaution, all passwords that had been used at NERSC for the previous three years were disabled. Users were emailed instructions on how to change their passwords and regenerate their SSH keys and Grid certificates, and they were advised to do the same on any systems outside of NERSC that may have been exposed to this security threat. NERSC staff generated and implemented a long list of action items for restoration of full service as well as long-term security improvements.

A week later, in an email from Washington to all NERSC staff, Bill Kramer forwarded the compliments of Office of Science officials on a job well done. "While the goal is to avoid such a thing, handling it well when it happens is another mark of excellence," Kramer wrote.

Computer security is often compared to an arms race, with each side constantly innovating to meet the opponents' challenges. Active participation in the cybersecurity community is one way that NERSC's security team stays up to date. For example, at the SC06 conference, a full-day tutorial on "Computing Protection in Open HPC Environments" was presented by Steve Lau, formerly of NERSC and now at the University of California, San Francisco; Bill Kramer and Scott Campbell of NERSC; and Brian Tierney of Berkeley Lab's Distributed Systems Department.

NERSC staff also contributed to a workshop titled "DOE Cybersecurity R&D Challenges for Open Science: Developing a Roadmap and Vision," held January 24–26, 2007, in Washington, D.C. The goal of the workshop was to identify the research needs and opportunities associated with cybersecurity for open science Brent Draney, Scott Campbell, and Howard Walter contributed a white paper on "NERSC Cyber Security Challenges that Require DOE Development and Support." Draney and Kramer participated in the workshop's deliberations, which were summarized in a report to DOE management.

CRAY XT4

SCIENCE-DRIVEN SERVICES

Enabling a broad range of scientists to use NERSC systems productively in their research

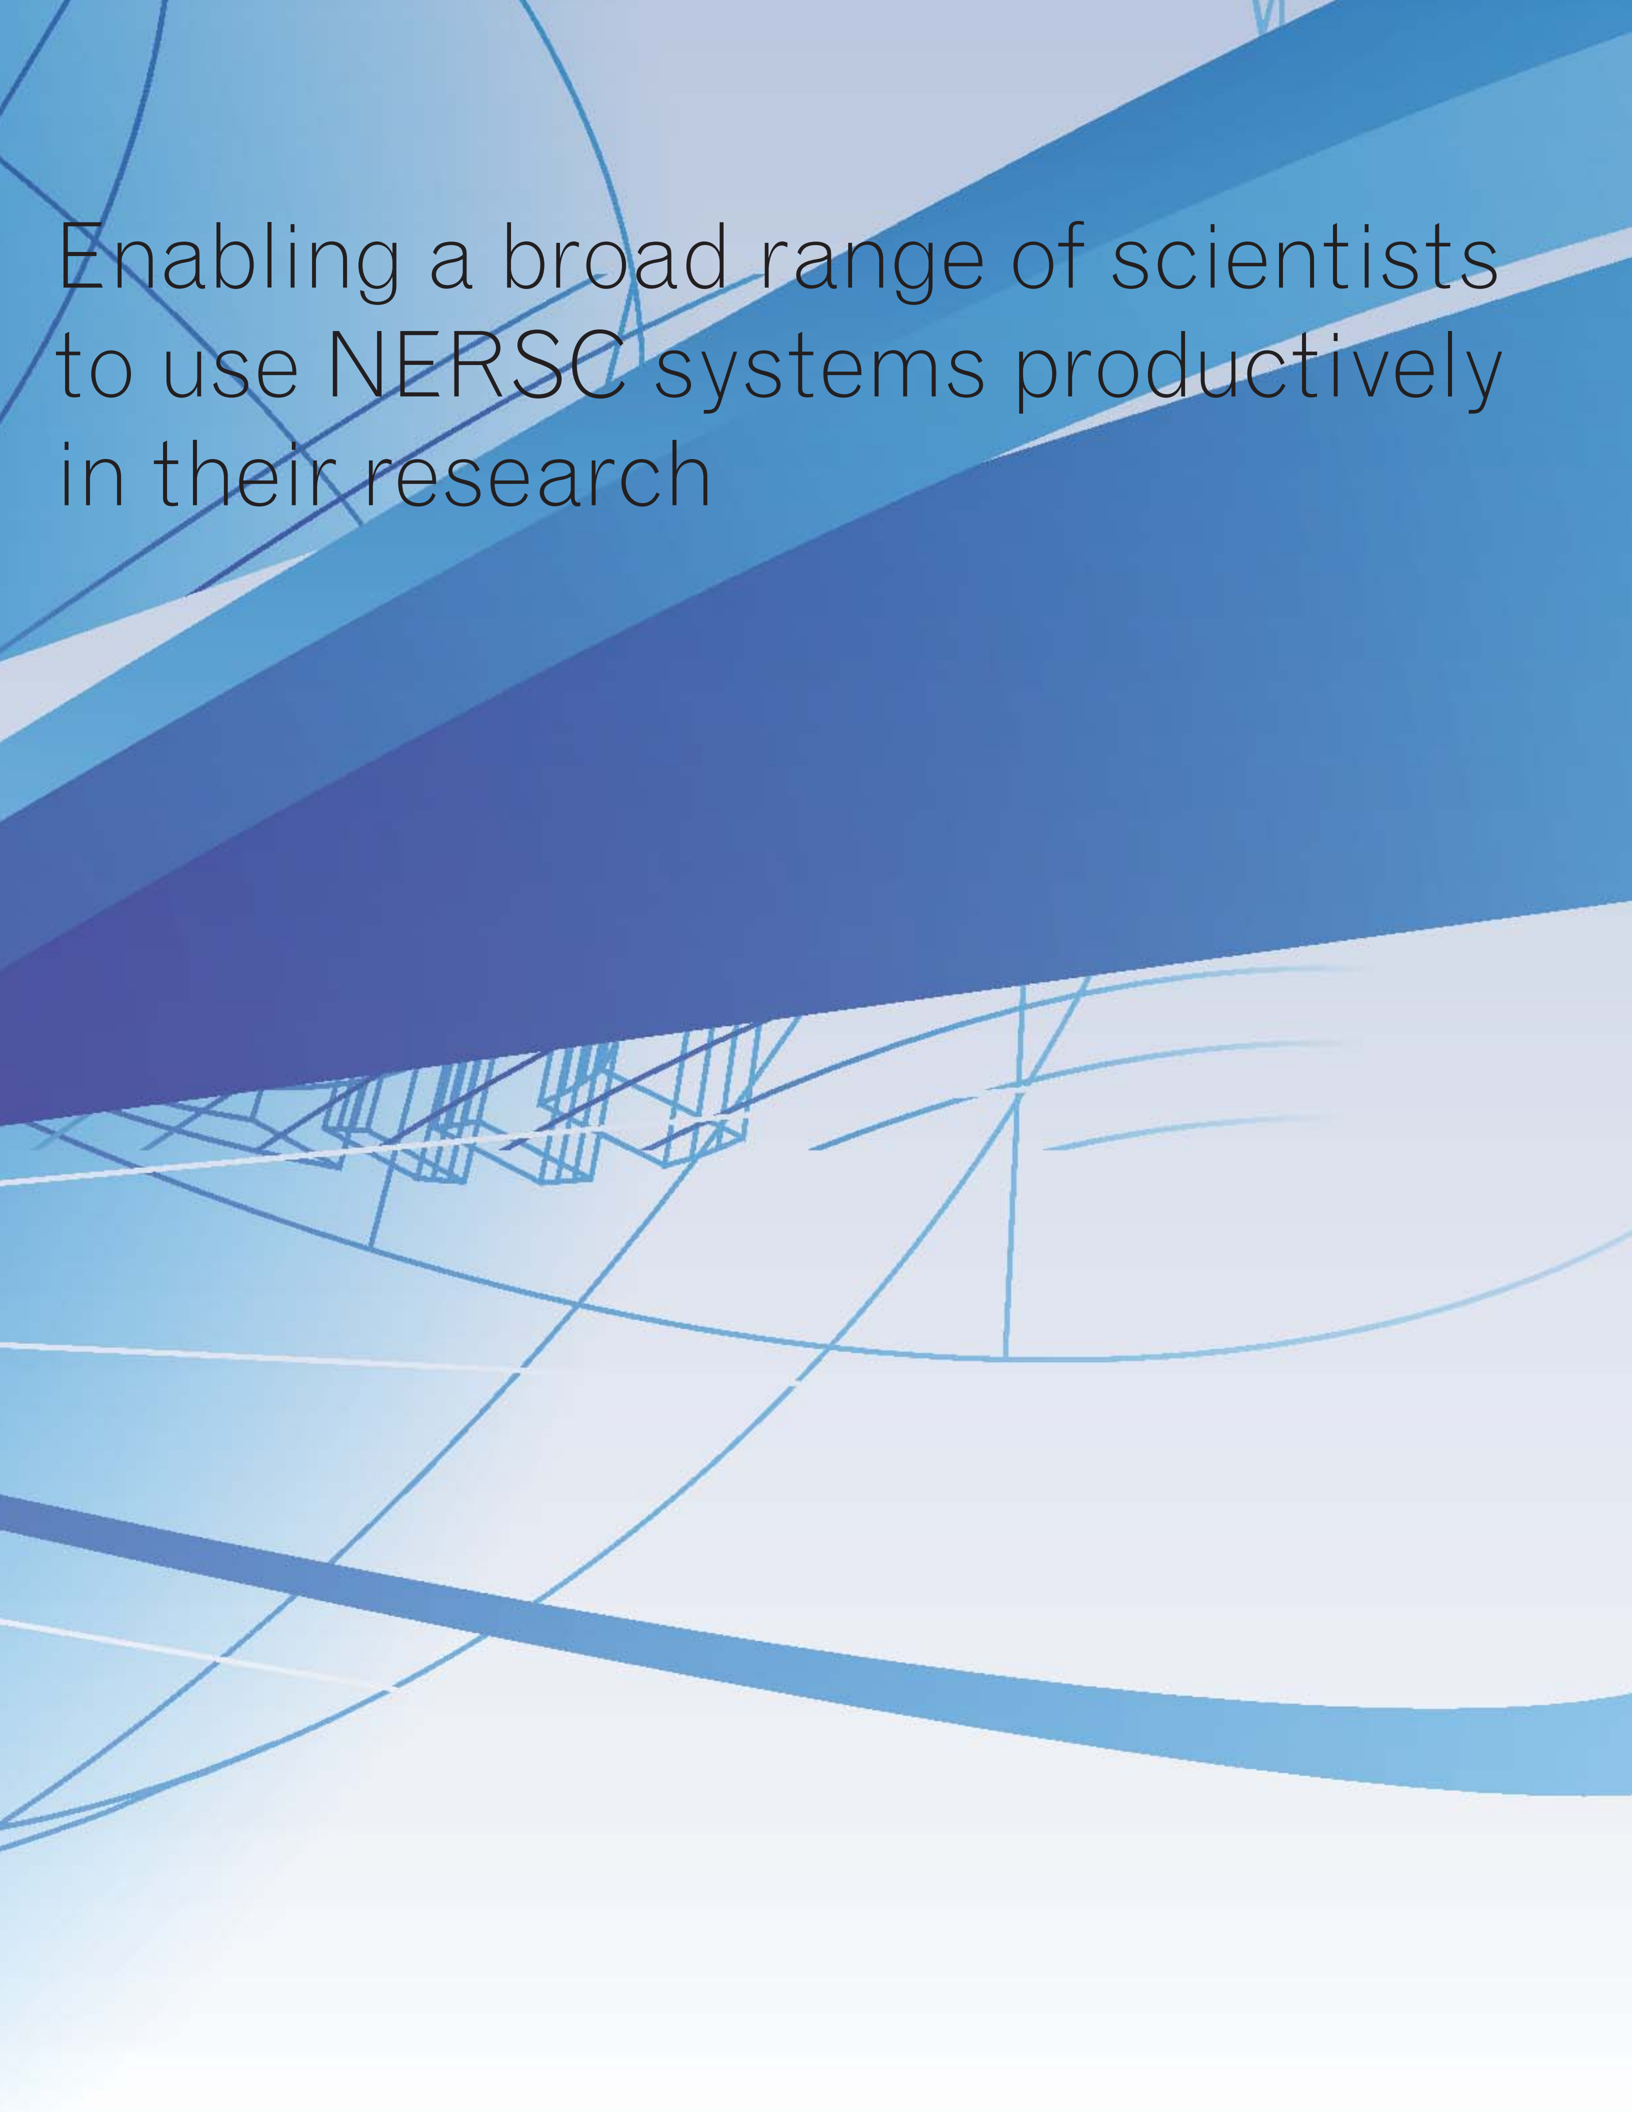

NERSC's Science-Driven Services strategy includes the entire range of support activities, from high-quality operations and user services to direct scientific support, that enable a broad range of scientists to effectively use NERSC systems in their research. NERSC concentrates on resources needed to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

NERSC MARKS FOURTH YEAR OF INCITE APPLICATIONS

When DOE Under Secretary of Science Dr. Raymond Orbach launched the Innovative and Novel Impact on Theory and Experiment (INCITE) program in 2003, NERSC was the only DOE HPC center able to support the high-impact computational science projects selected to participate.

The INCITE program seeks computationally intensive, large-scale research projects, with no requirement of current DOE sponsorship, that can make high-impact scientific advances through the use of a substantial allocation of computer time and data storage resources. From the three projects selected for NERSC allocations in each of the years 2004 and 2005, INCITE in 2007 comprises 45 projects with a total allocation of 95 million processor-hours at four Office of Science computing centers, including NERSC.

Under the 2007 INCITE awards, seven research projects will receive nearly 9 million processor hours at NERSC. The projects range from studying the behavior of a supernova to designing more energy-efficient cars. Application areas include turbulence, combustion, fusion energy, accelerator design, climate change, chemistry, and physics.

One of the projects receiving a 2007 NERSC allocation was a renewal of the 2006 INCITE-supported investigation into the chemistry of organolithium compounds led by Professor Lawrence Pratt of Fisk University in Tennessee (see page 20). According to Pratt, before receiving his INCITE award, most of his research was carried out using PCs.

"This severely limited the job throughput and in some cases even limited the kinds of calculations that could be performed," Pratt said. "The award has more than doubled my research productivity, resulting in two publications for 2006 and three papers accepted for publication in 2007 to date. Several more projects are nearing completion in 2007."

Indeed, INCITE recipients say the allocations, coupled with effective services from NERSC staff, have enabled them to accelerate their work. Donald Lamb, from the University of Chicago, for example, was able to announce a breakthrough in March 2007 thanks partly to the NERSC staffers' effort to make it possible to run large computing jobs under a tight deadline in January. Lamb's research demonstrated for the first time the ability to naturally detonate the explosion of a white dwarf star in a three-dimensional simulation. The modeling of this Type Ia supernova also confirmed previous beliefs that the white dwarf star detonates in a supersonic process resembling diesel-engine combustion.

Other INCITE researchers also expressed their gratitude.

"NERSC staff worked with us extensively to get our runs done," said Berkeley Lab fusion researcher Cameron Geddes, whose 2006 INCITE project created detailed three-dimensional models of laser-driven wakefield particle accelerators (see page 28). "Consultants and managers met with us to provide needed disk space and queue organization and allocations to get large runs done, and this was critical to the results."

Robert Harkness, a scientist at the San Diego Supercomputer Center and a co-investigator for the 2006 INCITE project on characterizing the shape, matter-energy contents, and expansion history of the universe, said he was impressed with Seaborg and HPSS. "Seaborg and HPSS have been quite simply outstanding in the high degree of reliability," Harkness said. "I am particularly impressed with the HPSS—I've written 80 to 100 terabytes without any trouble."

MAKING SHARED RESOURCES EASY TO USE

Staff in the Open Software and Programming Group are bringing new services to NERSC users through grid technologies provided by the Open Science Grid.

The Open Science Grid (OSG), funded by the DOE Office of Science and the National Science Foundation, is a distributed computing infrastructure for scientific research built by a nationwide community of universities and national laboratories. In January 2006, NERSC General Manager Bill Kramer was elected the first chair of the OSG Council, made up of representatives from each member institution or partner organization.

The OSG software stack provides distributed mechanisms for moving data and submitting jobs to computing resources on the OSG grid. Several of the frameworks have long been in use in the high energy physics community and are seeing increased use in new science areas.

“The standard model for computing at most parallel computing centers is to provide batch queues where users can log in and submit specific jobs,” said David Skinner, who became leader of the Open Software and Programming Group in May 2006. “More and more science teams expect to be able to streamline their computing needs into a workflow that can be accessed and operated in a distributed way.”

The OSG received an allocation of time at NERSC for the purpose of extending its user base and the range of research areas and applications using grid technologies on parallel systems. Discussions with the DOE Joint Genome Institute and with nanoHUB, a web-based resource for nanoscience and technology, have extended the OSG’s application base to bioinformatics and materials science.

David Skinner

Providing excellent services to NERSC’s research customers depends on both having useful capabilities and also configuring and organizing them in ways that make them simple for researchers to access. Grid technologies can be daunting in their complexity, and easing the adoption of grid services into new science areas is an active area of work in the Open Software and Programming Group.

So far, OSG grid services have been deployed on PDSF, DaVinci, Jacquard, and HPSS. These NERSC resources provide OSG with new capabilities in parallelism and filesystem performance. “Not all data endpoints are equal in terms of performance,” Skinner said. “We stood up these services first on DaVinci because it provides a much faster connection to the NERSC Global Filesystem. DaVinci is also a natural location for data-intensive analytics activities. By providing the full suite of Open Science Grid software and services, we hope to make DaVinci much more useful to our users.” The OSG software stack will be deployed on other NERSC systems in 2007.

DEPOSITING A SLICE OF THE UNIVERSE

Data from the world’s largest radio telescope, which scans the sky for signs of pulsars, black holes, and extraterrestrial civilizations, has found a new home in NERSC’s vast storage system.

Led by Dan Werthimer, a scientist at the University of California-Berkeley’s Space Sciences Laboratory, the “Berkeley High Resolution Neutral Hydrogen Sky Survey” recently began feeding hundreds of gigabytes of data daily to NERSC’s HPSS.

Collected by the Arecibo Observatory in Puerto Rico (Figure 2), the radio signal data will provide valuable information to astrophysicists, cosmologists, and other researchers who aim to solve some of the most intriguing mysteries about the universe. The Arecibo Observatory, a 305-meter dish, is operated by Cornell University under an agreement with the National Science Foundation.

FIGURE 2. Serving up the cosmos, the Arecibo dish in Puerto Rico detects radio signals emitted from space and helps researchers understand stars, dark energy, and everything in between. (Image courtesy of the NAIC Arecibo Observatory, a facility of the NSF.)

“Before we received assistance from NERSC, our sky survey data was stored on over 1,000 tapes. So it was very difficult to access,” Werthimer said. “NERSC will allow us to make our data accessible to the scientific community for a variety of research projects.”

For Werthimer and his research team, the data will help advance three projects: SETI@home, AstroPulse, and hydrogen mapping. SETI@home, launched in 1995, seeks evidence of extraterrestrial civilizations by searching for certain types of radio signals that, because of their narrow bandwidths, do not occur naturally.

The SETI (Search for Extraterrestrial Intelligence) project has created a global distributed computing network spanning 226 countries that invites anyone with an Internet-connected computer to join. So far, it has attracted more than 5.2 million participants, whose personal computers help to crunch the data from Arecibo.

AstroPulse, on the other hand, will look for dispersed radio pulses that are shorter than the ones sought by SETI@home. The microsecond pulses could come from rapidly spinning pulsars, black holes, or extraterrestrial intelligence.

Both SETI@home and AstroPulse run on BOINC, an open-source software. BOINC, developed by David Anderson and his team at the Space Sciences Lab, makes it easy to develop and manage distributed computing projects. Projects that take advantage of BOINC include Einstein@home at the University of Wisconsin and LHC@home at CERN.

The third project maps the galactic distribution of hydrogen. The high-resolution map will provide a wealth of three-dimensional information such as the density and temperature of interstellar objects, enabling researchers to better understand the structure and dynamics of our galaxy.

"We recognize the importance of Dan's projects to advance a wide-range of space research," said Francesca Verdier, Associate General Manager for Science Driven Services at NERSC. "The data from Arecibo will be put to good use by the scientific community worldwide."

Werthimer plans to store 200 terabytes of data over two years at HPSS. HPSS, which was put in place in 1998, currently has a theoretical capacity of 30 petabytes. It is capable of accepting data at 450 megabytes per second.

SURVEY GAUGES USER SATISFACTION

NERSC users gave kudos to HPSS, Jacquard uptime, and Bassi Fortran compilers in the results of a survey to solicit feedback about their experiences using NERSC resources in 2006.

The annual survey enables NERSC staff to gauge their successes and make improvements. The survey asked researchers to rate their overall satisfaction with the hardware, software, and services. It also posed more specific questions about these tools and services, and invited users to share their thoughts. On a scale of 1 to 7, with 7 being "very satisfied," survey respondents gave an average of 6.3 for the question about their overall satisfaction with NERSC, an improvement from 2005 and 2004.

"We take feedback from our users very seriously, whether it is in the form of day-to-day interactions, through the NERSC User Group meetings, or via our annual survey," said Bill Kramer, General Manager of NERSC. "We work very hard to bring any low scores up and to keep the overall satisfaction high in all areas. I want to thank all the users who took the time to provide this valuable feedback."

About 13 percent of the active NERSC users, or 256, took the online survey. The respondents represented all six programs within the DOE Office of Science.

Areas with the highest user satisfaction in 2006 included account and consulting services, DaVinci C/C++ compilers, and network performance within the NERSC center. The largest increases in satisfaction over the 2005 survey were for the Jacquard Linux cluster, Seaborg batch wait times and queue structure, NERSC's available computing hardware, and the NERSC Information Management (NIM) system.

Despite the improvements, however, Seaborg batch wait times received a low average score compared with assessments about other systems and services. Other areas that received lower scores included PDSF disk storage, interactive services and performance tools, analytics facilities, and Bassi and Seaborg visualization software.

The survey gave users an opportunity to pen their own thoughts about using NERSC. The question "What does NERSC do well?" drew answers from 113 users. Among them, 87 stated that NERSC gave them access to powerful computing resources without which they could not do their science; 47 mentioned excellent support services and NERSC's responsive staff; 27 highlighted good software support or an easy-to-use user environment; and 24 pointed to hardware stability and reliability.

"NERSC runs a reliable computing service with good documentation of resources," wrote one user. "I especially like the way they have been able to strike a good balance between the sometimes conflicting goals of being at the 'cutting edge' while maintaining a high degree of uptime and reliable access to their computers."

In previous years, queue turnaround and job scheduling issues were areas of the greatest concerns. NERSC has made many efforts to acquire new hardware, to implement equitable queuing policies across the NERSC machines, and to address queue turnaround times by adjusting the duty cycle of the systems.

These efforts have clearly paid off. In 2004, 45 users reported dissatisfaction with queue turnaround times. In 2005 this number dropped to 24. In 2006 only five users made such comments.

Improvements to Jacquard's computing infrastructure and the deployment of the NERSC Global File System also addressed shortcomings voiced by users in previous years. Although job scheduling remained a concern in 2006, NERSC users pointed to new challenges for NERSC staff. Researchers said they would like more compute cycles and software fixes.

The complete survey results can be found at http://www.nersc.gov/news/survey/2006/.

SCIENCE-DRIVEN ANALYTICS

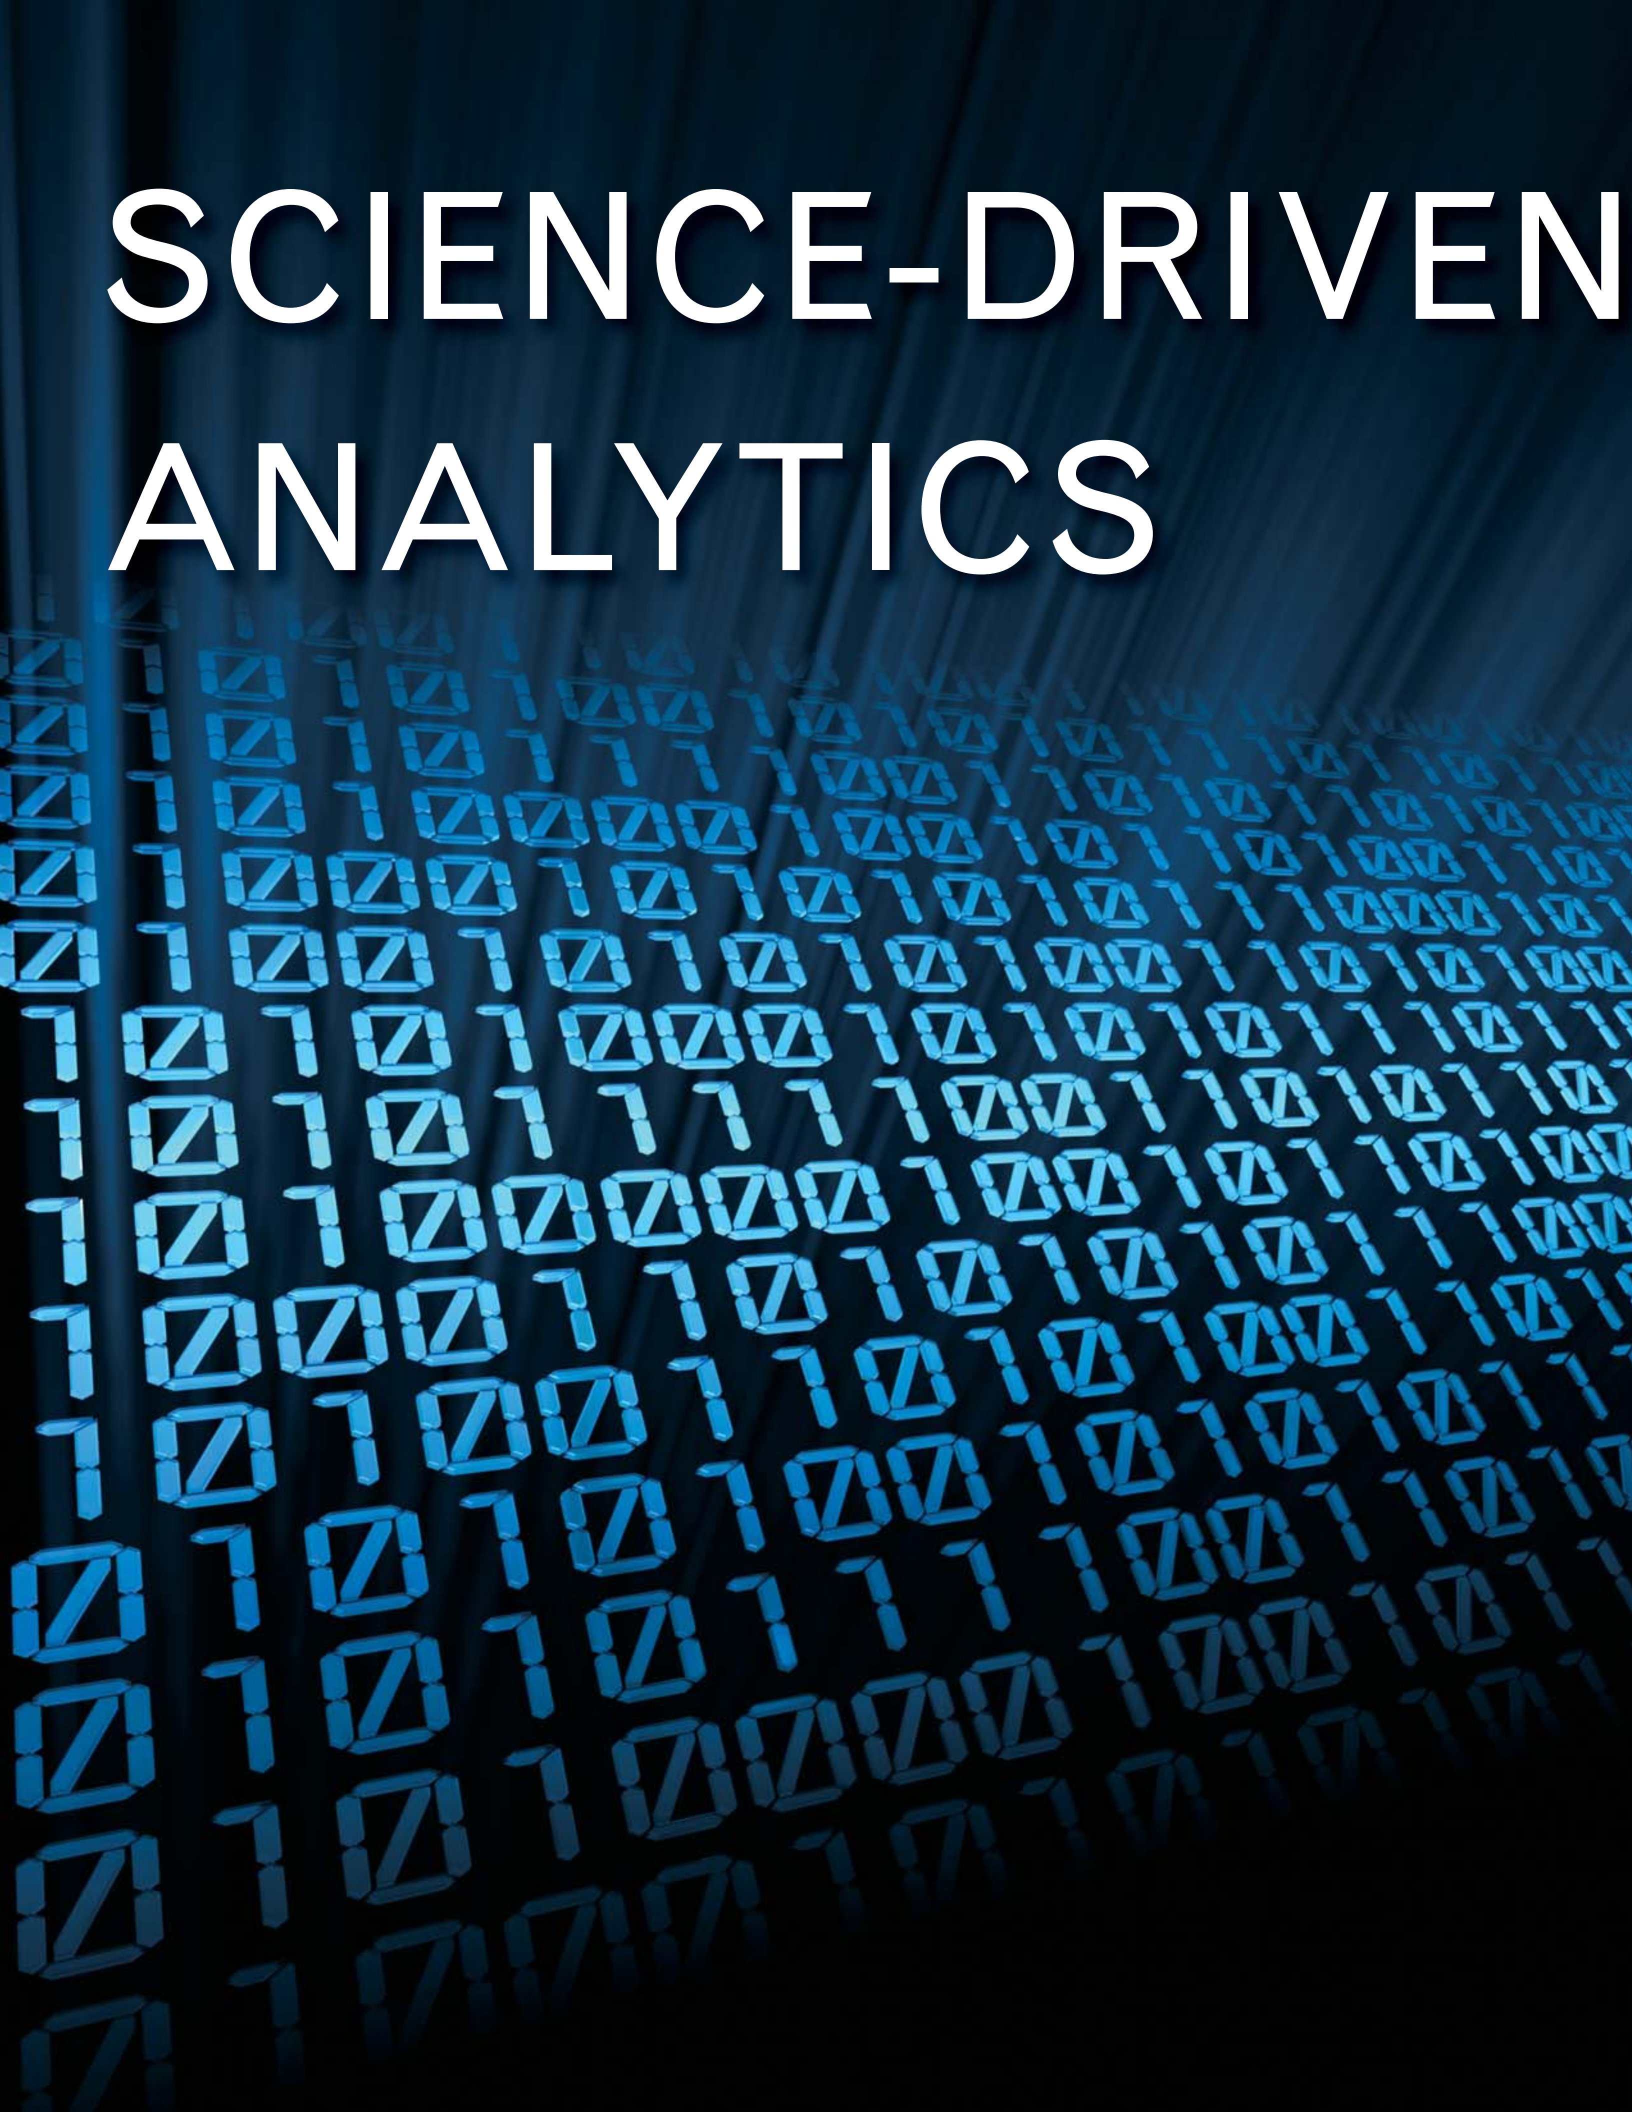

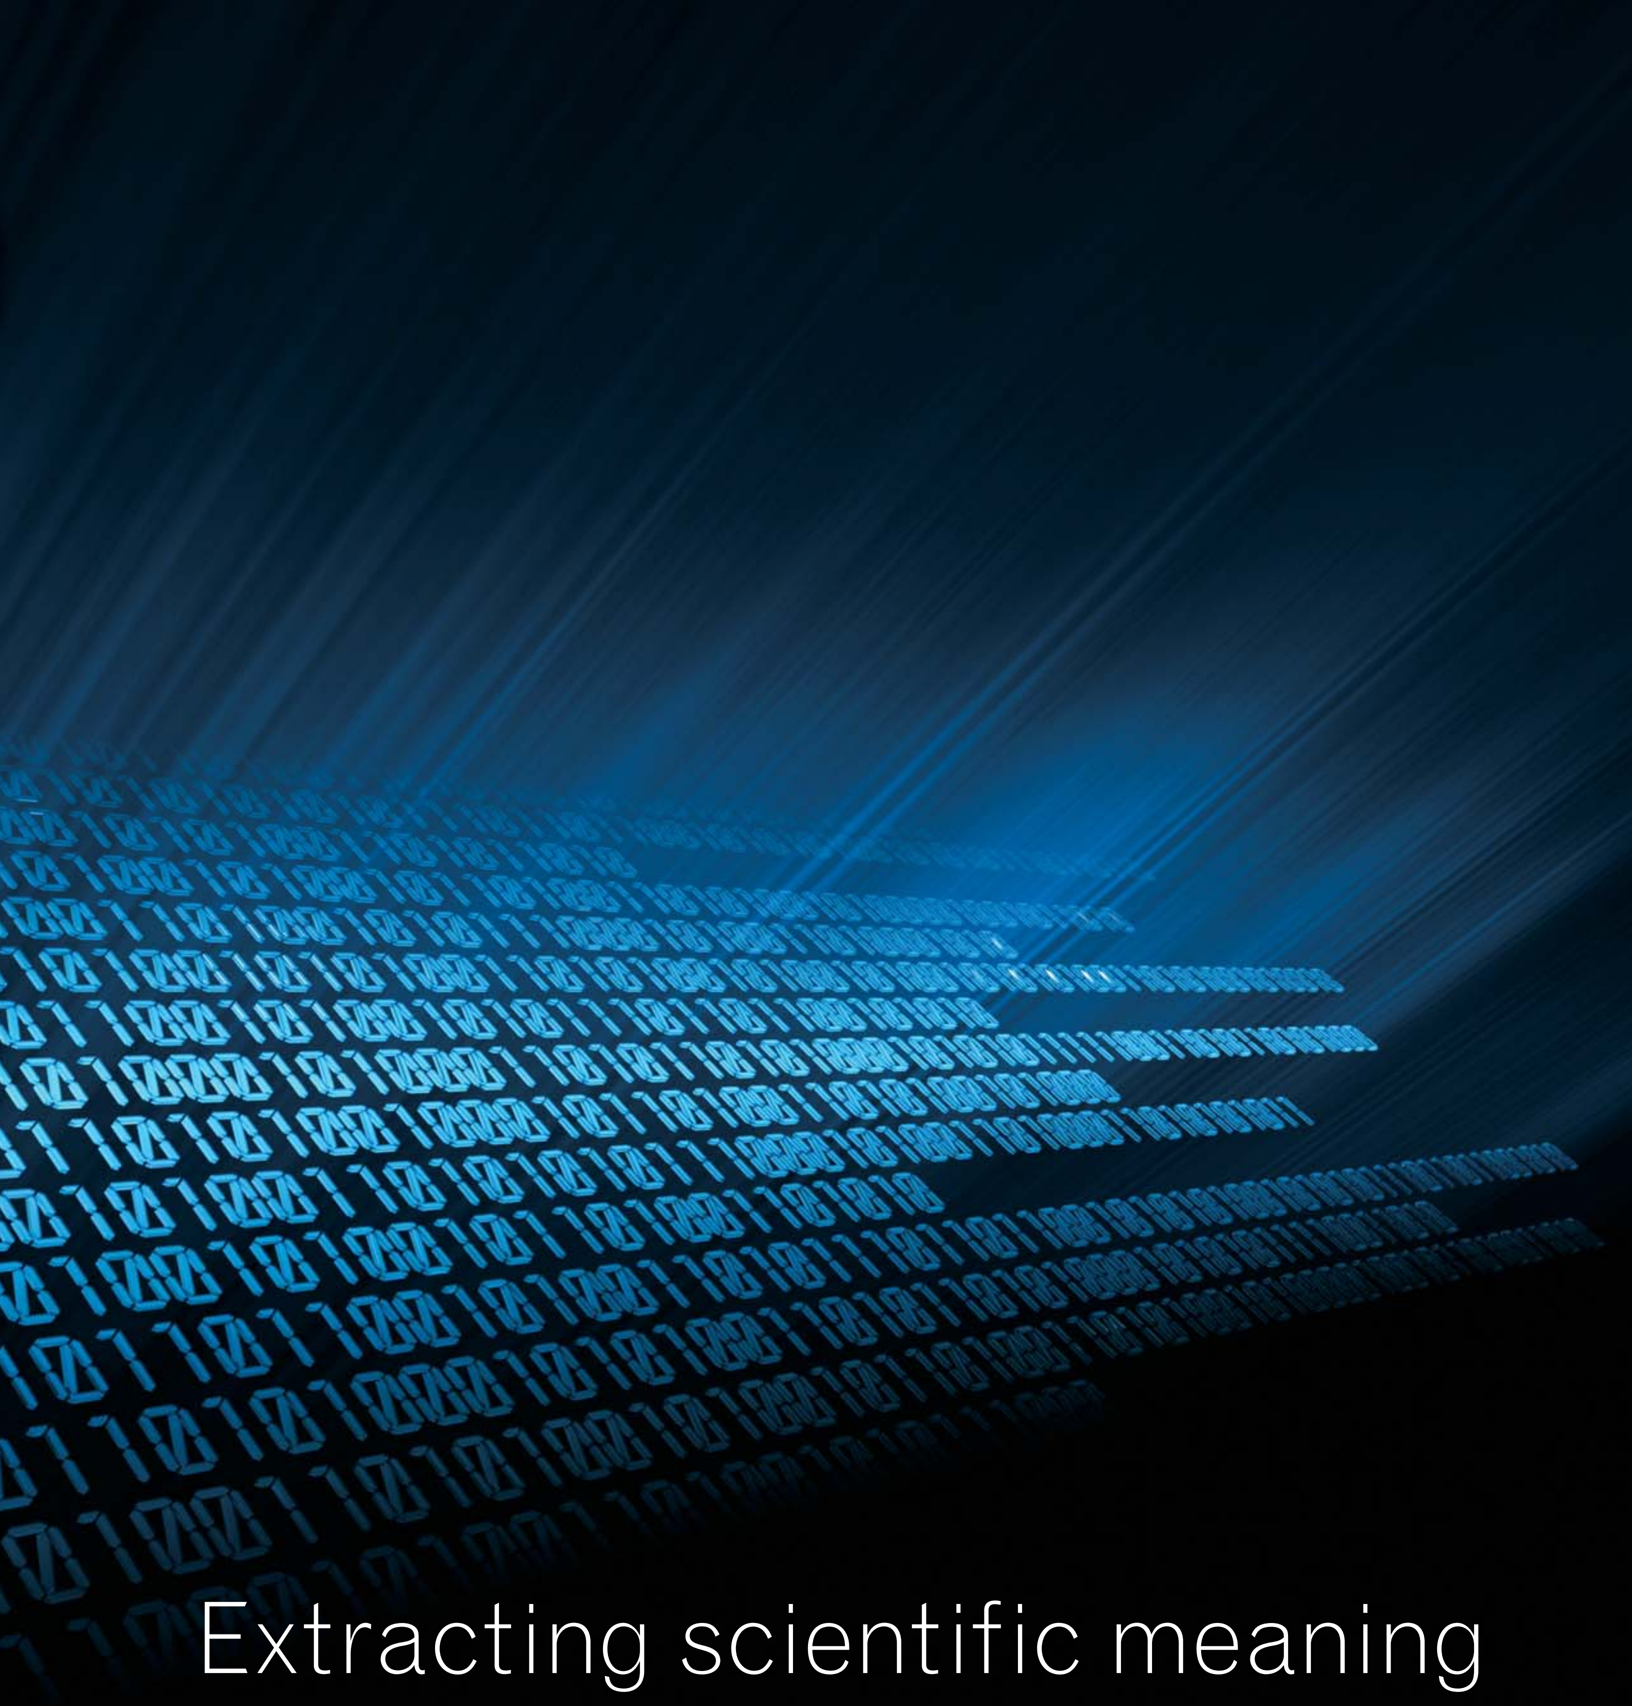

Extracting scientific meaning from massive quantities of data

NERSC's Science-Driven Analytics program provides 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.

BUILDING THE ANALYTICS INFRASTRUCTURE

In its first full year of activity, NERSC's Analytics Team has laid the foundation for an analytics infrastructure that combines software, development and application of analytics technologies, working with NERSC users, and hardware (DaVinci, an SGI Altix 350, which went into full production use in 2005). The key technologies that contribute to the analytics program are data management, data analysis and data mining; visual data exploration; and workflow management. The mission of the Analytics Team is to deploy and adapt these technologies—by combining them or extending them—to help NERSC users spend more time doing research and less time managing data and wrangling analytics software.

Scientific data management (SDM) refers to storage and retrieval of scientific data from various storage sources such as main memory, disk, and tape. SDM covers integration of data formats; data description, organization, and metadata management; efficient indexing and querying; and file transfer, remote access, and distributed data management across networks. Managing experimental data and files of simulation output, as well as converting files from one format to another, continue to be time-consuming aspects of high performance computing. Several groups in NERSC, as well as the Analytics Team, are working with LBNL's Scientific Data Management Center to deploy the Storage Research Manager (SRM) software NERSC-wide in order to provide transparent distributed data management. SRM provides a fault-tolerant mechanism for transferring files from one location to another, as well as uniform access to heterogeneous storage (e.g., disk, tape). In addition, members of the Analytics Team have been working with collaborators at the Paul Scherrer Institute in Switzerland on H5Part, a storage model and high-performance, parallel data I/O application programming interface (based on the hierarchical data format HDF5) for simplifying data exchange within the accelerator modeling community.

Data analysis and data mining are broad categories that include many techniques. Data analysis techniques include simple post-processing of experimental data or simulation output (e.g., removing noise or corrupt data, merging data from different sources, selecting data subsets), as well as using mathematical methods (applying statistical tests or optimization methods, filtering data, etc.). Data mining usually refers to the application of more advanced mathematical techniques such as dimensionality reduction, classification, clustering, time series analysis, pattern recognition, and outlier detection in order to find features or trends in data.

These analysis applications have been applied by the Analytics Team in areas such as climate modeling, astrophysics, and network security. In climate science, a member of the Analytics Team tested the application of two analysis methods for climate modeling. The first was the blind source separation (BSS) method, a statistical method for detecting unusual signals, to the output of a general circulation model to see whether the BSS method could detect the combination of simulation parameters and range of values that correspond to tropical storms without *a priori* defining them. BSS performed well for this application, and the features it detected as anomalies were variations on rotating low-pressure systems, a finding that is consistent with conventional definitions of tropical storms. A second example was the application of the multi-taper method (MTM), which combines spectral frequency analysis with a statistical model, to detect trends caused by model-induced drift in long-term simulations of surface air temperature. In order to determine spatial-temporal patterns of interest in the simulations, it is necessary to separate trends resulting from model drift from periodic oscillations due to natural variability. The MTM was able to determine the background trends, both on global and smaller spatial scales, so that the data from the surface air temperature simulations could be de-trended (Figure 1).

Analysis of simulation results also played a role in confirming that the physics in the FLASH model explains the transition from deflagration to detonation in thermonuclear supernovae. A member of the NERSC Analytics Team developed software to generate synthetic spectra from simulation results in order to compare simulation results with observational data from supernovae. The synthetic spectra capture the characteristics of the observed spectra as shown in Figure 2.

In another collaboration with NERSC users, analysis and visualization of supernova data and comparison with model data led to the detection of a new type of Type Ia supernova that has a mass greater than the Chandrasekhar mass limit of 1.4 solar masses.[1]

Working in collaboration with the Nearby Supernova Factory (described in more detail on see page 39), members of the Analytics Team developed improved algorithms for processing images of supernova candidates[2] and applied machine learning algorithms[3] to discover and classify supernovae more efficiently and accurately. Development and implementation of these analysis methods, together with workflow management (described below) have led to significant performance enhancements and savings in time and effort.

The success of astronomical projects that look for transient objects (e.g., supernova candidates in the case of the Nearby Supernova Factory) depends on having high-quality reference images for comparison with images of candidate objects. A member of the Analytics Team has developed a processing pipeline to co-add 2°-by-2° sections of one million existing images from sky surveys. The pipeline consists of a serial task, which determines the relative depth of each image with respect to the other images in the set and the physical scale of each image and its location on the sky, and a parallel part, which does the co-addition of hundreds of images to produce a single 2°-by-2° reference image. The outcome of this project, which is expected to take years to complete, will be a library of high-quality reference images available to the astronomical community.

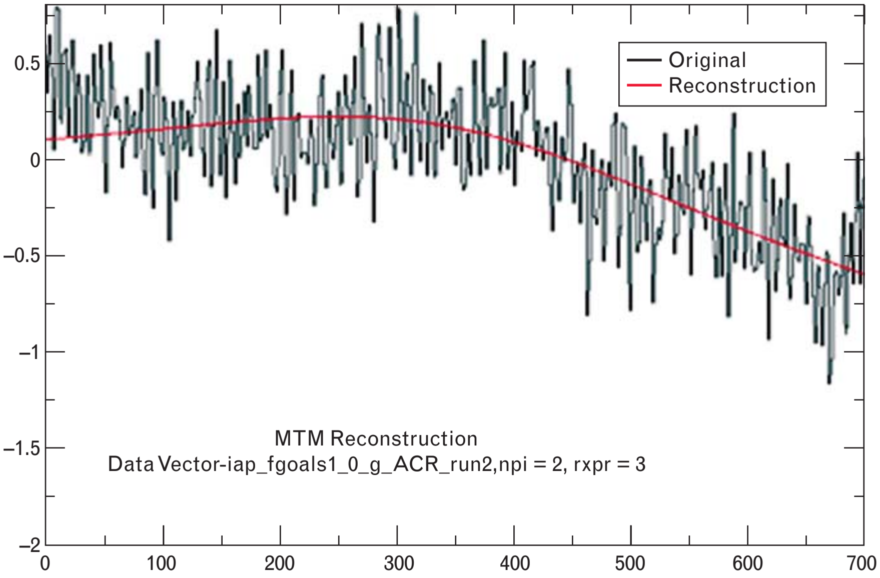

FIGURE 1. Application of the MTM to long-time simulations of surface air temperature. The black line shows the simulation results for changes in surface air temperature (°C) vs. time. The red line shows the trend line determined by the MTM method. (Figure provided by Raquel Romano, LBNL)

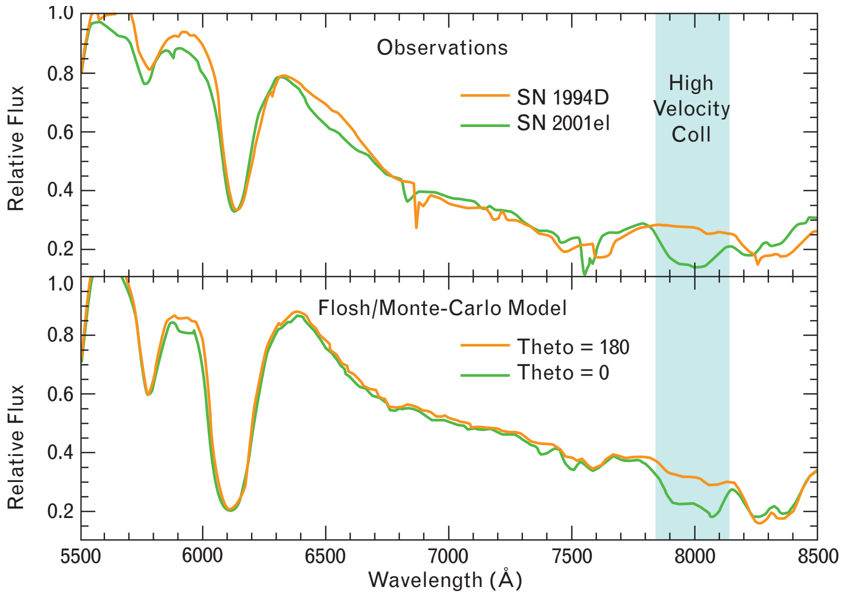

FIGURE 2. Relative flux from observations of two supernovae (top figure) and analysis of model results (bottom figure). (Figure by Peter Nugent, LBNL)

Visualization of data is the transformation of data—experimental data or simulation results—into images. Visualizing data is an invaluable tool for data exploration because it provides multiple views of the data (e.g., isosurfaces, volume rendering, streamlines) and facilitates searching for features or regions of interest. In addition, when presented as a series of images representing a change in time, parameter values, etc., data visualization provides a medium for sharing and communicating results. The NERSC Analytics Team includes members of the Berkeley Lab Visualization Group, whose mission is to assist researchers in achieving their scientific goals more quickly through creative

---

[1] D. Andrew Howell et al., "The type Ia supernova SNLS-03D3bb from a super-Chandrasekhar-mass white dwarf star," Nature **443,** 308 (2006).

[2] Cecilia R. Aragon and David Bradburn Aragon, "A fast contour descriptor algorithm for supernova image classification," Lawrence Berkeley National Laboratory technical report LBNL-61182 (2007).

[3] Raquel A. Romano, Cecilia R. Aragon, and Chris Ding, "Supernova recognition using support vector machines," Lawrence Berkeley National Laboratory technical report LBNL-61192 (2006).

and inspired visualization of their data. The NERSC Analytics Team/Visualization Group provides the expertise to help users focus on visualizing their data without having to invest the significant amount of time required to learn about visualization tools and techniques. Team members also work with users to develop new capabilities, such as data readers, to facilitate importing simulation or experimental data into visualization applications.

Recent collaborative visualization work with NERSC users includes visualization of results from accelerator physics, fusion energy, hydrodynamic, and molecular dynamics simulations. Several examples are shown in Figures 3–5.

The Analytics Team also collaborates with other NERSC groups to provide support for DOE computer science research programs. The resources that NERSC provides for such research—both in terms of hardware and expertise—are unique. Working together, staff from NERSC network and security, ESnet, LBNL's Scientific Data Management Center, and the NERSC Analytics Team developed a network traffic analysis application to detect the occurrence of "distributed scan" attacks. The application combines visual analytics with state-of-the-art SDM technology for indexing and querying data. A "hero-sized" data set consisting of 42 weeks of network traffic connection data was analyzed in order to discover and characterize a sophisticated distributed network scan attack and to identify the set of hosts perpetrating the attack. For details of the team's analysis techniques and data mining strategy, see the following section, "Query-Driven Visualization."

Workflow management is the last key technology that plays a role in the NERSC Analytics Program. The goal of workflow management is to automate specific sets of tasks that are repeated many times, thus simplifying execution and avoiding human errors that often occur when performing repetitive tasks. A large factor in the success of the Nearby Supernova Factory has been the application of a workflow management strategy by members of the Analytics Team

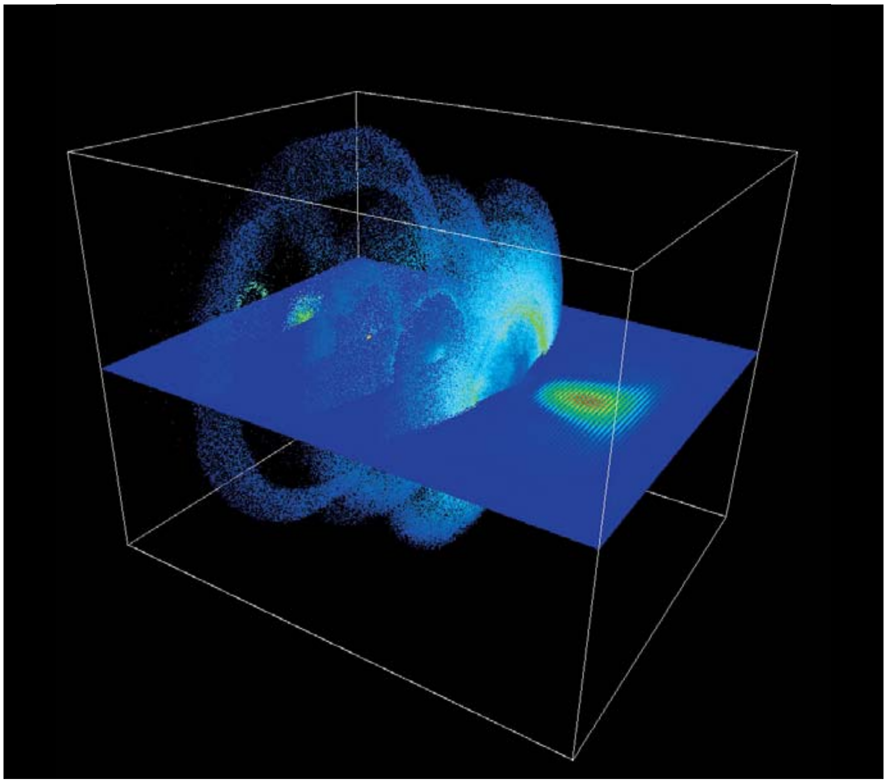

FIGURE 3. Visualization of electric field and high-energy electrons in a laser wakefield accelerator (LWFA). Experiments with LWFAs have demonstrated accelerating gradients thousands of times greater than those obtained in conventional particle accelerators. This image shows a horizontal slice through the electric field in LWFA. The electrons, displayed as spheres, are colored by the magnitude of the momentum. This example shows how several key types of data can be combined in one image to show their spatial relationship. (PI: Cameron Geddes, LBNL. Visualization by Cristina Siegerist, LBNL.)

PCCP

Physical Chemistry Chemical Physics

An international journal

FIGURE 4. The cover of the March 2006 Journal of Physical Chemistry/Chemical Physics shows a plot of the wide range of attainable molecular hydrogen binding affinities with simple ligands and metal complexes, together with a contour plot of the H_2 electrostatic potential in the background. (PI: Martin Head-Gordon, UCB/LBNL. Visualization by Rohini Lochan and Cristina Siegerist.)

Journal of Microscopy

Blackwell Publishing

Volume 223 Part 1 July 2006

FIGURE 5. The cover of the July 2006 Journal of Microscopy shows a slice from a volume-rendered tomographic reconstruction of the bacterium *Deinococcus radiodurans* showing similar high-density bodies as those seen in *D. grandis.* (PI: Luis Comolli, LBNL. Visualization by Luis Comolli and Cristina Siegerist.)

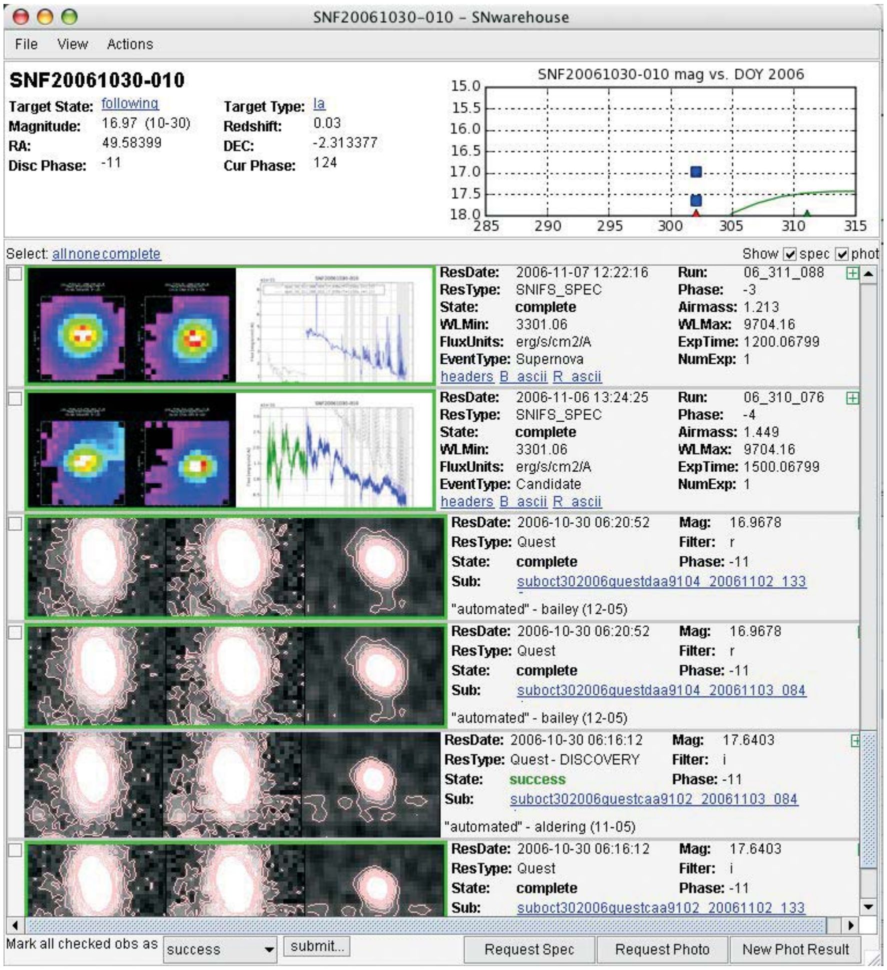

SNF20061030-010

Target State: following
Magnitude: 16.97 (10-30)
RA: 49.58399
Disc Phase: -11
Target Type: Ia
Redshift: 0.03
DEC: -2.313377
Cur Phase: 124

Select: allnonecomplete
Show ☑ spec ☑ phot

☐
ResDate: 2006-11-07 12:22:16
ResType: SNIFS_SPEC
State: **complete**
WLMin: 3301.06
FluxUnits: erg/s/cm2/A
EventType: Supernova
headers B_ascii R_ascii
Run: 06_311_088
Phase: -3
Airmass: 1.213
WLMax: 9704.16
ExpTime: 1200.06799
NumExp: 1

☐
ResDate: 2006-11-06 13:24:25
ResType: SNIFS_SPEC
State: **complete**
WLMin: 3301.06
FluxUnits: erg/s/cm2/A
EventType: Candidate
headers B_ascii R_ascii
Run: 06_310_076
Phase: -4
Airmass: 1.449
WLMax: 9704.16
ExpTime: 1500.06799
NumExp: 1

☐
ResDate: 2006-10-30 06:20:52
ResType: Quest
State: **complete**
Sub: suboct302006questdaa9104_20061102_133
Mag: 16.9678
Filter: r
Phase: -11
"automated" - bailey (12-05)

☐
ResDate: 2006-10-30 06:20:52
ResType: Quest
State: **complete**
Sub: suboct302006questdaa9104_20061103_084
Mag: 16.9678
Filter: r
Phase: -11
"automated" - bailey (12-05)

☐
ResDate: 2006-10-30 06:16:12
ResType: Quest - DISCOVERY
State: **success**
Sub: suboct302006questcaa9102_20061103_084
Mag: 17.6403
Filter: i
Phase: -11
"automated" - aldering (11-05)

☐
ResDate: 2006-10-30 06:16:12
ResType: Quest
State: **complete**
Sub: suboct302006questcaa9102_20061102_133
Mag: 17.6403
Filter: i
Phase: -11

FIGURE 6. Supernova Warehouse screenshot. The Supernova Warehouse is a Web-based visual analytics application that provides tools for data management, workflow visualization, and collaborative science. This screenshot shows several supernova or supernova candidate events, including processed images, spectra from the SuperNova Integral Field Spectrograph, and other observational data and information.

to create a new pipeline, the Supernova Factory Assembly Line (SUNFALL), for managing data and processing images of supernova candidates. Major components of SUNFALL are improved image processing and classification methods, a Web-based workflow status monitor, data management services, and the Supernova Warehouse, a visual analytics system for examining supernovae and supernova candidates (Figure 6). The Supernova Warehouse provides convenient Web-based access to project data and information, easy-to-use data annotation tools, and improved context awareness. SUNFALL, which went into production in fall 2006, has automated moving images and data through the analysis process and has significantly improved not only the speed with which supernova candidates are classified, but also the accuracy of the classifications. Thus, the new pipeline has significantly reduced both the time and labor involved in identifying supernova candidates and, as a result, has increased scientific productivity.

In addition, the Analytics Team is assessing use of the Kepler workflow system. Kepler is a grid-based scientific workflow system that can orchestrate complex workflows such as the pre-processing of simulations (preparation of input files and checking of simulation parameters), checking of intermediate results, and launching post-processing analyses. Kepler is supported in part by the SciDAC Scientific Data Management Center. Though Kepler may prove useful for other domains as well, the Analytics Team's current focus is on applying Kepler to projects in the accelerator modeling community.

Members of the Analytics Team in 2006—Cecilia Aragon, Wes Bethel, Peter Nugent, Raquel Romano, and Kurt Stockinger—worked with and provided support for NERSC users working in a dozen science domains (including accelerator physics, fusion energy, astrophysics, climate sciences, and life sciences) and at both national labs and universities. Collaborations between the Analytics Team and NERSC users have led to increased productivity for researchers, allowing them to focus more on science and less on managing data and simulation and analysis infrastructure, as well as development of methods and technologies that can be deployed at NERSC and elsewhere.

More information about the technologies described above, additional cases studies, and links to documentation on software applications installed on NERSC machines are available on the NERSC Analytics Resources Web pages at www.nersc.gov/nusers/analytics.

QUERY-DRIVEN VISUALIZATION: A CASE STUDY

NERSC and the NERSC Analytics Program offer resources uniquely suited to support research and development activities in the areas of scientific visualization, visual analytics, data analysis and other similar data-intensive computing endeavors. These resources offer capabilities impossible to replicate on desktop platforms and, in some cases, capabilities that cannot be met with "departmental clusters." The centerpiece of these resources is DaVinci, an SGI Altix with 32 1.4 GHz Itanium 2 processors and 192 GB of RAM in an SMP architecture.

These resources were used effectively in 2006 in support of ASCR-sponsored high performance visualization research efforts.

One such effort focused on performing analysis, data mining, and visualization of a "hero-sized" network traffic connection dataset. This work and the results, which include papers presented at IEEE Visualization 2006[4] and SC2006,[5] were made possible only by the unique resources at NERSC. The team consisted of staff from ESnet and NERSC's network and security group, visualization researchers from LBNL, and members of the NERSC Analytics Team.

A typical day of network traffic at an average government research laboratory may involve tens of millions of connections comprising multiple gigabytes of connection records. These connection records can be thought of as conversations between two hosts on a network. They are generated by routers, traffic analyzers, or security systems, and contain information such as source and destination IP address, source and destination ports, duration of the connection, number of bytes exchanged, and date/time of the connection. A year's worth of such data currently requires on the order of tens of terabytes or more of storage. According to ESnet General Manager Joe Burrescia, the traffic volume on ESnet, DOE's scientific production network serving researchers at national laboratories and universities, has been increasing by an order of magnitude every 46 months since 1990.[6] This trend is expected to continue into the foreseeable future.

The steady increase in network traffic volume increases the difficulty of forensic cybersecurity or network performance analysis. Current network traffic analysis toolsets rely on simple utilities like grep, awk, sed and gnuplot. Though sufficient for analyzing hours of network traffic data, these utilities do not scale nor perform up to the level needed for analyzing current and future levels of network traffic.

To address the need for rapid forensic analysis capabilities, the NERSC Analytics Team combines two complementary technologies to analyze data for a network traffic case study. The first technology is FastBit, a state-of-the-art scientific data management technology for data indexing and querying developed by the SciDAC SDM Center.[7] Data mining and knowledge discovery depend on finding and analyzing "interesting" data, so achieving maximum possible performance in data mining requires the best possible technology. Second, the team uses a query-driven visualization and analytics research application for formulating multi-resolution queries—specifically, multi-resolution in the temporal dimension—and for obtaining and displaying multidimensional histograms. A key concept in this work is that the team is computing and displaying data histograms and does not need to access the raw data directly.

Attackers often use sets of previously compromised hosts to collectively scan a target network. This form of attack, known as a distributed scan, is typically accomplished by dividing the target address space up among a group of "zombie" machines (systems that have been enslaved for the purpose of carrying out some action, usually malicious)and directing each zombie to scan a portion of the target network. The scanning results from each zombie are aggregated by a master host to create a complete picture for the attacker. An example attack would be a search for unsecured network services on a given port or range of ports. Identifying sets of hostile hosts under common control is helpful in that the group of hosts can

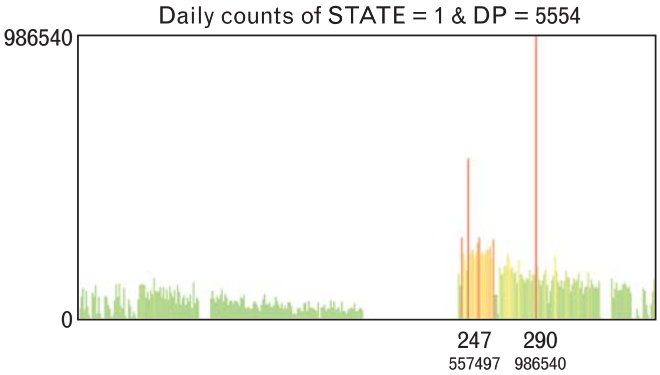

FIGURE 7. Unsuccessful connection attempts to port 5554 over a 42-week period shown as a one-dimensional histogram. While the source data are sampled at per-second resolution, the team created the histogram by requesting the number of connection attempts on a per-day basis, i.e., each histogram bin is one day in width. While the largest spike occurs on day 290, the range of elevated activity around day 247 is more interesting as it indicates what may be temporally coordinated activity. Each histogram bar is color-coded according to the number of standard deviations from the mean. Green bars are close to the mean number of unsuccessful connection attempts, while red bars indicate bins that have a count that are three or more standard deviations away from the mean number of per-bin counts.

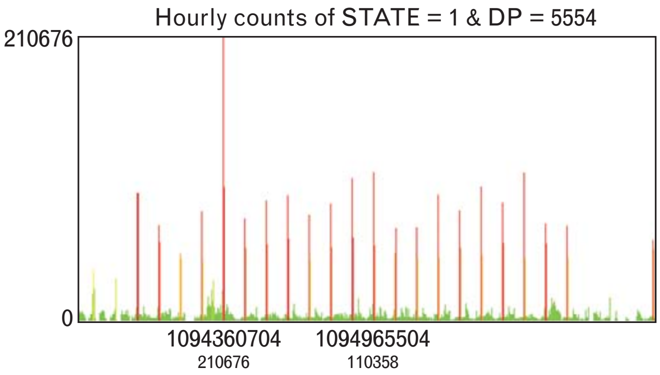

FIGURE 8. Histogram of unsuccessful connection attempts at one-hour resolution over a four-week period. This histogram shows that the suspicious activity is temporally periodic, with a period of approximately one day and repeating over a twenty-one day window.

[4] E. Wes Bethel, Scott Campbell, Eli Dart, Kurt Stockinger, and Kesheng Wu, "Accelerating network traffic analysis using query-driven visualization," Proceedings of the 2006 IEEE Symposium on Visual Analytics Science and Technology, Baltimore MD, November 2006, pp 115-122; Lawrence Berkeley National Laboratory technical report LBNL-59819 (2006).

[5] Kurt Stockinger, E. Wes Bethel, Scott Campbell, Eli Dart, and Kesheng Wu, "Detecting distributed scans using high-performance query-driven visualization," Proceedings of SC06; Lawrence Berkeley National Laboratory technical report LBNL-60053 (2006).

[6] Joseph Burrescia and William E. Johnston, "ESnet status update," Internet 2 International Meeting, September 19, 2005.

[7] K. Wu, E. Otoo, and A. Shoshani, "Optimizing bitmap indices with efficient compression," ACM Transactions on Database Systems **31,** 1 (2006).

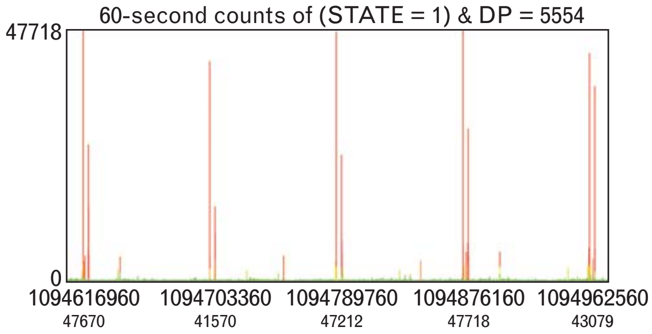

FIGURE 9. Histogram of unsuccessful connection attempts to port 5554 over a five-day period of time sampled at one-minute granularity. The histogram indicates a repeating pattern of unsuccessful connection attempts that occur on a regular twenty-four hour interval. Each primary spike is followed by a secondary, smaller spike fifty minutes later.

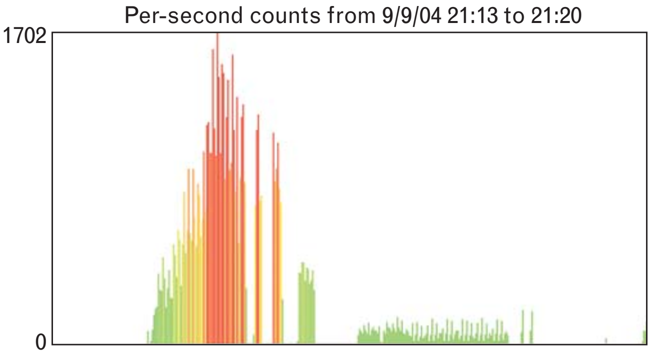

FIGURE 10. Histogram of unsuccessful connection attempts that occur at one-second temporal resolution within a seven-minute window. This seven-minute window corresponds to a primary daily activity spike within the four-week period of interest. The figure shows a ramp-up in activity that then declines and drops off.

be blocked from access to critical infrastructure, and the hosts can be reported to the larger community for further analysis or action.

To detect a distributed scan and identify the set of remote hosts participating in the scan, the team carried out the following sequence of data mining steps. Their initial assumption was that some type of suspicious activity is occurring on destination port 5554. Such an assumption is based on an intrusion detection system (IDS) alert. The first step is to obtain a global view—how many unsuccessful attempts occurred over a 42-week period on port 5554? Answering this question (Figure 7) helps the team begin to understand the characteristics of a potential distributed scan attack.

The large counts on days 247 and 290 are visible in Figure 7 as tall red spikes. Note that around day 247 there is a consistent increase in daily connection counts. The activity on day 290 looks different, since there appears to be a large increase in scanning activity, possibly indicating a new scanning tool or technique. The activity around day 247 appears at first glance to be careful work over time by a set of hosts, whereas the day 290 activity could easily be a single host scanning at a very high rate with the intent to "get in and get out" before its presence can be detected and its access blocked. Such high speed scans are quite common and often are part of the reconnaissance phase of a larger attack mechanism, or a combined scan and attack tool.

The team drills into the data by computing a histogram containing the number of unsuccessful connection attempts over a four-week period at one-hour resolution (Figure 8). There is no need to increase resolution at the global level since they are interested only in those events within a smaller time window. Here is where the power of query-driven visualization and analytics comes into play—focusing visualization and analysis efforts only on interesting data.

The next step is to drill into the data at a finer temporal resolution by posing a query that requests the counts of per-minute unsuccessful connection attempts over a five-day period within the four-week window (Figure 9). The purpose of this query is to further refine our understanding of the temporal characteristics of the potential attack.

With the one-minute view shown in Figure 9, a distinct pattern of increased unsuccessful connection attempts on precise 24-hour intervals can be seen. Decoding the UNIX timestamp reveals the event occurs daily at 21:15 local time. Each such spike is followed by a secondary spike that occurs about 50 minutes later.

Drilling yet deeper into the data, the team constructs a histogram showing the number of connection attempts over a seven-minute window at one-second temporal resolution (Figure 10). The seven-minute window is chosen to completely contain the primary daily activity spike that occurs on one day within the overall temporal region of interest.

At this point, the temporal characteristics of the suspicious activity can be determined—an organized scanning event is occurring daily at 21:15 local time within a four-week window of the 42-week data set. While there also appears to be a secondary event that occurs 50 minutes later within the same four-week period, to simplify the rest of this analysis, this case study focuses only on the daily event occurring at 21:15.

So far, the analysis has focused on understanding the temporal characteristics of the attack. Next the team wishes to establish that the attack is a scan, and then identify all the hosts that are perpetrating the attack in order to determine if the attack is from a single host, or from multiple hosts that are coordinating their effort. Their next steps will be to look at the destination addresses covered by the attack.

This phase of the analysis begins by constructing a 3D histogram to understand the extent of coverage in the destination address space (Figure 11). Two of the axes represent the destination C and D address octets; the third axis represents time. The time axis covers a two-hour window sampled at one-minute granularity.

At this point, the second of the three analysis questions has been answered—the suspicious activity appears to be a scanning attack,

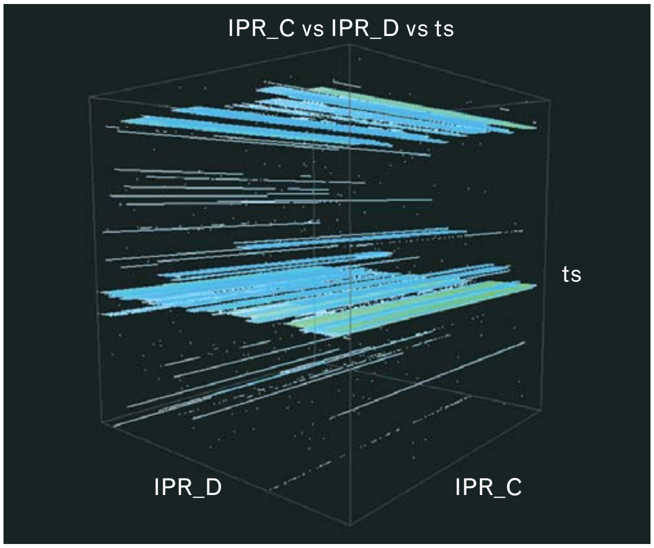

FIGURE 11. 3D histogram showing the coverage of the potential attack in the destination network. Two of the axes are the destination C and D address octets, and the third (vertical) axis is time. Here two sheet-like structures can be seen that correspond to the primary and secondary daily spikes in suspicious activity. The sheet structure indicates that the suspicious activity is occurring across the entire range of destination C and D addresses — such behavior is indicative of a scanning attack.

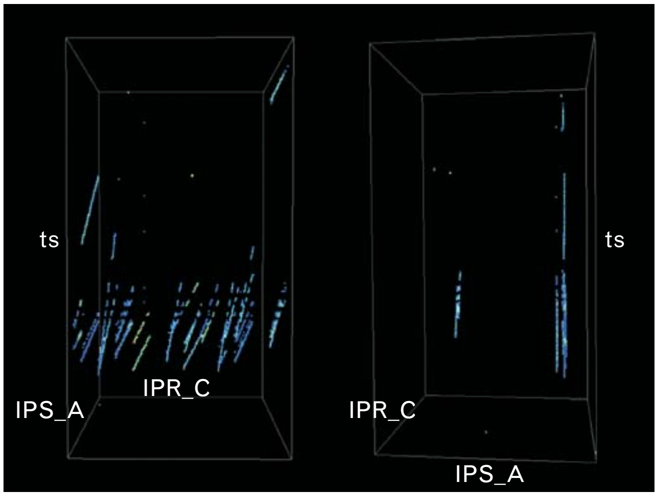

FIGURE 12. Different views of a 3D histogram showing the number of unsuccessful connection attempts to all destination C addresses from all source A addresses within a seven-minute window at one-second temporal resolution.

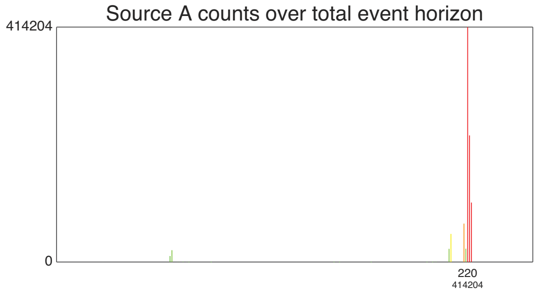

FIGURE 13. Histogram of unsuccessful connection attempts from each of the 255 addresses within the source A address octet within the same seven-minute window of time.

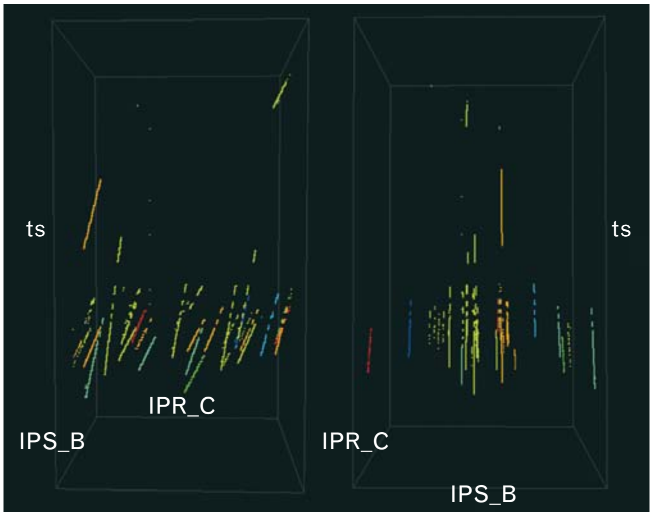

FIGURE 14. Different views of a 3D histogram showing the number of connection attempts to all destination C addresses from all source B addresses within a seven-minute window sampled at one-second temporal resolution.

as all C and D address octets within the destination network are being probed. The next step is to discover the set of remote hosts that are perpetrating the attack. To do so, a series of steps are performed to identify the A, B, C, and D address octets of the attacking hosts.

A set of histograms (Figures 12 and 13) are constructed to identify the A address octet. Figure 12 shows two different views of a 3D histogram in which two of the axes are the destination C and D addresses, and the third axis is time. The time axis encompasses a seven-minute window with a resolution of one second. The seven-minute window encompasses one of the daily primary activity spikes.

The 3D histogram shows structures that are indicative of scanning activity. The slanted lines in the image on the left in Figure 12 show that different destination C addresses are being scanned over a relatively narrow range of time. The image on the right in Figure 12 shows that such activity is confined to a fairly narrow range of source A addresses.

The histogram in Figure 13 shows exactly the source A addresses participating in the attack, with the source A address of 220 showing the greatest level of activity. The team's visual analytics application indicates precisely (not shown here) the top (i.e., highest frequency) bins in a histogram. In this case, a total of seven unique source A addresses are identified as participating in the scan. Because all hosts in the seven A addresses are engaging in probing activity at about the same time, the team assumes that they are part of a coordinated distributed scan.

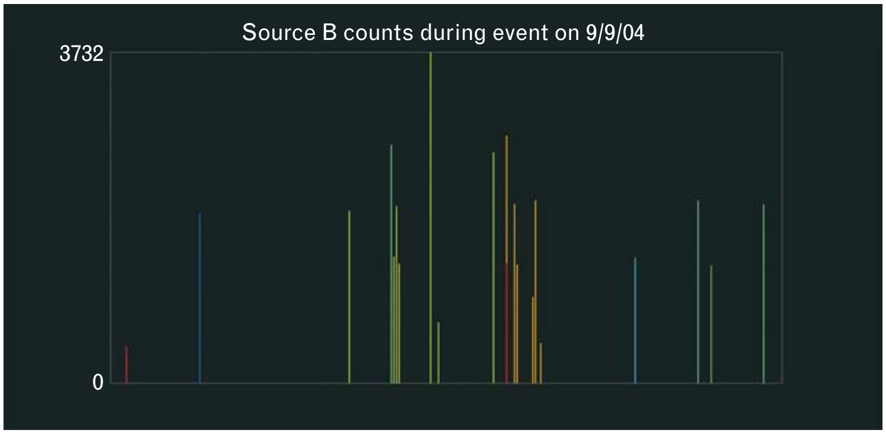

FIGURE 15. Histogram of unsuccessful connection attempts from addresses within the source B address octet within a seven-minute window of time.

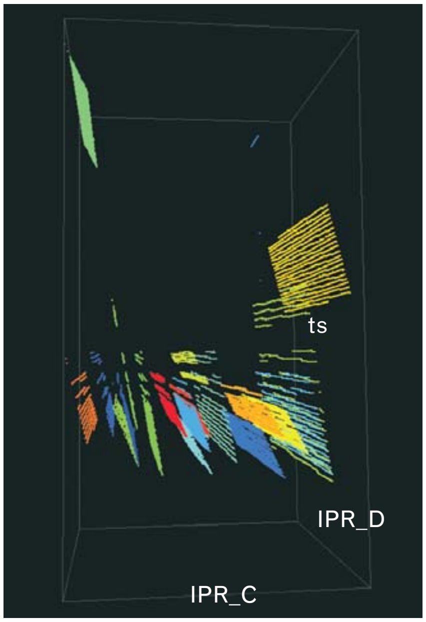

FIGURE 16. 3D histogram of coverage in the destination C and D addresses by all twenty hosts participating in the distributed scan over a seven-minute window sampled at one-second granularity. The visualization is color-coded by unique source address to show how each source host is attacking a different part of the destination address space.

The analysis is repeated, looking at source B addresses. The query iterates over the seven source A addresses identified in the previous step. The results are shown as two different views of a 3D histogram (Figure 14) and a one-dimensional histogram (Figure 15). In these histograms, the dots (Figure 14) and bars (Figure 15) are color-coded according to the source A address of the attacking host. Figures 14 and 15 show that different hosts are attacking different portions of the destination addresses. This type of behavior is indicative of a distributed scan, where the destination address space is divided among a group of zombie hosts.

This analysis step is repeated to identify the unique C and D addresses of the attacking hosts. As shown in Figure 16, the four analysis steps reveal that a total of twenty different hosts are participating in the distributed scan.

The data mining example above would not have been possible without the ability to quickly interrogate data to produce histograms. Generally speaking, the amount of time required to perform the queries varies according to the size of the source data (the size of the data indices, to be more precise), the complexity of the query, and the number of items returned by the query. The average time to perform the queries used in the analysis presented here is on the order of a few minutes and uses an iterative approach implemented in serial fashion on the NERSC analytics machine, DaVinci.

Several unique characteristics of this system enabled rapid progress on this research. First, by using an SMP with a large amount of memory, the team was able to quickly prototype and benchmark the performance of a family of parallel histogramming routines. Those results are described in more detail in the SC2006 paper.[5] Second, the large amount of memory in DaVinci enabled rapid prototyping and performance benchmarking of serial versions of the visual analytics interface early in the research process. These results, which include performance comparison with other serial tools for indexing and querying, are described in more detail in the IEEE Visualization 2006 paper.[4] None of these results would have been possible without a machine with a large SMP. Third, the architectural balance between fast I/O and large amounts of memory is especially well suited for data-intensive computing projects like this. The NERSC analytics machine has a large amount of scratch secondary storage (~25 TB) providing on the order 1 GB/s in I/O bandwidth—a capability that also was crucial to the success of this work.

As data size and problem complexity continue to increase, there are corresponding increases in the level of difficulty of managing, mining, and understanding data. Query-driven visualization represents a promising approach for gaining traction on the data mining and knowledge discovery challenges facing science, engineering, finance, security, and medical applications. The Analytics Team's approach to query-driven visualization combines state-of-the-art scientific data management technology for indexing and querying data, with visual analytics applications to support rapid drill-down, hypothesis testing, and displaying results.

This work is novel in several respects. First, it shows that it is possible to quickly analyze a large collection of network connection data to detect and characterize a complex attack. Previous works in network analysis focus on a few hours' or days' worth of data—this case study involved 42 weeks' worth of data. Second, it shows how a visual analytics application that combines state-of-the-art scientific data management technology with full-featured and straightforward visualization techniques can be brought to bear on a challenging data analysis problem. Third, the application shows how complex queries are formulated through an iterative, guided process that relies on statistics, rapid data access, and interactive visualization.

LEADERS

IN SCIENCE-DRIVEN CO

Helping
shape the future of
high performance
computing

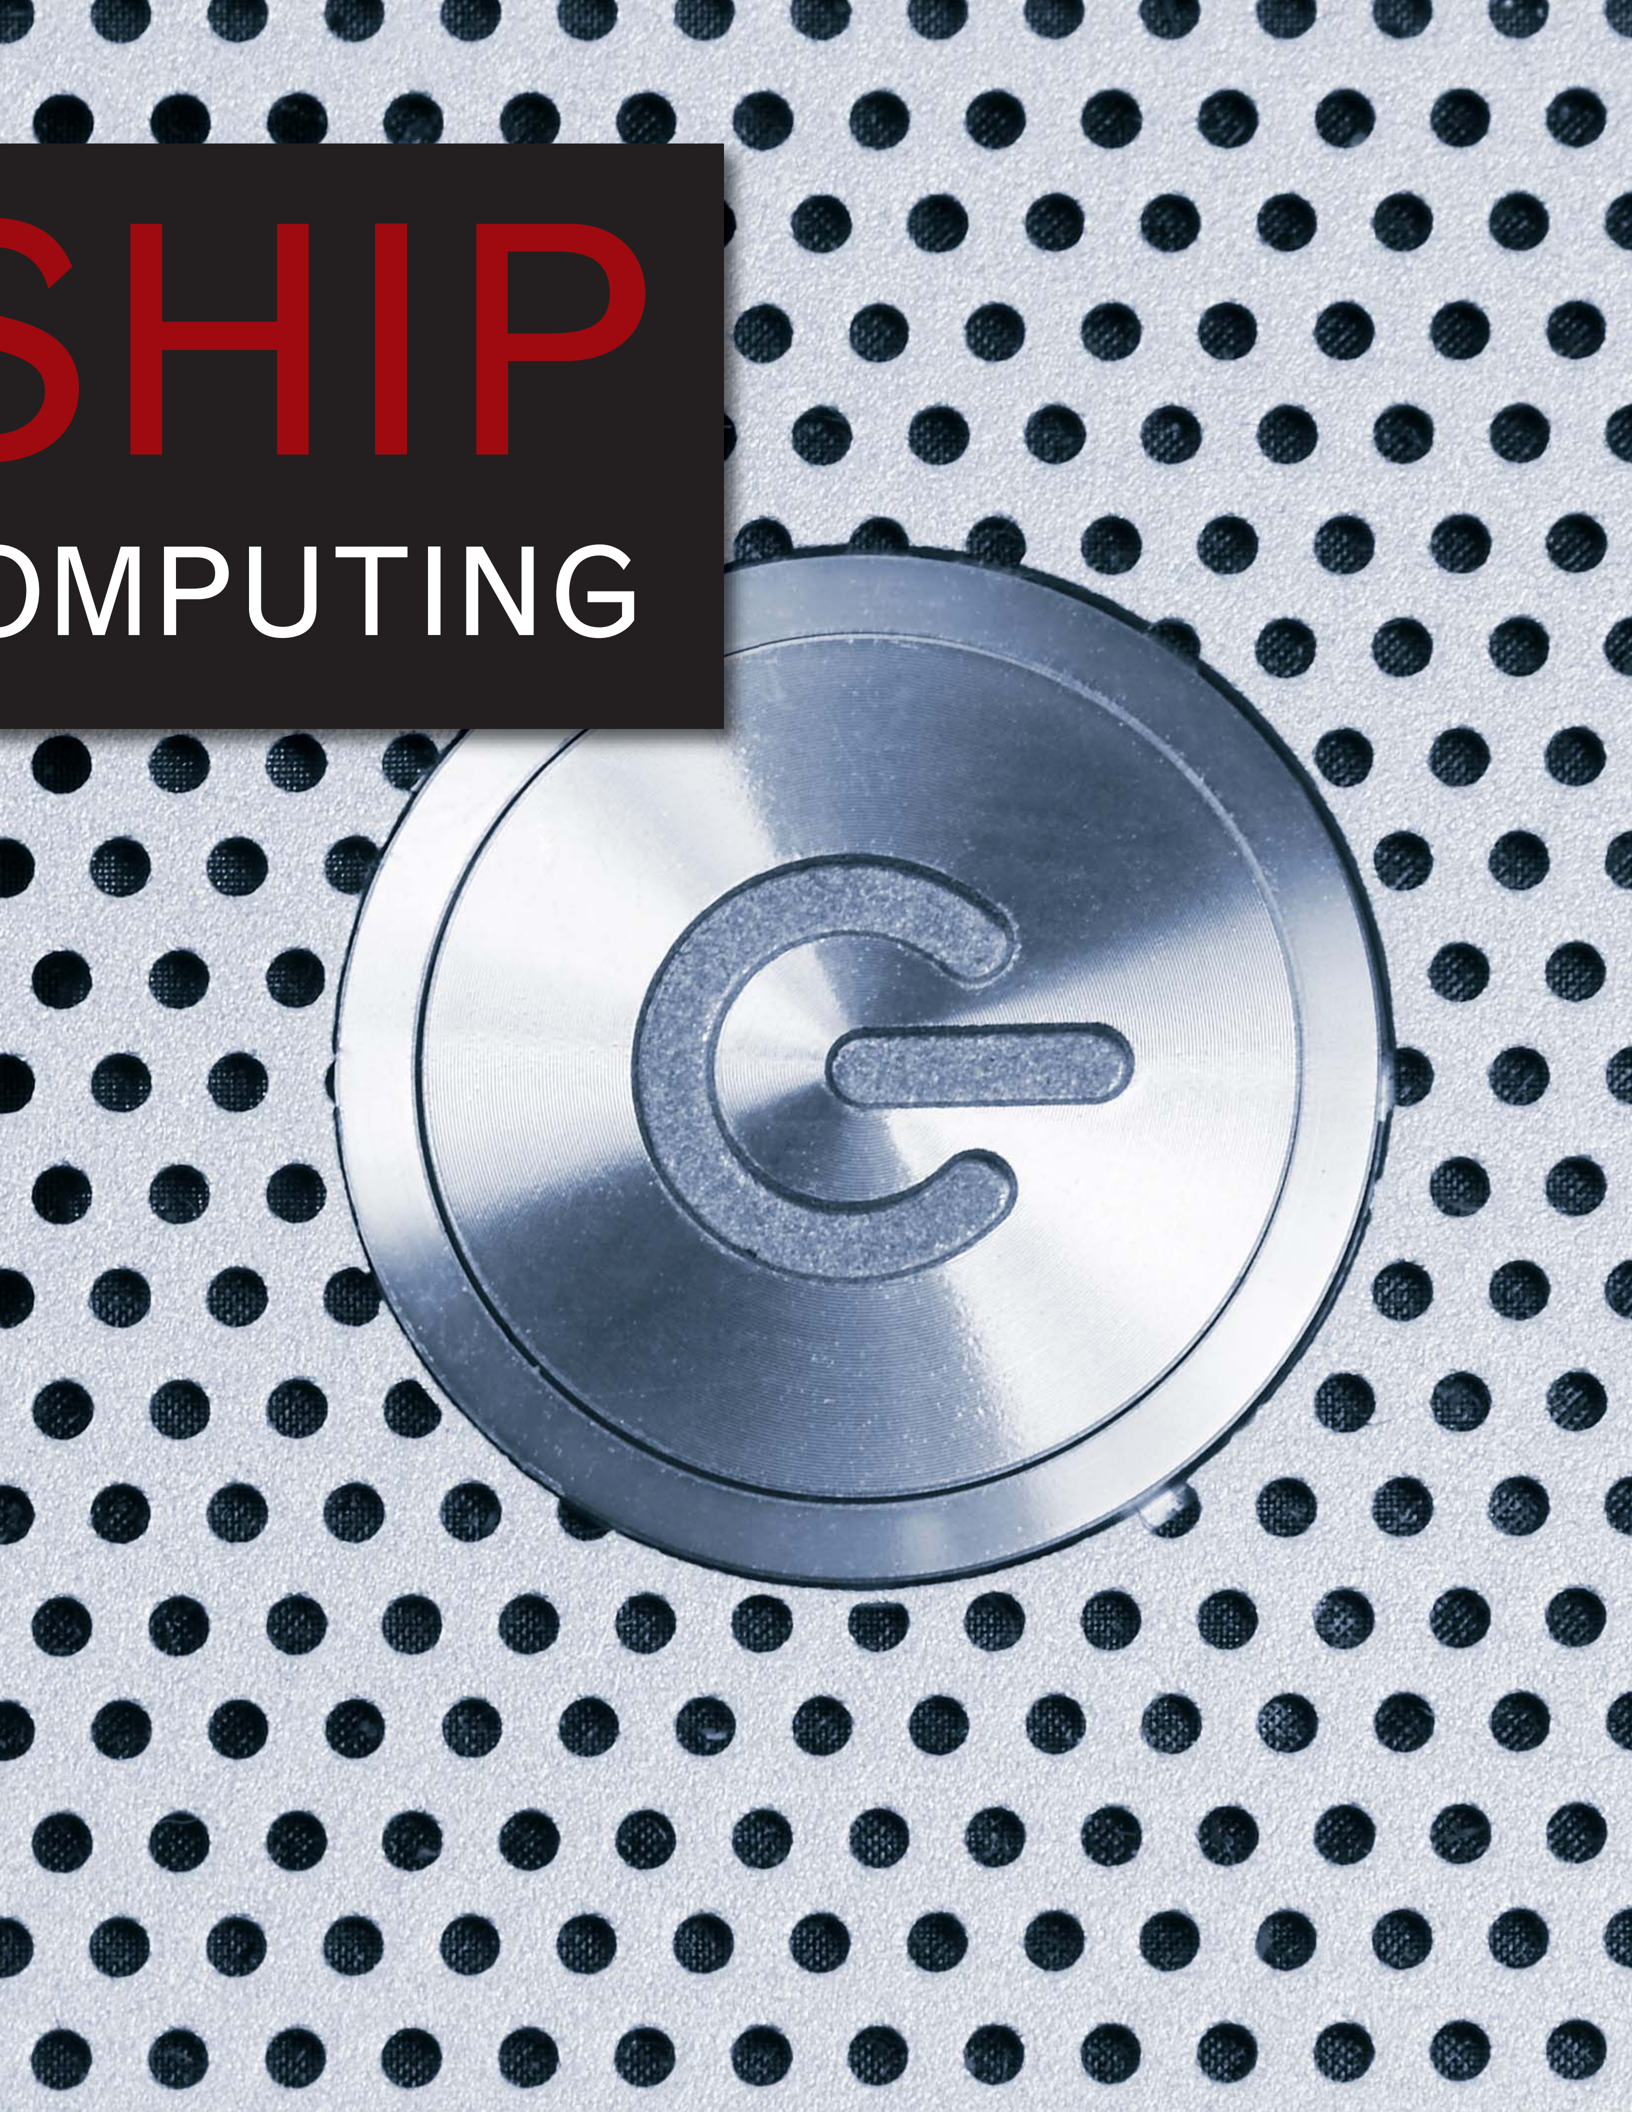

SHIP

OMPUTING

As the Department of Energy's flagship unclassified scientific computing facility, NERSC provides leadership and helps shape the field of high performance computing. NERSC is constantly trying new approaches to boost our clients' scientific productivity in an ever-changing technological environment. NERSC shares experience, knowledge, and technology not only with our clients, but also with other supercomputer sites and industry.

DEVELOPING METRICS FOR PETASCALE FACILITIES

In spring of 2006, Dr. Raymond Orbach, the Department of Energy Under Secretary for Science, asked the Advanced Scientific Computing Research Advisory Committee (ASCAC) "to weigh and review the approach to performance measurement and assessment at [ALCF, NERSC, and NLCF], the appropriateness and comprehensiveness of the measures, and the [computational science component] of the science accomplishments and their effects on the Office of Science's science programs." The Advisory Committee formed a subcommittee to respond to the charge, which was co-chaired by Gordon Bell of Microsoft and James Hack of the National Center for Atmospheric Research.

NERSC has long used goals and metrics to assure what we do is meeting the needs of DOE and its scientists. Hence, it was natural for NERSC to take the lead, working with representatives from the other sites to formulate a joint plan for metrics. Together with the other sites, NERSC then reviewed all the information and suggestions.

The committee report, accepted in February 2007, identified two classes of metrics—*control metrics* and *observed metrics*. Control metrics have specific goals which must be met, and observed metrics are used for monitoring and assessing activities. The subcommittee felt that there should be free and open access to the many observed metrics computing centers collect and utilize, but "it would be counter-productive to introduce a large number of spurious 'control' metrics beyond the few we recommend below."

The committee report pointed out, "It should be noted that NERSC pioneered the concept of 'project specific services' which it continues to provide as part of SciDAC and INCITE projects." Another panel recommendation is that the all centers "use a 'standard' survey based on the NERSC suser urvey that has been used for several years in measuring and improving service."

The final committee report is available at http://www.sc.doe.gov/ascr/ASCAC/ASCAC_Petascale-Metrics-Report.pdf.

SOFTWARE ROADMAP TO PLUG AND PLAY PETAFLOP/S

In the next five years, the DOE expects to field systems that reach a petaflop of computing power. In the near term (two years), DOE will have several "near-petaflops" systems that are 10% to 25% of a peraflop-scale system. A common feature of these precursors to petaflop systems (such as the Cray XT3 or the IBM BlueGene/L) is that they rely on an unprecedented degree of concurrency, which puts stress on every aspect of HPC system design. Such complex systems will likely break current "best practices" for fault resilience, I/O scaling, and debugging, and even raise fundamental questions about programming languages and application models. It is important that potential problems are anticipated far enough in advance that they can be addressed in time to prepare the way for petaflop-scale systems.

DOE asked the NERSC and Computational Research divisions at Lawrence Berkeley National Laboratory to address these issues by considering the following four questions:

1. What software is on a critical path to make the systems work?
2. What are the strengths/weaknesses of the vendors and of existing vendor solutions?
3. What are the local strengths at the labs?
4. Who are other key players who will play a role and can help?

Berkeley Lab responded to these questions in the report "Software Roadmap to Plug and Play Petaflop/s."[1] In addition to answering the four questions, this report provides supplemental information regarding NERSC's effort to use non-invasive work-

---

[1] William T.C. Kramer, Jonathan Carter, David Skinner, Lenny Oliker, Parry Husbands, Paul Hargrove, John Shalf, Osni Marques, Esmond Ng, Tony Drummond, and Kathy Yelick, "Software Roadmap to Plug and Play Petaflop/s," Lawrence Berkeley National Laboratory report LBNL-59999 (July 2006), http://www.nersc.gov/news/reports/LBNL-59999.pdf.

THE SOFTWARE CHALLENGES OF PETASCALE COMPUTING: AN INTERVIEW WITH KATHY YELICK

(Reprinted with permission from the November 10, 2006 issue of HPCwire)

In this HPCwire interview, Kathy Yelick, one of the world's leading performance evaluation experts and a member of NERSC's Science-Driven System Architecture Team, discusses software challenges related to petascale and other large-scale computing systems. Yelick is a professor of computer science at UC Berkeley, with a joint appointment in Lawrence Berkeley National Lab's Computational Research Division, where she leads the Future Technologies Group and the Berkeley Institute for Performance Studies.

HPCwire: Are petascale initiatives putting enough emphasis on software?

Yelick: No. Unfortunately, the race for each major performance milestone, in this case petascale, has resulted in a de-emphasis on software. Procurement teams and system developers vying for the first petascale platform need to put as much money as possible into hardware in order to be first. This leaves less funding for software. The situation has gotten worse over the past decade, when multiple agencies were supporting HPC software development.

HPCwire: Assuming the important goal isn't peak or Linpack petaflops performance, but sustained petaflops performance across a spectrum of applications, what software challenges need to be addressed?

Yelick: The major software challenge facing the petascale efforts is the explosion in hardware parallelism, which will require a complete redesign of applications, libraries, and algorithms to reach the level of parallelism needed to fully utilize a petascale machine. This parallelism increase is coming from the introduction of multi-core processors within the compute nodes and the trend towards building machines out of a larger number of smaller compute nodes. Other challenges include the hierarchical nature of these machines, the use of hardware accelerators such as SIMD units within compute nodes, and the trend toward lower degree networks. Full crossbars for petascale machines are unaffordable. Software needs to adapt to these features, and I believe it will. The question is how general the solutions will be and therefore how large the set of petascale applications will be. The reliability of these systems is also a major concern, and one that I think represents the largest risk for specific machines. We need to have better methods for handling hardware and software failures throughout the software stack.

HPCwire: Which of these software challenges are hardest? How much can be accomplished by the 2010 timeframe?

Yelick: Reliability is probably the hardest, because so far we have written user-level software assuming that the lower level system is mostly reliable. Checkpointing is the only commonly used technique, and the Berkeley Lab Checkpoint/Restart project is developing software for petascale systems; but this model is useful only as long as failures are not too frequent. There are research efforts to develop new ways of writing fault-tolerant software, but the solution is not yet clear. In the meantime, we need to do a very good job of testing systems software, in particular operating systems, to reduce the frequency of failures.

HPCwire: Can good programming languages and other software get around bad machines?

Yelick: No. There is nothing software can do to get around bad machine design. Global address space languages like UPC, CAF, and Titanium are in some sense giving good hardware an advantage over bad by trying to expose features such as low overhead communication or global address space support. That said, one of the goals of the Berkeley UPC compiler is to make UPC an effective language for a larger class of machines and for less sophisticated programmers. We have advocated language extensions such as non-blocking bulk data reads and writes to allow programmers to obtain the best possible performance on clusters, and are also working on compiler technology to automatically optimize programs written in a fine-grained style. This could make programs written for a global address space machine like the Cray X1E run reasonably well on generic clusters — not as well as on the X1E, but reasonably well.

HPCwire: What are the limits of MPI's usefulness? What would it be like relying on MPI for petascale computing?

Yelick: MPI is likely to be a very popular and effective programming model on petascale machines. There are two issues, one related to performance and the other to ease of use. For performance, the problem is that the two-sided protocol in MPI, which involves message matching, and the requirement of message ordering all slow down data transfer. The fastest mechanism on a machine with minimal RDMA support is to write data directly from one processor into another processor's memory. Fast implementations of MPI do use this mechanism, but it requires some protocol overhead, since the remote address is not known to the sending processor. As we've shown in our UPC work, one-sided communication can be used in bisection-limited problems, like global FFTs, to improve communication overlap and reduce running time. At a petascale, bisection bandwidth is going to be expensive, and MPI may not give the best utilization of the network or the best management of memory due to the need for buffering. From an ease-of-use standpoint, I think the issue with MPI is that the community of petascale programmers, like terascale programmers today, will be small, because the barrier to entry for an application code is high. There are many computational scientists today who are not using parallel machines at all. This will have to change with the shift towards multi-core, but the question is whether they will adopt a scalable programming model.

HPCwire: Talk about the importance of partitioned global address space, or PGAS, programming languages.

Yelick: Aside from my answers to the previous question, PGAS languages offer a real advantage over OpenMP for shared memory platforms, because they give programmers the opportunity to express locality properties of the data structures. This makes the PGAS models an alternative to the hybrid MPI/OpenMP model for hierarchical machines, which has proven difficult to use.

But aside from the specifics on PGAS languages, I think they represent an important step in HPC programming models, because they've demonstrated that new languages are still a viable option, in spite of the backlash that occurred when HPF failed to take hold. The PGAS languages are popular within some government agencies and labs, including the work at AHPCRC on CFD codes in UPC. We have also learned some important lessons in the UPC process: interoperability with other programming models (in particular MPI) and ubiquity across platforms are essential to success. We have new methods for analyzing and quantifying productivity; and found that performance is still critical to swaying the most elite of HPC programmers.

HPCwire: What's the status of these languages today, including Berkeley UPC?

Yelick: UPC has an active community consortium that meets regularly to work on language design issues, maintain the language spec, and exchange implementation and application experience. There is a UPC compiler of some form for nearly every serial and parallel platform, including vendor compilers from Cray, HP, and IBM, and open source compilers from Intrepid, Inc., Michigan Tech and Berkeley Lab. The Berkeley compiler has optimized implementations using native communication layers for the Cray XT3, Quadrics, Myrinet, Altix, and the IBM SP platforms. Co-Array Fortran is being adopted into the Fortran spec, and in addition to the Cray CAF compiler, there is an open source effort led by John Mellor-Crummey at Rice. That compiler is designed for portability; it uses a source-to-source translation model like Berkeley UPC, and there are plans to do work on porting and releases in the near future. Titanium is still primarily a Berkeley effort, but it is used outside Berkeley and the compiler runs on many parallel and serial platforms. Berkeley UPC, Intrepid's gcc-upc, Titanium, and at least one instance of the Rice CAF compiler all use our open source communication layer called GASNet, which helps leverage the porting effort.

HPCwire: Some people say that if there's a lot of pain involved, they won't switch to a new programming language. How can you motivate people to migrate to a more efficient new language?

Yelick: The key is that, because of interoperability, full applications do not need to be rewritten. Instead, individual components can be written in these languages as new algorithms are developed for the increasing machine scale. I am working with Parry Husbands and Esmond Ng, for example, on a fast sparse direct linear solver written in UPC. This may end up in an application without rewriting the rest of the code. And if the performance gains of new languages are significant, some people who care deeply about performance will switch. The harder argument is productivity, because while the community as a whole might save significant amounts of time and money in the long run by rewriting some code in new languages, this is difficult to quantify up front; and from the short term perspective of a 3-year or 5-year project, it is difficult to justify.

HPCwire: Do you want to say anything about scaling algorithms for petascale computing?

Yelick: I think we need to rethink our algorithms to look for all possible sources of parallelism, rather than the single-level view that we have used recently. For the SC06 conference, I co-authored a paper with Shan, Strohmaier, Qiang, and Bailey which is a case study in scaling and performance modeling, which reflects our ability to understand application performance on large machines. We look at a beam interaction application from accelerator modeling and develop a series of performance models to predict performance of this application code. Performance modeling, and therefore a simple understanding of performance, becomes increasingly difficult as the machines scale, as they become more hierarchical, and as network contention increases. All of these will be serious issues on petascale machines....

HPCwire: Is there anything important that we missed talking about?

Yelick: I think the introduction of parallelism into mainstream computing is both a challenge and opportunity for the HPC community. We must be able to handle the increase in parallelism along with everyone else, but if ever there was a time to innovate in parallel hardware, languages, and software, this is it. There are very likely to be new languages and programming models for multi-core programming, and the HPC community has the chance to take advantage of that software revolution by both influencing and using these innovations.

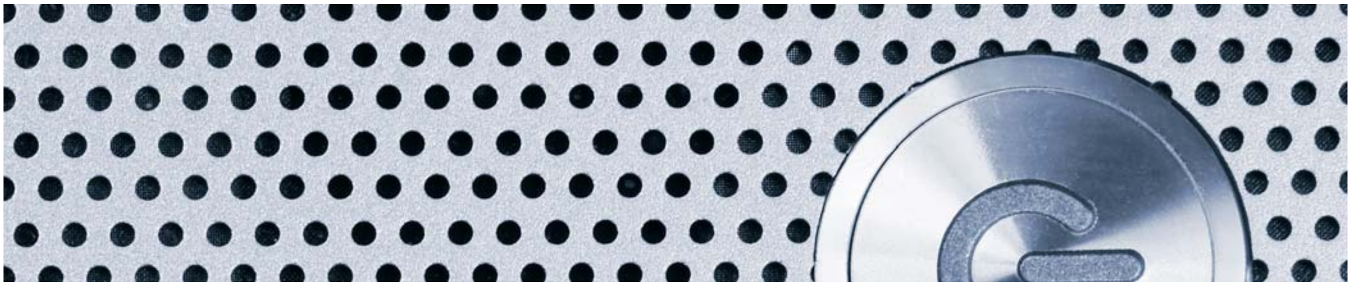

load profiling to identify application requirements for future systems; describes a set of codes that provide good representation of the application requirements of the broader DOE scientific community; and provides a comprehensive production software requirements checklist that was derived from the experience of the NERSC-3, NERSC-4, and NERSC-5 procurement teams. It presents a detailed view of the software requirements for a fully functional petaflop-scale system environment, and assesses how emerging near-petaflop systems conform or fail to conform to these requirements.

THE BERKELEY VIEW OF THE LANDSCAPE OF PARALLEL COMPUTING RESEARCH

The path to petascale computing will be paved with new system architectures featuring hundreds of thousands of manycore processors. Such systems will require scientists to completely rethink programming models. Among those computer scientists already looking to the petascale horizon are Science Driven Systems Architecture (SDSA) Team members John Shalf and Kathy Yelick, who are two of the co-authors of a white paper called "The Landscape of Parallel Computing Research: A View from Berkeley." Based on two years of discussions among a multidisciplinary group of researchers, this paper addresses the challenge of finding ways to make it easy to write programs that run efficiently on manycore systems.

The creation of manycore architectures—hundreds to thousands of cores per processor—demands that a new parallel computing ecosystem be developed, one that is very different from the environment that supports the current sequential and multicore processing systems. Since real-world applications are naturally parallel and hardware is naturally parallel, what is needed is a programming model, system software, and a supporting architecture that are naturally parallel. Researchers have the rare opportunity to re-invent these cornerstones of computing, provided they simplify the efficient programming of highly parallel systems. The paper provides strategic suggestions on how to accomplish this (see http://www.eecs.berkeley.edu/Pubs/TechRpts/2006/EECS-2006-183.pdf).

Another SDSA research collaboration is the RAMP Project (Research Accelerator for Multiple Processors), which focuses on how to build low cost, highly scalable hardware/software prototypes, given the increasing difficulty and expense of building hardware. RAMP is exploring emulation of parallel systems via field programmable gate arrays (FPGAs). Although FPGAs are slower than other types of hardware, they are much faster than simulators, and thus can be used to evaluate novel ideas in parallel architecture, languages, libraries, and so on.

PERFORMANCE AND POTENTIAL OF CELL PROCESSOR ANALYZED

Though it was designed as the heart of the new Sony PlayStation3 game console, the STI Cell processor also created quite a stir in the computational science community, where the processor's potential as a building block for high performance computers has been widely discussed and speculated upon.

To evaluate Cell's potential, LBNL computer scientists evaluated the processor's performance in running several scientific application kernels, then compared this performance with other processor architectures. The group presented their findings in a paper at the ACM International Conference on Computing Frontiers, held May 2-6, 2006, in Ischia, Italy. An article about the paper in the HPCwire newsletter was the most-read item in the history of the newsletter, according to editor Michael Feldman, and after a mention on SlashDot, the paper was viewed online by about 30,000 readers.

The paper, "The Potential of the Cell Processor for Scientific Computing," was written by Samuel Williams, Leonid Oliker, Parry Husbands, Shoaib Kamil and Katherine Yelick of Berkeley Lab's Future Technologies Group and by John Shalf, head of NERSC's Science-Driven System Architecture Team.

"Overall results demonstrate the tremendous potential of the Cell architecture for scientific computations in terms of both raw performance and power efficiency," the authors wrote in their paper. "We also conclude that Cell's heterogeneous multicore implementation is inherently better suited to the HPC environment than homogeneous commodity multicore processors."

Cell, designed by a partnership of Sony, Toshiba, and IBM, is a high performance implementation of software-controlled memory hierarchy in conjunction with the considerable floating point resources that are required for demanding numerical algorithms. Cell takes a radical departure from conventional multiprocessor or multicore architectures. Instead of using identical cooperating commodity processors, it uses a conventional high performance PowerPC core that controls eight simple SIMD (single instruction, multiple data) cores, called synergistic processing elements (SPEs), where each SPE contains a synergistic processing unit (SPU), a local memory, and a memory flow controller.

Despite its radical departure from mainstream general-purpose processor design, Cell is particularly compelling because it will be produced at such high volumes that it could be cost-competitive with commodity CPUs. At the same time, the slowing pace of commodity microprocessor clock rates and increasing chip power demands have become a concern to computational scientists, encouraging the community to consider alternatives like STI Cell. The authors examined the potential of using the forthcoming STI Cell processor as a building block for future high-end parallel systems by investigating performance across several key scientific computing kernels: dense matrix multiply, sparse matrix vector multiply, stencil computations on regular grids, as well as 1D and 2D fast Fourier transformations.

According to the authors, the current implementation of Cell is most often noted for its extremely high performance single-precision (32-bit) floating performance, but the majority of scientific applications require double precision (64-bit). Although Cell's peak double-precision performance is still impressive relative to its commodity peers (eight SPEs at 3.2 GHz = 14.6 Gflop/s), the group quantified how modest hardware changes, which they named Cell+, could improve double-precision performance.

The authors developed a performance model for Cell and used it to show direct comparisons of Cell with the AMD Opteron, Intel Itanium2 and Cray X1 architectures. The performance model was then used to guide implementation development that was run on IBM's Full System Simulator in order to provide even more accurate performance estimates.

The authors argue that Cell's three-level memory architecture, which decouples main memory accesses from computation and is explicitly managed by the software, provides several advantages over mainstream cache-based architectures. First, performance is more predictable, because the load time from an SPE's local store is constant. Second, long block transfers from off-chip DRAM can achieve a much higher percentage of memory bandwidth than individual cache-line loads. Finally, for predictable memory access patterns, communication and computation can be effectively overlapped by careful scheduling in software.

While their current analysis uses hand-optimized code on a set of small scientific kernels, the results are striking. On average, Cell is eight times faster and at least eight times more power efficient than current Opteron and Itanium processors, despite the fact that Cell's peak double-precision performance is fourteen times slower than its peak single-precision performance. If Cell were to include at least one fully utilizable pipelined double-precision floating point unit, as proposed in their Cell+ implementation, these speedups would easily double.

The full paper can be read at http://www.cs.berkeley.edu/~samw/projects/cell/CF06.pdf.

INTEGRATED PERFORMANCE MONITORING TOOL ADOPTED BY OTHER HPC CENTERS

Although supercomputing centers around the country operate different architectures and support separate research communities, they face a common challenge in making the most effective use of resources to maximize productivity. One method for doing this is to analyze the performance of various applications to identify the bottlenecks. Once identified, these performance speedbumps can often be smoothed out to get the application to run faster and improve utilization. This is especially important as applications and architectures scale to thousands or tens of thousands of processors.

In 2005 David Skinner, then a member of NERSC's User Services Group (he now leads the Open Software and Programming Group), introduced Integrated Performance Monitoring, or IPM. IPM is a portable profiling infrastructure that provides a performance summary of the computation and communication in a parallel program. IPM has extremely low overhead, is scalable to thousands of processors, and was designed with a focus on ease of use, requiring no source code modification (see http://www.nersc.gov/nusers/resources/software/tools/ipm.php).

Skinner cites the lightweight overhead and fixed memory footprint of IPM as important innovations. Unlike performance monitoring based on traces, which consume more resources the longer the code runs, IPM enforces strict boundaries on the resources devoted to profiling. By using a fixed memory hash table, IPM achieves a compromise between providing a detailed profile and avoiding impact on the profiled code.

IPM was also designed to be portable and runs on the IBM SP, Linux clusters, Altix, Cray X1, NEC SX6, and the Earth Simulator. Portability is key to enabling cross-platform performance studies. Portability, combined with IPM's availability under an open source software license, has also led to other centers adopting and adding to the IPM software. Current users include the San Diego Supercomputer Center (SDSC), the Center for Computation and Technology (CCT) at Louisiana State University, and the Army Research Laboratory.

After hearing a presentation on IPM by Skinner in 2005, SDSC staff began porting IPM to their machines. By spring 2006 it was in production.

"We decided to use it because there is nothing else out there that is as easy to use and provides the information in an easy to understand form for our users," said Nick Wright of SDSC's Performance Modeling and Characterization Lab. "It has helped our center assist users with their performance-related issues."

IPM is used quite extensively at SDSC to understand performance issues on both the IBM Power4+ system (Datastar) and the three-rack BlueGene system. In addition to regular usage by SDSC users, "IPM is used extensively by user services consultants and performance optimization specialists to diagnose and treat performance issues," Wright said. In fact, research using IPM contributed to two technical papers written by SDSC staff and submitted to the SC07 conference.

At Louisiana State University's CCT, the staff will soon start running IPM on a number of systems. Staff at the center first learned about the tool from NERSC's John Shalf, who has longstanding ties to the LSU staff, and a more recent visit by Skinner. According to Dan Katz, assistant director for Cyberinfrastructure Development at CCT, a number of CCT users are already familiar with IPM as they have run applications at both NERSC and SDSC.

LSU operates two sets of HPC resources. At LSU itself, mostly for local users and collaborators, CCT is in the process of installing a new ~1500-core Linux system, in addition to some smaller IBM Power5 systems and a few other small systems. For the Louisiana Optical Network Initiative (LONI), a statewide network of computing and data services, CCT is installing six Linux systems with a total of ~9000 cores, and six IBM Power5 systems across the state.

"In both cases, we plan to run IPM, initially on user request, and longer term automatically all the time," Katz said "We will do this for two reasons: First, to help users understand the performance of their applications, and therefore, to be able to improve performance. Second, it will help us understand how our systems are being used, which helps us understand what systems we should be investigating for future installation."

REDUCING THE RED TAPE IN RESEARCH PARTNERSHIPS

While research partnerships with other government agencies or academic institutions are the most common types of partnerships for NERSC, collaboration with commercial computer and software vendors has also produced valuable results over the years. For example, collaborations with IBM have made possible the use of GPFS software with systems from multiple vendors for the NERSC Global Filesystem, the evaluation of the Cell processor for scientific computing, and the development of the eight-processor node in the IBM Power line of processors.

Protecting intellectual property is always a priority in collaborative research, so partnerships with vendors typically involve negotiation of agreements on issues such as non-disclosure, source code licensing, or trial/beta testing. With an organization as large and dispersed as IBM, frequent collaboration can mean repeated negotiations on similar issues with different units of the organization—a time-consuming task.

NERSC staff thought there had to be a better way of setting up research agreements without starting from scratch every time. A meeting between Bill Kramer and Bill Zeitler, senior vice president and group executive of IBM's Systems and Technology Group, led to the formation of a task force to address the issue.

The Berkeley Lab team included Kramer, NERSC procurement specialist Lynn Rippe, and Cheryl Fragiadakis and Seth Rosen from the Lab's Technology Transfer Department, with input from Kathy Yelick, Lenny Oliker, and Jim Craw. Working with IBM's Doug Duberstein, they developed a set of master agreements between NERSC and IBM that could be used as starting points. IBM also set up a central repository where these agreements could be accessed throughout the organization.

For example, the master non-disclosure agreement contains the terms and conditions that the two organizations have negotiated over the years, along with a list of the types of disclosed information that have already been agreed on. When a new project involves some other type of information, the two parties just have to add a supplement that specifies that information, without having to renegotiate all the other terms and conditions. Or, to test a new trial software package, NERSC and IBM simply sign a supplement adding that software package to the master testing and licensing agreements.

These master agreements are expected to save a substantial amount of time for both organizations. NERSC is making the language in these agreements available to other laboratories so that they can simplify their negotiations too.

Appendices

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

In support of the DOE Office of Science's mission, the NERSC Center served 2,978 scientists throughout the United States in 2006. 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.

More than 1,400 scientific publications in 2006 were based entirely or in part on calculations done at NERSC; a list is available at http://www.nersc.gov/news/ reports/ERCAPpubs06.php.

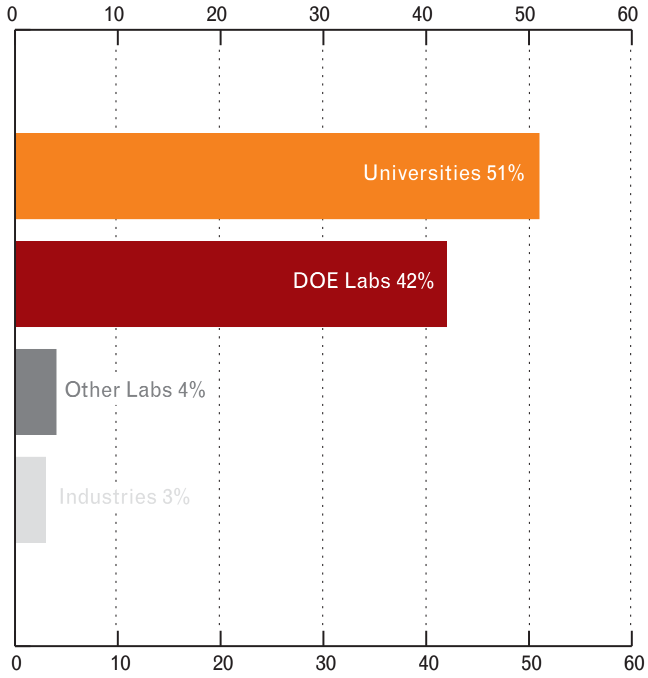

FIGURE 1. NERSC MPP usage by institution type, 2006.

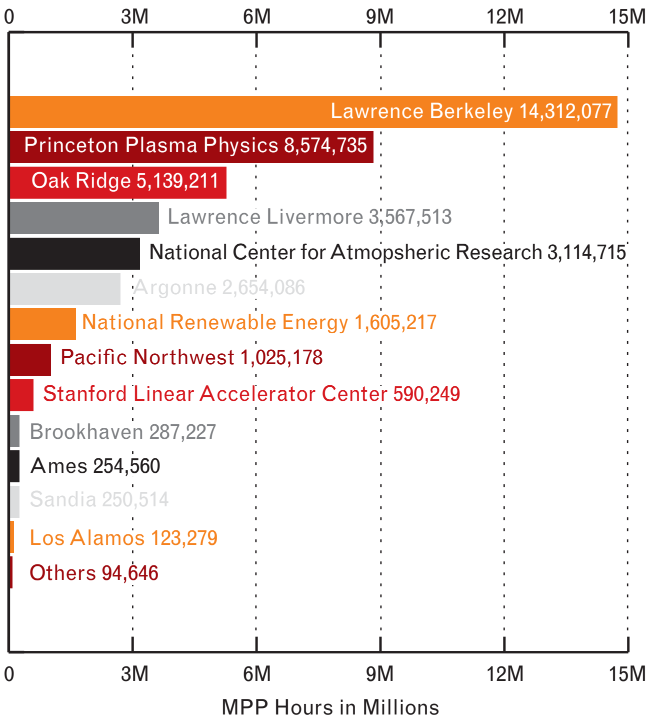

FIGURE 2. DOE and other Federal laboratory usage at NERSC, 2006 (MPP hours).

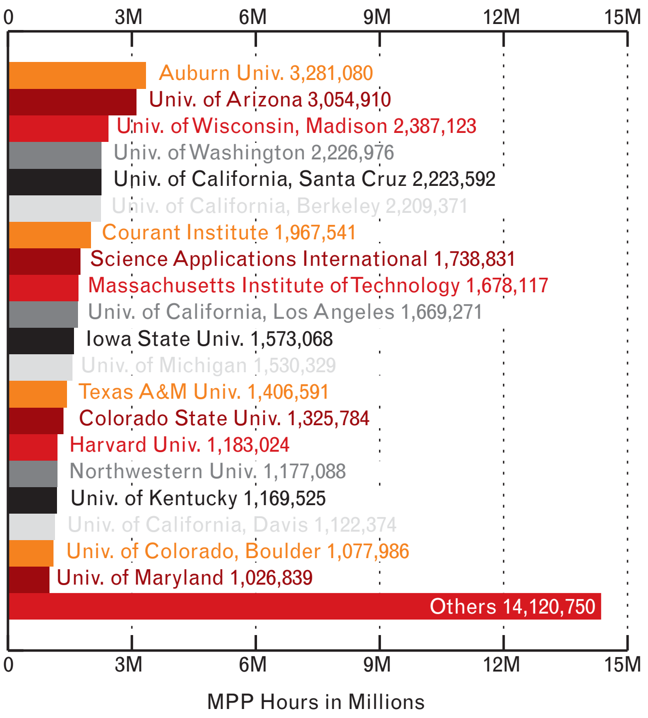

FIGURE 3. Academic and private laboratory usage at NERSC, 2006 (MPP hours).

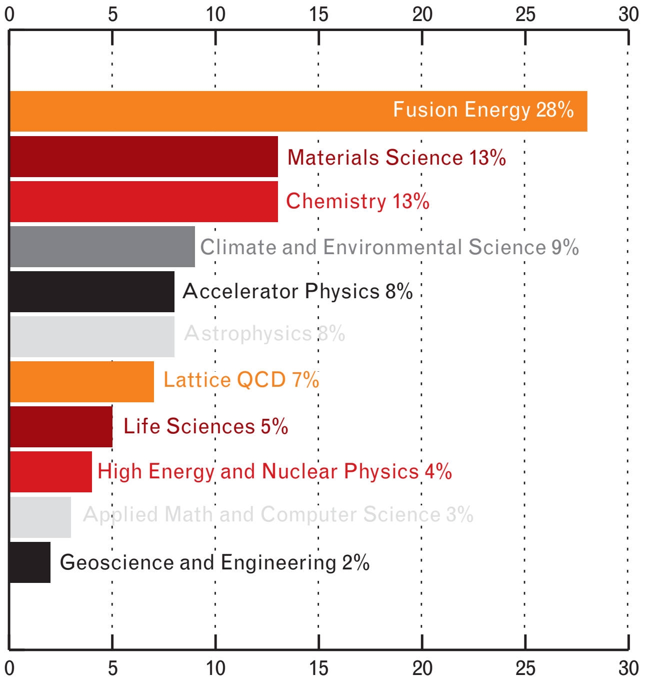

FIGURE 4. NERSC usage by scientific discipline, 2006.

Appendix C

NERSC Users Group Executive Committee

Members at Large

Yuen-Dat Chan
Lawrence Berkeley National Laboratory

Gerald Potter
Lawrence Livermore National Laboratory

Douglas Swesty
State University of New York at Stony Brook

Xingfu Wu
Texas A&M University

Office of Advanced Scientific Computing Research

Kirk Cameron
Virginia Polytechnic Institute and State University

Mike Lijewski
Lawrence Berkeley National Laboratory

Ravi Samtaney
Princeton Plasma Physics Laboratory

Office of Basic Energy Sciences

Bas Braams
Emory University

Eric Bylaska
Pacific Northwest National Laboratory

Thomas Miller
University of California, Berkeley

Office of Biological and Environmental Research

David Beck
University of Washington

Brian Hingerty
Oak Ridge National Laboratory

Adrianne Middleton
National Center for Atmospheric Research

Office of Fusion Energy Sciences

Stephane Ethier (Vice Chair)
Princeton Plasma Physics Laboratory

Andris Dimits
Lawrence Livermore National Laboratory

Alex Friedman
Lawrence Livermore and Lawrence Berkeley National Laboratories

Office of High Energy Physics

Olga Barranikova
University of Illinois at Chicago

Warren Mori
University of California, Los Angeles

Frank Tsung
University of California, Los Angeles

Office of Nuclear Physics

David Dean (Chair)
Oak Ridge National Laboratory

Patrick Decowski
Lawrence Berkeley National Laboratory

James Vary
Iowa State University

Appendix D

Office of Advanced Scientific Computing Research

The primary mission of the Advanced Scientific Computing Research (ASCR) program 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 research projects described in this report:

Michael R. Strayer
Associate Director, ASCR

Melea Baker
Administrative Specialist

Nancy White
Program Analyst

Julie Scott
Financial Management Specialist

Norman Kreisman
Senior Advisor

Jon Bashor
Senior Advisor

Daniel Hitchcock
Acting Director, Facilities Division

Barbara Helland
Computational Scientist and NERSC Program Manager

Sally McPherson
Program Support Specialist

David Goodwin
Physical Scientist

Vincent Dattoria
General Engineer

Walter Polansky
Acting Director, Computational Science Research and Partnerships (SciDAC) Division

Terry Jones
Program Support Specialist

Teresa Beachley
Program Support Assistant

Fred Johnson
Senior Technical Manager for Computer Science

Robert Lindsay
Computer Scientist

Christine Chalk
Physical Scientist

George Seweryniak
Computer Scientist

Yukiko Sekine
Computer Scientist

Gary Johnson
Computer Scientist

Anil Deane
Mathematician

Mark Sears
Mathematician

David Brown
Detailee

Appendix E

Advanced Scientific Computing Advisory Committee

The Advanced Scientific Computing Advisory Committee (ASCAC) provides valuable, independent advice to the Department of Energy on a variety of complex scientific and technical issues related to its Advanced Scientific Computing Research program. ASCAC's recommendations include advice on long-range plans, priorities, and strategies to address more effectively the scientific aspects of advanced scientific computing including the relationship of advanced scientific computing to other scientific disciplines, and maintaining appropriate balance among elements of the program. The Committee formally reports to the Director, Office of Science. The Committee primarily includes representatives of universities, national laboratories, and industries involved in advanced computing research. Particular attention is paid to obtaining a diverse membership with a balance among scientific disciplines, institutions, and geographic regions.

Jill P. Dahlburg, Chair
Naval Research Laboratory

Robert G. Voigt, Co-Chair
College of William and Mary

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

Horst D. Simon
Lawrence Berkeley National Laboratory

Ellen B. Stechel
Sandia National Laboratories

Rick L. Stevens
Argonne National Laboratory

Virginia Torczon
College of William and Mary

Thomas Zacharia
Oak Ridge National Laboratory

Appendix F

Acronyms and Abbreviations

ACM	Association for Computing Machinery
AEI	Albert-Einstein-Institut (Germany)
AFOSR	Air Force Office of Scientific Research
AHPCRC	Army High Performance Computing Research Center
Al	Aluminum
ALD	Associate Laboratory Director
ALCF	Argonne Leadership Computing Facility
AMD	Advanced Micro Devices, Inc.
AMR	Adaptive mesh refinement
API	Application programming interface
ASCAC	Advanced Scientific Computing Research Advisory Committee
ASCR	Office of Advanced Scientific Computing Research (DOE)
ATC	Alcoa Technical Center
BER	Office of Biological and Environmental Research (DOE)
BES	Office of Basic Energy Sciences (DOE)
BG/L	Blue Gene/L (IBM computer)
BOINC	Berkeley Open Infrastructure for Network Computing
BVSS	Best value source selection
CAF	Co-Array FORTRAN
CAM	Community Atmospheric Model
CCSM	Community Climate System Model
CCT	Center for Computation and Technology (Louisiana State University)
CFD	Computational fluid dynamics
CISM	Center for Integrated Space Weather Modeling
CMB	Cosmic microwave background
CMPD	Center for Multiscale Plasma Dynamics
CNRS	Centre National de la Recherche Scientifique (France)
COBE	Cosmic Background Explorer satellite
CPU	Central processing unit
CRD	Computational Research Division, Lawrence Berkeley National Laboratory
CScADS	Center for Scalable Application Development Software
DNA	Deoxyribonucleic acid
DOD	Department of Defense
DOE	U.S. Department of Energy
DRAM	Dynamic random access memory
EPSRC	Engineering and Physical Sciences Research Council (UK)
ERDC	Engineering Research and Development Center
ESP	Effective System Performance benchmark
FEMA	Federal Emergency Management Agency
FES	Office of Fusion Energy Sciences (DOE)
FFT	Fast Fourier transform
FFTW	Fastest Fourier Transform in the West (C subroutine library)
FPGA	Field programmable gate array
GASNet	Global Address Space Networking
GBMF	Gordon and Betty Moore Foundation
GeV	Giga (one billion) electron volts
Gflop/s	Giga (one billion) floating point operations per second
GHI	GPFS-HPSS Integration
GPFS	General Parallel File System (IBM)
GPS	Global Positioning System
GPSC	Center for Gyrokinetic Particle Simulations of Turbulent Transport in Burning Plasmas and Multiscale Gyrokinetics
GSFC	Goddard Space Flight Center (NASA)
GTC	Gyrokinetic Toroidal Code
HEP	Office of High Energy Physics (DOE)
HP	Hewlett-Packard
HPC	High performance computing
HPF	High Performance FORTRAN
HPSS	High Performance Storage System

IEEE	Institute of Electrical and Electronics Engineers
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)
INSU	Institut National des Sciences de l'Univers (France)
I/O	Input/output
IPM	Integrated Performance Monitoring
ISF	Israel Science Foundation
ITER	A multinational tokamak experiment to be built in France (Latin for "the way")
ITG	Ion temperature gradient
JGI	Joint Genome Institute (DOE)
JINA	Joint Institute for Nuclear Astrophysics
JSPS	Japanese Society for the Promotion of Science
KB	Kilobyte
LBNL	Lawrence Berkeley National Laboratory
LCF	Leadership Computing Facility, Oak Ridge National Laboratory
LCRC	Laboratory Computing Resource Center, Argonne National Laboratory
LLNL	Lawrence Livermore National Laboratory
LOASIS	Laser Optical Accelerator Systems Integrated Studies
LONI	Louisiana Optical Network Initiative
LSU	Louisiana State University
MATLAB	A numerical computing environment and programming language (short for "matrix laboratory")
MB	Megabyte
MPI	Message Passing Interface
NAIC	National Astronomy and Ionosphere Center
NASA	National Aeronautics and Space Administration
NCAR	National Center for Atmospheric Research
NCHC	National Center for High-Performance Computing (Taiwan)
NERSC	National Energy Research Scientific Computing Center
NGF	NERSC Global Filesystem
NIM	NERSC Information Management system
NLCF	National Leadership Computing Facility (Oak Ridge National Laboratory)
NOAA	National Oceanographic and Atmospheric Administration
nm	Nanometer
NOOMSRC	Naval Oceanographic Office Major Shared Resource Center
NP	Office of Nuclear Physics (DOE)
NSF	National Science Foundation
O	Oxygen
ONR	Office of Naval Research
OSF	Oakland Scientific Facility (LBNL)
OSG	Open Science Grid
OSKI	Optimized Sparse Kernel Interface Library
PC	Personal computer
PDSF	Parallel Distributed Systems Facility (NERSC)
PERCU	Performance, effectiveness, reliability, consistency, and usability
PERI	Performance Engineering Research Institute (SciDAC)
PGAS	Partitioned global address space
PGF	Production Genome Facility (JGI)
PIC	Particle-in-cell
PNAS	Proceedings of the National Academy of Sciences
PNC	Programme National de Cosmology (France)
RAMP	Research Accelerator for Multiple Processors
RDMA	Remote direct memory access
SC	Office of Science (DOE)
SciDAC	Scientific Discovery through Advanced Computing (DOE)
SDM	Scientific Data Management Center
SDSA	Science-Driven System Architecture (NERSC)
SDSC	San Diego Supercomputer Center
SEM	Scanning electron micrograph
SETI	Search for Extraterrestrial Intelligence
SIGGRAPH	ACM Special Interest Group on Computer Graphics and Interactive Techniques
SIMD	Single instruction, multiple data
SMP	Symmetric multiprocessing (or multiprocessor)
SPE	Synergistic processing element
SPR	Surface plasmon resonance
SPU	Synergistic processing unit
SSAI	Science Systems and Applications, Inc.
SSH	Secure Shell protocol
SSP	Sustained System Performance benchmark
SST	Sea surface temperature
STEM	Scanning transmission electron microscopy
STI	Sony, Toshiba, and IBM
TACC	Texas Advanced Computing Center
TB	Terabyte
teraflop/s	Tera (one trillion) floating point operations per second
Tflop/s	See teraflop/s
TOPS	Terascale Optimal PDE Simulations Center
UC	University of California
UCAR	University Corporation for Atmospheric Research
UIUC	University of Illinois, Urbana-Champaign
UPC	Unified Parallel C
UPS	Uninterruptible power supply
USACOE	U.S. Army Corps of Engineers
Y	Yttrium

For more information about NERSC, contact:

Jon Bashor
NERSC Communications
Berkeley Lab, MS 50B4230
1 Cyclotron Road
Berkeley, CA 94720-8148

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

NERSC's web site:
http://www.nersc.gov/

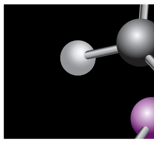

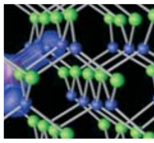

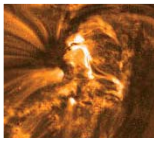

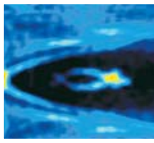

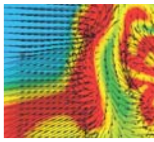

CSO 13541

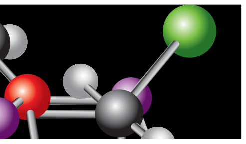

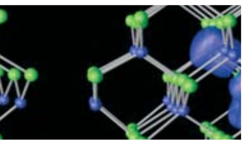

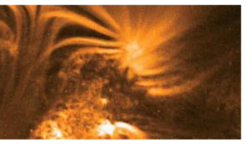

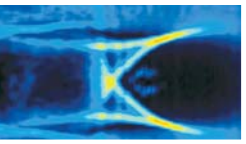

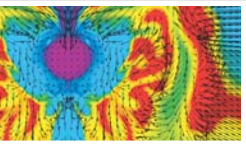

Published by the Berkeley Lab Creative Services Office in collaboration with NERSC researchers and staff.

Editor and principal writer: John Hules.

Contributing writers: Jon Bashor, Ucilia Wang, Lynn Yarris, Paul Preuss, and Wes Bethel (Berkeley Lab); David Hosansky, NCAR; and Laura Shea, Northeastern University.

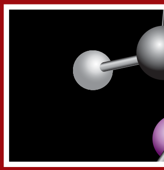

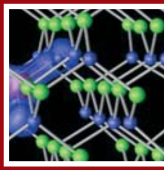

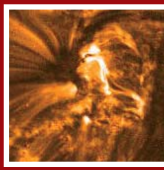

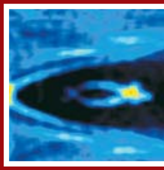

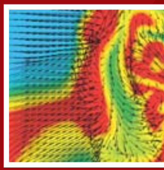

Lawrence Berkeley National Laboratory / 1 Cyclotron Road / Berkeley CA / 94720
LBNL-63342