PROJECTS IN SCIENTIFIC COMPUTING

1991

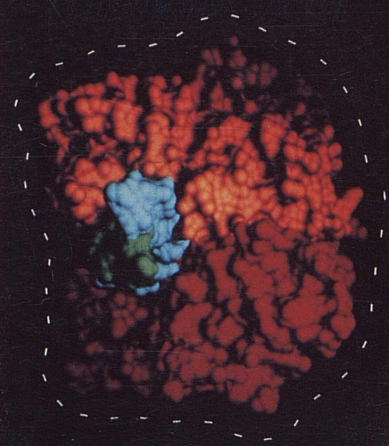

PITTSBURGH
SUPERCOMPUTING
CENTER

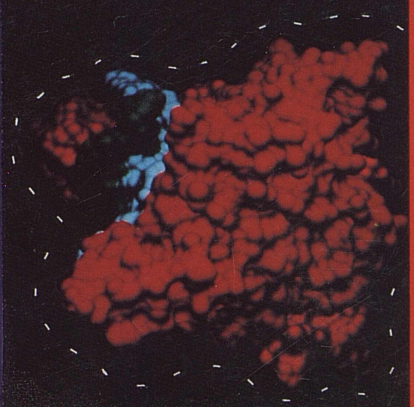

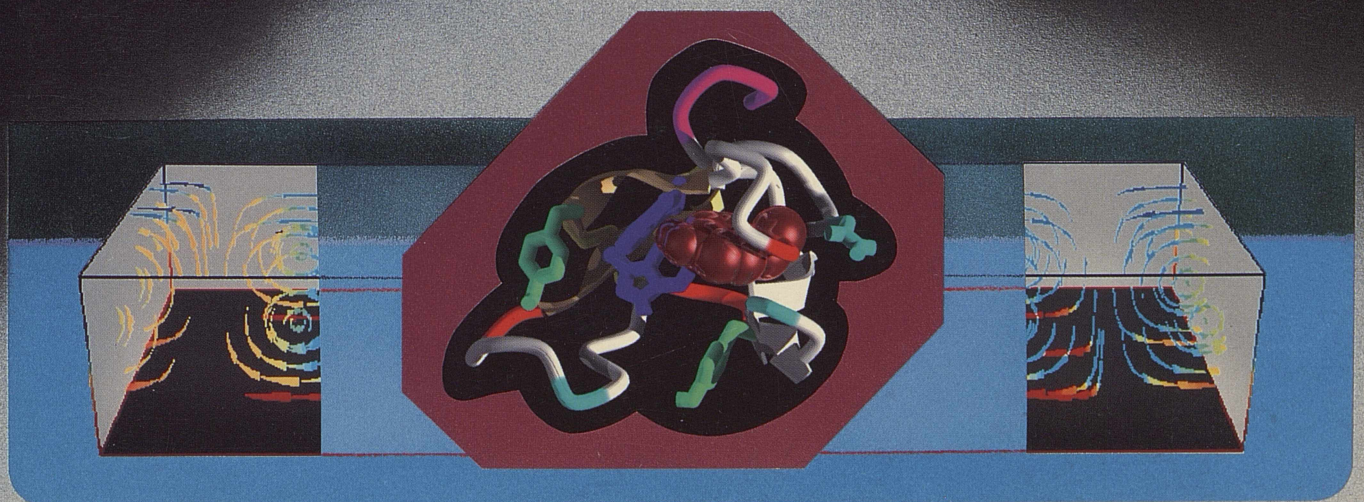

Projects in Scientific Computing

1991

The Pittsburgh Supercomputing Center is a joint project of Carnegie Mellon University and the University of Pittsburgh together with Westinghouse Electric Corporation. It has received 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, Incorporated
Digital Equipment Corporation
Harris Corporation
International Business Machines Corporation
Network Systems Corporation
Stardent Computer Incorporated
Storage Technology Corporation
Sun Microsystems, Incorporated
Thinking Machines Corporation
Ultranetwork Technologies
Bell of Pennsylvania

Published August 1991 by the Pittsburgh Supercomputing Center, 4400 Fifth Avenue, Pittsburgh, PA 15213. (412) 268-4960.

Credits

Editor: *Michael Schneider, Pittsburgh Supercomputing Center.* Design & Illustrations: *Steve Palmer, Pittsburgh.* Writing: *Michael Schneider, Bob Anderson, Heather Lum, Sam Peters.* Editorial Assistant: *Megan Irvine.* Photography: *Michael Haritan (7,12,14,16), Jeff Macklin (5,10), Larry Rippel (4,5), Jim Schafer (1,4,5,6,8,31,38).* Color Separation & Typesetting: *Digicon Graphics, Inc., Chicago.* Printing: *Westinghouse Printing.* Production Coordination: *Michael Hoeg, Westinghouse Printing.*

Page layout & design were done by Steve Palmer at the Pittsburgh Supercomputing Center on a Macintosh II computer using PageMaker 4.0 desktop publishing software. These digital layouts included spot color, encapsulated postscript graphics, and tag image file format (TIFF) scientific graphics produced by researchers at the Pittsburgh Supercomputing Center and translated to TIFF using the MacGPLOT graphics interpreter software developed by Phil Andrews and Grace Giras of the Pittsburgh Supercomputing Center graphics group. Thanks to Joel Welling, Anjana Kar and other members of the graphics group for their assistance.

Cover art: *Steve Palmer, Pittsburgh.*

Foreword

At five years, the NSF Supercomputing Centers Program is having major impact on science and engineering in the nation. More than 15,000 researchers from over 500 institutions from every state have used the facilities at the national Centers. More than 7,000 papers have already been published based on work at the Centers. Increasingly powerful networks link the national Centers, and they have been working together to present users across the nation with increased and unified capability. This capability is enabling the country's best scientists and engineers to attack the pressing computational problems, the Grand Challenges, whose solution is crucial to our national vitality and creativity. The success of the Centers has reinforced the perception of the national importance of High Performance Computing and Communications, which has become one of the three federal initiatives in science.

PSC has been a leader in developing this new capability, articulating the heterogeneous systems that will characterize future high performance computing systems. This year PSC was the first to install and use a high-speed link (HiPPi) between its CRAY Y-MP and its Connection Machine and to develop applications which ran across both machines. The paired system is often an order of magnitude more powerful than each component separately. PSC has been a leader in developing the Andrew File System for a high performance computing environment, in developing three-dimensional metafiles for visualization, and in implementing gigabit networks.

PSC's work in enhancing capability was recognized in two awards. Manfred Prammer from the University of Pennsylvania worked with PSC consultants to develop code which ran on PSC's Y-MP at over two gigaflops, for which he won a Cray Gigaflop Award. Work by PSC scientific specialists, systems and communications staff on genetic sequence analysis code for the distributed Cray-Connection Machine system was the leading non-European awardee in the SuParCup Competition, sponsored by the University of Manheim.

The federal High Performance Computing and Communications Program correctly emphasizes the importance of education. This year, PSC has broadened its outreach to include high schools, in addition to its usual efforts in undergraduate, graduate and postgraduate education.

The benefits of high performance computing for national competitiveness will be accelerated by closer linkage between the national Centers and the private sector. This year, DuPont and Chevron, American corporations with considerable supercomputing experience, chose to affiliate with PSC to benefit from its expertise.

The real justification for the Centers is science. In this booklet, we outline a sampling of the diverse projects that are being undertaken using the Pittsburgh Supercomputing Center facilities; they span Heaven and Earth, the practical and the esoteric, and deal with life, disease, death and even unemployment.

It is exciting to be at a Center of this activity. We look forward to even more exciting developments in the years to come.

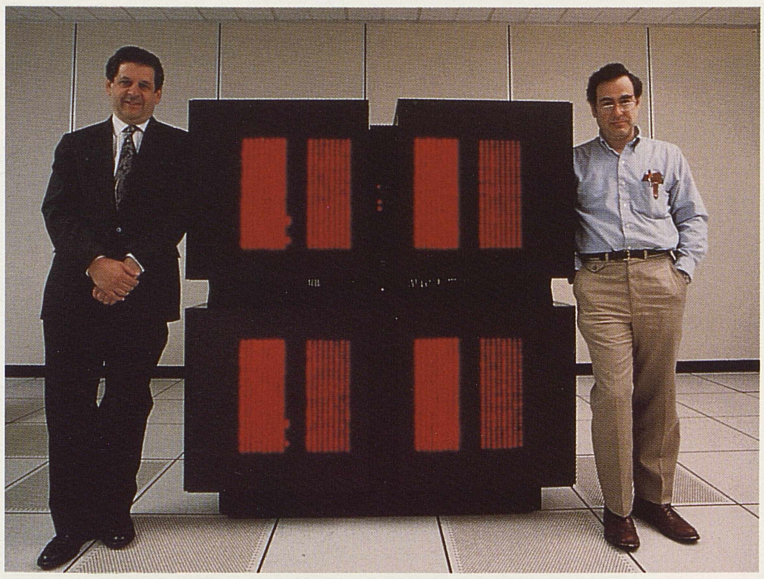

Michael Levine (right) and Ralph Roskies, Scientific Directors of the Pittsburgh Supercomputing Center, with the Center's Connection Machine, CM-2, at Westinghouse Energy Center, Monroeville, Pennsylvania.

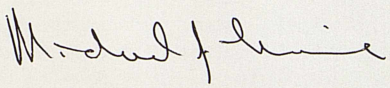

Michael J. Levine, Scientific Director

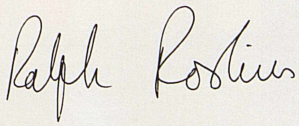

Ralph Z. Roskies, Scientific Director

PITTSBURGH SUPERCOMPUTING CENTER

Projects in Scientific Computing, 1991

CONTENTS

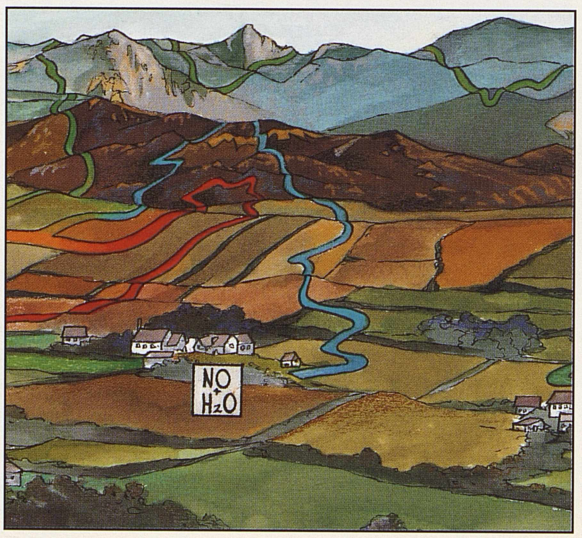

PITTSBURGH SUPERCOMPUTING CENTER

ASTRONOMICAL SCIENCES

MATHEMATICAL SCIENCE

BIOLOGICAL SCIENCE

GEOSCIENCES

SOCIAL SCIENCES

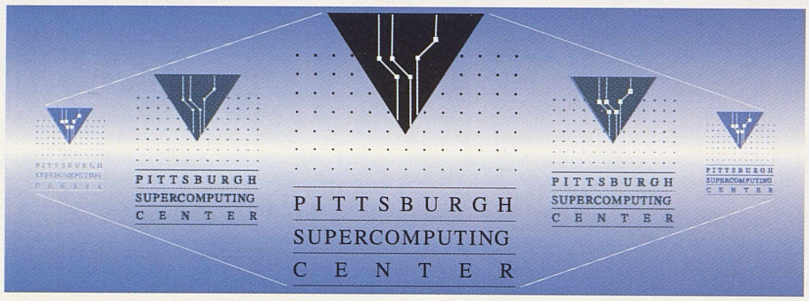

New Technology: Pittsburgh Steps Forward

This year, five years down the road from its 1986 opening, marked the emergence of the Pittsburgh Supercomputing Center as a leader in the development of new supercomputing technology. Two different research teams working on two separate projects demonstrated that the much talked-about concept of *heterogeneous computing* is more than a fancy idea. In Pittsburgh, heterogeneous computing works.

"For many of the important research efforts in this country, what we call the grand challenge problems," says Michael Levine, Scientific Codirector of the Center, "a major obstacle is that there's not enough computing available to get the answers we need."

How to soup up supercomputing? How to squeeze more performance from computers that already represent the limits of computing technology? A frequently proposed answer is heterogeneous computing — a computing environment in which computers of different capability can do what each does best while communicating with each other in an interconnected system.

Link, for instance, a top-of-the-line vector processor like the CRAY Y-MP and a massively parallel system like the Center's Connection Machine, CM-2. This can allow scientists to distribute their code between the two machines, so that inherently serial parts run on the CRAY while parts of the problem amenable to parallel solution exploit the CM-2.

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 Corporation.

To date, more than 3700 scientists and engineers at over 350 universities and research centers (red and green dots) in 48 states have used the Center's computing resources to advance their research. This work has resulted in nearly 1000 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.

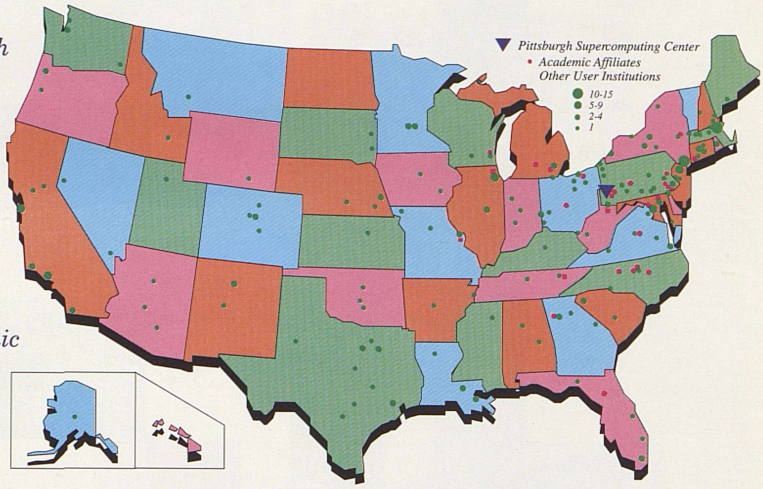

CRAY Y-MP Usage by Discipline

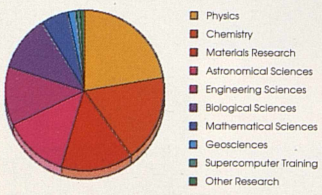

Beverly Clayton, Executive Director of the Pittsburgh Supercomputing Center. JS

Jim Kasdorf, Director of Supercomputing, Westinghouse Electric Corporation. Kasdorf and his staff maintain and operate the Center's supercomputers and associated hardware in the machine room at Westinghouse Energy Center, Monroeville, Pennsylvania. LR

Two Distributed Applications

This is what Greg McRae and Robert Clay did to reduce computing time for "the assignment problem." (See "Shared Assignment," p. 12). The two Carnegie Mellon researchers teamed with three Pittsburgh Supercomputing Center staff members to split the application between the Y-MP and CM-2. In February, this team successfully linked the two computers with a high-speed interconnection, the HiPPI (high-performance parallel interface) — the first time this much talked-about possibility has been accomplished. Distributing the application code between the two machines with data transfer by way of the HiPPI allowed test problems to be solved 10 times faster than is possible with either machine alone.

Another breakthrough in distributed computing came this year when a six-person team of Pittsburgh Supercomputing Center scientists and analysts distributed code for the dynamic programming algorithm between the Y-MP and CM-2. "This algorithm," explains Hugh Nicholas, Biomedical Scientific Specialist, who coordinated the effort, "allows rapid and rigorous comparison between particular sequences of genes and proteins and large gene and protein databases. It is the preferred procedure for comparing newly sequenced genes and proteins with previously known sequences."

Finding sequence similarities among biological molecules is one important way, for instance, of determining if diseases have a genetic origin, and computerized search techniques can save months of expensive trial-and-error experimentation. The heterogeneous computing approach developed by the Pittsburgh group takes 5 to 10 times less elapsed computing time than on the Y-MP alone. This project received international recognition as a 1991 prize-winner in the Mannheim SuParCup competition for innovative use of parallel computing, sponsored by the University of Mannheim, Germany.

LR

Dan Nydick, Research System Programmer, who works on the Andrew File System project.

Andrew File System

Another significant development project this year has been the effort to incorporate the Andrew File System (AFS) into the Center's production environment. To fully use the rich set of resources available, a Pittsburgh Supercomputing Center user must deal with several different file systems (UNICOS, VMS, CFS and Connection Machine among others). The Center has incorporated AFS into UNICOS running on the CRAY Y-MP, the first time AFS has been ported into this environment. Testing to meet the Center's rigorous production standards is currently underway.

AFS offers a number of advantages in dealing with the diversity of the Center's computing environment. First and foremost, it provides a single, uniform method to access a file regardless of where it physically resides. AFS makes all files appear to be part of the local file system of the user. By keeping copies of recently used files in a local disk cache, AFS improves supercomputer performance by reducing input/output bottlenecks and network traffic. AFS also has superior security compared to other distributed file systems.

JM

Donna Fantini (center) of the Center's Production Systems staff confers with Lori Smith (left) and Vasiliki Hartonas-Garmhausen of User Services and Yasunari Tosa, applications engineer for Thinking Machines Corporation assigned to the Pittsburgh Supercomputing Center. Fantini, Smith and Hartonas-Garmhausen have worked on bringing the Center's CM-2 into a production environment.

The Connection Machine in Production

In April 1990, the Center took a vigorous step forward in its plans for heterogeneous computing when, with support from the Defense Advanced Research Projects Agency (DARPA) and NSF, it acquired its Connection Machine, a 32,768 processor CM-2. The Center has moved rapidly to bring this massively parallel system into full production. After a six-month "friendly user" period, the CM-2 was opened up in December for research proposals from scientists and engineers nationwide.

Making the CM-2 available for network access by researchers at remote locations presented some difficult, new challenges to the Center's production staff. "The Connection Machine is typically in an open environment, like a departmental workstation" explains Janet Brown, Manager of Production Systems, "where the users know each other and communicate face-to-face to resolve priorities in resource scheduling." Brown and the Center's consulting staff have worked closely with analysts from Thinking Machines Corporation, who manufacture the CM-2, to develop software tools for an environment in which the users are, in effect, blind to each other.

Two new Thinking Machines' software products, says Brown, the network queuing system (NQS) and time-sharing, which controls access to the CM-2 for short debugging runs, have helped considerably. Both products have undergone testing on the Center's CM-2 and are now in production use.

JS

Rich Raymond, User Services Coordinator, has helped train the Center's consulting staff on the CM-2. "Programming a massively parallel machine," says Raymond, "requires a new perspective. In many cases, the speed of the code and success of the project are very sensitive to decisions made early on; choosing the right algorithm and developing the proper data structure can be critical. For this reason, our CM-2 consulting is personalized to the project, and we will make ourselves available from the beginning of the project's life cycle."

The Biomedical Initiative

In August 1990, the Center received a five-year, $6.1 million grant from NIH supporting the Center's effort to develop new capabilities in biomedical software, and to make supercomputing more widely available to biomedical scientists. This grant, which follows up on the $2.2 million grant NIH awarded the Center in 1987, will fund a series of workshops and several software development projects. Two of these projects, biological sequence analysis (see above) and molecular dynamics, use the Center's heterogeneous computing capability to improve performance.

"We've barely begun to realize the potential of what supercomputing can accomplish in biomedicine," says Scientific Codirector Ralph Roskies. "Because of the amount of molecular data involved and the processing required to analyze it, computational techniques have become important in research addressing a wide range of health problems, including heart disease, cancer and AIDS."

Since 1987, workshops at the Pittsburgh Supercomputing Center have introduced hundreds of biomedical scientists to the techniques of high-performance computing. Many of these researchers are now using these techniques at the Center to investigate such things as the structural dynamics of DNA, the relation between structure and biological function of proteins, and the fluid mechanics of blood circulation.

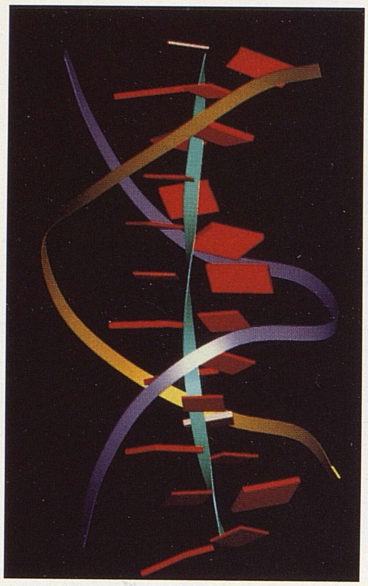

This three-dimensional representation of DNA provides an unambiguous way to visualize basepair alignment along the DNA central axis. It is derived from an algorithm developed by Biomedical Scientific Specialist David Deerfield, student intern Darrin York and Tom Darden of the National Institute of Environmental Health Services; the graphic was rendered by Pittsburgh Supercomputing Center consultant Grey Lorig.

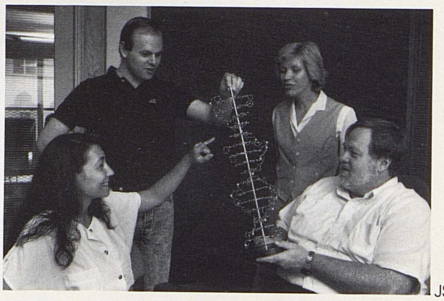

JS

Members of the Pittsburgh Supercomputing Center team that developed a heterogeneous computing approach to the dynamic programming algorithm examine a DNA model (left to right): Vasiliki (Vicki) Hartonas-Garmhausen, Alex Ropelewski, Grace Giras and Hugh Nicholas. Two members of this research team, Michael Kopko and Chris Maher, are not in the photo.

Networking: The Ties that Bind

This year the Center's communications group participated with NSF in upgrading the capability of the main network linking supercomputing centers to each other. The upgrade increases the "bandwidth" of data communications by a factor of 30, from 1544 Kbits (T1 lines) to 45,000 Kbits (T3). Pittsburgh is one of eight sites nationally (as of July 1991) connected to this new, greatly improved capability.

The intention of these upgrades, notes Gene Hastings, technical supervisor of the Center's communications group, is to implement a high degree of user connectivity. "Ubiquitous access is a key phrase in this effort. Increased connectivity is very important in advancing the grand challenge research efforts, and all these data highways run through Pittsburgh."

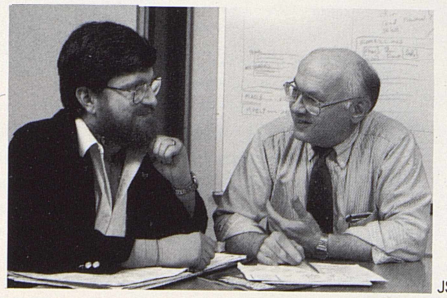

JS

Gene Hastings meets with User Services Manager Bob Stock.

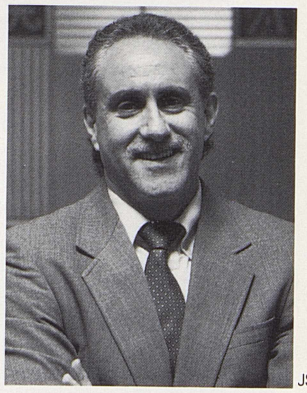

JS

Marvin Zalevsky, Assistant Director, Pittsburgh Supercomputing Center. "Our main business is assuring that users can be productive, and our solid infrastructure is the core of that effort. From the hotline consultants to our scientific support staff and our production systems and communications groups — these are talented professionals that hold things together on a day-to-day basis."

Spreading the Know-How of Supercomputing

Since the Center opened in 1986, it has conducted 139 workshops and seminars in which more than 4200 people have participated. Over the last five years, this education has become ever more central in the Center's effort to promote supercomputing. During the past year alone, under the direction of Education Coordinator Casey Porto, 42 Center-organized workshops and seminars drew more than 1200 participants.

A new offering at the Center this year was the four-day workshop on the computational chemistry software GAUSSIAN 90. This year also featured the Center's first workshop on the techniques of massively parallel supercomputing on the Connection Machine.

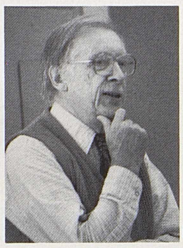
MH


MH

Professor John Pople, whose group of researchers in the chemistry department at Carnegie Mellon University developed the GAUSSIAN series of computational chemistry programs, lectures to participants in the Pittsburgh Supercomputing Center's first GAUSSIAN workshop.

Pittsburgh Supercomputing Center Workshops
(offered at least once during the past year)

Introduction to Supercomputing
Supercomputing Techniques: CRAY
Scientific Visualization
Code Optimization
Multitasking
Supercomputing Techniques: Connection Machine
Industrial Affiliates
GAUSSIAN 90
Advanced Computing Techniques for Biomedical Researchers
Biological Fluid Dynamics
Molecular Mechanics & Dynamics of Biopolymers
Nucleic Acid & Protein Sequence Analysis

Information at the Pittsburgh Supercomputing Center

Proposals for Computing Time	*Wendy Janocha* 412-268-5005
Biomedical Initiative	*Nancy Kiser* 412-268-5206
Workshops & Summer Institute	*Casey Porto* 412-268-7808
Corporate Affiliates Program	*Ira Hochman* 412-268-2776
Newsletter and Documentation	*Vivian Benton* 412-268-6355

These Pittsburgh Supercomputing Center staff members can also be reached via electronic mail.
Internet: *lastname*@psc.edu.
Bitnet: *lastname*@cpwpsca.

Outreach to High School Teachers

More than 100 high school teachers and administrators participated in the Center's March 8 workshop, "Supercomputing: An Innovative Tool for Teachers" — the kick-off event for the Center's new secondary education outreach program. "This program," says Jane Konrad, who coordinates the high school outreach effort, "is designed to excite the imagination of students and teachers by enabling them to use supercomputing in a science project."

Konrad, who joined the Center staff in December, taught high school biology for many years. Along with her work at the Center, she directs the Pittsburgh Regional Center for Science Teachers, a teaching resource center.

The program will provide ten teacher-student teams with the opportunity to use supercomputing for science projects at the high school level. Konrad believes this will foster better understanding of the role played by computing and simulation in scientific research. "We've designed this program with the aim of increasing interest in science as a career. Supercomputing can help bridge the gap between textbooks and the real-life process of 'doing science.'"


MH


MH


MH


MH

Supercomputing: An Innovative Tool for Teachers

Clockwise from top left: Grey Lorig of the Pittsburgh Supercomputing Center staff demonstrates ChemTool, an interactive software tool for visualizing chemical structure, with Jane Konrad, Secondary Education Outreach Coordinator, looking over his shoulder. Congressman William Coyne addresses the assembled group. Ralph Roskies chats with Donald Henderson, Provost, University of Pittsburgh. Beverly Clayton talks with Pennsylvania State Representative Ron Cowell.

"The Pittsburgh Supercomputing Center is an important resource for Pennsylvania," notes Clayton. "Thirty-five Pennsylvania universities and corporate research centers use our facilities, and the availability of supercomputing is one of the factors that helps promote economic development in this area. Through our high school outreach program we are now also helping to educate young people about science and the vital role computation plays in scientific research."

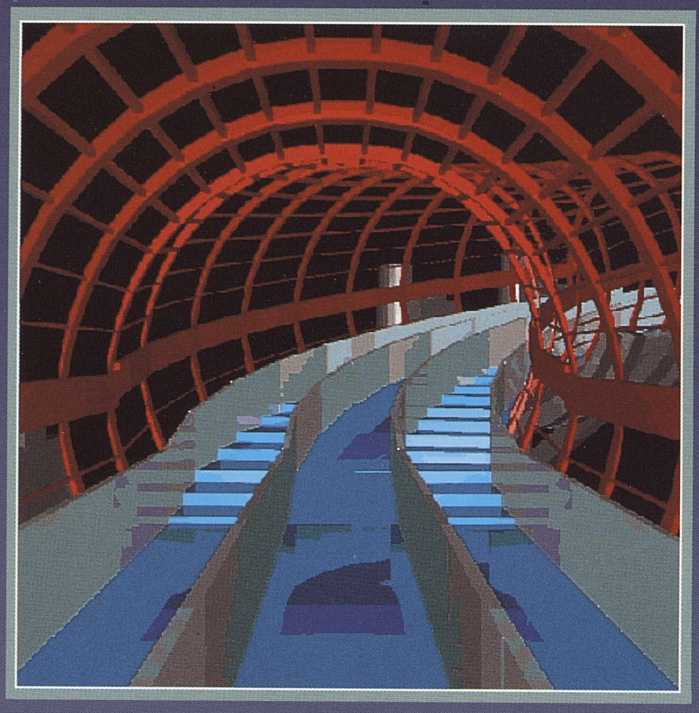

Joel Welling collaborated with Carnegie Mellon architectural student Alaa Eldin Ibrahim to produce a video "walkthrough" of this walkway, designed by Ibrahim.

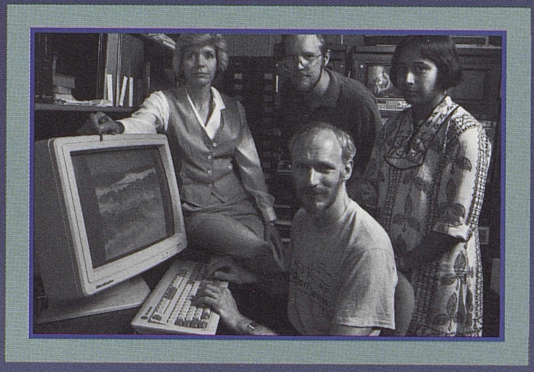

Members of the Pittsburgh Supercomputing Center's scientific visualization group (l to r): Grace Giras, Chris Nuuja, Anjana Kar & Joel Welling (seated).

A Good Look at Bad Tissue

Using a novel combination of laboratory techniques and computer visualization technology, Pittsburgh Supercomputing Center Scientific Specialist Joel Welling collaborated with pathologist Richard Siderits and media services director Cyril Evans of Pittsburgh's Shadyside Hospital to construct a 3-D image of cancer in a human lymph node. This is one of the first attempts to use computer imaging to examine the microscopic detail of biological tissue.

Three-dimensional reconstruction allows researchers to view an organ or tissue system comprehensively rather than painstakingly a slice at a time. "Traditionally, anatomic pathology is a two-dimensional science," says Siderits, "based almost entirely in peering through a microscope at thin sections mounted on glass slides. A 3-D reconstruction can reveal interrelations that wouldn't otherwise be noticeable. With this lymph node, for instance, we see that what looks like separate tumors when we look at consecutive slices may in fact be different branches of a single tumor."

Welling, Siderits and Evans constructed the 3-D image from a lymph node removed from a woman with colon cancer. Siderits cut the node into slices and projected every twentieth of these thin sections, 17 total, onto paper with a projecting microscope. Evans then prepared the images for capture and saved them to graphics files on an IBM XT with an image capture board.

Welling used the P3D graphics software, developed under his lead, to render the slices as a composite 3-D solid object on a Stardent graphics workstation. He then produced a video animation that rotates and turns the node, sheds the outer wall and takes the viewer on a journey into the malignant interior, in reality less than a half-inch in diameter. "Branches of the tumor seem to stream past the viewer," wrote science editor Byron Spice of the *Pittsburgh Post-Gazette*. "The scene is equal parts horror and beauty." It's also a harbinger of things to come in the diagnosis, treatment and understanding of human disease.

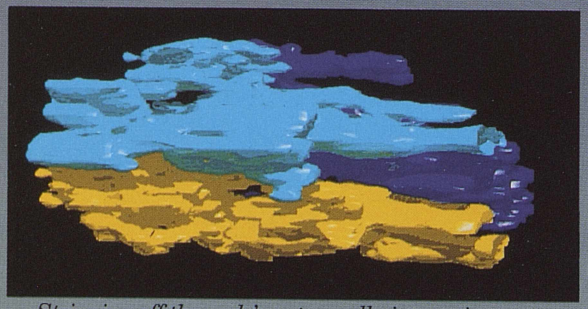

Stripping off the node's outer wall gives a view of the tumor within.

The white outer shell shows the lymph node surface, and the colored structures represent separate branching growths (metastases) of the cancer. Where the cancer extends through the lymph node surface, it anchors the node to surrounding tissue, a process which can be better observed with computer visualization.

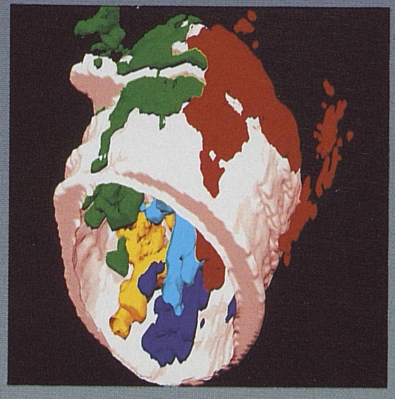

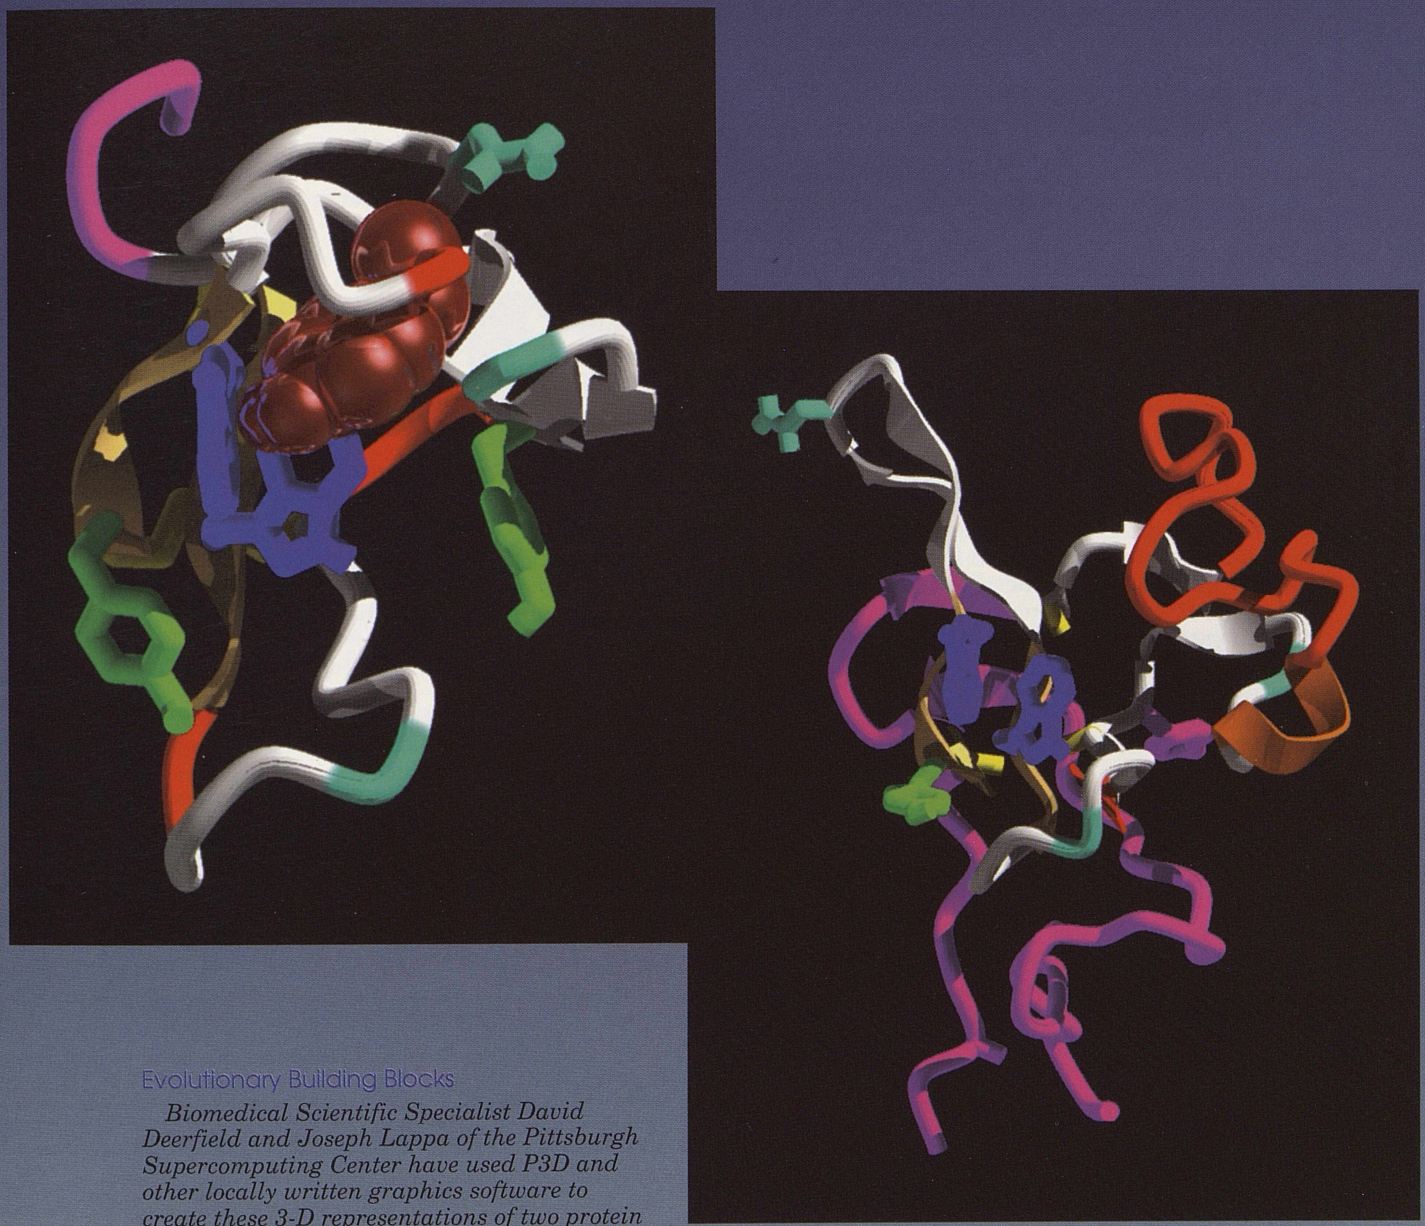

Evolutionary Building Blocks

Biomedical Scientific Specialist David Deerfield and Joseph Lappa of the Pittsburgh Supercomputing Center have used P3D and other locally written graphics software to create these 3-D representations of two protein domains. "Functional domains in proteins," explains Deerfield, "are the basic evolutionary building blocks of biological systems."

Deerfield has collaborated with Keith Constantine and Miguel Llinas in the Carnegie Mellon University chemistry department on research showing that these two domains, known as Type II (left) and Kringle (right), bear an evolutionary relation to each other. Related regions are shown by color matching in the two domains (white, blue, green and tan). Magenta represents portions of protein backbone that come before and after the structurally related regions. Red and orange structures represent deletions going from one domain to the other.

Technology Transfer

From Basic Research to New Products

Corporations and Supercomputing

To make the know-how of high-performance scientific computing readily available to American industry: This has been an objective of the National Science Foundation supercomputing program since its inception. Toward that end, to help meet the needs of any corporation looking to keep pace with state-of-the-art research capability, the Pittsburgh Supercomputing Center has a well developed Corporate Affiliates program. The basic idea: Make supercomputing easy to get to, not only the hardware but the knowledge that goes with it.

This year three of the Center's Corporate Affiliates, ALCOA, USX and Pfizer Inc, who have relied on the Center's resources as part of their research effort for a number of years, renewed their participation, and another major American corporation, Chevron Oil, came on board.

Chevron: Supercomputing and Oil

Far from being a babe in the woods when it comes to supercomputing, Chevron has had its own supercomputers for a decade and has developed computer modeling of oil reservoirs into a powerful engineering tool. "The question," says Chevron scientist Ernest Chung, who leads the company's simulation research group, "is how best to get the oil out of the ground. Computer simulations give us a way to account for the geological complexities of a reservoir and predict how it will respond to the stresses of drilling and related recovery methods. If the simulations help determine the recovery scheme that can get a few more percent of oil out of a large field, economically that's a tremendous gain."

For Chevron, the main attraction of affiliating with the Pittsburgh Supercomputing Center is to augment its in-house supercomputing with access to the Center's massively-parallel computer, the CM-2, and the opportunity to distribute the computing for reservoir simulations between different supercomputers. "Pittsburgh has the new technologies," says Chung, "and at the pace this technology is moving, we need to leverage our resources. We can't have it all in-house anymore."

In particular, Chevron researchers anticipate taking advantage of the Center's working high-speed link between its CRAY Y-MP and CM-2, and they expect to increase the level of detail in their oil reservoir simulations by a factor of 10 or more. Chevron researchers hope this improved modeling capability will allow them to better predict oil recovery under different production schemes, enhancing Chevron's ability to manage its oil and gas assets.

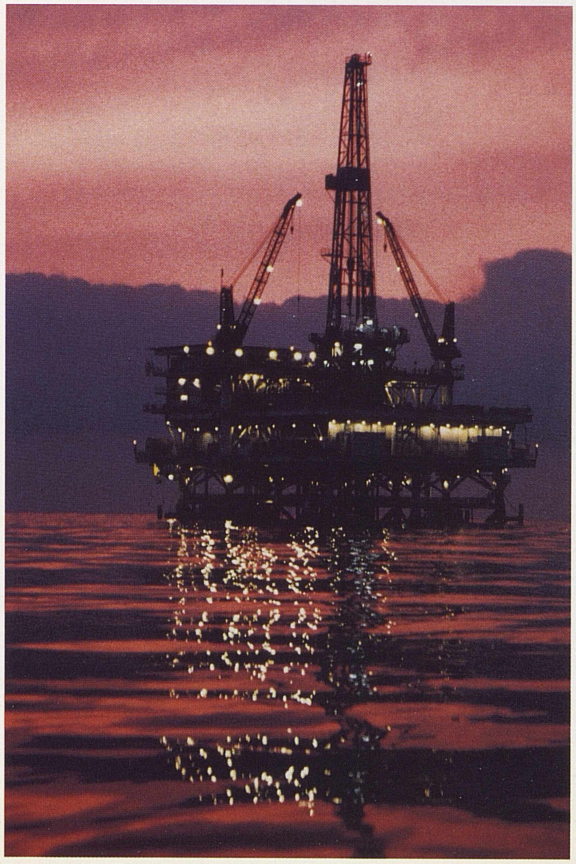

Larry Lee

Getting More Oil Out of the Ground

Obviously, the ability of Chevron (and other oil companies) to link their research with state-of-the-art high-performance computing has implications for national energy policy. Oil wells in general yield only 30% of the oil in the ground; the remaining 70% is not recovered. An increase of even one-percent in recovery from known oil fields in the United States will yield more oil than the largest domestic oil discovery in the last decade.

"Supercomputing is a tremendous national resource," says Beverly Clayton, Executive Director of the Pittsburgh Supercomputing Center. "It is one of our basic strengths in an increasingly competitive global economy. Our objective with the Corporate Affiliates Program is to harness this technology, and the knowledge base that goes with it, to the ingenuity of American business."

Technology Transfer Agreements

Increasingly, as the advantages of supercomputing become better known, more corporations are choosing to buy their own supercomputing hardware. For some of these organizations, the computing power they have can meet the demands of their research program. Nevertheless, researchers and systems personnel can be more effective when they have a knowledge base to turn to for problem solving and to stay abreast of new developments in the field.

Recognizing this need, the Pittsburgh Supercomputing Center this year implemented a new avenue for sharing its knowledge with industry: Technology Transfer Agreements. "Most corporate researchers who use supercomputing," explains Ira Hochman, Pittsburgh Supercomputing Center Corporate Liaison, "are too busy doing the research to stay on top of the technology. For them supercomputing is a tool; for us, what we do is supercomputing. It is our job to have the best tools available and know how to use them. With the Technology Transfer Agreement, we provide a structured program to make our expertise available."

In September 1990, E. I. DuPont de Nemours & Co. (Dupont) enrolled as the Center's first Technology Exchange Affiliate. Having upgraded its supercomputing from a CRAY X-MP to a Y-MP, Dupont was particularly interested in technical assistance with implementing a new operating system, UNICOS, on the Y-MP. They have also consulted with Center personnel regarding their graphics environment and workstation interfaces to the CRAY.

In March, Westinghouse became the second participant in the Center's Technology Transfer program. Westinghouse's fundamental interest is in maintaining access to the latest advances in file storage systems and heterogeneous systems.

DUPONT
REG. U.S. PAT. & TM. OFF.

Westinghouse

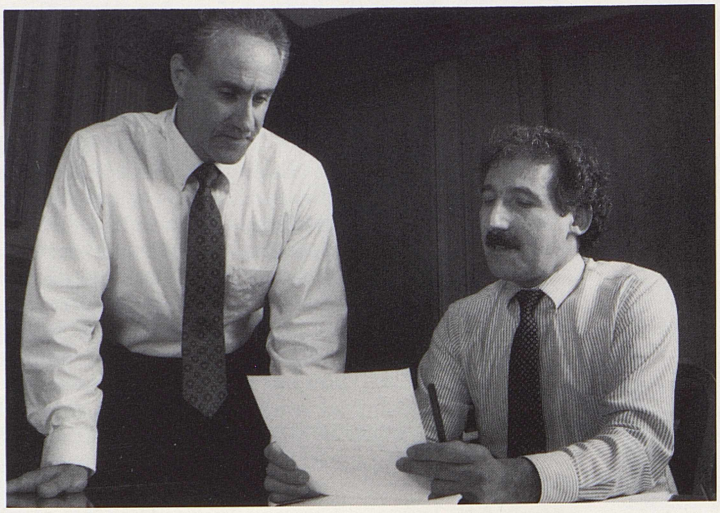

Ira Hochman (seated), Pittsburgh Supercomputing Center Corporate Liaison confers with Marvin Zalevsky, Assistant Director. "Part of the idea of our Corporate Affiliates Program," says Hochman, "is to ease the transition into supercomputing. We eliminate the guesswork about what to buy because it's part of our job to stay on top of this technology and to have the best resources currently available. With our workshops and staff, we also give companies the ability to avoid many of the problems that can occur in getting started with new technology. We give them the ability to hit the ground running."

Accelerating the Flow of Information

"University-based computational research," says Marvin Zalevsky, Assistant Director of the Pittsburgh Supercomputing Center, "is a fertile ground for the creation of sophisticated new techniques — algorithms and applications software, for instance — that can be very useful in industry. The problem traditionally has been *How do we make that knowledge available?* The supercomputing center is one answer.

"We're here not only to provide access to the best hardware, which we certainly do. But the hardware doesn't do any good unless people know what it can do for them and how to use it. That's what we provide through the workshops and consulting services we make available to corporations who affiliate with us. We're a channel for the flow of information from universities to industries who can use it directly to make new and better — and cheaper — products. We accelerate that flow of information."

Shared Assignment

A Distributed Computing Solution to the Assignment Problem
Gregory J. McRae, Carnegie Mellon University

When Two Supercomputers Are Better Than One

"Heterogeneous computing over high-speed networks can fundamentally change the way we do science and engineering," says Greg McRae, professor in chemical engineering and engineering and public policy at Carnegie Mellon. The basic idea of heterogeneous computing, linking different computers together to do things they couldn't do alone, is not new. What is new, however, is that a team of researchers from Carnegie Mellon University and the Pittsburgh Supercomputing Center has for the first time made it work between a CRAY Y-MP and a Connection Machine.

McRae and Robert Clay, a graduate student in chemical engineering, brought their knowledge of chemical plant process modeling into a working relationship with the network communications know-how of three members of the Pittsburgh Supercomputing Center staff — network engineers Wendy Huntoon and Matt Mathis and senior consultant Jamshid Mahdavi. In February, they linked the Center's CRAY Y-MP/832 and its 32,000 node Connection Machine CM-2, supercomputers of radically different design, with an interconnect for high-speed data transfer known as a HiPPI (high-performance parallel interface). Though the feasibility of doing this has been discussed in high-performance computing circles for some time, the Pittsburgh team are the first to make it work.

The research team that developed the first CRAY to Connection Machine HiPPI link is shown with their hands full of HiPPI; the black cable they are holding is a HiPPI cable like the one connecting the two supercomputers. Seated: Robert Clay and Wendy Huntoon. Standing (l to r): Jamshid Mahdavi, Gregory Mcrae and Matt Mathis.

McRae and Clay used the HiPPI link to do the computations for the "assignment problem," a classic and important problem in the field of practical mathematics usually called combinatorial optimization or combinatorial theory. Their code in effect uses the two supercomputers as if they were one linked system passing data back and forth between the CRAY's vector processing units and the CM-2's massively parallel array of processors. The advantage is that inherently serial parts of the problem can run on the CRAY and other parts of the problem, amenable to parallel solution, can exploit the massively parallel architecture of the CM-2.

"We've been able to reduce the elapsed time," says McRae, recipient of the first Forefronts in Computational Science award (given in 1989 for his work in large scale air-quality modeling), "by 10 times over what could be done on either the CRAY or the CM-2 alone."

Process Modeling and the Assignment Problem

For McRae and Clay, who came to do a Ph.D. at Carnegie Mellon after several years as lead engineer for Exxon's real-time optimization group, the assignment problem is important because it plays an essential role in scheduling the day-to-day operations of chemical plants. "It's a very demanding computational problem," says McRae, "with a multitude of very real practical applications, not only in chemical engineering but across the board."

As an example, McRae cites the airline industry. Assigning crews to routes can present difficult scheduling problems. A limited number of crews are available, and during any particular time frame each can fly only one route. The assignment problem asks what match between crews and routes, out of all the possibilities, will most reduce the overall cost (layover time, travel expense, per diem etc.) associated with each crew to route assignment.

In a modern chemical plant, such as an oil refinery, each separate stage of the production process, from raw materials to finished product, is mathematically modeled and controlled through process monitoring. The objective is to minimize the elapsed time between readout of control monitors and computations to adjust and optimize the production schedule.

"If you want to schedule the operation of a chemical plant," says McRae, "you have to solve the equations describing the chemical dynamics as well as devise the schedule faster than the plant is running or it's useless to you. We want to solve the problem 100 times faster than the plant itself is evolving. That way we can adapt and change, look at different schedules and figure out the best one, even as the plant is actually running."

A Distributed Solution: The Hungarian Algorithm & Initial Matching

The assignment problem is a significant chunk of the computing involved in plant process modeling, and McRae and Clay's distributed code attains exceptional performance. Their approach derives from a well-known solution of the assignment problem known as the Hungarian algorithm — after the Hungarian mathematicians, König and Egerváry, who developed it in the early 20th century.

"What's interesting," says McRae, "is that like most problems, there are many different ways to solve the

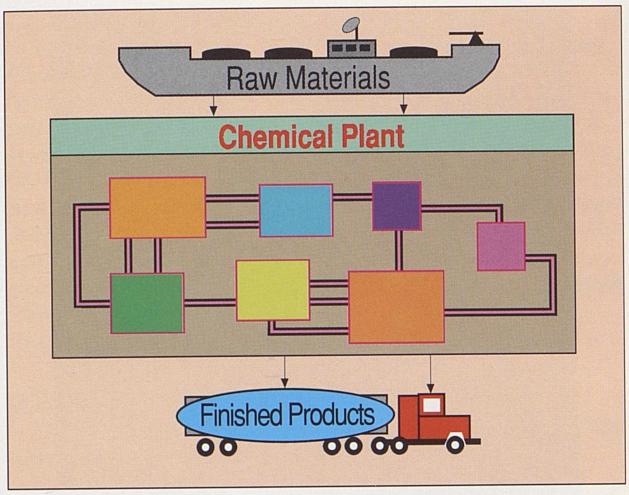

Chemical Process Modeling

assignment problem, and the Hungarian algorithm is one of the earliest. Like many efficient parallel algorithms, it comes from the days when numerical analysis was done by teams of people with desk calculators. As serial computers became available, different numerical strategies evolved, but now, with massive parallelism, some of these earlier approaches begin to look attractive again."

McRae and Clay experimented with various ways of mapping the Hungarian algorithm to the massively parallel architecture of the CM-2. Their optimum solution has the advantage that performance improves with increasing problem size, so that for large matrices their code significantly outperforms the existing state-of-the-art solution to the assignment problem (developed by Joseph Pekny of Purdue, a former McRae student). "This part of our code runs spectacularly well on the CM-2," says McRae.

The additional gain from distributing the problem between vector and massively parallel systems derives from the part of the solution called "initial matching." Each row of the assignment "cost matrix" has as its elements the cost associated with assigning each resource (e.g., airline crews) to one of the various possible demands (e.g., flights). This matrix — 8000 by 8000 elements in McRae and Clay's test runs — is loaded into the CM-2.

The first step, explains Clay, is a simple, fast initialization that establishes a large number of potentially optimum assignments; on the order of 80% or more of the demands are assigned to a resource. The CM-2 then transfers this much reduced dataset across the HiPPI to the Y-MP, where the initial matching algorithm finds "the maximum number of those assignments which are optimal at that stage."

"Initial matching," says McRae, "is an inherently serial procedure, and the algorithm works much better on the Y-MP, a very fast machine for serial calculations." After completing the initial matching, the Y-MP sends the results back to the CM-2, where the parallel Hungarian algorithm completes the solution; an iterative process creates additional optimal matches until all resources are assigned to a demand.

Proof of the Pudding

Probably the main significance of the Pittsburgh accomplishment is as a concrete indication that heterogeneous computing is not just a fascinating idea; it works. In a sense, the East and West of supercomputing have now been linked, and they are likely to stay linked. "This is proof of the pudding," says McRae, "that a heterogeneous computing environment works; it is no longer a debate about which supercomputer is best but how to exploit the best features of all of them on the network."

This research was supported by the National Science Foundation and the Defense Advanced Research Projects Agency under Cooperative Agreement NCR-8919038 with the Corporation for National Research Initiatives.

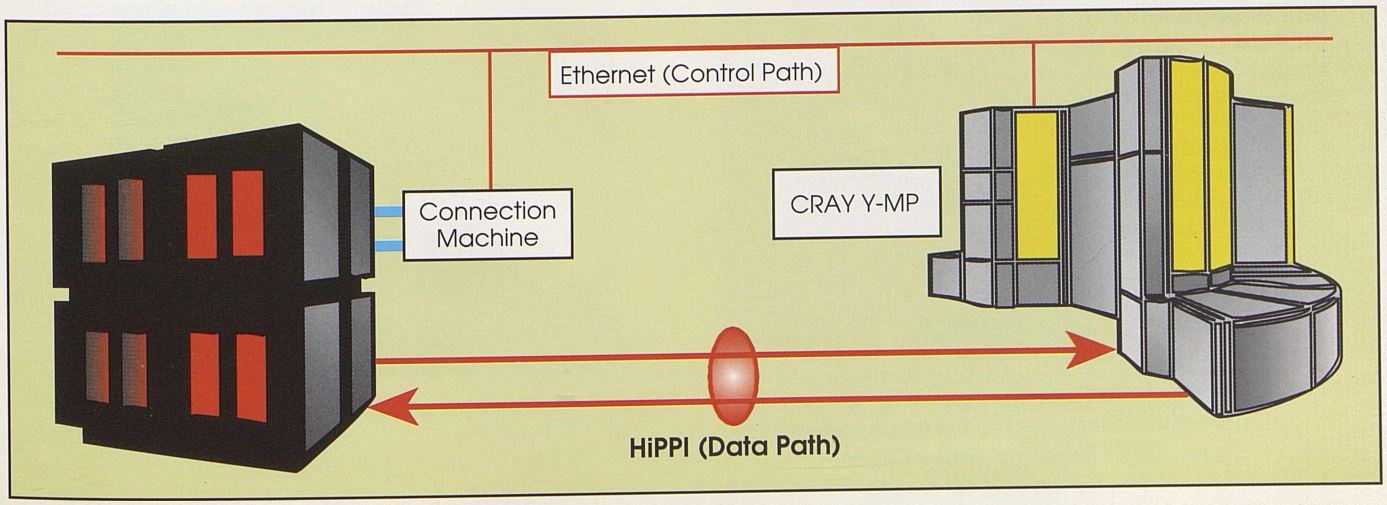

Shear Heat of Sudden Impact

Finite-Element Simulation of Adiabatic Shear Banding
P. C. Chou, Drexel University
Javad Hashemi, Texas Tech University

Success at Modeling High-Speed Failure

As the popular song says, when an irresistible force meets an immovable object, something's gotta give. Especially when the impact happens at high speed, as when a metal cutting-tool digs into a sheet of steel moving at thousands of feet per minute. Or when a rocket-propelled projectile crashes through the armor plating of a tank.

In engineering terms, the metal "fails," and when failure occurs as a result of high-speed impact it's often related to a phenomenon called *adiabatic shear banding*. "It's a very important dynamic mode of failure," says Javad Hashemi, "first identified in the early 1940s." As a post-doctoral researcher at Drexel, working under Professor P. C. Chou, Hashemi has helped develop code for the CRAY Y-MP that simulates adiabatic shear banding with more success than has been achieved before. As a result, their code — called DEFEL — is now being used by several defense contracting firms. DEFEL also offers promise as a valuable aid in designing safe, efficient cutting tools for high-speed metal removal processes used in many kinds of manufacturing.

Adiabatic Shear Banding?

"It's a thermoplastic instability phenomenon," explains Hashemi, meaning basically that heat is produced and the metal deforms as a result. "What happens is that due to the load from the impact, the material starts to strain harden, and this produces plastic work. That plastic work in turn generates heat, and if the process is happening fast enough, there's not enough time for the heat to diffuse through the material." As a result, the heat localizes in a narrow region, the adiabatic (which means no heat exchange) shear band. This localized heat raises the temperature in the shear-band region, decreasing the flow stress — often called the "yield point" — of the material, which means that the same amount of load causes more deformation than it would at lower temperature.

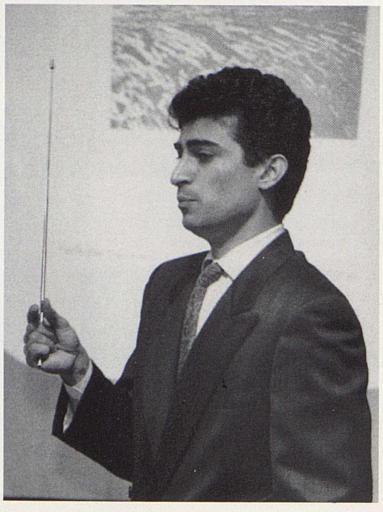

Javad Hashemi explains adiabatic shear banding in a seminar at the Pittsburgh Supercomputing Center.

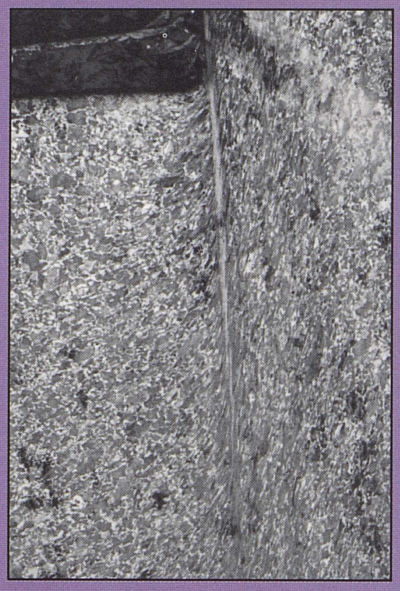

An adiabatic shear band in pearlitic 4140 alloy steel. In this 100 x magnified photograph, severe shear deformation is localized in a region about five micrometers wide [about .0002 inches]. "The temperature in the shear band goes so high," says Hashemi, "that it causes a phase transformation. When this transformed region is exposed to a specific kind of etching, it turns white."

Controlled Depth-of-Penetration Experiments

To create shear bands experimentally, Hashemi and Chou rely on a relatively simple method they call "controlled depth of penetration impact." The apparatus has a penetrator and a stopper plate in contact with a steel target. A projectile is dropped onto the penetrator, for low velocity impact, or it can be shot from a gun at ballistic velocities (between one and two kilometers per second). "With this apparatus," says Hashemi, "we can control the depth of penetration and generate shear bands of a specified length. It allows us to examine high strain-rate phenomena at relatively low-velocity impact."

Finite-Element Computations

Over the past eight years, Chou's group at Drexel has developed DEFEL, a finite-element code to computationally simulate shear banding and other metal-forming processes. "It's very difficult to experimentally measure the temperature within a shear band," says Hashemi, "and there's not much known about it. If we can accurately simulate this process, we'll learn a great deal."

DEFEL uses a material model, the Johnson-Cook constitutive equation, that includes terms for strain hardening of the material and for the softening brought about by temperature increase. The finite-element technique in effect lays a mesh-like grid over the metal surface, separating the computation into many discrete elements. To reflect the very small dimensions of the shear band phenomenon, the code uses an extremely fine mesh size — five microns — in the region just below the penetrator, where the shear band occurs.

"The dimension of these elements is so small that it's at the grain level," says Hashemi, "and the stability of the computation depends on these elements not being stretched beyond their limit." To control this, Hashemi's code automatically decreases the time increment for which strain and temperature in each element are computed as the mesh-size stretches to its limit.

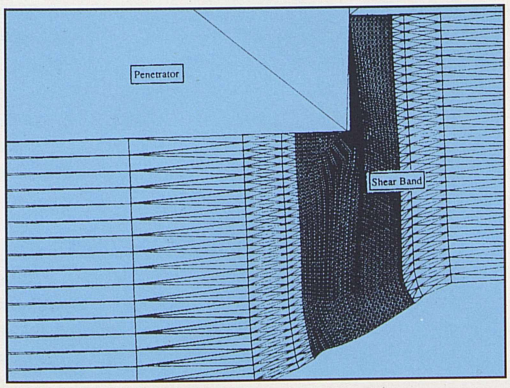

Shear band simulation using uniform mesh size of 0.005 mm.

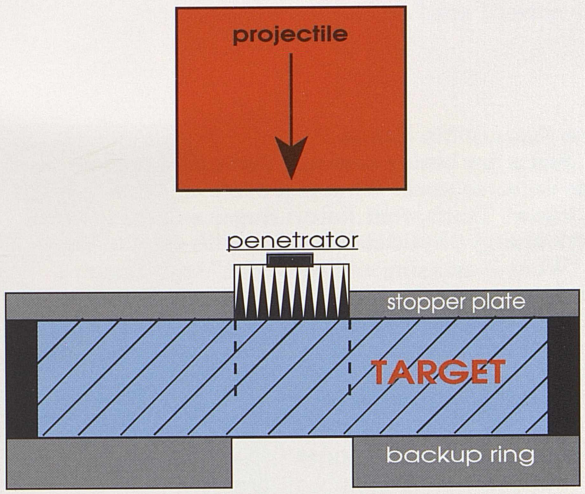

Controlled Depth-of-Penetration Impact Test Apparatus

The Y-MP computations simulating the controlled depth of impact studies give results that compare well with experiment. "You see hardly any deformation except in the localized column at the edge of the impact. This is an excellent simulation of shear banding." To verify that his code is actually simulating the physics of the problem, Hashemi also did computations that excluded the material softening from heat localization. These simulations didn't produce any localization or shear banding. "We found that, yes, the thermal softening effects are very important; this is what controls formation of the shear band."

High-Speed Machining & Explosive Detonation

Having established that their code realistically simulates shear-banding produced in controlled experiments, Hashemi and Chou have begun work on several real-world problems.

In high-speed machining, it's desirable that the metal be dislodged from the material surface in segmented chips rather than a continuous slice. Long continuous chips can bind in the cutting machinery, increasing wear and tear and creating a hazard to workers. Segmented chips form as a result of adiabatic shear banding, the onset of which varies with the properties of the metal being machined and the speed of the machining. In studies to date, DEFEL has successfully simulated the basic "brittle fracture" mechanism of the chip formation process.

Hashemi and Chou have also begun to simulate explosive detonation. Shear-banding is believed to be a factor in what's known as *insensitive munitions*. "For example," explains Hashemi, "when bombs are in storage, if one explodes by accident, the fragments will hit other bombs and can create a chain effect that's catastrophic. We want to understand the detonation mechanics so we can prevent that from happening."

References:

J. Hashemi, P. C. Chou & A. A. Tseng, "Thermomechanical Behavior of Adiabatic Shear Band in High Speed Forming and Machining," **International Symposium on Heat and Mass Transfer in Manufacturing Processes, Dubrovnic, Yugoslavia, August, 1990.**

This research is supported by Cray Research Inc. and the National Science Foundation.

Quark Soup

Simulating Hadron Thermodynamics with Massively-Parallel Computing
The High-Temperature Quantum Chromodynamics Collaboration

The Rules of the Game

Just a few basic questions, that's all. Like how does the universe work? What is it made of? What rules does it follow? That's what Robert Sugar and the group of high-energy physicists he works with want to know.

"What is amazing and beautiful," says Sugar, a professor at the University of California, Santa Barbara, "is that up to now we have always found that some very simple rules determine what we see in the physical world. With these few simple rules we can predict a huge amount. Discovering those rules is very interesting."

Robert L. Sugar

Sugar and his colleagues are using massively-parallel computing at the Pittsburgh Supercomputing Center to expand what we know about one of these "simple" rules. With the Center's Connection Machine, a CM-2 with 32,768 processors, they are tackling one of the grandest of the grand challenges identified in "The Federal High Performance Computing Program" (a 1989 report from President Bush's Office of Science and Technology Policy). The challenge is quantum chromodynamics, QCD as it is known in the trade — the theory of the strongest force in nature, the force that holds the nucleus of an atom together.

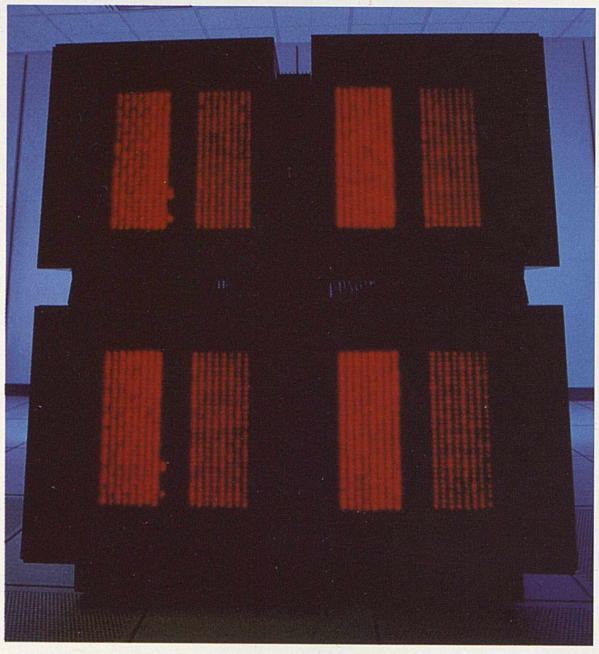

May the Force Be With You

Along with gravity, electromagnetism and the weak nuclear force — which governs nuclear decay, the nuclear strong force is one of twentieth-century physics' gang of four: the four forces that at root level determine everything in nature. Until about twenty-five years ago, physicists called protons and neutrons and their relatives (which include mesons, baryons and many others) elementary particles. These very small things inside the nucleus of an atom, generically called hadrons, were thought to be the indivisible, basic building-blocks of matter. QCD changed that by introducing something even more fundamental: quarks. Protons and neutrons, says QCD, are bundles of quarks, three in each bundle. The job of weaving quarks into these webs of energy we call matter is carried out by particles called gluons — because they act like the strongest imaginable glue.

QCD holds sway as the prevailing explanation for the smallest things in the universe even though no one has ever seen a quark or gluon. Under ordinary conditions, according to the theory, individual quarks or gluons cannot be separated from the bundles they form. Even today's powerful particle accelerators, descendants of the atom smashers of yesteryear, have not been able to break a hadron apart into its individual quark constituents.

"If we tried to pull one quark out of a proton," explains Sugar, "the strength of the force holding it in would require us to add so much energy that instead of freeing a quark we would create a new quark/anti-quark pair. It would be like pulling on a string, and suddenly the string would break, but instead of a quark in one hand and two quarks in the other, we would end up with a meson in one hand and a proton in the other."

Free Quarks & the Big Bang

Nevertheless, physicists think that under very special conditions it will be possible to "see" quarks. The trick is finding a way to reproduce the first second or so after the big bang. At that first instant of time, modern physics tells us, some 15 billion years ago, everything in the universe — everything that has since been rushing outward to form galaxies, stars, planets and people — was packed in a single spark of energy.

Under the unimaginably extreme conditions of those first few microseconds, so extraordinarily hot (about a trillion degrees Kelvin) that the strong nuclear force — the strongest force we know — becomes weak, the universe was a soup of free quarks. "One believes," says Sugar, "that at very high temperatures for very high densities there is a transition from ordinary hadronic matter to a state which can be described as a plasma of quarks and gluons."

Work is currently underway on several powerful new colliders designed to look for quark-gluon plasma, yet even if these experiments work as expected, free quarks will be difficult to detect. Information from computer simulations of QCD will help interpret these experiments, which is one of the main reasons why this work is recognized as a grand challenge. "Identifying a definite signal for this state is not easy," explains Sugar, "and one should understand this high temperature phase as well as one can."

"There are therefore agents in nature able to make the particles of bodies stick together by very strong attractions. And it is the business of experimental philosophy to find them out."

— Isaac Newton

The High-Temperature Quantum Chromodynamics Collaboration

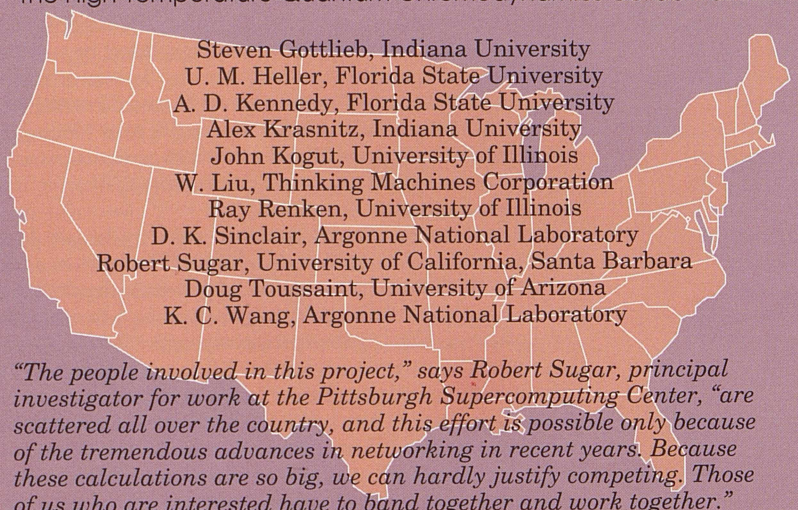

Steven Gottlieb, Indiana University
U. M. Heller, Florida State University
A. D. Kennedy, Florida State University
Alex Krasnitz, Indiana University
John Kogut, University of Illinois
W. Liu, Thinking Machines Corporation
Ray Renken, University of Illinois
D. K. Sinclair, Argonne National Laboratory
Robert Sugar, University of California, Santa Barbara
Doug Toussaint, University of Arizona
K. C. Wang, Argonne National Laboratory

"The people involved in this project," says Robert Sugar, principal investigator for work at the Pittsburgh Supercomputing Center, "are scattered all over the country, and this effort is possible only because of the tremendous advances in networking in recent years. Because these calculations are so big, we can hardly justify competing. Those of us who are interested have to band together and work together."

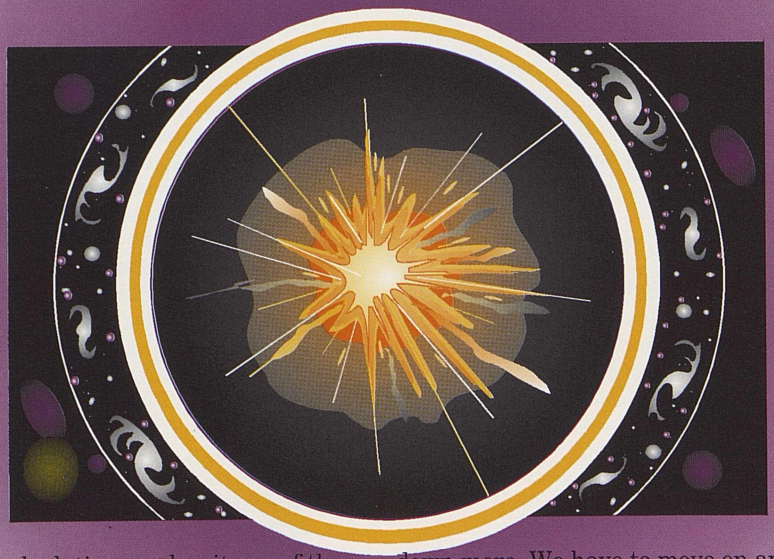

Looking for Quark Soup

The difficulty of QCD calculations makes it one of the most demanding computational projects in science. It is also one of the most promising avenues for progress in elementary particle physics. Because the problem is especially well suited to CM-2 architecture, the QCD code developed by Sugar and his collaborators will perform at the impressive rate of 4 billion calculations a second on a full-size CM-2 (65,536 processors). As a result, this research group has been able to simulate QCD under conditions (smaller particle masses and a bigger computational box or lattice) that come closer to a full treatment of the theory than has been possible before.

In preliminary calculations, Sugar and colleagues identified the range of thermodynamic properties where the phase transition will occur. "We have located the crossover," says Sugar, "but we would like to pin it down more. We have to move on and do more study of the high-temperature phase. These calculations take a massive amount of computing power. We're making progress, but we have a significant way to go."

The objective is to confirm that twentieth-century physics is on the right track. There have been tremendous advances in what we know about the invisible world inside the atom, and we have an incredibly successful explanation — the big bang — for what we see in the vast realms of the cosmos. Is modern physics consistent with itself? Do the theories of the smallest and biggest things jive? That is the scientific payoff hanging in the balance of the search for quark soup.

References:
S. Gottlieb, W. Liu, R. L. Renken, D. Toussaint & R. L. Sugar, "QCD Thermodynamics with Eight Time Slices," *Physics Review D* **41**, 622 (1990).

This research is supported by the Department of Energy and the National Science Foundation.

Putting the Squeeze on Hydrogen

Theory of Ground State and Excited State Properties of Solids
Marvin Cohen & Steven Louie, University of California, Berkeley

High-Pressure Predictions

Hydrogen: a gas, first in the periodic table, simplest and lightest of the elements. Fill a huge balloon with it and you can fly away. Most abundant material in the universe. It fuels the sun and makes incredibly powerful bombs. This much we learn in school.

But how about this? Squeeze hydrogen hard enough, about 2 million times atmospheric pressure, and it becomes a metal. Squeeze it even harder, 4 million times atmospheric pressure, and it becomes superconducting — electrical current flows with no resistance — at a temperature possibly twice as high (about -45° F) as any existing superconducting material. That's what Marvin Cohen, Steven Louie and their colleagues have predicted from quantum theory calculations at the Pittsburgh Supercomputing Center.

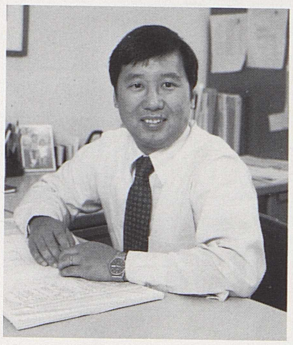

Steven G. Louie

Marvin L. Cohen

It may seem like a far-fetched idea, turning hydrogen into the most effective superconductor known by squeezing it at pressures like those in the core of the Earth, but stay tuned. A new technology, diamond anvil high-pressure cells, is making rapid strides. Laboratory scientists around the country are using the diamond anvils, working hand-in-hand with theorists like Cohen and Louie, to see how materials change under high pressure, and several groups are already working on making metallic hydrogen.

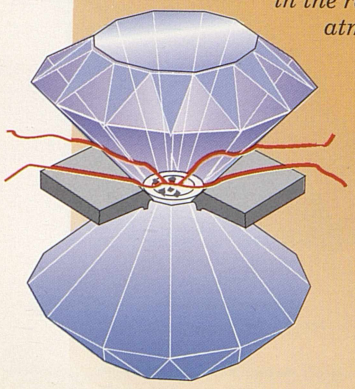

The diamond-anvil high-pressure cell. This device, which can exert pressures in the range of 2 to 3 million times atmospheric pressure, makes it possible for theorists like Cohen and Louie to verify their calculations. Two diamonds banded by steel are pressed together at the small ends, squeezing material between them. Laser beams and wires measure pressure and other characteristics of the material.

Calculating the Energy Gap

Over fifty years ago, in 1935, noted physicist Eugene Wigner and his Princeton colleague H. B. Huntington predicted that if you squeeze hydrogen hard enough it will become a metal. There was no way to test the idea at the time, but in 1979 two researchers at the Carnegie Institution, Peter Bell and Ho-Kwang "Dave" Mao, created solid hydrogen at about 60,000 times atmospheric pressure, strongly suggesting that Wigner and Huntington were on to something.

At this "low" pressure, however, solid hydrogen is an insulator — a poor conductor of electricity — not a metal. The difference has to do with what physicists call the "energy gap." In an insulator, electrons do not flow unless boosted by a jolt of added energy. "When you shine light on a system like this," explains Cohen, "you take an electron from a low energy level and you move it into a high energy level. With a metal, even the slightest bit of energy will move an electron to the next level. But in an insulator, you have to add enough energy to jump over a gap, and that's called the energy gap."

Cohen and Louie's calculations derive from the basic theoretical insight that squeezing a solid hard enough to rearrange the crystal structure can reduce the energy gap. Squeeze an insulator hard enough, in other words, and voilá — a metal.

Until about five years ago, there was no way, even with supercomputing, to calculate the energy gap. The reason, explains Louie, is that an electron in a solid is affected by the forces from billions of trillions (10^{22}) of other electrons. These complex interactions within the dense electron cloud in a solid make it extremely difficult to isolate the energy of a single electron, which is what is required to calculate the energy gap. "It wasn't until 1985," says Louie, "that we developed an approach to calculating the excitation energy of electrons in solids so that we are able now to calculate the energy gap."

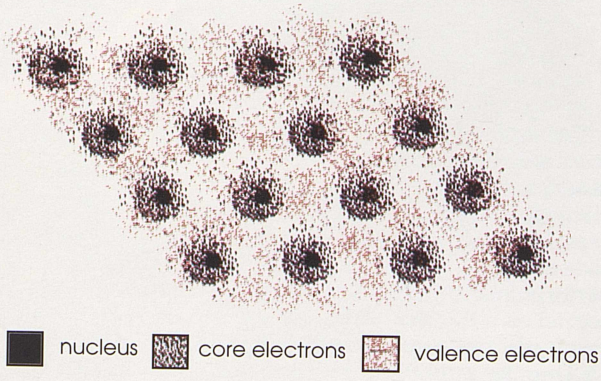

A solid is an array of atoms, each surrounded by a tightly held core of electrons. The array is itself embedded in a sea of free-floating valence electrons, electrons from the atom's outer shell which in a metal form the flow of electrical current.

Validating the Theory: Silicon & Xenon

Cohen's group at Berkeley achieved a major advance in 1985 when they applied their approach, called the total-energy pseudopotential, to silicon. Starting from "first principles," with only the atomic number and mass of the element, no experimental data, they determined that silicon, a semiconductor, would under extreme pressure become a hexagonal superconducting metal. These predictions have since been experimentally confirmed. "For the first time," notes Cohen, "a material and its superconductivity were predicted using a first-principles theoretical approach."

In 1989, another opportunity to see how well their theory worked presented itself to Cohen and Louie. Laboratory experiments showed that solid xenon under high pressure changed from an insulator to a metal. Using the CRAY Y-MP at Pittsburgh, Louie did calculations that closely agreed with the experimental data.

Two Kinds of Metallic Hydrogen

More computations on Pittsburgh's Y-MP led to the superconducting metallic hydrogen predictions, published in *Nature* in 1989. Hydrogen atoms tend naturally to form molecular pairs of two atoms each, and early theory suggested that these molecular pairs would have to split for a metallic solid to form. "Our first calculations," says Cohen, "were for that kind of model; we broke apart the molecules." The calculations searched for the most stable structure and arrived at one which was monatomic and hexagonal — a single atom at each of six points in the crystal structure.

Further calculations, however, revealed that metallic hydrogen can occur even when the hydrogen pairs stay together. "It turns out," says Cohen, "that when you squeeze hydrogen, you can actually make the energy gap smaller and smaller as a function of pressure. Our calculations show that if you squeeze hard enough, you could close the energy gap completely; the hydrogens would still be married at the lattice sites, but you would have a metal. The idea now is that there's two kinds of hydrogen metal. One is molecular, and one is an atomic periodic lattice. These two systems should both be metallic, and we think they should both be superconducting. We're working on that now."

References:
T.W. Barbee III, Alberto Garcia & Marvin Cohen, "First-principles prediction of high-temperature superconductivity in metallic hydrogen," *Nature* **340**, 369 (1989).
H. Chacham, X. Zhu & S.G. Louie, "Metal-Insulator Transition in Solid Xenon at High Pressures," *Europhysics Letters* **14**, 65 (1991).
Marvin L. Cohen, "Predicting New Solids and Superconductors," *Science* **234**, 549 (1986).

This research is supported by the National Science Foundation and by the Office of Energy Research, Office of Basic Energy Sciences, Materials Sciences Division of the US Department of Energy.

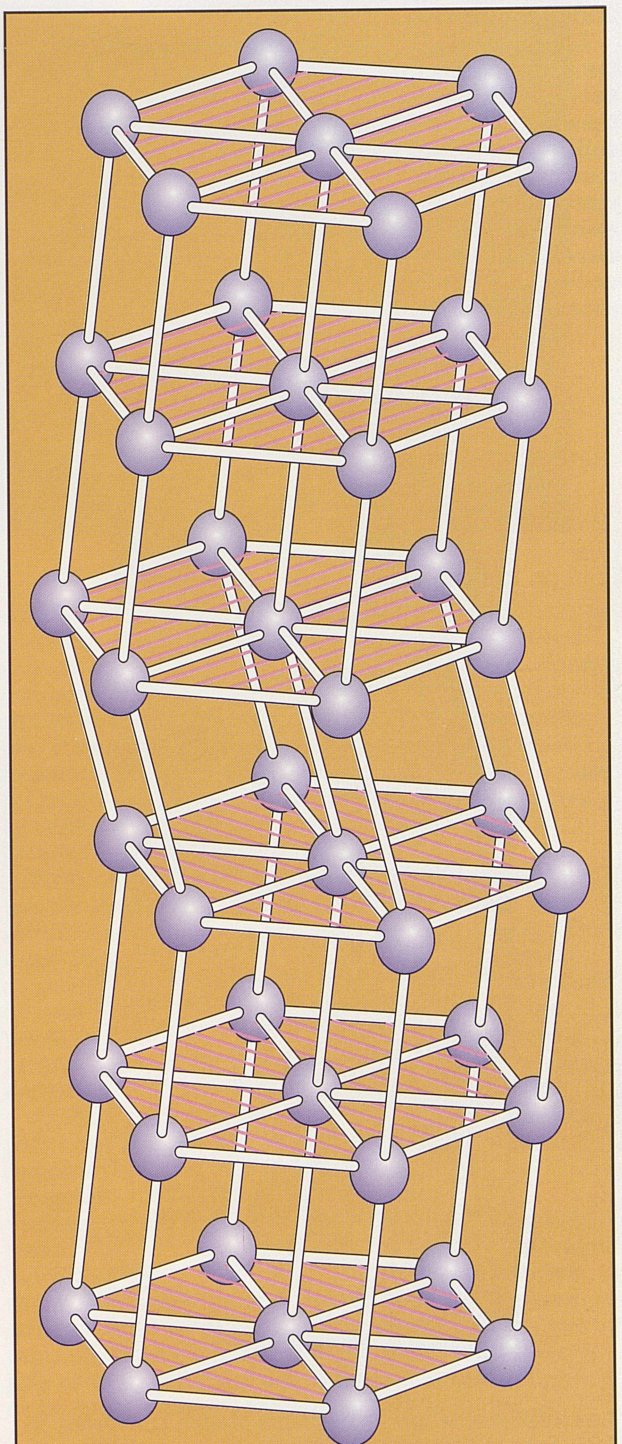

Monatomic metallic hydrogen as predicted by Marvin Cohen's calculations at the Pittsburgh Supercomputing Center. At high pressure, it forms this structure of three repeating planes.

Well Traveled Paths

Development and Application of Energy Derivative Methods in Molecular Orbital Calculations
H. Bernhard Schlegel, Wayne State University

Reaction Paths and GAUSSIAN

To hear Berny Schlegel talk about chemical reactions, you might think he was hiking in the Alps. "The nature of the mountain ridge separating two valleys — the elevation of the lowest mountain pass between the valleys and the route taken through this pass — controls the rate at which molecules move from the reactant valley to the product valley." He's talking about reaction paths, a concept that has become increasingly important to chemists as a way of better understanding the minute changes that occur during a chemical reaction.

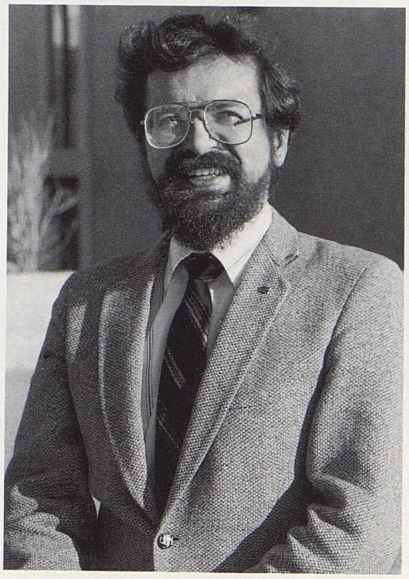

H. Bernhard Schlegel, Wayne State University. "In the last ten years, we've seen marvelous progress on the software and on the hardware, and theoretical chemistry is now coming to be an accepted partner with experimental chemistry. In the '70s, it was 'Let's see if we can do a few examples to show that theoretical chemistry is worthwhile.' By the late '80s, it was more 'Let's see if theory is a viable partner for some experiments.' I suspect that by the end of the '90s we'll see theoretical chemistry reliably predicting properties, energetics and reactivity in competition with experimental chemistry in certain areas."

As a visiting fellow at Carnegie Mellon University in 1977, Schlegel worked in John Pople's group and since then has played an active role in the evolution of GAUSSIAN, the package of programs — developed under Pople's direction — that has been called "the great workhorse of quantum chemistry." Most recently, Schlegel, in collaboration with his graduate student Carlos Gonzalez, has turned his attention to reaction paths. "We've developed a set of tools," he says, "that we think are the best path-following algorithms available for exploring reaction mechanisms."

Only with Supercomputing

Schlegel's work involves applying the insights of quantum theory to chemical reactions. In most cases, these calculations can be done only with supercomputers, and the results provide chemists with detailed information they can't get any other way. Laboratory methods are good at analyzing the starting materials and products of chemical reactions, but they tell little about the in-between stages. "In the lab, we have no tools to look at the structure of transition states," says Schlegel. "In other words, we have no good ways experimentally to get at the intimate details of how the atoms rearrange when molecule A collides with molecule B in the course of a chemical reaction.

"If we know those details," he continues, "and understand them, then we have a better grasp at predicting how similar reactions behave. For example, if molecule A reacts with molecule B, the calculations can tell us, for instance, whether that happens all at once, as a concerted reaction, or whether it goes step-by-step, through a series of transition states. We can then use our knowledge of that chemistry to design new reactions."

The key to getting these answers is computing efficiency. The quantum theory calculations required to map potential energy surfaces and trace reaction paths are enormously complex. Even with supercomputers, computing time is a severe limit on what can be accomplished.

"What really counts is how fast you can explore the terrain and trace out the path," says Schlegel, "and how few electronic structure calculations you have to do to prove that reactants connect to products along a particular pathway. Even simple, elementary chemical reactions can be very high dimensional systems with complex topography, making reaction paths difficult to follow." As compared to other algorithms for reaction path following, notes Schlegel, the methods that he and Gonzalez developed are able to take large step-sizes and still stay close to the reaction path.

Over the Pass from Reactant to Product

A reaction path is, simply, a trail that traverses a landscape-like "potential energy surface" — basically a topographical map of the energy of the reacting system. The ruggedness of the landscape varies with the complexity of the chemical reaction it maps. Valleys and the passes from valley to valley correspond to molecular structures at different stages of the reaction. The deepest valleys represent the relatively stable, low-energy "ground state" structures at the beginning and end of a reaction — the reactants and products.

The points of greatest interest along a reaction path

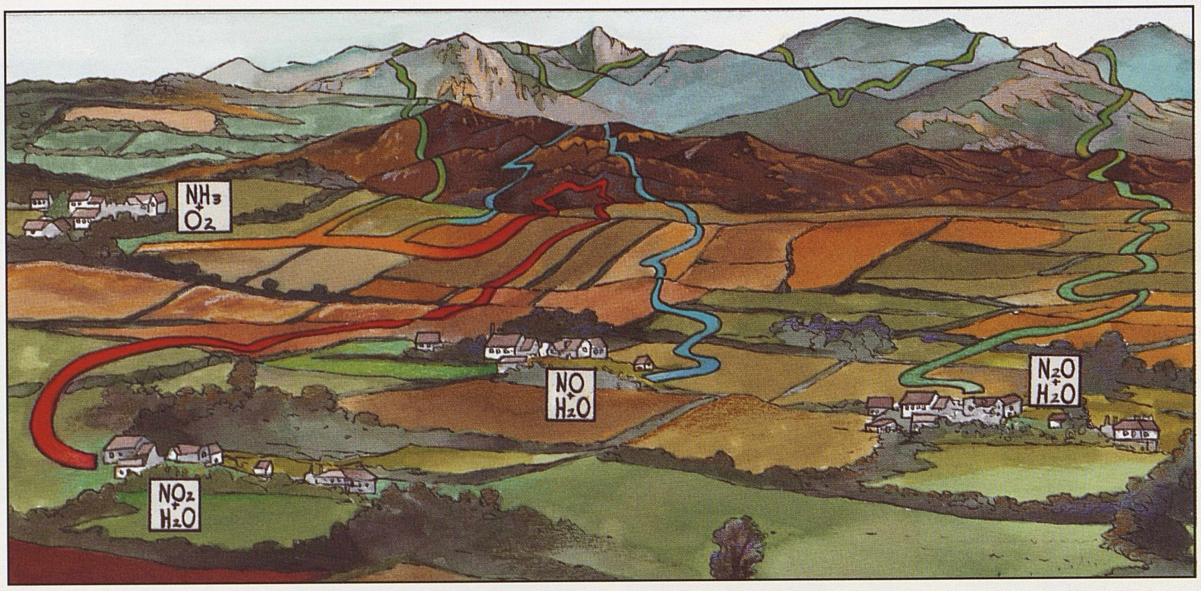

are the mountain passes that connect one valley with the next. These saddle points, also called transition states, correspond to high-energy, short-lived structures, often called activated complexes, that occur as the molecules in a reaction collide and rearrange.

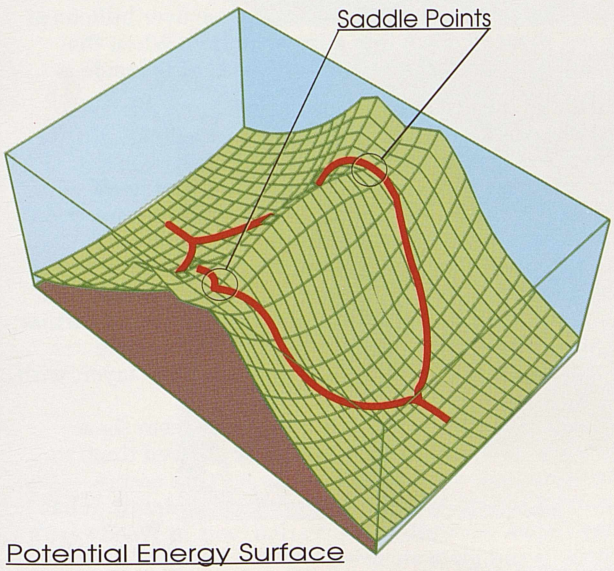

Potential Energy Surface

"If you visualize yourself as going up over a ridge," says Schlegel, "standing in one of these saddle points, then behind you is the reactant valley, where you've come from, and in front of you is the product valley. For a molecular system, the question is how do we get from the bottom of one valley to the other? The molecule has to get over that barrier, and generally it follows the path that uses the least energy — which means through a saddle point, one side or other of a peak."

Applied Reaction Path Following

As an example of what reaction path computations can accomplish, Schlegel cites work (with Wayne State colleague Robert Bach and others) on the transition structure for oxygen atom transfer from a hydroperoxide. Though oxygen atom transfer is one of the most important reactions in chemistry, significant in biological processes and in many oxidation reactions studied by organic chemists, it is not well understood at a detailed level.

"Prior to our computations," explains Schlegel, "people thought that an OH group moved when you oxidize something with hydrogen peroxide (H_2O_2). The calculations show that in the gas phase a hydrogen atom moves first and then a naked oxygen. That's quite a revolutionary concept. It might seem like a small detail, but the idea that this naked oxygen is moving rather than a larger group means that maybe we can design a different reagent to take advantage of this."

As these techniques become more efficient, reaction path computations can be applied to larger, more complex reactions that have practical as well as theoretical implications. "At the moment," says Schlegel, "we're limited to relatively small molecules, maybe up to a dozen atoms — that's roughly the realm in which we've been working. But that's the very low end of what experimental organic chemists are interested in. Still, we're looking at problems that are two to three orders of magnitude more difficult than we were able to study ten years ago. Supercomputers are giving us an opportunity we wouldn't otherwise have."

References:
Carlos Gonzalez & H. Bernhard Schlegel, "Reaction Path Following in Mass-Weighted Internal Coordinates," *Journal of Physical Chemistry* **94**, 5523 (1990).
Robert D. Bach, Amy L. Owensby, Carlos Gonzalez, H. Bernhard Schlegel & Joseph J. W. McDouall, "The Nature of the Transition Structure for Oxygen Atom Transfer from a Hydroperoxide. A Theoretical Comparison Between Water Oxide and Ammonia Oxide," *Journal of the American Chemical Society* (in press).

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

Cooking with the Stars

Computer Simulations of Compressible Convection in Stars
Paul Woodward, University of Minnesota

"Rage, rage against the dying of the light."

New Light on Stars

Burnout. It happens to rockets, people and stars. Even the sun. That bright orb of radiant energy whose appearance each morning is as constant as death and taxes will eventually use up the hydrogen that drives the nuclear reaction at its core.

No need to panic; scientists say there's at least five billion years remaining of sunrise as we know it. Maybe even six billion. Either way it is a very long time, though the range of difference in these estimates gives some indication of how little we really know about the inner workings of even the most familiar star.

Paul Woodward's research sheds new light on stars. Using the CRAY Y-MP, he and David Porter, a post-doctoral researcher at Minnesota, are doing detailed calculations on what happens in the layer of gas in the outer third of the sun's radius.

"What we want to study," says Woodward, "is convection in the outer layers of stars like the sun." Convection in stars is basically the same process that circulates heat in a radiator-heated room. The outward motion of hot gas balances with pockets of cooler (thus heavier) gas pulled back toward the center by gravity. "Fluid motion exists in the outer part of the sun," explains Woodward, "because the gas that makes up the sun, mostly hydrogen, is cooler near the outside and doesn't conduct heat well enough for the interior energy to escape by thermal conduction."

It is important to understand the details of this convection process, says Woodward, because it relates in a fundamental way to theories about the evolution of stars. "What is their structure? How long do they live? If some of the components in the outer parts of the sun mix back down, it could considerably change the rate of nuclear burning."

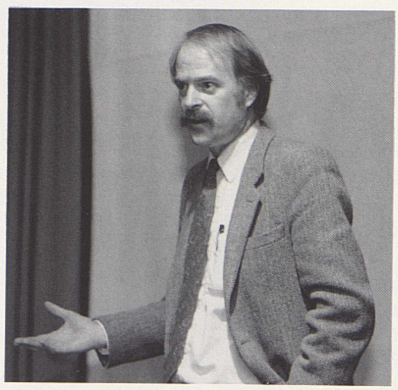

Paul R. Woodward

The Onion Model

Most theory about stars models them as spherical layers of gas, nested one inside the other, like the layers of an onion. At the extremely hot core (at least 14 million degrees Kelvin at the sun's core), hydrogen ions fuse into helium. The heat energy released by this reaction travels outward through cooler layers until at the surface it radiates away as light — with life on Earth as a byproduct. In large part, this model assumes that the gas in the outer regions stays where it is; it does not, in other words, mix down to refuel the furnace.

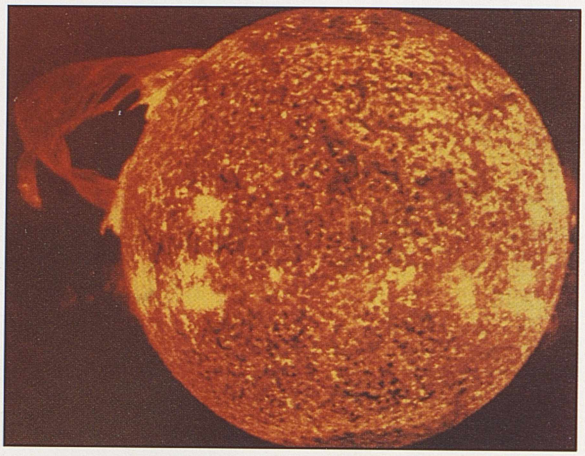

Woodward believes that this model is a reflection of limited computing capability. Until recently, it was not possible to do calculations that adequately tested whether convection in the outer region could send hydrogen into the sun's interior.

"So much of our thinking about stellar models is one-dimensional," explains Woodward, "because that is all we've been able to compute. You assume that all the material in the convectively unstable region out there at radius two-thirds has become thoroughly mixed over the lifetime of the sun. With this assumption, none of this material ever really has a chance, even over billions of years, to move down into a lower radius. When the model is a series of mass shells, one that is inside is always inside."

Visualizing Penetrative Convection

The possibility that the sun might refuel itself through what Woodward calls "penetrative convection" was originally suggested in earlier research (by Hurlburt, Toomre and Massaguer at the University of Colorado, Boulder). Woodward and Porter designed a calculation on the CRAY Y-MP at Pittsburgh to further test this possibility. They simulated an unstable convective layer, modeled on the sun's outer layer, with a stable layer of gas below it.

They created an extensive animation from these calculations, 7500 frames long, that allowed them to visualize the complex fluid dynamics that occurs between the two layers at a high level of detail. What they found, on careful observation, was in Woodward's words "a complete surprise."

They discovered that spinning vortices of turbulent gas tend to form pairs, with opposite spins joined together, and once paired these vortices "translate" — they travel through the surrounding fluid. Under the right conditions, when a descending vortex pair is moving close to the vertical, it will penetrate deep into the stable layer.

"You can see," says Woodward, "that these vortex pairs go down, and they don't come back up. This was a real discovery." These results suggest a mechanism by

which gas in the outer layer of stars can move into the star's interior, which would change current understanding about the rate at which stars exhaust their nuclear material.

Woodward emphasizes that since his calculations are two-dimensional the prediction of penetrative convection must at this point be considered tentative. "We don't know whether it is going to turn out to be important for real stars, which are three-dimensional, rotating, magnetized objects." In future calculations, Woodward plans to examine penetrative convection with three-dimensional simulations. "It is only by doing detailed multi-dimensional simulations like this," says Woodward, "that we will ever find out whether or not this can happen."

References:
David H. Porter, Paul R. Woodward & Qi Mei, "Simulations of Compressible Convection with the Piecewise-Parabolic Method (PPM)," *International Video Journal of Engineering Research* **1** (1991). (This is the first issue of this journal, published by John Wiley & Sons, which presents engineering research in video format.)

Epigraph: from "Do Not Go Gentle into that Good Night," Dylan Thomas.

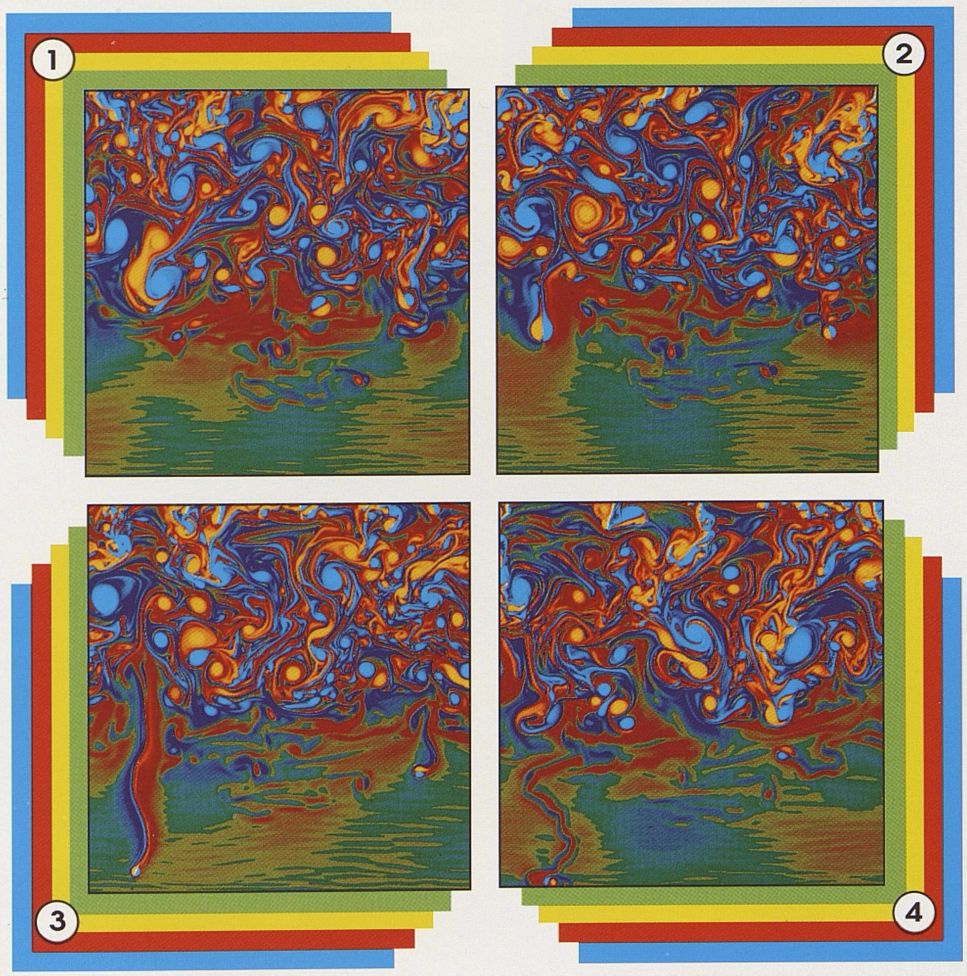

Penetrative convection, as simulated at the Pittsburgh Supercomputing Center by Paul Woodward and David Porter. Colors differentiate rotation and velocity of the fluid vorticity, yellow clockwise and blue counterclockwise. This sequence shows a vortex pair forming and translating downward from the unstable into the stable convection layer below it. "It's like trying to push a balloon under the ocean," says Paul Woodward, "but these vortex pairs don't just bob back up the way a balloon would because the lower region is more stable. Thermal conduction is more efficient down there. The temperature difference between the fluid in the vortices and the ambient fluid will even out, and that removes the buoyancy. This is how material from the upper layer can come down there and stay."

Taming the Whirlwind

The Study and Application of Chaotic Dynamics to Turbulent Flows

Lawrence Sirovich, Brown University

"A. A violent order is disorder, and
B. A great disorder is an order. These
Two things are one. (Pages of illustrations.)"

— Wallace Stevens, "Connoisseur of Chaos"

Where There's Smoke There's Disorder

Watch the plume of smoke that rises from a chimney on a still winter morning. It drifts upward in a smooth column and then, all at once, breaks apart into swirls and billows. At that point, order seems to break down completely. This phenomenon is called turbulence, and it is all around us, inherent in how gases and liquids act when they are in motion. Turbulence is part of the vagaries of weather, the flow of ocean currents, the combustion that powers engines.

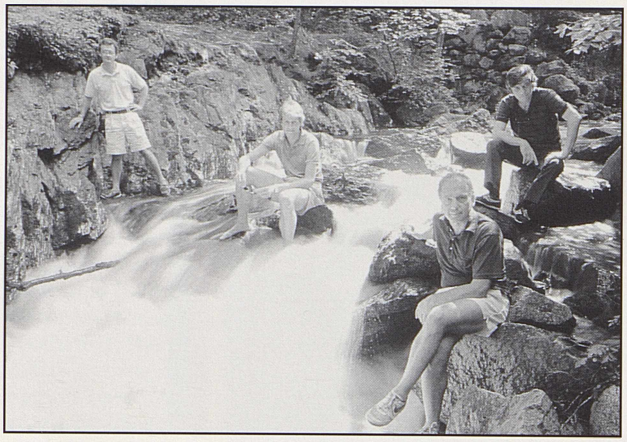

Lawrence Sirovich (front right) and his colleagues at the Brown University Center for Fluid Mechanics, (l to r) Martin Maxey, Eckart Meiburg and Fred Bisshopp, on the banks of a turbulent Rhode Island stream.

"Turbulence has been one of the hardest problems in science to crack," says Lawrence Sirovich, who directs Brown University's Center for Fluid Mechanics. "Many of the great luminaries in mathematics and mathematical physics have applied themselves to it and been turned back. Now, with large-scale computing, we have a new opportunity to look at these problems."

Understanding turbulence is not just knowledge for its own sake; the practical payoffs are tangible and significant. With airplane flight, for instance, turbulence in the air flow around the wings and fuselage of the plane increases drag, and reducing drag one-percent could save a billion dollars a year in the cost of fuel for commercial airline travel.

Recent developments in mathematics and physics suggest that lurking beneath the apparent chaos of turbulent flow is a predictable pattern, a kind of order, even simplicity, that if fully grasped would allow us to build faster planes, more efficient engines and in general to better understand something that influences our life in many ways. Using the CRAY Y-MP at Pittsburgh, Sirovich has applied some of these insights to develop a new computational approach to a classic turbulence problem, Rayleigh-Bénard convection.

Fingers of Hot Water

"The Rayleigh-Bénard problem," says Sirovich, "goes back to the mid-19th century. In simplest terms, it describes what happens when you heat a pan of water." The water at the bottom of the pan heats up first, and because hot water is lighter (less dense) than cold water, this heated water rises. As the hot water moves upward, it loses heat to the cooler water around it, and eventually, as the hot rising water cools, it loses its buoyancy, falls back down to the bottom, and the cycle repeats.

"As you increase the heat delivered to the pan," explains Sirovich, "this process becomes more and more violent, and you get turbulence. And this turbulence plays a role in transporting heat from the bottom to the top of the water."

Sirovich's computer simulations give a detailed picture of this turbulent heat transfer process, indicating that the lower half of the fluid has three distinct regions: In a thin layer at the bottom, heat moves from water molecule to water molecule (conduction) as well as through the upward movement of the heated water (convection). Above this layer, narrow plumes of heated water extend like fingers. These fingers grow and move around until they pinch off from the lower layer and accelerate upward into a turbulent mixing zone, where hot and cold waters swirl and blend together.

Soft and Hard Turbulence

These results are in general agreement with important experimental work on Rayleigh-Bénard turbulence done by a group at the University of Chicago led by Albert Libchaber. Both studies have identified a transition from one state of turbulence to another as heat is added to the system. The first state, called soft turbulence, has a random distribution of temperature fluctuations from region to region, while the second, hard turbulence, shows a structured pattern in the distribution of hot and cold regions, indicating that heat transport is less efficient.

"There is a profound difference between these two different forms of turbulence," says Sirovich. "We found the exponential distribution [hard turbulence], and that has been regarded as a strong confirmation of the Chicago experiment, and of course it is also strong confirmation of the simulation."

The uniqueness of Sirovich's approach to simulating turbulence lies in his method of simplifying. He has been able to reduce the number of equations needed to accurately characterize the system by using what is known as the Karhunen-Loéve procedure. "It goes back to the early part of this century," says Sirovich, "though until we worked on it, it was used only for one-dimensional systems." For one of Sirovich's three-dimensional Rayleigh-Bénard simulations, the Karhunen-Loéve procedure reduces the number of equations from 100,000 to a mere 400.

These simulations require 100-200 hours of CRAY computing time to generate the data, calculations that as a practical matter could not be done without access to supercomputing.

Numerical Experiments

For Sirovich, computer modeling is a way to learn things that can't be learned any other way, since laboratory approaches to turbulence are inherently limited in what they can measure. "In numerical experiments, you have a detailed space-time description. You can vary the parameters of the problem with ease; you can go through a range of parameters. You can vary geometry. Eventually, you'll be able to establish a flow in which all the parameters match what goes on in nature — something you can't possibly do in an experiment. That's for the future. There are very large-scale calculations that have to be done."

References:
S. Balachandar, M. R. Maxey & L. Sirovich, "Numerical Simulation of High Rayleigh Number Convection," *Journal of Scientific Computing* **4**, 219 (1989).
L. Sirovich, S. Balachandar & M. R. Maxey, "Simulation of Turbulent Thermal Convection," *Physics of Fluids A* **1**, 1911 (December 1989).

This research is supported by the Defense Advanced Research Projects Agency.

Photo: John Forasté, Brown University.

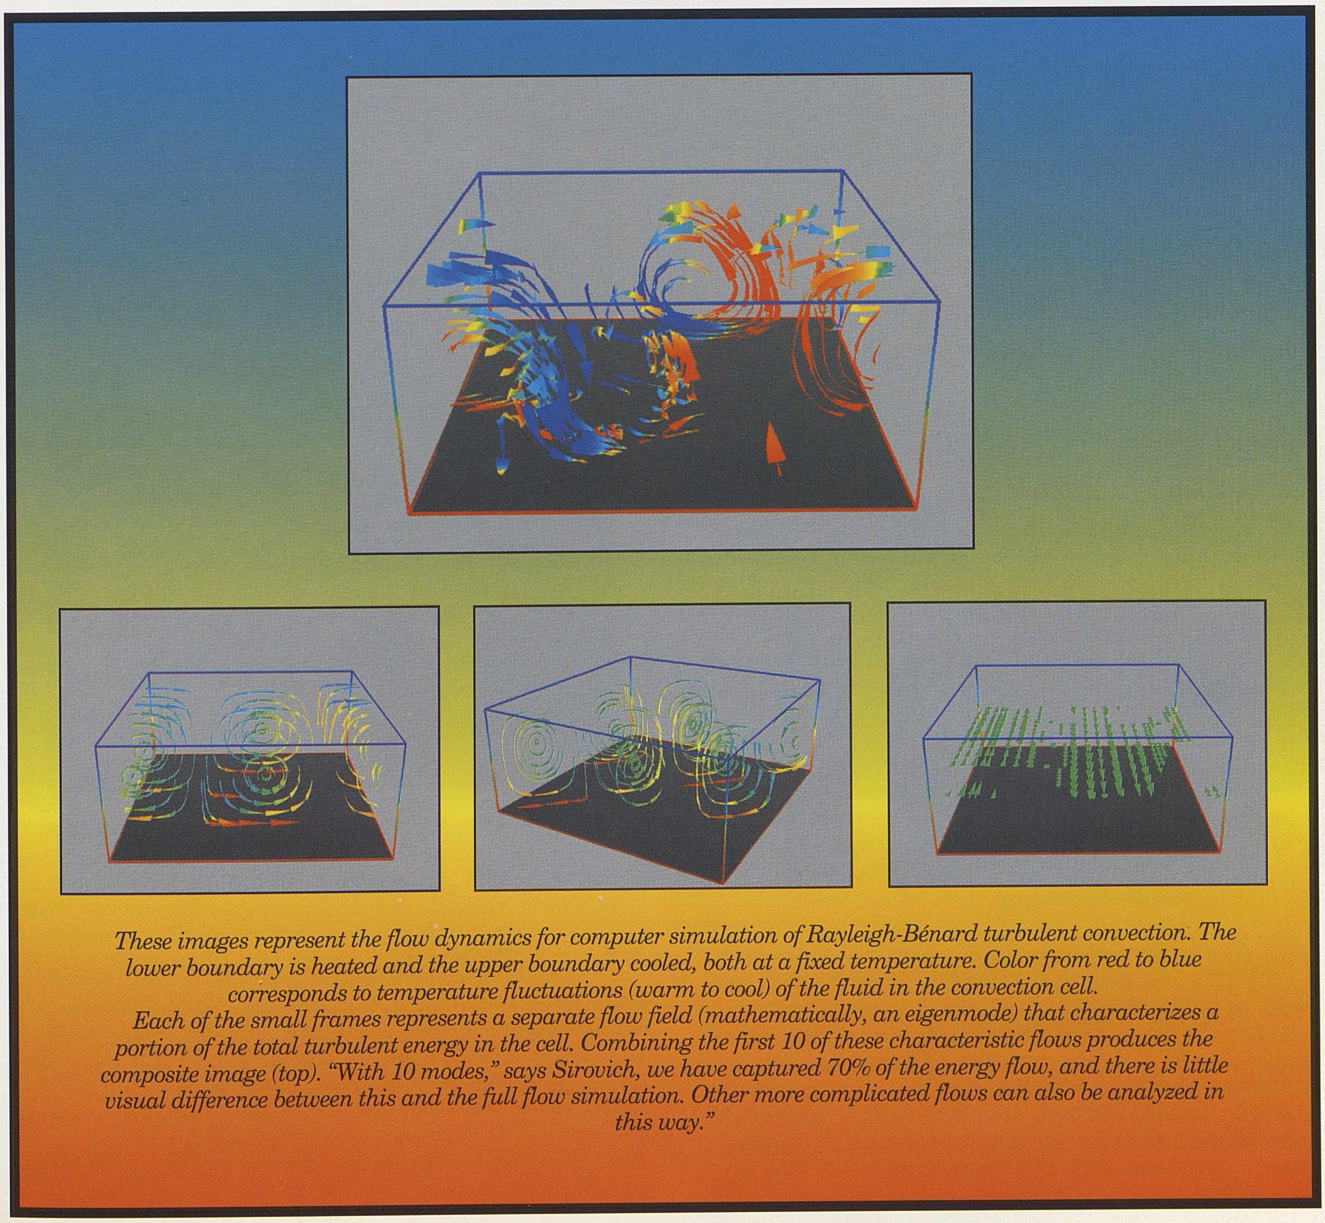

These images represent the flow dynamics for computer simulation of Rayleigh-Bénard turbulent convection. The lower boundary is heated and the upper boundary cooled, both at a fixed temperature. Color from red to blue corresponds to temperature fluctuations (warm to cool) of the fluid in the convection cell.
Each of the small frames represents a separate flow field (mathematically, an eigenmode) that characterizes a portion of the total turbulent energy in the cell. Combining the first 10 of these characteristic flows produces the composite image (top). "With 10 modes," says Sirovich, we have captured 70% of the energy flow, and there is little visual difference between this and the full flow simulation. Other more complicated flows can also be analyzed in this way."

Kinky DNA

Structure and Dynamical Simulations of the DNA-Eco RI Endonuclease Complex

John M. Rosenberg, University of Pittsburgh

"With thy sharp teeth this knot intrinsicate
Of life at once untie."
– Antony and Cleopatra, V, ii, 304-05.

The Double Helix

The discovery of the double helix structure of DNA, worked out in its rudiments by Francis Crick and James Watson in 1952, is without question one of the scientific triumphs of this century. DNA is the basic repository of genetic information, an amazingly intricate molecular code governing the biological processes we call "life." Yet, as it tends to go when we learn about the world around and within us, the more we know, the more we don't know.

Nearly everything that happens biologically from DNA happens because many different proteins somehow have the ability to isolate their activity at specific sites on the long helical strands that comprise DNA molecules. John Rosenberg uses the CRAY Y-MP to help understand these complex processes, usually called protein-DNA recognition.

"There are large classes of proteins," says Rosenberg, "that recognize specific sequences of DNA, and they do very important things — ranging from controlling which genes express when and how, to rearranging the structure of DNA itself. These are basic biological processes that we want to understand partly for their own sake, but also for two practical reasons: first, many disease processes may be related to aberrations in these events and, second, restriction enzymes are vital tools of the biotechnology industry."

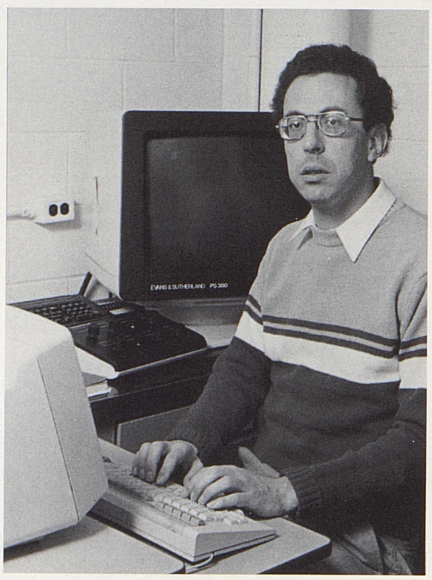

John M. Rosenberg

Nature's Exquisite Scalpels

For 15 years, Rosenberg and the scientists he works with at his University of Pittsburgh laboratory have studied the DNA-recognition mechanisms of a protein called Eco RI endonuclease. Eco RI is one among a class of enzymes, called *restriction enzymes*, that protect the DNA in bacterial cells by attacking DNA brought in by viruses and other foreign agents. Restriction enzymes haven't yet been isolated in human cells, but they are nevertheless of paramount practical importance because of their role in DNA analysis and genetic engineering.

"Eco RI is one of the most frequently used in what's called recombinant DNA technology," says Rosenberg, "or in the jargon 'cloning.' These enzymes recognize a particular sequence of bases of DNA and cut the DNA at those sites, breaking it into well-defined pieces that can be put back together in new combinations. Eco RI is the prototype, the first one of these enzymes to be understood and the first one used in this technology."

Largely as a result of Rosenberg's in-depth research on Eco RI, it is the paradigm example of these "exquisitely precise scalpels," as Stanford's Lubert Stryker calls restriction enzymes (in his textbook *Biochemistry*). "With Eco RI," explains Rosenberg, "we're trying to really get at the basic recognition process. One of the most intriguing questions in molecular biology today is whether general recognition principles will emerge from the details of these individual recognition mechanisms."

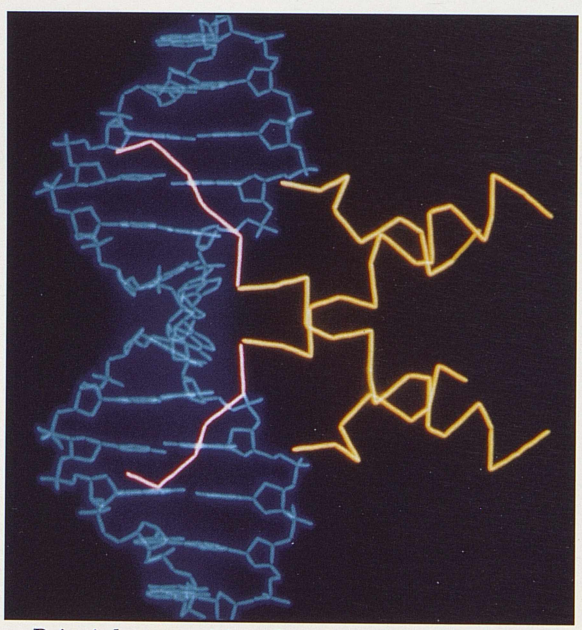

Principle recognition elements of Eco RI (pink & yellow) positioned in the "kinked" GAATTC site of DNA (blue).

The CRAY Biology Lab

The two helical strands that comprise DNA's "backbone" are linked together by bonded pairs of "bases" — four chemicals, guanine, cytosine, adenine and thymine (abbreviated G, C, A and T). The sequence of these bases define the basic genetic properties built into DNA. Like most of the many proteins involved in DNA processes, Eco RI recognizes a specific sequence — GAATTC in this case, and it attaches to DNA to do its surgical duty only at these locations.

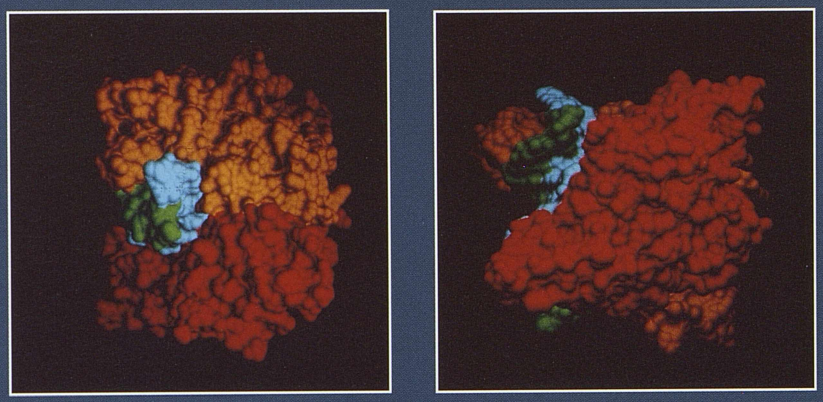

Space-filling model of Eco RI endonuclease bound to its target DNA, top-view (looking down the DNA axis) and front view. One subunit of ECO RI is shown in red, the other orange; the two helical strands of DNA are shown in green and blue. The twofold symmetry of the enzyme corresponds to the twofold symmetry of DNA's double helix.

Rosenberg and his graduate students, now Ph.Ds, John Grable and Yongchang Kim used X-PLOR, an efficient program for refining the structure of biological molecules (developed by Axel Brünger of Yale), to help identify the mechanisms by which Eco RI recognizes the GAATTC sequence. These computations resulted in clarifications to the structural model of Eco RI (published in *Science*). "That's a very involved calculation," says Rosenberg. "Just about all our crystallographic refinement was done on the CRAY. We couldn't have done it without supercomputing power."

This study elaborated on Rosenberg's earlier work showing that Eco RI's structure of paired identical subunits conforms closely to the structural pattern of a number of other proteins that bind to "nucleotide" chains. A series of folds allows the protein to wedge into the large groove, the so-called "major groove" between the two DNA strands and spread them apart. This gives better access to the base-pairs, located at the floor of the groove. The protein wraps almost completely around the DNA, as if embracing it with extended arms, and at the same time kinks the DNA at the center of the GAATTC sequence .

Rosenberg believes that this folded structure, called the nucleotide-binding fold, may be one of the keys to understanding protein-DNA recognition. "This architecture is connected very deeply to how proteins recognize nucleic acids. In an evolutionary sense, it's one of the ancient patterns."

Large-Scale Molecular Dynamics

Rosenberg has also used the Y-MP to do molecular dynamics simulations of DNA. He and recent physics Ph.D. Shankar Kumar collaborated with Peter Kollman of the University of California at San Francisco, creator of the AMBER molecular dynamics program. Kumar this year received a computational science fellowship from the National Science Foundation Division of Advanced Scientific Computing to support his work on this project. The researchers designed their simulations of the kinked DNA to examine whether the kink was intrinsic to the DNA or was caused by binding with Eco RI.

"It's a very pronounced kink," says Rosenberg, "and the question is what does it mean? We set up two hypothetical models, the first one being that the kink represents a structural form of DNA not previously identified. The other model, what we call 'molecular strain,' is that the protein pushed the DNA into a distorted state, like a set of compressed coiled springs. We were definitely leaning toward the first model, a new structural form of DNA." The computation simulated the kinked DNA by itself (not bound to Eco RI) to see if the kinked structure was stable. The answer: not stable at all. "The computations gave us a very definite answer, that the kink results from binding with the protein, and this leads to a whole series of ramifications about how the system works."

In other simulations with AMBER, Rosenberg and Kumar have stretched the limits of computational biology. Their simulations of DNA in solution (with 1970 water molecules for 81 picoseconds) represent one of the most extensive molecular dynamics computations on DNA molecules to date.

References:

G.A. Thomas, W.L. Kubasek, W.L. Peticolas, P.J. Greene, J.C. Grable and John M. Rosenberg, "Environmentally Induced Conformational Changes in B-Type DNA: Comparison of the Conformation of the Oligonucleotide d(TCGCGAATTCGCG) in Solution and Its Crystalline Complex in Restriction Nuclease Eco RI." *Biochemistry* **28**, 2001 (1989).

Y. Kim, J.C. Grable, R. Love, P.J. Green and John M. Rosenberg, "Refinement of Eco RI Endonuclease Crystal Structure: A Revised Protein Chain Tracing." Technical Comments, *Science* **249**, 1307 (1990).

This research is supported by the National Institutes of Health.

Good Vibrations

Dynamical Structure of Nucleic Acids and Proteins
David L. Beveridge, Wesleyan University

"The detail of the pattern is movement."
— T. S. Eliot, "Burnt Norton"

Molecules in Motion

You might say that biological molecules are the original workaholics. Never mind a vacation, these large molecules, the proteins and DNA essential to life, don't even take a break. Even when they look like they're standing still, they're moving, actually vibrating in place.

That's why the ball and stick models of protein and DNA structure in textbooks and articles can be misleading. They usually represent a time-averaged form of the molecule, essentially a freeze-frame snapshot. In its natural environment, as a functioning component in a living system, a biological molecule is flexible — a little like a loose-jointed, swivel-hipped dancer, and it fluctuates at an extremely rapid pace, shifting and changing from one picosecond (a trillionth of a second) to the next.

To better understand how proteins and DNA carry out their biological tasks, says David Beveridge, we have to understand these intricate movements: "The conventional view of a molecule is static, with atoms in a fixed geometrical relation to each other, like tinker toys. This works pretty well in thinking about the chemistry of small molecules. When you get to macromolecules, however, systems the size of DNA and proteins, some regions are very flexible. Here the concept of structure must be expanded to consider explicitly the idea of dynamical motion."

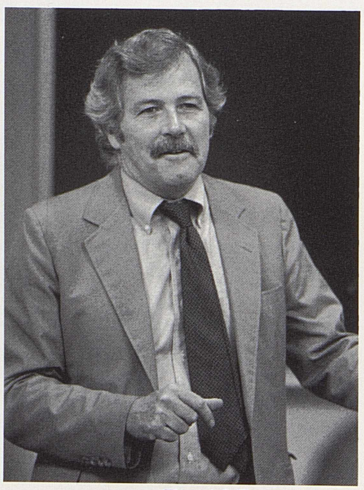

David Beveridge

Dynamical Simulations of DNA

Beveridge and his research colleagues at Wesleyan have been using the CRAY Y-MP at the Pittsburgh Supercomputing Center to address this problem. Using molecular dynamics (MD) code developed by S. Swaminathan, they have focussed on a single turn of the DNA double helix (a sequence containing the recognition site for Eco RI endonuclease, see previous article).

Their computations simulate how the DNA double helix moves in its watery environment, interacting with sodium ions and two layers of surrounding water (over 2000 water molecules). Their most recent calculation traced the fluctuations of this 12 basepair sequence over a timespan of 480 picoseconds, computations which are believed to be the most extensive molecular dynamics studies on DNA to date.

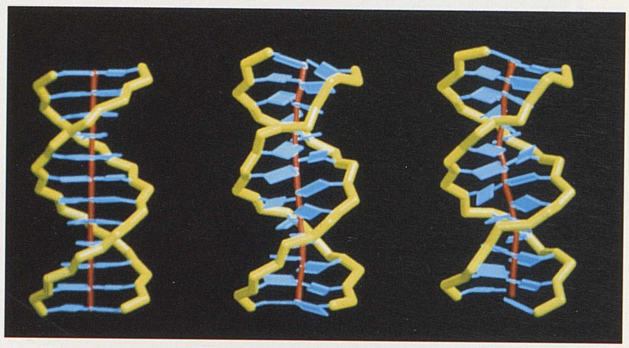

Three snapshots from an MD simulation of DNA with counterions and water (not shown). DNA's sugar-phosphate backbone (yellow) connects to nucleotide base-pairs (blue) that join along a central axis (red).

Another member of the Wesleyan research group, G. Ravishanker, has developed a comprehensive set of programs, the MD Toolchest, for analyzing the results of MD simulations. Using these tools, the group found deformations in the DNA helical axis that correspond closely to the DNA crystal structure solved several years ago by Richard Dickerson and colleagues at UCLA. This indicates that a propensity toward bending is intrinsic to the DNA, not merely a result of crystallizing the molecule.

In a related study, Jane Withka and Philip Bolton compared the results of the MD simulations to DNA structural data they obtained from nuclear magnetic resonance (NMR) spectroscopy. Using the MD data, Withka calculated nuclear Overhauser effects for all the protons in the 12 basepair sequence. She found that the MD structure agreed much better with the NMR results than with either the classic Watson-Crick structure, the so-called canonical B form of DNA, or the crystal structure alone. "This study is a step," says Beveridge, "toward understanding the structure of DNA in solution."

Domain Communication in a Key AIDS Enzyme

In other recent work, Beveridge, Swaminathan and Bill Harte from the computer-assisted drug design group at Bristol-Myers Squibb performed molecular dynamics simulation of an enzyme, HIV-1 protease, that plays a key role in reproducing the AIDS virus. They simulated the dynamics of this protein, composed of 198 amino acids, surrounded by 6000 water molecules, using nearly 100 hours of CRAY Y-MP time in the process. "We couldn't even have attempted this simulation," says Beveridge, "without the availability of the supercomputing center." The calculation gave an accurate account of the crystal structure of HIV-1, and the results are being used at Bristol-Myers Squibb in research aimed at AIDS drug therapy.

The researchers also applied a new technique, called a dynamical cross-correlation matrix, to analyze how the motions of one part of the HIV-1 enzyme correlate with that of other parts. The results revealed an unexpected teeter-totter effect, an up-down relation between two distinct domains widely separated in space.

"The idea of through-space interaction in proteins," says Beveridge, "has broad implications that may be related to enzyme mechanisms and protein folding. You would never know about this without supercomputer level molecular dynamics simulations. No experiment can ever give you all this information."

References:

S. Swaminathan, G. Ravishanker & D.L. Beveridge, "Molecular Dynamics of B-DNA Including Water and Counterions: A 140 psec Trajectory for d(CGCGAATTCGCG) Based on the GROMOS Force Field." J. Am. Chem. Soc. 114, 5027 (1991).

J. M. Withka, S. Swaminathan, D. L. Beveridge & P. H. Bolton, "Time Dependence of Nuclear Overhauser Effects from Molecular Dynamics Simulations," J. Am. Chem. Soc. 114, 5041 (1991).

W.E. Harte, Jr., S. Swaminathan, M.M. Mansuri, J.C. Martin, I.E. Rosenberg & D.L. Beveridge, "Domain communication in the dynamical structure of human immunodeficiency virus 1 protease." Proc. Natl. Acad. Sci. 87, 8864 (1990).

S. Swaminathan, W. E. Harte, Jr. & D. L. Beveridge, "Investigation of Domain Structure in Proteins via Molecular Dynamics Simulation: Application to HIV-1 Protease," J. Am. Chem. Soc. 113, 2712 (1990).

This research is supported by Grant GM 37909 from the National Institutes of General Medical Studies (NIH) and a Goodyear Cooperative Research and Development Grant from the State of Connecticut Department of Higher Education and Bristol-Myers Squibb.

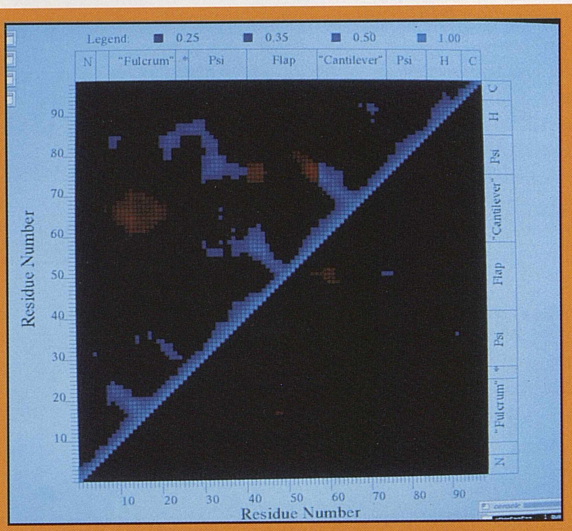

Dynamical cross-correlation matrix for HIV-1 protease. Structural subregions of the enzyme are indicated along the top and side. Residues moving in concert with one another correspond to color peaks on the matrix. Positive correlations are shown in the upper triangle; negative in the lower triangle. Blue shows correlations between residues in contact with each other, and red indicates correlations between residues separated in space, the significant finding of this study.

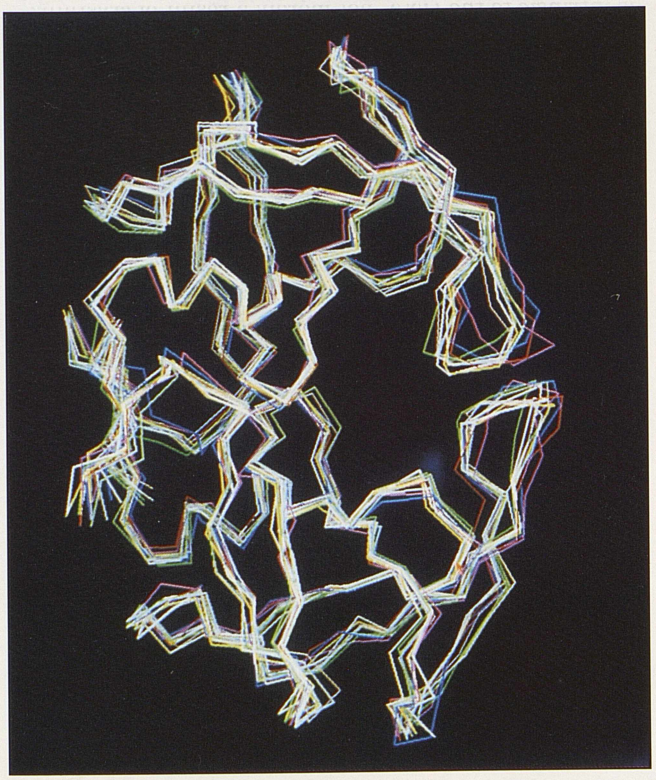

Nine snapshots from an MD simulation of HIV-1 protease are superimposed on the crystal structure (magenta).

Knifeless Biopsy

Ab Initio Simulation of Two-Dimensional Nuclear Magnetic Resonance

Manfred G. Prammer, University of Pennsylvania

Making the Body Transparent

In medical science, the 1980s was the NMR decade. The technology of nuclear magnetic resonance, developed in the 1940s, stepped out of the chemistry lab and moved into hospitals across the country, providing a powerful new capability to see inside the body. Proponents of NMR imaging describe it as the most important advance in medical science since X-rays. Using powerful magnets and radio-frequency (RF) pulses, NMR maps the distribution of certain atoms in the body, most commonly hydrogen, information that can then be translated into images that show the body's soft tissue — notably the brain, heart, liver and muscles — more clearly than X-rays. Since the RF pulses used in NMR are lower frequency than X-rays, it poses no known hazard to human tissue.

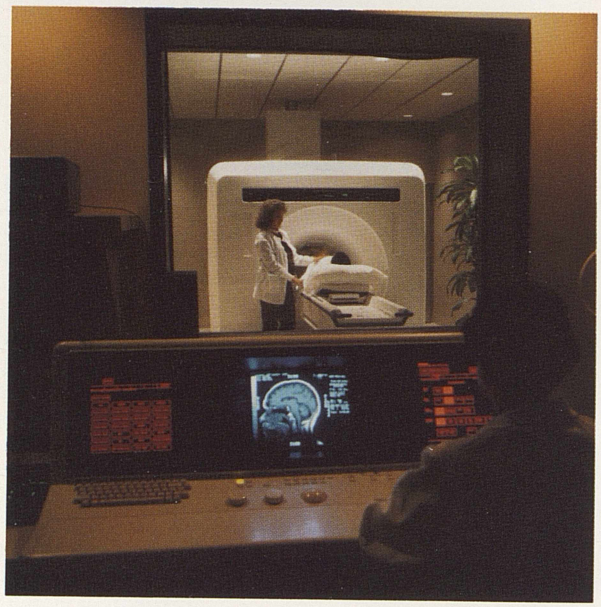

NMR imaging has caught on rapidly, but some scientists feel that a related use of the same technology, NMR spectroscopy, offers even greater potential as a diagnostic tool. NMR spectroscopy, which usually supplies graphs rather than X-ray-like images, can provide the kind of precise chemical information that normally requires laboratory analysis of tissue samples. In effect, it can be a non-invasive form of biopsy and has been described as having the capability of making the body transparent, a living museum of the chemical metabolism of life. Studies with NMR spectroscopy suggest it could be used to monitor the effect of drugs in treating heart attacks and in cancer therapy.

Overcoming Noise

Applying NMR spectroscopy to humans is far from straightforward. "We have to become very sophisticated," notes Manfred Prammer, "if we're going to fulfill the promise of a completely non-invasive diagnostic tool." Prammer, a professor in biochemistry and biophysics at the University of Pennsylvania, is working with a group of biomedical scientists on solving the problems of applying NMR spectroscopy to living organisms. To this end, he has developed a FORTRAN program for the CRAY Y-MP that numerically simulates what happens when NMR is applied to living tissue.

Prammer's code implements quantum theory with a "brute force" *ab initio* calculation of the extremely small energy transitions that occur in two-atom groups of biological molecules when they are subjected to RF pulses. Basically, this is equivalent to the changes that occur when the signal from an AM or FM radio station passes through the human body. NMR works because these signals can be adjusted to "resonate" with the nuclei of certain molecules. When applied to the body, however, the RF signals from resonating nuclei almost get lost in the background of "noise" — stray signals from the surrounding chemical environment.

To alleviate this problem, explains Prammer, researchers have developed several techniques. These include complicated modulation patterns for the RF pulses and interrogating two atomic species at once. "With these tricks of the trade, we can filter out the signals from coupled atoms, say a carbon-hydrogen group, without drowning the signal in the overwhelmingly stronger background of the hydrogen in water molecules."

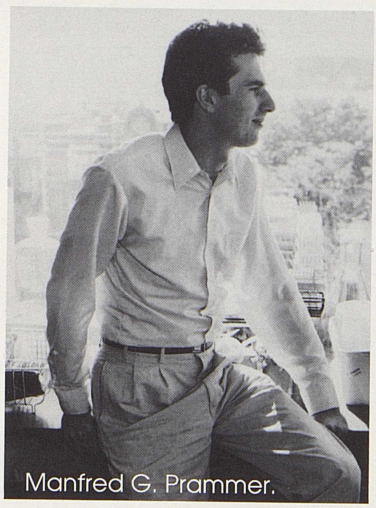

Prammer together with his research colleague, Katherine E. Vogele, were honored by Cray Research Inc. as recipients of a 1990 Gigaflop Performance Award. The award recognizes researchers worldwide who have solved practical scientific problems with programs that perform at levels above 1.5 gigaflops (1.5 billion floating-point operations per second). Prammer's code for simulating two-dimensional NMR spectroscopy runs at 2.082 gigaflops.

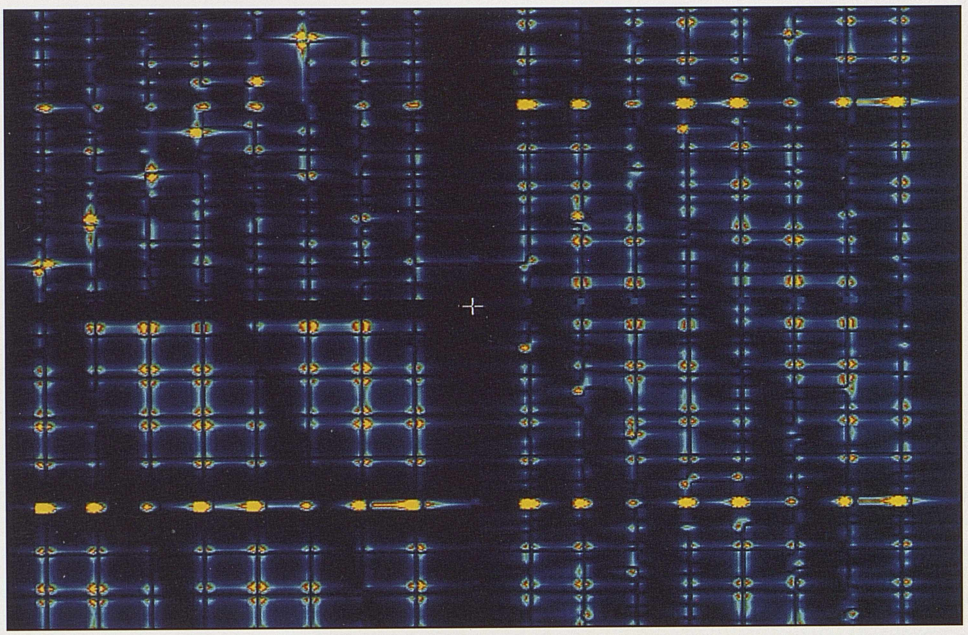

This graphic representation of a two-dimensional NMR simulation shows the response of a two-nuclei spin system to radio-frequency pulses of increasing energy. Each of the four quadrants represents a different RF pulse, with the power increased by 20% in each quadrant from upper left to lower right. "The response is highly non-linear," says Prammer, "and reflects the chemical bonds present around the spin system."

Simulating NMR at Two Gigaflops

The prototype problem that Prammer's program addresses is carbon-hydrogen bonds, a fundamental building block in DNA and proteins. The program allows researchers to analyze and fine-tune the RF pulse patterns so that NMR experiments can be set up to yield the detailed metabolic chemistry that will tell doctors if an organ is functioning the way it should.

The program, notes Prammer, is the first of its kind. "There's no direct parent code. We had several smaller programs, but we were never able to do a full simulation. We always had to take some shortcut to get results." To take maximum advantage of the vector- and parallel-processing capabilities of the CRAY Y-MP, Prammer reformulated the problem from scratch.

Running on all eight processors of the Center's Y-MP, Prammer's highly-optimized code zooms along at more than two billion floating-point operations per second (two gigaflops), close to theoretical maximum speed. One complete run does over 56 billion operations, and many runs are necessary to tune the RF signal so that an experiment will yield the information of interest. With Prammer's code, a simulation can be accomplished in minutes instead of days, making it much more valuable as a tool that can help realize the potential of NMR in diagnosing human disease.

This research is supported by the National Institutes of Health and by the James S. McDonnell Foundation.

NMR Imaging photograph: Courtesy West Penn Hospital, Pittsburgh.

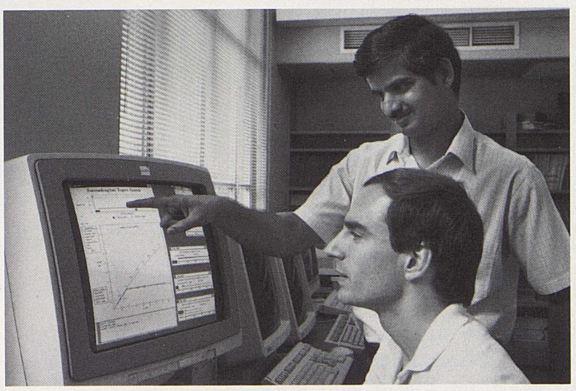

Raghurama Reddy and Jim Ellis (seated) of the Pittsburgh Supercomputing Center staff helped coordinate testing of Manfred Prammer's two gigaflop NMR code. "Since it involves matrix operations, this problem is really suited to optimizing on the Y-MP," notes Reddy, who specializes in code optimization. "To achieve this level of performance, you have to pay close attention to every detail of the problem." Vector instructions must be scheduled so that data flows in and out of memory in synchrony with the ability of the arithmetic elements to process it.

"DEDCPU [Dedicated CPU utility from UNICOS] is useful for debugging parallel code," says Ellis. "It allows us to dedicate a specific number of CPUs to a program without having to shutdown normal batch processing." By relying on four-processor runs with DEDCPU, Prammer debugged and tested his code with only two fully dedicated eight-processor runs.

Where Have All the Oysters Gone?

A Three-Dimensional Curvilinear-Grid Hydrodynamic Model for Estuaries, Lakes and Coastal Waters

Y. Peter Sheng, University of Florida

Exploding Algae and Suffocating Fish

Some summers in the Chesapeake Bay, you can see fish rushing toward the shore and leaping out of the water. It's not a spontaneous desire to evolve nor are the fish jumping for joy. They're suffocating in the bay's oxygen-depleted water.

Because of excess nutrients, nitrogen and phosphorous, that enter the bay from phosphate detergents, treated sewage, agricultural runoff and acid rain, algae in the water grow much more rapidly than usual. Biologists call this explosion in growth rate an algal bloom, and it uses up so much of the oxygen in the water that the fish, particularly bottom dwellers like crabs and oysters, don't have enough. In the '70s, Maryland fishermen caught over two million bushels of oysters a year in the bay. Currently, they catch less than half a million.

To help manage this problem, Peter Sheng is using the Center's CRAY Y-MP to develop a three-dimensional model for tracing the flow patterns in estuarine and tidal systems, sensitive areas where freshwater and saltwater ecosystems meet. Sheng has begun to apply his model to several Florida lakes and in the Chesapeake Bay. "Water management agencies want to control external loading of sewage, industry and agricultural waste, and they need a management tool to determine what level of loading reduction is necessary," says Sheng. Using a computer model of the circulation patterns in a water body, environmental agencies have a way to relate measured contaminant concentrations to pollutant sources. With this information, it's possible to initiate appropriate control strategies. "The model we're developing can provide the management tool that the agencies need."

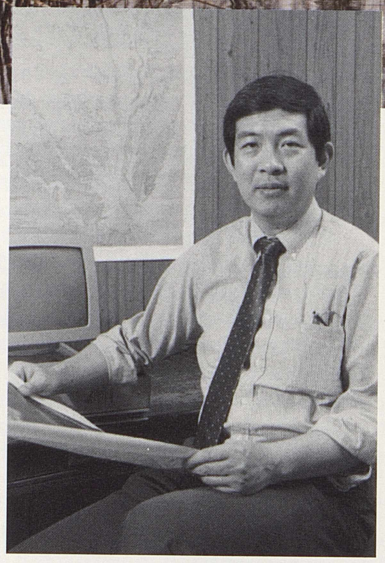

Y. Peter Sheng is a professor in the University of Florida's Coastal and Oceanographic Engineering Department.

A 3-D Curvilinear-Grid Model

To predict how contaminants circulate through an estuary or lake, scientists must first understand the flow patterns. Sheng's finite-difference, three-dimensional, curvilinear-grid hydrodynamic model — called "CH3D" — has achieved greater accuracy than prior approaches to the problem. The model solves the three-dimensional and time-dependent Navier-Stokes equations, the classic equations of fluid flow, for tens of thousands of grid points within the water body. Like different blocks that fit together to build a model house, CH3D can combine several grid types to map shoreline and bottom topography.

Because the various grids range in shape from rectangular to curved, they can map the intricate geometry of the 190-mile-long Chesapeake, with its hundreds of fingers and tributaries. Such grids also resolve bottom topography realistically so the model doesn't end up having a smooth, flat "bathtub" bottom. CH3D takes into account the ridges and gullies found on the bottom of most waters because they affect flow patterns in slight but significant ways.

Modeling the Chesapeake Bay

Circulation in estuaries like the Chesapeake is driven by wind, tide, waves, temperature changes, Earth's rotation, and concentration differentials for salt and other particles in the water as well as the water body's unique geography. A major challenge was to accurately simulate turbulent mixing in the water column. To this end, Sheng developed a simple, yet comprehensive model that accounts for tidal and wind-driven currents and particle buoyancy.

Early Chesapeake Bay simulation results indicated a migration of salt from deeper areas to adjacent shallow areas which doesn't happen in the actual bay. Sheng alleviated the problem by applying a different grid technique to troublesome deep areas and by incorporating special advective algorithms "based on state-of-the-art development in computational fluid dynamics."

In simulating wind effects on the Chesapeake, Sheng applied CH3D with 16,000 grid points and tested his computations against measured data for the month of

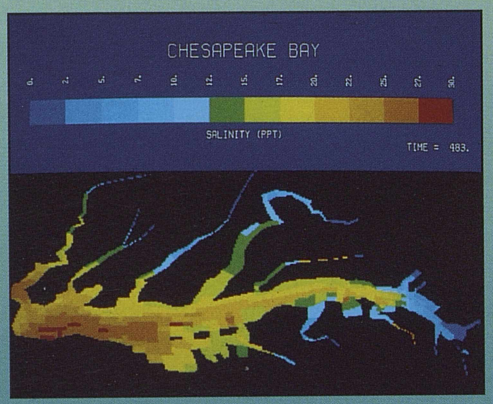

Salinity distribution in the Chesapeake Bay during September 1983 as simulated by CH3D. The color scale represents salinity in parts per thousand (ppt), increasing from blue to red. Freshwater from tributaries becomes increasingly saline as it mixes with ocean water entering at the mouth of the bay.

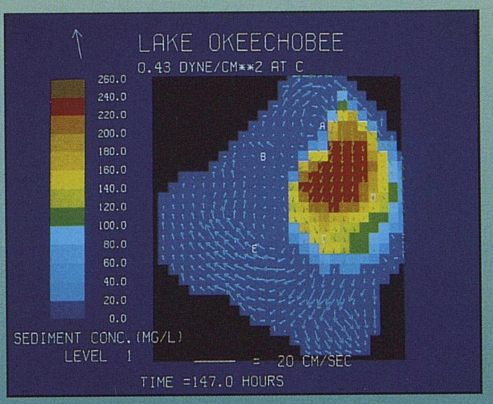

Wind-driven surface currents and suspended sediment concentration in Florida's Lake Okeechobee during June 1989 as simulated by CH3D. The color scale represents sediment concentration in milligrams per liter (mg/l), increasing from blue to red. Excess nutrients in Lake Okeechobee are a potential threat to the biological balance of the Everglades. The moving vegetation boundary of this lake has presented Sheng with a new challenge to his modeling ingenuity, one he has begun to tackle in test runs on the Center's CRAY Y-MP.

September 1983. The model used the initial salinity, temperature fields, wind forcing, atmospheric heating and cooling and ocean tide field data, and the results of the 30-day simulation were compared with water level, currents, salinity and temperature recorded at several stations in the bay. The model, in Sheng's words, "faithfully simulated" wind-induced mixing in surface waters as measured at one of several measuring stations on the bay. "This type of comparison," says Sheng, "is a very stringent test of the model's performance."

A Practical Tool for Improving Water Quality

Testing of a model like CH3D requires dozens of long-term simulation runs, and the project is only feasible on a supercomputer like the Center's CRAY Y-MP, where a month-long Chesapeake Bay simulation takes only 40 minutes. "Without vectorization and other supercomputer features," notes Sheng, "it's virtually impossible to perform long-term simulations within any reasonable amount of time."

Sheng's model is now undergoing additional development and testing by the Chesapeake Bay Modeling Study, a group project involving EPA, environmental agencies of surrounding states and the U.S. Army Engineers. The Chesapeake Bay group will use CH3D as a water-quality management tool. "With the model, you can calculate the flow, and then you can add on the water-quality parameters to determine how much to reduce the nutrient inflow," Sheng explains. "Then the model can tell you how long a system will take to respond to reduced nutrient levels, and whether you've achieved the desired level of water quality."

References:
Y. Peter Sheng, "Modeling Wind-Induced Mixing and Transport in Estuaries and Lakes," *Estuarine Water Quality Management* (W. Michaelis, ed.), Springer-Verlag (1989).
Y. Peter Sheng, "Evolution of a Three-Dimensional Curvilinear-Grid Hydrodynamic Model for Estuaries, Lakes and Coastal Waters: CH3D," *Estuarine and Coastal Modeling*, ASCE (1990).

This work is supported by the National Science Foundation, the U.S. Army Engineer Waterways Experiment Station and the South Florida Water Management District.

Photograph of Lake Okeechobee: courtesy Florida Department of Commerce, Division of Tourism.

Solar Winds and Magnetic Fields

A Study of the Solar Wind-Magnetosphere Interaction at the Dayside and Nightside Magnetopause

Lou-Chuang Lee & Joseph G. Hawkins,
University of Alaska Fairbanks

"Like a streamer of the northern morn,
Seen where the moving isles of winter shock
By night with noises of the northern sea."
— Tennyson

Hot Gas & Northern Lights

Like fireworks over the arctic circle, only ghostly silent and more beautiful than fireworks, a shimmering, swirling curtain of iridescence — the northern lights is one of the more dazzling special effects mother nature has devised. Scientists at the University of Alaska Fairbanks are well situated to study this phenomenon, *aurora borealis* (northern dawn) as it is formally named for the Roman goddess of dawn. With a three-year grant from NASA's space physics theory program, physicist Lou-Chuang Lee and electrical engineer Joseph Hawkins are using the supercomputing resources at Pittsburgh to investigate the complex electromagnetic processes that generate these majestic displays.

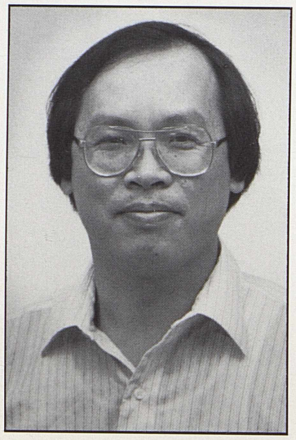

Lou Lee

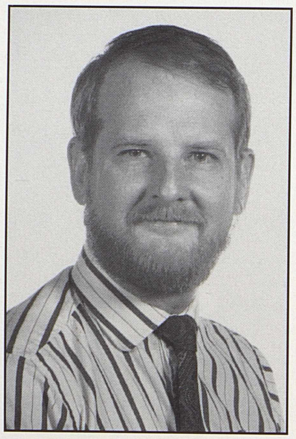

Joe Hawkins

Since the 1960s, satellites and interplanetary probes like *Mariner* have added a great deal to our knowledge of the northern lights and its south polar counterpart, *aurora australis*. It is now believed that the auroras are a visible result of energy transfer from the so-called "solar wind" as it blows into, around and through the magnetic field surrounding Earth.

Somewhat like a gigantic blast furnace (except unimaginably hotter than any blast furnace), the sun continuously spews gusts of ionized gas or "plasma" from its upper atmosphere, the corona, outward into space. Often thought of as a fourth state of matter (along with the more familiar solid, liquid and gas), plasmas form when gas is so hot that electrons, in effect, boil off from atoms, creating a diffuse mass of charged particles (electrons and ions).

"Auroral activity," says Hawkins, "is similar to what happens in your TV set. Electrons are accelerated and guided into Earth's atmosphere along magnetic field lines. As they pound into the atmosphere, they cause illumination."

A Beautiful Menace

The auroras, while beautiful, also pose serious problems. They occur as electromagnetic storms, related to cycles in solar activity, that generate surges of power of a million watts or more in Earth's upper atmosphere. These auroral substorms, as they are called, can disrupt radio and satellite communication, power transmission and defense systems. Scientists like Lee and Hawkins want, ultimately, to be able to predict when and how intensely these substorms occur — which is where supercomputing enters the picture.

"In space physics," says Lee, "computer simulation is extremely important." Satellites give limited information, he explains, since they can "see" only isolated points at any one time. "You can't get a global picture, even if you have several satellites." Because the processes are so huge, involving the transfer of tremendous amounts of energy, mass and momentum over vast distances of space, laboratory experiments don't help either. "A lot of speculation is involved in space research. We have to carry out computer simulations to show that these things can really occur the way we think they do, from what the observations suggest."

Wind in the Field: Magnetic Reconnection

Lee and Hawkins' work focuses on what happens at the meeting place where the solar wind, blowing at supersonic speeds, drives into the outer reaches of Earth's magnetic field. The magnetic field presses back and, in effect, carves out a long, blunt-nosed bubble in the solar wind — Earth's "magnetosphere." At the solar wind-magnetosphere boundary, the "magnetopause," some of the plasma breaks free from the deflected flow and enters the magnetosphere. It is this transfer of energy that ultimately generates the auroras.

While the large-scale outlines are fairly well understood, scientists still know relatively little about exactly how this transfer occurs. For a number of years, Lee has worked on computer simulations of the most prominent theory — magnetic reconnection. The solar wind carries a magnetic field of its own, called the interplanetary magnetic field (IMF). When the IMF at the magnetopause is directed southward, "anti-parallel" to the

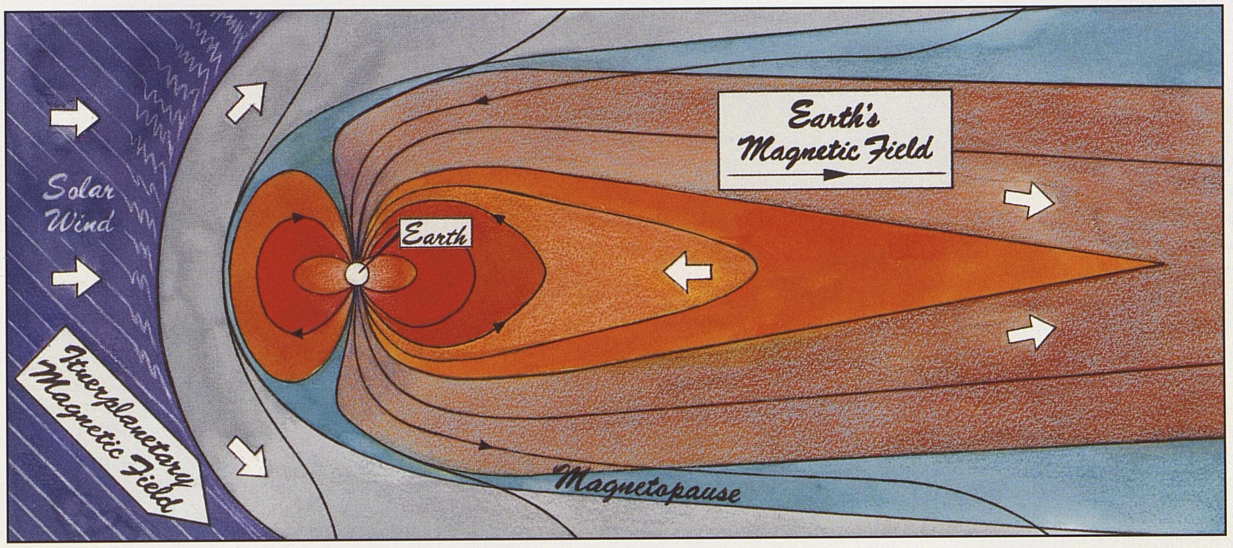

The solar wind compresses the sunward end (dayside) of Earth's magnetosphere into a blunt nose that extends about five Earth diameters out in space. Most of the solar wind streams around the magnetosphere, which on Earth's nightside stretches into a comet-like magnetotail millions of miles long. Solar plasma that enters the magnetosphere forms a sheet along the mid-line of the magnetotail.

northward orientation of the magnetosphere field, the opposed field lines can break, cross over and reconnect, making a pathway for solar plasma to flow into the magnetosphere.

When scientists first proposed this magnetic reconnection model for the solar wind-magnetosphere interaction, in the 1960s, the model included only a single "x-line" or point of reconnection. Lee's work has extended this relatively simple model to include multiple x-lines. In current work, he hopes to account for other important factors that, due primarily to computing limitations, have been beyond the scope of prior modeling.

"The field-line topology involved in magnetic reconnection," says Lee, "is inherently three-dimensional. It's much more complicated than people believed 10 years ago. We need a highly capable system to do this kind of simulation."

Impulsive Penetration

Lee and Hawkins are also using Pittsburgh's CRAY Y-MP to simulate another proposed explanation of how solar plasma can cross the magnetopause. "If solar wind pressure is high enough in a small area, due to a localized increase in plasma density," says Hawkins, "this irregularity could have enough momentum to pop through the magnetopause, and the magnetospheric field lines would then have to reconnect in such a way that an isolated blob of plasma could be stranded in the magnetosphere." This idea, impulsive plasma penetration, is relatively new (first proposed in the late 1970s), and many researchers in the field remain unconvinced.

Lee and Hawkins' recent simulations in two dimensions provide a new way of analyzing impulsive plasma penetration and indicate that under just the right conditions it can occur. They plan to further test the possibility with 3-D computations, and they believe that these more realistic models will show a stronger likelihood that impulsive penetration is at least one part of the complicated story of what happens when the solar wind blows into Earth's magnetic field.

References:

J.G. Hawkins, Z. W. Ma & L.C. Lee, "Impulsive Penetration at the Dayside Magnetopause," *Physics of Space Plasmas,* SPI conference proceedings and reprint series, No. 10, MIT (1990).

L. C. Lee, "Toward a Time-Dependent Magnetic Reconnection Model," *Eos* **69**, 1617 (1988).

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

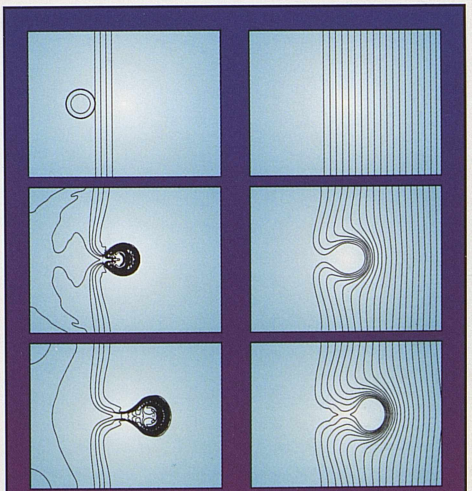

Time evolution of plasma density (left) and magnetic field topology (right) during impulsive penetration, as simulated in two-dimensions by Lee and Hawkins.

Measuring the Misery Index

New Evidence on Whether Unemployment and Out of the Labor Force are Distinct States

Füsun Gönül, Carnegie Mellon University

"A man loses his job and he loses himself like he's beat with a stick."

— Studs Terkel

Unemployed or Out of the Labor Force?

How many people who want jobs don't have them? It's probably our most basic yardstick of how well the economy functions. For nearly everyone, a job means food, clothing, shelter and self-esteem, and the reported unemployment rate, nationally and locally, is often front page news. If it's high, we expect economists to have ideas on how to stimulate the economy, and we expect elected officials and business leaders to act on these ideas.

The reported figures are compiled from surveys, and though it seems simple enough, the question of whether a person is unemployed is wrapped in ambiguity. As defined in federal law, unemployment is a narrower category than joblessness. To be unemployed, a person must be available for work and searching for a job at least four weeks prior to the survey week (or on layoff and waiting to be recalled or waiting to report to a new job within 30 days). A jobless person not searching for work is "out of the labor force" (OLF) and isn't counted as unemployed.

Are the two states really as distinct as the law implies? Many social service workers claim that reported figures understate the true dimension of the problem by ignoring "discouraged workers," people out of work, often for a long time, who for whatever reason — lack of economic or psychological resources — stop searching but nevertheless need a job no less than other jobless people. Some surveys suggest that jobless people often stop searching for a period of time, though still interested in a job, and then find work later — in effect going from OLF to employment without being "unemployed" as the government defines it.

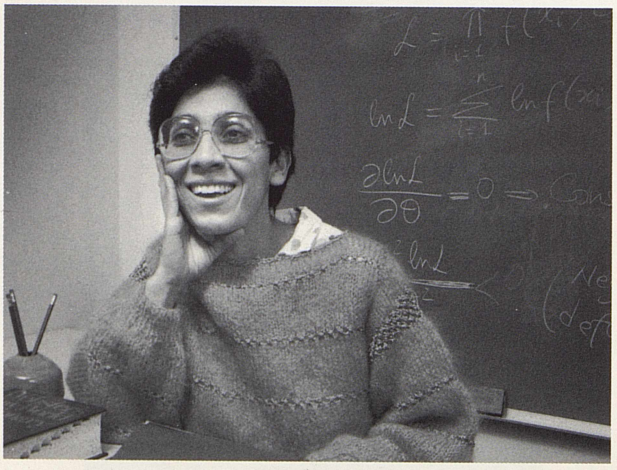

Füsun Gönül, Carnegie Mellon University, Graduate School of Industrial Administration. "In social science, we tend to have more data to process than in natural science, where the research is done in controlled laboratory settings. But not only do we have more data, we have more unobservables to control for because our 'experiments' are from real lives of real people."

"If the two nonemployment states are really the same," says Füsun Gönül, "then official unemployment statistics are misleading, and unemployment compensation schemes, which are based on the presumption of a distinct job search state, may be misguided." Gönül, an economics professor in Carnegie Mellon's Graduate School of Industrial Administration, has approached the question from the point of view of statistical analysis. Using a dataset from the National Longitudinal Surveys of Labor Market Experience (NLS), considered unique in its detailed job-history records, Gönül has compared the probability of getting a job when unemployed to getting a job when OLF. Using the CRAY Y-MP, she was able to include more data and more parameters — in effect, to do a more sophisticated analysis — than has been possible before.

A New Approach to the Problem

Despite the policy import of this research, only two prior studies have dealt with the question. "The subject has been avoided," says Gönül, "primarily due to incomplete data." The NLS survey is more complete than many employment surveys in that it records whether and how long the respondent searched for work during a jobless spell. In Gönül's view, however, this information isn't enough to do a statistically reliable assessment because it leaves out the sequence of transitions between states.

Two Examples of Transitions Between Labor Force States

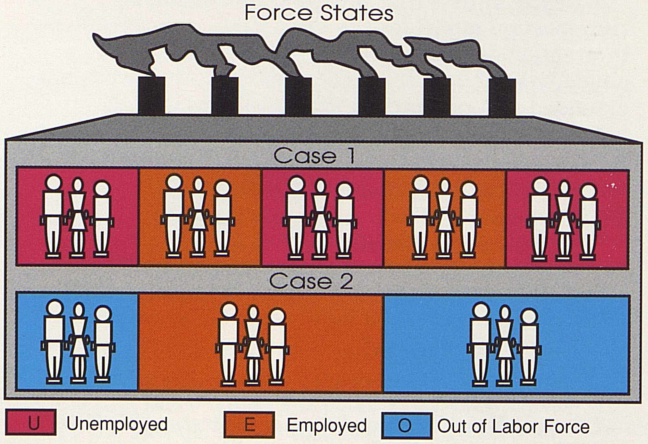

The diagram represents two cases for which the total duration of joblessness is the same ($U_1 + U_2 + U_3 = O_1 + O_2$), but the amount of time in a state and the number of state-to-state transitions are different. "Simply comparing the mean duration of the two states [U and O]," says Gönül, "is static and therefore inaccurate. Both cases would yield the same mean, which would be misleading. Mean duration is one aspect of a two-sided problem. The second aspect is the number of exits between states." An earlier study with the NLS data solved the problem of missing sequence data by including only samples who searched either for the entirety of a jobless spell or not at all. This is a less than desirable

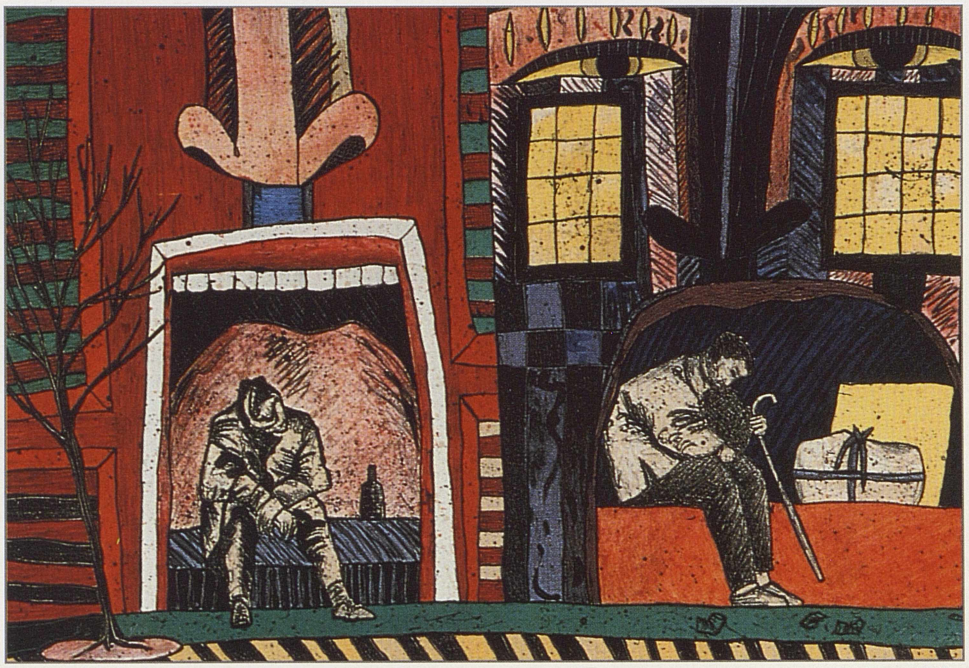

© *Elly Simmons 87*

approach, notes Gönül, because it eliminates all U to O and O to U transitions. "Deleting this data cuts the sample in half and potentially biases the results."

With access to the data-crunching capability of the CRAY Y-MP, Gönül devised an approach that allowed her to perform a "maximum likelihood analysis" using all of the NLS data. Generally considered to be a highly reliable statistical method, maximum likelihood analysis is a complex, computationally demanding technique to "fit" data to a probability distribution. For each sample lacking the sequence of U and O states, Gönül enumerated all the possible combinations of sequences and summed them, thereby arriving at the "likelihood expression" for each jobless spell. The CRAY then maximized this likelihood expression for all of the observable parameters — age, sex, race, father's education and others (24 in total) — associated with each individual in the survey.

Lifesaver Supercomputing

"There is no substitute for supercomputers when it comes to estimating a likelihood expression," says Gönül. "This is a large data set, over 300 people with weekly labor histories for a five-year period, and the function is nonlinear, which means the computation is iterative and very time-consuming. Analytical solutions are impossible given the complex nonlinearity of the problem.

"What makes this function even more complicated is that I take into account unobserved differences between people. Things like looks, personality and speaking ability are not measurable traits, yet we all know they affect the probability that one person will get a job rather than another." Gönül treats unobservable factors as a "noise" term, giving it a distribution across the entire sample. "I integrate this complicated function, the likelihood function, across all possible values of that noise." Asked how much time she saved using the CRAY Y-MP, Gönül laughed: "My whole life, I'm sure."

Results and Conclusions

The computations indicate that for men, unemployment and OLF are nearly indistinguishable. For women, on the other hand, the two states appear to be distinct. "For men," says Gönül, "it's not that they lack the ability to distinguish between the two states, it's that they don't behave measurably differently. Sometimes they actively search and sometimes they don't, and it isn't clear when they're in one or the other state."

Gönül believes that the difference she finds in the jobless behavior of men and women is a result of cultural factors. "Women have been trained to stay at home and care for children, and when they're at home, they're more likely not to be searching for work. When they put themselves in the job market, they know it; they've arranged their babysitters, settled the size of their family, et cetera."

Gönül's study suggests that policy makers should review the structure of unemployment compensation and official unemployment statistics. As they apply to men, existing standards appear to lack a rational basis. "It sounds sexist, but the data show that men and women act differently. We can assume they're equal, but by disregarding these differences, we're not being scientific."

References:
Füsun Gönül, "New Evidence on Whether Unemployment and Out of the Labor Force are Distinct States." *Journal of Human Resources*, May 24, 1990.

This research is supported by the National Science Foundation.

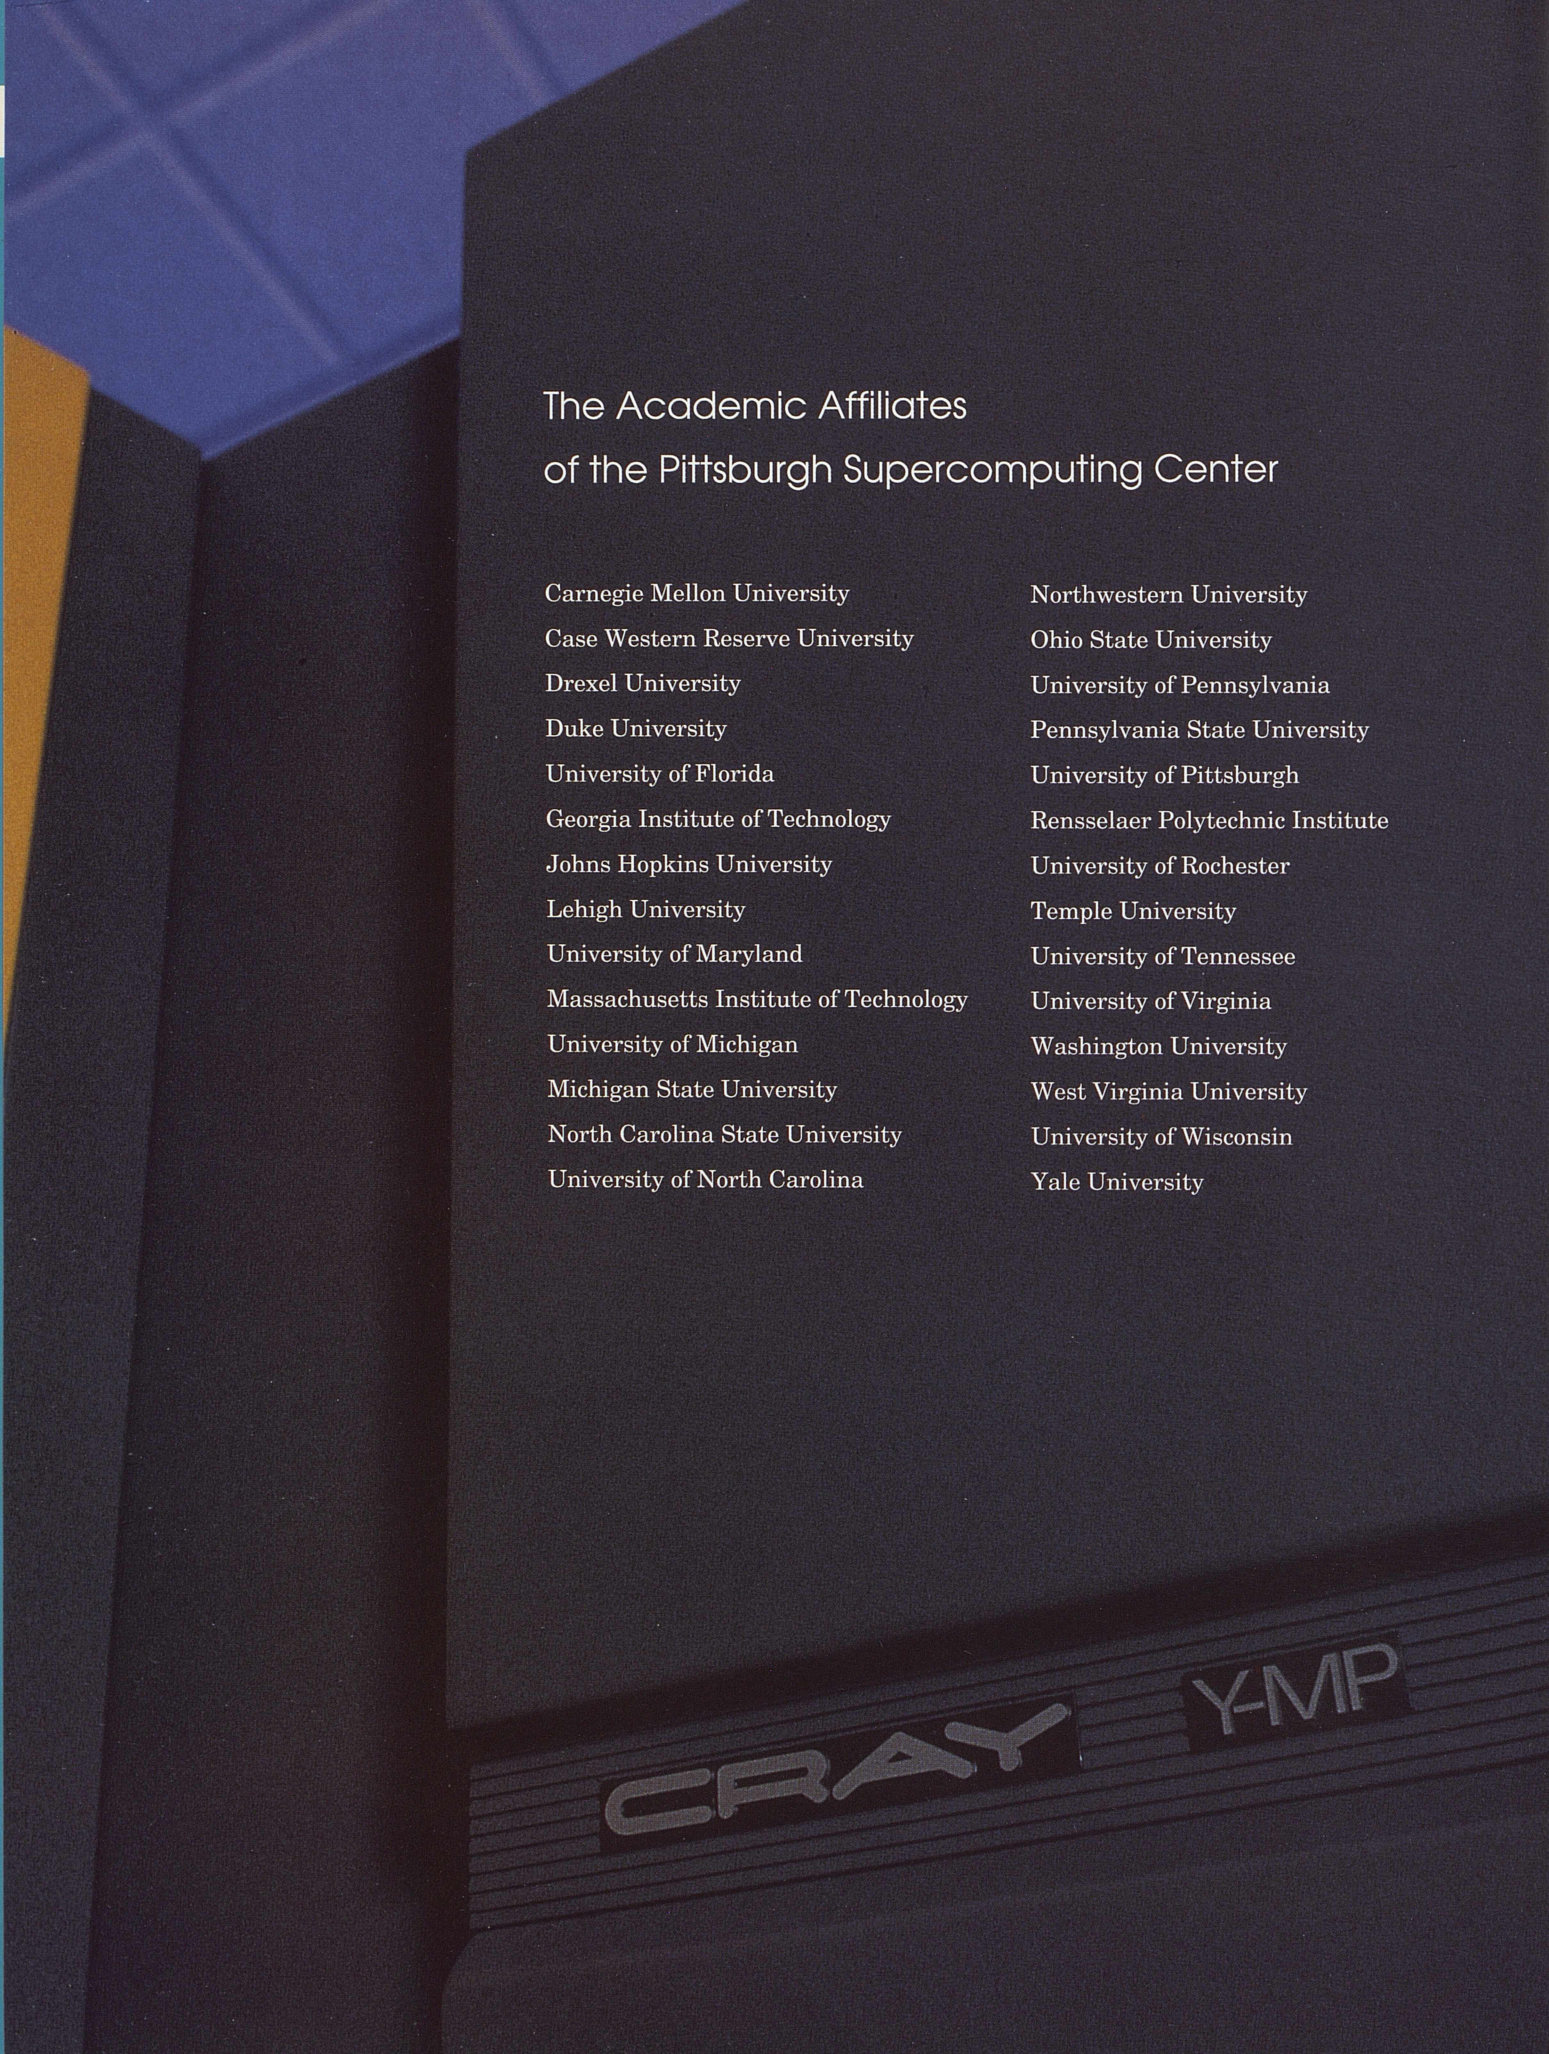

The Academic Affiliates of the Pittsburgh Supercomputing Center

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