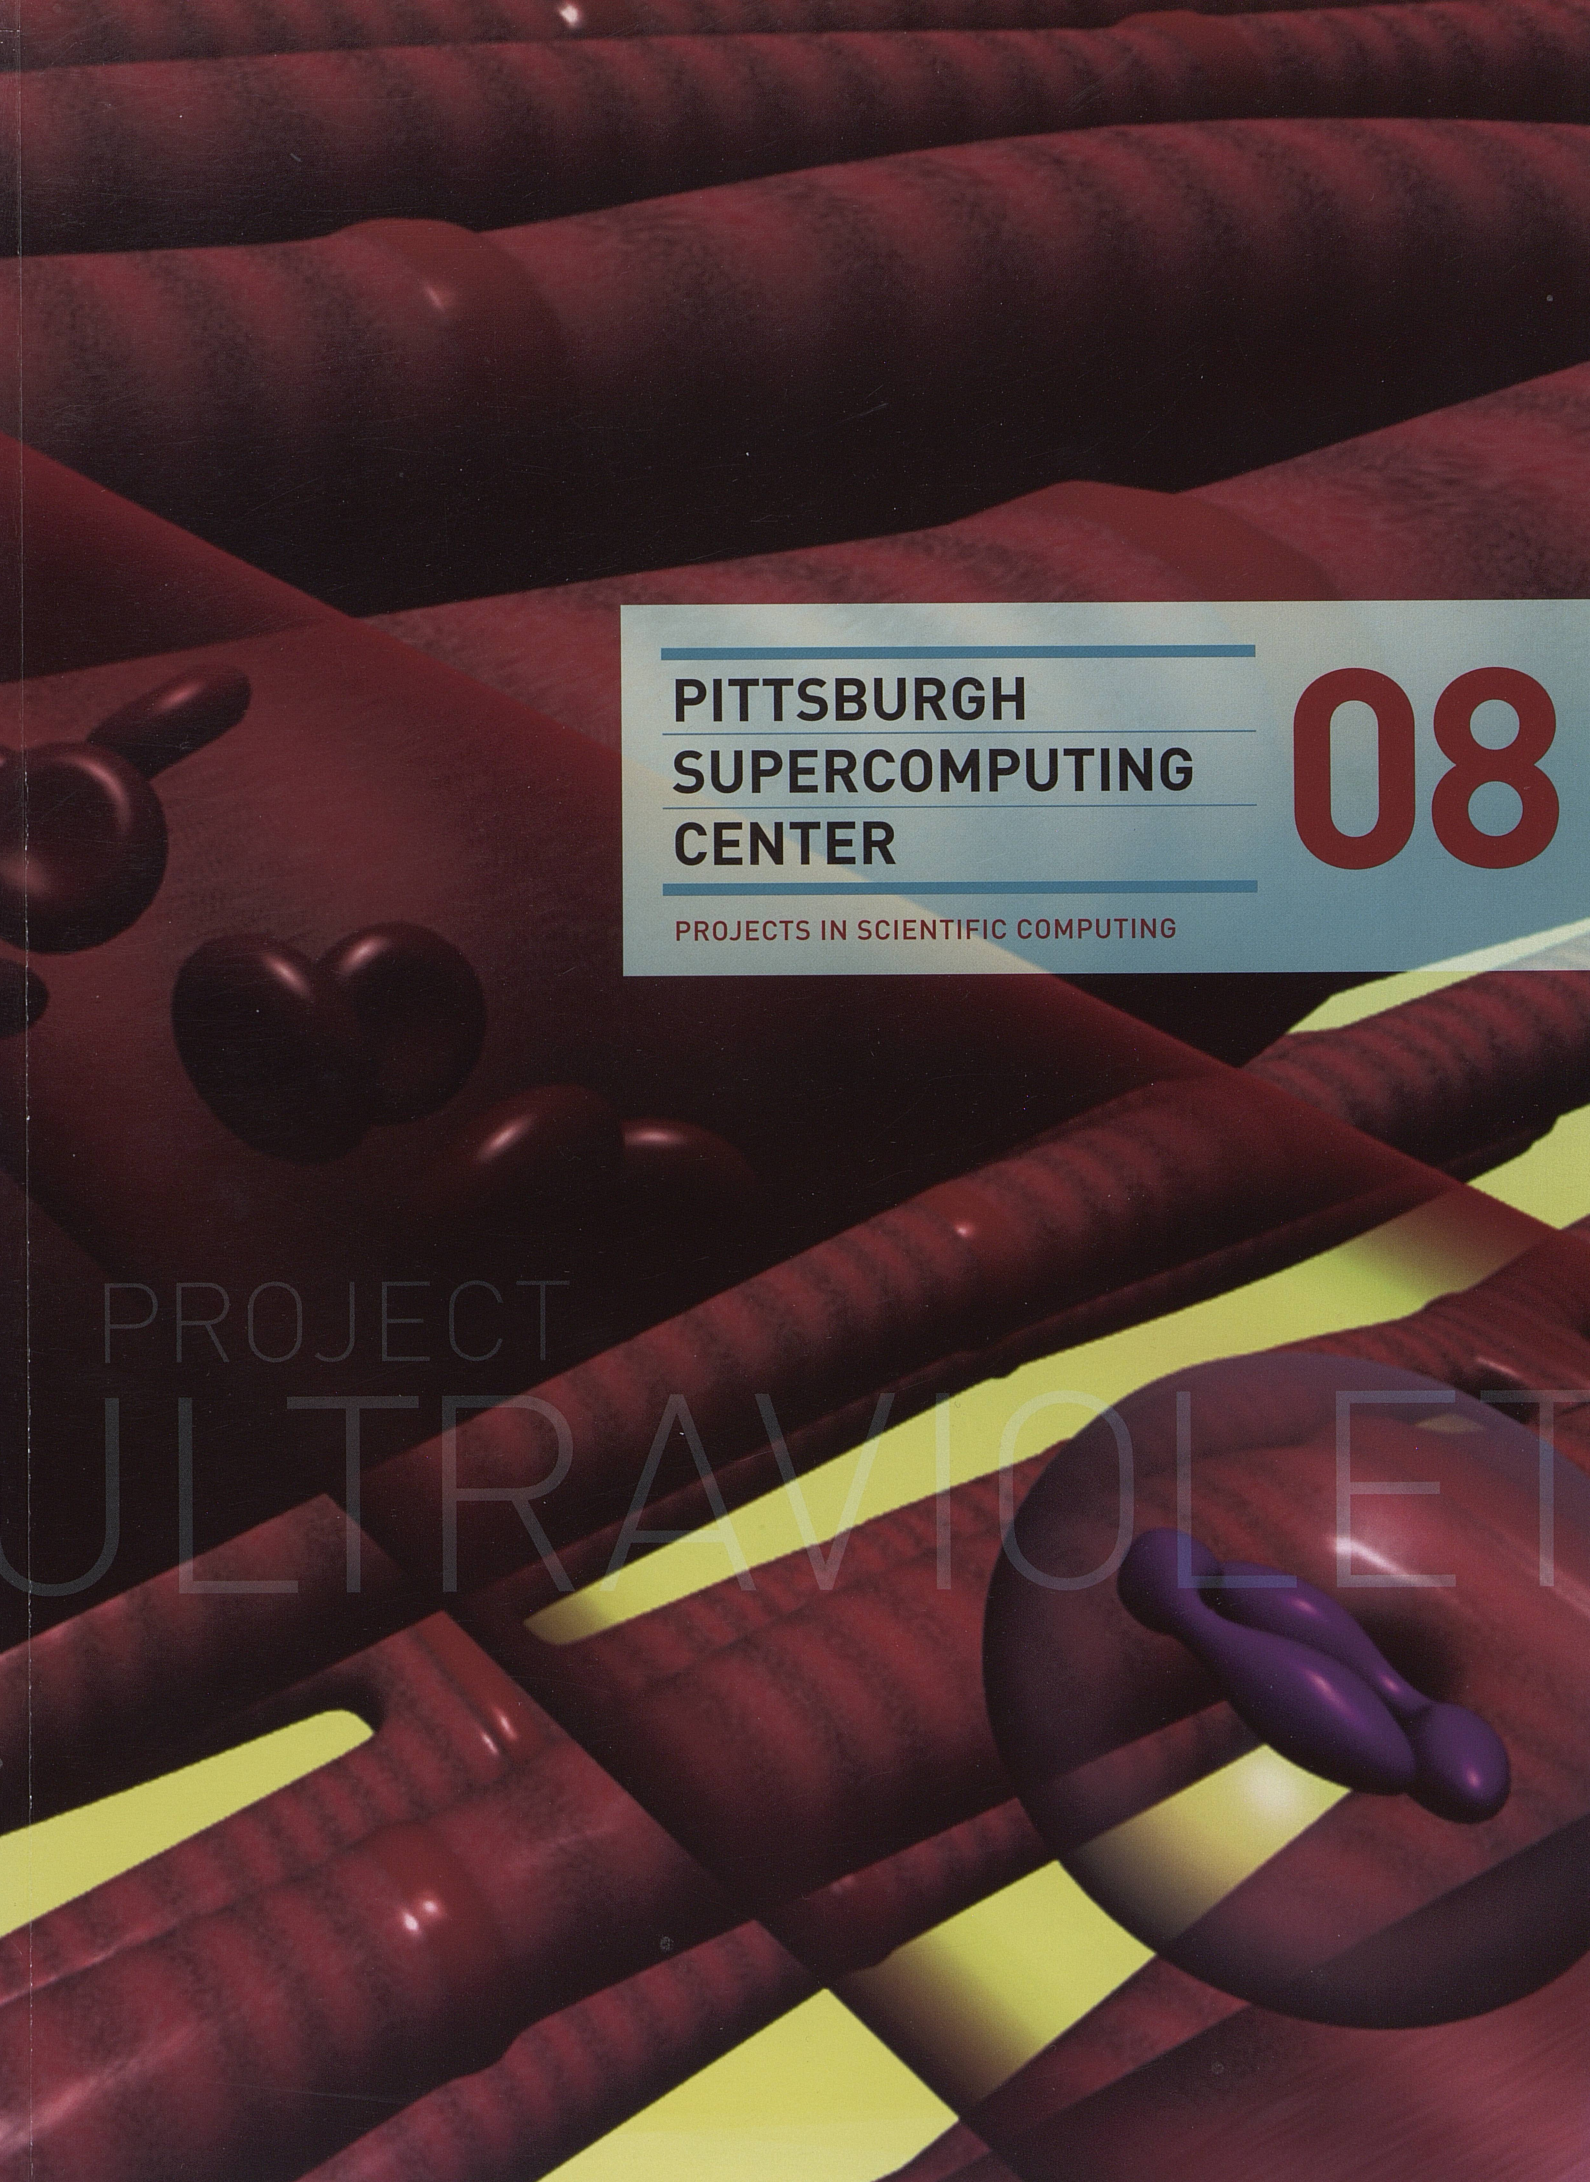
PITTSBURGH
SUPERCOMPUTING
CENTER
08
PROJECTS IN SCIENTIFIC COMPUTING
PROJECT
ULTRAVIOLET

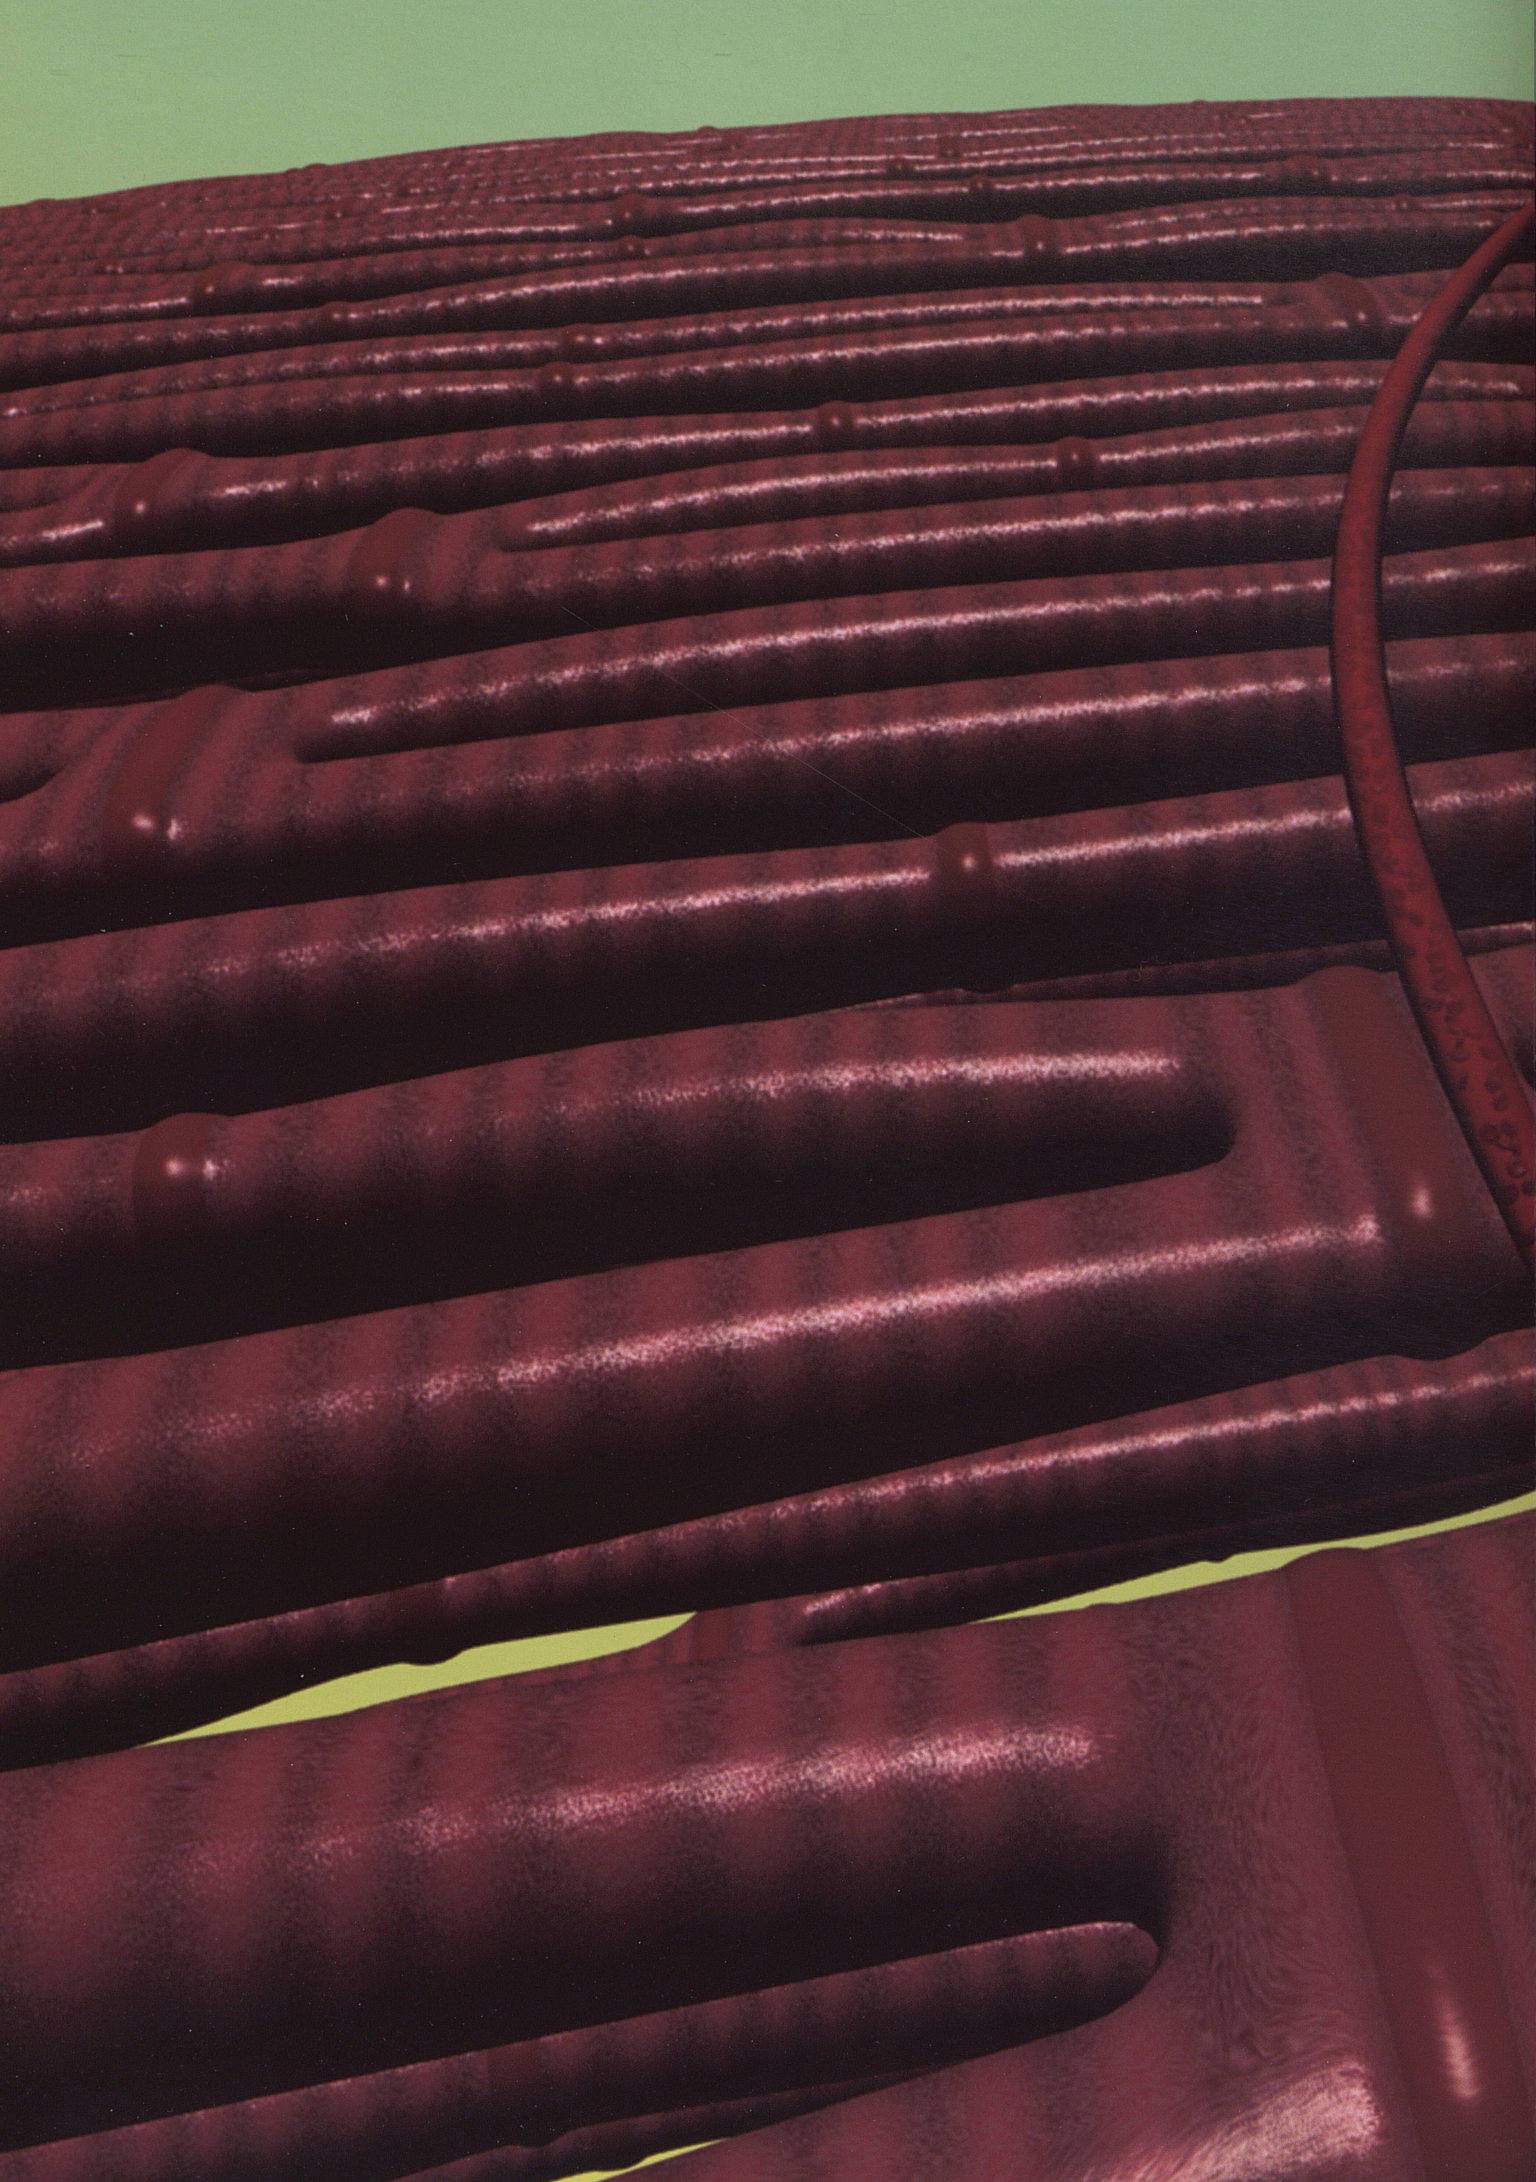

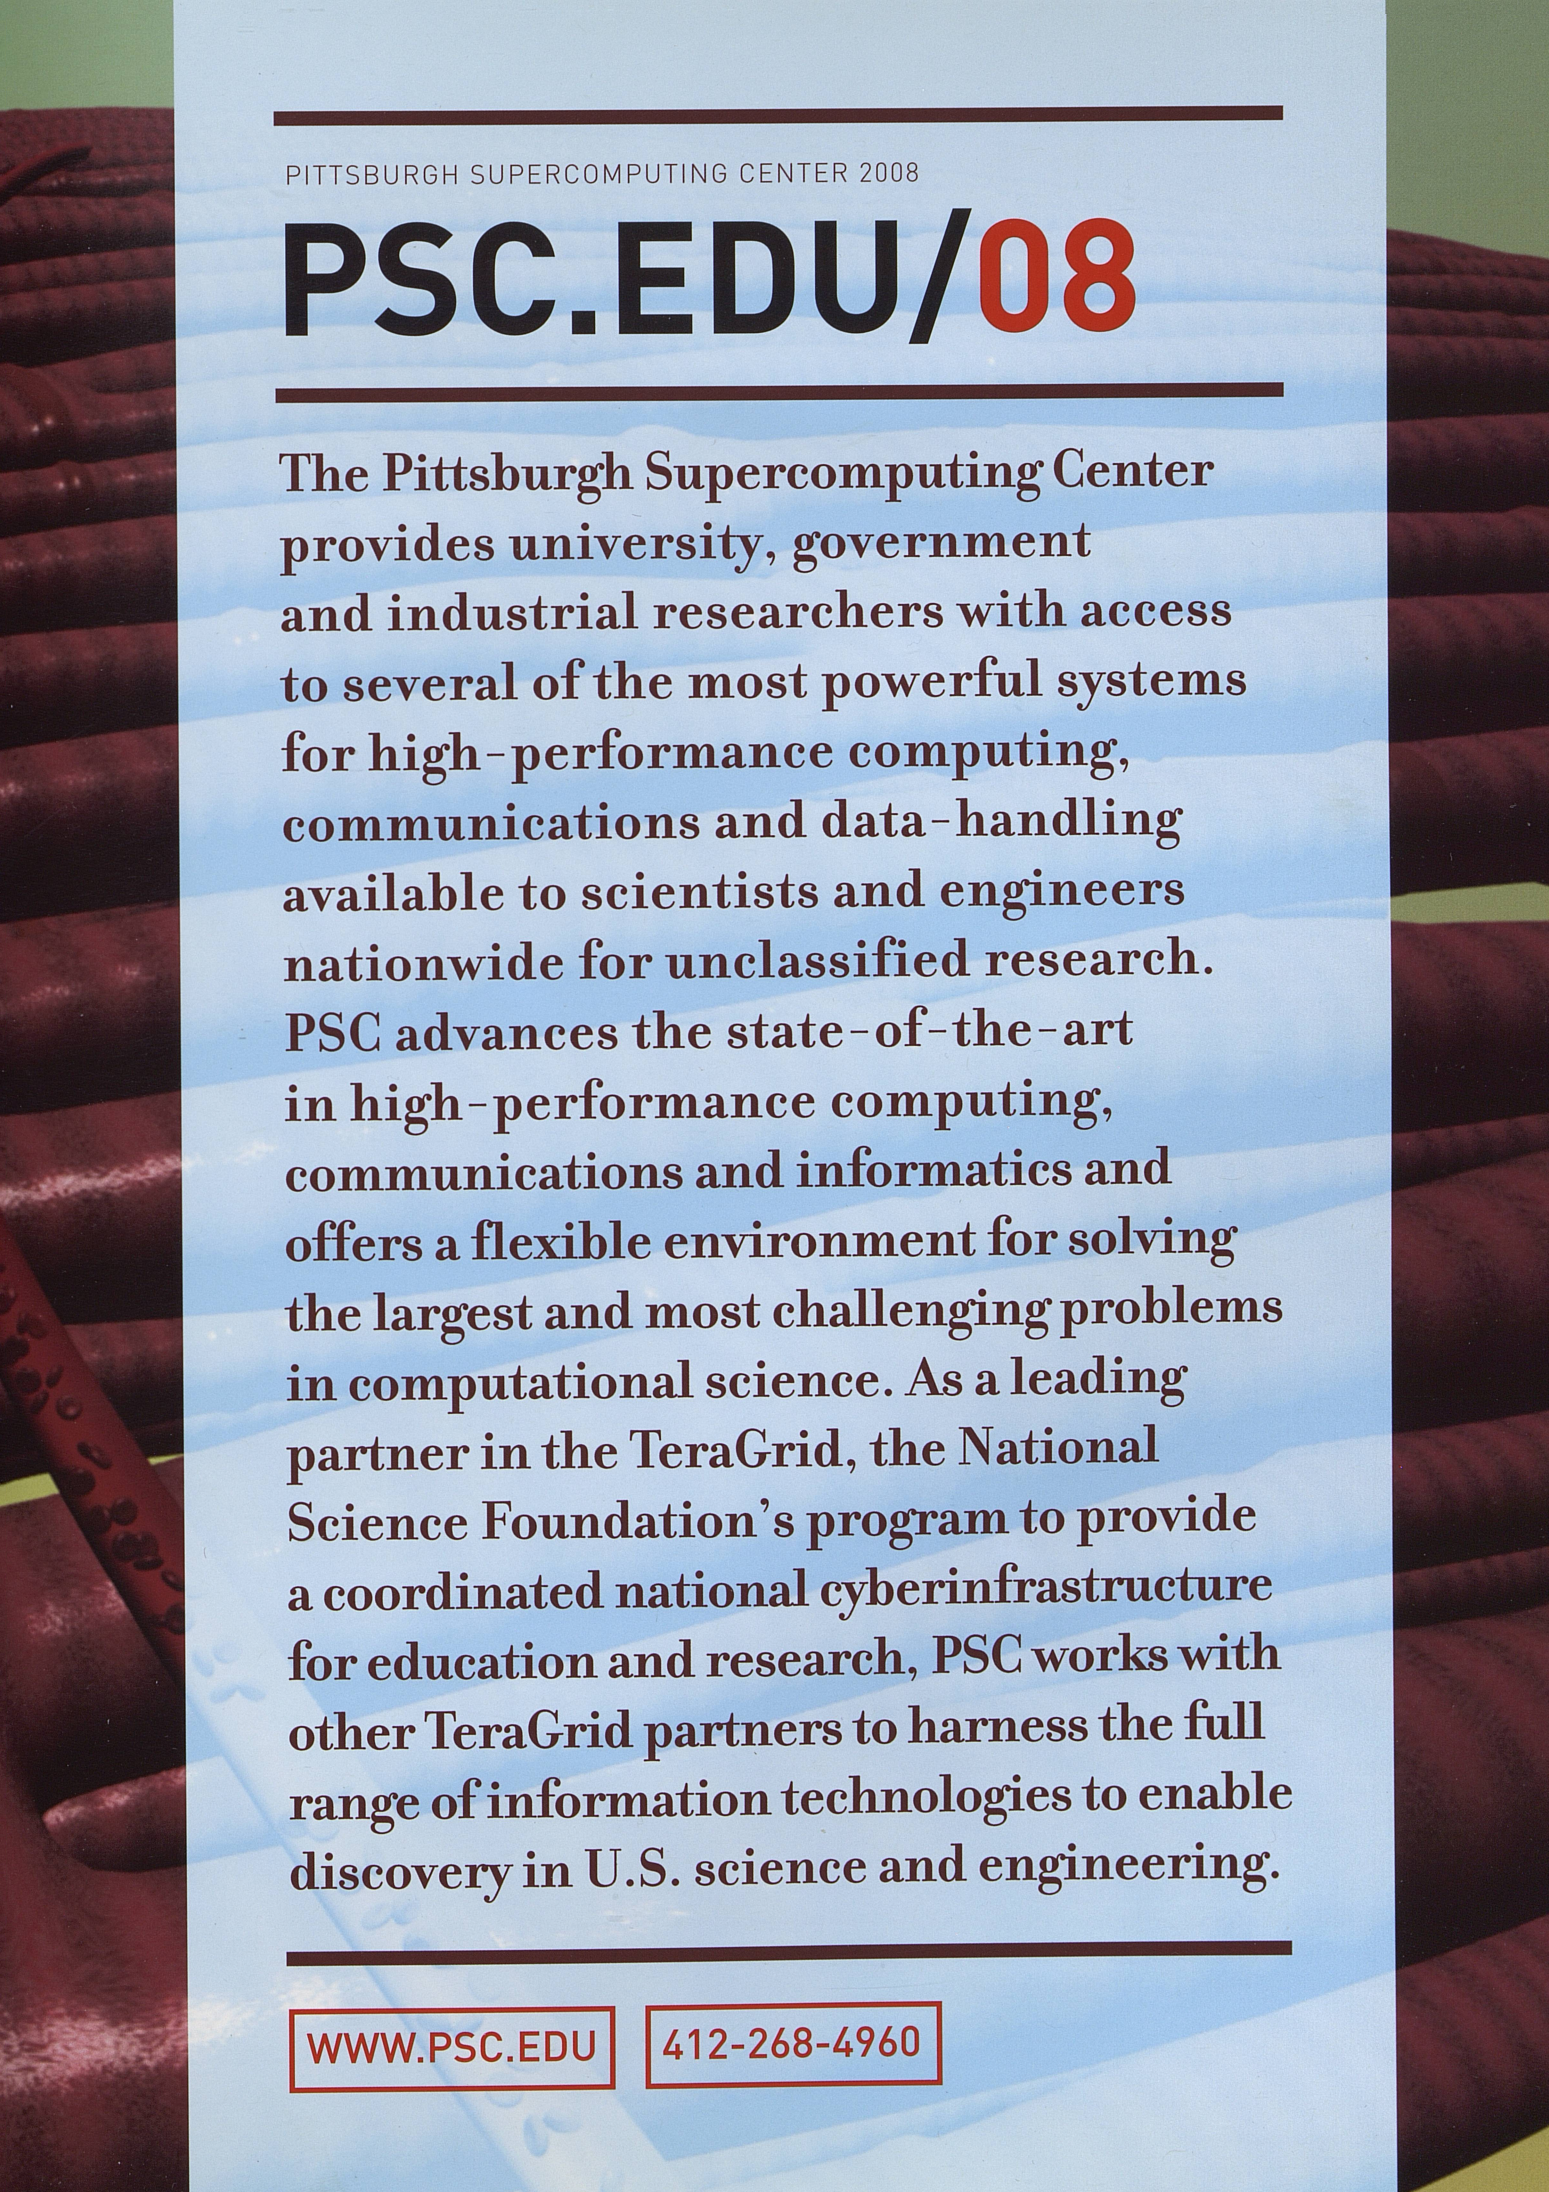

PITTSBURGH SUPERCOMPUTING CENTER 2008

PSC.EDU/08

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

Michael Levine → and Ralph Roskies, PSC co-scientific directors.

FOREWORD FROM THE DIRECTORS

This year's report on research at the Pittsburgh Supercomputing Center finds us on the cusp of a transformative leap in the technology of computational science. After installing a large shared-memory system in March, we learned in April that the National Science Board approved our proposal to implement a powerful, new petascale-level, shared-memory system (see p. 4). The next years will be exciting for all of us at PSC as we work to bring this latest of many PSC first-of-their-kind systems into being as a powerfully productive tool in the advancement of U.S. science and engineering.

Once again, we're pleased to highlight some of the recent work at PSC. Two structural biology projects exemplify the remarkable insights facilitated through computational simulation and modeling at the level of biomolecules. Alex MacKerell teamed up with his experimentalist colleague Paul Shapiro for some exciting work on ERKs (p. 18), enzymes involved in many aspects of cell function that have been implicated in various cancers. MacKerell's computations identified a list of compounds with potential to inhibit some of these cancer-inducing ERK-regulated pathways, and Shapiro's lab work confirmed that 10 of these compounds show ability to limit cancer growth in vivo. They have applied for patents on these compounds.

Maria Kurnikova at Carnegie Mellon models ion channels, and her recent insights into the glutamate receptor (p. 22), a structure that controls transmission of glutamate, the most prevalent neurotransmitter in the central nervous system, suggest possibilities for new drug design.

The Quake Group at Carnegie Mellon has evolved their abilities to simulate earthquake soil vibration as computational technology has advanced. In recent work, they created an award-winning animation from PSC computations, and in collaboration with the Southern California Earthquake Center they carried out an important validation study among three different strategies for earthquake modeling (p. 26).

Over the past 20 years, PSC has contributed to a number of corporate-research projects, including ALCOA's "lightweighting" of beverage cans and aluminum car design, turbine designs for Westinghouse, and tundish mixing for United States Steel. This year we report on two corporate projects in product development (p. 30) — a novel catheter design and sunglasses that change color automatically — singled out by the Council on Competitiveness as exemplifying how high-performance computing can accelerate time-to-market for new products.

The Joannopoulos group at MIT has done path-breaking work in physics at PSC for many years, making major contributions to the new field of photonics. In recent work (p. 34), they've developed a numerical approach that has allowed them to produce fascinating insights into a quantum property known as the Casimir force. This work has potential to facilitate breakthroughs in nano-engineering.

Alexei Kritsuk and colleagues at the University of California, San Diego did massive computations on PSC's BigBen and other TeraGrid resources. This work produced new understanding of turbulence as it affects compressible hypersonic flows, such as occur in the hydrogen-dense clouds of space where stars form (p. 38).

PSC continues to be an important resource for research in Pennsylvania (p. 6) and, through the Super Computing Science Consortium (p. 8), we help to promote economic development in southwest Pennsylvania and West Virginia and important work at NETL and elsewhere on development of clean-fuel technologies.

This publication testifies to the skill, experience and hard work of the PSC staff. We salute them in this work. All of us are 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.

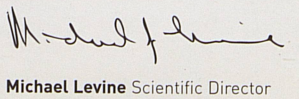

Michael Levine Scientific Director

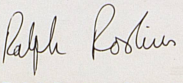

Ralph Roskies Scientific Director

TABLE OF CONTENTS

CREATING CYBERINFRASTRUCTURE

The TeraGrid is the world's most comprehensive distributed cyberinfrastructure for open scientific research. As a major partner in this National Science Foundation program, PSC helps to shape the vision and progress of the TeraGrid.

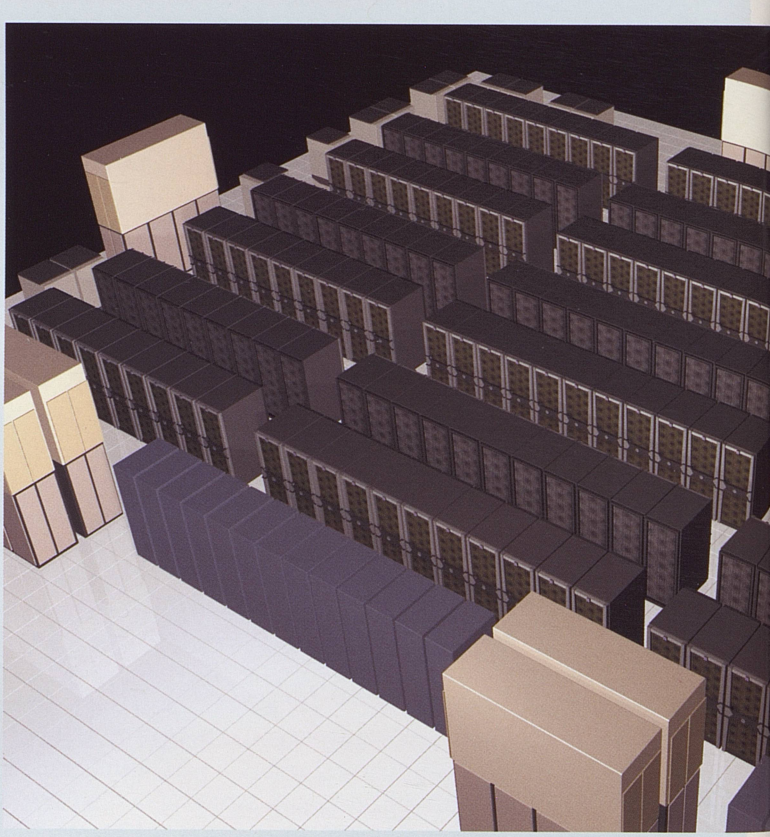

PSC's proposed new system, as rendered by PSC scientific visualization specialist Greg Foss.

ULTRAVIOLET: ALMOST VISIBLE

The National Science Board in April authorized NSF to award grants for PSC to acquire and operate a large-scale system based on Silicon Graphics' newest shared-memory architecture, called Project Ultraviolet. With this system — to be integrated into the TeraGrid — PSC will help lead U.S. scientists and engineers into the next generation of scientific computing.

PSC's Project Ultraviolet-based system will comprise more than 100,000 Intel next-generation cores — creating a shared-memory system of unprecedented scale (more than 100 terabytes), which will significantly extend TeraGrid capability for data-intensive and non-traditional applications, such as epidemiological modeling, machine learning and game theory.

An advanced Silicon Graphics high-bandwidth, low-latency interconnect called NUMAlink5 will link the processors, providing both shared-memory support and acceleration for message-passing among processors that will enhance scalability. The mass-storage system will incorporate an innovative disk/file system called ZEST, developed at PSC, (see sidebar, PSC Awards for Innovation).

The new system is one of three being implemented nationally through an NSF initiative to provide "petascale" computing for science and engineering by 2010. Petascale refers to supercomputers that can operate at "petaflop" levels of performance — a quadrillion (10^{15}) calculations per second, roughly equivalent to about 100,000 of the latest laptop systems. The other two new NSF systems are distributed-memory architecture. Because PSC's system is a shared-memory system, it will complement the others and extend the capability available to U.S. scientists and engineers.

Delivery of the new system will begin in fall 2009. Earlier in 2009, PSC plans to install a small prototype for testing and optimizing of application software to run on the new architecture.

Jim Kasdorf, PSC director of special projects

POPLE & SALK: NEW SHARED-MEMORY SYSTEMS

In March, PSC took delivery of two new Silicon Graphics Altix® 4700 systems. One, named "Pople" — for Nobel-Prize-winning chemist John Pople — features 768 processors, 1.5 terabytes of memory and peak performance of 5.0 teraflops. Pople became a TeraGrid production resource in July, substantially increasing the "shared memory" capability available through NSF for U.S. science and engineering research.

The other, named "Salk" for Jonas Salk, was acquired with support from NIH's National Center for Research Resources for NRBSC, PSC's biomedical program. Salk features 144 processors and 288 gigabytes memory and is devoted exclusively to biomedical research. "To make a system of this scale openly available for biomedical research is unprecedented," said NRBSC director Joel Stiles.

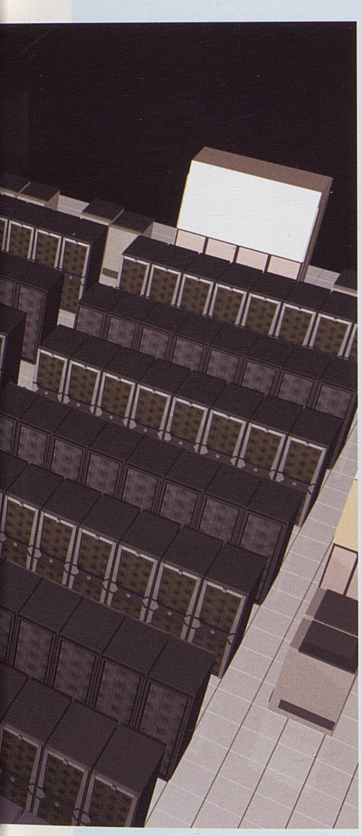

Both new systems feature shared memory, which means that the system's main memory can be directly accessed from all processors, as opposed to distributed memory (in which each processor's memory is directly accessed only by that processor). Because all processors share a single view of data, a shared memory system is relatively easy to program. The usability features of these two new systems have attracted new researchers in work involving data-analysis, computer science and other projects.

PSC AND TERAGRID

PSC is actively involved in TeraGrid leadership. Scientific director Ralph Roskies serves on the executive steering committee of the Grid Infrastructure Group that guides TeraGrid. Co-scientific director Michael Levine is PSC's representative to the TeraGrid Forum — TeraGrid's principal decision-making group.

Other PSC staff with TeraGrid leadership roles include Sergiu Sanielevici, Area Director for User Support and Jim Marsteller, head of PSC's security, who chairs the TeraGrid Security Working Group. Laura McGinnis plays a lead role in TeraGrid education, outreach and training (EOT) activities. PSC director of systems and operations, J. Ray Scott, leads the TeraGrid effort in Data Movement, and PSC director of strategic applications, Nick Nystrom, leads the TeraGrid Extreme Scalability Working Group, which fosters planning to meet the challenges of deploying extreme-scale resources into the TeraGrid.

PSC staff members serve on all of TeraGrid's working groups.

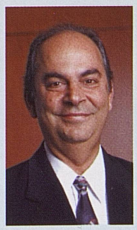

← **David Moses,** PSC executive director, oversees the day-to-day activities of PSC's scientific and technological staff.

← **PSC staff whose work contributes to TeraGrid include (l to r):**

Anirban Jana, Phillip Blood, Laura McGinnis (seated), Marcela Madrid, Nick Nystrom (seated), Shawn Brown, Kathy Benninger, Shandra Williams, James Marsteller, Derek Simmel, Sergiu Sanielevici;

(rear, l to r):

Rob Light, Ed Hanna, Joseph Lappa, R. Reddy.

Not in picture:

Michael Schneider, Josephine Palencia, David O'Neal, Rich Raymond & J. Ray Scott.

TERAGRID RESOURCE PROVIDERS

- Indiana University
- Louisiana Optical Network Initiative
- National Center for Supercomputing Applications
- National Center for Atmospheric Research
- National Institute for Computational Sciences
- Oak Ridge National Laboratory
- Pittsburgh Supercomputing Center
- Purdue University
- San Diego Supercomputer Center
- Texas Advanced Computing Center
- The University of Chicago/Argonne National Laboratory

PSC'S ZEST & NRBSC WIN AWARDS FOR INNOVATION

Last November at the Supercomputing 2007 conference in Reno, Nevada, HPCwire, a leading electronic news outlet for high-performance computing and communication, awarded two of its 2007 Reader's Choice Awards for innovation to PSC:

- The National Resource for Biomedical Supercomputing, PSC's biomedical research program (see p. 8), won for "Most Innovative Use of HPC in the Life Sciences," and
- ZEST, a PSC-developed file system that facilitates scientific computing on very large-scale (petascale) systems, won for "Most Innovative HPC Storage Technology or Product."

Developed by PSC's advanced systems group, ZEST is a distributed file-system infrastructure that vastly accelerates write bandwidth, achieving more efficient backup reliability (higher "checkpointing bandwidth") than any other program available.

SUPERCOMPUTING IN PENNSYLVANIA

With Commonwealth of Pennsylvania support, PSC provides education, consulting, advanced network access and computational resources to scientists and engineers, teachers and students across the state.

Pennsylvania State Representatives Dwight Evans, 203rd District, Philadelphia County (right center), and Jake Wheatley, 19th District, Allegheny County (right), visited PSC on August 13. PSC scientist Joel Stiles (left) and co-scientific director Ralph Roskies (right of Stiles) demonstrated CMIST, PSC's innovative program of animated teaching modules to engage high school and undergrad science students.

Cheryl Begandy, PSC manager of outreach, coordinates PSC's programs of corporate research, education & community outreach in Pennsylvania.

PSC's well-attended CAST workshops show high-school science teachers how to effectively incorporate computational science into their curricula. "Many of the teachers," says Begandy, "expressed strong interest in an advanced CAST training, and this year we introduced a two-day workshop for teachers who completed the introductory program."

K-12 SCIENCE EDUCATION

PSC's K-12 programs to help prepare technology-ready workers and a science-literate populace continued this year. PSC's K-12 workshops trained 120 science teachers in 32 school districts in Western Pennsylvania.

In December 2006, the Heinz Endowments Education Program awarded $150,000 to support PSC's CAST (Computation and Science for Teachers) workshops, which kicked-off the 2008-09 school year with a weeklong workshop in June. Thirteen teachers learned to use software that implements the basics of mathematical modeling in the classroom. Lunchtime talks featured PSC scientists Nathan Stone and Shawn Brown and a talk about the pressing need for science education from Carnegie Mellon vice-provost for education Indira Nair.

PSC also extended its CMIST (Computational Modules in Science Teaching) program this year with a new module on atomic and molecular movements. Developed by PSC's National Resource for Biomedical Supercomputing (p. 10), CMIST creates complete teaching modules in subject areas based on recent research at PSC (see p. 46). Together, CAST and CMIST offer an approach to secondary science that includes both specific computational modules (CMIST) and a general framework for computational science disciplines (CAST).

PSC MENTORS STUDENT-SCIENCE WINNERS

Several Pittsburgh-area high-school and college students received recognition for their science projects at the third annual conference of the TeraGrid, NSF's program of cyberinfrastructure for U.S. science and education, June 9-13 in Las Vegas.

Matthew Stoffregen, a junior at Woodland Hills High, won a competition called "The Impact of Cyberinfrastructure on Your World," which invited students to showcase ways in which cyberinfrastructure will affect human communities. Three of six entries

TG '08 Student Winners (l to r): Max Hutchinson, Shivam Verma, Cadeal Chase, Hari Seshadri, Bryan Bemley (front), Tyrell Ferguson and Jessica Travierso.

nationwide were from the Pittsburgh region, including two university students who work at PSC. Shivam Verma, of North Allegheny School District, placed second and Srihari Seshadri, a Franklin Regional Senior High student and PSC student employee, came in third.

In a second competition, "TeraGrid Student Research," students described the benefits of grid computing. Maxwell Hutchinson, a PSC student programmer and Carnegie Mellon undergrad, came in second for his work with solenoids.

"This an impressive performance," said PSC's Laura McGinnis, manager of data information and resource services, who recruited and mentored area students in the TeraGrid competitions. "It represents a head start for this area in producing cyber-savvy individuals who will help to generate technological innovations and future economic growth."

COMMUNITY OUTREACH

PSC staff have taken part in numerous community outreach programs, both locally and nationally. For the first time this year, PSC was a bronze sponsor of the Pittsburgh Regional Science and Engineering Fair, where Matthew Stoffregen of Woodland Hills and Srihari Seshadri of Franklin Regional, also recognized at TeraGrid 08, received PSC awards for effectively incorporating computation into their projects. PSC also exhibited at the SciTech Spectacular at the Carnegie Science Center in October 2007. This event fosters understanding of science and technology among middle and high school students. PSC's booth featured activities that demonstrate how supercomputing can generate interest in science.

Through SC² (p. 8) and the PSC networking group (p. 12), Evergreene Technology Park in Greene County provides companies with access to PSC resources. PSC coordinated economic development outreach programs through Keystone Innovation Zones in Indiana, Johnstown, Fayette Country and Waynesburg.

The networking group worked with eight Pennsylvania Intermediate Units in support of E-Fund applications to provide wide-area network connectivity and access to educational resources on Internet2. PSC's network staff also coordinated K-20 outreach activities with MAGPI, the Philadelphia-based network hub, leveraging their well-developed K-20 programs.

PRIVATE-SECTOR & UNIVERSITY RESEARCH

The Council on Competitiveness, a Washington, DC organization of business, labor, academic and governmental leaders who focus on private-sector competitiveness, this year featured three of PSC's corporate-partner research projects in case studies (funded by NSF) that promote the use of high-performance computing to facilitate innovation in product development (see p. 30).

Research by university scientists in Pennsylvania supported by PSC is exemplified by several projects in this booklet:

- Turbulent Combustion: Peyman Givi, University of Pittsburgh (p. 9).
- Understanding the glutamate receptor: Maria Kurnikova & Tatyana Mamonova, Carnegie Mellon University (p. 22).
- Earthquake soil vibrations: Jacobo Bielak & David O'Hallaron, Carnegie Mellon University (p. 26).
- Game theory: Tuomas Sandholm & Andrew Gilpin, Carnegie Mellon University (p. 43).
- Epidemiological Modeling: University of Pittsburgh Graduate School of Public Health & PSC (p. 44).

RESEARCH AT PENNSYLVANIA COLLEGES & UNIVERSITIES, 2007-2008

From July 1, 2007 through June 30, 2008, PSC provided 7.4-million processor hours to 779 Pennsylvania researchers. PSC workshops in high-performance computing reached 240 Pennsylvania grad and undergrad students. This usage includes more than 200,000 hours of new-user allocations on PSC's Ben system (now decommissioned) along with 6.9-million hours on PSC's lead system, the 4,096 processor Cray XT3, BigBen. If purchased from a commercial provider, this computing time would be valued, conservatively, at over $7-million. The following Pennsylvania universities and colleges used PSC resources during this period:

Allegheny-Singer Research Institute
Bryn Mawr College
Bucknell University
Cabrini College
Carnegie Mellon University
Cedar Crest College
Cheyney University of Pennsylvania
Drexel University
Duquesne University
Franklin and Marshall College
Immune Tolerance Network
Indiana University of Pennsylvania (all campuses)
Juniata College
Lehigh University
Lock Haven University
Pennsylvania State University (all campuses)
Shippensburg University of Pennsylvania
Temple University
University of Pennsylvania
University of Pittsburgh (all campuses)
University of the Sciences in Philadelphia
Ursinus College
Villanova University
Waynesburg College
Widener University

THE SUPER COMPUTING SCIENCE CONSORTIUM

Pennsylvania-West Virginia partners in development of clean power technologies.

↑
(SC)² co-chairs Lynn Layman, PSC (left) & Bob Romanowsky, NETL

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 and education.

Through (SC)², Evergreene Technology Park in Greene County provides a resource that supports and encourages companies to collaborate with local universities in southwest Pennsylvania and West Virginia and to have access to PSC.

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.

Turbulent combustion and propulsion predictions with (a) direct numerical simulation (DNS), (b) large-eddy Simulation (LES), and (c) Reynolds average simulation (RAS). Currently RAS is most commonly employed in industry, but its range of validity is limited. DNS is the most detailed, but is too computationally demanding for most realistic engineering problems. LES is a compromise between the two and provides excellent reliability and applicability. NASA has applied LES/FDF to research on hyper-velocity propulsion, and gave the NASA Public Service Medal to Givi for developing this method.

↓

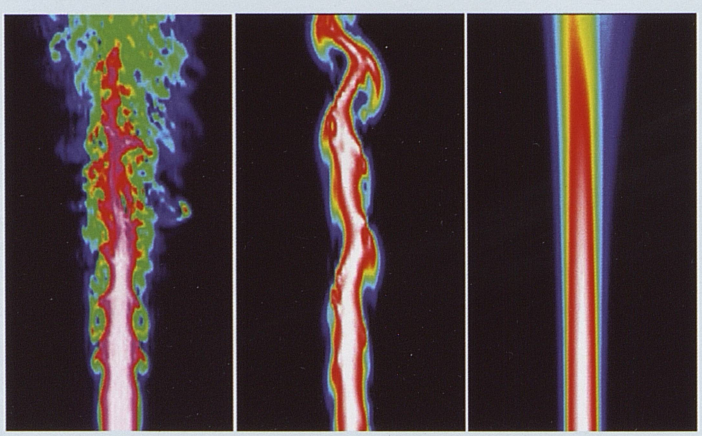

a b c

HIGH-FIDELITY SIMULATION OF TURBULENT COMBUSTION

Life in the 21st-century runs on electricity, which comes from turbines — jet engines bolted to the floor. In gas turbines, combustors ignite fuel and blast hot, pressurized gas to do the turning work that produces megawatts of electricity. High efficiency — as complete as possible conversion of raw energy into turbine rotation — is the key not only to low emission of pollutants, but also to the cost of electricity. Small gains that slightly reduce cost per megawatt translate to huge savings overall, and turbine engineers measure efficiency in tenths of a percent.

"High-fidelity simulation methodologies are indispensable to modern gas-turbine design," says Peyman Givi, William Kepler Whiteford Professor of Mechanical Engineering and Materials Science (MEMS) at the University of Pittsburgh. "Turbine engineers face significant challenges to meet performance targets and emission standards."

The MEMS department at Pitt has long been a leader in research on gas turbines, and through (SC)², Givi uses PSC resources to develop methods for realistic simulation of gas-turbine combustion. His group has implemented a proven numerical method, large-eddy simulation (LES), with an improvement called the filtered-density function (FDF). "FDF is a novel means of implementing LES in combustion systems," says Givi. "The primary advantage is that it accounts for the effects of subgrid-scale chemical reactions in an exact manner, regardless of the speed of reaction."

Givi's team has used PSC's BigBen, Rachel and Pople to address some of the most demanding problems in turbulent-combustion modeling. Givi cites crucial support from PSC scientist Raghurama Reddy. His group has applied LES/FDF to complex geometries along with realistic combustion-chemistry models. This approach is computationally expensive and has presented a serious challenge in parallel scalability. "Recently," says Givi, "we have innovated a parallelization strategy that overcomes the performance bottlenecks, and we have demonstrated scalability to massively parallel architectures. With this development, LES/FDF has proven to be a major tool for prediction of engineering combustion problems."

Byron Stauffer, executive director, Indiana Country Office of Planning and Development, welcomes participants in the KIZ event held on June 4 at Indiana University of Pennsylvania.

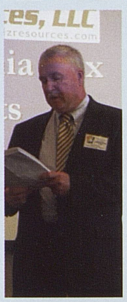

Don Chappel, executive director, Greene County Industrial Development Authority, delivers welcoming remarks at the June 11 KIZ outreach event at Waynesburg.

KEYSTONE INNOVATION ZONE

The Keystone Innovation Zone (KIZ) program is a Pennsylvania Department of Community and Economic Development initiative focused on college and universities and their surrounding areas. KIZ's provide tax incentives and funds to generate job creation through technology transfer and entrepreneurship.

PSC helped to organize two KIZ outreach efforts in June. The first, June 4 at Indiana University of Pennsylvania (IUP), was supported by the Indiana and Johnstown KIZs. Approximately 15 people from small businesses heard presentations on NETL licensing and tech-transfer opportunities, grant opportunities at PSC, Pennsylvania funding for small business IT projects and KIZ tax credits. Another outreach event, June 11 at the Waynesburg Center for Research and Economic Development, drew 44 participants.

PSC & (SC)²: RESEARCH FOR CLEAN ENERGY

Since the 1999 founding of (SC)², 51 (SC)² researchers have used PSC systems for a range of clean-energy related projects, using more than 5.4-million hours of computing time, over 500,000 hours within the past year.

THIS WORK INCLUDES:

- Modeling an Operational Clean-Coal Power Plant
  http://www.psc.edu/science/2007/coal/
- Clean Liquid Fuel from "Syngas"
  http://www.psc.edu/science/2006/sc2/
- Fuel-Quality Hydrogen from Fossil Fuels
  http://www.psc.edu/science/2005/sc2/
- Gas from Black Liquor
  http://www.psc.edu/science/2004/sc2/
- Fluidized-Bed Combustion of Silane
  http://www.psc.edu/publicinfo/netl/
- Lean-Fuel Mixes in Next-Generation Power-Generating Turbines
  http://www.psc.edu/science/Richards/clean_power.html
- A New Design for a Power-Generating Turbine
  http://www.psc.edu/science/cizmas2002.html

(SC)² Super Computing Science Consortium

(SC)² PARTNERS

National Energy Technology Laboratory
Pittsburgh Supercomputing Center
Carnegie Mellon University
Duquesne University
University of Pittsburgh
Waynesburg University
West Virginia University
NASA Independent Verification & Validation Facility
The West Virginia Governor's Office of Technology

MORE INFORMATION:
http://www.sc-2.psc.edu

The National Resource for Biomedical Supercomputing

National Leadership in High-Performance Computing for Biomedical Research

Established in 1987, PSC's National Resource for Biomedical Supercomputing (NRBSC) was the first external biomedical supercomputing program funded by the National Institutes of Health (NIH). Along with core research at the interface of supercomputing and the life sciences, NRBSC scientists develop collaborations with biomedical researchers around the country, fostering exchange among experts in computational science and biomedicine and providing computational resources, outreach and training. In October 2006, NRBSC received $8.5 million from NIH's National Center for Research Resources (NCRR) to renew its work for five years.

"Over the past decade, computing has become essential to almost all aspects of biomedicine," says PSC's Joel Stiles, director of NRBSC, a medical doctor who also holds a doctorate in physiology. "Here at the NRBSC, we're developing and distributing computational tools in simulation, visualization, and education that are helping to transform our understanding of life and disease."

MORE INFORMATION:
http://www.nrbsc.org

COMPUTATIONAL SERVICE & TRAINING

Since NRBSC's inception, PSC and NRBSC together have provided access to computing resources for more than 1,400 biomedical research projects involving more than 4,200 researchers at 274 research institutions in 46 states and two territories. Among these are several projects featured in this booklet (pp. 18 & 22).

NRBSC training activities reach hundreds of scientists each year. More than 3,500 researchers have participated in NRBSC workshops in such areas as spatially realistic cell modeling, volumetric

← The NRBSC team: (seated, l to r) Boris Kaminsky, Pallavi Ishwad, Jenda Domaracki, (standing) Art Wetzel, Jack Chang, Markus Dittrich, Troy Wymore, Christal Banks, Alex Ropelewski, Joel Stiles, Adam Kraut, Greg Hood. Not pictured: Hugh Nicholas and Jacob Czech

NRBSC BIOMEDICAL COLLABORATIONS

Albert Einstein College of Medicine
Carnegie Mellon University
Cornell University
Duke University
Harvard University
Howard University
Janelia Farm, Howard Hughes Medical Institute
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 Kansas
University of North Carolina, Chapel Hill
University of Pittsburgh
University of Pittsburgh School of Medicine
University of Puerto Rico, Medical Sciences Campus

RESEARCH

NRBSC research focuses on three areas of biomedicine that span many scales of space and time: spatially realistic cell modeling, large-scale volumetric visualization and analysis, and computational structural biology.

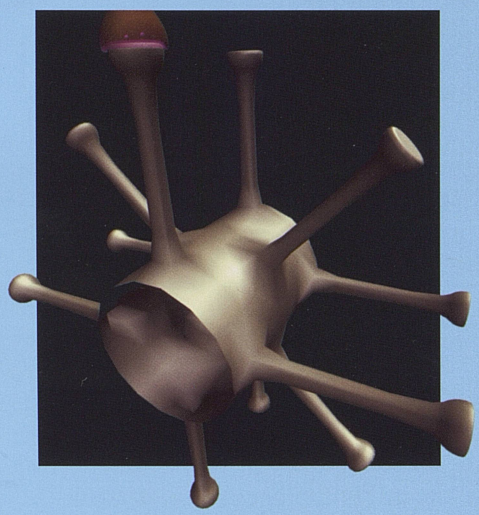

SPATIALLY REALISTIC CELL MODELING centers on realistic → 3-D simulations of movements and reactions of molecules within and between cells, to better understand physiological function and disease. *MCell*, *DReAMM* and *PSC_DX* software is developed at the NRBSC and used to model and visualize events such as this image, which represents neurotransmitter release in one dendritic spine.

data visualization and analysis, protein and DNA structure, genome sequence analysis and biological fluid dynamics."

NRBSC participates in a range of undergraduate and graduate training programs.

THESE INCLUDE:

- A joint Carnegie Mellon and University of Pittsburgh Ph.D. program in computational biology **(www.compbio.cmu.edu)**
- The Ray and Stephanie Lane Center for Computational Biology at Carnegie Mellon **(lane.compbio.cmu.edu)**
- The University of Pittsburgh Department of Computational Biology **(www.ccbb.pitt.edu)**
- The undergraduate Bioengineering & Bioinformatics Summer Institute **(www.ccbb.pitt.edu/bbsi)**, sponsored by NRBSC, Carnegie Mellon, the University of Pittsburgh, and Duquesne University, and funded jointly by NSF and NIH.

K-20 SCIENCE OUTREACH

The NRBSC and PSC have developed innovative Computational Modules In Science Teaching (CMIST) for high school and undergraduate biology, chemistry, physics, computer science and math. CMIST modules bring critical concepts to life in novel ways, using realistic models and simulations with visually appealing, scientifically accurate animations (see p. 46). NRBSC distributes the modules online and on DVDs. They include lecture slides, animations, lesson plans aligned to national and state standards, worksheets and answer keys. This year NRBSC developed a second CMIST module on atomic and molecular movements, bridging enormous space and time scales that are important to biological systems. A third module on enzyme structure and function is also currently under development.

VOLUMETRIC →
VISUALIZATION using the NRBSC's PSC_VB software enables multiple users to share, view and analyze extremely large datasets and time series obtained from light and electron microscopes, CAT and MRI scanners, etc. This image from an MRI dataset, viewed with PSC_VB's volume browser, shows detailed structure of the left ventricle of a mouse heart with an error range of less than 3-percent. Mice hearts are frequently used in cardiac research, and this high degree of accuracy facilitates extended duration studies of cardiac development and recovery after injury.

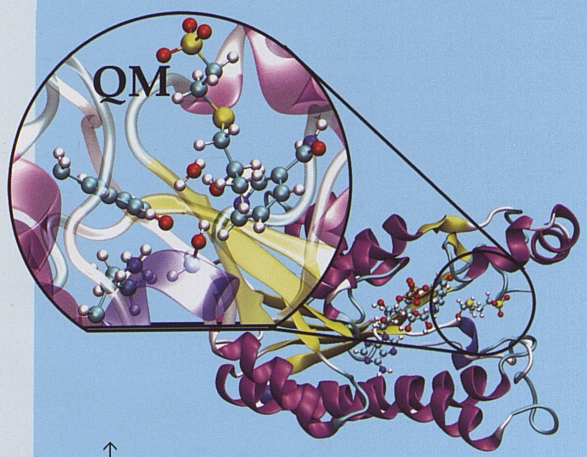

↑
NRBSC STRUCTURAL BIOLOGY focuses on computational tools used to determine the structure of proteins from their amino acid sequence and development of quantum-mechanical simulation methods for biomolecules such as enzymes. This image shows the 3-D structure of r-hpcdh, an enzyme that catalyzes a key reaction — a coupled proton/hydride transfer. PSC-developed software enables researchers to simulate enzyme reactions, to reproduce experimental reaction rates and gain new insight into enzyme function, which facilitates design of new therapeutic drugs.

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 the NPAD diagnostic server — has created valuable tools for improving network performance nationally.

MORE INFORMATION:
http://www.psc.edu/networking/

← **Tim Devlin,** program director of instructional media services, Allegheny Intermediate Unit, discusses educational impact of PSC's new PoP in Allegheny County.

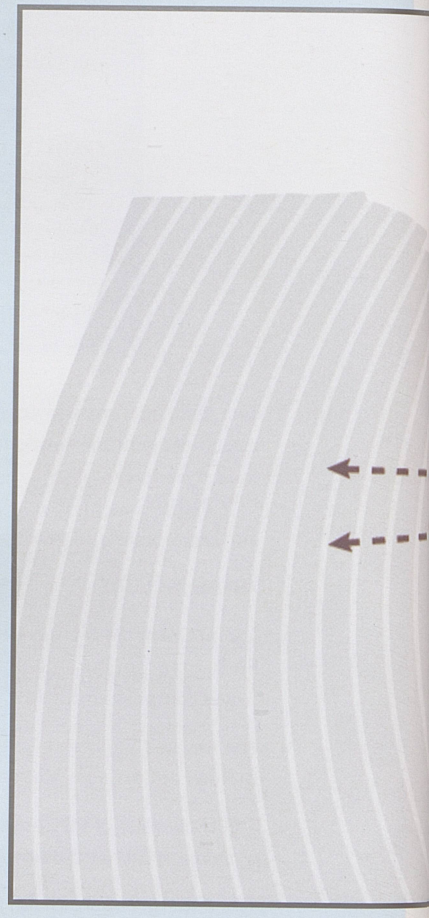

3ROX BRANCHES OUT: PSC AND FIBERTECH OPEN NEW POP

Through 3ROX, PSC connects universities and public schools in Pennsylvania and West Virginia to high-performance networks, such as Internet2, which links leading U.S. universities, corporations, government research agencies, and not-for-profit networking organizations.

This year 3ROX extended its connectivity in southwest Pennsylvania by opening a second PSC PoP (point of presence) in Fibertech's new collocation facility in Pittsburgh's Allegheny Center Mall. At an April kickoff event, PSC Director of Networking Wendy Huntoon noted in brief remarks that the new PoP is PSC's first off-campus location, making connectivity through 3ROX more cost-effective in providing underlying support for economic growth in the southwest Pennsylvania region.

Tm Devlin, program director of instructional media services, Allegheny Intermediate Unit, described the networking program of the Allegheny Intermediate Unit, 42 suburban school districts surrounding the city of Pittsburgh. Optical fiber will link all 200 school buildings of all 42 districts and, through the Allegheny Center Mall (ACM) PoP, connect back to 3ROX, providing access to state-of-the-art educational resources.

FiberTech Networks operates a Metropolitan Wide Area Network in Pittsburgh that uses fiber optic-based transport services. Available to local businesses, the Allegheny Center Mall facility will enable customers to meet their current and future connectivity requirements with unlimited scalability.

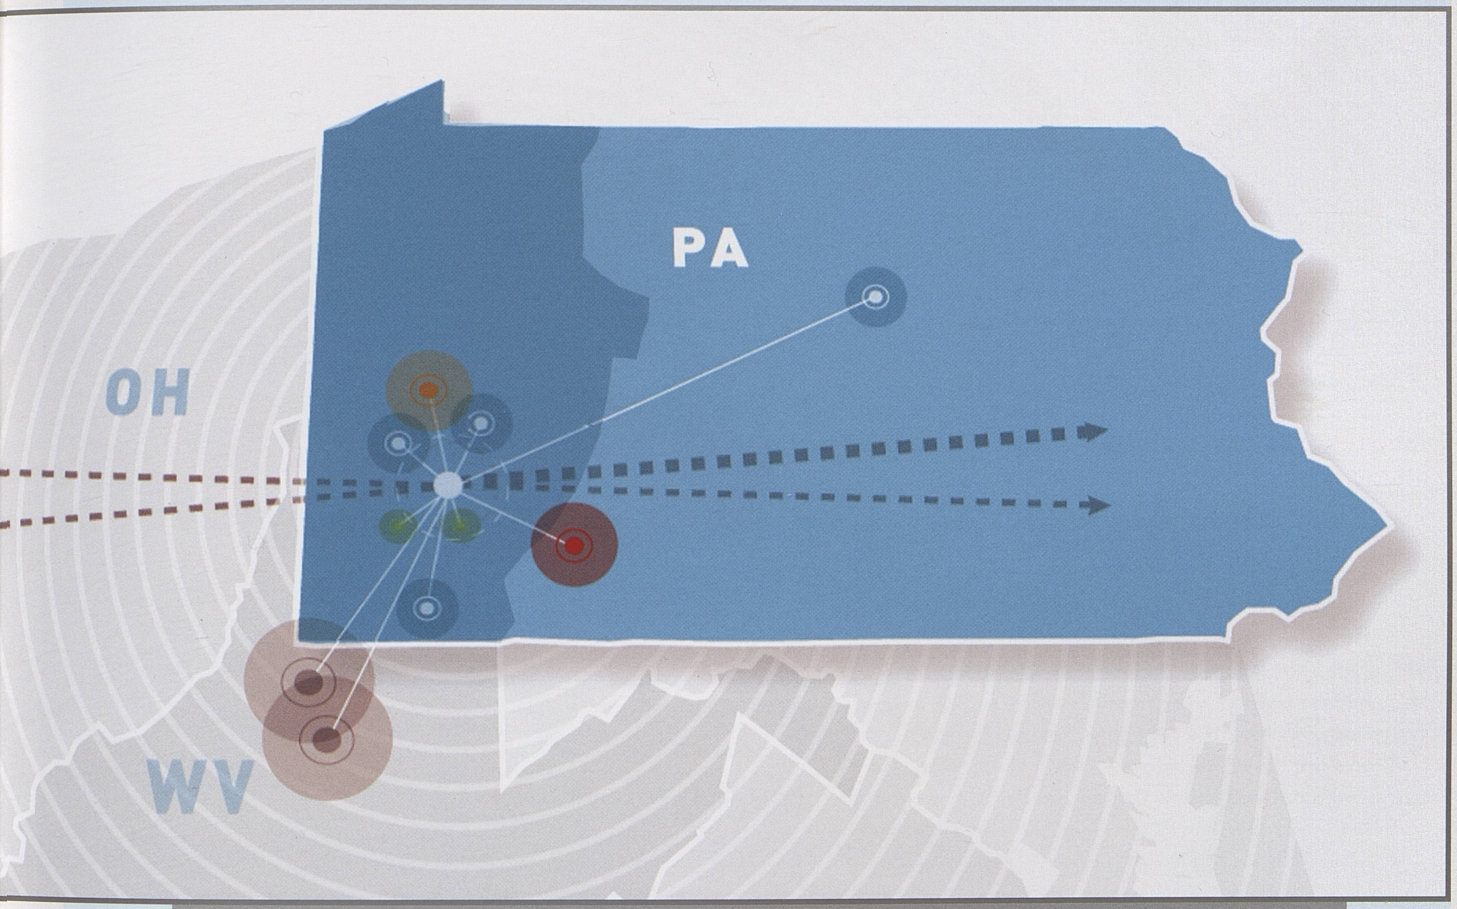

3ROX MEMBERS

Universities
Carnegie Mellon University, Pennsylvania State University, University of Pittsburgh, Waynesburg University, West Virginia University

NLR Member Institutions
PSC, University of Pittsburgh, Pennsylvania State University, Indiana University

K-12 Institutions
Allegheny Intermediate Unit (AIU3), Arin Intermediate Unit (IU28), Beaver Valley Intermediate Unit (IU27), Intermediate Unit One, Northwest Tri-County Intermediate Unit (IU5), Riverview Intermediate Unit (IU6), City of Pittsburgh School District (IU2), Woodland Hills School District

Business
Comcast, Westinghouse Electric Co.

Government Laboratory
The National Energy Technology Laboratory

Other
Computer Emergency Response Team

NETWORK CONNECTIONS

National Research Networks
Internet2 — 1 Gbps, ESnet — 1 Gbps, National LambdaRail PacketNet — 10 Gbps, TeraGrid Extensible Backplane Network — 30 Gbps

Other Network Connections
Southern Crossroads (SOX) — 1Gbps, TransitRail — 1Gbps

National Commodity Internet Networks
Global Crossing — 1 Gbps; Sprint — 1 Gbps

Pittsburgh Local Exchange Networks
Comcast

Note: Gbps: a billion (Giga) bits per second

Networking the Future (continued)

← **Matt Mathis,** PSC senior network engineering specialist. A paper he co-authored in 1997 won the ACM SIGCOMM Test of Time Award.

ADVANCED NETWORK RESEARCH: PSC NETWORK ENGINEER RECOGNIZED

Three organizations this year recognized PSC senior network engineering specialist, Matt Mathis, for his network research. Mathis has been a network engineer at PSC since 1988, helping to lead such projects as NPAD and Web100 and the related Net100.

In August, the Special Interest Group on Data Communication (SIGCOMM), of the Association for Computing Machinery (ACM), notified Mathis that he will receive their Test of Time Award. ACM SIGCOMM is the leading forum for professional discussion of data communications and computer networks. The award recognizes Mathis's 1997 paper, co-authored with former PSC staff members Jamshid Mahdavi and Jeff Semke and with Teunis Ott (then at Bellcore, now at the New Jersey Institute of Technology), "The macroscopic behavior of the TCP congestion avoidance algorithm," that was a foundation for Internet traffic control standards.

The paper was published in the ACM journal *Computer Communication Review*. The Test of Time Award honors papers from 10 to 12 years ago in *CCR* deemed by a committee to be "an outstanding paper whose contents are still a vibrant and useful contribution today."

In October, Cisco System's Collaborative Research Initiative presented an unrestricted gift of $65,500 to Carnegie Mellon University to support Mathis's research proposal titled "Rebalancing Internet Congestion Control." The project will explore the possibility of changing how the Internet manages traffic. "Our goal," says Mathis, "is to shift responsibility for allocating network capacity from the end-systems to the network itself, such that the network can support the safe operation of diverse control algorithms."

The Cisco-funded research is related to the earlier paper, explains Mathis: "The macroscopic model paper was key to establishing the 'TCP-friendly' principle that has guided Internet congestion control standards. The Cisco funding intends to move the Internet beyond this principle, to permit new standard network protocols that are not bound to politely share the network with legacy TCP."

In September, BBN Technologies funded a PSC proposal from Mathis for work on prototyping of the NSF-sponsored GENI suite of network research infrastructure. This work will proceed using a "spiral development" approach in which simultaneous trials will give rapid feedback to guide evolving designs. Multiple competing approaches are funded, of which Mathis's work is one.

↑ **PSC's directors and managers (l to r):** Katherine Vargo, manager of scientific computing systems; Cheryl Begandy, outreach manager; David Kapcin, director of financial affairs; Elvira Prologo, manager of administration; Joel Stiles, director of NRBSC; Bob Stock, PSC associate director; John Kochmar, manager of high-performance computing facilities; Richard Raymond, manager of user support; J. Ray Scott, director of systems & operations; Wendy Huntoon, director of networking; Sergiu Sanielevici, director of scientific applications & user support. **Not pictured:** Nick Nystrom, director of strategic applications; Laura McGinnis, manager of data & information resource services; Janet Brown, manager of networking.

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

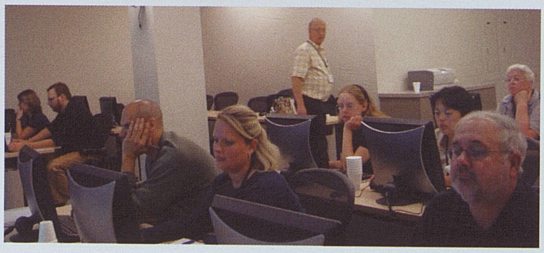

PITTSBURGH SUPERCOMPUTING CENTER WORKSHOPS (2007-2008)

Hybrid Programming for Shared-Memory and Clustered SMP Systems

Summer Institute in Bioinformatics (for minority-serving institutions)

Computational Methods for Spatially-Realistic Microphysiological Simulations

Parallel Computing

PROJECTS IN SCIENTIFIC COMPUTING 2008

PROJECTS 2008 CONTENTS

CORPORATE RESEARCH

A BETTER STROKE BUSTER & SMART SHADES

Accelerating time-to-market in new-product research and development

John Kalafut, Medrad, Inc.
Jun Deng & Michael Makowski, PPG Industries

QUANTUM PHYSICS

ATTRACTIVE REPULSION

Computation and visualization of Casimir forces in arbitrary geometries

John Joannopoulos, Steven Johnson & Alex Rodriquez, Massachusetts Institute of Technology

EVOLUTION & STRUCTURE OF THE UNIVERSE

Cooking Stars at Mach 6

The statistics of supersonic isothermal turbulence

Alexei Kritsuk, Mike Norman, Paolo Padoan & Rick Wagner, University of California, San Diego

IN PROGRESS

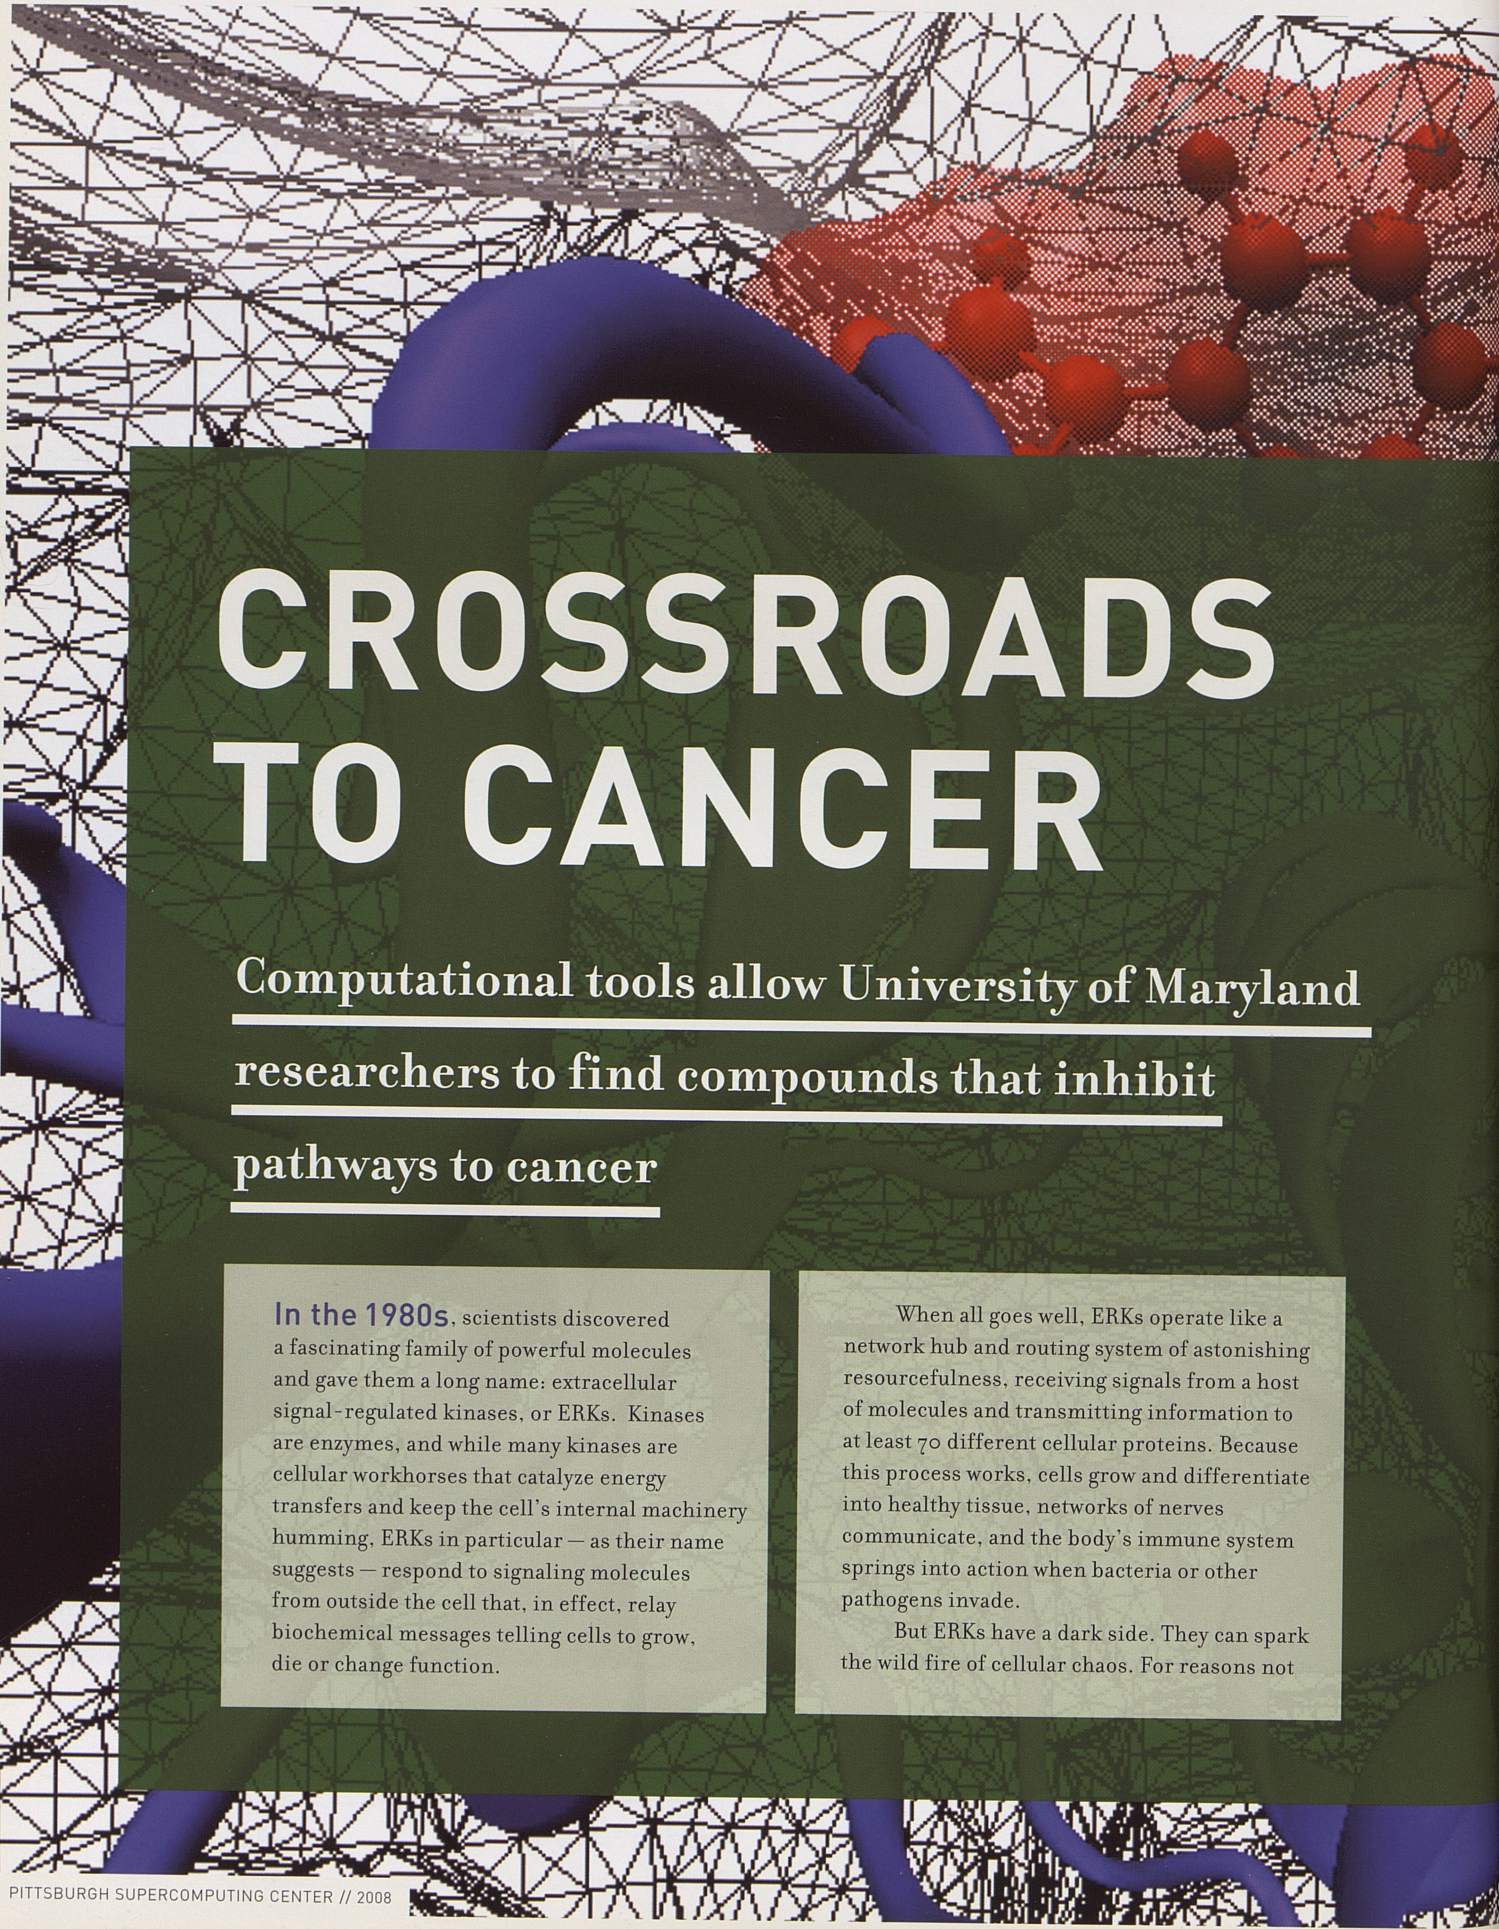

CROSSROADS TO CANCER

Computational tools allow University of Maryland researchers to find compounds that inhibit pathways to cancer

In the 1980s, scientists discovered a fascinating family of powerful molecules and gave them a long name: extracellular signal-regulated kinases, or ERKs. Kinases are enzymes, and while many kinases are cellular workhorses that catalyze energy transfers and keep the cell's internal machinery humming, ERKs in particular — as their name suggests — respond to signaling molecules from outside the cell that, in effect, relay biochemical messages telling cells to grow, die or change function.

When all goes well, ERKs operate like a network hub and routing system of astonishing resourcefulness, receiving signals from a host of molecules and transmitting information to at least 70 different cellular proteins. Because this process works, cells grow and differentiate into healthy tissue, networks of nerves communicate, and the body's immune system springs into action when bacteria or other pathogens invade.

But ERKs have a dark side. They can spark the wild fire of cellular chaos. For reasons not

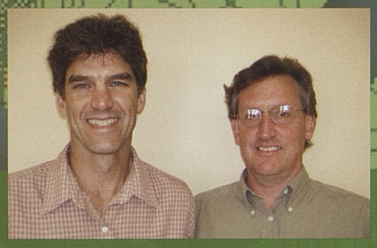

Alexander MacKerell, *right*

Paul Shapiro, *left*

fully understood — but often due to genetic mutations in the proteins that signal ERKs, some ERKs become over-energized and, in this state, activate biochemical pathways that trigger uncontrolled cell proliferation — disease states we experience as various forms of cancer. As scientists have learned more about these processes, ERKs have become immensely interesting to researchers; they are one of the most promising targets available for anti-cancer drugs.

The trick, notes University of Maryland, Baltimore (UMB) biophysicist Alex MacKerell, is to design "smart drugs" — drugs that can block the over-energized ERKs while allowing normal ERKs to go about their essential cellular business. MacKerell, who directs the Computer-Aided Drug Design Center of the University of Maryland School of Pharmacy, and his UMB biochemist colleague Paul Shapiro have attacked this challenge with a formidable one-two punch of computation and laboratory experiments.

Over the last four years, with multi-year support from the National Institutes of Health and tens of thousands of hours worth of computing on several PSC systems, starting with LeMieux and evolving to BigBen, MacKerell has applied powerful molecular-search techniques. From millions of possible candidates, he has identified about 100 drugs that have potential to stop only the bad ERKs.

In laboratory studies, Shapiro has taken the next step. Working from MacKerell's list of candidates, he found several compounds with the ability to stop malignant growth in a variety of human cancer cells, including breast and lung cancer. "We're very optimistic," says MacKerell. "We have patents pending on some of these compounds, and people in Paul's lab are testing more of them. Companies want to talk to us about licensing."

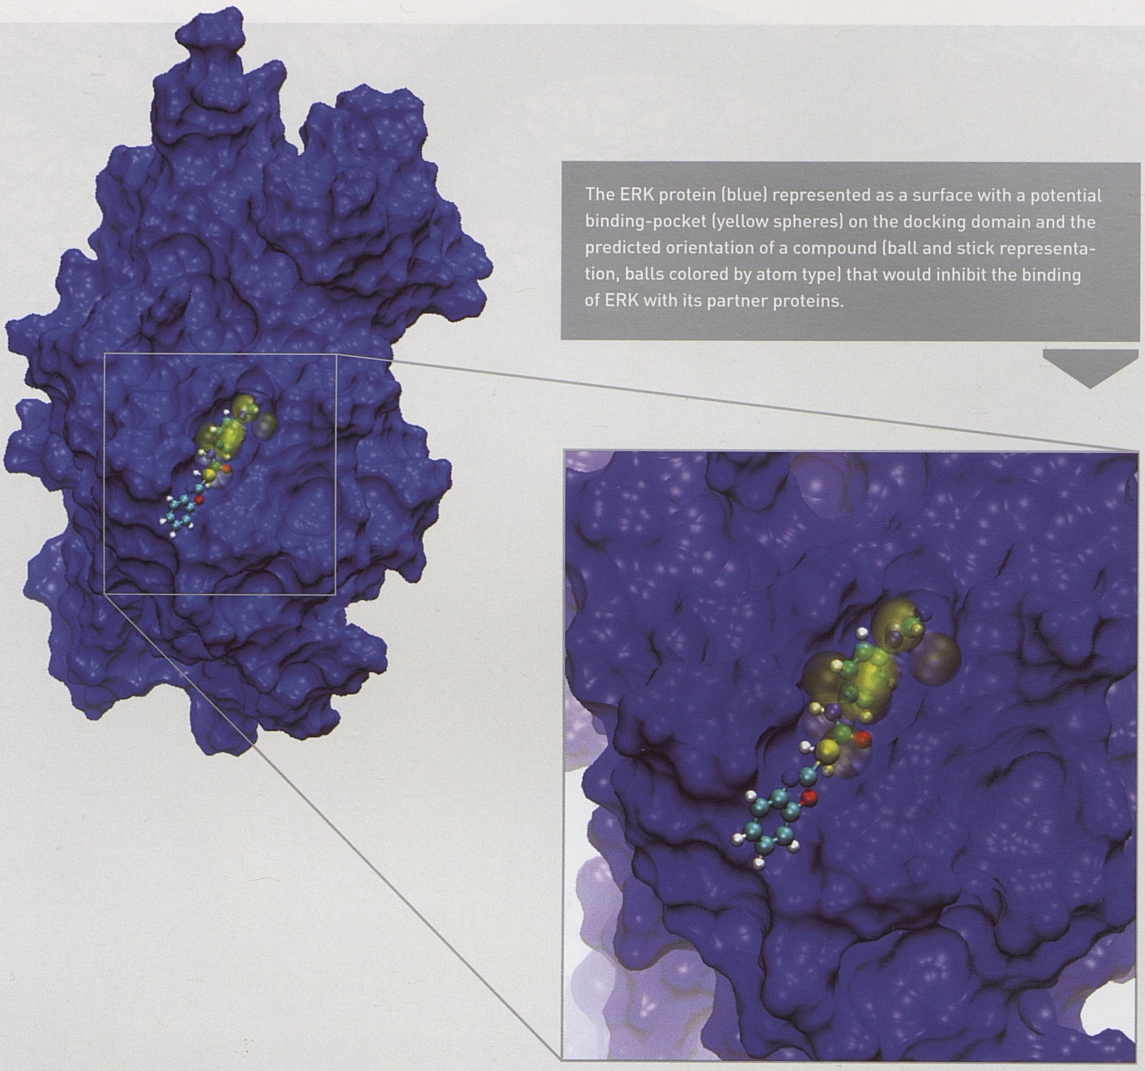

The ERK protein (blue) represented as a surface with a potential binding-pocket (yellow spheres) on the docking domain and the predicted orientation of a compound (ball and stick representation, balls colored by atom type) that would inhibit the binding of ERK with its partner proteins.

THE RIGHT HAND-HOLD

In recent years, researchers around the world have confirmed ERKs' role in the generation of cancer. In 2001, for example, scientists working in Chicago, New York and Israel found that one of the most used cancer-fighting drugs, Taxol, has its effect through interaction with ERKs. Such intriguing finds help to suggest the prospect of drugs that selectively target particular sites within the complexly coiled structure of ERKs.

You can think of a protein as a big machine with many functions, says MacKerell, and some of them may not be good. "You can break the machine in many places, but that might mean you'd lose some of the good things the machine does. So our job is to be extremely specific about how we break the machine, and to do that you have to understand exactly how the machine is built, its precise shape."

ERKs, notes MacKerell, can be thought of as a crossroad protein. They are activated by a range of different signaling molecules and in turn can communicate with (via a reaction called "phosphorylation") up to 70 different proteins.

"We want to block just a subset of those 70 proteins — so that we block one pathway and not others. This work falls into the realm of 'chemical biology,' an important area right now in biomedical research, and it helps us to understand in detail how ERKs function."

Across the atomic landscape of ERKs are specific places — "docking domains" — where other molecules can link up. For ERKs, a partner protein first docks and then a chemical change occurs at the ERK's "active site." The method to MacKerell's and Shapiro's madness is that experiments show different parts of the docking domain link with different proteins.

"It's like if you jump into a moving truck," says MacKerell. "You've got to latch onto it, and then you can pull yourself in. Basically, 70 different proteins can latch onto ERKs, and one part of an ERK is involved in the latching on, and another part handles the phosphorylation. And that hand-hold where proteins latch on involves different parts of the ERK docking domain for different proteins."

Each of the hand-holds — better known as binding pockets — represents the start of a distinctive

biochemical pathway that controls a specific task in cellular life. Finding these pockets is a first step. The next is to find a compound that can fill the right pocket and block a cancer-causing pathway. Ultimately, such compounds have the potential to become new drugs that can treat certain types of cancer.

THE PREDICTION BUSINESS

The trick is to find molecules that are the right size, shape and electronic charge at the right places to latch on and fill the pockets. This quest puts MacKerell squarely into the prediction business, and it would be impossible without computational tools. With "molecular dynamics" — simulations that track the atom-by-atom position of a molecule as it changes its shape over time, MacKerell identifies likely binding pockets.

The next step is to use this computer-drawn blueprint to search through a database of small organic compounds to find candidates that are likely fits to a binding-pocket target. To match a single molecule with a single binding-pocket within an ERK's docking domain requires screening nearly a million compounds one-by-one. When a promising compound is found, it takes up to 20 "docking runs" — computational screens that test to see how snug is the fit, how well the compound matches structurally and biochemically with the pocket. Each of these runs can require up to 100-billion calculations.

From 100 compounds, 10 showed promise as molecules capable of turning off taps of cancer at their source.

Without massively parallel systems — such as PSC's BigBen — that make it possible to employ hundreds of processors simultaneously, this work would not be feasible. "What PSC does," says MacKerell, "by allowing us access to so many processors, is make it possible for us to screen through our database multiple times in a very short period of time." Searches that otherwise would take months can be done in as little as a couple of hours.

When the data storm of the first round of screening settled, MacKerell had found about 100 compounds worth focusing on — his best picks for Shapiro to take into his laboratory for *in vivo* biological testing with cancer cells. "Alex gave me a list of 100 or so compounds," says Shapiro, "and I bought a small amount of each of them. In my lab, we have several cancer (cell) lines that we use for our preliminary studies and then the most biologically active compounds are tested in animal models. Here's where we evaluated whether the compounds Alex predicted could actually be useful by biological standards."

From the original round of 100 compounds, Shapiro found 10 that showed significant promise as ERK inhibitors, molecules capable of turning off taps of cancer at their source. "It's exciting," says Shapiro "that some of these compounds are showing they have an effect on stopping cancer cell proliferation."

Shapiro counts the computer-based methodology as highly "cost and person-power effective." His findings to date led two years ago to renewal of NIH support for his work with MacKerell's lab. "Computational modeling has identified a number of compounds with biological activity," says Shapiro. "Some of them are now under patent protection." (FS)

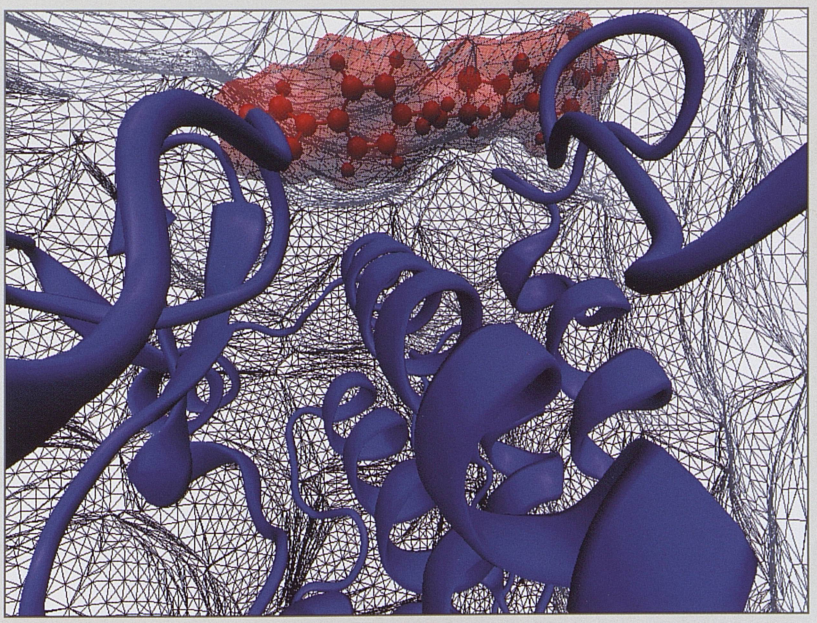

View from the interior of the ERK protein with a bound inhibitor (red), where the ribbons represent the protein backbone and the mesh represents the protein surface.

MORE INFORMATION

http://www.psc.edu/science/2008/erks.html

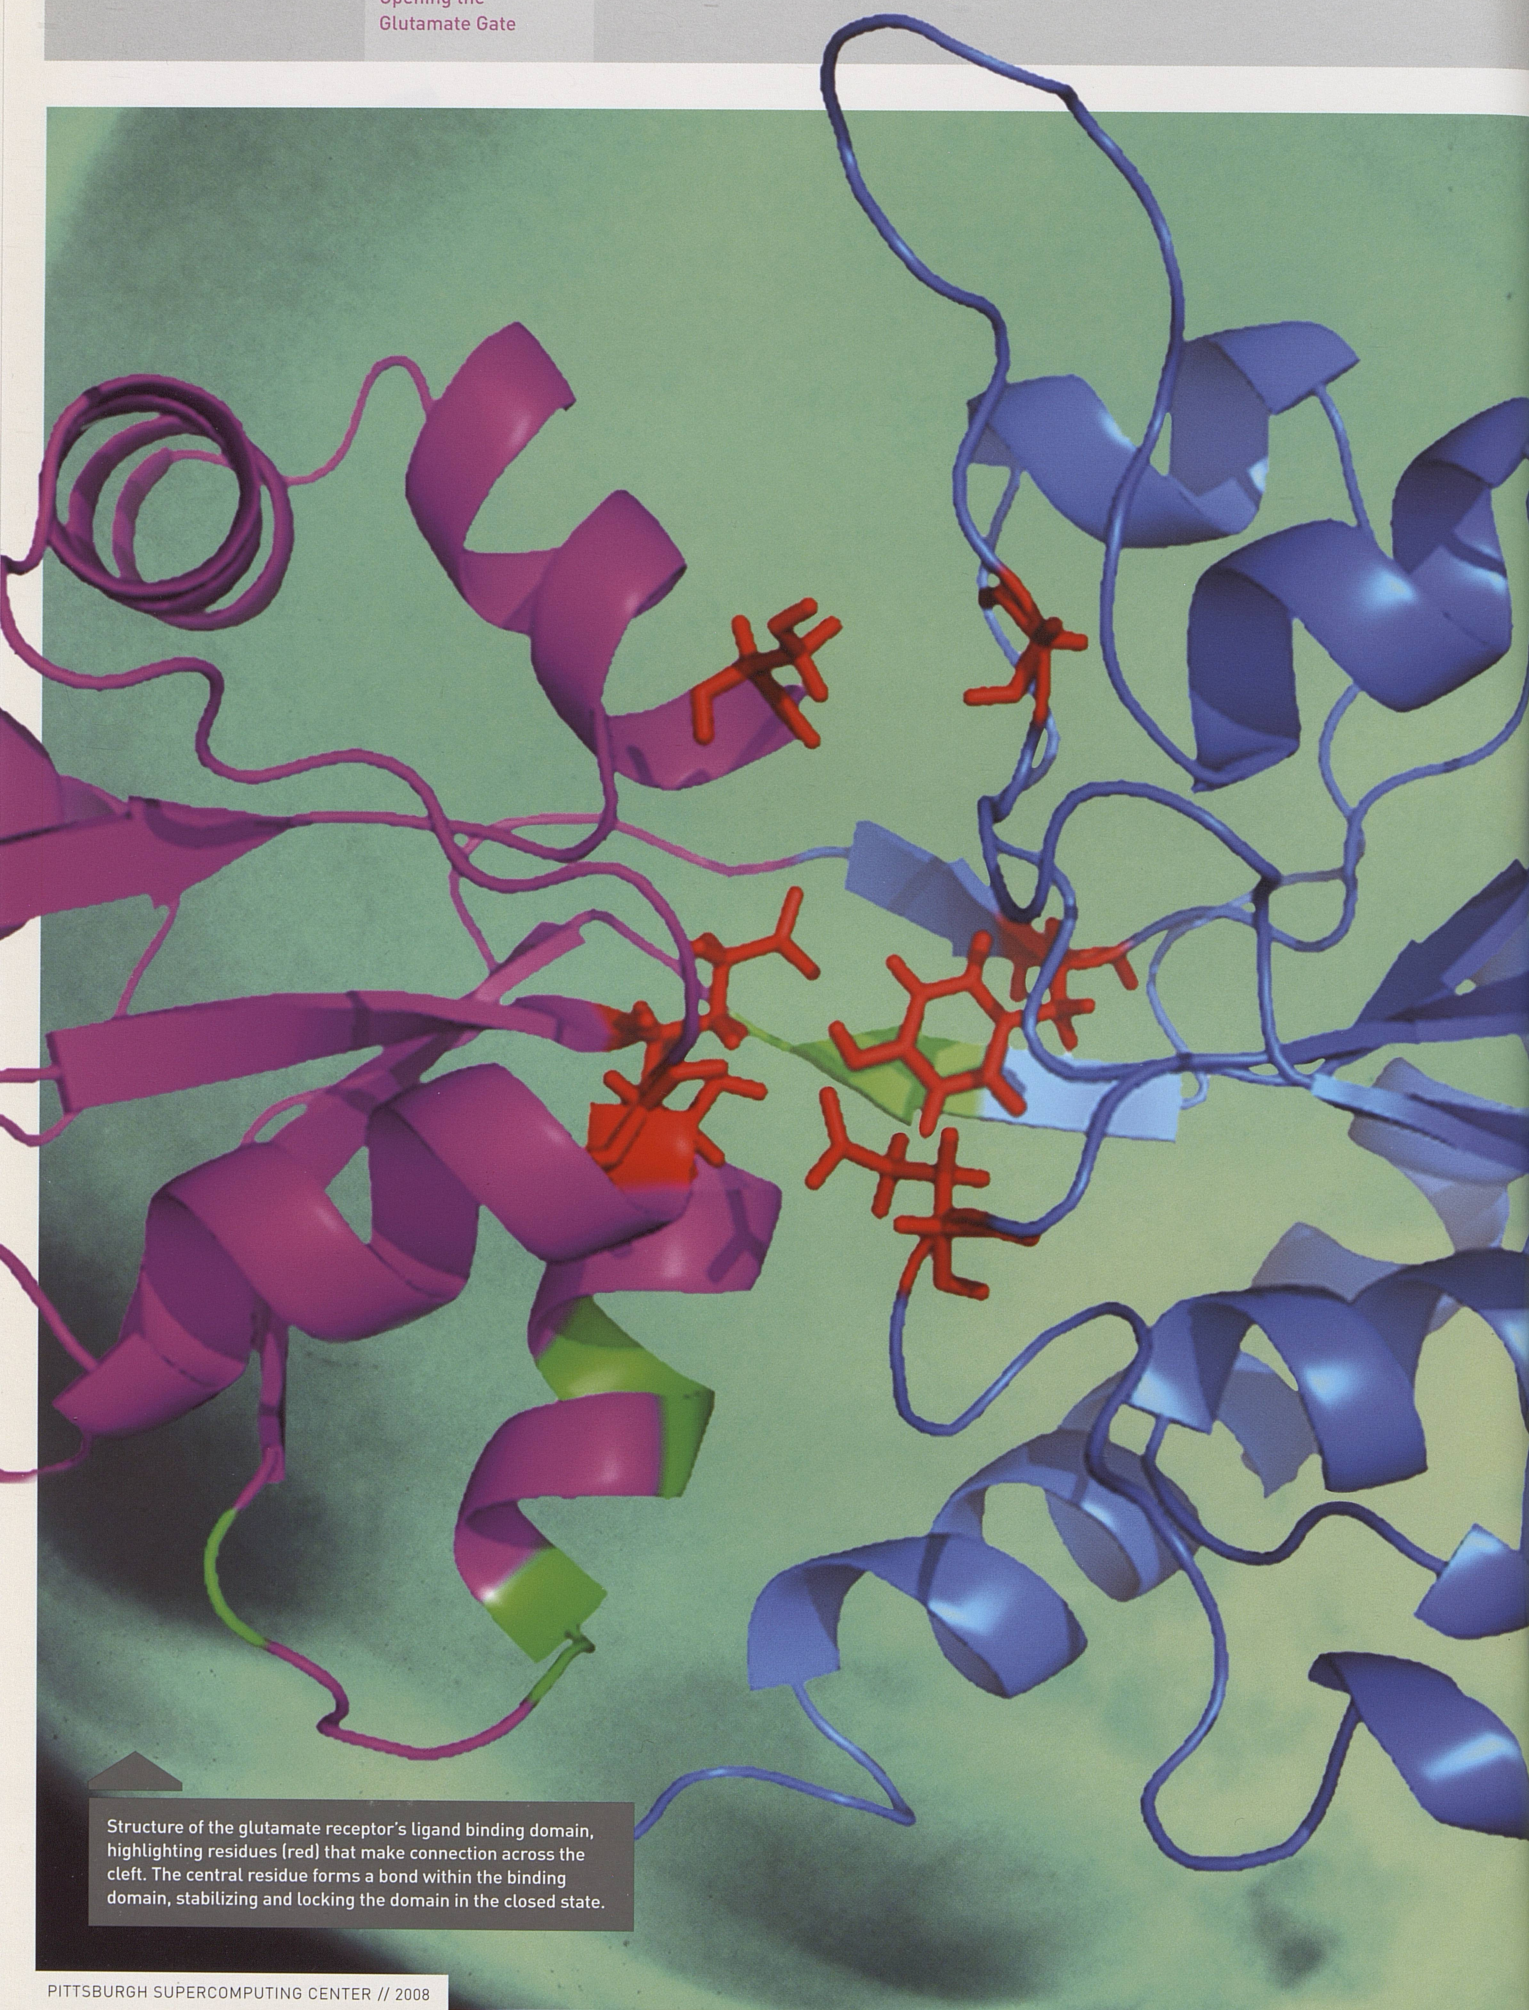

Structure of the glutamate receptor's ligand binding domain, highlighting residues (red) that make connection across the cleft. The central residue forms a bond within the binding domain, stabilizing and locking the domain in the closed state.

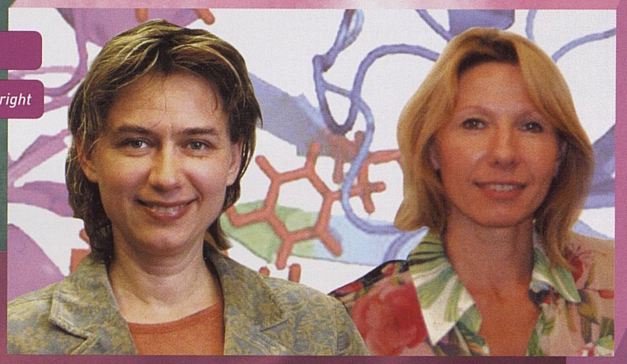

Maria Kurnikova, *left*

Tatyana Mamonova, *right*

OPENING THE GLUTAMATE GATE

Simulations of the most prevalent receptor in the brain provide new understanding of how it opens its "gate" to fire neurons

In the time it takes you to read this sentence, more than a million neurons will transmit electrical impulses in your nervous system — triggered by touching this page, or using your eyes to make sense of the letters and words. These impulses — set off by touch, sight, sound and other stimuli — are relayed neuron-to-neuron through the brain by tiny messenger molecules called *neurotransmitters*.

You could think of a neurotransmitter as a lightning-fast pony express rider leaving St. Joseph, Missouri to get to Elwood, Kansas, the next stop, in microseconds. The distance between the two towns is like the synapse, the gap between a sending and receiving neuron. When the rider gets to Elwood, he hands his mail pouch to the next rider, whom you could think of as a receptor, a specialized protein on the membrane of the message-receiving neuron. As a neurotransmitter docks at a receptor, the neuron becomes activated, signaling it to release neurotransmitters for relay to the next neuron. But the interaction between the neurotransmitter and receptor must be just right, or the message won't be relayed, the next rider won't leave Elwood to journey across the next synapse.

It's a crude, imperfect analogy for an extremely complicated process. Though incompletely understood, the complex interactions between neurotransmitters and receptors are the basis for the therapeutic action of many drugs. Prozac and its derivatives, for instance, treat clinical depression by binding to receptors for the neurotransmitter serotonin, increasing the amount of serotonin available in synapses.

Computational chemist Maria Kurnikova of Carnegie Mellon University uses supercomputing to study membrane proteins such as receptors and how they interact with neurotransmitters. In extensive work over the past three years, she and post-doctoral fellow Tatyana Mamonova used PSC's LeMieux and BigBen to gain new understanding of the receptor for an important neurotransmitter called glutamate.

"What we found is exciting," says Kurnikova, "because it's something nobody expected." Their findings go beyond laboratory experiments and identify the precise atom-by-atom changes that occur as glutamate binds to the glutamate receptor, a process that changes the receptor's shape and eventually leads to activation of the neuron — information that presents new possibilities for precisely designed therapeutic drugs.

THE GLUTAMATE RECEPTOR'S FLYTRAP

Although glutamate may not be as well known to most people as serotonin and some other neurotransmitters, it's the most prevalent neurotransmitter in the central nervous system and an important focus of medical and pharmaceutical research. The interaction between glutamate and glutamate receptors plays an essential role in memory and learning, and dysfunction of this interaction is related to a diverse list of central nervous system disorders, including Alzheimer's disease, epilepsy, schizophrenia and depression.

A key to better understanding of all these dysfunctions is precise knowledge of how the glutamate receptor works. While there are several different glutamate receptors, Kurnikova and Mamonova looked at one (known as AMPA) that is the most common and initiates the series of steps that lead to neural activation. AMPA is a four-part protein (a tetramer) that forms a pore-like "ion channel" in the neural membrane — like a valve that opens selectively to allow ions, biochemical electricity, to flow through the channel into the cell.

The valve-apparatus of the receptor — its "gating mechanism" — that causes it to shift from closed to open was in large part a mystery when Kurnikova and Mamonova began their work. What was known is that each of the four parts of the receptor has its own "binding domain" — a cleft-like structure that extends outside the membrane — where glutamate can attach. When glutamate binds, this cleft in the binding domain closes —like a Venus flytrap catching a fly — which in turn causes the transmembrane channel to open for ion flow.

This image shows electrostatic charge, positive (blue) and negative (red), of the protein and ligand, glutamate (right), illustrating electrostatic complementarity between the ligand and protein.

"Cleft closure controls opening of the pore," says Kurnikova. "And the degree of cleft closure is believed to relate to how much the pore opens. Experiments show that this seems to be correct. But exactly how it happens nobody knows. It would be really nice to know on the molecular level what controls the degree of cleft closure, but to separate these processes into two — docked glutamate first and then the closed binding domain — is experimentally impossible."

REVEALING NEW DETAILS

Kurnikova and Mamonova tackled this problem by, in effect, building an atom-by-atom model of glutamate and the glutamate receptor binding-domain inside the computer. They used their model, comprising 20,000 atoms, with PSC systems — initially LeMieux, PSC's terascale system, and more recently BigBen — to simulate the shift from the unbound to bound state of the receptor.

Using a method called "molecular dynamics," they tracked the forces among all the atoms as the receptor atoms shift position. They relied on software called AMBER (Assisted Model Building with Energy Refinement) and a special technique called "umbrella sampling."

In terms of the number of atoms, their problem was relatively modest. Their major challenge was to simulate the entire duration of the transition from unbound to bound state. With umbrella sampling, they divided the transition into 12 windows, each of which they simulated for 10 nanoseconds, a relatively long period in terms of biomolecular events. The simulations sliced time into femtoseconds, a million freeze-frame snapshots for each nanosecond. At each snapshot interval the simulation recalculated the forces and positions of every atom.

"Forcing a protein from one conformation into another is a challenging simulation," says Kurnikova. "Umbrella sampling is a technique where you bias the system to move slowly through its conformational space so you can sample the entire space. With mathematical techniques, you eliminate the biasing effects and calculate the free energy difference."

With their first series of simulations in 2004, the researchers compared results from their modeling to experimental data, gathered from infrared spectroscopy, measuring shifts in the vibrational frequencies of glutamate as it binds to the receptor. Their calculations closely reproduced the experimental results, validating their approach.

These simulations also provided insight into the positive/negative electronic polarities between the unbound receptor binding-domain — the open cleft

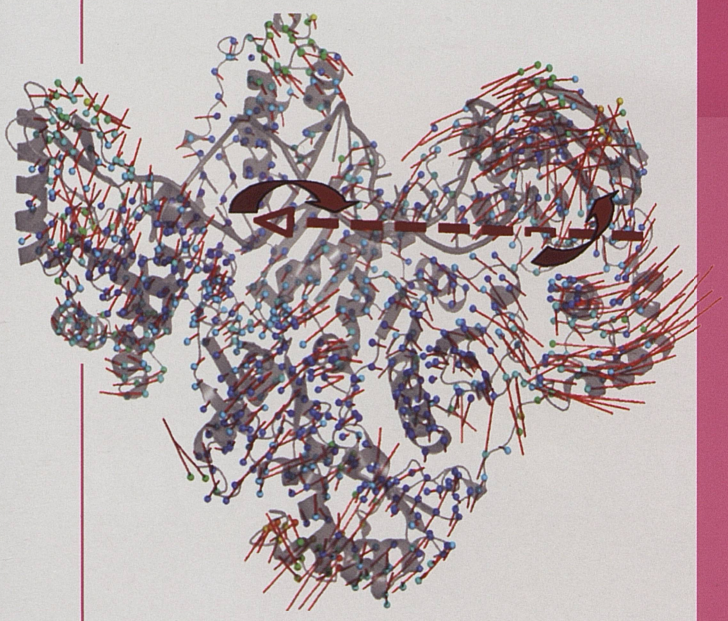

↑ This image from the simulation illustrates the axis along which RT seems to twist, a motion, says Madrid, that might allow the DNA to move across the protein and facilitate its replication.

Marcela Madrid

TWIST AND SLIDE: NEW CLUES TO KNOCKING OUT AIDS

In another project, PSC scientist Marcela Madrid collaborated with Maria Kurnikova and colleagues in a series of large-scale simulations of HIV-1 reverse transcriptase (RT). This multi-functional protein plays a critical role in the life-cycle of HIV, the virus that causes AIDS. It replicates HIV's DNA, essentially a copy-and-paste function, which is then incorporated into immune-system cells of the infected person. Because of this critical role, RT is the target of several FDA-approved anti-AIDS drugs.

Using PSC's LeMieux and the Cray XT3, Madrid and the Carnegie Mellon scientists simulated RT with and without DNA, simulations that involved 123,000 atoms. Notably, they extended their molecular dynamics simulation for 40 nanoseconds of biological time, much longer than any previous similar work. Because of this extended time, the simulations reveal motions and detailed interactions of DNA bound to RT not before observed. In particular, they show that the DNA undergoes a "twist and slide" motion, which may facilitate its positioning at RT's active site. Interfering with this motion could disrupt RT's function. "This work is important in understanding RT's function," says Madrid, "because it shows details of the motion that have not been observed before by any other computational technique."

of the Venus flytrap — and glutamate. They showed that the open cleft projects a positive electronic charge that strongly attracts the negatively charged molecular face of glutamate.

With the 2005 arrival of BigBen at PSC, the researchers were able to step up the pace of their work. "When we started," says Kurnikova, "running a simulation of one nanosecond would take several days. With BigBen, we were suddenly able to run 10 nanoseconds overnight."

The improved technology allowed them to do extensive simulations that tracked variant (mutant) glutamate receptor structures and the effect of the mutations on the transition from open to closed state of the glutamate binding domain. Ultimately, this allowed Kurnikova and Mamonova to identify precisely which molecular interactions controlled the binding-domain transition from open to closed. They identified several key interactions, but — most interesting — they found a bond that forms within the binding-domain itself once the cleft closes that, more than other interactions, stabilizes and locks the binding domain in the closed state.

"With BigBen, we were suddenly able to run 10 nanoseconds overnight."

"Nobody expected this from experimental analysis," says Kurnikova, "that there is this one interaction which is the most important and controls that whole transition." For drug design, this detailed knowledge is very important, she adds, and can guide pharmaceutical companies toward finding compounds with very specific effects — to either inhibit or enhance the gating mechanism of the glutamate receptor.

"That's the advantage of atomistic simulations," she says. "We can look at that experimentally and observe some effect and guess why it works, but with this modeling we can see what causes what we see experimentally — not just what happens, but how. We can understand that in great detail." (SP)

MORE INFORMATION

http://www.psc.edu/science/2008/glutamate.html

SHAKE, RATTLE AND ROLL

Quake Group simulations show areas of amplified ground-motion in potential large-impact Southern California quake

The earth shook on the day Jacobo Bielak was born and — whether you think of it as serendipity or random chance — earthquakes became his career. "It was a small one," says Bielak. In his home town of Mexico City, quakes are common due to the top layer of soft soil. "It's like a bowl of jello. The waves amplify. You develop an inner ear that tells you, 'Oh, this is a minor one,' and then once in awhile you say, 'Oh-oh, this one's for real."

The "for real" ones have occupied Bielak's workdays for many years. A professor of civil and environmental engineering at Carnegie Mellon University, Bielak and his Carnegie Mellon colleague, computer scientist David O'Hallaron, lead the Quake Group, one of the world's leading efforts at developing computational strategies to realistically simulate the soil vibrations that occur during earthquakes. Early this year, Bielak won a four-year, $1.6-million grant from the National Science Foundation to support work on simulations with colleagues from the University of California that will help predict seismic risks affecting the Los Angeles basin and other earthquake-prone urban areas.

Working in close collaboration with the Southern California Earthquake Center (SCEC), an inter-disciplinary community of hundreds of scientists worldwide, the Quake Group's guiding objective is "hazard analysis" — to understand the probability region-by-region within an earthquake basin that a certain level of ground-motion will occur. "The severity of shaking," says Bielak, "can vary significantly within relatively small areas — depending on geological characteristics. Our models can predict the ground motion of a prescribed quake, and engineers need this information to define building codes that provide for the safest possible structures at reasonable cost."

Jacobo Bielak, *left*

David O'Hallaron, *right*

Ricardo Taborda

Leonardo Ramírez-Guzmán

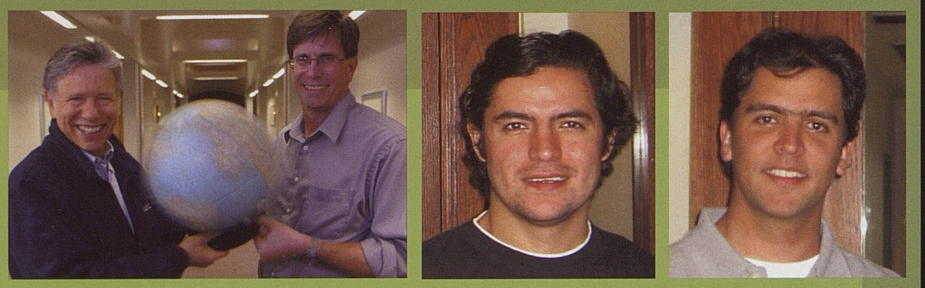

In 2006, the Quake team — in a project led by former grad student Tiankai Tu — won the HPC Analytics Challenge Award at SC06 in Tampa for "Hercules," their software that coordinates all the stages of large-scale earthquake simulation, from problem definition to final visualization. Using PSC-developed software called PDIO, Hercules can visualize results in real time as the simulation is running. With this unified framework, all tasks — building a software mesh that subdivides the quake volume, partitioning the job among hundreds or thousands of processors, the simulation itself, and visualizing results — are performed in place on the computing platform.

In recent Quake Group work, Ph.D. students Leonardo Ramirez and Ricardo Taborda (with consulting from PSC scientist John Urbanic and participation of former Ph.D. student Julio Lopez), ran Hercules on BigBen, PSC's Cray XT3, to simulate a major Southern California earthquake scenario called ShakeOut. Their simulation of this potential magnitude 7.8 quake shows how the largest ground motions occur within sedimentary valleys, where waves are trapped within the basin and amplified by the soft soil deposits.

Hercules simulations of ShakeOut, furthermore, have been part of an important cross-validation study with two other SCEC simulation groups. After three years, the results of this work — which rely partly on data-analysis routines developed by PSC scientist Joel Welling — show that the three schemes are consistent and accurate enough to rely upon for future work. "We've come a long way from where we started," says Bielak, "and this is an important point from which we can go forward with confidence in the validity of our simulations."

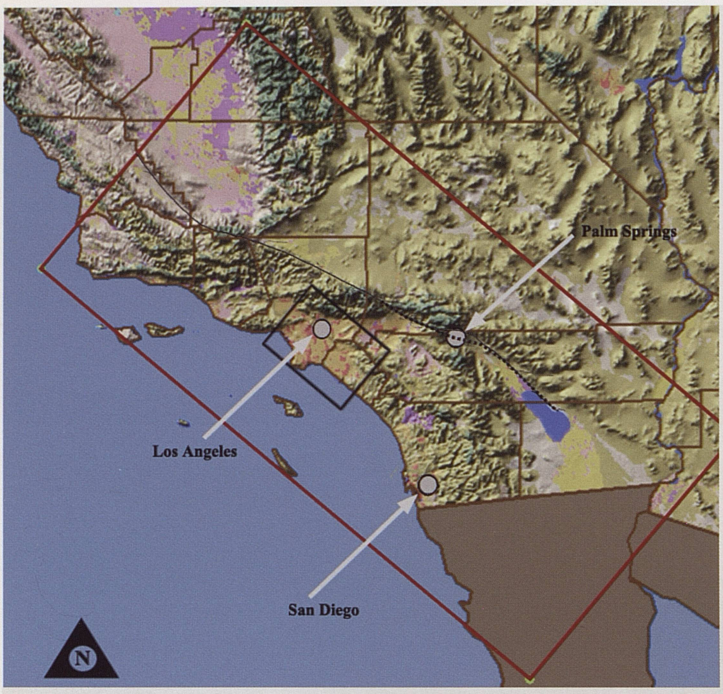

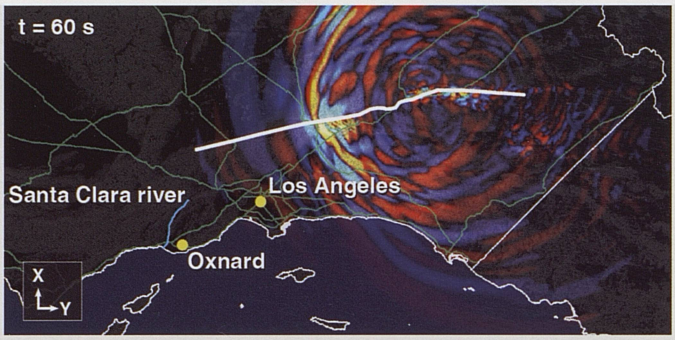

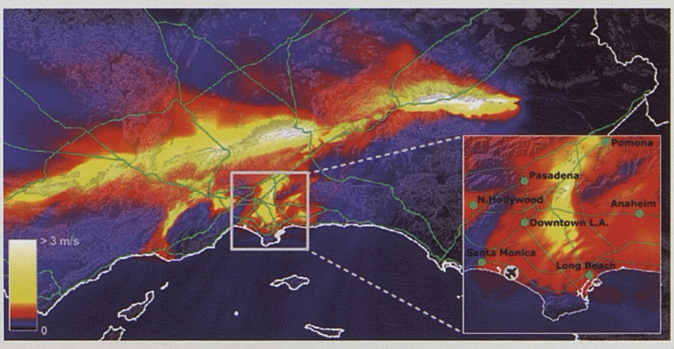

The red box shows the area of the ShakeOut quake scenario, with the dashed line representing the portion of the San Andreas fault where slip occurs. A snapshot from the simulation (top right) shows distribution of the horizontal ground velocity through the epicentral region 60 seconds after onset of the quake. The final graphic (bottom right) shows the distribution of maximum ground-velocity from the ShakeOut quake, with an inset showing the Los Angeles region. The largest motion (yellow to white) occurs within sedimentary valleys where waves are trapped and amplified.

THE GREAT SOUTHERN CALIFORNIA SHAKEOUT

The ShakeOut quake scenario — a project that involves more than 300 professionals — was defined by the U.S. Geological Survey and SCEC as a magnitude 7.8 quake set off by a rupture along more than 250 kilometers of the San Andreas fault. Scientists believe that such a quake — large enough to cause strong shaking over much of southern California — is inevitable. Estimates are that it will cause 2,000 deaths, 50,000 injuries, $200-billion in damage and other losses and long-lasting disruption.

Understanding the potential of these impacts is an important step in preparing for the event. In November 2008, the ShakeOut scenario will be the centerpiece of emergency-response and public-preparedness exercises in Southern California involving three-million people.

The ShakeOut volume domain — 600 kilometers long by 300 kilometers wide by 84 kilometers deep — encompasses the most prominent fault structures in the region, and includes all major cities in the Los Angeles basin. Simulations provide details of the ground motion, including variations within cities and regions, that help to identify a range of effects from direct physical impacts to long-term consequences, and assist in response planning.

At a February 2008 meeting of the Earthquake Engineering Research Institute (EERI), the Quake Group presented results from a ShakeOut simulation — using 1,024 BigBen processors — in which they modeled ground motion up to a frequency of 0.5 vibrations per second (Hz). Unlike other Shakeout simulations, Hercules employs a "meshing" method — developed and parallelized by O'Hallaron — that tailors the size of subvolumes according to soil stiffness. The advantage of this "adaptive mesh" is that, for a given frequency of vibration, wavelengths are shorter in softer soils, and Hercules adjusts the mesh-size to finer resolution that can accurately capture these shorter wavelengths. Hercules' meshing algorithm is highly efficient, and required only 70 seconds on BigBen to build a mesh of 81.5 million elements.

The Quake Group's video of this simulation, which won EERI's first annual graphics competition, shows distribution of peak motions, both for the entire region and within smaller regions — information of interest to engineers, who want to know where the largest ground-motions occur. "We show," says Bielak, "that away from the immediate epicentral region, the largest ground-motion occurs within the sedimentary valleys, where waves are trapped and amplified by the soft soil deposits."

DATA, DATA, DATA

An important goal of the Quake Group's work with the ShakeOut scenario has been to compare results with two other SCEC groups who also did ShakeOut simulations, both of which used an approach (finite difference) that offers tradeoffs to Hercules (finite element). All three groups simulated the same scenario using TeraGrid computing resources. While the Quake team used BigBen, a group from San Diego State University computed at SDSC and a group from URS Corporation used resources at TACC.

Early on in this work, which proceeded over several years, it was clear that direct comparisons would be difficult due to the large spatial grids and many time steps involved, and because of differences in the simulation algorithms. Wave fronts from the same initializing data, for instance, can propagate at slightly different rates.

To help with this effort, PSC scientist Joel Welling developed a computational routine that statistically compares overall results between two different simulations. The concept is drawn from other research (Kristeková et al.) adapted by Welling to the ShakeOut data. The routine is concerned with two wave-related variables of phase and amplitude (envelope) and provides a graphical representation of mismatches between two simulation outputs. "This provides an outstanding means," says Bielak, "for comparing between different simulation techniques, and greatly facilitated our verification efforts."

Similar statistical analysis is important to challenges posed by the huge datasets that will be produced by petascale computing. For the Quake Group, future work — some of it to be done with TeraGrid resources at TACC and NICS — aims at ShakeOut simulations from 1.5 Hz up to 3.0 Hz. These will be the highest resolution quake simulations ever done, entailing datasets in the range of 500 terabytes. To analyze such data-intensive results, PSC is developing the capability — with help from a large shared-memory system expected to be operational in 2010 (see p. 4) — to transpose the data from its time-varying orientation and represent it spatially. This powerful way of seeing earthquake simulation results will allow researchers to focus on specific locales, such as where ground motion is most intensive.

PSC developed a routine to compare overall results between two different simulations.

MORE INFORMATION:

http://www.psc.edu/science/2008/quake.html

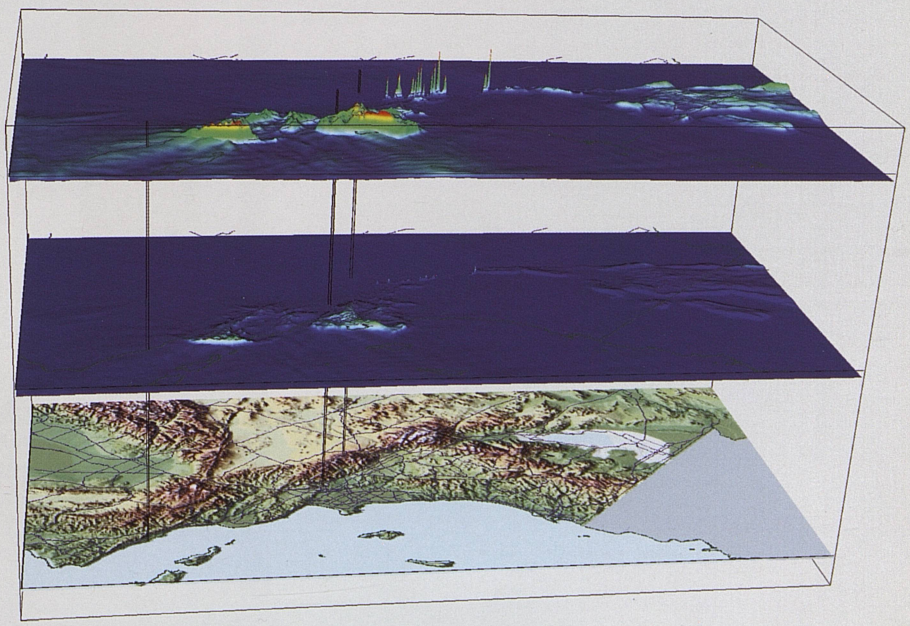

COMPARING BETWEEN SHAKEOUTS

The upper two planes of this display, developed by PSC scientist Joel Welling, show the envelope mismatch (top) and phase mismatch between the waveforms produced by two simulations. The vertical markers indicate seismogram location.

A BETTER STROKE BUSTER & SMART SHADES

PSC's corporate partnerships lead to improved high-tech sunglasses and a rapid feasibility decision for a catheter design

Through PSC's corporate affiliates program, the same computational resources and expertise that help university scientists advance knowledge in many fields are available to business. Since its beginnings in 1986, PSC has worked with a number of corporations in new product research and design. These include more efficient power turbines by Westinghouse, better mixing of molten steel in "tundishes" for continuous casting by U.S. Steel, and improvements in lightweight beverage cans and aluminum car parts by ALCOA.

Two recent PSC corporate projects — one involving PPG Industries and another involving Medrad, Inc. — were recognized by the Council on Competitiveness, a Washington, DC organization of business, labor, academic and governmental leaders who focus on private-sector competitiveness. PPG worked with PSC to improve a state-of-the-art technology for eyewear — sunglasses that automatically change from clear to dark in the presence of ultraviolet rays. Medrad, an Indianola, Pennsylvania company, collaborated with PSC and Carnegie Mellon University to study the feasibility of a novel catheter for safe, efficient removal of deep-vein blood clots.

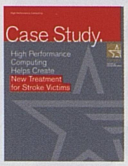

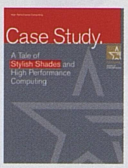

Through its High Performance Computing (HPC) Initiative, which facilitates usage of HPC across the private sector, the Council on Competitiveness recognized both projects as success stories — demonstrating how HPC drives corporate innovation and productivity. The Council's case studies about these projects were made possible by a National Science Foundation grant and are available on the Council website at **http://www.compete.org**

Both projects highlight how the availability of sophisticated HPC technologies can reduce costly trial-and-error design and accelerate time-to-market for innovative ideas.

A WINDOW OF TIME FOR FEASIBILITY

The leading cause of death in industrialized nations is blood clots — "thromboembolic disease" — a piece of blood clot (thrombus) blocking a vessel (embolism). Ruptured clots that lodge in a vessel can cause heart attack, pulmonary embolism or stroke.

For all these conditions, time is a factor. The sooner there's treatment, the better the chance of recovery without serious consequences. Recently, a patented but unmarketed, deep-vein catheter technology that could speed-up treatment of blocked vessels came to the attention of Medrad.

"The patented prototype device seemed like a good fit with Medrad's growth objectives, so we purchased the rights," says John Kalafut, principal research scientist at Medrad. "But before we could give the go-ahead to proceed with product development, we had a number of practical questions that needed answers, and we needed them quickly."

A leader in providing medical devices and services for diagnostic and therapeutic imaging of the human body, Medrad is an affiliate of Bayer Schering Pharmaceutical AG, Germany with annual revenues of around $500 million and 1,700 employees. Their products have captured 70 to 80-percent market share.

The task for Kalafut and Medrad's R&D group was to determine if the potential technology was actually feasible — will it work and, if so, what are its limitations? "Because we examine many different opportunities each year," says Kalafut, "we need to be able to quickly gauge the technology's efficacy and either start the development process, or turn it down and move on to the next. If the technology looks promising, we make an initial commitment and then, several years later, launch a multimillion-dollar project to commercialize it."

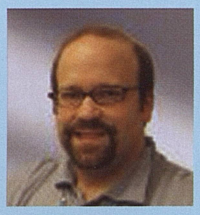

John Kalafut,
principal research scientist, Medrad, Inc.

The classic approach for a biomedical product involves building bench-top models, subjecting each one to a variety of trial conditions and then moving into animal and human testing. For a catheter, this costly, time-consuming approach had serious limitations in the ability to efficiently capture the complicated interactions between blood cells, vessel walls, the clot and the catheter.

Kalafut felt the cath eter was an opportunity to try numerical simulations. "Not only did we need to understand the physics, we also wanted to explore different design and manufacturing approaches. We felt that doing this computationally would be both more cost-efficient and faster than building lots of different physical prototypes."

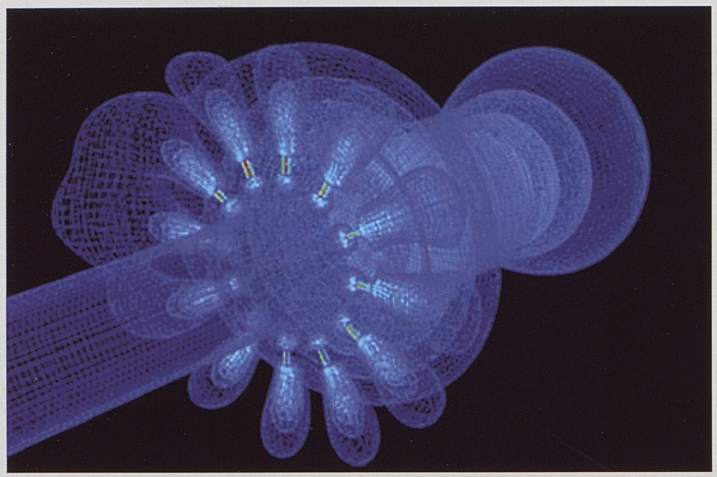

Simulated flow field from the prototype catheter as computed by 3D computational fluid dynamics software at PSC.

Supercomputing cut eight-to-ten months — a huge savings in time and money — off the feasibility decision.

The physics involved with a blocked circulatory system and the catheter includes complex fluid dynamics that can be represented mathematically — via techniques associated with a field of HPC called computational fluid dynamics (CFD). The R&D group's high-end workstations, however, lacked the horsepower for these complex simulations. They also didn't have the in-house expertise to develop or use detailed CFD codes.

Other challenges facing Kalafut's group had less to do with technology than with tradition and individual mindsets. Some of the company's engineers who had been building prototypes for decades felt that computer simulation wouldn't work — you had to build something real, something tangible, not just work with a bunch of equations. Kalafut and the R&D group convinced management that HPC was essential to capture first-to-market advantage if Medrad wanted to move forward.

Kalafut consulted with the Institute for Complex Engineered Systems at Carnegie Mellon University (CMU), and PSC supercomputing was enlisted in the R&D process. Medrad's work with PSC consultants Jun-Woo Lim and Dave O'Neal and its HP shared-memory systems, Jonas and Rachel, focused first on the physics — using CFD software to simulate the process of the catheter destroying clots, adjusting the parameters again and again to ensure that the phenomenon was repeatable. This work validated that the patent's theory was solid and that the device would do what its inventors claimed.

Medrad and PSC then did more modeling to refine the prototypes by simulating many different combinations of changes to arrive at the best design.

"Using PSC's systems," says Kalafut, "we have been able to look at multiple iterations of different design parameters without building numerous, expensive prototypes."

With the CFD studies as a foundation, a Medrad and CMU team demonstrated the physical principles of the catheter technology. Subsequent animal studies validated the operational principles and the CFD results, but clinical validation wasn't proven. Kalafut estimates that supercomputing cut eight-to-ten months off feasibility determination, a huge savings in time and money.

"If we weren't partnering with PSC and CMU," says Kalafut, "we'd probably be a year away from determining feasibility. It is critical to quickly and accurately evaluate new medical devices with numerical, bench-top and animal models. The resources at PSC allowed us to fully explore a concept months before expensive animal testing." In a competitive global market, a year can be the difference between market success or failure. And for victims of thromboembolic disease, it can be the difference between life and death.

THE QUANTUM SCIENCE OF STYLISH EYEWEAR

If you wear eyewear made with Transitions® lenses, you have the coolest shades in town — maybe even cooler than you realize. Riding on your nose is a product created with help from quantum chemistry and PSC's parallel supercomputing resources — mainly Rachel, PSC's shared-memory HP system, and some use of BigBen.

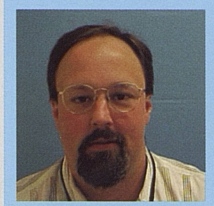

Michael Makowski, *leader of the computational chemistry research group, PPG Industries*

Jun Deng, *PPG Industries, specializes in quantum chemistry.*

Transitions lenses are made by Transitions Optical Inc., a joint venture formed in 1990 by PPG Industries of Pittsburgh and Essilor International of Paris. Based on proprietary photochromic technology, the lenses quickly change from clear to dark in the presence of

ultraviolet light and block 100-percent of harmful UVA and UVB rays. The transition is the result of photochromic dyes. When exposed to ultraviolet light, the dyes' molecular bonds break and the molecular structure changes, which in turn changes the lens color and provides UV protection. Remove the UV, and the lenses quickly return to a colorless state.

A calculation that might have taken a week on in-house machines runs in a matter of hours at PSC.

The development of successive generations of photochromic lenses has been based in large part on PPG research. Beginning several years ago, Jun Deng and Michael Makowski, scientists in PPG's computational-chemistry research group, faced a challenge in eyewear. "About five years ago, a major challenge we faced was meeting market demands for improved photochromic technologies for a high-growth segment within the ophthalmic lens market using new impact-resistant, high-index and polycarbonate materials."

One of the main differentiators among lenses that "transition" is the performance of the photochromic dyes — how fast the lenses shift from light to dark and back again, how dark the lenses can become when exposed to UV light and the color itself — brown and gray are good, hot pink or dark purple not so good. How long the product will last before the photochromic coating begins to lose its effectiveness is also key to product success. "With each successive generation," says Makowski, "you want to develop a product that has better performance, is robust to various substrates and processing, and has a lower price tag."

Essentially, it's a scientific problem. In order to advance photochromic dyes and coatings, PPG had to understand what was going on not only at the molecular level (between dyes and their matrix), but also at the quantum level (electronic structure), which dictates the dyes' behavior when they interact with light.

To physically synthesize and test a new dye can take weeks to months, and PPG needed to examine many materials. Computational research was key to fast time-to-market, but PPG soon found that the quantum-chemistry problems it wanted to tackle were too computationally intensive for their in-house systems.

By joining PSC's industrial affiliates program, PPG gained access not only to PSC computational resources but also to the expertise of PSC personnel. A complex calculation that might have taken a week on PPG's in-house machines — if it could be run at all, says Makowski, now takes only several hours.

This permits PPG to predict many performance characteristics of molecular structures under a variety of conditions, without having to construct a physical prototype. "We can computationally screen a whole series or family of new structures proposed by our organic chemists," says Makowski, "and weed out the 80-percent that will ultimately fail when they are tested experimentally."

With PSC resources doing quantum calculations, PPG brought its new fifth-generation photochromic technology to market ahead of competitors, gaining market share and increased sales and earnings. Ongoing work now continues and PPG's sixth-generation product is slated for launch within the next year.

"It's very difficult to calculate just how much time and money we are saving by working with PSC," says Makowski, "but the return has been very evident. We achieve faster, more accurate results with far fewer physical prototypes, which reduces cost. This not only means a more efficient R&D process, it also speeds up the time-to-market and enhances PPG's competitive position.

MORE INFORMATION

http://www.psc.edu/science/2008/stroke_shades.html

Transitions® lenses shift from dark to light automatically in the presence of harmful UV light.

Alex Rodriquez, *left*

Steven Johnson, *right*

John Joannopoulos

ATTRACTIVE REPULSION

An innovative numerical approach makes possible new insights into a fascinating quantum effect called the Casimir force

Attractive Repulsion

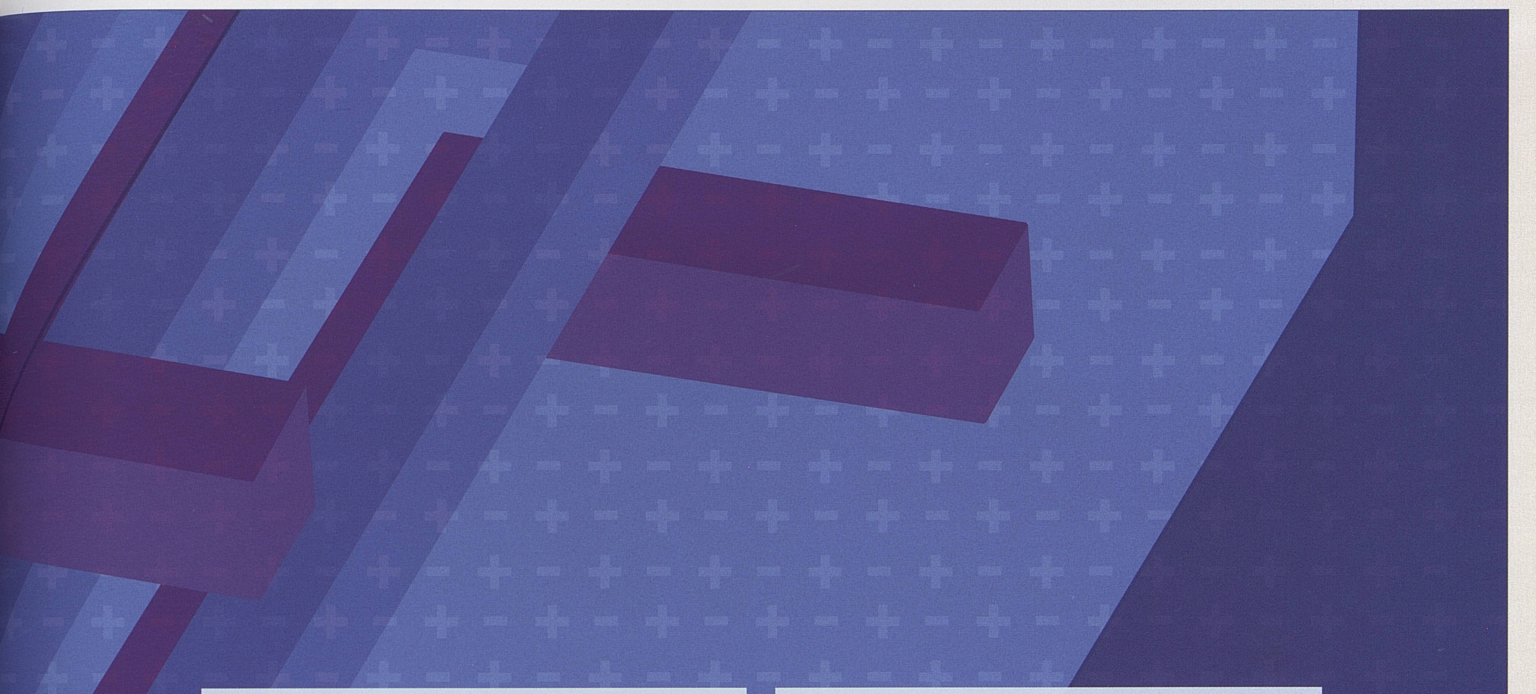

If opposite charges attract and like charges repel, what happens between neutral objects, ones that have no charge? They maintain their distance in a perpetual standoff, right? Not in the quantum universe, where space between objects bubbles with "virtual" particles and anti-particles of evanescent lifetimes that can cause even neutral objects — when they're close enough — to move nearer to each other. It's called "the Casimir effect" after Dutch physicist Hendrik Casimir, who in 1948 first calculated the attractive force between two neutral metal plates about a micrometer apart.

Since Casimir's original work, physicists have struggled to understand the forces between surfaces that aren't flat — a flat plate and a sphere, for instance, and beyond that to general geometries. "The problem has been so difficult," says Steven Johnson, MIT professor of mathematics, "that many researchers have employed drastic approximations to get some answer that they hope will be at least crudely correct."

In the last year, Johnson and physics grad student Alex Rodriquez and others who work with MIT professor John Joannopoulos have employed a numerical approach that is shattering long-held notions about the Casimir effect. Using BigBen, PSC's Cray XT3, they have carried out simulations that calculate an "exact" solution. Their method depends for its accuracy only on the availability of enough computational power. Moreover, it is independent of materials or geometry, so it can be used to study the Casimir forces in any system.

One of their important findings is that the Casimir force between two purely metal objects can be *repulsive*. In recent work with cylinders within cylinders, furthermore, they've for the first time demonstrated the possibility of stable, Casimir-force induced suspension between objects. "The possibility of repulsive forces is especially tantalizing," says Johnson, "because it raises the prospect of quantum levitation — suspension of mechanical parts in air or fluids purely from quantum fluctuations, without the device having to supply power." A levitating sphere of this sort, for instance, could serve as a frictionless bearing in a microelectromechanical system, aka MEMS.

To arrive at this fascinating finding has involved heavy-duty computation. "Computing these forces using *approximations* is hard," says Rodriguez, "but we're using *exact* Casimir forces, and the geometries we've been looking at lately are big." Referring to work at PSC between April and July 2008, says Rodriguez, "We essentially used up a year's worth of supercomputing in three months. They were very expensive calculations, and the results are very interesting."

A FROTH OF PARTICLES

For Casimir's original calculation — two parallel, flat-metal plates — the problem is, relatively speaking, easy. "There is a teeming froth of virtual particles," explains Rodriquez, "that cause these two plates to feel a pressure." You get a feel for this effect, he says, by thinking of the quantity of virtual particles outside the plates as exceeding those between the plates, so that pressure on the outside surfaces pushes the plates together.

Other geometries, however, become demonically complicated, and most previous attempts to model the Casimir force have used some version of an approximation known as "the proximity force approximation" (PFA) to make the calculations feasible. The PFA assumes any object is composed of tiny, discrete chunks of material, and these chunks interact only in pairs. In reality, however, other nearby chunks in the object also interact, and the PFA can drastically underestimate or overestimate the Casimir force.

In 1965, Russian physicist Evgeni Lifschitz developed a better approach. By carefully framing the problem in terms of electromagnetism, Lifschitz provided a mathematical tool that the Joannopoulos group turned to in their recent work.

Lifschitz's math involves solving Maxwell's electromagnetic equations for a point-source of electricity. Think of it, says Johnson, as calculating the electric field surrounding a tiny antenna through which an electric current is flowing. Physicists call this the Maxwell Green's function, a problem involving partial differential equations (PDEs). In this case, the supercomputer first solves many PDEs to obtain the electric-field value at a large number of points (tiny antennas) surrounding the surface of an object over a wide range of electromagnetic frequencies.

Summing these electric-field values yields an approximate value for the Casimir force. If this process is repeated many times, the solution eventually converges to the exact value. Accuracy depends only on the computational power and time available. "Unlike many earlier approaches to Casimir forces, we weren't forced to sacrifice generality for the problem to be solvable," says Johnson. "We were able to trust that the supercomputer could handle the most general possible problem and leave us free to explore geometries limited only by our imaginations."

BigBen is well suited, says Johnson, for this kind of calculation. "The Cray XT3 provided us with a good combination of many processors and relatively powerful individual processors."

REPULSIVE RESULTS

Letting their imaginations roam, the Joannopoulos group tried a geometry that looks like a zipper — two parallel plates with interleaving metal brackets. After setting the separation distance between the two plates, they concentrated 1,024 processors of the Cray XT3 on a circular space a few micrometers in diameter surrounding the surface of a bracket. Then the parallel plates of the zipper were moved closer together, and the process repeated, many times over.

The zipper model showed for the first time that the Casimir force could be repulsive between metals (see sidebar), a major result long suspected but impossible to prove without an exact numerical solution. Ironically, the repulsion results from attraction between sub-parts of the system. "No one had investigated this structure before," Rodriguez says, "nor demonstrated repulsive forces purely between metals."

Expanding on this discovery, the researchers recently studied an even more complex system — cylinders within cylinders, separated by a liquid (ethanol). The outer cylinder is metal, while the inner cylinder is silica, a semiconductor. To show that their method is independent of geometry, the team ran the simulation both with circular and square cylinders.

The results showed not only a repulsive force between the metal and the silica, but also that this force is stable — the inner cylinder is effectively "levitating" inside the metal cylinder. Such stable structures could act as frictionless bearings in a MEMS device, or as an oscillator in a switch.

New understanding of a remarkable direct manifestation of quantum fluctuations as a measurable force.

The greatest benefit of understanding the Casimir force may occur when MEMS systems evolve into nanoelectromechanical systems. The Casimir force in devices with parts separated by nanometers can cause moving parts to stick together and the device to grind to a halt — called *stiction*. "Casimir forces are thought to be a significant contributor to stiction in the smallest devices," says Johnson, "and a better understanding of these forces could lead to devices that operate more smoothly."

Even if there were no practical applications, the Joannopoulos group would continue to study

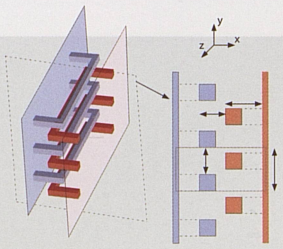

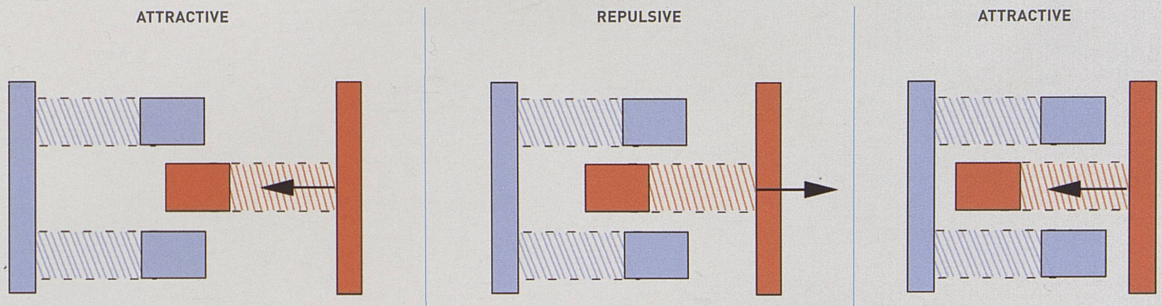

ZIPPER GEOMETRY: ATTRACTIVE REPULSION

The 3D schematic (upper left) shows the Casimir "zipper" geometry of interlocking metal brackets (in different colors for illustration only), along with a 2D *xy* cross-section. Dashed-lines indicate that the brackets attach to the plates as shown in 3D.

The three 2D boxes show the transition from attraction to repulsion for this zipper geometry. The repulsive force arises from attractive interactions of the bracket sub-parts of the system. With the parallel plates at relatively large separation, the Casimir force between them is attractive, almost as if the zipper brackets (squares) were not there. When the plates are closer together and the brackets overlap, a strong attractive Casimir force tries to pull the adjacent brackets toward each other, and this attraction resists the attraction between the plates, producing a repulsive force (indicated by the arrow). If the plates are closer yet together, however, the resistance of the brackets is overcome and the overall Casimir force is again attractive.

Casimir forces. They are interested in the physics of the "froth of virtual particles" that may be better understood with their ability to model Casimir forces exactly. "These forces," says Johnson, "are a remarkable direct manifestation of quantum fluctuations as a measurable force."

One of the possible phenomena they want to look for is a geometry that results in repulsive Casimir forces not due to attraction of sub-parts of the system. Such a purely repulsive geometry, if it exists, will require much more of BigBen's time. (TP)

MORE INFORMATION

http://www.psc.edu/science/2008/casimirforce.html

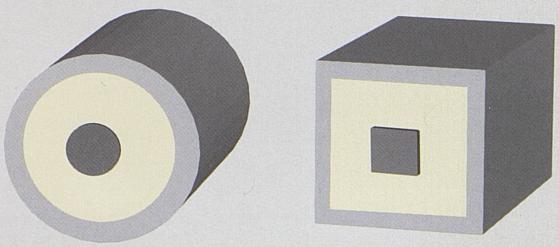

CYLINDERS WITHIN CYLINDERS: STABLE CASIMIR FORCE

Simulations for these two cylindrical geometries, circular and square, with a cylinder of silica (SiO_2) suspended in ethanol inside a metallic cylinder, showed a stable repulsive force between the two cylinders, so that the inner cylinder in effect levitates, presenting the possibility of frictionless bearing in a MEMS device.

Alexei Kritsuk, *top left*

Paolo Padoan, *top right*

Mike Norman, *middle left*

Rick Wagner, *bottom left*

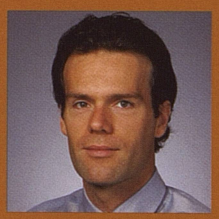

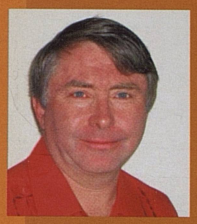

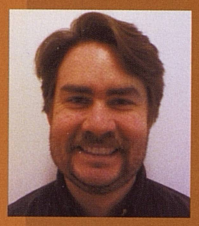

COOKING STARS AT MACH 6

Simulations open new understanding of hypersonic shock waves that contribute to the birth of stars

"In space, no one can hear you scream." Although this line from the movie "Alien" packs a dramatic punch, the physics is wrong. Sound travels in space, even in vast regions where the density of matter is extremely low. And if you were inside a molecular cloud — regions where hydrogen is dense and stars are born — you'd experience deafening booms of hypersonic sound waves.

Clouds of gas crashing inward under their own gravity produce sonic booms, and an occasional exploding star would add to the acoustical mayhem. "The place where stars are formed is not in the least quiet," says University of California, San Diego (UCSD) astrophysicist Alexei Kritsuk.

Kritsuk and UCSD colleagues Michael Norman, Paolo Padoan and Rick Wagner have used a range of TeraGrid computing resources to simulate hypersonic shock waves. They are especially interested in how these violent phenomena induce turbulence — one of the most important factors in a complex of factors that lead to the birth of new stars. But accurate simulation poses exceptional challenges because, while most turbulence studies deal with "incompressible" flow — having uniform fluid density — the shock waves in molecular clouds produce huge contrasts in density and the turbulence is highly compressible.

In 2007, Kritsuk and colleagues used systems at SDSC and NCSA for a series of simulations of turbulent flow at Mach 6. Their results, reported in the *Astrophysical Journal* (August 2007), showed that well-established characteristics of incompressible turbulent flows do not hold for supersonic compressible turbulence, and they proposed a more general theory that encompasses both flow regimes.

continued page 40

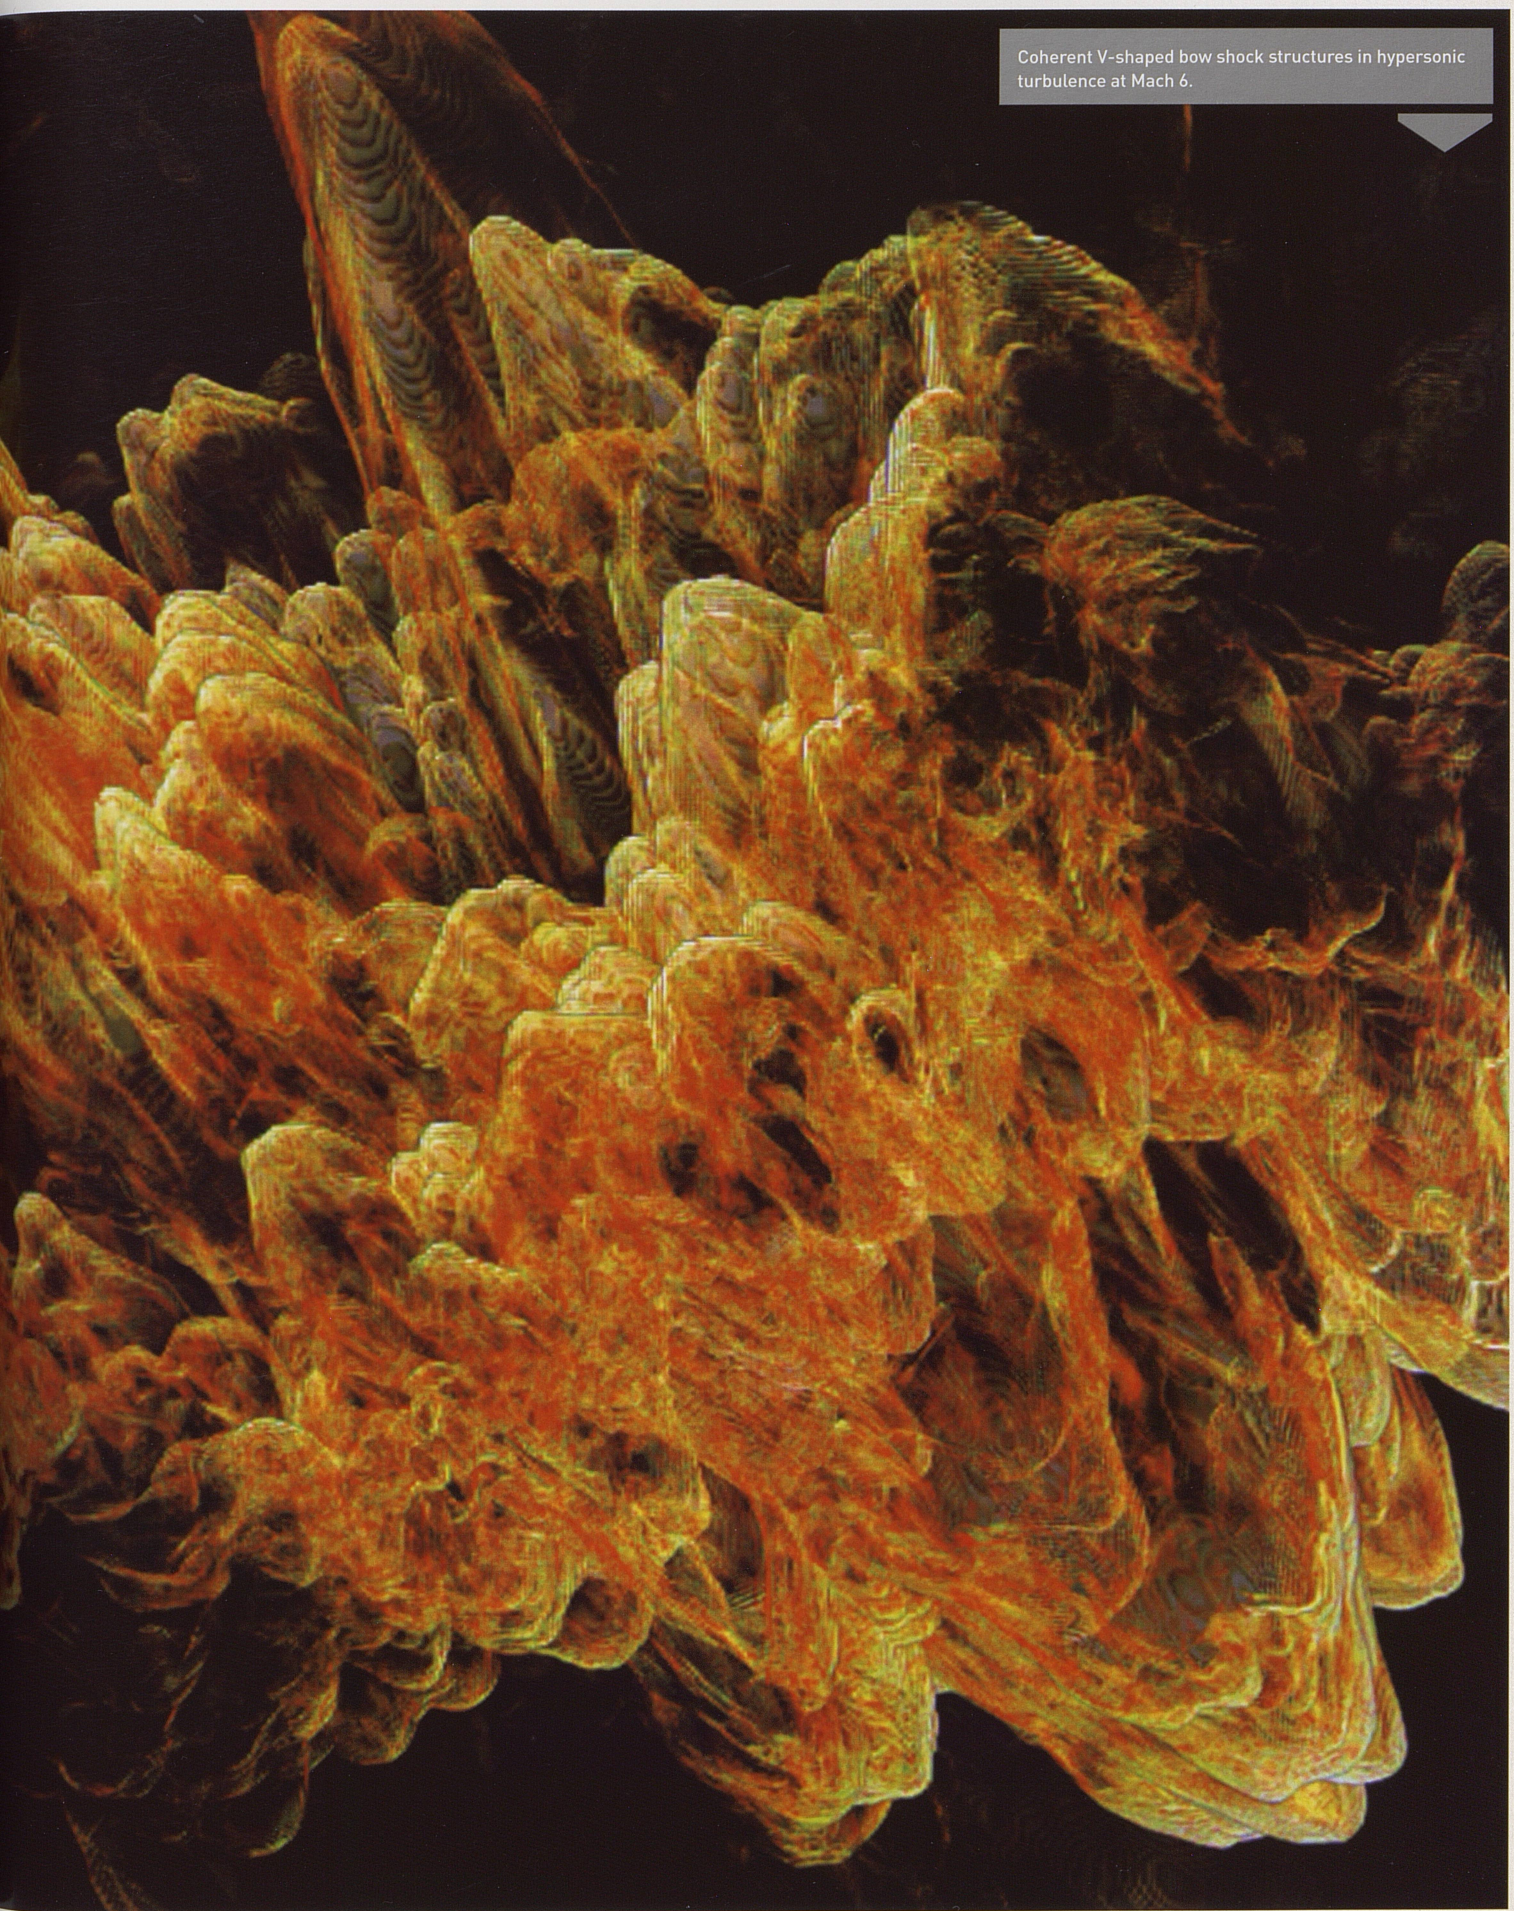

Coherent V-shaped bow shock structures in hypersonic turbulence at Mach 6.

continued from page 38

This year, the researchers extended this work with simulations at PSC (BigBen) and TACC (Ranger). This newer work, involving 1.3-million hours of computing, produced 50 terabytes of data. With analysis still underway (using SDSC's DataStar and PSC's Pople), preliminary results show that the new data confirm the prior studies. Beyond what they tell us about star formation, these studies may shed new understanding on the hypersonic aerodynamics of high-speed aircraft and re-entry space vehicles.

STIRRED, NOT SHAKEN

The space between stars — the interstellar medium — for the most part contains cold, neutral (no electrical charge) hydrogen atoms. "If you want to create a star," says Kritsuk, "you need to cook up a cloud of the cold, neutral medium, then find a dark space and compress it further, so molecular hydrogen (H_2) begins to form, and as it cools to lower temperatures and the densities go up, you will eventually get a collapsing protostellar core."

When shock waves occur in these clouds, they push the hydrogen molecules closer together, and at a certain critical density stars begin to form. It sounds simple enough, but much is unknown about these turbulent phenomena and their impact on star formation.

Most of what we know about turbulence is based on geophysical and engineering phenomena: mixing of hot and cold air in the atmosphere to form weather fronts, flow of oil through a pipeline, mixing of gas vapors and air in an engine's combustion chamber. The fluid is at virtually the same density throughout the turbulent region, with only a small amount of variation. Scientists call this uniform state an "incompressible fluid."

Shock waves moving at hypersonic speeds, however, create enormous density variations. One region of a gas cloud can be a million-times denser than a nearby region, and this "compressible" fluid behaves much differently than its incompressible cousin. "If you are sitting on a blob of gas that is a million-times denser than the surrounding medium," says Kritsuk, "you will just fly like a bullet."

To simulate gas clouds with such a density contrast, Kritsuk and his colleagues used ENZO, a program developed at the Laboratory for Computational Astrophysics at UCSD for simulations of cosmological structure. They modeled a cubic volume of space five parsecs (about 98-trillion miles) on a side, filled with hydrogen atoms at an initial density of 500 particles per cubic centimeter and a temperature of 10-degrees Kelvin. In the first round of simulations, they subdivided this volume with a grid (512^3 points) and zoomed in to a finer resolution (2048^3 points)

This graphic shows gas density in a thin slice through the simulation volume, increasing from very low to very high (black through red and yellow to white).

Simulating these clouds of molecular hydrogen led to a new fundamental understanding of turbulence.

to get a closer look when the density in a particular area increased by a factor of two or more. In the latest simulation, they used the higher resolution throughout the run to take advantage of the Cray XT3's larger memory capacity and speed-up the calculations.

The shock waves produced in these simulations churned cold clouds of molecular hydrogen into compressed regions, and led to a new fundamental understanding of turbulence.

RESTORING THE FORGOTTEN VARIABLE

In 1941, Russian mathematician Andrei Kolmogorov developed equations that offered a new insight into turbulent flow. In Kolmogorov's view, large vortices or eddies occurred first in a fluid in which turbulence formed, and these large vortices then formed smaller vortices, leading to a conception of turbulent flow as "whirls within whirls."

Velocity and density are both part of Kolmogorov's equation, but for a long time researchers concentrated only on velocity. This made sense, says Kritsuk, because mostly they were concerned with solving

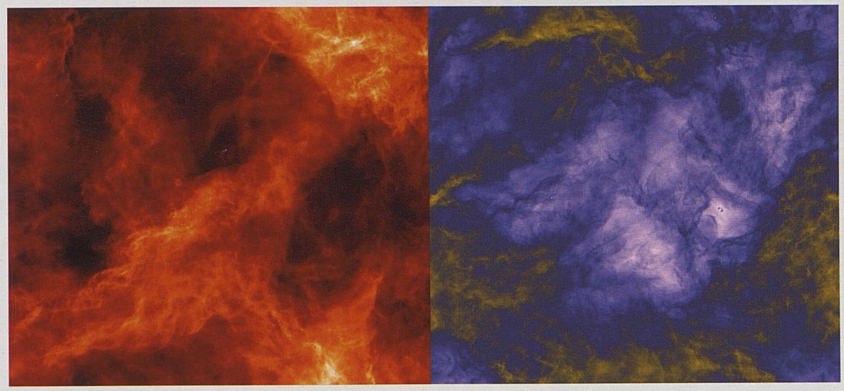

3D projections through the simulation volume represent density of the molecular cloud.

engineering problems. But this concentration also led to a blind spot. "People were working for so many years with incompressible turbulence that they almost completely forgot about density, and what we have done here is put the density variable back in place.

In restoring this forgotten variable, Kritsuk and his colleagues have extended turbulence theory to new levels. "We have shown that if you put density back in place you still get Kolmogorov's law, even for hypersonic flows," says Kritsuk. "When you have density contrasts of a million, it's really important to include density-velocity correlations — and then modeling turbulence becomes simple."

To make things even more simple, Kritsuk is analyzing the 50 terabytes of data using SDSC's DataStar and Pople, PSC's newest system whose shared memory and open multi-processing (OpenMP) architecture allows him to work with the huge 2048^3 data-cubes quickly and efficiently. Pople also shares a parallel file-system with BigBen that minimizes file-transfer time.

Besides demonstrating that Kolmogorov's theory applies universally to compressible and incompressible fluids, Kritsuk's work also confirmed that there is order in the apparent chaos of turbulence. Visual representations of the simulation reveal "coherent structures" in the molecular clouds. The most common structural elements are nested bow-shocks having a V-shaped structure, also known as "Mach cones." These cones form when supersonic shock waves coming from opposite directions collide in the molecular cloud. The large V-shaped bows are made up of sequentially smaller V-shaped bows to the current limits of resolution in these simulations, indicating the self-replicating, fractal nature of these structures.

Finally, in confirming that turbulence described by Kolmogorov's equation is universal, covering situations where the density of a fluid varies by a factor of a million, Kritsuk's group may also help astrophysicists explain another universal phenomenon: the uniform distribution of star masses in molecular clouds — called the "initial mass function."

"You basically have some small fraction of massive stars, then more stars of smaller masses," explains Kritsuk. "The peak of this mass distribution of newly formed stars is roughly one solar mass — the mass of our sun. Then there are brown dwarfs and such stars having very small masses." Everywhere you look in the universe this distribution is the same.

Kritsuk's collaborator Paolo Padoan and Åke Nordlund of the Neils Bohr Institute, Copenhagen have proposed that supersonic turbulence in dense molecular clouds is the process that underlies the initial mass function, fragmenting the interstellar medium and creating the initial conditions for star formation. Based on his latest simulations, Kritsuk tends to agree. "There are not that many candidates out there. We have high hopes that the universal process we all know —turbulence—will solve the mystery of the initial mass function."

Kritsuk next plans to collaborate with researchers at Purdue University who have a shock tube capable of achieving Mach 6 hypersonic flow — to validate the computational model experimentally. Then he wants to do a second study using BigBen to add magnetic fields to the model of star formation. "Of course, to really model star formation, you need to include magnetic fields," says Kritsuk. "Numerically it's a challenge, and there's a lot of new interesting physics there. Adding the magnetic field components is like opening a Pandora's box." It's a box for which Kritsuk and his colleagues are eager to start prying the lid. (TP)

MORE INFORMATION

http://www.psc.edu/science/2008/mach6.html

IN PROGRESS //

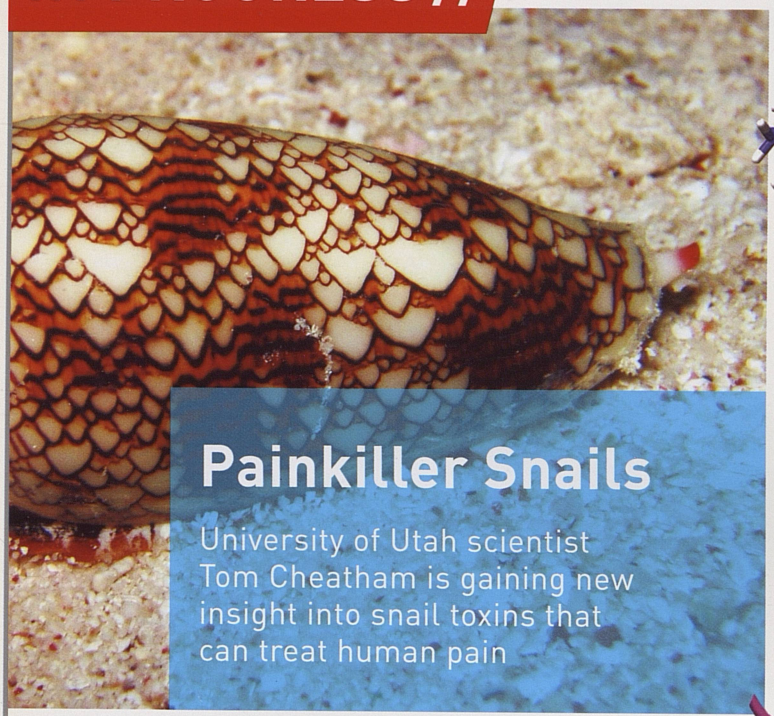

Painkiller Snails

University of Utah scientist Tom Cheatham is gaining new insight into snail toxins that can treat human pain

Although they come in pretty shells shaped like ice-cream cones and are popular for jewelry, cone snails are carnivorous hunters, and some of the larger species can be deadly to humans. They are equipped with a barbed, hollow tooth that launches like a harpoon. A neurotoxic venom paralyzes small fish almost instantly.

The good news is that these cone-snail venoms, known as conotoxins, show promise in humans as potent painkillers. One conotoxin-derived drug, Ziconotide, received FDA approval in 2004, and others are being tested as possible treatments for Alzheimer's disease, Parkinson's disease and epilepsy.

University of Utah pharmaceutical chemist Tom Cheatham and visiting student Pawel Gruszczynski from the University of Gdansk in Poland are using Pople, PSC's shared-memory SGI Altix, to investigate a family of conotoxins called Mu-conotoxins. "These toxins," says Cheatham, "are tissue specific and act by binding to voltage-sensitive sodium channels." Using long-timescale molecular dynamics (MD), with software called AMBER, Cheatham's group is working to predict and refine the structure of these conotoxins. With Pople, says Cheatham, the aggregate throughput was excellent. "We benefit from the fastest and most balanced machines."

Their results, analyzed in collaboration with University of Utah experimentalists Greg Bulaj and Toto Olivera, make it possible to map the pharmacophore — the essential molecular features responsible for the drug's biological activity. This work aims toward identifying the precise atom-by-atom details of how conotoxin molecules — short amino-acid chains called peptides — have their potential pain-killing effect.

These pictures of a Mu-conotoxin show representative structures from clustering that occur with 200 nanoseconds of MD simulation. "Flexibility in the side-chain is evident," notes Cheatham, "although the core backbone structure is largely intact."

Shared Memory Poker

Carnegie Mellon computer scientists are using PSC's Pople to scale up their champion poker program and solve other problems in game theory

You won't find them at the Vegas casinos and what they do is hard to call gambling, but it's fair to say that Tuomas Sandholm and grad student Andrew Gilpin of Carnegie Mellon's School of Computer Science are professional poker players. This July in Chicago — at the AAAI (Association for the Advancement of Artificial Intelligence) Computer Poker Competition, involving 19 programs from six countries — they walked away with no pile of cash but, nevertheless, were the biggest winners.

Their field, game theory, in which Sandholm's work is internationally recognized, describes conflict in which the payoff is affected by actions and counter-actions of intelligent opponents. Head-to-head poker between two players is a prominent example of what's called two-person zero-sum games: One player wins what the other loses.

In recent years, poker has emerged as an AI challenge much as chess was for many years, but poker is far more demanding, says Sandholm: "In poker, a player doesn't know which cards the other player holds or what cards will be dealt in the future."

Like many games, poker can be formulated mathematically, but the formulations are huge. Two-player poker has a game-tree of a billion-billion (10^{18}) nodes, notes Gilpin. "To solve that requires massive computational resources. Our research is on scaling up game-theory solution techniques to those large games, and new algorithmic design." The most computationally intensive portion of their algorithm is a matrix-vector product, where the matrix is the payoff matrix and the vector is a strategy for one of the players. This operation accounts for more than 99-percent of the computation, and is a bottleneck to applying game theory to many problems of practical importance.

To drastically increase the size of problems the algorithm can handle, Gilpin and Sandholm devised an approach that can potentially exploit massively parallel systems of non-uniform memory-access architecture, such as Pople, PSC's SGI Altix. By making all data addressable from a single process, shared memory simplifies a central, non-parallelizable operation performed in conjunction with the matrix-vector product. Sandholm and Gilpin have revised their code to run on Pople, and are doing experiments to learn how much parallelism can help, and possibly point to areas for further algorithmic improvement.

To Stop the Pandemic

Modeling with PSC shared-memory resources will improve decision-making in deploying vaccines to stop disease outbreaks

With support of a $10-million grant from the Bill & Melinda Gates Foundation, a team of researchers, including PSC scientist Shawn Brown, is using PSC resources to explore computational models that simulate the spread of infectious disease. The grant funds computer simulations of epidemics to show worst-case and best-case outbreak scenarios. Findings will be used to evaluate new vaccine technologies and modes of vaccine delivery. The goal is more informed decision-making in deploying vaccine technologies in pandemic outbreaks and, in general, to support World Health Organization Communicable Disease Control programs.

The work is coordinated through the Vaccine Modeling Initiative (VMI), headquartered at Pitt's Graduate School of Public Health (GSPH). VMI is a research partnership among infectious-disease modeling teams at the University of Pittsburgh, The Pennsylvania State University and Imperial College London. VMI also involves collaborations with leading infectious disease experts and public health officials at Johns Hopkins University, Médecins Sans Frontières Epicentre, University of Georgia, Resources for the Future and the World Health Organization. The project will exploit the shared-memory environment of Pople, PSC's SGI Altix, to build extensive agent-based models of disease spread in a number of third-world populations. The group expects that a large shared-memory system will improve their ability to complete full studies in a timely fashion.

"Infectious diseases create an enormous burden on the world's population, from both a human suffering and an economic development perspective," says Donald S. Burke, M.D., principal investigator of the grant and dean of GSPH. "One of the major challenges we face in stopping infectious disease outbreaks is predicting how control strategies, such as vaccines, will work. By using computer models to conduct 'epidemiology in silicon,' we will be able to test the impact of new candidate vaccine technologies and select the most effective strategies."

Initially, the project will focus on evaluation of new vaccine technologies for influenza, measles and dengue, a mosquito-borne infection, diseases that affect millions of people globally. Later, the project will develop vaccine models of epidemic pertussis, rotavirus, polio, pneumococcus, malaria and tuberculosis.

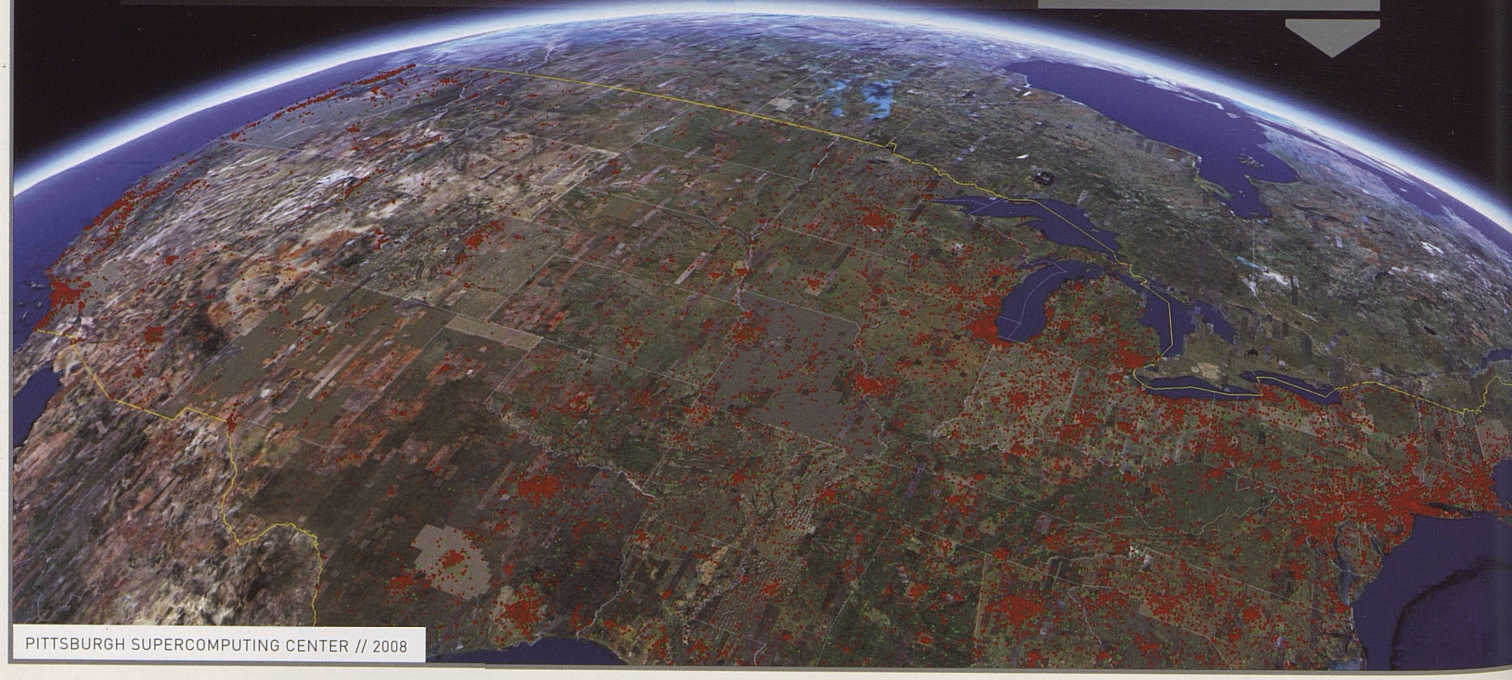

Day 35 of an influenza outbreak in the United States as simulated with an agent-based model (written by Neil Ferguson, Imperial College, London) on Pople, PSC's SGI Altix. The image shows the infected population (red) and population segments that have recovered (green). PSC scientist Shawn Brown created this visualization with a PSC-developed interface to Google Earth.

Hitting the Fight-or-Flight Target

With help from BigBen, Thomas Huber of Rockefeller University makes progress toward understanding the molecular mechanism of therapeutic drugs

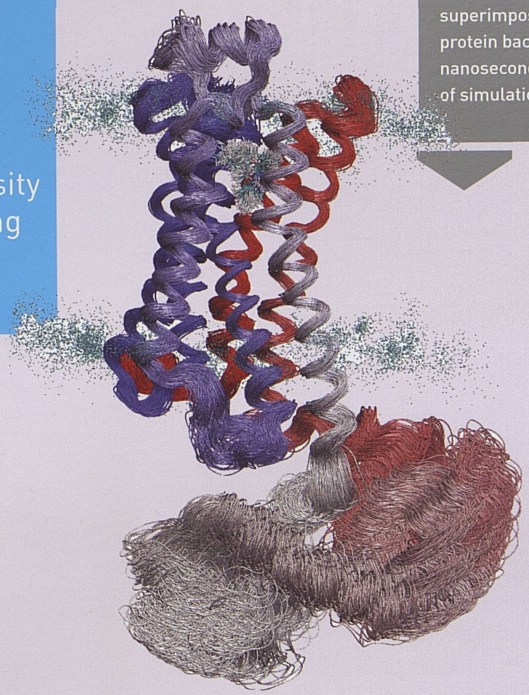

This representation of Huber's simulations shows a BAR structure (seven vertical helical coils) embedded within a bilayer membrane (cyan spheres represent the membrane-water interface), with a ligand called carazolol (amid the vertical coils) that blocks BAR activity. The overall model is about 50,000 atoms. The image shows superimposed structures of protein backbone snapshots every nanosecond over 60 nanoseconds of simulation.

Drugs have targets. They're called receptors — usually large proteins lodged in cell membranes. When another molecule of the right shape, size and electrical charge arrives — called a ligand, it binds to and activates the receptor, igniting a cascade of biochemical processes. Most ligands are from within the body (endogenous), such as neurotransmitters. Drug research involves finding other ligands, either to activate receptors or to block them from being activated by endogenous ligands.

One "superfamily" of receptors — called G protein-coupled receptors, GPCRs for short — is estimated to be the target for 27-percent of approved therapeutic drugs. Within the GPCR family, an important sub-group is adrenergic receptors (ARs), which interact with the "fight-or-flight" hormones epinephrine (adrenaline) and norepinephrine. Many cells have ARs, and binding of these hormones increases heart rate, diverts blood to muscle tissue, and mobilizes energy.

Thomas Huber of Rockefeller University has been using PSC's BigBen to better understand the biophysical principles of ligand binding to GPCRs. In recent work he has focused on one of the ARs — beta(2)-adrenergic receptor, BAR for short — the main target for drugs that treat asthma, by relaxing the pulmonary muscles. The crystal structure of BAR in 2007 was a long-awaited breakthrough in the field. Huber used molecular dynamics (MD) software to simulate BAR structures as soon as they became available for study. Working with PSC scientist Phil Blood to take advantage of BigBen's parallelism, he was able to run four independent 64-processor jobs at one time. "For the work I'm doing, this is the fastest machine available," says Huber. "The queues are relatively predictable and allow good job throughput."

On models of BAR embedded in a membrane with receptor that encompasses 50,000 atoms, he completed 600 nanoseconds of MD simulation. One set of simulations used the beta-blocker drug carazolol, which binds to and blocks activation of the receptor in pharmacological experiments. The second set used the activating drug epinephrine. Findings so far show details of how BAR's ligand-binding pocket changes when bound alternatively to these two drugs. While the structural changes are too small to account for receptor activation, the long simulations on BigBen demonstrate that these changes are statistically significant. "Information on this timescale," says Huber, "shows how receptor activating drugs pull on the receptor in an attempt to toggle its switch." This information can help point the way to new GPCR-targeted drugs.

Attracting Future Scientists

An innovative science-learning program called CMIST is helping teachers bring science to life for high-school students

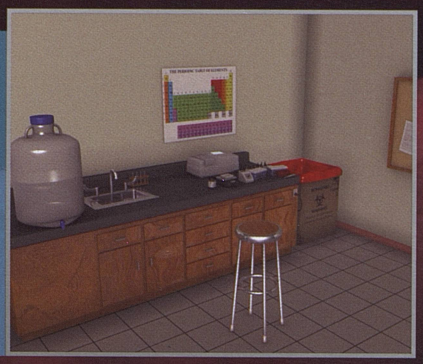

From the CMIST module "Enzyme Structure and Function": A virtual laboratory designed and rendered with Blender includes common lab equipment related to MCell simulations, which in this case show red blood cells in a heart vessel delivering oxygen to muscle cells.

"The visualizations allow them to understand the concepts as they actually occur," says Rebecca Day, who teaches high-school biology near Pittsburgh. "It's one thing to read about it, another thing entirely to see it happening," says Marian Opest, who teaches science in the Pittsburgh suburb of Penn Hills. They're talking about CMIST (Computational Modules In Science Teaching), an innovative science-learning program for high-school students developed by a team at PSC.

"CMIST addresses the challenge of science learning for the video-gamed, TV-nation, multi-tasking population of students." says Pallavi Ishwad, education outreach specialist for PSC's National Resource for Biomedical Supercomputing (NRBSC). Ishwad, NRBSC director Joel Stiles, scientific e-learning specialist Jenda Domaracki and visualization specialist Jacob Czech created and developed CMIST as a learning tool that could reach the majority of students, not just the self-motivated few. The hook is appealing content in an easily usable form.

In contrast to many other teaching tools, CMIST modules are produced with realistic modeling and simulation software, such as developed and used in research at PSC. The pilot module, "Molecular Transport in Cells," produced with software called MCell and DreAMM (co-authored by Stiles and used in research centers around the world) presents important principles of osmosis and diffusion with 3D examples. Recent new modules include "Brownian Motion" and "Enzyme Structure and Function." CMIST is distributed as ready-to-use DVDs that include complete lecture slides, animations, a lesson plan aligned to national and state standards, and worksheets with answer keys.

Funded in large part by NIH, the CMIST team has presented the program to regional high-school science teachers in half-day workshops at which teachers have given enthusiastically positive feedback. PSC has also presented CMIST to teachers and administrators at state and national meetings, and is introducing web-based distribution of the prototype module, while also planning additional modules on metabolism, synaptic physiology, meteorology, geophysics and astronomy, informed by some of the forefront research at PSC.

PSC scientist Shawn Brown demonstrates WiiMD at a National Science Foundation Open House in Arlington, Virginia. "The main goal" says Brown, "is providing a means of education and outreach that connects people to supercomputing and simulations in an engaging and familiar way."

Kids and Adults Say "Wow"

The intuitive technology of Nintendo's video console applied to molecular dynamics makes a game of science learning

With an innovative hybrid of game technology and supercomputing, a team at PSC is winning new friends, young and old, to the excitement of discovery and learning. Their interface, called WiiMD, integrates the controller of the popular Nintendo Wii (pronounced "we") game console, the WiiMote, with molecular dynamics simulation software. The result for a wide range of audiences is an experience both entertaining and informative.

The package of simulations includes, for instance, buckyball bowling, in which participants use the WiiMote to impel a "buckyball" carbon molecule to smash a 10-pin formation of other buckyballs. The team has demonstrated WiiMD at various venues, including SC07 in Reno, the National Science Foundation Open House in Arlington, VA, and at Pittsburgh's Carnegie Science Center SciTech Festival. Children are captivated by the colorful display of molecular shapes that move and change in response the WiiMote. Along with students from K-12, WiiMD has drawn enthusiastic attention from teachers and the general public.

At the heart of WiiMD are sophisticated biomolecular simulation technologies (NAMD and VMD) developed by Klaus Schulten and colleagues at the University of Illinois Urbana-Champaign and used by researchers world-wide on the largest supercomputing resources available. Jordan Soyke, a University of Pittsburgh student, adapted the WiiMote controller for use with VMD to enable fully 3-D control. Further work by PSC scientists Shawn Brown and Philip Blood — relying on PSC-developed software called PDIO — connected the programs with BigBen, PSC's CrayXT3.

With initial implementation of WiiMD complete, the team is developing a curriculum to go with it. Marylou Kunkle, like Soyke an undergrad at the University of Pittsburgh, is adapting the Linux-based program to run on desktop platforms, such as Windows and Mac OS X, in order to allow easier use by the general public.

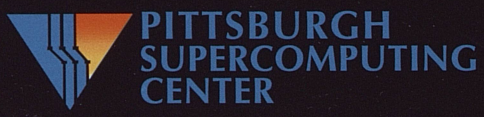

The Pittsburgh Supercomputing Center is a joint effort of Carnegie Mellon University and the University of Pittsburgh together with Westinghouse Electric Company. It was established in 1986 and is supported by several federal agencies, the Commonwealth of Pennsylvania and private industry.

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 National Archives and Records Administration
The U. S. Department of Defense
The U. S. Department of Energy
Cisco Systems, Inc.
Cray Inc.
DataDirect Networks
Hewlett-Packard Company
Silicon Graphics, Inc.
The Buhl Foundation
Bill and Melinda Gates Foundation
The Grable Foundation
The Heinz Endowments

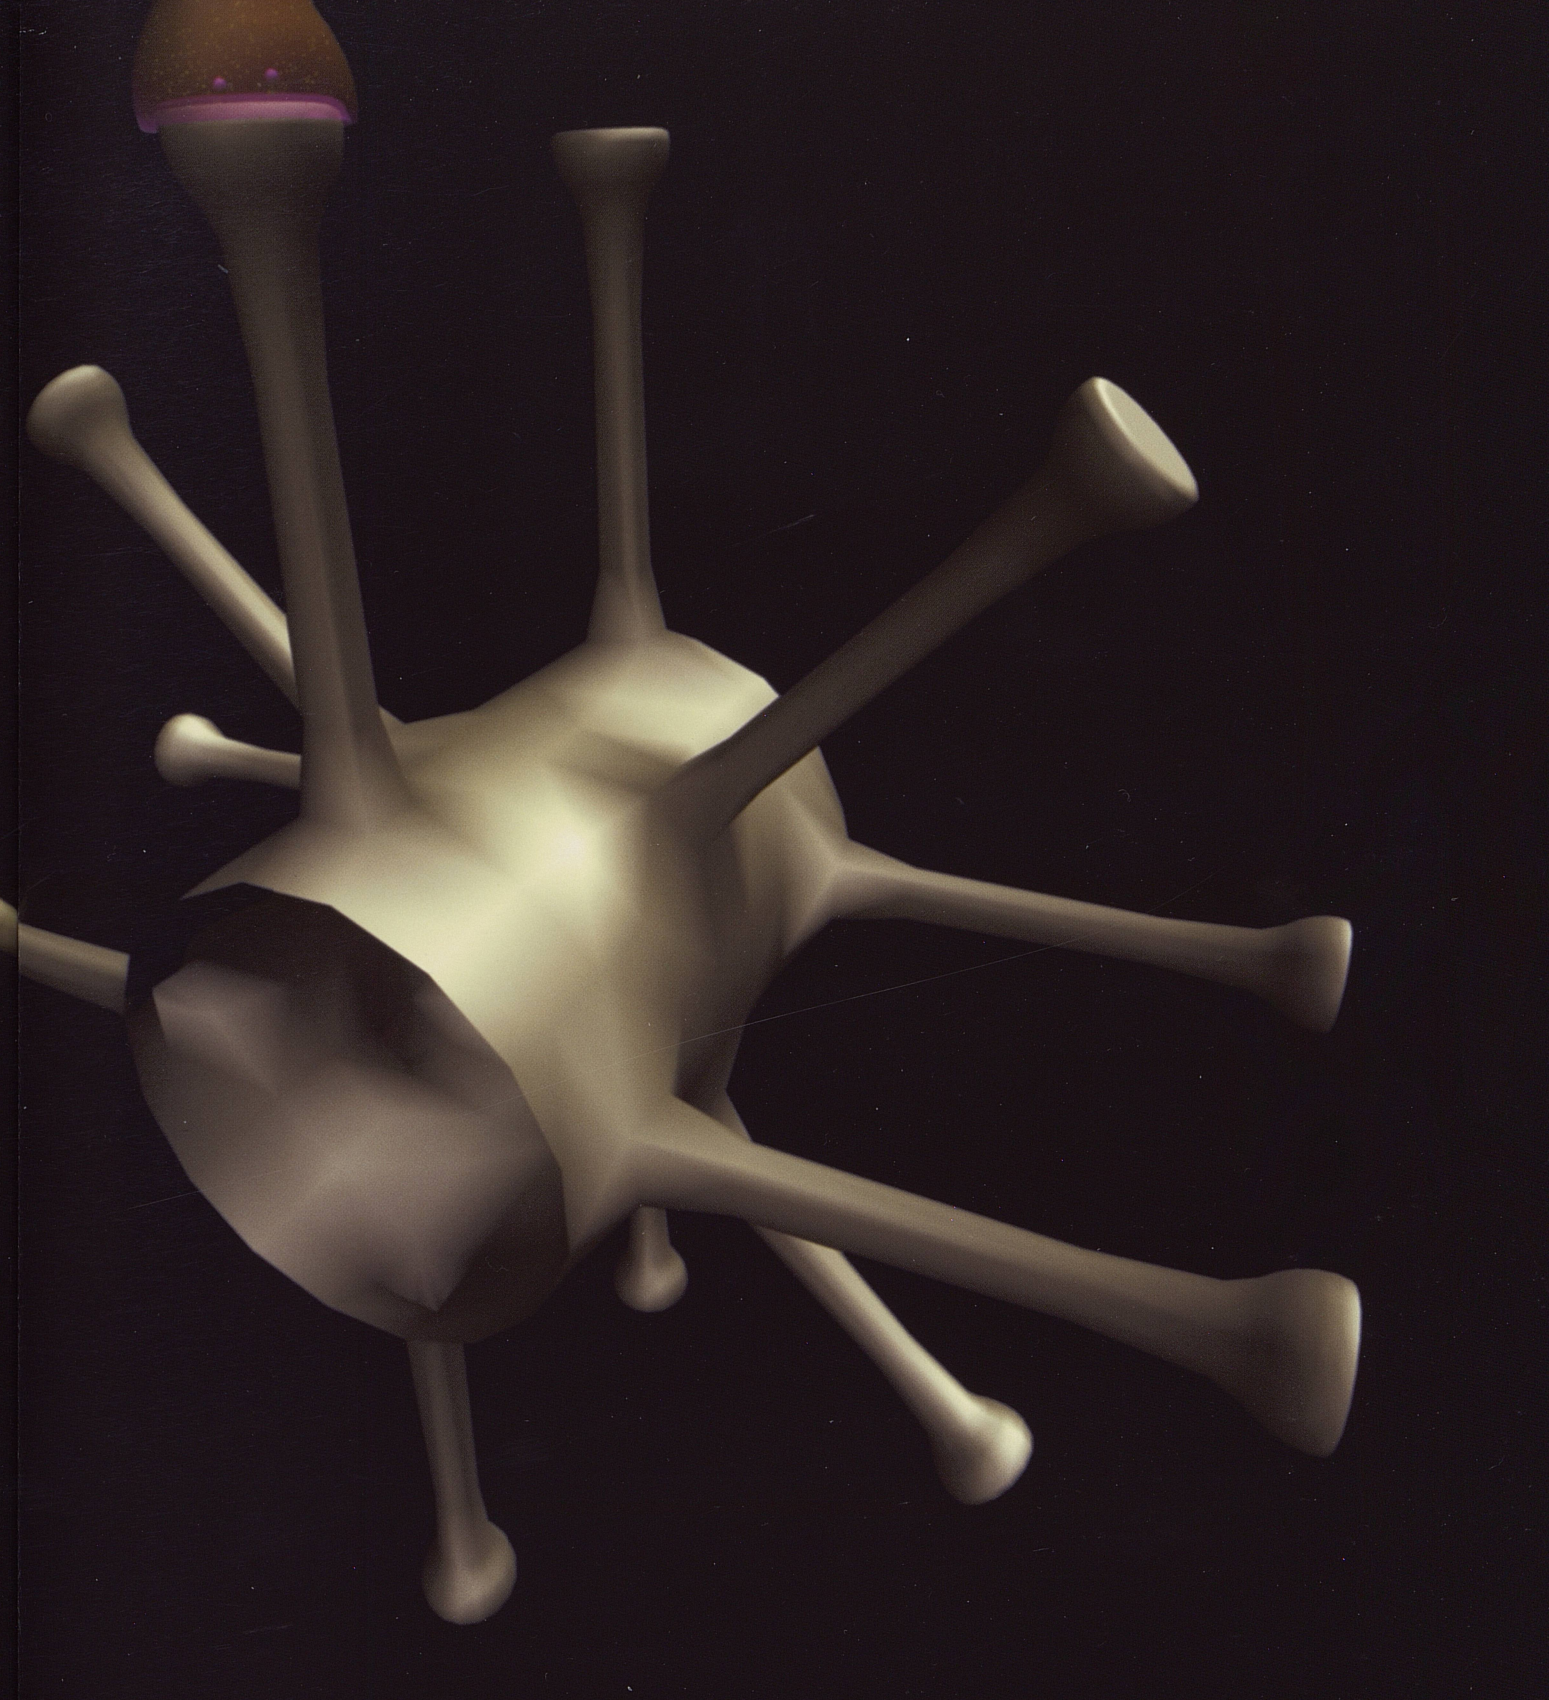

EDITOR/WRITER: Michael Schneider, PSC

CONTRIBUTING WRITERS: Tim Palucka, Sheila Prakash, Frank Stephenson

GRAPHICS RESEARCH, PHOTOGRAPHY DIRECTION & DESIGN COORDINATION: Shandra Williams, PSC

TRANSCRIBING: Ryan Omecene, Jessica Dailey, PSC

PHOTOGRAPHY: Photography & Graphic Services at Mellon Institute

GRAPHICS: Thanks to the researchers and to Shawn Brown, Greg Foss, Joel Stiles, Joel Welling, Art Wetzel, and Troy Wymore of PSC for graphics.

COVER GRAPHIC: Flow of red blood cells in a vessel delivering oxygen to muscle cells in the heart (myocytes), as rendered for CMIST (see p. 46) developed by PSC's National Resource for Biomedical Supercomputing.

DESIGN: Wall-to-Wall Studios, Inc.

PRINTING: Broudy Printing

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

PITTSBURGH SUPERCOMPUTING CENTER
300 S. CRAIG STREET
PITTSBURGH, PENNSYLVANIA 15213

PAID