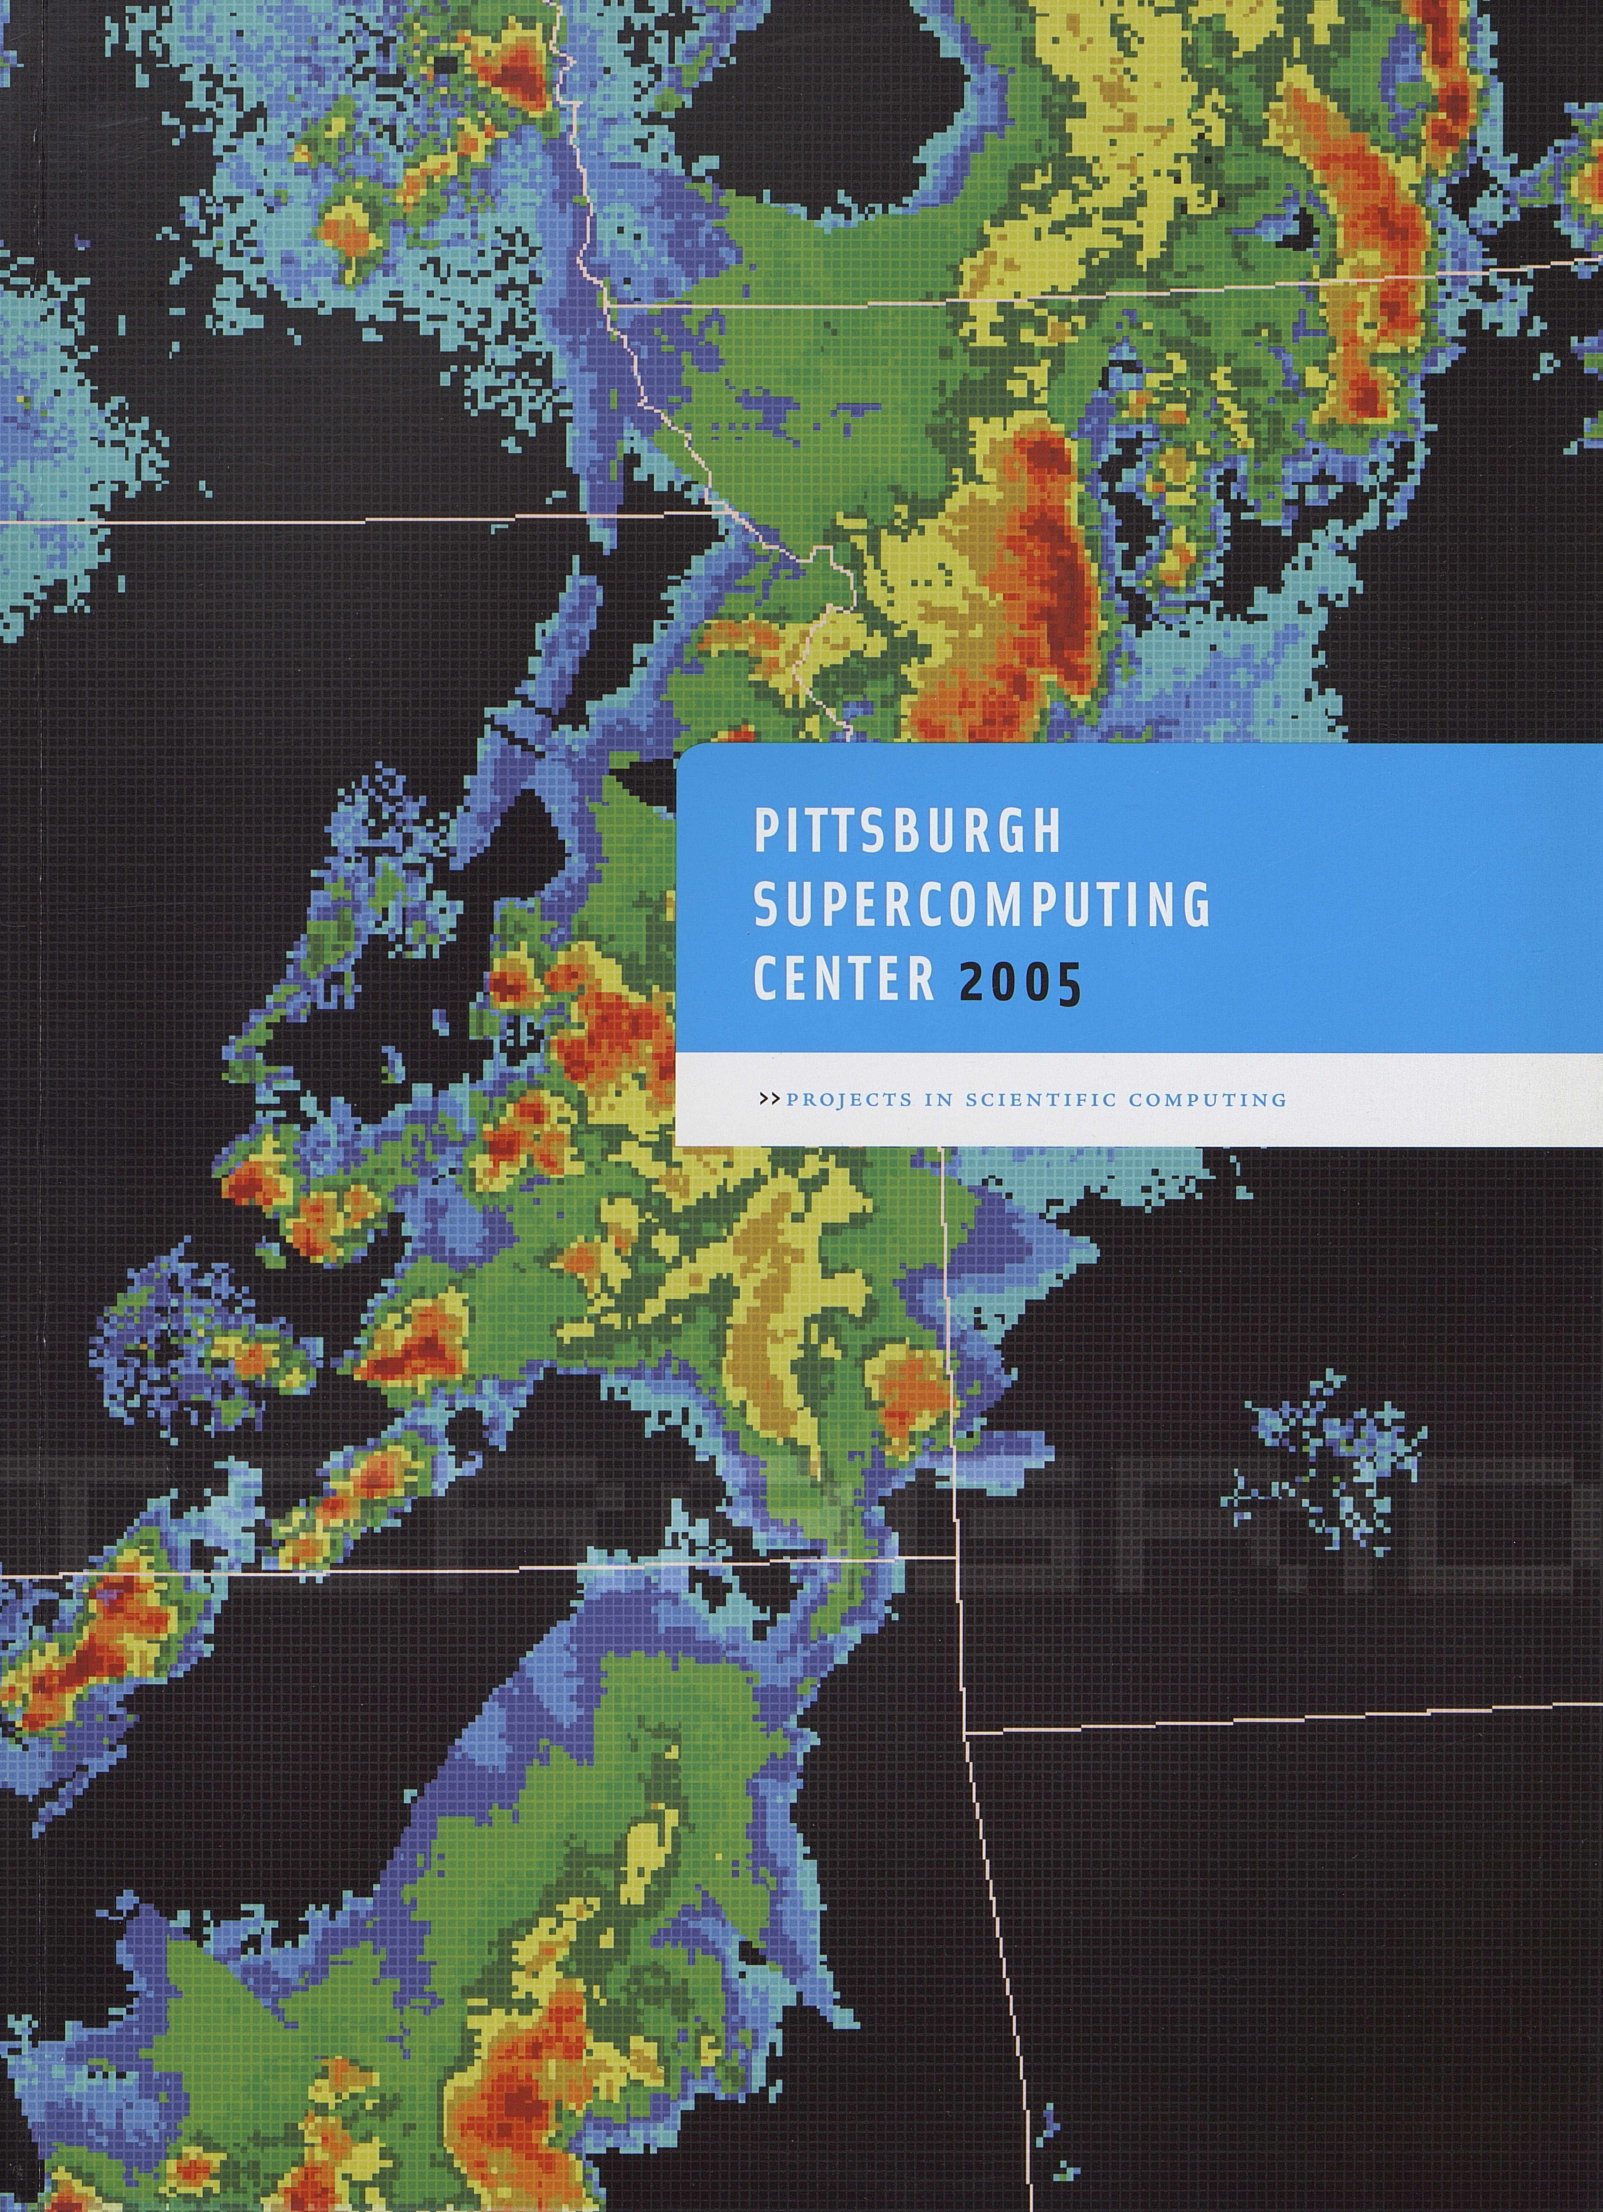

PITTSBURGH SUPERCOMPUTING CENTER 2005

>> PROJECTS IN SCIENTIFIC COMPUTING

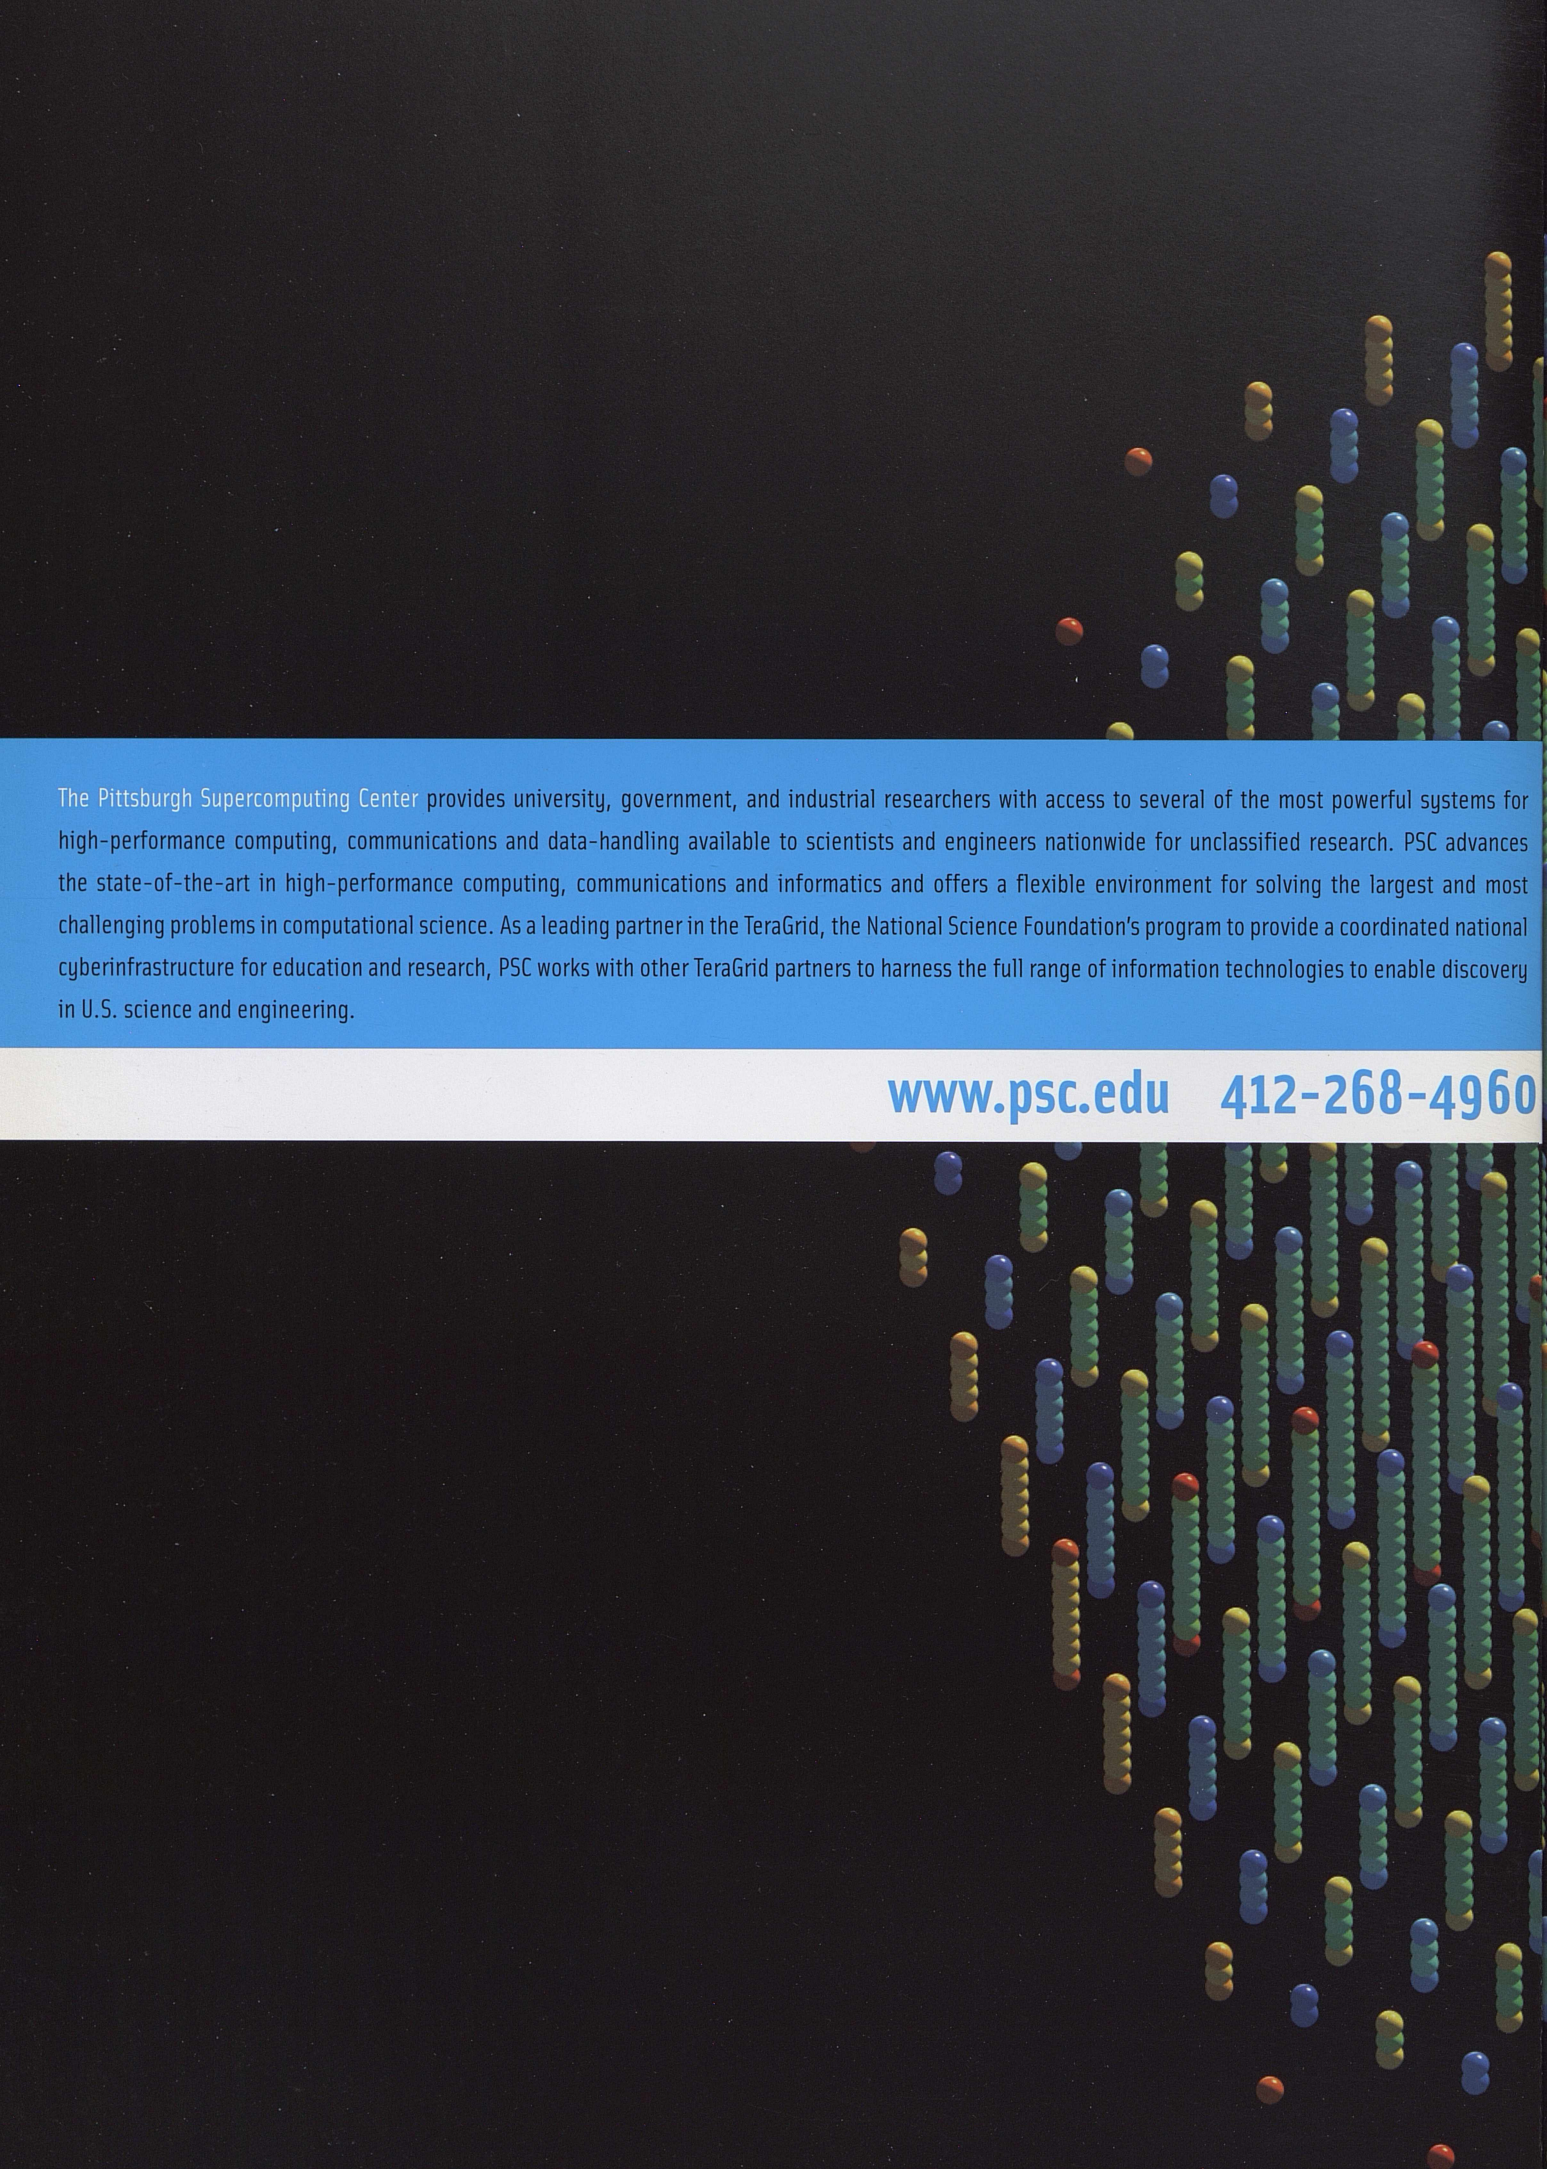

The Pittsburgh Supercomputing Center provides university, government, and industrial researchers with access to several of the most powerful systems for high-performance computing, communications and data-handling available to scientists and engineers nationwide for unclassified research. PSC advances the state-of-the-art in high-performance computing, communications and informatics and offers a flexible environment for solving the largest and most challenging problems in computational science. As a leading partner in the TeraGrid, the National Science Foundation's program to provide a coordinated national cyberinfrastructure for education and research, PSC works with other TeraGrid partners to harness the full range of information technologies to enable discovery in U.S. science and engineering.

www.psc.edu 412-268-4960

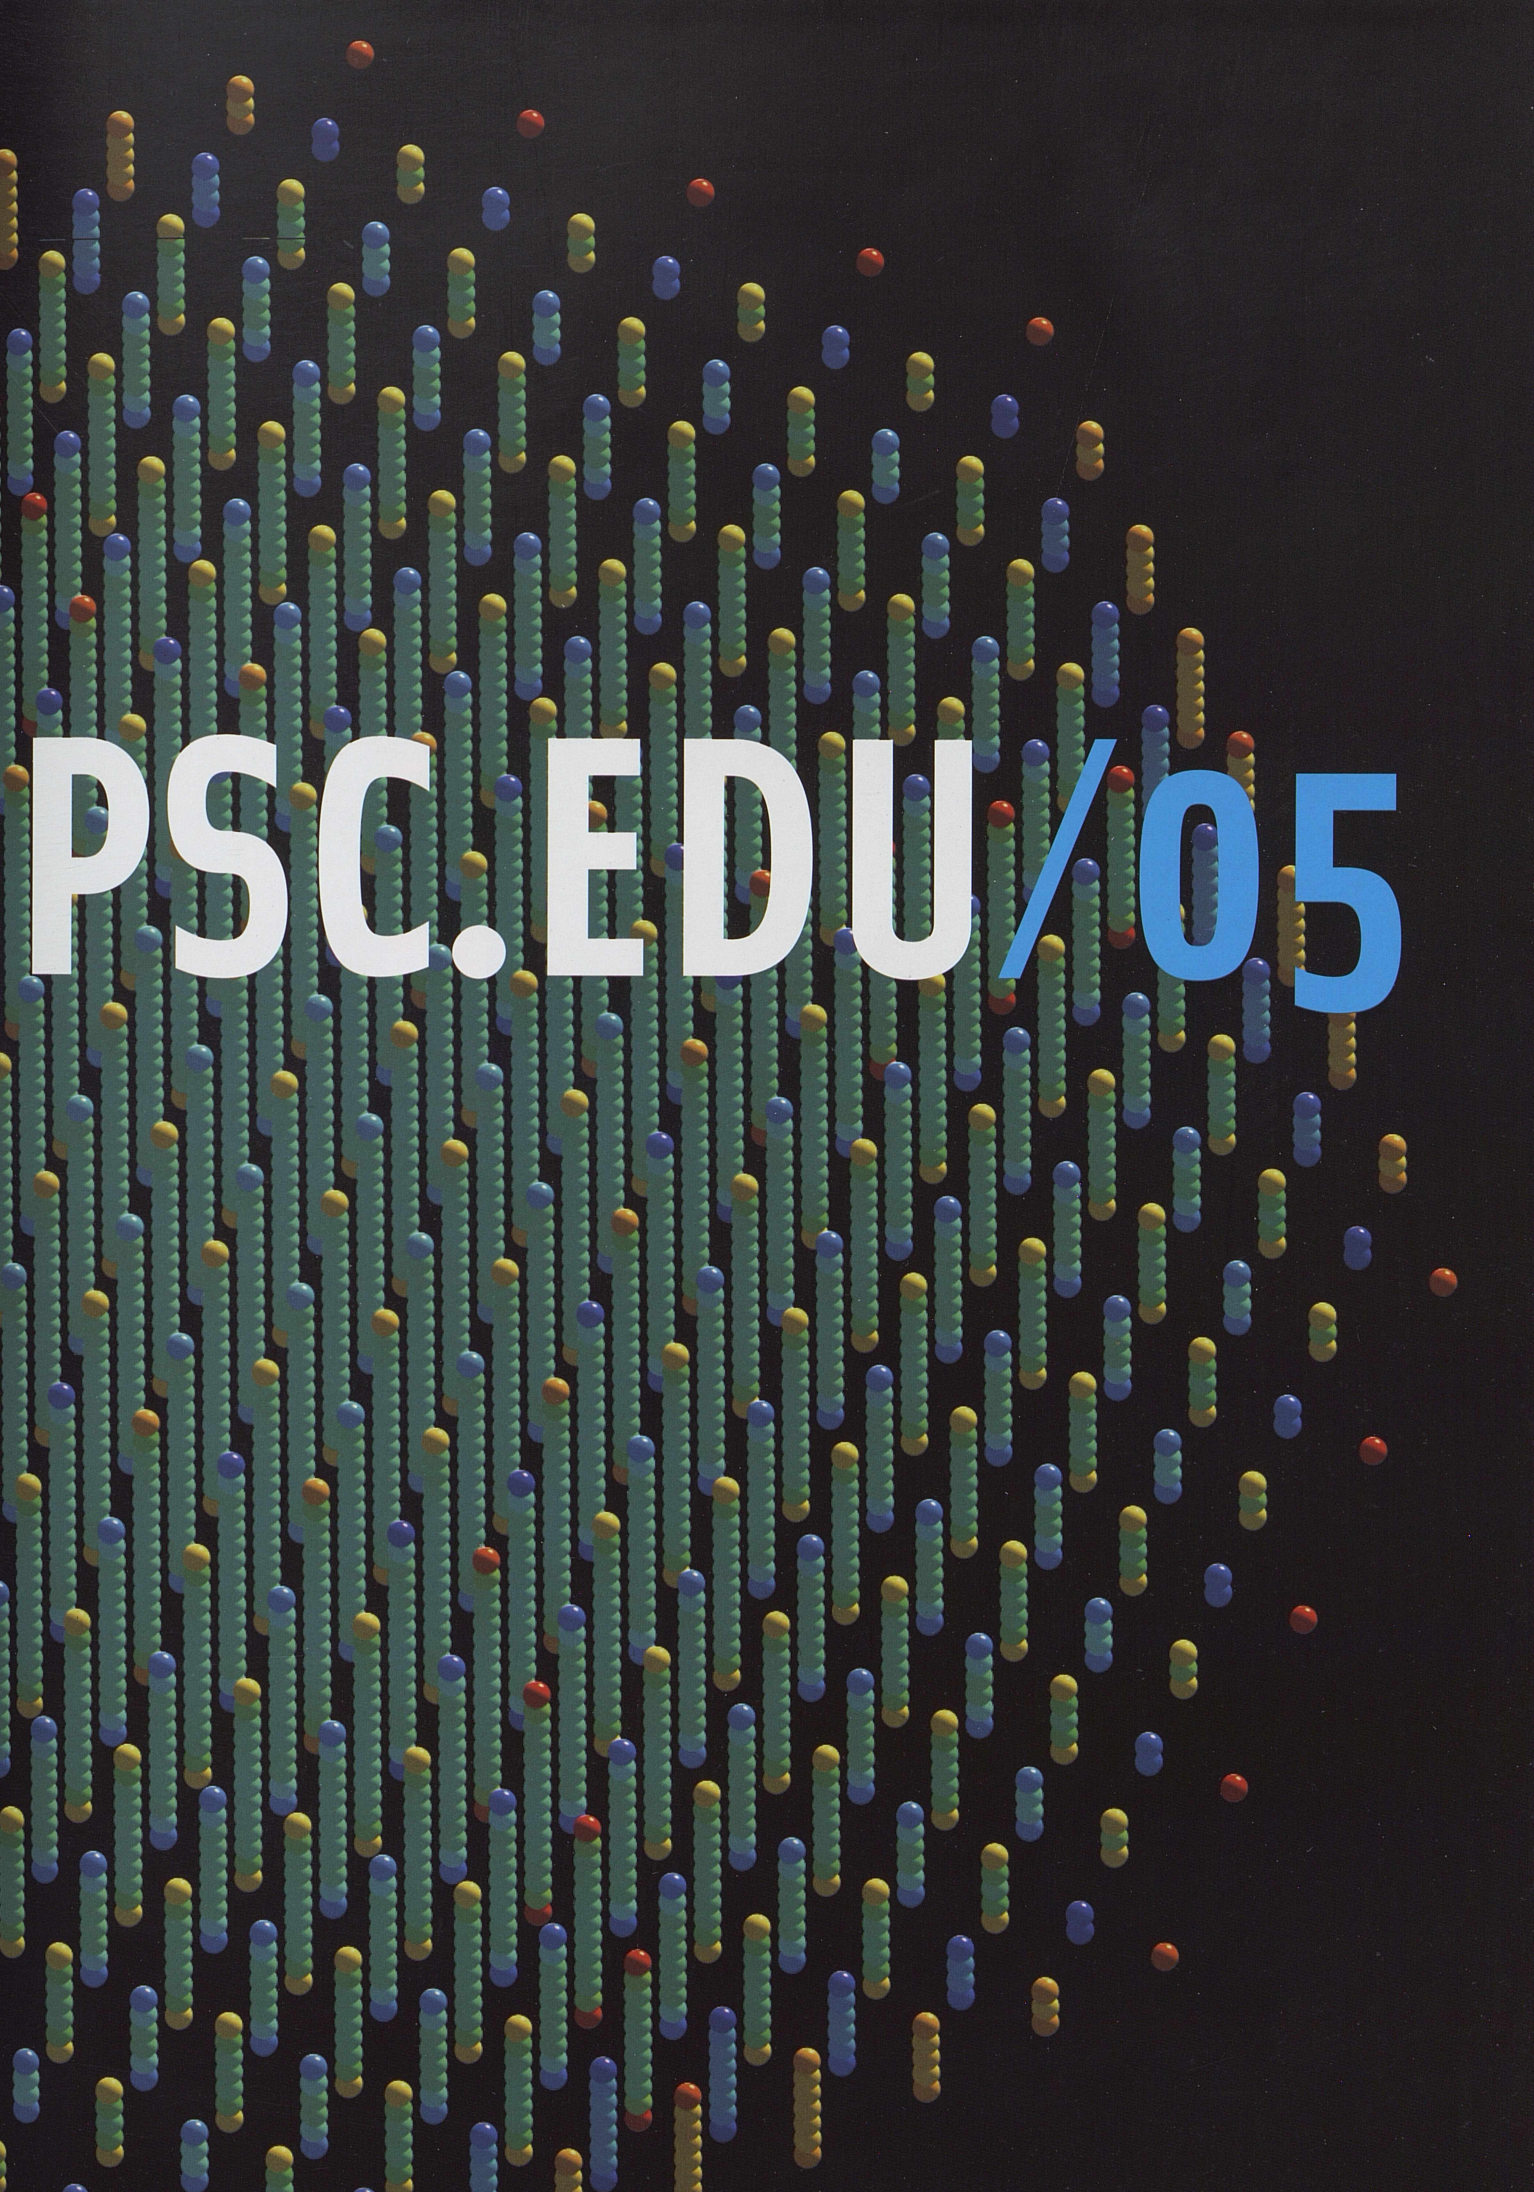
PSC.EDU/05

FOREWORD FROM THE DIRECTORS

In this booklet we highlight the past year's scientific progress at PSC. As with half a dozen prior state-of-the-art systems, PSC has demonstrated its ability to rapidly transform new technologies into productive research tools. The first XT3 system shipped from Cray, Big Ben (p. 4) has in the space of a year been installed, tested, optimized for research applications and is now a production resource on the TeraGrid. Such feats happen only because PSC's exceptional systems and networking staff (p. 12) possess the world-class knowledge, experience, creativity and energy to make them happen.

With Sandia National Laboratories and Cray Inc., PSC collaborated to envision an architecture, realized in the XT3, that represents the most advanced current thinking for a high-performance system to support very large-scale computations across many fields, and that is fully expandable to well over 100 teraflops. With PSC's active leadership, Big Ben became the first Cray XT3 to run applications, first to be made available to NSF researchers, and has already produced notable results in several fields, including turbulence (p. 5) and nanotechnology (p. 46).

Michael Levine (left) and **Ralph Roskies**, PSC co-scientific directors, with Big Ben, Cray XT3, serial # 1.

This entirely new architecture is now integrated into the TeraGrid, where it will become the resource best suited for very large-scale, demanding projects. The United States needs leading-edge capability to continue to advance in science and engineering, and we are pleased that the NSF has decided to push to a Petaflop computer by 2010.

The TeraGrid itself, NSF's visionary program to transform the nation's research capability, is now entering a second major phase, with PSC in a leadership role — supporting the research community in exploiting the TeraGrid as a resource to produce new knowledge, as we have done with PSC systems, and leading the TeraGrid's increasingly vital security efforts. With a major NSF award to support our TeraGrid-related operations (p. 4), we look forward to the challenges of moving this ambitious program forward.

Among many scientific accomplishments PSC has enabled, the advance in storm forecasting (p. 18), high-point of a decade-long collaboration with the Center for Analysis and Prediction of Storms, is at the forefront of work that demonstrates how high-performance computing contributes directly to improving our lives.

Understanding neural processes is an area where PSC marshals special excellence. One feature of PSC's biomedical program (p. 10), which engages HPC expertise with the life sciences, the MCell simulation program, co-authored by PSC scientist Joel Stiles, and supported with our outstanding biomedical visualization capabilities, has proven ability to realistically model extremely complex neurotransmitter events at synapses. Stiles' recent work on exocytosis (p. 22) is forcing a rethinking of some long-held views.

Photonic crystals represent a major success story for computational science, as exemplified by two projects reported here (p. 34). Shanhui Fan at Stanford found a new way to stop light in its tracks. The Joannopoulos team at MIT developed a new fiber that makes laser surgeries possible that used to be impossible. This progress depends on the ability to simulate how light behaves, and shows us, if we still need to be shown, that running models inside computers can tell us important things we have no other way to learn.

Ken Jordan's work (p. 30) with water clusters has led to several papers in *Science* in the past year, including recognition as one of the 10 most important science breakthroughs of 2004. P.K. Yeung's work (p. 40) is the largest-scale turbulence simulation yet. Molecular dynamics studies by Lynch and Reggio (p.26) may help in the development of new therapeutic drugs.

PSC's development work in many areas, such as PDIO (p. 5), facilitates the scientific progress reported here. This is exemplified also by our work with automated checkpoint-restart (p. 44), which will become more critical as we progress to systems with many more processors. We look forward to collaborating with Keshav Pingali and his team at Cornell to realize the advantages of CPR with Big Ben.

For all these achievements, credit is due to the creative, hard-working staff of PSC. We are most grateful for the support we receive from the National Science Foundation, the U.S. Department of Energy, the National Center for Research Resources of the National Institutes of Health, the Commonwealth of Pennsylvania and many others.

MICHAEL LEVINE, SCIENTIFIC DIRECTOR

RALPH ROSKIES, SCIENTIFIC DIRECTOR

CONTENTS

PITTSBURGH SUPERCOMPUTING CENTER, 2005

Creating Cyberinfrastructure

CRAY XT3 SERIAL #1: THE NEWEST STAGE IN EVOLUTION OF HPC TECHNOLOGY, A MAJOR BOOST FOR U.S. COMPUTATIONAL SCIENCE

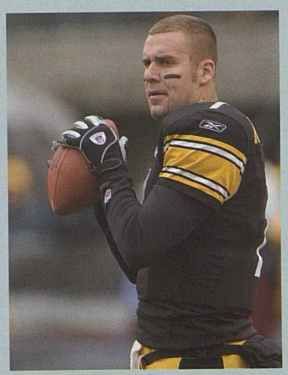

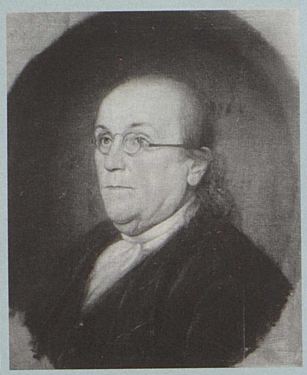

Installed in late 2004 and early 2005, officially unveiled in July, in full production by October, Big Ben supports research nationwide as part of the NSF TeraGrid. Acquired via a $9.7 million grant from the National Science Foundation in September 2004, Big Ben is the first of the most advanced line of HPC systems from Cray Inc., the XT3, to be shipped. It comprises 2,090 processors with an overall peak performance of 10 teraflops: 10 trillion calculations per second.

Named for both the Pittsburgh Steelers quarterback and for Ben Franklin, the nation's first great scientist, Big Ben is a leading-edge computing resource on the TeraGrid. PSC's current lead system, LeMieux, has been one of the most productive TeraGrid systems, and although LeMieux — 3,000 processors, six teraflops — remains a much-used resource, Big Ben has begun to take over LeMieux's role as the TeraGrid resource best suited for very large-scale, demanding projects.

On a per processor basis, Big Ben is 2.4 times faster than LeMieux. More than sheer processor speed, however, its most significant technological advance is "inter-processor bandwidth" — the speed at which processors share information with each other. Because of this, Big Ben has demonstrated performance nearly 13 times better than LeMieux on key applications when 1,000 or more processors are used. Many areas of research will benefit from this, including nanotechnology, design of new materials, protein dynamics studies that lead to new therapeutic drugs, modeling of earthquake soil-vibration, and severe storm forecasting.

BIG BEN RIBBON-CUTTING (JULY 20, 2005): NSF DIRECTOR COMMENTS

"With this system," said Arden Bement, director of the National Science Foundation, "we are fulfilling an important national goal — providing one of the fastest computing capabilities to the U.S. research community. We celebrate a significant leap in science and engineering research and education capacity. Richness of data, combined with powerful computing facilities and innovative people, promises a multitude of scientific breakthroughs. The Pittsburgh Supercomputing Center has all of these."

Jim Kasdorf, PSC director of special projects

NSF AWARDS $52 MILLION TO PSC FOR TERAGRID OPERATIONS

In August 2005, the National Science Foundation awarded $52 million over five years to support operations of the Pittsburgh Supercomputing Center (PSC) as a leading partner in the TeraGrid, NSF's program to provide national cyberinfrastructure for education and research. Built over the last four years, the TeraGrid is the world's largest, most comprehensive distributed cyberinfrastructure for open scientific research. The award to PSC is part of a five-year, $150 million NSF award to the eight

TERAGRID RESOURCE PROVIDERS

- Indiana University
- National Center for Supercomputing Applications
- Oak Ridge National Laboratory
- Pittsburgh Supercomputing Center
- Purdue University
- San Diego Supercomputer Center
- Texas Advanced Computing Center
- The University of Chicago/ Argonne National Laboratory

TeraGrid™

TeraGrid partner institutions.

Much as physical infrastructure such as power grids, telephone lines, and water systems enables modern life, cyberinfrastructure makes possible much of modern scientific research. Through high-performance network connections, the TeraGrid integrates high-performance computers, data resources and tools, and high-end experimental facilities at locations around the country.

"TeraGrid unites the scientific and engineering community so that larger, more complex scientific questions can be answered," said Arden Bement, director of the National Science Foundation. "Solving these larger challenges will, in turn, motivate the development of the next generation of cyberinfrastructure."

"This award represents an opportunity to play an important leadership role," said PSC co-scientific directors Michael Levine and Ralph Roskies. "We look forward to meeting the challenges ahead with TeraGrid in harnessing the full range of information technologies for coordinated, distributed, productive work enabling leading-edge science."

A LEADERSHIP ROLE IN CREATING NATIONAL CYBERINFRASTRUCTURE

Within the TeraGrid, PSC has taken leadership responsibility in user services and cyber-security as well as emphasizing capability computing, the ability to tackle the computationally most-demanding problems. PSC leads the TeraGrid user services working-group, which coordinates the effort to provide thousands of researchers nationally with consulting and other support to assure they can productively use the TeraGrid, and also leads the TeraGrid security working-group, which guides TeraGrid security policy.

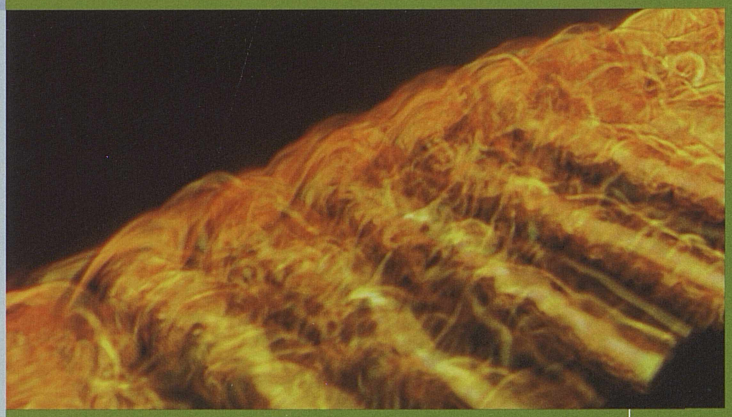

BIG BEN DELIVERS IN REAL TIME VIA TERAGRID

In September 2005, scientists used special PSC-developed software to demonstrate real-time access from a remote location to data from simulations running on Big Ben. The demonstration took place at iGrid2005 in San Diego, a conference on the scientific use of high-performance networks and grid computing. From the iGrid show floor, the researchers –Paul Woodward, David Porter and colleagues at the University of Minnesota – used Big Ben to simulate turbulent fluid dynamics in shear driven mixing layers.

As the numbers crunched in Pittsburgh, the researchers volume-rendered images from the data and displayed them in San Diego. Over the course of the 90-minute demonstration, Big Ben delivered a sustained 200 megabits per second via the TeraGrid, with burst rates nearing 800 megabits per second. To accomplish this, the researchers relied on a new PSC capability – Portals Direct I/O (PDIO) – that can route simulation data from Big Ben's processors in real-time to remote users anywhere on the network. Data fragments written by each Big Ben processor are assembled by PDIO into complete files at the receiving end. PDIO controls the rate of data transfer, even slowing down the application if necessary. With PDIO, Woodward and Porter tested their demonstration over slower networks before using the TeraGrid backbone. PDIO has shown absolute stability for both short (minutes) and long (hours) periods of use.

Supercomputing In Pennsylvania

WITH COMMONWEALTH OF PENNSYLVANIA SUPPORT, PSC PROVIDES EDUCATION, CONSULTING, ADVANCED NETWORK ACCESS AND COMPUTATIONAL RESOURCES TO SCIENTISTS AND ENGINEERS ACROSS THE STATE

Cutting the Ribbon for Big Ben
Beverly Clayton, PSC executive director, presided at a "ribbon cutting" for Big Ben, PSC's Cray XT3 system, on July 20, 2005 at Westinghouse Energy Center in Monroeville, Pa.

ECONOMIC DEVELOPMENT

PSC's high-performance computing and networking are world-class technology resources that officials and agencies can point to in promoting Pennsylvania as a location for business. During the past year, PSC provided Pennsylvania companies with over 15,000 hours of computing time. Among them, Pittsburgh-based PPG Industries uses LeMieux, PSC's terascale system for computational modeling in several aspects of its product lines as a global supplier of coatings, glass, fiberglass and chemicals.

PSC's networking group advises the Pennsylvania Department of Administration on how telecommunications can help attract new business to the state. Networking staff advised on developing a state-wide networking infrastructure, leading to a new agreement for K-20 organizations to participate in the Internet2 advanced network organization. They also consulted with the Center for Appalachian Network Access at Bedford Senior High School, a regional resource for deployment of internet access in rural communities.

From November 6-12, 2004, as a result of PSC's global presence in computational science and technology, the annual Supercomputing conference — SC '04 — was held in Pittsburgh. More than 7,900 people from around the world gathered at the David Lawrence Convention Center. This annual conference, in Pittsburgh for the

(l to r) **David Kiefer**, senior VP, Cray Inc.; **Sean Logan**, Pennsylvania senator, 45th District; **Dan Onorato**, chief executive, Allegheny County, Pennsylvania; **Peter A. Freeman**, NSF assistant director for Computer & Information Science & Engineering; **Jared L. Cohen**, president, Carnegie Mellon University; **Mark A. Nordenberg**, chancellor, University of Pittsburgh; **Bob Regola**, Pennsylvania senator, 39th District; **Arden Bement**, director, National Science Foundation; **Steve Tritch**, president & CEO, Westinghouse Electric Company; **Ralph Roskies**, co-scientific director, Pittsburgh Supercomputing Center; **Mike Levine**, co-scientific director, Pittsburgh Supercomputing Center.

second time in eight years, brings together scientists and engineers, hardware and software companies and many others to share ideas and assess new developments in the fields of high-performance computing, networking and storage. According to the Greater Pittsburgh Convention & Visitors Bureau, SC '04 produced an economic gain for the southwest Pensylvania region of $13 million.

OUTREACH & TRAINING

This year PSC presented the fourth in a series of annual technology-briefing days to staff from the Bechtel Bettis Atomic Power Laboratory in Pittsburgh. PSC consultants provided information on developing, managing and using a parallel distributed-computing environment.

PSC participates in a number of events each year to provide information on high-performance computing to companies and government agencies. In December 2004, PSC exhibited at the Annual Eastern Intergovernmental Technology Conference in Hershey, Pa. In August 2005, PSC exhibited at the TechTrends conference in Pittsburgh, an annual conference that highlights opportunities for government, academic and industrial collaboration in the four-state region of Pennsylvania, Delaware, Maryland and New Jersey.

RESEARCH IN PENNSYLVANIA

By supporting Pennsylvania university researchers, PSC resources help to attract research funds to the state. During the past year, more than 430 Pennsylvania researchers from 12 institutions used nearly 650,000 processor hours on PSC resources that are specifically targeted at Pennsylvania users. In addition, Pennsylvania researchers received allocations through the NSF process of nearly five million hours on LeMieux, PSC's terascale system, along with 169,000 hours on PSC's HP Marvel systems.

COOL DESIGN IN ELECTRONIC COMPONENTS

Cristina Amon is a specialist in cool design. Her research has pioneered the use of computational methods as a way to model the heat removed from micro-circuitry in portable electronic devices such as laptops and wearable computers. As director since 1999 of Carnegie Mellon's Institute for Complex Engineered Systems, she has promoted research-industry collaboration in western Pennsylvania. In projects using resources at PSC, she has developed numerical methods to model the heat produced by and removed from microcircuitry. In the early 1990s, she applied her approach to the thermal design of VuMan, a first-generation wearable computer.

Cristina Amon, Raymond J. Lane Distinguished Professor of Mechanical Engineering at Carnegie Mellon University and PSC senior scientist Marcela Madrid.

In more recent work, Amon collaborated with PSC scientist Marcela Madrid to develop an innovative method to calculate the thermal conduction in silicon "thin films". The predominant material used to manufacture the "chips" of micro-electronics, silicon in some devices is deposited in very thin layers — from 10 to several hundred nanometers. In many applications, these silicon thin films are deposited on dielectric materials that conduct heat poorly, requiring that the heat produced must be removed by heat flow within the thin film itself. Thermal properties of silicon in these very thin layers differ from bulk silicon, and because of this, existing computational approaches have had limited usefulness.

Madrid's specialty is an approach called "molecular dynamics," which tracks the forces and interactions among atoms over time. In collaboration with Amon, she helped to develop a "potential" — a mathematical expression — to realistically represent thermal conductivity for applying molecular dynamics to silicon thin films. Applying this potential at the surface of the thin film overcomes inaccuracies introduced in prior computational approaches. This innovative approach allows effective use of molecular dynamics to model thermal conductivity of silicon thin films, and the results predicted for films ranging in thickness from two to 200 nanometers show good agreement with experimental data.

The Super Computing Science Consortium

PENNSYLVANIA-WEST VIRGINIA PARTNERS IN DEVELOPMENT OF CLEAN POWER TECHNOLOGIES

Formed in 1999 and supported by the U.S. Department of Energy, the Super Computing Science Consortium is a regional partnership of research and educational institutions in Pennsylvania and West Virginia. $(SC)^2$ provides intellectual leadership and advanced computing and communications resources to solve problems in energy and the environment and to stimulate regional high-technology development.

PSC supplies training along with consulting support to $(SC)^2$ partners. Training includes workshops in parallel programming and visualization. In September 2005, PSC scientific-visualization specialist Kent Eschenberg presented a tutorial on ParaView, an application for visualization of large data sets, to NETL researchers. ParaView runs on distributed parallel systems such as PSC's LeMieux and Big Ben.

Since the spring of 2000, a high-speed network — the first fiber-optic service to Morgantown, West Virginia — has linked the National Energy Technology Laboratory (NETL) campuses in Morgantown and Pittsburgh with PSC, facilitating NETL collaborations. Researchers at NETL and WVU have actively used this link to tap PSC computational resources. Since the founding of $(SC)^2$, more than 40 $(SC)^2$ researchers have used PSC systems for a range of projects, using more than 2.6 million hours of computing time, 840,000 hours within the past year. Along with surface chemistry simulations described here (facing page), this work includes:

- Gas From Black Liquor
  http://**www.psc.edu/science/2004/sc2**
- Fluidized-bed Combustion Of Silane
  http://**www.psc.edu/science/2001/sc2**
- Lean-fuel Mixes In Next-generation Power-generating Turbines
  http://**www.psc.edu/science/Richards/clean_power.html**
- Industrial-scale Technology For Coal Gasification
  http://**www.psc.edu/publicinfo/2002/sc2**
- A New Design For A Power-generating Turbine
  http://**www.psc.edu/science/cizmas2002.html**

$(SC)^2$ co-chairs: **Lynn Layman**, PSC, and **Bob Romanowsky**, NETL. "The $(SC)^2$ partnership continues to advance research in both energy and the environment," says Layman. "Our regional partners are joining together to move forward our nation's knowledge in these vital areas."

$(SC)^2$ PARTNERS

National Energy Technology Laboratory
Pittsburgh Supercomputing Center
Carnegie Mellon University
Duquesne University
University of Pittsburgh
Waynesburg College
West Virginia University
Institute for Scientific Research
NASA Independent Verification & Validation Facility

>>More information: **http://www.sc-2.psc.edu**

$(SC)^2$ Super Computing Science Consortium

(SC)² AT SUPERCOMPUTING '04

In November 2004 for the first time, (SC)² sponsored a booth at Supercomputing, the annual International Conference for High Performance Computing, Networking and Storage. The conference, held in Pittsburgh for the second time in eight years, features displays and demonstrations by leading vendors and research organizations around the world. The booth featured posters highlighting research activities from (SC)² partners. In the photo is Barbara Kirby of Waynesburg College, who represents Waynesburg on the (SC)² steering committee and chaired the booth team.

TEAMING TO WIN

(SC)² joined other regional organizations at West Virginia's Teaming to Win conference June 1-2, at the Stonewall Resort, Walkersville, West Virginia. Sponsored by Congressman Allan B. Mollohan (WV, Dist. 1), the event provides a forum for West Virginia organizations to explore opportunities for collaborative projects. The photo shows Brian Dotson and Steve Zitney of NETL preparing a demonstration of their 3D visualization capability.

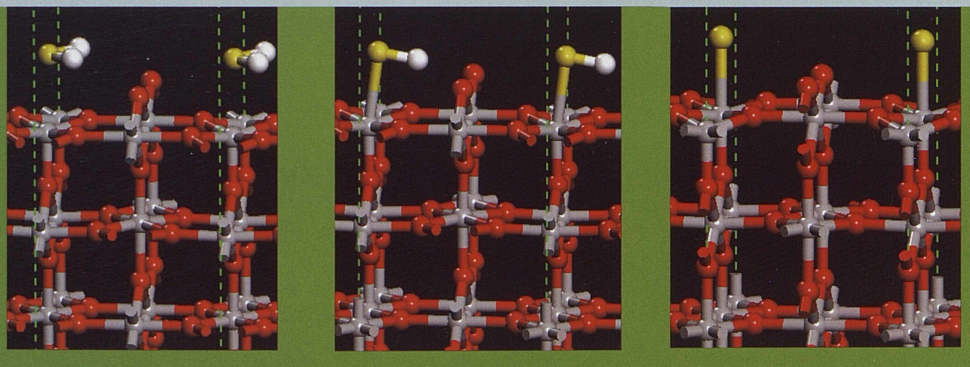

HYDROGEN FROM FOSSIL FUELS

The Bush administration's announcement in 2003 of a hydrogen fuel initiative has spurred interest in hydrogen's potential to play a major role in U.S. long-term energy planning. Hydrogen is clean, efficient and with hydrogen fuel-cell technology could power cars, airplanes, furnaces and air conditioners. One of the major challenges, however, in harnessing this renewable resource is finding affordable ways to produce fuel-quality hydrogen.

As a member of NETL's Computational Chemistry Team, NETL scientist Dan Sorescu has used PSC resources for a wide range of studies. In this work, he applies quantum-theory calculations from first principles to gain detailed molecular understanding of the chemistry and physics underlying fossil-fuel related processes. In work this year, he used Ben — PSC's 64-processor Alphaserver cluster — to explore a catalytic process that offers the potential of recovering hydrogen from hydrogen sulfide (H_2S), one of the principal byproducts from fossil fuel.

Titanium dioxide — a material found in beach sand and widely used as a white pigment — catalyzes the reaction by which hydrogen sulfide dissociates to hydrogen and sulfur. It does this by providing a surface to which H_2S chemically adheres — adsorbs — and from which H_2 subsequently desorbs. As these three frames show, Sorescu's computations with Ben map the sites on a titanium dioxide (gray & red) surface where H_2S (white & yellow) adsorbs, leading to sulfur formation and desorption of hydrogen. This kind of detailed understanding is essential to optimize the catalytic process.

The National Resource for Biomedical Supercomputing

NATIONAL LEADERSHIP IN HIGH-PERFORMANCE COMPUTING FOR BIOMEDICAL RESEARCH

Established in 1987 as one of the very first biomedical supercomputing programs in the country and renamed this year, the National Resource for Biomedical Supercomputing (NRBSC) pursues leading-edge research in the life sciences while exploiting high-performance computing and applying PSC expertise in computational science to collaborative biomedical research nationwide.

With support from NIH's National Center for Research Resources, NRBSC also conducts workshops and courses to train scientists in the use of high-performance computing for biomedical research, in areas such as spatially realistic cell modeling, volumetric data visualization and analysis, protein and DNA structure, genome sequence analysis and biological fluid dynamics. Since 1987, PSC biomedical training has reached more than 3,200 researchers, and PSC computational resources have supported more than 1,000 biomedical research projects involving more than 2,500 researchers at 218 research institutions in 48 states. Among these are several projects featured in this booklet (pp. 22-25, 26-29).

>>More information: **http://www.psc.edu/nrbsc**

NEURAL MODELING

PGENESIS is a parallelized version of widely used software for modeling biological neural systems such as the brain. Originally developed at PSC by biomedical researchers Nigel Goddard and Greg Hood, PGENESIS extends software called GENESIS to parallel systems and makes it possible to do multiple simulations and "parameter searches" — identify and zero-in on the most interesting among a set of simulations while they are running. This year, PSC released a new version of PGENESIS that integrates support for interprocessor communications and facilitates its use on workstation clusters and massively parallel systems such as LeMieux and Big Ben.

NRBSC BIOMEDICAL COLLABORATIONS

Albert Einstein College of Medicine
Carnegie Mellon University
Duke University
Hospital for Sick Children, Toronto
Howard University
Marine Biological Laboratory, Woods Hole
Morgan State University
North Carolina Central University
Rockefeller University
The Salk Institute
The Scripps Research Institute
University of California at Davis
University of California at San Diego
University of North Carolina, Chapel Hill
University of Pittsburgh
University of Pittsburgh School of Medicine
University of Puerto Rico, Medical Sciences Campus

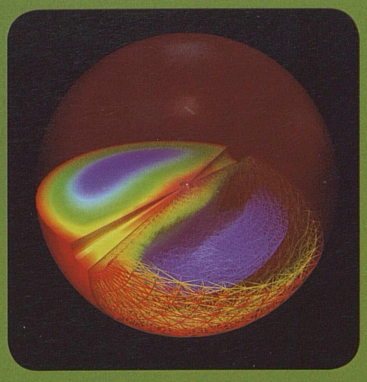

THE LIFE OF CELLS

With colleagues at the Salk Institute, scientists at PSC's Center for Quantitative Biological Simulation and NRBSC are developing MCell and DReAMM, software for simulation and visualization of physiologically realistic cellular models (see pp. 22-25). New versions of both programs include major new features. MCell can now simulate diffusion and reaction of multiple chemical species in solution, enabling much broader studies of metabolic, signaling and regulatory networks. DReAMM can visualize complex MCell models containing thousands of different objects and, as shown here, can also represent volumetric data. This example shows a finite-element simulation of electric field strength. Future versions of MCell will integrate Monte Carlo and finite-element algorithms for quantitative multiscale cellular simulations.

PREDICTING PROTEIN STRUCTURE

Genes are blueprints for proteins, and the flood of sequence data from genomic research challenges scientists to take the next step and use this massive data to deduce the 3D structures of proteins. NRBSC scientist Troy Wymore and University of Pittsburgh student Adam Marko participated in CASP6, an international experiment by which molecular biologists assess techniques to do this. Held every two years, CASP sets up blind tests with sequences for structures known but not published. In 2004, 166 research groups took part. The PSC group developed a method that combines comparative modeling, statistical potentials and replica-exchange simulations, and they achieved several high-scoring results. For a protein (POINTER) from a small flowering plant named arabidopsis, their proposed structure was the fifth-best worldwide.

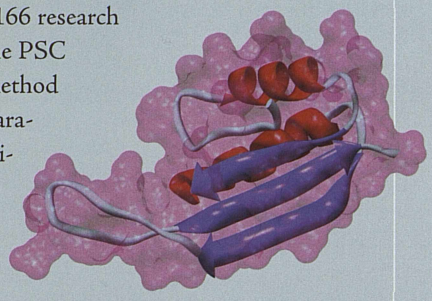

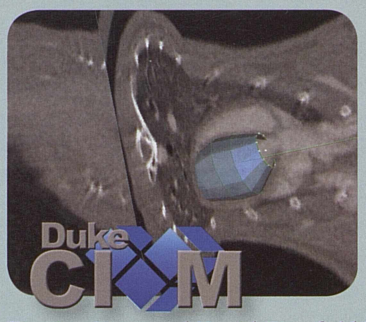

THE BEATING HEART OF A MOUSE

With support from the National Library of Medicine, NRBSC scientists Arthur Wetzel, Stuart Pomerantz, Demian Nave and David Deerfield have developed a new version of the PSC Volume Browser, PSC-VB II. The original PSC-VB, released in 2002, operated with large static anatomical volumes such as the NLM Visible Human datasets. PSC-VB II has many additional features needed for "gene knockout" research or, as shown here, studies of 3D motions in living animals. The image — from a micro-CAT scan analysis by NRBSC research partners at the Duke University Center for In Vivo Microscopy — shows a mouse heart. Producing three gigabytes of volume data, the scan captured eleven steps in the heartbeat cycle. Using PSC-VB II, researchers track the region of interest (blue) to study volume and other changes in the left ventricle.

TRACKING THE EVOLUTION OF GENES

Species evolve over time because their genes change. Some patches of DNA copy perfectly and others drop out or fail to replicate exactly, leading to changes in the molecular processes of the species. NRBSC biological chemist Hugh Nicholas specializes in computational methods to identify and track these changes and thereby to reconstruct the history of genes. In recent work, Nicholas is collaborating with Carnegie Mellon biologist Dannie Durand to trace the family tree for a large family of enzymes, called glutathione S-transferases, that help rid the body of toxins such as environmental pollutants. Better knowledge of this enzyme family could help in developing drugs to improve bioremediation efforts and chemotherapy treatment.

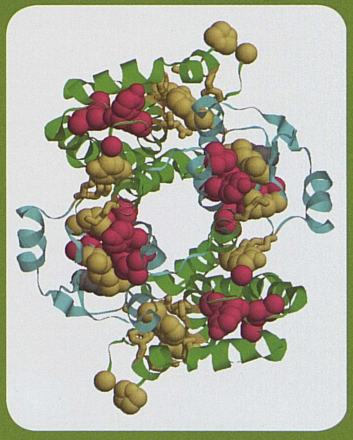

Networking The Future

ONE OF THE LEADING RESOURCES IN THE WORLD FOR NETWORK KNOW-HOW

PSC's Advanced Networking group is one of the leading resources in the world for knowledge about networking. Through 3ROX (Three Rivers Optical Exchange), a high-speed network hub, they provide high-performance networking for research and education. Their research on network performance and analysis — in previous projects such as Web100 and current work with HPN-SSH and pathdiag — has created valuable tools for improving network performance nationally.

>>More information: **http://www.psc.edu/networking**

Wendy Huntoon, PSC director of networking

HUNTOON DIRECTS OPERATIONS FOR NATIONAL LAMBDARAIL

In September 2005, PSC director of networking Wendy Huntoon was appointed director of operations for National LambdaRail, a major initiative of U.S. research universities and the private sector to provide infrastructure for research in networking technologies.

"The members of NLR will benefit immensely from Wendy's vast experience at the frontier of high-performance research and education networking," said Tom West, CEO of NLR. "With her longstanding leadership in the advanced networking community, she brings an extremely valuable set of skills to the task of finalizing the buildout of the NLR infrastructure."

NLR is deploying a nationwide optical network that provides up to 40 simultaneous light wavelengths, each capable of transmitting 10 gigabits per second. This infrastructure simultaneously supports cutting-edge network research and ambitious scientific projects that require guaranteed networking reliability and performance.

"It's a privilege to be able to help put this infrastructure in the hands of researchers who will create the next generation of networking technologies," said Huntoon. "As we improve the performance of our networks and experiment with new networking techniques, we'll be enabling a new breed of high-performance and network applications."

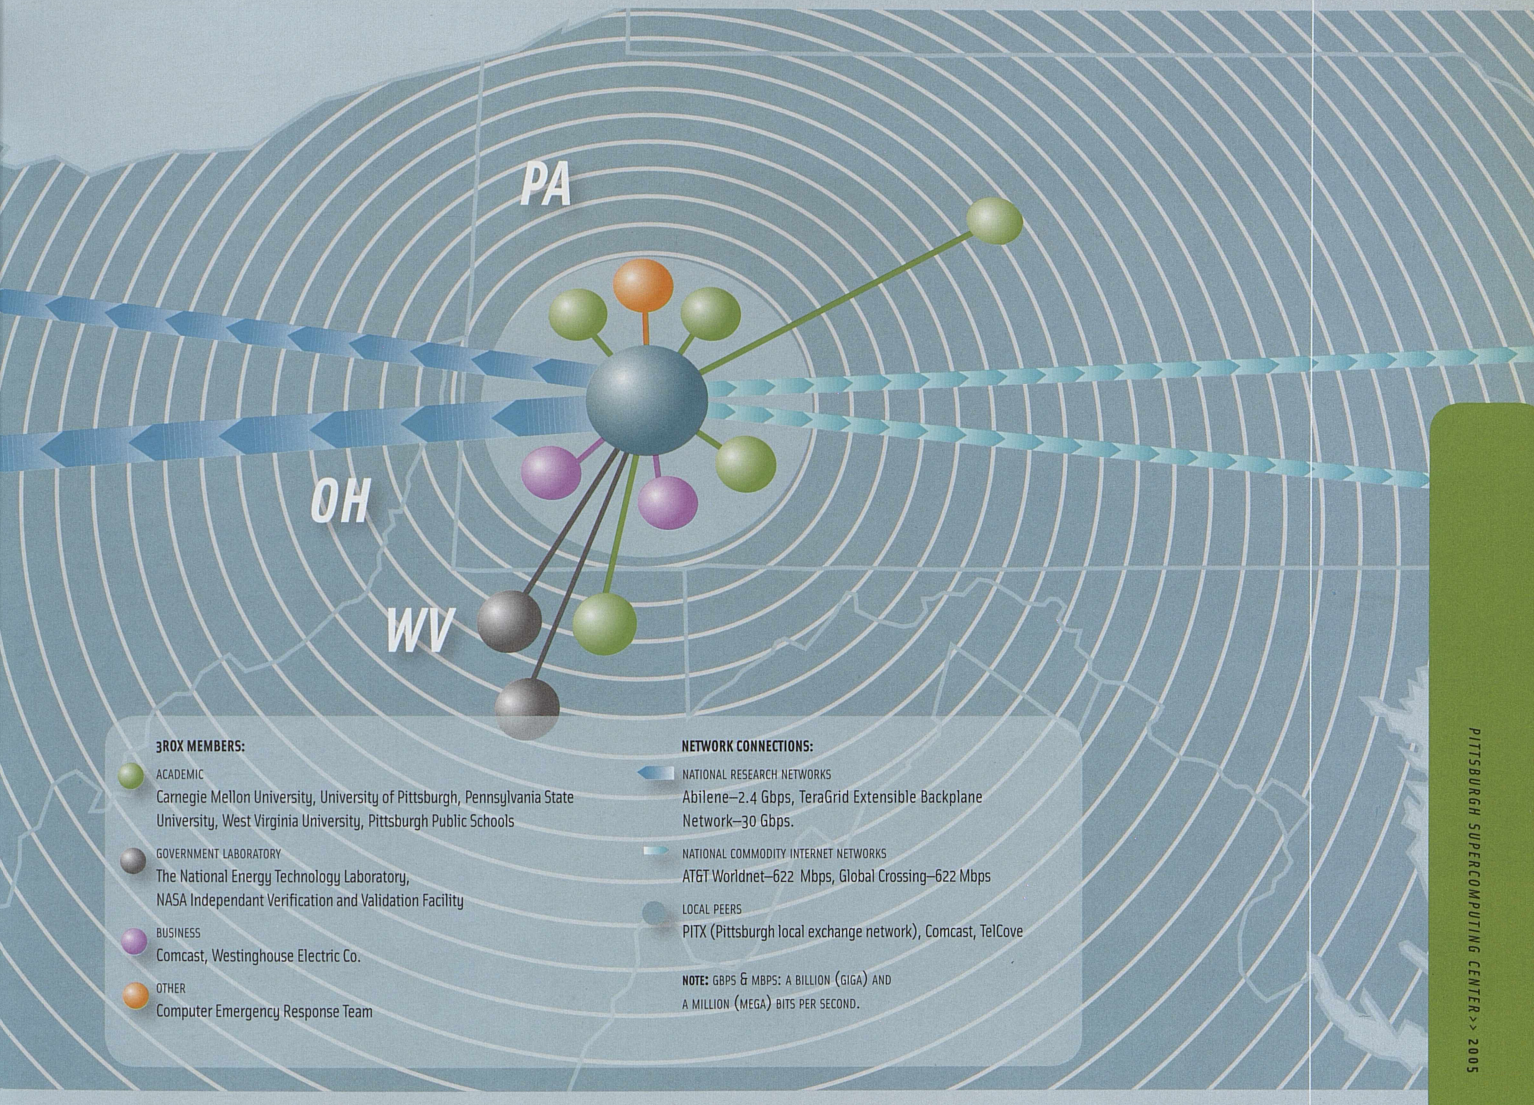

3ROX

Through 3ROX, a high-speed network hub that serves Carnegie Mellon, Penn State, the University of Pittsburgh, West Virginia University and the Pittsburgh Public Schools, PSC provides advanced network resources for education and research. 3ROX connects the universities and PPS to Abilene a high-performance network linking more than 250 U.S. universities and research organizations.

This year a new high-speed fiber linking TelCove, a major provider of telecommunication services, with 3ROX improved efficiency and speed of access for TelCove users in the Pittsburgh region. In July, 3ROX turned on the next-generation Internet protocol, called IPv6, with the commodity network Global Crossing, thereby becoming the first network hub to implement IPv6 on both research and commodity networks, a change which improves performance for 3ROX members connecting to IPv6 sites.

NOTES & HIGHLIGHTS, CONTINUED

NETWORKING RESEARCH & DEVELOPMENT AT PSC

UNPRECEDENTED CONNECTIVITY: Last November at Supercomputing '05 in Pittsburgh, PSC's network group accomplished a "first" in high-performance networking. Using two Cisco CRS-1 routing systems fitted with OC-768 interfaces, they demonstrated 40-Gigabit per second transmission over a single "lambda" — a light wave — in actual application processing.

The demonstration relied on PSC-developed protocol that coupled an application running on LeMieux to the showfloor network at SC '05. "This is the next step forward in networking," said PSC network director Wendy Huntoon. "We demonstrated PSC's ability to fully utilize this bandwidth with real application data."

ELIMINATING A NETWORK BOTTLENECK: SSH (Secure Shell) is a much used security-enhanced protocol for network communications, freely available for many operating systems and network products from Apple, Sun, HP, IBM, Cisco and Juniper. SSH, however, often limits the network performance that's potentially available on high-speed networks. Members of the PSC network group this year created a software patch to overcome this bottleneck. In many circumstances, this patch, called HPN-SSH, can give more than a 10-fold speedup in performance.

NASA has implemented this patch as part of their Solar Dynamics Observatory mission. "This is helping us tremendously," said NASA scientist Deepak Paul, "because it allows us to keep our real-time requirements and our security requirements. Previously, we could support only one or the other."

IMPROVED NETWORK DEBUGGING: Network performance debugging is difficult because most flaws show up with the same symptom: slowed performance. The nature of the problem causing slowed performance is often "masked" by self-correcting routines built into network protocol, and to identify the precise cause of the slowdown poses a difficult challenge. To deal with this, PSC engineers in collaboration with the National Center for Atmospheric Research (NCAR) developed innovative software — called pathdiag — that clearly and accurately diagnoses flaws and prescribes corrective action. This technology, which draws on tools PSC helped to develop in a project called Web100, is now in use at NCAR, the University of Michigan, Duke University and the Lawrence Berkeley National Laboratory.

PITTSBURGH SUPERCOMPUTING CENTER WORKSHOPS (2004-2005)

Introduction to the Cray XT3

Third Annual Workshop on Charm ++ and its Applications

Introduction to Computational Neuroscience

Introduction to Terascale Code Development

Developing Bioinformatics Programs

Parallel Programming Techniques

Nucleic Acid and Protein Sequence Analysis

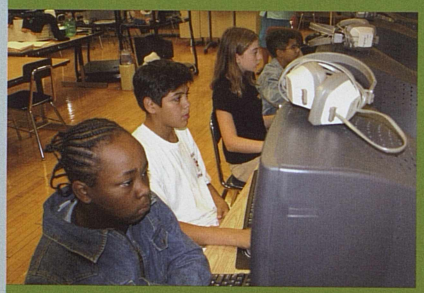

Students using the network at Frick International Studies Academy, Pittsburgh Public Schools

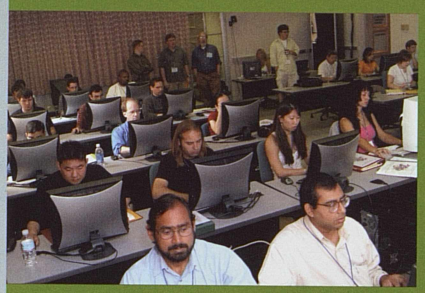

A workshop underway in the PSC Computer Training Center, equipped with 30 "dual-boot" workstations and a projector for overhead display of the instructor's desktop.

The PSC operational management team: (front, l to r) **David Kapcin**, manager, financial affairs; **Nick Nystrom**, director, strategic applications; **Rich Raymond**, manager, user support; **Sergiu Sanielevici**, director, scientific applications and user support; **Gwendolyn Huntoon**, director, networking; (back, l to r) **John Kochmar**, HPC facilities manager, systems and operations; **Joel Stiles**, director, Center for Quantitative Biological Simulation, National Resource for Biological Supercomputing; **J. Ray Scott**, director, systems and operations; **Bob Stock**, associate director. Not in photo: **Janet Brown**, manager, networking; **David Deerfield**, director, Biomedical Initiative; **Elvira Prologo**, manager, administration.

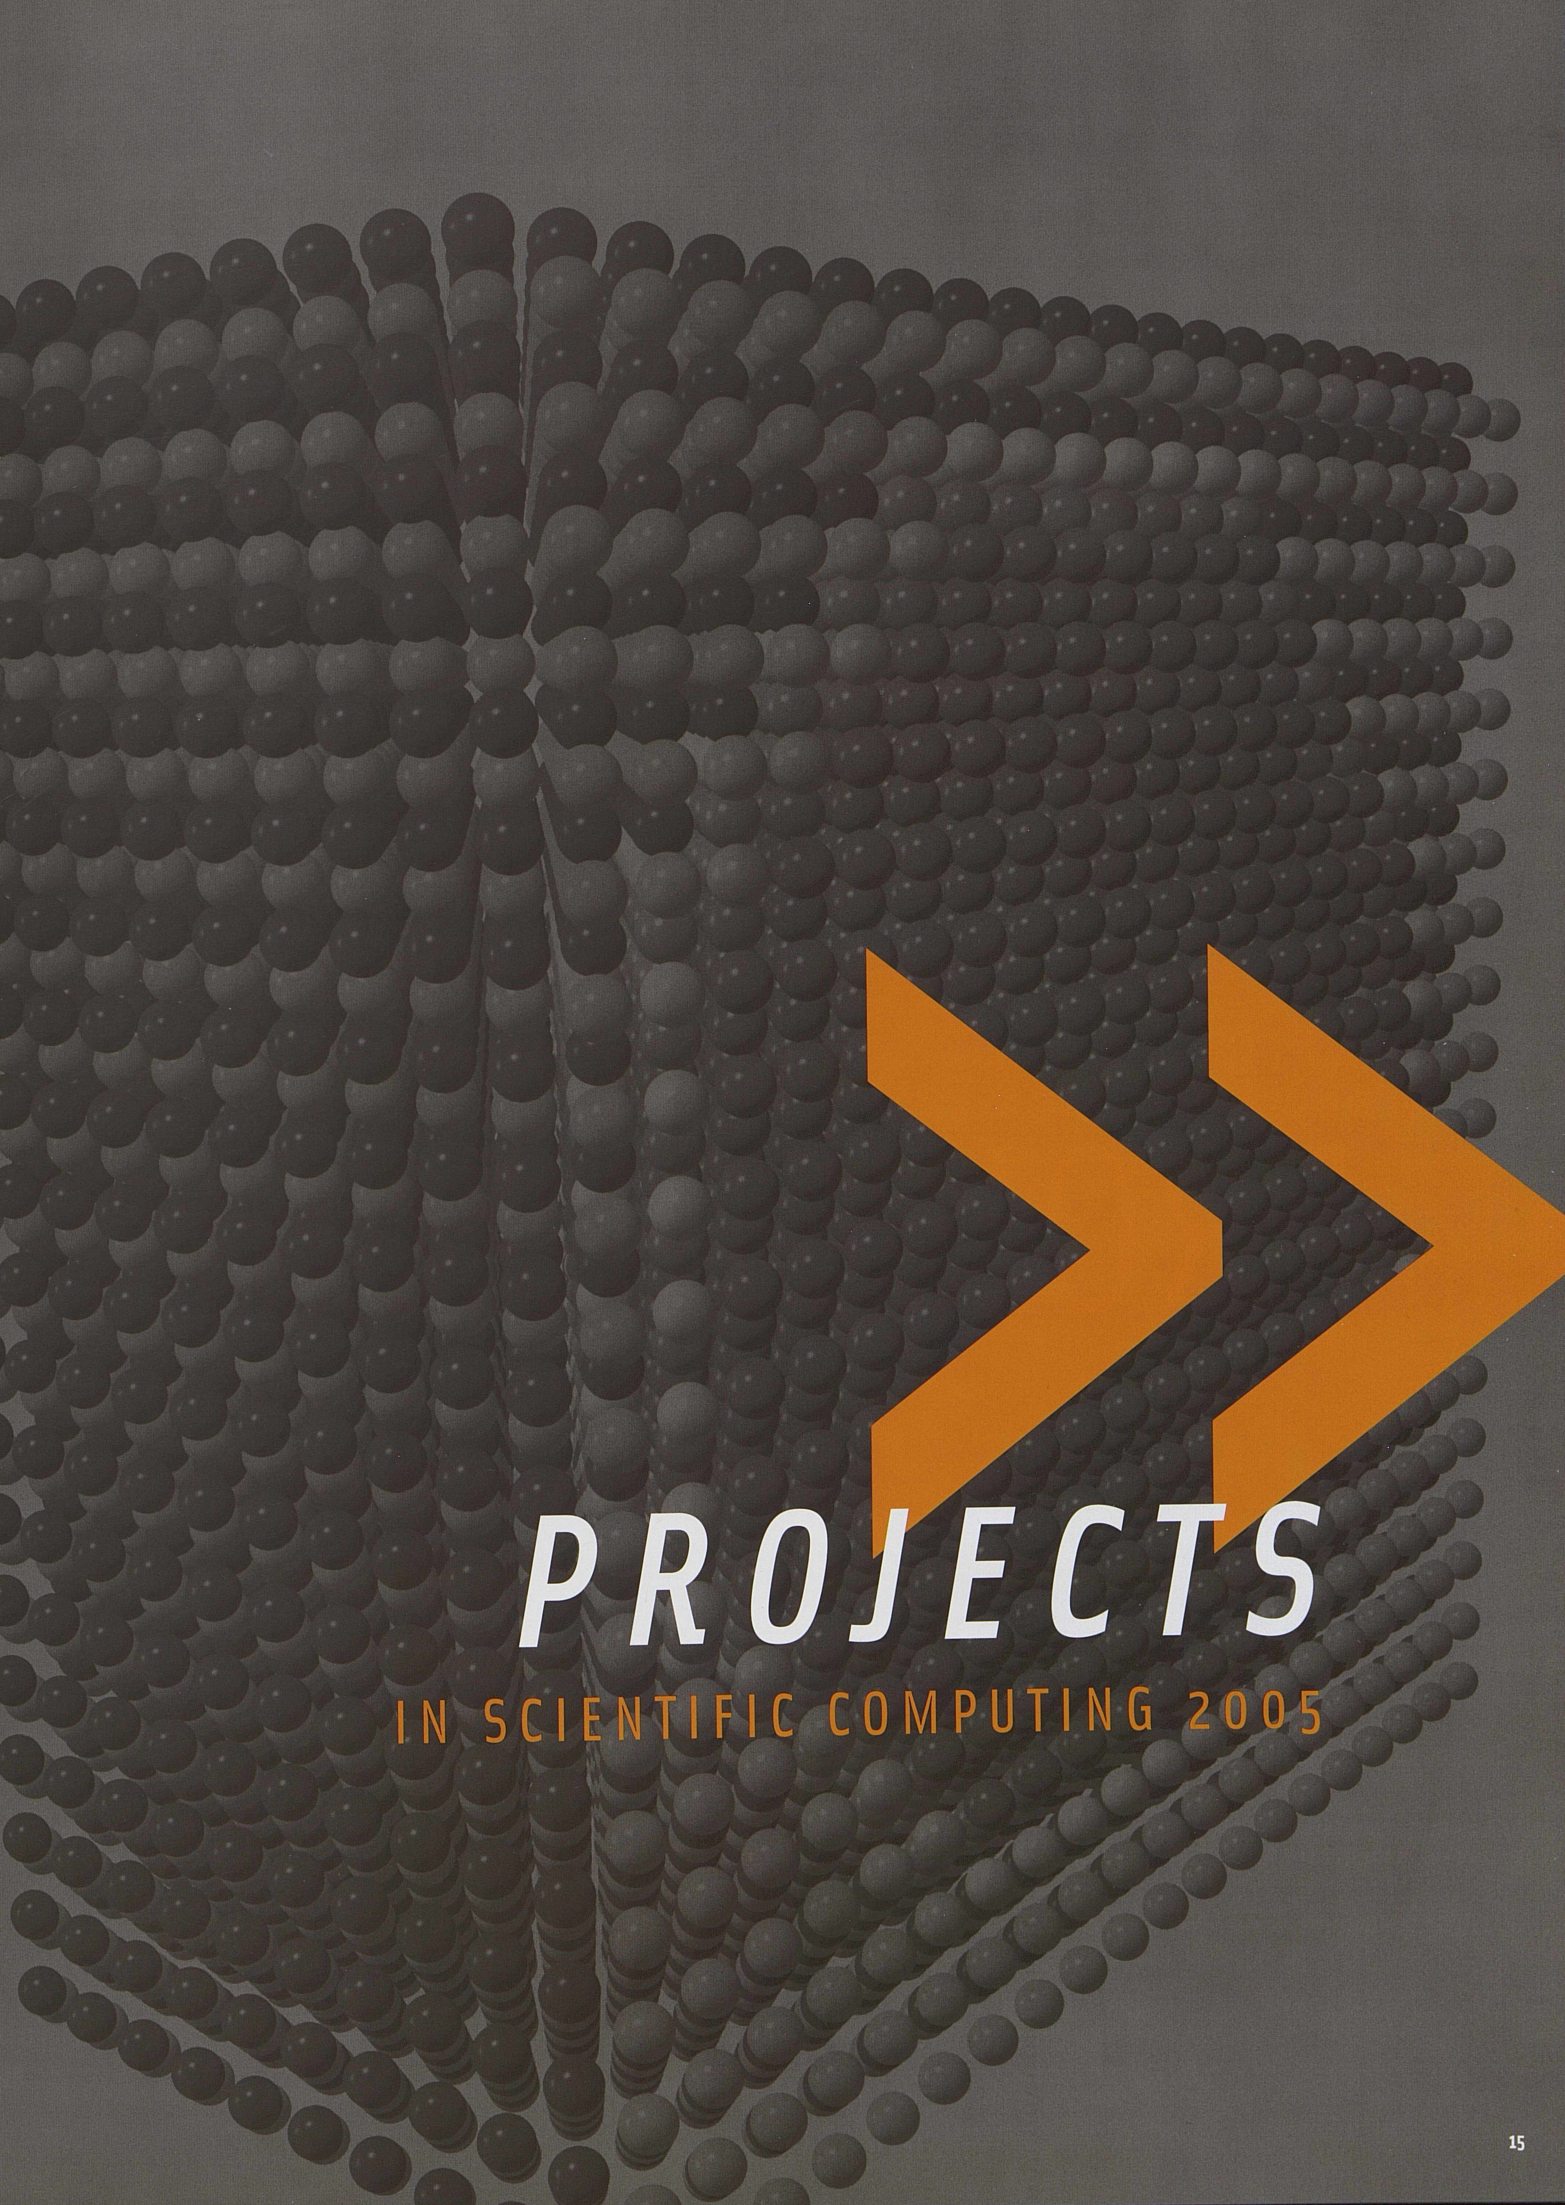

PROJECTS

IN SCIENTIFIC COMPUTING 2005

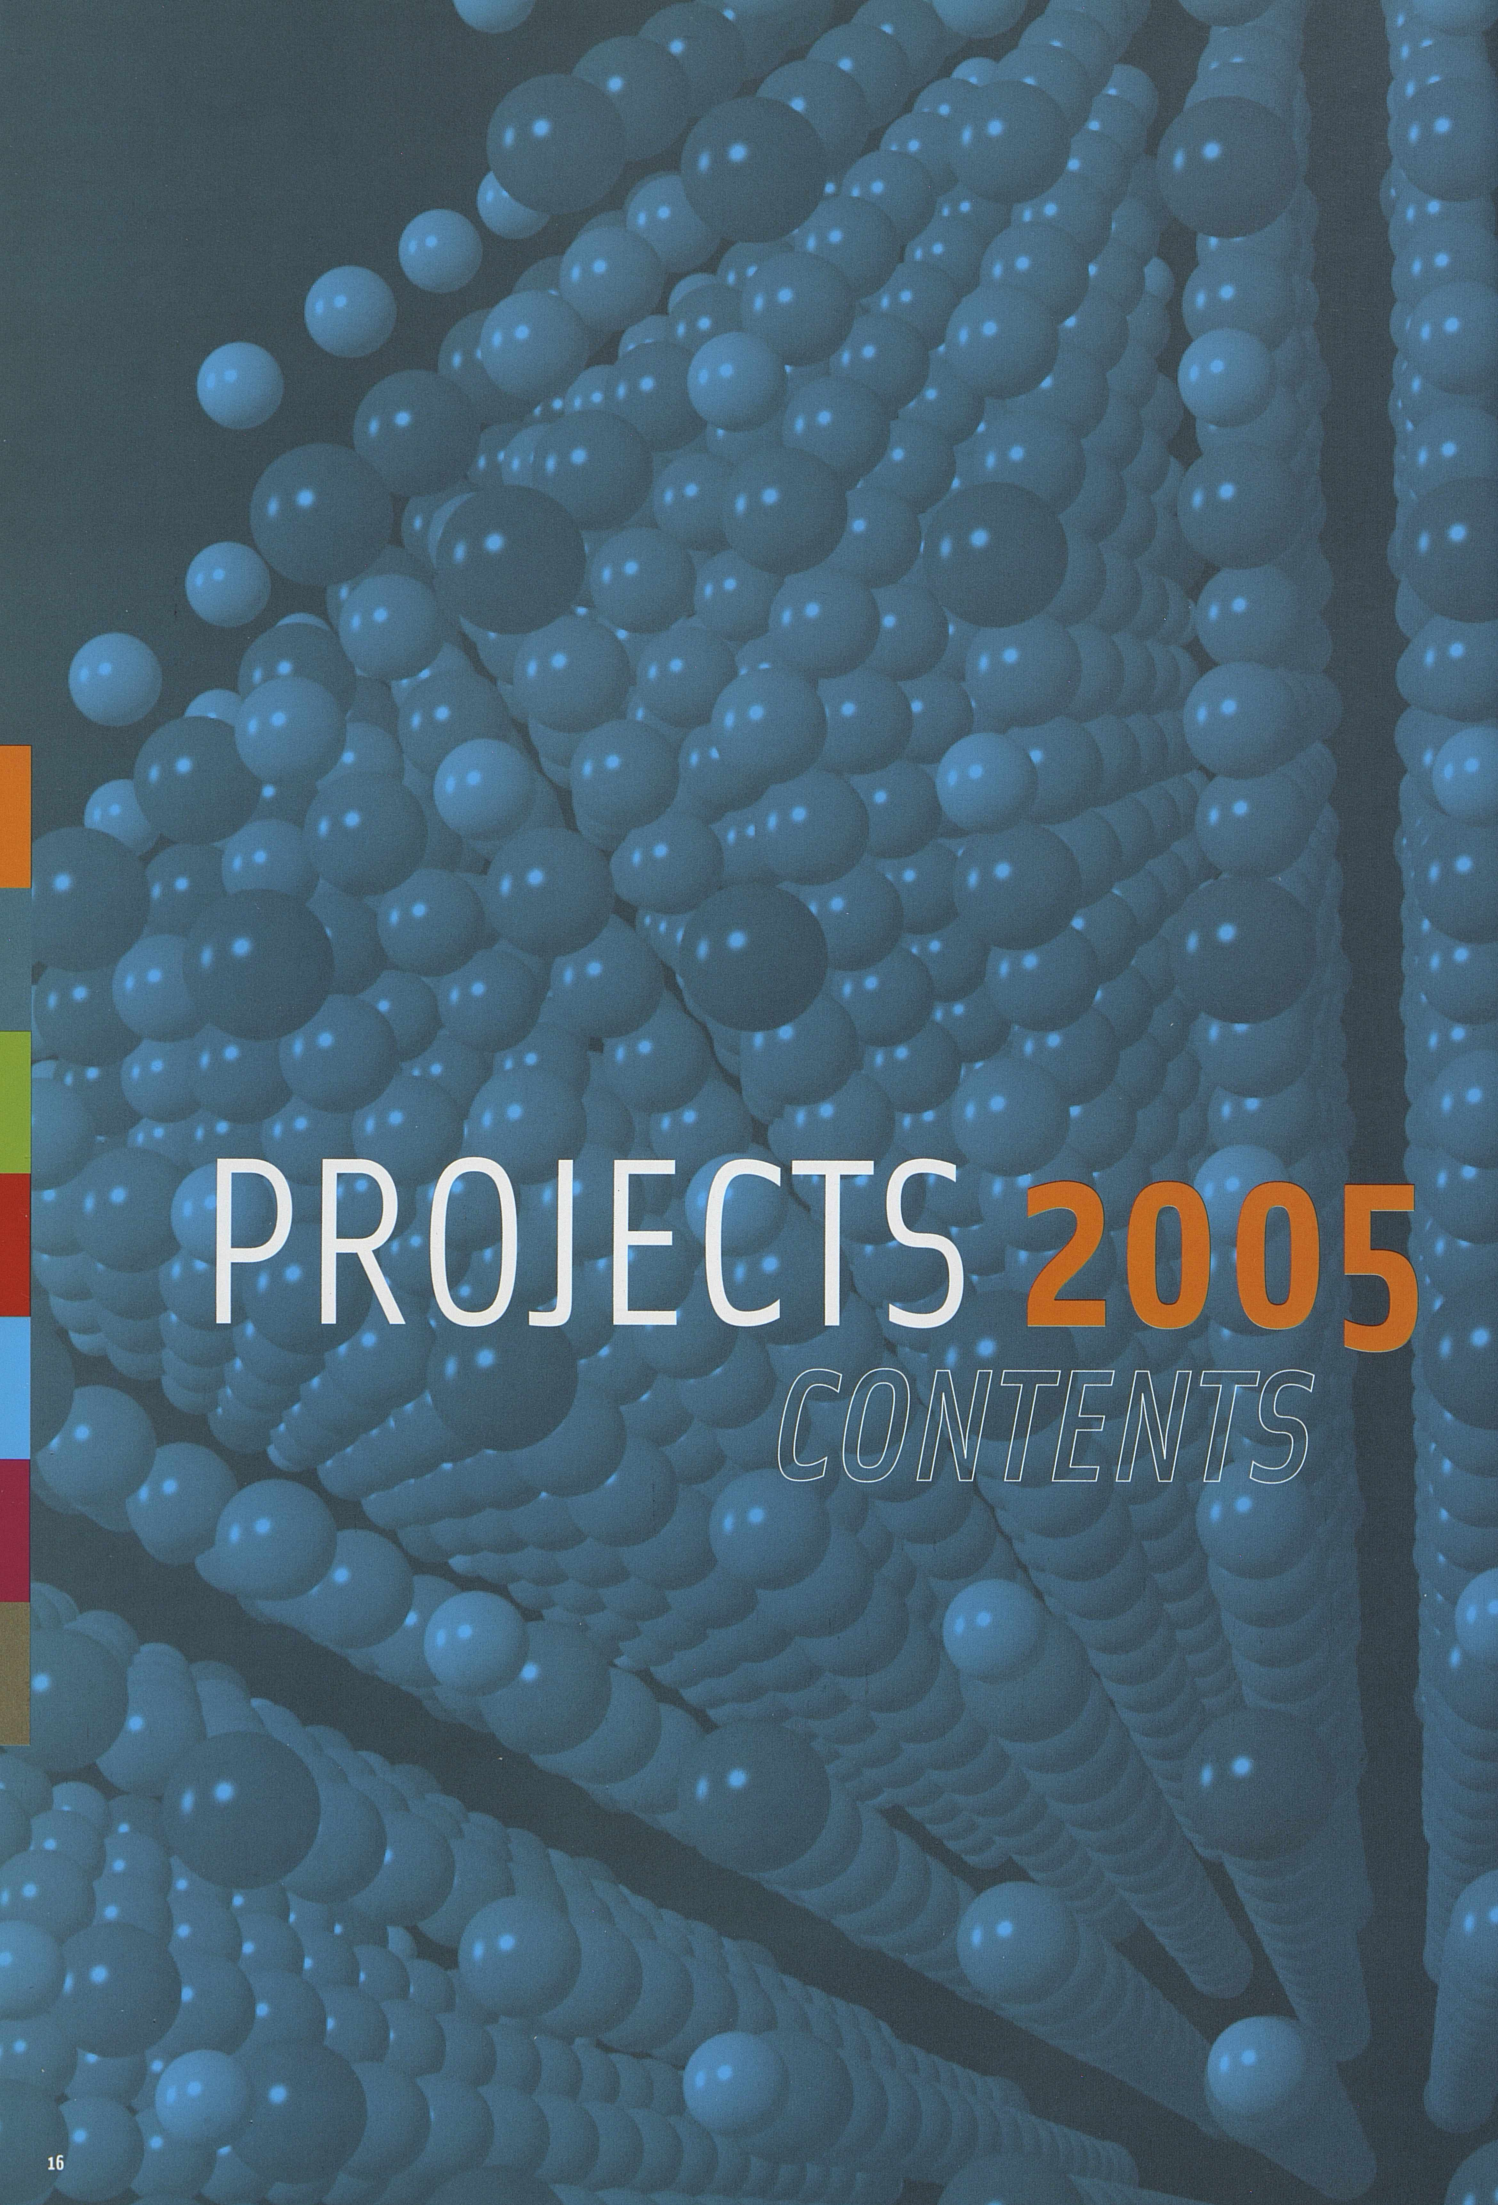

PROJECTS 2005

CONTENTS

ROLLING THUNDER PREVIEW

THIS SPRING, FOR THE FIRST TIME, FORECASTS CORRECTLY PREDICTED THE DETAILS OF THUNDERSTORMS 24 HOURS IN ADVANCE

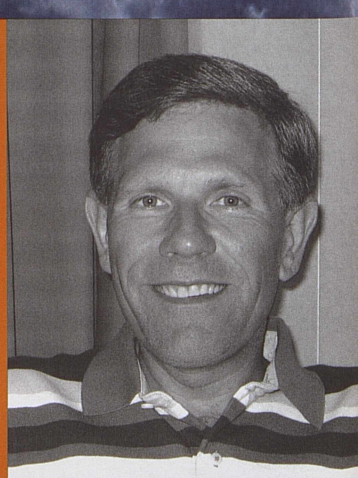

Kelvin Droegemeier, University of Oklahoma, Norman. "For nearly 20 years," says Droegemeier, "PSC has provided outstanding personalized support to large and challenging projects such as this spring program."

If anything is certain in 2005 — not counting death and taxes — it's that we're at the mercy of forces we don't control. Despite incredible advances in understanding nature — leading to amazing technologies our forebears couldn't imagine — our planet still unleashes furious energies that devastate communities and lives.

Even before Katrina, the U.S. loss due to extreme weather — hurricanes, floods, winter storms, tornados — averaged $13 billion annually. The human cost — nearly 1,000 people each year — is incalculable. Would better forecasting make a difference? No doubt. More to the point, is better forecasting possible?

You bet, says Kelvin Droegemeier, who directs the Center for Analysis and Prediction of Storms (CAPS) at the University of Oklahoma, Norman. Take thunderstorms, the nasty ones, with rotating updrafts — called supercells. They surge across the Great Plains each spring with the potential to spawn deadly tornados. How much would it be worth to know six hours in advance — instead of, as with current forecasting, a half-hour to an hour — that one of these storms is headed your way? And to have not just an ambiguous "storm warning" but precise information about when and where it will hit, how severe it is, how long it will last?

"We want to be able to say," says Droegemeier, "that over Pittsburgh this afternoon at 3:30 there'll be a thunderstorm with 30 mile-per-hour wind, golfball-sized hail, two-and-a-half inches of rain, and it will last 10 minutes, and to give you that forecast six hours in advance."

Since the 90s, CAPS has taken several strides toward demonstrating that, with sufficient resources in data-gathering and computing, it's possible to do this. This spring, they took another stride. In a major one-of-a-kind collaboration with NOAA (the National Oceanic and Atmospheric Administration), CAPS used resources of the NSF TeraGrid, in particular LeMieux, PSC's terascale system, to produce the highest-resolution storm forecasts yet attempted. On several occasions, CAPS predicted the occurrence of storms within 20 miles and 30 minutes of where and when they actually happened, and they did it 24 hours in advance.

"That type of result," says Droegemeier, "pretty much sets conventional thinking on its ear."

Are Thunderstorms Predictable?

In contrast to the daily weather reports on TV, which are generated from large-scale models that predict atmospheric structure over the continental United States, storm-scale forecasting involves a tighter focus — at the scale of a county or city. It requires observational data — temperature, pressure, humidity, wind speed and direction, and other variables — at a corresponding finer spatial resolution, and it demands the most powerful computing available, and then some, to run the models.

When CAPS began in 1988, the prevailing view about storm-scale forecasting was skepticism. Numerical weather prediction was not in question. Computers programmed with equations that represent the atmosphere and initialized with observational data were proven — since the 1970s — to be the best way, by far, to forecast weather. The question was more fundamental. Are thunderstorms predictable?

"The challenge we set ourselves to," says Droegemeier, "was, if you take the concept of computer forecast technology and apply it at this smaller scale, does the atmosphere possess any intrinsic fundamental predictability, or is at all turbulence? We had hopes, but we didn't know. With big help from the Pittsburgh Supercomputing Center, we resolved that question."

CAPS developed groundbreaking new techniques to gather atmospheric data from Doppler radar and to assimilate this data with other meteorological information. And they developed a computational model that uses this data to predict weather at thunderstorm scale.

"It all starts with observations," says Droegemeier, "because to predict we need to know what's going on right now." Data to feed weather models comes from many sources — upper air balloons, the national Doppler radar network, satellites, sensing systems on commercial airplanes. From these sources, a huge amount of information, computationally processed and spread across a 3D grid representing the atmosphere, becomes the initial conditions for National Weather Service forecasts. With grid spacing at 10 to 30 kilometers, the NWS operational models do well at showing high and low pressure areas and storm fronts that develop from them — weather that happens, roughly speaking, on the scale of states. Individual thunderstorms originate at smaller scales and to forecast them, says Droegemeier, requires much finer spacing, down to at least one to two kilometers.

The foundation of a storm-forecast model is 15 to 20 non-linear differential equations. They represent

"FROM A MODELING, COMPUTATIONAL AND COMMUNICATIONS PERSPECTIVE, THIS HAS NEVER BEEN DONE BEFORE."

the physical phenomena of the atmosphere and how it interacts with the surface of the Earth. To make a forecast involves feeding the 3D grid with initializing data, solving these equations at each position on the grid, and then doing it over again for the next time step, every five to ten seconds for 24 hours of weather. For a single forecast, this means solving trillions of equations. Each doubling of the number of grid points in 3D requires eight times more computing. If you also halve the time step — to capture corresponding finer detail in time — it's a 16-fold computing increase. For this reason and others, storm-scale forecasting poses an enormous computational challenge.

Since 1993, CAPS has run forecasting experiments during spring storm season. In 1995 and '96, using PSC's CRAY T3D — a leading-edge system at the time — for a limited region of the Great Plains, they successfully forecast location, structure and timing of individual storms six hours in advance — a forecasting milestone. For this accomplishment, CAPS and PSC won the 1997 Computerworld-Smithsonian award for science, and CAPS garnered a 1997 Discover Magazine award for technological innovation.

If there were lingering doubts about storm forecasting, that the question has shifted from scientific and technological feasibility to national policy — whether sufficient resources can be made available and when — this spring's storm forecast experiment should erase them.

Watershed Forecasts

As they have during many storm seasons over the past dozen years, CAPS and PSC this spring collaborated to produce real-time storm forecasts. The difference this year was that the forecasts covered two-thirds of the continental United States — from the Rockies east to the Appalachians. Using LeMieux, they successfully produced an on-time, daily forecast from mid-April through early June. "This was an unprecedented experiment," says Droegemeier, "that meteorologists could only dream of several years ago."

Conducted in collaboration with NOAA, the program included about 60 weather researchers and forecasters from several NOAA organizations — the Storm Prediction Center and the National Severe Storms Laboratory, both in Norman, and the Environmental Modeling Center in Maryland — and the NSF-sponsored National Center for Atmospheric Research in Boulder, Colorado along with CAPS.

This experiment offered an unprecedented chance for forecasters — as well as researchers — to work with advanced technology on a daily basis, technology that, says Droegemeier, may be five years from being incorporated in daily forecast operations at the resolutions used. Each evening, meteorologists in Norman transmitted new atmospheric conditions to Pittsburgh. By the next morning, LeMieux produced a forecast that covered the next 30 hours, and transmitted the forecast back to SPC and NSSL in Norman, where researchers turned the model output-data into images corresponding to what they see on radar. These model runs were conducted daily with virtually no problems.

Using several different versions of the Weather Research and Forecasting Model, an advanced model designed for research as well as operational use, the partners generated forecasts three times daily. EMC and NCAR used grid spacing of from four to 4.5 kilometers. With LeMieux at its disposal, running on 1,228 processors, CAPS went a step further. With grid spacing of two kilometers, more than five times finer than the most sophisticated NWS operational model — and requiring 300 times more raw computing power — their forecasts are the highest-resolution storm forecasts to date.

"Our daily WRF-model forecasts," said Droegemeier, "had twice the horizontal resolution and nearly 50-percent greater vertical resolution than the other two experimental products." This higher resolution meant that the forecasts were able to capture individual thunderstorms, including their rotation. On several occasions, when the 24-hour forecast showed development of thunderstorms, it proved to be accurate within 20 miles and 30 minutes.

Just as importantly, the computer model produced images that matched well in structure with what forecasters saw later on radar. "The computer forecasts looked very similar to what we see on radar," said Steven Weiss, SPC science and operations officer. "The structure you see on the screen is important in judging whether the storm is likely to produce tornadoes, hail or dangerous wind. These results were an eye-opener in many respects."

"Real time daily forecasts over such a large area and with such high spatial resolution," says Droegemeier, "have never been attempted before, and these results suggest that the atmosphere may be fundamentally more predictable at the scale of individual storms and especially organized storm systems than previously thought." Such results could lead potentially, he adds, to a revision of classical predictability theory put forth by Edward Lorenz, the now retired MIT professor, whose pioneering research led to chaos theory. Spring 2005 — the forecasting community is still absorbing the findings, but it may mark a watershed in the understanding of atmospheric predictability. (MS)

>>More information: **http://www.psc.edu/science/2005/storms**

Next Steps in Storm Forecasting

by Steven Weiss
Science and Operations Officer
NOAA Storm Prediction Center

The accomplishments of CAPS, PSC and NCAR this spring in running the two-kilometer model on a reliable daily basis were spectacular. From a modeling, computational and communications perspective, this has never been done before, and we have now crossed a major hurdle, namely, to show that it is possible to run very high-resolution, large-domain prediction models for daily forecast purposes.

Of course, considerable research remains to be done, including finding ways to better define initial conditions, and to learn how forecasters can interpret and use new high-resolution model output. The latter is of particular interest at the SPC. How can we know ahead of time when to believe the remarkable details produced by the high-resolution model; when to make adjustments to them or discount them? There is potential that very high-resolution ensemble prediction systems may help us in that area. As is commonly the case in groundbreaking science and technology, some very important questions get answered, but the new knowledge and capabilities also result in new issues to address.

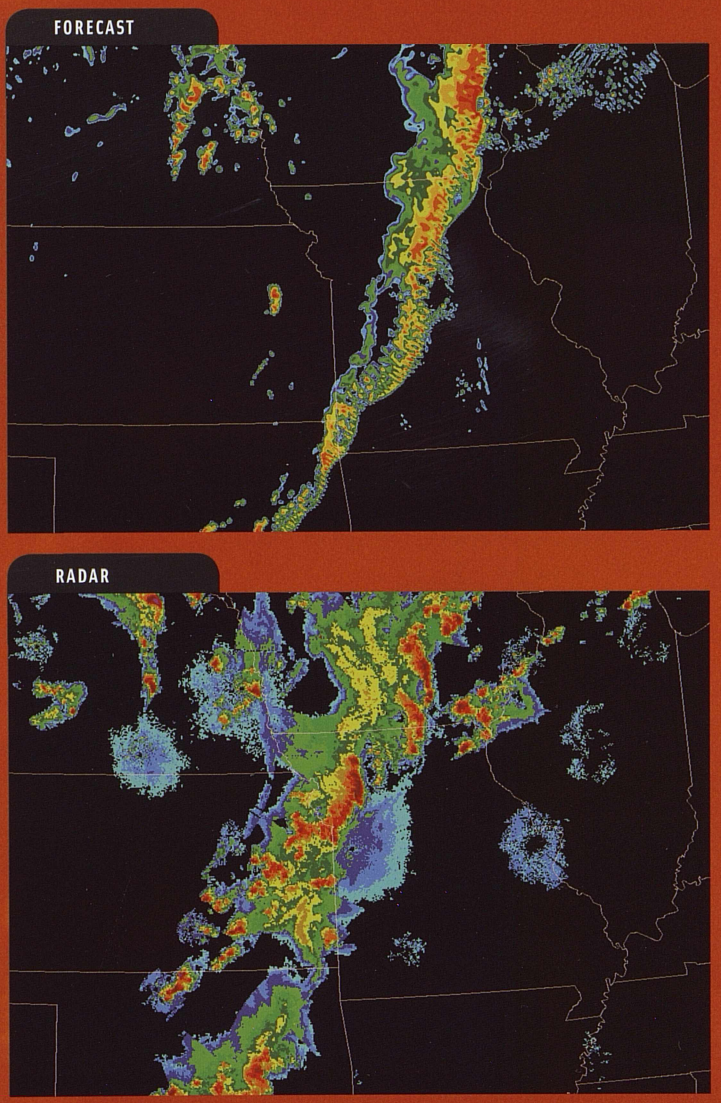

Predicting Storms 24 Hours in Advance

The CAPS computer forecast for the night of June 4 & 5, 2005 (run at PSC on June 3) predicted a storm line moving from Kansas into Missouri extending into Iowa. This forecast [above] for June 4, 10 p.m. CDT, compares well in structure and placement of the developing storms with the actual radar. Both images show radar reflectivity (increasing from blue through red), which is proportional to precipitation intensity. The smaller frames [below] show forecast vs. radar over a time span from 2 p.m. June 4 until 1 a.m. June 5. All the images shown here represent only a sub-domain of the complete Rockies to Appalachian computational domain.

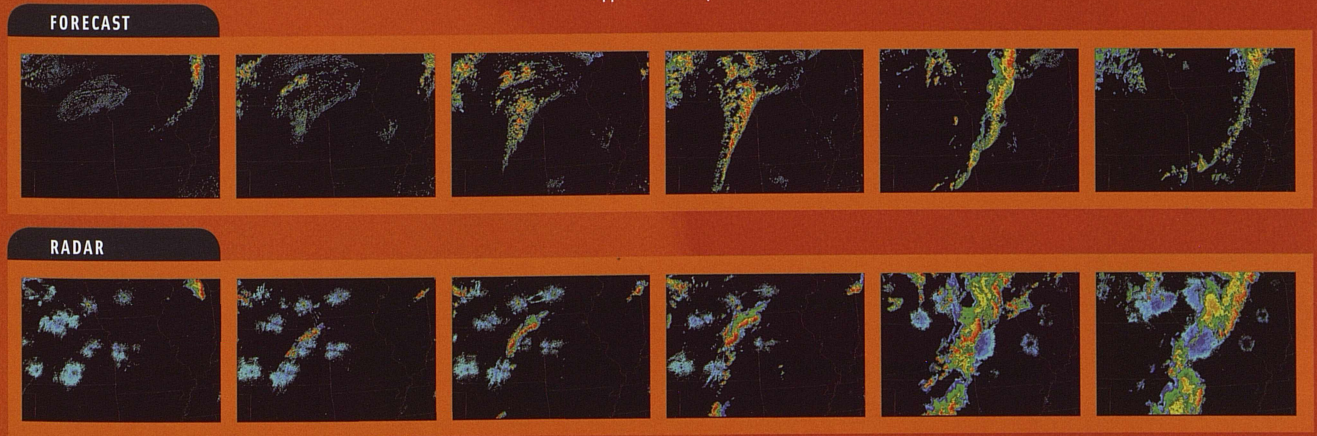

SIMULATIONS OF NEURAL SYNAPSES OVERTURN SOME LONGSTANDING ASSUMPTIONS

FIRE IN THE BRAIN

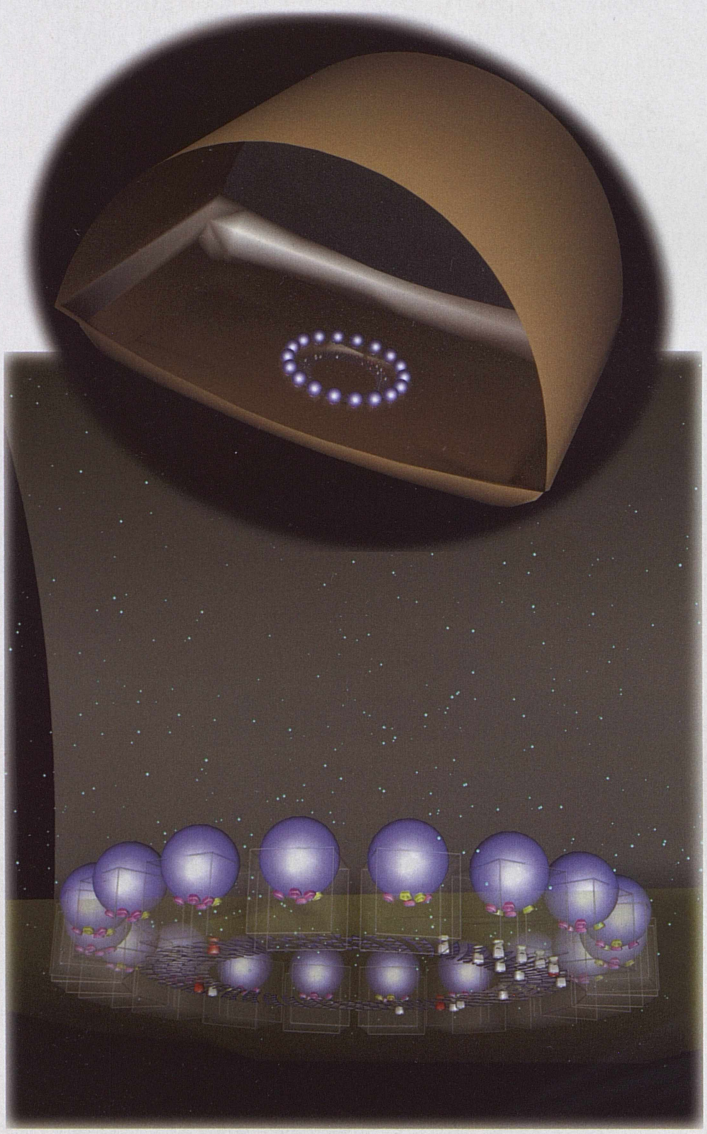

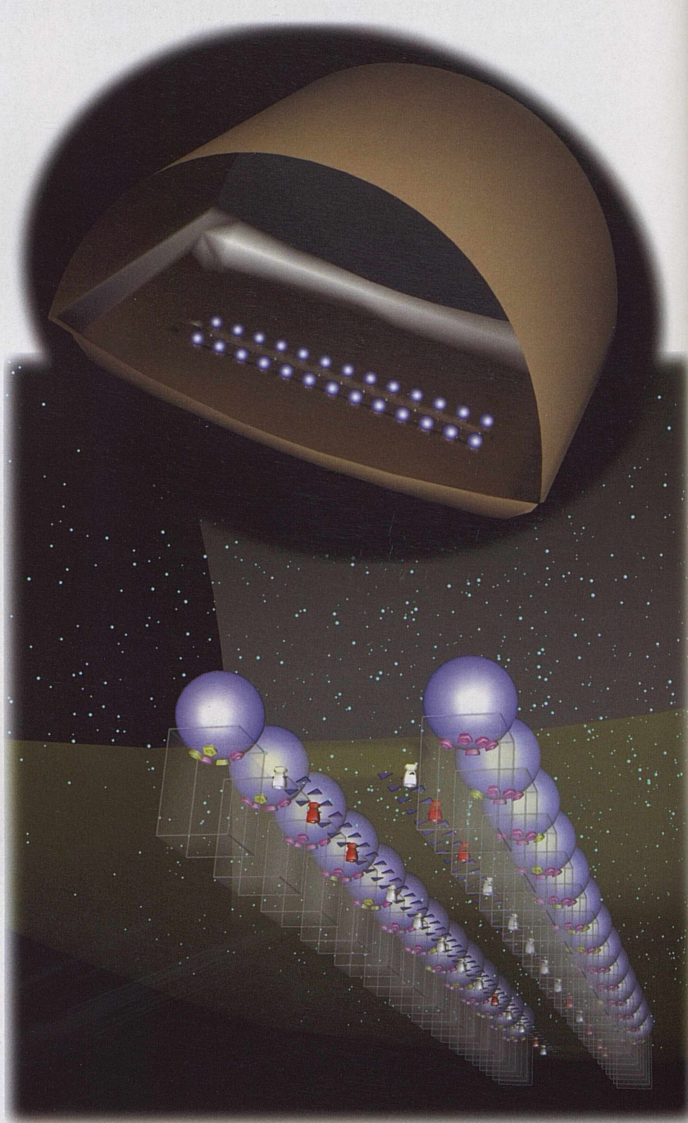

Each time Sergei Rachmaninoff sat down to perform his virtuosic Piano Concerto No. 3, he played more than 28,000 notes in about 30 minutes, his fingers, a blur of muscular energy, striking the keys with the precise force — at the precise instant — the maestro intended. With only faint electrochemical signals running through nerves of greatly varied length, how is it possible for the brain to initiate and control such rapid-fire movement, dozens of muscles simultaneously, with exquisite sensitivity and split-second timing?

Over the past century, scientists have taken many steps toward understanding the biology underlying such feats. The synapse — where nerve cells meet or connect to a muscle fiber — amplifies the tiny voltage generated in the brain into forces large enough to pound out a thunderous finale. Still, these intricately coordinated mind-body processes hold many mysteries.

PSC senior scientist Joel Stiles and his collaborators solved a few recently by using a supercomputing system to recreate the business end of a synapse. Their tools included MCell, powerful software — co-authored by Stiles and Thomas Bartol of the Salk Institute — for simulating the microphysiology of interacting cells, coupled with DReAMM visualization software developed in Stiles' lab at PSC. The computational work relied heavily on Jonas, PSC's 128-processor shared-memory HP system dedicated to biomedical research. Stiles and colleagues modeled two kinds of synapses in unprecedented detail, and their findings, published recently in *Science* and other journals, overturn some longstanding assumptions about neural communication, offer insight into a family of crippling diseases, and demonstrate the power of computation allied with experimental measurement.

A Barrage of Neurotransmitters

The secret to coordination, says Stiles, is predictability. As a child grows, the motor cortex in the brain develops increasingly sophisticated programs of voluntary muscle control. These programs must anticipate delays that occur as a command impulse travels to the spinal cord, out a peripheral nerve, and across the nanometers-wide synaptic cleft that separates a neuron from the muscle fiber it innervates.

Decades ago, biologists measured the signal delays introduced by nerve-muscle synapses, and found that the lag time is amazingly consistent, varying less than 30 millionths of a second from one synaptic firing to another. How does the body manage this feat of consistency?

Over the years, molecular biologists have filled-in much of the story. On its way from brain to limb, an impulse shoots down the nerve cell to its terminus and arrives in the form of a rapid change in sodium ion concentration. That creates a voltage. The voltage triggers thousands of sphincter-like proteins embedded in the cell wall of the neuron; these protein channels open, and some allow calcium ions to flow into the cell. As the calcium ions enter, they diffuse throughout the interior of the neuron, bumping into tiny spherical containers — called vesicles — arranged in neat rows close to the interior wall.

Loaded with the neurotransmitter acetylcholine, the vesicles are like a battery of fireworks poised to launch a neurotransmitter barrage into the synaptic cleft — where this amplified signal will in turn energize the adjacent muscle fiber. In the last step, what biologists call a "fusion event," the calcium ions trigger some of the vesicles to fuse with the cell membrane and open outward, allowing their payload to escape.

The details of this critical last step, however, have so far stumped neuroscientists. "We still don't know how calcium binds to receptor areas on the vesicles to lead to this fusion event," says Stiles. "Nor do we know how many protein channels need to open to admit the calcium, or how far the calcium can diffuse to find the vesicles and cause them to fuse. Does a calcium channel have to open right next to a vesicle, or can the ions come in through a variety of different gates and collectively trigger a variety of vesicles?"

<<Two Active Zone Architectures

Cutaway views into the interior of a nerve cell show low (top) and high magnification views of two different active-zone configurations simulated with MCell by Stiles and colleagues. The simulations include diffusing calcium ions (small blue dots) and synaptic vesicles (large purplish spheres) with calcium binding sites (red & yellow) and calcium channels (color-coded glyphs in blue triangles) in the nerve membrane. This model has a small number of binding sites on each vesicle. Translucent boxes are used to count calcium ions in different regions of space. (DReAMM image by Joel Stiles, PSC).

Joel Stiles, Pittsburgh Supercomputing Center, and **Stephen Meriney**, University of Pittsburgh.

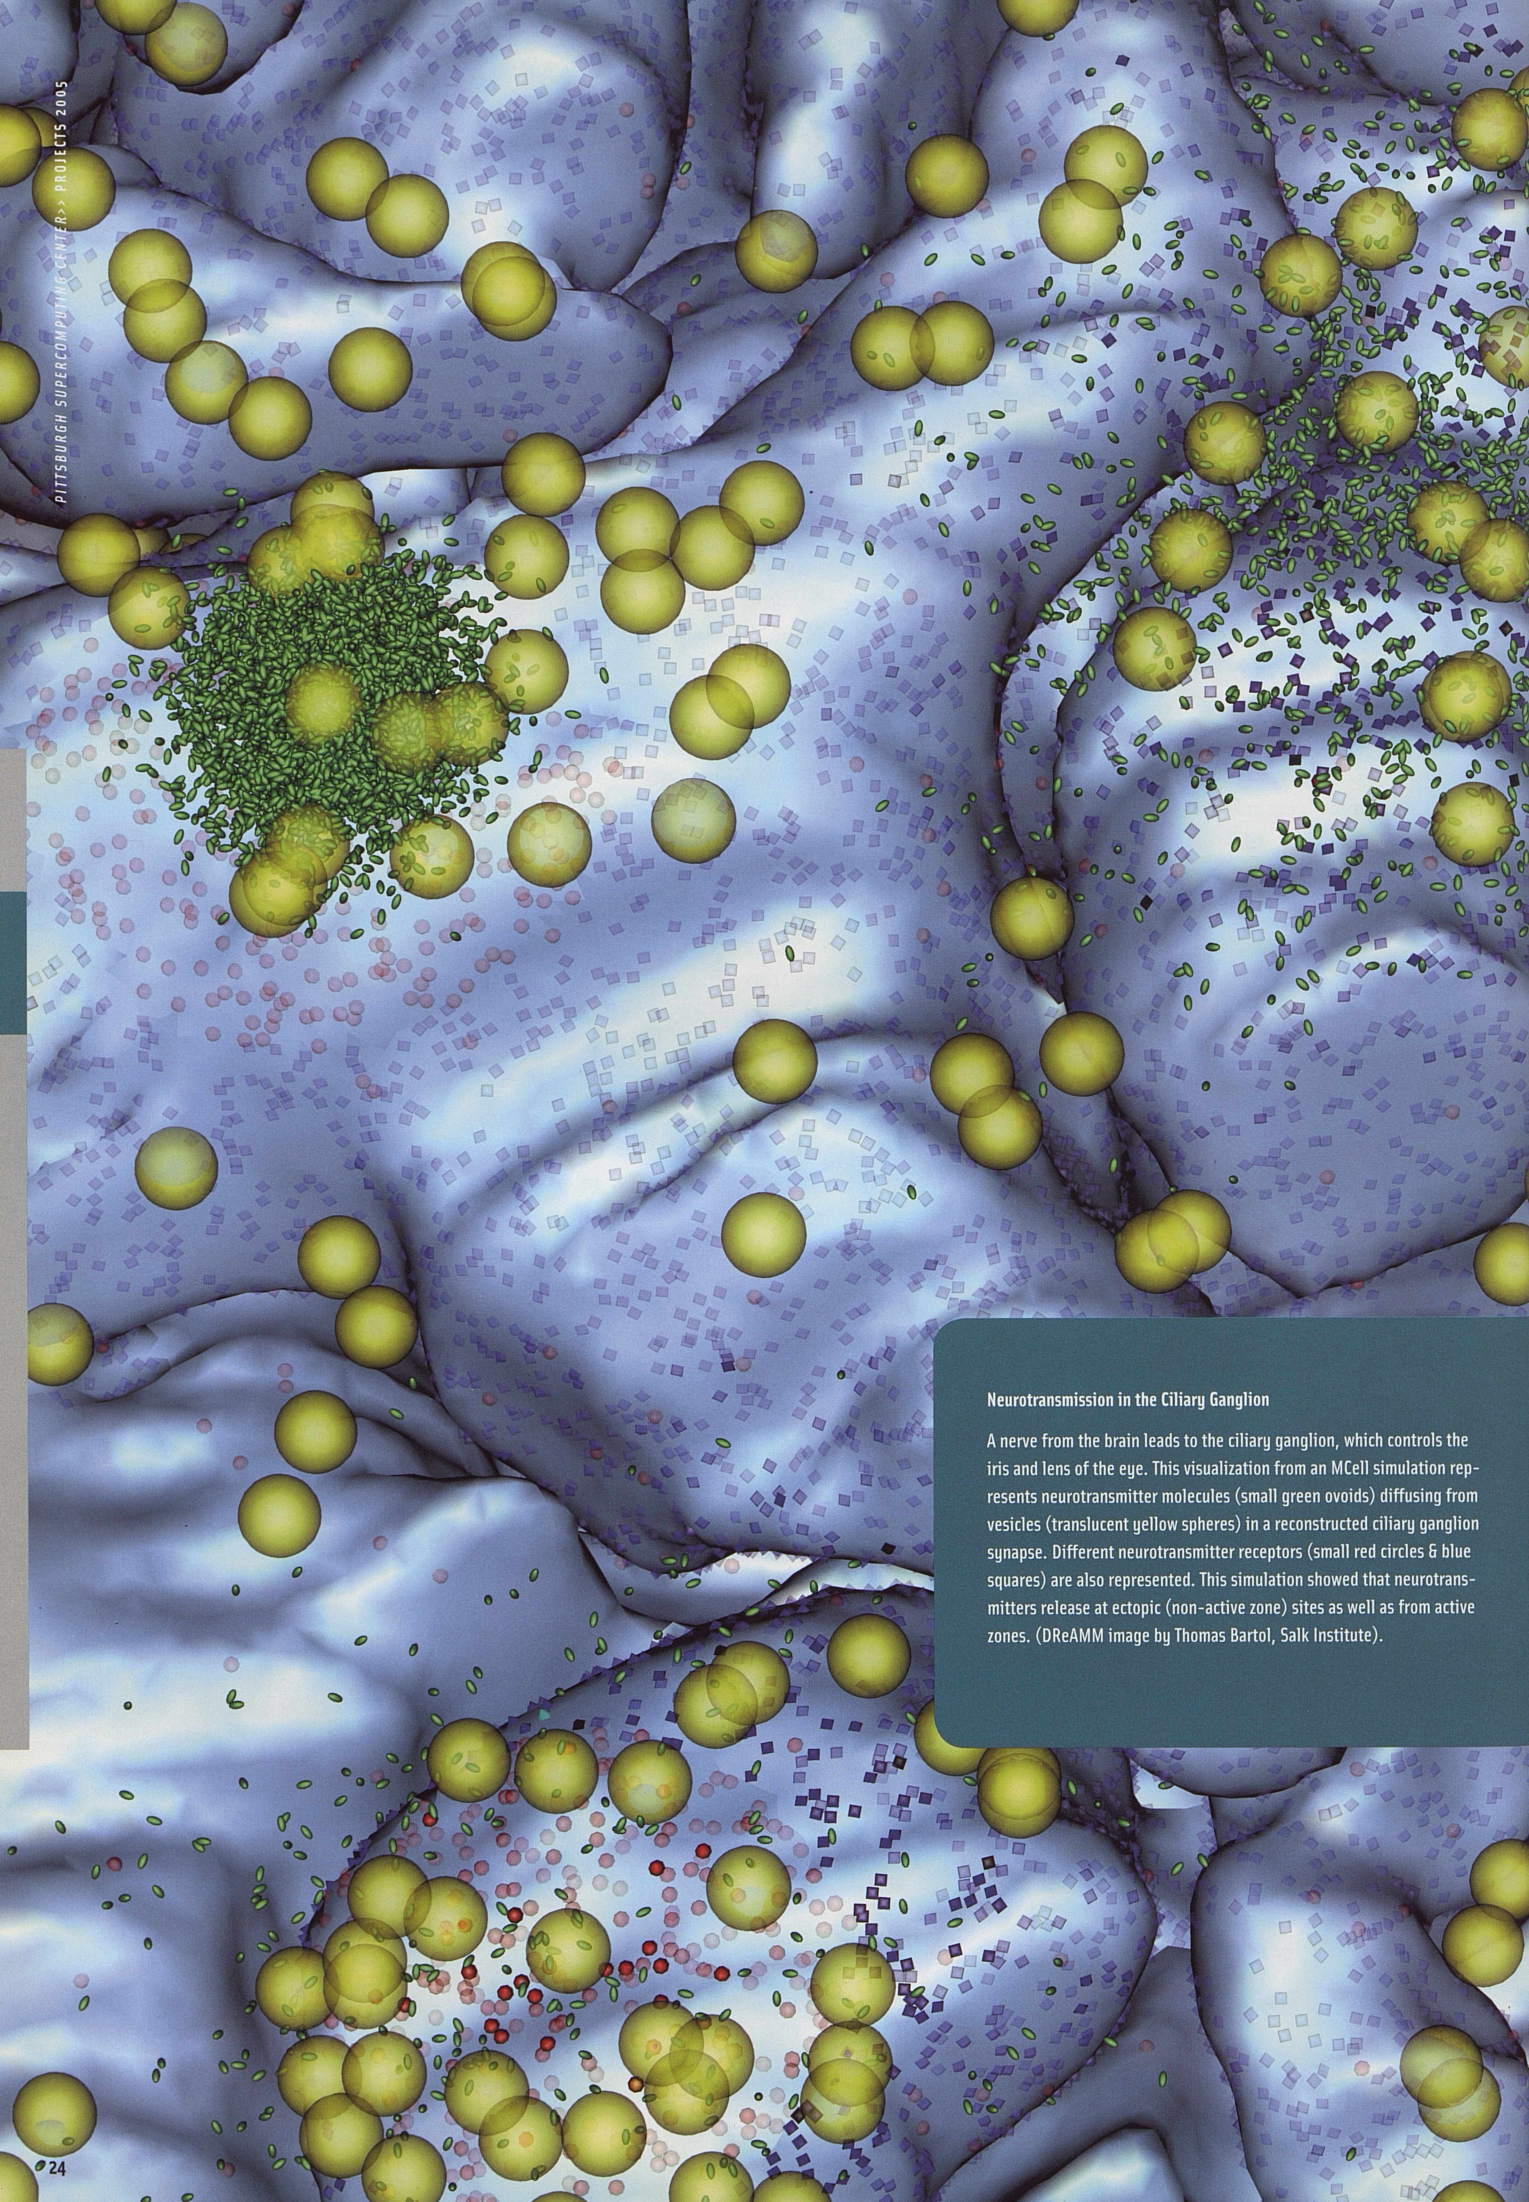

Neurotransmission in the Ciliary Ganglion

A nerve from the brain leads to the ciliary ganglion, which controls the iris and lens of the eye. This visualization from an MCell simulation represents neurotransmitter molecules (small green ovoids) diffusing from vesicles (translucent yellow spheres) in a reconstructed ciliary ganglion synapse. Different neurotransmitter receptors (small red circles & blue squares) are also represented. This simulation showed that neurotransmitters release at ectopic (non-active zone) sites as well as from active zones. (DReAMM image by Thomas Bartol, Salk Institute).

A New Picture of Fusion

To help answer these questions, Stiles teamed up with PSC researcher John M. Pattillo (now at Macon State College) and neurobiologist Stephen D. Meriney of the University of Pittsburgh. With support from the National Institutes of Health, the group was able to take a hybrid approach that combined classic empirical observation with a sophisticated Monte Carlo supercomputer simulation.

Stiles and Pattillo used computer-aided design tools to create a three-dimensional model of an entire "active zone," the region inside the neuron where dense arrays of vesicles dock and fuse. Then, using MCell they created a simulation that tracks each calcium channel, the calcium binding sites on vesicles, and thousands of diffusing calcium ions inside the micron-wide active zone for several milliseconds of the firing cycle.

In the laboratory, Meriney probed living nerve and muscle cells to record how calcium concentration spikes and then falls following an electrical impulse. "These and other experimental results constrain many of the variables," explains Stiles, "leaving us with only a few free parameters to play with in the simulations." Running the simulations over and over, Stiles and Pattillo looked for combinations that would reproduce the recorded behaviors of real synapses: for example, the short and consistent delay between the stimulating impulse and neurotransmitter release.

Long used in high-energy physics, astronomy and other areas of science, such a hybrid approach to modeling, notes Stiles, has been less used in biology. "In part, this is because biology is so complicated and difficult to measure on these scales, and in part because the computational cost is so high. We are only now getting to the point where we have the supercomputer power and the insight into biomolecular dynamics to do computational biology this way."

To thoroughly explore the plausible range of permutations, Stiles and Pattillo had to run roughly 500,000 simulations. Each run generated thousands of output files, so the group devised a compression scheme, analogous to the MPEG encoding used for DVD movies, that allowed them to efficiently store the results and mine them for insights.

After more than a year of patient work, the data mining struck gold — in a surprising place. Neuroscientists had for years guessed that each synaptic vesicle sports four binding sites for calcium ions, and that fusion occurs only when ions dock at all four. This was one of the first models Stiles and Pattillo tried. "The results made it immediately obvious," recalls Stiles, "that this wasn't right." The virtual neuron almost never released transmitter, and no amount of tweaking other variables could produce realistic behavior. "We scratched our head and said, 'OK, let's push up the number of binding sites on each vesicle and see what happens.'"

After testing many different combinations, the group finally discovered one that neatly reproduces the experimental data. In this model, each vesicle is dotted with 25 to 40 binding sites, and fusion occurs when calcium ions fill six to eight of those sites. "The latest data coming in from biochemists now suggest that there are good reasons to expect this is true," says Stiles. "So that is quite gratifying."

MCELL SIMULATIONS HAVE PROVEN ABILITY AT REVEALING THE DETAILS OF CELL-TO-CELL INTERACTIONS

The achievement builds on another project to which Stiles also contributed, in collaboration with Terrence J. Sejnowski, Thomas M. Bartol and others at the Salk Institute and the University of California, San Diego. That effort similarly constructed a 3-D model of a synapse, in this case one that connects two neurons. Using a model derived from microscope cross-sections of actual synapses, the simulations overturned the conventional view that vesicles release neurotransmitters only within the active zone. Fusion events, they concluded, must be occurring in other regions of the synapse as well.

Such detailed insights into the structure of synapses are especially relevant for a class of diseases, called myasthenias, that arise when synapses are malformed or attacked by the immune system, leading to weakness, motor dysfunction, even paralysis. Because there's a need, and because MCell has proven abilities, prospects are promising for this way of understanding cell-to-cell interactions. With a soon-to-be-released new version of MCell and DReAMM, even more precise answers will be possible. "In the new version, molecules are able to react chemically with each other," says Stiles, "as they diffuse through space. So much more general phenomena now become potential subjects for MCell simulations." (WG)

>>More information: **http://www.psc.edu/science/2005/stiles**

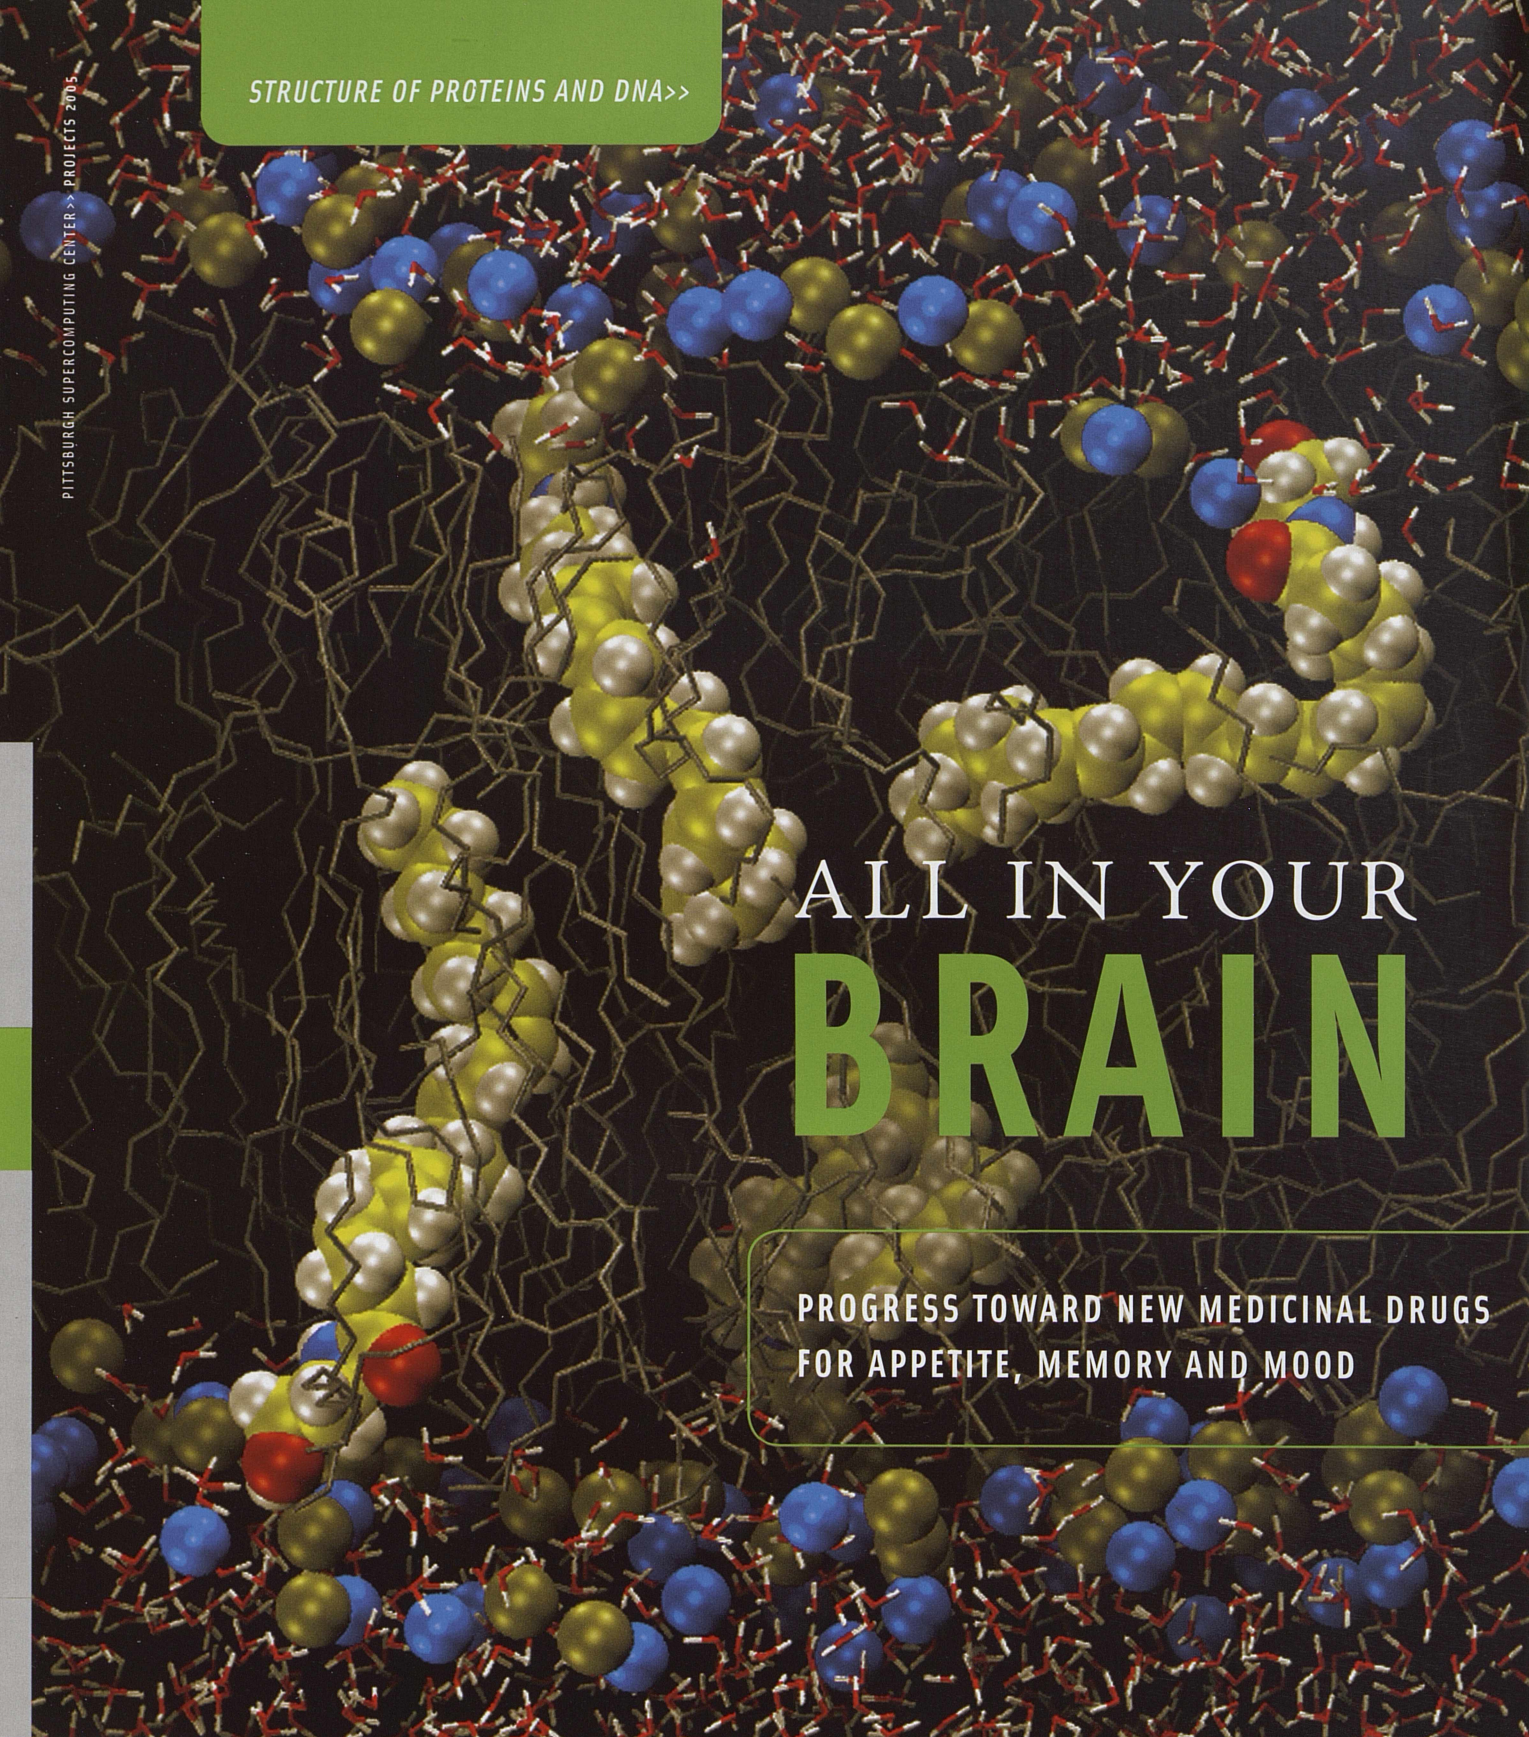

ALL IN YOUR BRAIN

PROGRESS TOWARD NEW MEDICINAL DRUGS FOR APPETITE, MEMORY AND MOOD

The model cell membrane for Lynch and Reggio's simulations was a bilayer sandwich of lipids – nitrogen (blue) and phosphorous (gold) heads facing outward on both sides of filamentary tails (gray). The membrane, embedded within water molecules (red & white), holds four anandamide molecules, with tails of carbon (yellow) and hydrogen (white) and heads containing oxygen (red). They used this model for preliminary simulations, without TMH6, to establish anandamide's arrangement, with its tail extended into the membrane core.

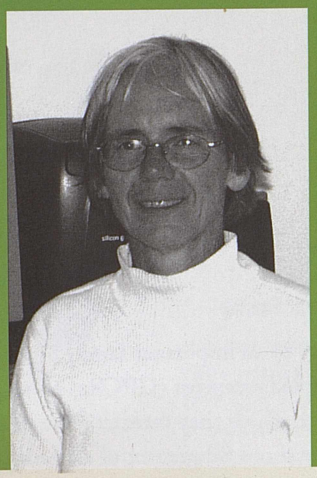

From left, **Patti Reggio**, University of North Carolina at Greensboro and **Diane Lynch**, Kennesaw State University

Your brain is loaded with receptors, billions of them. Each brain cell has thousands — molecules, usually proteins, like jigsaw puzzles with a missing piece. When another molecule of the shape and size to fit arrives, the receptor is activated, triggering a cascade of biochemical processes that, to name a few possibilities, can lift your mood, get rid of a headache or make you feel hungry.

The missing pieces, called ligands, can come from inside the body. Among these are the neurotransmitters — like the endorphins that produce "runner's high" or dopamine, which is involved in the so-called reward circuit associated with addiction. Other ligands, like nicotine or the active ingredients in aspirin, come from outside. Most therapeutic drugs have their effect by interacting with receptors, and much of drug research involves finding compounds of the right size and shape to activate receptors or to block them from other ligands.

Scientists postulated receptors in the first half of the twentieth century but didn't actually find them until 1972, when radioactive tagging led to finding the receptor for opium and related ligands. This spurred major efforts to find others, and in 1988 scientists found the brain's cannabinoid receptor — activated by tetrahydrocannabinol, the most active compound in marijuana, and called CB1.

Since the mid-80s, Patti Reggio, professor of chemistry and biochemistry at the University of North Carolina at Greensboro, has studied the cannabinoid receptor. When she started, scientists knew it existed and that it played a role in hunger, memory and appetite — making it a promising target for therapeutic drugs. But it wasn't until new tools came along — the ability to clone receptors — that scientists identified the amino-acid sequence of CB1. Even so, until recently scientists knew little about how ligands interact with CB1.

Reggio and her colleague Diane Lynch, a theoretical chemist at Kennesaw State University, used LeMieux, PSC's terascale system, to simulate part of CB1 interacting with a ligand called anandamide. Their results present the first picture of the atom-by-atom details of how a ligand initially becomes attached to CB1. This information, not available until now, is valuable to pharmacy companies searching for novel drugs that, for instance, could increase the appetite of chemotherapy patients or, as some have proposed, block painful memories associated with post-traumatic stress syndrome.

THE FIRST DETAILED PICTURE OF HOW THIS IMPORTANT RECEPTOR "GREETS" A LIGAND

A New Receptor Family

CB1 belongs to an important family of receptors, G-protein coupled receptors (GPCRs) — so named for the proteins with which they interact. "We now know," says Reggio, "that about 80 percent of all neurotransmitters and hormones work through GPCRs." Because GPCRs are involved in many pathological conditions — such as allergies, inflammation and depressed mood — they are the target of 40 to 50 percent of modern medicinal drugs.

Densely clustered in areas of the brain associated with mood, memory, and appetite, CB1 is an unusual receptor in several ways. First, it's ancient on the evolutionary scale. A remarkably similar variation has been found in all organisms except insects. Second, it has the ability to interact with other brain receptors, dampening them by what Reggio calls a "mellowing effect." Although the mechanism isn't well understood, it's known, for instance, that CB1 activation can reduce the presence of dopamine. "CB1 as a direct drug target is interesting," says Reggio, "but as an indirect target it's even more interesting because it modulates other receptors."

CB1 is also unusual, even among GPCRs, in that its internal ligand, anandamide, is synthesized from lipid, a fat-like material in the cell membranes. Most ligands of GPCRs are water soluble and approach the receptor from outside the cell. The discovery of a lipid as the natural ligand for CB1 in 1992 added a fascinating new twist to CB1 research. "The whole lifetime of these ligands is connected to the lipid environment," says Lynch. "As a result, if you want to model CB1 and its ligands, you have to accurately describe lipids. That's a large computational task, hence the need for high-performance computing."

While CB1 in totality is a bundle of seven helices, the region that changes shape when CB1 is activated is in one helix, TMH6 (transmembrane helix 6). From a PSC biomedical workshop in 2001, Lynch learned to use software called VMD (Virtual Molecular Dynamics) to build a computational model. With VMD, she constructed the "lipid bilayer" sandwich that forms the cell membrane, added water molecules — to realistically represent the cellular environment — and inserted anandamide and TMH6.

With LeMieux, Lynch first "equilibrated" the model, nearly 16,000 atoms, with a series of calculations — allowing the molecules to adjust to each other. With a molecular dynamics (MD) program called NAMD2 — developed by Klaus Schulten's group at the University of Illinois Urbana-Champaign — she ran a series of MD simulations, which track the atoms as they move and how the molecules move and change shape as they interact with each other over time.

Lynch calculated the position and energies of each atom every femtosecond (a millionth of a billionth of a second). Each simulation covered six nanoseconds (six billionths of a second) of molecular activity — a long time in biological terms. Lynch repeated this simulation eight times, to track six different trajectories of the interaction and better close-in on a statistically meaningful picture. Using 64 LeMieux processors, it took two weeks of computing to generate 48 nanoseconds of data.

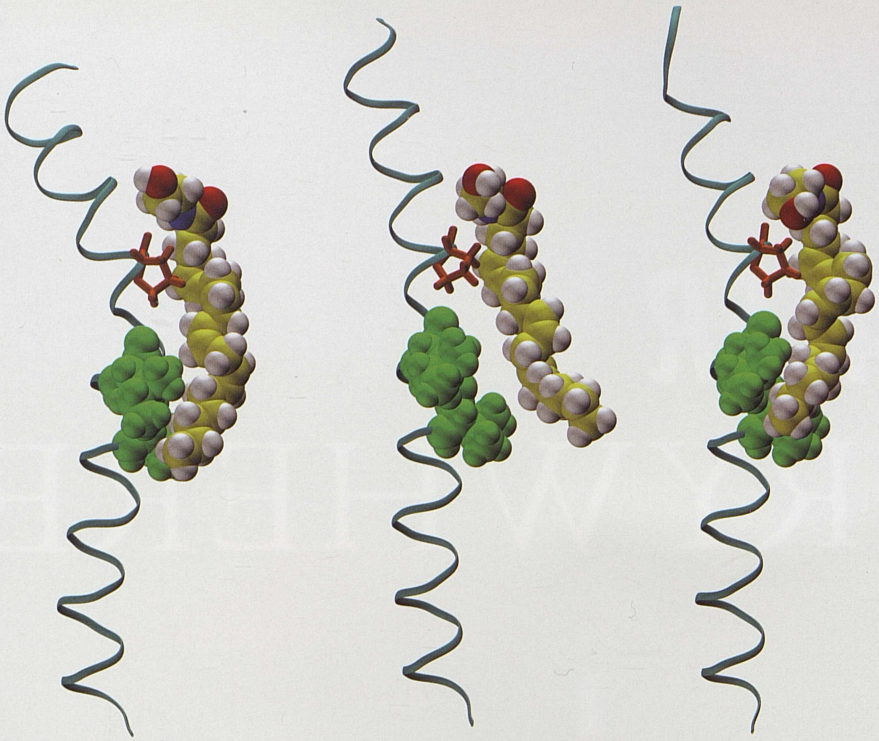

These three frames from the simulations show anandamide's hydrocarbon tail (white & yellow) inserted in the TMH6 groove, releasing and then reinserting, strong evidence that the two amino acids (isoleucine and valine, both green) which form a groove on the helix backbone (cyan) are CB1's "greeter" molecules, which initially recognize a ligand as it approaches.

In the Groove

The result — a short movie showing anandamide interacting with TMH6, short but loaded with information. In the opening scene, the tail of anandamide is recognized and "greeted" by two amino acids, side-chains on the TMH6 helical backbone, valine and isolecuine. They form a groove where anandamide briefly inserts, then breaks away, then reinserts — a sequence that appears to confirm that the initial recognition elements for anandamide/CB1 binding are in the TMH6 groove and that anandamide has the right shape to interact with this outward-facing lipid surface of CB1.

"This shows," says Reggio, "that anandamide is long enough in the membrane," says Reggio, "for its tail to reach the groove on the helix. It goes in and out of that groove and so is capable of touching the area of the two amino-acids and interacting with it." From this result, a collaborator of Lynch and Reggio — Zhao-Hui Song at the University of Louisville — did laboratory studies to confirm that the valine/isoleucine groove is the gateway by which anandamide and other cannabinoids attach to CB1 from within the cell membrane.

Perhaps more importantly, the movie suggests the signaling mechanism of the activation process. When anandamide interacts with TMH6, the helix "kicks out" to the side. "If you were sitting inside the cell and looking up at the full receptor," says Reggio, "you would see the receptor widen as a result. When it does this, it opens space for the G protein, which is sitting very nearby, to insert a segment into the bottom." The G protein changes shape and pieces of it come into contact with enzymes and other cellular components, sending a signal that starts the biochemical cascade.

A drug about to come on the market targets CB1. Sanofi-Aventis, the third-largest pharmaceutical company in the world, has developed Accomplia, a drug that reduces cravings for food and cigarettes by turning off CB1 in the hypothalamus. In the last stage of clinical trials, Accomplia may reach the market in 2006. Deeper knowledge of how CB1 interacts with ligands, such as provided by Lynch and Reggio's findings, will no doubt lead to further drugs.

For the next stage of their work with LeMieux, Reggio and Lynch plan to repeat their simulations with the entire seven-helix CB1. They also plan to look at how anandamide is destroyed in the cell membrane after it activates CB1. This will involve building a simulation cell containing fatty-acid amide hydrolase — the protein, says Lynch, that "chews up and spits out anandamide" — and will require five to ten times the number of atoms of the anandamide-TMH6 simulations, a much more demanding job for LeMieux.

"We're trying to figure out how something works that's going on in your brain right now," says Reggio. "That figuring out is fun for its own sake, and once you understand it, you can use that information to improve someone's life, to help them with a disease for which there's no drug at this point that works." (TP)

>>More information: **http://www.psc.edu/science/2005/brain**

WATER WATER EVERYWHERE

WITH COMPUTATIONS PROVIDING THE ABILITY TO INTERPRET INNOVATIVE LABORATORY STUDIES, WATER DIVULGED SOME OF ITS SECRETS

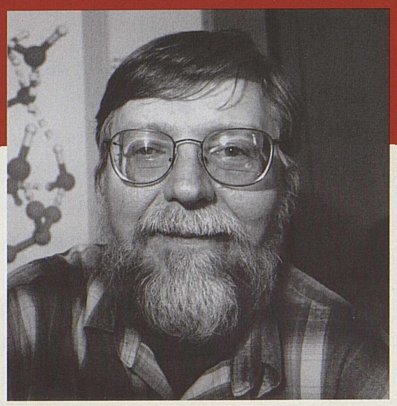

Ken Jordan,
University of Pittsburgh

PHOTO CREDIT:
Justin Merriman/*Pittsburgh Tribune-Review*

It's in the oceans, rivers and lakes, the sky, bathtubs and drinking glasses. For most of us, water is a taken-for-granted part of our lives. Nevertheless, after more than a century of scientific study, water still holds many secrets.

"Water has properties that differ from any other substance," says Ken Jordan, professor of chemistry at the University of Pittsburgh, "and these properties play an important role in enabling life as we know it. Just one example is that water has a very high ability to store heat energy, otherwise day-night and seasonal temperatures would fluctuate much more radically than they do. We still don't fully understand what makes water so effective at this."

Water divulged a few of its secrets in 2004, with *Science* magazine listing a flurry of scientific papers on water's structure and chemical behavior as among the top 10 scientific breakthroughs of the year. These new findings, said *Science*, "could reshape fields from chemistry to atmospheric sciences."

Jordan collaborated on a number of these projects, and also in follow-up studies reported this year. He's a specialist in theoretical and computational chemistry, well known for his work on water clusters — groups of water molecules linked together. "There are huge gaps in our understanding," he says, "of how water molecules interact with each other." For much of the work cited in *Science*, Jordan and his team at Pitt relied on Rachel, PSC's 128-processor HP Marvel, a system well suited for the quantum-level computations involved.

"There's two common themes in these projects," says Jordan. "One is water. The other is the power of computer modeling when coupled with state-of-the-art experimental studies." Several of these projects explore long-standing questions about what happens when an extra electron or proton interacts with water clusters, changes that can affect many chemical processes. In all these projects, Jordan's computational work has complemented and added to what can be learned in the laboratory.

<<The Magic Number

This graphic, representing recent findings by Jordan and collaborators, shows water's magic number cluster as a dodecahedron with the proton on the surface. Twenty H_2O molecules are bound together by hydrogen bonds (dotted lines). One H_2O (purple) is in the center of dodecahedral cage, and the excess proton is associated with an H_3O+ ion on the surface (blue).

Spectral Signatures

Vibrational spectroscopy is a powerful tool for identifying subtle differences in molecular structure. Much like middle-C on the piano vibrates at a higher frequency than the B below it, different arrangements of molecules vibrate at different frequencies.

In clusters of water molecules, vibrational spectroscopy can register subtle differences in how the H_2Os link up with each other. Clustering occurs through "hydrogen bonds" — links between hydrogen of one water molecule and oxygen of another. Weaker than the covalent bonds that yoke two hydrogens and an oxygen to form water, hydrogen bonds happen as electronic charges — positive for hydrogen, negative for oxygen — on different water molecules interact.

Inevitably some of the hydrogens are left dangling — and these, as you might expect, vibrate at higher frequency than the hydrogens involved in a hydrogen bond. Due to these differences, each of many possible configurations for the same number of water molecules has its own vibrational frequencies, its "spectral signature" — which can be determined with sensitive laboratory techniques.

BETTER UNDERSTANDING OF THE "MAGIC NUMBER" CLUSTER MAY HELP TO UNLEASH A MAJOR SOURCE OF UNTAPPED ENERGY

From a sample of water clusters of the same mass at low temperature, Jordan's collaborators at the University of Georgia and Yale University record spectral signatures. The challenge then is to interpret them. "Is the experiment seeing the global minimum?" says Jordan. "Or is it probing an ensemble of many low-energy structures?" To address that question, Jordan's team applies theory and calculates "the global minimum" — the geometrical arrangement that has the lowest potential energy. This structure is the most stable arrangement of the cluster and the one that most populates the vibrational frequency spectra at low temperature.

To find the global minimum, however, is easier said than done. For a cluster of 21 waters, Jordan estimates there are easily 10^{20} minima — that's 100 million-trillion possible structures that may form transiently, of which only one is the global minimum. Even years of computing on the world's most powerful system wouldn't be enough to calculate the frequencies of all these structures.

Jordan's team first uses "model potentials" to identify a reduced set of structures that comprise the likely candidates for global minimum. "You can't possibly examine all the minima," explains Jordan. "We have a research effort in our group to develop fast algorithms that survey the landscape." Within this reduced set, they then use quantum-chemistry methods to calculate frequencies. "We can then see which calculated spectrum agrees best with the experimental spectrum."

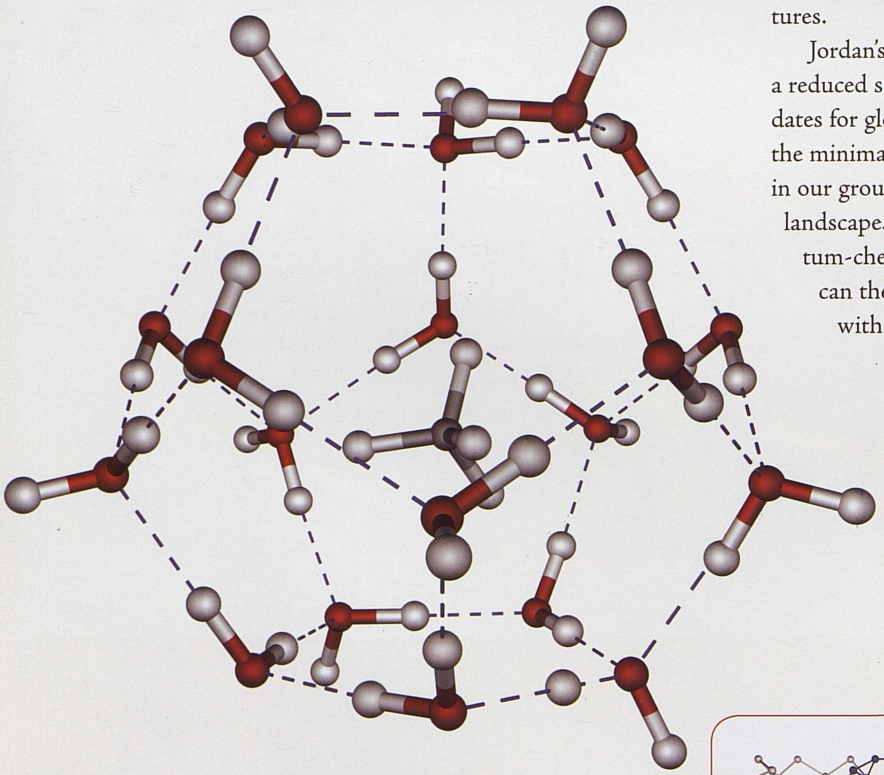

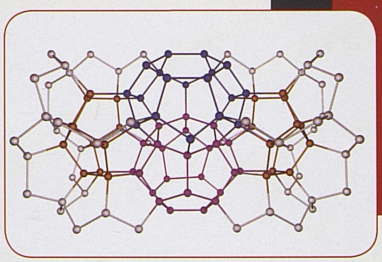

Methane Hydrate
Much of the natural gas on Earth is created from organic deposits on the ocean floor and held there in crystalline structures of frozen water up to 100 meters thick. Known as methane hydrate, these crystals form from dodecahedral clusters of water (red & white, dotted lines show hydrogen bonds), which create a cage around a single methane molecule (CH_4, gray & white). This same arrangement of waters forms the "magic number" cluster. The water lattice for these hydrates often includes two types of cages (blue & magenta).

Wet Electrons, Protons & Magic Numbers

This ability — practical only with advanced systems like Rachel — to look at a large number of minima and compute the associated vibrational spectra, which can then be related to laboratory data, has made possible many of the new findings about water. One of the questions that recent work has addressed is the nature of the "wet electron."

Also called the hydrated electron, this phenomenon — an extra electron added to water — has been widely studied because of its importance in "electron transfer" processes in photosynthesis and in the body, where electrons flit molecule-to-molecule, sparking the reactions of metabolism. Most of what's been known about the wet electron involves bulk water — as opposed to clusters. New work by several experimental groups used sophisticated spectroscopic methods to analyze an extra electron attached to water clusters. These studies developed information at a level of detail beyond what's been possible before. Some of this work relied on computational studies by Jordan's group.

Another of water's mysteries has been the structure of "magic number" clusters. Mass spectrometry shows that a cluster of 21 water molecules — with one extra proton (H+) — is much more stable than clusters with either 20 or 22 water molecules. "There's something imparting special stability," says Jordan, "and that's often associated with a special geometrical arrangement." Studies over the past 30 years have postulated a dodecahedron, a cage of 20 water molecules, with an H_2O in the middle. But where's the extra proton? Does it go with the central H_2O or on the surface of the cluster?

Using Rachel, among other computational resources, Jordan's student Richard Christie did calculations to complement laboratory teams at Yale and the University of Georgia. "There were these exciting experimental results," says Jordan, "and the question was how to explore the needed range of structures quickly enough to rapidly publish a joint experimental/theoretical paper." Running on 16 of Rachel's processors required about two days of computing to arrive at spectral signatures for one minima.

The results have settled the question. The magic number proton is on the surface. "This debate has been going on for many years," says Jordan, "but we've found pretty definitely that the proton sits on the surface of the dodecahedron." This answer, as is usual, brings new problems to resolve. "There's still a question about how fast the proton can move around. It's not a finished story."

Resolving questions about the magic-number cluster, notes Jordan, has implications for a major source of untapped energy. Methane hydrates — a structure that includes dodecahedral cages of water enclosing molecules of methane, aka natural gas. Research in the last decade has found that huge deposits of methane hydrate lie on the ocean floor, and there are major research efforts underway to find how to harness this methane. Jordan's group is working with four laboratory groups in California looking at relationships between water clusters and the similarly structured gas hydrates.

In 2005, as a follow-up to their 2004 report on the magic-number cluster, Jordan and his collaborators at Georgia and Yale published exciting results involving an extra proton in smaller water clusters, often referred to as the "hydrated proton." Research over many years identified two competing arrangements. "One," says Jordan, "is where the proton is associated with a single water molecule, and that gives H_3O+ — often called the Eigen form, for the Nobel scientist who proposed it. The other form is with the proton equally shared between two waters — H_5O_2+, called Zundel."

With innovative spectroscopic techniques and Jordan's calculations to interpret the data, the researchers for the first time identified clear spectral signatures for the two structures. By adding water molecules one by one, they found striking shifts in the frequencies — indications of movement back and forth between the Eigen and Zundel forms as well as the importance of structures that are intermediate between these two limiting forms.

"It gives us a handle," says Jordan, "on how sensitive the spectra are to the environment." The extra proton is fundamental to the chemistry of acids and this finding, which no one expected, has wide implications. What especially intrigues Jordan is that when water reveals secrets it seems to hint at even deeper ones. "To me that's the most interesting science — when you discover hidden questions you didn't anticipate when you started." (MS)

>>More information: **http://www.psc.edu/science/2005/jordan**

BRINGING LIGHT TO HEEL AND HEAL

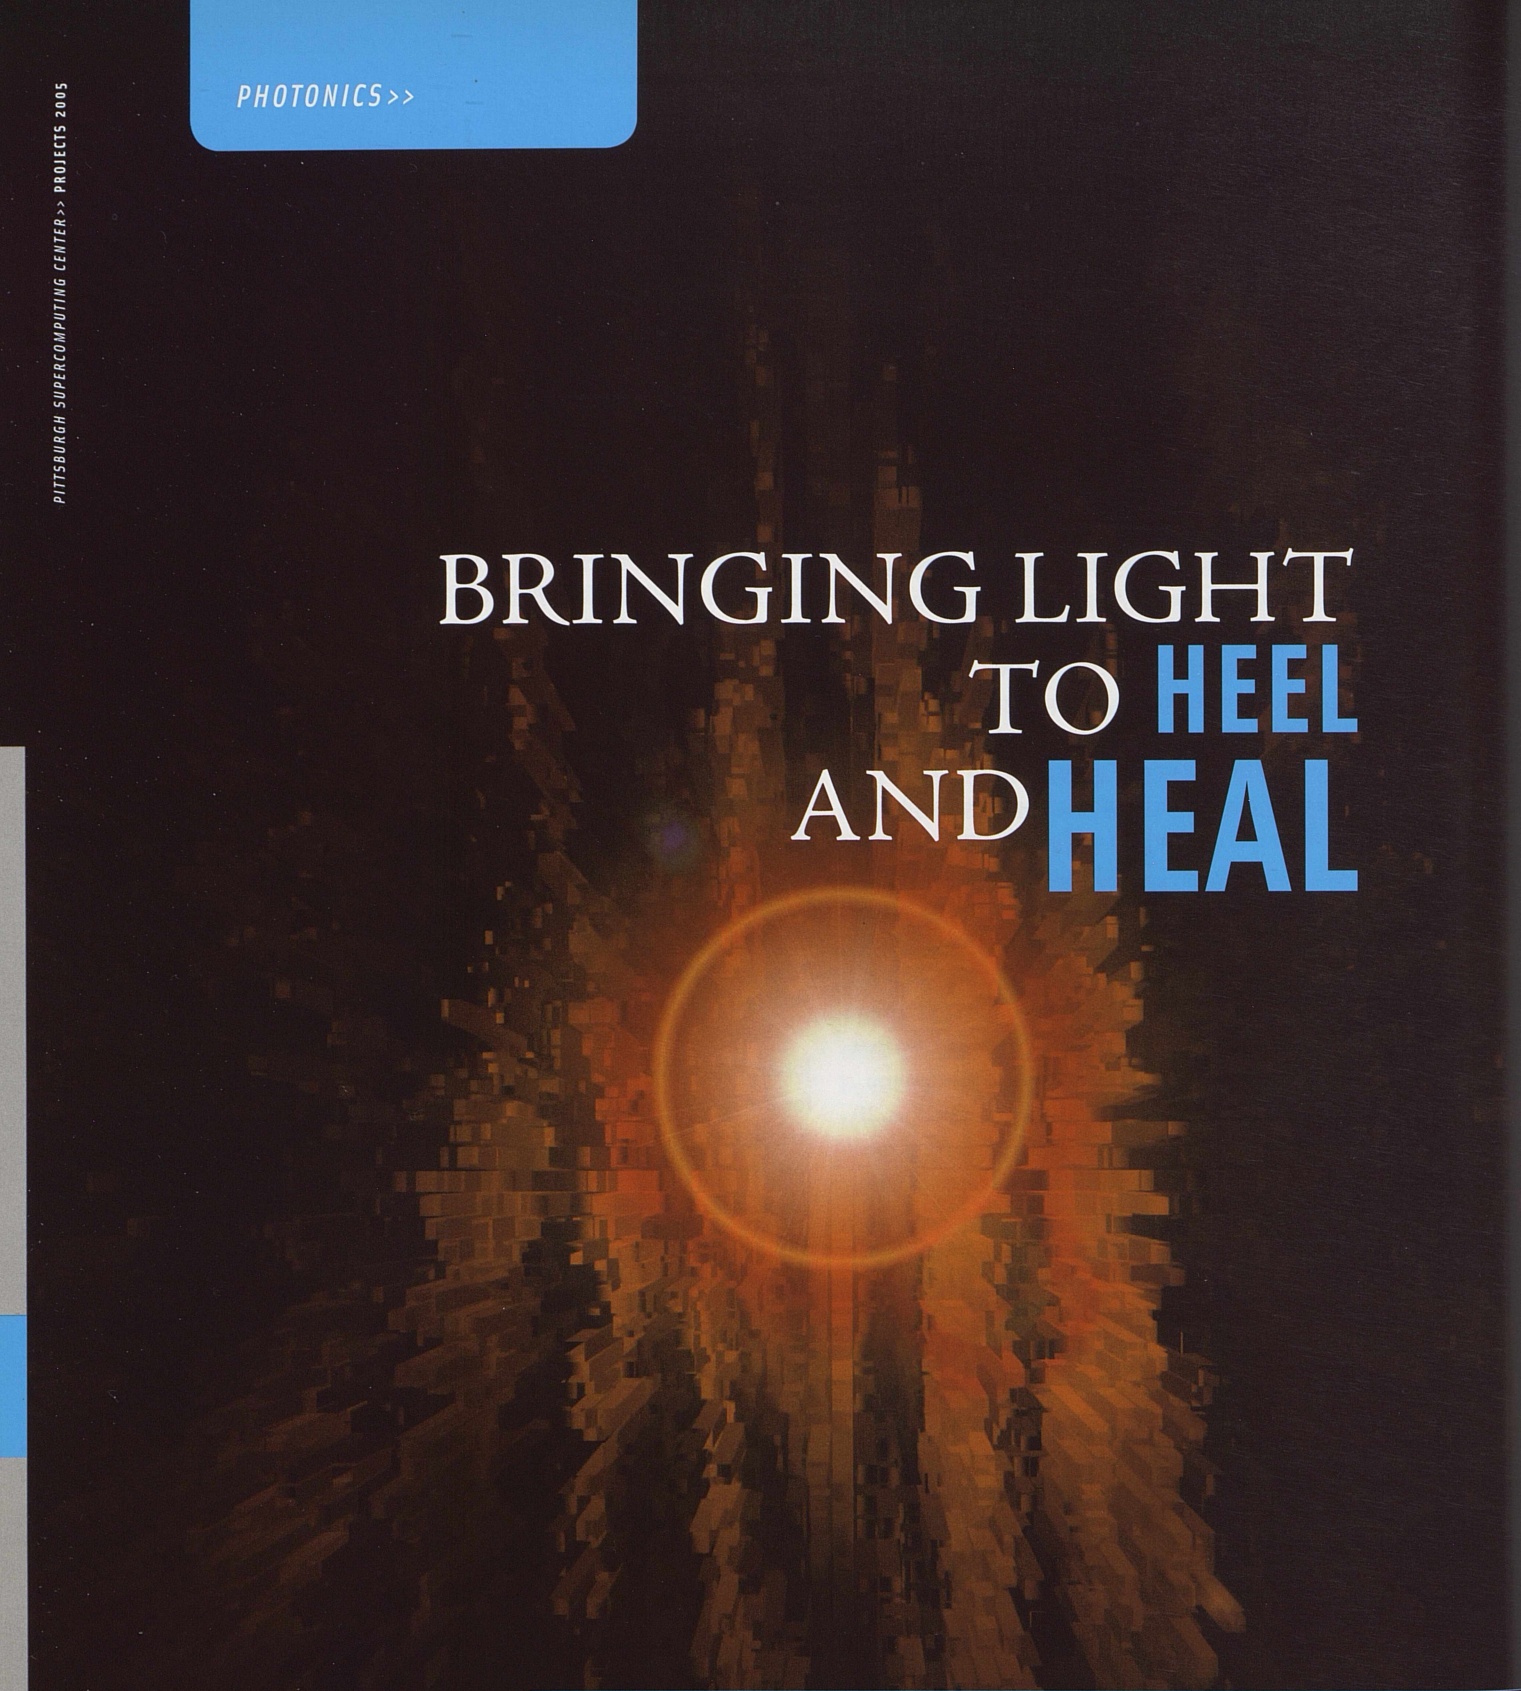

TWO PROJECTS – A NEW WAY TO STOP LIGHT IN ITS TRACKS AND A NEW FIBER FOR LASER SURGERY – EXEMPLIFY HOW COMPUTATIONAL SIMULATION DRIVES DISCOVERY

$$\begin{cases} \nabla \times \vec{E} = -\dfrac{\partial \vec{B}}{\partial t} \\ \nabla \times \vec{H} = \vec{J} + \dfrac{\partial \vec{D}}{\partial t} \end{cases}$$

Photonic crystals — not new age amulets, these fascinating materials may nevertheless herald a new age. They hold the promise of a technological revolution — computing and communications at the speed of light.

Tiny, crafted from semiconductor materials — usually silicon — they have unique abilities to trap, guide and control light. Their promise is to shift information technology from electronics to photonics — to use photons, the smallest lumps of energy in light, rather than electrons as the markers of digital 0s and 1s.

One of the most immediate applications is fiber optics. The ability to send multiple wavelengths at high speeds within fibers has transformed communications, but light could do better, much better, if it weren't hobbled by the electronic switches, routers and other devices of current optical communications technology.

Since they operate by converting optical signals to electronics and back again, these devices considerably reduce the efficiency of current optical networks. Is it possible to create all-optical circuitry — something analogous to the microcircuitry of "chips" but that doesn't require converting light to electrical current? It's a challenge many scientists worldwide are addressing, and photonic crystals loom large as an answer within reach.

A key to this work, driving it forward, has been computational simulations that predict — successfully and precisely — how photonic crystals will work in advance of actually making them. "Computation," says John Joannopoulos of MIT, a pioneer in this nascent field, "has played a dominant role in the study of photonic crystals." Photonics simulations in Joannopoulos' group recently led directly to a striking advance in the optical fibers used in laser surgery (see "Healing Light").

The beauty of photonics simulations, says physicist Shanhui Fan of Stanford, is the ability to use the full form of Maxwell's equations. This set of equations, named for James Clerk Maxwell, a 19th century Scottish physicist, governs most optical and electromagnetic phenomena. Not so long ago, limitations in computing technology required clever approximations to apply these equations. Systems like LeMieux, PSC's terascale system, however, have written a new script.

"With LeMieux," says Fan, "we have the ability to solve the entire set exactly." This means that computational experiments precisely mimic physical reality and give the researchers high confidence that their predictions can be realized in the laboratory.

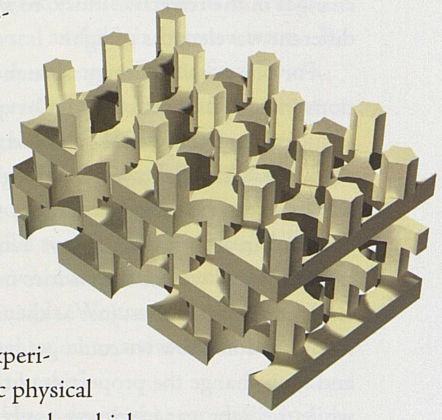

What looks like Swiss cheese with precisely spaced holes is a 3D photonic crystal designed by researchers at MIT. The distance from hole-to-hole is about half a micron, ten times thinner than an average human hair.

Fan and graduate student Mehmet Faith Yanik used LeMieux to simulate a new device that can stop light and hold it captured until a subtle shift in optical features releases it. They've teamed with a laboratory group at Stanford to build and demonstrate their device, which suggests it may be possible to corral complicated light pulses and, moreover, do it in a way that integrates easily with existing chip technology.

Optical Resonance Chambers

It made news in 2001 when researchers brought light to a standstill for the first time. Two groups at Harvard demonstrated a technique that captured light in clouds of gaseous atoms. But these systems of atomic gases are impractical for an all-optical circuit. Because the Stanford team's approach relies on photonic crystals, rather than gases, their device could operate at room temperature and be only microns in length, allowing it to easily integrate with traditional microcircuitry.

By careful design of irregularities within patterns of tiny cavities, photonic crystals can allow — or forbid — the passage of certain wavelengths of light on prescribed paths. Exactly which wavelength, or band of wavelengths, can travel through or not depends on the properties of the crystal.

Yanik stumbled on the light-stopping mechanism while using LeMieux to simulate the impact of changing one property of a crystal, the index of refraction — the ratio of light's speed in a vacuum (well established at 186,000 miles per second) to its speed in a medium, where it travels more slowly. His original goal was a tunable switch — a crystal that could be prompted, by small changes in the refractive index, to allow safe passage to different wavelengths of light.

For one possible design of such a switch, the simulations indicated the effect could be quite strong. Small changes in refractive index allowed a large change in the bandwidth of allowed wavelengths. And that wasn't all. "I saw an optical signature very similar to the ones observed in atomic media," says Yanik. "So the question became, could we use the cavities in the crystal to store electromagnetic pulses, just as they were stored in atomic media? If somehow we could get light into this structure, and then change the properties of the entire structure while the light was inside, we could change the properties of light as well and trap it."

The idea depends on a phenomenon called optical resonance, which is similar to why long and short pipes in an organ produce notes of different frequency. In an organ, each pipe is cut to the length required to amplify sound waves of a desired frequency. The sound energy bounces back and forth inside the pipe and establishes an unmoving wave pattern, or resonance, at the desired frequency. In the Stanford team's approach, the role of the organ pipe is played by a waveguide — either an empty channel or closely spaced cavities inside the crystal that allow light to propagate. Prior to this work, many groups had used optical resonators to trap light of a single wavelength. Optical communication, however, uses light pulses to encode and transmit information, with each pulse composed of many wavelengths. Trapping such a multi-wavelength pulse in a single resonator would lose the information carried by the pulse.

Yanik and Fan's idea, however, goes a crucial step further by tuning all of the wavelengths within a pulse to the same frequency and, at the same time, adjusting the crystal to resonate at that frequency. They do this by adjusting the index of refraction once the pulse has entered the crystal. As all the frequency components are collapsed to a single frequency, the information becomes encoded by the phase and intensity of light along the waveguide.

Changing the resonance of the crystal, Yanik explains, is like adjusting the spacing of stepping stones across a river. Shifting the crystal's index of refraction is similar to spreading the stones out, so that photons — the tiniest energy chunks of light — of a particular frequency can no longer hop from stone to stone. They have been trapped. When the pulse needs to be released, the index of refraction is shifted back, the stones move closer together, and the photons zip away.

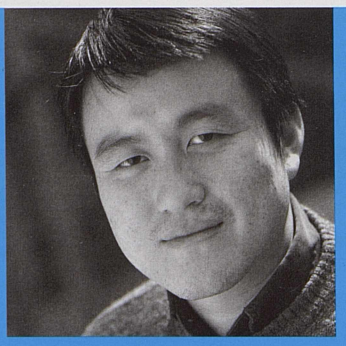

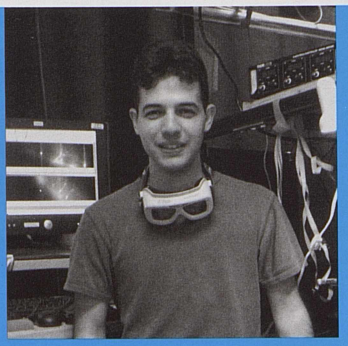

Shanhui Fan (left) & Fatih Yanik, Stanford University

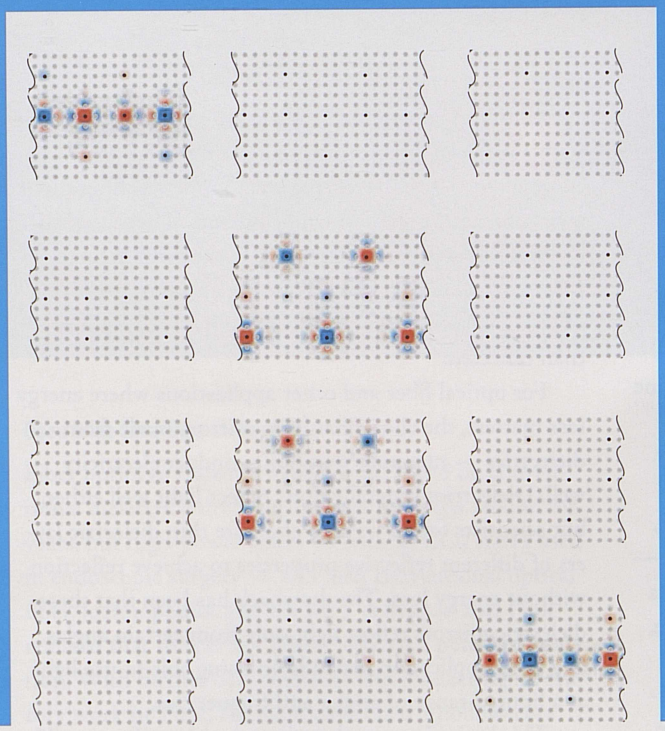

Capturing Light in a Photonic Crystal

This graphic from simulation shows snapshots of the positive (red) and negative (blue) electric fields as an optical pulse propagates (left to right) through a photonic crystal, shown in three segments at four times (top to bottom). Resonant frequencies of the cavities (black dots) are tuned to stop the pulse during the time interval shown in the second and third snapshot, when the cavities are detuned and the pulse is released.

A Dynamic Duo: Maxwell & LeMieux

"The entire idea," says Yanik, "from refractive-index switches to light-trapping devices, was first realized on a supercomputer." Once he and Fan identified the light-stopping possibility, Yanik adapted software he'd already written to simulate it. Using almost every one of LeMieux's 3,000 processors, they simulated a series of possibilities until arriving at a 100-micron waveguide with 120 side-cavities. "A hundred microns," says Fan, "fits on a chip, a small distance in practice, but a long distance to simulate."

To exploit the large-scale parallelism of LeMieux's 3,000 processors, Yanik's software parceled separate parts of the crystal waveguide to separate processors. It took 10 simulations to describe the light-trapping behavior, with each simulation of a light pulse entering the wave guide requiring two hours, which Yanik estimates as a year's worth of computing on a desktop PC.

The simulations showed that shifting the index of refraction around the pulse forces the wavelengths to adopt a single frequency, and traps the pulse in and between cavities. In the 100 micron, 120 side-cavity waveguide, a 1/10,000th shift in the index of refraction is enough to capture the information in commonly used pulses of light.

Another surprising result of the simulations, says Fan, is that if the index of refraction were tuned beyond the point where the light pulse screeches to a halt, the pulse would not merely stop, but reverse in its tracks, backing out of the crystal as though it were a train reversing direction to re-emerge, caboose first, from a tunnel. This time reversal effect, he says, might prove useful in repairing signal degradation.

Efforts to build the device in the lab, in collaboration with Stanford colleagues Martin Fejer and James Harris, are now running parallel to more simulations. "What we've done so far is a two-dimensional simulation," says Fan, "as a proof of principle. We are now extending it to a three-dimensional simulation to arrive at the exact structure the device needs to take."

For optical networks, a device that can catch and hold light for an arbitrary length of time offers promise to alleviate the congestion that happens when too many pulses arrive simultaneously at a network junction. Beyond that, there's the promise of quantum computing, the vision of transistors that manipulate single photons rather than electrons. It's a future, perhaps sooner than we think, in which circuits will be a thousand times smaller and faster. Yanik and Fan's simulations with LeMieux bring us a step closer. (KG)

>>More information: **http://www.psc.edu/science/2005/fan**

LIGHT BACKS OUT OF THE CRYSTAL LIKE A TRAIN FROM A TUNNEL, CABOOSE FIRST

HEALING LIGHT

In November 2004, a woman in North Carolina with potentially suffocating growths in her larynx and trachea had them removed by a high-power laser and went home the same day. This condition had never before been treated without anesthesia and operating-room surgery. Six years earlier, John Joannopoulos' team of physicists at MIT used supercomputers to learn something no one knew about mirrors.

The two events are linked. A new laser technology, developed from a startling insight into the physics of light, may have saved the woman's life and, at the least, promises huge savings in the treatment of her disease — recurrent respiratory papillomatosis — one that affects tens of thousands of people in the United States alone.

It's a success, furthermore, that exemplifies how supercomputing is no longer merely a supporting character, but with increasing frequency plays a lead role in scientific discovery. In 1998, John Joannopoulos and his team of researchers at MIT discovered what has come to be called a "perfect mirror." Their eureka moment came not in the laboratory or with pencil and paper working out of mathematical theory. It happened because a computational model produced results no one expected.

The Perfect Mirror

It may be the most significant advance in mirror technology, said the *New York Times*, since Narcissus fell in love with his own image in a pool of water. The perfect mirror is so called because it reflects light at any angle with virtually no loss of energy. As a result it makes possible a number of applications in optical technology, the most significant to date being flexible optical fiber that can transmit the high-powered CO_2 lasers used in endoscopic surgery.

Until the Joannopoulos team's 1998 finding, reported with a paper in *Science*, mirrors were understood to come in two basic flavors, both with inherent limitations. Everyone knows about metallic mirrors. They work all too well for seeing your own face in the morning, but they don't work for optical fiber because a large portion of the light leaks away, absorbed by the metal, rather than reflected.

For optical fiber and other applications where energy loss matters, the choice has been mirrors made from dielectrics — materials that don't conduct electricity well. Dielectrics generally don't reflect light well either, but scientists found ways to alternate thin dielectric layers of different reflective properties to achieve reflection without energy loss. The drawback has been that these dielectric mirrors reflect light only from certain angles, and their application depends on being able to use light at a limited range of angles and frequencies.

This limitation was thought to be a law of nature, like gravity — no way to get around it — until 1998, when Joannopoulos' team noticed anomalous results from a computational model of a photonic crystal mirror they were running at the San Diego Supercomputer Center. The light seemed to reflect at a much larger angle than was thought possible. "We saw some interesting results in the computation," says Joannopoulos. "Then came the theory to explain the computation, and then came a real experiment making something like this and testing it."

The result: a multi-layered dielectric mirror that reflects light from all angles without energy loss. Within a few years, the perfect mirror proved to be the solution for delivering a high-powered laser via flexible optical fiber.

A DREAM COME TRUE FOR ENDOSCOPIC SURGERY PROMISES LARGE COST SAVINGS

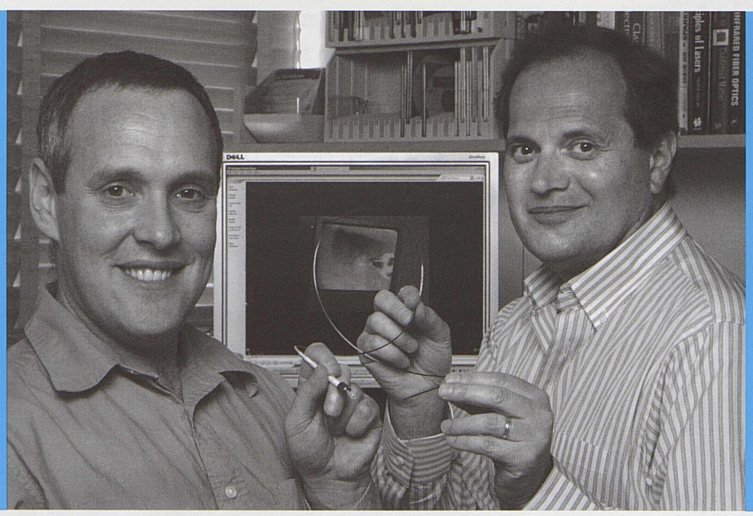

Yoel Fink (left) and John Joannopoulos hold a length of OmniGuide flexible optical fiber with a lens. The computer screen shows an open human throat.

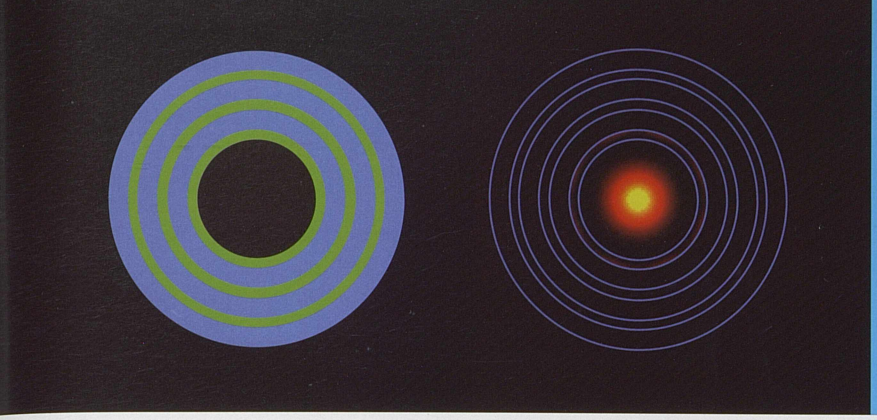

Hollow Fiber to Guide Light

These two cross-sectional images represent a schematic (left) of a model OmniGuide hollow core fiber and the first visualization (right) from a computational simulation by Joannopoulos of the same fiber. The "perfect mirror" photonic reflector consists of alternating concentric layers (green and blue) of dielectric with differing indices of refraction. The visualization shows boundaries between the dielectric shells (blue circles) and power density (increasing from red to yellow) of a light beam contained within the hollow core.

Open Wide for a High-Power Laser

Fiber optics to transmit visible light, based on conventional dielectric mirror technology, has been around for years, but high-power lasers — such as CO_2 lasers used in endoscopic surgery — will melt conventional optical fiber. Joannopoulos and his MIT colleague Yoel Fink realized that the perfect mirror offered a potential solution. With further computations and pioneering laboratory work by Fink, the team developed a hollow-core fiber — essentially a dielectric perfect mirror rolled up into a tube — designed in such a way, based on photonics, to transmit high-power lasers.

To take this idea beyond the laboratory into useful applications, Joannopoulos and Fink in 2000 helped to found a company, OmniGuide Communications, to develop and market the new hollow-core fiber. Further computations over the next few years — at San Diego, Illinois and Pittsburgh —explored other fundamental issues and phenomena of this new class of cylindrical photonic-crystal fiber.

In endoscopic surgery, the lack of a fiber for high-power transmission has meant that the laser had to be delivered via an apparatus with an articulated arm and large handpiece — precluding use of these precise lasers for many minimally invasive procedures. For this rea son, the surgery to treat RRP required dislocating the patient's jaw and general anesthesia, so that the laser could be brought close enough to the affected area.

UNTIL NOW, THIS SURGERY REQUIRED DISLOCATING THE PATIENT'S JAW

A test case for the OmniGuide hollow-core fiber came last year. In serious cases of RRP, the surgery often must be repeated to keep the breathing passage open. Dr. Jamie Koufman, director of the Center for Voice and Swallowing Disorders of Wake Forest University Baptist Medical Center, had a woman patient who had undergone several previous RRP surgeries, but once again had developed near-total obstruction of the larynx and trachea.

Koufman obtained FDA approval to use the prototype fiber. With a numbing topical spray in the throat and trachea, no anesthesia, and a CO_2 laser delivered via an OmniGuide fiber, she cleared the RRP growths. The patient, who went home that day, is doing fine.

"Unsedated, laryngeal laser surgery with the OmniGuide fiber is a dream come true for me as an endoscopic surgeon," said Koufman, "and the patient loved it because it was easy for her." Typical cost of RRP operating-room surgery with general anesthesia is $25,000. With expected FDA approval, the new procedure promises very large cost savings nationally. (MS)

>>More information: **http://www.psc.edu/science/2005/joannopoulos**

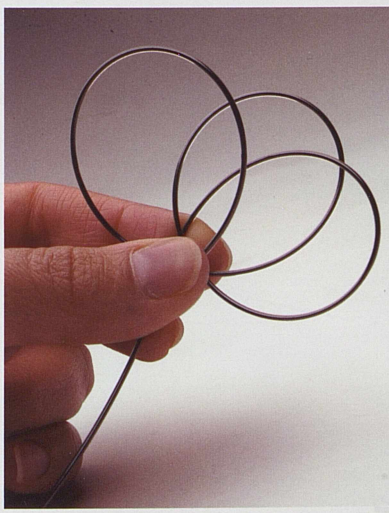

TOOLS OF DISCOVERY

"These novel optical fibers, based on photonic crystals," says Joannopoulos, "offer a new approach for medical lasers, making it possible to guide a CO_2 laser beam, which can cut tissue with high precision, into a patient's body through a very small incision. It will likely prove itself useful for many procedures."

"Computational science has come a long way over the past 20 years. Even well known equations can have remarkable unexpected consequences that we would never learn about without these powerful computational engines, such as LeMieux. This is just one advance that highlights how these machines are invaluable tools of discovery."

WITH 2,048 LEMIEUX PROCESSORS, P. K. YEUNG SET A MILESTONE IN TURBULENCE SIMULATION

TAMING THE WHIRLWIND

"We're experiencing a little turbulence, folks," the pilot says, as the plane plummets fifty feet like a car in a funhouse ride. White knuckles, churning stomach — on an airplane, a word you'd prefer not to hear is *turbulence*.

As P. K. Yeung is quick to tell you, however, turbulence is often beneficial and modern air travel couldn't exist without it. "But for the turbulent mixing of fuel and air in a jet engine, jet flight wouldn't be possible," says Yeung, a professor of aerospace engineering at Georgia Tech. He has applied his deep knowledge of turbulence to carry out some of the largest computational studies of this widespread and important phenomenon, and his recent work at PSC, using 2,048 processors of LeMieux, PSC's terascale system, sets a new milestone for large-scale turbulence simulation.

Part of our lives in many ways — from the cream we stir in our coffee to thunderstorms that ruin a night at the ballgame — turbulence defies easy definition, but is, roughly, a state of fluid flow in which the velocities at any point fluctuate randomly.

But for these random fluctuations, many important industrial chemical reactions would happen very slowly or not at all. On a larger scale, turbulent mixing in the lower atmosphere coupled with phenomena at high altitudes has a great effect on weather in the short term and climate in the long term. Turbulent mixing in ocean currents, such as the Gulf Stream, spanning thousands of miles helps to maintain the heat balance and ecology of the oceans.

Better understanding of turbulence, especially since the advent of supercomputers, has led to improvements in how we live, including better airplane wings, which lower the fuel-cost of air travel, and better artificial heart valves, which save lives. But it's an extremely complex phenomenon — one that Nobel Prize-winning physicist Richard Feynman once referred to as the "last unsolved problem in physics" — and many challenges remain.

One of the more pressing turbulence-related issues, says Yeung, is in preserving environmental quality. Where, for instance, will particles of pollutants from a smokestack end up minutes, hours, and days from now? "In order to maintain air quality, we need to understand the behavior of smoke emanating from pollution sources. Almost always the flow out of a chimney will be turbu-

lent. We can see that in the sky — the smoke follows an irregular path — and we want to be able to describe the motion of those pieces of fluid, which constitute that cloud. If a certain part of that fluid has been contaminated, we want to know where it goes."

Using a powerful method, called "direct numerical simulation," with the advanced parallel-processing capability of LeMieux, Yeung produced results that are a significant step toward this goal.

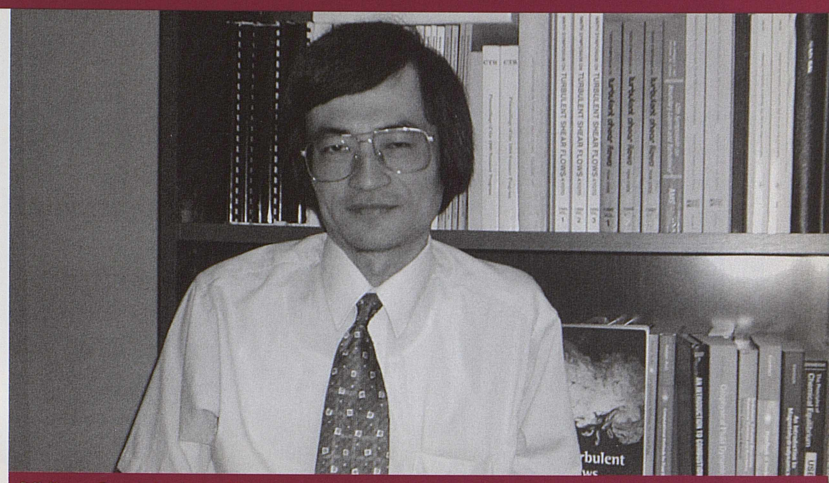

P.K. Yeung, Georgia Institute of Technology

Tracking the Particles

A major challenge in simulating turbulence is that the random fluctuations — the eddies and vortices — occur over a very wide range of scales, all of which must be taken into account in a realistic model. In the atmosphere, for instance, the swirls and eddies of air that make up the overall flow vary from several centimeters in diameter to thousands of kilometers, with every size in between. The ratio of scales can be in the thousands, with the number of variables — and thus the amount of computing required to keep track of them — increasing rapidly as the ratio increases. This imposes a daunting computational demand.

Yeung tackles this problem directly. Direct numerical simulation (DNS), starts with the fundamental equations of fluid flow and calculates speed and direction for each fluid particle. "Direct" means that velocities are calculated at each time step as the flow progresses, without reliance on experimental data to supply parameters. DNS tracks each particle — such as the particles in a plume of smoke — as it moves step-by-step within a high-resolution grid.

"We are acting as if we could measure the velocity everywhere in space and over a sustained period of time," Yeung explains. "With DNS, we are able to follow the irregular pathways or trajectories of fluid elements exiting a localized contaminant source."

Visualize a plume of smoke rising from a smokestack three feet in diameter. Any two smoke particles are separated by three feet, at most, as they exit the stack. DNS allows you to address the issue of how far apart these two particles will be after they wander about in the atmosphere for a sustained period of time. Do they separate or come together over time? To do this experimentally, to identify and keep track of a single pair of fluid particles, let alone all the particles in a large cloud, would be impossible. DNS — Yeung points out — provides more data, more accurately, than is possible to gather experimentally.

Scientists quantify the degree of turbulence in a fluid flow by the Reynolds number. Higher Reynolds numbers correspond to a wider spread in the range of eddy sizes, equivalent to higher levels of turbulence. Scientists have long been interested in simulating high Reynolds number flows, but have been limited by computing power. "The availability of computers like LeMieux allows us to increase the Reynolds number by expanding the number grid points," says Yeung, "and this allows us to simulate a wider range of scales."

A Turbulence Database & the Kolmogorov Constant

Over the past year, Yeung employed 2,048 LeMieux processors simultaneously solving the fundamental fluid equations in a three-dimensional grid with eight-billion grid points. This is the largest DNS ever done that tracks the path of particles over time. To get started on LeMieux required importing his software to a system where it hadn't run before, a major challenge. "We've had very capable and dedicated assistance from PSC consulting."

LeMieux's ability to communicate efficiently among processors has been an important factor in Yeung's ability to carry out his large-scale DNS work. His software efficiently exploits LeMieux, using thousands of processors with minimal added communication time involved in adding processors, a major advantage for his work.

With more than a million processor-hours of LeMieux time, Yeung's simulations produced terabytes of data that yield the highest Reynolds number ever calculated with the DNS approach. Previously, researchers had to extrapolate data from low turbulence simulations if they wanted to apply it in high turbulence situations, which led to uncertainties. "We will now be approaching the Reynolds numbers typical in applications more closely," says Yeung, "and if we still have to extrapolate we can do so with much greater confidence."

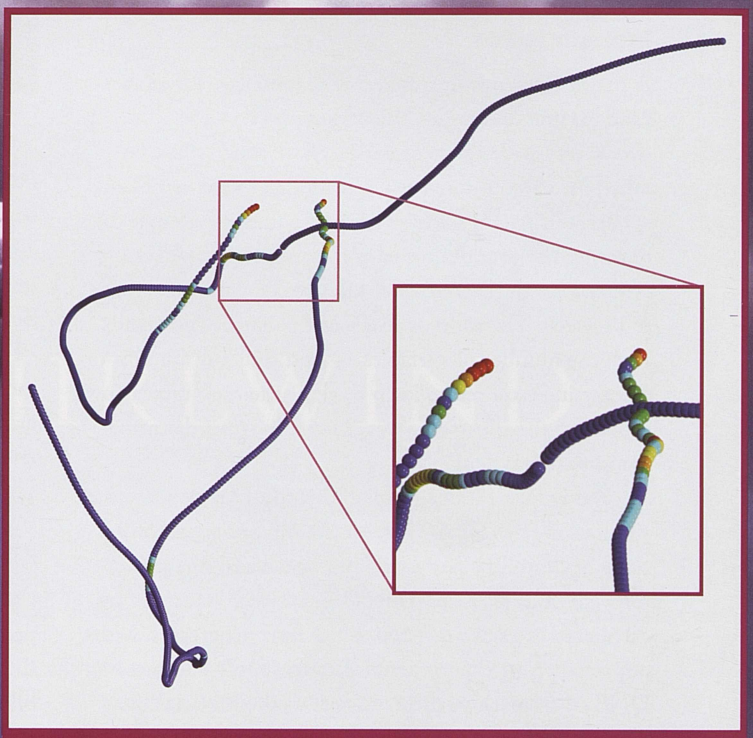

This image from Yeung's recent simulations at PSC with a resolution of 8.6 billion grid-points (2048^3) represents the trajectories of two fluid particles. They originate near each other at a location of intense local strain-rate (red decreasing to blue). This high strain-rate, which increases with Reynolds number, pulls the particles away from each other in time, a process that leads to efficient turbulent mixing and is crucial in the dispersion of contaminants, such as a plume of smoke dispersing from a smokestack.

Beyond his immediate goal of understanding pollutant dispersion for environmental purposes, Yeung's simulations create a valuable database — which can be made available at PSC to the wider community of turbulence scientists, who can use it to test their turbulence models. Because of the fundamental nature of his DNS simulations — free of assumptions derived from observation — the data is useful for turbulent flows in many different applications. For pollutant dispersion, such as a smoke plume, from a localized source, other researchers can compare their model results with Yeung's DNS data for a similar Reynolds number. "Using DNS we can obtain the fundamental data that would allow us to formulate those models more carefully," says Yeung, "and eventually to evaluate the performance of the model and suggest improvements."

In extending his DNS studies to higher Reynolds numbers, Yeung also is getting closer to pinpointing an elusive number called the Lagrangian Kolmogorov constant. In 1941, Russian mathematician A.N. Kolmogorov posited that at high enough Reynolds number small-

THE HIGHEST REYNOLDS NUMBER EVER CALCULATED BY DIRECT NUMERICAL SIMULATION

scale features of turbulence are independent of the large-scale flow geometry. This theory has been widely influential in turbulence research. "The Kolmogorov constant is of great interest because of the supposition that it is universal," says Yeung, "being the same for turbulent flows of various types of geometry as long as the Reynolds number is sufficiently high."

In the laboratory, with data from fixed measurement locations, it is straightforward to apply Kolmogorov's hypotheses, and this version of the constant is well established. Models of pollutant transport, however, use what's called a Lagrangian reference frame, which mimics an observer moving with the fluid flow — like a weather balloon that drifts with the wind. Research by Yeung and others has indicated that it takes very high Reynolds numbers simulations to establish this version of the constant. "This constant is very important to modeling," says Yeung. "Our group's large simulations on LeMieux have given quite clear evidence that the value is approaching a constant as the Reynolds numbers increases without limit." (TP)

>>More information: **http:www.psc.edu/science/2005/yeung**

This image represents instantaneous energy dissipation rates from an 8.6 billion grid-point (2048^3) simulation that P. K. Yeung's group carried out at PSC in 2004. The peaks (red) indicate points of intense local strain, such as are represented by the two particle-trajectories rapidly diverging. Such localized, short-lived bursts of intense activity increase with higher Reynolds number.

"This thing we call failure is not the falling down, but the staying down."
— Mary Pickford

AUTOMATED APPLICATION LEVEL CHECKPOINT-RESTART

WHEN BAD THINGS HAPPEN TO GOOD PROGRAMS

Living in the information age means, among other things, being aware that computers fail. Disks crash, viruses invade, bad things happen. With PCs, prudence dictates regular backups. If you're a scientist running simulations on a supercomputing system, which may take days, weeks or months to complete, prudence dictates checkpoint-restart — CPR for short.

CPR means that your "code" is designed and equipped with checkpoints, places where during execution the program periodically saves the results up to that point, so that if a crash occurs, you can restart from the last checkpoint with minimal loss. The concept is basic, if not always faithfully practiced.

To facilitate CPR on LeMieux, PSC's 3,000 processor terascale system, PSC senior research analyst Nathan Stone developed a very fast set of I/O routines — called TCS-CPR — to save checkpoint data to disk, back it up, keep track of it and automatically restart a program that crashes. "Automated restart," says Stone, "is a big sell for researchers who like to sleep."

TCS-CPR neatly takes care of everything associated with CPR except the fallible human part of manually writing checkpoints into the application code. This can be an onerous task and for some complex codes there's not necessarily an obvious place to insert checkpoints. At the same time, with the growing prevalence of large-scale parallel systems like LeMieux and the large-scale, long-running simulations they enable and, even more, the movement toward grid computing, the risks of failure are increasing, trends that make CPR more important than ever.

Recognizing this, Cornell University computer scientist Keshav Pingali and his team devised a scheme to save human time by automating the code-writing part of CPR. With their approach, called C^3 (Cornell Checkpoint pre-Compiler), a pre-compiler program takes uncheckpointed source code and transforms it to become self-checkpointing and self-restartable. In collaboration with Stone and other PSC staff, they have tested C^3 on several platforms including LeMieux and a Windows cluster at Cornell. Their results show that C^3 introduces little overhead into program execution and demonstrate the feasibility of their innovative approach.

Keshav Pingali, Cornell University and Nathan Stone, PSC

"MANY PEOPLE WOULDN'T BELIEVE THIS IS POSSIBLE, BUT THEY'VE DEMONSTRATED IT."

The Goal: Full CPR Support for Big Ben

CPR took on heightened importance in 2001 at PSC with the arrival of LeMieux, the National Science Foundation's first terascale system. Comprised of 3,000 processors, LeMieux provided NSF researchers with unprecedented capability to carryout very large-scale simulations. It also meant, for simulations that used hundreds or thousands of processors simultaneously, a statistically higher likelihood that something bad would happen on one of those processors during execution.

"Until recently," says Pingali, "most parallel processing was done on relatively reliable big-iron machines whose mean time between failures was much longer than the execution time of most programs. But many programs now are designed to run for days or months on even the fastest computers, even as the growing size and complexity of parallel systems makes them more prone to hardware failure."

Another pressing reason to pay increased attention to CPR is grid computing. As cyberinfrastructure such as the NSF TeraGrid gains maturity, it will become more common to run applications on geographically distributed systems, with different systems teaming on different parts of the job. "With grid computing," says Pingali, "the probability of failure becomes much higher."

With TCS-CPR, Stone's idea was to provide robustness and reliability in the checkpoint write and restart operations. Simply by replacing the open, write, close statements in their source code CPR routines with TCS statements, users gained a fast connection to disk, the reliability of checkpoint backup — by virtue of an offline process that replicates the checkpoint data and keeps track of it — and automated restart in the event of a crash.

Despite the advantages, users did not rush to implement TCS-CPR. Faced with this situation, Stone a few years ago became involved in a CPR working group at Global Grid Forum. There he met Paul Stodghill of Pingali's group at Cornell and realized their work complemented his with LeMieux. "They knew smart ways to automate writing CPR into code," says Stone, "and we knew how to save and keep the checkpoint data in a robust way and give it back to the application so that the user has no idea it was ever interrupted."

In the collaboration that ensued, Pingali's group used LeMieux as a test platform for C^3. They tested different application codes to evaluate the amount of overhead — added execution time — involved in bookkeeping code that C^3 inserts to implement CPR. On all among a large range of codes, this overhead is less than 10 percent, and in most cases around 2 to 3 percent. Further testing evaluated the overhead cost involved with saving checkpoints during execution. With checkpoints done hourly, this overhead is under 4 percent — a small cost for the security against failure gained from CPR.

A further objective of Pingali's CPR scheme is portability — being able to restart applications on a different platform than the one they start on, a facility that envisions grid computing. In one experiment, they applied their approach to this difficult problem — implemented with software they call PC^3 — by taking checkpoint files from an application running on LeMieux and restarting it on a Windows cluster at Cornell. The job ran to completion on the Windows cluster with identical results to the LeMieux process. "Nobody's ever done this before," says Stone. "Many people wouldn't believe this is possible, but they've demonstrated it."

Stone and others at PSC are now collaborating with Pingali's group to test and implement other checkpoint techniques on Big Ben, PSC's 10 teraflop Cray XT3 system and a lead resource on the TeraGrid for the most demanding large-scale parallel applications. "Our goal," says Stone, "is full production-level CPR support for the XT3. We're going for the five nines, 99.999 percent reliable."

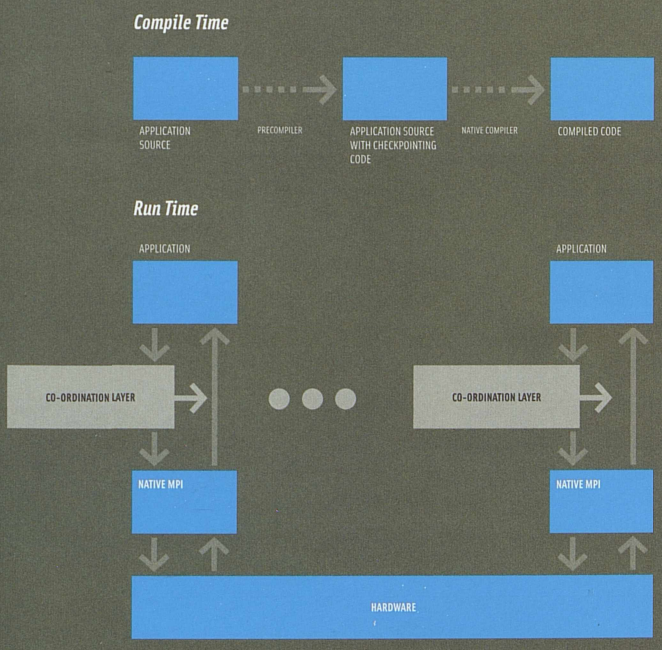

Overview of Cornell Checkpoint preCompiler

C^3 takes source codes in C (programming language) or Fortran with MPI (message-passing interface) and instruments them to do application-level CPR. The only requirement for programmers is to insert a simple statement – a "pragma" – in the source code to identify points where checkpoints might be taken. The C^3 run-time system determines whether to write checkpoints at those locations.

The precompiler output is compiled with the user's native compiler of choice for the hardware platform and linked with a run-time library that constitutes a "co-ordination layer," which sits between the application and MPI library and intercepts all calls from the application program to the MPI library. The coordination layer is designed in such a way that it is unnecessary to modify the underlying MPI library.

In Progress

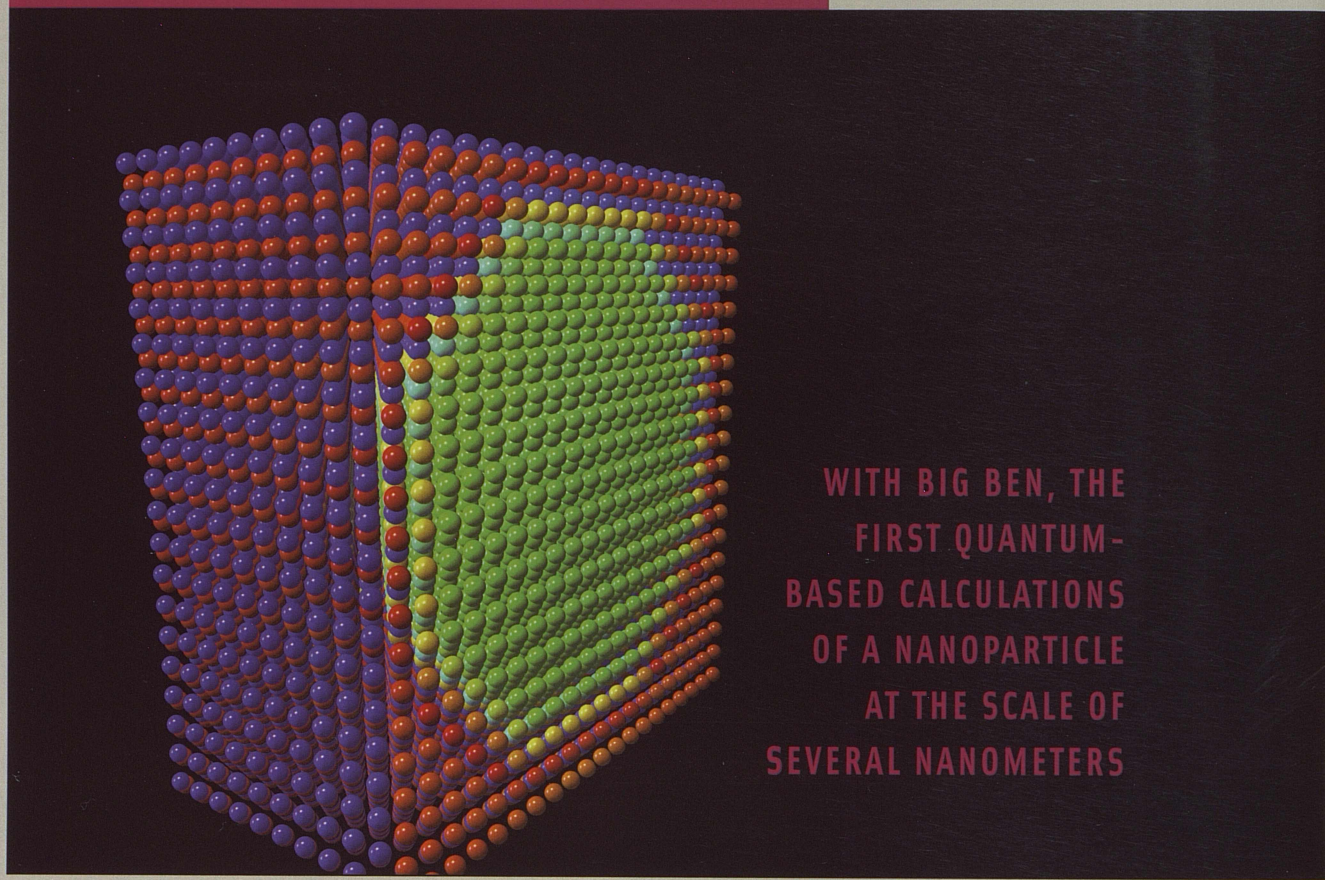

WITH BIG BEN, THE FIRST QUANTUM-BASED CALCULATIONS OF A NANOPARTICLE AT THE SCALE OF SEVERAL NANOMETERS

DWARF MATERIALS

Nanotechnology — a world of very small things (nano, Greek for dwarf) with big potential to affect our lives. Basically, nanotechnology means the ability to design materials at the scale of one to 100 nanometers. How small is this? Imagine a human hair sliced into 10,000 length-wise slivers. The diameter of each would be about one nanometer — a billionth of a meter.

Starting in the 1980s, new abilities to visualize and synthesize materials at the scale of individual atoms opened up possibilities for many applications, including supersmall, superfast electronics and molecule-sized particles to deliver insulin. Because the atomic structure of molecules governs a material's properties, nanotechnology creates the potential to craft materials atom-by-atom according to the properties — such as hardness, strength, electrical conductivity — desired. To take advantage of this potential, however, requires precise understanding of atom-to-atom interactions.

The ability to predict these interactions depends on quantum-theory computations, but even the simplest nanostructures involve thousands of atoms, imposing huge computational demands. One approach is highly efficient software called LSMS (locally self-consistent multiple scattering method). When implemented at PSC to take advantage of massively parallel processing on the Cray T3E, LSMS became the first research software to achieve sustained performance over one teraflop (a trillion calculations per second). For this accomplishment, a team from Oak Ridge National Lab, the National Energy Research Scientific Computing Center, the University of Bristol (UK) and PSC won the 1998 Gordon Bell Prize, given for the top achievement in high-performance computing.

In recent work, PSC senior scientist Yang Wang and colleagues at Oak Ridge and Florida Atlantic University used LSMS and Big Ben, PSC's Cray XT3, to calculate electronic and magnetic structure of an iron nanoparticle five nanometers in diameter. Composed of 4,409 atoms, the nanoparticle was embedded in an iron-aluminide (FeAl) matrix of 11,591 atoms — for a total of 16,000

atoms. A cross-sectional slice from the simulation shows charge distribution on the atoms (blue-positive to red-negative). In the nanoparticle itself, neutral Fe (green) is bounded by Fe atoms that lose electrons (light blue) and others that gain (yellow). Other boundary Fe atoms become more negative (red). In the matrix, Al atoms (blue) lose electrons to Fe atoms (orange). These are the first quantum-based calculations of a physical system several nanometers in scale.

SCHOOL VOUCHERS

What would happen if many families in the largest U.S. cities receive vouchers they could use to send their children to private school? This issue has been debated for many years, sometimes with considerable heat. There are many uncertainties. How might vouchers affect where families choose to live, thereby shifting populations, property values and the tax base of public-school funding? What would happen to the quality of public schools?

To bring better understanding to these and other uncertainties, Maria Ferreyra, assistant professor of economics at Carnegie Mellon's Tepper School of Business, turned to supercomputing. She created a "general equilibrium" model — a kind of economic model that accounts for a broad range of factors affecting individual decisions. Ferrerya's model went signficantly beyond previous work in this area, incorporating a number of factors — such as idiosyncratic taste for location and school choice, household religious preference, and Catholic and non-Catholic private schools — that give a deeper connection to the reality of the school-choice process.

The first part of Ferreyra's task was to develop a model without vouchers that would reproduce the reality of location and school-choice in U.S. metropolitan areas. Using census data, she focused on New York, Chicago, Philadelphia, Detroit, Boston, St. Louis and Pittsburgh. Because the model needed to be run thousands of times, this was a supercomputing-scale project. She turned to PSC's pool of 20 Intel Pentium 4s (2.4 GHz, 512 MB RAM). "The amount of computing involved was enormous," says Ferreyra, "This required about three weeks of computing on the PSC Condor pool and would have taken well over a year on a single desktop machine. I couldn't have done this without supercomputing."

After this process of "estimating" the model, Ferreyra simulated two different large-scale voucher programs in Chicago: one with universal vouchers and another with vouchers restricted to non-sectarian schools. This graphic of the Chicago area indicates boundaries between census tracts (fine lines) and school districts (colors). For the model estimation and predictions, Ferreya aggregated school districts into pseudo-districts (thick black lines) with neighborhoods of roughly the same number of housing units. The model predictions were surprising in several respects. Because vouchers allow more latitude to choose schools independent of household location, many households locate in less expensive neighborhoods than they would without vouchers and send their children to private schools. There is much movement and many private schools open in both the inner city and suburbs. Whether or not public schools improve depends on where they are located, with improvement that might not be anticipated occuring in neighborhoods of relatively low-cost housing.

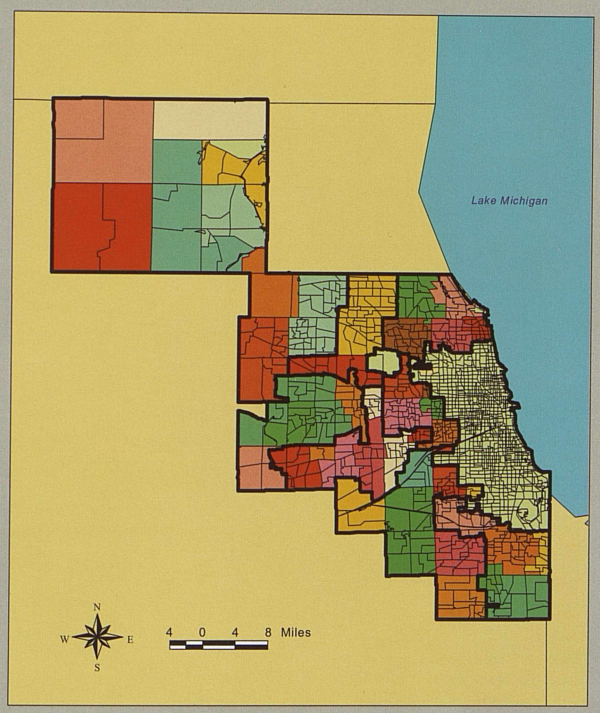

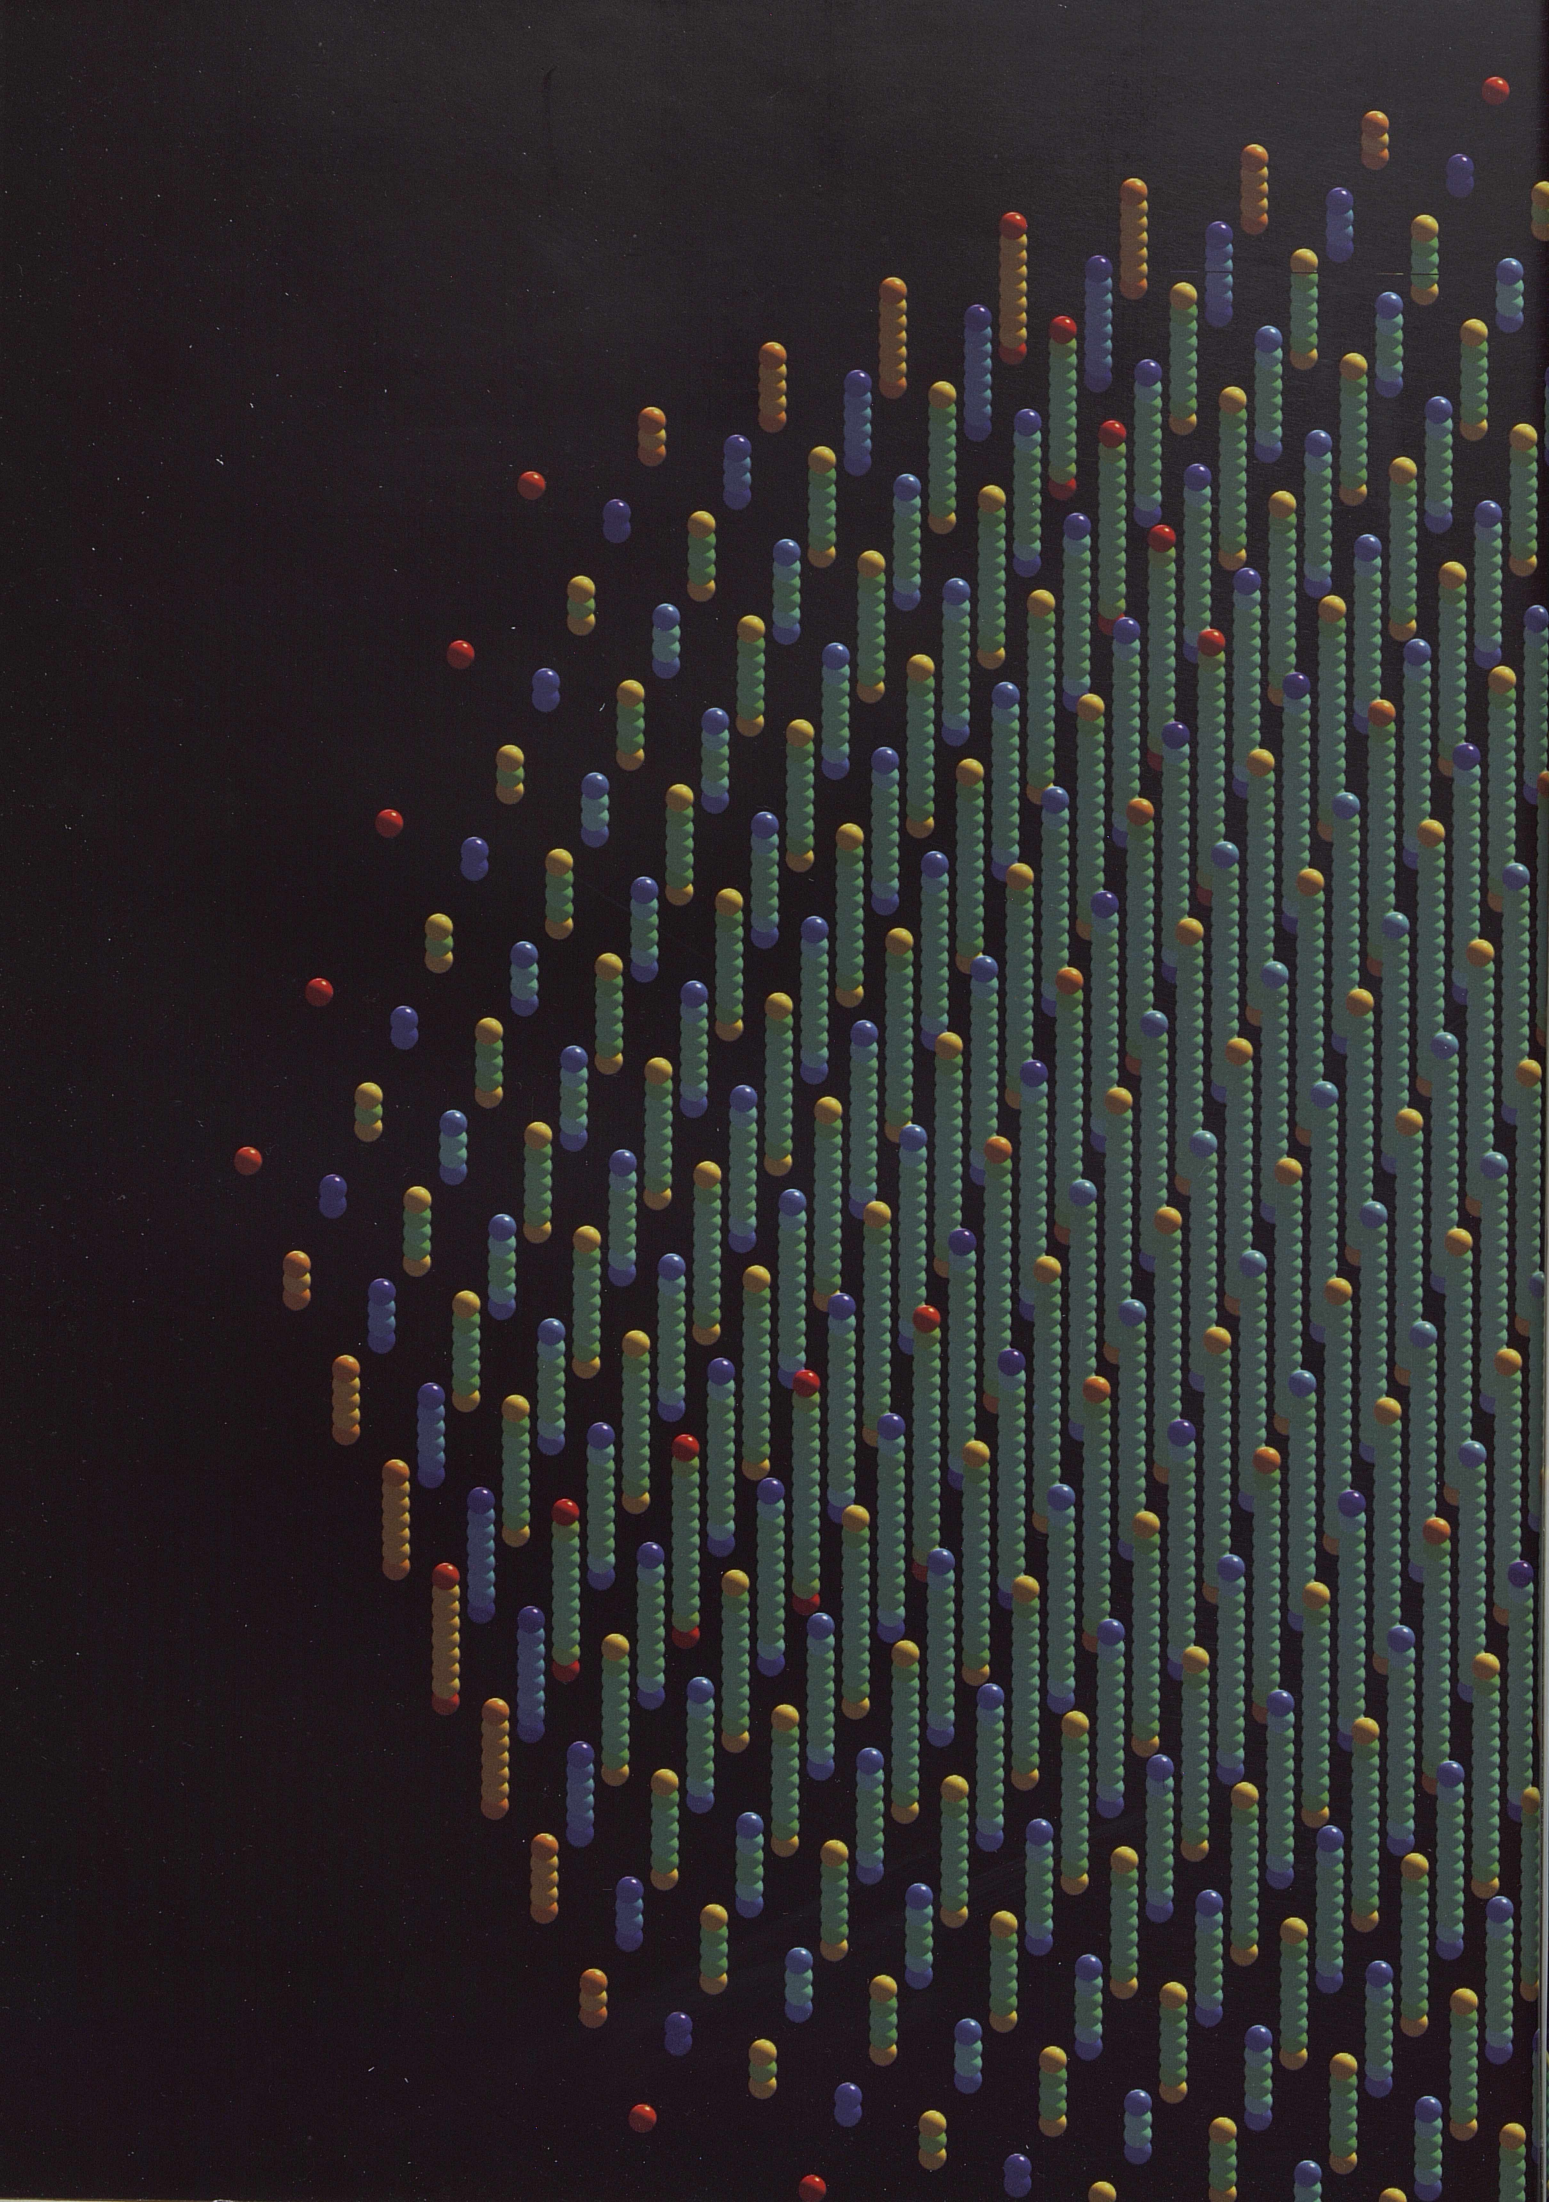

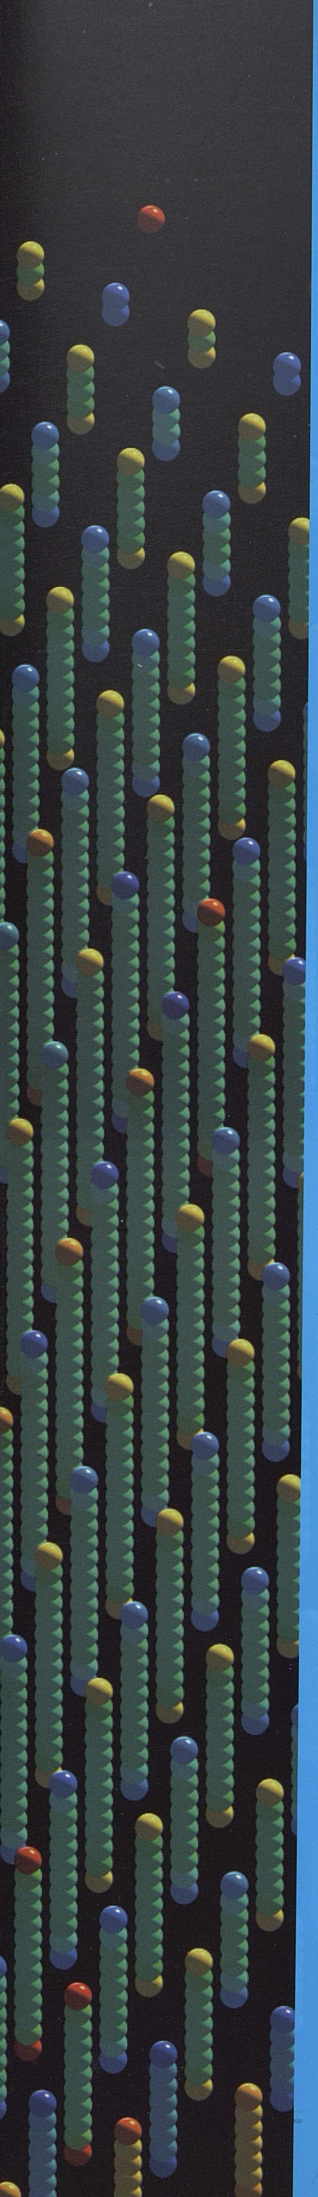

PITTSBURGH
SUPERCOMPUTING
CENTER

The Pittsburgh Supercomputing Center is a joint effort of Carnegie Mellon University and the University of Pittsburgh together with Westinghouse Electric Company. It was established in 1986 and is supported by several federal agencies, the Commonwealth of Pennsylvania and private industry.

PSC gratefully acknowledges significant support from the following:

The Commonwealth of Pennsylvania
The National Science Foundation
The National Institutes of Health
The National Energy Technology Laboratory
The U. S. Department of Defense
The U. S. Department of Energy
Cisco Systems, Inc.
Cray Inc.
Hewlett-Packard Company
The Grable Foundation

EDITOR/WRITER: Michael Schneider, PSC

CONTRIBUTING WRITERS: Wayt Gibbs, Katie Greene, Tim Palucka

GRAPHICS RESEARCH, PHOTOGRAPHY DIRECTION, COPY EDITING: Sean Fulton, PSC

TRANSCRIBING: Ryan Manzer, Courtney Machi, Megan Hackworth, PSC

PHOTOGRAPHY: Photography & Graphic Services at Mellon Institute. L. Barry Hetherington. Westinghouse Electric Company. Pittsburgh Public Schools. Pittsburgh Steelers.

GRAPHICS: Thanks to Greg Foss, PSC, for turbulence and nanoparticle graphics and for his help in general. And to Sean Fulton for re-renderings. P. Rhines, Geophysical Fluid Dynamics Laboratory, University of Washington (p. 40 image).

DESIGN: Wall-to-Wall Studios, Inc.

PRINTING: First Impression Printing

Printed on Sappi McCoy Gloss Paper, a premium sheet with 10 percent post-consumer waste fiber, with vegetable-based inks.

PITTSBURGH SUPERCOMPUTING CENTER
MELLON INSTITUTE BUILDING
4400 FIFTH AVENUE
PITTSBURGH, PENNSYLVANIA 15213

NONPROFIT ORG
U.S. POSTAGE
PAID
PITTSBURGH, PA
PERMIT NO. 251