PITTSBURGH SUPERCOMPUTING CENTER 2001

PROJECTS IN SCIENTIFIC COMPUTING

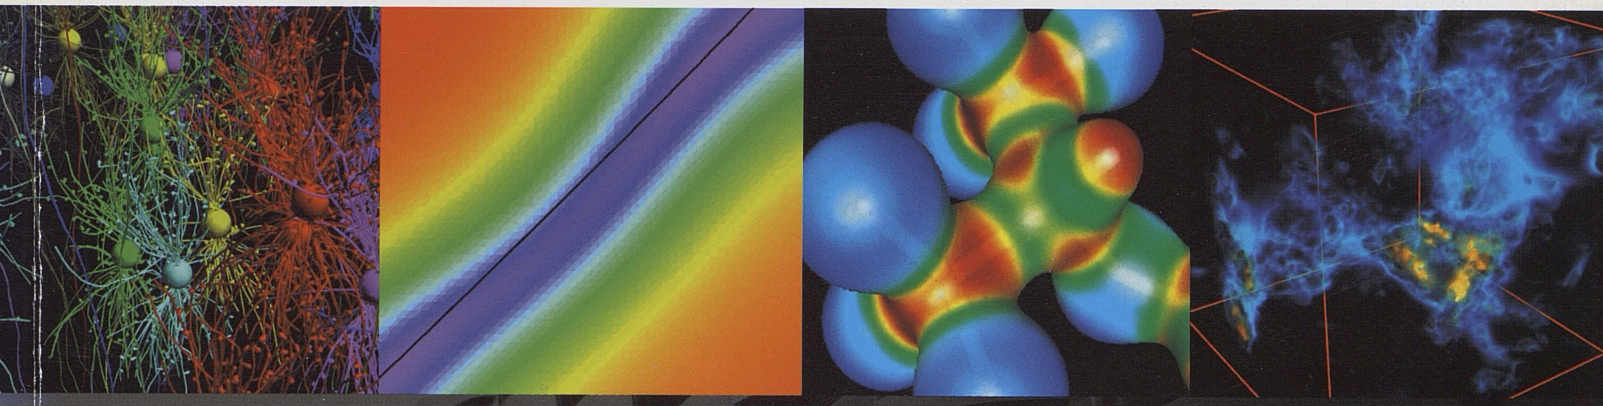

TERASCALE

CONTENTS

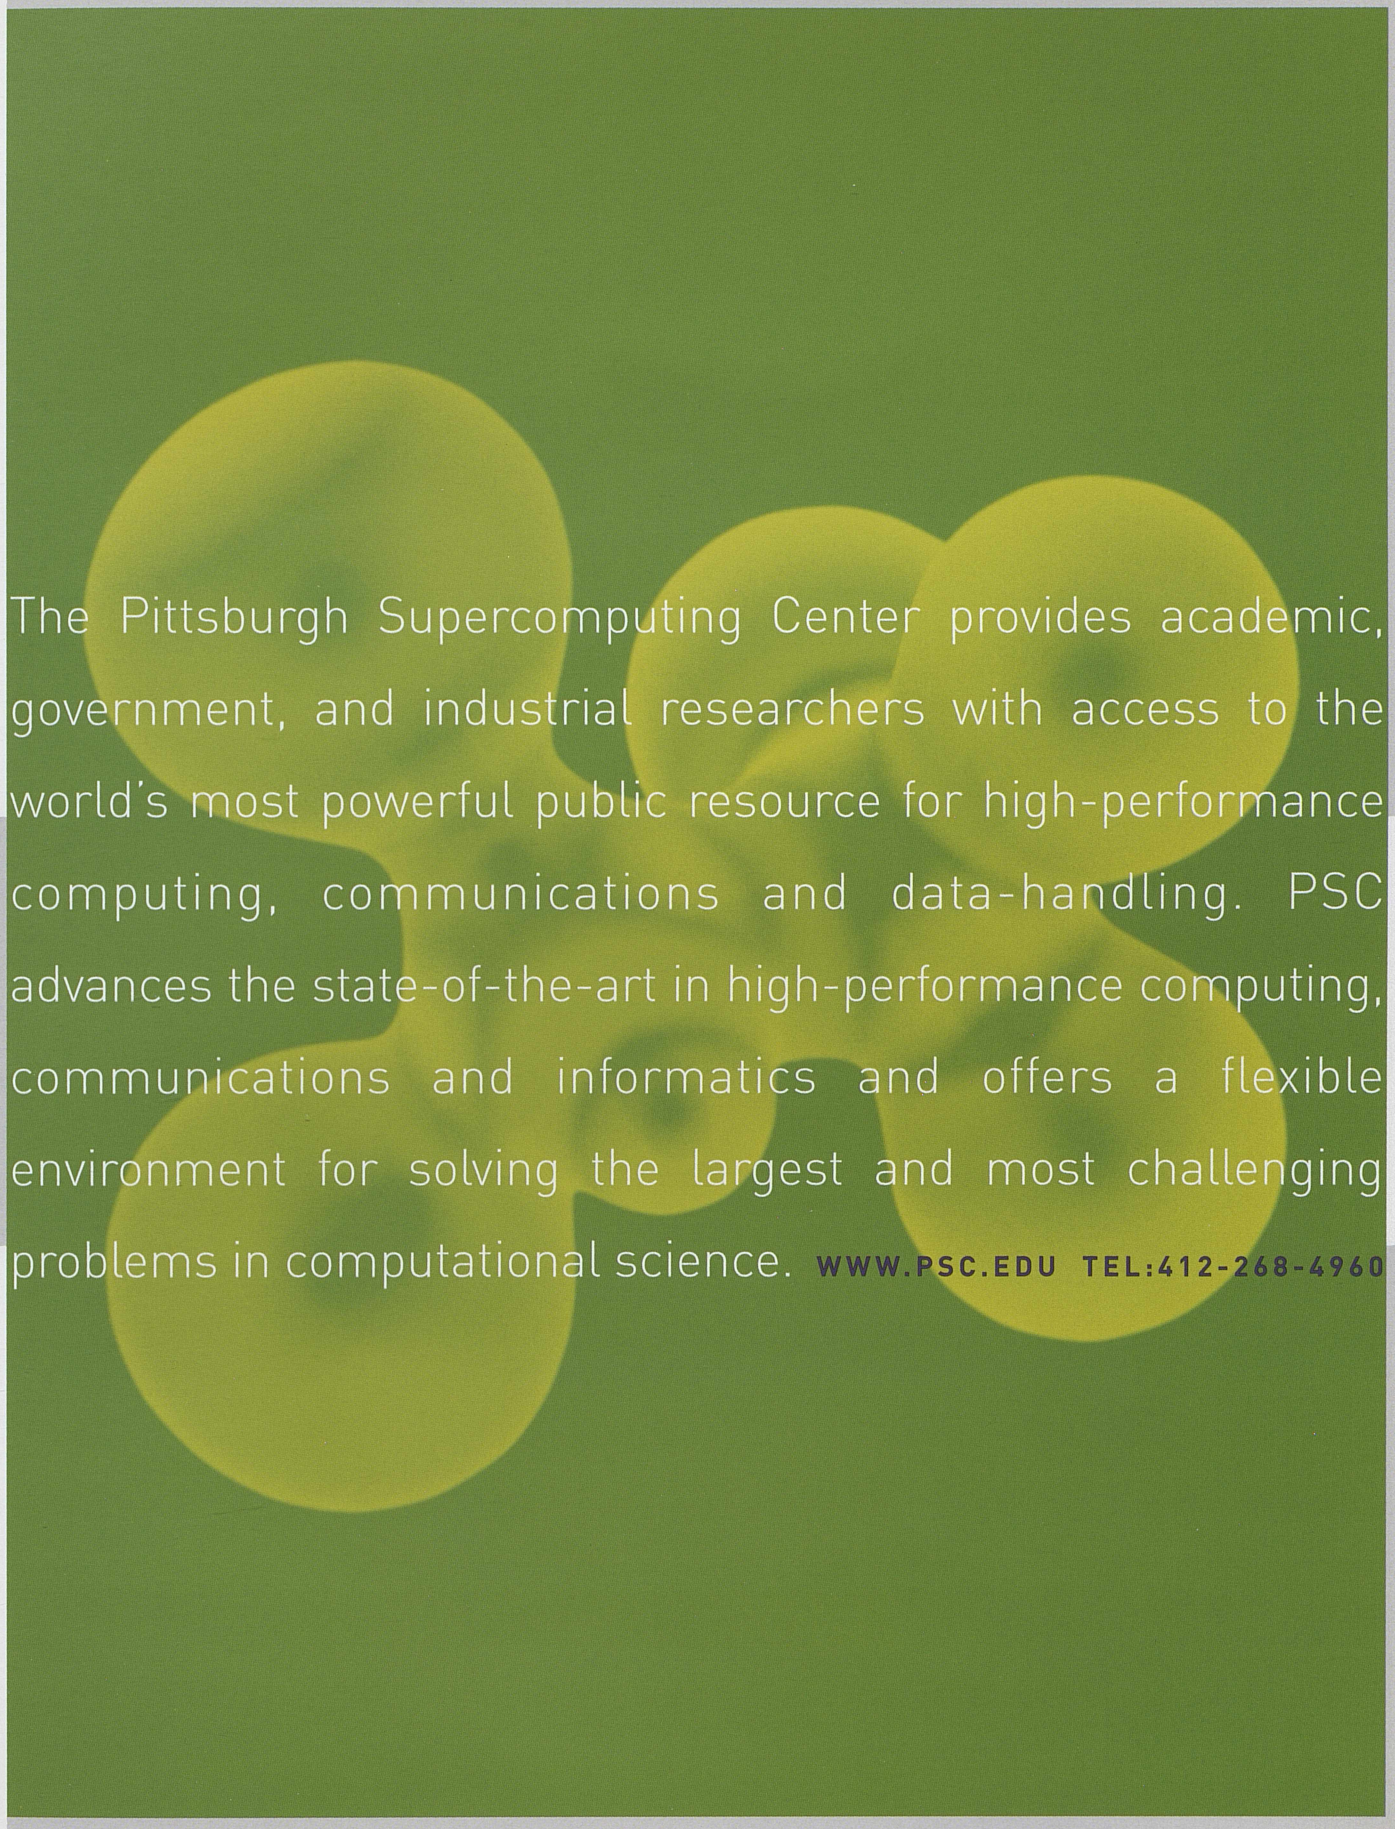

The Pittsburgh Supercomputing Center provides academic, government, and industrial researchers with access to the world's most powerful public resource for high-performance computing, communications and data-handling. PSC advances the state-of-the-art in high-performance computing, communications and informatics and offers a flexible environment for solving the largest and most challenging problems in computational science. **WWW.PSC.EDU TEL:412-268-4960**

Foreword from the Directors

The Terascale Computing System is in, up and running. Installation was completed on time (see p. 4), and preliminary testing shows actual performance that exceeds four teraflops. This most powerful of machines dedicated to open research is already demonstrating the extraordinary capability that will enable a new scale of applications for the nation's research community.

On-time installation represents an outstanding effort from Compaq Computer Corporation. It is Compaq's largest system to date, incorporating multiple new elements, and the install effort involved many dozens of people, at multiple U.S. sites and in Scotland and Ireland. The very substantial initial allocations for this system testify to the large number of researchers eager to use it. Many thanks, Compaq, on behalf of PSC and the entire research community.

During the past year, along with deploying this unprecedented system, PSC has strengthened its research partnership with the National Energy Technology Laboratory (see p. 12). Our biomedical initiative is opening new pathways for the use of high-performance computing in the life sciences (see p. 8). Our networking group has gained further national prominence through its leadership role in the Web100 project (see p. 6). We are contributing to regional and statewide high-technology development (see pp. 10-13). Members of the scientific staff have won support from such competitive programs as the Department of Energy's Scientific Discovery through Advanced Computing, the Department of Defense's Programming Environment and Training and NSF's Information Technology Research.

The research articles in this booklet illustrate the range and depth of work carried out at PSC and exemplify, with a small sampling, the kind of contribution to knowledge that will be enabled by the Terascale system. Four of the projects described here represent the fruits of time on the 256-processor prototype Terascale. In astrophysics, availability of this system enabled the most exacting simulation of star formation yet undertaken (p. 14), and the results provide the first ever picture from a computational simulation of emergent disk-like cores, the progenitors of "proplyds" seen by the Hubble Space Telescope.

The prototype Terascale also made possible new insight into an important class of proteins, the mechanosensitive ion channels (p. 20), and it made possible the most detailed simulation to date of anesthetic drugs acting within a cellular membrane (p. 24). Computational fluid-dynamics is a proven arena when it comes to computational simulation, and there also — in high-performance aircraft design — the Terascale system has shown (p. 30) that this level of computational capability will lead to significantly improved accuracy and reliability in aircraft structural design.

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. We look forward to seeing how the TCS will further unleash the creative potential of our nation's scientists and engineers.

MICHAEL J. LEVINE, SCIENTIFIC DIRECTOR

RALPH Z. ROSKIES, SCIENTIFIC DIRECTOR

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

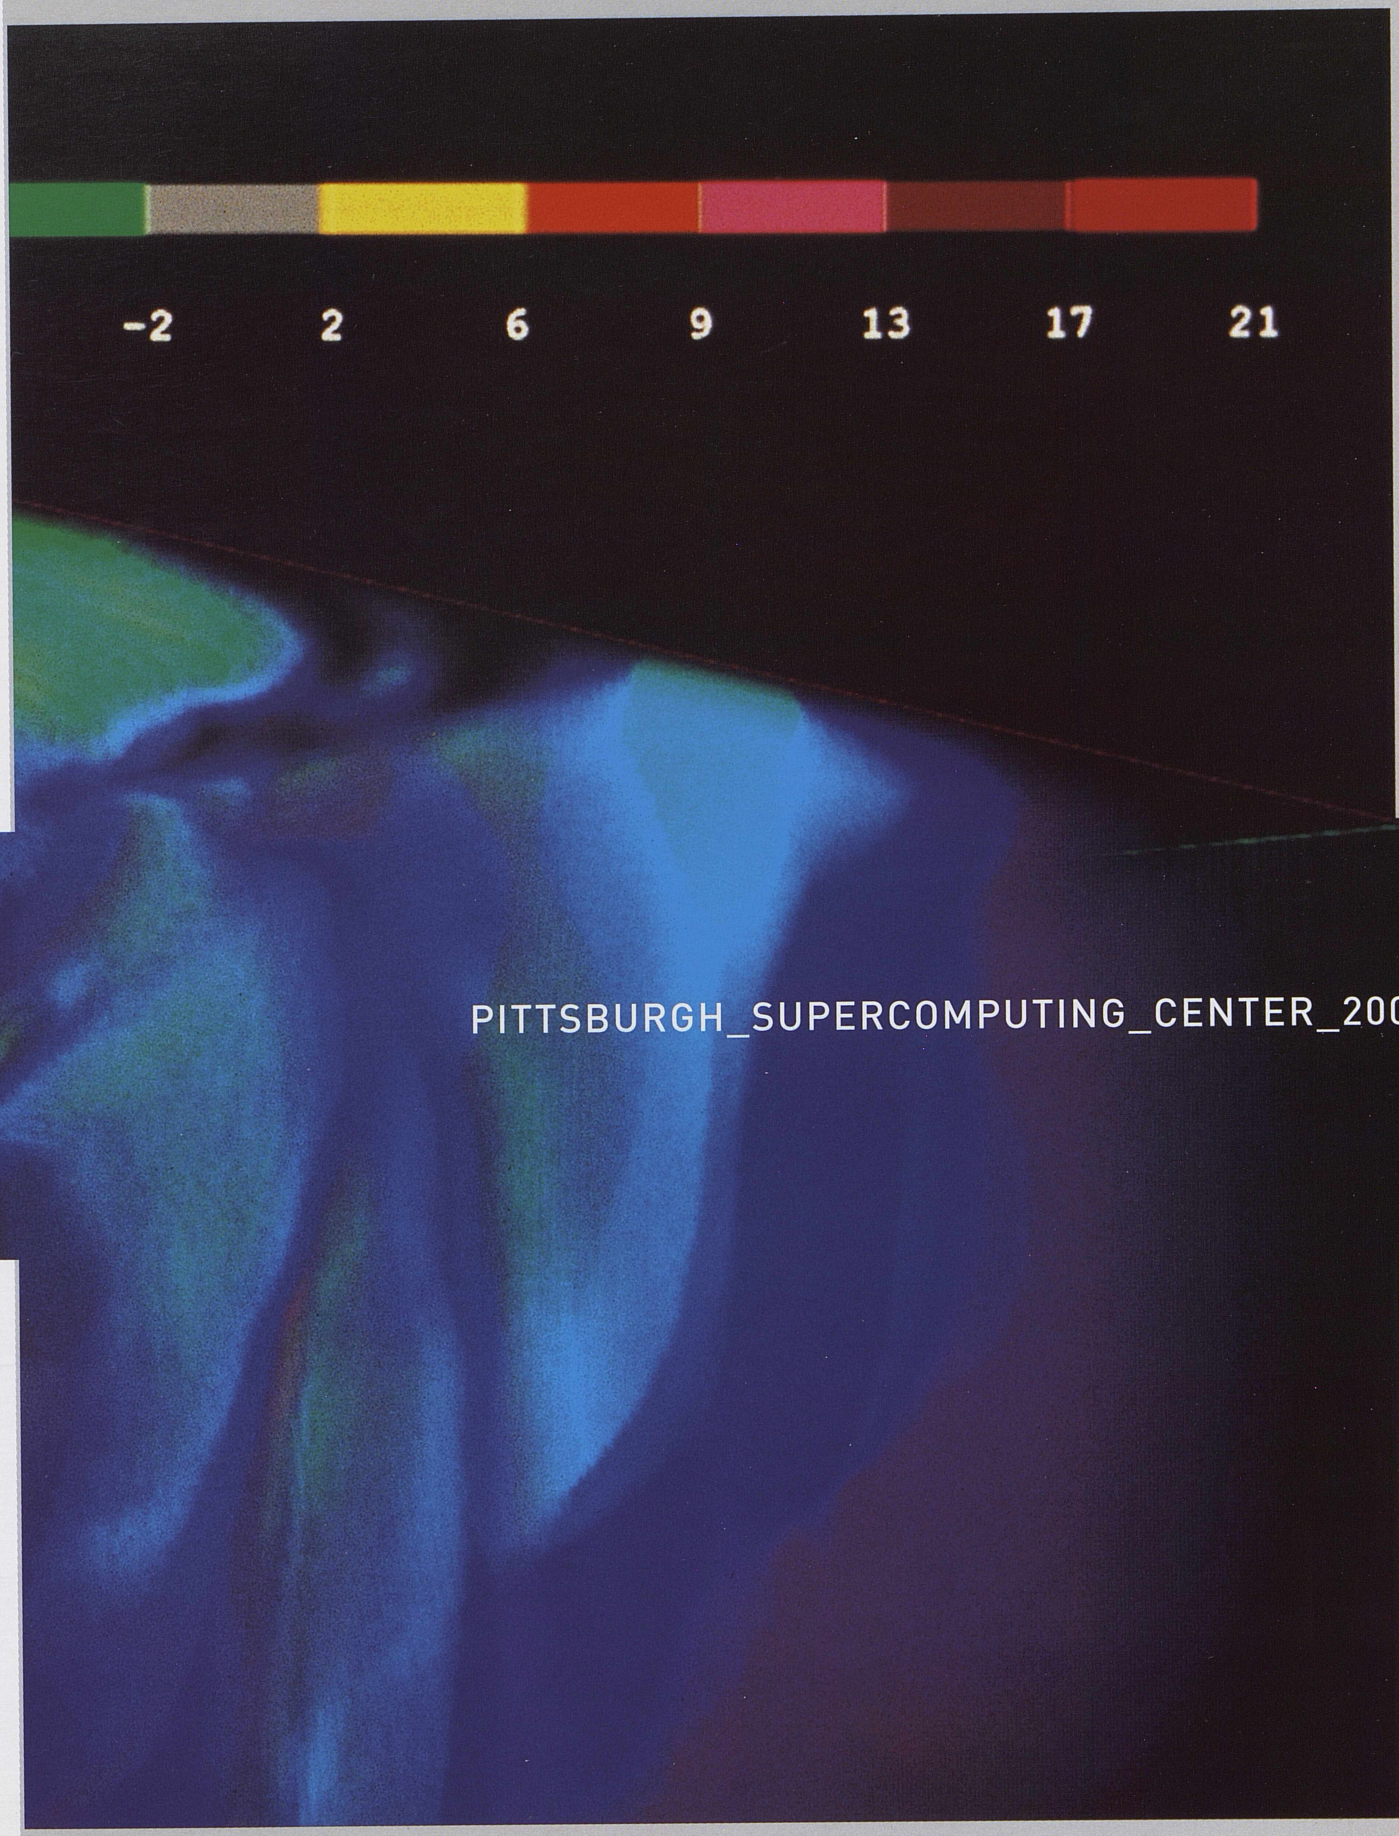

PITTSBURGH_SUPERCOMPUTING_CENTER_2001

Update: Big Iron in the Steel City

THE MOST POWERFUL COMPUTING SYSTEM IN THE WORLD FOR PUBLIC RESEARCH

As this publication goes to press, the PSC computer room at Westinghouse Energy Center in Monroeville, Pa. is a busy scene, the focus of activity for a team of engineers from PSC and Compaq Computer Corporation. Systematically, cable-by-cable, box-by-box, test-by-test, they are bringing into being an unprecedented computing system. When fully installed (scheduled for Oct. 1), it will be the most powerful system in the world available for public research.

With 3,000 Compaq Alpha EV 68 microprocessors, housed in 750 four-processor AlphaServer systems, the Terascale Computing System will provide six teraflops (six trillion calculations per second) of computational capability to U.S. engineers and scientists nationwide. While several other terascale systems are available for classified research, the PSC system will be the most capable to date provided as an open resource for scientists attacking a broad range of problems.

As it becomes a productive research tool, first with a "friendly user" testing period and then, by early 2002, integrated into the NSF's Partnerships for Advanced Computational Infrastructure (PACI) program, the TCS will be used to advance knowledge in many fields. These include earthquake modeling, storm-scale weather forecasting, global climate change, and protein genomics modeling that is integral to the development of new drug therapies.

Developed and implemented at PSC, in close collaboration with Compaq, pursuant to a three-year $45 million award to PSC from the NSF-PACI program, the TCS employs the very fast EV68 chip. This latest evolution of Alpha microchip technology, widely used in both commercial and scientific computing for 10 years, has peak floating-point capability of two gigaflops (two billion calculations per second) and boasts exceptional data-transfer rates.

Along with six teraflops of processing power, the TCS features three terabytes of memory, high-bandwidth, low-latency interconnections and remarkable capabilities for large-scale data handling, including the ability to write the entire memory to disk in under 40 seconds. This extremely short system-write time, developed through PSC software engineering, is critical to efficient checkpointing, needed to preserve research data in the event of component failure.

A prototype 256-processor TCS became operational in October 2000, well ahead of schedule. Since April, it's been an allocated component of the PACI program, demonstrating reliability of the TCS concept. Several projects featured in this booklet (pp. 16-21 & 28-31) represent significant scientific results made possible by this early-version TCS, which is just the beginning of fruitful relationships between the TCS and science.

In scale alone, the TCS pushes beyond where open-resource supercomputing technology has been before or would have gone without this award. Its storage capacity is 100,000 times that of most PCs, with 10 million times the communications capability.

PSC and Compaq have collaborated on numerous machine enhancements to improve the performance of this system, changes that range from the disk controller and file system to wiring optimizations.

While the immediate, direct beneficiaries of this formidable new tool will be academic scientists, the benefits will flow to the country as a whole, in practical ways we can't really forecast. We know that the span from basic research to practical impact is in the range of 10 years. We know that there's a major impact on the economy. And we know that U.S. leadership in basic research is a key factor in our economic strength.

— Michael Levine, PSC co-scientific director

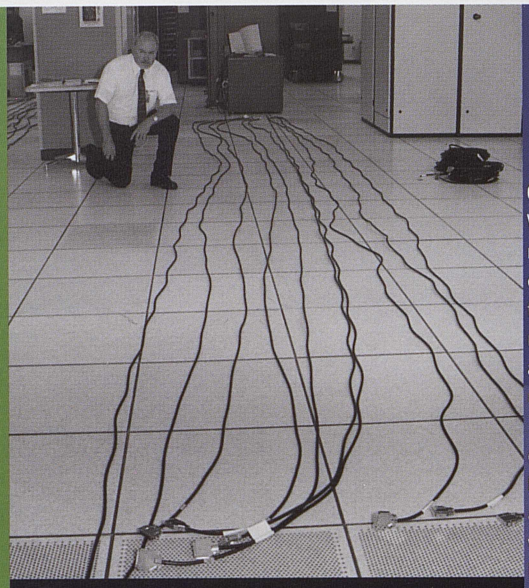

MILES OF CABLE

Lynn Layman (ABOVE), manager of supercomputing at Westinghouse Energy Center, who oversaw the TCS installation, contemplates some of the 14 miles of Quadrics interconnect cable by which 3,000 TCS processors talk to each other rapidly and without stopping for long breaks.

Total TCS floor space is roughly that of a basketball court. By careful site planning and redesign of the AlphaServer configurations, PSC engineers reduced the distance between processors, thereby also reducing the cabling and minimizing network latency. Power cabling is installed under a large signal-cable support grid just under the floor (BELOW, LEFT) that resembles a mattress frame. A similar grid supports additional cables on top of the AlphaServer cabinets.

Along with being a world-class computing system, the TCS is a powerful generator of heat. With a cooling load of 2.5 million BTUs, running the TCS is equivalent to burning 169 pounds of coal an hour. More than 600 feet of eight-inch cooling pipe, weighing 12 tons, circulate up to 900 gallons of water per minute. Twelve 30-ton air-handling units (white cabinets) provide cooling capacity equivalent to 375 room air conditioners.

OPENING NEW FIELDS

We are going through a revolution, or a series of revolutions. As we increase the capability of computational platforms, not only in raw computational power, but also in communications, in memory, the amount of data a system can hold and manipulate comfortably, we keep opening up new fields that were not feasible for work before.

In the biomedical area alone, the picture has changed dramatically since NSF supercomputing began, little more than 15 years ago. In structural biology, computational work was feeling its way and it's now making important contributions. Functional magnetic-resonance imaging work on mapping the brain was not underway, and now, with supercomputing and high-performance networks, we can do multi-modal, real-time imaging that involves complicated transformations from raw data. This is now confined to research, but it will lead to clinical applications.

In Earth science, we're understanding the geodynamo for the first time, why the magnetic fields reverse, because of computation. And we're on the verge of reliable storm-scale weather forecasts, which would be impossible without these computational advances.

From the CRAY X-MP, our first machine in 1986, to the TCS is a giant leap — 6,000 times more compute power with 40,000 times the memory. We're going to see new work in dynamic visualization technologies, event re-creation and simulated reality. We expect to see important new work with practical consequences in the area of power generation, where the simulation technology is ready to be a design tool that will improve the efficiency of power-generating turbines. A very small improvement in efficiency translates to significant reductions in the cost of power, easily billions of dollars over years. This kind of work requires this new level of computational capability.

— Ralph Roskies, PSC co-scientific director

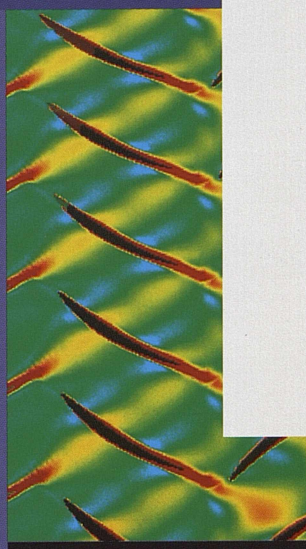

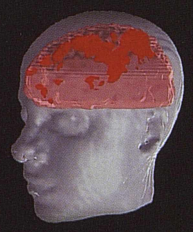

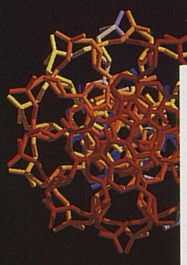

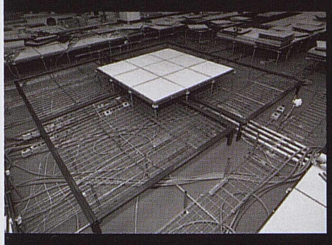

Research Notes & Highlights

WE ADD CAPACITY ON A REGULAR BASIS SO IT'S THERE BEFORE IT'S NEEDED.

NETWORKING THE FUTURE

As a resource for networking know-how, PSC's team of engineers has few peers. They provide engineering consulting for advanced networking nationally, and they conduct seminars that disseminate knowledge to engineers around the country. In projects such as Web100, they're actively involved in technology development. They are, in short, one of the leading groups in the world shaping the networks of the future.

Through the Pittsburgh GigaPoP, a high-speed network crossroads that serves Carnegie Mellon, Penn State, the University of Pittsburgh and West Virginia University, PSC provides advanced network resources for higher education and research. There are more than 20 GigaPoPs in the United States, and Pittsburgh's is among a select group. "In terms of the number of bits we can push, our installed infrastructure is matched only at a few other places," says Gwendolyn Huntoon, who directs the PSC team.

The GigaPoP connects all four universities to Abilene, a high-performance network linking more than 170 U.S. universities and research organizations. Data zooms along the Abilene backbone at 2.4 billion bits per second, fast enough to download the complete works of Shakespeare 436 times per second. With upgrades earlier this year, the GigaPoP link to Abilene improved fourfold to 622 million bits per second. "Demand for bandwidth is constantly increasing," says Huntoon. "What makes the GigaPoP unique is that we add capacity on a regular basis so it's there before it's needed."

Many research applications depend on high-performance networks. PSC staff are collaborating with scientists at the University of Pittsburgh Health System to develop technologies for "telemedicine," such as matching patient tissue samples against a database of already diagnosed tissue. At Carnegie Mellon, research on 3-D modeling of dynamic events, similar to but more sophisticated than the instant-replay technology at this year's Super Bowl, also requires high-speed networks.

More information: http://www.ncne.nlanr.net

University Corporation for Advanced Internet Development
Abilene

vBNS

THE PSC OPERATIONAL MANAGEMENT TEAM

(LEFT TO RIGHT): Bob Stock, associate director; Elvira Prologo, manager, administrative staff; David Kapcin, manager, financial affairs; Rich Raymond, manager, user support; J. Ray Scott, assistant director, systems and operations; Sergiu Sanielevici, assistant director, scientific applications and user support. Not in photo: David Deerfield, assistant director, biomedical initiative; Gwendolyn Huntoon, assistant director, networking; Janet Brown, manager, networking.

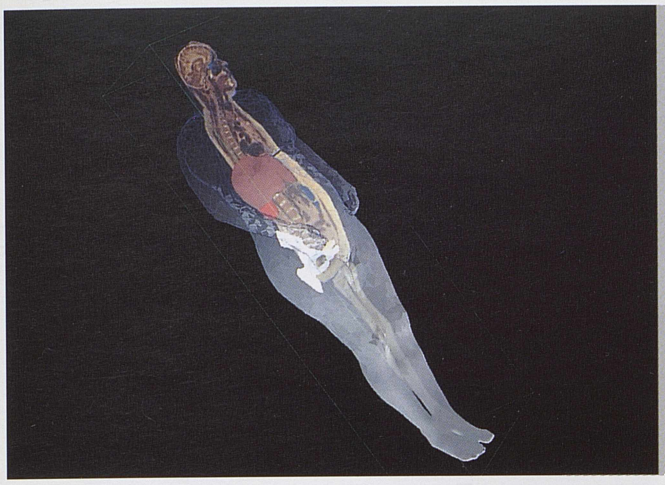

GETTING IN TUNE WITH WEB100

Most high-performance networks can transfer data at 100 million bits per second (Mbps) or faster. Why then do researchers who use them seldom realize rates above a few Mbps?

Good question, say network engineers at PSC, the National Center for Atmospheric Research (NCAR) and the National Center for Supercomputing Applications, who are doing something about the problem. With support from the National Science Foundation, they've mounted a research program, called Web100, to "tune" computer operating systems to better exploit available network bandwidth.

Most current operating systems have default configurations suited for low-bandwidth use, such as a home PC. But these settings often limit performance on high-bandwidth networks. And scientists shipping visualization data or interacting via a video-conferencing camera, for instance, need every available bit of performance.

The key to overcoming this limit is the Transmission Control Protocol, a "language" computers use to communicate across networks. With adjustments to TCP settings, network experts can "tune" the operating system to the network and optimize performance. Web100's goal is to eliminate the need for a human expert. They have refined TCP software in the Linux operating system to automatically achieve the highest possible transfer rate.

"Our goal is to make it easier for everyone to move data across networks at 100 Mbps or higher," says Matt Mathis, PSC network research coordinator. The Web100 team distributed the initial version of their software in March 2001. Forty-nine researchers at 26 institutions — including Stanford Linear Accelerator Center, Oak Ridge National Laboratory, Lawrence Berkeley Laboratory and Argonne National Laboratory — are testing this release.

In a related project called Net100, funded by the U.S. Department of Energy, PSC collaborates with NCAR, Lawrence Berkeley National Laboratory and Oak Ridge National Laboratory. Net100 expands on Web100 by including software to diagnose network status and feed this data to the auto-tuning capability of Web100.

TUNING THE VISIBLE HUMAN

Researchers at the University of Michigan have implemented Web100 as part of the Visible Human Project, a collaboration with PSC to develop new technologies for the National Library of Medicine's Visible Human data. The massive database, which fills more than 60 CD-ROMs, can now be viewed only in static 2-D cross-sections.

The Michigan-PSC team is developing tools to better discern 3-D shapes and spatial relationships. Using these tools through an easy-to-use graphical interface, the Visible Human will be more valuable as a training technology for medical students, and the Web100 software's ability to tune the network is essential.

In one of the enhanced capabilities implemented by PSC, a user may select a feature from a 2-D view, such as the liver section (ABOVE, RIGHT). By clicking on this section, the interface will outline the anatomical feature and automatically display a 3-D image (ABOVE, LEFT). This and other innovative capabilities will allow students to discern anatomical relationships that previous students could learn only by dissecting cadavers.

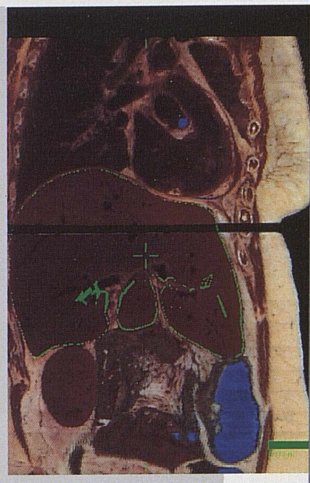

ROBOT CAMERAS FOR SUPER SUNDAY

This year's Super Bowl introduced an instant-replay technology, called Eye Vision, that let viewers see a replay as if time is frozen while a camera circles around the action. It's an impressive gee-whiz effect that can also help to resolve difficult calls. To develop Eye Vision, CBS turned to Carnegie Mellon University scientist Takeo Kanade, who in turn drew on PSC's parallel systems and visualization expertise.

Thirty cameras like the one shown here were mounted on remotely controlled pan-tilt heads custom-built by Mitsubishi Heavy Industries. A team of Carnegie Mellon and PSC scientists and engineers developed software to control the system. A human-operated master camera recorded pan-tilt angle, focus and zoom and fed this data to a central computer, which computed a control signal for each of the other cameras. All 30 cameras then simultaneously recorded an image and sent it to very fast videodisc, one for each camera. Along with the ability to select among 30 different angles for replay, the system also made it possible to cycle through all 30 discs, creating the illusion familiar to movie-goers who saw "The Matrix" — time radically slowed down as the viewpoint revolves in space.

Research Notes & Highlights [continued]

NATIONAL LEADERSHIP IN COMPUTATIONAL RESOURCES FOR BIOMEDICAL RESEARCH

BIOMEDICAL SUPERCOMPUTING

"For over 12 years, the Pittsburgh Supercomputing Center has provided national leadership in applying advanced computational resources to biomedical research," said Michael Marron, associate director for biomedical technology at NIH's National Center for Research Resources, last year when NCRR awarded $8.6 million to renew PSC's program in biomedical supercomputing. Since the program's inception in 1987, when it became the first extramural biomedical supercomputing programs in the country funded by NIH, PSC biomedical scientists have brought leading-edge computational resources and expertise together with experts in biology and medicine to solve important problems in the life sciences.

The program has provided access to computing resources for more than 800 biomedical research projects involving nearly 1,800 researchers in 43 states and the District of Columbia. Among these are several projects featured in this booklet (pp. 20-27).

PSC's workshops on computational biology have trained more than 2,000 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. "Our training activities reach hundreds of biomedical scientists each year," says biochemist David Deerfield, who directs the PSC program. "Techniques we've developed are helping scientists nationwide cope with the explosion of genome data."

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. Its researchers collaborate with scientists at many other institutions, including the University of Pittsburgh Medical School, Carnegie Mellon University, Scripps Research Institute, University of California at San Francisco, and Whitehead Institute.

More information: http://www.psc.edu/biomed/biomed.html

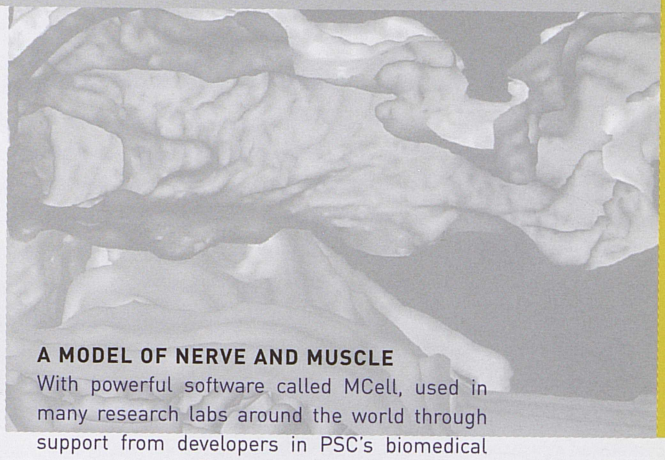

A MODEL OF NERVE AND MUSCLE

With powerful software called MCell, used in many research labs around the world through support from developers in PSC's biomedical group and at the Salk Institute, researchers can simulate the cascade of chemical reactions that occurs as molecules diffuse within, around and between cells. (See pp. 26-27.) This image represents a large-scale MCell reconstruction of a nerve terminal (gray) at the junction with the complexly folded membrane surface of a muscle cell (purple).

To build physiologically realistic models of the boundaries between cells and the space where they interact, MCell begins with scan data from electron micrography or tomography. From these datasets, the modeler traces membrane contours. By interpolating between contours of the physical sections, the model fills in 3-D geometry and adds color to define surfaces.

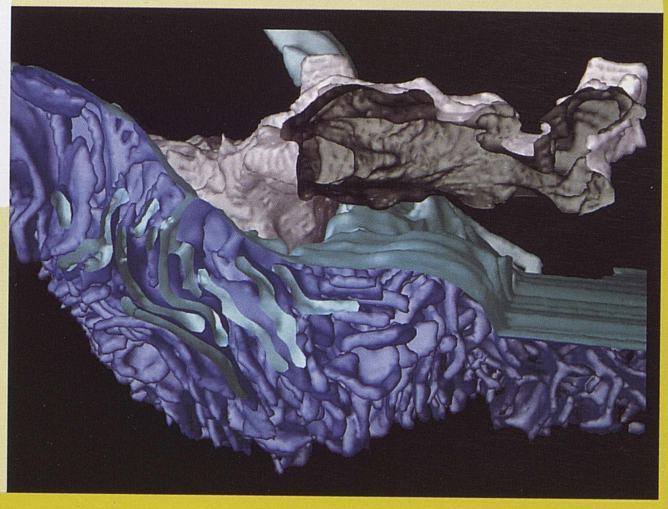

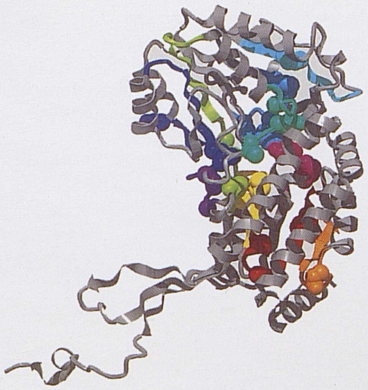

THE PROTEIN FAMILY TREE

How can the flood of human genomic data be gainfully employed? The first step is using sequence data to determine 3-D protein structure. But how can scientists wade efficiently through this massive information to solve hundreds of thousands of proteins that remain unsolved? A series of software tools, some of them developed at PSC, make it possible to search databases and classify proteins in family groups that reflect the evolutionary relationships that select for protein function.

PSC and University of Pittsburgh scientists have exhaustively analyzed the relationships among a family of enzymes called aldehyde dehydrogenase. Found in nearly every living thing, ALDH in mammals protects the body from toxic compounds. In this graphical representation of the ALDH molecular structure (ABOVE), colors represent amino-acid groups that are "highly conserved" — they remain essentially the same in nearly all species of the enzyme.

The logic of evolution holds that conserved residues are important in structure and function, and PSC scientists have developed algorithms to use conserved residues to identify sequence elements in related proteins that can predict crucial elements of 3-D structure and function. Glutathione S-Transferase is a very large enzyme family that protects cells from chemical toxicity. Using sophisticated sequence-alignment techniques, PSC researchers have classified GSTs into six subfamilies. This graphical representation of GST structure (BELOW) identifies molecular features that distinguish the subfamilies and predict the specificity of their biological function.

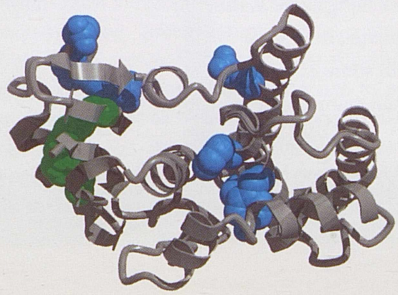

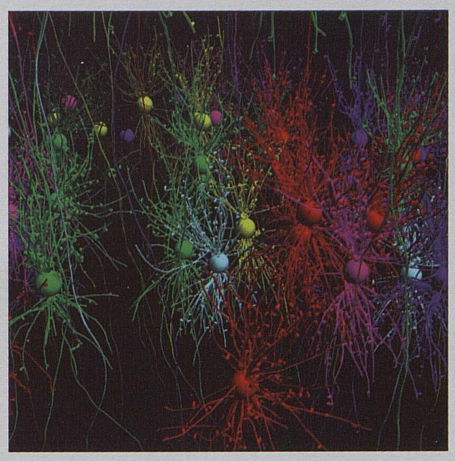

BRAINY SIMULATIONS

How do the human brain's trillion cells regulate and control how people perceive, learn, reason, communicate and act? Integrating research on the brain, probably the most complex system in existence, is the goal of the Human Brain Project, a multi-agency research initiative coordinated through the National Institutes of Health.

As part of this effort, researchers at PSC and the University of Edinburgh, Scotland are developing NEOSIM, a set of efficient, portable software tools for large-scale modeling of the nervous system. NEOSIM provides a framework to integrate computational models from the sub-cellular to whole-brain level. To effectively exploit advanced parallel computing, NEOSIM uses discrete event-simulation techniques to optimize the teamwork of many processors working on separable parts of a large-scale task. This image represents two layers of simulated neurons, with color indicating how NEOSIM would distribute them over six processing nodes.

PITTSBURGH SUPERCOMPUTING CENTER WORKSHOPS (2000-2001)

- Parallel Programming Techniques
- Nucleic Acid and Protein Sequence Analysis
- Building Computing Clusters for Biomedical Research
- Structure Determination Using Nuclear Magnetic Resonance
- Realistic Microphysical Simulations Using MCell
- Methods and Applications of Molecular Dynamics to Biopolymers

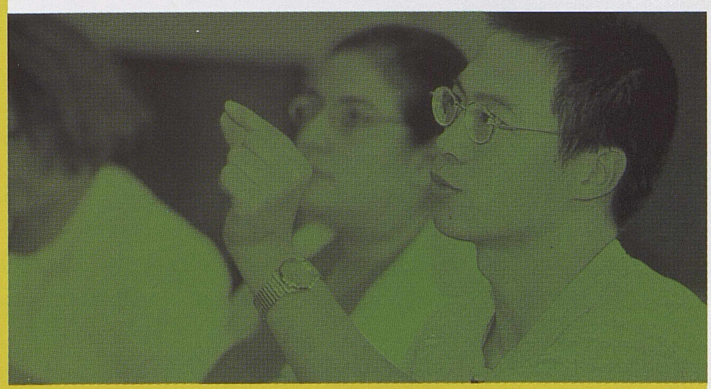

A workshop in progress at the PSC Computer Training Center.

Supercomputing in Pennsylvania

WITH SUPPORT FROM PENNSYLVANIA, PSC PROVIDES EDUCATION, CONSULTING, ADVANCED NETWORK ACCESS AND COMPUTATIONAL RESOURCES TO SCIENTISTS AND ENGINEERS.

Beverly Clayton, PSC executive director, coordinates PSC's program to provide advanced training and high-performance computing resources for researchers in Pennsylvania.

WORKFORCE DEVELOPMENT

PSC workshops provide training for industry researchers as well as for university faculty and students. They include extensive hands-on sessions, either in PSC's Computer Training Center or at corporate and academic sites around the state. During the past year, PSC presented off-site workshops in Parallel Programming Techniques at Drexel University and Penn State.

PSC trains Pennsylvania students through undergraduate internships. Since 1986, over 400 students have been interns at PSC, and many have gone on to find jobs in Pennsylvania. During the past year, PSC employed students from Carnegie Mellon, the University of Pittsburgh, Indiana University of Pennsylvania, Penn State, St. Vincent College and the Community College of Allegheny County.

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 network engineers have supported Pittsburgh high-technology conferences, including the August 2000 IEEE High-Performance Distributed Computing Workshop. They also worked with Marconi and AT&T Broadband Network Sevices to provide wide-area network connectivity — with access to the Abilene high-performance network — for the August 2000 Internet Engineering Task Force meeting at the Pittsburgh Convention Center. PSC engineers are also consulting on infrastructure planning for the new Pittsburgh Convention Center.

At the Bechtel Bettis Atomic Power Laboratory in Pittsburgh, PSC prepared and presented the first of two customized technology briefing days, with another scheduled for early 2002. PSC consultants provided Bechtel Bettis staff with information on input/output, scheduling, monitoring and visualization systems needed to develop and manage a parallel distributed-computing environment.

PSC outreach includes presence at conferences and science fairs, such as the 2001 Pittsburgh Regional Science and Technology Festival, shown here on the right. Over 20,000 people attended this week-long event at the Carnegie Science Center. PSC's booth included 3-D animated models of molecules and galaxies, and intriguing demonstrations of some of the ways computational models illuminate problems in science and engineering.

RESEARCH IN PENNSYLVANIA

During the past year, more than 900 Pennsylvania researchers from 19 institutions used more than a million CRAY T3E processor hours through PSC's Pennsylvania program. Many Pennsylvania researchers also gained access to the prototype Terascale Computing System through the National Science Foundation allocation process. The two projects represented here, along with other Pennsylvania projects featured in this booklet (pp. 22-27), exemplify how supercomputing plays a role in public health and safety as well as industrial technology.

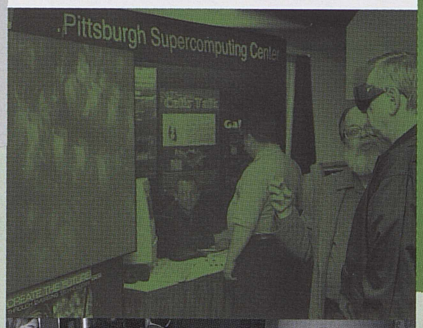

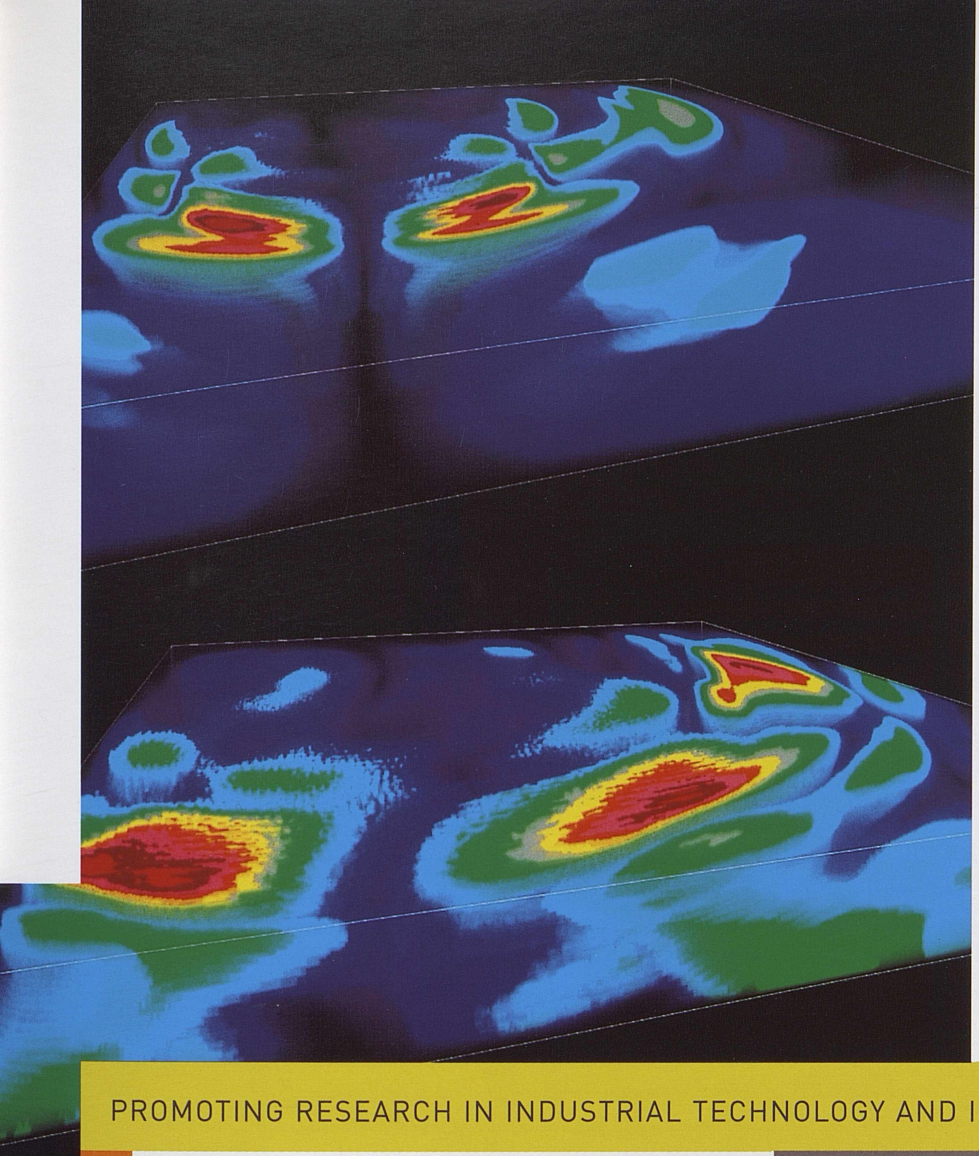

UNDERSTANDING EARTHQUAKES

Recent earthquake catastrophes have shown that the same seismic wave that flattens one block can leave the next unscathed. Local variations in factors such as basin shape and soil play a bigger role than previously thought. Through a major collaboration called The Quake Project, a team of engineers and computer scientists at Carnegie Mellon University and the University of California at Berkeley are working with seismologists from San Diego State University to better understand these factors. Led by CMU's Jacobo Bielak, the project aims to develop computational models that can reliably predict ground motion in large basins, information that can help in drafting building codes for safer, economical structures.

As part of this project, several earthquake research groups are working to jointly validate their modeling technologies. These visualizations represent Bielak's recent simulations pursuant to that effort. They model an earthquake in a vertical fault — called a strike-slip fault — typical of the San Andreas and other faults in California. The simulation represents the fault structure as a plane surface embedded within a 3-D volume. A rock layer, one kilometer thick, is supported on a stiffer rock mass. Color (increasing from violet to red) indicates the speed of ground motion in the horizontal direction. Comparing the first image (ABOVE) with the second (two seconds later in time) shows how the seismic wave travels rapidly away from the fault.

PROMOTING RESEARCH IN INDUSTRIAL TECHNOLOGY AND IN PUBLIC HEALTH AND SAFETY

UNDERSTANDING SUPERACIDS

University of Pennsylvania scientist Michael Klein and his colleague Dongsup Kim have used PSC's CRAY T3E and the prototype Terascale Computing System to develop quantum-level knowledge of superacids. Essential tools of the chemical industry, superacids have a powerful ability to break down raw petroleum, leading to such products as high-strength plastics and lead-free high-octane gasoline. Though industry uses hundreds of thousands of tons of superacid annually, there has been little understanding of how they work chemically.

Through simulations at PSC, Klein and Kim are shedding new light on why some superacids are stronger than others, showing that they form clusters with a ring-like structure, a previously unrecognized key to their chemistry. Their simulations also explained experiments showing that superacids conduct electricity better than can be accounted for by ionic diffusion, the normal process by which electrons in solution roam from ion to ion.

This graphic from Klein's computational study illustrates the details of a mechanism called "proton jumping" that underlies the electrical conductivity of superacids. In the strongest superacid, antimony pentafluoride, an excess proton (yellow) attaches to a fluorine atom (red), and they stay together as they "jump" along a linked chain of hydrogen-fluoride molecules (blue and green). The simulation shows that the proton and its fluorine partner travel across three bonds in 81.5 billion-millionths of a second (10^{-15} sec.).

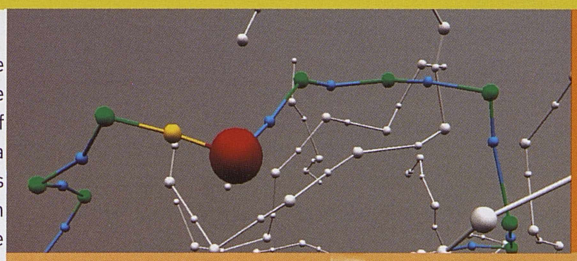

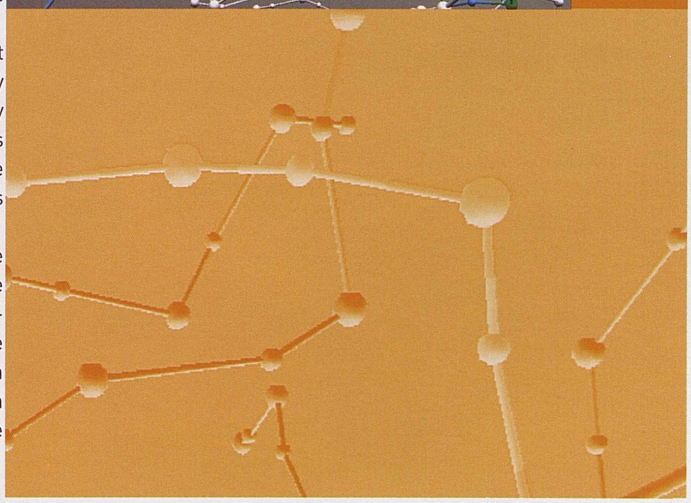

The Super Computing Science Consortium

PENNSYLVANIA-WEST VIRGINIA PARTNERS IN THE DEVELOPMENT OF CLEAN-POWER TECHNOLOGIES

Scientists at the National Energy Technology Laboratory and researchers at West Virginia University use PSC resources to address the critical challenge of developing clean, efficient fossil-fuel combustion. This research involves complex computer modeling in fluid dynamics, chemistry and geology. In recent work, NETL, the newest U.S. Department of Energy national laboratory, has relied on PSC resources to simulate fluidized-bed combustion and the use of lean-fuel mixes to develop clean and efficient next-generation power-generating turbines.

In January 2001, NETL awarded $2.5 million through the Super Computing Science Consortium to support energy research at PSC. Formed in 1999, the SCS Consortium is a regional partnership of NETL, PSC, Carnegie Mellon University, the University of Pittsburgh, West Virginia University and the West Virginia Governor's Office of Technology. The goal is to provide intellectual leadership and resources to apply high-performance computing and communication to problems in energy and the environment and to stimulate regional high-technology development.

"PSC has been an active partner with NETL in supporting regional initiatives," said Rita A. Bajura, director of NETL. "This agreement will make PSC's computing capabilities available to not only the region but the nation to further research in the efficient production and use of coal, oil, and natural gas — the resources that provide 85 percent of the nation's energy supply."

"We are very excited about the opportunities available through the Super Computing Science Consortium and the Pittsburgh Supercomputing Center," said John Weete, associate provost for research and economic development at West Virginia University. "Through this agreement WVU researchers will have new and enhanced tools to tackle complex problems requiring high-performance computing resources."

"Working with NETL, we've demonstrated that high-performance computing and communication are powerful tools for solving problems related to fossil-fuel combustion," said PSC scientific directors Michael Levine and Ralph Roskies in a joint statement. "By applying these technologies to the goals of the consortium, we are creating a stronger research and technology base for the region."

Additional resources will be provided through a "supercluster" of computers. Groups of linked computers at NETL, WVU and PSC will connect with each other via a high-performance fiber-optic network. With special software, PSC staff will coordinate and schedule the local clusters to act as a unified supercluster for advanced scientific applications.

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

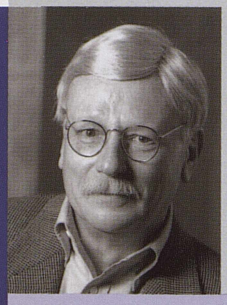

Jim Kasdorf, PSC director of special projects.

(SC)² Super Computing Science Consortium

On April 10, the Super Computing Science Consortium consortium broke ground for the EverGreen Technology Park in Waynesburg, Pa. With state and federal funding, this economic development project includes participation of all the school districts in Greene County. PSC has planned data network connections for EverGreen, to provide companies that locate there with access to PSC resources.

NATIONAL ENERGY TECHNOLOGY LABORATORY

"New advances in computer simulation technologies, in revolutionary new concepts for coal and natural gas fueled power plants, in carbon sequestration technologies, and in the production of ultra-clean transportation fuels and chemicals now allow us to envision a future in which the economic advantages of fossil fuel can continue to be enjoyed without environmental concerns."

— U.S. Secretary of Energy Bill Richardson

(FROM HIS DEC. 10, 1999 STATEMENT, DESIGNATING THE NATIONAL ENERGY TECHNOLOGY LABORATORY AS DOE'S 15TH NATIONAL LABORATORY)

SILANE REACTIONS IN A FLUIDIZED BED

An innovative technology with numerous potential applications, fluidized-bed combustion offers a highly efficient means to generate low-cost electricity from coal and other fuels with minimal environmental impact. It promises improved efficiency at substantially less cost to build than conventional non-fluidized bed combustion units used in power generation, and it offers a clean method to rely on coal as a fuel to meet growing energy needs.

Along with coal, fluidized-bed combustion can use waste materials as fuel — such as municipal trash and hospital medical waste — burning them to produce electricity while reducing the need for waste disposal. Because pollutants are incinerated or captured in the bed as they're generated, they aren't released to the air.

NETL carries out extensive programs of research in fluidized-bed combustion. Among the tools it has developed is simulation software called MFIX (Multiphase Flow with Interphase Exchanges). MFIX realistically models the gas-particle dynamics, chemical reactions and heat transfer processes that occur in fluidized-bed combustion.

In a recent MFIX study, NETL researchers used PSC's CRAY T3E to model the chemical reactions of silane, a chemical precursor of silicon, when injected into a fluidized bed. Carried out in collaboration with Dow Corning, this project investigated fluidized-bed technology as a production method for ultra-pure silicon, which deposits on the bed of particles as silane decomposes at high temperature.

This sequence of images, from an animation produced by PSC visualization specialist Greg Foss, depicts some of the simulation results. Color (decreasing from black through red to dark blue) corresponds to the fraction of gas versus solid material visible at a cross-sectional slice through the combustor. These results told the NETL researchers that the reaction occurs rapidly at the bottom of the combustor, where the silane is injected, and that "bubbles" of gas (black) tend to form along the walls of the combustor and migrate toward the center as they rise.

PROJECTS

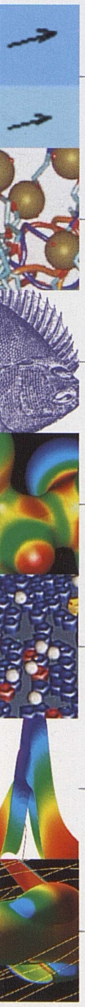

2001 CONTENTS

EVOLUTION_AND_STRUCTURE_OF_THE_UNIVERSE

In Search of Planetary Pancakes

There aren't enough stars in the Milky Way. Few of us notice, but it's a problem. We don't think about it when we look at the night sky, because there's 100,000 million stars in our Galaxy, which seems like enough. Still, physicists like Mike Norman want to know why there's not more.

The problem has to do with molecular clouds, huge gaseous conglomerations of matter — mostly molecular hydrogen (hence the name) along with helium and a trace of other elements. "You can think of molecular clouds as being a stage on the way to star formation," says Norman, an astrophysicist at the University of California, San Diego. "The more smoothly distributed gas in our Galaxy has to collect, become dense, and it has to cool. And just like hailstones form out of a thunder cloud when conditions are right, stars form out of a molecular cloud."

But something happens that physics hasn't yet been able to explain. "Given these big molecular clouds," says Norman, "which have up to a million solar masses of material — why are they forming stars at such a pitiful rate?"

The question, which astrophysicists have pondered for about 40 years, centers on a calculation called "free-fall time." Given the force of gravity and the observed density of molecular clouds, they should collapse into stars almost 10 times faster than they do. "You've got this enormous reservoir of gas," says Norman, "and you'd expect it to collapse in a million years and convert a large fraction of its mass to young stars. And it doesn't. Molecular clouds appear to be at least 10 million years old, and maybe 20 to 30 million years."

It's a perturbing gap in understanding, not least of all since the process of how stars form slides across the border from physics into biology. Carbon-based life, the only kind we know, requires planets, and planets are born from the aura of gas and dust swirling around young stars. The processes that underlie star formation and that, at least once anyway, led to a metal-laden planet with an oxygen-rich atmosphere are, in other words, intrinsic to the origins of life.

What holds up molecular clouds? That's the question in a nutshell. One idea has been that magnetic fields in the swirling, turbulent gas clouds are strong enough to counteract gravity. Or maybe the turbulence itself is enough counterforce, or some combination. Supercomputing offers a means to explore what's happening and fill-in the gaps between observation and theory, but it's a monstrously challenging modeling task. "Right off the bat," says Norman, "you're confronted with a three-dimensional, magneto-hydrodynamic, self-gravitating medium — a lot of adjectives. This is why it's an unsolved problem."

This spring, taking advantage of the prototype Terascale Computing System at PSC, Norman and his colleagues Pakshing Li of the National Center for Supercomputing Applications (NCSA), Mordecai-Mark Mac Low of the American Museum of Natural History and Fabian Heitsch of the Max Planck Institute for Astronomy mounted the most exacting simulation of star formation in a molecular cloud ever attempted. With resolution eight-times greater than heretofore possible, these computations provide clear pictures of gas clumps collapsing into dense, disk-like cores — the progenitors of stars. It's the first time dense cores with disk-like structure, resembling the pancake-shaped protoplanetary disks observed by the Hubble Space Telescope, have formed in a simulation.

PROPLYDS

A new word was coined in 1994 when the Hubble Space Telescope revealed swirling disks of dust around young stars in the Milky Way. Rice University astronomer Robert O'Dell, who discovered the disks, dubbed them "proplyds," for protoplanetary disks — the raw material out of which planets form.

Before the Hubble discovery, dust disks had been confirmed around only four stars. O'Dell and his colleague Zheng Wen surveyed 110 stars in the Orion Nebula, only 1,500 light-years from Earth, and found proplyds around more than half. All contain enough mass and necessary ingredients — carbon, silicates and other metals — to make Earth-like planets, and their abundance reinforces the belief that planetary systems are common in the universe.

It's an important discovery because proplyds, which are much harder to see than stars, have been a missing link in the story of how planets form, bridging the gap between molecular clouds and stars. "This gap is now starting to be filled observationally," says Norman, who did his Ph.D. thesis 20 years ago on star formation, "but it's never been filled in computationally, because there's such a vast range of scales that need to be encompassed — from interstellar distributed regions to protostellar concentrations, and we haven't had the computational capability."

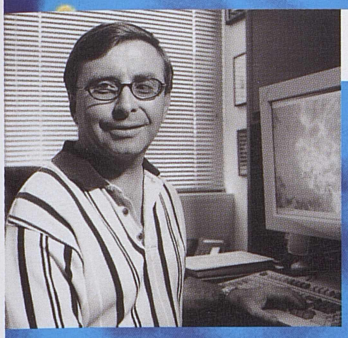

Mike Norman, University of California, San Diego

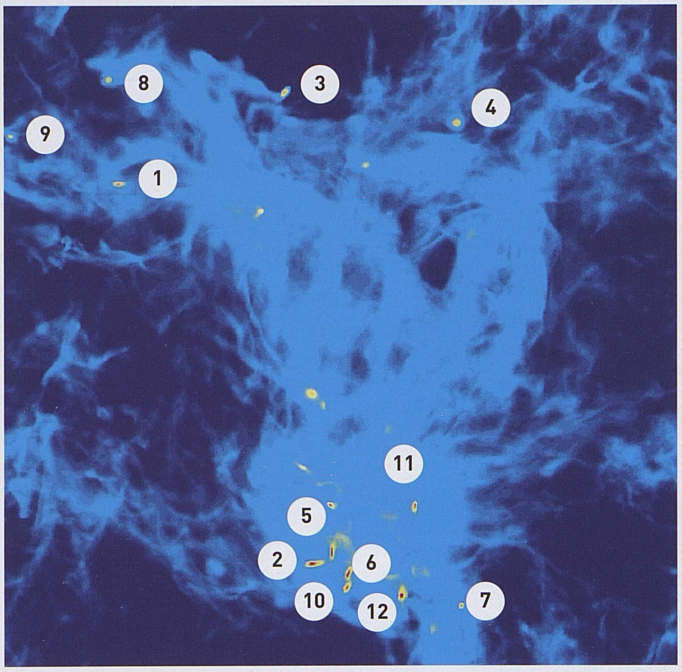

Density of the molecular cloud as simulated (increasing from dark blue to red) at one plane through the cubic grid. The simulation first establishes the molecular cloud as a turbulent medium. Then gravity is switched on. By 4.25 freefall times, the end of the simulation, 12 disk-like cores have formed (RIGHT). Grid resolution is insufficient to resolve further evolution of these cores.

THE MOST EXACTING SIMULATION OF STAR FORMATION YET ATTEMPTED REVEALS THE DISK-LIKE PROGENITORS OF PLANETS.

→ PAGE 17

O'Dell's observations include identifying features — frisbee-like spin of the disk material and jets of gas shooting out from the centers like the axis of a top — that provide a profile, in effect a picture of what to look for with simulations, which can add detail and physical understanding to the observations. "It's the picture we've had for years as a mental construct," says Norman. "As gas collapses out of the molecular cloud into a dense core, it conserves angular momentum, so that the core spins up into a disk around the forming star." With simulations at PSC, Norman evolved that picture further than has been possible till now.

RESOLUTION, RESOLUTION, RESOLUTION

Within the last two years, a series of simulations took aim at identifying the processes that hold up star-formation in molecular clouds. Using several SGI Origin 2000 systems, Mac Low, Heitsch and Dutch astrophysicist Ralf Klessen simulated gravitational collapse in turbulent molecular clouds using astrophysics modeling software called ZEUS, developed by Norman and his students at NCSA's Laboratory for Computational Astrophysics.

Essentially, these simulations put a supersonic turbulent medium with a magnetic field and gravity into a three-dimensional box. ZEUS sets the cloud into motion, governed by the known physics, and the simulation proceeds in time to see what happens. These large-scale Origin 2000 simulations indicate, first, that the turbulence of molecular clouds dissipates too rapidly to support them against gravitational collapse.

What about magnetic fields? "Turbulence decays rapidly," explains Norman, "but we think magnetic fields in molecular clouds don't dissipate. Like a lump under a rug, if you push on it here, it goes over there. Maybe in some parts of the molecular cloud the magnetic field is weak. Couldn't this be the way in which dense cores form?"

Heitsch, Klessen and Mac Low's most recent simulations looked at this question. Can pockets of gas form that become gravitationally unstable and collapse to form dense cores? The answer appears to be yes. The simulations showed cores beginning to form, but the 3D box — a cubic grid with 256 points in each dimension — lacked resolution to track the collapse far enough to see disk-like structure.

With this immediate background, Norman realized that the "friendly user" testing period on the prototype Terascale Computing System presented an opportunity. Preliminary runs using a new massively parallel implementation of ZEUS showed that the TCS ran three times faster per processor than the Origin system, with sustained per-processor performance of 500-600 million arithmetic operations per second. This performance, with full-system runs of 256 processors, made it feasible to double the grid-size — 512 points in each dimension, a total of 134 million gridpoints, an eight-fold improvement in the ability to see structural detail.

After a total running time of about 450 hours, the molecular cloud collapsed into 35 dense cores. Twelve of them show disk-like structure. Analysis of these 12 disk-like cores shows a magnetic field perpendicular to the disk plane that sweeps into a corkscrew at the spin axis on both sides of the disk, fitting with the proplyd picture of gas jets shooting out from the center. "The magnetic fields are bipolar," says Norman, "torqued into this helical structure that's often invoked, but never before produced in a self-consistent simulation."

It's the largest star-formation simulation ever done, and the results confirm that resolution can make all the difference in representing complex physical reality. Unlike proplyds, the disk-like cores in the simulation, says Norman, represent only the genesis of star formation. "Propylds are an end-point of the collapse of these cores to form new stars in protoplanetary systems, and our simulations don't go that far. Because of limited resolution, we have to stop well before proplyd disks are formed. But if you're going to make a star with a centrifugally supported disk of material around it, you're going to have to start with a collapsing core with some angular momentum. This is what we found."

For "Star Cores" the sequel, to see if these early-stage disk-like cores evolve further along the pathway to proplyds, stay tuned. With the full-scale Terascale Computing System, Norman plans to again double the resolution in each dimension, to 1,024 grid points. "The history of science is based on smart people attacking simple problems," he says, "leaving the hard problems for the next generation. The great thing about supercomputing is you can attack the hard problems."

More information: http://www.psc.edu/science/norman.html

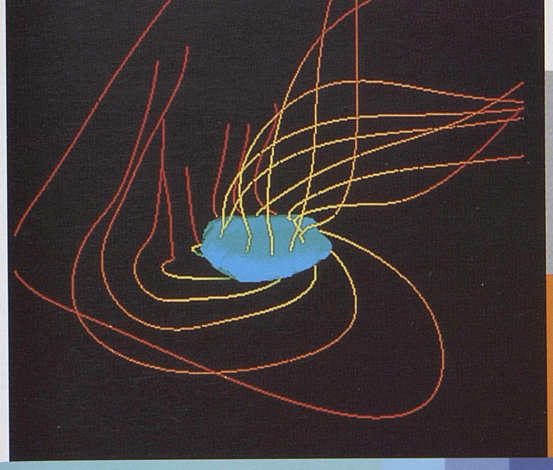

MAGNETIC FIELD LINES AND ROTATIONAL VELOCITY FOR CORE FIVE
Magnetic field lines (LEFT) indicate twisting from disk-like core rotation, with lighter color showing stronger magnetic field. For the velocity field (BELOW), vector length corresponds to magnitude and color represents density.

STRUCTURE_OF_PROTEINS_AND_DNA

Touchy Proteins

Run your fingertips across a tabletop. Is it hard? Can you feel texture? Our sense of touch gives us this information. But how? For that matter, how do we sense the gorgeous music of a symphony or the throb of an electric bass or cocktail conversation in the next room?

It may come as a surprise that for all of the above the answer is proteins. While it's well known that hundreds of thousands of different proteins are ceaselessly busy in the body, there's an important group we're only recently learning about in detail. They're called mechanosensitive channels, and they're found in all living organisms, where they reside in the membranes that form cell walls. Like other membrane-channel proteins, MS channels open and close to provide a pathway for molecules to exit or enter the cell. The flow of ions — calcium, sodium and others — through membrane channels creates electrical signals that regulate neural and muscular activity. The special trait of MS channels, however, is their ability to open in response to mechanical stress — such as the pressure of a fingertip on a tabletop or vibrations in the air — and thereby trigger neural processes like touch and hearing.

An estimated 30 percent of the proteins in cells are membrane proteins, yet we've been slow to learn about them because it's difficult to determine their structure. "While we know about 10,000 proteins," explains biophysicist Klaus Schulten of the University of Illinois, Urbana-Champaign, "only about 20 of these are membrane proteins — a very small fraction."

Until quite recently, none of these were MS channels. In 1998, however, structural biologists at Caltech determined the structure of a bacterial MS channel, called the bacterial large conductance mechanosensitive channel, or MscL. This ground-breaking work became the raw material for Schulten and physics graduate student Justin Gullingsrud to carry out a series of simulations.

"For this protein especially," says Schulten, "you want to know the dynamical function, how it opens — the details of what happens when you apply strain. The only way to learn this in detail is to take the structure and simulate its response to strain."

OPENING THE GATE

Experiments have yielded valuable information about MscL, but give no picture of how the molecular structure changes as strain on the membrane opens the channel. As a step in that direction, Schulten and his colleagues first constructed a computational model of a section of cellular membrane containing MscL. To realistically simulate the cellular environment, they "hydrated" the membrane — placing it within a bilayer of 7,387 water molecules — yielding a molecular system of 55,666 atoms.

Using the CRAY T3E, they first simulated this membrane-bilayer system with realistic conditions of temperature and pressure to allow it to "equilibrate," to fluctuate and find its natural state. Results revealed that MscL is structurally stable in its closed state, in agreement with experiments. The simulation also showed that the helices that form the narrowest part of the channel are its stiffest part, in agreement with experiments measuring the flexibility of different parts of the protein.

For this first simulation of an MS channel residing in a hydrated membrane bilayer, the equilibration computation alone was so demanding that the researchers chose not to put strain on the channel-membrane system. Instead, they simulated the protein alone, applying surface tension directly to MscL. Results showed an opening of 30 Angstroms as the channel helices flattened in response to pressure, in good agreement with experimental measurements of the "conductance pore."

UNCORKING THE BOTTLE

"Of course, this isn't totally satisfactory," explains Schulten. "We wanted to do the same simulation but with stretching of the membrane." For this, Schulten and Gullingsrud in early 2001 were among the first researchers to productively use the early model, 256-processor version of the Terascale Computing System.

They rebuilt the membrane environment, to accommodate the dynamics of adding strain, with 242 lipids and 16,148 water molecules, a system of 88,097 atoms. Their pioneering, very large-scale simulation confirmed the earlier results and looked more closely at the process of channel opening. The new simulations show a staged process and show, further, that the narrowest part of the channel functions like a cork in a bottle.

"First you get the widening of the protein," explains Schulten, "through the flattening of the helices, like an iris shutter on a camera, but the 'cork' is still sitting there. Then this opens up radially also, and the whole channel becomes accessible to the molecules that pass through it." The simulations give a detailed picture of which parts of the protein move and how they move as the channel goes from open to closed, details that elaborate upon experimental data and enrich understanding of this nascent branch of protein research.

More information: http://www.psc.edu/science/schulten2001.html

IN EARLY 2001, RESEARCHERS USED THE PROTOTYPE TERASCALE COMPUTING SYSTEM FOR A PIONEERING SIMULATION OF A MECHANOSENSITIVE MEMBRANE PROTEIN.

(RIGHT) Top and side view representation of MscL embedded in the hydrated cellular membrane. Colored rods represent helices of the ion channel, with colors corresponding to subunits of the channel's pentameric structure. Lipid molecules (light blue strands) that form the membrane include "headgroups" (green spheres) where they meet surrounding water.

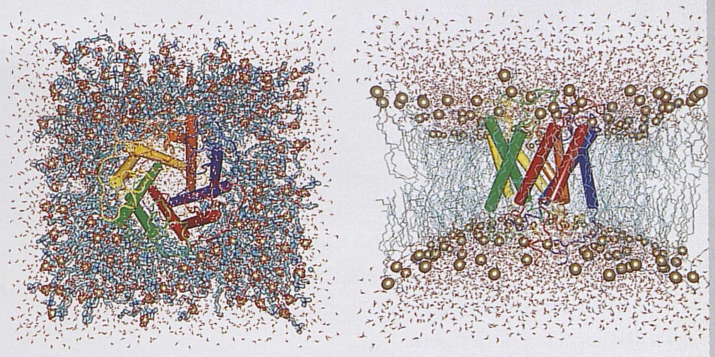

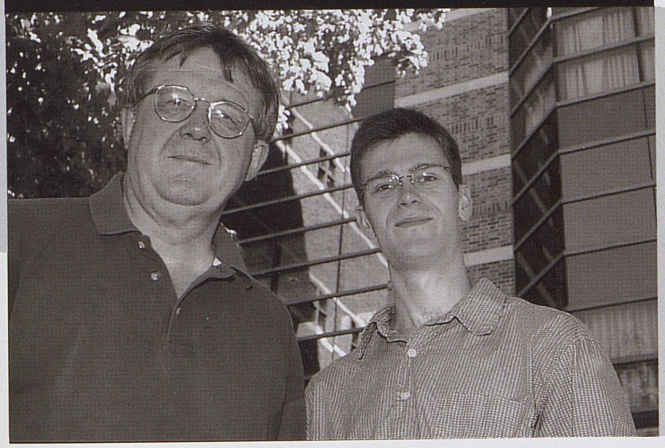

Klaus Schulten (LEFT) and Justin Gullingsrud, University of Illinois at Urbana-Champaign. Schulten directs the Theoretical Biophysics Group at the University of Illinois Beckman Institute for Advanced Science and Technology.

(BELOW) This image represents the full structure of MscL, including the helices shown extending below the membrane, which form the narrowest part of the channel. Simulations show that this sub-structure functions like a cork in a bottle, blocking the channel even as strain opens other parts of the passage until the cork finally lets go and the channel fully opens.

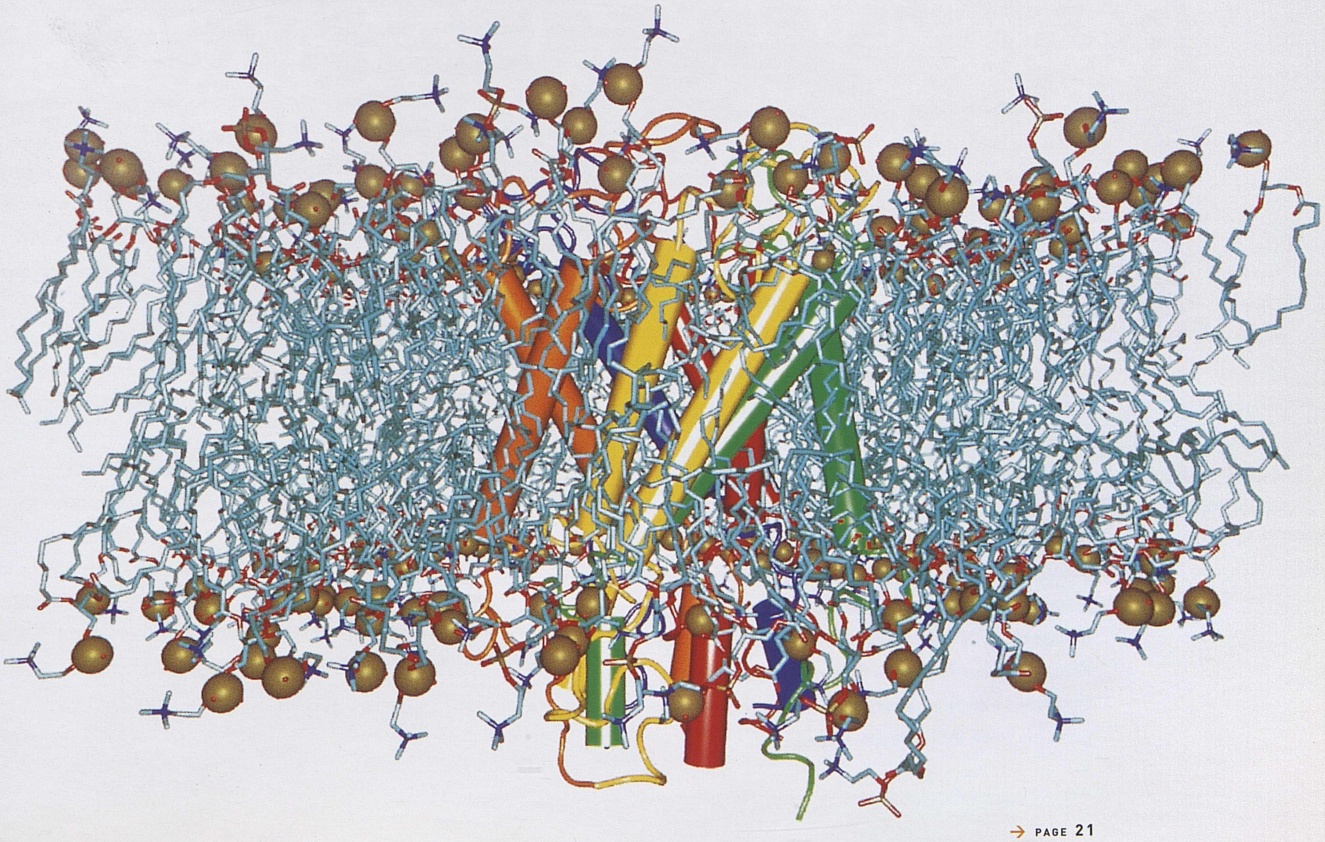

→ PAGE 21

STRUCTURE_OF_PROTEINS_AND_DNA

Fishy Proteins

From simple questions, marvelous and useful things sometimes transpire. In the 1950s, Canadian scientist P. F. Scholander wanted to know how Arctic fish can swim in water colder than the freezing point of fish blood. Eventually, his experiments showed that the blood of some northern fish contains "antifreeze." In the late 1960s, animal biologist Arthur DeVries investigated several Antarctic fish and isolated a protein that accounted for this antifreeze effect. By the 1990s, biotech companies began staking out patent rights.

So far biologists have found antifreeze proteins in Arctic and Antarctic fish, plants and the yellow mealworm beetle. While there are several variant types of AFPs, all show similar ability to interact with incipient ice crystals and inhibit crystal growth. This interaction lowers water's freezing point two to five degrees Celsius. Compared to other solutes, like salt, AFPs lower the freezing point hundreds of times more effectively.

Commercial possibilities are bounded only by the imagination. Scientists have already used AFPs to produce laboratory strains of transgenic tobacco, tomatoes and potatoes with improved resistance to frost damage. AFPs may eventually safely prevent freezer burn in ice cream, stop leafy vegetables like celery and lettuce from going limp in the fridge, reduce highway damage from de-icing with salt, and protect frozen blood from the effects of thawing.

What's not known is how AFPs work. "What we think," says Duquesne University biochemist Jeffry Madura, who has done AFP-related research for a decade, "is it modifies the shape of the ice crystal, and that modification inhibits further crystal growth." A prominent theory holds that this antifreeze effect arises from hydrogen bonds between the proteins and the ice. Recent experimental work by Madura's research colleague, A.D.J. Haymet at the University of Houston, however, refutes this hydrogen-bond formulation.

With Haymet's experiments as background, Madura last year turned to PSC's CRAY T3E, which he used to carry out the most realistic simulations to date of AFPs interacting with ice and water. His results support Haymet's findings and suggest an improved way to understand what happens as these proteins work their wizard-like effects.

TESTING THE HYDROGEN-BOND THEORY

For about ten years, scientists interested in AFPs have understood that these proteins accumulate where ice crystals form and that their presence inhibits crystal growth. The question has been what happens in this interaction that accounts for the inhibiting effect? Aside from pure science, the payoff is the ability to make synthetic AFPs that, more cheaply than proteins from living organisms, could be used in a range of applications.

The hydrogen-bond theory arose from experiments looking at how the AFP from winter flounder aligns with the ice-crystal surface. As a way of testing this theory, Haymet and his colleagues in effect turned off the hydrogen bonding in flounder AFP and tested to see if this mutant protein would also inhibit ice growth. His experiments showed no change. The modified AFP also had the antifreeze effect.

To further investigate this result, which negates the hydrogen-bond theory, Madura undertook a "molecular dynamics" simulation, a computational approach which, in effect, puts the molecules together and sets them in motion to go about their biochemical business while recording what happens. For Madura, the challenge was to set up his simulation to mimic the protein-ice interaction as realistically as possible.

ICE-WATER REALISM

Most prior AFP simulations have represented AFPs in interaction with ice alone, no liquid water. With computational practicality as a major consideration, excluding the water has been an acceptable approximation. The protein-ice interaction appeared to be the crux of the mater, and since the results in general have supported the hydrogen-bond theory, this approximation appeared to work.

Studies by Haymet, however, as well as computations by Madura show that ice crystals form in a complex region of transition between solid and liquid water. With access to the CRAY T3E, Madura constructed a simulation that included the ice-water interface region, about 15 Angstroms across — with 2,841 molecules of water in the ice phase and 5,484 molecules of liquid water — with a smoothly varying transition in density between liquid and solid. With the winter flounder AFP, the simulation included 25,000 atoms, larger than any prior ice-protein-water simulation.

Also in contrast to prior simulations, Madura for the first time represented the full ice-water interface with no imposed limits on the interactions. The first step was to allow the ice-water interface to find its natural equilibrium between solid and liquid. After this preliminary step, itself a major computation, Madura included the protein and collected data from 510 picoseconds (trillionths of a second) of simulated molecular dynamics, more than twice as long as prior studies.

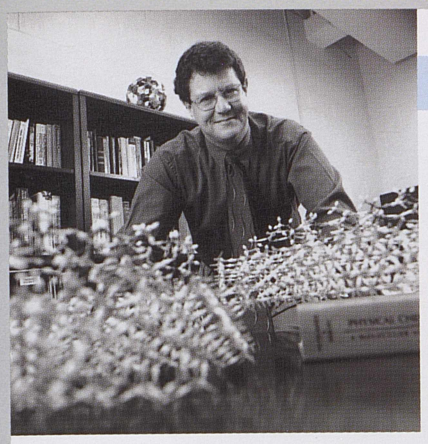

Jeffry Madura, Duquesne University.

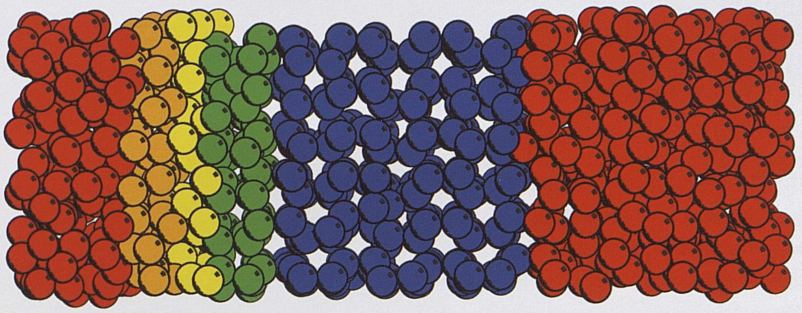

NEW SIMULATIONS HELP REVISE THINKING ABOUT ANTIFREEZE PROTEINS.

(BELOW) Snapshot of the winter flounder antifreeze protein simulated at the ice-water interface. The protein is a helical backbone (blue). Oxygen atoms (red) indicate ice and water molecules.

(ABOVE) This snapshot represents the ice-water interface region. Water molecules are shown as spheres centered on the oxygen atom with bulk liquid regions (red) to the left and right and bulk ice (blue) in the center. The left-hand interfacial region is highlighted with color gradation to illustrate layered variation in average diffusion coefficient.

The results support Haymet's experimental findings refuting the hydrogen-bond theory, and they show that accurate simulations of this interaction depend on realistically including the ice-water interface. If not hydrogen bonding, what interaction accounts for the protein's ability to inhibit ice-crystal growth? "It's vital that any alternative hypothesis account for the fact that the protein interacts with water molecules in between liquid and crystal form." With availability of PSC's Terascale system, he expects to get some answers.

More information: http://www.psc.edu/science/madura.html

STRUCTURE_OF_PROTEINS_AND_DNA

The Road to La-La Land

General anesthesia is one of the wonders of medicine. The ability to induce a deeply unconscious, immobile state makes possible life-saving procedures no one could have imagined 150 years ago, when surgeons were "saw-bones" and tooth extraction or amputation of a limb was excruciating beyond description. An estimated 15 million Americans undergo general anesthesia annually. Despite advances in pharmacology and wide use, however, no one can tell you how anesthetics do what they do.

"General anesthesia is one of the most important tools of medicine, and yet it remains mysterious," says Pei Tang, a physical chemist and assistant professor of anesthesiology and pharmacology at the University of Pittsburgh School of Medicine (UPSM). "Despite years of research, we still don't understand the molecular mechanism."

It's an alluring mystery. Understanding the molecular details would likely lead to better anesthetics, with fewer side effects, but perhaps even more importantly it should tell us something about consciousness itself, one of the grandly intriguing questions in science.

Tang is part of a UPSM team that uses a range of techniques to investigate how anesthetics work. In collaboration with Pittsburgh Supercomputing Center scientists Marcela Madrid and Troy Wymore, she has used computational methods to simulate how the drugs interact with the cellular membranes where they have their effect. The results of these studies challenge accepted thinking and offer support for an emerging new hypothesis.

A TALE OF TWO THEORIES

In the normal, undrugged state of consciousness, sense perceptions trigger a chain of events that releases electrical signals in the form of ion flow — sodium, calcium, potassium and other ions — through channels in the cell walls, better known to biologists as membranes. General anesthetics appear to exert their effect by changing ion flow through these membrane channels, either slowing it down or speeding it up. Thinking about how this happens falls generally into two schools: the lipid theory and the protein theory.

The lipid theory, the older point of view, says that general anesthetics work by their ability to dissolve in lipids, the fat that forms the cell membrane and seals it against the watery environs inside and outside the cell. Some experiments suggest that anesthetics make the lipid more fluid, a structural loosening that relaxes and changes the shape of the channels that control ion flow.

The protein theory is more recent and says, on the other hand, that general anesthetics interact directly with the channels, which are complex proteins, rather than indirectly through the lipids. Experiments have shown that a range of anesthetics have the ability to radically slow down flashes of light from a protein called luciferase, which makes the tails of fireflies glow. Luciferase doesn't exist in cell membranes, so in this case at least, the effect is direct. Still the question is how.

In experiments with membranes, as opposed to isolated proteins, it's extremely difficult to test the protein theory because it's nearly impossible to disentangle direct from indirect action. For that, computational simulations, which track the atom-by-atom details of the drug-membrane-channel interactions, have the potential to break through the theoretical logjam. "Only new techniques like large-scale simulations," says Tang, "that permit analysis at or near atomic resolution can test these theories."

LIPIDS, DRUGS & PROTEIN IN A WATER SANDWICH

Cellular membranes are complex molecular assemblies involving tens of thousands of atoms, and only in recent years has computational capability evolved to make it feasible to simulate these structures. Starting in 1999, as the first step in a staged process, Tang used PSC's CRAY T3E to construct and test a computational model of a cellular membrane called DMPC. The results gave structural parameters that agreed well with experimental data. For the next step, she computed the structural details and electronic properties of two frequently used general anesthetics, halothane and sevoflurane.

With this groundwork in place, Tang simulated 10 halothane molecules inside the DMPC membrane, which itself included a protein molecule, gramicidin, as an ion channel through the lipid bilayer. This very large-scale computation included 38,724 atoms and tracked the molecular movements for two nanoseconds (a billionth of a second) with a freeze-frame picture of the system every femtosecond (a millionth of a billionth of a second).

After about 240 hours of processing (on 128 CRAY T3E processors), the results show, contrary to accepted understanding, that halothane in the center of the lipid moves away from this hydrophobic environment toward the water. This result supports an emerging

RESEARCH IN GENERAL ANESTHESIA WILL HELP IN UNDERSTANDING THE BIOCHEMICAL NATURE OF SELF AWARENESS, MAYBE THE ESSENTIAL TRAIT OF BEING HUMAN.

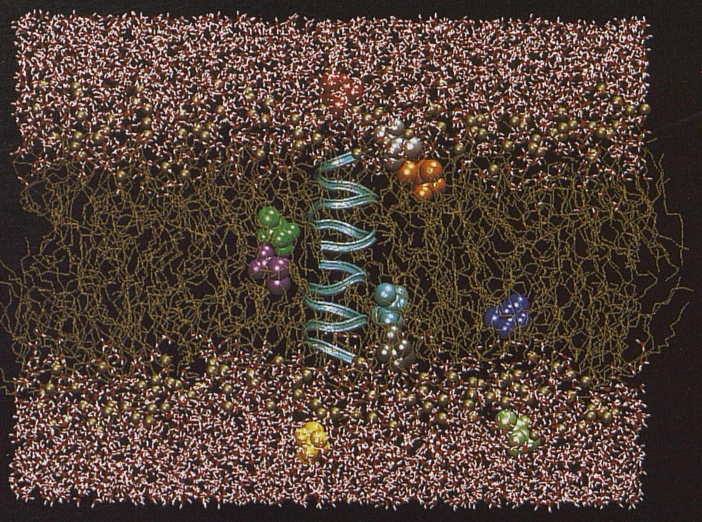

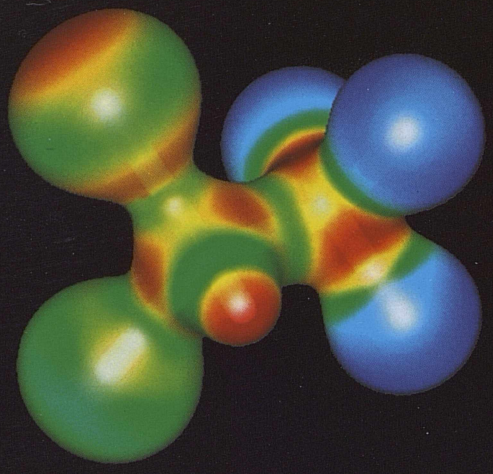

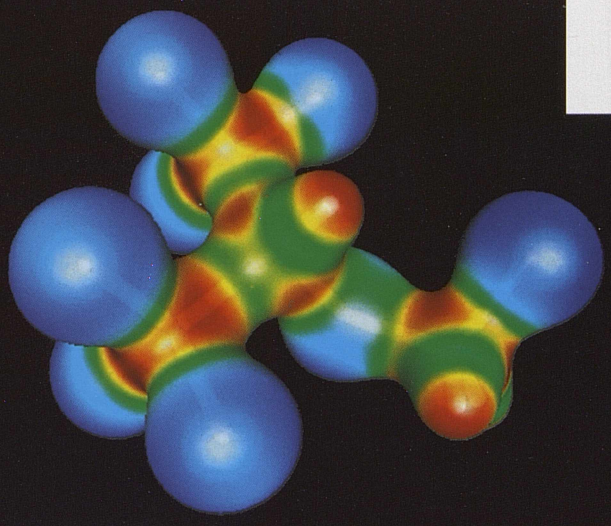

(LEFT) Halothane in membrane with channel. In this image from the simulation, ten halothane molecules (clustered balls) interact with gramicidin (blue coil), a protein inserted in the DMPC membrane as an ion channel. The simulation shows that halothane molecules move toward the entry of the channel.

(BELOW) Electrostatic potential, a measure of the replusive energy that electrons would feel at that point, mapped onto an electron-density surface for halothane (LEFT) and sevoflurane (RIGHT). Color (increasing from blue to red) indicates relative electrostatic potential. The skeleton of molecular structure is also visible.

hypothesis that bridges between the two competing theories and suggests that general anesthetics act at the channel-water interface.

To further test the interface hypothesis, Tang looks forward to PSC's Terascale Computing System. She'd like to simulate other anesthetics and extend simulation time into the millisecond range. She also plans to simulate compounds structurally very similar to anesthetics but that produce no anesthetic effect. These studies, she expects, will help pinpoint the molecular events that lead to general anesthesia. "With these simulations," says Tang, "I believe we'll be able to draw some conclusions that will lead us closer to solving this mystery."

More information: http://www.psc.edu/science/tang.html

Pei Tang, University of Pittsburgh School of Medicine.

CELLULAR_BIOLOGY

Where Nerve and Muscle Meet

From birth, John's muscles didn't work the way they should. At a month old, he couldn't hold up his "floppy" head. At six months, his eyelids were droopy. By age nine, he needed a wheelchair to get around.

Now 15, John — not his real name — still lives with the difficulties of his condition. His symptoms, progressive muscle weakness and fatigue, characterize a neurological disorder called slow-channel congenital myasthenic syndrome. But a series of lab tests on biopsies from his muscles revealed something puzzling.

The electrical current that triggers muscle cells to fire continued longer than it should before shutting off — the "slow-channel" response that gives SCCMS its name. But even at 16 months, the current was also very weak, like a much more advanced case. And the structure of his muscle cells at the nerve-muscle junction was inconsistent with the weak current. "What was unusual," says PSC senior scientist Joel Stiles, "is that he presented early in life with currents that were long and small, but the synaptic structure wasn't terribly abnormal."

SCCMS is a catch phrase for related genetic defects that lead to mutations in a protein, AChR, which plays a key role in muscle movement. A receptor protein in muscle cells, AChR works like a combination watchman and gate. It receives chemical signals from the nerve cells that tell it to open and flips the switch to become a channel for ions — mainly sodium — to flow into the muscle cell, a bioelectrical current that triggers movement. Lab studies showed a mutation in John's AChR different from any other documented case. Could this novel mutation lead to unprecedented AChR behavior that might explain John's unusual condition?

A medical doctor with a Ph.D. in physiology, Stiles specializes in computational neuroscience. With Thomas Bartol of the Salk Institute, he authored software called MCell that's used in two dozen labs around the world to simulate the microphysiology of nerve cells interacting with other cells. John's situation led Stiles to use MCell to test his hunch that the receptor was doing something other than just closing slowly.

From a series of simulations, he deduced that the AChR receptors were also slow to open. "The modeling led us to say 'Now wait a minute,'" says Stiles. "There has to be something beyond the classical picture of this disease. It has to have trouble getting into the open state as well." Lab work by Stiles's collaborators confirmed this insight into a previously unknown disease mechanism — invaluable new knowledge that can help in arriving at appropriate drug therapy as well as in research to develop new and better treatments.

SWEEPS IN PARAMETER SPACE

MCell is a powerful research technology that fills a gap between smaller-scale simulations that model molecules atom-by-atom and larger-scale approaches that model whole cells or groups of cells. It starts with a highly realistic 3-D model of the cellular space — reconstructed from electron microscope scans. It simulates molecular diffusion by means of algorithms that model random motion, and at each succeeding slice of time — often a microsecond or less — it uses statistical methods to test for all the possible reactions.

With an Information Technology Research grant from the National Science Foundation, Stiles and Bartol are collaborating with computer scientists at the Salk Institute, the University of California, San Diego and the University of Tennessee — under the direction of Francine Berman of UC San Diego — to expand MCell's usefulness. The idea is to create a "virtual instrument," software with an easy-to-use interface that can run on a "grid" of many different computers at the same time and keep track of thousands of simultaneous computations.

"What's embedded in this kind of work," say Stiles, "is the need to examine the parameter space the model sits in. Many unknowns go into this detailed a model — for example, the reaction between the neurotransmitter and the receptor and other proteins. If you have a model with 20, 50 or 100 input parameters, they may all need to be varied in some systematic way, and the modeler needs to examine how output varies with changes in input. This can lead to thousands of simulations going on simultaneously, what we call gigantic parameter sweep scenarios."

Stiles and Bartol and PSC systems engineer Stuart Pomerantz are working on MCell project-design issues and on a graphical user interface that will allow researchers to control these parameter-sweep simulations. At the same time, the team of computer scientists is developing sophisticated ability to "steer" them. With software that can monitor simulations in progress, the researcher will be able to shift midstream, with no loss of data, to refine parameters and test possibilities suggested by interim results.

More information: http://www.psc.edu/science/stiles.html

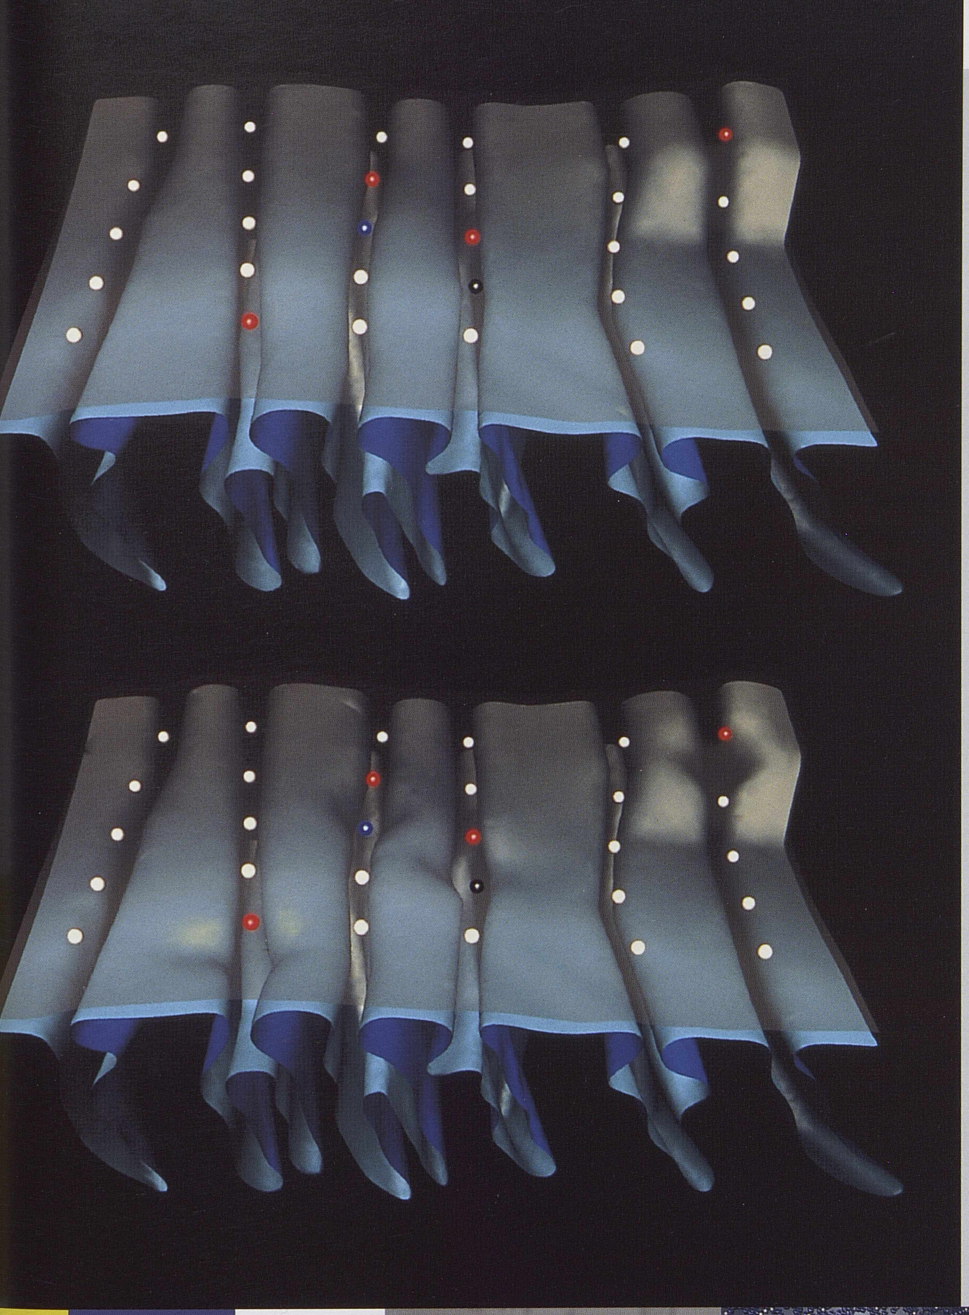

MODEL NEUROMUSCULAR JUNCTION
Normal architecture (TOP) contrasts with the focally deformed structure (BELOW). The localized depressions of the deformed structure represent the microphysiology of the novel SCCMS . In both cases, the model includes a nerve membrane (semi-transparent sheet) and 30 synaptic vesicles, five arrayed above each cleft in the muscle membrane. Four of these (red) are release sites for neurotransmitter molecules. Two other release sites (blue and black) are adjacent to focal deformations.

AN MCELL CLOSEUP
Neurotransmitter molecules are releasing from a vesicle (red) into a globally deformed neuromuscular junction, representing severe, late SCCMS. Another vesicle (black) isn't releasing. The simulation tests for neurotransmitters at two sampling regions (shown as boxes). The colored markers indicate AChR receptors, in five different states of interaction with the neurotransmitters, ranging from unbound (blue), to intermediately bound (red & purple), fully bound (green) and activated (yellow). Other markers (white and black) represent an enzyme, AChE. "In a normal muscle membrane," says Stiles, "yellow would predominate because a large fraction of the receptors would be open and making current."

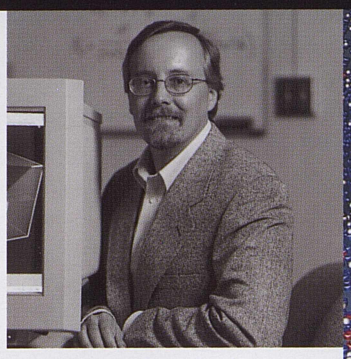

Joel Stiles, senior scientist,
Pittsburgh Supercomputing Center.

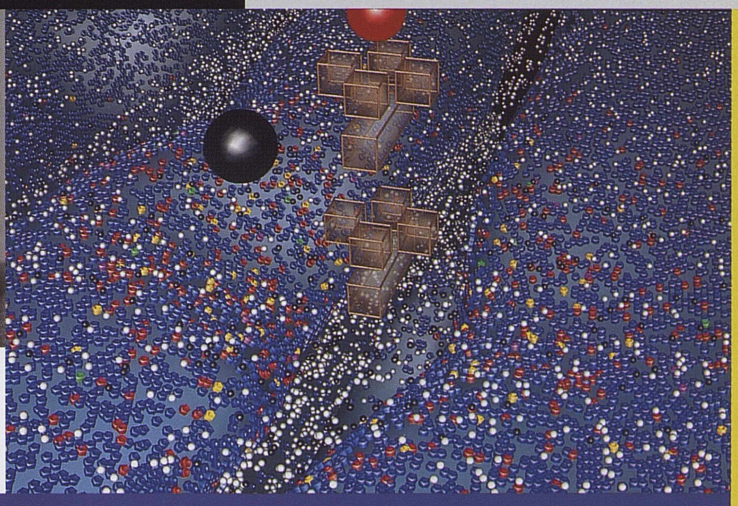

MODELING THE NERVE-MUSCLE JUNCTION IDENTIFIES A PREVIOUSLY UNKNOWN DISEASE PROCESS.

AT_THE_FRONTIER_OF_PHYSICS_AND_CHEMISTRY

Electronic Nirvana

In January 1986, Georg Bednorz and Karl Müller found that a novel copper-oxide compound chilled to 30 degrees above absolute zero allowed electricity to flow without resistance. Their discovery, which won a 1987 Nobel Prize, brought an exotic, quantum phenomenon into public consciousness and awakened a dream of technological nirvana — room-temperature superconductivity.

To transmit electrical current without the slightest loss of energy is magic without trickery, perhaps the closest thing to a free lunch Mother Nature offers. A material that is superconducting at room temperature would likely lead to high-speed trains that levitate on superconducting magnets, practical electric cars and superfast networks and computer chips.

Although room-temperature superconductivity remains an elusive quest, the 1986 breakthrough jump-started research around the globe that continues today. Within a few years, scientists found other copper-oxide materials and soon pushed the critical temperature, T_c, where resistance drops, well above 100 degrees Kelvin (100 K).

Along with furious laboratory efforts to find ever higher T_c materials, the 1986 breakthrough stirred intensive theoretical work. One leader in the field, David Pines, staff scientist at Los Alamos National Laboratory, says that understanding high-temperature superconductivity is "arguably the major problem in physics today" with thousands of published papers a year contributing to the effort.

"If we can arrive at a complete theoretical explanation of high-temperature superconductivity," says solid-state physicist Mark Jarrell of the University of Cincinnati, "then we should be able to design and synthesize a room-temperature superconductor, which would have tremendous technological implications."

Superconductivity is a quantum phenomenon in the solid state, and theoretical formulations to describe it depend on high-performance computing to solve the equations. The solid state, which includes metals, semiconductors and insulators, is a densely packed, regularly spaced lattice of atoms with electrons moving among them. The electrons and electron states that must be accounted for are, like fish in the sea, essentially infinite, and it's not possible, therefore, even with the most powerful supercomputers, to exactly calculate all the interactions that bear on the electronic properties. The theoretical challenge is to develop computational approaches that can reasonably approximate the complex physics and produce reliable predictions.

Within the past few years, Jarrell has developed an original approach that overcomes a serious limitation of another approach. Using the prototype Terascale Computing System at PSC, he and his colleague, post-doctoral fellow Thomas Maier, carried out computations with a theoretical model, the two-dimensional Hubbard model, that has gained general acceptance as a theoretical framework for high-T_c materials. Structurally, the high-T_c materials are a series of copper-oxide planes, with apparently no interaction between the planes, so they can be modeled as 2-D systems. Jarrell's computations indicate, nevertheless, that the 2-D Hubbard model is incomplete as a description of high-temperature superconductivity.

SOMETHING HAPPENING HERE

BCS theory, which described the electron-pairing phenomenon that underlies low-temperature superconductivity, doesn't work for high-Tc materials, although some version of electron pairs still appears to be the joy juice of the new superconductivity. In the words of a 60s song, "There's somethin' happening here. What it is ain't exactly clear."

Fifteen years of prodigious work on high-T_c materials has established that they constitute a new realm of solid-state physics. In virtually every respect, their normal state — behavior above T_c — differs markedly from conventional superconductors. The big job of finding a theory that pulls this exotic new solid-state world into a coherent picture has gone in many directions, but the most widely accepted starting place has been the 2-D Hubbard model.

"It's the simplest possible model you could construct," says Jarrell. Despite its relative simplicity, the 2-D Hubbard model has shown an ability to accurately calculate many of the strange properties associated with high-T_c superconductors, at least in the normal state. The main problem has come in finding a way to solve the Hubbard model under conditions that replicate the transition to superconductivity.

Solution of the Hubbard model for the infinite number of electrons in a solid-state lattice requires an approximation scheme. An approach called the Dynamical Mean Field has proven useful in many calculations, but is inherently inadequate for the high-T_c transition because it's "localized." It accounts for interactions between electrons at one atomic site, while other sites in the lattice are in effect averaged as a mean field. Studies have shown, however, that a fundamental characteristic of high-temperature superconductivity is that the pairing interactions are non-localized — electrons from neighboring atoms, rather than the same atom, interact strongly.

PAGE 28 ←

IMPROVED COMPUTATIONAL CAPABILITY INTRODUCED A DECISIVE CHANGE IN THE THEORETICAL PICTURE OF SUPERCONDUCTIVITY.

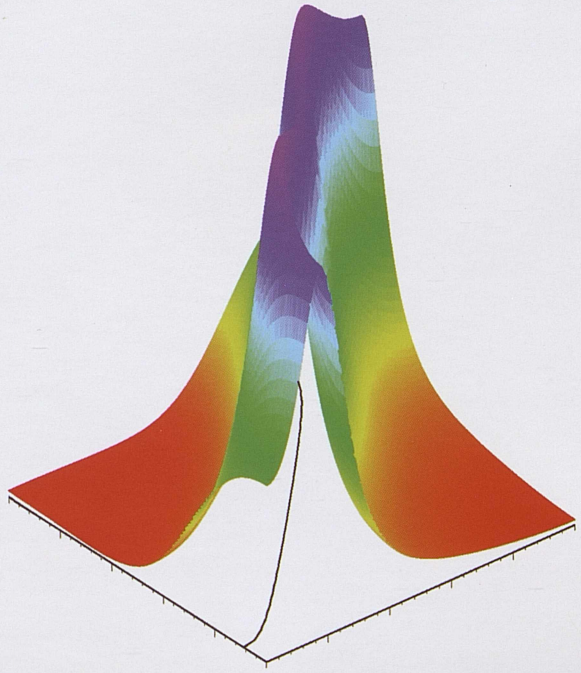

Spectral data of a cuprate superconductor in the normal (non-superconducting) state as calculated from the 2-D Hubbard model. These pictures from Jarrell's computations agree well with angle-resolved photo emission studies, an experimental method that measures electron distribution and maps the "Fermi surface" — a geometrical representation of a solid's electrical conductivity. The peak in the data (blue-violet) shows where there are many electrons at the Fermi energy. Agreement with ARPES results exemplifies that although the 2-D Hubbard model alone may not adequately describe the cuprate superconducting state, it describes many unusual properties of the normal state. Thomas Pruschke of the University of Augsburg produced these figures using software he and Jarrell developed.

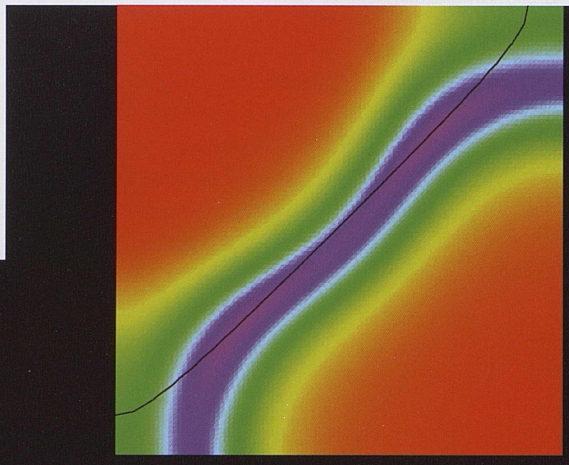

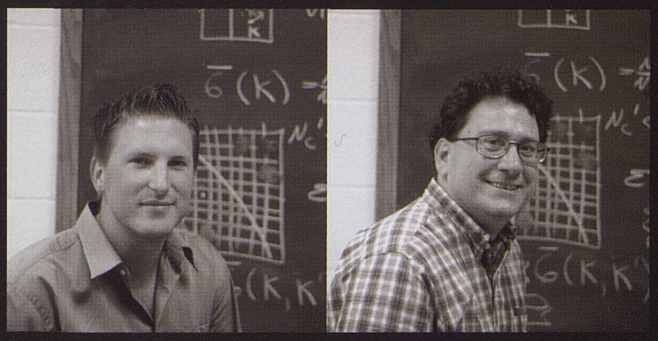

Mark Jarrell **(RIGHT)** and Thomas Maier, University of Cincinnati

REDRAWING THE MAP

A good deal of activity has gone toward developing non-localized extensions to the DMF approximation. Jarrell has developed a sophisticated approach, the Dynamical Cluster Approximation, that incorporates non-local corrections by mapping the problem onto a cluster of sites, which is itself embedded within the mean field. In 2000, he used the DCA approach to solve the 2-D Hubbard model on a CRAY T3E at Ohio Supercomputer Center. With a four-site cluster, the smallest possible, his results showed properties in good agreement with high-T_c materials, including transition to the superconducting state.

The computational load increases dramatically with larger clusters, and Jarrell was temporarily precluded from looking at the effect of larger cluster sizes. In spring 2001, he and Maier gained access to the prototype Terascale Computing System, facilitating a series of calculations with cluster sizes up to 16. The increase in computational capability introduced a decisive change in the theoretical picture. "As we systematically increased cluster size," says Jarrell, "superconductivity systematically went away. That tells you something fundamental. We believe that the 2-D Hubbard model is not sufficient by itself. Something new has to be introduced."

Among several possibilities, Jarrell notes that a fully accurate model of superconductivity may need to add coupling between copper-oxide planes. Another possibility is "chemical disorder," variation in the number of oxygen atoms from region to region. It's impossible to predict how much computing it will take to thoroughly explore these questions, says Jarrell, but the full-scale TCS will allow him to get started.

More information: http://www.psc.edu/science/jarrell.html

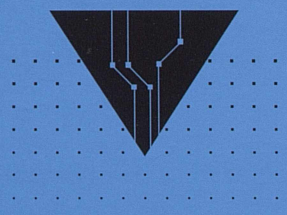

PITTSBURGH
SUPERCOMPUTING
CENTER

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

The Commonwealth of Pennsylvania
The National Science Foundation
The National Institutes of Health
The National Energy Technology Laboratory
Sandia National Laboratories
The U.S. Department of Defense
The U.S. Department of Energy
The Science Applications International Corporation
The Howard Hughes Medical Institute
Compaq Computer Corporation
Cisco Systems
Cray Research, Incorporated
Intel Corporation
International Business Machines Corporation
Microsoft Corporation
Silicon Graphics Incorporated
Sun Microsystems, Incorporated
Storage Technology Corporation
Dell Marketing L.P.
Gateway Companies Inc.
Hewlett Packard Co.

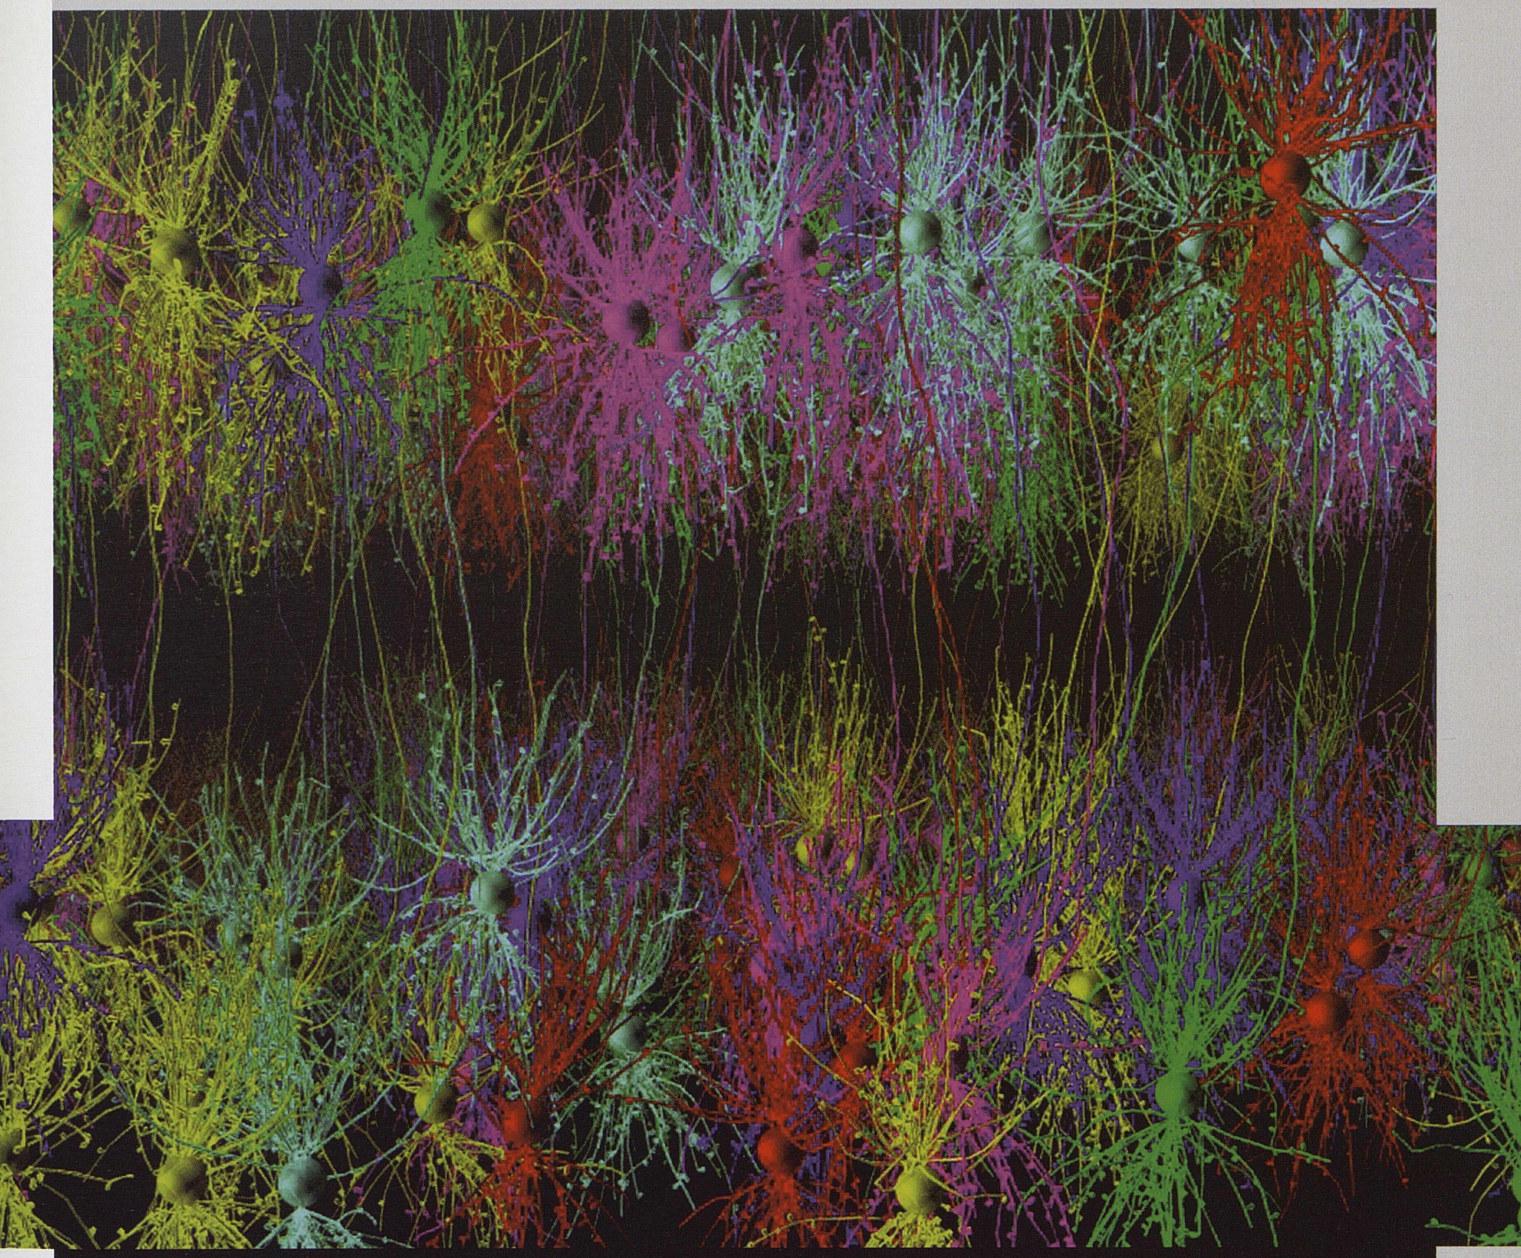

CREDITS

EDITOR/WRITER:
Michael Schneider, PSC

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

DESIGN:
Wall-to-Wall Studios, Inc.

PHOTOGRAPHY:
Photography & Graphic Services at Mellon Institute,
Michael Haritan, Jim Schafer, Robert Ruschak

PRINTING:
Hoechstetter Printing, Inc

♻ Printed on Strobe Paper, a premium sheet made with 10% post-consumer waste fiber, with vegetable-based inks.

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

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