
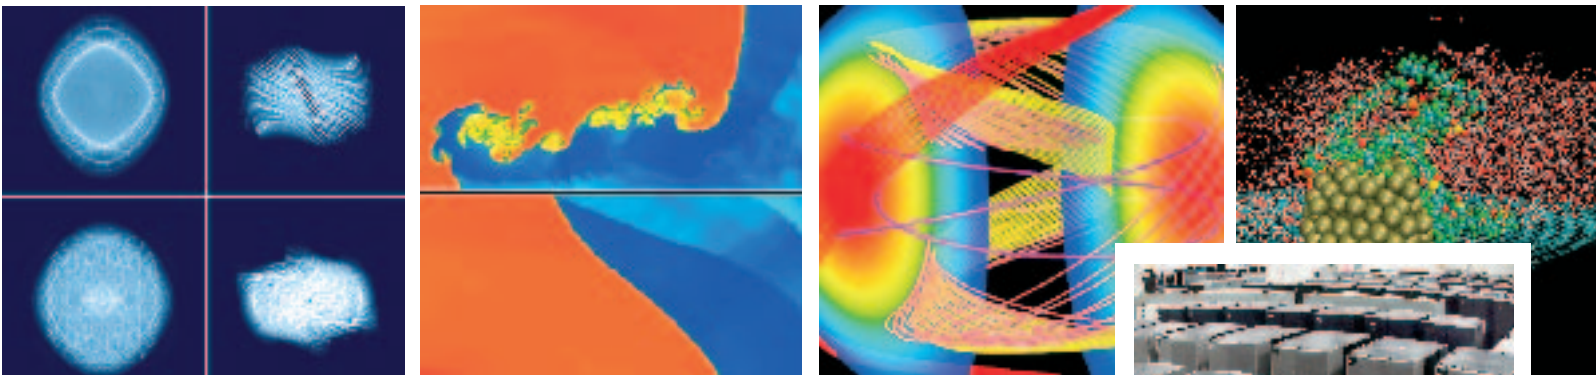


NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

2003 ANNUAL REPORT







NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

2003 ANNUAL REPORT



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-AC 03-76SF00098.

LBL-54267, January 2004

TABLE OF CONTENTS

| | | | |
|--|----|--|----|
| THE YEAR IN PERSPECTIVE | 2 | RESEARCH NEWS | 47 |
| ADVANCES IN COMPUTATIONAL SCIENCE | 5 | Accelerator Physics | 47 |
| Astrophysics: Computing the Big Picture | 7 | Applied Mathematics | 48 |
| Universe: The Movie | 8 | Astrophysics | 50 |
| Simulation Matches Historic | | Chemical Sciences | 52 |
| Gamma-Ray Burst. | 11 | Climate Change Research. | 55 |
| Nearby Supernova Factory Churns | | Computer Science | 56 |
| Out Discoveries. | 14 | Fusion Energy Sciences. | 57 |
| Fusion Energy: Improved Algorithms Shift | | High Energy Physics. | 62 |
| Plasma Simulations into High Gear. | 17 | Materials Sciences | 63 |
| Full Waves, Faster Runs for AORSA | 18 | Medical and Life Sciences. | 66 |
| SuperLU Solver Supercharges | | Nuclear Physics | 68 |
| NIMROD. | 20 | THE NERSC CENTER | 71 |
| AMR Methods Accelerate | | Clients, Sponsors, and Advisors | 71 |
| MHD Simulations | 22 | Capability and Scalability | 73 |
| Biology: Deciphering Nature's | | Comprehensive Scientific Support. | 78 |
| Assembly Manuals. | 25 | Connectivity. | 81 |
| ProteinShop: Solving Protein | | Advanced Development. | 83 |
| Structures from Scratch | 26 | <i>Appendix A: NERSC Policy Board</i> | 87 |
| Mapping the Universe of Protein | | <i>Appendix B: NERSC Computational</i> | |
| Structures | 29 | Review Panel | 88 |
| DOE Facility Plan Sees Major Role for | | <i>Appendix C: NERSC Users Group</i> | |
| Computing. | 32 | Executive Committee. | 90 |
| Nanoscience: Modeling a Realm | | <i>Appendix D: Supercomputing Allocations</i> | |
| Where Everything Is Different | 33 | Committee | 91 |
| Making Fuel Cells More Efficient | | <i>Appendix E: Office of Advanced Scientific</i> | |
| and Affordable. | 34 | Computing Research | 92 |
| Climate Modeling: Predicting the Results | | <i>Appendix F: Advanced Scientific Computing</i> | |
| of a Global Experiment. | 37 | Advisory Committee | 93 |
| Fingerprints in the Sky | 38 | <i>Appendix G: Organizational Acronyms</i> | 94 |
| Applied Mathematics: The Bridge between | | | |
| Physical Reality and Simulations. | 41 | | |
| Algorithms in the Service of Science | 42 | | |

THE YEAR IN PERSPECTIVE

Supercomputing continues to evolve at a rapid rate, and events at NERSC in 2003 demonstrate this quite clearly. For a few months we could claim to run the largest unclassified system in the U.S. But as we all know, these superlatives don't last. What does last are the scientific accomplishments obtained using large systems.

It is with great satisfaction that we write this introduction to a report that describes some truly astounding computational results that were obtained with the help of NERSC in 2003.

In July 2003 Dr. Raymond Orbach, Director of the DOE Office of Science, while testifying about high-end computing before the U.S. Congress, explicitly mentioned NERSC's role in the discovery that the expansion of the Universe is accelerating, as well as the recent U.S. decision to rejoin the international collaboration to build the ITER fusion reactor. "What changed," he said, "were simulations that showed that the new ITER design will in fact be capable of achieving and sustaining burning plasma." From accom-



Horst D. Simon

plishments such as these, it is clear that we really have entered the new era where computational science is an equal partner with theory and experiment in making scientific progress.

At NERSC the biggest change in 2003 was the upgrade of our current Seaborg platform from 5 to 10 Tflop/s peak performance. This created one of the largest systems in the U.S. dedicated to basic computational science. It now gives NERSC users around-the-clock access to an unprecedented 6,556 processors, coupled with one of the largest memory systems anywhere. DOE-supported computational scientists continue to have access to one of the best possible resources to further the DOE mission in the basic sciences.

The upgrade to 10 Tflop/s went very smoothly, and the new system

came online one month earlier than planned. The first applications test results on the upgraded Seaborg were very impressive: not only could users quickly scale several applications to 4,096 processors, but they were also able to achieve very high sustained performance rates, more than 60% of peak in a few cases. The speed of the upgrade process and the new applications results are a testimony to the expert staff at NERSC. We have again demonstrated our capability to quickly field and use a new system productively.

The expanded computational power of Seaborg made it possible to initiate a number of new projects. In the first half of 2003, NERSC started the scalability program, which gave users the opportunity to evaluate the scaling of their applications.

For many users, Seaborg became the first platform on which they could explore calculations using several thousand processors.

New opportunities for computational science breakthroughs will hopefully arise in 2004 through the new Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program at NERSC. The goal of the program is to support a small number of computationally intensive, large-scale research projects that can make high-impact scientific advances through the use of a substantial allocation of computer time and data storage at NERSC. We just completed the selection of the first three INCITE projects, and we hope to announce breakthrough results from these projects in next year's annual report.



William T. C. Kramer

One of the signposts of change in 2002 was the launch of the Earth Simulator system in Japan, which reenergized the high-end computing community in the U.S., resulting in a number of new activities in 2003 that made a profound impact on computational science and high performance computing. Both the High End Computing Revitalization Task Force (HECRTF) report and the Science-Based Case for Large-Scale Simulation (SCaLeS) report demonstrated that the high-end computing community in the U.S. is well positioned for future developments and growth in supercomputing. The SCaLeS report gives us a glimpse of what simulations are possible at sustained speeds in the range of tens of teraflop/s, and the HECRTF report lays out the critical research issues that we need to address in order to reach petascale computing performance. This is exciting news, because we are thinking big again in supercomputing. We


are thrilled that the DOE computational science community, in particular our NERSC users, are ready to take the lead in moving to the next level of computational science.

One small step towards this bright future will happen at NERSC in the next few months. We will have the opportunity to acquire a new computing system that will take over some of the existing user workload from Seaborg, making even more time available on Seaborg for large-scale capability computing. This will be just an intermediate step towards our next big acquisition with NERSC 5. We are looking forward to another year of both scientific and computing accomplishments at NERSC. As always, this progress would not be possible without the NERSC staff, who continue to tirelessly dedicate their time, skill, and effort to make NERSC the best scientific computing resource in the world. Our special thanks to all of you.

Horst D. Simon
NERSC Center Division Director

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

ADVANCES IN COMPUTATIONAL SCIENCE



ASTROPHYSICS

COMPUTING THE BIG PICTURE

From ancient times, people have tried to interpret what they see in the night sky as a way of understanding their place in the cosmos. Today astrophysicists rely on high-end computers, networks, and data storage to manage and make sense of the massive amounts of astronomical data being produced by advanced instruments. Computational simulations performed on NERSC computers are answering questions as general as how structure developed in the Universe, and as specific as how supernovae produce gamma-ray bursts. NERSC's real-time integration of data gathering with data analysis is enabling collaborators like the Nearby Supernova Factory to step up the pace of discovery. The projects described in this section demonstrate how modern instruments and computers have ushered in a golden age of astrophysics.

UNIVERSE: THE MOVIE

When we look up into a clear night sky, we're looking into the past. Even with the naked eye, we can see the Andromeda galaxy as it was 2.3 million years ago. Earth- and space-based telescopes provide us with a plethora of snapshots of the history of the Universe. But none of these snapshots show how the Universe or any of the objects within it evolved.

“Turning the snapshots into movies is the job of theorists,” says Joel Primack, Professor of Physics at the University of California, Santa Cruz. “With computational simulations, we can run the history of stars, galaxies, or the whole Universe forward or backward, testing our theories and finding explanations for the wealth of new data we’re getting from instruments on the ground and especially in space.”

Primack is one of the originators and developers of the theory of cold dark matter, which has become the standard theory of structure formation in the Universe because its predictions generally


match recent observations. His research group uses NERSC’s Seaborg computer to create simulations of galaxy formation. Their work here has two major thrusts: simulations of large-scale structure formation in the Universe, and simulations of galaxy interactions.

How did the Universe evolve from a generally smooth initial state after the Big Bang (as seen in the cosmic microwave background radiation) to the chunky texture that we see today, with clusters, filaments, and sheets of galaxies? That is the question of large-scale structure formation. In the cold dark matter theory, structure grows hierarchically

by gravitational attraction, with small objects merging in a continuous hierarchy to form more and more massive objects.

However, “objects” in this context refers not only to ordinary “baryonic” matter (composed of protons and neutrons), but also to cold dark matter—cold because it moves slowly, and dark because it cannot be observed by its electromagnetic radiation. In fact, the latest simulation results¹ suggest that galaxy clustering is primarily a result of the formation, merging, and evolution of dark matter “halos.”

The concept of dark matter



halos came from the realization that the dark matter in a galaxy extends well beyond the visible stars and gas, surrounding the visible matter in a much larger halo; the outer parts of galaxies are essentially all dark matter. Scientists came to this conclusion because of the way stars and galaxies move, and because of a phenomenon called gravitational lensing: when light passes through the vicinity of a galaxy, the gravity of the dark matter halo bends the light to different wavelengths, acting like imperfections in the glass of a lens.

In the view of Primack and his colleagues, the large-scale structure of the Universe is determined primarily by the invisible evolution and interactions of massive halos and subhalos, with ordinary matter being swept along in the gravitational ebbs and flows. They have used Seaborg to create the highest-resolution simulations of dark matter halos to date, producing 10 terabytes of output and models with

around 100,000 halos. These models can be used to predict the types and relative quantities of galaxies, as well as how they interact. Figure 1 shows the simulated distribution of dark matter particles and halos at two different stages of structure formation. Interestingly, there are more subhalos in this model than there are visible objects in corresponding observational data. Observations from the DEEP (Deep Extragalactic Evolutionary Probe) and ESO/VLT (European Southern Observatory/Very Large Telescope) sky surveys should soon confirm or refute the predictions of this model.

Elliptical galaxies, galactic bulges, and a significant fraction of all the stars in the Universe are thought to be shaped largely by galaxy interactions and collisions, with dark matter still playing the major role but with bright matter interactions adding some details. For the first half of the Universe's history, stars were formed mostly in

FIGURE 1

Simulated distribution of dark matter particles (points) and halos (circles) at redshift 3 (left) and redshift 0 (right). (Higher redshift signifies greater distance and age.) Image: Kravtsov et al., *astro-ph/0308519*.

¹ A. V. Kravtsov, A. A. Berlind, R. H. Wechsler, A. A. Klypin, S. Gottlöber, B. Allgood, and J. R. Primack, "The dark side of the halo occupation distribution," *Astrophysics Journal* (submitted); *astro-ph/0308519* (2003).

spectacular “starburst” events resulting from the compression of gases by the gravitational tides of interacting galaxies. A starburst is a relatively brief event on a cosmic scale, lasting perhaps 100 million years, while the merger of two galaxies might take 2 to 3 billion years. The Hubble space telescope in 1996 captured spectacular images of two colliding galaxies and the resulting starburst (Figure 2).

By creating galaxy interaction simulations (Figure 3) and comparing them with images and data from a variety of telescopes and spectral bands, Primack’s research team tries to understand the evolution of galaxies, including phenomena such as starbursts. They recently focused their attention on how much gas is converted to stars during the course of a major galaxy

merger, and found that previous studies had overestimated the amount of starburst. Their results showed that galaxy mergers produce large outflows of gas that inhibit star formation for a while after the merger, an effect that they are adding to their models of galaxy formation.

Like many physicists and astronomers, Primack believes that we are now living in the golden age of cosmology. “Cosmology is under-

going a remarkable evolution, driven mainly by the wonderful wealth of new data,” he says. “The questions that were asked from the beginning of the 20th century are now all being answered. New questions arise, of course, but there’s a good chance that these answers are definitive.”

Research funding: [HEP](#), [NSF](#), [NASA](#)
(Organizational acronyms are spelled out in Appendix G)

FIGURE 2

The Hubble telescope uncovered over 1,000 bright, young star clusters bursting to life in the heart of a pair of colliding galaxies. The picture in the upper left provides a ground-based telescopic view of the two galaxies, called the Antennae because a pair of long tails of luminous matter, formed by the gravitational tidal forces of their encounter, resembles an insect’s antennae. The galaxies are located 63 million light-years away in the southern constellation Corvus. Hubble’s close-up view provides a detailed look at the “fireworks” in the center of the collision. The respective cores of the twin galaxies are the orange blobs, crisscrossed by fila-

ments of dark dust. A wide band of chaotic dust stretches between the cores of the two galaxies. The sweeping spiral-like patterns, traced by bright blue star clusters, are the result of a firestorm of star birth that was triggered by the collision. *Image: Brad Whitmore (Space Telescope Science Institute) and NASA.*



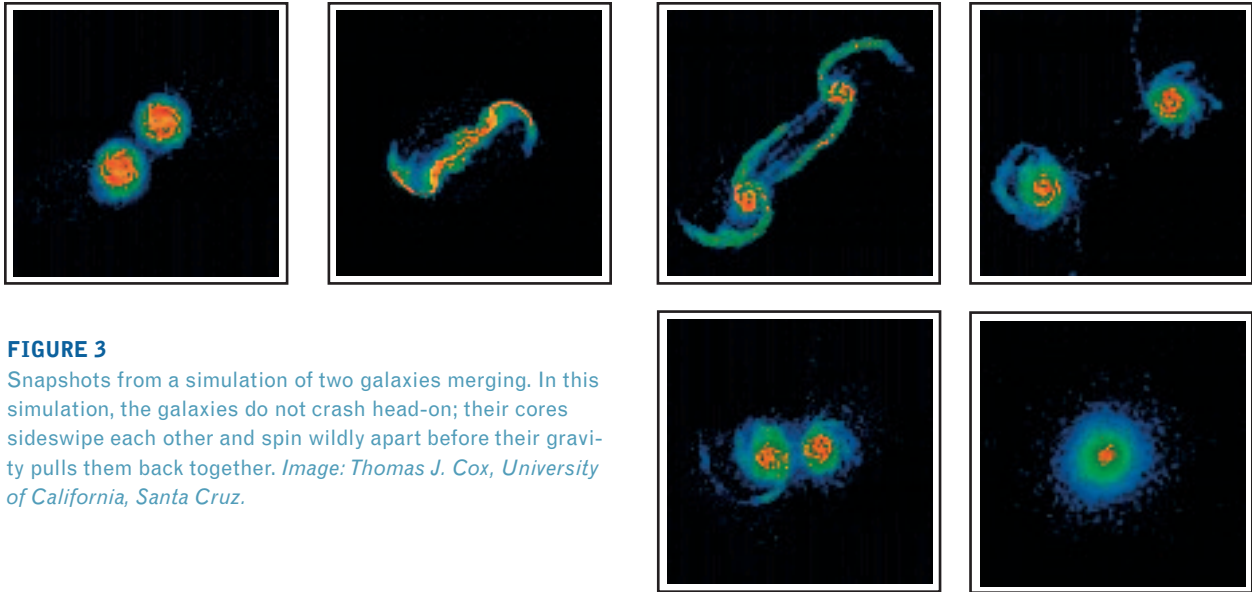


FIGURE 3
 Snapshots from a simulation of two galaxies merging. In this simulation, the galaxies do not crash head-on; their cores sideswipe each other and spin wildly apart before their gravity pulls them back together. *Image: Thomas J. Cox, University of California, Santa Cruz.*

SIMULATION MATCHES HISTORIC GAMMA-RAY BURST

After three decades of scientific head-scratching, the origins of at least some gamma-ray bursts (GRBs) are being revealed, thanks to a new generation of orbiting detectors, fast responses from ground-based robotic telescopes, and a new generation of computers and astrophysics software.

A GRB detected on March 29, 2003 has provided enough information to eliminate all but one of the theoretical explanations of its origin. Computational simulations based on that model were already being developed at the NERSC Center at when the discovery was made.

Gamma radiation from outer space is blocked by the Earth's atmosphere. But if it were not, and

if we could see gamma rays with our eyes, about once a day, somewhere in the world, people would see a spectacular flash, a hundred million times brighter than a supernova. These gamma-ray bursts, first discovered in the 1960s by satellites looking for violations of the Nuclear Test Ban Treaty, are the most energetic events in the Universe, but they are also among the most elusive. GRBs appear ran-

domly from every direction. They do not repeat themselves. And they last from only a few milliseconds to a few minutes—astronomers consider a GRB long if it lasts longer than 2 seconds.

With data so hard to pin down, scientists for a long time could only speculate about what causes GRBs. Speculations ranged from the intriguing (comet/anti-comet annihilation) to the far-fetched (inter-

stellar warfare). By 1993, 135 different theories on the origin of GRBs had been published in scientific journals. But in recent years, three theoretical models emerged as front-runners—one involving colliding neutron stars, and the other two involving supernovae, the “supernova” and “collapsar” models.

In the collapsar model (introduced in 1993 by Stan Woosley, Professor and Chair of Astronomy and Astrophysics at the University of California, Santa Cruz), the iron core of an aging star runs out of fuel for nuclear fusion and collapses of its own weight, creating a black hole or a dense neutron star. Material trying to fall onto this object forms a hot swirling disk and a narrow jet, which shoots out of the star in less than ten seconds at nearly the speed of light. When this jet erupts into interstellar space, it creates a “fireball” of gamma rays—the highest energy, shortest wavelength form of electromagnetic radiation. The rest of the star explodes as a supernova, but that event is eclipsed by the brighter GRB.

In the late 1990s, GRB research took a great leap forward when a new generation of orbiting detectors were launched, including the Italian-Dutch satellite BeppoSAX and NASA’s High-Energy Transient Explorer (HETE). These satellites were designed to provide quick notification of GRB discoveries for immediate follow-up observations by ground-based instruments. In

1997 astronomers discovered that long GRBs leave an afterglow of lower-energy light, such as X rays or visible light, that may linger for months, allowing researchers to pinpoint where the GRB originated and providing important clues about the event that produced it. The presence of iron in the afterglow light strongly suggested star explosions. In 1998, the supernova theories got an additional boost when a GRB and a supernova appeared in the same vicinity at roughly the same time; but the data was inconclusive and some scientists remained skeptical about the connection. Several similar suggestive but inconclusive events in subsequent years divided astronomers into two camps on the issue of associating GRBs with supernovae.

The skepticism was dispelled on March 29, 2003, when HETE detected an unusually bright and close GRB—only 2.6 billion light years from Earth instead of the typical 10 billion. The discovery triggered a swarm of observations that found the unmistakable spectral signature of a supernova in the afterglow, as reported in the June 19, 2003 issue of *Nature*.² This event, named GRB030329 after its detection date, was dubbed the “Rosetta stone” of GRBs because it conclusively established that at least some long GRBs come from supernovae, and it singled out the collapsar model as the one that best fits the data collected so far. The March 29 burst’s afterglow was so bright that it

allowed astronomers to study the event in unprecedented detail; they even joked about its casting shadows. The Hubble Space Telescope will be able to study the afterglow for another year, and radio telescopes may be able to track it longer.

The first 3D computational simulations of jet formation and breakout in the collapsar model (Figures 4 and 5) were already being conducted on Seaborg by Woosley and Weiqun Zhang, a Ph.D. candidate at UC Santa Cruz, under the sponsorship of the DOE Office of Science’s SciDAC Supernova Science Center (<http://www.super-sci.org/>), one of two SciDAC (Scientific Discovery through Advanced Computing) programs focusing on developing new computational methods for understanding supernovae. The goal of the Supernova Science Center is to discover the explosion mechanism of supernovae through numerical simulation.

Zhang, Woosley, and colleagues had been modeling jet formation after core collapse in both two and three dimensions for a few years when observations from GRB030329 validated their work and made it newsworthy. “We weren’t utterly surprised,” Woosley said, “because the evidence associating GRBs with supernovae had been accumulating for several years. But we weren’t totally confident, either. When the light spectrum from the March 29 burst confirmed

² J. Hjorth, J. Sollerman, P. Møller, J. P. U. Fynbo, S. E. Woosley, et al., “A very energetic supernova associated with the γ -ray burst of 29 March 2003,” *Nature* **423**, 847 (2003).

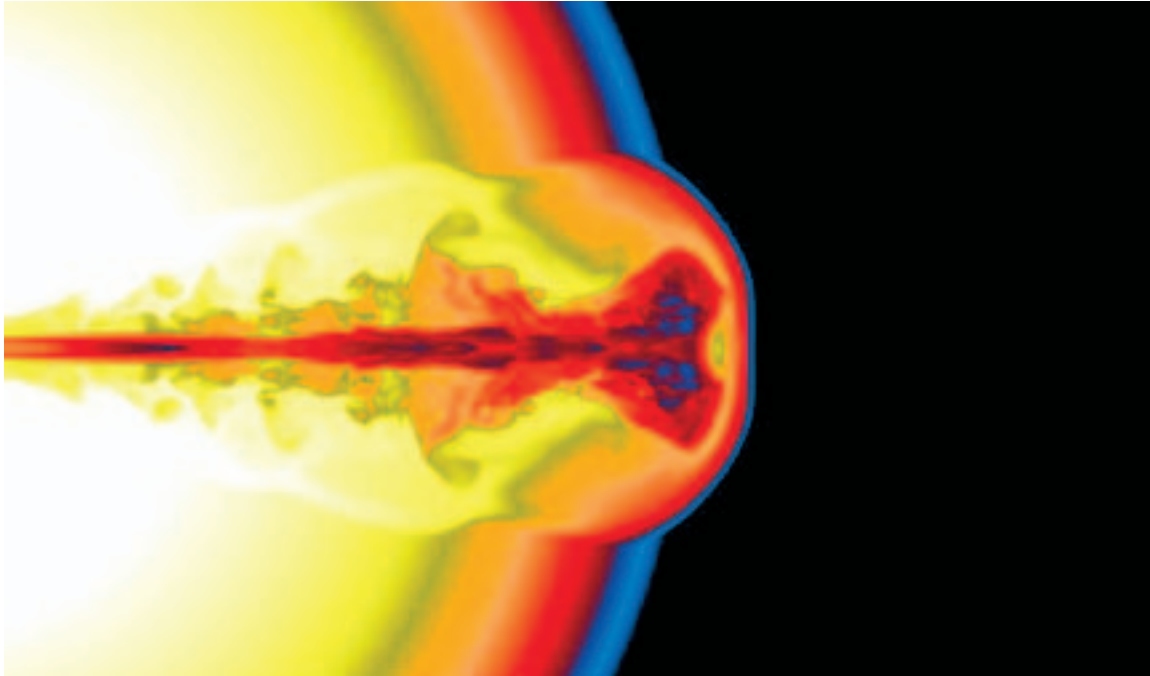


FIGURE 4

This image from a computer simulation of the beginning of a gamma-ray burst shows the jet 9 seconds after its creation at the center of a Wolf-Rayet star by the newly formed, accreting black hole within. The jet is now just erupting through the surface of the star, which has a radius comparable to that of the sun. Blue represents regions of low mass concentration, red is denser, and yellow denser still. Note the blue and red striations behind the head of the jet. These are bounded by internal shocks. Image: Weiqun Zhang and Stan Woosley.

that it came from a supernova, that felt good.”

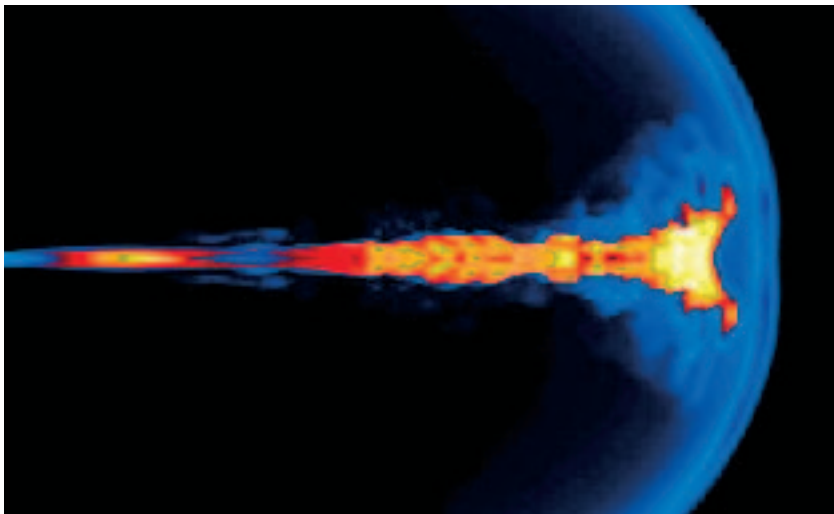
Woosley emphasized, however, that “the data is still way ahead of the theory.” The computational simulations are still incomplete, and the resolution needs to be improved. Upcoming work will include higher-resolution 3D studies of jet stability—which will help scientists interpret the differences in

GRB observations—as well as simulating the full star explosion that accompanies the GRB and the core collapse that precedes it. The project will be challenging, even on a computer as fast as Seaborg.

“This does not mean that the gamma-ray burst mystery is solved,” Woosley added. “We are confident that long bursts involve a core collapse, probably creating a black

FIGURE 5

This image from a computer simulation shows the distribution of relativistic particles (moving near light speed) in the jet as it breaks out of the star. Yellow and orange are very high energy and will ultimately make a gamma-ray burst, but only for an observer looking along the jet (\pm about 5 degrees). Note also the presence of some small amounts of energy in mildly relativistic matter (blue) at larger angles off the jet. These will produce X-ray flashes that may be seen much more frequently. Image: Weiqun Zhang and Stan Woosley.



hole. We have convinced most skeptics. We cannot reach any conclusion yet, however, on what causes short gamma-ray bursts.”

Aside from explaining the mysterious long GRBs, the study of supernovae will fill an essential gap in our knowledge of the Universe,

because supernovae are currently thought to be the source of all the elements heavier than iron.

Because supernovae cannot be recreated in a laboratory, numerical simulation is the only tool available for interpreting observational data and for developing a detailed

understanding of the physical processes involved. NERSC’s IBM supercomputer is currently being used by several research groups studying supernovae.

Research funding: [NP](#), [SciDAC](#)

NEARBY SUPERNOVA FACTORY CHURNS OUT DISCOVERIES

When the Supernova Cosmology Project got up to speed in the mid-1990s, its team of astrophysicists, led by Saul Perlmutter of Lawrence Berkeley National Laboratory, were looking forward to finally getting decent measurements of how much the expansion of the Universe was slowing down.

The team had figured out how to accelerate the discovery of Type Ia supernovae by systematically searching the same patches of sky at intervals, then subtracting the images from one another, following up with more detailed measurements of the remaining bright spots. By comparing the distance of these exploding stars (which all have the same intrinsic brightness) with the redshifts of their home galaxies, researchers could calculate how fast the Universe was expanding at different times in its history.

As so often happens in science and in life, things did not turn out as expected. The researchers got

their high-quality measurements all right, but the data showed that the expansion of the Universe was speeding up, not slowing down as everyone had assumed. The data were analyzed on NERSC computers using different sets of cosmological assumptions, but the results were inescapable. And Australia’s High-z Supernova Search Team, working independently, had reached the same conclusion.

So in 1998, the two teams of researchers made the startling announcement that the expansion of the Universe is accelerating. Before long a new term entered the vocabulary of astrophysics: dark

energy, the unknown force driving the expansion, which accounts for about two-thirds of the total energy density in the Universe. To find out what dark energy is, scientists need more detailed measurements of the Universe’s expansion history. Thus they need even more observations of Type Ia supernovae, and they need to reduce some remaining uncertainty about the uniform brightness of these exploding stars.

Reducing the uncertainty and improving the calibration of Type Ia supernovae are among the chief goals of the Nearby Supernova Factory (SNfactory), a spinoff of the Supernova Cosmology Project, led

by Berkeley Lab physicist Greg Aldering. At present, astronomers can determine the distance to a well-studied Type Ia supernova with an accuracy of 5 percent, an accuracy the SNfactory expects to improve. “Our motive is to establish a much better sample of nearby Type Ia supernovae, against which the brightness of distant supernovae can be compared to obtain relative distances,” says Aldering.

What researchers want to measure, says Aldering, also includes “the intrinsic colors of a Type Ia at every stage, so we’ll know the effects of intervening dust. And we want to know what difference the ‘metallicity’ of the home galaxy makes—that is, the presence of elements heavier than helium.”

To accomplish these goals, the SNfactory needs to discover and make detailed observations of 300 to 600 low-redshift supernovae, many more than have been studied so far. Automation and tight coordination of the search and follow-up stages are absolutely essential to achieving these numbers. During its first year of operation, the SNfactory found 34 supernovae. “This is the best performance ever for a ‘rookie’ supernova search,” Aldering says. “In our second year, we have shown we can discover supernovae at the rate of nine a month, a rate other searches have reached only after years of trying.”

This remarkable discovery rate is made possible by a high-speed data link, custom data pipeline software, and NERSC’s data handling and storage capacity. So far the SNfactory has processed a quarter-



million images and archived 6 terabytes (trillion bytes) of compressed data at NERSC—one of the few centers with an archive large enough to store this much data. “The SNfactory owes much of its success to NERSC’s ability to store and process the vast amounts of data that flow in virtually every night,” says Aldering.

The SNfactory’s custom-developed data pipeline software, developed by UC Berkeley graduate student Michael Wood-Vasey, manages up to 50 gigabytes (billion bytes) of data each night from wide-field cameras built and operated by the Jet Propulsion Laboratory’s Near Earth Asteroid Tracking program (NEAT). NEAT uses remote telescopes at Mount Palomar Observatory in Southern California and at the U.S. Air Force’s Maui Space Surveillance System on Mount Haleakala.

With the help of the High Performance Wireless Research and

FIGURE 6

A special link in the High Performance Wireless Research and Education Network (HPWREN) transmits images from the Near Earth Asteroid Tracking program (NEAT) at Mount Palomar to NERSC for storage and processing for the SNfactory.

Education Network (HPWREN) program at the San Diego Supercomputer Center (SDSC), the SNfactory was able to establish a custom-built, high-speed link with Mount Palomar (Figure 6). Existing links to Maui and to SDSC through the DOE's Energy Sciences Network (ESnet) complete the connection to NERSC.

NEAT sends images of about 500 square degrees of the sky to NERSC each night. The data pipeline software automatically archives these in NERSC's High Performance Storage System (HPSS). NEAT's telescopes revisit the same regions of the sky roughly every six days during a typical 18-day observing period. When a supernova appears in one of those galaxies, the SNfactory can find it using image subtraction software

that can sift through billions of objects. This analysis is done using NERSC's PDSF Linux cluster.


Eventually the pipeline will automate the entire discovery and confirmation process. Once a supernova is discovered from the Palomar or Maui images, follow-up observations will be obtained via remote control of a custom-built dual-beam optical spectrograph (being completed by SNfactory collaborators in France) mounted on the University of Hawaii's 88-inch telescope on Mauna Kea. The Hawaii observations will be shipped by Internet for image processing at a supercomputing center in France and then sent to NERSC for analysis.

A recent major discovery made possible by the SNfactory was the first detection of hydrogen in the form of circumstellar material

around a supernova—in this case, SN 2002ic, discovered near maximum light.³ Researchers have been looking for hydrogen to discriminate between different types of possible progenitor systems to Type Ia supernovae. Large amounts of hydrogen around SN 2002ic suggest that the progenitor system contained a massive asymptotic giant branch star that lost several solar masses of gas in a “stellar wind” in the millennia leading up to the Type Ia explosion. Discoveries like this, which provide a more detailed understanding of the physics of supernovae, will improve the usefulness of supernovae as distance markers on the journey back through the history of the cosmos.

Research funding: [HEP](#), [LBNL](#), [FBF](#), [CNRS](#), [IN2P3](#), [INSU](#), [PNC](#)

³ M. Hamuy, M. M. Phillips, N. B. Suntzeff, J. Maza, L. E. González, M. Roth, K. Krisciunas, N. Morrell, E. M. Green, S. E. Persson, and P. J. McCarthy, “An asymptotic-giant-branch star in the progenitor system of a type Ia supernova,” *Nature* **424**, 651 (2003).



FUSION ENERGY

IMPROVED ALGORITHMS SHIFT PLASMA SIMULATIONS INTO HIGH GEAR

“The computer literally is providing a new window through which we can observe the natural world in exquisite detail.”¹

Nowhere is that statement more true than in the field of plasma science, where for nearly three decades, computational simulations have contributed as much as experiments to the advancement of knowledge. The simulation of a magnetic fusion confinement system involves modeling of the core and edge plasmas, as well as the plasma-wall interactions. In each region of the plasma, turbulence can cause anomalies in the transport of the ionic species in the plasma; there can also be abrupt changes in the form of the plasma caused by large-scale instabilities. Computational modeling of these key processes requires large-scale simulations covering a broad range of space and time scales.

DOE's SciDAC program has brought together plasma physicists, computer scientists, and mathematicians to develop more comprehensive models and more powerful and efficient algorithms that enable plasma simulation codes to take full advantage of the most advanced terascale computers, such as NERSC's Seaborg IBM system. Dramatic improvements in three of these codes are described in this section.

¹ J. S. Langer, chair of the July 1998 National Workshop on Advanced Scientific Computing.

FULL WAVES, FASTER RUNS FOR AORSA

DOE's magnetic fusion energy research program aims to produce energy by the same process that takes place in the sun.

However, on Earth a fusion reaction requires that the fuel be heated to hundreds of millions of degrees, even hotter than in the sun. At these astronomical temperatures, the fuel atoms are torn apart into their constituent electrons and nuclei, forming a state of matter called plasma.

One of the ways to get the fuel this hot is to use intense electromagnetic waves, much as a microwave oven is used to heat food. Waves are also used for other important purposes such as to drive electric currents affecting the plasma magnetic configuration, to force mass flow of the plasma, affecting its stability, and for other plasma control tasks. Therefore, it is important to have a good theoretical understanding of wave behavior and to be able to calculate it accurately. The goal of the SciDAC project “Terascale Computation of Wave-Plasma Interactions in

Multidimensional Fusion Plasmas” is to develop the computational capability to understand the physical behavior of waves in fusion plasmas and to model these wave phenomena accurately enough that their effect on other plasma processes, such as stability and transport of particles and energy, can be understood.

According to project leader Don Batchelor of Oak Ridge National Laboratory, “Because the particles are so hot, they move at speeds almost the speed of light and can travel a distance comparable to a wavelength in the time of a few oscillations of the wave. This motion makes it difficult to calculate how the plasma particles will respond to the waves and how much heating or electric current they will produce.”

“Another challenge,” Batchelor adds, “is that at a given frequency, several kinds of plasma waves can

exist with very different wavelengths and polarizations. A wave launched into the plasma can in a short distance completely change its character to another type of wave, a process called mode conversion. To study these effects, the computer model must have very high resolution to see the small-scale structures that develop, which means that very large computers are needed to solve for the very large number of unknowns in the equations. Also, the computers must be extremely fast in order to obtain the solutions in a reasonable time.”

Until now researchers wishing to calculate the effects of waves in plasma have been forced by computational feasibility to make a difficult choice between restricting consideration to a single dimension or simplifying the physics model. The first choice, treating the plasma geometry as one-dimensional, is akin to tunnel vision. The wave is

computed along a single line through the plasma, but one does not get a picture of what is occurring in the whole plasma cross-section. The second choice eliminates from consideration many of the most important wave processes for today's fusion experiments, which require high frequencies and can have very short wavelengths in some regions of the device.

In a partnership between plasma physicists and computer scientists, this project has increased the resolution and speed of plasma wave solvers to the point that it is possible for the first time to study the mode conversion process in the complicated geometry and the full scale of real fusion devices. Collaborators have developed new wave solvers in 2D and 3D called All-Orders Spectral Algorithm (AORSA) solvers based on a more general formulation of the physics called an integral equation. It is now possible to compute plasma waves across an entire plasma cross-section with no restriction on wavelength or frequency (Figure 1). In this approach the limit on attainable resolution comes only from the size and speed of the available computer.

AORSA uses a fully spectral representation of the wave field in a Cartesian coordinate system. All modes in the spectral representation are coupled, and so computing the solution requires calculating and inverting a very large dense matrix. For example, with 272×272 Fourier modes in two dimensions or $34 \times 34 \times 64$ modes in three dimensions, the system

processes about 220,000 coupled complex equations that require about 788 GB of memory. AORSA-3D took about 358 minutes to run on 1,936 processors of NERSC's Seaborg computer. The efficient computation achieved over 60 percent of available peak performance (over 1.9 teraflop/s or 1,900 billion operations per second).

A new formulation of AORSA-3D transforms the linear system from a Fourier basis to a representation in the real configuration space. This alternative representation presents new opportunities to reduce the overall number of equations used. In one example of $34 \times 34 \times 64$ modes, the number of equations was reduced by more than a factor of 5—from 220,000 to about

40,000. The new reduced system required only 26 GB of memory and was completed in about 13 minutes, with a 100-fold speedup in the linear solver and a 27-fold speedup in overall runtime. The efficiency gained from the real space formulation of AORSA-3D allows higher resolution and more design analysis for future experiments such as the Quasi-Poloidal Stellarator.

"Funding from the SciDAC project has enabled us to assemble a team from multiple institutions to work on a unified set of goals in a way that was not possible before SciDAC," Batchelor says. "Working with computer scientists funded to work directly with our project has made it feasible to study new regimes of wave physics which previously

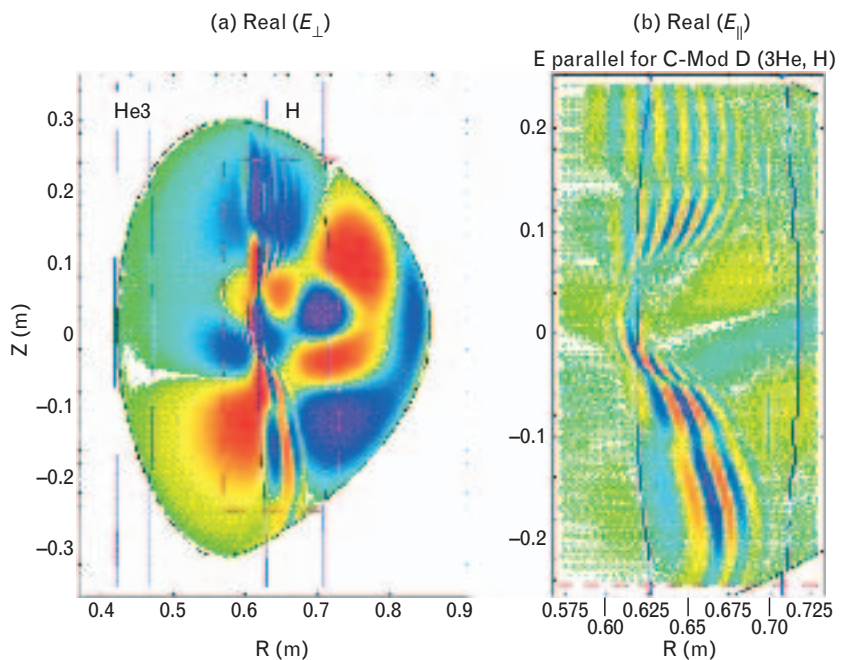


FIGURE 1 Mode conversion in a plasma cross section computed with AORSA. From E. F. Jaeger et al., "Sheared poloidal flow driven by mode conversion in tokamak plasmas," *Phys. Rev. Lett.* **90**, 195001 (2003).

could only be treated in one dimension or with approximate techniques, and to carry out multiple code runs at high resolution for scientific case studies where before it was only possible to perform a single run at minimal resolution.”

“Working with applied mathe-

maticians,” he added, “we are developing new mathematical representations for the wave fields that are more data efficient than the traditional spectral techniques presently employed and offer promise to substantially reduce the size of the computational load to solve plasma wave

problems. And we are making scientific progress on problems of importance to the fusion program by applying the codes we have developed and improved.”

Research funding: [FES](#), [SciDAC](#), [ORNL](#)

SUPERLU SOLVER SUPERCHARGES NIMROD

Developers of the NIMROD code, which is used to simulate fusion reactor plasmas, collaborated with members of the SciDAC Terascale Optimal PDE Simulations Center to implement the SuperLU linear solver software within NIMROD.

As a result, NIMROD runs four to five times faster for cutting-edge simulations of nonlinear macroscopic electromagnetic dynamics—with a corresponding increase in scientific productivity.

The NIMROD project, funded by the DOE Office of Fusion Energy Sciences and the SciDAC Center for Extended Magneto-hydrodynamic Modeling (CEMM), is developing a modern computer code suitable for the study of long-wavelength, low-frequency, nonlinear phenomena in fusion reactor plasmas. These phenomena involve large-scale changes in the shape and motion of the plasma and severely constrain the operation of fusion experiments. CEMM also supports a complementary plasma simulation code, AMRMHD, which uses different mathematical

formulations (see below).

By applying modern computational techniques to the solution of extended magnetohydrodynamics (MHD) equations, the NIMROD team is providing the fusion community with a flexible, sophisticated tool which can lead to improved understanding of these phenomena, ultimately leading to a better approach to harnessing fusion energy. Since the beginning of the project, the NIMROD code has been developed for massively parallel computation, enabling it to take full advantage of the most powerful computers to solve some of the largest problems in fusion. The team’s primary high-end computing resource is Seaborg at NERSC.

NIMROD’s algorithm requires solution of several large sparse matrices in parallel at every time

step in a simulation. The stiffness inherent in the physical system leads to matrices that are ill-conditioned, since rapid wave-like responses provide global communication within a single time-step. The preconditioned conjugate gradient (CG) solver that has been used was the most computationally demanding part of the algorithm, so Carl Sovinec of the NIMROD team consulted with the SciDAC Terascale Optimal PDE Simulations (TOPS) Center to find a replacement.

Conversations with David Keyes of Old Dominion University, Dinesh Kaushik of Argonne National Laboratory, and Sherry Li and Esmond Ng of Lawrence Berkeley National Laboratory pointed to SuperLU as a possibly more efficient matrix solver for NIMROD. SuperLU is a library of software for

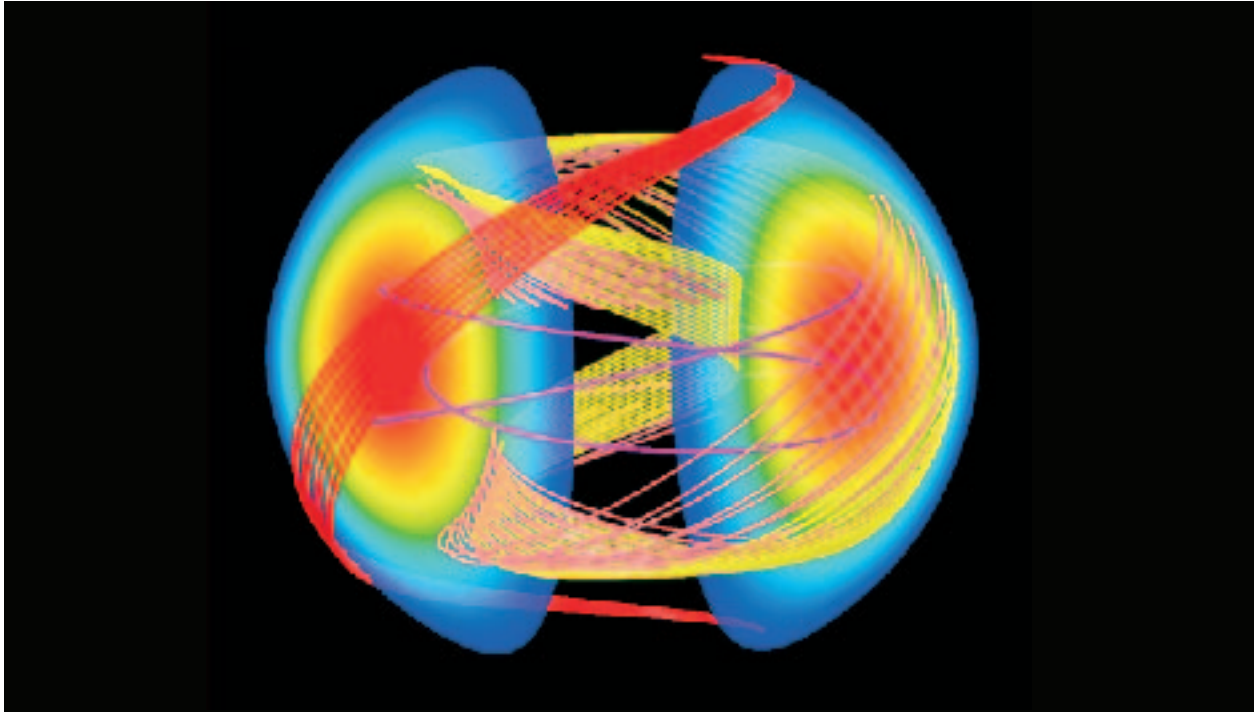


FIGURE 2

Full 3D numerical simulation of plasma particle drift orbits in a tokamak. From Charlson C. Kim, Scott E. Parker, and the NIMROD team, "Kinetic particles in the NIMROD fluid code," 2003 Sherwood Fusion Theory Conference, Corpus Christi, Texas.

solving nonsymmetric sparse linear systems in parallel using direct methods.

It took less than a month to implement, test, and release SuperLU into the full NIMROD production code, and the performance improvements were dramatic. For two-dimensional linear calculations of MHD instabilities, NIMROD runs 100 times faster with SuperLU than it does with the CG solver. While linear calculations do not require supercomputer resources, they are extremely useful for preliminary explorations that help determine which cases require in-depth nonlinear simulations. For linear problems, this performance improvement makes NIMROD competitive with special-purpose linear codes.

For cutting-edge three-dimen-

sional, nonlinear tokamak simulations (Figure 2), which require supercomputing resources, NIMROD with SuperLU runs four to five times faster than before. This improved efficiency yields four to five times more physics results for the same amount of time on the computer—a major improvement in scientific productivity.

The nonlinear simulations accumulate relatively small changes in the matrix elements at each time step, but there are no changes to the sparsity pattern. The NIMROD implementation allows SuperLU to reuse its factors repeatedly until the matrix elements accumulate a significant change and refactoring becomes necessary. Refactoring makes the performance improvement less dramatic in nonlinear simulations than in linear calcula-

tions, but it is still very significant.

The net result is that computationally demanding simulations of macroscopic MHD instabilities in tokamak plasmas, which until now were considered too difficult, have become routine.

Research funding: FES, SciDAC

AMR METHODS ACCELERATE MHD SIMULATIONS

The annual Supercomputing conference held every November is well known as a sort of international watering hole, where many of the world's leading experts in high-performance computing gather for a week to take stock of the competition, exchange ideas, and make new connections.

At SC2000 in Dallas, Phil Colella, head of the Applied Numerical Algorithms Group at Lawrence Berkeley National Laboratory, and Steve Jardin, co-leader of the Computational Plasma Physics Group at Princeton Plasma Physics Laboratory, were both scheduled to give talks in the Berkeley Lab booth. Colella discussed “Adaptive mesh refinement research and software at NERSC,” where Jardin has conducted his scientific computing for years. Jardin gave a presentation on “A Parallel Resistive MHD Program with Application to Magnetic Reconnection.”

The scientist and the mathematician got to talking between their presentations and one thing led to another. They began an

informal collaboration which was soon formalized under the auspices of SciDAC—Jardin is principal investigator for the Center for Extended Magnetohydrodynamic Modeling (CEMM), while Colella is PI for the Applied Partial Differential Equations Center (APDEC). Jardin’s group was able to incorporate the CHOMBO adaptive mesh refinement (AMR) code developed by Colella’s group into a new fusion simulation code, which is now called the Princeton AMRMHD code. “Using the AMR code resulted in a 30 times improvement over what we would have had with a uniform mesh code of the highest resolution,” Jardin says.

The AMRMHD code, devel-

oped in conjunction with Princeton researcher Ravi Samtaney and Berkeley Lab researcher Terry Ligocki, is already producing new physics results as well. It powered the first simulation demonstrating that the presence of a magnetic field will suppress the growth of the Richtmyer-Meshkov instability when a shock wave interacts with a contact discontinuity separating ionized gases of different densities. The upper and lower images in Figure 3 contrast the interface without (upper) and with (lower) the magnetic field. In the presence of the field, the vorticity generated at the interface is transported away by the fast and slow MHD shocks, removing the driver of the instability. Results are shown for an effective

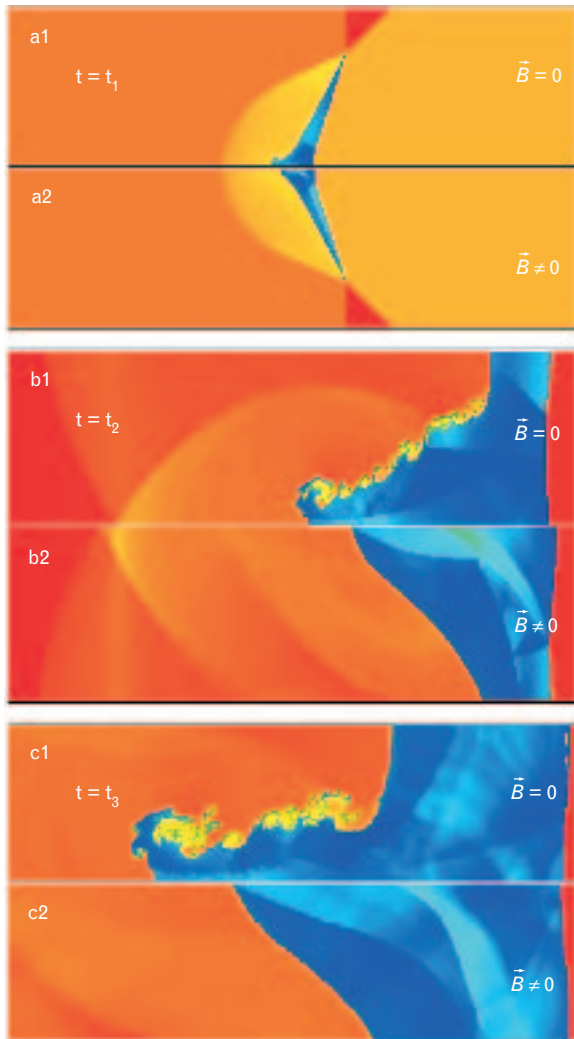


FIGURE 3
Stabilization of Richtmyer-Meshkov instability by a magnetic field. From R. Samtaney, "Suppression of the Richtmyer-Meshkov instability in the presence of a magnetic field," Princeton Plasma Physics Laboratory report PPPL-3794 (submitted to Phys. Fluids).

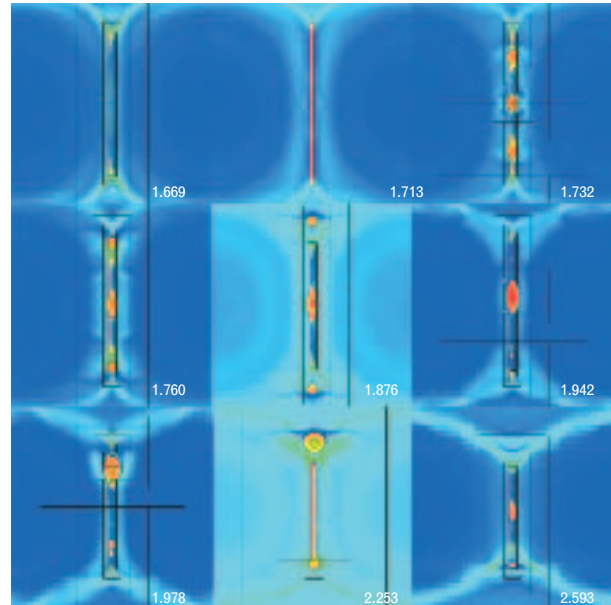


FIGURE 4
Current bunching and ejection during magnetic reconnection. From R. Samtaney, S. Jardin, P. Colella, and T. Ligocki, "Numerical simulations of magnetic reconnection using AMR," Sherwood Fusion Theory Conference, Rochester, NY, April 2002.

mesh of $16,384 \times 2,048$ points which took approximately 150 hours to run on 64 processors of Seaborg—25 times faster than a non-AMR code.

Another new physical effect discovered by the AMRMHD code is current bunching and ejection during magnetic reconnection (Figure 4). Magnetic reconnection refers to the breaking and reconnecting of oppositely directed magnetic field lines in a plasma. In the process, magnetic field energy is converted to plasma kinetic and thermal energy.

The CEMM project has been

collaborating with other SciDAC software centers in addition to APDEC. For example, in a collaboration that predated SciDAC, the group developing the M3D code was using PETSc, a portable toolkit of sparse solvers distributed as part of the ACTS Collection of DOE-developed software tools. Also in the ACTS Collection is Hypr, a library of preconditioners that can be used in conjunction with PETSc. Under SciDAC, the Terascale Optimal PDE Solvers (TOPS) Center worked with CEMM to add Hypr underneath the same code interface

that M3D was already using to call the PETSc solvers. The combined PETSc-Hypr solver library allows M3D to solve its linear systems two to three times faster than before.

According to Jardin, fusion plasma models are continually being improved both by more complete descriptions of the physical processes, and by more efficient algorithms, such as those provided by PETSc and CHOMBO. Advances such as these have complemented increases in computer hardware speeds to provide a capability today that is vastly improved over what

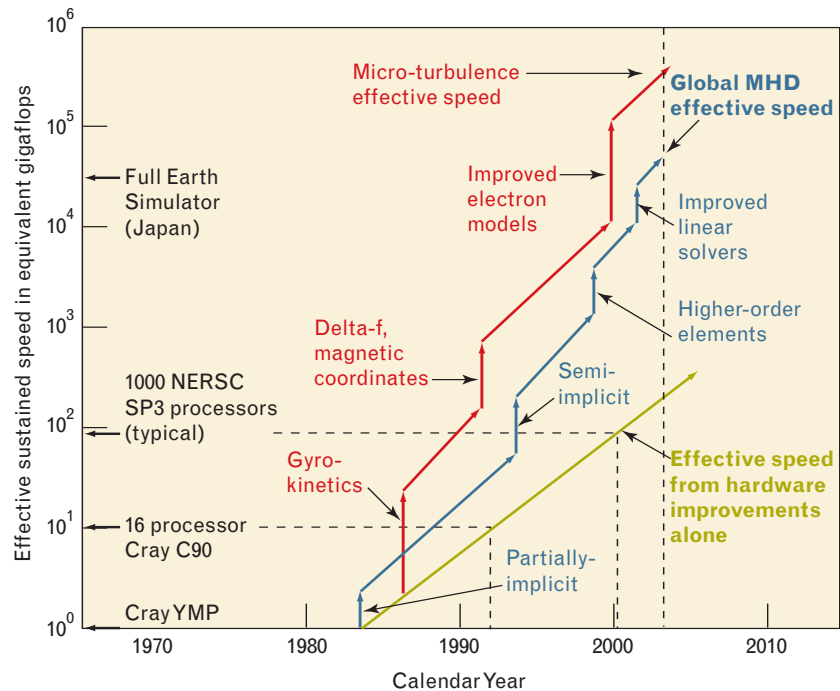


FIGURE 5
 “Effective speed” increases in fusion codes came from both faster hardware and improved algorithms. *Image: Steve Jardin.*

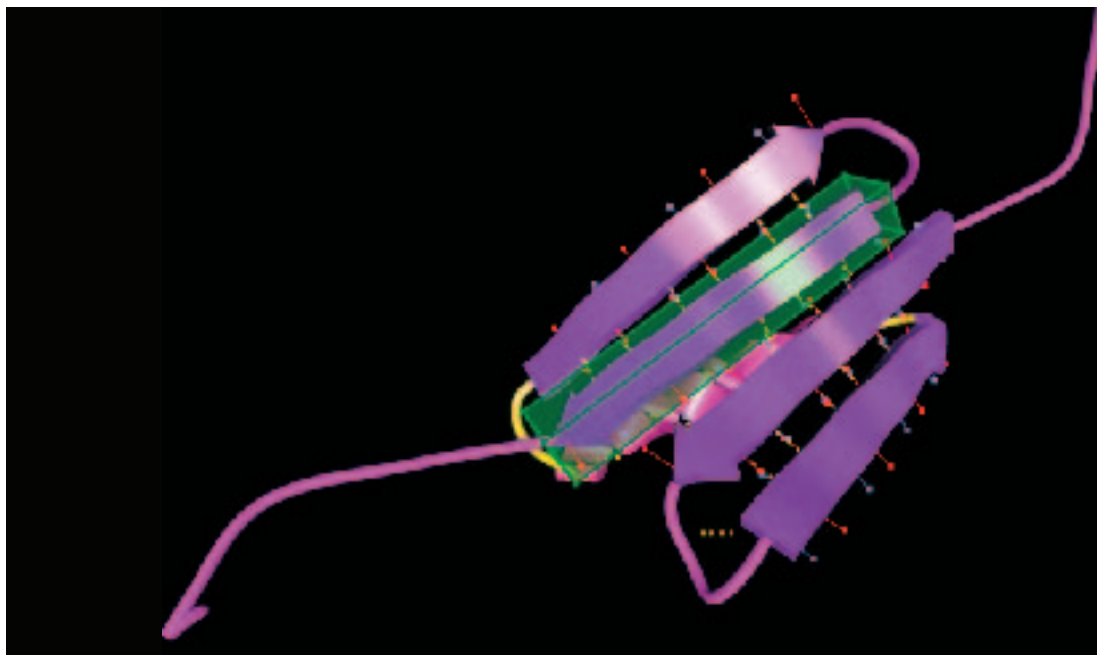
was possible 30 years ago (Figure 5). This rate of increase of effective capability is essential to meet the anticipated modeling demands of fusion energy research, Jardin says.

“Presently, we can apply our most complete computational models to realistically simulate both nonlinear macroscopic stability and

microscopic turbulent transport in the smaller fusion experiments that exist today, at least for short times,” Jardin says. “Anticipated increases in both hardware and algorithms during the next five to ten years will enable application of even more advanced models to the largest present-day experiments

and to the proposed burning plasma experiments such as ITER [the International Thermonuclear Experimental Reactor].”

Research funding: [FES](#), [SciDAC](#)
 For further discussion of AMR in computational science, see page 42.



BIOLOGY

DECIPHERING NATURE'S ASSEMBLY MANUALS

Now that a working draft of the human genome has been created, scientists are busy identifying the genes within the sequences of those 3 billion DNA base pairs of nucleotides. The complete genomes of many other organisms—from microbes to plants to animals—are also available, with new genomes being completed every week. Scientists are now facing the opportunity to understand not just the human body but the fundamental processes of life itself.

Once a gene's boundaries have been established, the next steps are to determine the structures of the proteins encoded in the gene, then to determine the functions of those proteins. NERSC users and collaborators are developing innovative computational tools to help determine protein structures and functions, as described in this section. These tools bring us closer to the goal of utilizing genomic discoveries to improve human health and to protect and restore the environment.

PROTEINSHOP: SOLVING PROTEIN STRUCTURES FROM SCRATCH

ProteinShop, a computer visualization tool for manipulating protein structures, is helping biologists close in on one of their cherished goals: completely determining an unknown protein's shape from its gene sequence.

Silvia Crivelli of the Visualization Group in Berkeley Lab's Computational Research Division says a major step forward came when "we copied concepts from robotics. When you move a robot's arm, you move all the joints, like your real arm." After a year of work on ProteinShop, Crivelli says, "we were able to apply the same mathematical techniques to protein structures."

It's just one of the components the ProteinShop vehicle uses to help competitors in the hottest race in biological computation—the "Grand Prix of bioinformatics" known as CASP, the Critical Assessment of Structure Prediction. In 1994, the challenge of predicting finished protein structures from

gene sequences gave rise to the unique CASP scientific competition. Teams of biologists and computer scientists from around the world try to beat one another to the fastest and most accurate predictions of the structures of recently found but as-yet unpublished proteins. The number of competitors has more than quintupled, from fewer than three dozen groups in 1994 to 187 in 2002's CASP5.

At CASP5, a team led by Teresa Head-Gordon of Berkeley Lab's Physical Biosciences Division, who is also Assistant Professor of Bioengineering at the University of California at Berkeley, employed the Global Optimization Method developed by Head-Gordon and her

colleagues, including Crivelli from Berkeley Lab and Bobby Schnabel, Richard Byrd, and Betty Eskow from the University of Colorado.

Instead of depending heavily on knowledge of protein folds (3D sets of structures) already stored in the Protein Data Bank (PDB), the Global Optimization Method uses initial protein configurations provided by ProteinShop, which was created by Wes Bethel and Crivelli of Berkeley Lab along with Nelson Max of Lawrence Livermore National Laboratory and the University of California at Davis, and UC Davis's Bernd Hamann and Oliver Kreylos.

With a ProteinShop jumpstart, Head-Gordon's CASP5 team was

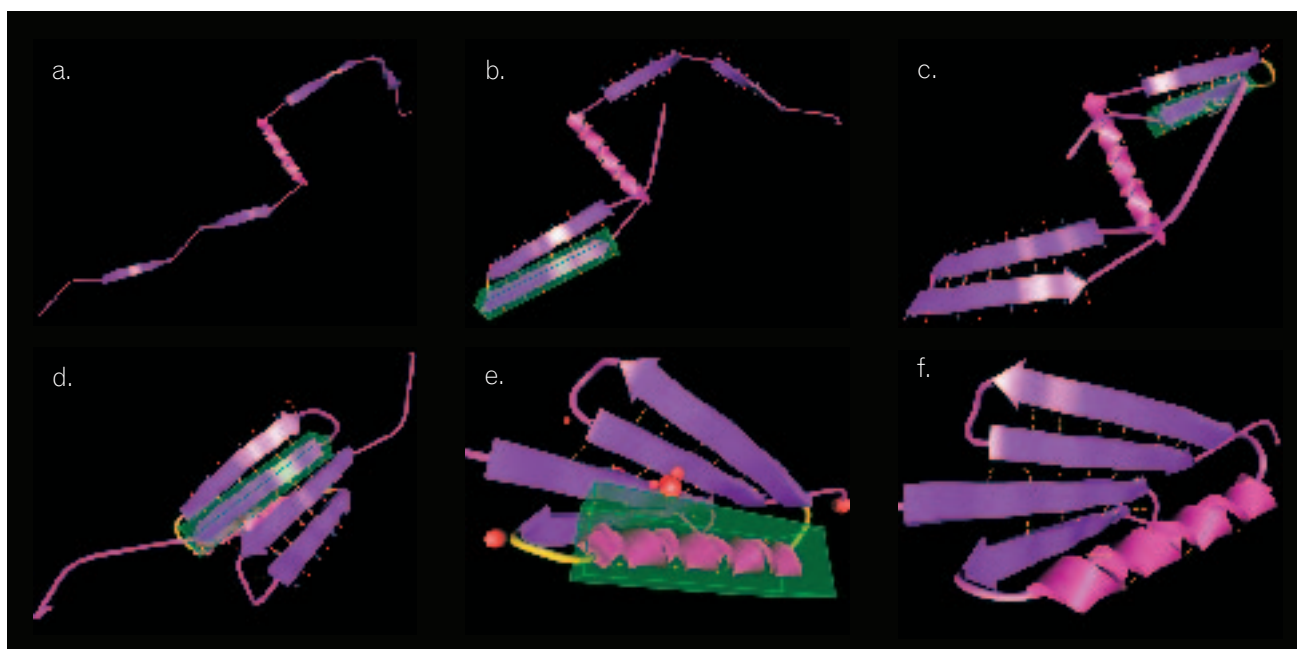


FIGURE 1
SETUP PHASE OF THE GLOBAL OPTIMIZATION METHOD.

a. ProteinShop reads an unknown protein's string of amino acids and assembles the predicted secondary structures: in this case, "uncoiled" coils alternate with beta strands and a central alpha helix. Even for large proteins, this step is virtually instantaneous.

b. The user tells ProteinShop to align two beta strands into a partial beta

sheet, forming hydrogen bonds between selected amino acid residues from each strand. In the yellow coil region, inverse kinematics calculates the correct dihedral angles as the strands are bent together.

c. To assemble a second beta sheet, ProteinShop semiautomatically aligns hydrogen bonds in the third and fourth beta strands.

d. The two partial sheets are pulled together and stabilized by hydrogen bonds between beta strands one and four.

e. Now the user selects the central helix and manipulates it to reduce interference (where the orange spheres mark atom collisions). Meanwhile, inverse kinematics calculates the angles in the flanking coil regions, and ProteinShop keeps all structures intact while the helix is moved.

f. The final 3D protein structure, ready for optimization.

Images: Oliver Kreylos

able to predict the configurations of 20 new or difficult protein folds ranging from 53 to 417 amino acids in length, which, says Crivelli, "is great for a method that doesn't use much knowledge from the PDB."

A protein's shape is what determines its function. Yet a one-dimensional string of amino acid residues, as specified by a gene's coding sequences, does not on the face of it reveal much about 3D shapes. Swimming in a watery environment, the amino-acid string—the protein's primary structure—quickly twists into familiar shapes known as alpha helices and beta sheets, striving to

reduce attractive and repulsive forces among the linked amino acids. The result is local regions where the energy required to maintain the structure is minimized.

The secondary structures swiftly fold up into a 3D tertiary structure; for most proteins this represents the "global" minimum energy for the structure as a whole. (Some complicated proteins like hemoglobin, assembled from distinct components, have a quaternary structure as well.)

By studying thousands of known proteins, biologists have learned to recognize sequences of amino acids that are likely to form alpha helices,

plus those that form the flat strands that arrange themselves side by side in beta sheets. This secondary-structure information can be stored in data banks and applied to unknown protein folds.

But the so-called coil regions that link secondary structures in an unknown protein are not so easily predicted. The more amino acids in the protein, the more local energy minima there are, and the harder it becomes to calculate the global energy minimum.

The Global Optimization Method tackles the problem in two distinct phases. During the setup

phase (Figure 1a–f), the program generates likely secondary structures from a given sequence of amino acids, which are combined into several approximate configurations. During the next phase, the overall energy of the structure is reduced step by step. Global optimizations are performed on subsets of the angles among the amino acid residues (dihedral angles) which are predicted to make up coil regions. Eventually, because the energy function is not perfect, several tertiary structures are proposed.

Before ProteinShop, the setup phase of the Global Optimization Method was immensely time-consuming. Because a protein's primary structure contains dozens or hundreds of acid residues, thus thousands of atoms, the initial phase required computer runs of hours or days before settling on secondary structures. Moreover, while the prediction of alpha helices was straightforward, beta sheets were hard to assemble from beta strands, and their configurations were less certain. In CASP4, in 2000, the team predicted the structures of only eight folds, the longest containing 240 amino acids.

The Visualization Group developed ProteinShop in preparation for CASP5. ProteinShop incorporates the mathematical technique called inverse kinematics—well known not only in robotics but also in video gaming. Inverse kinematics is used to predict the overall movements of structures consisting of jointed segments, for example the fingers, arms, and shoulders of a robot or an animated figure. By taking into account the degrees of freedom permissible

at each joint, contortions that don't break limbs or penetrate bodies can be predicted.

“The difference is that with a robot you have maybe 10 or 20 joints, but in a coil we often have long regions, 80 amino acids,” Crivelli says, “and we want all of the dihedral angles among them to move in a concerted way.” In ProteinShop, the secondary structures and coils are built up by adding amino acids to the structure one at a time, treating each as a jointed segment of a flexible structure. Within seconds a “geometry generator” module incorporates predicted secondary structures, or fragments of them, in the string. “The whole thing looks like you could just move it around like spaghetti,” says Crivelli. “But before we incorporated reverse kinematics, if you tried to move a protein configuration, it broke.”

Now the process works fast enough to be truly interactive, allowing the user to alter the dihedral angles between individual amino acids. In difficult cases like the assembly of beta strands into sheets, the user can manipulate the conformation to achieve the best, least energetic fit. Moreover, the user can play with the entire “preconfiguration,” dragging whole secondary assemblies into new relationships without breaking previous structures. “Once we have the alpha and beta structures, we want to leave them alone,” Crivelli says. “Mainly we work on the coil regions.”

The ProteinShop program allows for various styles of illustrating the fold and its parts, and it helps out the user by sending plenty

of signals. In one mode, little orange spheres pop up when the user has forced atoms to collide: the more energy-expensive the collision, the bigger the orange sphere. Manipulation can also be simplified by hand-selecting residues for automatic bonding.

The second phase of the Global Optimization Method—seeking the global energy minimum—remains a computational challenge. The basic approach to finding the optimal energy for a tertiary structure is to repeatedly tweak the energy budgets of many smaller regions, incorporating each improvement until there is no further change. This “physics-based” method depends only on energy calculations, not on knowledge of existing folds.

In practice, to do such calculations for every part of a structure would take much too long. The sophisticated procedures developed by Head-Gordon's team optimize only randomly sampled, small regions. Thus it's not certain that the end result is the true global optimum. Even so, global optimization of an unknown protein of moderate size may require weeks of computer time.

“We are working out new methodology that will combine our physics-based approach with knowledge-based methods,” says Crivelli. “By recognizing structures and fragments that are known to work, we won't have to calculate every angle from scratch. The tool will be highly interactive, displaying energies and saving minima as the user finds them. It will be organized as a guided search through a dynamically evolving

tree, basing new structures on previous ones that have been shown to work for the fold.”

To enable this high degree of interaction, still higher performance will be required from the Global Optimization Method’s already

high-performance parallel-processing codes. The results will have implications for scientific problems far beyond the successful solution of unknown protein structures. Other complex optimization problems might be amenable to this

combination of local and global optimization approaches, such as transportation and scheduling, structural optimization, chip design, and machine learning.

Research funding: [BER](#), [ASCR-MICS](#), [NSF](#)

MAPPING THE UNIVERSE OF PROTEIN STRUCTURES

The protein universe, the vast assemblage of biological molecules that are the building blocks of living cells and control the chemical processes which make those cells work, has finally been mapped.

Researchers at Berkeley Lab and the University of California at Berkeley have created the first 3D global map of the protein structure universe (Figure 2). This map provides important insight into the evolution and demographics of protein structures and may help scientists identify the functions of newly discovered proteins.

Sung-Hou Kim, a chemist with a joint appointment at Berkeley Lab and UC Berkeley, led the development of this map. An internationally recognized authority on protein structures, he expressed surprise at how closely the map, which is based solely on empirical data and a mathematical formula, mirrored the widely used Structural Classification System of Proteins (SCOP), which is based on the visual observations of scientists who have been solving protein structures.

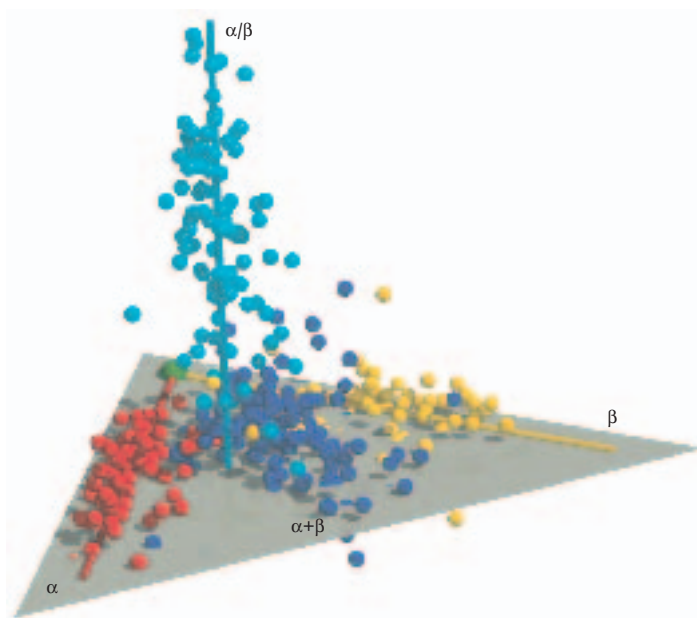


FIGURE 2

This 3D map of the protein structure universe shows the distribution in space of the 500 most common protein fold families as represented by spheres. The spheres, which are colored according to classification, reveal four distinct classes: α (red), β (yellow), $\alpha+\beta$ (blue), and α/β (cyan). It can be seen that the α , β , and $\alpha+\beta$ classes share a common area of origin, while the α/β class of protein structures is shown to be a late bloomer on the evolutionary flagpole. From J. Hou, G. E. Sims, C. Zhang, and S.-H. Kim, “A global representation of the protein fold space,” *PNAS* **100**, 2386 (2003).

“Our map shows that protein folds are broadly grouped into four different classes that correspond to the four classes of protein structures defined by SCOP,” Kim says. “Some have argued that there are really only three classes of protein fold structures, but now we can mathematically prove there are four.”

Protein folds are recurring structural motifs or “domains” that underlie all protein architecture. Since architecture and function go hand-in-hand for proteins, solving what a protein’s structure looks like is a big step towards knowing what that protein does.

The 3D map created by Kim along with Jingtong Hou, Gregory Sims, and Chao Zhang, shows the distribution in space of the 500 most common protein fold families as represented by points which are spatially separated in proportion to their structural dissimilarities. The distribution of these points reveals a high level of organization in the fold structures of the protein universe and shows how these structures have evolved over time, growing increasingly larger and more complex.

“When the structure of a new protein is first solved, we can place it in the appropriate location on the map and immediately know who its neighbors are and its evolutionary history, which can help us predict what its function may be,” Kim says. “This map provides us with a conceptual framework to organize all protein structures and functions and have that information readily available in one place.”

With the completion of a working draft of the human genome, in which scientists determined the sequences of the 3 billion DNA bases that make up the human genome, the big push now is to identify coding genes and the molecular and cellular functions of the proteins associated with them. Coding genes are DNA sequences that translate into sequences of amino acids, which RNA assembles into proteins.

The prevailing method for predicting the function of a newly discovered protein is to compare the sequence of its amino acids to the amino acid sequences of proteins whose functions have already been identified. A major problem with relying exclusively on this approach is that while proteins in different organisms may have similar structure and function, the sequences of their amino acids may be dramatically different. “This is because protein structure and function are much more conserved through evolution than genetically based amino acid sequences,” Kim says.

Kim has been a leading advocate for grouping proteins into classes on the basis of their fold structures and using these structural similarities to help predict individual protein functions. While the protein universe may encompass as many as a trillion different kinds of proteins on Earth, most structural biologists agree there are probably only about ten thousand distinctly different types of folds.

“A smaller number of new pro-

tein folds are discovered each year, despite the fact that the number of protein structures determined annually is increasing exponentially,” Kim says. “This and other observations strongly suggest that the total number of protein folds is dramatically smaller than the number of genes.”

The rationale behind this idea is that through the eons, proteins have selectively evolved into the architectural structures best suited to do their specific jobs. These structures essentially stay the same for proteins from all three kingdoms of life—bacteria, archaea, and eukarya—even though the DNA sequences encoding for a specific type of protein can vary wildly from the genome of one organism to another, and sometimes even within the same organism.

In the map created by Kim and his colleagues, elongated groups of fold distributions approximately corresponding to the four SCOP structural classifications can be clearly seen. These classifications, which are based on secondary structural compositions and topology, are the “alpha” helices, “beta” strands, and two mixes of helices and strands, one called “alpha plus beta” and the other “alpha slash beta.” The map reveals that the first three groups share a common area of origin, possibly corresponding to small primordial proteins, while the “alpha slash beta” class of proteins does not emerge until much later in time.

“It is conceivable that, of the primordial peptides, those containing fragments with high helix

and/or strand propensity found their way to fold into small alpha, beta, and alpha-plus-beta structures,” Kim says. “The alpha-slash-beta fold structures do not appear until proteins of sufficient size rose through evolution and the formation of supersecondary structural units became possible.”


Since understanding the molecular functions of proteins is key to understanding cellular functions,

the map developed by Kim and his colleagues holds promise for a number of areas of biology and biomedical research, including the design of more effective pharmaceutical drugs that have fewer side effects. “This map can be used to help design a drug to act on a specific protein and to identify which other proteins with similar structures might also be affected by the drug,” Kim says.

For the next phase of this research, Kim and his colleagues plan to tap into NERSC’s supercomputers to add the rest of the some 20,000 and counting known protein structures to their map. They also plan to set up a Web site where researchers can submit for inclusion new protein structures they have solved.

Research funding: [BER](#), [NIH](#), [NSF](#)

DOE FACILITY PLAN SEES MAJOR ROLE FOR COMPUTING



In a speech at the National Press Club on November 10, 2003, U.S. Energy Secretary Spencer Abraham outlined the Department of Energy's Office of Science 20-year science facility plan, a roadmap for future scientific facilities to support the department's basic science and research missions. The plan prioritizes new, major scientific facilities and upgrades to current facilities.

"This plan will be the cornerstone for the future of critical fields of science in America. These facilities will revolutionize science—and


society," said Abraham. "With this plan our goal is to keep the United States at the scientific forefront."

Priority 2 in the plan is an Ultra-Scale Scientific Computing Capability (USSCC), to be located at multiple sites, which would increase by a factor of 100 the computing capability available to support open scientific research. Most existing supercomputers have been designed with the consumer market in mind; USSCC's new configurations will develop computing capability specifically designed for science and industrial applications. These facilities will operate like Office of Science light sources: available to all, subject to proposal peer review.

The USSCC will involve long-term relationships with U.S. computer vendors and an integrated program of hardware and software investments designed to optimize computer performance for scientific and commercial problems by creating a better balance of memory size, processor speed, and interconnection rates. A very similar strategy was spelled out in the white paper that NERSC together with seven other DOE labs submitted to the High End Computing Revitalization Task Force in June 2003 (see <http://www.neresc.gov/news/HECRTF-V4-2003.pdf>).

An upgrade of the NERSC facility is tied for Priority 7. This upgrade will ensure that NERSC, DOE's premier scientific computing facility for unclassified research, continues to provide high-performance computing resources to support the requirements of scientific discovery.

NERSC will continue to provide the core scientific computing needed by the research community, and will complement the "grand challenge" approach pursued under the USSCC, according to the facility plan. Reaffirming the goals of the NERSC Strategic Proposal for 2002–2006, the plan describes the NERSC upgrade as using Grid technology to deploy a capability designed to meet the needs of an integrated science environment combining experiment, simulation, and theory by facilitating access to computing and data resources, as well as to large DOE experimental instruments. NERSC will concentrate its resources on supporting scientific challenge teams, with the goal of bridging the software gap between currently achievable and peak performance on the new tera-scale platforms.



NANOSCIENCE

MODELING A REALM WHERE EVERYTHING IS DIFFERENT

The nanoscale—between 10 and 100 billionths of a meter—is the scale at which the fundamental properties of materials and systems are established. Melting temperature, magnetic properties, charge capacity, and even color are dictated by the arrangement of nanoscale structures. The realm of molecular biology also operates largely at the nanoscale.

While much is known about the physical properties and behavior of isolated molecules and bulk materials, the properties of matter at the nanoscale cannot necessarily be predicted from those observed at larger or smaller scales. Nanoscale structures exhibit important differences that cannot be explained by traditional models and theories.

Research at the nanoscale aims to provide a fundamental understanding of the unique properties and phenomena that emerge at this length scale as well as the ability to tailor those properties and phenomena to produce new materials and technologies in fields as diverse as electronics, biotechnology, medicine, transportation, agriculture, environment, and national security. But a quantitative understanding of nanoscale systems requires development of new computational methods. This section focuses on one application central to DOE's mission: promoting clean and efficient energy production.

MAKING FUEL CELLS MORE EFFICIENT AND AFFORDABLE

Fuel cells sound like an ideal power source—they are efficient and quiet, have few moving parts, run on hydrogen, and emit only water vapor.

It's no wonder that President Bush proposed new research into hydrogen-powered automobiles in his 2003 State of the Union address. But for this dream to become reality, scientists and engineers need to make fuel cells less expensive.

While there are several types of fuel cells, proton exchange membrane (PEM) cells are the leading candidates for powering light-duty vehicles and buildings and for replacing batteries in portable electronic devices. The PEM is a polymer electrolyte that allows hydrogen ions (protons) to pass through it. As hydrogen (H_2) is fed to the anode side of the fuel cell, a catalyst (usually platinum or a platinum-ruthenium alloy) encourages the hydrogen atoms to separate into electrons and protons. The electrons travel from anode to cathode along an external

circuit, creating a usable electric current. Meanwhile, on the cathode side, oxygen molecules (O_2) are being separated by the catalyst into separate atoms in a process called oxygen reduction. The negatively charged oxygen atoms (adsorbed on the catalyst) attract the positively charged hydrogen ions through the PEM, and together they combine with the returning electrons to form water molecules (H_2O).

While splitting the hydrogen molecules is relatively easy for platinum-alloy catalysts, splitting the oxygen molecules is more difficult. Although platinum catalysts are the best known for this task, they are still not efficient enough for widespread commercial use in fuel cells, and their high price is one of the reasons why fuel cells are not yet economically competitive.

Finding more efficient and affordable catalysts is one of the goals of a research group led by Perla Balbuena, Associate Professor of Chemical Engineering at the University of South Carolina. Using computational chemistry and physics techniques, they are trying to predict the properties of electrodes consisting of bimetallic and trimetallic nanoparticles embedded in a polymer system and dispersed on a carbon substrate. They are working to identify the atomic distribution, vibrational modes, and diffusion coefficients of a variety of alloy nanoparticles.

“The rationale behind the use of alloy nanoparticles,” Balbuena says, “is that each of the elements, or particular combinations of them, may catalyze different aspects of the reaction. The large surface-to-volume

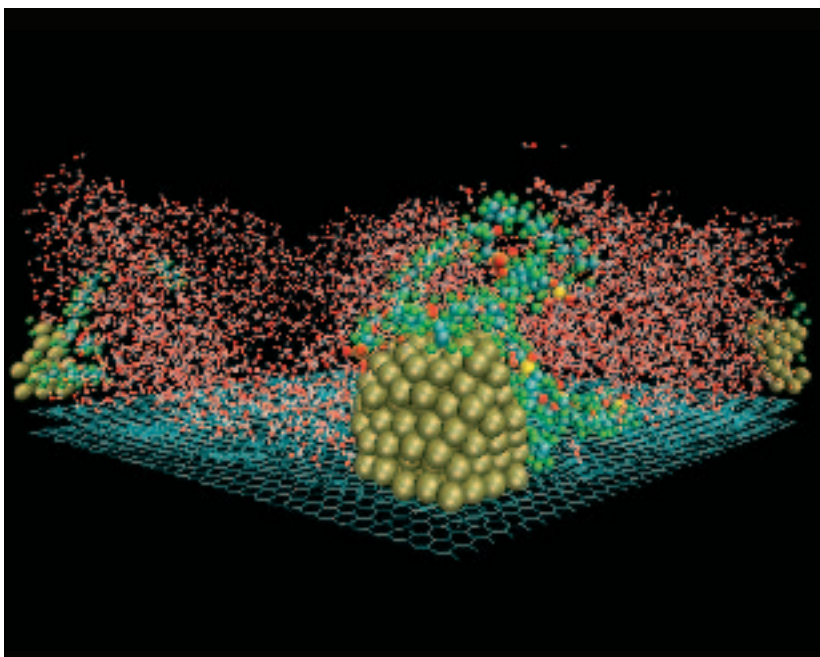
area of nanoparticles may also contribute to their effectiveness.”

Experimental data indicates that alloys can be as effective or better than platinum catalysts, but nobody fully understands why this is so.

Little is known about the physical, chemical, and electrical properties of alloy nanoparticles, and in particular about the atomic distribution of the various elements on the catalyst surface, which is difficult to determine experimentally. Computational studies play an essential role in answering these questions.

For example, Figure 1 shows a test simulation of a system involving only platinum nanoparticles on a carbon substrate and surrounded by a hydrated polymer. The goal of this simulation was to test the stability of the nano-size catalyst in such an environment—that is, how the polymer groups, water, and substrate influence the shape and mobility of the nanoparticles, and how they influence the exposed surfaces where the oxygen molecules will be adsorbed and dissociated into oxygen atoms.

Balbuena and her colleagues recently developed a new computational procedure to investigate the behavior of active catalytic sites, including the effect of the bulk environment on the reaction. They demonstrated this procedure in simulations of oxygen dissociation on platinum-based clusters alloyed with cobalt, nickel, or chromium, and embedded in a platinum matrix. To determine the optimal alloy, they studied a variety of compositions and geometries for alloy ensembles embedded into the bulk metal (Figure 2).



The researchers concluded that cobalt, nickel, and chromium make platinum-based complexes more reactive because they enable a higher number of unpaired electrons.

Their computations yielded extensive new information that would be very hard to obtain experimentally, such as showing for the first time how alloying results in changes in the density of electronic states, and how those changes influence the oxygen reduction process. One of the questions they addressed was whether nickel, cobalt, and chromium might become oxidized and act as “sacrificial sites” where other species can adsorb, leaving the platinum sites available for oxygen dissociation. Their study shows that this may be the case for nickel, but that alternative mechanisms are possible, especially for cobalt and chromium, which instead of being just sacrificial sites, act as active sites where oxygen dissociates, in some cases more efficiently than on

FIGURE 1

Platinum nanoparticles (gold-colored atoms) are deposited on a graphite support. The long chains are Nafion fragments composed of fluorine (green), carbon (blue), oxygen (red), and sulfur (yellow) atoms; some of these chains become adsorbed over the metallic clusters. The background contains water molecules (red for oxygen and white for hydrogen). Nafion, widely used for proton exchange membranes, has both hydrophilic and hydrophobic chemical groups, so some degree of water segregation is observed. *Image: Eduardo J. Lamas.*

platinum sites. They found that the most effective active sites that promote oxygen dissociation are O_2CoPt , O_2Co_2Pt , O_2CrPt , and O_2Cr_2Pt , while ensembles involving nickel atoms are as catalytically effective as pure platinum.

The Balbuena group's research is now focusing on the time evolution of the complete reaction on the best catalytic ensembles, including the combination of oxygen atoms with protons leading to water formation. Future work will address the possibilities of reducing further the catalyst size, nowadays in the order

of tens of nanometers, as well as developing a full understanding of the role of the rest of the environment (polymer, water, and substrate) on the outcome of the reaction. The ultimate goal of these studies is to enable the design of new catalysts from first principles rather than trial-and-error experiments, reducing the cost of catalysts for fuel cells through the use of optimum alloy combinations requiring only minimal amounts of expensive materials.

Research funding: [BES](#)

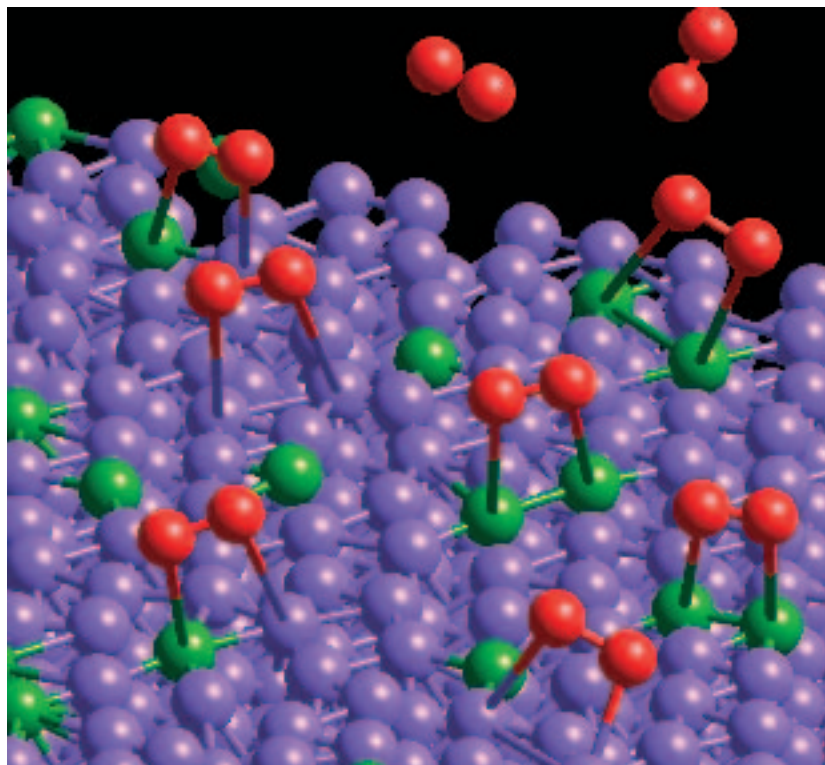


FIGURE 2

Oxygen molecules (red) adsorb on a bimetallic surface of platinum (purple) and cobalt (green). The oxygen molecules most frequently adsorb on "bridge" sites composed of two cobalt atoms or one cobalt and one platinum atom. The oxygen molecules will subsequently dissociate into individual atoms. From P. B. Balbuena, D. Altomare, L. Agapito, and J. M. Seminario, "Theoretical analysis of oxygen adsorption on Pt-based clusters alloyed with Co, Ni, or Cr embedded in a Pt matrix," *J. Phys. Chem. B* (in press).



CLIMATE MODELING

PREDICTING THE RESULTS OF A GLOBAL EXPERIMENT

Ever since the beginning of the industrial era, when humans began dramatically increasing the amount of carbon dioxide in the atmosphere by burning large quantities of fossil fuels, we have unintentionally been conducting a massive experiment on the Earth's climate. Unfortunately, this uncontrolled experiment does not meet the criteria of the scientific method—we do not have a parallel “control” planet Earth where we can observe how the climate would have evolved without the addition of CO₂ and other greenhouse gases. One of the goals of computational climate modeling is to provide such control scenarios so that human contributions to climate change can be quantified. This section describes how climate researchers studied the rising boundary between the troposphere and stratosphere and determined that human activity is the principal cause.

FINGERPRINTS IN THE SKY

For a physicist and atmospheric scientist, Ben Santer of Lawrence Livermore National Laboratory has an unusual specialty—he is a fingerprint expert.

Like a police detective, Santer looks for unique identifying marks that can be used to determine responsibility. But instead of looking for criminals, Santer is looking for specific, quantitative evidence of human influence on changes in the Earth's climate. The “fingerprints” he tries to match are patterns of data from observations of the atmosphere and from computer simulations.

“We examine the output from the computer models and compare this with observations,” Santer says. “In the model world (unlike the real world!) we have the luxury of being able to change one factor at a time, while holding the others constant. This allows us to isolate and quantify the effect of individual factors.”

One of the most dramatic appli-

cations of this “fingerprinting” technique recently answered the question of why the tropopause has been rising for more than two decades.

The tropopause is the boundary between the lowest layer of the atmosphere—the turbulently mixed troposphere—and the more stable stratosphere. The “anvil” often seen at the top of thunderclouds is a visible marker of the tropopause, which lies roughly 10 miles above the Earth's surface at the equator and 5 miles above the poles.

The average height of the tropopause rose about 200 meters between 1979 and 1999. In their first comparison of observed data with computer models, Santer and his colleagues concluded that the increase in tropopause height was

driven by the warming of the troposphere by greenhouse gases and the cooling of the stratosphere by ozone depletion.¹ But to quantify the influences (or “forcings” in climate jargon) even further, they considered three anthropogenic forcings—well-mixed greenhouse gases, sulfate aerosols, and tropospheric and stratospheric ozone—as well as two natural forcings—changes in solar irradiance and volcanic aerosols—all of which are likely to influence tropopause height.²

The observational data were compared with seven simulation experiments using the Parallel Climate Model (PCM). In the first five experiments, only one forcing at a time was changed. In the sixth experiment, all five forcings were

¹ B. D. Santer, R. Sausen, T. M. L. Wigley, J. S. Boyle, K. AchutaRao, C. Doutriaux, J. E. Hansen, G. A. Meehl, E. Roeckner, R. Ruedy, G. Schmidt, and K. E. Taylor, “Behavior of tropopause height and atmospheric temperature in models, reanalyses, and observations: Decadal changes,” *J. Geophys. Res.* **108**, 4002 (2003).

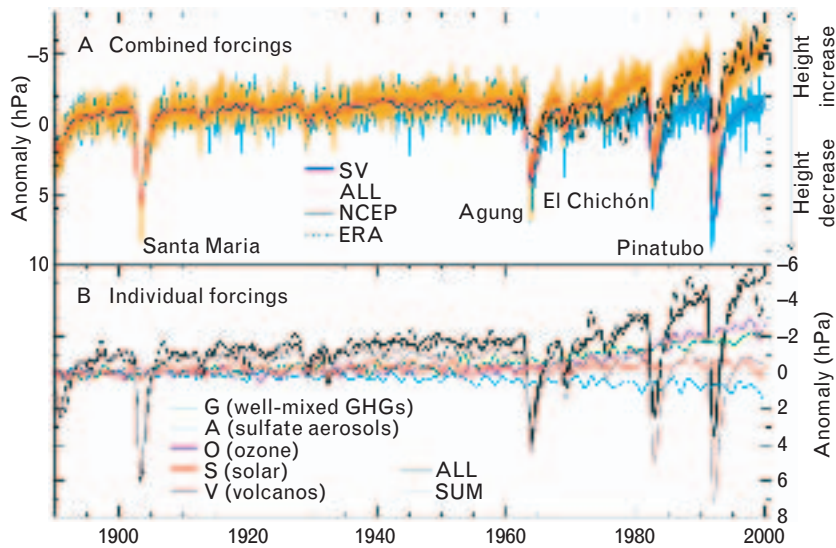


FIGURE 1
Effect of different forcings on tropopause height. Image: B.D. Santer, M. F. Wehner, T. M. L. Wigley et al.

included simultaneously. The seventh experiment used only the two natural forcings. To improve the reliability of the results, each experiment was performed four times using different initial conditions from a control run. This methodology allowed the researchers to estimate the contribution of each forcing to overall changes in atmospheric temperature and tropopause height.

The results of the analysis showed that in the mid-1980s, the combined five forcings (ALL) diverged from the natural forcings (SV), showing the dominance of human influences (Figure 1A). Of the five forcings, greenhouse gases (G) and ozone (O) played the biggest role (Figure 1B). Major volcanic eruptions can be seen to lower tropopause height sharply by cooling the troposphere and warm-

ing the stratosphere, but this effect is short-lived. The analysis concludes that about 80% of the rise in tropopause height is due to human activity. The model-predicted fingerprint of tropopause height changes was statistically detectable in two different observational data sets, NCEP and ERA (Figure 2).

Michael Wehner of Berkeley Lab's Computational Research Division, a co-author of the Science paper, points out that it would have been very difficult to estimate the relative contributions of different forcing mechanisms without the very large ensembles of climate model experiments conducted by DOE-funded researchers at the National Center for Atmospheric Research and made possible by NERSC and other DOE and NSF supercomputing facilities. Wehner

² B. D. Santer, M. F. Wehner, T. M. L. Wigley, R. Sausen, G. A. Meehl, K. E. Taylor, C. Ammann, J. Arblaster, W. M. Washington, J. S. Boyle, and W. Brüggemann, "Contributions of anthropogenic and natural forcing to recent tropopause height changes," *Science* **301**, 479 (2003).

manages the largest single collection of publicly available climate model output, the DOE Coupled Climate Model Data Archive (http://www.neresc.gov/projects/gcm_data/), which is stored in NERSC's HPSS system.

Looking at the tropopause data as part of the bigger climate picture, Santer said, "Tropopause height is an integrated indicator of human-induced climate change. It reflects global-scale changes in the temperature structure of the atmosphere. It is also consistent with results from other studies that have identified anthropogenic fingerprints in a range of different climate variables, such as ocean heat content, surface temperature, sea-level pressure patterns, and Northern Hemisphere sea-ice extent."

"The challenge for the future," Santer added, "is to take a closer look at the internal consistency of all these climate change variables. Our present work suggests that this story is a consistent one, but there are lots of details that need to be worked out."

Santer received the Department of Energy's 2002 Ernest Orlando Lawrence Award for "his seminal and continuing contributions to our understanding of the effects of human activities and natural phenomena on the Earth's climate system."

Research funding: [BER](#), [NOAA](#)

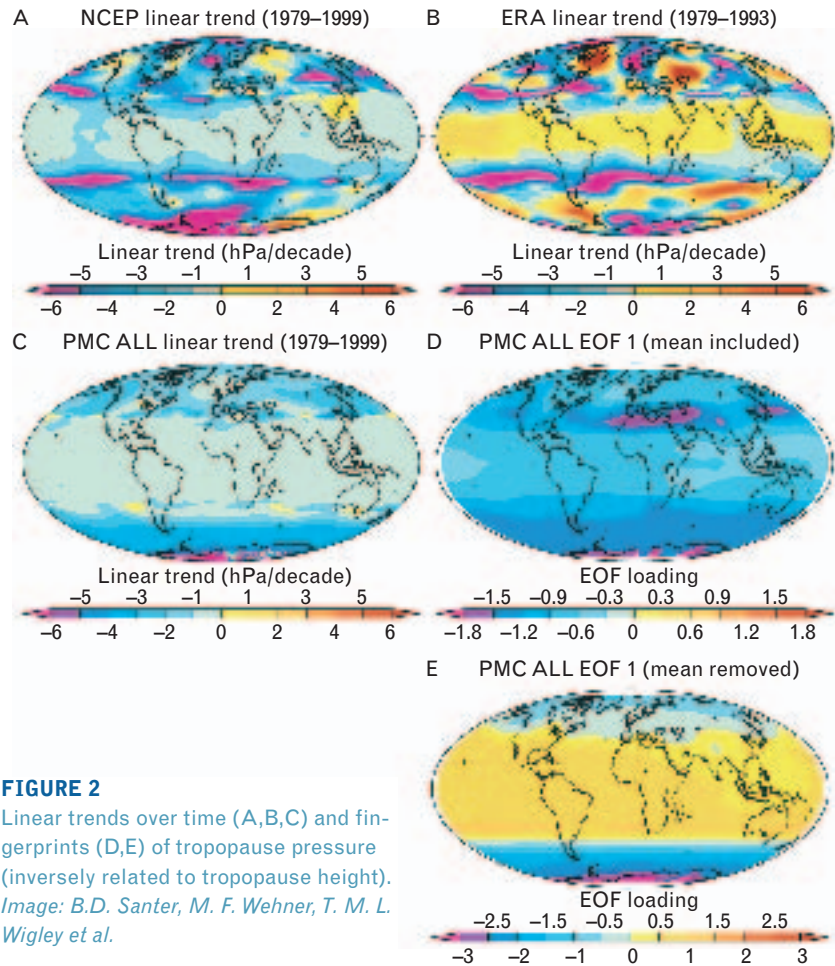
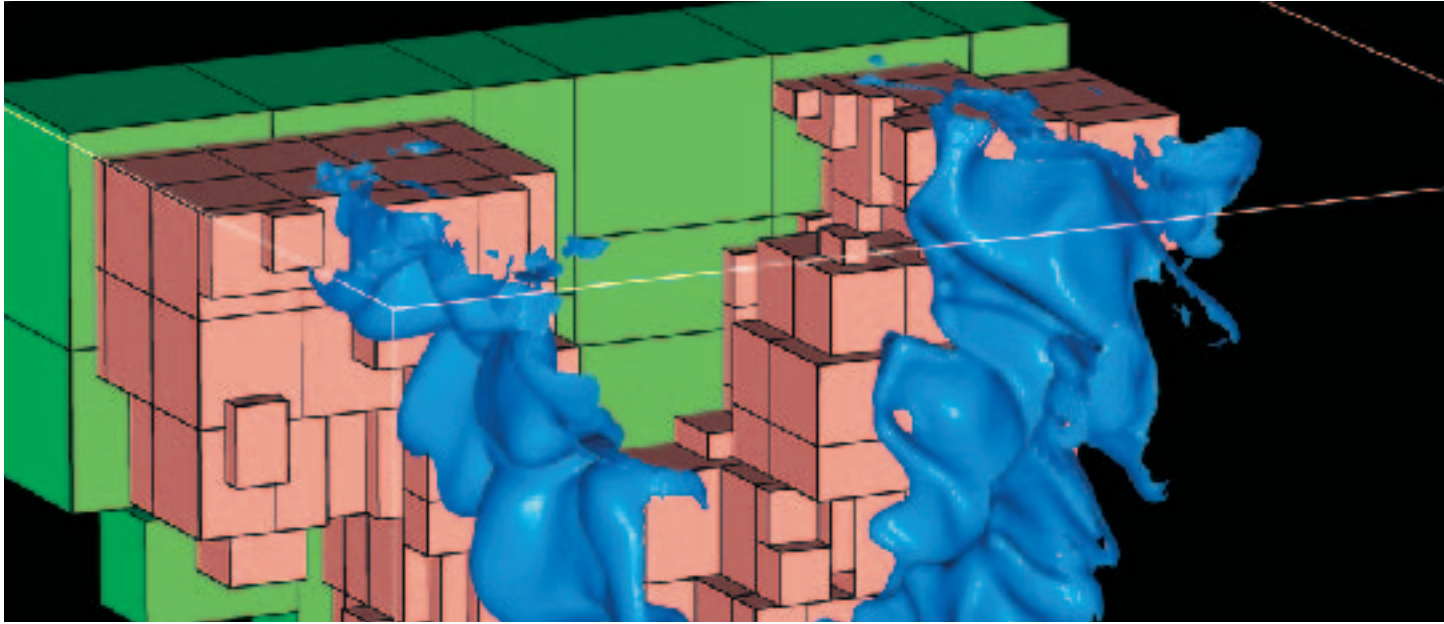


FIGURE 2 Linear trends over time (A,B,C) and fingerprints (D,E) of tropopause pressure (inversely related to tropopause height). Image: B.D. Santer, M. F. Wehner, T. M. L. Wigley et al.



APPLIED MATHEMATICS

THE BRIDGE BETWEEN PHYSICAL REALITY AND SIMULATIONS

“Mathematics is the bridge between physical reality and computer simulations of physical reality. The starting point for a computer simulation is a mathematical model, expressed in terms of a system of equations and based on scientific understanding of the problem being investigated, such as the knowledge of forces acting on a particle or a parcel of fluid.”¹

One of the challenges facing many model developers is managing model complexity. Many physical systems have processes that operate on length and time scales that vary over many orders of magnitude. These systems require multi-scale modeling, which represents physical interactions on multiple scales so that the results of interest can be recovered without the unaffordable expense of representing all behaviors at uniformly fine scales. This section describes the award-winning work of two research groups specializing in multiscale modeling.

¹ “A Science-Based Case for Large-Scale Simulation,” report to the DOE Office of Science, July 30, 2003, p. 31.

ALGORITHMS IN THE SERVICE OF SCIENCE

John Bell and Phillip Colella, applied mathematicians at Lawrence Berkeley National Laboratory and long-time NERSC users, were named co-recipients of the 2003 SIAM/ACM Prize in Computational Science and Engineering, awarded by the Society for Industrial and Applied Mathematics (SIAM) and the Association for Computing Machinery (ACM).

According to SIAM President Mac Hyman, the prize “is awarded in the area of computational science in recognition of outstanding contributions to the development and use of mathematical and computational tools and methods for the solution of science and engineering problems.” This is the first year the prize has been awarded.

Algorithms developed by Bell, Colella and their research groups at Berkeley Lab are used for studying complex problems arising in fluid mechanics and computational physics. The mathematical concepts, models, and methods they have developed have been applied in such diverse areas as shock physics, turbulence, astrophysics, flow in

porous media, and combustion.

Bell is head of the Center for Computational Sciences and Engineering (CCSE), while Colella heads the Applied Numerical Algorithms Group (ANAG), both within Berkeley Lab’s Computational Research Division. Both groups are participants in the SciDAC Applied Partial Differential Equations Center (APDEC).

In 2003, CCSE has been modeling turbulent reacting flow in methane flames, and applying some of the information gleaned in those experiments toward supernova research and modeling. ANAG has been involved in developing more sophisticated geometric tools for its Chombo software toolkit, and

applying it in areas such as beam dynamics in accelerator designs, gas jets in laser-driven plasma-wakefield accelerators, simulations of fusion reactors, and biological cell modeling.

In order to deal with problems in such diverse and complex sciences, CCSE and ANAG create customized software components that are able to translate scientific ideas and empirical insights into precise, 3D models that can be represented on a computer. This involves computational research that bridges the gap between mathematical, physical and computational sciences. By breaking up each problem into its fundamental mathematical parts, and designing optimal algorithms to handle each

piece, they are able to model a wide range of scientific phenomena.

In describing the cell modeling work he recently started with Berkeley Lab biologists, Colella explained that “in order to get better predictive models, you have to keep trying things out. The process becomes a very experimental, empirical activity, and that’s why the software tools are so important here.”

The software tools are based on adaptive mesh refinement (AMR), a grid-based system that uses partial differential equations to describe information in the most efficient manner possible—applying a fine mesh, and maximum computing power, to the areas of most interest, and a coarse mesh to areas where data is not as significant. AMR is particularly useful in problems where there is a wide range of time scales, and adaptive methods must be used to track movements in the mesh over time.

One of Bell and CCSE’s main projects in 2003 has been creating 3D simulations of turbulent methane combustion with swirling flow, which has produced the first detailed simulations in this area of combustion science (Figure 1). In the course of the project, two things have been of primary interest: how turbulence in the fuel stream affects the chemistry, and how emissions are formed and released. By studying the reacting flow of turbulent flames, CCSE has been able to develop computational models that dissect the detailed chemical reactions occurring in the flame. The group has also developed predictive

models that can be used to analyze flame properties and characteristics of turbulence. Even with a simple fuel like methane, there are some 20 chemical species and 84 reactions involved; the reaction zone, where the adaptive mesh is concentrated, occurs in an area that is 150 microns, while the whole grid consumes 12 centimeters. All of this

creates a problem size that demands some hefty computing power.

In order to create the flame visualizations, computations were performed on 2,048 processors of NERSC’s IBM SP, “Seaborg.” CCSE has created its own data analysis framework with specific algorithms to reformulate equations and accelerate the computations.

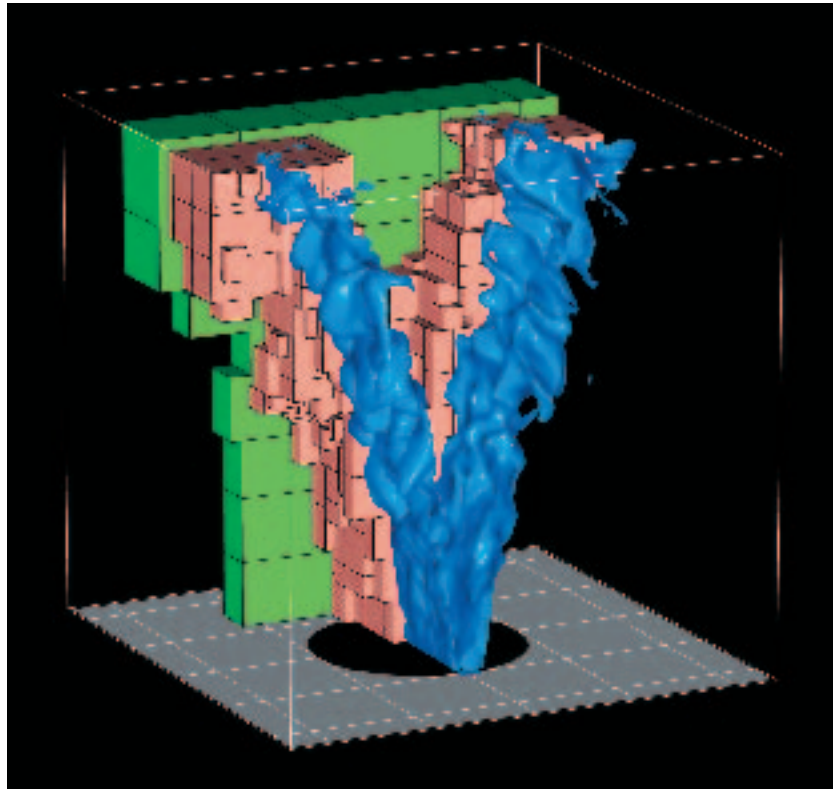


FIGURE 1

Instantaneous flame surface from a turbulent premixed combustion simulation. CCSE carried out the calculation using a low Mach number flow model with an AMR solution algorithm. The coarsest, base grid is invisible but fills the outlined cube. This cutaway view displays as green and pink boxes the containers that hold two successive levels of refined computational cells. These containers are dynamically reconfigured to follow the evolving flame surface, and are distributed over the parallel processors of the supercomputer. *Image: Marc Day.*

The result is a computational cost that is reduced by a factor of 10,000 compared with a visualization using a standard uniform-grid approach.

While a fundamental goal of the project is a better understanding of turbulence in flames, it could also have practical ramifications, as in engineering new designs for more efficient turbines, furnaces, incinerators, and other equipment held to strict emission requirements. A side benefit of the project is that it is stretching the group's software to the limit, putting new demands on the algorithmic and geometric elements of the toolset.

"The methane flame happens to be the one application that encompasses most of the pieces of the software. They are all folded into this application," said Marc Day, a staff scientist at CCSE. "You have to have applied math and computing hardware, together with good computational science, in order to accomplish what we have. Without any one of those pieces, it doesn't work."

The methane flame project falls into a class of problems called "interface dynamics," where something happens to cause an instability in an interface between two fluids of different chemicals and density. This instability causes the interface to grow and change form. The modeling goal then becomes to track the evolution of the new interface. AMR is able to tackle these types of problems, where sophisticated and powerful tools are needed to scale and quantify the complicated energy dynamics. Equally important is a skilled, dedicated group of

researchers to frame and calculate the problems.

To a great degree, CCSE and ANAG rely on their own computational scientists to apply the software to new projects and problem sets. But one of their goals has been to develop software components that other scientists can utilize more easily.

"We really want to support geometry and we need to support lots of different applications in different ways," said Colella. "So how can we re-factor the software to do that? For one, we're learning a lot about software engineering and managing the development process, using a 'spiral design' approach. When we add a new capability and develop applications in it, we get it out there for review. As we get feedback, we make changes to the software and repeat the process over again, so it becomes an ongoing spiral."

Another one of the group's goals in software development is to create algorithmic tools for doing computations in complex geometries. So if a problem demands hyperbolic, elliptic, and parabolic solver components, each of those pieces can be integrated while preserving the integrity of the model as a whole.

Both groups' software development is an ongoing process, and new capabilities are constantly being added to make them more flexible and applicable in broader arenas of science.

Computational biology is one of the major growth areas, said Colella. Current research is creating enormous amounts of data on how the development and behavior

of cells is governed by a complex network of reacting chemical species that are transported within a cell and among cellular communities. ANAG is working with the Arkin Laboratory in Berkeley Lab's Physical Biosciences Division and UC Berkeley's Bioengineering and Chemistry departments to develop the software infrastructure to turn this data into quantitative and predictive models. These models will contribute to a deeper understanding of cellular regulation and development, more efficient discovery of cellular networks, and a better ability to engineer or control cells for medical, agricultural, or industrial purposes.

The California Department of Water Resources has also asked for ANAG's help in creating high-quality simulation software for modeling water flows from the Sacramento Delta to the San Francisco Bay (Figure 2). These models are valuable for researchers interested in better understanding fresh water flows, and other environmental and water quality issues.

One of the reasons these mathematical software tools are so unique and powerful is that they can be applied to a wide array of scientific research, in a remarkably flexible fashion.

After analyzing the results of the methane flame project, for instance, CCSE scientists started to apply some of this newfound knowledge toward supernova research. Scientists have long been interested in fluid dynamics and energy transfers in supernovae, and trying to determine how and why they burn out. Using

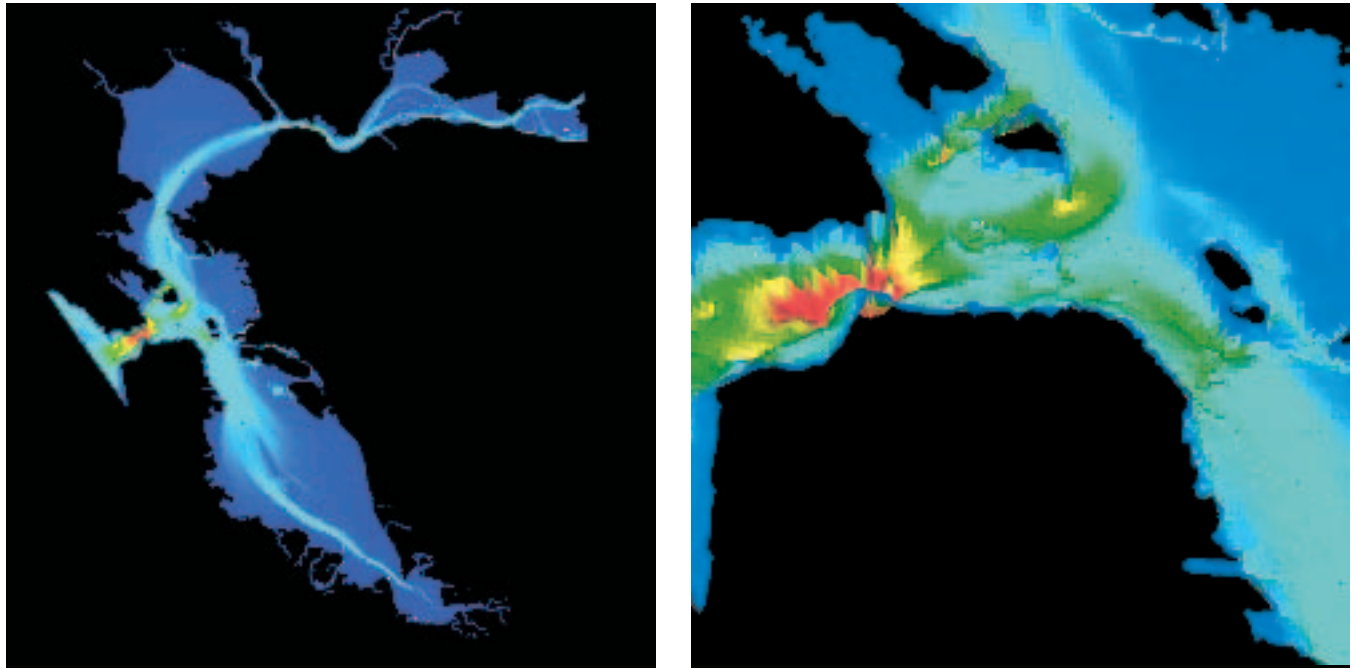


FIGURE 2.

One of the goals of the ANAG software effort is generating grids for complex geometries. The images here show a cut-cell grid generated to compute flow in the San Francisco Bay. At grid cells intersecting the bottom boundary, the numerical representation of the boundary is colored by the z-coordinate (i.e., depth) of the boundary in each cell. The left image is of the entire bay, the right one of the area around the Golden Gate. Image: Mike Barad (UC Davis) and Peter Schwartz (Berkeley Lab).

algorithms developed during its flame turbulence project, CCSE created 3D models of supernovae that shed new light on this problem, and have produced insights into the various stages of supernova explosions.

“You wouldn’t think that supernovae have anything in common with Bunsen burners, but algorithmically they were very similar,” said Day. “There was just a small number of pieces that we had to change out in order to make the chemical

combustion turn into a nuclear flame. So we were able to make significant progress over a short period of time in the area of supernova physics because we already had this stuff laying around.”

All of this makes you wonder what else CCSE and ANAG have lying around, and where it might lead them next.

Research funding: [ASCR-MICS](#), [SciDAC](#), [NASA](#), [CDWR](#)

RESEARCH NEWS

This section of the Annual Report provides a representative sampling of recent discoveries made with the help of NERSC's computing, storage, networking, and intellectual resources. News items are organized by science category, with funding organizations indicated by acronyms in [blue type](#). Organizational acronyms are spelled out in Appendix G.

ACCELERATOR PHYSICS

OPTIMIZING LUMINOSITY IN HIGH-ENERGY COLLIDERS

The beam-beam interaction puts a strong limit on the performance of most high-energy colliders. The electromagnetic fields generated by one beam focus or defocus the beam moving in the opposite direction, reducing luminosity and sometimes causing beam blowup. Qiang et al. performed the first-ever million-particle, million-turn, “strong-strong” (i.e., self-consistent) simulation of colliding beams (in the Large Hadron Collider). Instead of a conventional particle-mesh code, they used a method in which the computational mesh covers only the largest of the two colliding beams, allowing them to study long-range parasitic collisions accurately and efficiently.

J. Qiang, M. Furman, and R. Ryne, “Strong-strong beam-beam simulation using a Green function approach,” *Phys. Rev. ST AB* **5**, 104402 (2002). [HEP, SciDAC](#)

ELECTRON CLOUDS AND PARTICLE ACCELERATOR PERFORMANCE

Electron clouds have been shown to be associated with limitations in particle-accelerator performance in several of the world's largest circular proton and positron machines. Rumolo et al. have studied the interaction between a low-density electron cloud in a circular particle accelerator with a circulating charged particle beam. The particle beam's space charge attracts the cloud, enhancing the cloud density near the beam axis (Figure 1). This enhanced charge and the image charges associated with the cloud and the conducting wall of the accelerator may have important consequences for the dynamics of the beam propagation.

G. Rumolo, A. Z. Ghalam, T. Katsouleas, C. K. Huang, V. K. Decyk, C. Ren, W. B. Mori, F. Zimmermann, and F. Ruggiero, “Electron

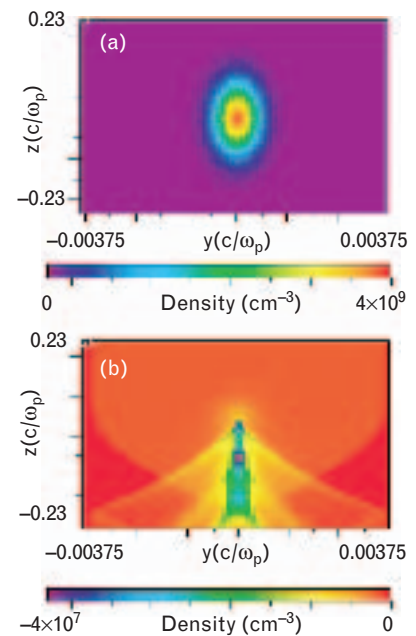


FIGURE 1 (a) Positively charged beam density. (b) The corresponding cloud density. This figure shows cloud compression on the axis of the beam.

cloud effects on beam evolution in a circular accelerator,” *Phys. Rev. ST AB* **6**, 081002 (2003). [HEP, SciDAC, NSF](#)

DEVELOPING FUTURE PLASMA-BASED ACCELERATORS

High-gradient acceleration of both positrons and electrons is a prerequisite condition to the successful development of a plasma-based linear collider. Blue et al. provided simulation support for the E-162 Plasma Wakefield Accelerator experiment at the Stanford Linear Accelerator Center. This work was successful at demonstrating the high-gradient acceleration of positrons in a meter-scale plasma for the first time. Excellent agreement was found between the experimental results and those from 3D particle-in-cell simulations for both energy gain and loss.

B. E. Blue, C. E. Clayton, C. L. O'Connell, F.-J. Decker, M. J. Hogan, C. Huang, R. Iverson, C. Joshi, T. C. Katsouleas, W. Lu, K. A. Marsh, W. B. Mori, P. Muggli, R. Siemann, and D. Walz, "Plasma-wakefield acceleration of an intense positron beam," *Phys. Rev. Lett.* **90**, 214801 (2003). [HEP](#), [SciDAC](#), [NSF](#)

ELECTRON CLOUD DEVELOPMENT IN PROTON STORAGE RINGS

Pivi and Furman have simulated electron cloud development in the Proton Storage Ring (PSR) at Los Alamos National Laboratory and the Spallation Neutron Source (SNS) under construction at Oak Ridge National Laboratory. Their simulation provides a detailed description of the secondary electron emission process, including a refined model for the emitted energy spec-

trum and for the three main components of the secondary yield, namely, the true secondary, redifused, and backscattered components.

M. T. F. Pivi and M. A. Furman, "Electron cloud development in the Proton Storage Ring and in the Spallation Neutron Source," *Phys. Rev. ST AB* **6**, 034201 (2003). [HEP](#), [SNS](#)

APPLIED MATHEMATICS

HIGH-PERFORMANCE SPARSE-MATRIX ALGORITHMS

Sparse-matrix problems are at the heart of many scientific and engineering applications of importance to DOE and the nation. Some applications are fusion energy, accelerator physics, structural dynamics, computational chemistry, and groundwater simulations. Teranishi et al. studied the effect of data mapping on the performance of triangular solutions. They showed that a careful choice of data mapping, coupled with the use of selective inversion, can significantly reduce the amount of communication required in the solution of a sparse triangular linear system.

K. Teranishi, P. Raghavan, and E. G. Ng, "Reducing communication costs for parallel sparse triangular solution," in *Proceedings of the SC2002 Conference* (2002). [ASCR-MICS](#), [SciDAC](#), [NSF](#)

OPTIMIZING MEMORY HIERARCHY

The BeBOP toolkit automatically tunes scientific codes by applying statistical techniques to produce highly tuned numerical subroutines. Vuduc et al. used the BeBOP tool SPARSITY to investigate memory hierarchy issues with such sparse matrix kernels as Ax (matrix-vector products) and $A^T Ax$ (normal equation products). They evaluated optimizations on a benchmark set of 44 matrices and four platforms, showing speedups of up to a factor of 4.2.

R. Vuduc, A. Gyulassy, J. Demmel, and K. Yelick, "Memory hierarchy optimizations and performance bounds for sparse $A^T Ax$." U.C. Berkeley Technical Report UCB/CS-03-1232, February 2003. [ASCR-MICS](#), [SciDAC](#), [NSF](#), [Intel](#)

MODELING THE COMPONENT STRUCTURES OF FLAMES

Combustion in turbulent flows may take the form of a thin flame wrapped around vortical structures. Marzouk et al. used the flame-embedding approach to decouple

computations of the “outer” non-reacting flow and the combustion zone by breaking the flame surface into a number of elemental flames, each incorporating the local impact of unsteady flow-flame interaction (Figure 2). Results show that using the flame leading-edge strain rate gives reasonable agreement in cases of low-strain-rate and weak-strain-rate gradient within the flame structure. This agreement deteriorates substantially when both are high. Agreement between the 1D model and the 2D simulation improves greatly when the actual strain rate at the reaction zone of the 1D flame is made to match that of the 2D flame.

Y. M. Marzouk, A. F. Ghoniem, and H. N. Najm, “Toward a flame embedding model for turbulent combustion simulation,” *AIAA Journal* **41**, 641 (2003). [ASCR-MICS](#)

IMPROVING THE ACCURACY OF FRONT TRACKING IN FLUID MIXING

Acceleration-driven fluid-mixing instabilities play important roles in inertially confined nuclear fusion, as well as in stockpile stewardship. Turbulent mixing is a difficult and centrally important issue for fluid dynamics that affects such questions as the rate of heat transfer by the Gulf Stream, resistance of pipes to fluid flow, combustion rates in automotive engines, and the late-time evolution of a supernova. To provide a better understanding of these instabilities, Glimm et al. have developed a new, fully conservative front tracking algorithm that results in one-order-of-accuracy improvement over most finite difference simulations and shows good agreement with experimental results.

J. Glimm, X.-L. Li, Y.-J. Liu, Z. L. Xu, and N. Zhao, “Conservative front tracking with improved accuracy,” *SIAM J. Sci. Comp.* (in press). [ASCR-MICS](#), [FES](#), [ARO](#), [NSF](#), [LANL](#), [NSFC](#)

PREDICTING THE ENERGIES OF QUANTUM WELL STATES

Understanding the electronic structures of quantum-well states (QWSs) in metallic thin films can lay the groundwork for designing new materials with predictable properties. An et al. compared the experimental electronic properties of QWSs in a copper film grown on a cobalt substrate with *ab initio* simulations. The calculations confirm the concept of the quantization condition inherent in the phase-accumulation model (PAM) to predict the energies of QWSs as a

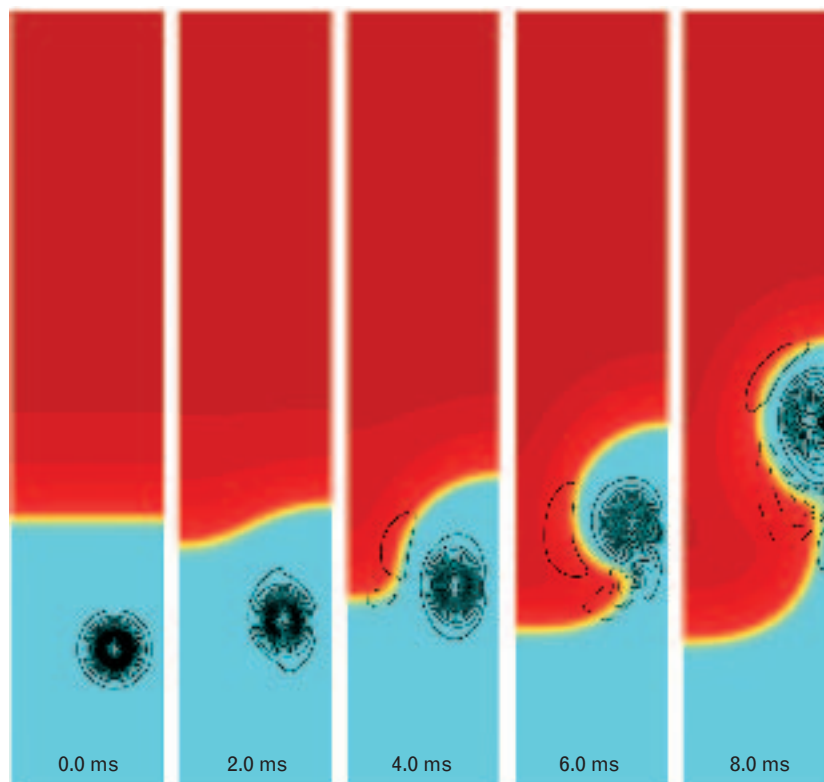


FIGURE 2 Premixed flame interaction with a counter-rotating vortex pair. The vertical right-hand edge of each frame is the centerline of the vortex pair, only half of which is shown. Colored contours indicate gas temperature, with darker shading corresponding to burned combustion products. Solid/dashed contours delineate levels of positive/negative vorticity.

function of their thickness, and to provide new insight into their nature. The results also show that band structures and reflection phases obtained from either experiment or *ab initio* theory can quantitatively predict QWS energies within the PAM model.

J. M. An, D. Raczkowski, Y. Z. Wu, C. Y. Won, L. W. Wang, A. Canning, M. A. Van Hove, E. Rotenberg, and Z. Q. Qiu, "The quantization condition of quantum-well states in Cu/Co(001)," *Phys. Rev. B* **68**, 045419 (2003). [ASCR-MICS](#), [NSF](#)

REDUCED-DENSITY MATRICES FOR ELECTRONIC STRUCTURE CALCULATIONS

Accurate first-principles electronic structure calculations are of special importance for spectroscopy and for chemical-reaction dynamics.

Reduced-density matrices (RDM) that have linear scaling for large N offer a first-principles alternative to density-functional and semi-empirical calculations for large systems. Zhao et al. reformulated the existing RDM method from a primal to a dual formulation, substantially reducing the computational cost. They computed the ground state energy and dipole moment of 47 small molecules and molecular ions, both open and closed shells, achieving results more accurate than those from other approximations.

Z. Zhao, B. Braams, M. Fukuda, M. Overton, and J. Percus, "The reduced density matrix method for electronic structure calculations and the role of three-index representability," *J. Chem. Phys.* (submitted, 2003). [ASCR-MICS](#), [NSF](#)

ASTROPHYSICS

CALCULATING ELECTRON CAPTURE RATES DURING CORE COLLAPSE

Supernova simulations to date have assumed that during core collapse, electron captures occur dominantly on free protons, while captures on heavy nuclei are Pauli blocked and are ignored. Langanke et al. have calculated rates for electron capture on nuclei with mass numbers $A = 65$ –112 for the temperatures and densities appropriate for core collapse. They found that these rates are large enough so that, in contrast to previous assumptions, electron capture on nuclei dominates over capture on free protons. This result leads to significant changes in core collapse simulations.

K. Langanke, G. Martinez-Pinedo, J. M. Sampaio, D. J. Dean, W. R. Hix, O. E. B. Messer, A. Mezzacappa, M. Liebendorfer, H.-Th. Janka, and M. Rampp, "Electron capture rates on nuclei and implications for stellar core collapse," *Phys. Rev. Lett.* **90**, 241102 (2003). [NP](#), [SciDAC](#), [DRA](#), [JHIR](#), [MCYT](#), [NASA](#), [NSF](#), [PECASE](#)

SIMULATING BINARY STAR FORMATION

Developing a comprehensive theory of star formation remains one of the most elusive goals of theoretical astrophysics, partly because gravitational collapse depends on initial conditions within giant molecular cloud cores that only recently have been observed with sufficient accuracy to permit a realistic attack on

the problem. Klein et al. have simulated the gravitational collapse and fragmentation of marginally stable turbulent molecular cloud cores




FIGURE 3

Simulation of the birth of a binary star system. Each member of the binary appears to be surrounded by a disk, with one component of the binary showing a two-arm spiral. The binary appears to form inside the structure of a filament in the cloud core.

PROBING THE CHEMICAL EVOLUTION OF THE UNIVERSE

Type II supernovae have a very large spread in their intrinsic brightness—greater than a factor of 500—because of their diverse progenitors. Despite this diversity, Baron et al. have shown that their atmospheres can be well understood with detailed synthetic spectral modeling that can determine the stellar compositions, degree of mixing, and kinetic energy of the explosion. These results make Type II supernovae attractive probes of chemical evolution in the Universe and potentially useful independent checks on the cosmological results derived from Type Ia supernovae.

E. Baron, P. Nugent, D. Branch, P. Hauschildt, M. Turatto, and E. Cappellaro, “Determination of primordial metallicity and mixing in the Type IIP supernova 1993W,” *Astrophysical Journal* **586**, 1199 (2003). [NP](#), [IBM SUR](#), [NASA](#), [NSF](#)

CHECKING THE FOREGROUND OF THE COSMIC MICROWAVE BACKGROUND

Observations of the cosmic microwave background (CMB) can be contaminated by diffuse foreground emission from sources such as galactic dust and synchrotron radiation, but an international team of researchers has developed a technique for quantifying this effect. They applied this technique to CMB data from the MAXIMA-1

experiment and found that the effect on CMB power spectrum observations is negligible.

A. H. Jaffe, A. Balbi, J. R. Bond, J. Borrill, P. G. Ferreira, D. Finkbeiner, S. Hanany, A. T. Lee, B. Rabbii, P. L. Richards, G. F. Smoot, R. Stompor, C. D. Winant, and J. H. P. Wu, “Determining foreground contamination in CMB observations: Diffuse galactic emission in the MAXIMA-1 field,” *Astrophysical Journal* (in press); [astro-ph/0301077](#). [HEP](#), [NASA](#), [NSF](#)

SIMULATING BLACK HOLE MERGERS AND GRAVITATIONAL WAVES

Simulations of the gravitational waves resulting from the collision of two black holes will provide patterns

that new gravitational wave detectors can look for as they attempt to verify the predictions of General Relativity. A research team led by Edward Seidel of the Albert Einstein Institute/Max Planck Institute for Gravitation Physics in Potsdam, Germany, has determined effective computational and theoretical (gauge) parameters needed to carry out binary black hole evolutions for time scales never before achieved. Another research group, led by Carlos Lousto of the University of Texas at Brownsville, has achieved a full numerical computation of the gravitational radiation generated by the evolution of a series of binary black hole configurations with aligned and counter-aligned spins.

M. Alcubierre, B. Bruegmann, P. Diener, M. Koppitz, D. Pollney, E. Seidel, and R. Takahashi, "Gauge conditions for long-term numerical black hole evolutions without excision," *Phys. Rev. D* **67**, 084023 (2003). [NP](#), [AEI](#), [EU](#), [LSU](#)

J. Baker, M. Campanelli, C. O. Lousto, and R. Takahashi, "The final plunge of spinning binary black holes," *Phys. Rev. D* (submitted); [astro-ph/0305287](#) (2003). [NP](#), [NASA](#), [NRC](#), [NSF](#)

CHEMICAL SCIENCES

CALCULATING THE INSTABILITY OF NITROGEN SALTS

Can the combination of a stable polynitrogen cation, such as N_5^+ , with a stable polynitrogen anion, such as N_3^- , lead to a stable ionic nitrogen allotrope, such as $N_5^+N_3^-$? Dixon et al. used high-level *ab initio* molecular orbital theory to conclusively determine whether the lattice energies of $N_5^+N_3^-$ and $N_5^+N_5^-$ are sufficient to stabilize them as solids. Their calculations showed that neither salt can be stabilized and that both should decompose spontaneously into N_3 radicals and N_2 . This conclusion was experimentally confirmed for the $N_5^+N_3^-$ salt.

D. A. Dixon, D. Feller, K. O. Christe, W. W. Wilson, A. Vij, V. Vij, H. D. Brooke Jenkins, R. M. Olson, and M. S. Gordon, "Enthalpies of formation of gas phase N_3 , N_3^- , N_5^+ , and N_5^- from *ab initio* molecular orbital theory, stability predictions for $N_5^+N_3^-$ and $N_5^+N_5^-$, and experimental evidence for the instability of $N_5^+N_3^-$," *J. Am. Chem. Soc.* (in press). [BES](#), [BER](#), [SciDAC](#), [DARPA](#), [AFOSR](#), [NSF](#)

A MORE ACCURATE CALCULATION OF SILICON DICARBIDE

Unlike many typical chemical systems, in which errors due to incomplete treatment of the one-electron basis and electron correlation are balanced and largely cancelled at modest levels of theory, silicon dicarbide has a barrier to linearity that requires both

prodigious basis sets and high-order correlation treatments for accurate predictions. Kenny et al. calculated molecular partial-wave expansions for the T-shaped and linear forms of SiC_2 using the MP2-R12/A method. Their results are the first of sufficient accuracy to call for revisitation of the spectroscopically determined barrier of 5.4 ± 0.6 kcal mol⁻¹, moving it upward to 6.3 kcal mol⁻¹.

J. P. Kenny, W. D. Allen, and H. F. Schaefer, "Complete basis set limit studies of conventional and R12 correlation methods: The silicon dicarbide (SiC_2) barrier to linearity," *J. Chem. Phys.* **118**, 7353 (2003). [BES](#), [SciDAC](#)

HOW ELECTRON IMPACT EXCITES A MOLECULE

A major goal of studying electron-molecule collisions is to explore the mechanisms that control the flow of energy from electronic to nuclear degrees of freedom. McCurdy et al. have completed the first calculation on vibrational excitation of a molecule (CO_2) via coupled resonant states (Figure 4). The multidimensional nuclear dynamics on these coupled surfaces was necessary to explain quantitatively, for the first time, the experimental observations on this system. These methods have been applied in calculations on electron-impact dissociation of water, which is one of the principal mechanisms suspected to initiate radiation damage of biological tissue in environments with ionizing radiation.

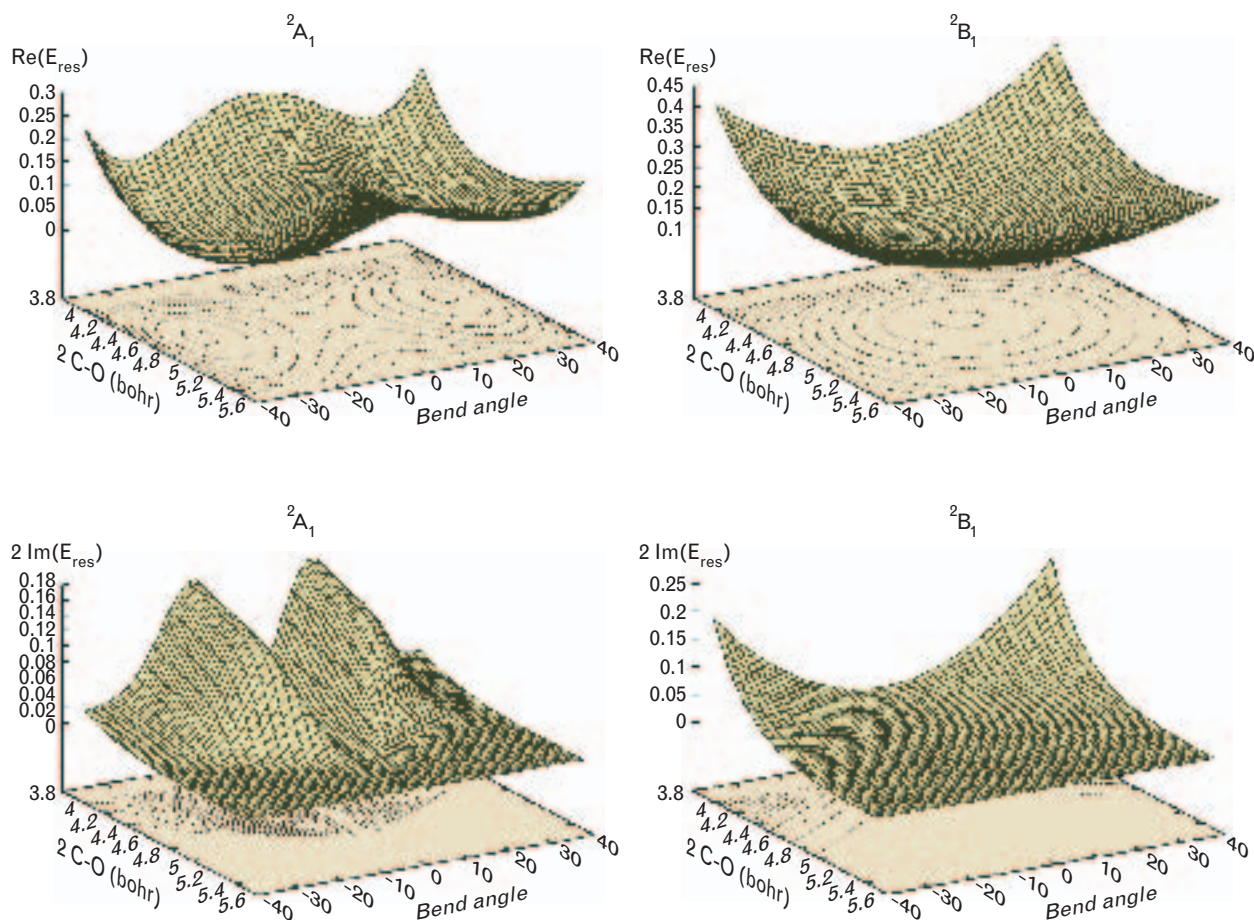


FIGURE 4
Complex potential surfaces for both components of the resonance. Top row, real part of resonance energies; bottom row, widths. Energies are in hartrees, and the bend angle, defined as π minus the O—C—O bond angle, is in degrees.

C. W. McCurdy, W. A. Isaacs, H.-D. Meyer, and T. N. Rescigno, "Resonant vibrational excitation of CO₂ by electron impact: Nuclear dynamics on the coupled components of the $^2\Pi_u$ resonance," *Phys. Rev. A* **67**, 042708 (2003). [BES](#)

DESIGNING NANO-SCALE SHOCK ABSORBERS

Carbon nanotubes are the fundamental building blocks in many nanosystems. Rivera et al. performed molecular dynamics computational "experiments" on double-walled carbon nanotubes of varying lengths

and diameters in which they pulled the inner nano-tube out of the outer nanotube to several different distances and allowed the nanotube to retract and subsequently oscillate. They found that the oscillation is damped (hence acting like a nano shock absorber) and can be predicted by a very simple mechanical model that takes into account the sliding friction between the surfaces. The oscillation was around 1 GHz, consistent with prior experimental and theoretical calculations.

J. L. Rivera, C. McCabe, and P.T.

Cummings, "Oscillatory behavior of double nanotubes under extension: A simple nanoscale damped spring," *Nano Letters* **3**, 1001 (2003). [BES](#), [NSF](#)

SOLVING THE SCHRÖDINGER EQUATION FOR LARGE SYSTEMS

Exact solutions of the Schrödinger equation within a given one-particle basis set, lie at the heart of quantum chemistry, providing invaluable benchmarks by which the accuracy of more approximate methods may be judged. Chan and Head-Gordon

used a newly developed density matrix renormalization group algorithm to compute exact solutions of the Schrödinger equation for water at two geometries in a basis of 41 orbitals, demonstrating that it is now possible to obtain numerically exact solutions for systems considerably larger than can be contemplated using traditional full configuration interaction algorithms.

G. K. L. Chan and M. Head-Gordon, "Exact solution (within a triple zeta, double polarization basis set) of the electronic Schrödinger equation for water," *J. Chem. Phys.* **118**, 8551 (2003). [BES](#)

QUANTUM MONTE CARLO STUDY OF THE ELECTRONIC EXCITED STATES

Accurate computational predictions of molecular electronic excited states have proved more difficult to obtain than ground states. Ethylene is the prototypical π -electron system whose photochemical behavior is important in chemistry, biology, and technology. El Akramine et al. carried out a theoretical study of the transition between the ground state and the lowest triplet state of ethylene using the diffusion Monte Carlo method. The ground state atomization energy and heat of formation were consistent with experimental results, and the triplet-state atomization energy and heat of formation were also predicted.

O. El Akramine, A. C. Kollias, and W. A. Lester, Jr., "Quantum Monte Carlo study of the singlet-triplet transition in ethylene," *J. Chem. Phys.* **119**, 1483 (2003). [BES](#)

OXYGEN DISSOCIATION ON COPPER STEPS

Copper and copper alloy catalysts are at the heart of important industrial processes, and copper also plays an important role in the microelectronics industry. Dissociation of dioxygen (O_2) is the first step of oxygen chemistry on the surface of metals such as copper. Xu and Mavrikakis studied the adsorption and dissociation of dioxygen on copper steps using periodic self-consistent density functional theory calculations, and found that the adsorption of both atomic and molecular oxygen is enhanced on the stepped surface.

Y. Xu and M. Mavrikakis, "The adsorption and dissociation of O_2 molecular precursors on Cu: The effect of steps," *Surface Science* **538**, 219 (2003). [BES](#), [DOD](#), [NSF](#), [3M](#)

A NEW METHOD FOR WAVE-PACKET PROPOGATION

Wang and Thoss have presented a method for wave-packet propagation that is based rigorously on a variational principle and is also sufficiently efficient to be applicable to complex dynamical problems with up to a few hundred degrees of freedom. This method can treat substantially more physical degrees of freedom than the original method, and thus significantly enhances the ability to carry

out quantum dynamical simulations for complex molecular systems. The efficiency of the new formulation was demonstrated by converged quantum dynamical simulations for systems with a few hundred to a thousand degrees of freedom.

H. Wang and M. Thoss, "Multilayer formulation of the multiconfiguration time-dependent Hartree theory," *J. Chem. Phys.* **119**, 1289 (2003). [BES](#), [NSF](#), [DAAD](#)

ADDING HYDROGEN TO LEAN PREMIXED NATURAL GAS

Hawkes and Chen performed direct numerical simulation of lean premixed methane-air and hydrogen-methane-air flames to understand how the addition of hydrogen can affect flame stability and pollutant formation in gas turbines. They found a greater rate of heat release in the enriched flame, which is explained by a difference in flame areas and influences of turbulent stretch on the local burning rates. They found that there may be a pollutant tradeoff between NO and CO emissions when using hydrogen-enriched fuels.

E. R. Hawkes and J. H. Chen, "Turbulent stretch effects on hydrogen enriched lean premixed methane-air flames," 3rd Joint Meeting of the U.S. Sections of the Combustion Institute, Chicago, IL (2003). [BES](#)

A MILLENNIUM OF CLIMATE CHANGE

Researchers at the National Center for Atmospheric Research continued their benchmark 1,000-year Community Climate System Model (CCSM) control run and extended the Parallel Climate Model (PCM) control run to nearly 1,500 years (Figure 5). Science experiments with the PCM included historical 20th century climate simulations using various combinations of specified greenhouse gases, atmospheric sulfate concentrations, solar and volcanic forcing, and tropospheric and stratospheric ozone. This set of

experiments represents the most extensive set of 20th century forcing climate simulations ever attempted.

A. Dai, A. Hu, G. A. Meehl, W. M. Washington, and W. G. Strand, "North Atlantic Ocean circulation changes in a millennial control run and projected future climates," *Journal of Climate* (in press). [BER](#), [SciDAC](#), [NSF](#)

MULTIRESOLUTION CLIMATE MODELING

Baer and colleagues have developed the Spectral Element Atmospheric Model (SEAM) to incorporate local mesh refinement within a global model. This makes it feasible to

model the interaction between coarse-scale and fine-scale phenomena. The method implies improved climate and turbulence simulations by allowing increased resolution in a few dynamically significant areas, such as fronts or regions of dense topography. Similarly, regional climate simulations may be improved by allowing regional resolution to be incorporated within a global model in a two-way interactive fashion.

F. Baer and J. J. Tribbia, "Sensitivity of atmospheric prediction models to amplitude and phase," *Tellus* (submitted, 2003). [BER](#), [SciDAC](#)

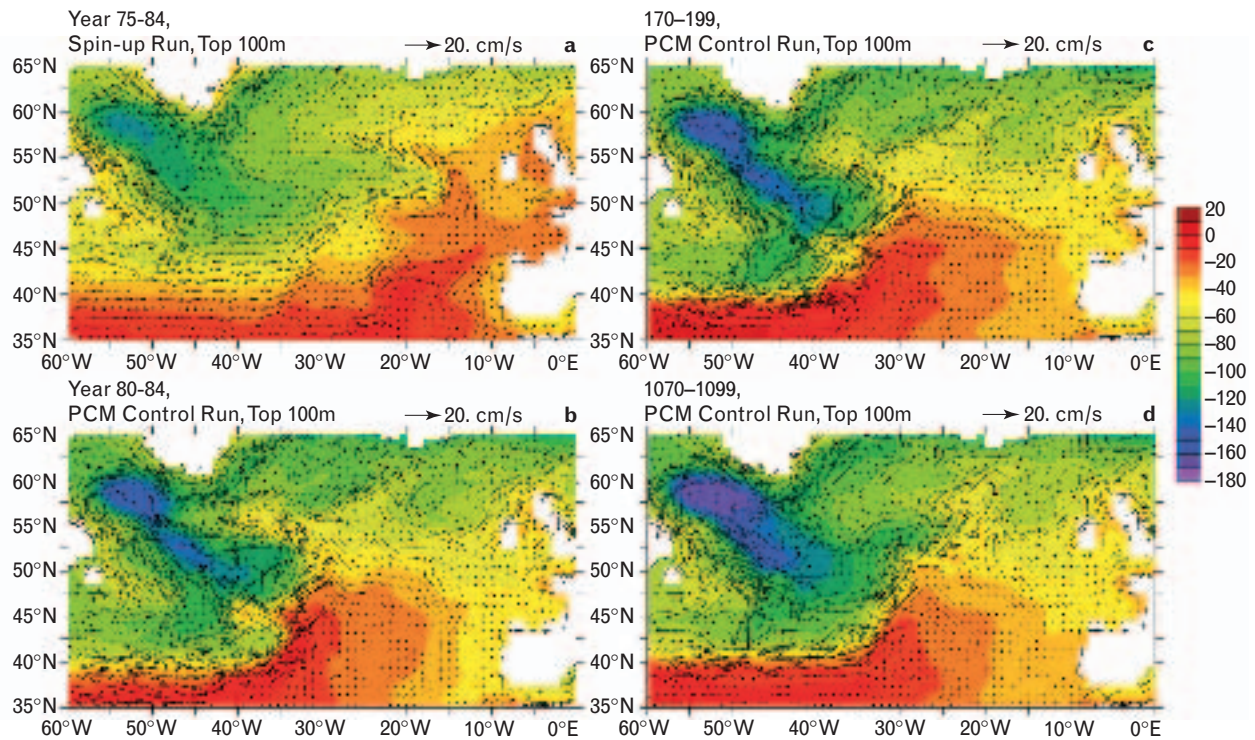


FIGURE 5 Sea surface height (color, in cm relative to a motionless ocean) and top 100m-averaged ocean currents (arrows) in the North Atlantic Ocean simulated by the PCM (a) at the end of the ocean-sea ice alone run (right before the coupling), and during the control run at (b) years 80–84, (c) 170–199, and (d) 1070–1099.

SORTING OUT THE SOURCES OF TROPOSPHERIC OZONE

Ozone in the troposphere is both formed *in situ* and transported from the stratosphere, which has much higher ozone levels. The importance of each source may vary both regionally and seasonally. Atherton et al. analyzed the results of a simulation in Phoenix from April 1 through October 31, 2001. Their results showed that there are periods of time during the “ozone season” in which higher levels of surface ozone may be due to transport of stratospheric ozone from above, rather than created by ground-level (and energy-related) processes.

C. Atherton, D. Bergmann, P. Cameron-Smith, P. Connell, D. Rotman, and J. Tannahill, “Large scale atmospheric chemistry simulations for 2001: An analysis of ozone and other species in central Arizona,” American Meteorological Society 83rd Annual Meeting, Long Beach, California, February 2003. [BER](#)

DO SOOT AND SMOKE CONTRIBUTE TO CLIMATE CHANGE?

Soot and smoke contain black carbon which absorbs solar radiation. These aerosols have been thought to enhance climate warming by absorbing radiation that might otherwise be scattered and by reducing overall cloudiness. Penner et al. evaluated forcing and changes in atmospheric temperature and clouds associated with both direct and indirect effects of soot and smoke using an atmospheric-climate model coupled to their chemical-

transport model. They found that climate forcing by soot and smoke was not as large as previously claimed, because the adjustment of the atmospheric temperature and clouds to aerosols tends to decrease the forcing, particularly if the aerosols reside at higher levels within the atmosphere.

J. E. Penner, S.Y. Zhang, and C. C. Chuang, “Soot and smoke aerosol may not warm climate,” *J. Geophys. Res.*, in press (2003). [BER](#)

CARBON-CYCLE FEEDBACK IN CLIMATE CHANGE

Taking into account the feedback effects of climate change on absorption of carbon by the ocean and terrestrial biosphere will allow better predictions of future climate; however, these effects are ignored in present U.S. climate models. Thompson et

al. have obtained the first results from a fully coupled, non-flux-corrected, multi-century climate-carbon simulation. They have performed three complete simulations: a control case with no emissions, a standard case, and a case where the uptake of CO₂ by vegetation is limited. The limited case shows that substantially more CO₂ stays in the atmosphere when carbon uptake by vegetation is reduced, thus giving a range of uncertainty in current carbon-climate modeling.

S. Thompson, B. Govindasamy, J. Milovich, A. A. Mirin, and M. Wickett, “A comprehensive GCM-based climate and carbon cycle model: Transient simulations to year 2100 using prescribed greenhouse gas emissions,” Amer. Geophys. Union Mtg., San Francisco, CA, December 2002; Lawrence Livermore National Laboratory technical abstract UCRL-JC-149761. [BER](#), [LLNL](#)

COMPUTER SCIENCE

REDUCING THE PERFORMANCE GAP

The SciDAC Performance Evaluation Research Center (PERC) is working toward reducing the gap between peak and sustained performance on high-end computers by developing hardware and software tools to monitor, model, and control performance. In the past year, PERC has analyzed and optimized application codes in astrophysics, lattice gauge theory, climate modeling,

applied mathematics, and other disciplines, significantly improving the performance and scalability of the codes. For example, improved compiler annotation technology (Active Harmony) achieved a 15% reduction in run time on the POP ocean modeling code.

D. H. Bailey, B. de Supinski, J. Dongarra, et al., “Performance technologies for peta-scale systems,” <http://perc.nersc.gov/docs/PERC-white-aes15.doc>. [ASCR-MICS](#), [SciDAC](#)

INTEGRATING THE GLOBUS TOOLKIT WITH FIREWALLS

The Globus Toolkit is a collection of clients and services that enable Grid computing. Each component of the toolkit has its own network traffic characteristics that must be considered when deploying the toolkit at a site with a firewall. Members of the Globus Project and the DOE Science Grid engineering team have developed requirements and guidance for firewall administrators, describing the network traffic generated by the Globus Toolkit services, how the Toolkit's use of the ephemeral port range can be controlled, and firewall requirements for client and server sites. NERSC provided a firewall testbed that was used for verifying much of the information in this document.

V. Welch et al., "Globus Toolkit Firewall Requirements Version 5," July 22, 2003, <http://www.globus.org/security/firewalls/>.
SciDAC

PRECONDITIONING INEXACT NEWTON ALGORITHMS

Inexact Newton algorithms are commonly used for solving large sparse nonlinear system of equations $F(u^*) = 0$. Even with global strategies such as line search or trust region, the methods often stagnate at local minima of $\|F\|$, especially for problems with unbalanced nonlinearities, because the methods do not have built-in machinery to deal with the unbalanced nonlinearities. Cai and Keyes studied a nonlinear additive Schwarz-based parallel

nonlinear preconditioner; they showed numerically that the new method converges well even for some difficult problems, such as high Reynolds number flows, where a traditional inexact Newton method fails.

X.-C. Cai and D. E. Keyes, "Nonlinearly preconditioned inexact Newton algorithms," *SIAM J. Sci. Comp.* **24**, 183 (2002). [ASCR-MICS](#), [SciDAC](#), [NSF](#), [NASA](#)

DEVELOPING AN ALTERNATIVE PROGRAMMING MODEL

While MPI is currently the de facto standard for programming supercomputers, several federal agencies are investing in alternative models to simplify programming. Yelick and

collaborators are investigating UPC, a "partitioned global address space" language. They developed and released an open-source version of the full UPC compiler, which translates UPC to C with calls to a newly developed UPC runtime layer. The open-source compiler has low overhead for basic operations on shared data; it is competitive with, and sometimes faster than, the commercial compiler.

W. Chen, D. Bonachea, J. Duell, P. Husbands, C. Iancu, and K. Yelick, "A performance analysis of the Berkeley UPC compiler," 17th Annual International Conference on Supercomputing (ICS), 2003. [ASCR-MICS](#), [DOD](#), [NSF](#)

FUSION ENERGY SCIENCES

GEODESIC ACOUSTIC MODES IN PLASMA TURBULENCE

Researchers in the SciDAC Plasma Microturbulence Project have observed a coherent oscillation in the poloidal flow of density fluctuations that shows remarkable similarity to predicted characteristics of geodesic acoustic modes with beam emission spectroscopy in DIII-D plasmas. Two methods of analyzing the data, one based on wavelets and one based on time-resolved cross-correlation, reveal similar features in the resulting flow field. These experimental observations

and the corresponding simulations provide compelling evidence that geodesic acoustic modes are active in the turbulence in the outer regions of tokamak plasmas and play an integral role in affecting and mediating the fully saturated state of turbulence and resulting transport.

G. R. McKee, R. J. Fonck, M. Jakubowski, K. H. Burrell, K. Hallatschek, R. A. Moyer, D. L. Rudakov, W. Nevins, G. D. Porter, P. Schoch, and X. Xu, "Experimental characterization of coherent, radially-sheared zonal flows in the DIII-D tokamak," *Phys. Plasmas* **10**, 1712 (2003). [FES](#), [SciDAC](#)

COMPUTATIONAL ATOMIC PHYSICS FOR FUSION ENERGY

Colgan and Pindzola have applied the time-dependent close-coupling theory to the study of the electron-impact ionization of helium from the excited ($1s2s$) configuration. They made the calculations in an effort to resolve the discrepancy between theoretical calculations and existing experimental measurements for electron scattering from the metastable states of helium. The relative difference between the non-perturbative close-coupling and the perturbative distorted-wave results grew larger as the principal quantum number was increased. This difference has important implications in the collisional-radiative modeling of many astrophysical and laboratory plasmas that use perturbation theory.

J. P. Colgan and M. S. Pindzola, "Time-dependent close-coupling studies of the electron-impact ionization of excited-state helium," *Phys. Rev. A* **66**, 062707 (2002). [FES](#), [SciDAC](#)

BALLOONING INSTABILITY IN THE EARTH'S MAGNETOTAIL

The Earth's magnetotail, the main source of the polar aurora, is the part of the magnetosphere that is pushed away from the sun by the solar wind. The ballooning instability of the magnetotail interests researchers because of its possible relevance to the dynamics of magnetic substorms. Zhu et al. studied this ballooning instability within the framework of Hall magnetohydrodynamics, which

enter the stability analysis through changes introduced in the plasma compressibility.

P. Zhu, A. Bhattacharjee, and Z. W. Ma, "Hall magnetohydrodynamic ballooning instability in the magnetotail," *Phys. Plasmas* **10**, 249 (2003). [FES](#), [SciDAC](#), [NSF](#)

SIZE SCALING OF TURBULENT TRANSPORT

Lin et al. have developed a nonlinear model for turbulence radial spreading based on the modified porous-medium equation. The model offers a phenomenological understanding of the transition from Bohm to gyro-Bohm scaling when the parameter (minor radius/gyroradius) is larger than 500. They also estimated the role of the trapped electron nonlinearity in zonal flow generation in trapped-electron-mode turbulence in the context of parametric instability theory.

Z. Lin, T. S. Hahm, S. Ethier, W. W. Lee, J. Lewandowski, G. Rewoldt, W. M. Tang, W. X. Wang, L. Chen, and P. H. Diamond, "Size scaling of turbulent transport in tokamak plasmas," *Proceedings of 19th IAEA Fusion Energy Conference*, Lyon, France (2002). [FES](#), [SciDAC](#)

ELECTRON-IMPACT IONIZATION OF OXYGEN IONS

Loch et al. made experimental measurements of the ionization cross section for O^{q+} , where $q = 1-4$, performed with a crossed-beams apparatus, and compared them with theoretical calculations. For O^+ , the experimental measurements are in very good agreement with configura-

tion-average time-dependent close-coupling calculations. For the remaining oxygen ions, the experimental measurements are in good agreement with time-independent distorted-wave calculations. As expected, the accuracy of the perturbative distorted-wave calculations improves with increasing ion charge.

S. D. Loch, J. P. Colgan, M. S. Pindzola, W. Westermann, F. Scheuermann, K. Aichele, D. Hathiramani, and E. Salzborn, "Electron-impact ionization of O^{q+} ions for $q = 1-4$," *Phys. Rev. A* **67**, 042714 (2003). [FES](#), [SciDAC](#), [DFG](#)

SIMULATION OF INTENSE BEAMS FOR HEAVY-ION FUSION

A multibeamlet approach to a high-current-ion injector, whereby a large number of beamlets are accelerated and then merged to form a single beam, offers a number of potential advantages over a monolithic single-beam injector. These advantages include a smaller transverse footprint, more control over the shaping and aiming of the beam, and more flexibility in the choice of ion sources. A potential drawback, however, is a larger emittance. Grote et al. are studying the merging of beamlets (Figure 6) and how this merging determines emittance. Their results suggest that a multibeamlet injector can be built with a normalized emittance less than 1π mm mrad.

D. P. Grote, E. Henestroza, and J. W. Kwan, "Design and simulation of a multibeamlet injector for a high-current accelerator," *Phys. Rev. ST AB* **6**, 014202 (2003). [FES](#)

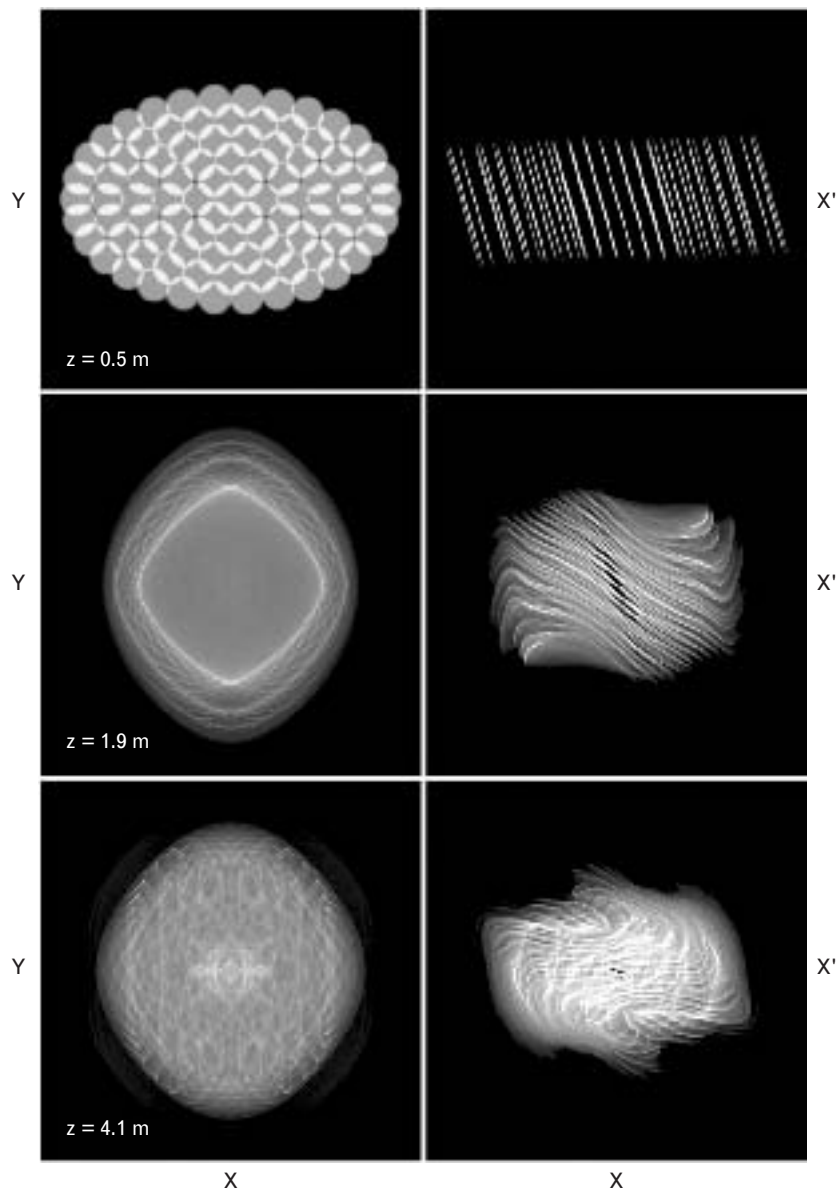


FIGURE 6

The merging of the beamlets creates complex space-charge waves that propagate across the beam and mix. This figure shows contours of density at several locations in configuration space (x - y) and in phase space (x - x'). They show that the individuality of the beamlets is quickly lost, leaving a nearly uniform beam. The z is the distance from the end of the preaccelerator column.

MODELING THE PLASMA-DENSITY LIMIT

Understanding the plasma-density limit is crucial for projecting the performance of future fusion reactors. Xu et al. have developed a model for the density limit in a tokamak plasma based on the edge simulation results from BOUT and UEDGE codes. Simulations of turbulence in tokamak boundary plasmas show that turbulent fluctuation levels and transport increase with collisionality. The simulations also show that it is easier to reach the density limit as the density increases while holding pressure constant than holding temperature constant.

X. Q. Xu, W. M. Nevins, T. D. Rognlien, R. H. Bulmer, M. Greenwald, A. Mahdavi, L. D. Pearlstein, and P. Snyder, "Transitions of turbulence in plasma density limits," *Phys. Plasmas* **10**, 1773 (2003). [FES, MIT, GA](#)

MICROSTRUCTURE EVOLUTION IN IRRADIATED MATERIALS

The promise of fusion as a viable energy source depends on the development of structural materials capable of sustaining operation in harsh radiation conditions. The structure and mobility of self-interstitial atom

SIMULATING THE THERMAL STRUCTURE OF SOLAR STORMS

Solar storms can potentially disrupt the operation of orbiting satellites, endanger astronauts, and cause failure of long-distance power grids on land. Understanding the physical mechanisms that originate these storms is of fundamental importance if they are to be forecasted. Mok et al. simulated the thermal structure of the solar atmosphere in

the neighborhood of a sunspot group in 3D for the first time. They found that the strong, structured magnetic field from the sunspots guides the heat flow and plasma flow, resulting in fine structures in temperature and density because of the complicated field-line topology.

Y. Mok, R. Lionello, Z. Mikic, and J. Linker, "Calculating the thermal structure of solar active regions in three dimensions," *Astrophys. J.* (submitted, 2003). [FES, NSF, NASA](#)

(SIA) clusters has profound significance on the microstructural evolution of irradiated materials. Marian et al. have studied the effect of substitutional copper solute atoms on the mobility of self-interstitial clusters in Fe-Cu alloys. The effect of this oversized substitutional solute is to enhance the three-dimensional character of small-cluster diffusion and to affect general cluster diffusion properties.

J. Marian, B. D. Wirth, A. Caro, B. Sadigh, G. R. Odette, J. M. Perlado, and T. Diaz de la Rubia, "Dynamics of self-interstitial cluster migration in pure alpha-Fe and Fe-Cu alloys," *Phys. Rev. B* **65**, 144102 (2002). [FES](#)

DESIGNING THE NATIONAL COMPACT STELLERATOR EXPERIMENT

The U.S. fusion community is constructing a compact, moderate-cost experiment, the National Compact Stellarator Experiment (NCSX), to investigate 3D confinement and stability. The configurations being studied are compact, are passively stable to known instabilities, and are designed to have a "hidden" symmetry to give good transport properties. In addition, because of their enhanced stability, they should be free of disruptions. These configurations are designed to combine the best features of the tokamak and the stellarator. The NCSX has achieved a successful conceptual design review.

G. H. Neilson and the NCSX Team, "Physics considerations in the design of NCSX," *Proceedings of 19th IAEA Fusion Energy Conference*, Lyon, France (2002). [FES](#)

STABILIZING MICRO-TURBULENCE IN PLASMAS

Bourdelle et al. have shown that microturbulence can be stabilized in the presence of steep temperature and density profiles. High values of $|\beta'|$ have a stabilizing influence on drift modes. This may form the basis for a positive feedback loop in which high core β values lead to improved confinement, and to further increase in β . In high- β spherical tokamak plasmas, high $|\beta'|$ rather than low aspect ratio is a source of stabilization. Therefore, the effect of high $|\beta'|$ should be stabilizing in the plasmas of the National Spherical Torus Experiment.

C. Bourdelle, W. Dorland, X. Garbet, G. W. Hammett, M. Kotschenreuther, G. Rewoldt, and E. J. Synakowski, "Stabilizing impact of high gradient of β on microturbulence," *Phys. Plasmas* **10**, 2881 (2003). [FES](#)

TRAPPED ELECTRON MODE TURBULENCE

Gyrokinetic simulations of turbulent particle transport in the Alcator C-Mod by Ernst et al. have led to an understanding of the observed spontaneous particle transport barrier formation, and its control using radio-frequency heating in the core. As the density gradient steepens, nonresonant, collisionless, trapped electron modes (TEMs) are driven unstable. Gyrokinetic stability analysis shows that the modes are driven by the density gradient, have wavelengths around twice the ion gyroradius, and are partially damped by collisions. After the TEM onset, the turbulent outflow strongly increases with the

density gradient. This outflow ultimately balances the inward Ware pinch, leading to a steady state.

D. R. Ernst et al., "Role of trapped electron mode turbulence in density profile control with ICRH in Alcator C-Mod," International Sherwood Fusion Theory Conference, Corpus Christi, Texas, April 28–30, 2003. [FES](#)

DESIGNING COMPACT TOKAMAK-STELLARATOR HYBRIDS

Ware et al. described for the first time the complete physics properties of compact, drift-optimized tokamak-stellarator hybrids with a high-shear tokamak-like rotational transform profile and $|B|$ that is quasipoloidally symmetric (Figure 7). The rotational transform in these hybrid configurations is produced primarily by a self-consistent bootstrap current. The nonaxisymmetric components of $|B|$ yield a bootstrap current lower than that in an axisymmetric device, which leads to good magnetohydrodynamic stability properties without the need for a conducting wall. The neoclassical confinement improves with increasing β , leading to a new form of configurational invariance in these stellarators.

A. S. Ware, S. P. Hirshman, D. A. Spong, L. A. Berry, A. J. Deisher, G. Y. Fu, J. F. Lyon, and R. Sanchez, "High- β equilibria of drift-optimized compact stellarators," *Phys. Rev. Lett.* **89**, 125003 (2002). [FES](#)

UNDERSTANDING NEO-CLASSICAL TEARING MODES

Tearing modes are magnetic islands formed by the topological rearrange-

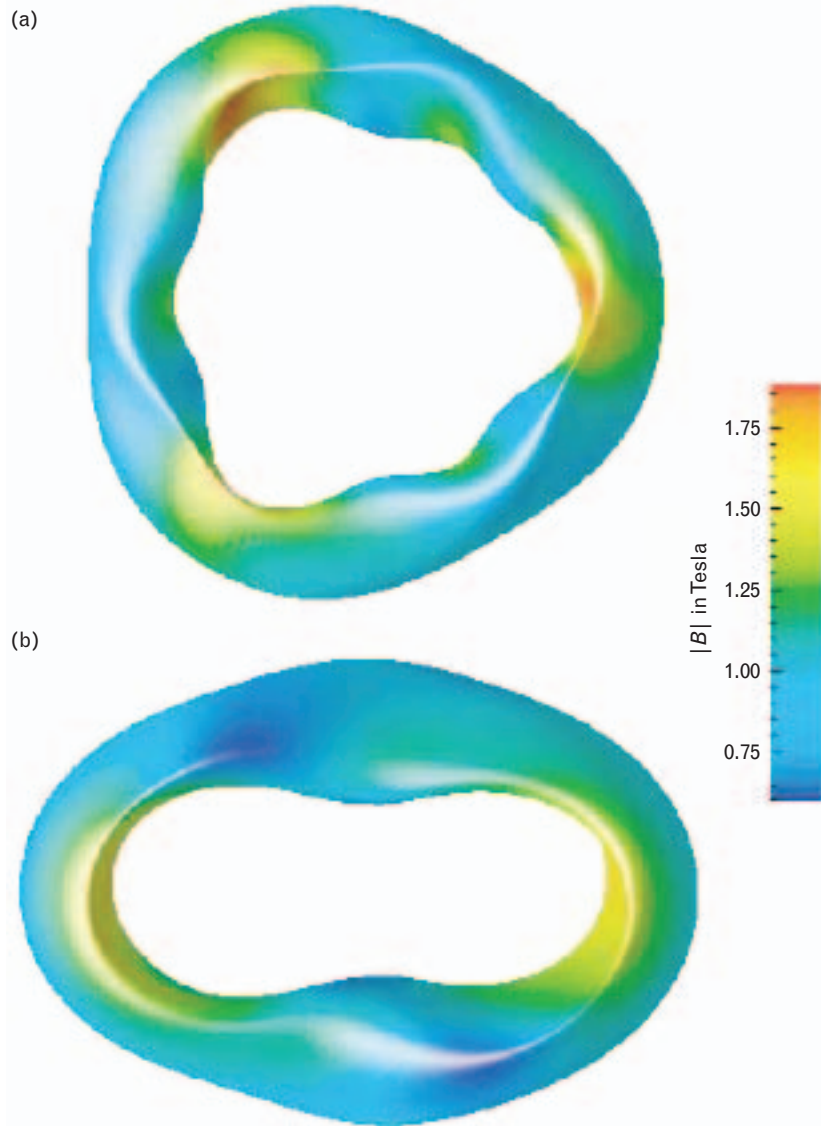


FIGURE 7

Last closed flux surface for two compact stellarator configurations: (a) QPS3, with three field periods, $A = 3.7$, $\beta = 15\%$; and (b) QPS2, with two field periods, $A = 2.7$, $\beta = 5\%$. Color indicates the variation in $|B|$.

SIMULATING PSEUDO-CHAOTIC DYNAMICS

Nonlinear dynamics and chaos simulations are important for understanding the microscale of plasma turbulence, as well as for understanding general nonlinear dynamical effects. A family of random systems with zero Lyapunov exponent is called *pseudochaos*. Zaslavsky and Edelman have shown how the fractional kinetic equation can be introduced for pseudochaos and how the main critical exponents of fractional kinetics can be evaluated from the dynamics. Pseudochaotic dynamics and fractional kinetics can be applied to the behavior of streamlines or magnetic field lines.

G. M. Zaslavsky and M. Edelman, "Pseudochaos," in *Perspectives and Problems in Nonlinear Science*, edited by E. Kaplan, J. E. Marsden, and K. R. Sreenivasan (Springer-Verlag, New York, 2003), pp. 421–443. [FES](#), [ONR](#), [NSF](#)

TURBULENT TRANSPORT WITH KINETIC ELECTRONS

Chen et al. have developed a new electromagnetic kinetic electron-simulation model in 3D toroidal flux-tube geometry that uses a generalized split-weight scheme, where the adiabatic part is adjustable. This model includes electron-ion collisional effects. For H-mode parame-

ment of magnetic field lines through reconnection. The prevention of neoclassical tearing modes (NTMs) in tokamak plasmas is a major challenge because they can degrade plasma confinement. Ideal modes can seed NTMs through forced reconnection, yet in sawtooth discharges it is not well understood why a particular sawtooth crash seeds an NTM after several preceding sawteeth did not. Also, tearing modes sometimes appear and grow without an

obvious ideal mode, causing a seed island. Based on theoretical and experimental results, Brennan et al. have proposed and tested a new mechanism for tearing-mode onset that explains these puzzling observations.

D. P. Brennan, R. J. La Haye, A. D. Turnbull, M. S. Chu, T. H. Jensen, L. L. Lao, T. C. Luce, P. A. Politzer, E. J. Strait, S. E. Kruger, and D. D. Schnack, "A mechanism for tearing onset near ideal stability boundaries," *Phys. Plasmas* **10**, 1643 (2003). [FES](#)

ters, the nonadiabatic effects of kinetic electrons increase linear growth rates of the ion-temperature-gradient-driven (ITG) modes, mainly due to trapped-electron drive. The ion heat transport is also increased from that obtained with adiabatic electrons. The linear behavior of the zonal flow is not significantly affected by kinetic electrons. The ion heat transport decreases to below the adiabatic electron level when finite plasma β is included due to finite- β stabilization of the ITG modes.

Y. Chen, S. E. Parker, B. I. Cohen, A. M. Dimits, W. M. Nevins, D. Shumaker, V. K. Decyk, and J. N. Leboeuf, "Simulations of turbulent transport with kinetic electrons and electromagnetic effects," *Nucl. Fusion* **43**, 1121 (2003). [FES](#), [NASA](#)

LATTICE QCD CONFRONTS EXPERIMENT

For almost thirty years precise numerical studies of nonperturbative QCD, formulated on a space-time lattice, have been stymied by an inability to include the effects of realistic quark vacuum polarization. The MILC, HPQCD, UKQCD, and Fermilab Lattice collaborations have presented detailed evidence of a breakthrough that may now permit a wide variety of high-precision, nonperturbative QCD calculations, including high-precision B and D meson decay constants, mixing amplitudes, and semi-leptonic form factors—all quantities of great importance in current experimental work on heavy-quark physics. The breakthrough comes from a new discretization for light quarks: Symanzik-improved staggered quarks. The researchers compared a wide variety of nonperturbative calculations in QCD with experiment, and found agreement to within statistical and systematic errors of 3% or less.

C. T. H. Davies, E. Follana, A. Gray, G. P. Lepage, Q. Mason, M. Nobes, J. Shigemitsu, H. D. Trotter, M. Wingate, C. Aubin, C. Bernard, T. Burch, C. DeTar, S. Gottlieb, E. B. Gregory, U. M. Heller, J. E. Hetrick, J. Osborn, R. Sugar, D. Toussaint, M. Di

Pierro, A. El-Khadra, A. S. Kronfeld, P. B. Mackenzie, D. Menscher, and J. Simone, "High-precision lattice QCD confronts experiment," *Phys. Rev. Lett.* (submitted, 2003), hep-lat/0304004. [HEP](#), [NSF](#), [PPARC](#)

TWO-COLOR QCD WITH FOUR CONTINUUM FLAVORS

Recently there has been a re-evaluation of the old idea that quark pairs might condense, giving rise to a transition to a color superconducting state at high baryon-number density. Kogut et al. approached this question by examining the spectrum of two-color lattice QCD with four continuum flavors at a finite chemical potential (μ) for quark-number, on a $12^3 \times 24$ lattice. They found evidence that the system undergoes a transition to a state with a diquark condensate, which spontaneously breaks quark number at $\mu = m_\pi/2$, and that this transition is mean field in nature. They then examined the three states that would be Goldstone bosons at $\mu = 0$ for zero Dirac and Majorana quark masses.

J. B. Kogut, D. Toublan, and D. K. Sinclair, "The pseudo-Goldstone spectrum of 2-color QCD at finite density," Argonne National Laboratory report ANL-HEP-PR-03-022 (May 2003), hep-lat/0305003. [HEP](#), [NSF](#), [HFS](#)

BIMETALLIC SURFACE SEGREGATION OF NANOPARTICLES

To better understand catalysts in fuel cells, Wang et al. used a Monte Carlo code built around the embedded atom method to investigate the segregation of platinum atoms on the surfaces of platinum-nickel nanoparticles. They used four kinds of nanoparticle shapes (cube, tetrahedron, octahedron, and cubo-octahedron) terminated by {111} and {100} facets to examine the extent of Pt segregation. Regardless of the shape and composition of the nanoparticles, the Pt atoms segregated preferentially to the facet sites, less to edge sites, and least to vertex sites in the outermost atomic layer of the nanoparticles.

G. Wang, M. A. Van Hove, and P. N. Ross, "Segregations in Pt-Ni catalyst nanoparticles," *Surface Science* (submitted, 2003).
BES, SciDAC

ELECTRONIC INTERACTION OF TOTA⁺ WITH DNA

Investigating the electronic interactions of compounds intercalated in DNA is important because of the biological and medical roles played by intercalators. Reynisson et al. made first-principles quantum mechanical calculations of trioxatriangulenium ion (TOTA⁺) binding to duplex DNA by intercalation. The electronic structure calculations revealed no meaningful charge transfer from DNA to TOTA⁺, but showed the importance

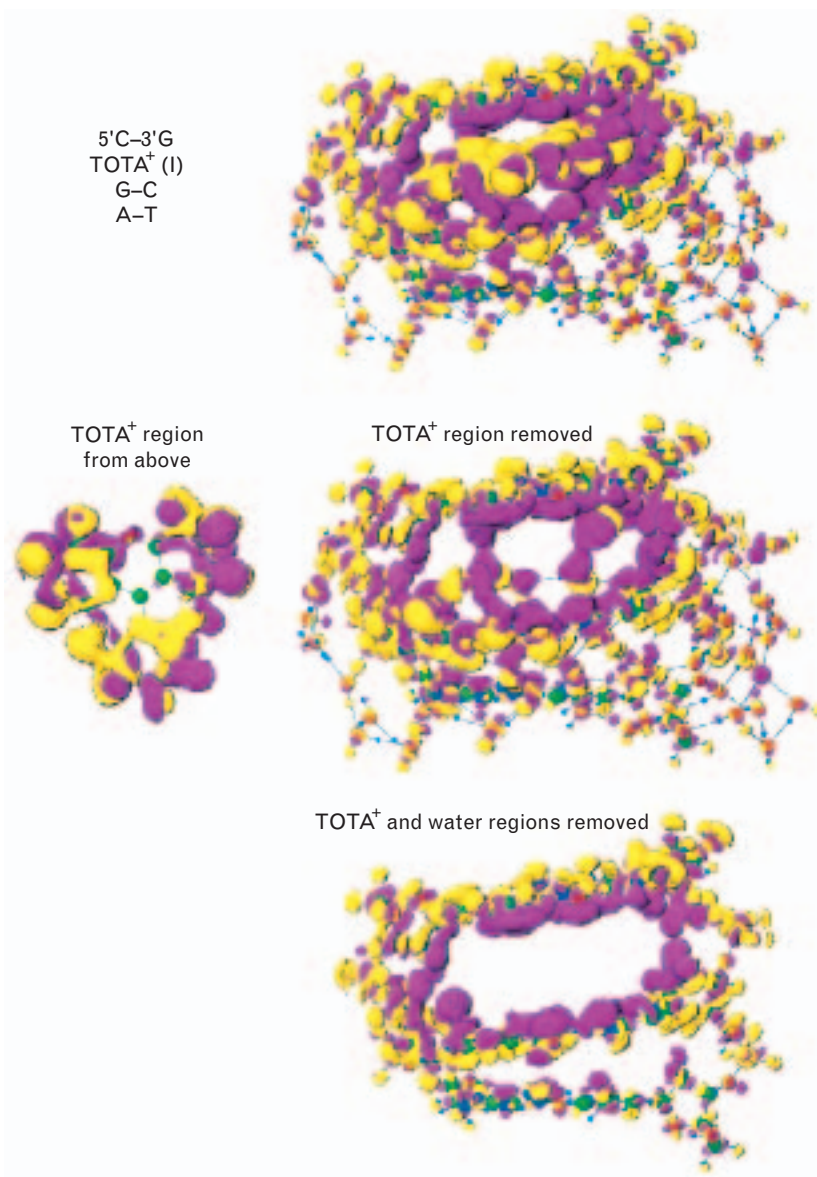


FIGURE 8

Polarization-charge ($\delta\rho$) isosurfaces calculated for the hydrated TOTA⁺(I) 3-bp duplex DNA segments. Magenta regions correspond to those where excess electronic charge is induced by the intercalation process, and yellow regions denote those where depletion of charge ensued. The upper panel shows $\delta\rho$ for the entire hydrated TOTA⁺(I) 3-bp duplex segment; in the middle (right) panel, the TOTA⁺ intercalator region is excluded; and in the bottom panel, both the TOTA⁺ intercalator and hydration shell regions are excluded. The latter illustrates clearly the electronic polarization that is induced by the intercalation of the TOTA cation.

of backbone, water, and counterion interactions, which shift the energy levels of the base and the intercalated TOTA⁺ orbitals significantly.

The calculations also showed that the inserted TOTA⁺ strongly polarizes the intercalation cavity, where a sheet of excess electron density surrounds the TOTA⁺ (Figure 8).

J. Reynisson, G. B. Schuster, S. B. Howerton, L. D. Williams, R. N. Barnett, C. L. Cleveland, U. Landman, N. Harrit, and J. B.

Chaires, "Intercalation of trioxatrianguleneium ion with DNA: Binding, electron transfer, x-ray crystallography, and electronic structure," *J. Am. Chem. Soc.* **125**, 2072 (2002). [BES](#), [NSF](#), [NCI](#)

ELECTRONIC STRUCTURE OF C₆₀ MONOLAYERS

Exhibiting properties such as high transition temperature superconductivity and antiferromagnetism, C₆₀-based fullerides are ideal model compounds for exploring key conceptual issues in strongly correlated physics. Yang et al. studied the structure and electronic structure of electron-doped C₆₀ monolayers. The overall shape of the calculated band structure showed good agreement with the high-resolution angle-resolved photoemission experiments, indicating a reduction of the occupied bandwidth by about 40 percent due to electron-phonon interaction.

W. L. Yang, V. Brouet, X. J. Zhou, H. J. Choi, S. G. Louie, M. L. Cohen, S. A. Kellar, P. V. Bogdanov, A. Lanzara, A. Goldoni, F. Parmigiani, Z. Hussain, and Z.-X. Shen, "Band structure and Fermi surface of electron-doped C₆₀ monolayers," *Science* **300**, 303 (2003). [BES](#), [ONR](#), [NSF](#)

NUCLEATION OF SINGLE-WALLED CARBON NANOTUBES

The nucleation pathway for single-walled carbon nanotubes has been investigated using first-principles density-functional calculations. Fan et al. have shown that incorporation of pentagons at an early stage of the nanotube growth reduces the number of dangling bonds and facilitates

bonding to the metal. As a result, the formation of a capped single-walled nanotube is overwhelmingly favored compared to any structure with dangling bonds (such as graphite sheets) or to a fullerene.

X. Fan, R. Buczko, A. A. Puretzky, D. B. Geohegan, J. Y. Howe, S. T. Pantelides, and S. J. Pennycook, "Nucleation of single-walled nanotubes," *Phys. Rev. Lett.* **90**, 145501 (2003). [BES](#), [ORNL](#), [NSF](#), [MEVU](#)

INVESTIGATING THE CONNECTION BETWEEN ORBITAL AND SPIN ORDERING

Medvedeva et al. studied the complex nature of the fundamental connection between orbital ordering phenomena and spin ordering alignment on a microscopic scale using their recently developed average spin state calculation scheme. This scheme allowed them to treat a paramagnetic state and to describe successfully the experimental magnetic and orbital phase diagrams of LaMnO₃, KCuF₃, La_{0.5}Sr_{1.5}MnO₄, and LaSr₂Mn₂O₇. For the first time, they investigated from first principles the stability of charge and orbital ordering upon onset of a magnetic transition in La_{0.5}Sr_{1.5}MnO₄ and LaSr₂Mn₂O₇.

J. E. Medvedeva, M. A. Korotin, V. I. Anisimov, and A. J. Freeman, "Charge and orbital ordering melting in layered LaSr₂Mn₂O₇," (in preparation, 2003). [BES](#)

CALCULATING THE PROPERTIES OF AN ORGANIC SEMICONDUCTOR

There has been increasing interest in organic semiconductors such as

pentacene for applications in electronic devices. Tiago et al. have studied the optical and electronic properties of crystalline pentacene, using a first-principles Green's function approach. They investigated the role of polymorphism on the electronic energy gap and linear optical spectrum by studying two different crystalline phases: the solution-phase structure and the vapor-phase structure. Charge-transfer excitons were found to dominate the optical spectrum. Excitons with sizable binding energies were predicted for both phases.

M. L. Tiago, J. E. Northrup, and S. G. Louie, "Ab initio calculation of the electronic and optical properties of solid pentacene," *Phys. Rev. B* **67**, 115212 (2003). [BES](#), [NSF](#)

RELATIVITY AND THE STABILITY OF INTERMETALLIC COMPOUNDS

Why are some intermetallic compounds stable while others are not? Wang and Zunger showed that the existence of stable, ordered 3d-5d intermetallics CuAu and NiPt, as opposed to the unstable 3d-4d isovalent analogs CuAg and NiPd, results from relativity. First, in shrinking the equilibrium volume of the 5d element, relativity reduces the atomic size mismatch with respect to the 3d element, thus lowering the elastic packing strain. Second, in lowering the energy of the bonding 6s,p bands and raising the energy of the 5d band, relativity enhances (diminishes) the occupation of the bonding (antibonding) bands. Raising the energy of the 5d

band also brings it closer to the energy of the 3*d* band, improving the 3*d*-5*d* bonding.

L. G. Wang and A. Zunger, "Why are the 3*d*-5*d* compounds CuAu and NiPt stable, whereas the 3*d*-4*d* compounds CuAg and NiPd are not," *Phys. Rev. B* **67**, 092103 (2003). [BES](#)

ELECTRONIC STRUCTURES AND PROPERTIES OF VITAMIN B12

Methodologies developed for the materials sciences are being extended to the study of biomaterials and biomolecules. Kurmaev et al. have calculated the electronic structure of the vitamin B12 molecule and its co-enzymes: cyanocobalamin, methylcobalamin, and adenosyl cobalamin. For the first time, all the side chains in these complex molecules were included in the *ab initio* calculation.

E. Z. Kurmaev, A. Moewes, L. Ouyang, L. Randaccio, P. Rulis, W. Y. Ching, M. Bach, and M. Neumann, "The electronic structure and chemical bonding of vitamin B12," *Europhys. Lett.* **62**, 582 (2003). [BES](#), [NSF](#), [EU](#), [RFBR](#), [NSERC](#), [NEDO](#), [MIUR](#), [CMR](#)

DENSITY FUNCTIONAL CALCULATION OF LEAD ON SILICON

Lead and silicon can form a well-defined interface (Pb/Si) which is ideal for studying two-dimensional behavior. But despite extensive work in the literature, the structure of the different Pb phases formed on Si(111) is still not clear, especially when the Pb coverage is more than

one monolayer (ML). Using first principles calculations, Chan et al. have determined the atomic configuration, energy stability, and electronic structure of the different phases of Pb/Si (111). Their calculations will help experimenters identify the atomic model and domain wall arrangement for the controversial structures of Pb/Si(111) at the Pb coverage of 1 to 4/3 ML.

T.-L. Chan, C. Z. Wang, M. Hupalo, M. C. Tringides, Z.-Y. Lu, and K. M. Ho, "First-principles studies of structures and stabilities of Pb/Si(111)," *Phys. Rev. B* **68**, 045410 (2003). [BES](#)

GROWTH AND FORMATION OF SILICON NANOCLUSTERS

Silicon nanoclusters have been the focus of intense theoretical and experimental interest, because the combination of the unique optical properties of silicon quantum dots and their compatibility with existing silicon-based and nanobiological technologies demonstrates great promise for numerous applications. Draeger et al., using density-functional and quantum Monte Carlo calculations, have demonstrated that the optical gap is significantly reduced by the presence of oxygen and fluorine atoms on the surface. This is an important result for experimentalists, as oxygen is a common contaminant in the procedure for synthesizing nanoparticles.

E. Draeger, J. C. Grossman, A. J. Williamson, and G. Galli, "Influence of synthesis conditions on the structural and optical properties of passivated silicon nanoclusters," *Phys. Rev. Lett.* **90**, 167402 (2003). [BES](#)

MOLECULAR-DYNAMICS SIMULATIONS OF KINETIC COEFFICIENTS

Asta et al. have demonstrated a new equilibrium molecular-dynamics simulation technique for calculating the kinetic coefficient (i.e., the proportionality constant between interface velocity and undercooling) of solid-liquid interfaces in elemental metals. This work represents a necessary first step towards studying dynamic properties of alloy solid-liquid interfaces near equilibrium.

M. Asta, D. Y. Sun, and J. J. Hoyt, "Role of atomistic simulation in the modeling of solidification microstructure," in *NATO-ASI Proceedings: Thermodynamics, Microstructure and Plasticity*, ed. by A. Finel (in press). [BES](#)

SOLVING THE MYSTERY OF EXCHANGE BIAS

Exchange bias, a shift in magnetization that occurs when two different kinds of magnet come into contact, is used in numerous electronic applications such as magnetic multilayer storage and read head devices, but the effect is still not properly understood. Canning et al. have performed first-principles spin-dynamics simulations of the magnetic structure of iron-manganese/cobalt (FeMn/Co) interfaces (Figure 9). These quantum mechanical simulations, involving 2,016-atom supercell models, reveal details of the orientational configuration of the magnetic moments at the interface that are unobtainable by any other means, offering important clues to solve the mystery of exchange bias.


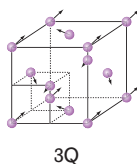


FIGURE 9

Visualization of the exchange-bias system initial state, viewed from the direction of the y-axis. The arrows represent the orientation of the magnetic moments. The gold balls represent the Fe atoms, the grey ones the Mn atoms, while the green and pink balls stand for the relaxed and unrelaxed Co atoms. On the left, the 3Q antiferromagnetic ordering is schematically displayed within the conventional fcc structure.

A. Canning, B. Ujfalussy, T. C. Schulthess, X.-G. Zhang, W. A. Shelton, D. M. C. Nicholson, G. M. Stocks, Y. Wang, and T. Dirks, "Parallel multi-teraflops studies of the magnetic structure of FeMn alloys," *Proc. Int. Parallel and Distributed Processing Symposium (IPDPS'03, April 2003, Nice, France)*, p. 256b. [BES](#), [ASCR-MICS](#)

FAST LARGE-SCALE NANOSTRUCTURE CALCULATIONS

Fast first-principles calculation of thousand-atom systems is a challenge in computational nanoscience. Li and Wang have found a method to solve this challenge, using the charge-patching method for the charge density and an idealized surface passivation. It takes only about one hour on 64 processors to calculate a thousand-atom system with first-principles accuracy. Thousand-atom quantum dots for various materials were calculated for the first time.

J. B. Li and L. W. Wang, "First principle thousand atom quantum dot calculations," *Phys. Rev. Lett.* (submitted, 2003). [BES](#)

UNDERSTANDING ATOMIC ORDERING OF SEMICONDUCTOR ALLOYS

It has long been understood that atomic ordering, widely observed in certain epitaxially grown semiconductor alloys, is driven by surface thermodynamics and/or by growth kinetics, but not by bulk thermodynamics. Batyrev et al. have proposed a microscopic growth model for the ordering of gallium arsenide antimonide (GaAsSb) to account for the observed 3D ordered structure. For

the first time, an interplay between ordering and surface reconstruction during the growth was proposed to explain the experimental observations. The model is qualitatively different from the widely accepted surface reconstruction- and dimerization-induced ordering models that explain only the in-plane 2D patterns.

I. G. Batyrev, A. G. Norman, S. B. Zhang, and S.-H. Wei, "Growth model for atomic ordering: The case for quadruple-period ordering in GaAsSb alloys," *Phys. Rev. Lett.* **90**, 026102 (2003). [BES](#)

MEDICAL AND LIFE SCIENCES

COMPUTATIONAL MODELING OF PROTEIN SWITCHES

The Src-family kinases and their close relatives, such as Abl, function as molecular switches and are important targets for therapeutic intervention; for example, the c-Abl inhibitor Gleevec is used to treat chronic myelogenous leukemia. The development of specific

inhibitors is complicated by the fact that human cells contain ~500 different protein kinases, all of which catalyze the same chemical reaction and contain highly conserved active sites. Because crystal structures provide only static views of certain states of each kinase, Nagar et al. used molecular dynamics simulations to help understand specific switching mechanisms. The simula-

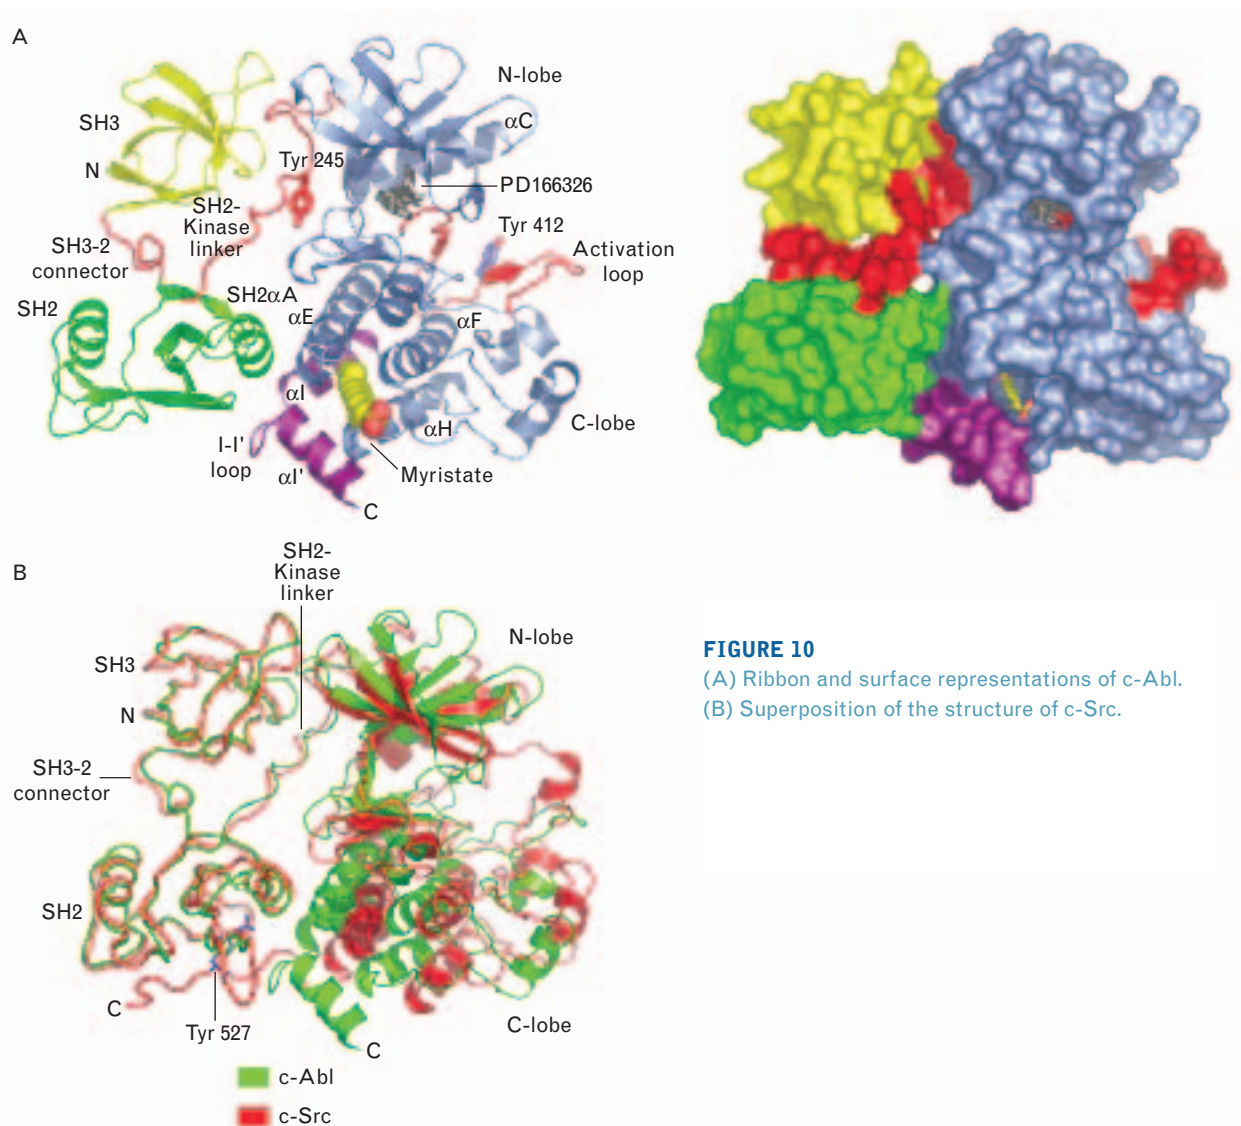


FIGURE 10
 (A) Ribbon and surface representations of c-Abl.
 (B) Superposition of the structure of c-Src.

tions confirmed that the regulatory mechanism of the proto-oncogenic Abl protein shares an important feature with the regulatory mechanism of the related Src-family tyrosine kinases (Figure 10). This result was verified by *in vitro* experiments.

B. Nagar, O. Hantschel, M. A. Young, K. Scheffzek, D. Veach, W. Bornmann, B. Clarkson, G. Superti-Furga, and J. Kuriyan, "Structural basis for the autoinhibition of c-Abl tyrosine kinase," *Cell* **112**, 859 (2003). [BER](#), [HFSP](#), [EMBL](#), [EMBO](#), [AF](#), [GNMF](#), [CAG](#), [DRCRF](#), [BWF](#)

THE IMPACT OF CARCINOGENS ON DNA STRUCTURE

Understanding the 3D structure of DNA that has been modified by chemical carcinogens is important in understanding the molecular basis of cancer and chemical toxicology. Peterson et al. studied [POB]dG, which is produced from metabolically activated tobacco-specific nitrosamines. This compound modifies the base guanine, posing a significant cancer risk. The results

of nuclear magnetic resonance spectra for [POB]dG were used as inputs to the DUPLEX code, which produced a minimum energy structure that agreed with the NMR data.

L. A. Peterson, C. Vu, B. E. Hingerty, S. Broyde, and M. Cosman, "Solution structure of an *O*⁶-[4-oxo-4-(3-Pyridyl)butyl]guanine adduct in an 11mer DNA duplex: Evidence for formation of a base triplex," *Biochemistry* ASAP, DOI: 10.1021/bi035217v (2003). [BER](#), [NIH](#), [ORNL](#)

MODELING ENZYME STRUCTURE AND KINETICS

Although information on the function of enzymes has been obtained experimentally, there is still an essential missing link between the structures and biomolecular dynamics and function. Yang et al. studied the binding change mechanism of F_1 -ATPase, the enzyme responsible for most of the ATP synthesis in living systems. By computing the chemical potential, they identified β_{TP} as the tight binding site for ATP and β_{DP} as the loose site. The half closed site was identified as the binding site for the solution ADP and P_i in ATP synthesis; it is different from the empty binding site for ATP hydrolysis. Based on this result, a consistent structural and kinetic model for the binding change mechanism is being developed for F_1 -ATPase.

W. Yang, Y. Q. Gao, Q. Cui, J. P. Ma, and M. Karplus, "The missing link between thermodynamics and structure in F_1 -ATPase," *PNAS* **100**, 874 (2003). [BER](#), [NIH](#)

INTERACTIONS BETWEEN AMINO ACIDS AND NUCLEOBASES

There is a paucity of theoretical information about interactions between amino acids and nucleobases. Dabkowska et al. made electronic-structure calculations concerning the simplest amino acid-nucleobase complex, i.e., the dimer of glycine and uracil. Glycine is the smallest amino acid, and uracil is a

building pyrimidine nucleobase of RNA. They demonstrated that the most stable complexes between uracil and glycine are formed when the carboxylic group of glycine is bound through two hydrogen bonds to uracil.

I. Dabkowska, J. Rak, and M. Gutowski, "Computational study of hydrogen-bonded complexes between the most stable tautomers of glycine and uracil," *J. Phys. Chem. A* **106**, 7423 (2002). [BER](#), [KBN](#)

NUCLEAR PHYSICS

QUANTUM MONTE CARLO STUDIES OF FERMI GASES

Determining the properties of Fermi gases is an intriguing topic for many-body physics, with applications to phenomena such as neutron stars. Carlson et al. conducted quantum Monte Carlo calculations of superfluid Fermi gases with short-range two-body attractive interactions with infinite scattering length. The energy of such gases is estimated to be (0.44 ± 0.01) times that of the non-interacting gas, and their pairing gap is approximately twice the energy per particle.

J. Carlson, S.-Y. Chang, V. R. Pandharipande, and K. E. Schmidt, "Superfluid Fermi gases with large scattering length," *Phys. Rev. Lett.* (in press), arXiv:physics/0303094 (2003). [NP](#), [NSF](#)

DATA ANALYSIS AND SIMULATIONS FOR KAMLAND

The primary goal of the Kamioka Liquid Scintillator Anti-Neutrino

Detector (KamLAND) is to search for the oscillation of antineutrinos emitted from distant power reactors. KamLAND's first results indicate that it is measuring only 61% of the reactor anti-neutrino flux expected in the no-oscillation hypothesis. This result (using CPT invariance) excludes all solutions to the solar-neutrino problem except for the "large mixing angle" (LMA) region (Figure 11).

K. Eguchi et al. (KamLAND Collaboration), "First results from KamLAND: Evidence for reactor antineutrino disappearance," *Phys. Rev. Lett.* **90**, 021802 (2003). [NP](#), [MEXT](#)

EXCITED BARYON PHYSICS FROM LATTICE QCD

Unraveling the nature of Roper resonance (the first excited state of the nucleon with the same quantum number) has a direct bearing on understanding the quark structure and chiral dynamics of baryons, one of the primary missions at laboratories

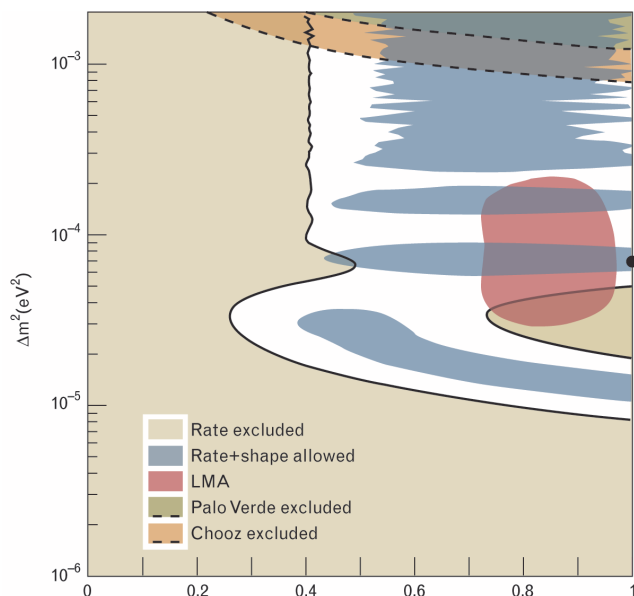


FIGURE 11
The neutrino oscillation parameter region for two-neutrino mixing of solar neutrinos. The solid circle shows the best fit to the KamLAND data.

such as Jefferson Lab. Dong et al. have calculated the correct level-ordering between the Roper (1440) resonance and the S_{11} (1535) resonance for the first time on the lattice. This result verifies that the Roper state is a radial excitation of the nucleon with three valence quarks, and indicates that spontaneously broken chiral symmetry dictates the dynamics of light quarks in the nucleon.

S. J. Dong, T. Draper, I. Horvath, F. X. Lee, K. F. Liu, N. Mathur, and J. B. Zhang, "Roper resonance and S_{11} (1535) in lattice QCD," Phys. Rev. Lett. (submitted, 2003), hep-ph/0306199. [NP](#)

ELECTRONIC STRUCTURE OF SUPERHEAVY ELEMENTS

Malli investigated the effects of relativity on the electronic structure and bonding in the tetroxides of the superheavy element hassium and its lighter congener osmium, with two major conclusions: (1) Relativistic effects lead to a dramatic increase

of ~225% and 185% in the predicted atomization energy for HsO_4 and OsO_4 , respectively. (2) Mulliken population analysis of the Dirac-Fock self-consistent field wave functions, in contrast to the nonrelativistic Hartree-Fock wave function, predicts the HsO_4 to be less volatile than the lighter congener OsO_4 , and this prediction is in accord with the first experimental work on hassium.

G. L. Malli, "Dramatic relativistic effects in atomization energy and volatility of the superheavy hassium tetroxide and OsO_4 ," J. Chem. Phys. **117**, 10441 (2002). [NP](#)

EXPLORING NUCLEON STRUCTURE USING LATTICE QCD

An important question about the nucleon and its excited state, Δ , is whether they are spherical or deformed. Recent experiments have accurately measured the electric and Coulomb quadrupole and magnetic dipole multipoles of the nucleon-to- Δ transition form factor,

which directly reflect the presence of deformation. Alexandrou et al. subsequently calculated this form factor using lattice QCD. The calculated magnetic dipole form factor and electric quadrupole amplitude were consistent with experimental results, but systematic errors due to lattice artifacts prevented a determination of the Coulomb quadrupole form factor. Further study of these lattice artifacts is needed for better control of systematic errors.

C. Alexandrou, P. de Forcrand, T. Lippert, H. Neff, J. W. Negele, K. Schilling, and A. Tsapalis, "N to Δ electromagnetic transition form factor from lattice QCD," hep-lat/0307018 (2003). [NP](#), [ESOP](#), [UC](#)

HFB CONTINUUM PROBLEM SOLVED FOR DEFORMED NUCLEI

One of the fundamental questions of nuclear structure physics is, What are the limits of nuclear stability? Terán et al. have solved, for the first time, the Hartree-Fock-Bogoliubov (HFB) continuum problem in coordinate space for deformed nuclei in two spatial dimensions without any approximations. The novel feature of the new Vanderbilt HFB code is that it takes into account high-energy continuum states with an equivalent single-particle energy of 60 MeV or more. In the past, this has only been possible in 1D calculations for spherical nuclei.

E. Terán, V. E. Oberacker, and A. S. Umar, "Axially symmetric Hartree-Fock-Bogoliubov calculations for nuclei near the drip-lines," Phys. Rev. C **67**, 064314 (2003). [NP](#)

MEASURING THE FATE OF PARTON FRAGMENT JETS

In collisions of heavy nuclei at high energies, a new state of matter consisting of deconfined quarks and gluons at high density is expected. Large transverse momentum (p_T) partons in the high-density system result from the initial hard scattering of nucleon constituents. After a hard scattering, the parton fragments to create a high-energy cluster (jet) of particles. The STAR Collaboration measured two-hadron angular correlations at large transverse momentum for $p + p$ and Au + Au collisions. These measurements provided the most direct evidence for production of jets in high-energy nucleus-nucleus collisions and allowed the first measurements of the fate of back-to-back jets in the dense medi-

um as a function of the size of the overlapping system. The back-to-back correlations were reduced considerably in the most central Au + Au collisions, indicating substantial interaction as the hard-scattered partons or their fragmentation products traversed the medium.

C. Adler et al. (STAR Collaboration), "Disappearance of back-to-back high p_T hadron correlations in central Au + Au collisions at $\sqrt{s_{NN}} = 200$ GeV," Phys. Rev. Lett. **90**, 082302 (2003). [NP](#), [HEP](#), [NSF](#), [BMBF](#), [IN2P3](#), [EPSRC](#), [FAPESP](#), [RMST](#), [MEC](#), [NNSFC](#)

CALCULATING THE ENERGY SPECTRA OF LIGHT NUCLEI

The absence of stable five- or eight-body nuclei is crucial to both primordial and stellar nucleosynthesis, enabling stars such as our sun to

burn steadily for billions of years. Wiringa and Pieper demonstrated that the binding energies, excitation structure, and relative stability of light nuclei, including the opening of the $A = 5$ and 8 mass gaps, are crucially dependent on the complicated structure of the nuclear force. They calculated the energy spectra of light nuclei using a variety of nuclear-force models ranging from very simple to fully realistic, and they observed how features of the experimental spectrum evolve with the sophistication of the force. They found that the absence of stable five- and eight-body nuclei depends crucially on the spin, isospin, and tensor components of the nuclear force.

R. B. Wiringa and S. C. Pieper, "Evolution of nuclear spectra with nuclear forces," Phys. Rev. Lett. **89**, 182501 (2002). [NP](#)

NERSC CENTER

CLIENTS, SPONSORS, AND ADVISORS

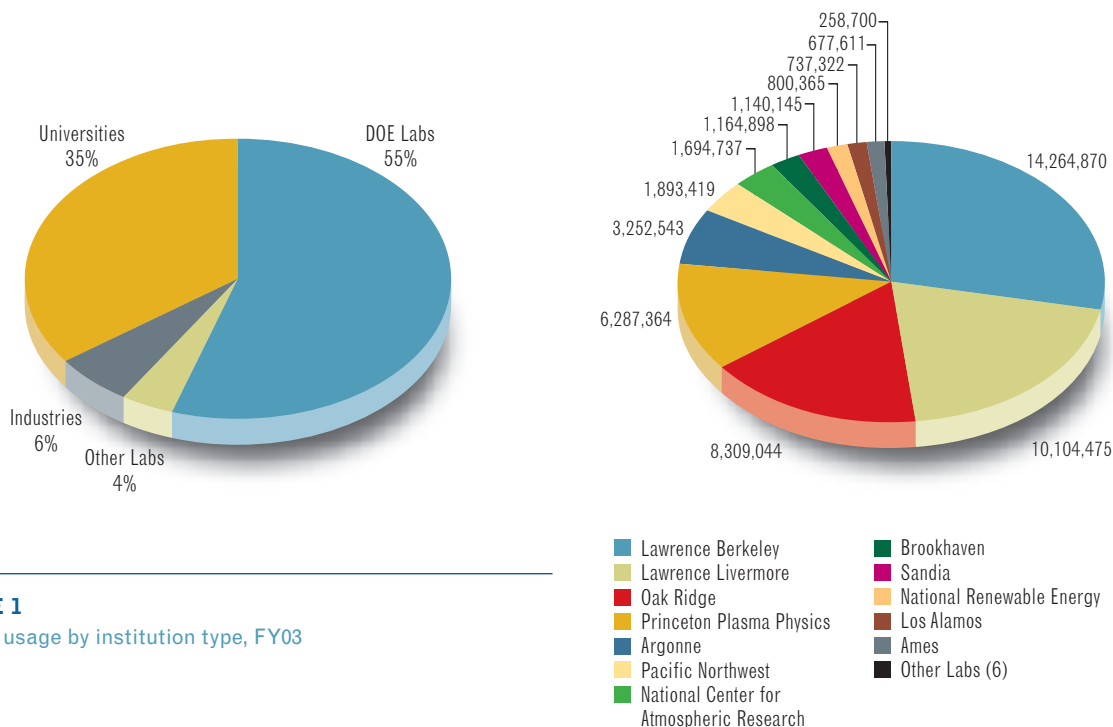


FIGURE 1
NERSC usage by institution type, FY03

FIGURE 2
DOE and other federal laboratory usage at NERSC, FY03

NERSC served 2,323 scientists throughout the United States in FY 2003. 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. Allocations of computer time and archival storage at NERSC are

awarded to research groups, regardless of their source of funding, based on an annual review of proposals. As proposals are submitted, they are subjected to peer review by the Computational Review Panel (CORP, Appendix B) to evaluate the computational approach and the readiness of the applicant’s codes to

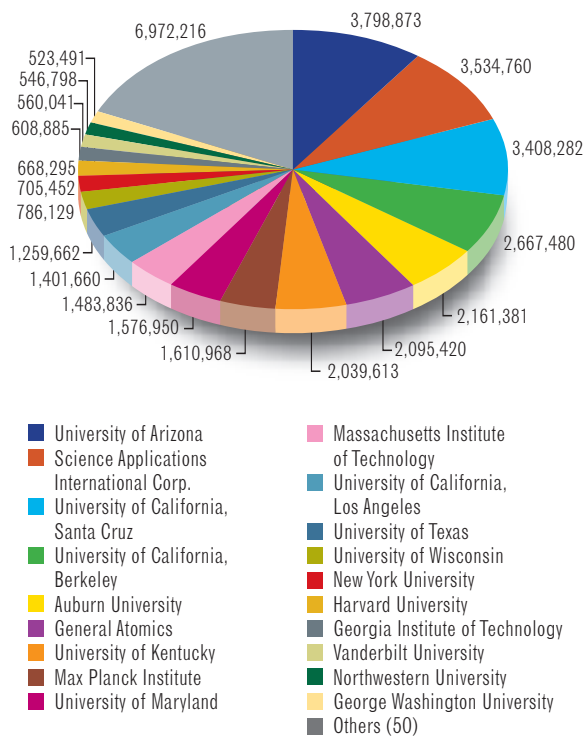


FIGURE 3
Academic and private laboratory usage at NERSC, FY03

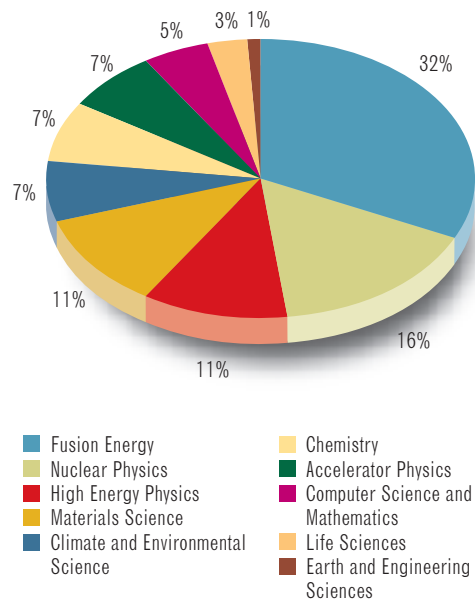


FIGURE 4.
NERSC usage by scientific discipline, FY03

fully utilize the computing resources being requested. DOE Office of Science program managers and the Supercomputing Allocations Committee (SAC, see Appendix D) review and make award decisions for all “base award” requests, reflecting their mission priorities. SciDAC awards are made by the SciDAC Resource Allocation Management Team, while allocations for small startup projects are determined by NERSC.

For FY 2004, a new program entitled Innovative and Novel Computational Impact on Theory and

Experiment (INCITE) awarded 10% of NERSC resources to three computationally intensive research projects that were judged capable of making high-impact scientific advances through the use of a large allocation of computer time and data storage. The three projects are: Thermonuclear Supernovae: Stellar Explosions in Three Dimensions; Fluid Turbulence and Mixing at High Reynolds Number; and Quantum Monte Carlo Study of Photoprotection via Carotenoids in Photosynthetic Centers.

Two other groups provide gener-

al oversight for the NERSC Center: the NERSC Policy Board (Appendix A) advises the Berkeley Lab Director on the policies that determine the impact and performance of the NERSC Center, and the NERSC Users Group (Appendix C) advises the NERSC management and provides feedback from the user community. DOE program management is provided by the Office of Advanced Scientific Computing Research (Appendix E), with advice from the Advanced Scientific Computing Advisory Committee (Appendix F).

CAPABILITY AND SCALABILITY

The past year was marked by a dramatic increase in NERSC's computational capability, resulting in an impressive upsurge in the productivity of many research projects. But in traditional NERSC fashion, these changes took place almost seamlessly, with little disruption to users' daily work.

In 2002, after requesting proposals to replace NERSC's soon-to-be decommissioned Cray systems, the procurement team, lead by Deputy Director Bill Kramer, made the decision to increase the capability of the existing IBM SP Power 3 system, rather than purchase an entirely new one. This expansion gave the user community an unexpected increase of 30 million MPP hours in FY 2003, and doubled the hours that were allocated for FY 2004. "The bottom line was choosing a

system that would have the most cost-effective impact on DOE science," Kramer said when the decision was announced. The new contract includes five years of support for the combined system, named "Seaborg."

The new Seaborg has 416 16-CPU Power 3+ SMP nodes (6,656 processors) with a peak performance of 10 teraflop/s, which made it the most powerful computer for unclassified research in the United States, with the largest memory capacity, at the time of its installation. The system has 7.8 terabytes of aggregate memory and a Global Parallel File System with 44 terabytes of storage.

A team of NERSC staff (Figure 1) worked for six months to configure, install, and test the system. The system was available more than 98 percent of the time during testing, and benchmarks ran at 72 percent



FIGURE 1

The Seaborg expansion procurement and installation teams doubled NERSC's computing resources 20 months earlier than expected and for 75% of the cost expected, without interrupting service to the research community.

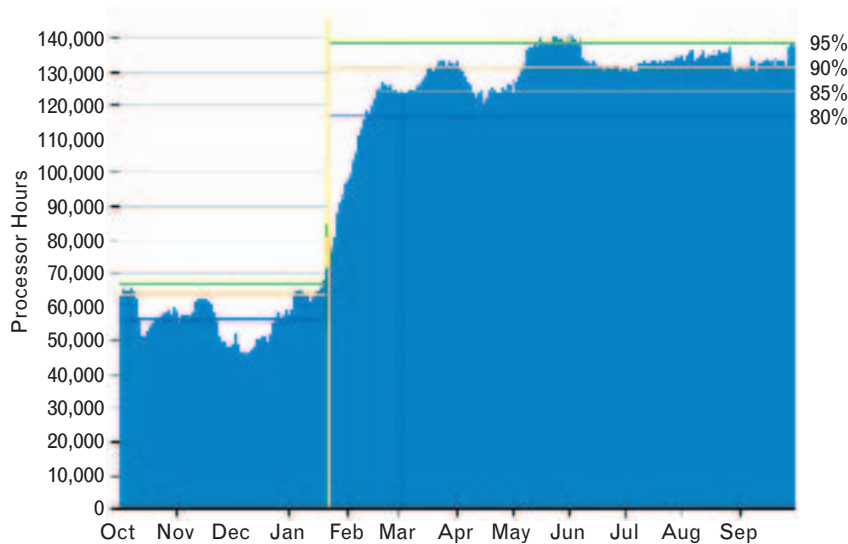


FIGURE 2
NERSC strives to maintain a high rate of processor utilization, balancing this goal with good job turnaround times.

of peak speed, much higher than that achieved on similar parallel systems. The team resolved major configuration issues, such as whether to operate Seaborg as one or two systems (they decided on one), and brought it into production status on March 4, 2003, one month earlier than planned. The expanded system quickly achieved the high availability and utilization rates that NERSC demands of itself (Figure 2).

Initial results from the expanded Seaborg showed some scientific applications running at up to 68% of the system’s theoretical peak speed, compared with the 5–10% of peak performance typical for scientific applications running on massively parallel or cluster architectures. Performance results for four science-of-scale applications are summarized in Table 1. (More details are available at <http://www.nersc.gov/news/newNERSCresults050703.pdf>.)

To learn about how Seaborg performs as a production environment for a larger sample of scientific applications, NERSC conducted a

| Project | Number of processors | Performance (% of peak) |
|---|----------------------|-------------------------|
| Electromagnetic Wave-Plasma Interactions | 1,936 | 68% |
| Cosmic Microwave Background Data Analysis | 4,096 | 50% |
| Terascale Simulations of Supernovae | 2,048 | 43% |
| Quantum Chromodynamics at High Temperatures | 1,024 | 13% |

TABLE 1
Performance Results for Science-of-Scale Applications

survey of user programs that run on a large number of nodes, as well as programs that run on a smaller number of nodes but account for a significant fraction of time used on the system. The report “IBM SP Parallel Scaling Overview” (<http://www.nersc.gov/computers/SP/scaling>) discusses issues such as constraints to scaling, SMP scaling, MPI scaling, and parallel I/O scaling. The report offers recommendations for choosing the level of parallelism best suited to the characteristics of the code.

Improving the scalability of users’ codes helps NERSC achieve goals set by the DOE Office of Science: that 25% of computing

time in FY 2003 and 50% in FY 2004 be used by jobs that require 512 or more processors. Job sizes for FY 2003 are shown in Figure 3.

The scaling report notes that as Seaborg doubled in size, the number of jobs increased correspondingly, and over time the number of jobs then decreased as the parallelism of jobs increased (Figure 4). The waiting time for large-concurrency jobs also changed in accordance with changes in queuing policies (Figure 5).

Another measure of Seaborg’s performance is Sustained System Performance (SSP). The SSP metric was developed at NERSC as a composite performance measurement of

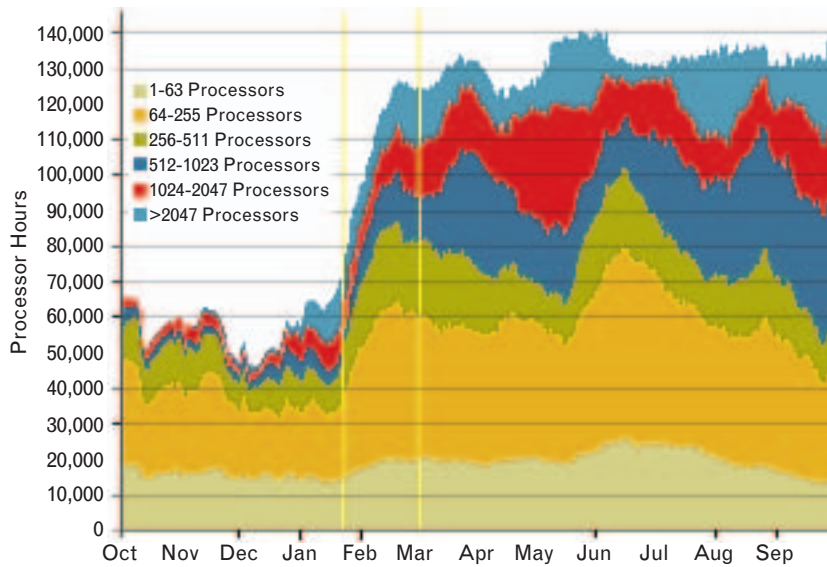


FIGURE 3
Job sizes on Seaborg over the course of FY 2003. DOE's goal for NERSC is to increase the percentage of jobs that use 512 or more processors.

codes from five scientific applications in the NERSC workload: fusion energy, material sciences, cosmology, climate modeling, and quantum chromodynamics. Thus, SSP encompasses a range of algorithms and computational techniques and manages to quantify sys-

tem performance in a way that is relevant to scientific computing.

To obtain the SSP metric, a floating-point operation count is obtained from microprocessor hardware counters for the five applications. A computational rate per processor is calculated by dividing

the floating-point operation count by the sum of observed elapsed times of the applications. Finally, the SSP is calculated by multiplying the computational rate by the number of processors in the system.

The SSP is used as the primary performance figure for NERSC's contract with IBM. It is periodically reevaluated to track performance and monitor changes in the system. For example, the performance impact of a compiler upgrade is readily detected with a change in SSP. The expansion of Seaborg more than doubled the SSP, from 657 to 1357.

During the acceptance period for the Seaborg expansion, there were reports of new nodes performing slightly better than old nodes, which was puzzling since the hardware is identical. All of these reports involved parallel MPI codes; no performance differences were detected on serial codes. The performance differences appeared to

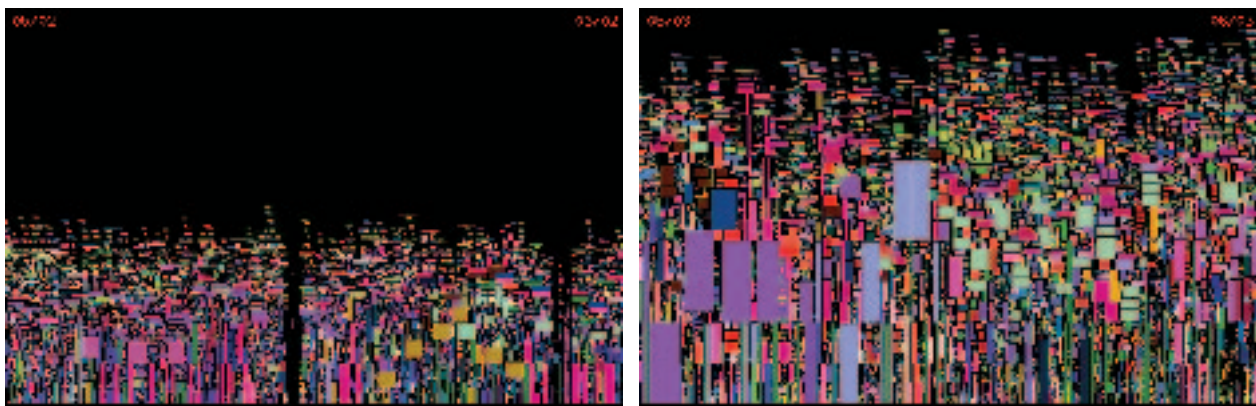


FIGURE 4
Number and size of jobs run on Seaborg in May 2002 (left) and May 2003 (right). The horizontal axis is time. The colored rectangles are jobs, with start and stop time depicted by width, and level of parallelism (number of nodes) depicted by height. Colors signify different users. The rectangles are arranged vertically to avoid overlapping.

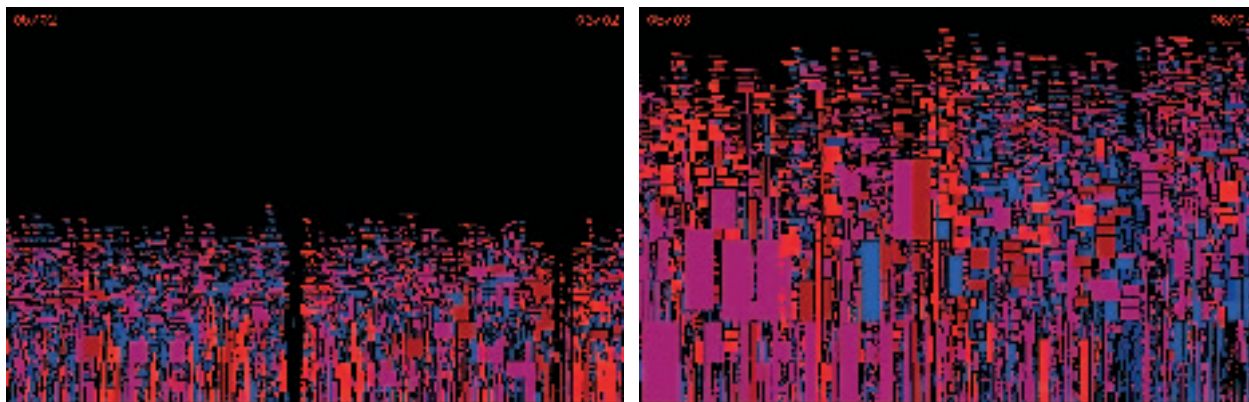


FIGURE 5
Queue waiting time for jobs run on Seaborg in May 2002 (left) and May 2003 (right). Jobs are depicted as in Figure 3 except for colors. Here the color range signifies waiting time, with blue depicting a shorter wait and red a longer wait.

depend on the concurrency and the amount of synchronization in the MPI calls used in the code, but periodic observations and testing provided inconsistent results, and at times no difference could be measured. So in October 2003, groups of old and new nodes were set aside for in-depth testing.

A critical insight occurred when the control workstation happened to be inoperable during the node testing. During that time, the performance of the old nodes improved to match the new nodes. NERSC staff discovered that the control workstation ran certain commands from its problem management subsystem up to 27 times more often on old nodes than on new nodes. Once the problem was pinpointed, it was quickly resolved by deleting four deactivated problem definitions. Resolving this issue improved the synchronization of MPI codes at high concurrency, resulting in faster and more consistent run times for most NERSC users.

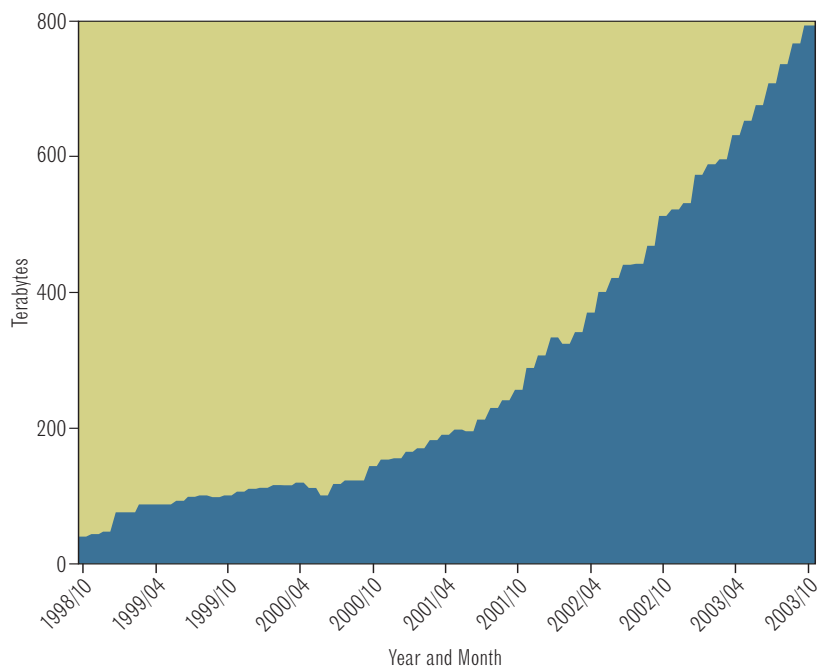


FIGURE 6
Cumulative storage by month, 1998–2003.

In addition to the Seaborg expansion, NERSC also upgraded its PDSF Linux cluster and High Performance Storage System (HPSS).

The PDSF, which supports large data capacity serial applications for the high energy and

nuclear physics communities, was expanded to 412 nodes with a disk capacity of 135 terabytes (TB). New technologies incorporated this year include Opteron chips and 10 TB of configurable network attached storage (NAS).

At NERSC, the data in storage doubles almost every year (Figure 6). As of 2003, NERSC stores approximately 850 TB of data (30 million files) and handles between 3 and 6 TB of I/O per day (Figure 7). To keep up with this constantly increasing workload, new technology is implemented as it becomes available. NERSC currently has two HPSS systems, one used primarily for user files and another used primarily for system backups. The maximum theoretical capacity of NERSC's archive is 8.8 petabytes, the buffer (disk) cache is 35 TB, and the theoretical transfer rate is 2.8 gigabytes per second. Implementation of a four-way stripe this year increased client data transfers to 800 megabytes per second.

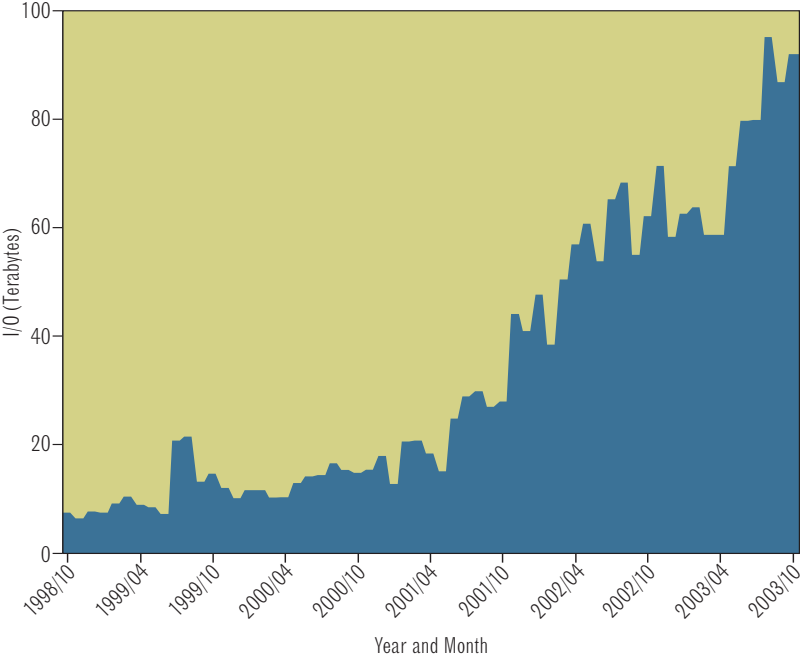


FIGURE 7
I/O by month, 1998–2003.

Many of the important breakthroughs in computational science are expected to come from large, multidisciplinary, multi-institutional collaborations working with advanced codes and large datasets, such as the SciDAC and INCITE collaborations. These teams are in the best position to take advantage of terascale computers and petascale storage, and NERSC provides its highest level of support to these researchers. This support includes specialized consulting support; special service coordination for queues, throughput, increased limits, etc.; specialized algorithmic support; special software support; visualization support; conference and workshop support; Web sever and Web content support



for some projects; and CVS servers and support for community code management.

Providing this level of support sometimes requires dogged determination. That was the case when the Center for Extended MHD Modeling, a SciDAC project, reported that one of their plasma simulation codes was hanging up intermittently. NERSC consultant Mike Stewart (Figure 1) tested the code on Seaborg in both heavily and lightly loaded scenarios, at various compiler optimization levels, on different combinations of nodes, and on NERSC's SP test system. He submitted the bug to IBM as a "severity 2" problem; but after a week passed with no response, he asked the local IBM staff to investigate. They put Mike in touch with an IBM AIX developer, at whose request he ran more test cases. Mike even wrote a FORTRAN code to mimic the C test case, and it displayed the same problem.

Using TotalView debugging software, Mike finally pinpointed the location of the bug in an MPI file read routine. With this new information, IBM assigned the problem to an MPI developer, who worked with Mike for several hours by phone to fix the bug. IBM then

FIGURE 1
NERSC consultant Mike Stewart's persistence paid off in the timely fix of an obscure software bug.



FIGURE 2
Richard Gerber has made significant contributions to several research projects by improving code performance and helping scientists use NERSC resources effectively.

released a revised MPI runtime library that allowed the plasma code to run without interruption. Without Mike's persistence, the researchers' work might have been delayed for months.

NERSC staff are not content to wait for users to report problems. Consultant Richard Gerber (Figure 2) has taken extra initiative to strengthen support for SciDAC and other strategic projects by staying in touch with principal investigators

and by attending project meetings, where he discusses the state of their research efforts and ways in which NERSC can contribute to their success. From these discussions, Richard has undertaken a number of successful projects to improve code performance and to resolve problems, such as finding a way to install a logistical networking node that meets NERSC security requirements.

One particularly difficult problem involved the SciDAC Accelerator Science and Technology project's 3D beam modeling code, which was scaling poorly on the expanded Seaborg. Richard's initial analysis showed that the code's runtime performance depended strongly on the number of tasks per node, not on the number of nodes used. He then pinpointed the problem as the default threading strategy used by IBM to implement the FORTRAN random number intrinsic function, and he enlisted the help of other NERSC consultants to find a solution. Jonathan Carter and Mike

Stewart found two obscure references to this function in IBM documentation that pointed the way to a solution. Richard used this information to conduct some tests, and found that a runtime setting that controls the number of threads used in the intrinsic function can increase code performance by a factor of 2 when running 16 tasks on one node. Using this setting, the beam modeling code is now performing better than ever.

With the variety of codes that run on NERSC systems, NERSC tries to provide users with flexibility and a variety of options, avoiding a "one size fits all" mentality. For years NERSC and other high performance computing (HPC) sites have been urging IBM to support multiple versions of compilers, since not all codes get the best performance from the same compiler version. When IBM provided scripts to install versions newer than the default compiler—but not older versions—NERSC staff saw an opportunity they could take advan-

tage of (Figure 3). David Skinner analyzed these complicated scripts and with great insight was able to engineer this process for older compilers as well. Majdi Baddourah built extensive test cases and showed ingenuity in resolving the problems these test cases exposed. David Paul integrated the multiple compilers into the system administration structure. Together they made the use of multiple compilers fail-safe for users, and they gave helpful feedback to IBM on improvements in compiler support.

To maintain quality services and a high level of user satisfaction, NERSC conducts an annual user survey. In 2003, 326 users responded, the highest response level in the survey's six year history. Overall satisfaction with NERSC was 6.37 on a 7-point satisfaction scale. Areas with the highest user satisfaction (6.5–6.6) were HPSS reliability, timely consulting response, technical advice, HPSS uptime, and local area network. Areas with the lowest user satisfaction (4.7–5.0) were Access Grid classes, visualization software and services, and training.

In answer to the question "What does NERSC do well?" respondents pointed to NERSC's good hardware management practices, user support



FIGURE 3

David Skinner, David Paul, and Majdi Baddourah (not pictured) helped NERSC users get the best performance out of their codes by implementing support for multiple compilers.

and responsive staff, documentation, and job scheduling and throughput. The question “What should NERSC do differently?” prompted concerns about favoring large jobs at the expense of smaller ones, as well as requests for more resources devoted to interactive computing and debugging, more compute power overall, different architectures, mid-range computing support, vector architectures, better documentation, and more training. Complete survey results are available at <http://hpcf.nersc.gov/about/survey/2003/first.php>.

During the past year, NERSC instituted several changes based on the results of the 2002 survey:

- Run-time limits were extended so that jobs could run longer, especially jobs running on more than 32 nodes.
- Interactive jobs were given a higher priority than debugging jobs.
- Two new queue classes were added to facilitate pre- and post-processing and data transfers to HPSS.
- Software was enhanced to make the SP environment easier to use.
- The NERSC user web site (<http://hpcf.nersc.gov>) was restructured with new navigation links that make finding information faster and easier.
- The presentation of SP documentation was streamlined.
- More training was provided on performance analysis, optimization, and debugging.
- More information on initial account setup was added to the New User Guide, which was also reformatted for ease of use.

CONNECTIVITY

As a national facility whose users routinely transfer massive quantities of data, NERSC needs fast networks as much as fast computers. DOE's Energy Sciences Network (ESnet) provides high performance connections between DOE sites as well as links to other networks on the Internet. This year NERSC upgraded its border router connection to ESnet to OC-48 (2.4 Gb/s). Two one-gigabit Ethernets channeled together link the router with NERSC's internal network. The router itself was upgraded to Jumbo Frames capability, enabling NERSC to send 9,000-byte data packets across the Internet instead of the previous 1,500-byte packets.

Integration of NERSC's computing and storage systems into the DOE Science Grid is almost complete. All production systems have

GridFTP data transfer functionality, including Seaborg, HPSS, PDSF, and the math and visualization servers. Globus libraries have been installed on Seaborg's interactive nodes and will be running on all nodes in early 2004, when the Globus Gatekeeper interface will also be installed, providing basic data and job submission services. A Grid interface to the NERSC Information Management (NIM) system will make it easier for users to get Grid authentication certificates. The active intrusion-detection program Bro is being modified to monitor Grid traffic. And a Grid-enabled Web portal interface called VisPortal is being developed to deliver visualization services to Grid collaborators (see page 85 below).

In the past year, viruses and worms with names like Nimda,

SoBig, Klez, Slammer, Blaster, and Bugbear were wreaking havoc across the Internet. In the case of SoBig, over one million computers were infected within the first 24 hours and over 200 million computers were infected within a week. Nevertheless, NERSC had no major security incidents, thanks to its alert security staff and intrusion detection by Bro, which is updated regularly to defend against emerging threats. In addition, NERSC's server team upgraded the email server and added a spam filter.

When DOE officials realized that they were unable to categorize

FIGURE 1

Eli Dart, Scott Campbell, Brent Draney, and Howard Walter developed a methodology to categorize network traffic that has been adopted across the DOE complex.





FIGURE 2

This team won top honors for the highest performing application in the SC2002 Bandwidth Challenge competition, moving data at a peak speed of 16.8 GB/s. The team included staff from Berkeley Lab's Computational Research (CRD), Information Technologies and Services (ITSD), and NERSC Center divisions. Shown here are: (*back row*) Al Early (ITSD), Mike Bennett (ITSD), Brian Tierney (CRD), John Christman (ITSD); (*middle row*) Jason Lee (CRD), Wes Bethel (CRD), Cary Whitney (NERSC); (*front row*) John Shalf (CRD) and Shane Canon (NERSC). Chip Smith and Eli Dart of NERSC are not shown.

and measure different types of network traffic on ESnet, NERSC's networking and security group, led by Howard Walter, stepped in to meet the challenge. In less than a month, using network monitoring data from Bro, routers, and databases, Brent Draney, Eli Dart, and Scott Campbell produced data that clearly categorized more than 95% of all network traffic (Figure 1). The categories they developed—bulk data, database, Grid, interactive, mail, system services, and Web—are now in use at all DOE laboratories and facilities to monitor and measure network traffic. The most important outcome of this effort was the demonstration that ESnet is very different from commercial Internet service providers, and that DOE network

traffic is clearly driven by science.

Although the Berkeley Lab/NERSC Bandwidth Challenge team (Figure 2) retired from the competition after winning three years in a row at the annual SC conference on high performance computing and networking, NERSC's networking and security staff continued to play a major role at SC2003 in Phoenix, helping to create an infrastructure providing the 7,500 attendees with high-speed wired and wireless connections. Bro was used to monitor the conference network for security breaches, and passwords sent unsecurely over the network were displayed on a large plasma screen in the exhibits hall. The display continually attracted viewers, looking to see if their passwords had

been grabbed. Another screen showed the number of complete and incomplete network connections emanating from the conference. The strong interest in Bro was demonstrated by a standing-room-only turnout for a presentation on Bro in the Berkeley Lab booth shortly before the end of the conference.

In the field of high-performance computing, there is a saying that if you are standing still, you are really falling behind. On the other hand, just grabbing the newest technology in the hope of keeping ahead is not necessarily a prudent course. Knowing what is looming just over the horizon—and helping shape it before it comes completely into view—has long been a hallmark of the NERSC staff.

By working in close partnership with leading developers of high-performance computing systems and software, NERSC staff share their collective expertise and experience to help ensure that new products can meet the demands of computational scientists. Vendors benefit from this relationship in that NERSC serves as a real-world testbed for innovative ideas and technologies. This section presents current examples of how NERSC is working in partnership to advance the state of scientific computing.

CREATING SCIENCE-DRIVEN COMPUTER ARCHITECTURE

In 2002 Lawrence Berkeley and Argonne national laboratories, in close collaboration with IBM, proposed a strategy to create a new class of computational capability that is optimal for science. The white paper “Creating Science-Driven Computer Architecture: A New Path to Scientific Leadership” (<http://www.nerosc.gov/news/blueplanet.html>) envisioned development partnerships

in which teams of scientific applications specialists and computer scientists would work with computer architects from major U.S. vendors to create hardware and software environments that will allow scientists to extract the maximum performance and capability from the hardware. The paper included a conceptual design for a system called Blue Planet that would enhance IBM’s Power series architecture for scientific computing.

The first implementation of this strategy in 2003 involved a series of meetings between NERSC, IBM, and Lawrence Livermore National Laboratory (LLNL). IBM had won a contract to build the ASCI Purple system at LLNL, the world’s first supercomputer capable of up to 100 teraflop/s, powered by 12,544 Power 5 processors. LLNL and IBM scientists thought that they might enhance the performance of ASCI Purple by incorporating key elements of the Blue Planet system architecture, particularly the node configuration, and so NERSC staff joined in discussions to work out the details. The resulting architecture enhancements are expected to benefit all scientific users of future IBM Power series systems.

MAKING ARCHIVES MORE ACCESSIBLE

As a long-time collaborator in HPSS development, NERSC made two major contributions this year. The first was the Grid authentica-

tion code for the Parallel FTP (PFTP) file transfer protocol; this enhancement to PFTP is now being distributed by IBM. The second contribution, in collaboration with Argonne National Laboratory, is a testbed for the Globus/Grid interface with HPSS.

NERSC is also collaborating with the DOE’s Joint Genome Institute (JGI) on a project to optimize genomic data storage for wide accessibility. This project will distribute and enhance access to data generated at JGI’s Production Genome Facility (PGF), and will archive the data on NERSC’s HPSS. The PGF is one of the world’s largest public DNA sequencing facilities, producing 2 million trace data files each month (25 KB each), 100 assembled projects per month (50–250 MB), and several very large assembled projects per year (~50 GB). NERSC storage will make this data more quickly and easily available to genome researchers. Storage will be configured to hold more data online and intelligently move the data to enable high-speed access.

NERSC and JGI staff will work closely to improve the organization and clustering of DNA trace data so that portions of an assembled dataset can be downloaded instead of the entire project, which could contain millions of trace files. The Web interface being developed will for the first time correlate the DNA trace data with portions of the assembled sequence, simplifying the selection of trace data, reducing

the amount of data that must be handled, and enhancing researchers' ability to generate partial comparisons. Cached Web servers and file systems will extend the archive to the wide area network, distributing data across multiple platforms based on data characteristics, then unifying data access through one interface. When this project is completed, its storage and data movement techniques, together with its data interfaces, could serve as a model for other read-only data repositories.

SHAPING THE FUTURE OF FILE STORAGE

In 2001, NERSC began a project to develop a system to streamline its file storage system. In a typical HPC environment, each large computational system has its own local disk along with access to additional network-attached storage and archival storage servers. Such an environment prevents the consolidation of storage between systems, thus limiting the amount of working storage available on each system to its local disk capacity. The result is an unnecessary replication of files on multiple systems, an increased workload on users to manage their files, and a burden on the infrastructure to support file transfers between the various systems.

NERSC is developing a solution using existing and emerging technologies to overcome these inefficiencies. The Global Unified Parallel File System (GUPFS) project aims to provide a scalable, high-performance, high-bandwidth, shared-disk file system for use by all

of NERSC's high-performance production computational systems. GUPFS will provide unified file namespace for these systems and will be integrated with HPSS. Storage servers, accessing the consolidated storage through the GUPFS shared-disk file systems, will provide hierarchical storage management, backup, and archival services. An additional goal is to distribute GUPFS-based file systems to geographically remote facilities as native file systems over the DOE Science Grid.

When fully implemented, GUPFS will eliminate unnecessary data replication, simplify the user environment, provide better distribution of storage resources, and permit the management of storage as a separate entity while minimizing impacts on the computational systems.

The major enabling components of this envisioned environment are a high-performance shared-disk file system and a cost-effective, high performance storage-area network (SAN). These emerging technologies, while evolving rapidly, are not targeted towards the needs of high-performance scientific computing. The GUPFS project intends to encourage the development of these technologies to support HPC needs through collaborations with other institutions and vendors, and through active development.

During the first three years of the project, the GUPFS team (Figure 1) tested and evaluated shared-disk file systems, SAN technologies, and other components of the GUPFS environment. This investigation included open source




FIGURE 1 Rei Lee, Cary Whitney, Will Baird, and Greg Butler, along with (not shown) Mike Welcome (Computational Research Division) and Craig Tull, are working to develop a high performance file storage system that is easy to use and independent of computing platforms.

and commercial shared-disk file systems, new SAN fabric technologies as they became available, SAN and file system distribution over the wide-area network, HPSS integration, and file system performance and scaling.

In 2003, following extended trials with multiple storage vendors, NERSC selected the YottaYotta NetStorager system to address the scalability and performance demands of the GUPFS project. A major goal is to provide thousands of heterogeneous hosts with predictable and scalable performance to a common pool of storage resources with no increase in operational costs or complexity.

ENABLING GRID-BASED VISUALIZATION

The Berkeley Lab/NERSC visualization group is developing a Web portal interface called VisPortal that will deliver Grid-based visualization and data analysis capabilities from an easy-to-use, centralized point of access. Using standard Globus/Grid middleware and off-the-shelf Web automation, VisPortal hides the underlying complexity of resource selection and application management among heterogeneous distributed computing and storage resources. From a single access point, the user can browse through the data, wherever it is located, access computing resources, and launch all the components of a distributed visualization application without having to know all the details of how the various systems are configured. The portal

can also automate complex workflows like the distributed generation of MPEG movies or the scheduling of file transfers. The VisPortal developers are currently working with NERSC users to improve the workflow management and robustness of the prototype system.

A TESTBED FOR CLUSTER SOFTWARE

Three years ago, Berkeley Lab purchased a 174-processor Linux cluster, named Alvarez in honor of Berkeley Lab Nobel Laureate Luis Alvarez, to provide a system testbed for NERSC and other staff to evaluate cluster architectures as an option for procurement of large production computing systems. Recently the system has also proven useful for vendors looking to test and develop software for large-scale systems, such as Unlimited Linux and operating system components for Red Storm.

Unlimited Scale, Inc. (USI) used Alvarez to test a beta version of Unlimited Linux, an operating system developed to provide production capabilities on commodity processors and interconnects that match or exceed those of the Cray T3E supercomputer, once the workhorse of the NERSC Center. USI asked NERSC to assess Unlimited Linux's operational and administrative functionality, performance, desired features and enhancements, and user functionality. A few minor problems were discovered in adapting Unlimited Linux to the Alvarez network address scheme, and USI

will change its configuration tools to make it easier to install the system on existing hardware configurations. Performance benchmarks showed that Unlimited Linux compares favorably to the IBM system software that runs on Alvarez.

Cray Inc. was given access to Alvarez to test the operating system under development for the Red Storm supercomputer being developed by Cray Inc. and DOE's Sandia National Laboratories for the National Nuclear Security Administration's Advanced Simulation and Computing (ASCI) program. The first round involved testing system software and administration, and system scalability. The Red Storm Linux-based software allows simulation of multiple virtual processors per physical processor, and simulations of up to 1,000 processors were run on Alvarez. A set of scripts for system configuration and easier job launch were developed, and experiments provided some initial data on launch time for small and large executable size programs. The team found and fixed approximately 20 system bugs, from the portals layer up through MPI and system launch.

In the second stage of testing, Cray's I/O team is working with NERSC to use the GUPFS testbed platform, and GUPFS joined together with Alvarez, to do scalability testing of two potential Red Storm file systems, PVFS and Lustre. They will also collaborate with NERSC staff who are working on global file systems that will eventually be deployed within the NERSC Center.

NERSC POLICY BOARD

ROBERT J. GOLDSTON

Princeton Plasma Physics
Laboratory

STEPHEN C. JARDIN

(ex officio, NERSC
Computational Review Panel)
Princeton Plasma Physics
Laboratory

SIDNEY KARIN

University of California, San
Diego

WILLIAM J. MADIA

Battelle Memorial Institute

PAUL C. MESSINA

Argonne National Laboratory

ALBERT NARATH

Lockheed-Martin Corporation
(retired)

DANIEL A. REED

University of North Carolina

ROBERT D. RYNE

(ex officio, NERSC Users
Group Chair)
Lawrence Berkeley National
Laboratory

TETSUYA SATO

Earth Simulator Center/Japan
Marine Science and Technology
Center

STEPHEN L. SQUIRES

Hewlett-Packard

MICHAEL S. WITHERELL

Fermi National Accelerator
Laboratory

NERSC COMPUTATIONAL REVIEW PANEL (CORP)

COMPUTATIONAL SCIENCE

HIGH ENERGY PHYSICS

Robert Sugar
University of California, Santa
Barbara

ACCELERATOR PHYSICS

Robert Ryne
Lawrence Berkeley National
Laboratory

NUCLEAR AND ASTROPHYSICS

Douglas Swesty
State University of New York at
Stony Brook

CHEMISTRY

Robert Harrison
Oak Ridge National Laboratory

MATERIAL SCIENCES

David Turner
Ames Laboratory

Michael Weinert

Brookhaven National Laboratory

EARTH AND ENGINEERING SCIENCES

David Alumbaugh
University of Wisconsin

FUSION ENERGY

Stephen Jardin
Princeton Plasma Physics
Laboratory

Donald Spong
Oak Ridge National Laboratory

CLIMATE RESEARCH AND ENVIRONMENTAL SCIENCES

Ferdinand Baer
University of Maryland

Douglas Rotman
Lawrence Livermore National
Laboratory

LIFE SCIENCES

Brian Hingerty
Oak Ridge National Laboratory

COMPUTER SCIENCE AND APPLIED MATHEMATICS

David Bailey
Lawrence Berkeley National
Laboratory

APPLIED MATHEMATICS

Jesse Barlow
Pennsylvania State University

Christopher Beattie
Virginia Polytechnic Institute and
State University

Ralph Byers
University of Kansas

Marc Day
Lawrence Berkeley National
Laboratory

Eric de Sturler
University of Illinois at Urbana-
Champaign

Tony Drummond
Lawrence Berkeley National
Laboratory

Joseph Grear
Lawrence Berkeley National
Laboratory

Van Henson
Lawrence Livermore National
Laboratory

Ilse Ipsen
North Carolina State University

Andrew Knyazev
University of Colorado at Denver

Robert Ward
University of Tennessee

Chao Yang
Lawrence Berkeley National
Laboratory

COMPUTER SCIENCE

Jim Ahrens
Los Alamos National Laboratory

Robert Armstrong
Sandia National Laboratories

David Bernholdt
Syracuse University

Brett Bode
Ames Laboratory

Sung-Eun Choi
Los Alamos National Laboratory

Erik DeBenedictis
Sandia National Laboratories

William Gropp
Argonne National Laboratory

James Kohl
Oak Ridge National Laboratory

John Mellor-Crummey
Rice University

Shirley Moore
University of Tennessee

Robert Numrich
University of Minnesota

Allan Snavely
San Diego Supercomputer
Center

Erich Strohmaier
Lawrence Berkeley National
Laboratory

Rajeev Thakur
Argonne National Laboratory

Mary Zosel
Lawrence Livermore National
Laboratory

NERSC STAFF REVIEWERS

HIGH ENERGY AND NUCLEAR
PHYSICS

Iwona Sakrejda

LATTICE QCD

Parry Husbands

ACCELERATOR PHYSICS

Andreas Adelmann

ASTROPHYSICS

Julian Borrill

Peter Nugent

CHEMISTRY

Jonathan Carter

David Skinner

MATERIAL SCIENCES

Andrew Canning

Lin-Wang Wang

EARTH AND ENGINEERING
SCIENCES

Osni Marques

FUSION ENERGY

Thomas DeBoni

Jodi Lamoureux

CLIMATE RESEARCH AND
ENVIRONMENTAL SCIENCES

Chris Ding

Helen He

Michael Wehner

LIFE SCIENCES

Adrian Wong

Jonathan Carter

APPLIED MATHEMATICS AND
COMPUTER SCIENCE

Esmond Ng

NERSC USERS GROUP EXECUTIVE COMMITTEE

MEMBERS AT LARGE

Ahmet Y. Aydemir,* University of Texas at Austin

David J. Dean, Oak Ridge National Laboratory

Stephen C. Jardin, Princeton Plasma Physics Laboratory

Ricky A. Kendall, Ames Laboratory, Iowa State University

Gregory Kilcup,* The Ohio State University

Donald Sinclair,** Argonne National Laboratory

Donald A. Spong,** Oak Ridge National Laboratory

John A. Taylor, Argonne National Laboratory

OFFICE OF ADVANCED SCIENTIFIC COMPUTING RESEARCH

David Bernholdt,** Oak Ridge National Laboratory

Ahmed Ghoniem,* Massachusetts Institute of Technology

David Keyes, Old Dominion University

Xiaolin Li,* State University of New York, Stony Brook

David Turner,** Ames Laboratory

OFFICE OF BASIC ENERGY SCIENCES

Feng Liu,** University of Utah

G. Malcolm Stocks, Oak Ridge National Laboratory

Theresa L. Windus, Pacific Northwest National Laboratory

OFFICE OF BIOLOGICAL AND ENVIRONMENTAL RESEARCH

Carmen M. Benkovitz,* Brookhaven National Laboratory

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

Douglas Rotman (NUG Vice Chair), Lawrence Livermore National Laboratory

Vincent Wayland, National Center for Atmospheric Research

OFFICE OF FUSION ENERGY SCIENCES

Stephane Ethier,** Princeton Plasma Physics Laboratory

Alex Friedman, Lawrence Livermore and Lawrence Berkeley National Laboratories

Jean-Noel Leboeuf,* University of California, Los Angeles

Carl Sovinec,** University of Wisconsin, Madison

Donald A. Spong,* Oak Ridge National Laboratory

OFFICE OF HIGH ENERGY AND NUCLEAR PHYSICS

Kwok Ko, Stanford Linear Accelerator Center

Douglas L. Olson, Lawrence Berkeley National Laboratory

Robert D. Ryne (NUG Chair), Lawrence Berkeley National Laboratory

* Outgoing members

** Incoming members

SUPERCOMPUTING ALLOCATIONS COMMITTEE

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

COMMITTEE CHAIRMAN

David Goodwin

**OFFICE OF ADVANCED
SCIENTIFIC COMPUTING
RESEARCH**

Daniel Hitchcock

**OFFICE OF BASIC ENERGY
SCIENCES**

Robert Astheimer

**OFFICE OF BIOLOGICAL AND
ENVIRONMENTAL RESEARCH**

Michael Riches

**OFFICE OF FUSION ENERGY
SCIENCES**

Michael Crisp

**OFFICE OF HIGH ENERGY
PHYSICS AND OFFICE OF
NUCLEAR PHYSICS**

David Goodwin

OFFICE OF ADVANCED SCIENTIFIC COMPUTING RESEARCH

Mathematical, Information, and Computational Sciences Division

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

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

C. EDWARD OLIVER

Associate Director, Office of Advanced Scientific Computing Research, and Acting MICS Director

ALAN J. LAUB AND WILLIAM H. KIRCHOFF

SciDAC

WALTER M. POLANSKY

Acting MICS Director (until March 2003)

WILLIAM H. (BUFF) MINER, JR.

NERSC and Applications Program Manager

DANIEL HITCHCOCK

ASCR Senior Technical Advisor

FREDERICK C. JOHNSON

Computer Science Research

GARY M. JOHNSON

Advanced Computing Research Testbeds

THOMAS D. NDOUSSE-FETTER

Networking

CHARLES H. ROMINE

Applied Mathematics and Statistics

MARY ANNE SCOTT

Collaboratory Research

GEORGE R. SEWERYNIAK

ESnet

LINDA TWENTY

ASCR Program Analyst

JOHN R. VAN ROSENDALE

Computer Science Research

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.

CO-CHAIRS

Margaret H. Wright
Courant Institute of
Mathematical Sciences
New York University

John W. D. Connolly
Center for Computational
Sciences
University of Kentucky

Roscoe C. Giles
Center for Computational
Science
Boston University

James J. Hack
National Center for
Atmospheric Research

Helene E. Kulsrud
Center for Communications
Research

William A. Lester, Jr.
Department of Chemistry
University of California, Berkeley

Gregory J. McRae
Department of Chemical
Engineering
Massachusetts Institute of
Technology

Thomas A. Manteuffel
Department of Applied
Mathematics
University of Colorado at
Boulder

Karen R. Sollins
Massachusetts Institute of
Technology

Ellen B. Stechel
Ford Motor Company

Stephen Wolff
CISCO Systems

MEMBERS

Jill P. Dahlburg
Naval Research Laboratory

David J. Galas
Keck Graduate Institute

ORGANIZATIONAL ACRONYMS

| | |
|-----------|--|
| DOE | U.S. Department of Energy |
| SC | Office of Science |
| ASCR-MICS | Office of Advanced Scientific Computing Research, Mathematical, Information, and Computational Sciences Division |
| BER | Office of Biological and Environmental Research |
| BES | Office of Basic Energy Sciences |
| FES | Office of Fusion Energy Sciences |
| HEP | Office of High Energy Physics |
| NP | Office of Nuclear Physics |
| SciDAC | Scientific Discovery through Advanced Computing Program |
| 3M | 3M Company |
| AEI | Albert Einstein Institute/Max Planck Institute for Gravitation Physics (Germany) |
| AF | Aventis Foundation |
| AFOSR | Air Force Office of Scientific Research |
| ARO | Army Research Office |
| BMBF | Bundesministerium für Bildung und Forschung (Germany) |
| BWF | Burroughs Wellcome Fund |
| CAG | Cellzome AG |
| CDWR | California Department of Water Resources |
| CNR | Consiglio Nazionale delle Ricerche (Italy) |
| CNRS | Centre National de la Recherche Scientifique (France) |
| DAAD | German Academic Exchange Service |
| DARPA | Defense Advanced Research Projects Agency |
| DOD | Department of Defense |
| DRCRF | Damon Runyon Cancer Research Foundation |
| DFG | Deutsche Forschungsgemeinschaft |

APPENDIX G

| | |
|---------|--|
| DRA | Danish Research Agency |
| EMBL | European Molecular Biology Laboratory |
| EMBO | European Molecular Biology Organization |
| EPSRC | Engineering and Physical Sciences Research Council (UK) |
| ESOP | European Electron Scattering Off Confined Partons Network |
| EU | European Union |
| FAPESP | Fundação de Amparo à Pesquisa do Estado de São Paulo, Brazil |
| FBF | France-Berkeley Fund |
| GA | General Atomics |
| GNMF | German National Merit Foundation |
| HFS | Holderbank Foundation Switzerland |
| HFSP | Human Frontier Science Program |
| IBM SUR | IBM Shared University Research Grant |
| IN2P3 | Institut National de la Physique Nucléaire et de la Physique des Particules (France) |
| INSU | Institut National des Sciences de l'Univers (France) |
| Intel | Intel Corporation |
| JHIR | Joint Institute for Heavy Ion Research |
| KBN | Polish State Committee for Science Research |
| LANL | Los Alamos National Laboratory |
| LBNL | Lawrence Berkeley National Laboratory |
| LLNL | Lawrence Livermore National Laboratory |
| LSU | Louisiana State University |
| MCYT | Spanish Ministry of Science and Technology |
| MEC | Ministry of Education of China |
| MEVU | William A. and Nancy F. McMinn Endowment at Vanderbilt University |
| MEXT | Japan Ministry of Education, Culture, Sports, Science, and Technology |

APPENDIX G

| | |
|--------|--|
| MIT | Massachusetts Institute of Technology |
| MIUR | Ministero dell'Istruzione, dell'Università e della Ricerca (Italy) |
| NASA | National Aeronautics and Space Administration |
| NCI | National Cancer Institute |
| NEDO | New Energy Development Organization (Japan) |
| NIH | National Institutes of Health |
| NNSFC | National Natural Science Foundation of China |
| NOAA | National Oceanographic and Atmospheric Administration |
| NRC | National Research Council |
| NSERC | Natural Sciences and Engineering Research Council of Canada |
| NSF | National Science Foundation |
| NSFC | National Science Foundation of China |
| ONR | Office of Naval Research |
| ORNL | Oak Ridge National Laboratory |
| PECASE | Presidential Early Career Award for Scientists and Engineers |
| PNC | Programme National de Cosmology (France) |
| PPARC | Particle Physics and Astronomy Research Council (UK) |
| RFBR | Russian Foundation for Basic Research |
| RMST | Russian Ministry of Science and Technology |
| SNS | Spallation Neutron Source Project |
| UC | University of Cyprus |



Published by the Berkeley Lab Technical and Electronic Information Department in collaboration with NERSC researchers and staff.
JO#8798

EDITOR AND PRINCIPAL WRITER: John Hules
CONTRIBUTING WRITERS: Rich Albert, Jon Bashor, Christopher Jones, Paul Preuss, Lynn Yarris

DISCLAIMER

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

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

