1992

PROJECTS IN SCIENTIFIC COMPUTING

PITTSBURGH SUPERCOMPUTING CENTER

Projects in Scientific Computing 1992

Published July 1992 by the Pittsburgh Supercomputing Center, 4400 Fifth Avenue, Pittsburgh, PA 15213. (412) 268-4960.

The Pittsburgh Supercomputing Center is a joint project of Carnegie Mellon University and the University of Pittsburgh together with Westinghouse Electric Corp. It is pleased to acknowledge significant support from:

The National Science Foundation
The Commonwealth of Pennsylvania
The National Institutes of Health
The Defense Advanced Research Projects Agency
The Ben Franklin Technology Center of Western Pennsylvania
Buhl Foundation
Cray Research, Inc.
Digital Equipment Corp.
International Business Machines Corp.
Network Systems Corp.
Storage Technology Corp.
Sun Microsystems, Inc.
Thinking Machines Corp.
Ultranetwork Technologies
Bell of Pennsylvania

Credits

Editor: *Michael Schneider, Pittsburgh Supercomputing Center.* **Design & Illustrations:** *Steve Palmer, Pittsburgh.* **Writing:** *Steve Eisenberg & Michael Schneider, Pittsburgh Supercomputing Center.* **Science Writing Intern:** *John Lyczko, University of Pittsburgh.* **Photography:** *Jim Schafer (1), Bill Redic (4), Jeff Macklin (5-9), Michael Haritan (6, 8).* **Printing:** *Westinghouse Printing.* **Cover Masthead Design:** *Jamie Christner, Pittsburgh Supercomputing Center.*

Page layout & design were done by Steve Palmer at the Pittsburgh Supercomputing Center on a Macintosh Quadra 700 using PageMaker 4.2 desktop publishing software. These digital layouts included spot color, encapsulated postscript graphics, and tag image file format (TIFF) scientific graphics created from calculations by researchers at the Pittsburgh Supercomputing Center. Special thanks to Joel Welling, Anjana Kar and other members of the Pittsburgh Supercomputing Center graphics group for their assistance.

Cover: *A representation of hemoglobin, the protein that carries oxygen in the blood, produced by Joseph Lappa and David W. Deerfield II of the Pittsburgh Supercomputing Center biomedical group using "graphx," software they developed that creates representations of protein secondary structure from structural coordinates in the Protein Data Bank. The red balls represent the "heme" units, which give blood its red color. Hemoglobin has four protein subunit structures, and the hemoglobin helices are colored according to which of these four subunits contain them — there are two alpha units (lavender and blue) and two beta units (green and yellow).*

Back Page: *Joel Welling of the Pittsburgh Supercomputing Center scientific visualization group produced this image, a frame from the video animation of Alaa Eldin Ibrahim's fantasy space colony (see p. 11).*

Printed on recycled paper.

Projects in Scientific Computing 1992

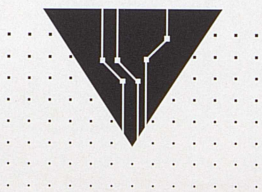

PITTSBURGH
SUPERCOMPUTING
CENTER

FOREWORD

This booklet briefly describes some of the developments in the past year at the Pittsburgh Supercomputing Center. Of paramount importance are the scientific advances which our users have made. Even the small sample that we can report on demonstrates that their significance is profound and their scope breathtaking.

We are in the midst of a technological revolution. New hardware and software capabilities are enabling new science. Several articles, including an industrial application, report on work carried out on our massively parallel Thinking Machines CM-2. Just as it has become a production platform, we are bringing into use the latest in massively parallel technology, the CM-5.

The paradigm of heterogeneous computing — bringing the most appropriate computational resource to bear on different pieces of a particular scientific problem — which we pioneered, has been harnessed to achieve new results in molecular biology.

Most of the scientific results of the past year were achieved on our CRAY Y-MP. Demand for Y-MP resources greatly exceeds its capacity. It will soon be upgraded to Cray's most recent product, the C90, close to six times as powerful as the Y-MP.

Michael Levine (left) and Ralph Roskies, scientific directors of the Pittsburgh Supercomputing Center.

The ongoing computing revolution, backed by new federal initiatives, is encouraging computational scientists everywhere to raise their sights and tackle new grand challenges, computational problems whose solution has significant social benefit. Among those discussed in this booklet are protein structure and drug design, how the heart and brain work, properties of new materials and saving the ozone layer.

The Pittsburgh Supercomputing Center is working with its sister institutions in the NSF Supercomputing Centers Program at Cornell, Illinois and San Diego to push forward the frontiers of high-performance computing and to support the national research community with the best tools for attacking these important challenges. Cooperative efforts, designed to improve national scientific capability, are burgeoning as the four centers are knitting together more tightly in what is becoming known as the Metacenter. These cooperative efforts will allow the community to benefit from the diversity of emerging architectures and to realize on a national scale the benefits of heterogeneous computing already illustrated in this booklet.

While the hardware and software are important for all these advances, nothing would have been accomplished without our talented and dedicated staff. We salute them for their efforts.

Michael J. Levine, Scientific Director

Ralph Z. Roskies, Scientific Director

Contents Projects in Scientific Computing, 1992

Pittsburgh Supercomputing Center 1992

PITTSBURGH SUPERCOMPUTING CENTER

What's new in supercomputing? Blink your eyes and you'll miss it. In this rapidly evolving field, last year's state-of-the-art hardware and software may be today's relic of times gone by. It's part of the Pittsburgh Supercomputing Center's job not only to keep up with these changes, but to lead the way, and this year the center continued its aggressive approach to the development of new technology with big steps forward on several fronts:

- The CRAY Y-MP C90, the most powerful vector supercomputer now available, is scheduled for October installation.
- An early model of a pioneering massively parallel system, the Connection Machine CM-5, was installed in January and has already run production codes.
- A new platform for fast scalar computation, a "SuperCluster" of 10 DEC DS5000 workstations, became operational.
- A new, multi-tiered high-speed network structure is providing the center's users with both speedier response and improved reliability.
- Developmental efforts in several areas, including file systems, video conferencing and documentation, in coordination with other National Science Foundation supercomputing centers, presage the advent of a "national machine room" linking the resources of all four NSF centers.

Established in 1986 with a grant from the National Science Foundation supplemented by the Commonwealth of Pennsylvania, the Pittsburgh Supercomputing Center is a joint project of Carnegie Mellon University and the University of Pittsburgh together with Westinghouse Electric Corp.

To date, more than 4,500 scientists and engineers at over 370 universities and research centers (red and green dots) in 49 states and the District of Columbia have used the center's computing resources to advance their research. This work has resulted in over 1,200 published papers in professional science and engineering journals.

Researchers connect to the center via regional electronic networks that feed into NSFnet, a high-speed pathway that links NSF supercomputing centers. Twenty-eight universities (red dots) are Pittsburgh Supercomputing Center Academic Affiliates (see back page). Representatives from these campuses form the center's main advisory body.

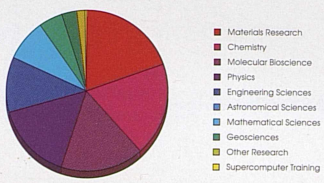

CRAY Y-MP Usage by Discipline

Pittsburgh Supercomputing Center
User Distribution

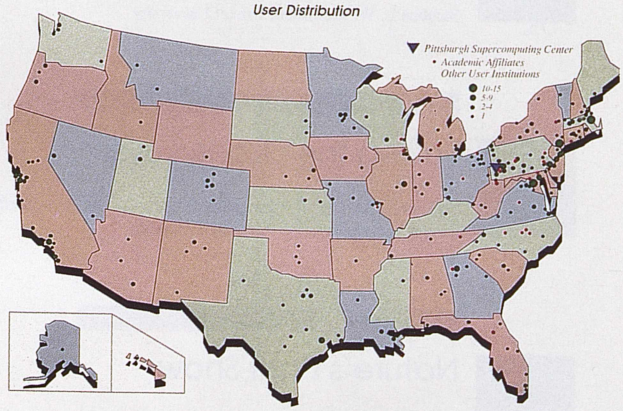

Bill Redic, © 1992

Jim Kasdorf, director of supercomputing, Westinghouse Electric Corp., and Beverly Clayton, executive director of the Pittsburgh Supercomputing Center, with the center's CRAY Y-MP at Westinghouse Energy Center. Kasdorf and his staff at Westinghouse maintain and operate the center's supercomputing equipment. In June, Westinghouse received a Computerworld Smithsonian award, an international award recognizing beneficial uses of information technology, for its role in support of university-based supercomputing.

To Clayton, impressive hardware is only one part of what makes the Pittsburgh Supercomputing Center tick; even more impressive, she says, is the talent and professionalism of its staff: "Hardware doesn't run by itself. More than the speed and memory of our machines, this center's success comes from the dedication, hard work and focus of its people. Our important work in software development, our outstanding user education, the stability and reliability of our networks and production systems, and our high school outreach program — all this and more must be credited directly to the effort of the staff."

The Newest, Mightiest CRAY

To meet the growing requirements of scientists and engineers nationwide for supercomputing, the center in November 1991 signed an agreement with Cray Research, Inc. to replace the center's Y-MP with Cray's newly announced Y-MP C90. The newest and most powerful of Cray's line of supercomputers, the C90 runs six times faster at peak speed and has eight times more memory than the Y-MP it will replace — which means users of the center's Cray will be able to tackle substantially larger problems in scientific research.

"We're now forced to turn down more than half the computing time requested on our Y-MP because the machine can't meet the demand," says Ralph Roskies, scientific co-director. "The C90 will encourage researchers to address problems of a range and depth they couldn't attempt before, simply because the computing to do the job wasn't available."

Delivery is scheduled for October, and as with the current Y-MP, Pittsburgh will be the first non-government site in the U.S. to receive the newest, mightiest Cray. The center and Cray also will collaborate on the development of new supercomputing technologies. "We have had a mutually beneficial relationship with Cray Research since we began," says scientific co-director Michael Levine, "and this agreement further enhances our ability to provide the most advanced computing tools possible to the research community nationwide."

Neither the hardware nor software of the CM-5 is in its final form yet, and this early model was installed in January so that the center's staff and users, in collaboration with the National Center for Supercomputing Applications (NCSA) in Illinois, could begin to explore its potential. "The true capabilities of this machine remain to be seen," says Levine, "and will depend on the ingenuity of our staff working in conjunction with Thinking Machines and the scientific user community on real applications."

Meanwhile, as the center is moving toward bringing the CM-5 into its existing heterogeneous computing environment, its initial venture into massively parallel computing, the Connection Machine, CM-2, has become a full-fledged production system. Over 800 researchers have used this resource to advance their work, a highlight of which is recent results achieved by John Joannopoulos of MIT.

Joannopoulos completed "first principles" calculations that map the full complexity of an extremely complex silicon crystal structure known as the 7 x 7 Takayanagi reconstruction. Running on 16,384 CM-2 processors, his code performed at the impressive speed of 600 million calculations per second (Mflops). This work, reported extensively in major science publications, portends the great potential of massively parallel computing for solving grand challenge problems in science.

Pioneering Massive Parallelism

The Connection Machine, CM-5, looks like the future. Its sleek, black, six-foot tall cabinets suggest the mysterious slab in "2001: A Space Odyssey." It is also a major step in massively parallel computing, moving beyond its predecessor, the CM-2, in a number of ways, but perhaps most significantly in its ability to be both a lock-step or SIMD (single instruction, multiple data) data-parallel machine, like the CM-2, in which all processors run identical instructions, and also a MIMD (multiple instruction, multiple data) machine, in which each processor can apply its own set of instructions to its own data.

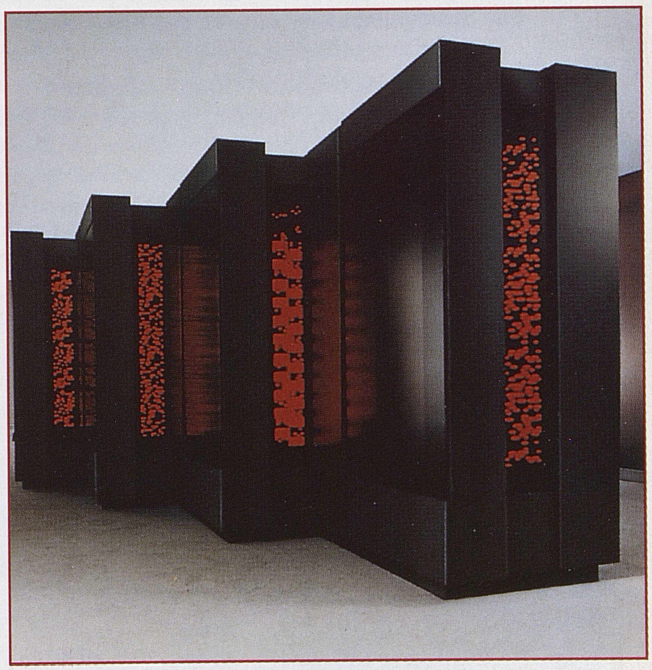

Jeff Macklin

David O'Neal (left), senior user consultant, and John Urbanic, parallel systems consultant, are serving as "point men" for users doing research on the center's CM-5. A group of high-energy physics researchers already is running production code for simulating hadron thermodynamics.

SuperCluster: A Third Form of Supercomputer

With its vector processing system, the CRAY Y-MP, jam-packed with usage, the center has looked for ways to ease the load. One of the answers is SuperCluster. With support from Digital Equipment Corporation, the center has acquired 10 DEC DS5000 workstations with over 18 billion bytes of disk storage. By linking these machines together with high-speed interconnections, the center can provide its users with a new and powerful platform for scientific computing.

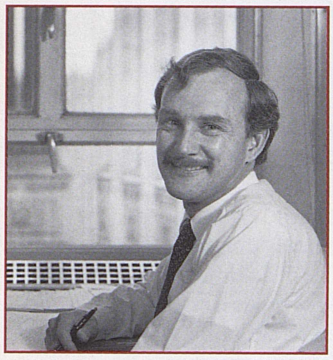

Jeff Macklin

Rob Pennington, research systems programmer, who leads the SuperCluster project.

"The need for this type of resource became apparent when we saw that a sizable fraction of the jobs on the Y-MP ran in the 10 Mflop range," explains research systems programmer Rob Pennington, "whereas highly vectorized codes operate in the hundreds of Mflops range." SuperCluster will take on these essentially scalar computing jobs with no appreciable loss in performance — freeing time on the Y-MP for jobs that make the best use of its powerful vectorizing capability. To users of SuperCluster, all the work of distributing processing among 10 workstations will be invisible, accomplished by networking software and the Andrew File System (AFS).

The Edge of High-Speed Communications

Over a year's work in planning, preparation and testing culminated in August 1991 when the center's high-speed network went into production. The three-tiered structure of the center's network is composed of the high-speed data transfer link (HiPPI) between the Y-MP and the CM-2, a primary production network path and a backup path. The primary path includes a T3 (45 million bits per second) connection to the National Science Foundation "backbone" network (NSFnet) and a high-speed fiber (100 Mbps) internal connection. The center's previous production path, a 1.5 Mbps link to NSFnet and 1.5 Mbps internal connections, remains in place as the backup path.

Michael Haritan

Wendy Huntoon, acting communications manager, and network engineers Matt Mathis (right) and Jamshid Mahdavi demonstrated the capability of the center's new network at Supercomputing '91, the annual gathering of the supercomputing industry and research community, held at Albuquerque this past November. Working with a group of equipment vendors, they implemented a demonstration network, called SCinet '91, modeled on the center's network. One of the coordinators of the project, Bill Boas of Ultra Network Technologies, wrote in Supercomputing Review *that conference participants viewed SCinet's performance as "the edge of communication."*

Jeff Macklin

Marvin Zalevsky, managing director of the Pittsburgh Supercomputing Center. "Our multi-tiered network structure has proven itself as an effective production tool, and our work on adapting the Andrew File System for use in a supercomputing environment has helped foster AFS's acceptance not only as the industry standard in workstation computing, but also its growing use in supercomputing. With these projects and others — DHSC, SuperCluster, our visualization and on-line documentation work — our center has demonstrated a commitment to developing the kind of practical tools that make supercomputing more broadly available.

"In these efforts, our staff is doing something much different than when we started six years ago. Then, the job was to implement systems that were used other places. Now, we're playing a highly visible, active role in the research environment in this country, at the forefront of this exciting and important technology — supercomputing. We're creating better ways of doing things, and our staff is showing we have the ability and dedication to break new ground. This isn't an easy thing to do, and I'm very proud of what we've been able to accomplish."

"The new network benefits our users in two significant ways," says Wendy Huntoon, acting communications manager. "First, the network hierarchy gives us a very robust network with no single point of failure. Second, users are seeing significantly less delay." Average delay over the old production network, she explains, was about 20 milliseconds; with the new network, the average delay is only four milliseconds.

The center's communications and user services staff continued work on the Distributed High-Speed Computing (DHSC) software developed in 1991. DHSC facilitates high-speed data transfer across the HiPPI that connects the Y-MP and CM-2 — making it possible (for the first time) to distribute computing jobs between these two systems. This year users began testing DHSC approaches to several new applications, including molecular dynamics (see p. 14) and computational fluid dynamics.

Innovation in File Systems & the Metacenter

The Andrew File System (AFS) is now implemented as a production file system on the Y-MP. "AFS provides a common file space accessible from all the computing platforms in our environment," says Janet Brown, production systems manager. "It has proven to be highly stable and functions smoothly. We are very pleased, and we have begun to encourage our CRAY users to become regular AFS users."

Other supercomputing centers have increasingly recognized the advantages of AFS, and the four NSF supercomputing centers have begun to evaluate it as the basis for development of an integrated national file system. This concept is a key feature in plans underway among the NSF centers for a "national machine room" or Metacenter. Users would be able to run their research code on whatever computer is most appropriate for the job, regardless of where their data happens to be stored.

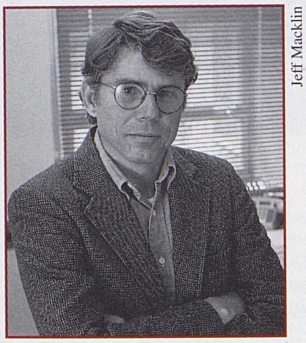

Jeff Macklin

Frank Wimberly, scientific applications coordinator, who oversees the center's work in applications software. He helped organize the center's research workshop in epidemiological modeling. "Part of what we try to do is to stimulate interest in applying supercomputing in areas where we perhaps haven't seen an appropriate level of activity."

Supercomputing in Biomedical Research

For the second time in three years, health-related research at the Pittsburgh Supercomputing Center has received the Forefronts of Large-Scale Computation award, the premier national prize for accomplishment in computational science. First given in 1989 to Gregory McRae's air pollution modeling, the 1991 award went to an interdisciplinary, multi-university collaboration led by University of Pittsburgh biologist John Rosenberg. Their Y-MP computations on how proteins recognize and bind with DNA provided detailed new information about this fundamental process.

It was especially satisfying to the center that the prize-winning collaboration developed directly from initial discussions at a 1988 workshop organized by the center's biomedical program. With funding from the Biomedical Research Technology Program of the National Institutes of Health, this is the only NIH-supported program of its kind in the country. The objective is to encourage the use of supercomputing in biomedical research, and educational activities this year included several workshops not offered before.

The Human Genome project funded a week-long workshop in August 1991 on "Nucleic Acid and Protein Sequence Analysis." Participants learned to use rigorous, high-speed sequencing techniques available at the center which can be valuable in mapping the human genome.

In December, the center invited fifteen leading epidemiology researchers to Pittsburgh. The goal: identify "bottleneck" problems in epidemiological modeling and assess how supercomputing might be used to overcome them — with particular focus on the AIDS epidemic. Brainstorming among the participants identified several problems, and a collaborative research effort was formed.

In April, the biomedical group used the teleconferencing facilities of the North Carolina Supercomputing Center to present a week-long workshop in biomedical supercomputing to 142 researchers. The teleconferencing capability allowed two-way communication with the participants, who were at 10 North Carolina colleges and universities as well as several government laboratories and private organizations.

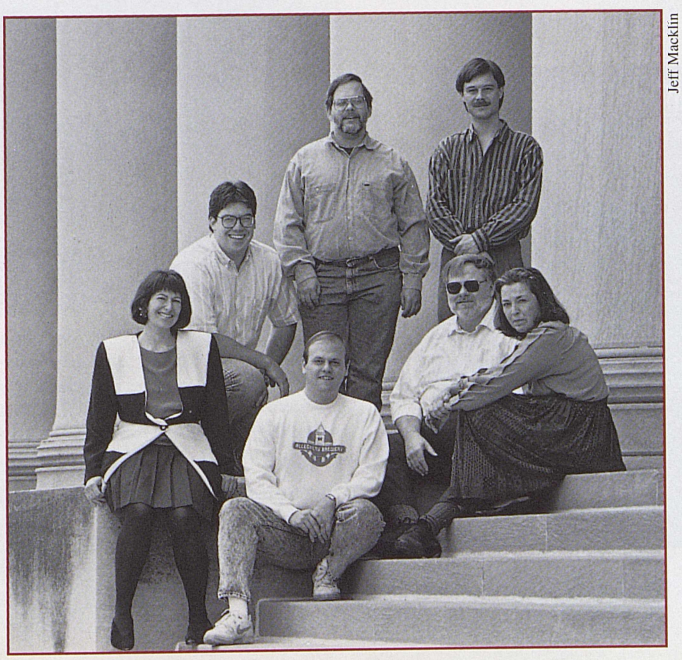

Jeff Macklin

The biomedical group of the Pittsburgh Supercomputing Center: clockwise from lower right, Nancy Blankenstein, Hugh Nicholas, Alex Ropelewski, Marcela Madrid, Doug Balog, David Deerfield, Nigel Goddard.

Supercomputing Education and Training

"The course was excellent: very organized and professionally done; good instruction; excellent availability of resources." This comment from a participant in the 1991 Summer Institute, an intensive two-week course in supercomputing attended by 23 science and engineering students and faculty, exemplifies the reputation for high-quality training that the center has earned. During the past year, Casey Porto coordinated the workshops and related programs that help spread the know-how of supercomputing — a task the center sees as a major part of its work.

Jeff Macklin

Casey Porto, education liaison, now leads the center's High School Initiative. This year alone, 17 center-organized workshops for researchers using supercomputing drew more than 511 participants. "We actively solicit comments and suggestions for improvement, and we take them seriously."

Supercomputing Goes to High School

To excite the imagination of high school students and teachers by giving them the chance to apply supercomputing in projects of their own choosing — this is one of the goals of The High School Initiative, the center's program of secondary education outreach, now beginning its second year. The program this year gave 10 teacher-student teams access to supercomputing through an on-site workstation donated by Digital Equipment Corp. Center staff analysts visited the high schools several times during the year, and each team worked with a mentor, drawn from university faculty and industrial researchers in the Pittsburgh area, with whom they communicated by electronic mail. Forty teachers and 150 students were actively involved, enriching their skills and gaining invaluable confidence in their ability to accomplish complex tasks. (MS)

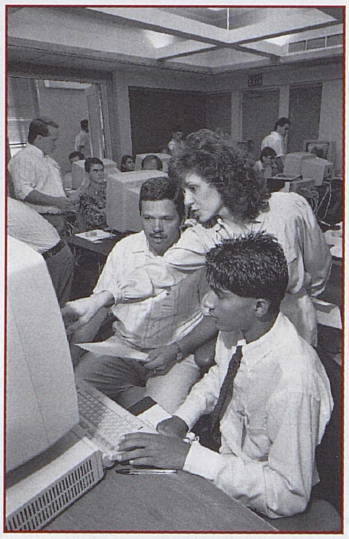

Michael Haritan

Lori Smith, staff consultant, makes a point with Charles Rodkey, science teacher, and Hitesh Soneji of South Fayette High School. Rodkey and Soneji used supercomputing to investigate Pittsburgh-area water quality.

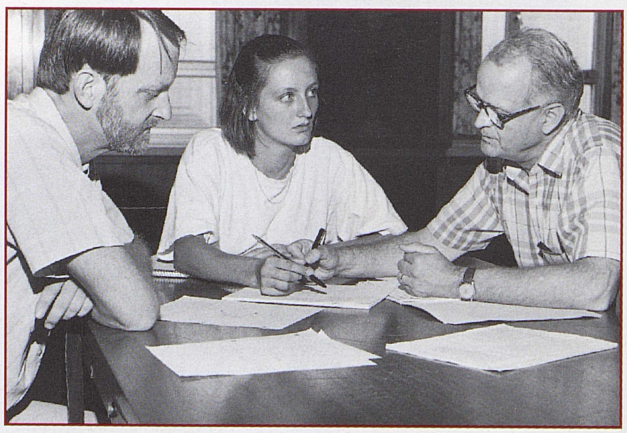

Michael Haritan

Robert Schumacher (right), physics professor at Carnegie Mellon University and mentor in the High School Initiative, confers with teacher Mark Walker and student Parke Wiegman of The Ellis School in Pittsburgh, on their project — the formation of rainbows.

Pittsburgh Supercomputing Center Workshops

(offered at least once between July 1991 & June 1992)

Biological Fluid Dynamics
Nucleic Acid & Protein Sequence Analysis
Fast Processes in Protein Folding Dynamics
Epidemiological Modeling Research
CRAY Supercomputing Techniques
Connection Machine Supercomputing Techniques
Connection Machine Techniques for Biomedical Researchers
Summer Institute of Supercomputing
GAUSSIAN Users Workshop

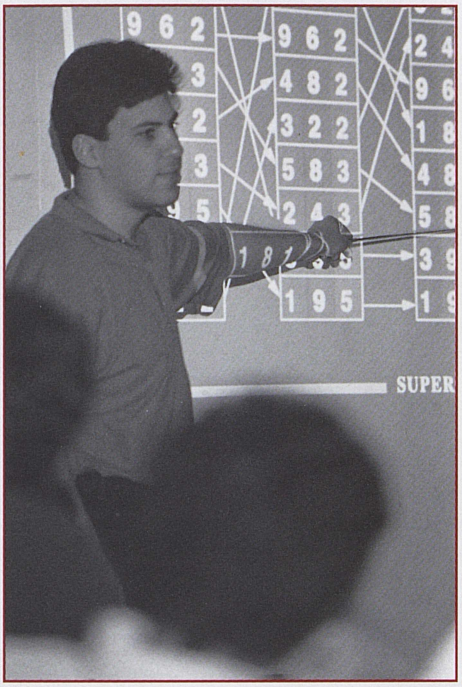

Jeff Macklin

Marco Zagha, fourth-year graduate student in computer science at Carnegie Mellon University, received the Supercomputing '91 award for best student paper for "Radix Sort for Vector Multiprocessors." Co-authored with Professor Guy Blelloch, the paper describes their algorithm for a generalized sort routine for the CRAY Y-MP. The routine runs five times faster on large problems than Cray's optimized library sort. Here, Zagha is explaining his sort algorithm at a "science for lunch" meeting at the Pittsburgh Supercomputing Center.

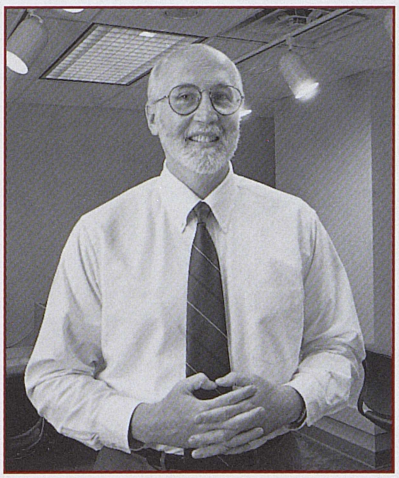

Jeff Macklin

Terry McGillen, corporate liaison. "To make the resources of this supercomputing center available to industry, not only the hardware, but the knowledge that goes with it, is a major aim of the National Science Foundation supercomputing program. Our corporate affiliates are getting some very valuable results (pp. 18 & 42) from their research here."

Information at the Pittsburgh Supercomputing Center

Proposals for Computing Time
Susan Lambl

Biomedical Initiative
Nancy Blankenstein

Workshops & Summer Institute
Alice Conniff

High School Initiative
Casey Porto

Corporate Affiliates Program
Terry McGillen

Newsletter and Documentation
Vivian Benton

These Pittsburgh Supercomputing Center staff members can be contacted by telephone (412-268-4960) or electronic mail. Internet: *lastname*@psc.edu. Bitnet: *lastname*@cpwpsca.

Scientific Visualization

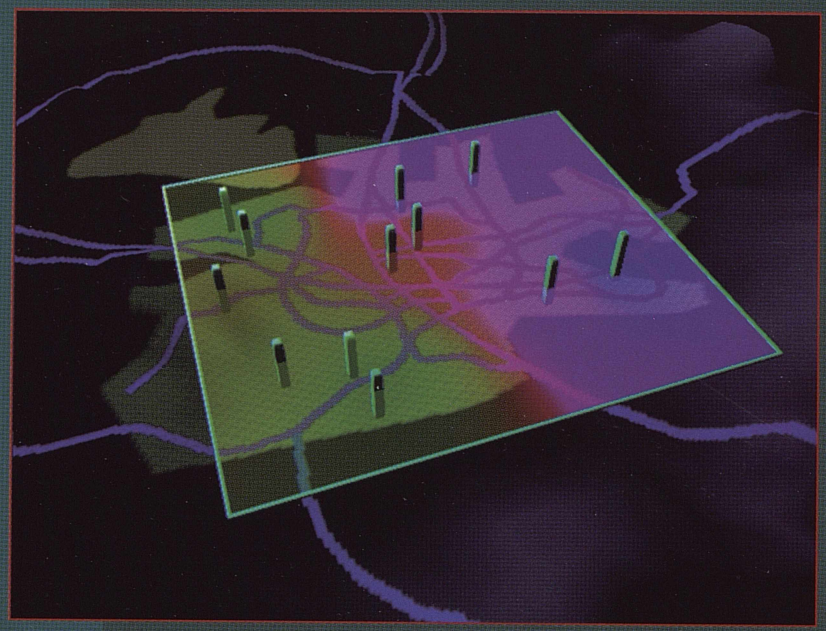

Ground Level Pollution in Mexico City

This image, from an animation produced by Chris Nuuja and Gregory McRae, shows how ground-level ozone varies across Mexico City. The data derives from hourly measurements at 12 locations denoted by the purple matchsticks. The height of each matchstick above the colored plane indicates the amount of pollutant at that point at a particular time.

The colored plane shows how pollution varies across the whole city, based on measurements at the 12 sites. When the ozone concentration at a sample point exceeds 0.1 parts per million, a matchstick pokes through the plane. "It's interesting to note," says Nuuja, "that the animation reveals short, pronounced surges in ground level ozone each day about 11 a.m., with little change in overall concentration throughout the rest of the day."

The Pittsburgh Supercomputing Center scientific visualization group: (l to r) Phil Andrews, Joel Welling, Chris Nuuja, Grace Giras, Anjana Kar. For its work on "Air Pollution in Los Angeles," a video produced from research by Gregory J. McRae of Carnegie Mellon University, the graphics staff won second prize at the 1991 Computer Animation Festival in Aizu, Japan. This animation, and a newer one showing air pollution in Mexico City (above), were included in "The New Explorers" episode of "The New Language of Science," a series produced by WGBH in Boston for public television.

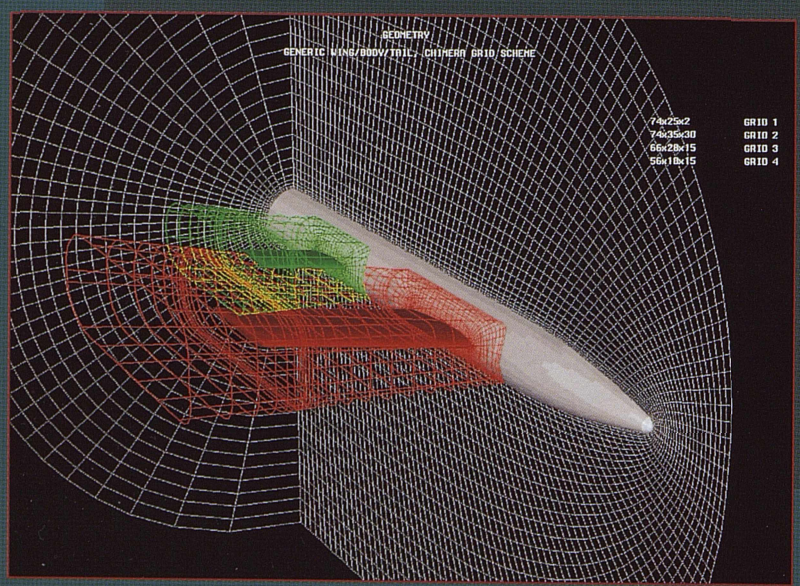

Airplane Surface with Computational Grids

Susheel Chitre, scientific programmer, generated this image using the P3D graphics libraries developed at the Pittsburgh Supercomputing Center. The dataset was produced with PLOT3D, a computational fluid dynamics plotting program created at NASA Ames. Originally, PLOT3D ran only on Silicon Graphics workstations. Chitre modified the code to use P3D libraries, so that it will now run on the CRAY Y-MP and display on any X-windows workstation.

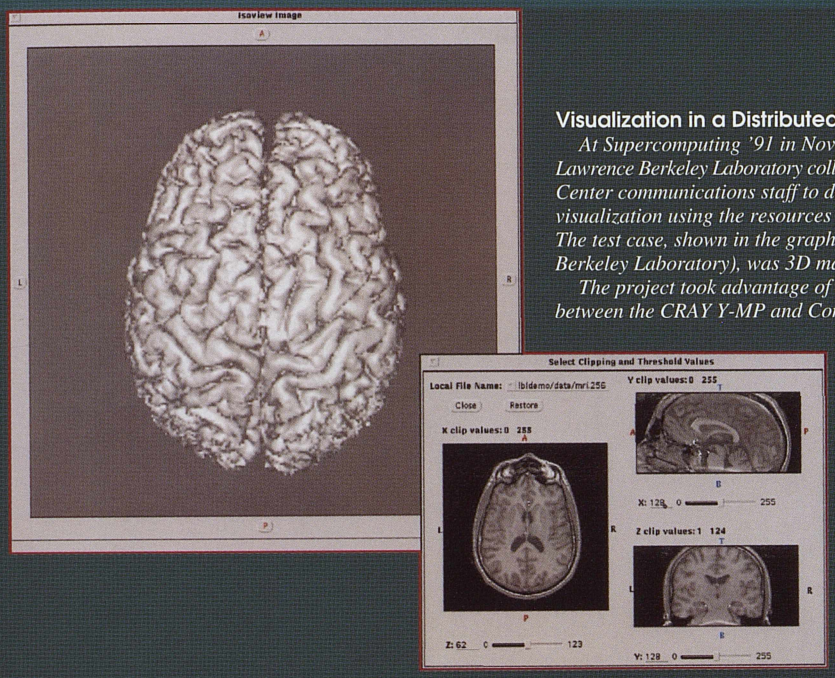

Visualization in a Distributed Computing Environment

At Supercomputing '91 in November at Albuquerque, researchers from Lawrence Berkeley Laboratory collaborated with the Pittsburgh Supercomputing Center communications staff to demonstrate the possibilities of interactive visualization using the resources of a heterogeneous computing environment. The test case, shown in the graphic (provided by Brian Tierney, Lawrence Berkeley Laboratory), was 3D magnetic resonance imaging (MRI).

The project took advantage of the center's high-speed connection (HiPPI) between the CRAY Y-MP and Connection Machine, CM-2. The CM-2 generated a set of 3D points from MRI data and sent this information across the HiPPI to the Y-MP. The Y-MP then rendered the image and sent it over the Internet for display at the on-site workstation. The project demonstrated that high-speed networks and supercomputers make it possible for scientists with only an inexpensive color workstation to use highly sophisticated visualization techniques.

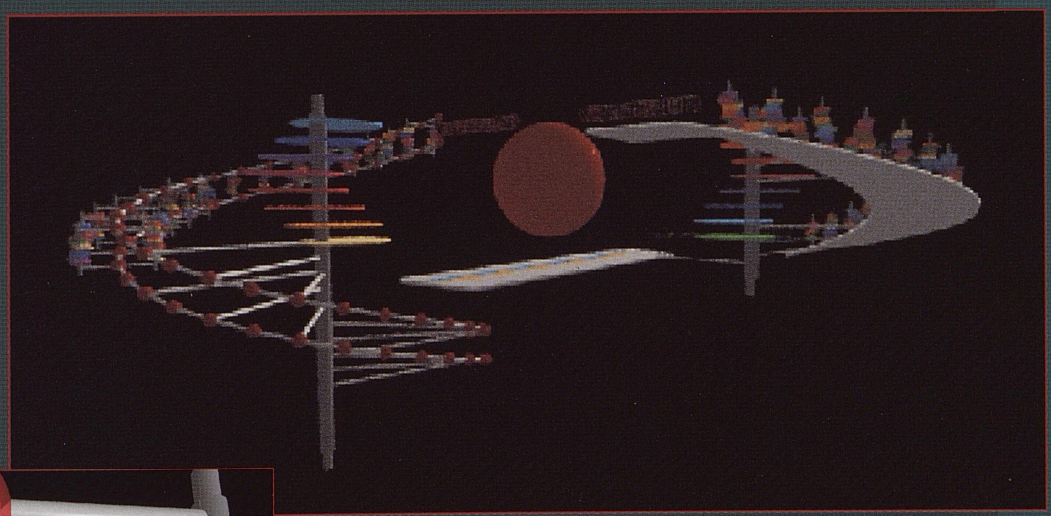

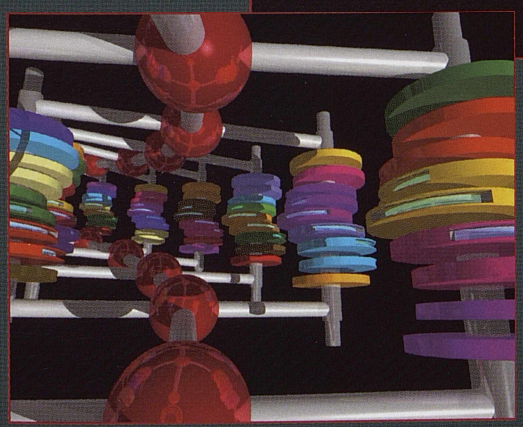

Fantasy Space Colony

These images are from a video walkthrough of a "fantasy space colony" designed by Carnegie Mellon architect Alaa Eldin Ibrahim. Joel Welling of the center's graphics staff collaborated with Ibrahim to produce the video.

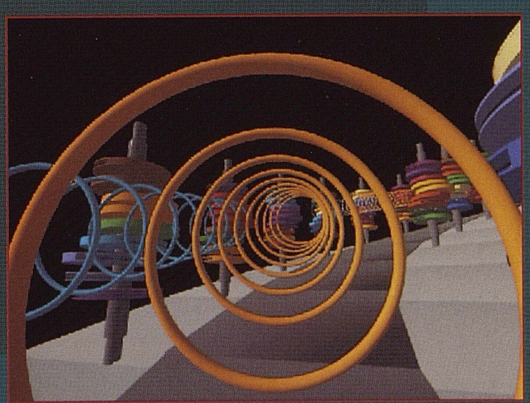

Grand Challenges

"The development of high performance computing and communications technology offers the potential to transform radically the way in which all Americans will work, learn and communicate in the future. It holds the promise of changing society as much as the other great inventions of the 20th century, including the telephone, air travel and radio and TV."

High-Performance Computing Becomes Law

With these words, on Dec. 9, 1991, President Bush signed the High Performance Computing Act. The law gives the stamp of authority and an allocation of $803 million to the federal High Performance Computing and Communications (HPCC) Program. Based on several years of planning by federal agencies in collaboration with scientists and managers from U.S. industry, universities and laboratories, HPCC maps strategy to advance high-performance computing and communications technology in this country over the next decade.

One of HPCC's main objectives — as stated in "Grand Challenges 1993," the FY 93 planning document from the Office of Science and Technology Policy, is to "accelerate development of a thousand-fold improvement in useful computing capability and a hundred-fold improvement in available computer communications capability by 1996." HPCC recognizes that great increases in computing power are within the foreseeable horizon of current technological feasibility and that creative use of this power will give scientists and engineers the ability to solve a wide range of "grand challenge" problems.

What are these Grand Challenges?

And how will solving them change society, as President Bush and his science advisors forecast? It should come as no surprise to those accustomed to the "super," "ultra," "hyper" product labels and one-up-manship marketing style of the high-performance computing industry that HPCC would be well stocked with grand claims. What, one might ask, does a grand challenge mean to me?

By any measure, global change questions are among the most pressing challenges we face as a society. The threats of vast climactic change from global warming and increased mortality from depletion of the ozone layer are among the "gifts" technology has brought us. Supercomputing, some say, is technology coming to rescue us from ourselves.

Scientists can take everything they know about what makes weather and climate and build a model inside a supercomputer. But all the computing power we have now, say the scientists who work on these models, isn't enough to give reliable predictions. With 100 or 1,000 times more capability, scientists may be able to tell us with reasonable certainty about the consequences we face and — even more important — what we can do to avert them.

Take the ozone layer, for instance — or what's left of it. Joseph Francisco is using supercomputing at Pittsburgh (p. 38) to find out if there are viable alternatives to the chlorofluorocarbons (CFCs) in refrigerants and insulation that are eating up the protective mantle of our upper atmosphere. His work is beginning to suggest that the alternatives pose problems of their own. But with existing computing, the complex quantum-theory calculations he needs to do can address only limited-size molecules. HPCC should advance the day when he can get the answers.

The discovery of the DNA double helix in the 1950s opened a vital new field of research, molecular biology, that has already begun to transform disease treatment, and scientists have only scratched the surface. One of the fundamental challenges of molecular biology is protein folding. This March the Pittsburgh Supercomputing Center organized a workshop, "Fast Processes in Protein Folding Dynamics," that brought together leading researchers on this problem for three days of intense interaction.

Even partial understanding of how a protein's amino-acid sequence determines its three-dimensional "folded" structure could give us the ability to design new proteins for specific biological needs. Imagine, for instance, protein therapy to improve digestion or lower blood pressure. Even more dramatic possibilities involve blocking the disease-producing action of viruses such as AIDS. Yet even today's state-of-the-art molecular dynamics methods for tracking protein movements, such as Charles Brooks has helped pioneer (p. 14), are thousands to millions of times short of being able to fully simulate some of the most scientifically interesting proteins. Drug companies (p. 18) also have begun to exploit the power of supercomputing to reduce the time involved in the immense amount of trial-and-error research required to bring new drugs into the marketplace.

What would it mean to heart research if a realistic, three-dimensional computer model of blood flow in the heart and major vessels were available? How many people could benefit from the insights into heart functioning and circulation that such a model could provide? How many lives could be saved? Charles Peskin's heart model (p. 20), as it now stands, involves manipulation of 11 million variables for every simulated heartbeat, requiring 10 hours of CRAY Y-MP computing. Even with the severe limitations of this model, Peskin and his colleague, David McQueen, have designed an artificial mitral valve that is much improved over prior prosthetic valves, reducing the risk of heart attack and death from valve disfunction.

The complexity of detail in a single integrated-circuit "chip" one centimeter square is equivalent to the map features, at a city-block level, of the entire Eastern United States, and the level of complexity is increasing. Applying high-performance computing to technology design makes it possible to explore alternative approaches at low cost in money and time. Research by Jerzy Bernholc (p. 24) and Paulette Clancy (p. 26) addresses potential improvements in the design and production of semiconductors, a technology that has already transformed our manufacturing-based economy and driven us into the much-heralded post-industrial era.

Globular star clusters provide a unique laboratory for studying the evolution and structure of galaxies, information that is precious in getting at fundamental questions about how the universe came to be the way it is. Steve McMillan's star cluster simulations (p. 36) push available computing to the limit, and his results tend to confirm intelligent speculations about the importance of binary stars. Yet today's supercomputers aren't capable of the detailed star-by-star simulations needed to fully understand the events that influence star cluster evolution.

The Pittsburgh Supercomputing Center has been active in nucleating research teams to address grand challenges, and each of the projects described in the following pages contributes to the understanding of a grand challenge problem and to carrying forward the far-sighted impetus of HPCC. The solution of these problems and many others that researchers are attacking with the resources of supercomputing will undoubtedly change our lives. (MS)

Parting the Waters

Distributing Molecular Dynamics Calculations Across Two Supercomputers
Charles Brooks and William Young, Carnegie Mellon University

Birth of a Protein

In the beginning, there is only a loosely coiled chain of amino acids. In this stretched out, gangly form of its infancy, a protein isn't ready to do the biological work it's cut out for. But in the space of seconds — a long time in protein biochemistry, the long chain coils, twists and folds into a compact bundle. In this twisted, globular form, replete with nooks and crannies that help it latch onto other molecules, the protein begins to perform its vital role, whatever it may be — from digesting sugar to carrying oxygen through the blood.

How did it happen? Why does a particular linear sequence of amino acids fold into the particular complex three-dimensional shape that enables the protein to carry out life-sustaining tasks? It's called the protein-folding problem, and it is one of the important, unsolved problems in molecular biology. The solution has immediate practical consequences — biotechnology companies design genes to synthesize certain proteins, and failure to fold properly is a common production concern.

Pathfinder Supercomputing

"It's a problem near and dear to my heart," says Charles Brooks, much of whose recent research has centered on protein folding. A professor of chemistry at Carnegie Mellon, Brooks has been using supercomputers to study proteins at the Pittsburgh Supercomputing Center since it opened in 1986, and he has helped develop the CHARMM package of molecular dynamics software into a valuable tool in protein and DNA research and drug design.

In a current project, Brooks is collaborating with Ph.D. student Bill Young to develop ways to involve the new, massively parallel supercomputers — in particular, the Pittsburgh Supercomputing Center's Connection Machine, CM-2 — in protein folding and other "grand challenges" in molecular biology. "It's an experiment," says Brooks. "We're looking for the most efficient routes." Using the center's high-speed data link (HiPPI) between the CM-2 and CRAY Y-MP, they have developed code that couples the two supercomputing systems together, opening a new and potentially useful avenue for molecular dynamics calculations.

Charles Brooks (left) and Bill Young.

Molecular Dynamics Simulations of Apomyoglobin

This image, from CRAY Y-MP calculations by Charles Brooks, represents the three-dimensional structure of apomyoglobin — essentially myoglobin without the "heme" unit that gives blood its red color. Apomyoglobin is an oxygen-carrying protein that occurs widely in muscle tissue. Each color in the image represents one of the eight helices, A through H, that make up the protein — red for A, blue for B, around to tan for H.

Brooks was investigating how the motion of apomyoglobin — in effect a partially unfolded form of myoglobin — differs from "native" myoglobin. His calculations showed a large amount of movement in the F helix (orange), where the heme usually attaches, and also in the C-D region (light blue and yellow). This information, says Brooks, will aid in interpreting NMR experiments underway on this protein.

Apomyoglobin is composed of 153 amino acids and the simulation included 4,299 water molecules and tracked 500 picoseconds (a picosecond is a trillionth of a second) of motion, a very large-scale study with current techniques.

Water Water Everywhere

A limiting factor for protein molecular dynamics has been the large amount of computing needed to account for the water molecules that surround a protein in its cellular environment and influence its shape. Even a relatively small protein (100-150 amino acids) requires about 5,000 water molecules to enclose its folded structure. To examine the folding process (which means unfolding and stretching out the protein) could easily take more than 16,000 water molecules — far too many as a practical matter to handle with existing techniques.

Brooks and Young hit on the idea of breaking out the part of the computing that involves only water molecules and giving it to the CM-2. You can think of molecular dynamics computations, which involve calculating how the forces and energies between atoms

"There was a long tumultuous shouting
sound like the voice of a thousand waters."
— Poe, *The Fall of the House of Usher*

change as the atoms move, as a three-part problem, explains Brooks: protein-protein — interactions between atoms of the protein itself with each other; protein-water — interactions between protein atoms and water molecules; and water-water.

Unlike the first two parts of the job, water-water interactions are inherently homogeneous. "One water molecule is basically the same as another," says Young, and this homogeneity lends itself naturally to the CM-2 — thousands of processors that can do the same computations simultaneously on different data. Furthermore, by addressing the water-water part of the job, the potential for saving computing time is great. "The water-water interactions," says Brooks, "take 80% or more of the overall time in molecular dynamics simulations."

A Logical Solution

Young's effort was to develop an algorithm that gives a one-to-one correspondence between each water molecule and a CM-2 processor. To accomplish this, he had to contend with irregularities in a volume of water — differing distances between water molecules — that prevent a simple spatial arrangement of the calculation. To solve this problem, he used a "logical mapping" technique that keeps track of each molecule's nearest neighbors. "It provides an extremely regular communication pattern" says Young, "with a very efficient way of doing the calculation."

Working closely with Pittsburgh Supercomputing Center staff analysts, Young and Brooks used the HiPPI to distribute several test computations between the two systems. At this stage, the approach is not yet viable for large-scale problems, says Brooks, but he is confident the work will pay off. "This is a novel paradigm for these calculations, and our work is only in its beginnings."

Young is now adapting his approach to Pittsburgh's CM-5 — Thinking Machines' newest, more flexible, massively parallel system, and he has already seen some improved performance. Such explorations of how best to use a heterogeneous computing environment like the Pittsburgh Supercomputing Center are an important part of computational research. In this case, they provide tools that hasten the day when the protein-folding problem will no longer be a problem. (JL, MS)

References:

Charles L. Brooks III, William S. Young & Douglas J. Tobias, "Molecular Simulations on Supercomputers," *International J. of Supercomputer Applications* **5**, 98 (1991).

Computer graphics were produced by Dave Deerfield and Joe Lappa of the Pittsburgh Supercomputing Center from calculations by Brooks and Young.

This research is supported by the National Science Foundation, Cray Research, Inc. and the National Institutes of Health. Charles Brooks is a recipient of an A.P. Sloan Research Fellowship (1990-91).

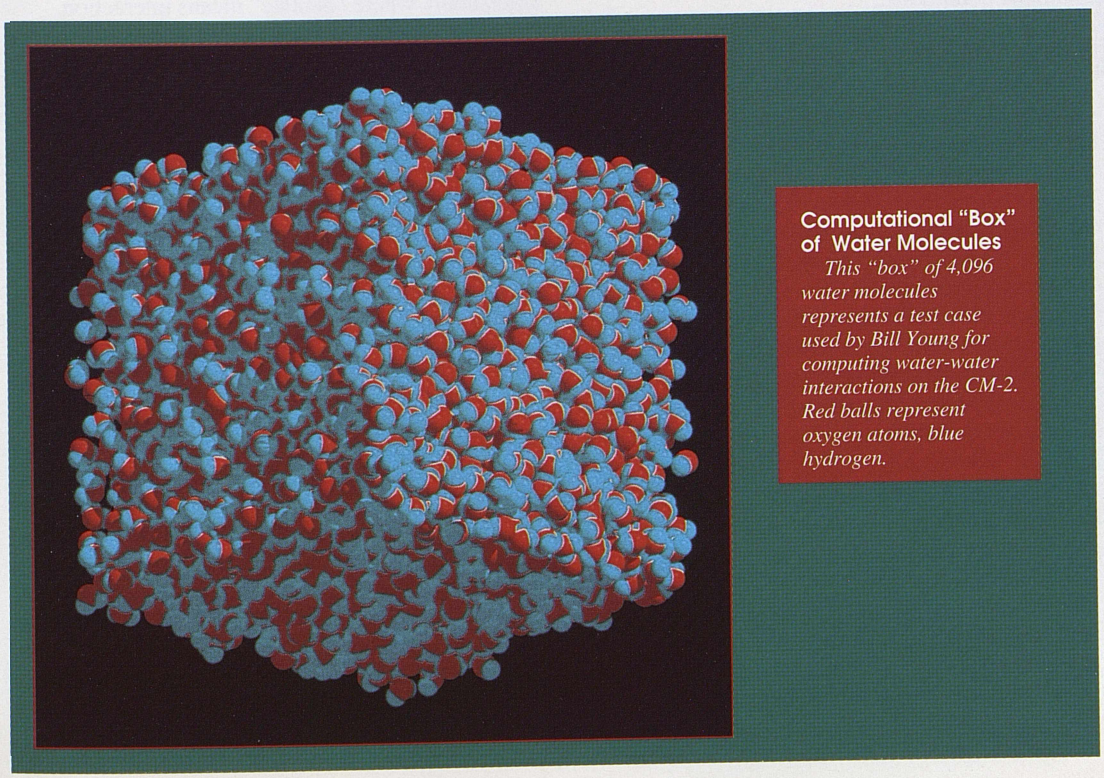

Computational "Box" of Water Molecules

This "box" of 4,096 water molecules represents a test case used by Bill Young for computing water-water interactions on the CM-2. Red balls represent oxygen atoms, blue hydrogen.

Mysteries of Clotting

Molecular Mechanisms of Blood Coagulation
Lee G. Pedersen, University of North Carolina at Chapel Hill
David W. Deerfield II, Pittsburgh Supercomputing Center

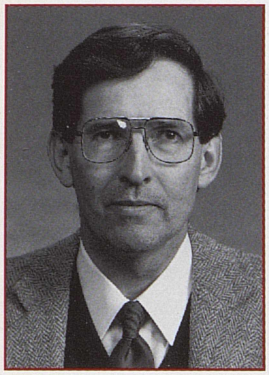

Lee Pedersen

The Strange Case of Calcium & Magnesium

Divide a blood plasma sample into three test tubes. Now, remove all metallic ions from the plasma and restore calcium to one test tube, magnesium to the second and both metals to the third. Stir vigorously and watch carefully. No clotting occurs in the magnesium sample, the plasma clots just fine in the calcium sample, and — most interesting — the plasma with both metals clots about 50 percent faster.

"Calcium can fill all the roles necessary for clotting to occur although magnesium apparently can fill a portion of those roles just as well, but not all of them," says chemist David W. Deerfield II, a scientific specialist at the Pittsburgh Supercomputing Center. For several years, Deerfield and Lee G. Pedersen, a chemist at the University of North Carolina at Chapel Hill, have been using supercomputers to study clotting on a molecular level, attempting to unravel the mysteries surrounding calcium and magnesium.

Blood contains similar concentrations of both metals, and because both are located in the same column on the periodic table — the alkaline earth metals — they have similar properties. But they behave differently in blood and understanding why, say the researchers, may help design drugs to assist clotting or, conversely, to thin the blood — important in treating high blood pressure, strokes and other conditions.

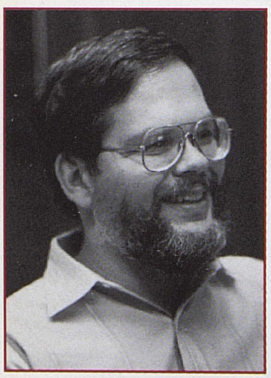

David Deerfield

The Clot Thickens

Clotting is the result of a multi-stage cascade of chemical reactions, with products from one reaction fueling the next until a protein known as fibrin is formed. An important product near the end of the cascade is prothrombin, an enzyme that anchors to the blood platelets — cell-like structures that become exposed following an injury.

"This is where calcium is important," says Pedersen. Calcium behaves like a bridge between the platelet surface and an amino acid on prothrombin known as gamma-carboxyglutamic acid (Gla). It orients the Gla molecules so they can bind with the platelet membrane. Calcium also acts as a bridge in another way. Both prothrombin and the membrane are negatively charged and would repel each other, but the positively charged calcium ion gets in the middle and allows them to interact.

This complex — the prothrombin, calcium, platelet pieces and other assorted proteins — then undergoes a chemical reaction and becomes the next enzyme in the clotting cascade. Meanwhile, magnesium appears to be an interested bystander.

Unidentate or Bidentate?

Unraveling the inner workings of that complex and the ensuing reaction is the focus of Deerfield and Pedersen's research. They have used a quantum mechanics approach — specifically, the GAUSSIAN software package at the Pittsburgh Supercomputing Center, a computationally intensive technique that allows researchers to examine the inner workings of chemical reactions. This is important because laboratory scientists can observe only reactants and products in chemical reactions. "Quantum mechanics is a theoretical procedure that allows you to predict experimental level results," Pedersen says. "It's a supercomputing type problem, and it lets you pose molecular questions that you couldn't do otherwise."

Although researchers already had known that a calcium or magnesium ion can bind in two different positions relative to the Gla side chain — with either one or two carboxylate groups (CO_2^-) — Deerfield and Pedersen used GAUSSIAN to study the geometric preferences for both positions. In the single carboxylate study, the calculations revealed that the preferred binding interaction ran counter to previous computational studies. They showed that water molecules — also known as the waters of hydration — filled five of six binding sites on a magnesium ion, leaving only one open for a carboxylate group. This is known as a unidentate orientation, where "dentate" means interaction.

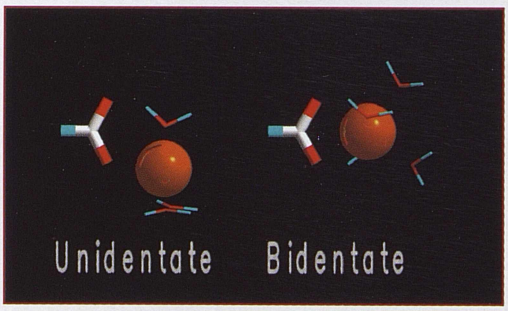

In this comparison graphic, red, white and blue represent atoms within carboxylate groups, blue and red represent water molecules, and the orange sphere is a magnesium ion. In the unidentate orientation, the ion interacts with only one oxygen atom (red) of a carboxylate group. In the bidentate orientation, the ion interacts with both oxygen atoms.

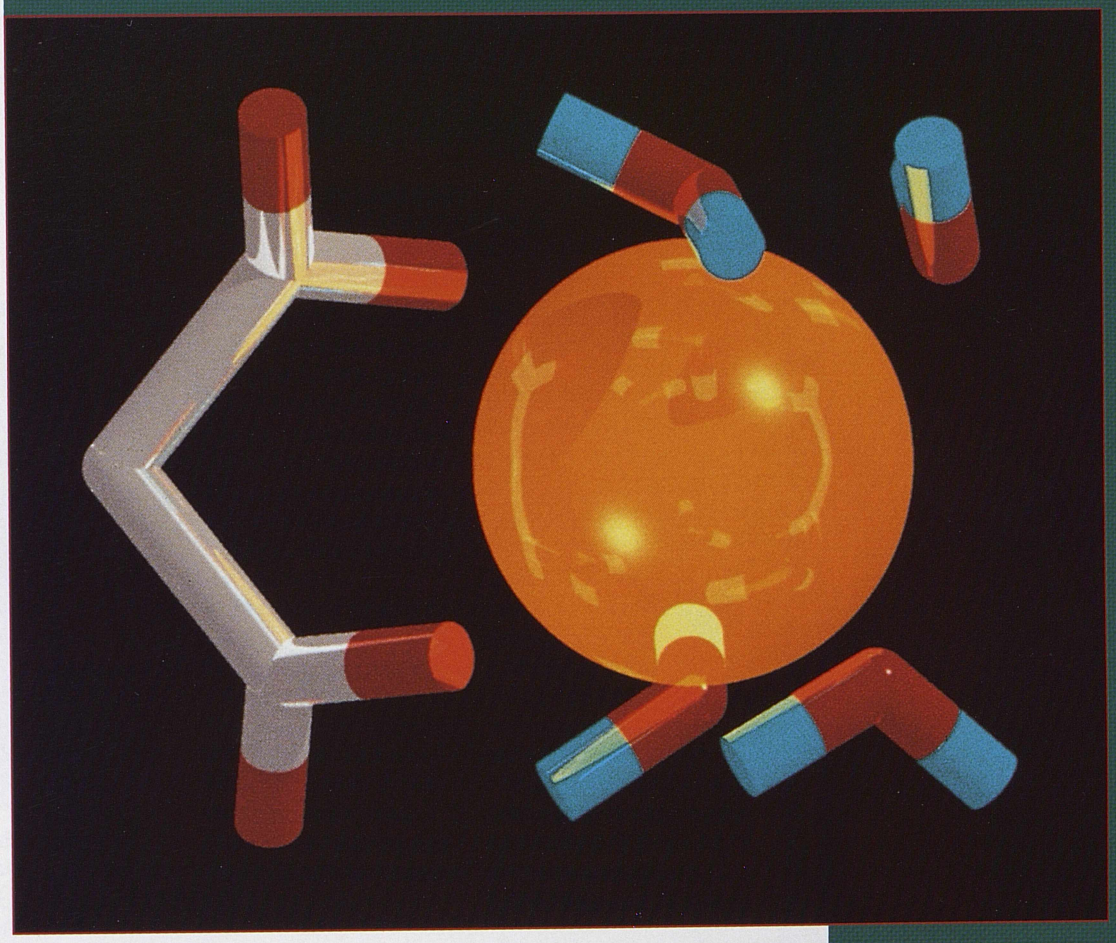

A calcium ion (orange) binding with oxygen atoms (red) in two carboxylate groups (red and white) on the Gla side chain, with water molecules (red and blue) filling in the other four of the six available binding sites on the ion.

"We have shown that water is instrumental in making the unidentate orientation preferred for both magnesium and calcium," Deerfield says. "Other researchers didn't have adequate computer resources to make large enough models. The CRAY Y-MP allowed us to model the waters of hydration." Deerfield and Pedersen's results apply whenever metal ions interact with single carboxylate groups, which are key components of all proteins.

The findings are potentially significant, Deerfield says, because genetic engineers insert metal ions into manufactured proteins to maintain structure. "Nature often uses these metal ions to help define the three-dimensional structure of proteins, and this may help design metal ion binding sites for these genetically engineered proteins."

The researchers also studied how a calcium or magnesium ion binds with two carboxylate groups on Gla, which is equally important to the coagulation cascade. The studies confirmed previous structural findings from laboratory experiments, which involved waters of hydration, and explained how this complex interacts with other substances.

More Answers to Come

Once the researchers figure out other pieces of the puzzle — for instance, how a calcium or magnesium ion binds with the platelet membrane, the trick will be putting together those pieces with how the two metals bind with Gla and discovering why calcium acts as a bridge and magnesium doesn't in the prothrombin complex.

That, however, doesn't mean that magnesium — especially considering the test tube example — is a non-player. Right now, Deerfield suspects that magnesium acts as a helper to calcium. That suspicion will be tested on the supercomputer. (SE)

References:

D.W. Deerfield II, D.J. Fox, M. Head-Gordon, R.G. Hiskey and L.G. Pedersen, "The Interaction of Calcium and Magnesium Ions with Malonate and the Role of the Waters of Hydration: A Quantum Mechanics Study," *J Am Chem Soc* **113,** 1892 (1990).

D.W. Deerfield II, H.B. Nicholas Jr., R.G. Hiskey and L.G. Pedersen, "Salt or Ion Bridges in Biological Systems: A Study Employing Quantum and Molecular Mechanics," *Proteins* **6,** 168 (1989).

D.W. Deerfield II, M.A. Lapadat, L.L. Spremulli, R.G. Hiskey and L.G. Pedersen, "The Role of Hydrated Divalent Metal Ions in the Bridging of Two Anionic Groups: An *ab initio* Quantum Chemical and Molecular Mechanics Study of Dimethyl Phosphate and Formate Bridged by Calcium and Magnesium Ions," *Journal of Biomolecular Structure & Dynamics* **6,** 1077 (1989).

This research is funded by the Biomedical Research Technology Program at the National Institutes of Health.

Looking for New Drugs

Molecular Dynamics Simulations of Interleukin-8 with Solvent
Beryl W. Dominy & James P. Rizzi, Pfizer Research

Faster than a Speeding Bullet

What's so "super" about supercomputing? Sure, it's state-of-the-art scientific computing, amazingly fast and more powerful than even the most advanced workstations, but what can it do that less expensive computers can't? Is the most high-tech computer on the block really needed if all it does is the same jobs only faster?

It's the sort of question people were asking at Pfizer Inc, a research-based company with worldwide pharmaceutical sales approaching $3 billion. To get the answer, Pfizer in May 1990 became a corporate affiliate of the Pittsburgh Supercomputing Center.

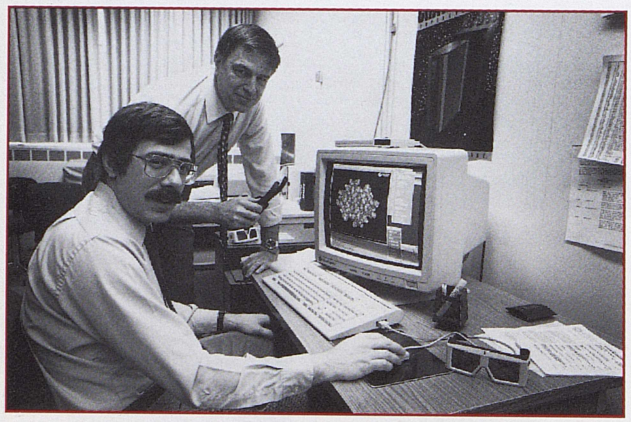

Beryl Dominy (standing) and Jim Rizzi of Pfizer Research. "The Pittsburgh Supercomputing Center," says Dominy, "provided us with the means to evaluate supercomputing at relatively low cost. We had excellent support from people on their staff, and they had the programs we wanted to use already available on the CRAY."

"We were intrigued with the possibilities of using a supercomputer," explains Beryl Dominy, who heads Pfizer's computational research group, "but we had no experience to know what it could really offer as a tool for discovering new drugs." Dominy and his colleague, Jim Rizzi, knew that the CRAY Y-MP would do molecular modeling faster than their in-house workstations. They had the benchmark studies to prove it, but that didn't answer the bigger questions. "Management wanted to know what else it would do," says Dominy. "They wanted to know what we could learn with a supercomputer that we couldn't with our existing system."

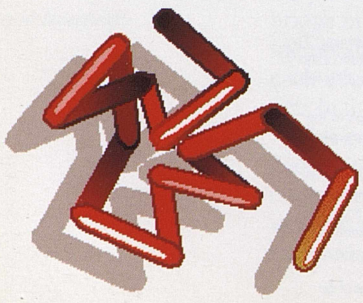

New Horizons for Protein Modeling

Dominy and Rizzi believed that supercomputing might give them not just faster, but also better results. "We had hopes," says Dominy, "that it would change the way we view a modeling project. When we set up large calculations on a workstation, we're aware of how long it will tie up the computer, and we're always balancing computational time against complexity of the model."

A prime example of this balancing act, explains Dominy, is size of the molecule. Often, for proteins typical of drug research, it was necessary to do calculations that addressed only a part of the protein. "A model compound can quickly get too big for the calculation, so you tear it apart and hope what you're doing accurately reflects the whole molecule."

Similarly, Pfizer's modeling usually left out solvent. In the living environment, proteins are surrounded by thousands of water molecules. Parts of the protein are loose and flexible, and the surrounding water can affect their structure. Because of computing limitations, however, Pfizer usually could account for this effect only indirectly. "You can do calculations on two different proteins that are related," explains Dominy, "and you hope that whatever influence solvent might have on one protein would be the same on the other, and that therefore the effects would cancel out. But you can't really depend on that."

A Case Study: Interleukin-8

Affiliating with the Pittsburgh Supercomputing Center gave Dominy and Rizzi a chance to see how supercomputing could affect a real project. Their case study was interleukin-8 (IL-8), a protein involved in the complex biochemistry of inflammation.

When a part of the body is injured, irritated or infected, cells in the affected tissue release signalling proteins that result in increased blood flow to the area. Although inflammation is part of the body's system of defenses (it facilitates the arrival of white blood cells for one thing), its immediate result is redness, swelling and pain. As one among the many signalling proteins involved in promoting inflammation, IL-8 provides a potential target for drug therapy to reduce the pain of inflammation, such as that associated with arthritis.

Affiliating with the Pittsburgh Supercomputing Center gave Dominy and Rizzi a chance to see how supercomputing could affect a real project.

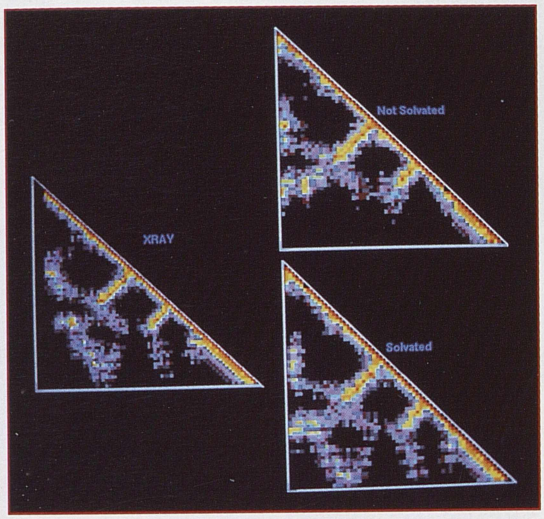

C-alpha Distance Maps for IL-8

These maps, which use the color scale (red to violet) to show relative proximity between carbon atoms in the protein backbone, give an indication of how the protein is folded. The simulation with solvent is similar to the X-ray crystal structure.

To design a drug that competes with IL-8 and blocks its biological function, the first and necessary step is a precise understanding of IL-8's molecular structure. As a starting point, Rizzi turned to platelet factor 4, a structurally related protein — with similar amino acid sequencing — whose three-dimensional structure was known. Working from this family resemblance, he developed a rough model of IL-8 and used molecular dynamics computations with a program of proven reliability (CHARMM) to test its validity.

Supercomputing Proves Its Worth

Rizzi's initial runs on workstations left out solvent. Even without water molecules, these computations took hours of workstation computing, and they gave results that didn't make sense to the researchers. "Over the course of the simulation," explains Rizzi, "the IL-8 began to unravel from its starting structure. One end of the protein twisted away from the rest of the molecule."

Rizzi and Dominy thought the strange results might be due to the absence of solvent, but test runs indicated that adding solvent would require about 400 hours of workstation computing, an impractical amount of time. Could supercomputing make a difference? Rizzi transferred his model to Pittsburgh's CRAY Y-MP, and a simulation with thousands of water molecules (20 hours of Y-MP computing) solved the problem; the IL-8 structure remained intact as it should.

With further Y-MP simulations, says Rizzi, "we created a final, low energy structure of IL-8 that we felt was a reasonable representation of how it appears in solution." A few months later, the crystal structure of IL-8 was determined from experimental methods, and Pfizer's model was "very accurate, especially in a region that is considered to be important for its biological activity."

"This project and others convinced us," says Dominy, "that one of the most useful aspects of using a supercomputer in drug discovery is the relevance of the models it allows us to generate." More realistic models give higher quality information — much more valuable, in Dominy's view, than simply creating lower quality information faster. Based on the success of their experience at Pittsburgh, Pfizer purchased a CRAY Y-MP2E last year. (MS)

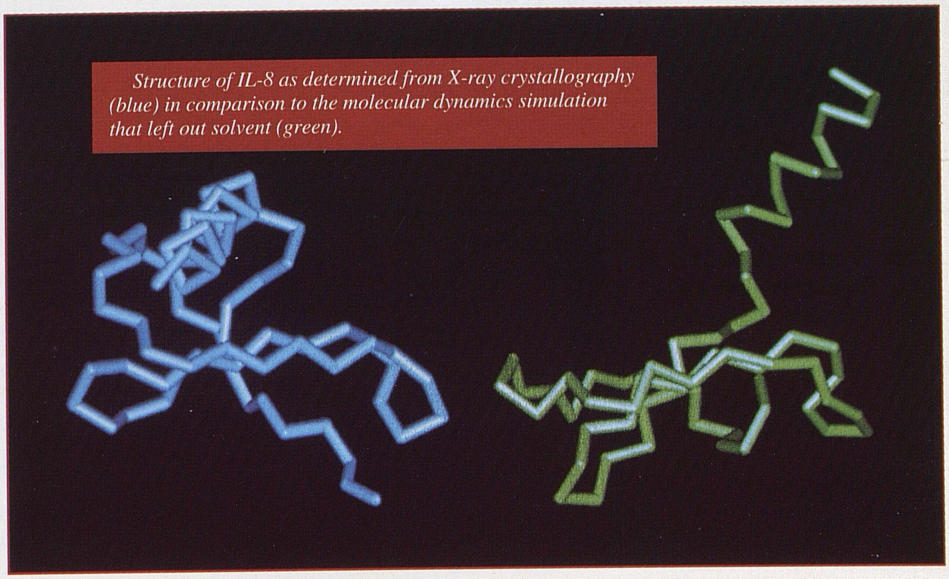

Structure of IL-8 as determined from X-ray crystallography (blue) in comparison to the molecular dynamics simulation that left out solvent (green).

Going with the Flow

Computational Fluid Dynamics with Immersed Elastic Boundaries
Charles S. Peskin and David M. McQueen, New York University

To Build a Better Heart Valve

Sails, parachutes, flags, kites, blood flow in the heart. To most of us, there's an oddball in this list. To Charles Peskin, however, seeing what these phenomena have in common is a key insight. It has enabled him to pioneer using computers to design better artificial heart valves, and other researchers have extended his method to many other biological flow problems — from fluids in the inner ear to how fish swim.

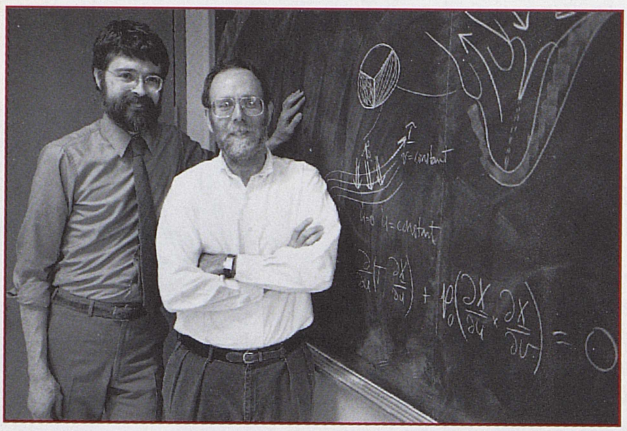

Charles Peskin (right) and David McQueen. In two three-day workshops at the Pittsburgh Supercomputing Center (April 1990 & July 1991), Peskin and McQueen presented their techniques for modeling flows with flexible boundaries to 40 other researchers, many of whom are now applying these techniques in their own research.

For about 15 years, Peskin and his colleague David McQueen, both of NYU's Courant Institute of Mathematical Sciences, have been developing a computer model of blood flow in the heart. Computer modeling of fluid flows, such as the way air flows around an airplane wing, has been used for years and has become a powerful tool in the design of airplanes, cars, pipelines and many other fluid flow processes. The trick for Peskin and McQueen, however, was that unlike the rigid surfaces of airplane wings and car bodies, the walls of the heart are muscle fiber, elastic and flexible, and they move with the flow of blood.

The Moving Boundary Problem

"The flow controls the boundary that controls the flow," says Peskin. In technical terms, the heart is a "coupled system"; the contracting muscles of the heart wall move the blood, and the moving blood in turn exerts pressure that moves the heart walls. This dynamic interaction makes it extremely difficult to use the standard techniques of computational fluid dynamics, as this field of computer modeling is called.

To get around this difficulty, Peskin devised a mathematical approach — "the immersed boundary technique" — that makes it possible to think of the moving heart wall as, rather than a wall *per se*, a special part of the fluid where additional forces happen to be applied. This innovation imposes regularity on the problem and makes it possible to use well developed methods (called "finite difference methods") that involve laying a grid-like mesh over the fluid, dividing the computation into small, discrete areas.

Peskin and McQueen first applied their technique in a two-dimensional model, and this work led to an improved design for an artificial mitral valve, the valve that controls blood flow between the left atrium and left ventricle. Existing prosthetic valves sometimes failed and often caused blood clots to form on the valve, as a result of blood flow stagnation. Peskin and McQueen used computer simulation to find a design less likely to induce clots, and their mitral valve has been patented and licensed to a corporation developing it for clinical use.

The Heart in Three Dimensions

In current work at the Pittsburgh Supercomputing Center, Peskin and McQueen have upped the ante of their research by moving from two to three dimensions. The ambitious goal is a realistic computer model of the heart, its valves and the nearby major vessels. Such a model could be a great boon to heart research by making it possible to study normal and diseased heart function without the limitations of using human and animal subjects.

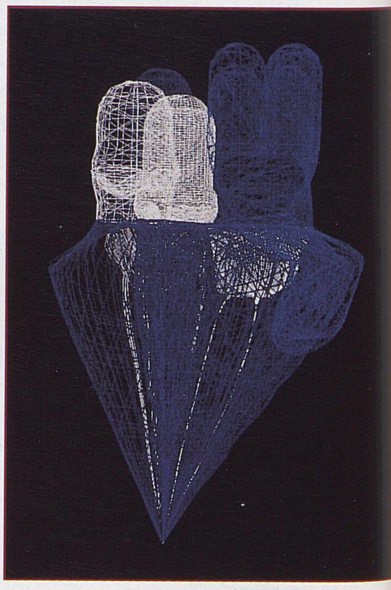

Three-Dimensional Heart Model

This graphic shows the entire 3-D heart model developed by Charles Peskin and David McQueen. White fibers indicate valves and, for the two outflow valves, associated great vessels — the aorta and pulmonary artery. Blue fibers indicate heart walls and other vessels. The left ventricle and atrium are on the right side of the image and vice-versa. Even by supercomputing standards, the computational demands of this model are enormous; one heartbeat requires about 10 hours of CRAY Y-MP computing.

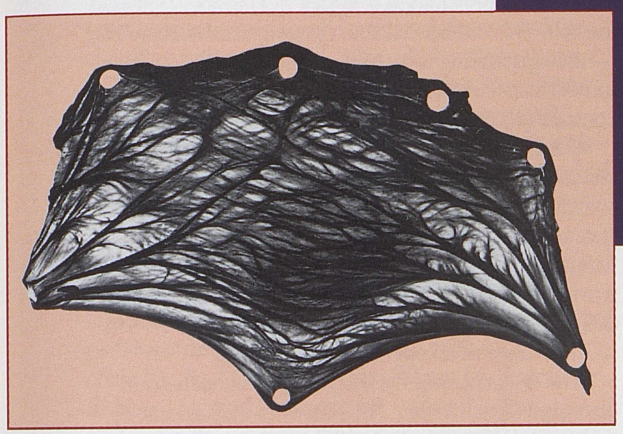

Single Leaflet of Actual Aortic Valve

This microphotograph shows the collagen fibers of a single aortic valve leaflet. The leaflet attaches to the aortic wall at the lower left and right corners. (Reproduced from A.A.H.J. Sauren, The mechanical behavior of the aortic valve, *Ph.D. thesis, Eindhoven Technical University, The Netherlands, 1981.)*

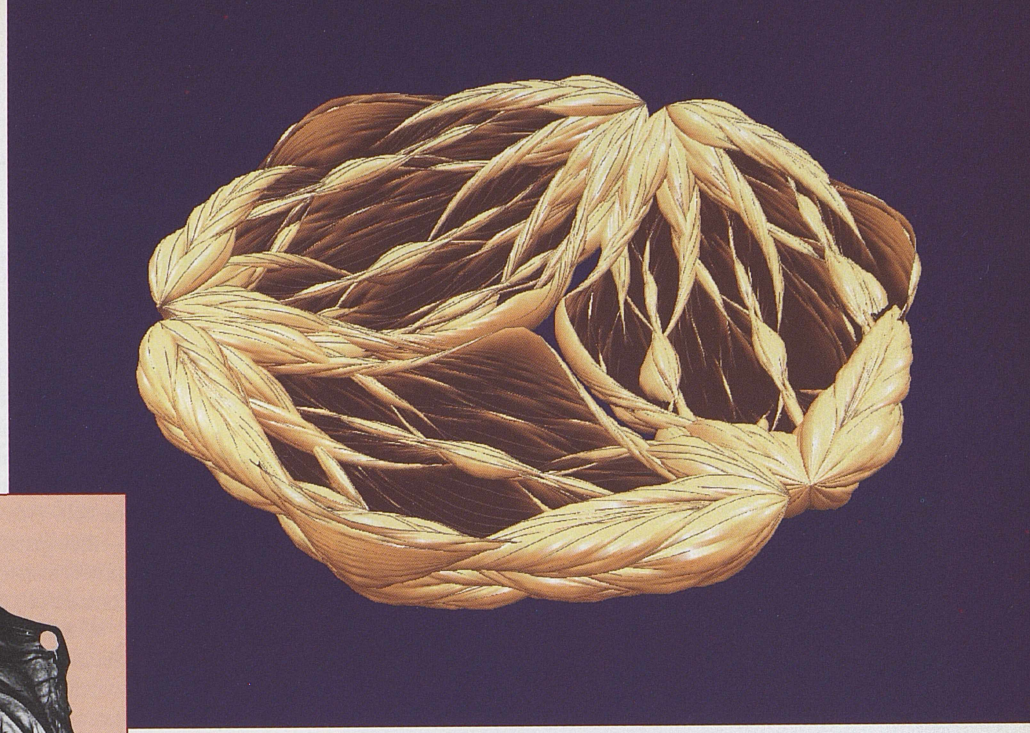

Computer Modeled Aortic Valve

This image, generated by David McQueen, shows Peskin and McQueen's computer model of the aortic valve in 3-D perspective with some of the fibers (one of every 16) superimposed on the surface. "Each of the three leaflets resembles half a walnut shell," says McQueen, "with the concave side visible." The small gap between the leaflets shows the valve as if it were just beginning to open.

Recently, they have focused on a difficulty with one part of their model — a leaky aortic valve. This valve regulates blood flow out of the heart (from the left ventricle) into the aorta, the major vessel carrying oxygen-enriched blood to the body. "We needed to model all the valves," notes Peskin, "but one of them, the aortic valve, in the original form of the model wasn't working. When it was supposed to be closed, it would flip over and let blood flow back into the heart."

To fix the leak, they went back to the drawing board. "Our approach," says Peskin, "was to try to understand the function of the valve and to derive how it should be constructed from first principles of engineering design." How should the valve's connective fibers be laid out to best support the load from blood in the aorta pressing back against the closed valve?

Fixing a Leaky Valve Mathematically

If you were to look at a closed aortic valve from inside the left ventricle, explains Peskin, you would see a three-sectored circle — like a wheel with three spokes — each sector being a separate leaflet of the valve. When the ventricle contracts, the leaflets open — letting blood into the aorta, and as it relaxes, they flip back to meet and close the valve.

From mechanical measurements by other researchers, Peskin and McQueen knew that the aortic valve fibers (which are made of collagen) are about ten times stronger along the arc of the circle than along the radius. This suggested they could assume a "one-parameter family of fibers," meaning, says Peskin, "that at any given point the fibers are basically running in one direction; they aren't crisscrossing in different directions." From this assumption, they derived a set of differential equations.

The results from their CRAY Y-MP computations took them by surprise. "From the nature of the mathematics, which involves smooth differential equations," says Peskin, "we thought we would get a smooth arrangement of fibers. Imagine our astonishment when the results showed up on our workstation screen. We saw complicated structures with branching and braiding that look very much like the collagen fibers that support the leaflets of the real valve." What especially interests the researchers is that a purely mathematical approach to design gave a result closely resembling actual anatomy. "This valve looks like it jumped off the pages of a biology book," says Peskin, "and yet it came from mathematics, not from biological data."

The new valve is now incorporated in the 3-D model and testing indicates the leak is fixed, which means that the researchers can turn their attention to other components of their computer heart. "Stiffness of the muscle, how fast and how slowly the muscle relaxes and contracts — because each computation is substantial," says Peskin, "we can't progress except by trial and error." (MS)

References:

Charles S. Peskin & David M. McQueen, "Mechanical Equilibrium Determines the Fiber Architecture of the Aortic Valve Leaflets," preprint, 1992.

C. S. Peskin & D. M. McQueen, "A three-dimensional computational method for blood flow in the heart," *Journal of Computational Physics* **81**, 372 (1989); **82**, 289 (1989).

This research is supported by the National Institutes of Health and the National Science Foundation.

Wild at Heart

Investigations of Regions of Instability in the Electrical Activation of Cardiac Cells
C. Frank Starmer, Duke University Medical Center

Killer Rhythms

Most of the time, like a good drummer, our heart keeps a steady rhythm. If it skips a beat, as it does now and then when we're excited, we have an "arrhythmia." Mild forms aren't much of a problem, but the most severe form — fibrillation — is a killer that takes hundreds of thousands of lives a year. The synchronization that keeps the heart muscles working together goes awry, and the normally steady rhythm becomes a wild, chaotic flutter. Unless normal heartbeat is restored, brain-wave activity stops in minutes.

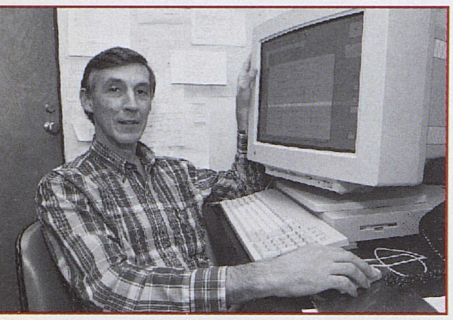

C. Frank Starmer

Emergency room doctors often administer drugs, called antiarrhythmic drugs, to bring these wild rhythms back to normal. They work fairly well, though how they work is not completely understood, and some recent research suggests that the same class of drugs that suppresses fibrillation can, under certain circumstances, trigger it.

"The mechanism that makes these drugs antiarrhythmic," says Frank Starmer, "simultaneously makes them proarrhythmic." For about 10 years, Starmer, a professor of computer science and experimental medicine, has studied these drugs with colleagues at Duke and, more recently, with research groups in Chicago, New York and Europe, including an active collaboration with theorist Valentin Krinsky in Pushchino, Russia. With the aid of supercomputing, he has used a model of the heart's electrical activity to analyze how they work. Starmer's "computer experiments" combined with laboratory research related to them have led to new understanding that could alter clinical practice in the treatment of fibrillation.

Action Potentials in the Heart

Antiarrhythmic drugs work by blocking channels in the membrane of heart muscle cells. These channels, formed by pore-like proteins, are in effect gates controlling the flow of sodium ions (Na^+) into the cell.

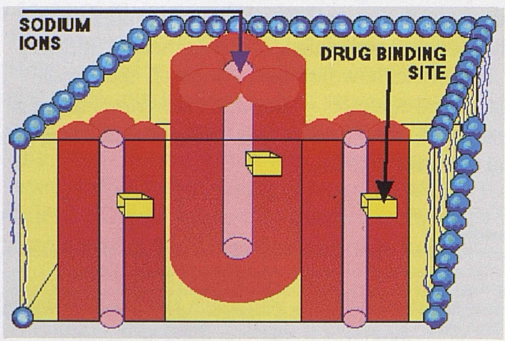

Sodium channels in heart-muscle cell membrane.

They open when stimulated by the local difference in electrical potential, called the "action potential," between the inside and outside of the membrane. When the channels open, sodium ions outside the cell, where the concentration is higher than inside, flow into the cell. This current leads to opening nearby channels — thus generating a "wave" of electrical activity that moves from cell to cell, triggering the muscle contractions that pump blood.

When ions rush in, the voltage across the open channel reduces until it once again closes. This switching back and forth between open and closed states, which occurs in millisecond intervals, is a key factor in understanding how antiarrhythmic drugs work, points out Starmer, because a molecule of the drug has to be in the right place at the right time — in striking range of an open channel.

"The drugs operate by keeping the channels from opening," says Starmer, "and the way they appear to do it is by getting inside the channel when it's open and plugging it. So these drugs behave in a way that is fundamentally different from drugs that access continuously accessible receptors. They depend on the frequency that the cell is being stimulated, and the interaction between the drug and the channel thus reflects 'pulse chemistry' in contrast to the usual continuous chemistry. If the stimulation frequency is once a second, the drugs have a millisecond shot at plugging this hole. If it is five times a second, they have five times as much chance to shut it down."

Modeling the Vulnerable Window

For his computational model of heart action potentials, Starmer worked from a previously developed model (the Beeler-Reuter model), extending it from one cell to an infinitely long series of connected cells. This allowed him to look at action potentials as a wavelike movement from cell to cell, approximating the reality of action potentials in the heart, and the difference was crucial.

"If you are looking at a single cell," explains Starmer, "it is always either excited (open) or recovering (closed). Once we were able to watch this wavefront move, suddenly we had an environment that was not homogeneous. Cells were in different states — inexcitable, partially excitable, totally excitable, and we were able to ask questions about the role drugs play in either amplifying or attenuating vulnerability to a rhythm that might lead to sudden cardiac death."

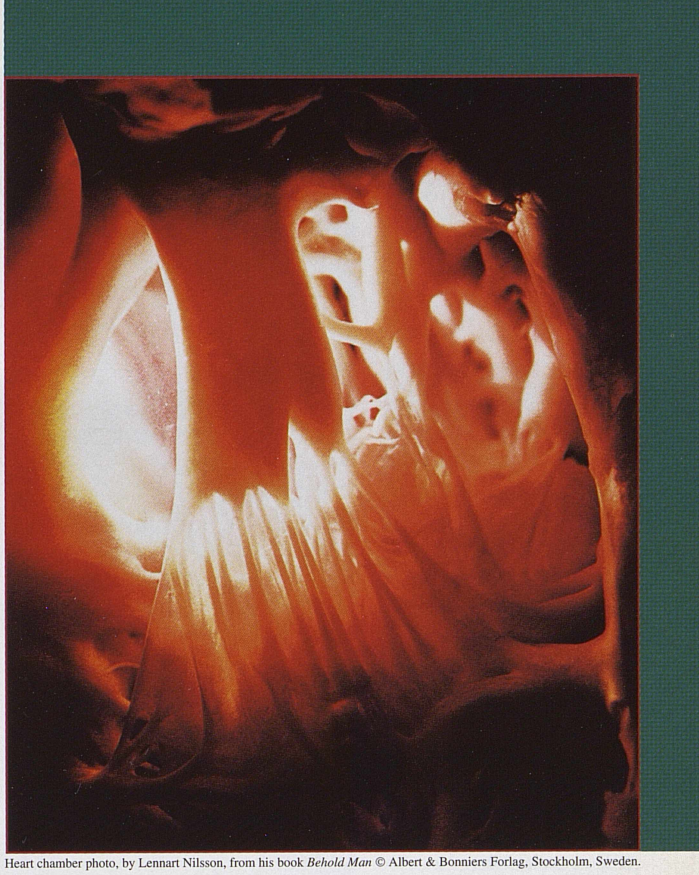

Heart chamber photo, by Lennart Nilsson, from his book *Behold Man* © Albert & Bonniers Forlag, Stockholm, Sweden.

"Anything like the sound of a rat
Makes my heart go pit-a-pat!"
— Browning, "The Pied Piper of Hamelin"

Clinical experience has led researchers to believe that fibrillation can be triggered by an extra heartbeat during a vulnerable period, when the heart may be sensitive to extra stimulation, and Starmer's model tended to confirm this surmise. He looked closely at what happens when a cell is restimulated toward the end of an action potential, when all the channels are not closed. His results showed a "vulnerable window," a several millisecond period of time during which restimulation can generate an action potential that can circulate forever — in effect a computational model of fatal arrhythmia.

A Double-Edged Sword

The next step was to include drugs in the model. Laboratory studies provided data on the rate at which particular drugs bind and unbind to open sodium channels. By factoring this into his model, Starmer found — as might be expected — that the drug prolonged the time (the "refractory interval") during which a previously stimulated channel will not respond to restimulation. This is the antiarrhythmic effect of blocking the sodium ion channel. What he also found, however, is that the drug prolonged the vulnerable window.

"It's a double-edged sword," explains Starmer. "The more antiarrhythmic a drug is, the more proarrhythmic it is. And you can't get away from it. This class of drugs, sodium channel blockers, was thought to be very safe, and we now have theoretical evidence suggesting that in different degrees of severity they are not."

Fellow researchers initially greeted Starmer's modeling results with skepticism, so he turned to the laboratory and conducted experiments in live tissue that reproduced the proarrhythmic effect with three sodium channel blocking drugs — propoxyphene (Darvon), cocaine and lidocaine.

Starmer credits his supercomputing studies with leading the way for the laboratory findings. "The model told us how to do the experiments. We wouldn't have known, for instance, how to look for the vulnerable window because it is so narrow — only three to five milliseconds. But our numerical experiment said *This is where it has got to be. Go look there.* And we found the effect." (MS)

Working in the Lab Without Walls

For Frank Starmer, access to supercomputing and the ability to network electronically with colleagues around the world is something to be excited about. In 1987 he traveled to Moscow and began collaborating with a research team there that shared his interest in antiarrhythmic drugs. Through the Moscow group, Starmer also became connected with a lab in Paris and began collaborations there. Through the Internet, he also maintains connections with researchers in Germany, Chicago and Syracuse.

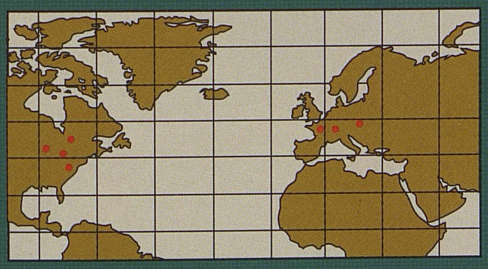

"I have some of the best minds in the world helping me, and all of us can be working on or thinking about these problems more or less simultaneously. We move data back and forth, and we move ideas back and forth with reckless abandon. It's much better than sitting in the lab by myself struggling with these things. It's exciting because the Internet and access to supercomputing has created literally a laboratory without walls."

References:

C. Frank Starmer, Anselmo A. Lastra, Vladislav V. Nesterenko, & Augustus O. Grant, "Proarrhythmic Response to Sodium Channel Blockade," *Circulation* **84**, 1363 (1991).

C. F. Starmer, A. R. Lancaster, A. A. Lastra & A. O. Grant, "Cardiac instability amplified by use-dependent Na channel blockade," *American Journal of Physiology* **262**,1305 (1992).

This research is supported by the National Heart, Lung, and Blood Institute, NIH; the joint US-USSR scientific exchange in problem areas 5, sudden death; and the Division of Research Resources, NIH.

Getting the Dope on Diamond

Structure and Dynamics of Semiconductors and Clusters
Jerzy Bernholc, North Carolina State University

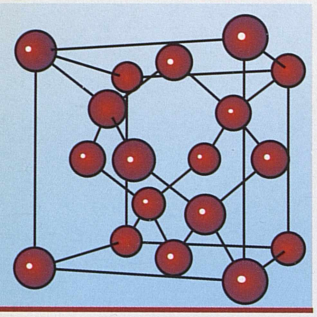

In diamond (above), each carbon atom bonds with four neighbors in a rigid three-dimensional framework, and a diamond crystal is, in effect, a single giant molecule. In graphite (below), on the other hand, each atom bonds to three others in a layer. The layers are only weakly connected to each other, so that at slight pressure they slide over each other, which is why graphite is a good lubricant.

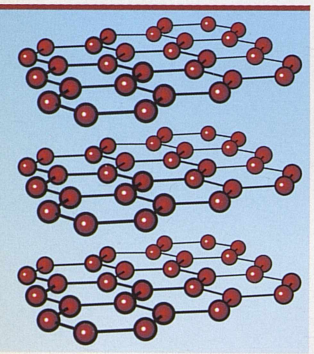

Mother Nature's Clever Atomic Architecture

How is it that two solid materials can be as different as night and day and yet fundamentally the same? Consider the "lead" in your pencil, for instance, and a diamond ring. One is black and so soft you can smear it on paper. The other is colorless, shines brilliantly and is so hard it scratches steel. One costs pennies a pound, the other millions of dollars. Yet both graphite and diamond are carbon.

The answer, as you might expect, is that Mother Nature is wondrously clever. Using the same atoms and simply changing the way they're put together, she creates an extraordinary range of properties.

As a solid-state physicist, Jerzy Bernholc works at understanding Mother Nature's architectural skills. The aim of his research is "an atomic level description of material properties." In other words, he wants to know where the atoms are in a material — its characteristic structure — and how these microscopic details relate to how hard it is, how flexible or brittle, how it conducts heat and electricity, and all similar properties that make it useful or not for a particular purpose.

Quantum Molecular Dynamics

An essential tool of Bernholc's trade is supercomputing. With complex quantum theory calculations, it is possible to compute the forces and energies that mold atoms together, and once the characteristic structure is well understood, it becomes possible, with more quantum-theory computing, to predict virtually everything about the material. The ultimate goal of this work — the vision for the distant, but foreseeable future — is to design new materials according to a wish-list of desired properties.

The obstacle is that the calculations are enormously complex. They have such an appetite for computing time that even with today's most advanced supercomputers, like the CRAY Y-MP, only relatively small structures can be studied. But the times are changing. With a new technique called quantum molecular dynamics (QMD), it is now possible, says Bernholc, to attack problems involving tens and, in some cases, hundreds of atoms.

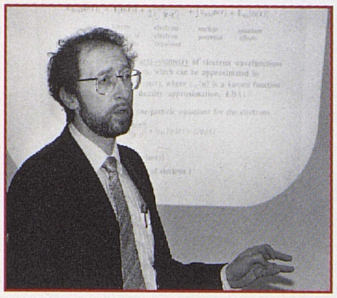

Jerzy Bernholc. At a seminar at the Pittsburgh Supercomputing Center in April, Bernholc described quantum molecular dynamics and its applications in his research.

QMD (developed by European physicists Car and Parinello) is a true breakthrough, says Bernholc: "It's a very fast technique, and because it is so fast one can apply it to complex systems." A big advantage of QMD, as compared to the Newtonian molecular dynamics used in protein and DNA research, is that it puts electrons on an equal footing with atoms. In effect, Bernholc explains, the method lets the electrons move with time and the atoms move with them, so that atom-atom interactions are computed directly from electronic forces — far more accurate and reliable than the Newtonian approach.

Bernholc's research group is one of only a dozen or so in the world currently using QMD, and they are applying it at the Pittsburgh Supercomputing Center in three projects: disordering of superlattices (structures used in creating high-speed "chips"), the structures of large metal clusters, and impurities in diamond. With recent work on the last of these three, Bernholc and his colleagues have explained a long-standing problem regarding a distortion in the crystal structure of diamond.

Nitrogen in Diamond

Nitrogen is a naturally-occurring impurity in diamond that gives it a slight blue tint. Since 1959, researchers have known that this impurity causes a shift in the crystal structure — the nitrogen atom, substituting for a carbon, pushes away from one of the four carbons around it. The phenomenon is unusual in that similar impurities — such as phosphorous in silicon — don't distort the crystal, and research for 30 years has failed to give a satisfactory explanation.

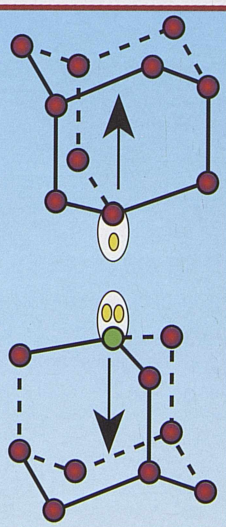

This schematic shows a single diamond crystal with a nitrogen atom (green) substituting for one of the carbons. Since nitrogen has five electrons in its outer orbitals, versus four for carbon, one of the four surrounding carbon atoms is left with one dangling bond, and this single unbonded electron repels the already paired electrons in nitrogen. Bernholc's calculations showed that this is why the structure distorts. The ovals indicate electron orbitals involved in this repulsive nitrogen-carbon interaction.

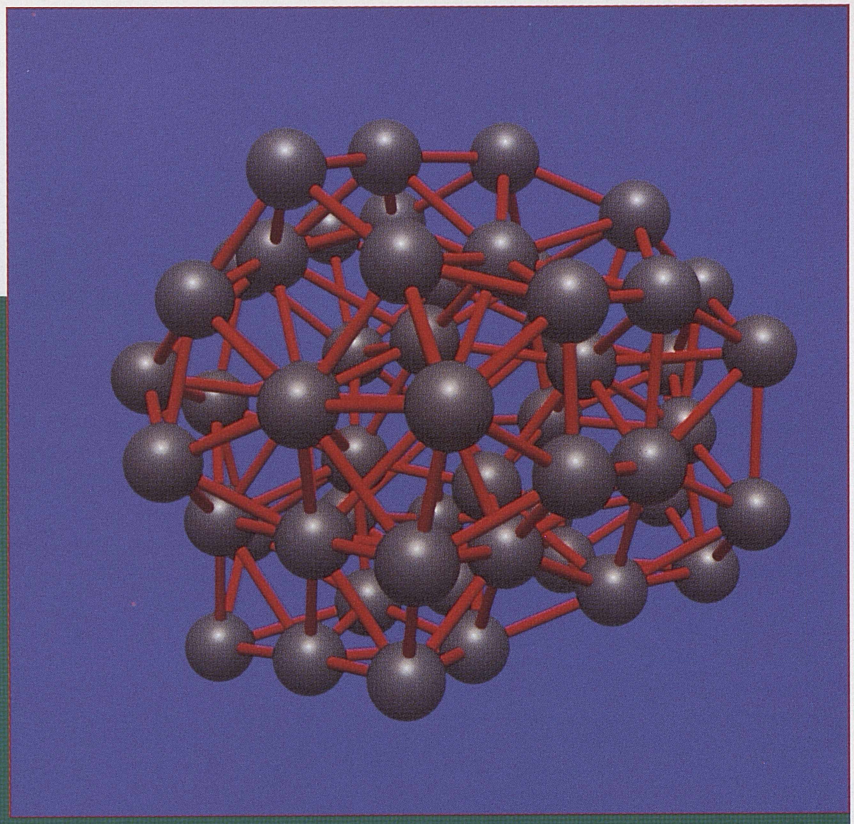

This image, from calculations by Jerzy Bernholc, represents a 55-atom cluster of aluminum after a computer simulation, using quantum molecular dynamics, that raises it to a high temperature and lets it slowly cool — a process called simulated annealing. With such calculations, Bernholc is investigating structural distortions in metal clusters. Joel Welling of the Pittsburgh Supercomputing Center generated the image from Bernholc' data.

"The question is why," says Bernholc, and his QMD calculations at Pittsburgh gave the answer. "We came up with a model of why nitrogen wants to distort and entered that as a starting configuration. The nitrogen distorted and didn't want to go back, showing that the model was correct. The computed atomic positions are in excellent agreement with experiment."

Doping Diamond

Along with resolving a 30-year-old theoretical issue, Bernholc's diamond research has important implications for industry. Like silicon, diamond is a semiconductor, and it can make an extremely fast "chip" for electronic applications — easily twice as fast as the fastest existing chips — once the problems of producing diamond thin films are solved. To make a diamond transistor, however, requires that diamond be "doped" — the term for adding impurities that control current flow.

"The problem," explains Bernholc, "is that one can dope diamond p-type (to create a positive charge), but it has been extremely difficult to dope it n-type (for a negative charge)." Though nitrogen would be a natural candidate as an n-type dopant (because it has one more electron than carbon), it doesn't work — for reasons explained by Bernholc's calculations on the nitrogen distortion. With additional calculations, Bernholc has been able to predict that phosphorous, lithium and sodium will work as n-type dopants in diamond. Several research groups are now working on laboratory experiments to confirm these predictions. (MS)

References:

S. A. Kajihara, A. Antonelli, J. Bernholc and R. Car, "Nitrogen and Potential n-type Dopants in Diamond," *Physical Review Letters* **66**, 2010 (1991).

Jae-Yel Yi, Dirk J. Oh, and J. Bernholc, "Structural Distortions in Metal Clusters," *Physical Review Letters* **67**, 1594 (1991).

This research is supported by the Office of Naval Research and the National Science Foundation.

Coping with Stress

Impurity Segregation at a Moving Solid-Liquid Interface
Paulette Clancy, Cornell University

Cracking an Egg

Semiconductors are the backbone of the electronics and computer industry. As components of integrated circuits, they amplify electrical signals and also switch them on and off. Of all the semiconductor materials, silicon is the most widely used. However, it must be combined with other materials — commonly known as dopants — to alter its electronic properties, and that usually requires implanting dopants, such as arsenic or boron, into silicon's crystalline structure. Because technicians can control the depth of the dopants, they can produce a variety of semiconductors with desirable properties.

However, the implantation technique "is a bit like using a sledge hammer to crack an egg,'' says Paulette Clancy, a chemical engineer at Cornell University. "You create a lot of damage in the crystal, and until you get rid of the damage, the material cannot function as a (semiconductor) device." The common solution in industry is putting the doped silicon into an oven for several hours and allowing the heat to anneal the damage.

Researchers also want to know how dopants move in semiconductor materials during the repair process. But instead of using ovens, they may use lasers to study the effect of the energy input on dopant movement. For instance, a low-energy laser anneals damage without melting semiconductor material. At higher energies, however, the laser melts the material, which then regrows as a crystalline structure. Researchers can control the melting and resolidifying stages so the recrystallization occurs quickly or slowly — a factor that affects dopant movement. If the material regrows slowly, the dopants move to the material's surface, and if it regrows quickly, dopants become trapped in the material — where they were before any melting had occurred.

Paulette Clancy

Playing a Game of Pool

"The goal is to understand the dopant's relationship with the semiconductor material during crystallization and also how that material relieves the stress created by the dopant atoms," Clancy says. But to examine the underlying causes of those phenomena, it's necessary to probe on the atomic level. "We got involved at this point to see if we could use a supercomputer to simulate the melting and subsequent regrowth of silicon after being bombarded with a laser burst," Clancy says, referring to the CRAY Y-MP at the Pittsburgh Supercomputing Center. "This may allow us to predict the conditions necessary to produce material with desired properties."

Clancy's three-dimensional simulations begin at the implantation stage and end at crystallization. Typically, they contain 5,000 to 10,000 silicon atoms arranged as in the crystal structure of diamond (and interacting according to the Stillinger-Weber model for the interatomic forces) — believed to be the most stable structure for silicon. In addition, they usually contain several hundred impurity atoms, such as arsenic. "The atoms move around for 300,000 to 400,000 time-steps under the influence of their neighbors, as well as simulated manufacturing conditions," Clancy says, "so that involves a big computational burden. That's why we need the supercomputer," adding that the entire pocketless pool game simulates only 600 to 800 picoseconds of real time (a picosecond is one-trillionth of a second).

Simulated Laser Bursts

After the impurities have found their favored locations in the simulations, it's time for the laser bursts and examining the transition from solid-to-liquid-to-solid. "I had a graduate student who finished up last year, and his research was the effect of stress on the segregation of impurities when a system is melted and regrown. We used to jokingly refer to his project as 'coping with stress,' which is really what we are looking at. We want to know how silicon copes with the stress of having particles in it that are different from the majority of the surrounding particles" — for instance, a size discrepancy.

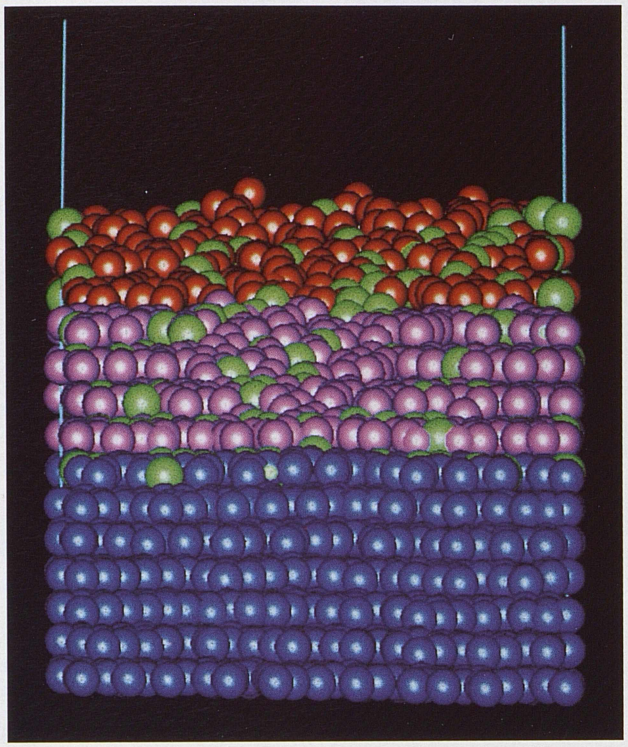

A cutaway view of a hypothetical alloy that has melted and begun to solidify, where the impurity atoms (green) are 10 percent larger than the semiconductor atoms. The impurity atoms have segregated toward the surface, which is still liquid (red). Some impurity atoms also are in the liquid/solid interface (magenta), and virtually none are located in the solid phase (blue). There are 3,500 atoms represented here, including about 260 impurity atoms.

From this work, Clancy has determined how segregation varies with the size difference between impurity and semiconductor atoms. At a critical stage of size difference, the increase in strain creates defects in the semiconductor. Whether this prevents the material from acting as a semiconductor depends on the number and type of the defects, says Clancy, which may vary from a space between atoms to a missing row of atoms. If size difference becomes too large, the semiconductor will no longer accommodate the stress, and it becomes amorphous. Although laboratory researchers had known these results previously, Clancy's simulation provides a detailed explanation of the mechanism.

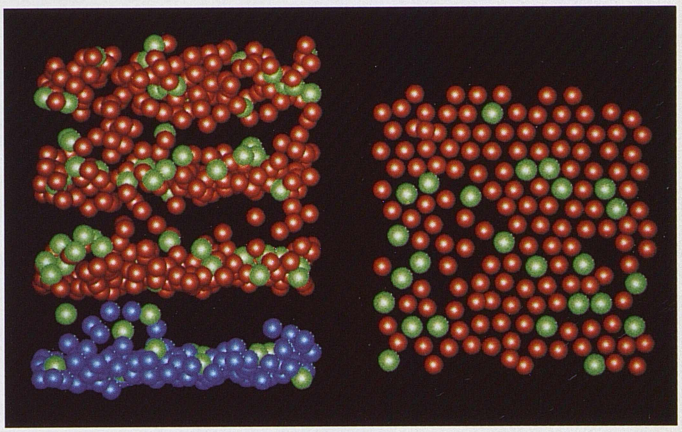

The left side represents four planes of the solidifying semiconductor represented above, where green spheres represent impurity atoms. The top three layers occur in the liquid/solid phase, becoming progressively more ordered as they begin to solidify. The blue layer is the last liquid/solid layer, being closest to a solid. On the right-hand side of the figure, the disordered region is clearly shown in the "top down" view of the top plane.

But there's a catch. Because experimental laser bursts and subsequent regrowth occur over a time period too long to simulate, Clancy needed to speed up those actions to make them manageable. "So our 'laser bursts' tend to be more like 20 picoseconds (of real time) with regrowth taking several hundred picoseconds," Clancy says. "We believe that the mechanisms remain the same even though the time scales are quite different."

Clancy says she now better understands how impurity atoms move during resolidification, whether the growth front is moving slowly or rapidly. "We have learned more about the interaction of the impurities and the substrates and how it determines the quality of the regrown material," she says. "This will have implications for predicting the properties of suggested new materials." (SE)

References:

Paulette Clancy, "Computer Simulation of Crystal Growth and Dissolution in Metals and Semiconductors," *The International Journal of Supercomputer Applications* **5**, 10 (1991).

Stephen J. Cook and Paulette Clancy, "Solute Trapping at a Rapidly Moving Solid/Liquid Interface for a Lennard-Jones Alloy," *Molecular Simulation* **5**, 99 (1990)

This research is funded by the National Science Foundation and the Semiconductor Research Corp.

Turbulence

Zooming in on Vortices

Visiometrics for Vortex Dominated and Turbulent Flows
Norman Zabusky, Rutgers University

A Swirling Puzzle

Walk down to a river bank and look at the bridge posts near the water. The water careens off the post at a certain velocity and spins into a rotating pattern, known as a vortex or eddy.

Now look at a high-flying jet and notice the two white trails that lag behind. They're actually ice crystals from the engines' combustion process that have formed vortex patterns. Over time, these vortices change and they connect with each other, forming turbulent regions in the atmosphere. "This connection process is important because it's one of the fundamental aspects of turbulence," says Norman J. Zabusky, a professor of computational fluid dynamics at Rutgers University.

Norman Zabusky (left), Deborah Silver and Richard Pelz.

For the last few years, Zabusky and his colleagues, Richard Pelz and Deborah Silver, have studied this connection process at the Pittsburgh Supercomputing Center, and they're discovering how it occurs. "It's a piece of the puzzle in turbulence and if more pieces are revealed, then you could hope to predict or control certain aspects of turbulence," Zabusky says. For instance, imagine looking at an airplane from its front when it's flying. Off the right wing, a vortex spins clockwise, and off the left wing, another vortex spins counterclockwise. The unstable vortices approach each other and connect. "One factor affecting the rate at which large airplanes land at an airport is the strength of the vortices generated from the airplanes landing previously. A small aircraft may have to wait longer to land because of the vortices," Zabusky says.

This process also occurs in combustion engines. There, it's important to control the turbulent mixing of air and gas to improve performance.

The Unpredictability of Vortices

To make the appropriate computer models, Zabusky and his colleagues represent vortices in three dimensions and use the classic equations for studying fluid flow, the Navier-Stokes equations. These non-linear equations describe the behavior of vortices, and "supercomputers are needed," says Zabusky, "because mathematicians cannot solve these non-linear equations with pencil and paper because they're too difficult."

The three-dimensional computer grids for the vortex models contain one million sites and require massive computer capability. For instance, in real time, two vortices swirl around and may connect within a fraction of a second, but because of the event's complexity, a supercomputer requires at least three hours to simulate it.

The Visual Key to Turbulence

After the computations at Pittsburgh, Zabusky transfers the results to his Rutgers laboratory, where he produces stunning graphics. "We found that in order to understand three-dimensional problems, it's essential to use a visualization and quantification process that we call visiometrics," he says.

By manipulating the graphics on a computer terminal's display, Zabusky's group has learned how vortices connect with each other. Essentially, it's a two-step process, where two vortices, such as those that come off a plane's wing tips, begin to drift toward each other. "It turns out that vortices spinning in opposite directions want to attract each other. They fall toward each other and while doing that, they're being stretched by themselves, which is a consequence of their attraction. At the same time, the region between them, which contains small amounts of vorticity, is being stretched. A bridge forms, and the two vortices connect. Then, the intense middle region is spit out in the form of a distorted ring and adds to the turbulence," he says.

As the middle section is produced, the two original vortices reform and move away. And the reconnection process occurs over and over again. However, Zabusky says he hasn't been able to reproduce consecutive reconnections on the computer. In addition, the properties of the middle region have remained elusive because it dissipates too rapidly. "Understanding the properties of the middle segment would give more reliable insight into how larger vortices divide into smaller vortices," Zabusky says, adding that it's possible that more than one such region may be spit out and add to the overall turbulence.

To obtain those results, Zabusky says he must improve the grid's resolution by upping the number of grid sites to 100 million, but that requires increasing the computer's memory. "Right now, we're limited by the computer's capability, and we are beginning the move to massively parallel computers like the Thinking Machines' CM-2 at the Pittsburgh Supercomputing Center because it has a larger memory than the CRAY Y-MP." (SE)

References:

N.J. Zabusky, O.N. Boratav, R.B. Pelz, M. Gao, D. Silver, and S.P. Cooper, "Reconnection in Orthogonally Interacting Vortex Tubes: Direct Numerical Simulations and Quantifications," *Physics of Fluids A* **4,** 581 (1992).

N.J. Zabusky, O.N. Boratav, R.B. Pelz, M. Gao, D. Silver, and S.P. Cooper, "Emergence of Coherent Patterns of Vortex Stretching during Reconnection: A Scattering Paradigm," *Physical Review Letters* **67,** 2469 (1991).

Francois J. Bitz and Norman J. Zabusky, "David and 'Visiometrics': Visualizing and Quantifying Evolving Amorphous Objects," *Computers in Physics* **4**, 603 (1990).

This research is supported by the National Science Foundation and the Office of Naval Research.

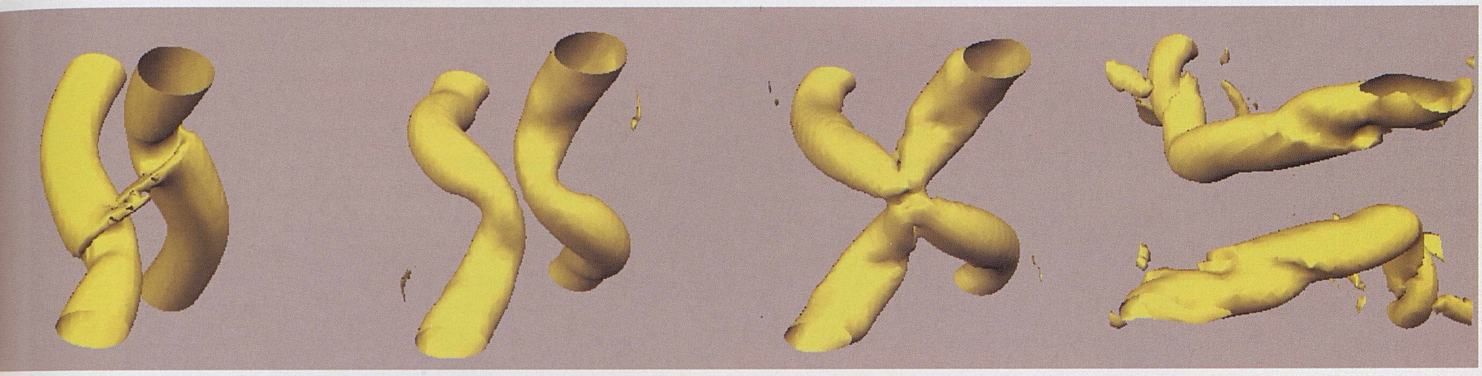

"Bursting" Reconnection of Two Vortex Tubes

These images show four snapshots in the evolution of two tubelike vortex regions of equal strength, as simulated in calculations by Olus Boratav and Richard Pelz. Initially, the vortex tubes are isolated, one beside the other crossing at a 90° angle. The tubes collapse toward each other and form an X-like configuration. In this bridging region, the local vorticity becomes intense. The tubes then "reconnect" forming a new vortex pattern — the outgoing state of the event. Visualization by M. Gao and D. Silver.

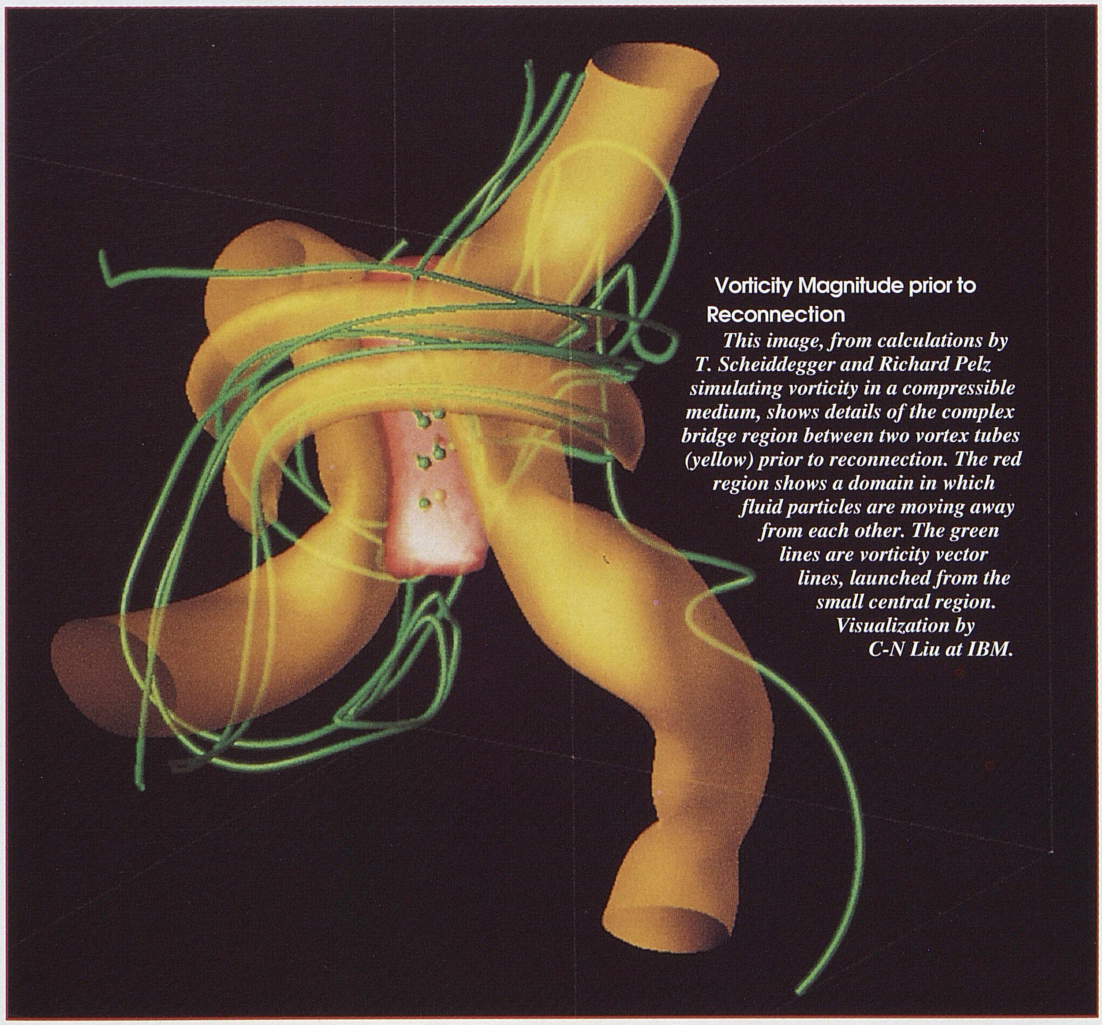

Vorticity Magnitude prior to Reconnection

This image, from calculations by T. Scheiddegger and Richard Pelz simulating vorticity in a compressible medium, shows details of the complex bridge region between two vortex tubes (yellow) prior to reconnection. The red region shows a domain in which fluid particles are moving away from each other. The green lines are vorticity vector lines, launched from the small central region. Visualization by C-N Liu at IBM.

Boundary Dynamics

A Sticky, Slippery Problem

Origin of Stick-Slip Motion in Boundary Lubrication
Mark Robbins, The Johns Hopkins University

Cacophony

He slowly opens the door, which creaks as the novice violinist practices inside. He sits down and listens to the squealing music, but as the violin bow continues to rub across the strings, he gets up and shouts, "Stop! Stop! I can't take it anymore." Rushing to the front of the room, he runs a piece of chalk over the blackboard, creating yet another obnoxious sound.

Creaking joints, violin music and squeaky chalk — all are examples of stick-slip motion, where two substances don't glide smoothly over each other, but instead move jerkily, alternately sticking and then slipping.

Researchers have known about the phenomenon for centuries but were unable to probe its microscopic origins. The first clues came from recent experiments at the University of California-Santa Barbara and the University of Illinois, where researchers studied friction between plates of mica separated by a film of lubricating fluid. When the spacing exceeded 10 molecular diameters, the plates glided effortlessly, but as the spacing decreased below four or five diameters, the liquid film sometimes acted like a solid.

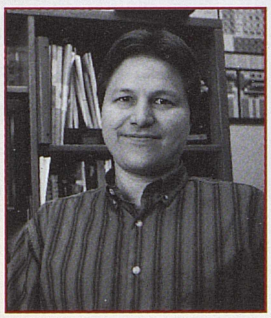

Mark Robbins

"Once the film solidifies, the plates become welded together," says physicist Mark Robbins. "The top plate doesn't move when you tug on it gently, but when you pull hard enough, it suddenly begins to move. Then, if you pull slowly, the plate alternately sticks and lurches forward, but if you pull faster, it slides smoothly." These observations motivated Robbins to study stick-slip motion on the molecular level at the Pittsburgh Supercomputing Center, where he used the CRAY Y-MP to simulate the plates and lubricating fluid.

Pushing a Sled

All explanations of stick-slip motion, which has technological significance because it greatly increases wear on moving parts, start from the observation that you must apply more force to start an object sliding than to keep it moving. This is probably familiar to anyone who has dragged a heavy object or pushed a sled. "You push harder and harder on a sled without budging it," says Robbins, "then suddenly it rushes forward and leaves you behind."

While the old model of stick-slip motion assumed that friction on the moving sled decreased continuously as the sled's velocity increased, Robbins believed that there was a sudden drop in friction, accompanied by an instantaneous change in how the two surfaces — the sled and the ground — interacted with each other at the microscopic level.

"When two surfaces are stationary, the molecules at the interface have time to order themselves in the most energetically favorable way. Then if you want the sled to move, you have to make those molecules less energetically happy, and they are going to complain. You have to overcome their resistance to make the sled start sliding. Once that happens, the surface molecules don't have enough time to order themselves. They cannot fit into the ground's microscopic roughness and cannot make chemical bonds," creating a less-ordered sliding state, which exerts less friction.

A Discontinuous Transition

At the Pittsburgh Supercomputing Center, Robbins and former graduate student Peter Thompson, now at the Exxon Research and Engineering Company in Annandale, N.J., created model systems including two solid walls and a lubricating liquid of spherical molecules, ranging in thickness from 1 to 20 layers. Their computations showed that when the spacing decreased to 2 to 4 layers, the lubricant became crystalline, and the molecules were neatly stacked in rows, creating the lowest energy state.

"The film resists motion of the top plate just like a normal solid would," Robbins says. But when the force on the top plate increases enough, the film can no longer resist. "Once the top plate slides forward by only an atomic spacing, the order in the film is lost, and the molecules begin dancing around. The top plate can then slide forward more easily." This "melting" of the film is discontinuous, like the abrupt transition from ice to water.

Robbins and Thompson also explained why stick-slip motion occurs only when the top plate moves slowly. If it moves slowly enough, its kinetic energy can be converted into potential energy in the film. The plate then stops, and the film recrystallizes. However, when the velocity is too high, the plate's kinetic energy is too large, and the film lacks the strength to stop it. As a result, no sticking occurs.

"This necessary condition for stick-slip and the idea of a discontinuous change in the frictional force with velocity are markedly different from previous theories," Robbins says. Two papers detailing Robbins and Thompson's simulations appeared in *Science*. Based on his studies and those at UC-Santa Barbara, Robbins says the ideal lubricant for ultrathin spaces should contain branched molecules, such as some hydrocarbons, instead of spherical ones, because it's more difficult to order them and the lubricant will tend to remain liquid. (SE)

References:

Mark O. Robbins and Peter A. Thompson, "Critical Velocity of Stick-Slip Motion," *Science* **253**, 916 (1991).

Peter A. Thompson and Mark O. Robbins, "Origin of Stick-Slip Motion," *Science* **250**, 792 (1990).

Research funded by the National Science Foundation and the Sloan Foundation.

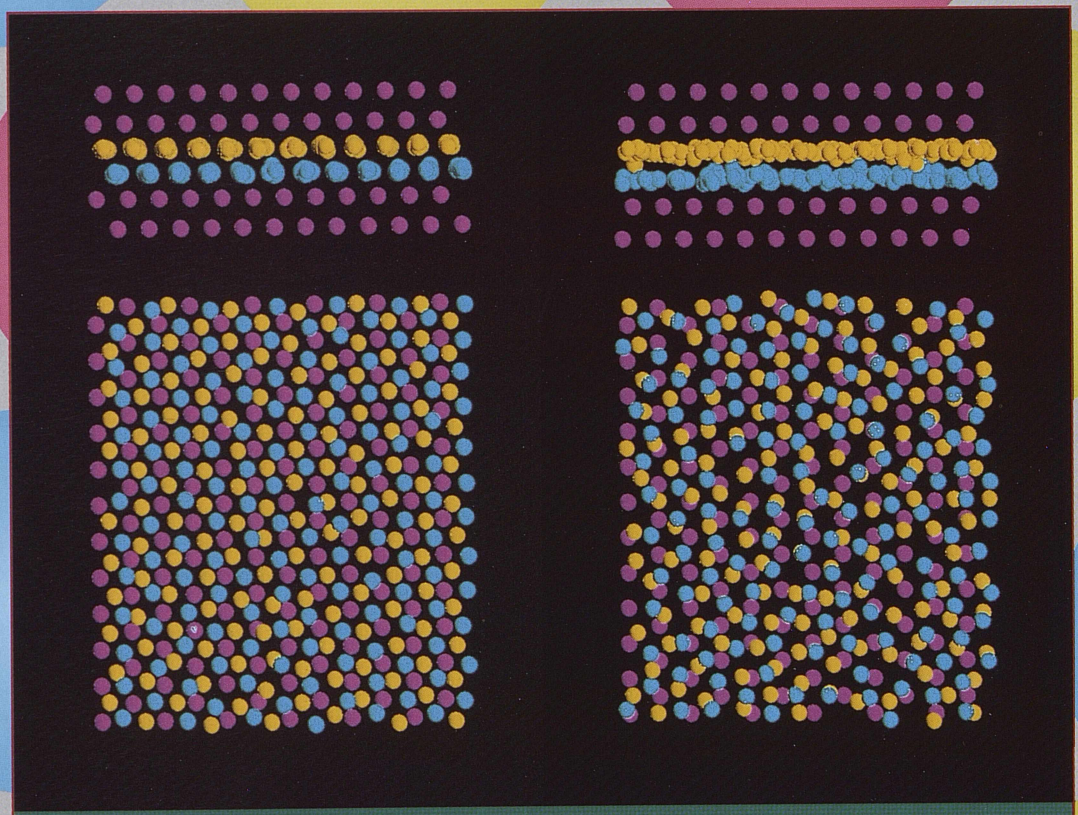

A thin film (blue and yellow) is confined between two solid plates (magenta) in the stuck (left) and sliding (right) phases of stick-slip motion. Top panels in both frames show cross sections of the plates and film, and the bottom panels show registry between layers as seen looking up toward the bottom layer of the top plate.

In the stuck state, the thin film molecules are aligned in layers, showing order characteristic of a crystal. The bottom panel shows that molecules in each layer of the film center themselves between three molecules in the plate surface.

In the sliding state, the film is more characteristic of a disordered fluid, although there still is some tendency for molecules to align in layers.

A Closer Look at Fluids

Molecular Dynamics of Fluid Interfaces
Jayanth Banavar, Penn State University & Joel Koplik, City College of New York

Tracking Down Molecules

It's a common problem in industry. Technicians want to separate two liquids — for instance, when they inject water into an oil reservoir to force out the black stuff — but one liquid, oil in this case, tends to divide into droplets, and it's often necessary for the drops to coalesce before the overall separation can occur. Small drops of oil, for example, can be trapped in the porous rocks of reservoirs, which makes recovery difficult.

Just the opposite problem arises during another industrial situation — when it's beneficial to divide drops into smaller drops. During the manufacture of cosmetic creams, for instance, various fluids must be well mixed to produce a uniform product.

Traditionally, scientists have used the Navier-Stokes equations to solve these fluid-flow problems, but they cannot provide crucial information about how drops' boundaries interact on a microscopic level. "In many such cases, it is not clear either from laboratory experiments or from reasoning based on macroscopic equations how the microscopic region should be treated," says physicist Jayanth Banavar, who, along with fellow physicist Joel Koplik, has explored the problem at the Pittsburgh Supercomputing Center.

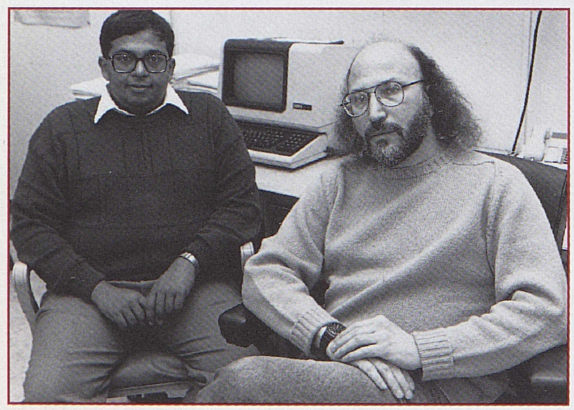

Jay Banavar (left) and Joel Koplik

Their computations on coalescence and fission of fluids are the first three-dimensional molecular models of these processes. "If I know where all the molecules are today, and I know what forces are acting on them, I can predict — within computer error — where they will be tomorrow," Banavar says, referring to the models, some of which contain as many as 32,000 molecules. "That is the basic idea of our work. It is very simple, but nevertheless, without a supercomputer, you really can't go very far because you need a large enough collection of molecules."

The No-Slip Boundary Problem

Banavar and Koplik, in collaboration with physicist Jorge Willemsen at the University of Miami, began using molecular simulations to study another microscopic boundary situation, and their current research builds on that work. Since the 19th century, scientists had assumed a fluid's velocity in a pipe was greatest at its center, but dropped to zero at its wall, or boundary — a phenomenon commonly known as the no-slip boundary condition.

"But there isn't a completely convincing theoretical analysis of this condition," says Koplik, "because it's a complicated problem involving many molecules in the liquid and in the solid. Such things are usually hard to solve, so in order to understand whether this boundary condition is true, we carried out some molecular simulations of the fluid moving past the solid wall."

In this case, the computer results supported the assumptions, but Banavar and Koplik next studied what happened when one fluid pushes another past a solid surface, such as when liquid pushes air up a straw when a person sips a soft drink. According to traditional assumptions, the velocity of the contact line — where the air, liquid and solid meet — would be zero. Their computations showed, however, that this no-slip boundary condition does not hold true.

"If I walk along the interface between the two fluids to the contact line," says Banavar, "it's clear that the fluids are moving at the wall. It's this inconsistency that leads to some physically absurd results when you solve (the large-scale) equations with the boundary condition that the velocity was zero everywhere along the solid wall."

"The idea is that the equations describe motion on relatively large scales," adds Koplik, "and to get under it, so to speak, you have to go to the next smaller scale — the scale of molecules — and extract the behavior on that scale, which is either no-slip or slip depending upon the circumstances. Then you can return to the large scale and use the equation with the correct boundary condition."

Fission and Fusion of Drops

Since Banavar and Koplik's work on the no-slip boundary problem, they have extended the insight and method to other fluid boundary phenomena. Their CRAY Y-MP calculations on coalescence and fission of fluids represent the first molecular level studies of these interactions, and their simulations reveal a level of detail not previously accessible.

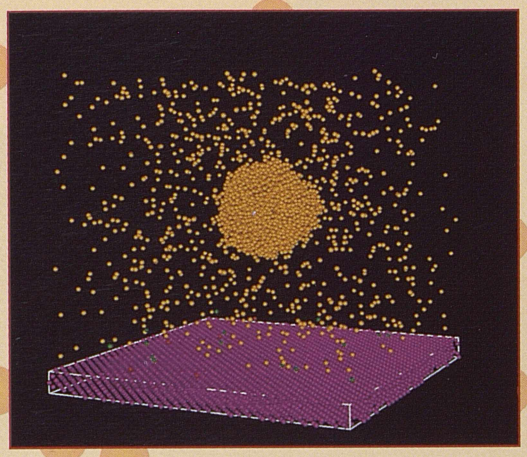

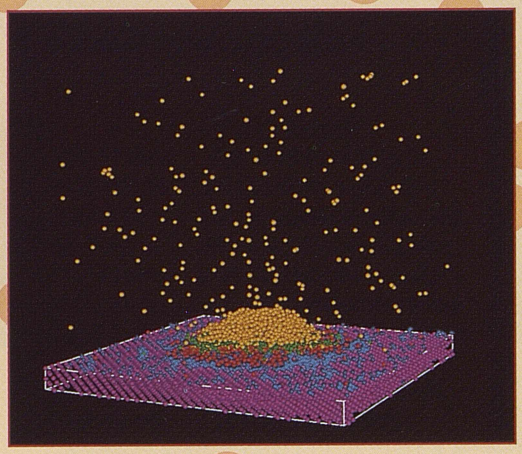

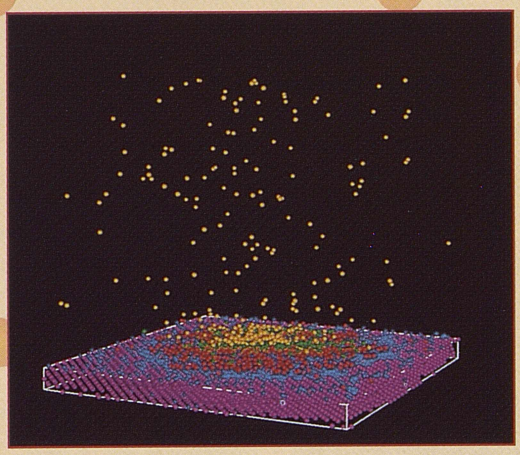

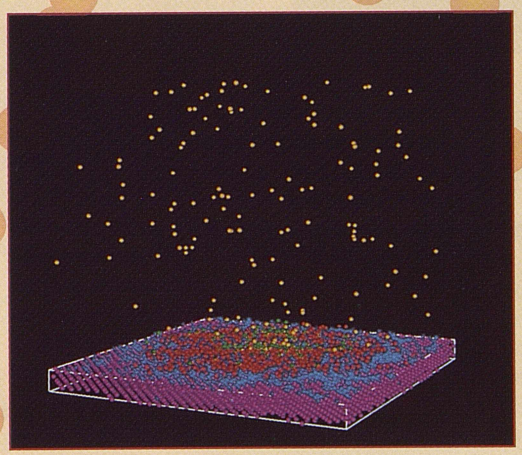

When a drop divides into smaller drops, the molecules jiggle randomly before they breakup, and as the drop distorts into a dumbbell shape, the molecules withdraw gradually to one side or another. Then, the dumbbell's handle becomes thinner and thinner until it runs out of molecules, and two separate drops remain.

Coalescence of two drops is slightly more complicated. Just outside one drop, one or two of its molecules may hang out — fairly typical for drops in liquid. As another drop approaches, the outcast molecules — instead of being attracted to their home base — develop an affinity for the new arrival. Soon, other molecules from the first drop follow the outcast molecules, and a thin connecting filament forms between the drops. As more and more molecules make the pilgrimage, the filament thickens, until the two drops are sucked together and smoothly merge.

"The simulations tell you that fission and coalescence of liquids are smooth, gradual events and not rapid, violent rearrangements of molecules," Koplik says. "On the small scale, we can follow the molecular motion and describe it. However, we have not been able to quantify the average behavior on small scales because of limitations in computer power."

The massive number of molecules is only one challenge for today's supercomputers; the simulations also must be three dimensional because "realistic drops and boundaries are three dimensional," Koplik says. A three-dimensional simulation of coalescence requires much more time — 10 hours at 100 million calculations each second — on the CRAY Y-MP than a two-dimensional approach to the same problem. On a high-end workstation, the same simulation might require 500 hours. (SE)

References:

W. Ma, A. Maritan, J.R. Banavar and J. Koplik, "Dynamics of Phase Separation of Binary Fluids," *Phys Rev A* **45**, R5347 (1992).

J.-x. Yang, J. Koplik and J.R. Banavar, "Molecular Dynamics of Drop Spreading on a Solid Surface," *Phys Rev Lett* **67**, 3539 (1991).

J. Koplik, J.R. Banavar and J.F. Willemsen, "Molecular Dynamics of Fluid Flow at Solid Surfaces," *Phys Fluids A* **1**, 781(1989).

Research funded by the National Science Foundation and the NASA Microgravity program.

Drop Spreading on a Solid Surface

In this sequence, from computations by Banavar and Koplik, a drop of fluid surrounded by vapor (yellow) settles onto a solid surface (purple) and spreads into several layers. Vapor molecules adsorbed onto the solid are blue, and the first, second and third layers of the spreading drop are red, green and yellow, respectively. The solid has 9,000 molecules, and the liquid and vapor have 4,000. The simulation covers one-billionth of a second in real time.

Experimental studies had indicated that drops sometimes spread on a surface in terraced layers. Banavar and Koplik's molecular dynamics simulations, however, revealed another level of detail. "Because we did the simulation at the molecular scale, we could study where all the molecules were going — something that couldn't be done in the experiments," Koplik says. "And we could see that although there were well-defined layers, they were not rigid. Molecules were diffusing back and forth and up and down." By the third frame, for instance, molecules previously in the yellow-colored layer begin moving downward to the red and green layers.

Koplik says the results, which were the first of their kind, surprised him because the structured layering found in the previous experiments had suggested that only a small amount of diffusion occurred in either direction. The computer-generated findings, he says, may have industrial implications. "If the molecules were not diffusing, then the layers would be like a impermeable rubber sheet and would prevent gases from getting at the solid very easily."

Excited States

Wavepacket Methods for Many-Body Quantum Dynamics
Rob D. Coalson, University of Pittsburgh

Quantum Mechanics in Action

What happens in a chemical reaction, really? Start with compound A, add something — heat, light, another compound — and it changes to compound B. But what happens along the way? How does one stable combination of atoms (a molecule) break apart, and how do the fragments regroup into a different structure?

Traditional laboratory methods tell about the beginning and end points of a reaction, but theoretical chemists, like Rob Coalson, are turning their attention increasingly to the in-between stages. For these short periods of time, often measured in femtoseconds (a millionth of a billionth of a second), chaos reigns in the molecular realm; electrons — "excited" by added energy — shake free from nuclei, chemical bonds break down and new ones form. Using quantum theory and supercomputing, Coalson's research opens a window onto this microscopic world, where a chemical reaction is a revolution.

"We're interested in the dynamic transformations produced by disturbing an isolated system," says Coalson, whose research group at the University of Pittsburgh includes five graduate students, two post-doctoral fellows and two foreign scholars. In quantum terms, they investigate the chemistry of excited states. They develop algorithms for computing the dynamics of chemical reactions, and they are among relatively few research groups in the world who, by doing this, are pushing applied quantum theory into uncharted territory.

Rob Coalson with research colleague Michael Messina (seated) and University of Pittsburgh chemistry graduate students Paul Harmon (left) and Val DeVito (right). Coalson's computations to date have focused on solid environments, such as LiF, and he sees liquid environments, reactions with molecules in solvent, as the future challenge of his work.

Beyond Electronic Structure

Most quantum chemistry involves calculations of "electronic structure." Powerful tools such as molecular orbital theory and the GAUSSIAN package of software, developed at Carnegie Mellon University, make it possible to solve Schrödinger's equation — the fundamental law of quantum theory — with reasonable accuracy for the electronic structure of many small and medium-sized molecules. The outcome is a microscopically precise description of the stable molecular structure.

Coalson's research begins with this information and then finds ways to solve Schrödinger's equation for the much more complicated situation where added energy disrupts the complex web of spinning electrons that holds the molecule together. "As useful as electronic structure is," says Coalson, "it can't tell the whole story about chemical reactions and other intrinsically time-dependent processes. Anything can happen," he adds. "The system can break up; it can skitter into little pieces. It's like taking a piece of chalk and throwing it against the wall. It may just bounce off, and it may explode. You don't know until you do the calculation."

Dirac's Ghost

The principal workhorse that allows Coalson to do quantum dynamics calculations on large systems is a solution method called the time-dependent Hartree method (TDH). "It is the best and, in some respects, only algorithm currently capable of yielding the desired information," says Coalson. Though TDH has been applied in computational quantum dynamics research only within the last ten years, it harks back to one of the pioneers of quantum mechanics, Paul Dirac.

Years before electronic computing, Dirac recognized that the real challenge of quantum mechanics was not reliability of the theory but getting solutions from the equations. "The ghost of his words," says Coalson, "still haunts us.

"The underlying physical laws necessary for the mathematical theory of a large part of physics and the whole of chemistry are thus completely known, and the difficulty is only that the exact application of these laws leads to equations much too complicated to be soluble."

— Paul Dirac (1902-1984)

LiF & CH_3Br

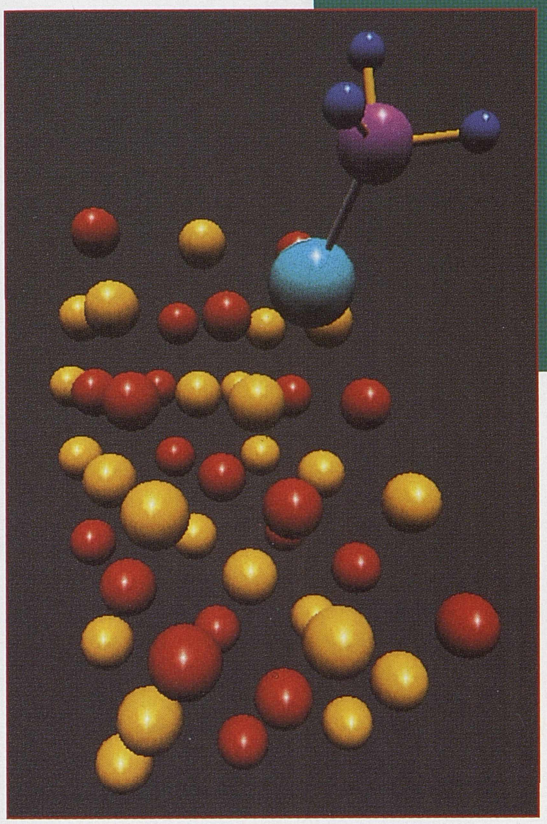

This model depicts methyl bromine (CH_3Br) bonded onto a lithium fluoride (LiF) surface. Lithium atoms are yellow, fluoride red. Rob Coalson's research group has done quantum theory calculations that track the vibrational states of the CH_3 fragment when it photodissociates from the Br atom (aqua).

This remains a challenging endeavor in 1992." The theory, in principle, makes it possible to exactly predict the outcome of any chemical event, yet even with today's most powerful supercomputers, like the CRAY Y-MP, the equations can be solved only in limited cases and usually only by approximating the full statement of Schrödinger's equation.

Around 1930, Dirac worked out an approach for treating electron motion in a time-dependent fashion, and this, basically, is the approach Coalson has turned back to. Put simply, TDH represents the wave function — the mathematical statement of quantum properties — for the entire excited-state molecular system as a collection of wave functions for the nuclei of each atom and lets each of these evolve separately in time. For relatively simple calculations (two mechanical degrees of freedom), the TDH approximation reduces computing by a factor of ten over an exact solution calculation, says Coalson, and this computational gain increases with more complex systems.

Computational Experiments

Take methyl bromine (CH_3Br) chemically bonded onto a lithium fluoride (LiF) surface and zap it with a laser beam of ultraviolet light. The added energy sets up vibrations in the LiF and the CH_3Br. Experimental studies of this system, using optical spectroscopy, show that one quantum unit of light (a photon) can break (photodissociate) the CH_3Br into a CH_3 fragment and a Br atom (that remains attached to the LiF surface). Using TDH, Coalson and graduate student Michael Messina completed calculations on this system that reproduce much of the observed data and, furthermore, predict the distribution of vibrational states for the dissociated methyl fragment, information that was not previously available for this system.

Coalson credits another research group (Kosloff, Gerber, Ratner and co-workers at Northwestern University) with being the first to implement TDH, but his group has gone beyond other work in developing an algorithm for processes that take place on electronically excited molecules, such as CH_3Br photodissociation.

Coalson sees his research as "computational experiments" that form a vital interactive loop with laboratory work. "Experiment feeds theory, and theory feeds experiment." Measured data forms the starting point of the CH_3Br computations, for instance, and the results suggest further experimental work, such as using laser spectroscopy to analyze vibrational states of the dissociated molecule. (MS)

References:

Janet R. Waldeck, José Campos-Martínez & Rob D. Coalson, "Application of a coupled-surface time-dependent Hartree grid method to excited state optical spectroscopy," *Journal of Chemical Physics* **94**, 2773, **95**, 1430 (1991).

M. Messina & R. D. Coalson, "Time-Dependent Hartree Wavepacket Dynamical Techniques for Computation of Electronically Excited Optical Spectra of Many-body Quantum Systems," *Journal of Chemical Physics* **90**, 4015 (1989).

Epigraph: P.A.M. Dirac, *Proceedings of the Royal Society* (London) **123**, 714 (1929).

This research is supported by the National Science Foundation, the Petroleum Research Fund, the Alfred P. Sloan Foundation, and the Camille and Henry Dreyfus Foundation.

Great Balls of Fire

Realistic Simulations of Star Cluster Dynamics
Stephen L. W. McMillan, Drexel University

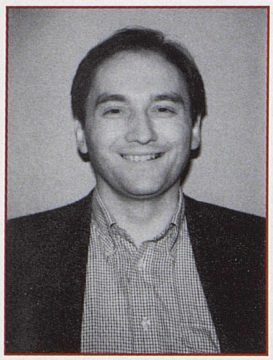

Steve McMillan. His code for N-body simulations of globular clusters runs at 90 to 100 million floating-point operations per second on one CRAY Y-MP processor.

Star Clusters

Twentieth century astronomy has revealed that stars are grouped in immense collections — galaxies. Like scattered islands in the vast sea of space, galaxies are the main components of large-scale structure in the universe. In recent years, many astrophysicists and astronomers wanting to understand how galaxies formed have begun looking closely at one of their most striking features: globular star clusters — great ball-shaped formations of thousands to hundreds of thousands of stars bound together by gravity. About 200 globular clusters orbit the outer framework — the halo — of our own galaxy, the Milky Way.

"Globular star clusters," says Steve McMillan, "are the oldest star systems in the galaxy. As such, they're giving us information about conditions very early in galaxy formation — a stage that is the most difficult, from a theoretical standpoint, to understand." For his 1983 dissertation in astronomy at Harvard, McMillan developed computer code to simulate the "core collapse" of globular clusters, and his continuing work on clusters — including three current projects at the Pittsburgh Supercomputing Center — has gone far to establish computer modeling as an invaluable tool in this field.

Married by Gravity

McMillan's CRAY Y-MP simulations of star cluster dynamics are the first studies to conclusively show that binary stars play an important role in the evolution of globular clusters. Until about five years ago, he explains, most people thought all the stars in globular clusters were single stars, like the sun. Improved observational techniques from ground-based telescopes, however — now supplemented by data from the Hubble Space Telescope — are providing strong evidence that many of the stars in clusters are binaries: two stars circling around each other, married by the mutual attraction of their own gravity.

If binary stars were present in globular clusters from the time of their formation, as the new observations suggest, McMillan's computations show that it could dramatically affect how the cluster has evolved — over roughly 15 billion years — to its present form. "It completely changes the dynamics of a cluster," says McMillan, "and the simulations we've done so far demonstrate that when you start off with 10 to 20 percent of the stars as binaries instead of all the stars being single, the cluster evolves much differently."

A Massively Parallel Approach to the Three-Body Problem

In a related project, McMillan and doctoral student Patricia Boyd are examining a famous problem in mathematical physics — the three-body problem. The moon, Earth and sun is a well-known example. "Since Newton," says McMillan, "many mathematicians and astronomers have contributed to our understanding of the great complexity of motion possible when three bodies interact only via their mutual gravitational attraction, but no practical general solution has ever been found, and none is anticipated."

The three bodies of particular interest to McMillan and Boyd are a binary star and a single star — an interaction that (as explained above) bears on the question of core collapse in globular clusters. Their calculations investigate the "initial value space" of this interaction — how the escape angle of the exit star varies according to the angle and position at which the single star approaches the binary.

The image depicts results from 21,931 individual three-body interactions. Each point (pixel) in the image represents a separate initial condition. The color scale, from blue through red to green, corresponds to the escape angle — from 0° to 360° — of the exiting star relative to the approach of the single star. The solid-color regions of this image represent regular regions of initial value space — that is, where the result (escape angle) varies smoothly with changes in initial values. The mottled regions are irregular or chaotic, where small changes in initial values produce large swings in the result.

Boyd is currently developing code to perform these calculations on the Connection Machine, CM-2 at Pittsburgh. Because each set of initial conditions is an independent calculation, the problem, says Boyd, is a natural for massively parallel architectures, and she expects to achieve a significant performance improvement on the CM-2 over the workstation environment at Drexel.

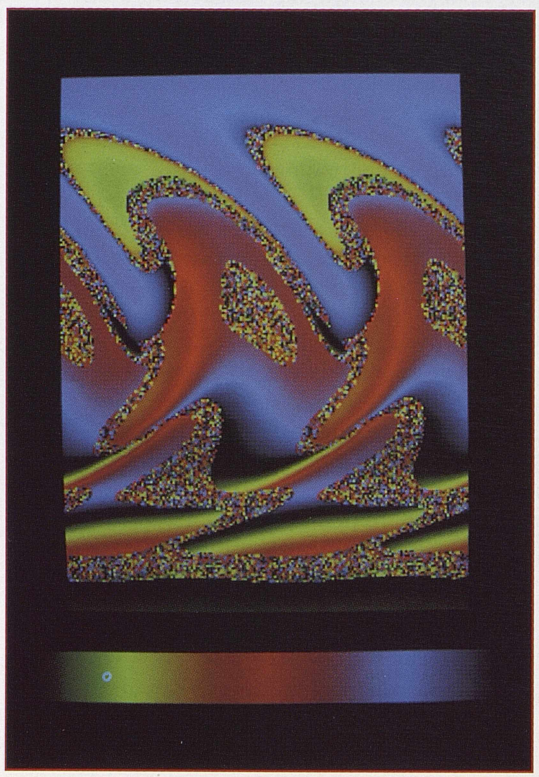

"The Galaxy is an unexplored continent filled with exotic beings of stellar dimensions."
— Carl Sagan, Cosmos

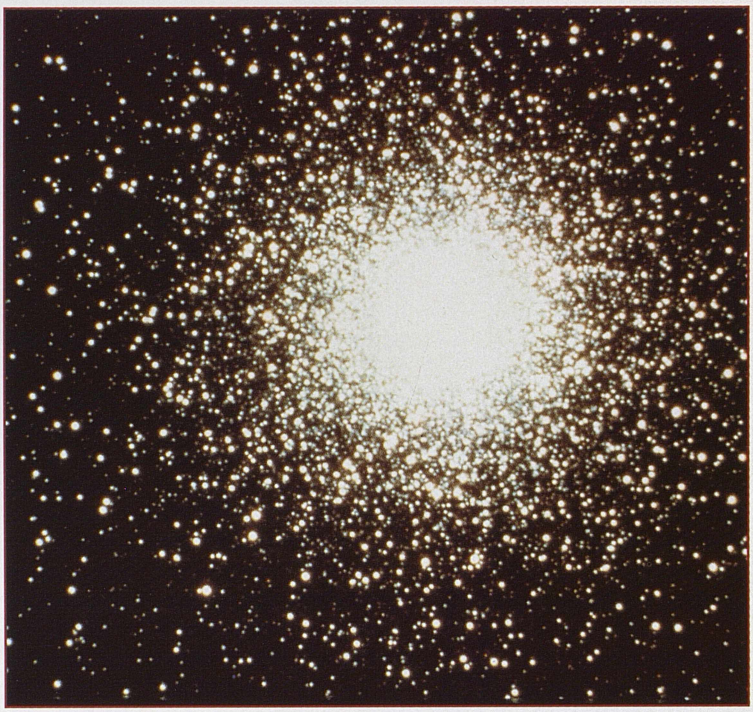

Globular Cluster M15, in the constellation Pegasus, as observed at Mauna Kea Observatory, University of Hawaii, © 1984.

Binary Stars and Core Collapse

The difference has to do with the way central regions of a cluster draw together over time — a phenomenon known as "core collapse." This process in a globular cluster is analogous to what happens during the formation of a single star. Starting with a diffuse gaseous cloud, gravity draws this material together, tighter and tighter, until the critical stage when hydrogen atoms fuse. Once fusion turns on, the energy it produces counteracts gravity and prevents the star from collapsing, at least until its nuclear fuel is used up.

Similarly, in star clusters gravity draws the central stars into a dense, hot, bright core — thousands of stars in close proximity — until an energy source asserts itself and retards the process. Exactly where this energy comes from in star clusters, however, has been an open question. Among the prime candidates are binary stars, and it is this question McMillan addressed in recent work.

"The basic mechanism by which binaries prevent clusters from collapsing," explains McMillan, "is that they act as a heat source." When a third star passes close to a binary pair, several different things can happen. One is that the third star can receive a boost in kinetic energy from interaction with the binary and fly off with more speed than it approached. If enough of these interactions occur in the collapsing core of a globular cluster, the energy released can substantially offset the contraction from gravity.

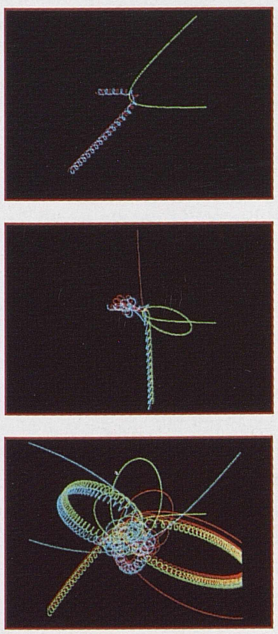

These images, from computations by Steve McMillan and Patricia Boyd, show several instances of the complex dynamics that occur when a single star interacts gravitationally with a binary pair. The initial binary is shown by the red and blue trajectories that begin on the left. The light green trajectory represents the single star. It approaches from the right at varying angles and positions of intersection relative to the binary.

McMillan's simulations of globular cluster dynamics — in collaboration with Piet Hut of the Princeton Institute for Advanced Study — are the first "N-body" simulations (calculations that track the movement over time of every star in a model cluster) to include the effect of binaries. These simulations showed that including binaries resulted in a core significantly larger than comparison runs without binaries. Furthermore, notes McMillan, the result with binaries is "much more in line with observations of clusters." This work demonstrates for the first time that binaries make a significant difference in cluster evolution, and it suggests that they may be the primary energy source retarding core collapse in globular clusters. (MS)

References:

Piet Hut, Stephen L. W. McMillan and Roger W. Romani, "The Evolution of a Primordial Binary Population in a Globular Cluster," *Astrophysics Journal* **389**, 527 (1992).

Stephen L. W. McMillan, "The Role of Binaries in Cluster Dynamical Evolution," and Stephen L. W. McMillan, Piet Hut and Junichiro Makino, "Star Cluster Evolution with Primordial Binaries," in *The Formation and Evolution of Star Clusters* (San Francisco: Astronomical Society of the Pacific, 1991), p. 324 & p. 421.

This research is supported by the National Science Foundation and the National Aeronautics and Space Administration.

Epigraph: C. Sagan, *Cosmos* (New York: Ballantine, 1980), p. 199.

Saving the Ozone Layer

Theoretical Studies of Atmospheric Oxidation Mechanisms of Alternative Halocarbons
Joseph S. Francisco, Wayne State University

Thin-Skinned Planet in Peril

During the last few winters, scientists have detected a hole in the ozone layer over the South Pole. And during this past winter, NASA scientists detected increased levels of ozone-destroying chlorine scattered elsewhere in the upper atmosphere, meaning similar holes may soon appear over parts of North America, Asia and Europe. The concern, scientists say, is that decreased amounts of ozone may expose people and wildlife to higher levels of ultraviolet radiation, potentially leading to more cases of skin cancer, eye cataracts and immune system disorders.

Scientists know that chlorine-containing chlorofluorocarbons, or CFCs — commonly found in refrigerants and foam insulation — are long lived and will continue to destroy ozone at their current levels in the upper atmosphere, and they are working to curtail their use. Several industrial countries, including the United States, have proposed to eliminate all CFCs before the year 2000.

Two proposed solutions are developing better maintenance methods for CFC-dependent machinery and also devising materials totally unlike CFCs to replace them. Those techniques, however, will cut CFC use by only 70 percent, and the remaining percentage will require CFC alternatives closely resembling the guilty party and its unique thermal properties.

Unfortunately, the alternatives may be just as destructive to the ozone layer if they're not properly designed, cautions chemist Joseph S. Francisco, who is studying the chemistry of CFC alternatives in ongoing work at the Pittsburgh Supercomputing Center: "Everyone is searching for suitable alternatives, and we're working on guidelines for selecting alternatives that would minimize the number of byproducts that could destroy ozone and could have environmental consequences. There are 55 choices for the compounds, and we need to look at their toxicity, environmental impact and other factors. If we give careful thought at the beginning, we can minimize the production cost for alternatives."

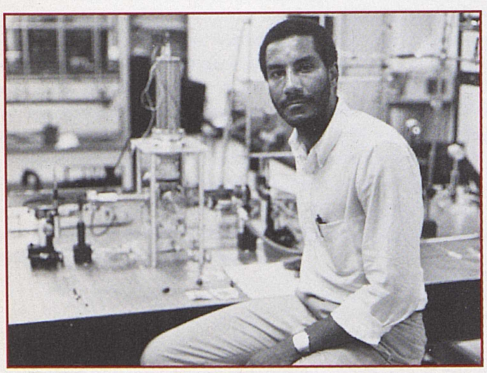

Joseph S. Francisco

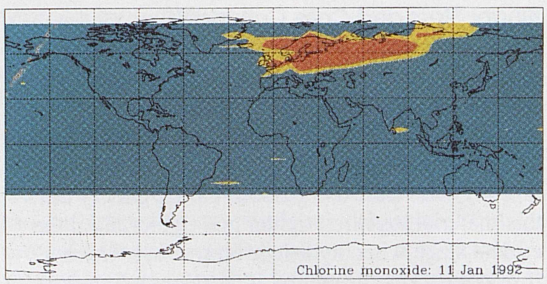

The red patch shows record-high concentrations of chlorine monoxide (ClO) in the northern hemisphere, as seen by NASA's Upper Atmosphere Research Satellite in January 1992. Chlorine monoxide, a byproduct of CFC decomposition, is the chief agent in destroying atmospheric ozone.

A Bridge to CFC Alternatives

Since 1974, researchers have known that CFCs are unreactive until 15 to 50 years after their release, when they rise to the stratosphere, about 12 to 30 miles above the ground. There, ultraviolet radiation bombards them and breaks them into chlorine atoms and chemical fragments. The chlorine destroys ozone — a ring-shaped molecule comprised of three oxygen atoms, and until now, scientists were concerned only with chlorine's role — each atom can destroy up to 100,000 ozone molecules. But Francisco wants to know what happens to the fragments. What is the chemical degradation from the fragments and what role do they play in the stratosphere's chemistry? "This has provided the framework for looking at the alternatives, because there is a bridge between the chemistries for them and the CFCs." In both cases, says Francisco, the reactants and products from the fragments are similar.

Until recently, researchers had not studied the fragments because of experimental difficulties. "There weren't accurate diagnostic tools to study these fragments," says Francisco, "and some of the resulting species from the fragments hadn't been detected in the laboratory." While laboratory researchers can study the starting materials and products of a chemical reaction, they can't explore the in-between stages.

To bypass this problem, Francisco is taking a quantum chemistry approach, which means analyzing the minute details of how the fragments are formed and eventually break down in the atmosphere. To do this, he's using the GAUSSIAN package of computer programs at the Pittsburgh Supercomputing Center — a package that's ideally suited for the complexity of Francisco's research and that's run quickly on

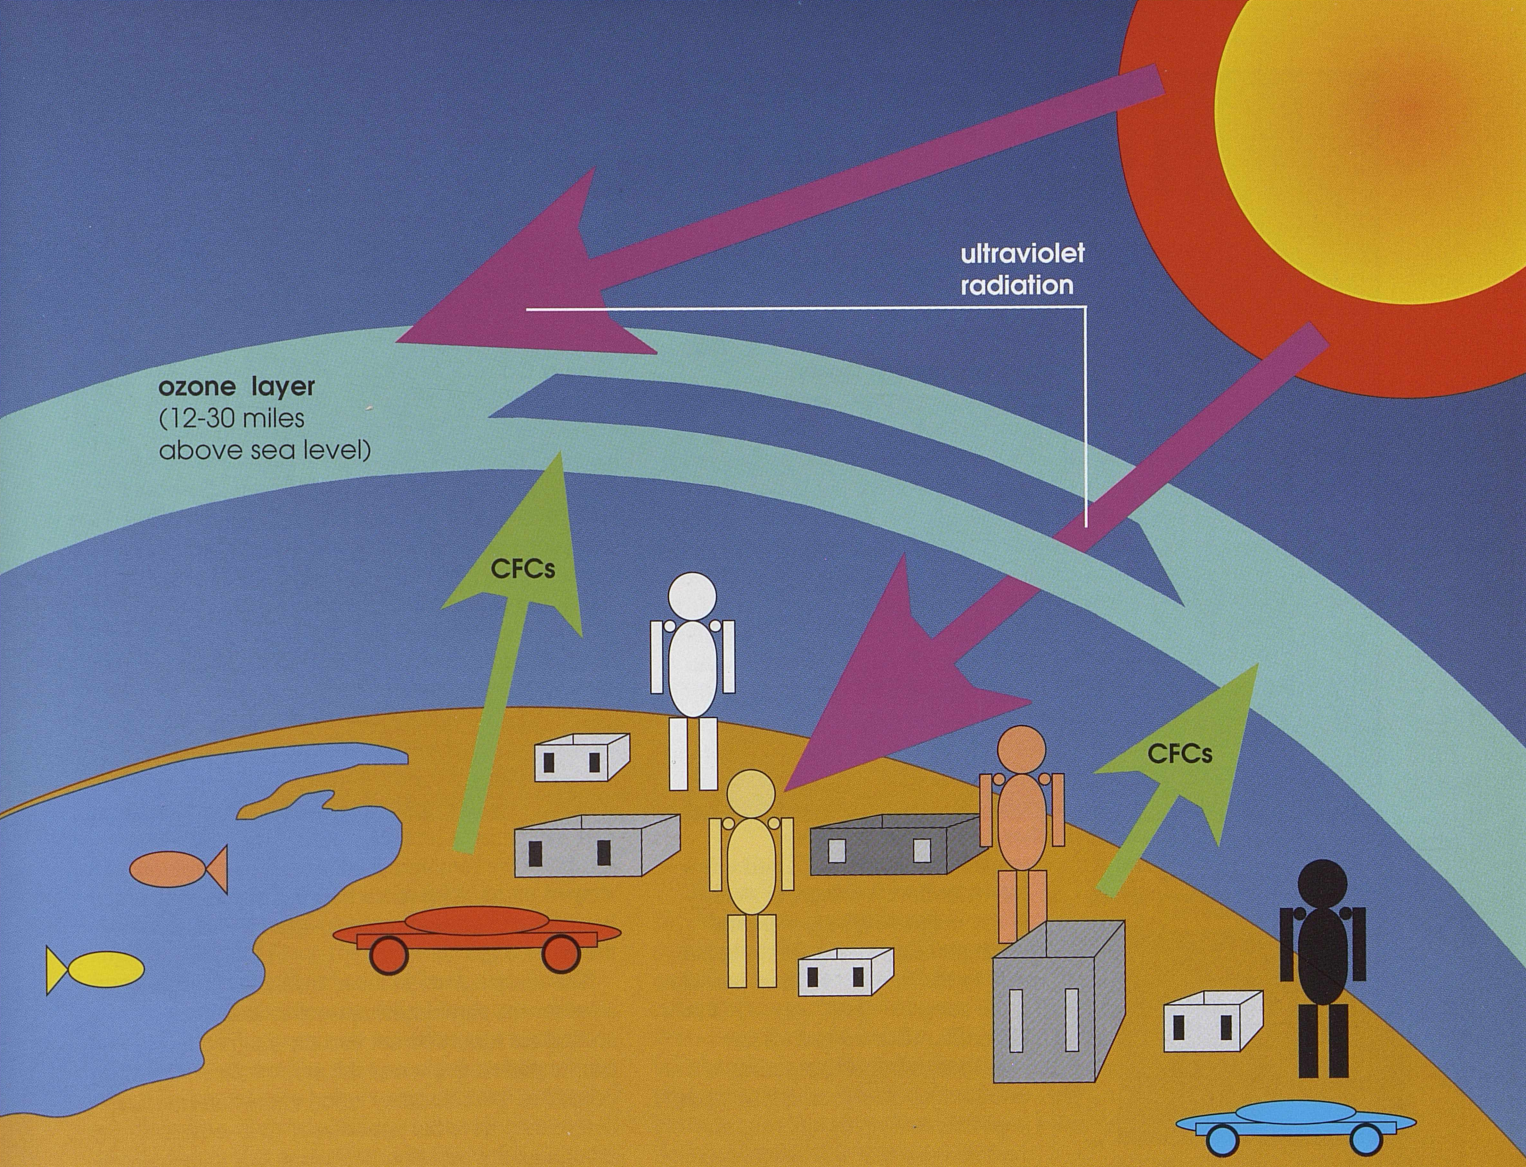

supercomputers. "With GAUSSIAN, we're trying to make some predictions," Francisco says. "What is the oxidation process for the fragments, such as how do they react and break down with oxygen and also ozone? What new species arise from these reactions and what is their impact on the chemistry in the atmosphere?"

Computational Experiments & the Laboratory

So far, Francisco has found several new species with his research at Pittsburgh, and he's trying to assess whether they're key players in the stratosphere. He has used the GAUSSIAN program to look at their structure and composition and has reproduced them in the laboratory. This, he says, will allow researchers to develop new tools to probe for the species in the laboratory and atmosphere. The challenge, however, is daunting because the species' lifetimes are fractions of seconds. "We're doing computational chemistry in conjunction with experiments," says Francisco. "Having that dual approach is what's needed to develop safe alternatives in a short time."

In future work, Francisco will apply the computational resources at Pittsburgh to detailed study of two particular alternatives — hydrochlorofluorocarbons (HCFCs) and hydrofluorocarbons (HFCs). Researchers are considering them because of their CFC-like qualities and also their short lifetimes. Unlike conventional CFCs, they will react with an OH radical in the troposphere, the lowest layer of the atmosphere, and form water and another chemical fragment.

"This in principle minimizes the ozone loss resulting from the original HFC or HCFC," says Francisco. "But will the fragments be hazardous to the stratosphere even though they appear to break down in the troposphere? Will their byproducts, such as chlorine-containing products, be stable and not harm the stratosphere? The goal is to make suggestions about the placement of halogens (chlorine and fluorine) on the alternatives to minimize the breakdown of deleterious byproducts in the atmosphere." (SE)

References:

Joseph S. Francisco, Avery N. Goldstein, Zhuangjie Li, Yao Zhao & Ian H. Williams, "Theoretical Investigation of Chlorofluorocarbon Degradation Process: Structures and Energetics of XC (O) Ox, Intermediates (X = F, Cl)," *Journal of Physical Chemistry* **94**, 4791 (1990).

L.E. Manzer, "The CFC-Ozone Issue: Progress on the Development of Alternatives to CFCs," *Science* **249**, 31 (1990).

This research is supported by the National Science Foundation, A. P. Sloan Foundation and the Camille and Henry Dreyfus Foundation.

Nature's Light Show

Precipitation of Fast Ion Beams from the Plasma Sheet Boundary Layer
Maha Ashour-Abdalla, University of California at Los Angeles

What Causes the Northern Lights?

They shoot to the ground as streams of light, but their energy source has long puzzled scientists. Now, a UCLA physicist has shown how the aurora borealis in the northern hemisphere and its mate down under — the aurora australis — form.

For years, everyone has agreed that the auroras occur when charged particles from the solar wind — electrons and protons injected into space from the sun — and other particles in Earth's magnetosphere (see illustration) become energized and spiral inward along the planet's magnetic field lines toward its upper atmosphere. There, the charged particles excite neutral gas molecules, such as nitrogen and oxygen, and because the gases crave a low-energy state, they emit light — thus, the northern and southern lights.

But what energizes the particles? Until a few years ago, scientists believed waves, or energy disturbances, in the magnetosphere were the source. "They thought that wave-particle interactions would result in the precipitation of particles and eventually, the auroral lights," says UCLA's Maha Ashour-Abdalla, who worked on her findings at the Pittsburgh Supercomputing Center. "However, as the measurement of waves became more sophisticated on spacecraft, we realized that the waves were too weak for the precipitation."

So researchers needed a way to explain how particles from the magnetosphere streamed toward Earth's upper atmosphere. The solution was their inability to hold sharp curves at high speed, and Ashour-Abdalla used supercomputing to show that this "scattering" hypothesis, as it is called, is correct.

Maha Ashour-Abdalla

Losing Control Around Curves

Extending from tens of thousands of miles to more than a million miles from Earth, there's a region known as the plasma mantle, where solar wind particles enter the Earth's magnetosphere. The plasma mantle is part of the magnetosphere's elongated section, the magnetotail, which extends downstream from the sun on the Earth's "nightside" and is completely unlike its pudgy counterpart — the "dayside" magnetosphere.

The particles, travelling at speeds exceeding 250 miles per second, usually spiral along the tail's magnetic field lines, which begin at symmetrical points along the Earth. While some lines continue for an infinite distance — for instance, those north of the northern lights region and their mirror images south of the southern lights — the remaining ones exist as elliptically-shaped paths, with more circular ones sandwiched between the auroral regions — an area known as the inner edge.

The sharp curves of the ellipsoids create problems for the particles because they lose control around the turns and fall into the magnetotail's central region, known as the current sheet. On other occasions, particles may spiral to the current sheet before they have the opportunity to lose control around a turn; those particles may fall to the central region after the magnetosphere's cross-tail electric field — created when the solar wind rushes by the magnetotail — pushes inward on them.

Plunging toward Earth

In the current sheet, particles begin their plunge toward Earth's upper atmosphere, which may happen several ways. For instance, in the illustration, the two particles behave quite differently. The bottom particle strikes near the x-line, where the magnetic field is weakest and the electric field is strongest, and as a result, the particle gains energy and accelerates toward the Earth, causing the discrete, and more intense, auroral region. This particle and others like it comprise 10 percent of the particles flowing from the plasma mantle and travel to Earth through a region known as the plasma sheet boundary layer, or PSBL, which includes the outermost, elliptically-shaped magnetic field lines.

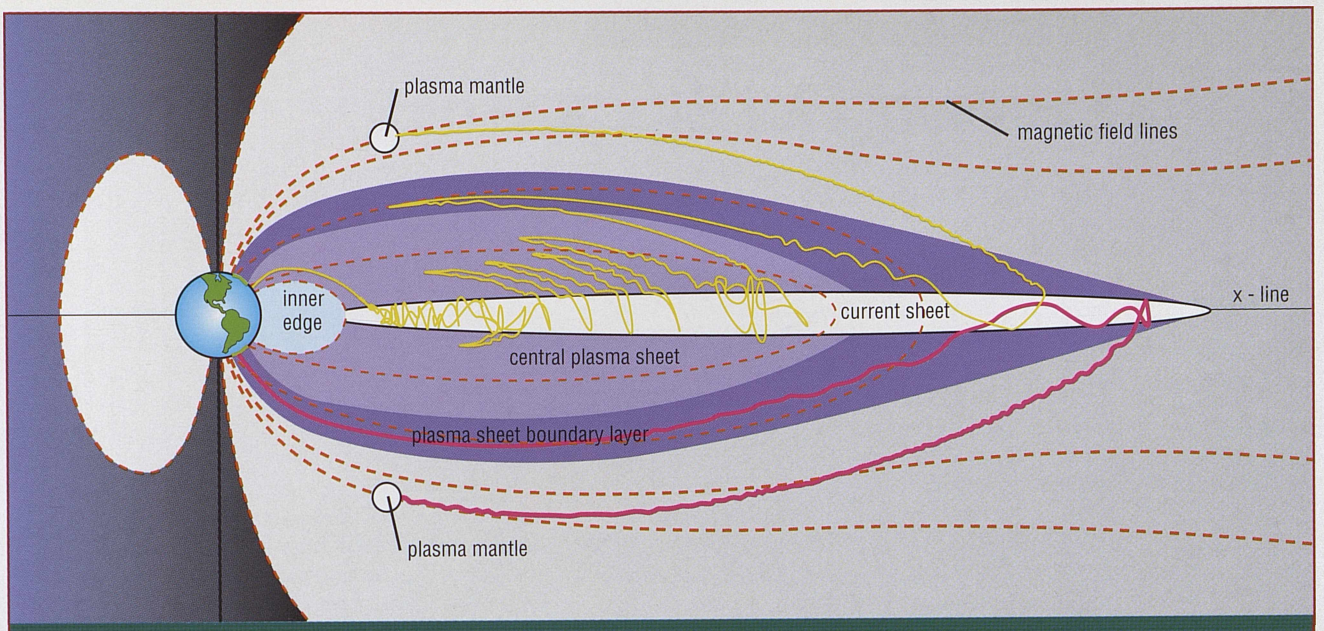

This illustration depicts part of Earth's magnetosphere, in effect a tear-shaped bubble millions of miles long, formed where the solar wind flows around Earth's magnetic field. The long tail (to the right of Earth in the illustration) is known as the magnetotail. Some particles from the solar wind mix into the magnetosphere, causing the auroras. Green areas in the thin layer of atmosphere around Earth represent the discrete auroral regions, where the color displays are most intense.

Meanwhile, the top particle may strike the current sheet with less energy or perhaps at a different angle and cannot reach the upper atmosphere on the first attempt. Instead, it bounces back to the current sheet, where it ricochets back and forth — similar to a ball stuck between two bumpers in a pinball machine — until it hits the upper atmosphere, where its low energy helps cause a diffuse auroral region, with a hazy appearance in the sky. On its way to Earth, this particle travels through a region known as the central plasma sheet, or CPS, which includes the innermost, elliptically-shaped magnetic field lines.

Simulating the Auroras

Using the CRAY Y-MP at Pittsburgh, Ashour-Abdalla simulated about 100,000 plasma mantle particles entering Earth's atmosphere. "Since we knew the average number of particles that enter the mantle each second, we took 100,000 particles and scaled it so each of the 100,000 represents a larger number, like millions and millions of particles," she says. For her simulations, she and her colleagues launched particles in groups ranging from 3,000 to 10,000.

The results supported the scattering hypothesis and also explained experimental observations. They even explained why satellites had detected echoes of energy bursts at certain latitudes in the auroral region. The simulations showed that some high-energy particles in the PSBL struck the Earth's upper atmosphere and then bounced back to that region — sometimes for several round-trip voyages, and for each voyage, they gained more energy.

Ashour-Abdalla's findings may provide more information about the ongoing mysteries surrounding the auroras, whose spectacular displays disrupt radio and satellite communication, defense systems and power transmission. The eventual goal is to predict when these displays will occur and with how much strength. (SE)

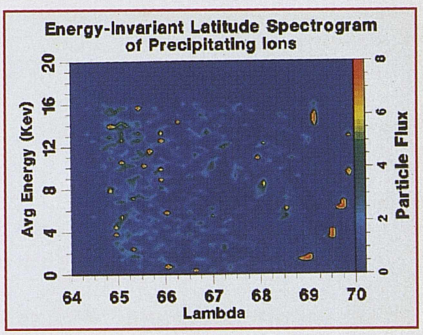

Results from the simulation at 1,000 kilometers above the Earth's surface show the echo scatterings among PSBL particles between 69 and 70 degrees latitude (lambda) and the diffuse, or spread out, energy readings of particles coming from the central plasma sheet, or CPS, in the lower latitudes, which cause the diffuse auroral regions.

References:

M. Ashour-Abdalla, L. Zelenyi, J.M. Bosqued and R.M. Kovrazhkin, "Precipitation of Fast Ion Beams from the Plasma Sheet Boundary Layer,"*Geophysical Research Letters* **19**, 617 (1992).

This research is supported by NASA.

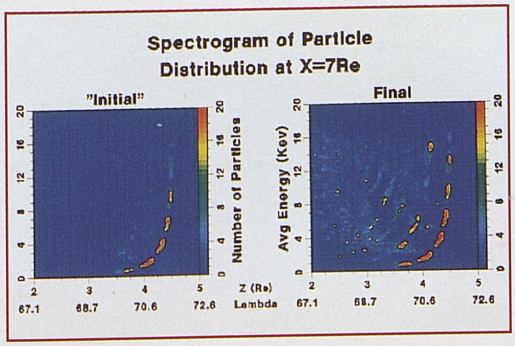

Results from Ashour-Abdalla's simulations of particles from the PSBL crossing a plane perpendicular to the Earth-sun line at a distance of seven Earth radii (7Re). Based on the distribution of particles on this plane, Ashour-Abdalla's model simulates particles entering the Earth's atmosphere. The left graph represents particles that reach the atmosphere on initial entry, with average energy in thousands of electron volts (Kev) shown on the vertical axis. These particles help create the discrete aurora — between 69 and 71 degrees latitude (lambda). The right graph shows the echoes that occur when some of the particles bounce back and forth between the Earth's upper atmosphere and the PSBL. For each voyage, the particles gain energy.

Oil and Gas Recovery

Looking for Black Gold

Using a Massively Parallel Computer for Reservoir Simulation
Ernest Chung, Chevron Oil Field Research Co.

Billions More Barrels

Oil recovery. It's similar to sticking tubes in a brick soaked with viscous goo and hoping the stash will ooze upward — not a simple task. Typically, oil is found in the minute pores of sandstone and limestone, whose volume — believe it or not — is up to 35 percent liquid. Of that amount, oil comprises about 50 to 70 percent and, on the average, only 30 percent of it is recoverable.

To maximize yields, oil companies must pinpoint the elusive stuff, and for that task, supercomputer-generated simulations have proven their usefulness. "Simulations help you decide where to drill, such as determining how the placement of recovery wells affects overall fluid flow in the reservoir, so that we will recover the most amount of oil at the lowest cost," says Ernest Chung of the Chevron Oil Field Research Co., in La Habra, Calif.

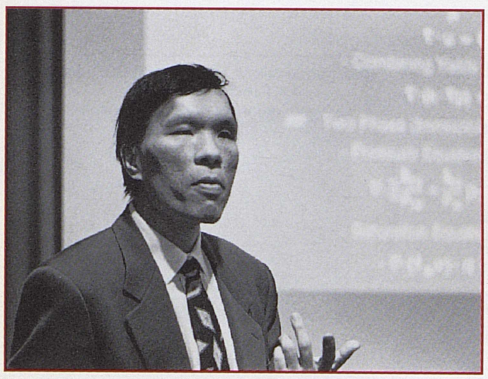

Ernest Chung

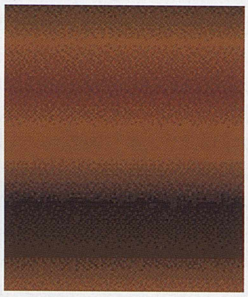

Today, in the United States, oil companies use simulations primarily to target existing reservoirs, because looking for oil is expensive and environmental standards have become more stringent. If companies increased productivity by only 1 percent at all known U.S. oil fields, they would produce more oil than the largest domestic oil cache found in the last 20 years — Prudhoe Bay in Alaska, which has 10 billion barrels.

High Resolution Modeling

Chevron currently is using an on-site CRAY supercomputer for its reservoir simulations, but at the Pittsburgh Supercomputing Center, the company is testing computer programs on the latest generation of massively parallel computers — the Connection Machine CM-2 and soon the CM-5 — to determine whether those computers also are practical for the simulations. The CMs, Chung says, potentially can obtain more detailed and quicker results than the CRAY, which has fewer processors; the CRAY's eight processors need to cycle repeatedly through such variables as fluid flow and rock location at hundreds of thousands of computer grid sites representing an oil reservoir. Meanwhile, a massively parallel machine would assign many sites to each of its thousands of processors and run the calculations simultaneously.

The trick, however, is learning to program the relatively new computers so they produce the best simulations. "It requires communication among grid points, such as how a pressure change or fluid flow at one point may affect the surrounding points. We have more than 10 years of experience with machines like the CRAY," says Chung, "but not many people have done simulations on massively parallel machines."

Benefits of NSF Supercomputing

At Pittsburgh, so far, test simulations on the CM-2 have compared favorably with actual ones already produced on the CRAY, Chung says, but the CRAY still is more efficient and can make models of higher resolution. "What we're trying to accomplish is to develop optimal programs on the CM-2 and demonstrate its use so it may become an in-house option at some point to produce simulations," he says.

Soon, Chevron will turn to the CM-5, which not only will be quicker than the CM-2, but also will provide greater modeling flexibility. The CM-5 architecture, unlike its predecessor, allows multiple processors to simultaneously execute different instructions — flexibility that fits better with the geological variability of an oil reservoir, which can include vastly different fluid characteristics in different parts of the same reservoir.

Besides monitoring oil recovery in a reservoir, Chevron is using simulations to determine whether secondary and tertiary recovery techniques, which are extremely costly, would be successful. The most common approaches are injecting high-pressured water, steam or carbon dioxide into the rock formations, where the injected fluids displace the oil.

Chung says Chevron gains a number of benefits from using the supercomputing resources at Pittsburgh: "We have proved to ourselves that using facilities like the Pittsburgh Supercomputing Center to experiment with the latest computer hardware is very cost effective. It is a great way to try the best computers without committing ourselves to a single direction." (SE)

References:

J.M. Rutledge, D.R. Jones, W.H. Chen and E.Y. Chung, "The Use of a Massively Parallel SIMD Computer for Reservoir Simulation," *Proc. SPE Symposium on Reservoir Simulation* (1991).

"We have proved to ourselves that using facilities like the Pittsburgh Supercomputing Center to experiment with the latest computer hardware is very cost effective. It is a great way to try the best computers without committing ourselves to a single direction."

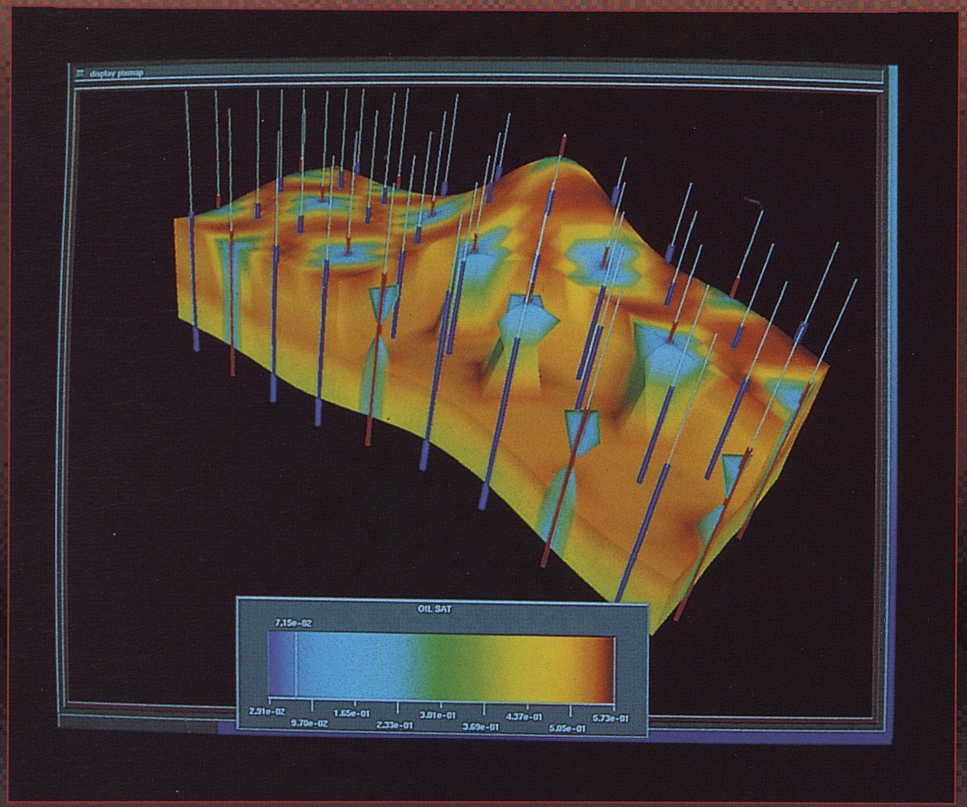

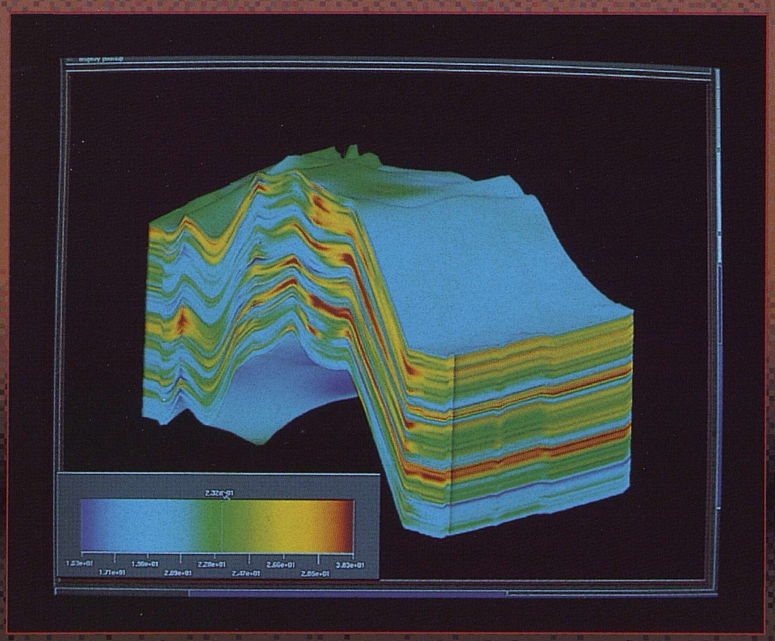

For each reservoir, Chevron produces numerous simulations, each depicting a different geologic condition. The graphic on the left shows oil saturation levels, with color representing the range of saturation — blue indicating that 3 percent of the sandstone's pores contain oil, red indicating 60 percent oil. This reservoir contains two types of wells; production wells are orange-red, and injector wells are surrounded with a blue haze, which represents the injected fluid.

In the graphic below, from a simulation that contains 250,000 grid points, Chevron is looking at the percentage of empty space, or porosity, in the sandstone, with blue representing 15 percent and red, 30 percent. Porosity strongly affects other rock properties, such as permeability, which measures a fluid's ease in moving inside rock. This model, which simulates an actual reservoir containing a few billion barrels, is one in a series that Chevron has produced, each with a different number of gridpoints and different ways to represent the properties at every grid point. "These models are being run to determine how much geologic detail is needed to yield accurate results," Chung says. "Also, we're trying to develop techniques to approximate the geologic detail so we could get almost the same accuracy but at a much lower cost."

Teaching a Computer to Read

A Parallel-Distributed Processing Model of Word Recognition and Naming
James McClelland, Carnegie Mellon University

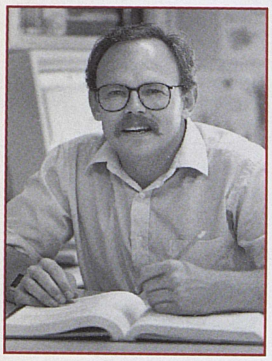

James McClelland

Spelling, Meaning and Context

If you're walking in the wilderness, you might see a sign plastered against a tree: "Read this! Fishing for bass is prohibited." And after a fiction reading, you might see a review saying: "He read in a loud, bass voice." In both cases, "read" and "bass" are pronounced differently and with "bass," the meaning changes. "In English, we think there are some rules about how spellings relate to sounds," says James McClelland, a cognitive psychologist at Carnegie Mellon University, "but there are so many exceptions. The same sequence of letters is often pronounced differently, and it all depends on the context."

The research challenge, as proposed by McClelland, is discovering not only how the brain differentiates between words, but also how we read. To accomplish that, he's using a supercomputer at the Pittsburgh Supercomputing Center, thus becoming part of an expanding scientific field — neural network modeling — in which computers model the activity of the plethora of nerve cells, or neurons, in the brain. "We don't notice that we use context and meaning when we read, but we rely on these sources of information all the time," says McClelland. "However, people haven't been able to build machines that perform those functions. My hope is that we could — maybe before I retire — have a machine that could really read, incorporating sound, spelling, meaning and context."

Among the significant implications of this research, McClelland says, is that it can help determine the best methods to teach children how to read: "Previously, researchers thought you had to stress either phonics or meaning and context. But recent research, based partly on neural network models, stresses the importance of teaching kids to use sound and also meaning and context together."

Neurons and Chemical Squirts

In the human brain, there are about 10 billion neurons in distinct groups, all of which transmit and receive information to accomplish various tasks — such as recognizing a word's spelling and pronunciation and determining its meaning. Each neuron has a cell body and an axon, from which an electrical signal is sent when the neuron is active. At the connections between neurons, these signals are turned into chemical squirts that tend to activate or inhibit other neurons.

"The brain learns to perform a certain job by adjusting the effects of these signals," says McClelland. "So if we want people to learn that a particular spelling goes with a particular pronunciation, we must get them to produce the neural activity patterns representing each one during the learning process and then let the brain adjust the connections between the activated neurons."

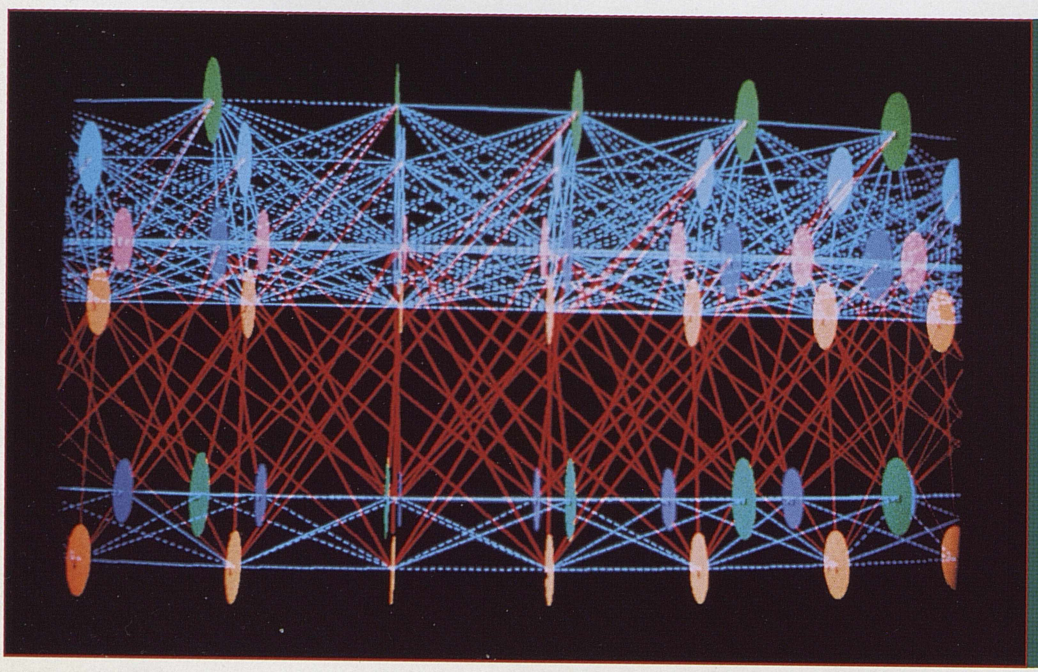

A simulated neural network used for speech processing. Units are indicated by circles, and connections are indicated by lines. (Photo courtesy of Jeff Elman, Department of Cognitive Science, University of California at San Diego.)

A Supercomputer that Reads

McClelland's computer modeling aims to mimic these neural processes of the brain. He presents spelling and sound pairs to a simulated neural network and attempts to teach the network to produce the right sound when it sees a particular spelling. McClelland has reproduced 1,000 to 2,000 neuron-like units on Pittsburgh's CRAY Y-MP, and he enters the spellings on a keyboard as a sequence of letters, which are translated into an internal representation. One word, for instance, activates a set of units — about 20 units for each letter, indicating its identity and placement — and each unit is updated many times until the pattern of activity stabilizes. Next, the neural network signals other units that represent sounds in the word's pronunciation, indicating such characteristics as long and short vowel sounds.

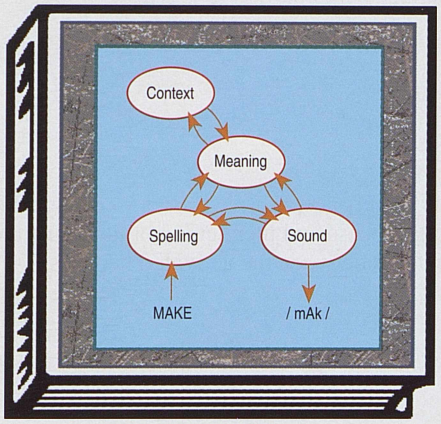

The computer tallies all the inputs that a neuron-like unit receives from other units and then updates the unit's activation. "Adding up all the inputs is what takes most of the computer time," says McClelland, "because each unit receives input from as many as 400 other units. Thankfully, supercomputers are optimized for this kind of computation."

After the network takes a stab at the answer, the researchers provide the correct one, like a teacher might prompt a child learning to read. Then, the necessary adjustments are made, until the network gets it right. "We've taught our network to read a vocabulary of 3,000 words," says McClelland, "and it has taken 20 to 100 presentations of each word to do that." Meanwhile, today's workstations would be able to simulate only 100 words. The goal, he says, is a full vocabulary of 40,000 words using more advanced supercomputers in the future.

Looking Ahead

To build a complete vocabulary, it will be necessary to incorporate context and meaning; otherwise, the word "bass," for instance, always would yield just one of its pronunciations. The task is difficult because there is no established system for representing either type of information. Although McClelland has not attempted either approach on supercomputers, he plans to allow the neural network to discover for itself how to represent context and meaning by teaching it to answer questions after it reads sentences describing scenes and events.

When McClelland's group develops representations for meanings and context— a task that may take a decade — they hope to match spellings with meanings, meanings with sounds and context with meanings, as well as sounds with spellings. The trick is getting the four subsystems to communicate. "In my view," says McClelland, "all of the parts of the system are in continuous, interactive communication." (SE)

References:

Mark S. Seidenberg and James L. McClelland, "A Distributed, Developmental Model of Word Recognition and Naming," *Psychological Review* **96**, 523 (1989).

J.L. McClelland, D.E. Rumelhart and G.F. Hinton, "The Appeal of Parallel Distributed Processing" in *Parallel Distributed Processing: Explorations in the Microstructure of Cognition* (Cambridge, Mass.: MIT Press, 1986).

This research is supported by the National Science Foundation and the National Institute of Mental Health.

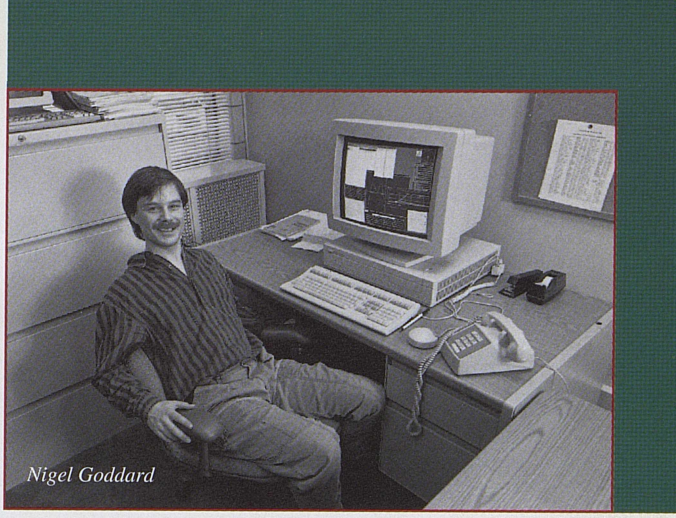
Nigel Goddard

Supercomputing Support for Neural Modeling

"We want to encourage, facilitate and organize the use of supercomputers," says Nigel Goddard, "by researchers for whom supercomputing isn't endemic to the culture." Goddard leads the Pittsburgh Supercomputing Center's new program to support neural modeling and neural network research in all of its varied forms. To date, he has focused on surveying researchers in the field to see what kinds of support can be most helpful, and he has also worked on making some useful software more readily available.

"One goal," says Goddard, "is to have simulation packages that look the same to users whether they're using a SUN workstation or a CRAY." Several packages, both for neural network simulation and for biologically-based neuronal modeling, are now running on the center's Y-MP; and a Thinking Machines package — that includes "back propagation" routines for teaching the network to learn from its errors — is running on the center's CM-2 and CM-5. "The new massively parallel architectures are very well suited to these kinds of networks."

Goddard's background spans different but related areas of neural modeling and includes work with several researchers widely known in the field. On the way to writing his 1991 computer science dissertation (with Jerry Feldman at the University of Rochester), Goddard worked a year in Christof Koch's neuroscience lab at Caltech. His dissertation, completed while a visiting scholar at Carnegie Mellon, is "a cognitive model of high-level parts of the motion pathway in the visual system." (MS)

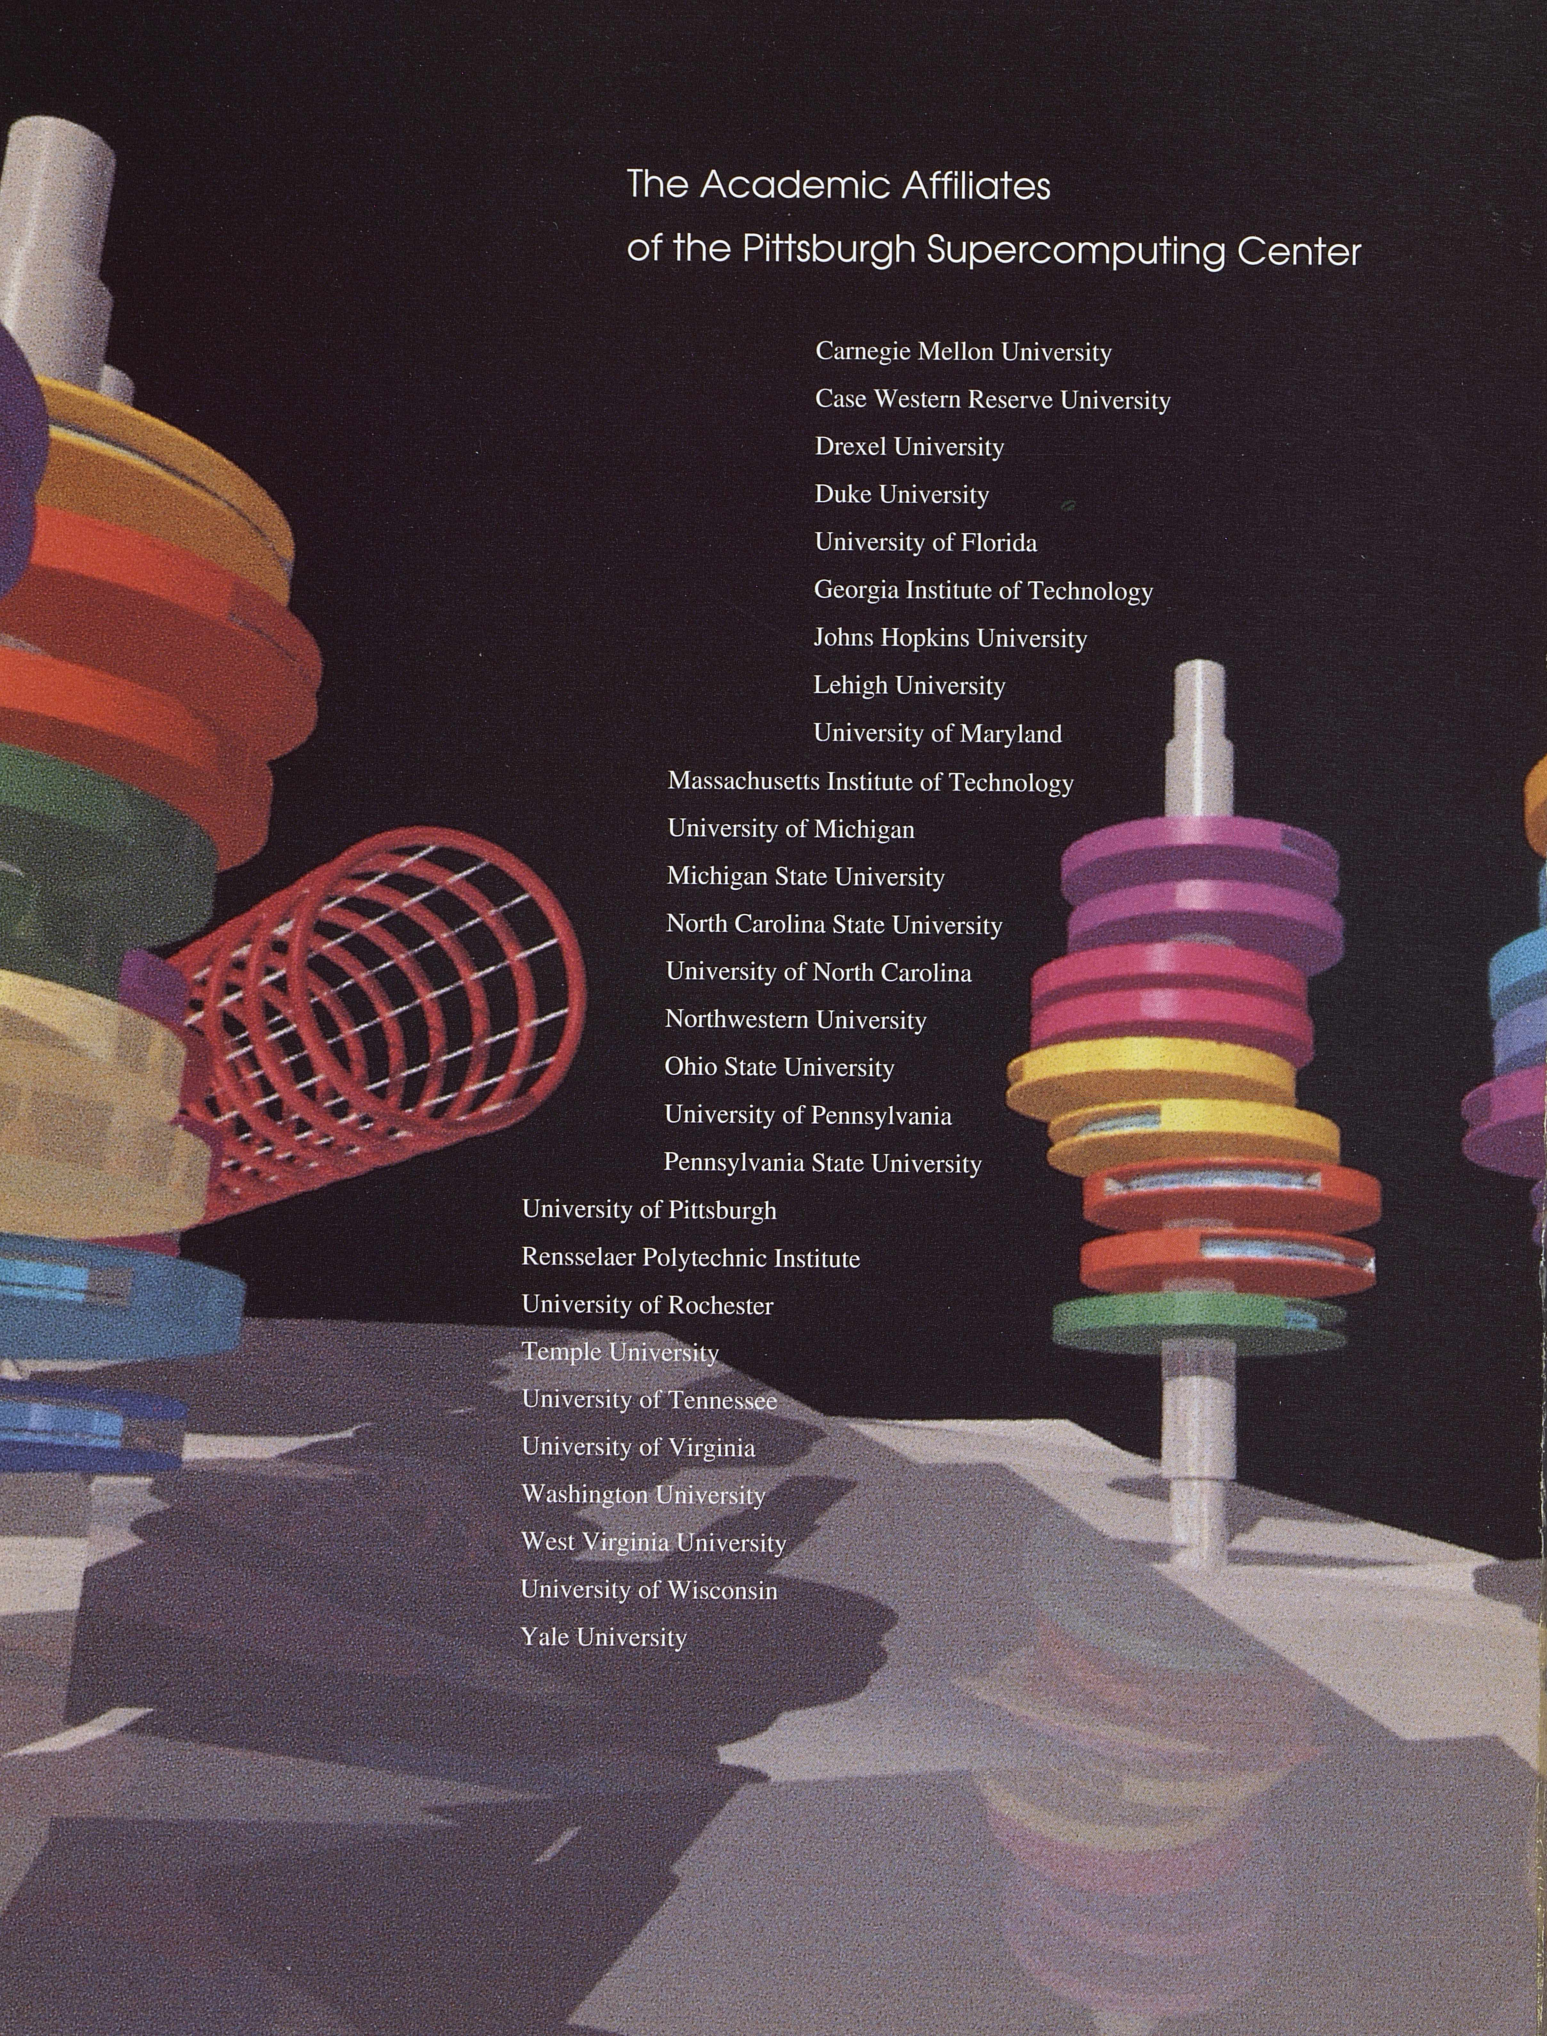

The Academic Affiliates of the Pittsburgh Supercomputing Center

Carnegie Mellon University
Case Western Reserve University
Drexel University
Duke University
University of Florida
Georgia Institute of Technology
Johns Hopkins University
Lehigh University
University of Maryland
Massachusetts Institute of Technology
University of Michigan
Michigan State University
North Carolina State University
University of North Carolina
Northwestern University
Ohio State University
University of Pennsylvania
Pennsylvania State University
University of Pittsburgh
Rensselaer Polytechnic Institute
University of Rochester
Temple University
University of Tennessee
University of Virginia
Washington University
West Virginia University
University of Wisconsin
Yale University

Pittsburgh Supercomputing Center
Mellon Institute Building
4400 Fifth Avenue
Pittsburgh, PA 15213

Nonprofit Org.
U.S. Postage
PAID
Pittsburgh, PA
Permit No. 251