access

NCSA | Vol 20 No 1 | Spring 2007
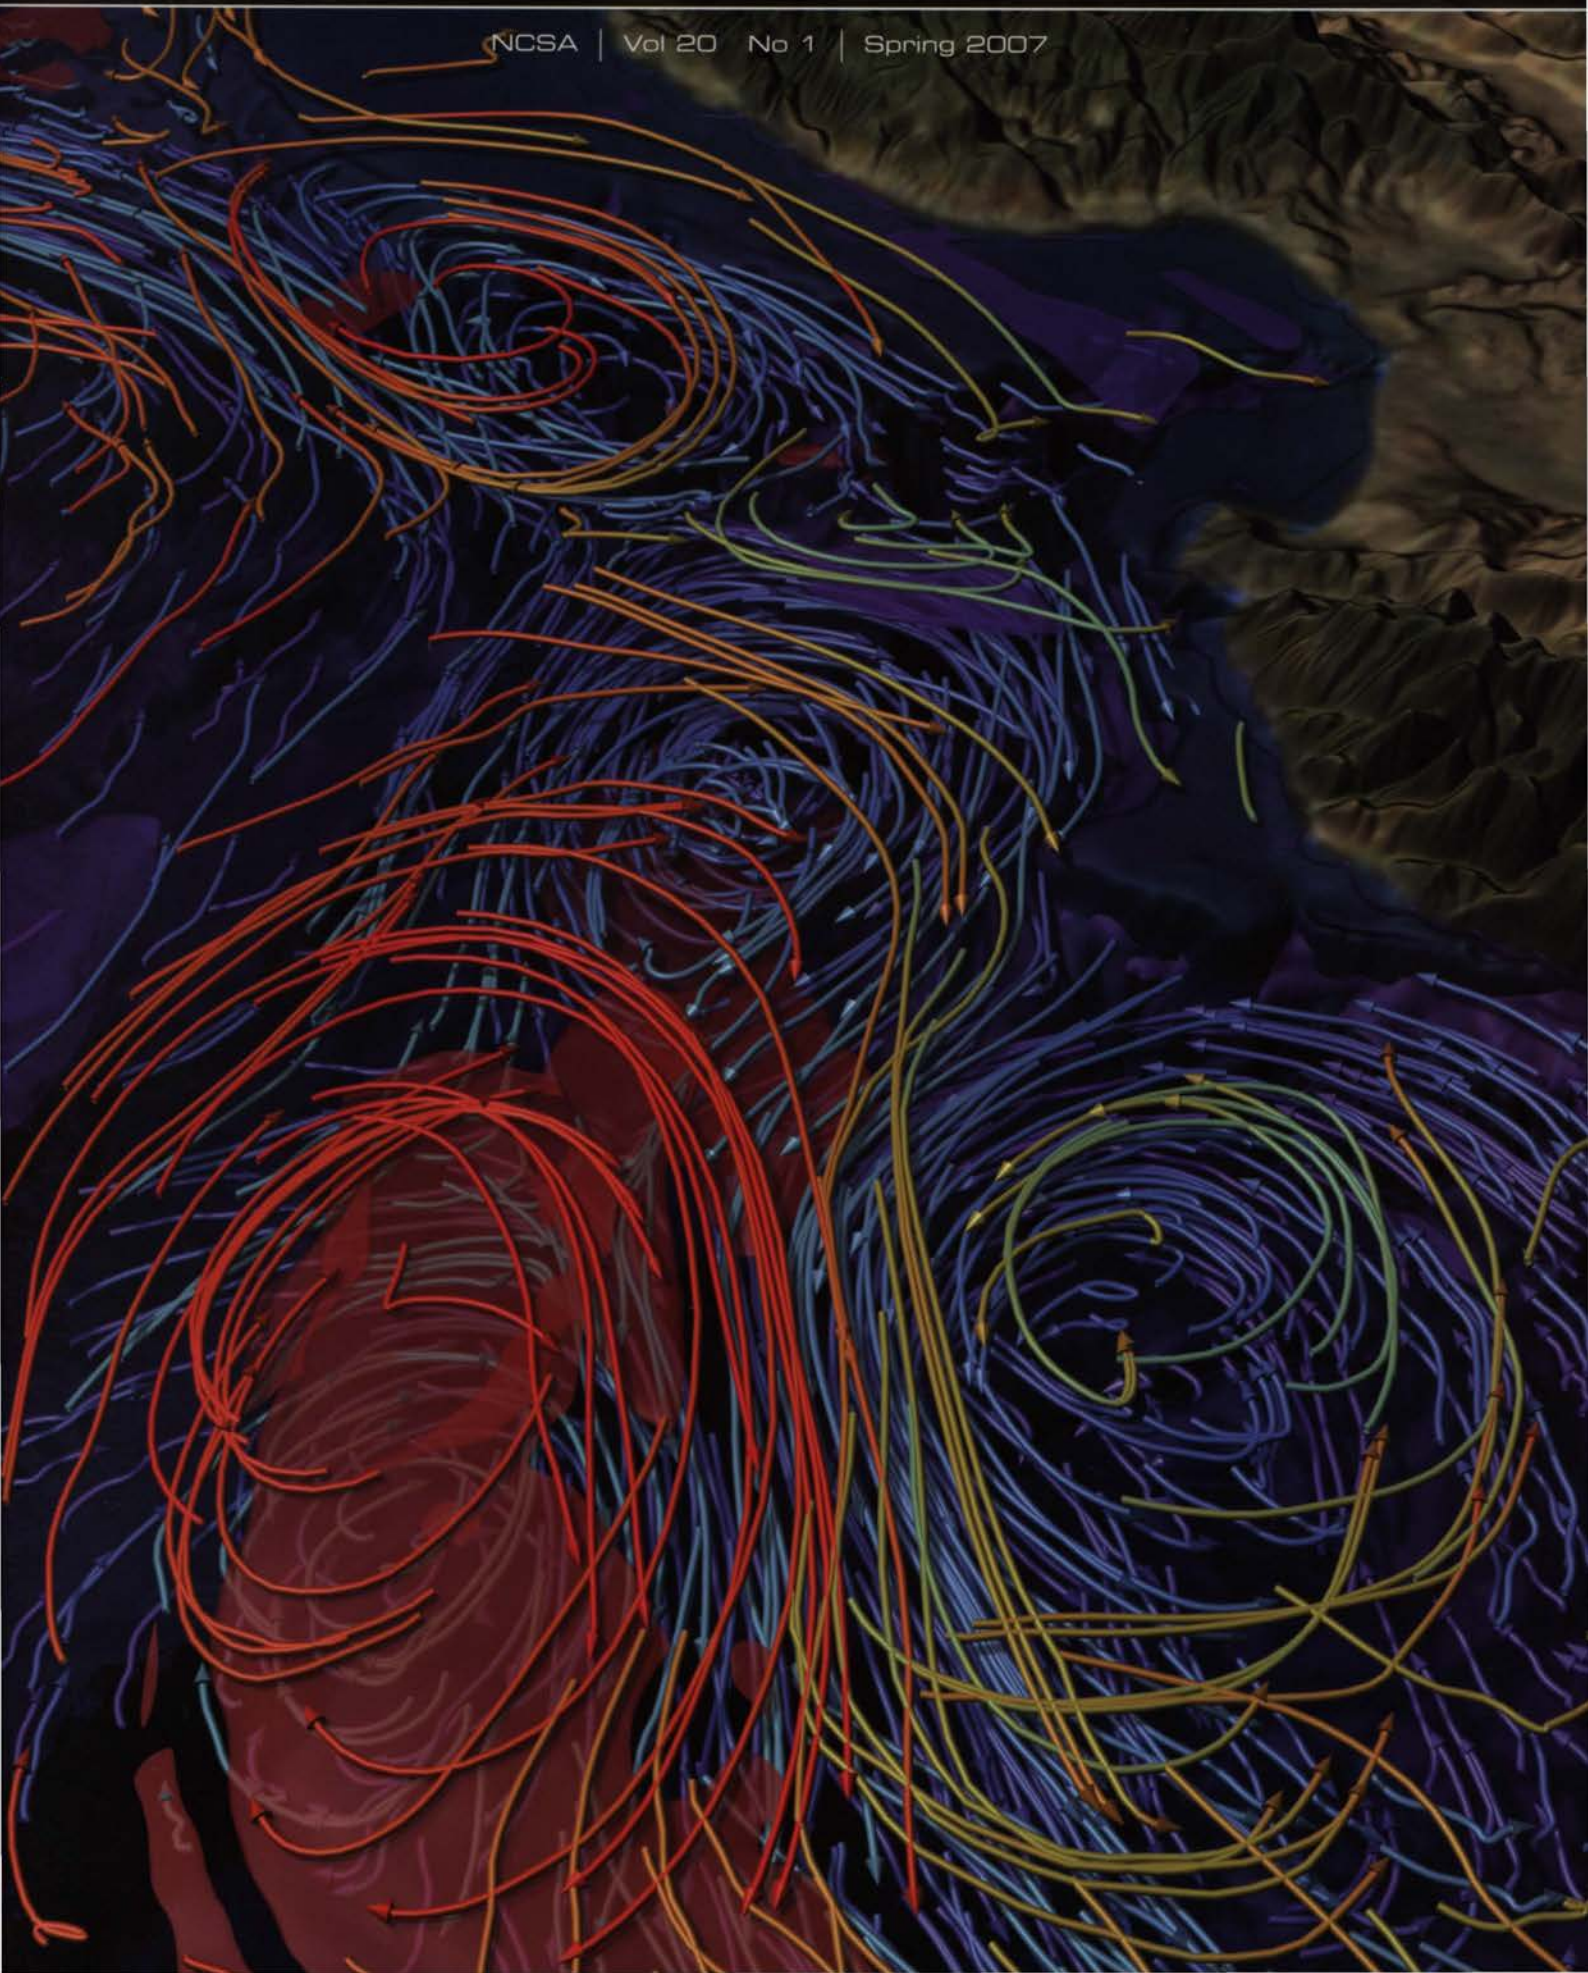
On the Web
News from NCSA is always available on the Web. Sign up for our RSS feed or email service for the latest developments from NCSA and its partners.
http://www.ncsa.uiuc.edu

National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign
1205 West Clark Street
Urbana, IL 61801
217-244-0072

Who we are
The University of Illinois at Urbana-Champaign's National Center for Supercomputing Applications (NCSA), one of the five original centers in the National Science Foundation's Supercomputer Centers Program, opened its doors in January 1986. Over the past twenty years, NCSA has contributed significantly to the birth and growth of the worldwide cyberinfrastructure for science and engineering, operating some of the world's most powerful supercomputers and developing the software infrastructure needed to efficiently use them. Today the center is recognized as an international leader in deploying robust high-performance computing resources and in working with research communities to develop the new computing and software technologies needed to take full advantage of the rapidly expanding national and international cyberinfrastructure.

The center focuses on three themes. Cyberenvironments will give research communities the means to fully exploit the extraordinary resources available on the Internet (computing systems, data sources and stores, and tools). NCSA's cyber-resources ensure that computing, data, and networking resources are available to solve the most demanding science and engineering problems and that the solutions are obtained in a timely manner. Finally, innovative systems research explores the path to petascale computing—testing and evaluating the performance of emerging computing systems for scientific and engineering applications.

NCSA is a key partner in the National Science Foundation's TeraGrid project, a $100 million effort to offer researchers remote access to some of the fastest unclassified supercomputers as well as an unparalleled array of visualization tools, application software, sensors and instruments, and mass storage devices. NCSA also leads the effort to develop a secure national cyberinfrastructure through the National Center for Advanced Secure Systems Research, a project funded by the Office of Naval Research.

The center leaves its mark through the development of networking, visualization, storage, data management, data mining, and collaboration software as well. The prime example of this influence is NCSA Mosaic, which was the first graphical Web browser widely available to the general public. NCSA visualizations, meanwhile, have been a part of productions by the likes of PBS's NOVA and the Discovery Channel. Through its Private Sector Program, top researchers explore the newest hardware and software, virtual prototyping, visualization, networking, and data mining to help U.S. industries maintain a competitive edge in the global economy.

Major support for NCSA is provided by the National Science Foundation. Additional funding comes from the state of Illinois, industrial partners, and other federal agencies. For more information, see
http://www.ncsa.uiuc.edu/.

| 1 | 5 | 10 | 15 | 19 | -2 SD | +2 SD |

ROMS Temperature (Celsius)   **Temperature Deviation**

On the cover
NCSA's Advanced Visualization Laboratory (AVL) works with scientists to develop and use the cyberinfrastructure to create high-end, high-fidelity, and high-resolution data-driven scientific visualizations. These visualizations are technically developed and aesthetically designed to support scientific narratives. The goal is to communicate, inspire, and reach out to large, non-expert audiences.

In creating these scientific productions, the AVL team develops new cybertechnologies and advanced visualization tools and software pipelines. All of the team members play unique roles and contribute a variety of skills to the process and productions; expertise includes advanced graphics and visualization software development, artistic design and choreography, multimedia production, data management, and render wrangling.

The AVL collaborated with the Monterey Bay Aquarium Research Institute (MBARI) and the Jet Propulsion Laboratory/Caltech to visualize Regional Ocean Modeling System (ROMS) simulation data. The image shown here is a visualization of ocean currents and temperature. The stream tubes, colored by temperature, represent 24 hour flow trajectories.

NCSA | I

access contents

John Towns
NCSA Director of Persistent Infrastructure
page 02

SUPERCOMPUTING IN THE 21ST CENTURY

As information technology and the emerging concept of cyberinfrastructure permeate science and engineering research endeavors, the notion of a traditional supercomputing center becomes dated. NCSA is working to keep ahead of these changes and is evolving into an institution whose *raison d'etre* is more than big iron. The supercomputing center of the 21st Century will be the lynchpin in an enterprise focused on the concept of empowering scientists and engineers.

Cyberinfrastructure is still, and will remain, indispensable. This is true both for us as a center and for our users who need the incredible capacity and capability that cyberinfrastructure provides. But in a new world we have to think in new ways about the things that remain the same. At NCSA, the shift in how we think about ourselves has led to a shift in the way we deploy and manage even our longest-standing services.

The major usage mode of the past 25 years for high-end computing—allocation of time to a single investigator, submission of jobs to a queue for execution, post-processing and analysis of the output later by the researchers—is giving way to other usage

modes. While this will remain a critical process for some users, as the community of scientists and engineers expands (both in number and into new disciplines) we find many users have new needs. Needs such as rapid turnaround to accelerate the science process—through interactive science and to allow for timely progress in development activities at scale. Others need dedicated resources for extended periods to ensure the timely completion of their projects or to assure that data can be handled when it becomes available from various data sources. Some even need on-demand computing to react to unforeseeable events. NCSA is working closely with the scientists and engineers we support to ensure that our cyber-resources and policies match these needs. The shift can already be seen along a number of dimensions.

First, we're working more closely with researchers in particular communities during the planning and acquisition process. Through our ongoing interactions, we continuously gather the evolving requirements; we then design, build, and operate the system expressly for them. The deployment of our newest resource, Abe, is a direct response to several requirements noted by the user community: Researchers will be able to reserve blocks of Abe for days or even weeks, allowing them to accelerate their research. We are hopeful about the currently submitted proposals for the NSF "Track 1" and "Track 2" solicitations, proposals born of the in-depth discussions we have had with both current and expected users of these resources. The systems are truly designed-to-order, reflecting users' technical needs better than systems of this magnitude ever have to provide capability computing given increasing general availability of capacity computing resources.

We're also allocating our existing cyber-resources—open to any peer-reviewed project that is awarded an account—so that they are more conducive to real-world use. For example, our Abe system has a peak speed of about 90 teraflops, making it the largest system supported by the National Science Foundation to date. The operation of Abe will build on capabilities we have developed on Tungsten to provide what we call tailored allocations. In a tailored allocation, specific pieces of the machine are turned over to particular users for given periods of time. These periods are planned in advance so that users know when they're going to get access, how long they're going to have it, and how much computing oomph they're going to have available. These dedicated runs give users the capability they need to complete crucial computations that must be done in a specific time frame, that require an unusually large number of processors, or

that otherwise give the queuing system fits. We are excited to be able to expand this service while still leaving room for the traditional capacity-oriented use, with smaller jobs passing through the queue and running unimpeded. Policy issues are the primary means of changing the services provided to scientists and engineers; these are driven by their needs and our desire to support their success.

NCSA's superior support of scientists and engineers is reflected in the 2005 and 2006 Cyberinfrastructure Partnership user surveys; an overwhelming majority of our users and users at other NSF-funded sites were pleased with the performance of our systems and the user support they receive, giving the highest marks in these areas to NCSA. But we won't rest there. We are increasing the services NCSA offers not only to provide a superior level of service to our traditional users as they continue to scale their applications to higher levels, but to also address the needs of researchers with less proficiency accessing cyber-resources via novel cyberenvironments. Both the complexity and scale of the resources continue to dramatically increase. As a result, we are developing the expertise in application scaling and performance optimization and in supporting cyberenvironments. Thus we can support the most demanding applications and deploy the cyberenvironments to support the needs of broad communities requiring low thresholds to access. To proceed in any other manner is unthinkable.

Cyber-resources will remain crucial to our success. They energize our ideas. They are the core that we build out from. They must be as powerful as our ingenuity can make them. But making them larger and faster is no longer our organizing principle. This fact pushes us outside our comfort zone and forces us to think in new ways. It's an exhilarating time at NCSA as we prepare to define this world turned upside down.

John Towns
Director of Persistent Infrastructure
NCSA

Where the rubber hits the road

In 2006, Dan Atkins was named head of the National Science Foundation's newly formed Office of Cyberinfrastructure (OCI), after chairing a blue ribbon panel on cyberinfrastructure for the foundation. Atkins discussed cyberinfrastructure with Access' J. William Bell in early 2007. Among other topics, they touched on a new vision document on cyberinfrastructure coming out of NSF, the social dynamics involved in infrastructure development, and relationships within the NSF's TeraGrid program, of which NCSA is a part.

Q: What forces within the discipline sciences are influencing cyberinfrastructure development today?

A: More and more fields have intrinsically complex, multiscale problems to confront. These problems require greater interdisciplinarity, often interinstitutional organizations with distributed players...One of the reasons that cyberinfrastructure is so critical to NSF in particular is so that it can take full advantage of the interdisciplinarity across all of the fields that it supports. The supercomputer makes more and more comprehensive, multiscale simulations and models feasible and requires more collaborative activity and complementary expertise than has been required in the past. There's a positive feedback loop between interdisciplinarity and increased power of computing environments.

Q: How do we translate that feedback loop into real dividends for cyberinfrastructure users?

A: Cyberinfrastructure is shifting the boundary between when it's in your self-interest to collaborate and when it's in your self-interest to compete. You're seeing, for example, more communities coming together to create community codes and share various functional models, and similar things happening around databases and production codes.

Speaking as a former dean and academic administrator, it's important that institutions make sure that they recognize these growing, intrinsic demands of science and engineering research and try to relax as many constraints in these areas as possible.

Q: You headed up a blue ribbon panel that helped give rise to NSF's Office of Cyberinfrastructure. Now that you've been in your OCI position for six months, what's struck you about cyberinfrastructure development?

A: I've been impressed with the diversity of activity that's underway. It's also been further reinforced that [cyberinfrastructure development] is an organic process that's here to stay. It's intrinsically a bottom-up, community-driven aspiration. I think the vision that

NSF has established for the potential transformative impact of cyberinfrastructure is pretty ambitious and reflects that. [A final version of this vision, "Cyberinfrastructure Vision for the 21st Century Discovery," is currently at http://www.nsf.gov/pubs/2007/nsf0728/index.jsp]

[It's also shown me] we're going to have to give more attention to how to make this a true multistakeholder system and to figuring out how to be a lot more clever in how we make, leverage, and find common ground between the activities of various fields and the investments being made by various stakeholders.

Q: It strikes me that the word "infrastructure" in the name of the office is pretty fraught with meaning. Help us understand why that infrastructure mindset is particularly important on these sorts of issues.

A: You're right, we've taken the infrastructure word quite seriously as opposed to random acts of applying information technology...[This mindset] affects the role of OCI in trying to invest its money in things that are useful to a variety of fields. A layered approach where we provide resources and middleware and tools that allow communities to tailor specific resources that they need for their particular work and rapidly build customized science gateways or collaboratories or cyberenvironments or virtual research environments. We're trying to build complementarity between what we invest in and the particular, discipline-area investments that the directorates are making. And to do that systematically.

Q: What goals does the vision document that you mentioned raise for your office and for the people building the technology?

A: It starts with a very bold statement of NSF's intent to provide leadership in this area. Then it talks about four intersecting components of activity and impact.

The first area, [high-performance computing,] places a high priority on science-driven, research- and education-driven cyberinfrastructure. The need, for example, for higher and higher

Questions & Answers

levels of computational performance comes from the increasing complexity and multiscale nature of the problems that contemporary scientists are taking on and, in some cases, the massive amounts of data associated with doing work in that area...The commitments to supporting a multitiered system of petascale and terascale computation is pretty much well developed and a funding priority has been established for that.

With the second area, data, the points to get across are the incredible increase in the volume of data that modern science needs to contend with and the increased need for improved stewardship, coming in part from the enhanced value of data. As you know, you can extract new knowledge from data than what was originally intended, so you run a greater risk from throwing data away. You might lose opportunity.

Then the third area builds upon the first two—virtual organizations. We're moving in the direction of more and more functionally complete collaboratories and virtual organizations, so that through a web-based portal you can reach the primary colleagues, data, computational models, and the other things they need to do their work. When you get to that point, you quite dramatically relax the constraints of distance and time, and this ability for broad participation in authentic scientific research communities is part of this excitement in the democratization of science.

It fits with the Web 2.0 movement—the idea that the web and the Internet are not just about information. They're about participation...

The fourth area is what's being called learning and workforce development. [This is] both the application of cyberinfrastructure to enhance learning, and in particular to provide hands-on experience with science to greater numbers of people, as well as the educational and workforce demands to empower more people to use and create cyberinfrastructure.

Q: The document talks about both provisioning the hardware and having a robust software and applications environment. Historically, that's been a challenging relationship. How do we make sure those keep up with one another?

A: Realistically, at least part of the approach needs to be providing mechanisms for the stout-hearted who want to do these codes initially. There's going to be a lot of hard work and hand crafting to get these codes to scale up to the peta regime.

It's also important to have mechanisms that look for common needs across fields and across disciplines and to try to extract those out and make them available so you're not constantly reinventing the wheel. So you're leveraging off the experiences and the developments of others...

The activities of the supercomputer centers and the various TeraGrid resource providers are going to be critical to that. That's really where the rubber hits the road between particular application communities and project teams that have need for or sign up to be pioneers in the migration to petascale. They're really going to be the ones who figure out how we do this.

Q: You mentioned the concept of Web 2.0. Are there hints of attitudes within our community moving in that direction?

A: The open courseware movement that started at MIT [has] now extended to 100 major universities. There's kind of a feeling that resources such as curriculum and so forth should be available to the world. That's quite a significant social shift.

These fundamental shifts in social architecture and microcontribution shift the boundaries between competition and cooperation. That could actually lead to quite different approaches to how work is done...

[Web 2.0 and] those shifting paradigms of creativity and shifting notions of intellectual property...are an exemplar in the popular culture of the social impact that these technologies are having. Analogous things will happen around education and research.

Q: One of the interesting tensions here is that infrastructure implies stability, while the Web 2.0 attitudes imply something very active.

A: There are sustainability and predictability issues involved in infrastructure, but you actually do want it to evolve—in a predictable way...I think it's important to realize that we're not talking about creating the cyberinfrastructure, infrastructure for science as one homogeneous system. We're talking about a more organic process where there are all types of systems of various sorts that are being built. To make it an infrastructure we need to deal with how we get interoperability and commonality and gateways. How we accommodate the heterogeneity of systems and create a more pervasive infrastructure. We're not trying to impose a standard, we're trying to figure out a way to accommodate heterogeneity.

You can have an overall infrastructure composed of systems that are at different levels of maturity and performance. We don't want to think of TeraGrid and Open Science Grid as being competing alternatives. One is best for some classes of jobs, the other is best for some classes of jobs.

From the point of view of a functioning scientist, they should think of and be able to use both of those grids in a complementary, ideally seamless way...Let's not think of TeraGrid as one cyberinfrastructure and OSG as another. But rather a TeraGrid system and an Open Science Grid system as components of an overall cyberinfrastructure for science.

Q: And in order to come to that, you mentioned being sure we look at it in the right way and you mentioned building the system at other levels of abstraction. What else can we at the centers-level do to encourage that, whether technical or social?

A: I think part of it has to do with finding a comfort zone in the space between cooperation and competition. And finding ways of all of us humbling ourselves to the collective opportunity that stands before us. The kinds of things that we're trying to work out in the context of the TeraGrid—the balance between the autonomy of and the spires of excellence of individual resource providers but also making these sources add up to an overall system for the science community. Being positive minded about such activities.
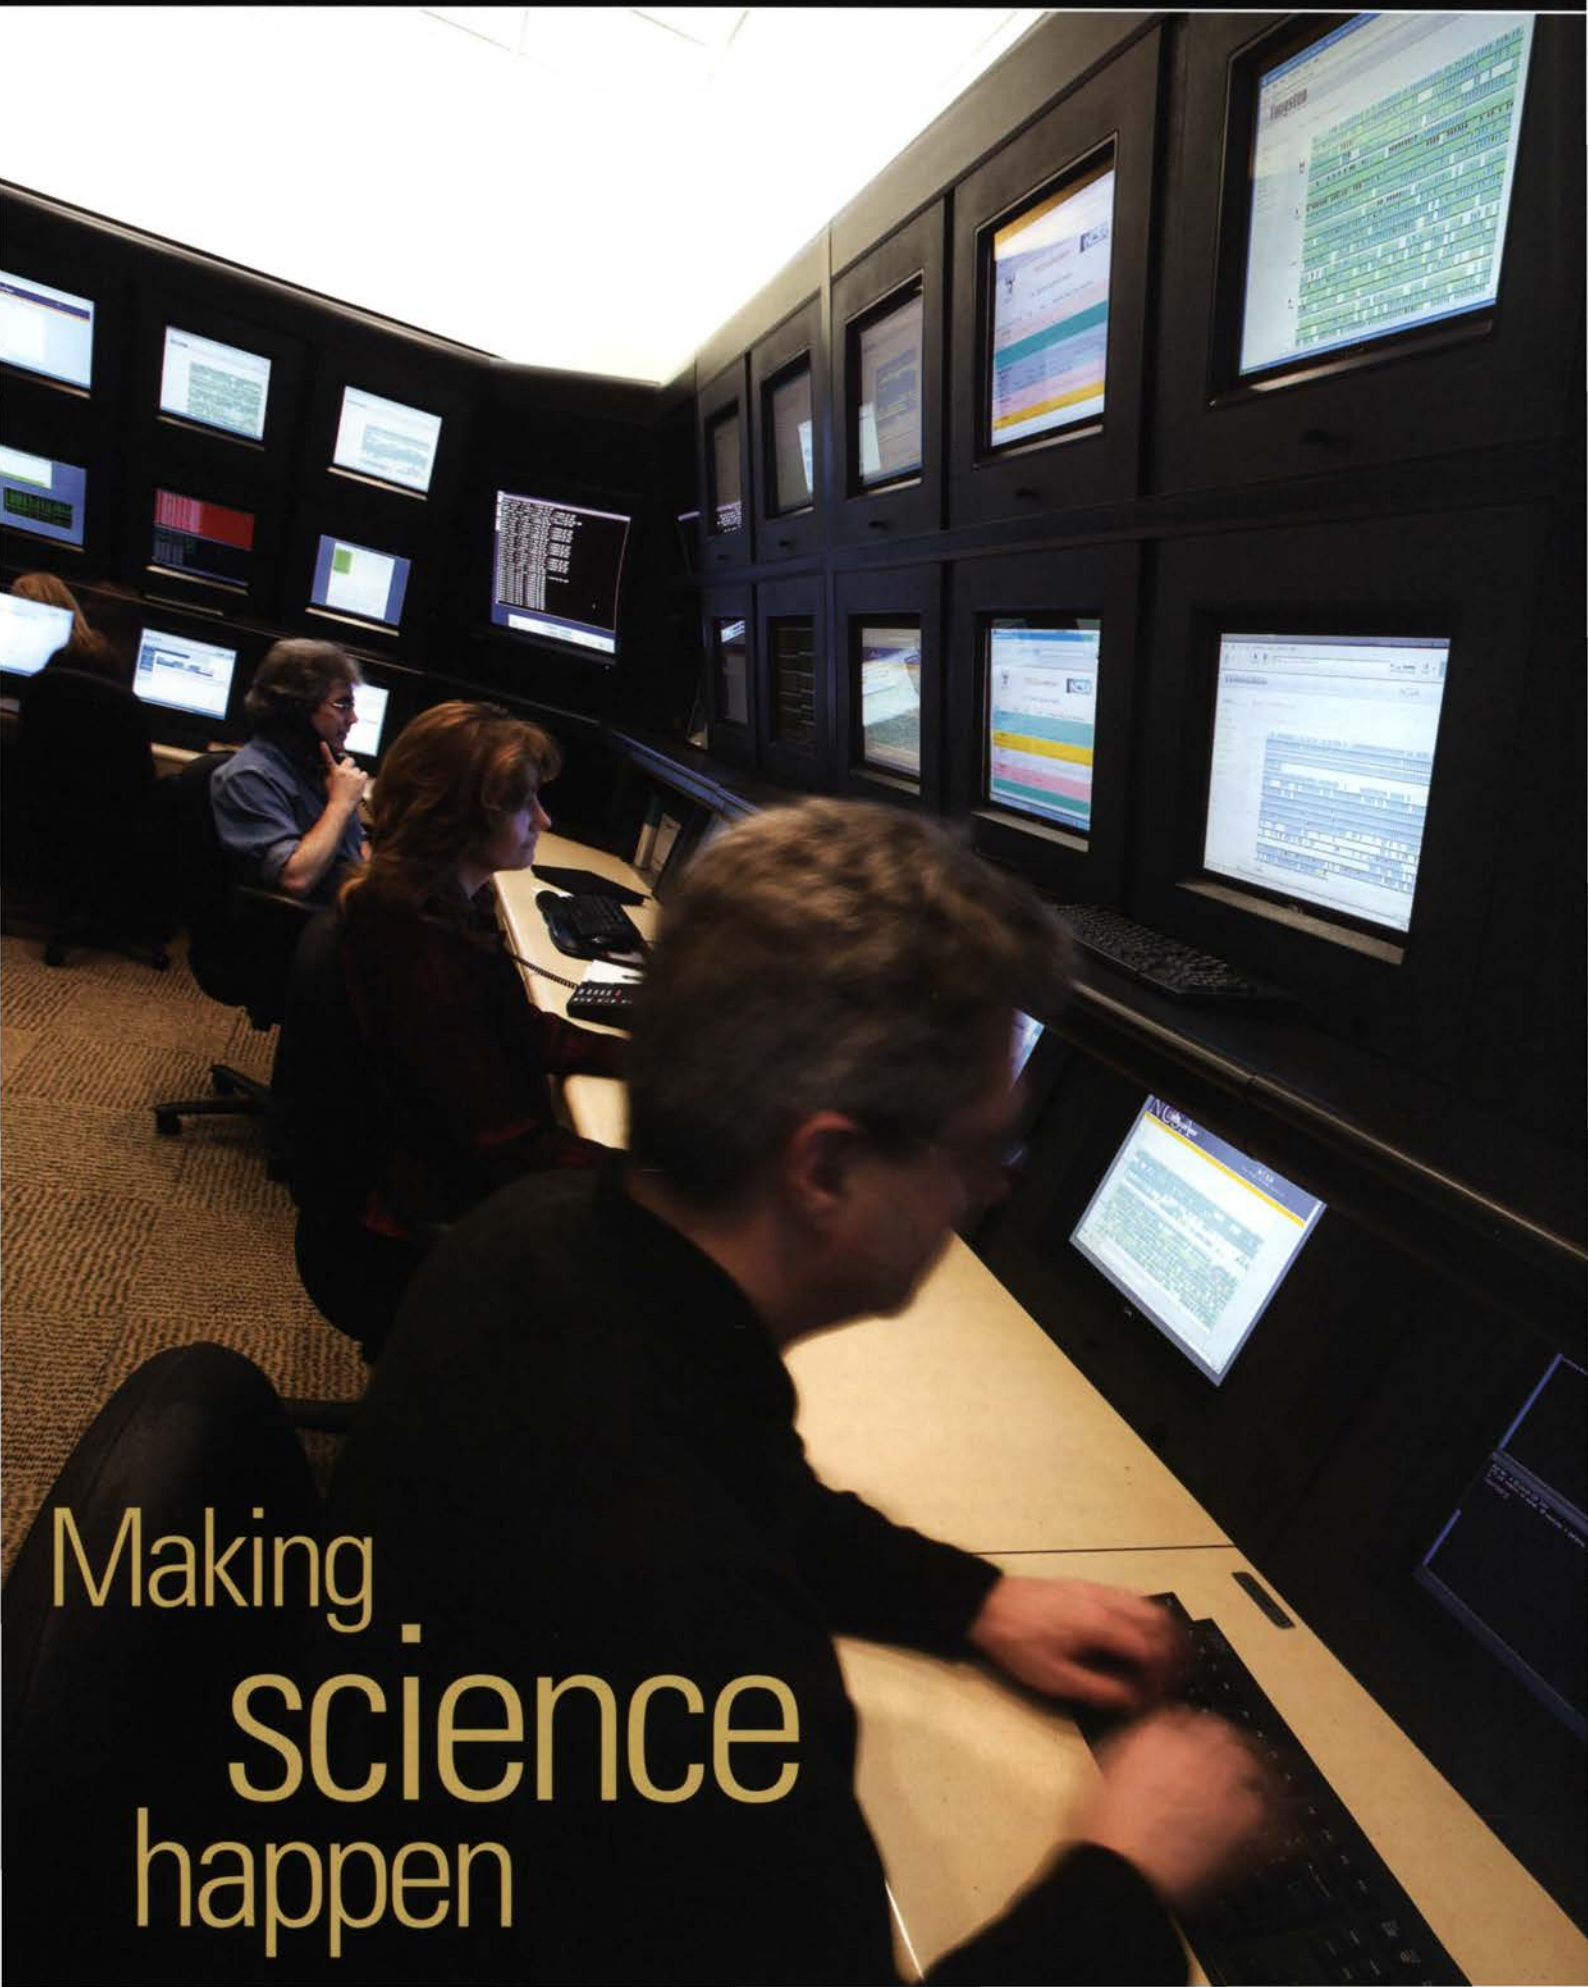
Making science happen

by Barbara Jewett

The user support offered by NCSA is unsurpassed. From answering the simplest question to assisting with complicated codes, advanced visualizations, or dedicated runs, NCSA has developed a level of user support that keeps science moving forward.

Microphysical processes in tropical cyclones. Thermal conduction in nanofluids. Formation of galaxies and supermassive black holes. Risk assessment of aerosol dispersion in a city. Determining pavement damage owing to different axle configurations. These are just a few of the research topics investigated using NCSA compute resources last year. In fact, nearly 1,400 scientists, engineers, and students racked up more than 717 million normalized units of computing time. And they once again gave NCSA top marks for user support services in the 2006 Cyberinfrastructure Partnership user survey because they get quick responses to their simple questions and the technical assistance they require for the more advanced ones.

Support for projects big or small

Support at NCSA takes many paths. Here's a brief look at just a few of the many ways researchers are assisted, whether their project is running solely on NCSA machines or involves TeraGrid resources.

The first route is the help desk, which provides phone and email support around the clock, every day of the year. For slightly more difficult issues, the consulting office is the next source of support. The consulting office aids all users of NCSA resources and services. They help users learn how to run their jobs on NCSA machines, including running batch jobs, or providing some programming assistance, either in general or for a particular application. The fields of computational chemistry, computational biology, structural mechanics, computational fluid dynamics, math, and physics require specialized expertise; the consulting office staff includes people proficient in the software programs these researchers use with the know-how to capably handle other issues.

"The NCSA consulting office was extremely helpful to me," says Renyue Cen, a senior research astronomer at Princeton University. "They implemented a required faster, parallel I/O scheme into the TIGER code and helped debug the code. As a result, we were able to make the world's largest cosmological hydrodynamic simulation with 8 billion cells."

Aside from the scientific results Cen obtained, this simulation was visualized by NCSA's Advanced Visualization Laboratory and featured in the PBS NOVA program "Monster of the Milky Way" that aired last October (http://access.ncsa.uiuc.edu/Stories/MilkyWay).

Getting researchers past all the barriers and making science happen is the objective of NCSA, and nowhere is that more evident than in the area of advanced support. Providing the missing link in a research project sometimes requires that user support experts focus on a particular researcher's stumbling block to success. To move the research forward might mean creating custom software

The NCSA help desk provides round-the-clock phone and email support to users at NCSA and throughout the TeraGrid. Staff also monitor and maintain the machines.
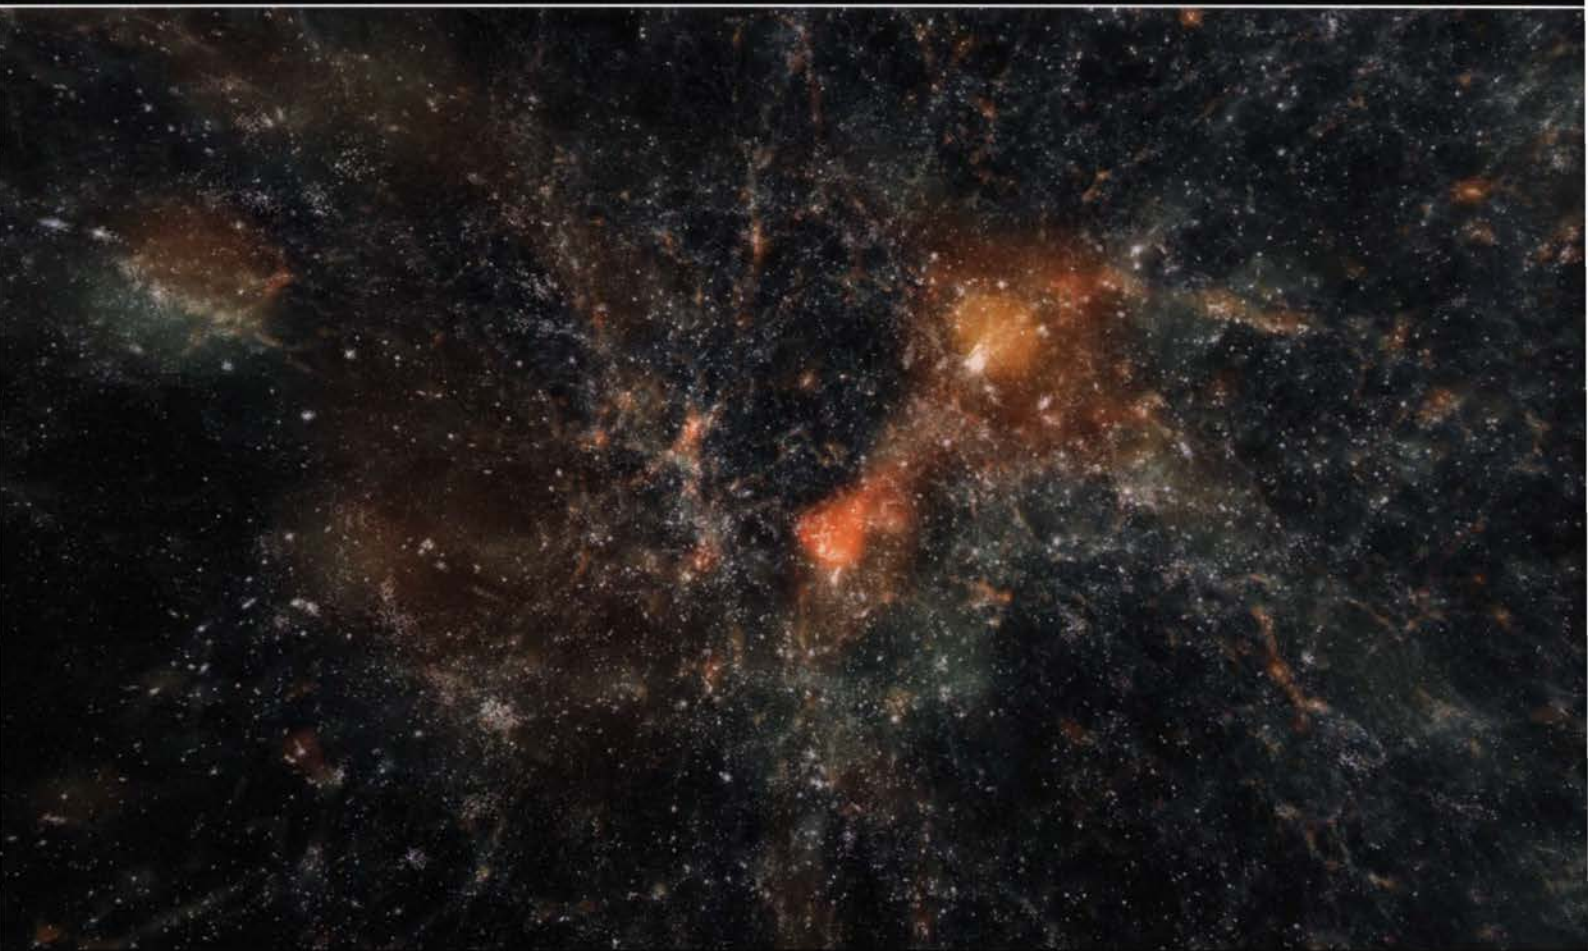
This visualization created by NCSA's Advanced Visualization Laboratory represents the nonlinear evolution of the universe, from 20 million to 14 billion years old. The simulation was created by Princeton researchers Renyue Cen and Jerimiah P. Ostriker using NCSA's computational resources.

or benchmarking systems. Dedicated services range from short-term efforts in understanding code on a particular processor to months-long dealings to achieve goals in areas such as scalability, large dataset management, and advanced visualization. Current projects include assisting with simulating novel biofuel liquid combustors, developing computational methods for large-scale simulation of buildings subjected to earthquakes and benchmarking systems at NCSA and the San Diego Supercomputer Center (SDSC).

Capability computing is a key component to helping researchers achieve success. For example, David Baker of the University of Washington developed a code, called Rosetta, to model the structure of proteins (http://access.ncsa.uiuc.edu/Stories/rosetta). Baker's team worked with SDSC to adapt the code to run on high-performance computing systems, then turned to NCSA for compute power. It was quickly apparent that in order for Baker to make progress with his research he would need substantial fractions of available resources, so NCSA set aside 128 nodes of the Tungsten cluster for several months for his use.

Baker also developed a web portal, called Robetta, which allows other researchers to run Rosetta jobs. Initially, Robetta jobs were run on a small cluster in Baker's lab, but when submissions exceeded capacity, NCSA was able to provide a solution. Working with Baker's team and the Condor development team from the University of Wisconsin, NCSA was able to integrate the Condor workload management system supporting Robetta with the center's Condor cluster, effectively mirroring the Robetta server at NCSA.

But Baker is not alone in his need for dedicated resources. At any given time last year, up to 60 percent of Tungsten's capacity was in use for dedicated allocations. And the requests for dedicated allocations have risen so dramatically that NCSA has installed a new almost 90 teraflop system, Abe, to meet the need. The addition of Abe also allows the center to further expand our capability computing responsiveness, whether it is dedicated allocations or on-demand scheduling for running simulations, such as in disaster situations.

We have data...now what?

Large project or small, all research utilizing NCSA has some basic parameters regardless of subject matter: Researchers generate data, next they need to process it, move it, and store it.

Answering the questions of how to deal with data—which, as in the case of astrophysics data from telescopes can be in the realm of tens of terabytes—is another area in which NCSA supports scientists

and engineers. Users at NCSA have access to high-speed parallel file systems on each of our platforms to support the creation and analysis of large data sets. There is more than a petabyte of raw disk capacity distributed among our machines as well as a long-term data storage system with several petabytes of archival capacity.

For instance, the Large Synoptic Survey Telescope (LSST) project will generate an estimated 15 terabytes of raw data and 100 terabytes of processed data every night when it begins operating in 2013. To lay the foundation for the LSST cyberenvironment, NCSA and the National Optical Astronomy Observatory (NOAO) are developing a prototype data pipeline using the vast stores of data generated by the ground-based observatories NOAO oversees (http://access. ncsa.uiuc.edu/Stories/lsst/Data1.html). NCSA created a mirror of the NOAO data archive in Urbana. The archive replication system is built on the Storage Resource Broker middleware developed at SDSC and a transfer queuing system NCSA developed for its archive of data from the Berkeley-Illinois-Maryland Array radio telescope.

Mirroring NOAO's archive not only gives astronomers a high-bandwidth site from which to access data, it is also part of a strategy of security through redundancy, ensuring that data will survive at one site even if a catastrophic event hits another site. NCSA's robust mass storage system archives more than 3 petabytes of researchers' data, with data added at a rate of 40 to 60 terabytes each month. However, the data management expertise of the center's staff—which encompasses experience running large file systems, parallel file systems, many storage architecture types, HPC storage, and database storage—is as much of an asset as the storage infrastructure.

As challenging as the tasks of capturing, moving, processing, and storing data are, they are just the preliminary steps. The real excitement begins when researchers can access and analyze the data. The center provides software to manipulate the data, and NCSA experts assist as needed with data management and visualizations.

For LSST, community access will be provided through a web-based virtual observatory (VO). NOAO and NCSA are working to develop the VO model and the VO tools, including an authentication and authorization framework for the NOAO portal, an online tool that enables users to find, access, and analyze the data available through multiple public archives, such as the Sloan Digital Sky Survey, Canadian Network for Observational Cosmology, Chandra X-ray Observatory Center, and others.

Access Online: http://access.ncsa.uiuc.edu/CoverStories/usersupport

Abe to serve research, T3 to support industry

NCSA is giving a powerful boost to researchers in academia and industry with the addition of two new supercomputing clusters. The two systems, Abe and T3, will provide 110 teraflops of computing power, more than doubling NCSA's total capacity to over 146 teraflops.

Abe, named in honor of 16th president Abraham Lincoln, has a peak performance of about 90 teraflops. The system was purchased with funds from the state of Illinois and will be shared by researchers at the University of Illinois' new Institute for Advanced Computing Applications and Technologies, headed by NCSA Director Thom Dunning, and by the national science and engineering communities served by NSF.

Continuing a model that has proven very popular with scientists using NCSA's Tungsten cluster, the center will allow researchers to reserve blocks of Abe (at a minimum of 1,000 cores) for days or even weeks, allowing them to accelerate their research. Abe's users will even be able to reserve access to all 1,200 nodes (9,600 cores) of the cluster for the most demanding simulations.

Abe, which will soon enter production, is a Dell blade system with 1,200 PowerEdge 1955 dual-socket, quad-core Intel Xeon 2.3 GHz processors, and InfiniBand and GigE connections. Each processor has 4 gigabytes of memory, providing a system total of 9.6 terabytes. The InfiniBand interconnect provides high-speed communication capability, enabling users to run tightly coupled applications that achieve high levels of scaling. The GigE control network allows for system diagnostic and machine-management functions without interrupting application communications. Abe will offer 170 terabytes of storage in a Lustre file system, providing 7.5 GB/s peak I/O performance.

T3 is the successor to T2, a Dell cluster that has been used extensively by NCSA's industrial partners since it came online late in 2004. Now in use, T3 is a Dell blade system with 1,040 dual-core 2.66 GHz Intel processors, an InfiniBand interconnect, 4.1 terabytes of total memory, and a 20 terabyte Lustre filesystem. The peak performance of T3 is 22.1 teraflops.

ONE TREE
TO BIND THEM

by J. William Bell

A comprehensive history of bumble bee evolution—completed with the

help of NCSA—cuts against the conventional wisdom on how the insects'

color patterns and social behavior developed.

You've probably been admonished not to miss the forest for the trees. But for those building phylogenies, trees recover the ancient evolutionary story of how organisms are related to one another. According to entomologist Sydney Cameron, adding branches (that is, species) to these trees increases the value of these systematic arrangements of species, based on their genetic makeup. "The inferences are much more powerful with more species included. Having all the species of an entire genus—or virtually all—is phenomenal," says Cameron, a professor at the University of Illinois at Urbana-Champaign.

In an upcoming issue of *Biological Journal of the Linnean Society*, Cameron and her collaborators reveal the first comprehensive phylogeny of the genus *Bombus*, also known as the bumble bee. Collaborators include graduate student Heather Hines and Natural History Museum of London entomologist Paul Williams. Their tree includes nearly 90 percent of all 250 species represented in the genus worldwide.

"This is one of the few, if any, insect DNA phylogenies that includes almost every species in a large genus. A decade ago, it would have been a pipe dream," she says. "The interplay between [gene] sequencing technology, its cost, the algorithms involved, and the computational power—these have all come together recently to allow such a large-scale study."

The computational power, and the expertise to harness it, came from NCSA. The team used the center's IBM p690 system to estimate the phylogeny of 220 species covering about 35 million years of evolution by comparing DNA sequences from five genes for each species. NCSA staff, meanwhile, assisted with debugging and writing code to manage the simulations.

The bumble bee tree is only now hitting the presses, but it has already stimulated new research on the timing of the bees' genetic divergence and ancient movements across continents. It has also led to new research into the evolution of their color patterns, which converge across distantly related species as a means of protective mimicry.
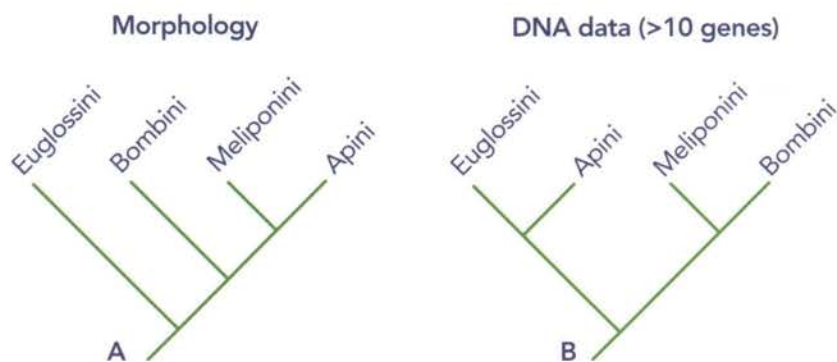
Bees as social climbers

Less than three percent of the approximately 20,000 known species of bees are classified as highly social; you see them on the quintessential Discovery Channel documentaries, with their strict division of labor centered around an egg-laying queen. The bumble bees of interest to the Cameron team are less hierarchical. They're classified as intermediately social, and the hives are less focused on the queen than are honey bees. But they're not the kind of bee to just lay their eggs and leave, never to see their offspring once they've hatched, as solitary bees do. This position in the middle of the social ladder between highly social and solitary bees makes them a particularly useful target of study.

Traditionally, highly social behavior was thought to be the pinnacle of evolutionary development. This implied a solitary common ancestor that divided into related but distinct species. Over millions of generations, some species maintained their solitary ways, while new species slowly accrued increasingly complex social behavior. Thus, the highly social bees were seen as the apex, having passed as earlier species through previous grades of sociality.

The Cameron team's studies of bee relationships, including the comprehensive bumble bee tree, contradict that thinking. The bumble bee branch does not sprout from the trunk between the solitary

A *Bombus lucourm* from Uppsala, Sweden. Cameron and her collaborators created the first comprehensive phylogeny of the genus, also known as the bumble bee.

Morphology

Euglossini   Bombini   Meliponini   Apini

A

DNA data (>10 genes)

Euglossini   Apini   Meliponini   Bombini

B

Relationships between four tribes of corbiculate bees (honey bees, stingless bees, orchid bees, and bumble bees). This illustration shows the differences between phylogenies based on morphological data and those based on DNA data. The bumble bee branch (Bombini) does not attach between the solitary Euglossini, near the base, and highly social honey bees (Apini) and stingless bees (Meliponi) at the top, as previous theories looking at morphology would imply. Instead, according to multiple independent DNA studies, the highly social stingless bees (Meliponini) are most closely related to the intermediately social bumble bees (Bombini) and relatively distantly related to the highly social honey bees (Apini).

bees near the base and highly social honey bees and stingless bees together at the crown, as previous theories looking at morphology would imply. Instead, according to the DNA studies, the highly social stingless bees are most closely related to the bumble bees and least closely related to their highly social cousins, the honey bees.

"This gives us a completely different perception of the evolution of social behavior," says Cameron. "Until recently, studies of the bees' physical traits and their behaviors showed us a progressive series of steps toward a pinnacle of social behavior. Now we see it's not a single pinnacle at all. There's no single origin of highly social behavior. Maybe it can happen many times. It certainly has happened twice here." Moreover, in a paper currently in press in the *Proceedings of the National Academy of Sciences of the USA*, Hines, Cameron, and colleagues have concluded that there is a similar pattern of social evolution among wasps, showing that sociality also evolved twice independently in stinging wasps.

From these insights, the team and others like them are beginning to explore what makes something social. Environment, given the bees' various locations around the world and how that relates to their social behavior, is thought to be significant. But entirely different pressures also likely play a role. Whole-genus and higher-level phylogenetic trees like those developed by the Cameron team are key to accurately teasing out gains and losses of traits, including different traits of social behavior and what might have caused them.

Upcoming work at NCSA will combine into a single large analysis of comprehensive data from phylogenies of several genera—bumble, honey, orchid, and stingless bees. "We can see interesting things in smaller tests. But as these phylogenies become larger, the evidence is that much more accurate," says Cameron.

A clever mimic

Bumble bees are also of interest to entomologists because of the striking degree to which they exhibit Müllerian mimicry.

With some forms of mimicry, harmless species adapt to take on the physical traits of a species that is noxious. A tasty plant may come to look like a poisonous plant over time, for example. In Müllerian mimicry, harmful species—different species of bumble bees

with their stingers—look like one another. This increases the pool of bees marked as dangerous to their predators and speeds the rate at which those predators figure out the threat.

It might be reasonable to assume that insects that share a color pattern are closely related genetically; that they all came from a common ancestor from which other traits diverged, but that they kept their similar markings. The comprehensive phylogeny of bumble bees shows conclusively that those with near identical striping patterns are usually distant relatives.

A particular black and orange pattern, for instance, "crops up all over the tree [in distant relatives]. But when you look at a distribution map of where all the bees with that black and orange color pattern are found, you see that they cluster in a geographic region," according to Cameron. Entomologists are looking at other selection pressures, besides protection from predation, that might influence their coloring. For example, differences in climate and other environmental factors might be involved.

Their Bayesian best

In Cameron's studies, millions of possible trees with different branching patterns are constructed during the computer analyses of species relationships, representing different possible ways all the species in a study could be related. Out of these millions of possible trees, the one thought to represent the most accurate tree based on a variety of criteria is chosen.

Until relatively recently, the data used in phylogenetic studies were mostly morphological characters, and the datasets were small. In Cameron's studies, the data include nucleotides from the DNA sequences her lab generates. The datasets are much larger, comprising thousands of nucleotides, compared to the few hundred characters generated by studies of morphology.

Today, computational power—like that available at NCSA—and advanced algorithms allow researchers to look at meaningfully sized collections of species, such as the whole bumble bee genus. Many researchers, like Cameron's team, now use model-based approaches, which incorporate specific theories concerning how nucleotides of DNA change over time.

For example, based on statistical analyses of DNA sequences, researchers know that nucleotide changes in some regions and at some sites of a sequence change much more rapidly than others, and that different genes may change at different rates. They can account for these variations by using specific models with the appropriate model-based methods of phylogenetic inference.

Powerful computers and algorithms also allow the Cameron team to use Bayesian analysis, in which observations of the emerging most-likely tree are used to infer the probability that the next proposed tree in the analysis will be the best tree. With these inferences, the algorithm can concentrate on the most likely candidates and reduce the number trees that are considered. These Bayesian calculations are taxing computationally.

With Bayesian algorithms and hundreds of hours of NCSA computing time required for a large analysis, there were challenges for Cameron's team to overcome. NCSA's Sudhakar Pamidighantam stepped in to help.

Pamidighantam worked with the team over the course of months to get the Bayesian algorithms to work properly and effectively on NCSA's systems. He also wrote scripts to make the code run in parallel across many processors at a time, allowing multiple analyses. "We would not have this comprehensive tree today without him," says Cameron.

This research is supported by the National Science Foundation's Division of Environmental Biology and the U.S. Department of Agriculture.

Access Online: http://access.ncsa.uiuc.edu/Stories/bombus

Team members

Sydney A. Cameron
Bjorn Cederberg
Heather M. Hines
Pierre Rasmont
Paul H. Williams

Worldwide phylogeny, as developed by the Cameron team. The phylogeny is estimated from mixed model Bayesian analyses of combined sequences from five gene fragments (16S, opsin, ArgK, Ef-1 alpha, and PEPCK). Subgeneric clades are individually color-coded and labeled with the subgeneric name. Values above branches are Bayesian posterior probabilities; values below branches are parsimony bootstrap values. Alternative resolution from parsimony analysis is shown as dotted lines. Outgroup branches (in grey) are represented by dashed lines and have been shortened for visual purposes. Graphic assistance with tree design came from Ben Grosser, Director, Imaging Technology Group, Beckman Institute.
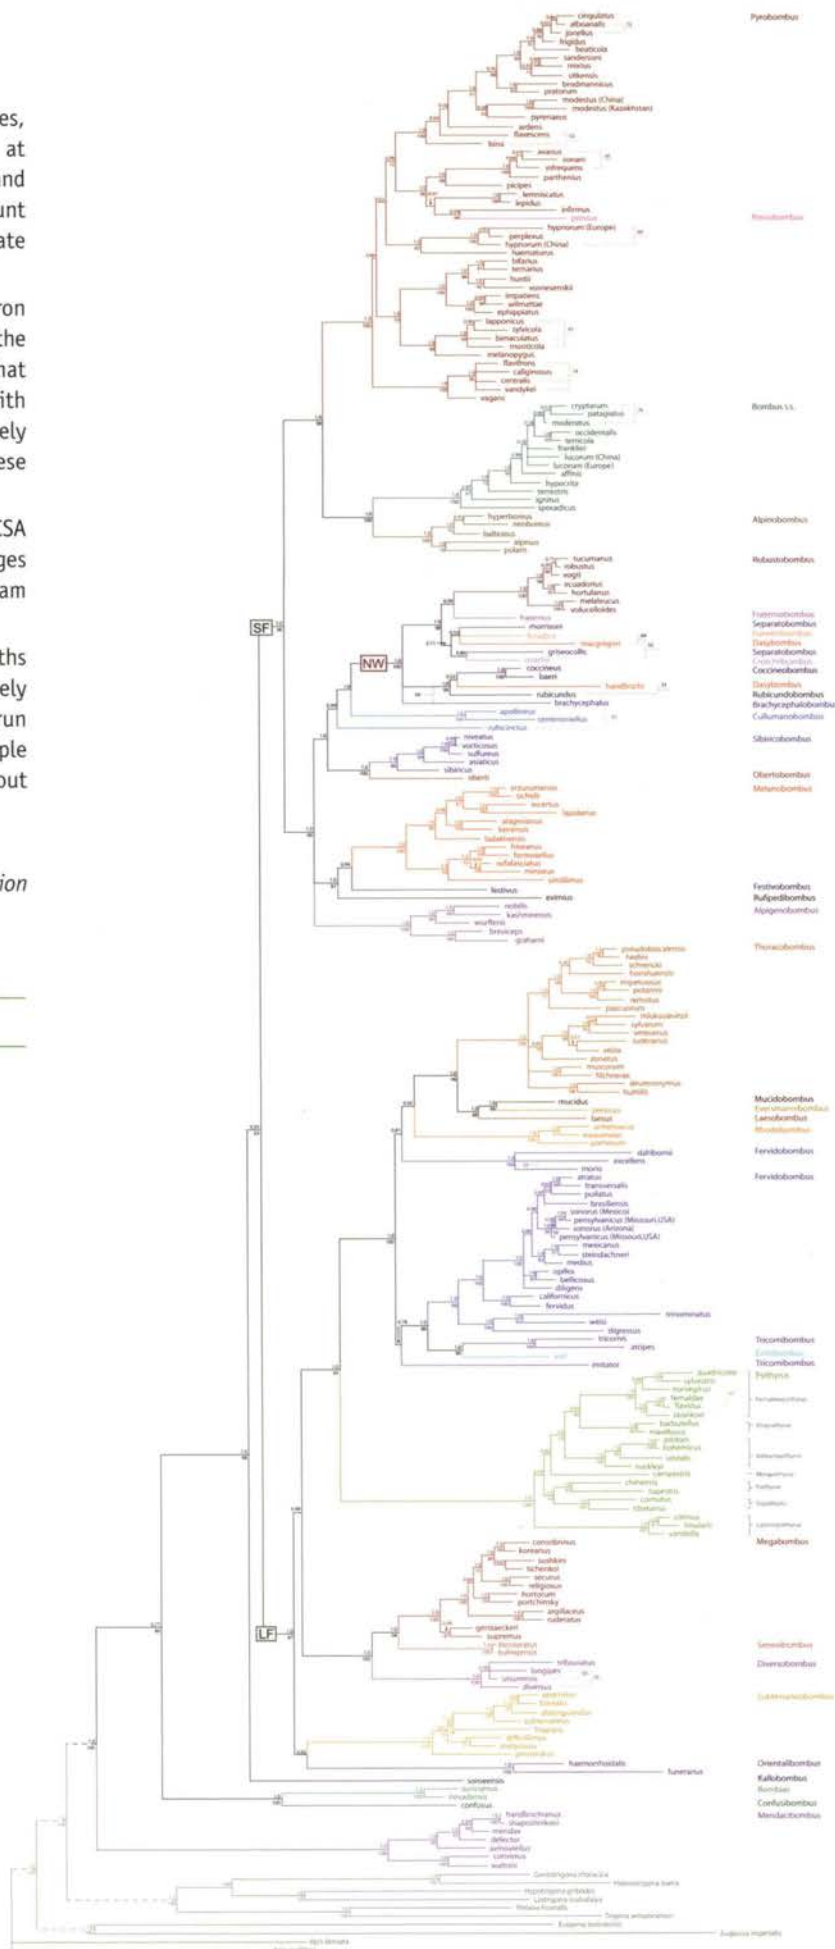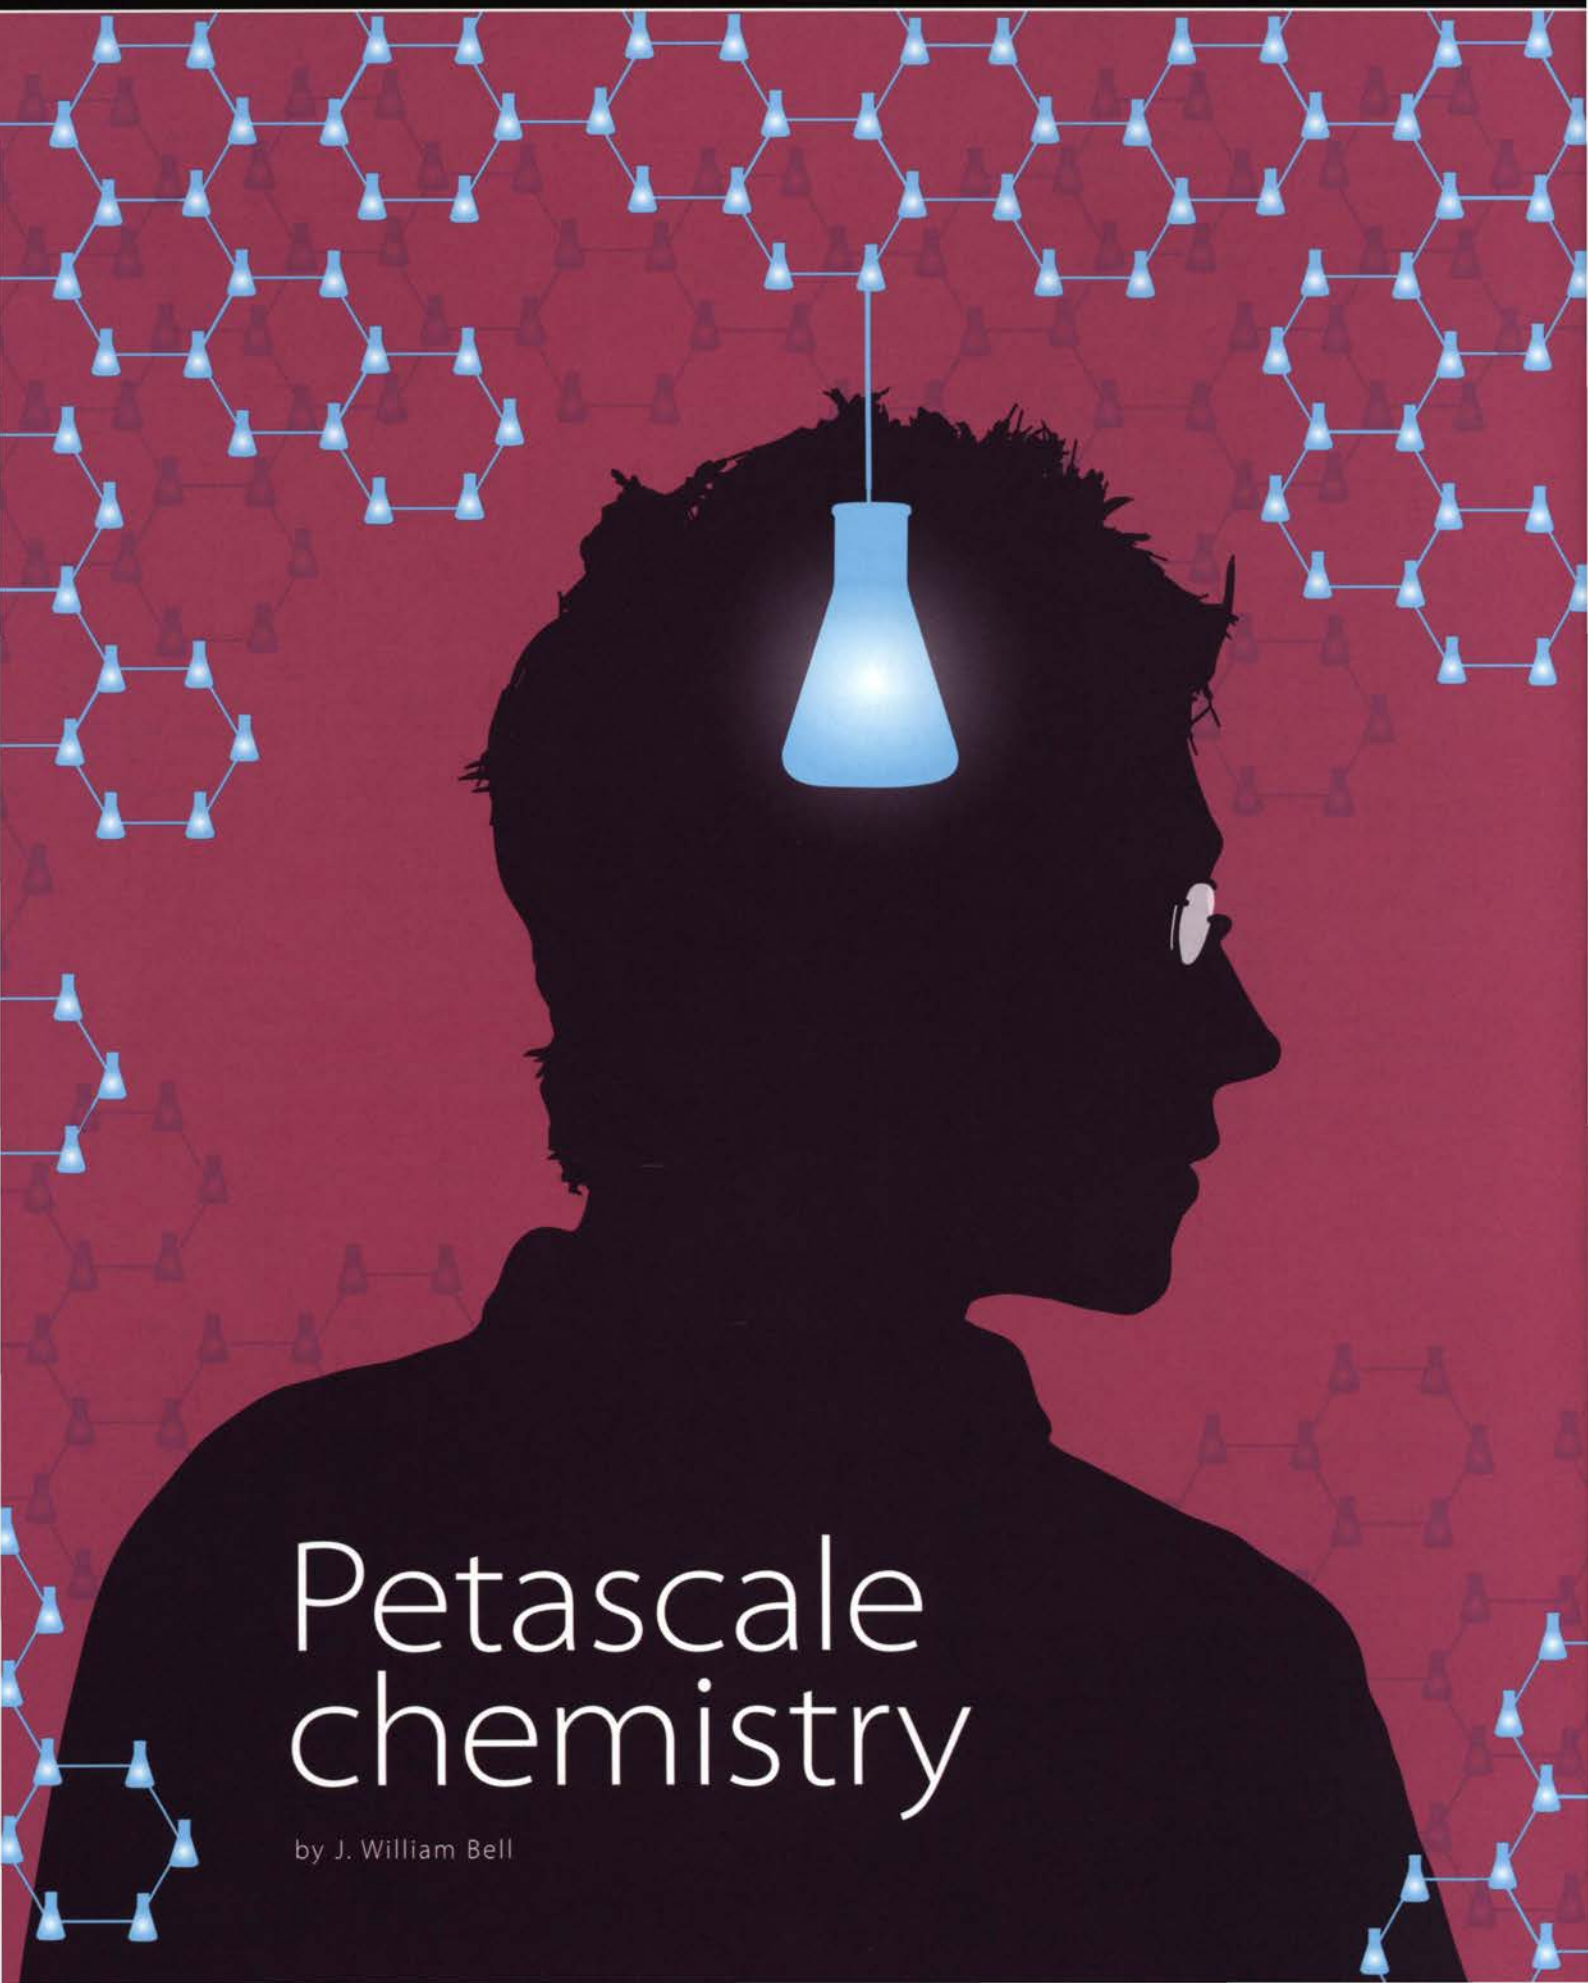
Petascale chemistry

by J. William Bell
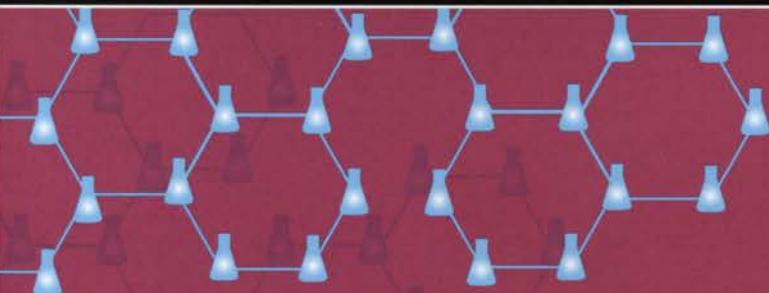
Petascale computing is rapidly approaching. The Department of Energy's Office of Science plans to install a computing system with a peak performance of one petaflop at Oak Ridge National Laboratory in 2009, and the National Science Foundation plans to install a computing system with a sustained performance of one petaflop by 2011. Potential breakthroughs from petascale computing abound—the NSF solicitation for the sustained petaflop system listed more than 30 science and engineering problems expected to benefit.

"Chemistry is one area of science that needs petascale computing to improve current molecular models and tackle more complex molecular phenomena, including nanoscale systems. Prediction of the structures, energetics, and reactivity of molecules is computationally intensive, and chemists have long been in the vanguard of high-performance computing," says Thom Dunning, NCSA's director. But, tapping the power of petascale computers will require the development of a new generation of scalable, parallel chemistry codes as petascale computers will obtain their power by harnessing hundreds of thousands of processors, not the thousands present in today's high-end systems.

Chemists tackled a similar problem in the 1990s, when "massively" parallel computers had a few hundred processors. This effort led to the identification of a few key concepts, which will be critical as applications scale to sustained petascale performance, says Robert J. Harrison, who leads the Computational Chemical Sciences Group at Oak Ridge National Laboratory and is a member of the chemistry faculty at the University of Tennessee, Knoxville.

One of these concepts is modularity. Codes being built or revised to attain petascale performance will need to accommodate new models and algorithms. The cost of development—in terms of person hours and dollars spent—is so high that tomorrow's applications scientists will have to be able to leverage investments that have already been made. This is more easily done if codes are modular.

As an example of modularity, Harrison points to NWChem, a chemistry application for parallel computers initially developed in the early 1990s at the Department of Energy's Environmental Molecular Sciences Laboratory and widely used by NCSA researchers. Harrison was NWChem's chief architect. NCSA's Dunning instigated and oversaw the project.

"NWChem creates a path that allows smaller groups to implement codes more quickly," Harrison says.

"It's important to think of NWChem as a framework for chemical computation rather than as just an application," Dunning says. "That's an ideal that's been with NWChem from the outset and that similar projects—now and going forward—will benefit from."

But modularity does not solve the problem of scaling to hundreds of thousands of processors. To address this issue, new algorithms will be needed. Chemistry codes make use of standard mathematical algorithms (for example, BLAS, SCALAPACK, and PEIGS) whenever possible, so they will be able to leverage investments made in improving the scalability of these packages. However, other algorithms are specific to chemistry and will require investments in chemistry-specific teams of mathematicians and computer scientists to advance the scalability of these algorithms. In some cases, entirely new algorithms may be required.

Another concept for applications developers to keep in mind as they conceive of petascale codes is object-based virtualization. Virtualization refers to representing multiple resources or processes as single entities. That might mean automating the method by which applications are spread out across the hundreds of thousands of processor cores so that users don't have to address the problem directly, or it might mean hiding the true complexity behind how data is called up from memory during a calculation. Here, Harrison calls out Charm++.

"We need to place a greater emphasis on changing programming models. The complexity of a petascale machine is much greater. [Those using petascale applications] will need to be working at much higher levels of abstraction than they're allowed to today," Harrison says.

Charm++ is a parallel programming system that can be used with a variety of high-performance codes in disciplines like computational fluid dynamics, biophysics, and cosmology. It runs on some of the largest systems at NCSA and around the world. Charm++ supports virtualization activities like dynamic load balancing, which migrates calculations among processors so that they work together more efficiently. It also provides a means of hiding memory latency.

With Charm++, "a processor is allocated to an object only when a message for the object is received. This means when a processor is waiting for a message, another object may execute on it. It also means that a single processor may wait for any number of distinct messages and will be awakened when any of these messages arrives. Thus it's an effective way of scheduling a processor in the presence of potentially large latencies," explains Laxmikant Kale, a computer science professor at the University of Illinois who leads the development of Charm++.

The performance advantages to be gained by using Charm++ in chemistry codes can be substantial. The LeanCP (also known as OpenAtom) project, a collaboration with Glenn Martyna of IBM's TJ Watson Research Laboratory and Mark Tuckerman of New York University, is devoted to extreme scaling for Car-Parrinello *ab initio* molecular dynamics simulations. LeanCP runs on various NCSA computers. It efficiently scales 256-molecule water simulations to all 40,960 Blue Gene/L processors of the machine at the TJ Watson lab. The team plans extensions for petascale computers.

Kale works closely with NCSA in planning for the deployment and use of emerging petascale systems and applications. His research team is also developing BigSim, a system simulator and emulator that allows scientists to develop, debug, and predict the performance of applications on petascale machines before the machines are available. "This way applications can be ready when the machine first becomes operational," Kale says.
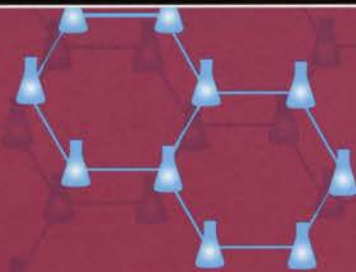
THE
STEPS
BETWEEN

by Kathleen Ricker

Below the nanometer level, flaws in the

surfaces of crystalline silicon yield unexplored

worlds ripe for colonization with data.

Silicon, in various forms, is the second most abundant element on Earth. It's found mostly as part of a compound in the form of sand and rock. To achieve the high-purity semiconductor form of silicon on which the global electronics industry is based, silicon rods are immersed in a silicon-containing compound gas at extremely high temperatures. The resultant chemical reaction deposits silicon from the gas onto the rods, producing a crystal with the same atomic structure as a diamond.

It's this form of silicon in which Cristian Ciobanu, an assistant professor of mechanical engineering and materials science at the Colorado School of Mines, is most interested. What happens, he wants to know, when you make cuts in the surface of a solid chunk of silicon that do not follow the "grain" of the solid? One might first imagine a flat edge along the cut, which is what would happen in wood or in steel at a macroscopic scale. But that's not what it looks like at the atomic level, says Ciobanu, because when the atoms on the surface are disrupted by the (mis)cut, they will try to rearrange themselves to form as stable a structure as possible. Ciobanu wants to discover what kinds of structures those might be.

He is interested in the problem of how atoms on a silicon surface reconfigure their structure primarily because it's interesting and difficult, he says, but emphasizes that it is a problem that also has practical applications. On some silicon surfaces, he explains, miscuts form steps and terraces that are very straight and evenly spaced over large areas of the silicon substrate at distances of around one nanometer. The steps can then be used as tracks for writing and storing information at a factor a million times denser than the storage capacity of a conventional CD-ROM.

A wafer of pure silicon cut from a crystal and ready for processing. Cristian Ciobanu, a materials scientist at the Colorado School of Mines, is investigating how silicon atoms reorganize themselves to form new surfaces that could, among other things, be used to store larger amounts of data.

A question of balance

Apples stacked in a supermarket display will hold the shape if stacked correctly, says Ciobanu, thanks to gravity. If half the apples in the front layer are removed carefully, the remaining apples, assisted by gravity, will reorganize to form a flat terrace followed by a step to another flat layer just behind it.

But the apples require close proximity to bond with each other or interact with each other. In contrast, when silicon atoms rearrange themselves, what happens to them depends not only on their immediate neighbors but also on other nonadjacent silicon atoms both within their layer and some layers deep. "A lot of atoms have to be happy," says Ciobanu, "not just the ones at the surface."

To achieve a stable structure, silicon atoms on the surface tend to bond so that each one has four neighbors—the number of neighbors they have in bulk (subsurface layers). Theoretically, one could look for surface and step structures by trying to increase the number of neighbors of each atom—or in the language of chemistry, by reducing the number of dangling bonds (broken covalent bonds) at the surface. But lowering the number of dangling bonds per surface area too far would result in lengthening and distorting the new surface bonds and stressing the surface.

Thus, the stable surface structure has to be achieved through a balance between reducing the density of dangling bonds and the stress generated by this reduction process. The physical quantity that incorporates both stress and chemical bonding is the surface energy, or the excess energy that the atoms have, per unit area, with respect to their bulk form. When the bulk is disrupted, the reconstruction of the atomic structure of the surface is determined by lowering this surface energy. Disrupting a surface means creating a step whose atomic structure can be found by lowering the step formation energy (excess energy with respect to the surface). Furthermore, disrupting a step creates a kink in it, the structure of which can be found by lowering the kink formation energy. Figure 1 illustrates the surface, step, and kink in relationship to one another.

Recently Ciobanu used NCSA supercomputing resources to study the atomic structure of surfaces; he is now studying the structure of steps on surfaces. In the latter case, the problem is that even when the surface structure is known, the positions of the steps on that surface—specifically, the location at which the step disrupts the atomic pattern (the reconstruction) on the terrace—are not known beforehand. Neither is known, *a priori*, how wide the step is,

Top view, along [114]

Upper terrace   Step region   Lower terrace

Side view, along [-110]

Parent A   Child C   Parent B

Padding zone

Support

Supercell

Figure 2. (a) Step region (shaded) for which the number of atoms, their positions, and the location of the step edge are to be determined. The step region, which as shown corresponds to a down step oriented along the [-110] direction, is surrounded by reconstructed Si(114) terraces with the surface unit cell marked by the dashed rectangles. **(b)** The genetic operation (cross-over) used by the optimization algorithm to search the configuration space of the step region. The step formation energy calculations involve relaxing a padding zone (in addition to the step region).

Figure 1. The russian-doll sequence of surface science. A surface or terrace is obtained when the bulk is cut with a plane (left). If the surface is disturbed, line defects or steps (middle) can often be seen. If a step is deviated from its straight direction, then the formation of kinks in the step (right) can be observed. The atomic structure of surfaces, steps, and kinks is in many cases unknown, requiring the use of complex search algorithms.
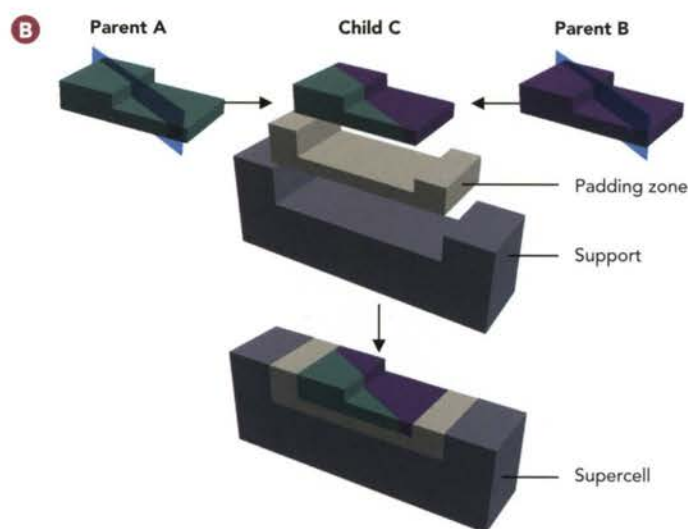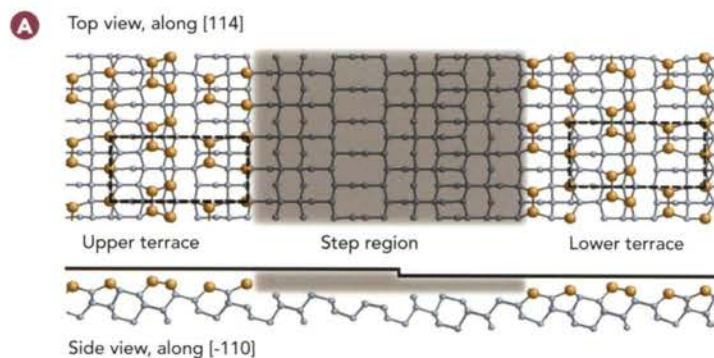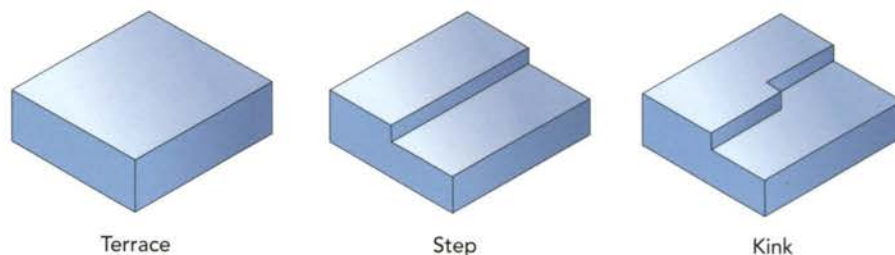
Terrace   Step   Kink

or exactly how many atoms the step region has. To deal with these complex structural issues, Ciobanu and his student, Ryan Briggs, have adapted a genetic algorithm for solving the atomic structure of steps on silicon surfaces, while at the same time finding the width of the step and its location (Figure 2).

Evolution of a step structure

Ciobanu is using NCSA's Tungsten cluster to run genetic algorithms to determine the way the silicon atoms are arranged at the steps that form between known surface structures (terraces). The genetic algorithm optimization begins with a given number of step structures in which the step-zone atoms are positioned randomly. This "gene pool," or Generation Zero, is then continuously optimized by cross-over operations (such as that shown in Figure 2b) in which two "parent" structures are picked arbitrarily and mixed to create a new structure, the "child." If the child is "fit," that is, if it has a sufficiently low surface energy without being identical to an existing structure, then it is included in the pool. With the addition of a new child, the worst structure (corresponding to the highest step formation energy) is discarded from the pool. Repeated application of the cross-over operation leads to a decrease in the formation energies (or to an increased fitness) of all the step structures in the pool, and as the genetic algorithm progresses, both the best and the average formation energies also decrease (Figure 3).

Interestingly, if the selected step zone (Figure 2a) is wide enough, the step can negotiate its best location with respect to the terraces around it. At the same time, the number of atoms in the steps and their location are being optimized, which is an important and versatile improvement in an otherwise well-known algorithm. As a result of this algorithm, Ciobanu has succeeded in accumulating a large number of simulated step structures that could prove useful in the future for applications such as increased data storage capacity, and applications used by experimental groups such as those at the University of Illinois and the Naval Research Laboratory who work on the Si(114) surfaces (the terraces chosen by Ciobanu for his study). "It's not too far-fetched," says Ciobanu, "to speculate that one day some of our simulations could be confirmed by experiments or could be useful in interpreting experimental data."

Ciobanu's calculations require a great deal of computing power because in order to find the best step structures he must evolve the algorithm through thousands of cross-over operations, only discarding unfit structures after their step formation energy has been computed. CSM, one of the older schools in Colorado, is a public research university devoted to engineering and applied science. In the past, the school has not had to provide extensive computational facilities for research. However, that could soon change, stimulated by the successful usage of national high-end computing facilities such as those provided by NCSA and other supercomputing centers. CSM is presently making investments to acquire high-performance computational facilities on campus, to keep up with the increasingly demanding research done by its faculty.

"By providing rapid and substantial allocations of computing time and reliable support for Tungsten," says Ciobanu, "NCSA has been of tremendous help in my research."

This research is supported by the Colorado School of Mines.
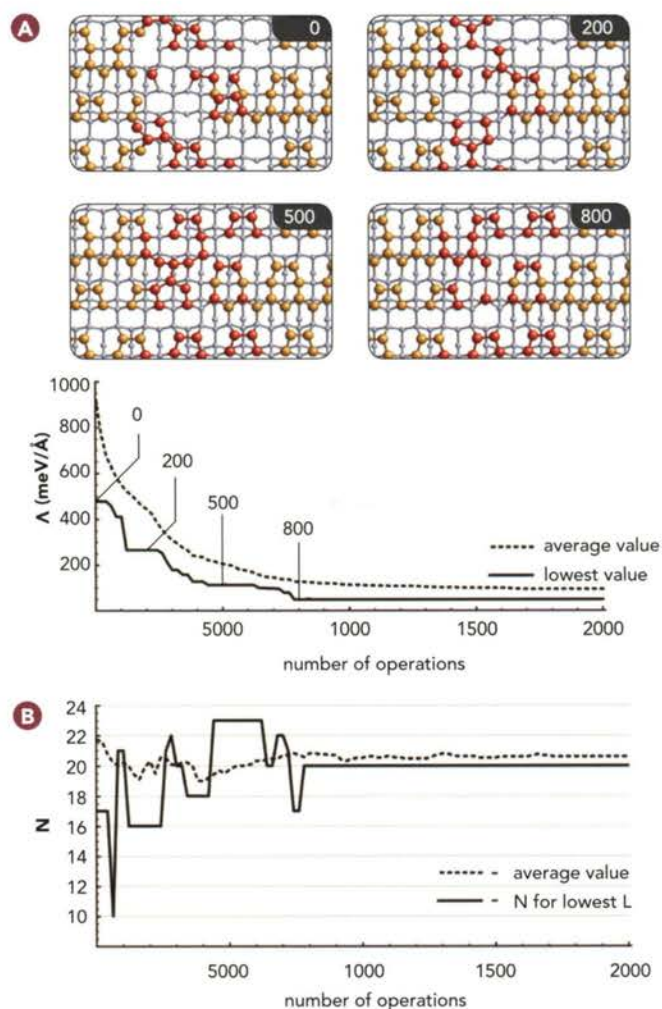
Figure 3. Finding the structure of steps oriented along the [22-1] direction on the Si(114) surface. **(a)** Formation energy of the lowest-energy structure (solid line) and averaged across the pool (dashed line) during the genetic evolution. The lowest energy structure is shown after 0, 200, 500, and 800 cross-over operations. The atoms subjected to optimization are shown as red spheres in the insets, the terrace atoms are yellow, and the bulk atoms are shown as small gray spheres. **(b)** Evolution of the average number of step-zone atoms across the pool (dashed line) and of the atom number corresponding to lowest-energy member (solid line).

A delicate balance

by Trish Barker

In order to effectively harness the energy created by networks of nitrogen atoms, scientists need to add stability to these systems without sacrificing too much energy output.

An old saying insists that "where there's smoke there's fire." And basic chemistry asserts that where there's fire, there must be oxygen. Everything from car engines to rockets relies on the presence of oxygen to spark the release of energy. So what happens in environments—from deep space to the deep ocean—where oxygen is scarce or entirely absent?

While carbon fuels are powerless in oxygen-starved environments, nitrogen fuels face no such restriction. A network in which nitrogen atoms are joined by single bonds readily decomposes in an exothermic reaction, releasing a burst of energy. Most of these compounds are highly unstable, however, making it impossible to harness their energy to create viable high energy density materials (HEDMs).

"If it's unstable, you have one application, and that's to create an explosion," says Douglas Strout, an associate professor of chemistry at Alabama State University.

Strout's research team, which includes undergraduate students gaining valuable hands-on experience, uses computational simulation in an effort to identify nitrogen molecules that provide a substantial energy payout but that are stable enough to be practical power sources. The team recently used NCSA's IBM p690 system, called Copper, to simulate the decomposition of various molecules in which nitrogen atoms were replaced with potentially stabilizing atoms, such as carbon and halogens (chlorine and fluorine). Their simulations identified promising compounds for experimental study, and their results were published in *The Journal of Physical Chemistry A* in July 2006.

Balancing energy with stability

Any molecule composed solely of nitrogen (N_x) is potentially a rich source of energy. However, theoretical modeling of many N_x molecules has shown that they dissociate too easily and are therefore too unstable to be practical fuel sources. So Strout and other researchers have turned their efforts toward finding molecules that combine energy-rich nitrogen with other stabilizing atoms.

"If you trade away nitrogen content, you're trading energy. Hopefully you can find high energy content and high stability," Strout says.

In previous computational simulations, Strout and his collaborators examined a variety of all-nitrogen molecules—N_{12}, N_{14}, N_{16}, N_{18}, N_{24}, N_{30}, and N_{36}. While the simulations provided clues about which structures add stability to nitrogen molecules (triangles seem to help, and cylindrical shapes are better than spheres), the molecules still weren't stable enough to be practical HEDMs.

So with the most recent calculations, Strout started with an all-nitrogen molecule, N_{12}, and altered it in several ways that might increase its stability (Figure 1 shows the structure). The six variations Strout and his students examined were $N_{14}H_4^{2+}$ (N_{12} plus two nitrogen atoms), $N_{14}F_4^{2+}$ (fluorine), $N_{14}Cl_4^{2+}$ (chlorine), $N_{12}C_2H_4$ (carbon), $N_{12}C_2F_4$ (carbon and fluorine), and $N_{12}C_2Cl_4$ (carbon and chlorine). They calculated the stability and energy-output of these six molecules using two methods (Hartree-Fock and the more accurate fourth-order perturbation theory) and NCSA's computational resources.

"We're talking about molecules that don't even exist yet in the laboratory," Strout says, so theoretical modeling saves the time, expense, and hazard of synthesizing and analyzing these substances in the lab.

Calculating energies

All of the molecules Strout's group analyzed have a similar cage structure of nitrogen-to-nitrogen bonds at their heart. The team calculated the energies for the breaking of each N-to-N and C-to-N bond for each molecule; the bond that can be broken with the lowest energy is the weak link where the explosion of energy from the molecule is likely to begin. (Table 1 shows the calculated energies).

Bond-breaking energies

Method	Bond	$N_{14}H_4{}^{2+}$	$N_{14}F_4{}^{2+}$	$N_{14}Cl_4{}^{2+}$
Hartree-Fock	NN1	+3.7	+0.1	-5.1
	NN2	+9.2	+4.1	+9.6
	NN3	+3.0	-1.3	+5.4
	NN4	+42.7	+38.5	+42.3
Fourth-order perturbation theory	NN1	+30.6	+28.5	+23.6
	NN2	+36.1	+30.8	+36.2
	NN3	+34.1	+30.3	+34.5
	NN4	+65.4	+62.4	+64.0

Method	Bond	$N_{12}C_2H_4$	$N_{12}C_2F_4$	$N_{12}C_2Cl_4$
Hartree-Fock	CN	+57.0	+56.0	+42.7
	NN2	+21.9	+19.7	+18.7
	NN3	+8.3	+5.9	+8.4
	NN4	+53.8	+52.7	+53.0
Fourth-order perturbation theory	CN	+81.5	+81.0	+69.9
	NN2	+48.7	+46.2	+45.5
	NN3	+37.1	+35.1	+36.7
	NN4	+77.9	+77.8	+76.7

Table 1. All calculations carried out with cc-pVDZ basis set; energies in kcal/mol

Figure 1. Structure of $N_{14}H_4{}^{2+}$ and other simulated molecules, showing bonds and the locations of hydrogen/halogen substitutions. Blue atoms are nitrogen for all the simulated molecules.

For the hydrogen-nitrogen compound and the two halogen-containing molecules, calculations using the fourth-order perturbation theory pointed to the NN1 bonds as the weakest, with the lowest dissociation energies. The bond-breaking energy for $N_{14}H_4{}^{2+}$ (30.6 kcal/mol) indicates that it might be stable enough to be a high-energy density material, while the even lower energies found for the halogen ions might make them too unstable to be practical.

Next, Strout's team considered whether adding carbon atoms to these molecules might create greater stability without the loss of too much energy. For the three carbonated ions, the researchers found that the weakest bond was now the NN3; in all three cases, the bond-breaking energy for the NN3 link was greater than 35 kcal/mol, meaning the molecules should have the necessary stability for high-energy density materials.

Now that Strout had found several molecules with sufficient stability, the next step was to look at their potential energy output. Calculations showed that the substance that provided the best energy return per unit of mass (2.2 kcal/gram) was the $N_{12}C_2H_4$; the halogens in the other molecules added mass without providing more energy, or greater stability.

While Strout's research has provided valuable insight into how networks of nitrogen atoms can be combined with other stabilizing substances, he cautioned that there is still more work to be done.
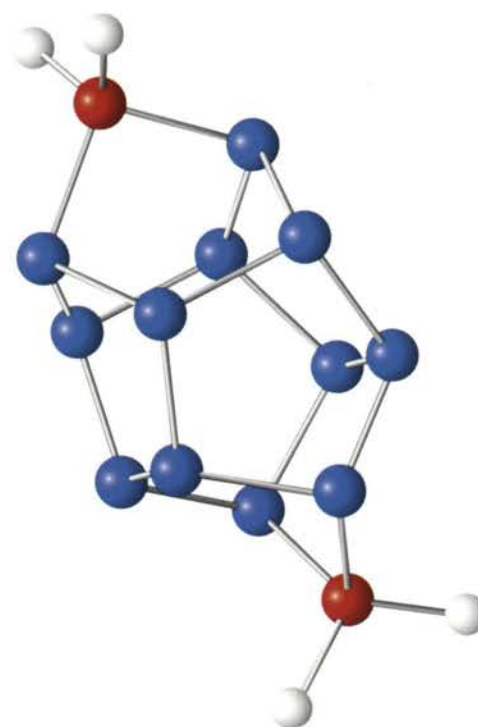
"Don't expect to be putting these molecules in your gas tanks any time soon," he said. "The applications will be less commonplace and more obscure than that."

Now that promising molecular structures have been identified through computational simulation, experimental chemists can direct their efforts toward synthesizing and testing the most promising potential HEDMs.

Access Online: http://access.ncsa.uiuc.edu/Stories/DelicateBalance

Team members

Roshawnda Cottrell

Jacqueline Jones

Ami Gilchrist

Danielle Shields

Douglas Strout

Creating the next generation of researchers

by Barbara Jewett

Cybereducation helps bridge the gap between research and education. Towards that end, NCSA's Cybereducation Division conducts a wide variety of activities and programs, including activities that disseminate advanced technologies to new and established communities, encouraging them to adopt innovative computational technologies as learning tools in pre-college settings, colleges, universities, and the workplace. The division is currently working on undergraduate and K-12 projects in the areas of astronomy and cosmology, biology, chemistry, humanities, public health, and the geosciences.

Plug into chemistry

Funded by a $5 million dollar grant from the National Science Foundation, NCSA, the College of Medicine and the Department of Chemistry at the University of Illinois at Urbana-Champaign, A-C Central Community Unit School District 262 in Chandlerville, Illinois, and the Illinois Regional Office of Education #38 in Lincoln are joining forces to develop a statewide Institute for Chemistry Literacy and Computational Science. The intensive, multi-year summer institute is providing professional development for 120 rural high school chemistry teachers who will become intellectual teacher-leaders and accomplished practitioners for the 21st century. This project uses the tools of computational chemistry to illuminate four large themes that among them comprise all of modern chemistry: molecular structures, covalent interactions, non-covalent interactions, and kinetics and thermodynamics. Within each of these themes, computational tools for the task of enriching high school chemistry will be made accessible to teachers.

The first group of 60 teachers will begin their institute participation in June 2007. In addition to the summer institute, teachers will receive extensive online communication during the following school year.

The goals for the institute are to strengthen teachers' understanding of chemistry in the context of the most up-to-date research and practice; to increase teachers' comfort with and use of computational and visualization tools in the classroom; and to train teacher-leaders who can support their colleagues and advocate for excellence in science education.

NCSA has extensive experience with communication and collaboration technologies and will aid participating schools in effectively using tools like the Access Grid. The virtual community aspects of the institute are particularly important for rural teachers at small schools, who are often the only science (and math) teacher in their district.

"We intend to make teachers competent and confident in using computational resources to both excite students and prepare them for the workforce," says Edee Norman Wiziecki, coordinator of education programs for NCSA and one of the co-principal investigators for the project.

The next generation

Creating the next generation of researchers in science and engineering is a focus of the cybereducation staff. Wiziecki says they are using summer workshops to show colleges and universities how to incorporate computational biology and computational chemistry into their courses. And they are working with grades K-12, as well, with the creation of programs such as GEMS—Girls Engaged in Math and Science.

The GEMS program for girls in grades 6-8 has been active in the local community since 1994. GEMS was created as a way to encourage girls to consider mathematics-oriented and science-oriented careers, to gain confidence in doing mathematics and science, and to take advanced level mathematics and science courses in high school. The GEMS program has provided participants with visits to research laboratories, guest speakers, after-school activities, a network of female role models, and opportunities to work side-by-side with a mentor in laboratories across the community.

One of the very first GEMS activities was creating a website for the program. The site was created by a team of eighth-grade girls using NCSA Mosiac, the world's first graphical web browser. Over the years, as the GEMS program has grown, so have its activities, including the use of emerging technologies and communication tools. The Girls on the Grid component of GEMS utilizes Access Grid technology to link girls to peers, scientists, and leading women in science and mathematics worldwide.

Productive connections

by Barbara Jewett

This mural is one way Chicago's Mexican immigrants beautify their neighborhoods and convey beliefs.

A new study finds that recent Mexican immigrants in Chicago possess a wealth of artistic, cultural, and networking assets and that those assets contribute to the social, cultural, and economic well being of many neighborhoods, organizations, and institutions in the Chicago area. In fact, the artistic and cultural activities of Mexican immigrants who arrived in Chicago after 1994 stimulate the local economy, especially in the music industry and service sectors.

So say the authors of the "Creative Networks: Mexican Immigrant Assets in Chicago" study, which is a joint project of The (Chicago) Field Museum's Center for Cultural Understanding and Change and the Science of Networks in Communities (SONIC) research group at NCSA. The Rockefeller Foundation funded the yearlong study through its Creativity and Culture Program.

Researchers began by examining selected suburbs and city neighborhoods. Later, the researchers incorporated comments from Mexican immigrants throughout the city. The aims of the study were to identify the cultural, artistic, and networking practices and capacities of post-1994 migrants from Mexico; to analyze how these practices act to buffer challenges or obstacles immigrants face as they "traverse the transnational landscape"; and to understand how changes in immigrants' cultural practices and network formation affect their identity, community building, and creative potential. The report recommends increasing networking opportunities through the arts.

Noshir Contractor, director of the SONIC group at NCSA, was co-primary investigator. Hank Green, a research scientist on the SONIC team, was a co-author of the study report.

For more information: http://www.fieldmuseum.org/creativenetworks

AVL develops improved software to render challenging AMR data

Adaptive mesh refinement (AMR) is a powerful technique for simulating phenomena—from the development of hurricanes to the formation of galaxies to the 3D mixing of fluids—at an incredibly wide range of temporal and spatial scales while maximizing computational efficiency. Visualizing this wide-ranging data is usually a challenge, but NCSA visualization programmer Matthew Hall, with assistance from Lorne Leonard, developed innovative software, named Amore, which substantially improves the rendering of AMR data.

Amore renders both particles and volumes at the high quality required for outreach venues such as planetarium domes and high-definition television programs as well as the far more challenging AMR data. Another advantage of Amore is that it runs on high-performance computational resources.

Using Amore, NCSA's Advanced Visualization Laboratory produced data-driven visualizations of the evolution of the universe for the high-definition NOVA program "Monster of the Milky Way," which debuted in October 2006. And visualizations of AMR data from detailed cosmological simulations were produced and shown in NCSA's booth at SC06.

Development of Amore was supported by funding from NSF and the Office of Naval Research through the Technology Research, Education, and Commercialization Center (TRECC) at NCSA.

Private Sector Program earns HPCwire Editors' Choice Award

The Private Sector Program at NCSA earned the 2006 HPCwire Editors' Choice Award for best collaboration between government and industry.

For 20 years, NCSA has partnered with leading-edge companies, tackling the real-world challenges faced by American industry and helping the center's collaborators gain the technological edge on their competitors. These partnerships span a wide breadth of projects—including high-performance computing, security, visualization, and data mining—and a wide variety of industries. Current NCSA partners include Abaqus, ACNielsen, Boeing, Caterpillar, Dell, Eclipse Energy Inc., ExxonMobil, IBM, Innerlink, John Deere, JPMorgan Chase & Co., Motorola Labs, Research Triangle Institute, and State Farm, with new partners joining the program each year.

For more information on NCSA's Private Sector Program, go online to: www.ncsa.uiuc.edu/AboutUs/Directorates/psp.html.

HPC wire

EDITORS' CHOICE AWARDS

BEST HPC COLLABORATION
BETWEEN
GOVERNMENT & INDUSTRY

NCSA
PRIVATE SECTOR PROGRAM

2006

Sharing the word about NCSA

Participating in exhibitions at conferences and meetings provides a forum for NCSA to spread the word to scientists, engineers, and others who have yet to tap into the resources the center offers. In the past few months NCSA could be found at exhibitions of the American Chemical Society, SC06, and the American Geophysical Union.

In addition, NCSA's expertise in scientific visualization was demonstrated in a first-ever public open house at NSF headquarters February 5, 2007. The 15 exhibitors were selected by their representative NSF directorate/office with NCSA representing the office of Cyberinfrastructure.

NCSA's scientific visualizations and other projects were showcased later in February at the American Association for the Advancement of Science (AAAS) as part of the NSF booth. The booth also featured a presentation by NCSA collaborator Thanh Truong on CSE-Online (a cyberenvironment for computational chemistry). During the meeting, NCSA's Donna Cox participated in a panel discussion describing the capabilities provided by the TeraGrid with TeraGrid Director Charlie Catlett and Phil Maechling from the Southern California Earthquake Center.

Donna Cox, leader of NCSA's Advanced Visualization Laboratory, narrates a demonstration of scientific visualizations shown on the center's high-definition stereo theater as part of NSF's booth at the AAAS exhibition in February. Photo: Bob Patterson

NCASSR builds framework for investigating cyberattacks

FBI Major Incident 216 is the kind of nightmare scenario no system administrator wants to face alone. In the summer of 2004, a foreign attacker gradually accumulated usernames and passwords, then invaded and compromised the integrity of a large number of hosts in government, higher education, and the private sector, both in the U.S. and abroad, with the ultimate goal of building a network of these subverted servers for some unknown purpose.

This year, the National Center for Advanced Secure Systems Research (NCASSR) is integrating many of the technologies its researchers have developed over the years into a cyberenvironment framework designed to help system administrators and law enforcement professionals collaboratively investigate attacks like that incident. Investigators will be able to work together in protected collaborative workspaces on the framework to safely upload data, analyze it, visualize the results, and communicate securely about how to handle the intrusion and prevent similar attacks in the future.

Since its launch in 2003, NCASSR has focused on basic and applied research leading to the development of next-generation information security technologies such as data mining, intrusion detection and analysis, secure communications, and secure grids, sensor networks, and infrastructure. Led by Randy Butler and Von Welch of NCSA, the center is a collaboration between NCSA, Pacific Northwest National Laboratory in Richland, Washington, and the Naval Postgraduate School in Monterey, California. NCASSR recently received fourth-year funding of nearly $2.7 million from the Department of Defense Office of Naval Research to support development of the framework, as well as several other projects, including a secure infrastructure for emergency communications, a training simulator for DOD system administrators, and systems for automated security testing, detection of rare events in data streams, and performance testing and evaluation of cybersecurity tools.

For more information about NCASSR and its current and past research efforts, see www.ncassr.org.

| Five in Focus |

Enabling computational and experimental chemists to easily do computational chemistry by providing an open source cyberinfrastructure scientists can access from their own desktop computers is the driving force behind the Computational Chemistry Grid (CCG). CCG is a virtual organization that provides access to high-performance computing resources for computational chemistry with intuitive interfaces and measurable quality of service. User support and user services are also provided. Here are five ways CCG helps chemists.

1 **Client inferface.** Running Java on their desktop computer, CCG users utilize GridChem, an easy-to-maneuver client interface, to access the network through a secure Internet portal. The GridChem interface integrates the hardware, software, and middleware resources necessary to solve computational chemistry problems using grid technologies.

2 **Manage jobs.** GridChem allows users to submit and monitor jobs run on any computer on which they have an allocation, without having to log in to each system or to remember the unique features of each system. Files are staged automatically.

3 **Resource allocations.** Last year, 170 researchers used more than 335,000 service units of computing time. Jobs are automatically dispatched to available CGG resources.

4 **Download results and post-process review.** Results are downloaded to the user's desktop machine. After the computation is completed, users need to extract information from the output files. Usually, the output information is determined by research parameters such as software applied, type of calculation, theory level, and so forth.

5 **Educational benefits.** The simplified interface of the GridChem tools makes it easy to build models or create molecular simulations. Tools can be combined via web pages to create virtual "lab assignments" for courses.

1 Computational chemistry a click away

GridChem uses well established grid technologies to provide chemists access to cyberinfrastructure (software, hardware, and services). A self-contained desktop client for input editing, job submission to remote HPC sites, and job management and post-processing for computational chemistry applications, GridChem is portable, adaptable to other applications, easy to install, and does not require installation of any grid software on the desktop system. Clients include expert, novice, and non-traditional users in both computational and experimental chemistry.

By using GridChem and the computing resources of the CCG, Robert Bach and Olga Dmitrenko, researchers at the University of Delaware, were able to complete a series of model theoretical calculations that suggests a new mechanism for the oxidation step in enzymatic cytochrome P450 hydroxylation of saturated hydrocarbons. "Such calculations involving structures with 500-600 basis functions are simply not feasible utilizing local work stations available at my university," says Bach. "The potential for several electronic states (doublet vs. quartet) increases the amount of computing power necessary to study such complex problems." Their work is supported by NSF.

P450 is an enzyme in the liver that breaks down carbohydrates, fats, and other things that contain hydrocarbons, such as steroid drugs, into metabolites. How P450 breaks down a drug affects how it works. Learning more about how P450 acts can lead to better, more effective medicines.

The CCG/GridChem project is an NSF three-year National Middleware Initiative program now in its third year. NCSA was a key player in the initial development of the GridChem client software and middleware. It is the middleware that makes this project feasible. "The middleware hides all the complexity of the grid, the distributed cyberinfrastructure, and provides an intuitively simple interface to the end-user chemist," explains Sudhakar Pamidighantam, an NCSA research scientist involved in the development of GridChem. He now provides support to GridChem users and helped Bach and Dmitrenko solve some problems with computational chemistry applications integrated into GridChem.

Other project partners are: the Ohio Supercomputing Center (Ohio State University), the Center for Computational Sciences (University of Kentucky), the Center for Computation and Technology (Louisiana State University), and the Texas Advanced Computing Center (University of Texas).

Oceanic forecasting

NCSA's Advanced Visualization Laboratory (AVL) collaborated with the Monterey Bay Aquarium Research Institute (MBARI) and the Jet Propulsion Laboratory/Caltech to visualize Regional Ocean Modeling System (ROMS) simulation data. The simulation data visualized here were generated as part of the Autonomous Ocean Sampling Network II program (AOSN II). Fleets of robotic vehicles are enabling the observation and prediction of ocean processes by providing an adaptive observation system for the ocean interior. In 2003, during the AOSN II field experiment, the first large-scale deployment of vehicles was used to predict the evolution of episodic wind-driven upwelling in the environs of Monterey Bay. Twelve different institutions contributed to the effort, which was led by MBARI. The observing system included a communication framework that allowed observations to be transmitted to two real-time oceanographic models. The resulting system provided the oceanic equivalent of atmospheric weather prediction, with all the advantages

that prediction entails. The models generated nowcasts and forecasts of ocean conditions. A key to the effort was 'adaptive sampling,' the development of control strategies to command the mobile vehicles to places where their data would be most useful. The image shown here is a visualization of ocean currents and temperature as realized by the ROMS ocean model, run by the Jet Propulsion Laboratory under the direction of Yi Chao.

The AVL team worked with James Bellingham, MBARI chief technologist, to develop novel 3D visualization techniques for the ROMS data. "Computer visualization empowers us with the ability to explore the ocean's inner workings," says Bellingham.

The AVL team advected particles to create three-dimensional flow trajectories colored by temperature scalar values. The temporal length of each trajectory represents a 24-hour period. Isosurfaces were extracted to identify water regions that are colder or warmer
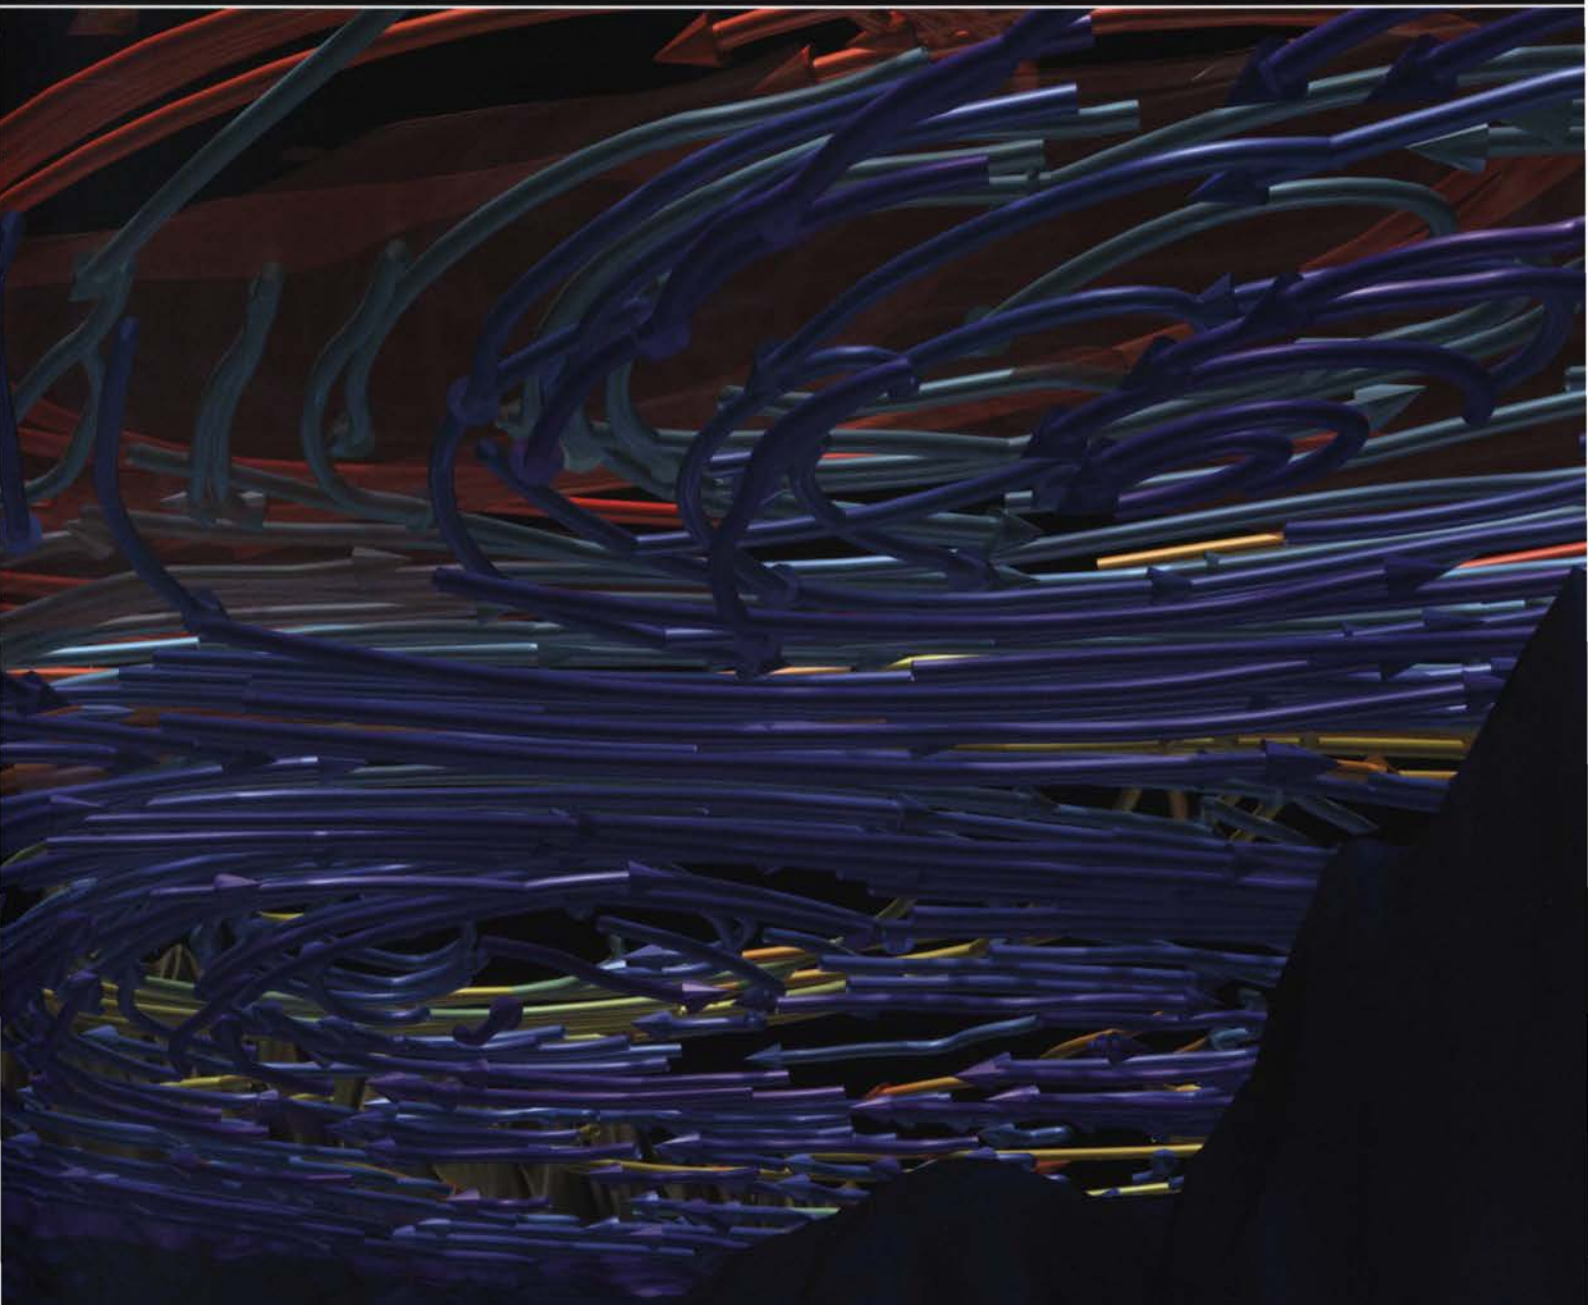
than normal by two standard deviations from the depth adjusted mean. Elevation and satellite data from multiple sources were ortho-rectified to create a digital elevation model for the Monterey Bay region. The AVL team also created an animation of the simulation that covers a one-month time period.

NCSA's data visualization plugin for Maya/Mental Ray was used to read the trajectory and isosurface data and to set mappings between simulation data and rendering parameters. This visualization research was funded in part by the Laboratory for the Ocean Observatory Knowledge INtegration Grid (LOOKING) National Science Foundation grant and by the Technology Research, Education, and Commercialization Center (TRECC), a program of the University of Illinois at Urbana-Champaign, funded by the Office of Naval Research (ONR) and administered by NCSA.

JPL ROMS research is funded by NASA and ONR.
AOSN was funded by ONR.

Team Members

Scientific simulation
Yi Chao

Visualization
Alex Betts
Donna Cox
Matt Hall
Lorne Leonard
Stuart Levy
Robert Patterson

NCSA

I