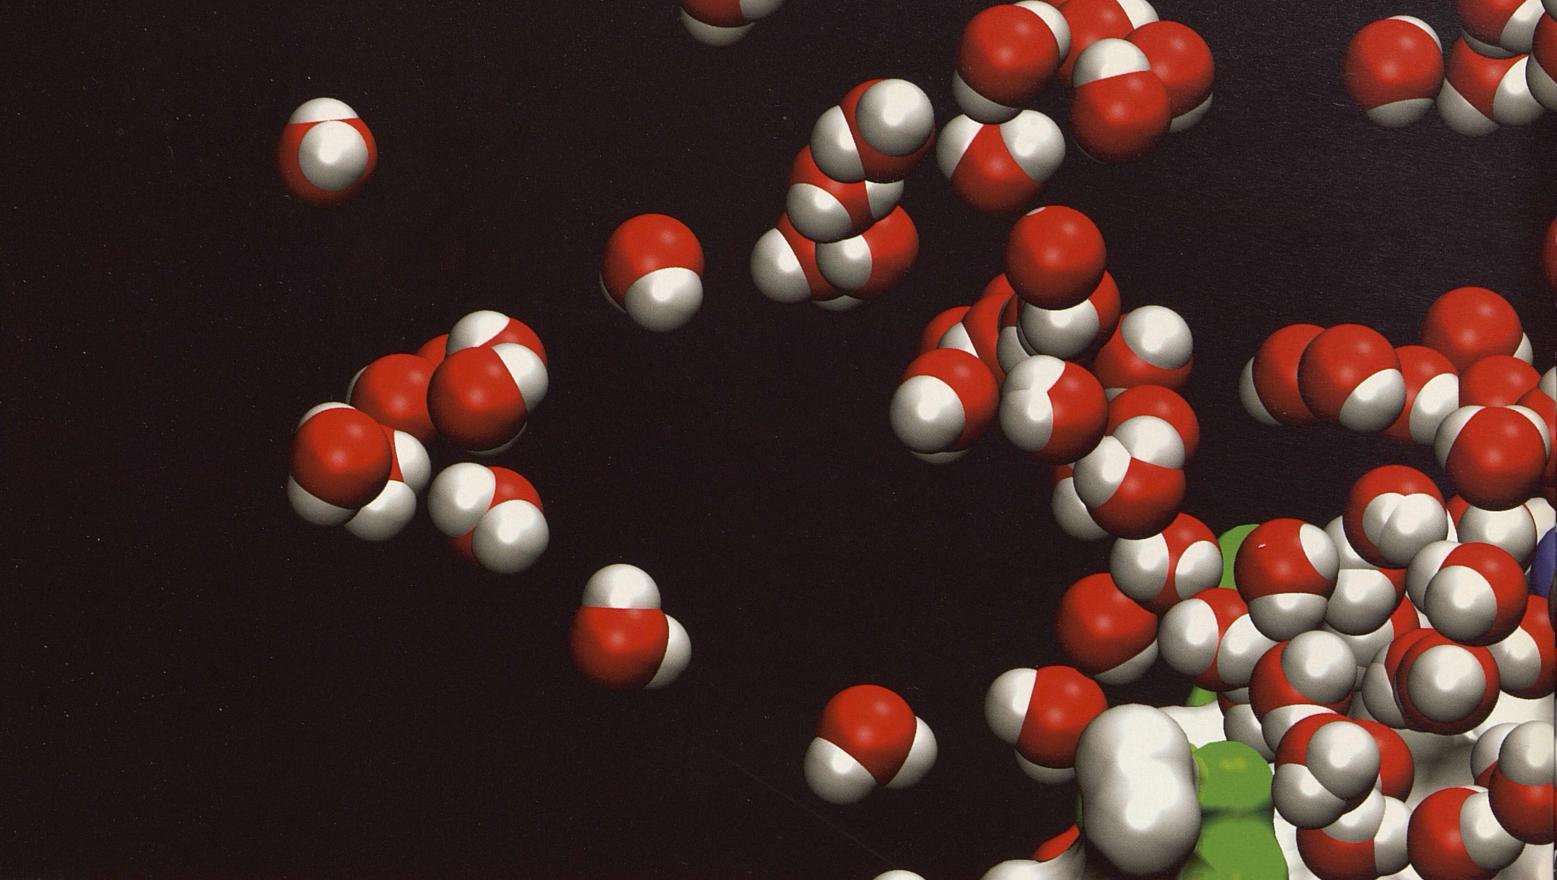


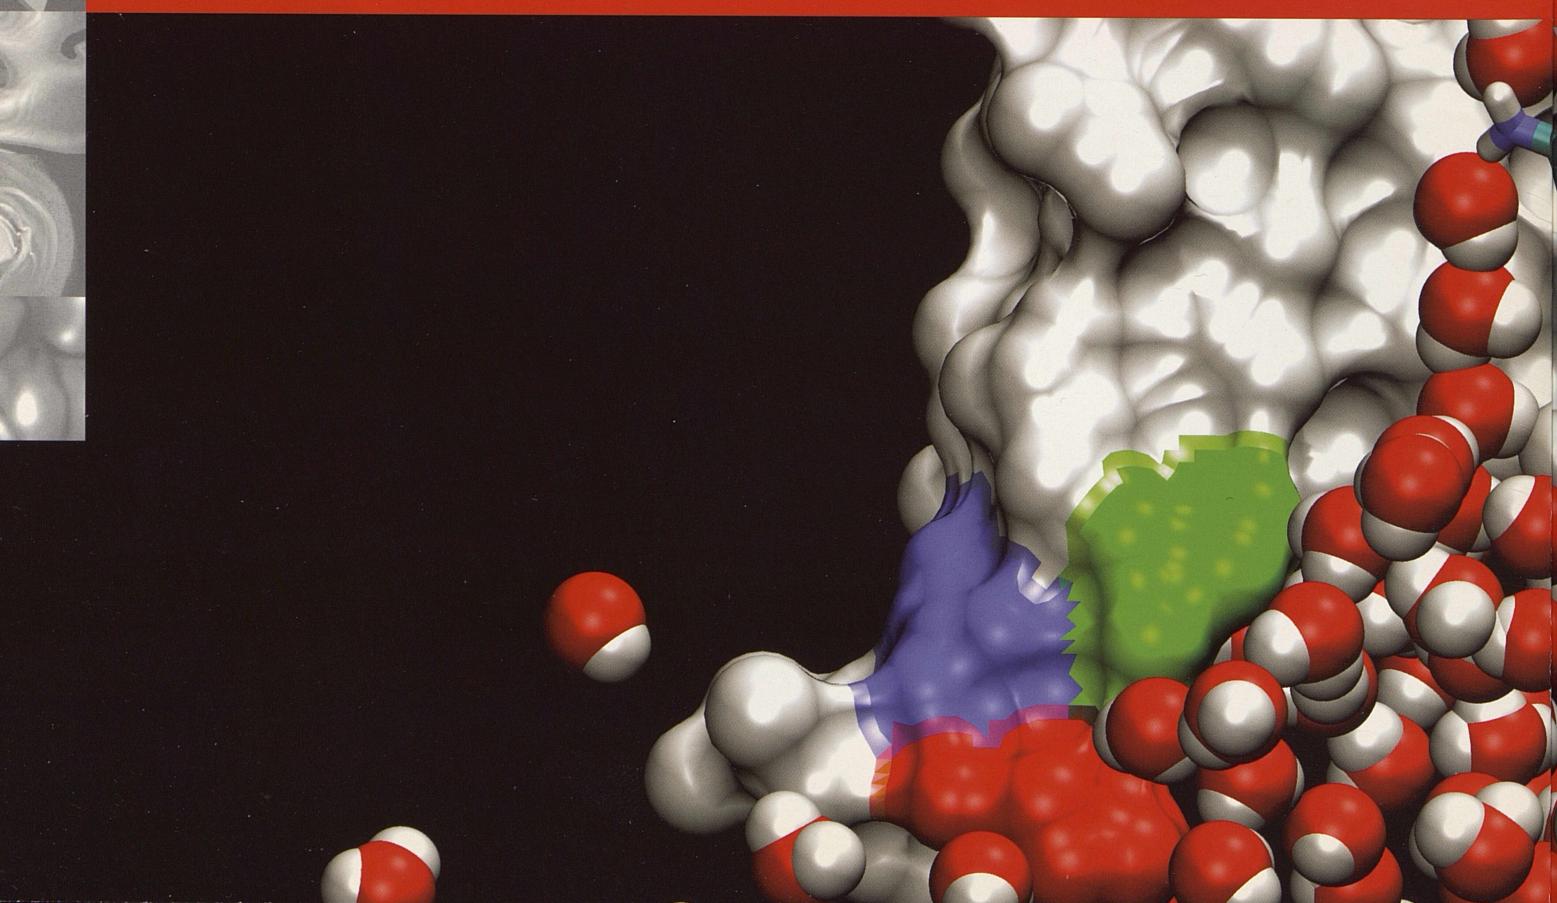
PITTSBURGH
SUPERCOMPUTING
CENTER 2003

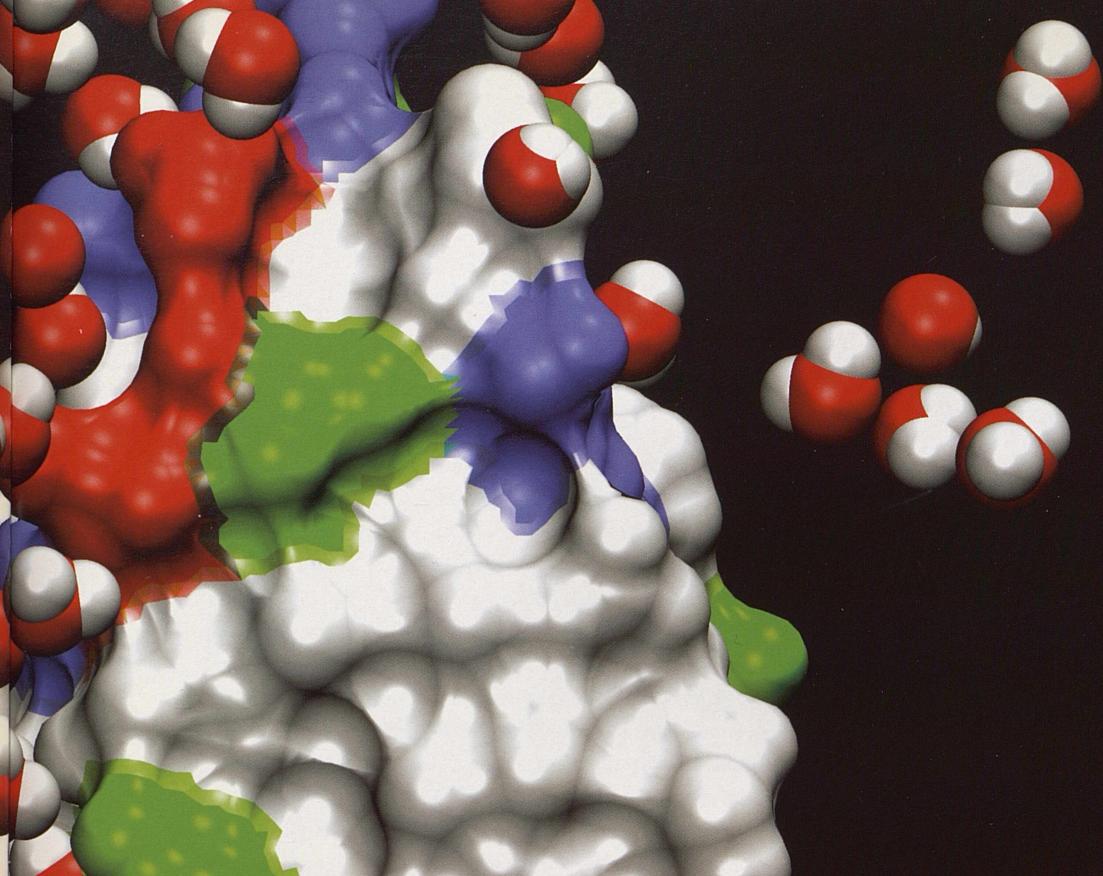
→ PROJECTS IN SCIENTIFIC COMPUTING



PSC.EDU/03

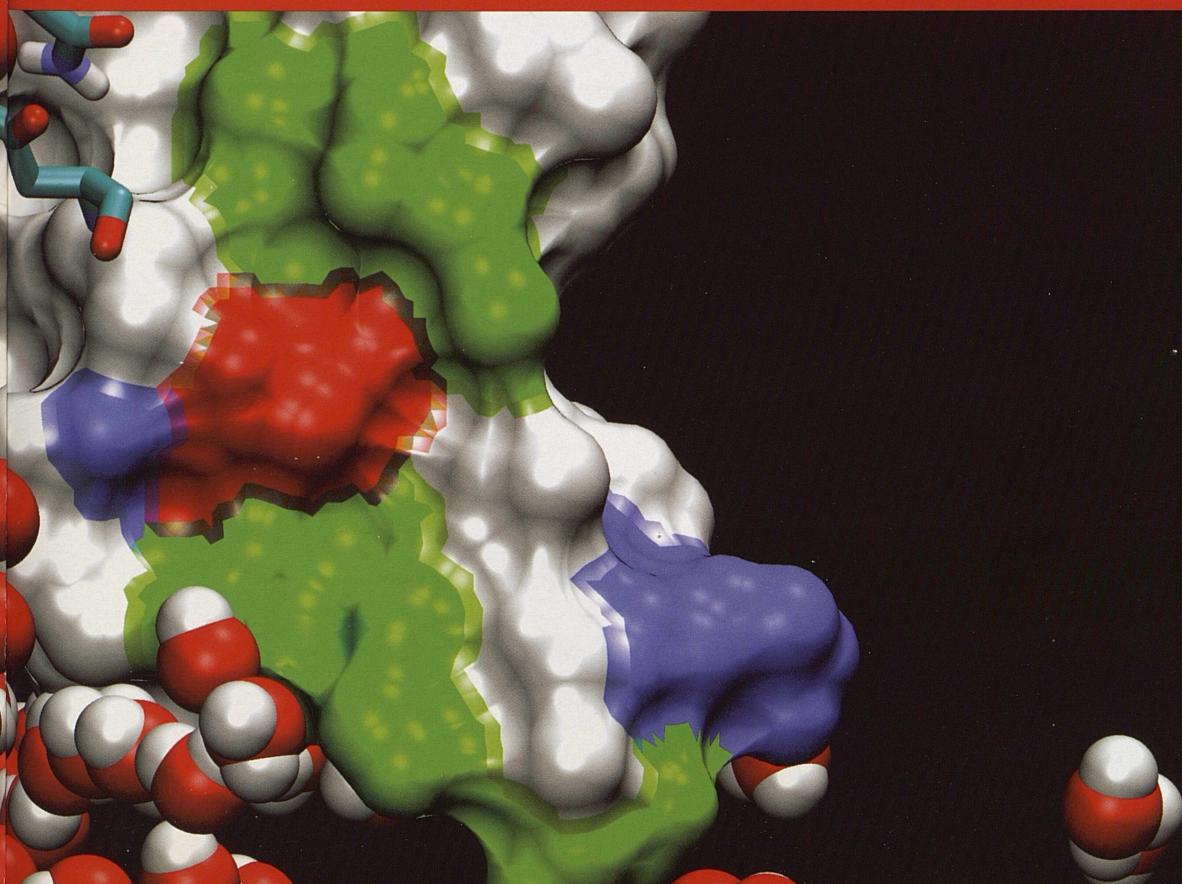
PITTSGURGH _ SUPERCOMPUTING _ CENTER / 2003 →





The Pittsburgh Supercomputing Center provides university, government, and industrial researchers with access to the most powerful academically-based resource in the United States for high-performance computing, communications and data-handling. 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. WWW.PSC.EDU TEL: 412-268-4960

PSC_2003



FOREWORD_FROM_THE_DIRECTORS



Michael Levine (right) and Ralph Roskies,
scientific directors, Pittsburgh Supercomputing Center

led the way in a rapidly emerging field, photonics, that holds immense promise for next-generation materials science and also high-performance computing and communications technology. Access to LeMieux has been crucial to their progress. LeMieux has also made possible important new findings in planet formation (p. 34) and solar physics (p. 38), and there's important work in DNA-enzyme interactions (p. 26) and new understanding of ventricular fibrillation (p. 30), a condition that's often fatal.

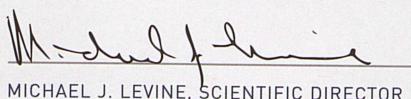
We're pleased that some of the earliest work on LeMieux, last year's simulations by Klaus Schulten and colleagues at the University of Illinois, Urbana-Champaign (p. 47), contributed to the pioneering research on membrane channels by Peter Agre and Roderick MacKinnon that received the 2003 Nobel Prize in Chemistry.

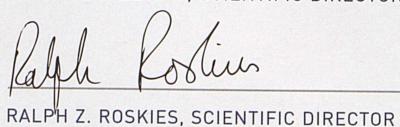
Complementing the computation, visualization and information manipulation capabilities of LeMieux, is a new, enormously more powerful, multi-petabyte, shared data storage and access facility replacing PSC's venerable but outdated data system.

As a prototype of large-scale cyberinfrastructure, the TeraGrid (see p. 5) represents a vision of the future and a laboratory in which to explore realistically the possibilities of geographically distributed computing, storage and visualization linked through high-speed networks of unprecedented bandwidth — soon to interconnect over a half-dozen of the nation's most advanced research centers. Interoperability — the ability to accommodate heterogeneous hardware, operating systems and network protocols seamlessly — is one of the TeraGrid's great challenges. PSC is proud to coordinate the TeraGrid effort to meet this and other challenges.

The Japanese Earth Simulator has helped to bring the continued importance of high-performance computing within cyberinfrastructure into focus, and to spur a national re-assessment of the future of supercomputing. In many fields, a 100-fold boost in computing power would allow scientists to achieve new milestones in their work. At PSC, we devote our energies toward progress on all these fronts — supercomputing, networking, storage, interoperability and others — that contribute to the creation of a national cyberinfrastructure.

We gratefully acknowledge our support 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 J. LEVINE, SCIENTIFIC DIRECTOR


RALPH Z. ROSKIES, SCIENTIFIC DIRECTOR

Cyberinfrastructure is a relatively recent word, one that doesn't exactly roll off the tongue, but that denotes the opportunities and challenges that face us in harnessing for coordinated, distributed, productive work the full range of information and information technologies. Much as physical infrastructure — roads, bridges, power grids, telephone lines, water systems — enables modern society, cyberinfrastructure is the totality of inter-related technologies — computing, networking, storage, visualization, collaborative tools, instruments, data repositories and other components — that enables much of modern scientific research.

With a generation of amazing progress behind us, we are entering an age in which it is possible to envision a comprehensive cyberinfrastructure that will empower and revolutionize the conduct of research for decades to come. The recent National Science Foundation blue-ribbon study, the Atkins report, urges the federal government to seize the moment and invest generously now to push forward from the foundation of current resources — such as the TeraGrid, NSF's multi-year grid project, and LeMieux, NSF's terascale system at PSC.

This volume presents a collection of recent scientific and technical advances at the PSC and indicates the direction of future advances. LeMieux — in production less than two years — has already brought forth important gains in scientific understanding. The Quake Group at Carnegie Mellon (p. 18), for example, has achieved ten-fold improvement in simulations of earthquake ground motion. John Joannopoulos' team at MIT (p. 22) has



CONTENTS

Foreword from the Directors



02

PITTSBURGH SUPERCOMPUTING CENTER, 2003

Creating National Cyberinfrastructure: Marvels Arrive, LeMieux Joins the TeraGrid	→	04
Supercomputing in Pennsylvania	→	06
The Super Computing Science Consortium	→	08
Research Notes & Highlights	→	10

PROJECTS 2003: CONTENTS	→	15
--------------------------------	---	----

EARTHQUAKE_MODELING Big City Shakedown	→	18
---	---	----

ADVANCES_IN_OPTICS Guiding Light	→	22
-------------------------------------	---	----

STRUCTURE_OF_PROTEINS_&_DNA When DNA Flips Out	→	26
---	---	----

BLOOD_CIRCULATION_&_THE_HEART Hearts Gone Wild	→	30
---	---	----

EVOLUTION_& STRUCTURE_OF_THE_UNIVERSE How to Cook a Giant Planet	→	34
---	---	----

ADVANCES_IN_SOLAR_PHYSICS Seeing Spots	→	38
---	---	----

IN_PROGRESS Clean Energy From Coal / Protein Landscapes / Inner Earth Cosmic Structure / Better Networks / Protein Border Guards	→	42
--	---	----

Creating National Cyberinfrastructure: Marvels Arrive, LeMieux Joins the TeraGrid

NEW SYSTEMS COMPLEMENT LEMIEUX, THE MOST POWERFUL
ACADEMICALLY-BASED SYSTEM IN THE COUNTRY

Two new systems established residence in the PSC machine room this year. Processor for processor, these systems — Hewlett-Packard GS1280 AlphaServers — are even more powerful than LeMieux, PSC's workhorse 3,000 processor terascale system, which delivers much of the public-domain computational research in the country.

During development, HP referred to the GS1280 as the Marvel, and PSC's twin Marvels are among the first to roll out of HP production. Both house 128 of the newest generation of the powerful Alpha processor, the EV7. Both provide 512 gigabytes of shared memory — each soon to be configured as two 256-gigabyte, 64-processor units.

The only difference between these capable Marvel twins is that one of them, dubbed Jonas, is dedicated to NIH biomedical research, while the other, Rachel, supports NSF science and engineering. Together, they complement LeMieux with new capabilities for scientists and engineers nationwide.

Along with a "shared memory" architecture, both new systems have exceptional "memory bandwidth" — the speed at which data transfers between hardware memory and the processor.

Benchmark tests have demonstrated that the GS1280 memory bandwidth is five to ten times greater than comparable systems.

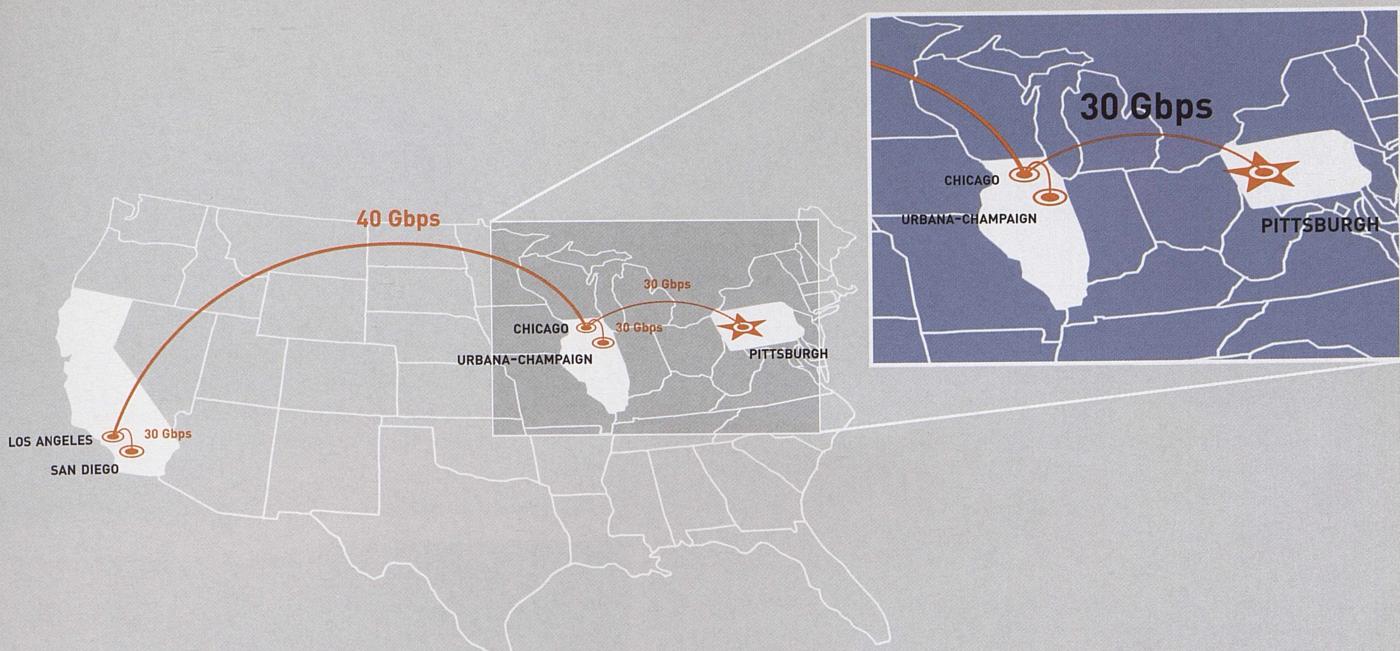
"Such a shared memory, high-performance system is not now available to biomedical researchers," said Marjorie A. Tingle, director of the Shared Instrumentation Grant Program at NIH's National Center for Research Resources, which awarded \$1.3 million to PSC for Jonas. "This new resource will enable them to study particularly demanding data-intensive, memory-intensive and compute-intensive problems that are currently beyond reach."

Funding for Rachel is part of a \$35 million NSF grant to PSC and its two sister supercomputing centers in October 2002. This grant supports continued development of the TeraGrid (see facing page) and provides for harnessing LeMieux and Rachel into this evolving national grid.

"The memory structure of this system represents a significant advance for scientific computing," said PSC scientific directors Michael Levine and Ralph Roskies in a joint statement. "It will make a difference in several important areas, including quantum chemistry and genomics."



Jonas Salk



PSC CONNECTS TO THE TERAGRID

The nation's most powerful, academically-based computing resource — LeMieux — is now linked with supercomputers in Illinois and California as part of the TeraGrid, an integrated national system of cyberinfrastructure. PSC network engineers in August implemented and tested a light pipeline that connects PSC to the Chicago hub of the TeraGrid's high-speed "backplane," which joins Chicago and Los Angeles.

Funded by NSF, the TeraGrid is a multi-year effort to deploy the world's fastest, most comprehensive distributed-computing infrastructure for open scientific research. Analogous to an electrical power grid, the TeraGrid will make computational power available to scientists and engineers nationwide as a seamless resource, without regard to physical location of computing systems.

Implemented in February 2003, the TeraGrid's Chicago-Los Angeles backplane moves data at 40 gigabits per second—600,000 times faster than typical Internet dial-up — the fastest network in the world. It is referred to as a backplane — a circuit board with plug-in

slots for other devices — because the TeraGrid is conceived of as a single national machine for science and engineering research, with powerful systems in different places that plug-in to the backplane. The Pittsburgh to Chicago link currently operates at 10 gigabits per second and will triple this fall.

LeMieux, different in both processor and operating system than existing TeraGrid systems, poses the challenge of "interoperability" — to create a grid environment integrating heterogeneous systems. The TeraGrid has established a PSC-led Interoperability Working Group to develop the software needed to meet this challenge.

"Many of us in many places are working together to create the TeraGrid," said Rick Stevens of Argonne National Laboratory and the University of Chicago, TeraGrid project director. "This work will empower U.S. research in science and engineering for years to come. The Pittsburgh to Chicago connection is a major step toward bringing this vision of integrated national cyberinfrastructure into reality."

PSC_2003

APPLICATION GATEWAYS

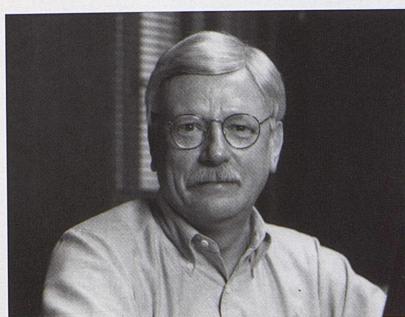
PSC network engineers have developed a unique and effective approach to link the TeraGrid backplane and LeMieux. The fiber-optic line — called a lambda — connects directly to a router, supplied by Cisco Systems, in the PSC machine room. From there, however, rather than configuring all 750 of LeMieux's AlphaServers with a direct connection, a much smaller number of machines (from 20 to 30) will bridge from the router to LeMieux. Called an "Application Gateway," this interface will run software, developed by PSC, that translates between the TeraGrid network and LeMieux's internal network.

WHAT'S IN A NAME?: MEET JONAS & RACHEL

Since Mario, the CRAY C90, from the early 1990s, PSC tradition has been to name its systems, and the newly arrived Marvels presented an opportunity. This time, PSC turned from sports to science. The twin Marvels are Rachel and Jonas, for Rachel Carson, a southwest Pennsylvanian who jump-started modern environmentalism with her 1962 book *Silent Spring*, and Jonas Salk, whose polio vaccine, developed at the University of Pittsburgh Medical School, in 1955 halted the spread of this crippling disease.

A LIGHT PIPELINE TIES THE EAST TO THE WEST OF NATIONAL CYBERINFRASTRUCTURE.

Jim Kasdorf, PSC director of special projects



Supercomputing in Pennsylvania

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

ECONOMIC DEVELOPMENT

PSC's high-performance computing and networking resources help to boost the competitiveness of Pennsylvania business and industry. During the past year, PSC focused its Industrial Affiliates efforts in presenting information on PSC training and consultation with several start-up companies.

In December 2002, Pittsburgh-based PPG Industries became a member of the PSC Industrial Affiliates Program. PPG's Scientific Computing Group uses LeMieux, PSC's terascale system, and the quantum-chemistry program GAUSSIAN for computational modeling in several aspects of its product lines as a global supplier of coatings, glass, fiber glass and chemicals. "Our approach is to use modeling to reduce discovery and development cycle times," says Michael Makowski, who leads PPG's computational chemistry group, "and to gain fundamental understanding of our core technologies and competencies, gain a competitive advantage, and ultimately reduce costs."



PSC outreach includes presence at numerous conferences and science fairs. At this year's BIO, the annual national convention for the life sciences industry, PSC shared a booth with the University of Pittsburgh and Carnegie Mellon University and joined over 100 representatives of Pennsylvania government along with industrial and academic organizations in representing statewide biotechnology.



Beverly Clayton,
PSC executive director,
coordinates PSC's program
for Pennsylvania researchers.



TEACHER TECH '03

From August 4-8, PSC sponsored a weeklong workshop for Pittsburgh Public Schools science teachers. Presented in collaboration with Rice University's Center for Equity and Excellence in Education and the National Science Foundation, the workshop introduced new technologies for helping to teach science and raised awareness about teacher roles in shaping and encouraging the next generation of scientists.

Support for TeacherTECH 03 came from the National Science Foundation, the Grable Foundation and the Henry C. Frick Educational Fund of the Buhl Foundation.

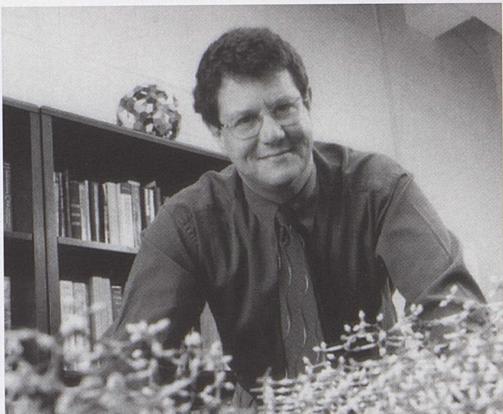
Three Pittsburgh Public Schools science teachers are using a Calculator Based Laboratory, a hand-held, data-collection system that works with various sensor technologies. Here, with a voltage sensor, the CBL is measuring battery voltage.

RESEARCH IN PENNSYLVANIA

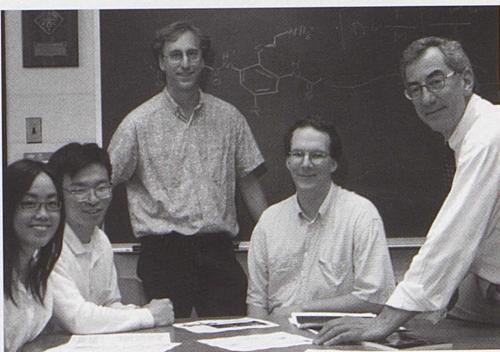
By supporting Pennsylvania university researchers, PSC resources help to attract research funds to the state. During the past year, more than 700 Pennsylvania researchers from 14 institutions used nearly a million CRAY T3E processor hours through PSC's Pennsylvania program. In addition, Pennsylvania campus researchers used 4.5 million processor hours on the Terascale Computing System through the National Science Foundation allocation process. The projects represented here (see also p. 18), along with many other research projects by Pennsylvania researchers using PSC resources, exemplify how supercomputing plays a role in scientific and engineering research in Pennsylvania.



This graphic from Madura's simulations shows the structure of RT (colored ribbon) bound with nevirapine (gray) in comparison to simulations of RT (superposed colored lines) without nevirapine.



Jeffry Madura, Duquesne University,
chair of chemistry and biochemistry



Michael Klein (right) and colleagues
at the University of Pennsylvania

AIDS & MOLECULAR ARTHRITIS

Jeffry Madura of Duquesne University in Pittsburgh has employed PSC systems in several research projects. In recent work, he and graduate student Zhigang Zhou used LeMieux to simulate a biomolecule that is a target for drug treatment of AIDS.

While drug therapies have helped enormously in treating AIDS, their success has been limited to managing the disease. AIDS remains among the most serious health threats around the world, especially in Africa and other developing countries, and the research goal remains to find a knockout punch. One of the best prospects is drugs that bind with and inhibit the functioning of an enzyme known as HIV-1 reverse transcriptase (RT). RT's role is to "transcribe" HIV's genetic code from single-stranded RNA into double-stranded DNA, a necessary step in HIV's ability to live and reproduce within a host organism.

Molecular dynamics simulations track the detailed atom-by-atom movements of a molecule, and in this case, Madura's interest was to track how a drug called nevirapine affects RT motion. The RT structure resembles a hand, the fingers and thumb of which normally move back-and-forth in relation to each other with a wide range of motion. The simulations showed, however, that when nevirapine binds with RT, the result is "molecular arthritis." The fingers and thumb become much more static, which helps to explain how nevirapine slows the effect of the virus. "This work suggests," says Madura, "that a good inhibitor should be one that can efficiently enter and maximally occupy the binding pocket."

MORE INFORMATION: <http://www.psc.edu/science/madura.html>

NEW WEAPONS FOR THE GERM WARS

Research by Professor Michael Klein and colleagues at the University of Pennsylvania was a finalist in the 2003 Computerworld Honors in Science, an international awards program recognizing innovations in information technology. Klein's research exemplified the scientific gains made possible by development and implementation of LeMieux, PSC's terascale system.

Klein's group modeled a new compound — an anti-microbial polymer — that offers great potential to create germ-resistant materials, such as bandages that kill bacteria, antiseptic operating tables, surgical gowns and many others. With hospital-acquired infection as the fourth-leading cause of death in the United States, these materials could save countless lives. Based on the success of this work, the University of Pennsylvania filed for several patents and created a company, POLYMEDIX, to exploit the possibilities for useful applications.

MORE INFORMATION: <http://www.psc.edu/science/klein2002.html>

The Super Computing Science Consortium

PENNSYLVANIA-WEST VIRGINIA PARTNERS
IN DEVELOPMENT OF CLEAN POWER TECHNOLOGIES

Formed in 1999, 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.

During the past year, Waynesburg College, a small liberal arts college in Greene County, Pennsylvania, between Pittsburgh and Morgantown, became a new (SC)² partner.

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

TECHNOLOGIES FOR CLEAN, AFFORDABLE POWER

Since the spring of 2000, a high-speed network — which provided 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 West Virginia University have used this high-speed channel to tap the computational resources at PSC for a range of projects involving state-of-the-art simulation technology. This work includes:

- 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
- industrial-scale technology for coal gasification,
<http://www.psc.edu/publicinfo/2002/sc2/>
- a new design for a power-generating turbine.
<http://www.psc.edu/science/cizmas2002.html>

In recent work (p. 42), NETL researchers have used LeMieux, PSC's terascale system, to make significant new progress toward a realistic simulation technology for industrial-scale coal gasification.

WORKFORCE DEVELOPMENT & EDUCATION

Economic development for the Pennsylvania and West Virginia region was the central topic of the (SC)² Seminar 2003: Regional Business Opportunities in Supercomputing. This July 14 event at Waynesburg College brought together leaders in technology, business, higher education and government. Participants included over 130 people representing 31 companies, 15 colleges and universities, seven economic-development agencies and six government agencies.

"With the most advanced technological resources, as well as through partnerships with local technology leaders, our Commonwealth's small businesses will be able to face current challenges,

achieve their long-term goals, and set a positive course for the future," said Pennsylvania governor Edward G. Rendell in a letter commending the (SC)² event.

In another event, held June 23-27, ten middle and high-school teachers from Pennsylvania and West Virginia participated in the first annual (SC)² parallel-computing workshop. Developed and presented by Duquesne University, Penn Hills School District, Fluent Corporation and NETL, the workshop featured presentations on topics that included parallel versus sequential processing and scientific visualization as well as a hands-on practicum.



Twenty-seven students from 11 regional colleges presented posters on a range of topics related to technology and computational science.



The (SC)² Seminar 2003 provided many opportunities for discussion. Thomas McGraw (left) of the Pennsylvania Department of Community and Economic Development talks with Bob Romanosky, NETL manager of Power Systems Advanced Research, and Lynn Layman (right), (SC)² co-director.



(SC)² PARTNERS

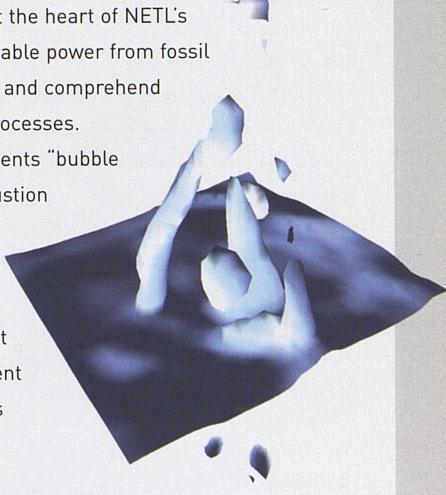
Carnegie Mellon University
Duquesne University
The Institute for Scientific Research
The National Energy
Technology Laboratory
The Pittsburgh Supercomputing Center
The Pennsylvania State University
The University of Pittsburgh
Waynesburg College
West Virginia University
The West Virginia Governor's Office
of Technology

MAVIS: SEEING THE INNER FLAME

During the past year, PSC scientific-visualization specialist Kent Eschenberg collaborated with NETL researchers to develop software for interactive 3D visualization. The program, called Mavis, operates with NETL data sets produced by a powerful simulation technology called MFIX (Multiphase Flow with Interphase Exchanges). Developed over years of research, MFIX realistically models the complex processes — gas and particle dynamics, chemical reactions and heat transfer — involved in combustion.

Detailed understanding of combustion lies at the heart of NETL's research objective: environmentally clean, affordable power from fossil fuel. Mavis provides a quick, effective way to see and comprehend the results from computational study of these processes.

This frame from a Mavis visualization represents "bubble coalescence" in a fluidized-bed reactor, a combustion process in which air is injected from below at high speed into a bed of particles. The particles collectively behave like fluid and, among other fluid-like effects, form bubbles, which affect heat distribution. In this simulation, two jets of different velocity are injected into the bed. Larger bubbles formed above the faster jet (left) attract bubbles formed from the slower jet.



Research Notes & Highlights

NETWORKING THE FUTURE

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

PSC's network-engineering group, the National Center for Network Engineering (NCNE), is one of the leading resources in the world for network know-how. They provide engineering consulting for advanced networking nationally, and they conduct seminars that disseminate knowledge to engineers around the country. Since 1998, when an NSF grant established NCNE, their training activities have reached more than 3,000 people.

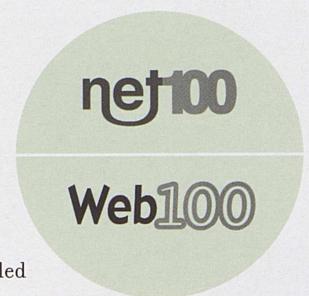
MORE INFORMATION: <http://www.ncne.org>

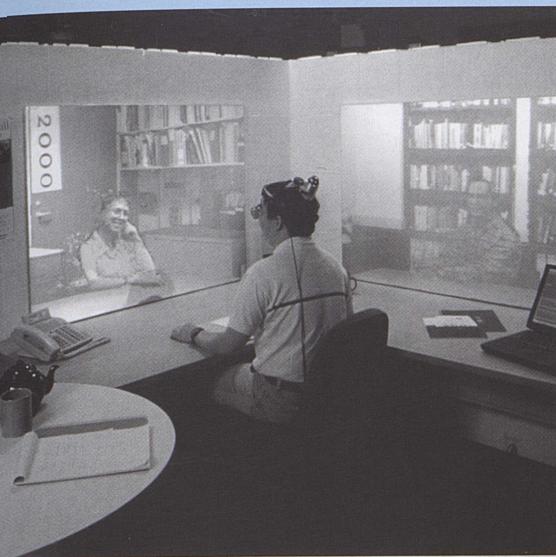
GETTING IN TUNE WITH WEB100 & NET100

With funding from the National Science Foundation and Cisco Systems, PSC network researchers — in collaboration with the National Center for Atmospheric Research (NCAR) and the National Center for Supercomputing Applications — have developed software to "tune" computer operating systems to better exploit available network bandwidth. Called Web100, this software is now released publicly and used in many research projects nationwide.

Continuing PSC work on Web100 during the past year has added "autotuning" to the publicly released software. Collaboration with developers of the Linux operating system will lead to incorporating Web100 features in Linux. Also as part of Web100, PSC engineers have developed and documented a set of tools, called Pathprobe, to analyze and characterize a network path.

In a related project called Net100, funded by the U.S. Department of Energy, PSC's network research group is collaborating with NCAR, Lawrence Berkeley National Laboratory and Oak Ridge National Laboratory to create software tools that allow operating systems to tune themselves in response to changing traffic conditions on the network. PSC developed a real-time graphical Net100 demonstration for the Supercomputing 2002 conference and has incorporated selected Net100 enhancements into Web100.





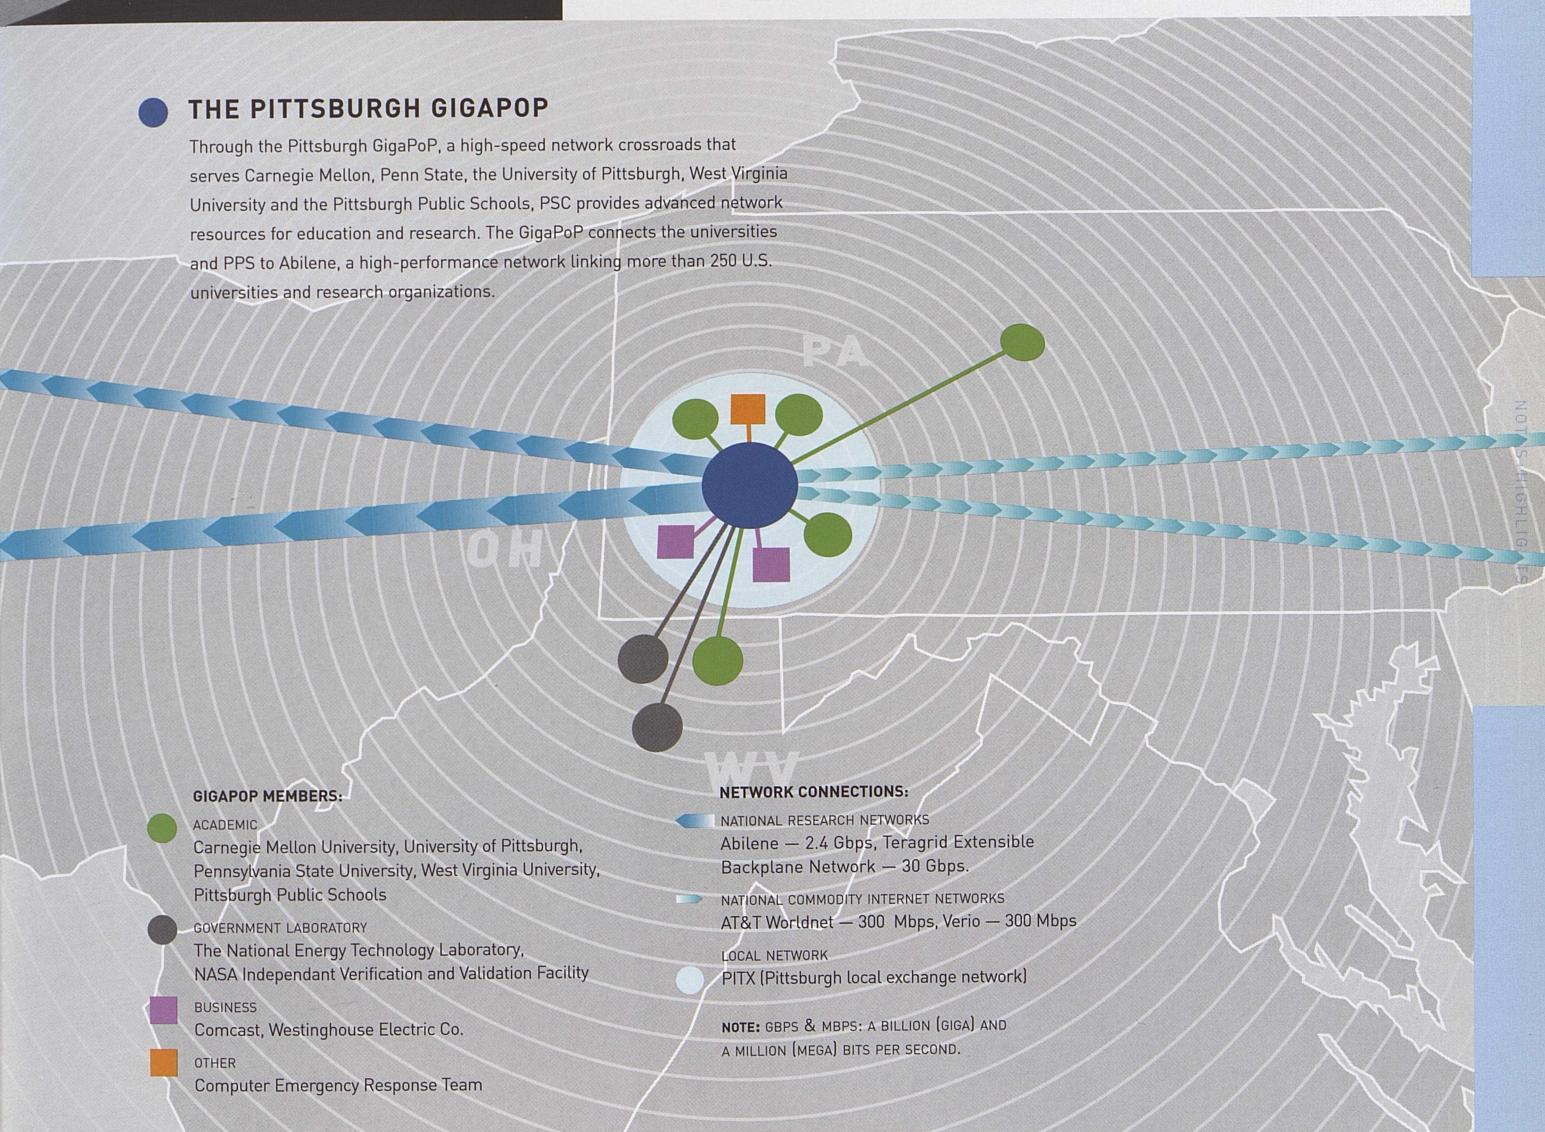
HUMAN INTERACTION AT A DISTANCE

Computer vision and graphics systems that present a movie-like illusion of being in the same physical space as others who are hundreds or thousands of miles away — this developing technology is called tele-immersion. Similarly to the telephone a hundred years ago, tele-immersion holds potential to transform human interaction.

With support from the National Science Foundation Information Technology Research program, PSC scientists are collaborating with scientists at the University of Pennsylvania and the University of North Carolina, Chapel Hill to develop a prototype tele-immersion system. The real-time emphasis of tele-immersion presents a novel challenge for computing and requires the most advanced networking. During the past year, the PSC team, in collaboration with Penn and UNC researchers, carried out tele-immersion experiments using over 1,500 LeMieux processors simultaneously, achieving more realistic video rates and higher-quality images than previously possible, a step toward what will ultimately — in a decade or so — become conventional technology.

THE PITTSBURGH GIGAPOP

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



Pittsburgh Public School Students

GIGAPOP NEWS

Upgrades to hardware and fiber-optic cable in December 2002 boosted bandwidth between LeMieux and Abilene tenfold, dramatically improving access to LeMieux for scientists and engineers nationwide. "These upgrades," said NCNE director Gwendolyn Huntoon, "allow LeMieux to support applications with much higher distributed data requirements, including large single-stream transfers or large transfers from multiple users."

In August 2003, the Pittsburgh Public Schools became a new member of the Pittsburgh GigaPoP. One of the first Pennsylvania school districts with a connection to Abilene, PPS now has the bandwidth necessary to take advantage of initiatives in distance learning and specialized courses delivered via videoconferencing.

BIOMEDICAL SUPERCOMPUTING

NATIONAL LEADERSHIP IN COMPUTATIONAL RESOURCES FOR BIOMEDICAL RESEARCH

In 1987, the PSC biomedical program became the first extramural biomedical supercomputing program in the country funded by NIH. Since then, with support from NIH's National Center for Research Resources (NCRR), PSC has fostered exchange between PSC expertise in computational science and experts in biology and medicine to solve important problems in the life sciences.

PSC workshops and courses on computational biology have trained more than 2,500 researchers in the use of high-performance computing for biomedical research, in such areas as sequence analysis in genome research, the structure of proteins and DNA, and biological fluid dynamics.

Since its inception, PSC's biomedical program has provided computing resources for more than 900 biomedical research projects involving thousands of researchers in 43 states and the District of Columbia. Among these are several projects featured in this booklet (pp. 26-33, 43 & 47).

In addition to training and access to computational resources, the biomedical group carries out research in structural biology, protein and nucleic-acid sequence analysis, computational neuroscience and microphysiology and biomedical visualization.

MORE INFORMATION: <http://www.psc.edu/biomed>

PSC BIOMEDICAL COLLABORATIONS

Albert Einstein College of Medicine
Carnegie Mellon University
Duke University
Scripps Research Institute
University of California at San Francisco

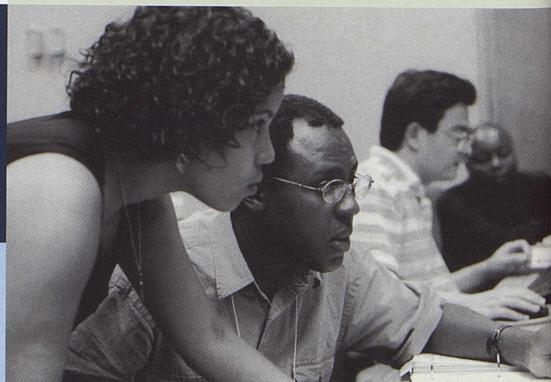
University of Michigan
University of Pittsburgh Medical School
University of North Carolina at Chapel Hill
Whitehead Institute



DAVID DEERFIELD ↗

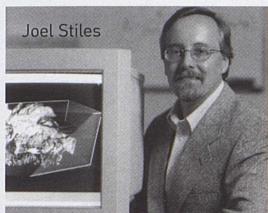
David Deerfield (above right), director of PSC's biomedical initiative, explains a concept to participants in this year's Nucleic Acid and Protein Sequence Analysis, at PSC from Aug. 10-15. PSC also presents this popular workshop, supported by the NIH National Human Genome Research Initiative, on-site at universities and research centers.

"Our training reaches hundreds of biomedical scientists each year," says Deerfield. "Techniques we've developed are helping scientists nationwide cope with the explosion of genome data."



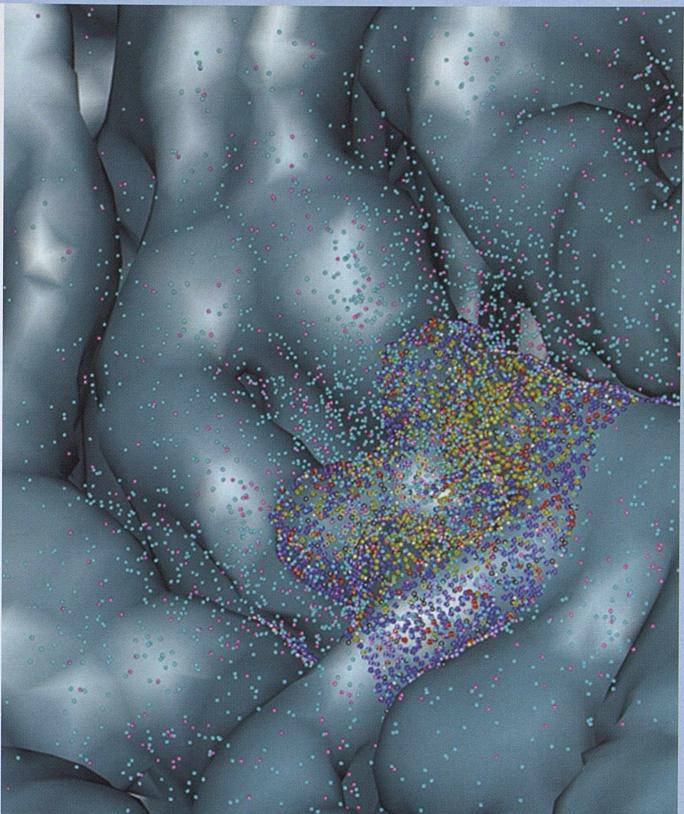
THE CENTER FOR QUANTITATIVE BIOLOGICAL SIMULATION

In recognition of his outstanding research in computational microphysiology, PSC scientist Joel Stiles in July was appointed to the faculty of Carnegie Mellon University as associate professor in the Mellon College of Science. Stiles is co-developer of MCell and DReAMM, software used in many laboratories around the world to simulate microcellular physiology.



Along with contributing his expertise to the PSC Biomedical Initiative, Stiles will lead a new research group within PSC, the Center for Quantitative Biological Simulation. With support from multiple sources at NIH and NSF,

the new center will integrate collaborations at Carnegie Mellon and the University of Pittsburgh. "The goal," says Stiles, "is to dramatically extend the range of biomedical problems amenable to quantitative simulation, a key step toward detailed understanding and targeted interventions in today's world of complex multifactorial health issues."



THE PSC VISIBLE HUMAN BROWSER

Released in October 2002, PVB is an interactive graphical interface and rapid-retrieval system. When used with the National Library of Medicine's Visible Human, it provides medical students with a versatile, self-teaching resource in human anatomy. For doctors, it's a tool to plan surgery and radiation therapy.

Developed over a three-year collaboration between PSC and the University of Michigan, PVB was designed to display the Visible Human, a 3D digital reconstruction of human anatomy, and can also be used with other datasets. Most viewing technologies limit the Visible Human to right-angle plane views and lack the ability to manipulate large-volume data at reasonable speed. With funding from NLM and also from NCRR, the PSC and UM researchers overcame these limitations.



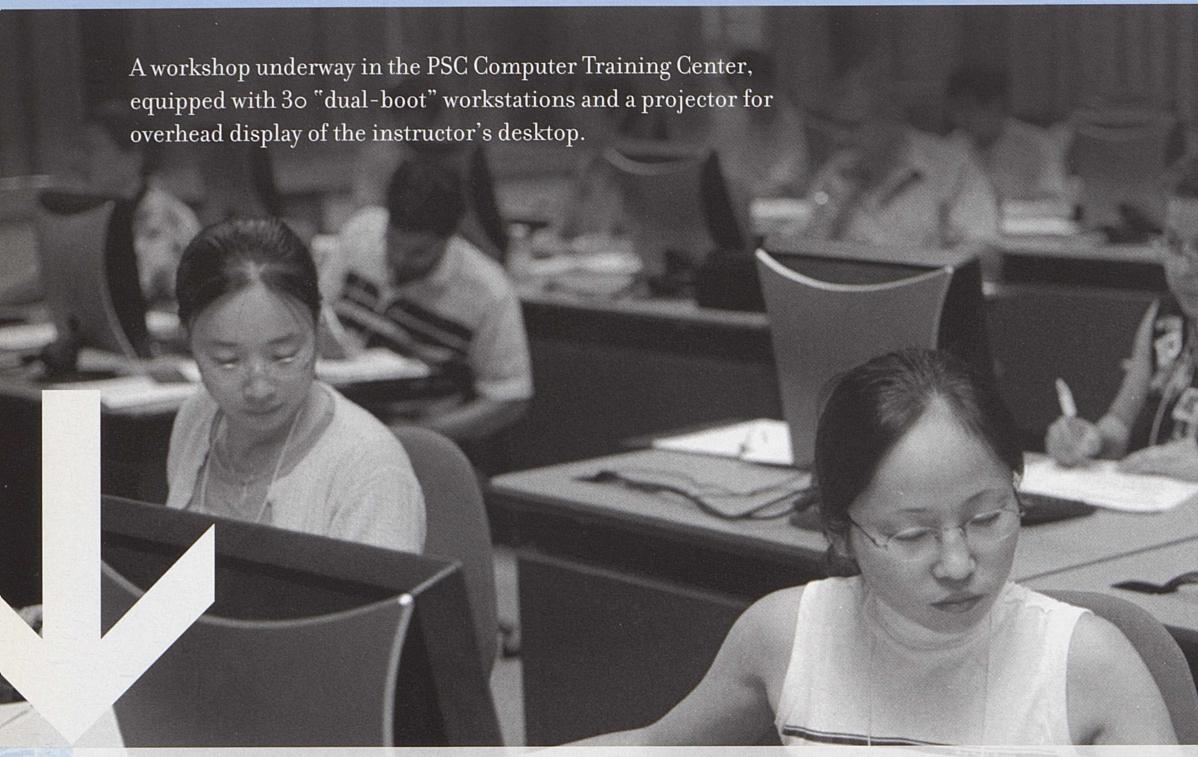
To gain speed, PVB employs innovative data-compression that, for the most compressible data, achieves a ratio greater than 30:1. Response is also enhanced with a network-delivery system that allows up to 40 people to navigate independently through the data in near real-time. The compressed and restructured dataset is served from a high-speed network server at PSC.

"PVB is a highly interactive system," says PSC scientist Arthur Wetzel, who leads the PVB team. "It allows users to navigate easily through the Visible Human or other volumetric data and to select 2D slices for viewing with complete freedom — as suggested by the orientation of an object, not limited to top, front and side views. The design of PVB allows this to happen quickly enough to serve a large number of simultaneously navigating users."

MORE INFORMATION: <http://www.psc.edu/biomed/research/VH/PVB>

WITH THIS NEW TECHNOLOGY, STUDENTS AND SURGEONS CAN VIEW HUMAN ANATOMY RAPIDLY AND INTERACTIVELY FROM ANY ANGLE.

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



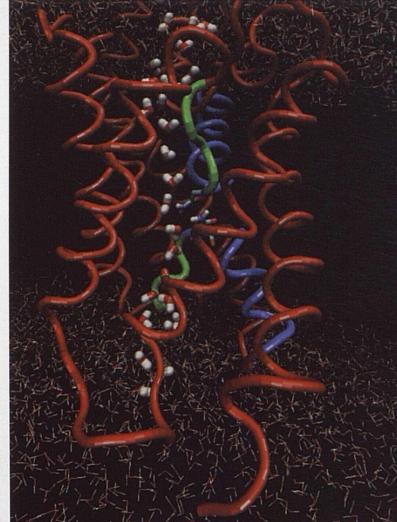
PITTSBURGH SUPERCOMPUTING CENTER WORKSHOPS (2002-2003)

- Developing Bioinformatics Programs
- MCell/DReAMM Modeling for Drosophila Research
- From Structure to Function: Frontiers of Biological Ion Channels
- Parallel Programming Techniques
- Nucleic Acid and Protein Sequence Analysis

HIGH-PERFORMANCE MOLECULAR DYNAMICS WITH LEMIEUX

Not just any molecular dynamics program — that's how NAMD got its name. Having sustained performance of more than a teraflop (a trillion calculations a second) on LeMieux, PSC's terascale system, NAMD has proved to be well christened.

Molecular dynamics, which tracks the atom-by-atom movements of proteins, is one of the most widely used applications in computational biology. NAMD, developed at the University of Illinois, Urbana-Champaign, under the direction of Klaus Schulten, uses modern computer-science techniques to effectively exploit large-scale parallel systems. The UIUC team used all 750 of LeMieux's nodes — each node has four processors — to achieve the teraflop mark, which for molecular dynamics sets a record.



THE PSC OPERATIONAL MANAGEMENT TEAM

(front, l to r) Elvira Prologo, manager, administrative staff; Sergiu Sanelevici, assistant director, scientific applications and user support; Bob Stock, associate director; (2nd row) Rich Raymond, manager, user support; Gwendolyn Huntoon, assistant director, networking; J. Ray Scott, assistant director, systems and operations; (back row) Jim Marsteller, manager, Pittsburgh GigaPop; David Kapcin, manager, financial affairs; David Deerfield, assistant director, biomedical initiative.



PROJECTS_IN_SCIENTIFIC_COMPUTING_2003



PROJECTS 2003 CONTENTS

EARTHQUAKE_MODELING		
Big City Shakedown	→	18
High-Resolution Forward and Inverse Earthquake Modeling		
Jacobo Bielak, Omar Ghattas, David O'Hallaron, <i>Carnegie Mellon University</i>		
ADVANCES_IN_OPTICS		
Guiding Light	→	22
Polarization-Independent Linear Waveguides in 3D Photonic Crystals		
E. Lidorikis, M. L. Povinelli, S. G. Johnson, & J.D. Joannopoulos, <i>Massachusetts Institute of Technology</i>		
STRUCTURE_OF_PROTEINS_&_DNA		
When DNA Flips Out	→	26
Protein-facilitated Base Flipping in DNA by Cytosine-5-Methyltransferase		
Niu Huang, Nilesh K. Banavali & Alexander D. MacKerell, Jr., <i>University of Maryland</i>		
BLOOD_CIRCULATION_&_THE_HEART		
Hearts Gone Wild	→	30
Suppression of Alternans Despite Steep Restitution		
Flavio Fenton & Elizabeth Cherry, <i>Beth Israel Medical Center, New York City</i>		
EVOLUTION_&_STRUCTURE_OF_THE_UNIVERSE		
How to Cook a Giant Planet	→	34
Formation of Giant Planets by Fragmentation of Protoplanetary Disks		
Thomas Quinn, <i>University of Washington</i>		
ADVANCES_IN_SOLAR_PHYSICS		
Seeing Spots	→	38
Simulations of Near-Surface Solar Magnetoconvection		
Juri Toomre, <i>University of Colorado, Boulder</i>		
Marc DeRosa, Neal Hurlburt, <i>Stanford-Lockheed Institute for Space Research</i>		
IN_PROGRESS		
Clean Energy from Coal	→	42
Protein Landscapes	→	43
Inner Earth	→	44
Cosmic Structure	→	45
Better Networks	→	46
Protein Border Guards	→	47

EARTHQUAKE_MODELING

PROJECTS 2003

1

BIG CITY SHAKE DOWN

INSIDE LEMIEUX, THE 1994 NORTHRIDGE EARTHQUAKE SHAKES ALMOST LIKE THE REAL THING

People in greater Los Angeles won't soon forget January 17, 1994, 4:30 a.m. For eight seconds that may have seemed like eight hours, as if Mother Earth herself was restless in bed, nearly everyone across a 2,500 square-mile area jerked awake, adrenaline pumping.

Striking the densely populated San Fernando Valley of northern LA, the Northridge earthquake was the second time in 60 years that the earth ruptured directly beneath a major U.S. urban area. By the time the dust settled and officials counted the toll, 57 people were dead, more than 1,500 seriously injured. Collapsed freeways choked traffic for days. Over 12,000 buildings and 170 bridges sustained moderate to severe damage. Total economic loss was estimated at \$20 billion.



Photo courtesy of Federal Emergency Management Agency (FEMA)

Except for building codes that require earthquake-resistant structures, fatalities and damage would have been much worse. Still, the loss was enormous, and one of the main lessons of Northridge, as well as other urban earthquakes of recent decades, is the need for better information about where and how much the ground will shake.

"We've learned that the severity of ground shaking and consequent damage patterns vary significantly within relatively small areas," says Jacobo Bielak, professor of civil and environmental engineering at Carnegie Mellon University. "Even from one block to the next, the level of shaking can change dramatically due to types of subsurface soil and rock and other geological characteristics and the nature of the seismic waves."

Bielak and his Carnegie Mellon colleagues Omar Ghattas and David O'Hallaron lead the Quake Group, a large collaborative research team. Using sophisticated computational methods, they work to create realistic 3D models of earthquakes in geologically complex basins. Their objective is to provide accurate forecasts of earthquake ground motion as a necessary step toward creating building codes that provide for the safest possible structures at reasonable cost.

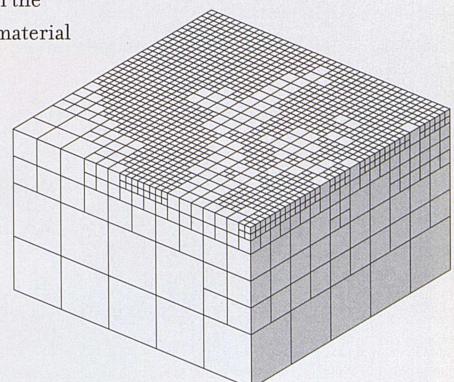
They have used LeMieux, PSC's terascale system, to great advantage, taking big steps forward in their work. "We've benefited enormously from having this powerful system at PSC," says Ghattas, "and we've developed algorithms that maximize our ability to use it well."

Using as many as all 3,000 LeMieux processors at one time with high efficiency, they have carried out the most detailed, accurate simulations yet of the Northridge quake — at twice the frequency of prior models. They've also made major inroads on an important problem called "the inverse problem," the goal of which — magical as it may seem — is to determine subsurface geology by working backward from seismic measurements on the surface.

WAVELLENGTH TAILORING

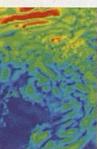
Ancient advice says it's better to build your house on rock than sand. If you live in an earthquake basin, that ancient advice, generally speaking, still holds.

From soft soils near the surface to hard rock deeper down and in the mountains, subsurface material varies tremendously in stiffness, the property that dictates seismic wavelength. For a given frequency, the softer the material, the



WAVELLENGTH-TAILORED MESH

This closeup shows how grid spacing is smaller in regions of slower seismic waves. The width of the smallest elements is about 50 meters.



1

shorter the length of the seismic waves, which means finer resolution, more computing, to accurately model the shaking.

"We've found," says Bielak, "that even within a few hundred meters, the variability in soil properties — and therefore ground motion — can be very substantial. Because of this, similar buildings located near each other can experience significantly different levels of damage."

To accurately capture this wide range of ground vibration in a large earthquake basin like Los Angeles poses enormous challenges for earthquake modeling. One of the Quake Group's key strategies has been to tailor their computational mesh — which divides the basin into millions of subvolumes — to soil stiffness.

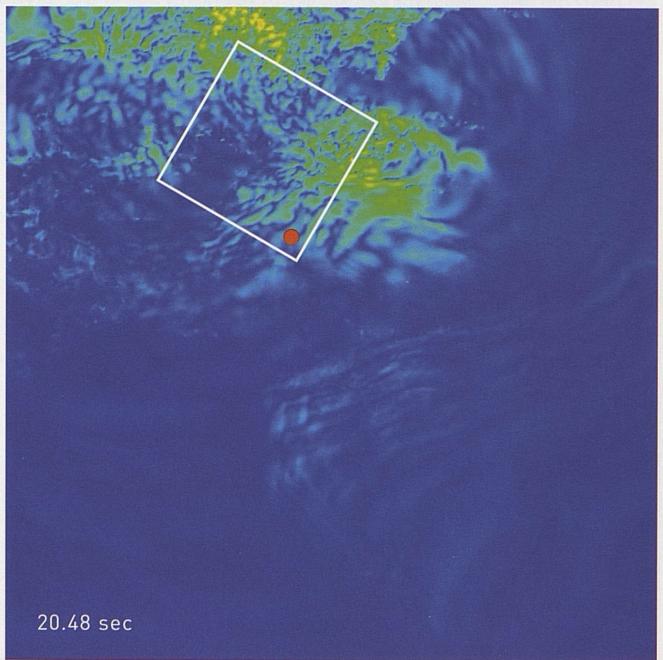
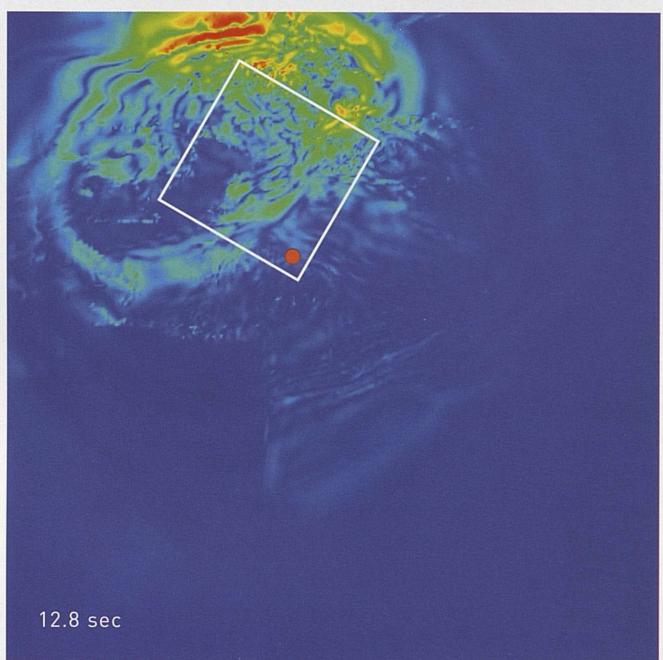
They generate their LA Basin computational model from a geological model created at the Southern California Earthquake Center (SCEC). Where the SCEC model indicates softer soils, therefore shorter wavelengths, the mesh-generating software creates a denser mesh.

"By using disk space instead of computer memory," says O'Hallaron, "our out-of-core algorithms can generate an extremely large mesh." For the recent simulations, they represented the basin as 80 kilometers on each side by 30 kilometers deep. Within this volume, their irregular mesh maps more than 100 million subvolumes, making their computations with LeMieux the largest unstructured mesh simulations ever done.

"These are the most highly resolved LA Basin earthquake simulations to date," says Ghattas, "and they are made possible by our adaptive meshes and their low memory requirements. To achieve similar accuracy with a uniform mesh would require 1,000 times more computing power."

SNAPSHOTS OF PROPAGATING WAVES

These images from the simulation show seismic waves propagating on the LA Basin surface, outward from the epicenter (red dot) and fault plane (rectangle) at three successive times, with color (blue to red) corresponding to frequency of vibration (0 to 1 Hz). "The direction of the ground motion," says Bielak, "and the concentration of motion near the fault corners follows a pattern observed during the actual earthquake."



REALISTIC FREQUENCIES & THE INVERSE PROBLEM

Using 2,048 LeMieux processors, they simulated the Northridge quake for more than 30 seconds of shaking. Their "wave-propagation" software sustained exceptional performance — nearly a teraflop (a trillion calculations a second) over six hours of computing time. And it ran at nearly 90 percent parallel efficiency, a measure of how well the software uses many processors at the same time.

Ground-motion frequency is a key factor for building design, and this simulation accounted for shaking up to one vibration cycle per second (1 Hz) — double the previous high (.5 Hz). Earthquake modeling has been limited in its ability to simulate higher frequencies, from 1 to 5 Hz, that present the greatest danger to "low-rise" structures — which include most city buildings, predominantly two-to-five stories — because each doubling of frequency means a 16-fold increase in computing.

"Our challenge is to attain realistic frequencies," says Bielak. "Now, for the first time, we're in the range that engineers need to know about. Typically, they want to see results up to 4 Hz, which points to the need for more computational power."

The simulation reproduced ground motion of the Northridge quake more accurately than possible until now, but — not surprisingly — at some locations it failed to reproduce significant shaking. These discrepancies, notes Ghattas, are inevitable considering that the geological model is inherently incomplete. "Because of uncertainties in what we know about earthquake source and basin material properties, a critical challenge facing us is to obtain these properties by source inversion from records of past earthquakes."

USING LEMIEUX, THE QUAKE GROUP HAS PROVEN THE FEASIBILITY OF SOLVING THE INVERSE PROBLEM, ONE OF THE IMPORTANT CHALLENGES OF COMPUTATIONAL SCIENCE AND ENGINEERING.

This problem — the inverse problem — is one of the important challenges of computational science and engineering, with potential applications in many fields, and it is key to the Quake Group's goals. Ghattas and his former students Volkan Akcelik and George Biros won the best paper award last year at Supercomputing 2002 for their inverse wave-propagation algorithm that exploits parallel systems like LeMieux.

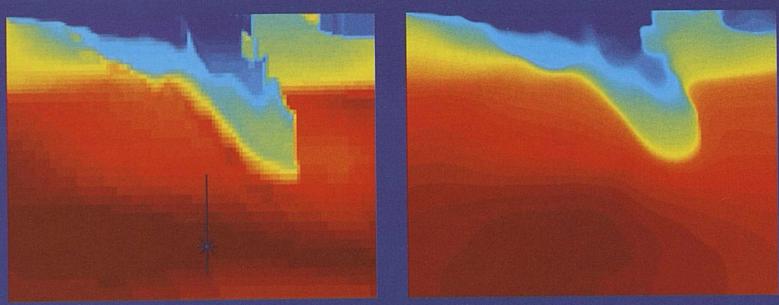
Using a sophisticated mathematical approach, their algorithm makes it possible to ascertain deep geological features based on seismic recordings on the surface. The deep geology is often not well understood and is known to play an important role in surface shaking. With LeMieux, for the first time, the Quake team solved a test case in two dimensions that proves the feasibility of this inverse approach.

"The inverse problem is orders of magnitude more difficult than the forward problem," says Ghattas. "Large parallel systems and powerful algorithms are crucial." One of the Quake Group's near-future plans for LeMieux is to further test their inverse approach with the added difficulty of three dimensions.

MORE INFORMATION: <http://www.psc.edu/science/earthquake.html>



[l to r] Omar Ghattas, Jacobo Bielak, David O'Hallaron



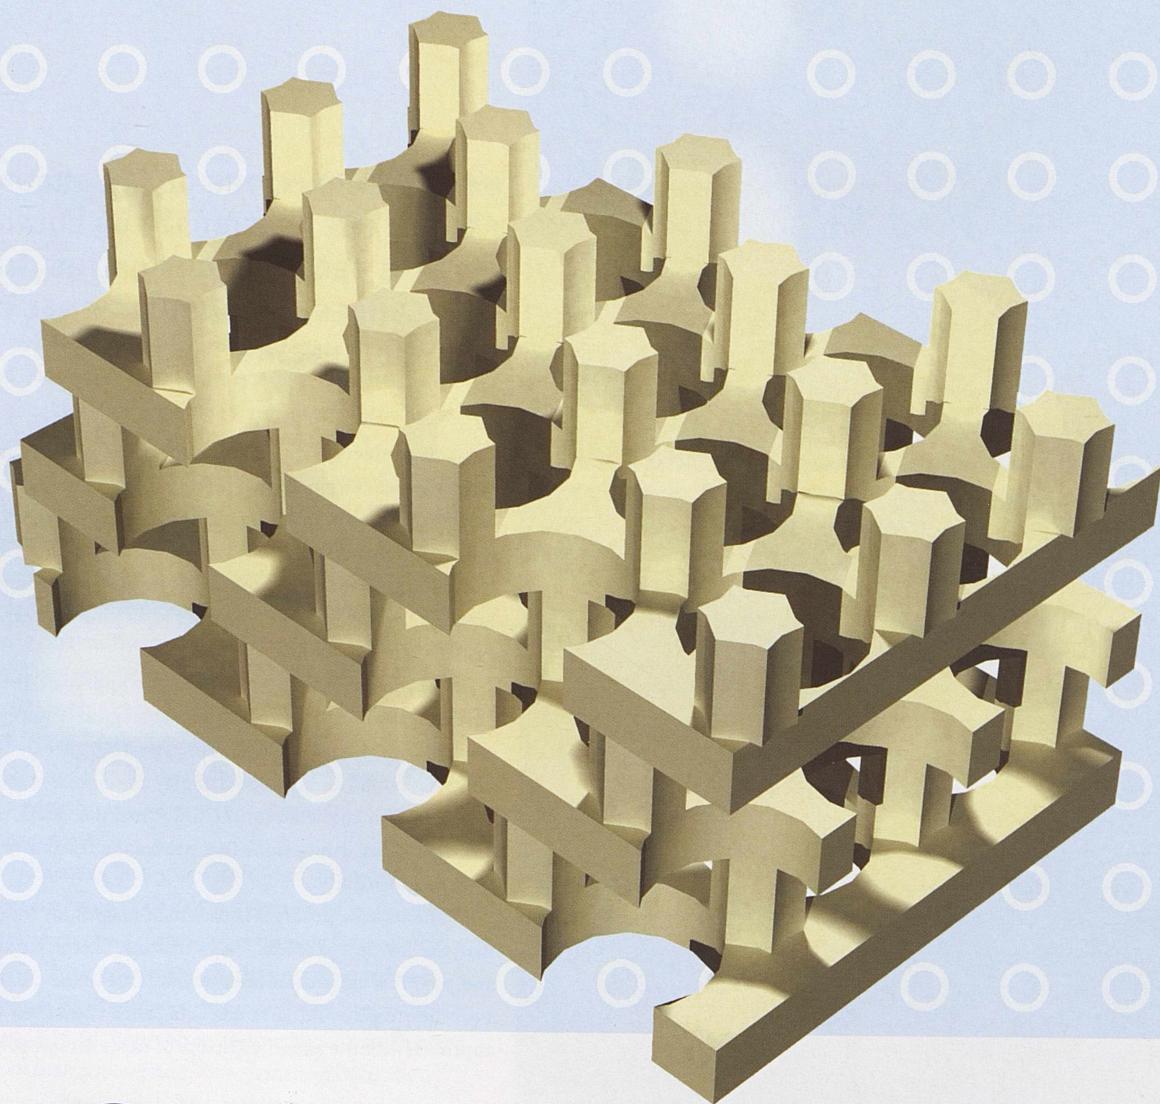
SOLVING THE INVERSE PROBLEM

As a first test for the extreme computational challenge of the inverse problem, the Quake Group chose a 2D shear-wave velocity distribution (left) in a 35 km x 20 km cross-section of the LA Basin model as a synthetic target. Starting with ground-motion measurements on the surface (64 points distributed evenly), the inverse algorithm (with a 257 x 257 grid) arrived at material properties for the cross-section that gave a velocity distribution (right) in close match to the target.

ADVANCES_IN_OPTICS

PROJECTS 2003

2



GUIDING

PHOTONICS MAY BE THE WAVE OF THE FUTURE IN COMPUTING AND COMMUNICATIONS, WITH BIG HELP FROM LEMIEUX

For MIT physicist John Joannopoulos, success means defects and imperfections. His research draws on the vocabulary of post-apocalypse fantasy, with explorations of "degenerate bands" and "forbidden zones." His findings, however, aren't fiction. His work on innovative materials called "photonic crystals" presages an impending technological revolution — computing and

communications at the speed of light. Tiny, honeycomb-like, crafted from layers of silicon, photonic crystals have unprecedented ability to trap, guide and control light. Their promise is to manipulate photons, the tiniest lumps of energy in light, with the same precision that semiconductors make possible with electrons. To move from electronics to photonics — using photons rather than electrons as the markers of digital 0s and 1s — is the objective.

← Designed by Joannopoulos' group, this photonic crystal is made from alternating 2D layers, one with dielectric (silicon) rods in air, the other with air holes in the dielectric.

The result will be networks that move data at trillions of bits per second, a thousand times faster than today. For computing, along with a radical leap in processing speed, photonics should shrink by a thousand fold the size and power needs of circuitry. It's a rapidly emerging technology, burgeoning with promise, with likely impacts in many fields including medical and chemical scanning and sensing devices.

Over the last decade, Joannopoulos and his MIT colleagues have done much of the work that has generated a buzz about this new field, with several patents to their credit. They've developed software designed especially to simulate, predict and explain the remarkable properties of these materials, and they work closely with MIT laboratory researchers to fabricate their designs.

When LeMieux, PSC's terascale system, became available as a production resource in 2002, Joannopoulos tapped into this massive new supply of computational power. "This machine is a fantastic creation," he says. "It has enabled us to do calculations that weren't possible before. We've made great progress on several projects."

With LeMieux, Joannopoulos' group last year simulated a photonics crystal phenomenon called "superlensing," demonstrating that it's feasible in three-dimensional crystals. They had previously demonstrated superlensing in two dimensions, but until LeMieux were stymied with the more challenging 3D computation.

They've also applied LeMieux to an important problem with polarized light. Photonic crystal "waveguides" are able to guide a stream of photons like a canal channels water, but to date they've been selective in

how they transmit light's two polarities, the horizontal and vertical component of the wave. In recent work, Joannopoulos and colleagues Elefterios Lidorikis, Michelle Povinelli and Steven G. Johnson appear to have solved this problem, and their new findings are a big step toward integrated photonic networks.

BRIGHT IDEAS

Fiber-optic cables have for years been the information-carrying medium of choice for high-performance networks, but fiber optics is inherently limited. Among other drawbacks, all the routers and switches that decode and sort the information are electronic; they convert photons to electrons and back again, creating bottlenecks.

Enter the photonic crystal, known also, more descriptively, as the photonic band-gap crystal. They're built from clusters of material arrayed in a regular pattern with holes between the clusters. Think of microscopic Swiss cheese with precisely spaced holes.

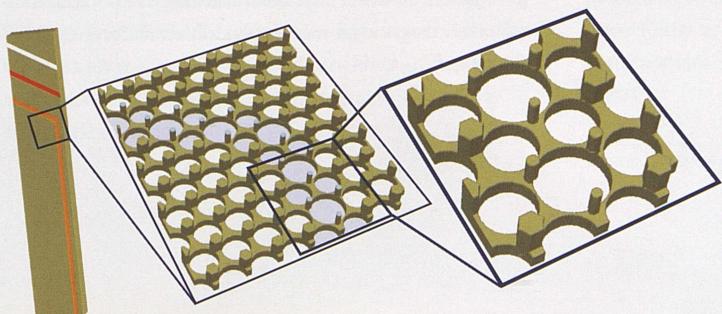
The material, often silicon, is of a different refractive index than the holes, usually air. When this geometry is designed just right, it creates an environment where certain wavelengths of light can't enter. Most frequencies along the light spectrum pass through freely, but for a forbidden few — called the band gap or forbidden zone — the crystal acts like a mirror. It's the optical equivalent of a semiconductor, which controls electrons by way of built-in energy gaps.

Connoisseurs of Oriental carpets know that a flaw woven into the pattern is the weaver's way to say that only

God can create perfection. Somewhat similarly, the real beauty of photonic crystals is their deliberate defect. An enlarged hole, for instance, of just the right placement and dimension within the otherwise regularly patterned structure can trap band-gap photons. "Pretend you're a photon in the defect," says Joannopoulos. "You can see the perfect crystals, but you know better than to move there because if you do, you'll cease to exist."

L I G H T

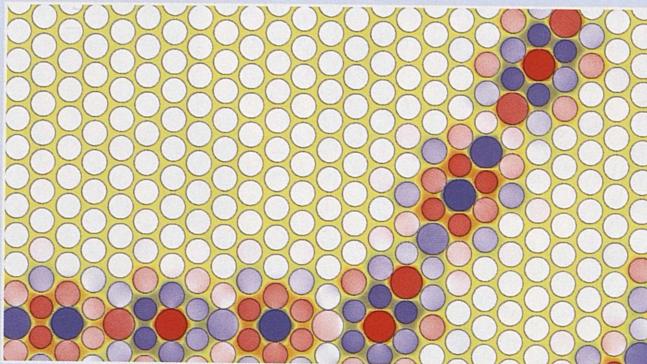
This schematic shows the polarization-independent waveguide bend (orange) and linear waveguide (red) structure simulated by Joannopoulos and colleagues. The blowups show the geometry of stacked layers, a layer of pillars alternating with a layer of "Swiss cheese" dielectric. The rightmost blowup shows the defects that create the waveguide: varying radii of holes and pillars. The distance from one hole to the next is approximately half a micron, ten times thinner than an average human hair.



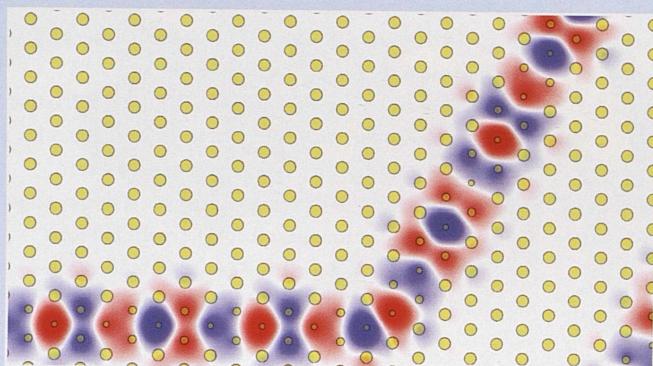
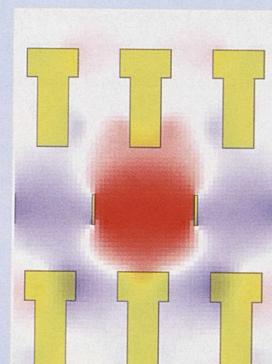
POLARIZED LIGHT GOES AROUND THE BEND

These images from the simulations show the two polarizations (modes) of light in the waveguide, each in adjacent layers of the crystal. The top panels show lateral and vertical images of the TE (transverse electric) mode in the Swiss-cheese layer. The lower panels show the TM (transverse magnetic) mode in the pillar layer. Color (red & blue) indicates positive and negative values for the magnetic and electric fields.

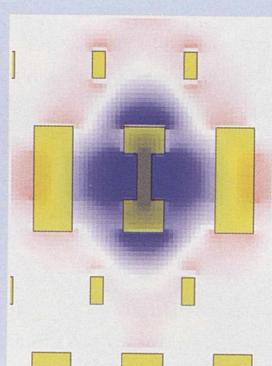
"Both modes are very well localized with the waveguide area," notes Lidorikis, "and share a similar mode profile. This is important for using the crystal as a component in an optical network."



TE: Transverse Electric



TM: Transverse Magnetic



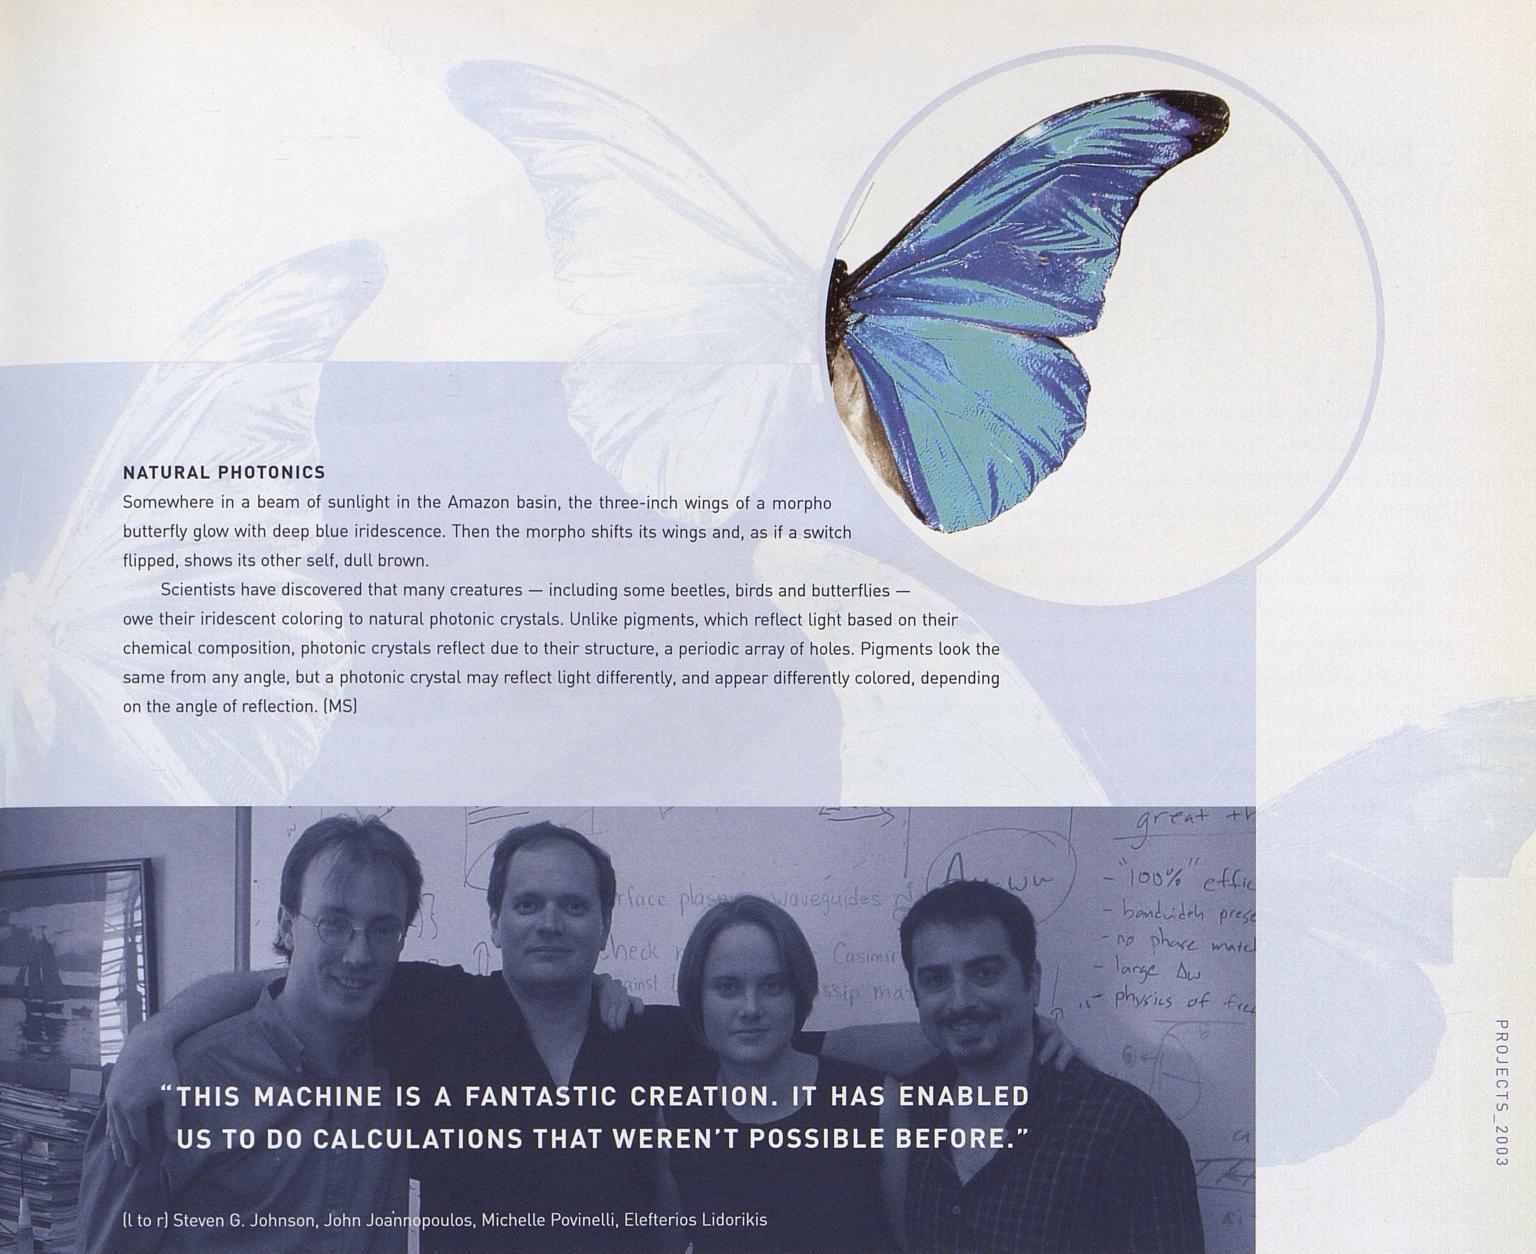
Joannopoulos' work has shown that with relative ease these defects can be fine tuned to trap desired wavelengths, such as, for instance, to capture the red light from a white beam. By lining up a series of defects, the crystal becomes a "waveguide." Like traveling in a tunnel, photons have no choice except to stay on the path defined by the defects.

"We can do things with photonic crystals that we've never been able to do with light before," says Joannopoulos, "and with that ability all sorts of possibilities arise, all new fields open up."

DEGENERATE LIGHT

Fishermen know that polarized sunglasses make it easier to see fish. The reason is that light vibrates both horizontally and vertically as it travels, and under some conditions takes on more of one or the other polarity. The sun's glare off water's surface tends to polarize horizontally. Vertically polarized glasses filter out this horizontal vibration, and the sea appears more transparent.

Joannopoulos' group at MIT has demonstrated the feasibility of photonic crystal waveguides that provide tight cornering without energy loss, a significant advance over fiber optics. A common drawback for these photonic waveguides, however, has been that they are polarization selective; they transmit one polarization differently than the other.



NATURAL PHOTONICS

Somewhere in a beam of sunlight in the Amazon basin, the three-inch wings of a morpho butterfly glow with deep blue iridescence. Then the morpho shifts its wings and, as if a switch flipped, shows its other self, dull brown.

Scientists have discovered that many creatures — including some beetles, birds and butterflies — owe their iridescent coloring to natural photonic crystals. Unlike pigments, which reflect light based on their chemical composition, photonic crystals reflect due to their structure, a periodic array of holes. Pigments look the same from any angle, but a photonic crystal may reflect light differently, and appear differently colored, depending on the angle of reflection. (MS)

"THIS MACHINE IS A FANTASTIC CREATION. IT HAS ENABLED US TO DO CALCULATIONS THAT WEREN'T POSSIBLE BEFORE."

(l to r) Steven G. Johnson, John Joannopoulos, Michelle Povinelli, Eleftherios Lidorikis

Joannopoulos' team set out to demonstrate that, by joining two different styles of photonic-crystal slabs into a single crystal, it should be possible to create waveguides that are polarization independent. To test their idea, without the prohibitive expense of trial-and-error fabrication in the laboratory, they turned to LeMieux. First, they worked on a linear waveguide, a straight-line tunnel through the crystal, analyzing various defect designs to find a combination that resulted in "degeneracy" — a condition in which the two polarizations behave in the same way.

They found a winning combination and went further, testing their design on a waveguide with a 60-degree bend. Because the curve disrupts the symmetrical arrangement of defects, the computational task is much more challenging. To realistically evaluate the light transmission requires clear separation between incident, reflected and transmitted pulses, which in turn requires a very large simulated structure.

Using 256 LeMieux processors, they simulated light pulses of different polarizations. With an extensive series of computations — each requiring about six hours — to completely map all possible wavelengths, the researchers found that, within a frequency range, light comes out of the bend with an efficiency of about 95 percent, maintaining virtually the same polarization over its entire journey. "This is proof of principle," says Lidorikis. "With further tuning, it should be possible to achieve 100 percent degenerate transmission through sharp bends."

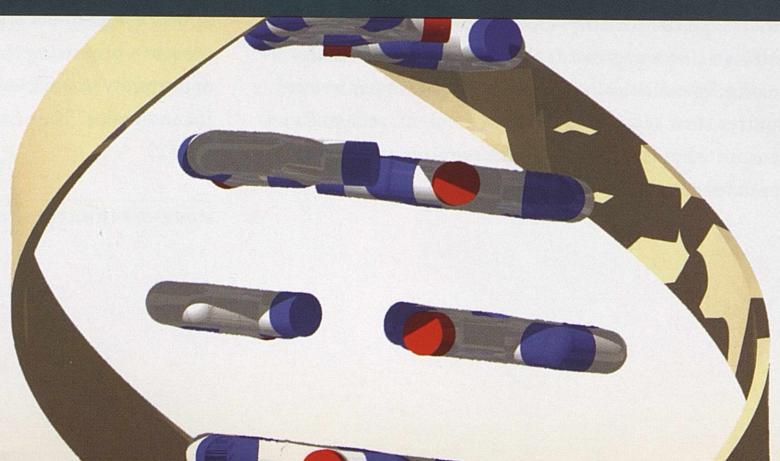
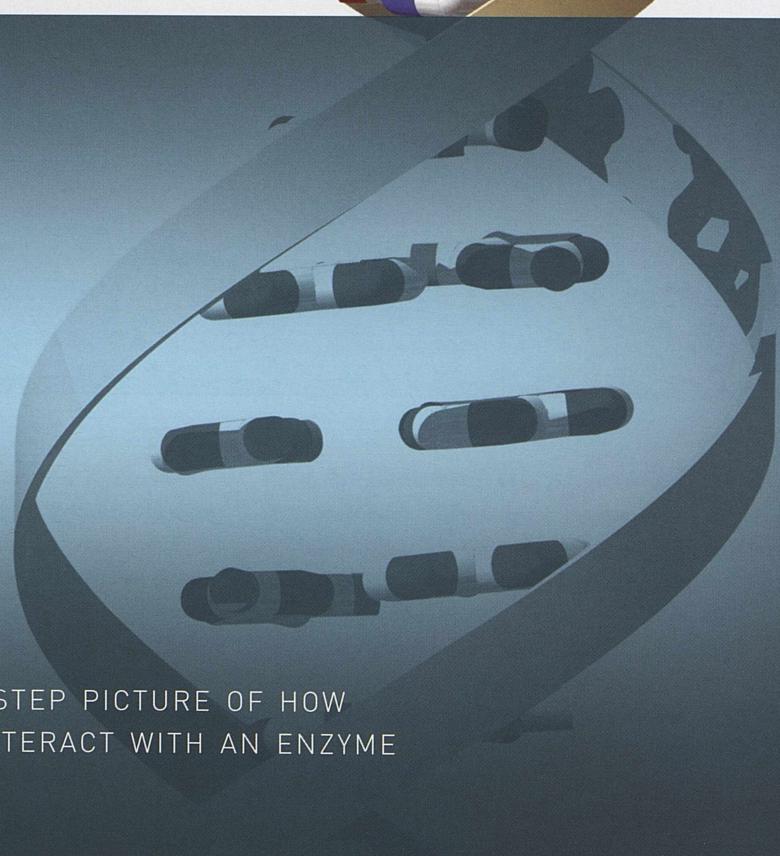
It's a breakthrough that marks the first optical-circuit design to control light to this degree. Miniaturization of optical components is a major goal, and photonic crystals present a promising choice. "Designers will have the opportunity to work with smaller components," says Joannopoulos, "free from concerns about polarization." (MMF)

MORE INFORMATION: <http://www.psc.edu/science/joannopoulos.html>

STRUCTURE_OF_PROTEINS_AND_DNA

3
WHEN
DNA
FLIPS
OUT

THE FIRST STEP-BY-STEP PICTURE OF HOW
DNA OPENS UP TO INTERACT WITH AN ENZYME



Feb. 28, 1953: Two young men walk into a dingy pub in Cambridge, England called the Eagle. To a lunchtime crowd they announce that they've discovered the secret of life. They have.

The two young men were Watson and Crick. Fifty years later, their success at deciphering the structure of DNA stands as the founding event of molecular biology. The elegant spirals of this structure and the phrase that denotes it, "the double helix," have become ingrained in our culture. School lessons teach us that the rungs of the DNA spiral staircase are bonded pairs of chemical bases — A and T, C, G — letters that shape the destiny of all known forms of life.

As this new science has progressed, we've learned much about how the sequence of bases exerts its mighty influence, and we know that DNA doesn't act alone. Enzymes are the *deus ex machina* of DNA drama, coming into the scene and inciting change. Enzymes interact with the bases to facilitate cell division and protein-making, and as a first step in these processes the base pairs must fold out from their sheltered space inside the double helix, a structural shift called base flipping.

"The information in DNA is hidden," says University of Maryland biophysicist Alex MacKerell, "and for DNA to perform its biological function, DNA has to open up so the information can be accessed. Base flipping is a simple structural change that may be the first step in replication and transcription of DNA and is essential for other processes in which enzymes interact with the bases."

Laboratory studies have shown the structure of flipped-out DNA, but laboratory work tells virtually nothing of exactly what happens to initiate the shift and what intermediate states occur along the way. With the availability of LeMieux, PSC's terascale system, MacKerell and his research team tackled these questions with an extensive series of simulations. Their results — reported in the *Proceedings of the National Academy of Sciences* (Jan. 7, 2003) — provide the first atom-by-atom, step-by-step picture of enzyme-facilitated DNA base flipping.

WHICH GROOVE?

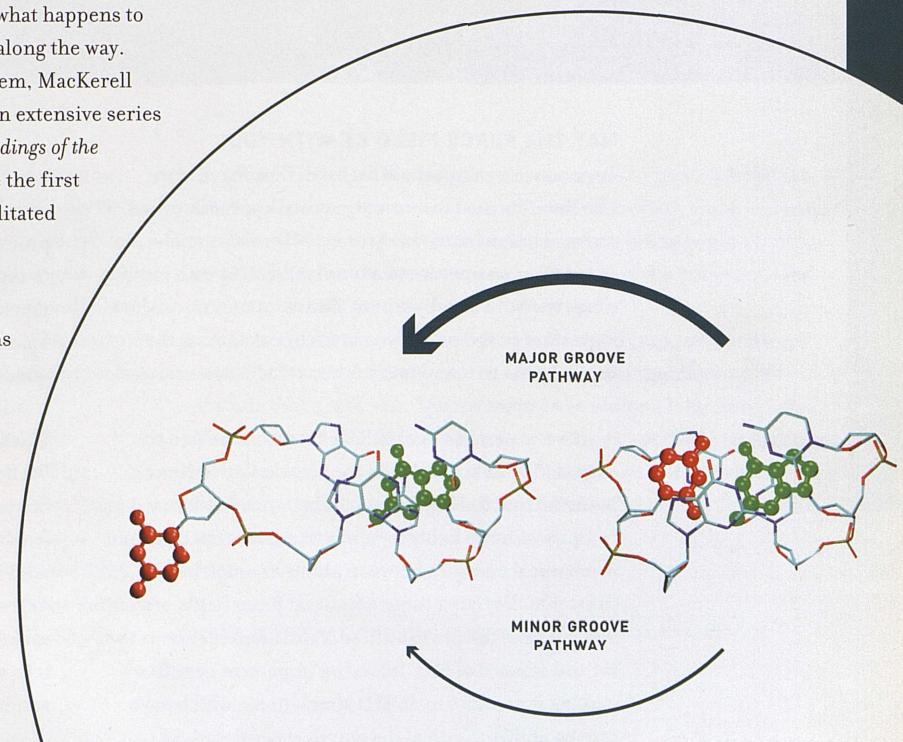
Although DNA base flipping happens in all organisms from plants to people, researchers first confirmed it in bacteria. Laboratory studies have shown that an enzyme called methyltransferase attaches to cytosine, the C of A, T, C and G, and chemically changes it, by adding a methyl group (CH_3-). This relatively simple chemical change, called methylation, is thought to be widespread in DNA interactions. "We're starting to understand," says MacKerell, "that chemical modifications of certain bases are involved in the regulation of the expression and transcription of DNA."

The base has to flip out for methylation to occur, and the flipped-out DNA structure has been identified in laboratory work. Still, what was known before MacKerell's work was a bare outline of the process, like seeing the opening scene of a romantic movie and falling asleep until the noisy wedding at the end.

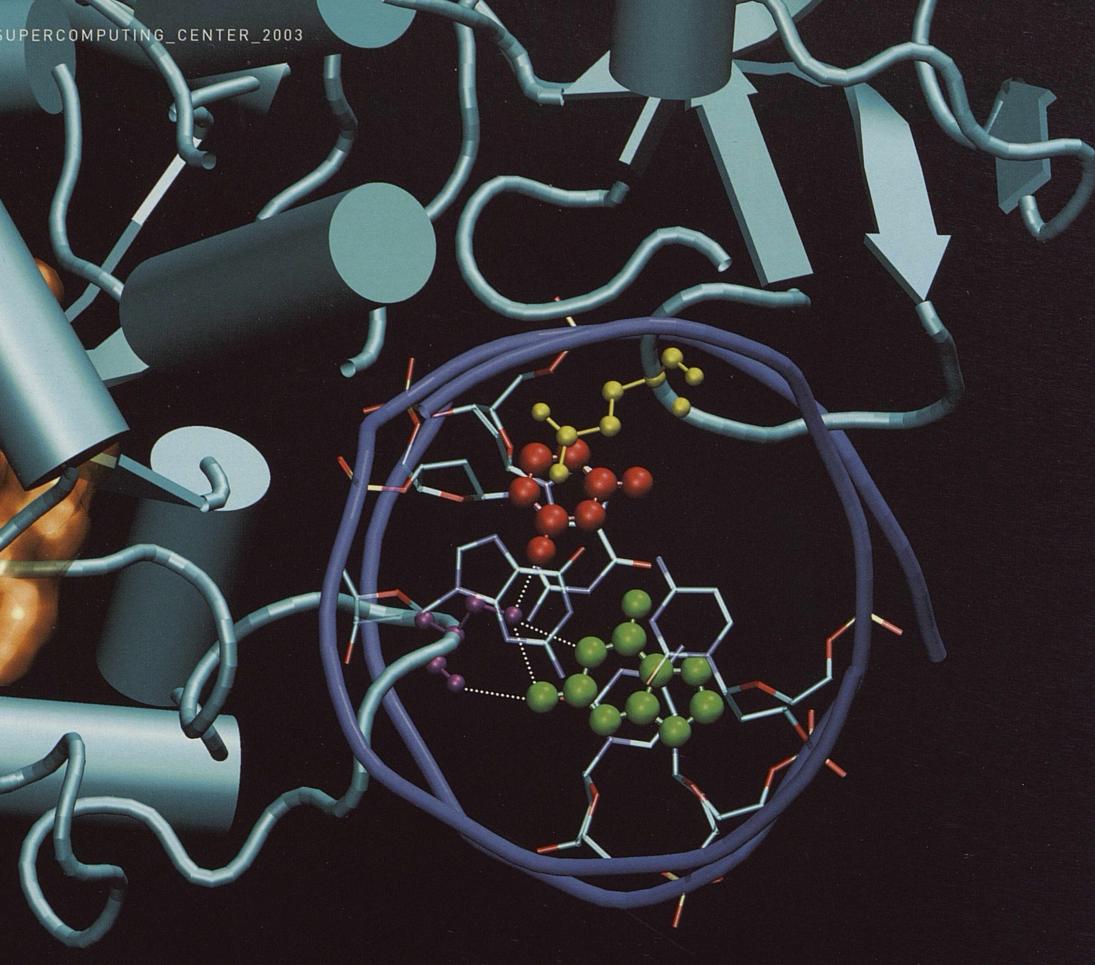
A central unanswered question had to do with how much the enzyme is involved in the base flipping. Does the enzyme help the base to flip out, or does it bind after it's already flipped? Experiments gave no clear answer. "How do we understand," asks MacKerell, "going from the normal duplex DNA shape to the flipped out shape?"

Another question had to do with DNA's grooves. Because of the way paired bases stack up, an intact double helix of DNA has a groove on each side, one smaller than the other, aptly named the minor groove, and a larger one called the major groove. Through which of these grooves do the bases turn as they flip outward? Structural evidence suggested the minor groove, but some experimental evidence suggested the major groove.

"This is where the computer is invaluable," says MacKerell, "because it allows us to systematically change the structure and look at events, which in experimental time frames happen so fast that you can't see them. In the computer, using our mathematical models, we can see what happens."



This top-down view represents the unflipped DNA helix (right) compared to the flipped state (left), with the target base for methylation, cytosine (atoms in red), fully turned out. The "orphaned" guanine (green atoms) remains within the helix.



WHAT HAPPENS WHEN THE ENZYME ARRIVES

MacKerell's simulations show that the methyltransferase binding site — the part of the enzyme that interacts with the DNA — has a dramatic impact on the free energy of base flipping. When the enzyme (blue rods and coils) pairs with the DNA helix, amino acids within the enzyme's binding-site loop (dark blue) interact directly with the cytosine (red) and guanine (green). The first frame (left) shows a serine (purple) competing for the hydrogen bonds (dotted lines) that bind the two bases. The second frame (right) shows the early flipped state, stabilized by the serine and a glycine (gold).

3

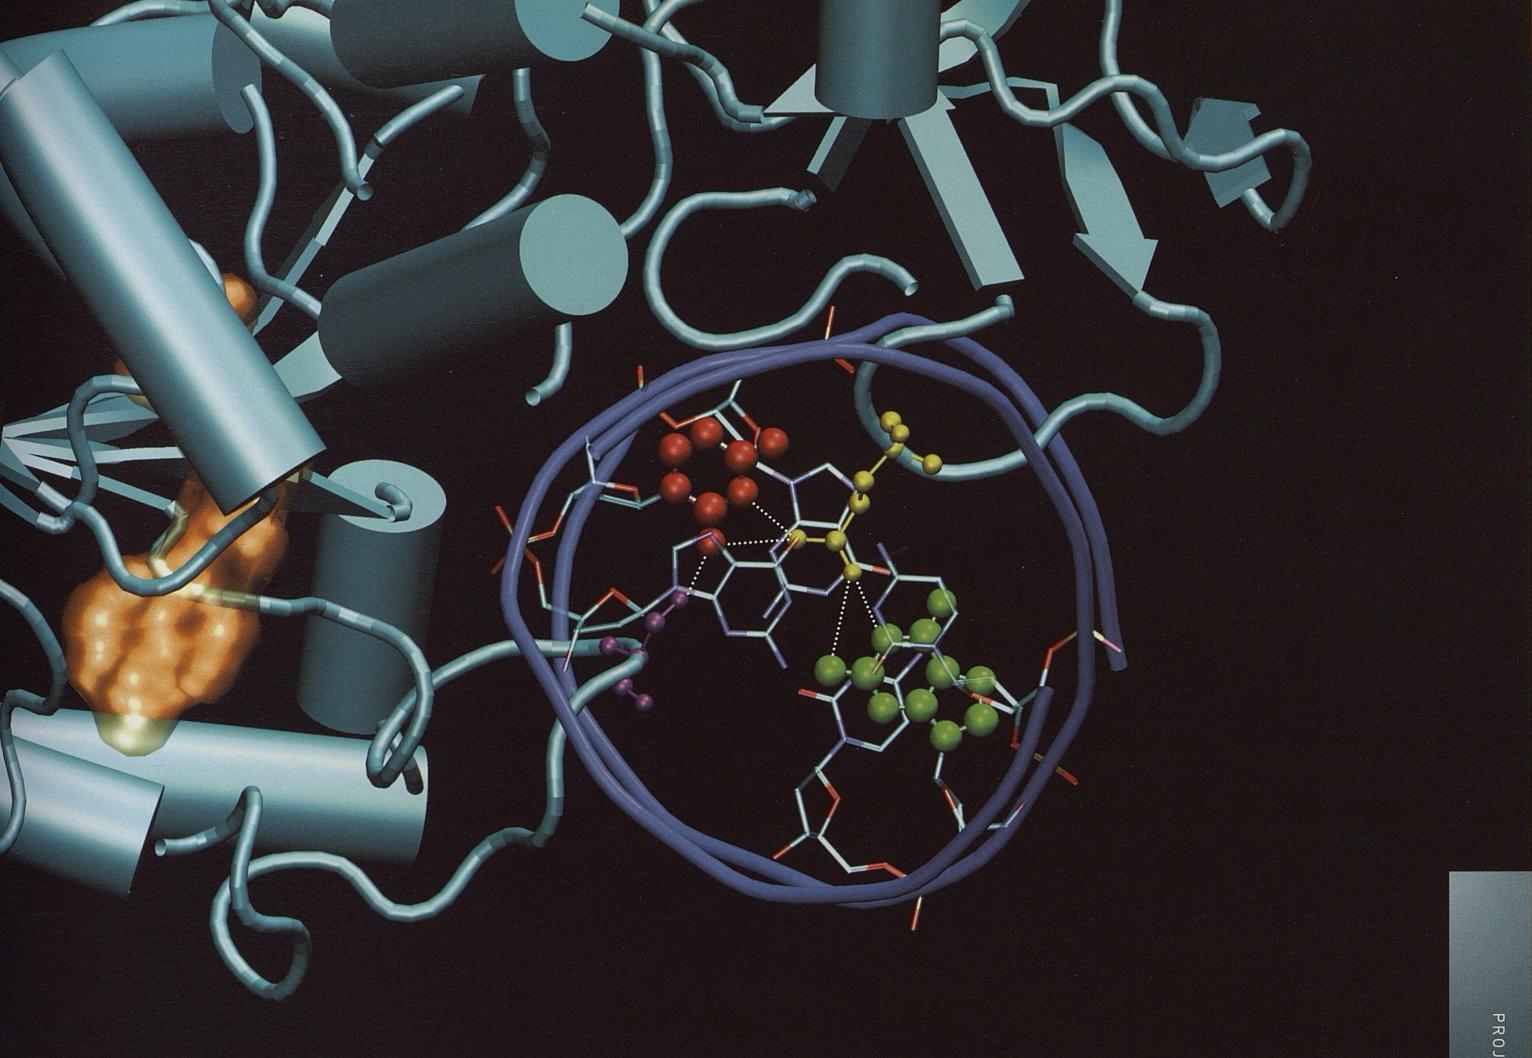
MAY THE FORCE FIELD BE WITH YOU

To produce a comprehensive base-flipping picture, MacKerell turned to a computational approach called molecular dynamics. In essence, MD treats a molecule as a dynamic structure of atoms interacting with each other and with nearby atoms. The computer tracks how each atom in the molecule moves by calculating the forces between it and every other atom at successive slices of time.

Over a period of years, MacKerell has helped to extend MD, first developed for proteins, to become a powerful tool for DNA. Much of his work has focused on "empirical force fields" — a way to express the quantum-mechanical energies between atoms as empirical constants. Deriving these empirical force fields, which approximate the probabilities of quantum theory, is the art and science of MD. It has the important benefit of making it possible to do MD simulations, which have proven ability to reveal the atomic-level details of biomolecular processes.

"To get these parameters to treat the chemical system accurately," says MacKerell, "is a continual process in which we optimize the empirical force field to reproduce experimental data. We also use quantum mechanical data as part of the target data. The empirical force fields have become more sophisticated and more accurate with time."

To arrive at his full-story picture of base flipping, MacKerell broke the process down into chunks he called "simulation windows." Each window is a scene from the full scenario. All possible configurations of the DNA, from closed-to-flipped-to-closed, are represented as a circle, like a clock face, which is sub-divided into 72 five-degree arcs. Within each of these windows, MacKerell calculated relative free-energies, key information that tells what shape the molecule prefers, since it tends to assume the shape that requires the least expenditure of energy.



"THE COMPUTER IS INVALUABLE, BECAUSE IT ALLOWS US TO LOOK AT EVENTS, WHICH IN EXPERIMENTAL TIME FRAMES YOU CAN'T SEE. IN THE COMPUTER WE CAN SEE WHAT HAPPENS."

Alex MacKerell (left) and Niu Huang, University of Maryland

Using eight LeMieux processors for each free-energy window, MacKerell simulated four different configurations of a 12-base DNA sequence: An unflipped helix in a water solution (17,700 atoms), a flipped helix with methyltransferase in two different positions, and a flipped helix with methyltransferase and a third molecule, called a cofactor. For each five-degree window, he simulated 160 picoseconds (a trillionth of a second) of movement—with a snapshot of the action every two femtoseconds, 80,000 time slices per window.

With fourteen months of computing time, 80,000 single-processor hours, and much careful analysis, MacKerell and his colleagues had answers where before there was only mystery. The enzyme initiates flipping, and the base flips through the major groove pathway. "The presence of the enzyme destabilizes the DNA," says MacKerell, "and then the base interacts further with the enzyme, until the enzyme-cofactor complex stabilizes the fully flipped state."

These findings, MacKerell believes, suggest a process by which DNA and the enzyme are in cross talk with each other, like a molecular *pas de deux*. The enzyme, arms spread, approaches to begin binding, and the DNA in turn starts to open, which draws the enzyme closer, until it stabilizes the DNA in the flipped state.

Overall, it's a result that highlights the power of computational methods to uncover the details of DNA-enzyme interactions, a field of study that's still new. "Everyone has known for a long time," says MacKerell, "that DNA has to change its shape to perform its function. We've been able to show for the first time how an enzyme actually facilitates the conformational change. And we've been able to see the atomic details of how it does that." (LK)

MORE INFORMATION: <http://www.psc.edu/science/mackerell.html>

BLOOD_CIRCULATION_AND_THE_HEART

NEW UNDERSTANDING OF WHEN WE'RE
AT RISK FOR VENTRICULAR FIBRILLATION

HEARTS GONE

Ventricular fibrillation — two big words that mean a frenzied, irregular heartbeat that kills. In the United States alone, this condition claims hundreds of thousands of lives each year.

Cardiac arrhythmia is a generic phrase that covers many kinds of heartbeat irregularities — from the occasional skipped beat, which happens to many of us, to various flutters and throbs. Ventricular fibrillation is a cardiac arrhythmia, but in a class by itself for rapid fatal consequences. The heart's built-in synchronization, which keeps millions of muscle cells firing in coordinated wave-like rhythm, goes awry. Instead of a muscular blood pump that sustains life, the heart suddenly becomes a wriggling lump of fibrous tissue. Unless normal heartbeat is restored, blood pressure goes to zero in minutes.

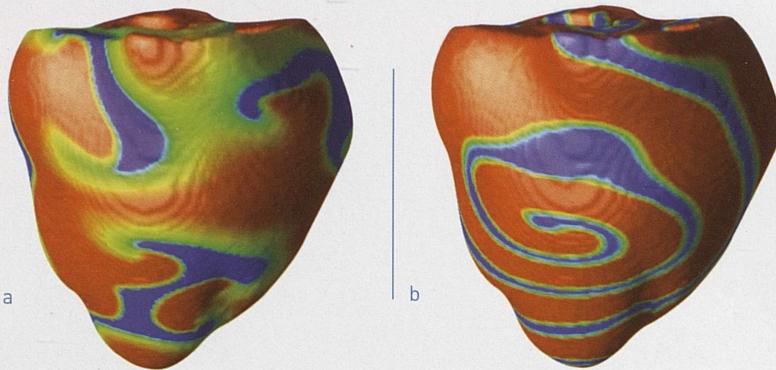
What happens? "After decades of research," says Flavio Fenton, director of electrophysiology at the Heart Institute, Beth Israel Medical Center, New York City, "we still have only an incomplete understanding of how ventricular fibrillation initiates and evolves."

To bring a systematic approach to understanding these problems, Fenton leads a research program that combines clinical, experimental and theoretical work with an emphasis on computer simulation. Using LeMieux, PSC's terascale system, Fenton and his colleagues have incorporated the effects of arrhythmia-inducing drugs into whole-tissue models, the first time this has been done, obtaining excellent agreement with experimental results. And in other recent work, they show that a popular hypothesis for predicting the likelihood of fibrillation is incomplete.

"Despite the seriousness of the problem," says Fenton, who also works with the Center for Arrhythmia Research at Hofstra University, "treatment for some arrhythmias at this time is in general unsatisfactory." Implantable defibrillators help many patients, once they're diagnosed as at risk, but many deaths from fibrillation occur in people with no prior symptoms. Anti-arrhythmic drugs work with varying success, although how they work isn't clearly understood, and some



WILD



CELL-TO-CELL COUPLING IN A 3D RABBIT VENTRICLE

These images show cell-to-cell coupling effects in a 3D computational model of rabbit ventricles. Colors represent electrical potential, showing excited tissue (orange) distinct from quiescent tissue (dark blue) and intermediate voltages (yellow & green). In one model (a), spiral waves break into multiple waves. In the other model (b), with identical restitution properties, cell-to-cell effects suppress alternans and prevent breakup of the spiral wave.

research has shown that the same drugs that suppress fibrillation can, in other circumstances, trigger it.

Fenton and his colleagues have developed models that combine cellular details of electrical transmission with realistic heart anatomies. In this, their work is at the forefront of computational cardiac electrophysiology.

"The heart is complicated," says Fenton, "because there are so many variables. Computational modeling is an invaluable tool as a complement to traditional experiments because it allows us to ask and answer questions that we otherwise can't investigate."

4

"COMPUTATIONAL MODELING ALLOWS US TO ASK AND ANSWER QUESTIONS WE OTHERWISE CAN'T INVESTIGATE."

WAVES OF BIOELECTRICITY

The steady lub-dub we call the heartbeat is regulated by a built-in pacemaker — the sinus node, a small piece of tissue at the top of the right atrium, which sends out electrical pulses about once a second. These pulses stimulate channels in nearby cardiac cells to open. In this "excited state," the cell allows ions to flow in and out, which in turn excites nearby cells, generating a cell-to-cell electrical wave that flashes in milliseconds from the atria to the ventricles and becomes the concerted muscular contraction that pumps blood.

Research shows that many dangerous arrhythmias occur due to "reentrant waves" — in effect, a mistimed electrical pulse that interferes with the normal pattern and loops through the heart-muscle fibers at high frequency. These rapid arrhythmic pulses tend to circle back on themselves, like a dog chasing its tail, and form spiral waves. With a single spiral wave, the resulting fast heartbeat is known as tachycardia. When a single spiral wave becomes unstable and breaks down into many smaller waves, tachycardia transforms into potentially lethal fibrillation.

What's known about these reentrant pulses represents a bare outline, from which Fenton's team works to fill-in the blanks. "The complicated structure of cardiac tissue, as well as the complex ionic currents in the cell," says Fenton, "has made it extremely difficult to pinpoint the detailed dynamics of these life-threatening, reentrant arrhythmias."

Over the past several years, they have developed a set of computational models that allow them to do basic research on how arrhythmias get started and evolve, both as a function of the electrical properties of heart cells and larger-scale heart anatomy. "To represent anatomical complexity," says Fenton, "the cells must be formed into realistic geometrical structures, including several distinct cell types that influence arrhythmia dynamics."

The models draw on the branch of mathematical physics called nonlinear dynamics, which analyzes the evolution in time of systems whose components compete and interact in ways that make the whole more complex than simple cause-and-effect relations among its parts. Because of the great range of scales involved in realistic heart anatomy — both spatially (single cell to full-size heart) and over time (microseconds to minutes) — this kind of modeling demands terascale-level systems such as LeMieux.

ALTERNANS AND THE RESTITUTION HYPOTHESIS

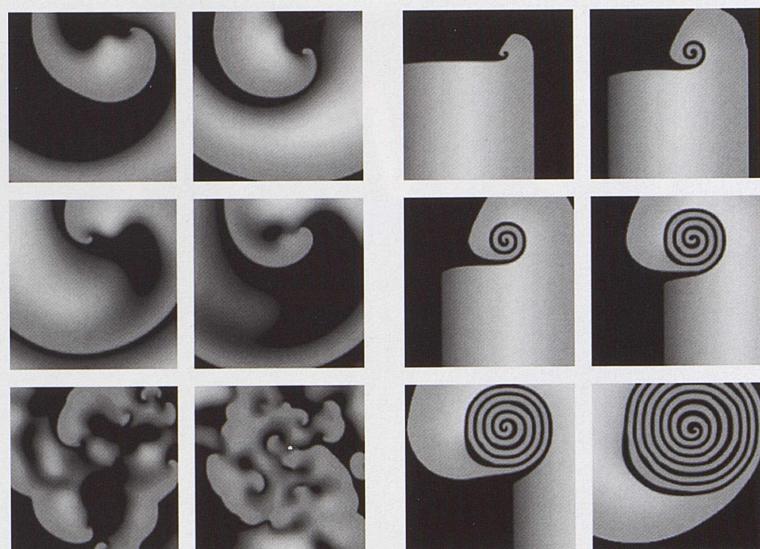
In one recent project, Fenton's team followed up on experiments carried out at the University of Alabama, Birmingham, using two drugs, each of which produced a different arrhythmia — tachycardia or fibrillation — in rabbit ventricles. From the experimental data for each drug, they built separate mathematical models of cellular electrical activity, and they integrated the cellular models into an anatomical model of rabbit ventricles (developed at the University of California, San Diego).

With LeMieux, using up to 500 processors, they simulated the effects of each drug. "We found that two different arrhythmias were induced," says Fenton, "that matched the experimental results not only qualitatively but also quantitatively, with excellent agreement in the dominant frequencies of each arrhythmia."

In another study, Fenton and his colleague Elizabeth Cherry addressed a phenomenon called "alternans," the term for an every-other-heartbeat variation that shows up on electrocardiograms, and which is a well recognized warning sign for ventricular fibrillation. Research with single cells has shown that alternans occurs when, as heart rate quickens, the length of time a cell stays in the excited state becomes shorter more quickly than time in the unexcited state.

Using LeMieux, Fenton and Cherry constructed a series of models to test this relationship — called the "restitution hypothesis" — in the more realistic situation in which cells are in 2D sheets of tissue and also in a realistic 3D heart anatomy, where other physiological effects come into play. In both these situations, they show that accurate predictions of alternans depend not only on the restitution hypothesis, but also on the effects of cell-to-cell coupling and a property called short-term memory, which describes how cells adapt to rate changes.

These results help to explain experimental findings that don't fit the restitution hypothesis, and they are a step forward in understanding the complex processes involved in ventricular fibrillation. "With computational modeling," says Fenton, "we have shown for the first time in whole tissue how arrhythmias may not occur, even when the restitution hypothesis predicts that they will. Our results support the idea that designing antiarrhythmic drugs targeting restitution may not be appropriate because restitution alone is not under all conditions an appropriate predictor of alternans and arrhythmia." (MMF)



SIMULATION OF SPIRAL WAVES

A spiral wave was simulated in 2D sheets of tissue with two different models, differing only in cell-to-cell coupling current due to differences in cell repolarization. In the first model (left), alternans develops and the wave breaks up. In the larger-current model (right), alternans is suppressed, and the induced spiral wave remains stable. Excited tissue (white and light gray) is differentiated from quiescent tissue (black).



Flavio Fenton and Elizabeth Cherry

MORE INFORMATION: <http://www.psc.edu/science/fenton.html>

EVOLUTION_AND_STRUCTURE_OF_THE_UNIVERSE

SIMULATIONS ON LEMIEUX SHOW THAT PLANETS CAN FORM IN DRAMATICALLY LESS TIME THAN PREVIOUSLY THOUGHT

HOW TO COOK

Imagine that you're nearsighted. You get new glasses and look out your window and, all of a sudden, the backyard is vividly in bloom with crocuses, daffodils and marigolds you couldn't see the day before.

For astronomers who gaze deeply into the Milky Way, it's like that. For millennia, as far as we knew, ours was the only solar system, the Sun the only star, among billions visible, that had planets. Now suddenly, within the last eight years, our galaxy is blooming with planets.

Since 1995, astronomers have found more than 100 extrasolar planets, and almost certainly that's just the beginning. The new pair of glasses is Doppler spectroscopy, a method that infers the presence of planets, which only faintly reflect light, from the wobble their gravitational tug induces in their host star.

The method has been so successful, especially as employed by the California team of Geoffrey Marcy and Paul Butler, that as Marcy said three years ago, "Planet hunting has morphed from the marvelous to the mundane." With a data sample that's suddenly 10 times larger, it's not surprising that accepted ideas about planets, including how they come into being, are up for review.

How do planets form? Recently, a research team led by University of Washington astrophysicist Thomas Quinn revisited this old question. Using a powerful new tool — LeMieux, PSC's terascale system — and advanced software called GASOLINE, developed by Quinn's group, they simulated the process by which "gas giant" planets, like Jupiter and Saturn, coalesce into solid mass from the swirl of gas that surrounds a young star.

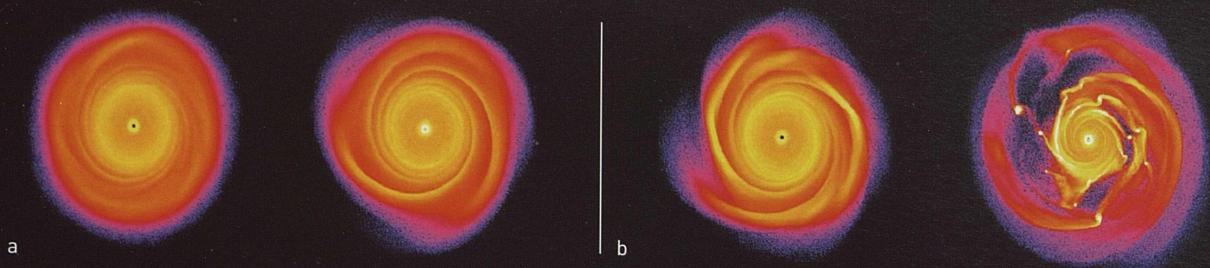
Because of LeMieux, they were able to include 10 times more gas particles in their simulation than previous similar work. Their results with this increased resolution — reported in *Science* (Nov. 29, 2002) — mount a convincing challenge to accepted thinking. "We used a new model of planet formation," said Quinn, "that couldn't adequately be tested without this kind of computing power, and we found that these giant planets can form in hundreds of years, rather than the millions that the standard model predicts."

A GIANT PLANET



PROJECTS_2003

This is an artist's concept of a giant planet recently discovered orbiting the sun-like star 79 Ceti, 117 light-years away in the constellation Cetus the sea monster. The planet, in elliptical orbit about 79 Ceti, was indirectly detected by the slight wobble its gravitational pull induces on the star. The illustration, by artist Greg Bacon of the Space Telescope Science Institute, is courtesy of NASA's Office of Space Science.



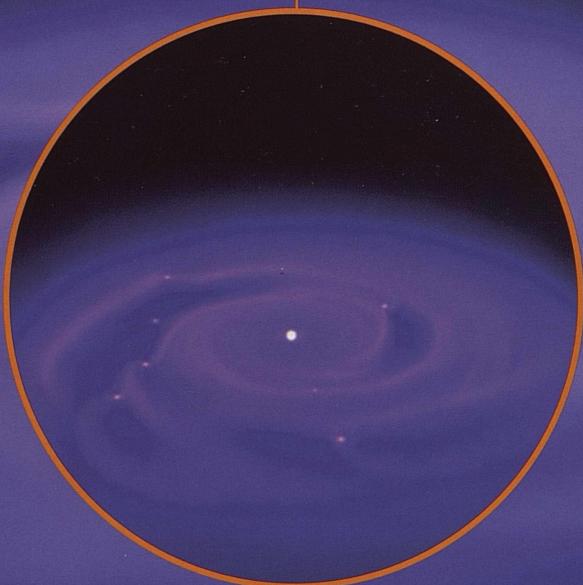
PROTOPLANETARY DISCS

These face-on snapshots from 3D simulations by Quinn and colleagues show density (increasing from dark blue to magenta, red and yellow) of two differently constituted protoplanetary disks (a & b), both at two points of evolution (160 and 350 years). The first disk (a) contains about 10 percent less gaseous mass than the other. Both disks revolve around a solar point mass (black dot) with density shown to a radius of 20 times the Sun to Earth average orbit (about 93 million miles).

"After about 150 years," says Quinn, "the disks develop trailing spiral arms. In the lighter disk, the spirals grow in amplitude up to about 300 years, and then settle into a nearly stationary pattern. In the more massive disk, a two-armed mode grows up to the point, where after about 200 years, fragmentation occurs along the arms, and more than one distinct clump appears."

PLANET FORMATION IS A CURRENT HOT TOPIC AMONG ASTRONOMERS, AND QUINN'S FINDINGS HAVE HELPED TO SHIFT THINKING TOWARD THE FAST-COOKING SCENARIO.

5



This image from simulations by Quinn and colleagues shows how a protoplanetary disk surrounding a young star begins, in a relatively short time, to fragment and form gas giant planets with stable orbits.

PLANET FORMATION ISN'T SPECIAL

For at least 20 years, the textbook view has been that it takes a long time to cook a planet. Little by little — says the old recipe — clumps of solid mass form within the gas, dust and ice that whirl like pizza dough around a young star. Gravity gathers bits of dust, which merge into boulder-sized chunks, which in turn coalesce into bigger rocks.

In about a million years — according to this recipe — these rocks become planets, and over the next few million years, gas from the disk settles around some of these young planets to grow Jupiter-size gas giants, which have most of their mass in a gaseous envelope surrounding the solid core. According to this view, called the "core-accretion model," it would take as much as ten million years for a gas giant planet to form.

The core-accretion model first came up for rethinking, says Quinn, in the early 90s, when the Hubble Space Telescope was able to see these gaseous protoplanetary disks (or proplyds). Astronomers realized that these disks can't last for a long time, roughly a million years, because the gas is rapidly cooked away by radiation from nearby stars.

Still, core accretion held up as a theory, because — so the thinking went — massive planets like Jupiter and Saturn are relatively rare. The Marcy and Butler team carefully surveyed about 1,000 stars, however, and showed that roughly one out of 10 has a planet, probably more, with most of them gas giants, from roughly the size of Jupiter to ten times larger.

"Now," says Quinn, "we see that planet formation is not particularly special, and this doesn't jibe with the standard model. If a gas giant planet can't form quickly, it probably won't form at all."

FAST COOKING WITH GRAVITATIONAL INSTABILITY

Another way for gas giants to form was proposed in 1997 by astrophysicist Alan Boss of the Carnegie Institution. Using computational simulations, Boss found that gravity could act suddenly to form planet-size clumps out of instabilities in the swirling gas disk, and that the clumps formed this way were massive enough to build up a gas envelope. This model, called gravitational-instability, in theory would form a gas giant in much less time.

Until the work by Quinn's team with LeMieux, however, no one had been able to do the simulations with enough precision and length of time to make the case effectively. "The main criticism people had," says Quinn, "is that this model wasn't ready. Nobody was making predictions with it. But that's because they didn't have enough computational horsepower."

The major obstacles in such a calculation have to do with two factors. One is the inherent non-linearity of gravity. As mass starts to collect, gravity gets stronger and stronger. The other factor is dynamic range. "We have to go from gas densities in the disk of less than a microgram per cubic centimeter," says Quinn, "to planetary densities of a gram per cc, a million-fold increase in density."

Quinn's team had spent about a decade developing software, called GASOLINE, designed to simulate cosmology. Within the last four years, they added the ability to accurately capture the evolution and structure of galactic gas. Quinn has long been interested in planet formation and realized that the gas dynamics capability of GASOLINE would also work, with relatively minor modification, to simulate planet formation.

"Developing parallel software," notes Quinn, "is a non-trivial task." In this case, the software was ready at a time when the gravitational-instability model presented itself as a challenge, and the confluence of software and a powerful system, LeMieux, with a problem that needed to be effectively addressed was propitious.

Running on 32 LeMieux processors, over about 30,000 hours of computing, the simulations traced the gas dynamics using a million particles to represent the protoplanetary gas disk, about 10 times more resolution than had previously been attempted. The result was a distribution of masses and orbits comparable to observed extrasolar planets, formed in only hundreds of years, not millions.

Authors of the article in *Science*, along with Quinn, are Lucio Mayer, a post-doctoral researcher with Quinn who's now at the University of Zurich, Joachim Stadel, University of Zurich, and James Wadsley of McMaster University.



Thomas Quinn, University of Washington

Quinn's findings have helped to stir debate among astronomers. Planet formation recipes were a hot topic at the May 2003 meeting of the American Astronomical Society, and some of the experts have shifted their allegiance from core accretion toward gravitational instability (*Science*, June 6, 2003).

"This work shows," says Quinn, "that gravitational instability can actually form self-gravitating protoplanets and that long-lived systems with masses and orbits consistent with those of extrasolar planets arise."

MORE INFORMATION: <http://www.psc.edu/science/quinn.html>

ADVANCES_IN_SOLAR_PHYSICS

PROJECTS 2003

SEEING SPOTS

DETAILED STUDY WITH LEMIEUX SOLVES A PUZZLE ABOUT WHAT
SOLAR PROCESSES GOVERN THE DECAY OF SUNSPOTS

To compare the Sun to spaghetti doesn't occur to many of us, but spaghetti is what comes to mind when Juri Toomre thinks about solar magnetism.

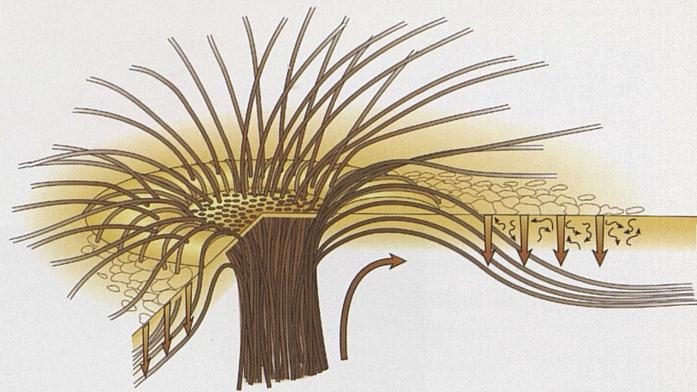
"One of the great mysteries is how the Sun builds and rebuilds its magnetic fields," says Toomre, professor of astrophysics and planetary science at the University of Colorado. He has led several "grand challenge" research efforts on solar turbulence and magnetism and has spent decades working to better understand solar phenomena such as sunspots, the surface blotches that are still a puzzle despite hundreds of years study. What processes cause them? Why do they show up, like clockwork, every 11 years? The answers, says Toomre, have to do with spaghetti.

Blasts of ionized gas inside the Sun create powerful magnetic fields that twist through the solar interior like bundles of spaghetti. Occasionally, the ends of these churning bundles push their way to the Sun's surface, where they appear to us as sunspots. Over time, the Sun's rotation and convection and the effects of other nearby magnetic fields wrench the bundles apart. When this happens, like spaghetti cooked too long, the magnetic field turns to mush and its corresponding sunspot fades from view.

When sunspots break up, they do it with a flourish that can create problems on Earth. With the relatively sudden release of magnetic energy, the Sun also can eject huge blobs of ionized gas into space. At about a million miles an hour, a few days and 93 million miles later, these solar flares can cripple Earth telecommunications and military operations. "These innocent-looking sunspots," says Toomre, "are the footprints for big magnetic structures in the corona, and if a coronal mass ejection comes toward Earth, this is 200 billion tons of plasma, enough to bring down a power grid."

"Being able to forecast these storms is a major area of research activity," says Marc DeRosa, a physicist at the Stanford-Lockheed Institute for Space Research, Palo Alto, California, "and by understanding the evolution of the surface magnetic fields, we hope to gain insight into how these major eruptive events occur."

To that end, DeRosa, who participates in Toomre's research group, and his Stanford-Lockheed colleague Neal Hurlburt, turned to LeMieux, PSC's terascale system, for detailed simulations of how these spaghetti-like, magnetic-field bundles get cooked by turbulent fluid inside the Sun. With a series of computations, they appear to have resolved a vexing question about what solar processes are in charge when the bundles come undone.



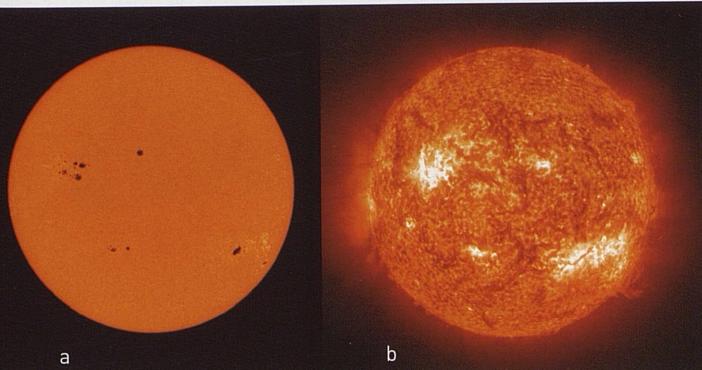
Magnetic field structure of a sunspot. Credit: Nic Brummell/*Nature*

ROILING AND BROILING

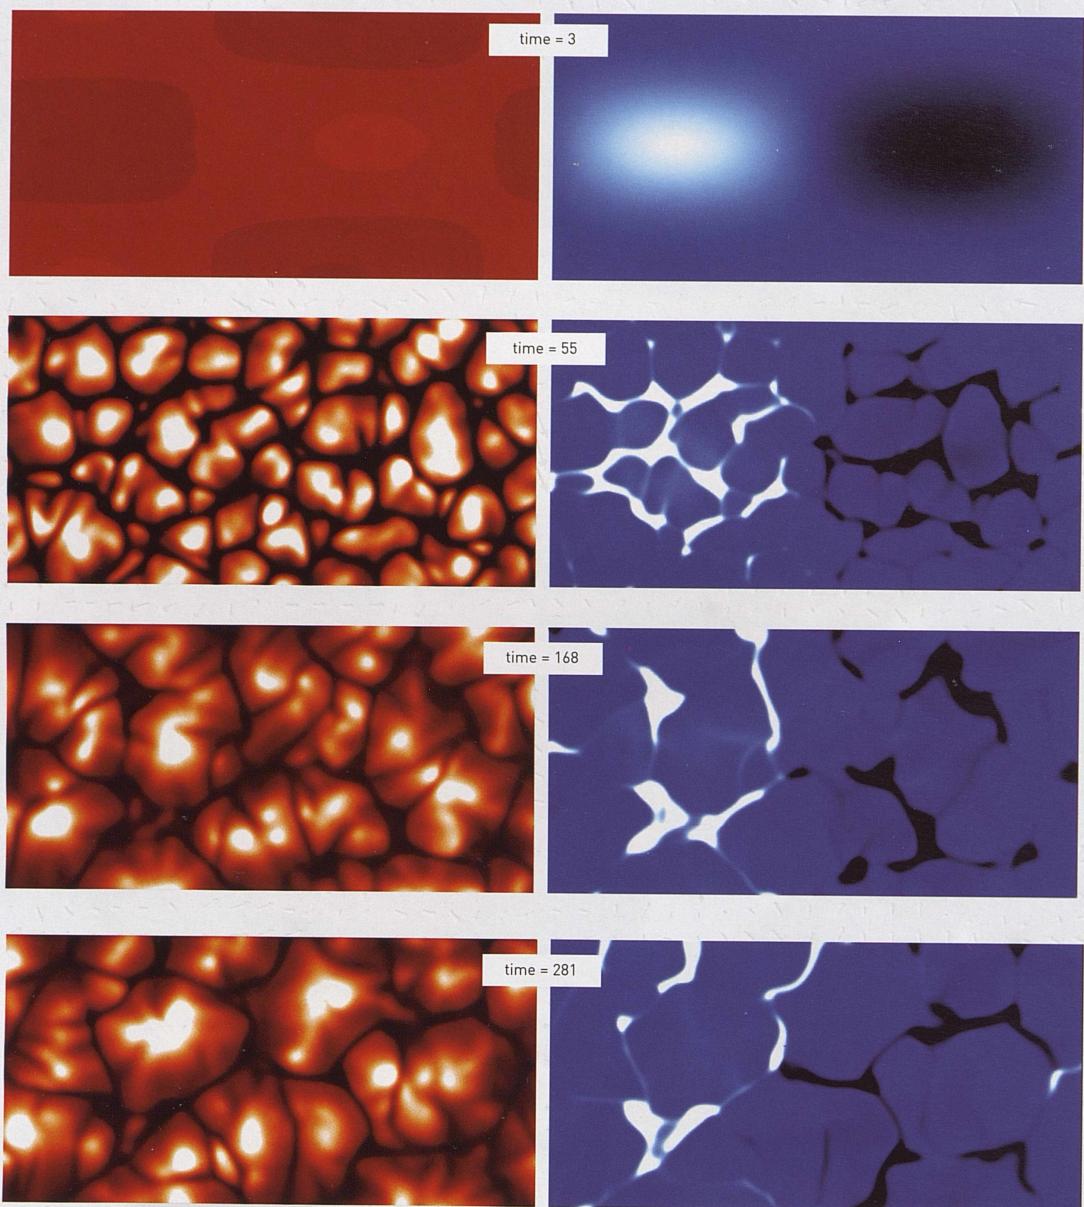
For nearly five centuries, since Galileo and others first observed them around 1610, sunspots have attracted the attention of scientists. We've learned that they appear every 11 years, usually in matched pairs of opposed magnetic-field polarity. Because the polarities switch every 11 years, a complete cycle is 22 years. "Despite being embedded in the roiling and broiling of the Sun's outer layers," says DeRosa, "sunspots display a remarkable amount of order over long time scales."

Sunspot evolution is governed by the "solar dynamo" — the Sun's convective, magnetic churning processes that continually shear and twist magnetic fields. At the Sun's core, a 16-million-degree centigrade nuclear furnace forges helium from hydrogen atoms. Energy slowly diffuses outward until, about two-thirds of the way to the surface, it reaches the "convection zone." There, buoyant gas carries the radiant energy to the solar atmosphere, where it releases as light and the fluid cools. Repeated hot upwellings and cooler sinkings of this mixture appear on the Sun's surface as tightly packed, cell-like granules.

At the inner border of the convection zone is the tachocline, a layer of violent shearing that gives birth to solar magnetic fields. The tachocline organizes magnetic fields, says Toomre, much like a cotton gin combs cotton into fibers. These magnetic-field "fibers" are called flux, and bundles of parallel flux — which have a negative and positive pole — are flux tubes.



These images from NASA's SOHO satellite show sunspots in early September 2001. Because they are cooler than the rest of the Sun's surface, sunspots are dark in visible light (a). An ultraviolet image (b) from about the same time, shows sunspots as visible traces of magnetically active regions.



FLUID VELOCITY AND MAGNETIC FIELD ↑

The vertical flow velocity (left) and magnetic field (right) from the same region show gradual decay of the magnetic field. "The network of downflow lanes changes continuously as new convection cells emerge and old cells die out, causing the magnetic network to evolve," says DeRosa. "As fragments of opposite polarity collide, the field cancels."

THE SIMULATION DOMAIN →

These perspective views represent a radial slab near the top of the solar convection zone. In this view, the flow (top) has evolved past the initial startup and several convection cells have established themselves, with broad upflow centers (white) surrounded by a network of narrow downflow lanes. The magnetic fields (lower) show a cross-section of two opposite-polarity flux tubes. "At regions of strong magnetic flux," says DeRosa, "the flow velocities are noticeably lower."

Within the convection zone, the flux tubes and the up-down thermal churning of convection are pitted against each other. "Convection continually buffets flux tubes and gradually breaks them up," says DeRosa. "At the same time, the strong magnetic field of a flux tube works to inhibit convective motions."

How sunspots maintain their coherence as long as they do — several days or weeks — has puzzled scientists, given that the flux tubes are embedded in the extreme turbulence of the convection zone. The expectation, says DeRosa, is that "turbulent decay" would govern, and the magnetic fields would decay faster than they do.

"THESE ARE ELEGANT CALCULATIONS THAT TAKE US CLOSER TO UNDERSTANDING THE MYSTERY OF THE SOLAR DYNAMO."

Instead, the magnetic fields decay more in accord with "Ohmic decay," a slower process in which the flux tube's field diffuses over time through interaction with flux tubes of opposite polarity. Until now, computational simulations of this process, which were limited to 2D, failed to solve the puzzle, since they indicated, in accord with expectation, that turbulence — rather than Ohmic decay — governed the process.

SEEING INSIDE THE SUN

Since there's no way to step inside the Sun and look at the solar dynamo, computational simulations are essential for detailed understanding. "Most of our knowledge of convection within the Sun," says DeRosa, "is limited to what is measured at the solar surface. Because we can't see below the Sun's atmosphere, we have only a general picture of how the fluid is moving."

With access to LeMieux, it became possible for DeRosa and Hurlburt to undertake much more computationally demanding — and realistic — 3D simulations. "LeMieux's fast interprocessor communication," says DeRosa, "was critical to this work."

To get an in-close look, they modeled a curved, spherical slab near the top of the convection zone, a solar domain large enough for a bipolar paired set of flux tubes, the surrounding magnetic field and convective motion. Using 32 LeMieux processors, they ran a series of three simulations, testing different parameters, with each simulation requiring about 700 hours of computing time.

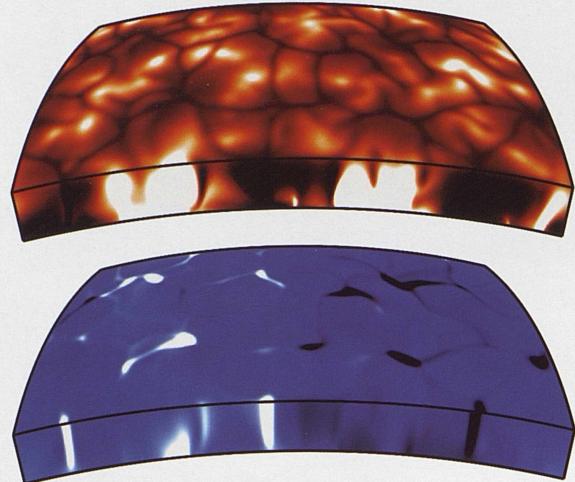
The results: Contrary to 2D simulations, the more realistic 3D study shows magnetic-field decay that matches more closely with observed decay. "Despite the presence of turbulent convection," says DeRosa, "the decay is closer to the Ohmic rate. This helps to explain why solar active regions persist for several days to weeks."

"These are elegant calculations," says Toomre, "that take us closer to understanding the mystery of the solar dynamo." DeRosa thinks the Ohmic rate may predominate because convection, contrary to expectation, is strongest between the opposed-polarity flux tubes and may help to keep them separate, inhibiting turbulent decay.

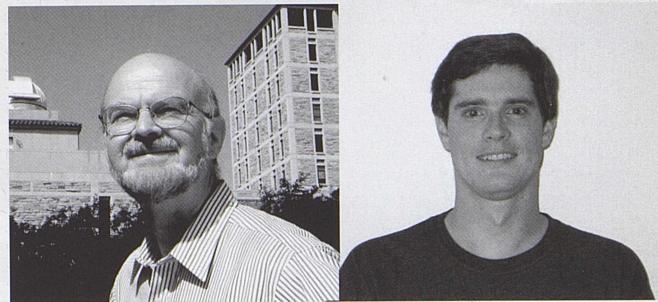
With future simulations on LeMieux, he plans to extend the spherical slab upward to include the solar corona and to couple the interactions between the convection zone and the corona, the part of the Sun we see from Earth. In the corona, as compared to the convection zone, the magnetic field is much stronger in relation to the energy of fluid flow. To combine detailed models of these two physically different regions is a numerically challenging problem that no one has tried, and that presents another test for LeMieux. (LW)

MORE INFORMATION:

<http://www.psc.edu/science/sunspots.html>



Juri Toomre (left) and Marc DeRosa



In Progress

CLEAN ENERGY FROM COAL

Environmentally clean, affordable power from fossil fuel is the goal of research at the National Energy Technology Laboratory, and one of the most promising technologies is coal gasification. Processes to convert coal with 100-percent efficiency into clean-burning natural gas exist. The challenge is to do it in a commercially viable way at industrial scale.

In a program of collaboration between field-test experiments and computer simulation, NETL researchers are working to solve the design problems. NETL consulting engineers Chris Guenther and Madhava Syamlal and U.S. Department of Energy project manager James Longanbach are collaborating with on-site engineers to simulate coal gasification in a plant-sized Kellogg Brown & Root, Inc. fluidized-bed reactor.

SIMULATIONS CAN TEST DESIGN CHANGES AT FAR LESS COST THAN EXPERIMENTS.

Using LeMieux, they have carried out an extensive series of 3D simulations with data from the KBR facility. The goal is a realistic computational model, which will make it possible to test design changes at a fraction of the cost of scale-model experiments. The recent simulations represent major progress in this direction. "This study is unique," says Guenther, "in the scale of our 3D simulations that include both chemistry and heat transfer. We've obtained excellent agreement with the experimental data. These computations could not have been attempted in a reasonable amount of time without access to LeMieux."

From simulations of coal-gasification in an 80-foot tall experimental reactor, this graphic shows concentrations of three gas species — hydrogen (blue), carbon monoxide (green) and carbon dioxide (yellow) — with pathlines (ribbons) to indicate particle dynamics.



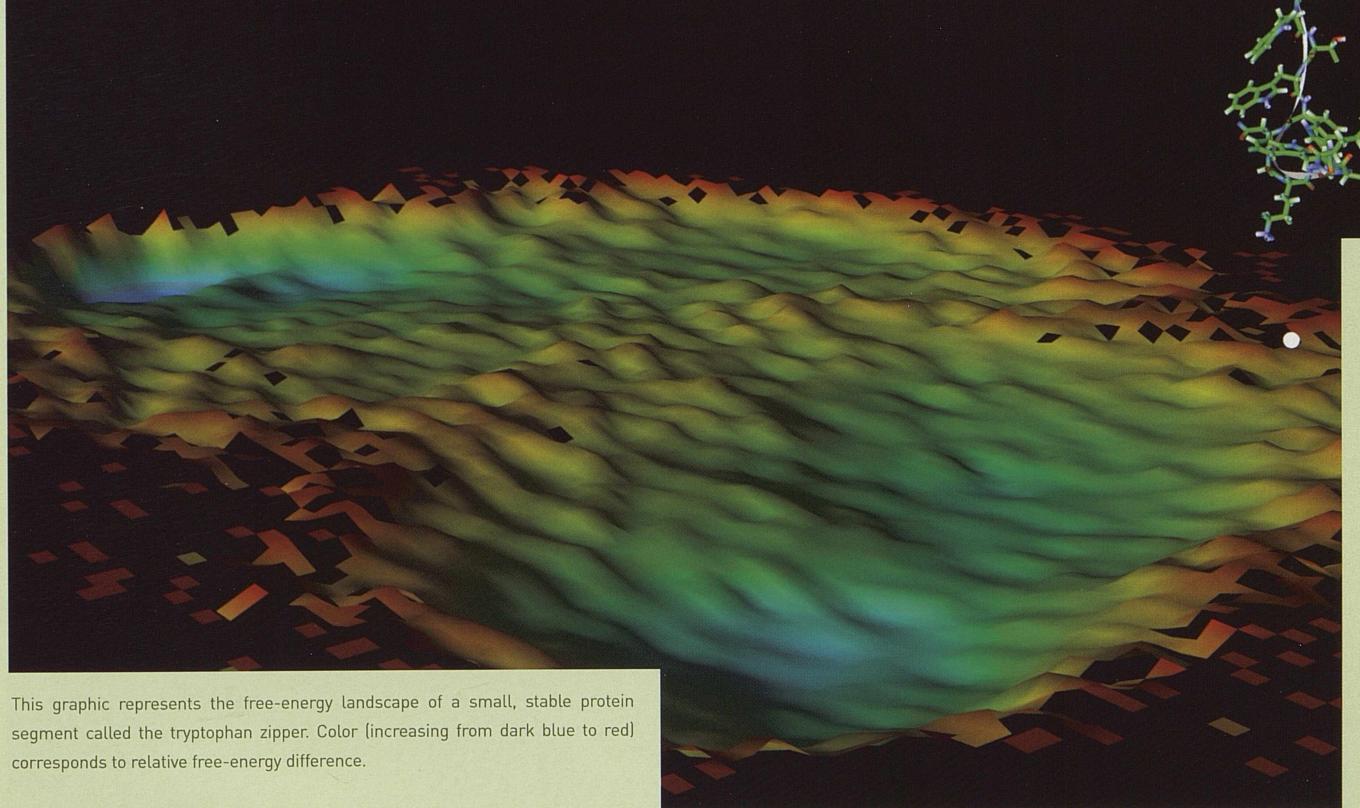
PROTEIN LANDSCAPES

Before it can carry out its biological role, a protein must transform from its newborn form as a droopy string of amino acids into just the right folded shape. How it does this — in seconds or less — is of great interest to molecular biologists.

"We want to understand the folding process," says Carlos Simmerling of SUNY at Stony Brook, "not just the final structure, but how it gets there and what can go wrong, and how this may be related to disease." In recent work, Simmerling has used LeMieux to determine the "free-energy landscape" — a map of how energy interactions among the atoms change as molecular structure changes — for several small proteins.

Each landscape represents many inter-related computations. To carry out these studies, Simmerling has implemented the software for a method, called "replica exchange," that exploits massively parallel systems like LeMieux with high efficiency. As many as 50 simulations at a time, running on up to 32 processors each, communicate with each other periodically to map the landscape's ridges and valleys.

MAPPING THE RIDGES AND VALLEYS OF PROTEIN ENERGY.



INNER EARTH

Why is it that instead of fading away, as magnetic fields eventually do when new energy isn't added, Earth's magnetic field is still going strong — guiding ships, planes and Boy Scouts — after billions of years? Einstein is said to have considered it one of the most important unsolved problems in physics.

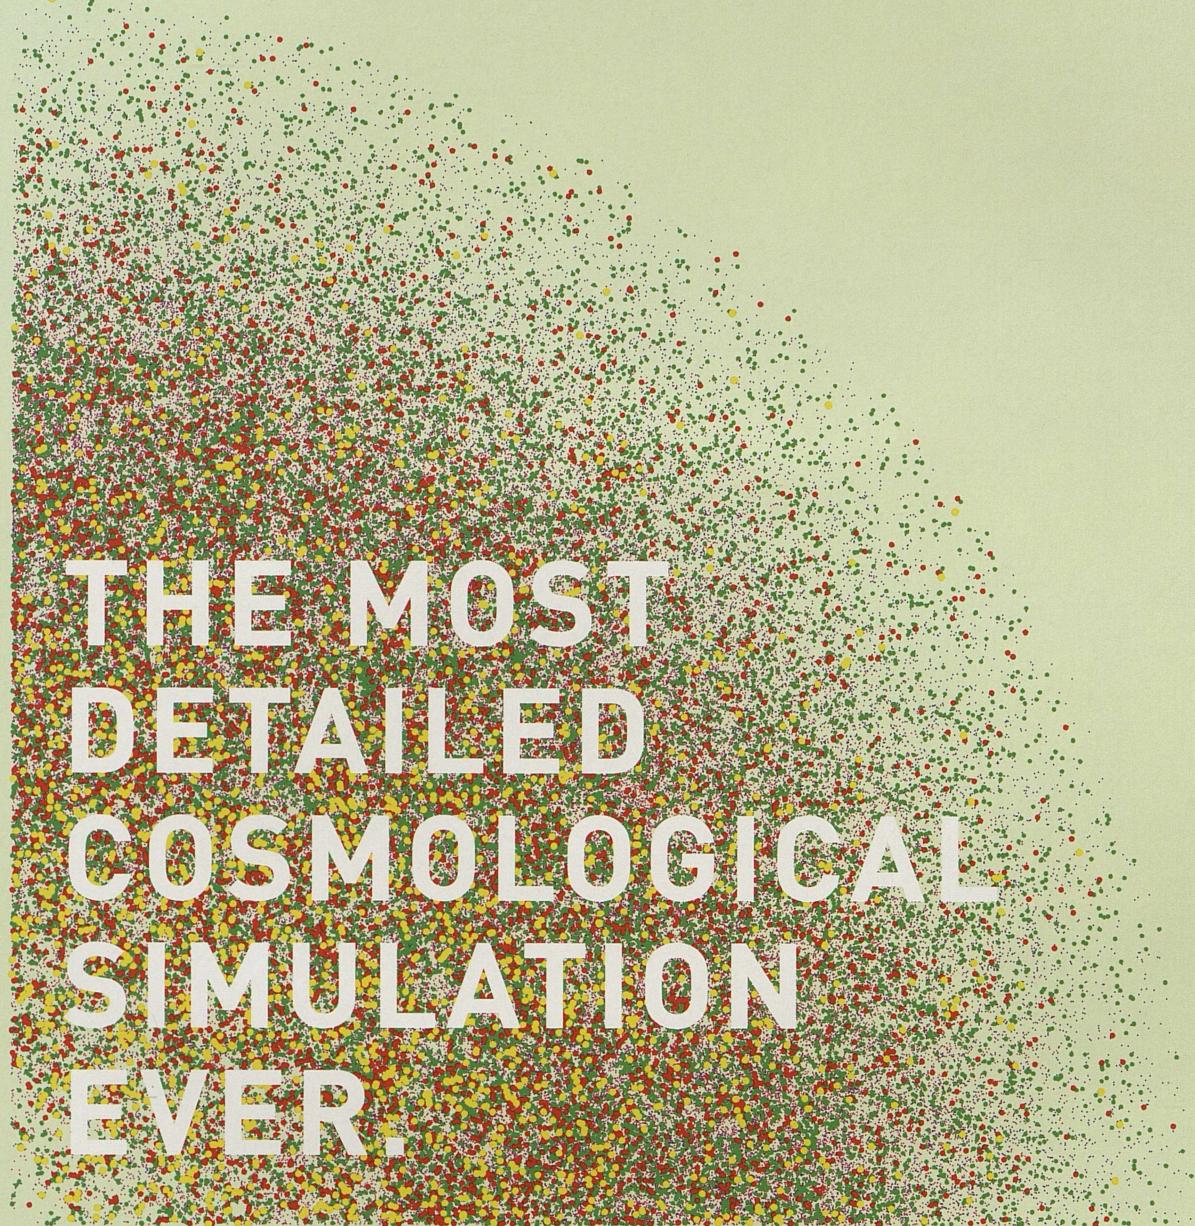
Many scientists believe the answer lies in the turbulent, convective motion of fluids in the Earth's molten outer core. Buoyancy forces, it's thought, that arise from the solid, hot inner core help to drive this fluid motion, which in turn generates new magnetic field.

Gary Glatzmaier of the University of California, Santa Cruz and his colleague Paul Roberts of UCLA developed the first computational model of these geodynamic processes that evolves on its own — self consistently. This model has successfully simulated many features of Earth's magnetic field, including magnetic-field reversal, a recognized phenomenon that has happened many times over Earth's history. In recent work with LeMieux, Glatzmaier and graduate student Darcy Ogden have investigated heat from radioactive decay in the inner core as a buoyancy source that may help to drive the geodynamo.

MORE DETAILED LOOKS AT EARTH'S MAGNETIC FIELD.

Snapshot from a geodynamo simulation showing inward-directed magnetic field lines (blue) distinct from outward-directed lines (red). In the model, Earth's rotation axis is vertical. The intense, complicated field of the core becomes smoother at the Earth's mantle. Field lines extend outward two Earth radii.

This image from the simulation represents what an observer in the lower-left corner might see in a 15×90 degree cone of the sky with an X-ray telescope. Each dot is a cluster containing anywhere from 100 to thousands of galaxies, with color corresponding to mass, increasing from violet to blue, green, red, yellow.



IN_PROGRESS

COSMIC STRUCTURE

Start with a detailed picture of the oldest light in the universe, then simulate 13 billion years of time. The outcome is a distribution of mass over the history of the universe until now. "We can look at that," says Princeton University astrophysicist Paul Bode, "and say 'Does this look like the present-day universe?'"

Jeremiah Ostriker of Princeton and Bode used LeMieux to carry out the most detailed simulation of this type ever done, with two-billion particles to represent cosmic mass. They used data from NASA's WMAP (Wilkinson Microwave Anisotropy Probe) satellite — a detailed map of the cosmic "microwave

background" at 379,000 years after the Big Bang — to initialize the simulation, which took five days of computing time on 420 LeMieux processors.

They are now analyzing the data, work that will produce catalogues of halos — from hundreds to thousands of galaxies in clusters held together by gravity — out to the farthest expanding edges of the universe. These catalogues will allow scientists to compare the simulation's cosmological model — spatially flat, containing both dark energy and dark matter — with astronomical observations, and they will help in planning new observations.

NETWORK SIMULATIONS A HUNDRED TIMES FASTER WITH LEMIEUX.

From a network simulation of 10 subnets (blue circles) of 538 nodes each (yellow dots). The red lines trace links between nodes

BETTER NETWORKS

How long does it take for your e-mail to make the journey to its recipient? Seconds, minutes, or — when traffic is clogged — days? How many crossroad stops along the way? How can networks become faster and use available resources more efficiently?

One of the most powerful tools to answer these questions is packet-level network simulation, which uses computational modeling to track individual parcels of bits as they make their way from sender to receiver across complex webs of routers, switches, nodes and other network components. Because this kind of design-and-analysis tool places great demands on computing time, however, its use normally is limited to relatively small networks of a few hundred components.

Using LeMieux, researchers at Georgia Tech have shown that faster, much expanded simulations are feasible. A team of computer scientists — Richard Fujimoto, George Riley, Kalyan Perumalla and Mostafa Ammar — simulated networks of more than five-million elements and modeled more than 10^6 -million packet transmissions in one second of clock time — a hundred times faster than similar prior simulations. "This offers new capability for engineers and scientists," says Fujimoto, "to design networks with improved speed, reliability and security."

PROTEIN BORDER GUARDS

Two things that don't go together well at border crossings are fast traffic and high security. Biological cells need to be highly selective in what molecules they allow to pass through their walls, and they accomplish this with membrane proteins, which reside in the wall and, like a border guard, control what passes. How they exercise this control, in many cases, however, remains a mystery, the solution of which requires detailed knowledge of how the proteins interact with the molecules that seek to cross.

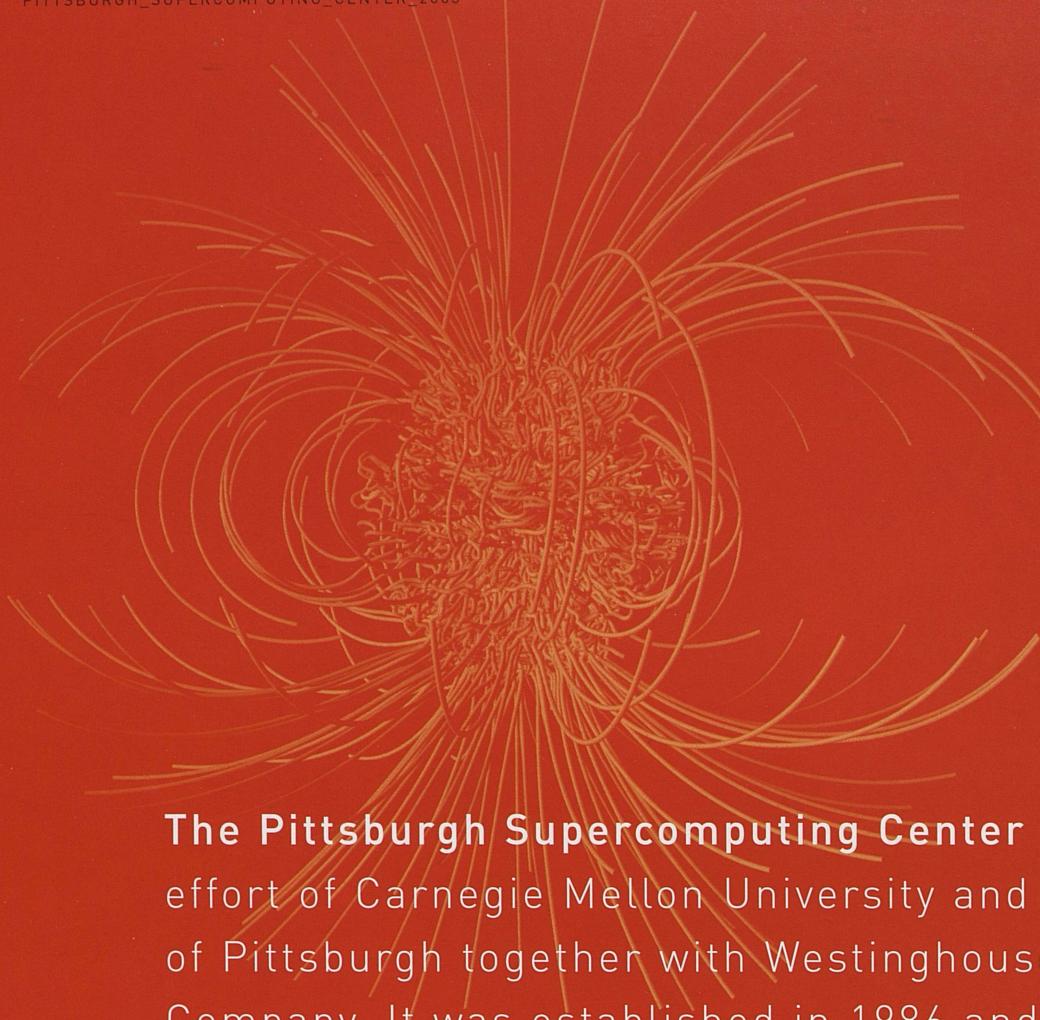
NEW TECHNOLOGY ALLOWS SCIENTISTS TO "FEEL" PROTEINS.

To help solve these puzzles, scientists led by biophysicist Klaus Schulten at the University of Illinois at Urbana-Champaign have developed a new research technology called "interactive molecular dynamics." Using scientific visualization and molecular dynamics simulation — which tracks how atoms in a molecule move — along with computer-game technology, a researcher can, in effect, interact with the protein and "feel" mechanical resistance as an ion, for instance, passes through the channel.

In initial work, they've used LeMieux to run the simulations. One family of membrane proteins they have studied extensively is aquaporins, an important family in many organisms, including people. Last year, they answered a big question about aquaporins' selectivity www.psc.edu/science/schulten2002.html. In a recent study, they explained how some aquaporins transport sugars in a highly selective manner. In ongoing work, they hope to uncover how aquaporins block ions while, at the same time, allowing water to pass freely.

This frame from a simulation of aquaporin represents water molecules (red and white) passing from outside the cell (top) through the aquaporin channel in single file to the cell interior.

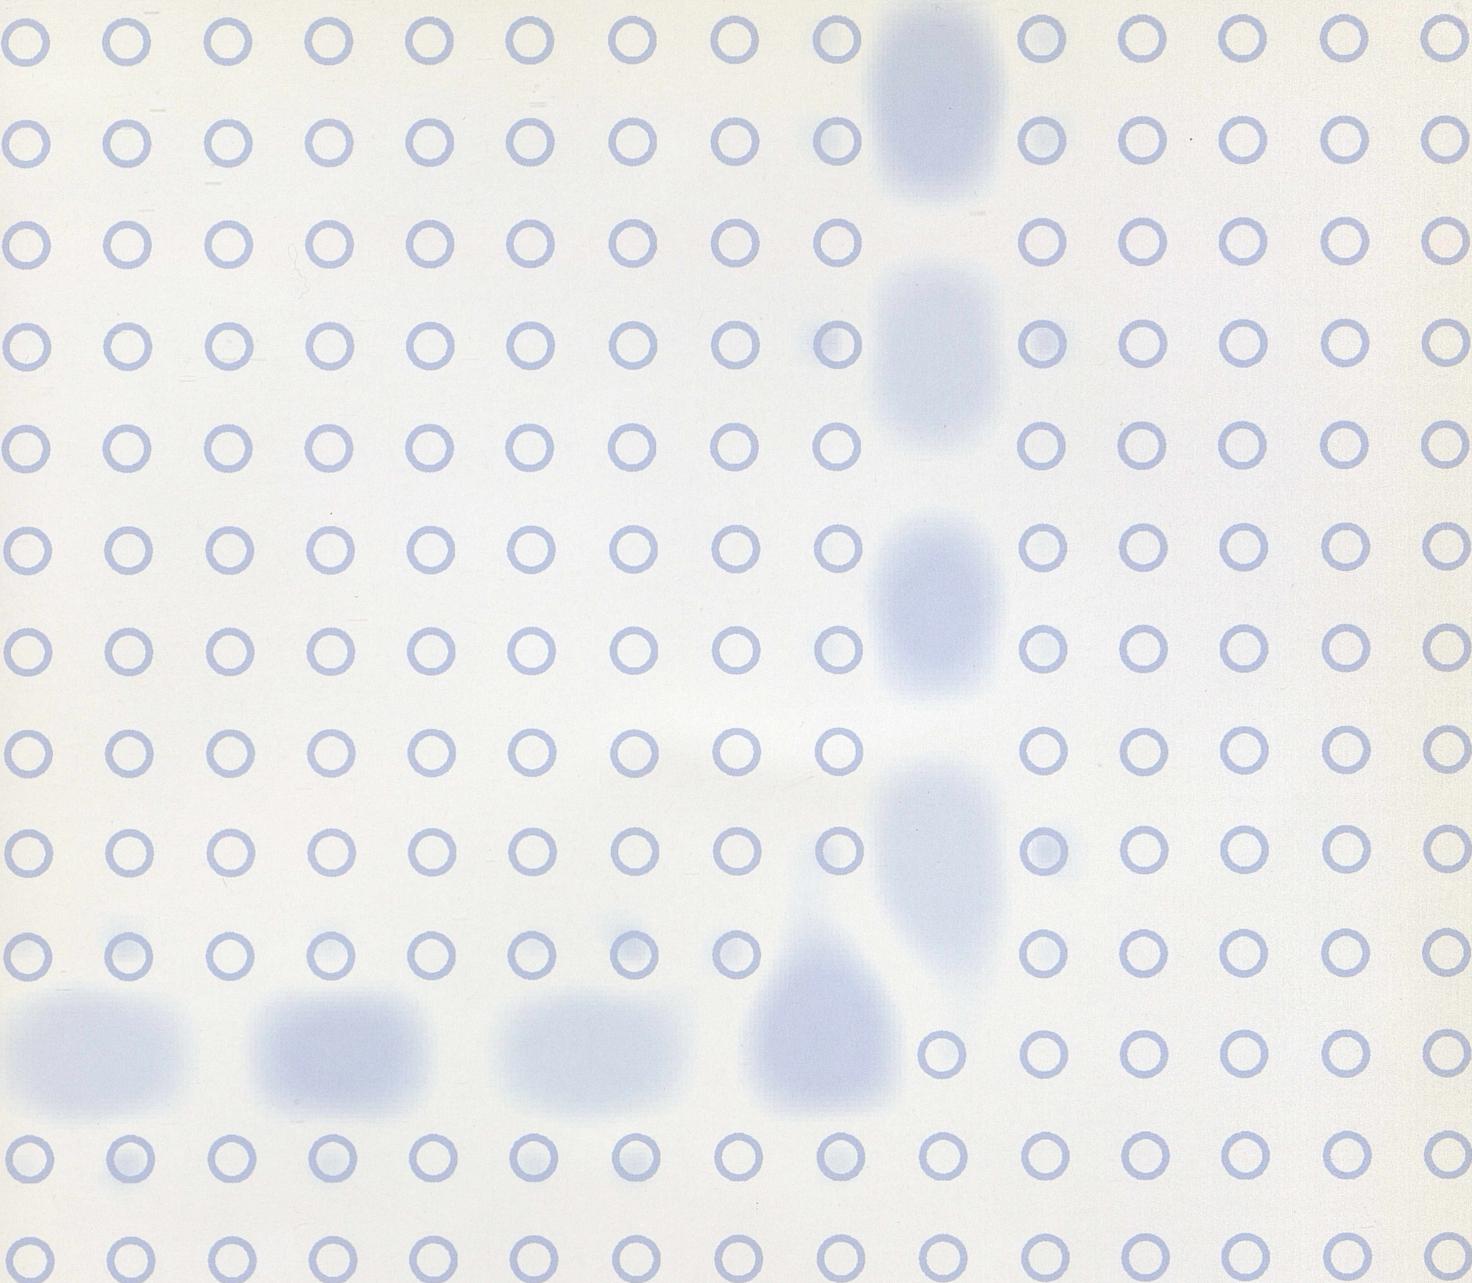
IN PROGRESS



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



The Commonwealth of Pennsylvania
The National Science Foundation
The National Institutes of Health
The National Energy Technology Laboratory
The U. S. Department of Defense
The U. S. Department of Energy
The Howard Hughes Medical Institute
Cisco Systems, Inc.
Cray Inc.
Intel Corp.
International Business Machines Corp.
Silicon Graphics, Inc.
Storage Technology Corp.
Sun Microsystems, Inc.
Hewlett-Packard Company
The Buhl Foundation
The Grable Foundation



EDITOR/WRITER: Michael Schneider, PSC

CONTRIBUTING WRITER: Mary Megliola Franzen, Leah Kauffman, Lauren Ward

GRAPHICS RESEARCH, PHOTOGRAPHY DIRECTION,

COPY EDITING: Sean Fulton, PSC

TRANSCRIBING: Kevin Krisfalusy & Josh Haskins, PSC

PHOTOGRAPHY: Photography & Graphic Services at Mellon Institute. Jonas Salk photo: University Archives, University of Pittsburgh. Photo of Rachel & Jonas the computers, Derek Simmel. Special thanks to the Carnegie Museum of Natural History, Section of Invertebrate Zoology, for the morpho butterfly photo.

GRAPHICS: Thanks to David Deerfield, PSC, for the DNA graphic, Joel Stiles, PSC, for the MCell image, Art Wetzel & Stu Pomerantz, PSC, for the Visible Human graphics & Greg Foss, PSC, for coal gasifier, earthquake, protein landscape and aquaporin graphics and for his help in general. Thanks also to Kostas Daniilidis, University of Pennsylvania, for the tele-immersion graphic; Niu Huang, University of Maryland, for graphics of DNA base flipping; Kalyan Perumalla, Georgia Tech, for the packet simulation graphic.

DESIGN: Wall-to-Wall Studios, Inc.

PRINTING: Broudy Printing Inc.

PITTSBURGH SUPERCOMPUTING CENTER

MELLON INSTITUTE BUILDING

4400 FIFTH AVENUE

PITTSBURGH, PENNSYLVANIA 15213

NONPROFIT ORG

U.S. POSTAGE

PAID

PITTSBURGH, PA

PERMIT NO. 251

