PROJECTS IN SCIENTIFIC COMPUTING

1989-90

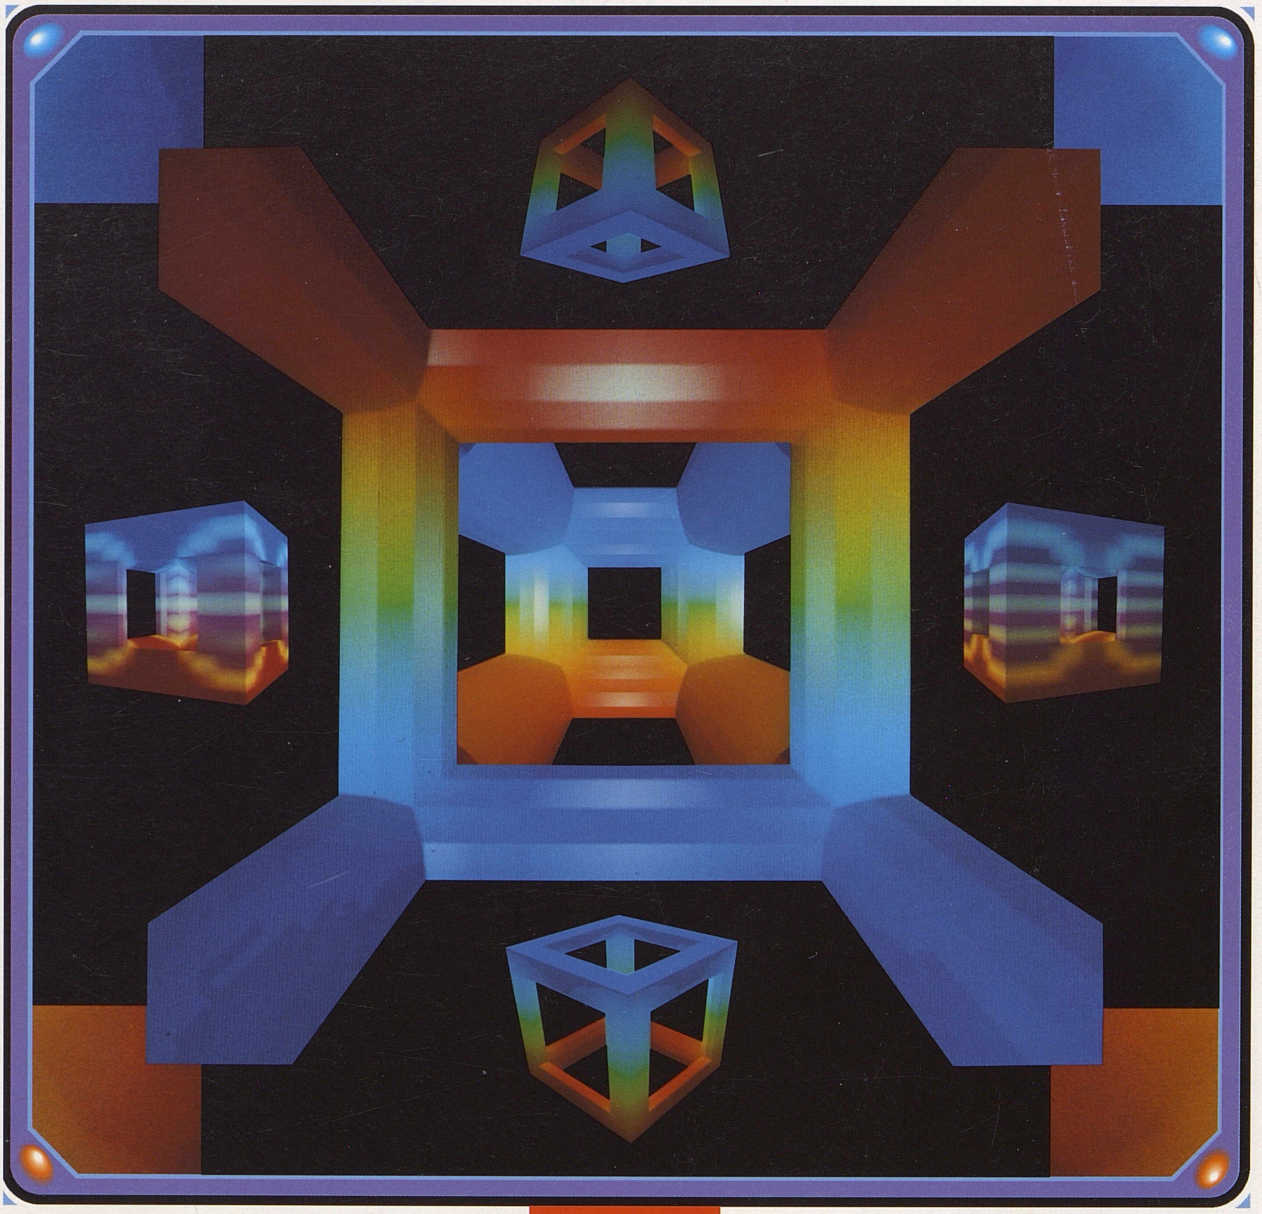

PITTSBURGH SUPERCOMPUTING CENTER

Projects in Scientific Computing

1989-90

Editor:
Mike Schneider
Pittsburgh Supercomputing Center

Writing:
Mike Schneider
Tom Anfuso
Karen Carnathan
Janet Vietmeier

Research Assistant:
Allison Robinson

Editorial Assistant:
Megan Irvine

Color Separation & Typesetting:
Digicon Graphics, Inc.

Printing:
Broudy Printing Inc.

Design:
Steve Palmer
Pittsburgh

Purchasing:
Vivian Benton

Illustrations:
Steve Palmer (17, 23, 25, 27, 28, 29, 36)
Fred Carlson (10, 23, 30)
Bruce Cridlebaugh (PostScript art: 2, 3, 4, 8, 9, 12, 13, 21, 22, 25, 27, 28, 33, 35, 36, 37)
Brian Seifert (13)
Trish Wilds (25)

Photography:
Ken Andreyo (7, 22, 35)
David Ashkenas (5, 22)
Michael Haritan (6, 7)
Megan Irvine (6, 7, 10, 28, 32, 38)
Jerry Voron (5)
Larry Rippel (4, 5, 7)
Jim Schafer (1, 5, 6, 15)

Thanks also to the helpful staff at Mellon Institute Library and especially to the scientists and engineers who, along with doing the research, gave generously of their time to help prepare the articles in this booklet.

The Pittsburgh Supercomputing Center is a joint project of Carnegie Mellon University, the University of Pittsburgh and Westinghouse Electric Corporation. It has received significant support from:

The National Science Foundation
The Commonwealth of Pennsylvania
The National Institutes of Health
The Ben Franklin Technology Center of Western Pennsylvania
Cray Research, Incorporated
Digital Equipment Corporation
Harris Corporation
International Business Machines Corporation
Stardent Computer Inc.
Storage Technology Corporation
Thinking Machines Corporation
UltraNetwork Technologies
The Defense Advanced Research Projects Agency
Bell of Pennsylvania

Published July 1990 by the Pittsburgh Supercomputing Center
4400 Fifth Avenue
Pittsburgh, PA 15213
(412) 268-4960

Copyright © 1990, Pittsburgh Supercomputing Center. All rights reserved.

Cover: A composite, constructed by Steve Palmer and Bruce Cridlebaugh, using 3-D images of idealized ice crystals developed by Pat Burns and Mark Christon, Colorado State University, from computations at the Pittsburgh Supercomputing Center.

Foreword

This has been an exciting year for the Pittsburgh Supercomputing Center and for high-performance computing in general. In August 1989, the National Science Board renewed support for the Pittsburgh Supercomputing Center for the next five years. The National Science Foundation's vision of the potential contribution of supercomputing to scientific research was confirmed in September when Dr. Allan Bromley, the President's Science Advisor, issued *The Federal High Performance Computing Program*. This report underscores the national importance of supercomputing and indicates the Bush administration's commitment to expand efforts in this area.

In November, Professor Gregory McRae, a major user at our Center, was honored by the Society for Industrial and Applied Mathematics as the first recipient of the Forefronts in Computational Science award. His work on air quality in Los Angeles is one of the more notable examples of how supercomputing can help solve important national problems. The economic relevance of supercomputing also came to the fore this year. USX, the largest producer of American steel, and Pfizer Inc, an international pharmaceutical firm, decided that affiliating with the PSC would help them to achieve their corporate goals.

This year also witnessed evolution in the paradigm of supercomputing, with increasing recognition that distributing computing tasks in a heterogeneous environment offers great possibilities for advancing scientific research. The PSC took major steps in this direction by acquiring a Connection Machine, installing a gigabit network, improving its visualization hardware and software, developing software that helps users manage this more complex environment and undertaking research that promises to solve severe problems in file servers and archiving.

The scientific and engineering results of supercomputing are the bottom-line. This booklet, which features a small sample of the projects underway at the Pittsburgh Supercomputing Center, illustrates the diversity of disciplines that benefit from access to supercomputing. At this juncture, four years into the business of providing crucial supercomputing resources to the nation's research scientists and engineers, it is gratifying that the importance of high-performance computing is being broadly recognized, and that the Pittsburgh Supercomputing Center is helping to define what it will mean in the 1990s.

Mike Levine

Michael J. Levine, Scientific Director

Ralph Roskies

Ralph Z. Roskies, Scientific Director

Ralph Roskies (left) and Mike Levine in the machine room at Westinghouse Energy Center, site of the Pittsburgh Supercomputing Center's CRAY Y-MP/832.

Projects in Scientific Computing

1989-9

Contents

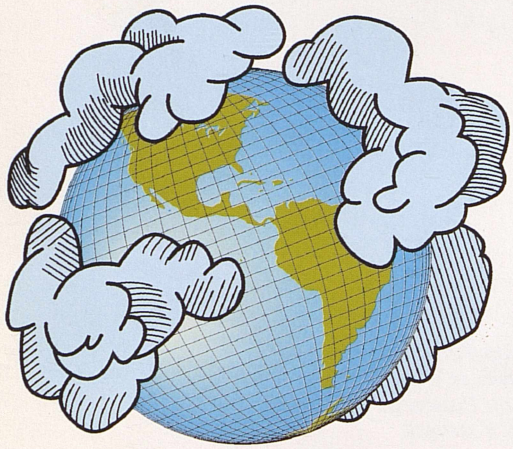

Last year (1988-89) the Pittsburgh Supercomputing Center took some vigorous steps forward in its effort to support the high-performance computing needs of the national research community: a new operating system (UNICOS) and a new supercomputer (the first CRAY Y-MP in the country at a non-government site). These changes can now be viewed as the culmination of the first phase of the Center's operation. In July 1989 the National Science Board approved a grant supporting plans for the next five years, and this year the Center began to implement these plans—several major projects directed toward developing a *heterogeneous computing environment.*

PITTSBURGH SUPERCOMPUTING CENTER 1990

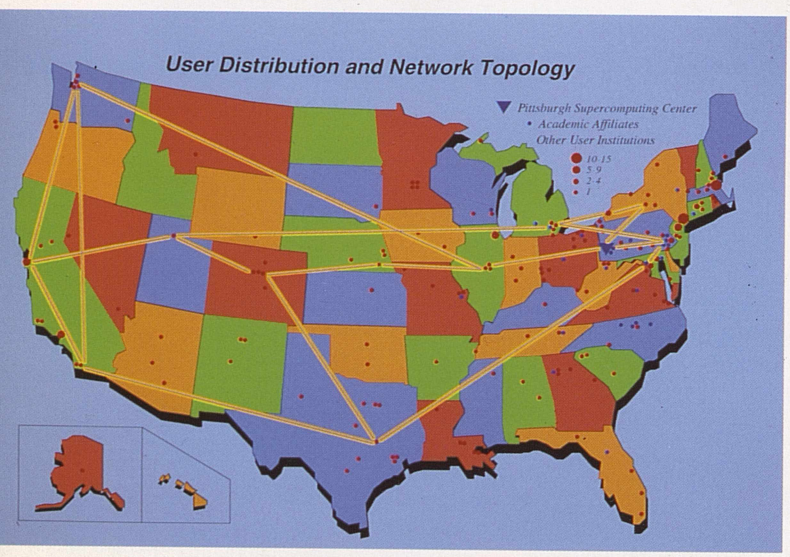

In 1986 a grant from the National Science Foundation supplemented by the Commonwealth of Pennsylvania established the Pittsburgh Supercomputing Center. Since then more than 2300 scientists and engineers at 257 universities and research centers in 47 states have relied on the Center to help advance their projects. This work has produced nearly 700 papers in science and engineering journals.

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

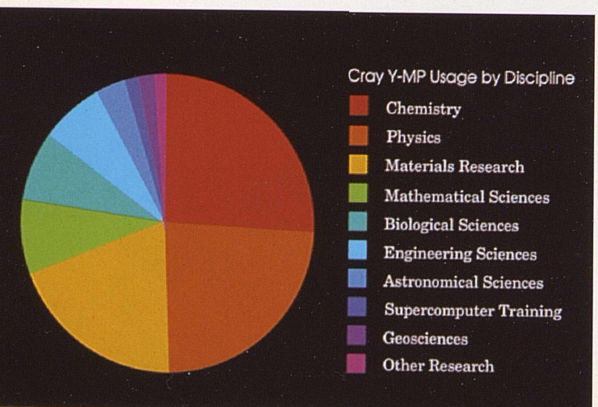

Heterogeneous Computing: Stepping into the Future

An environment in which multiple, independent but cooperative computers can work on a single application, that's the basic idea, explain Scientific Directors Mike Levine and Ralph Roskies. "For the really large, grand-challenge problems such as global climate, gene sequencing and protein structure," says Levine, "we need more computing capability. Heterogeneous computing expands the classes of research problems that can be feasibly attacked."

Beverly Clayton, Executive Director of the Pittsburgh Supercomputing Center. "Through our workshops and on-line documentation, and especially through the commitment and problem-solving skill of our staff, we've gained a reputation for excellent service to users."

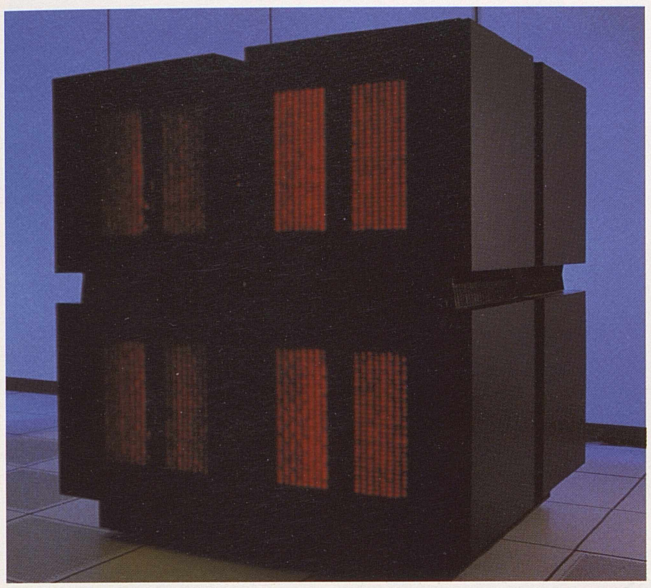

Pittsburgh Supercomputing Center's Connection Machine (CM-2).

With support from the Defense Advanced Research Projects Agency (DARPA) and the National Science Foundation (NSF), the Center plunged into this wider world of supercomputing in April when it installed a CM-2 model Connection Machine. "The CM-2 is a major component in our plans for heterogeneous computing," said Roskies. "When linked with the Y-MP, so that users can distribute applications between the two systems, it will give a substantial boost to the research effort in this country."

With 32,000 separate processing units, the CM-2 is a "massively-parallel" computer, in a sense the antithesis of the CRAY Y-MP, which has eight extremely powerful processors. Each CM-2 processor is less powerful than a personal computer, but for appropriate problems it attains supercomputer performance via the team approach: All 32,000 processors compute simultaneously, working on independent segments of the job.

Several applications are already running on the CM-2, and work is underway on a direct high-speed link (800 million bits per second) between the CM-2 and Y-MP. This linked system will allow users to divide tasks between the two machines, taking advantage of the CM-2's massively-parallel architecture for appropriate parts of a problem. "Hooking these computers together is a tough challenge," notes Roskies, "but we have a talented staff, and development projects like this, aimed at making better tools available for scientific research, are one of the most important things we do."

Jim Kasdorf, Director of Supercomputing, Westinghouse Electric Corporation, shows the Center's CRAY Y-MP to participants in a "Supercomputing Fundamentals" workshop, held August 14-18, 1989. Participants included representatives from a number of historically black colleges and universities. Kasdorf and Bob Stock, Manager of User Services, organized this workshop to broaden awareness of supercomputing in undergraduate science programs. The Minority Research Initiative of the National Science Foundation provided support.

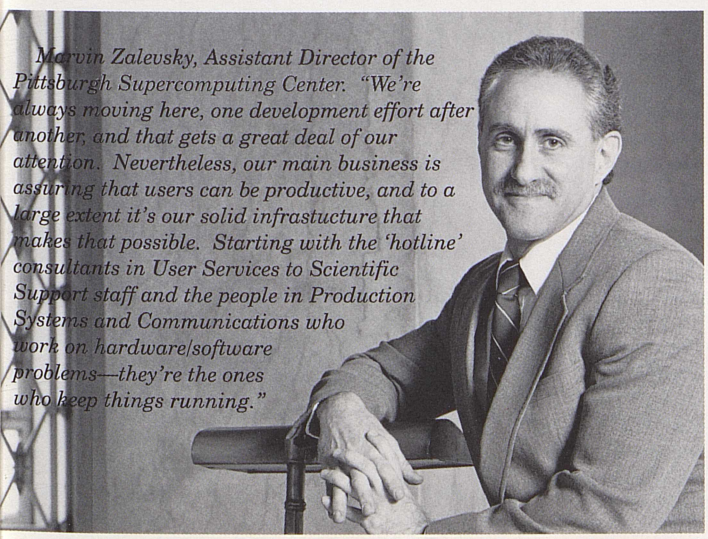

Marvin Zalevsky, Assistant Director of the Pittsburgh Supercomputing Center. "We're always moving here, one development effort after another, and that gets a great deal of our attention. Nevertheless, our main business is assuring that users can be productive, and to a large extent it's our solid infrastucture that makes that possible. Starting with the 'hotline' consultants in User Services to Scientific Support staff and the people in Production Systems and Communications who work on hardware/software problems—they're the ones who keep things running."

The Ties that Bind: Networking

Reliable, high-speed interconnections among hardware components is a necessity in a heterogeneous environment. To that end the Center is participating with Carnegie Mellon's Department of Computer Science, Bell of Pennsylvania and Bellcore as one of five national testbeds for gigabit (a billion bits per second) networking. Transferring data at this rate means a 20-fold increase over current speed. The project is coordinated by the Corporation for National Research Initiative and receives support from NSF and DARPA.

The Center is also installing very high-speed links between the computer room (at Westinghouse Energy Center) and its research facilities at the Mellon Institute building. UltraNet hubs (Ultra Network Technologies) capable of gigabit data transfer are situated at both ends of this connection. Extensive testing is now underway—to satisfy the Center's stringent production standard. When complete, this link will greatly enhance visualization capability; an UltraNet "frame buffer" in the Center's graphics lab will allow users to view color images interactively at animation rate.

Software to Manage Heterogeneous Computing

To help manage the complexity of a heterogeneous environment, the Center has participated in a joint software project with Digital Equipment Corporation (DEC). The product of this effort, the Distributed Code Manager (DCM), was announced this year. "Connections between high-powered workstations, file servers and supercomputers can enormously increase the productivity of a researcher," says Chris Maher, Manager of Scientific Support, who coordinates the DCM project. "But there's substantial overhead involved in this kind of environment."

Chris Maher (seated) and Grace Giras of the Pittsburgh Supercomputing Center along with Dennis Carleton of DEC developed the Distributed Code Manager.

Distributing applications among different machines requires switching between different operating systems, and incompatibilities can be a major headache. DCM lifts much of this burden from the user. It will automatically migrate files from one machine to another, start processes on remote machines and communicate results to the user. It eliminates the need to log onto different systems and copy files across the network. DCM is now being tested at DEC's Marlboro, Massachusetts facility and at the Pittsburgh Supercomputing Center.

New Visualization Tools: P3D & expanded GPLOT

The Center's continuing development effort in computer graphics bore some exciting new fruit this year. Pittsburgh Three-Dimensional, P3D, is a standard format for three-dimensional graphics developed by Joel Welling, Chris Nuuja and Phil Andrews. "We've had great success with two-dimensional graphics," says Welling, who coordinates the Center's graphics facilities, "because of our system for supporting the standard format provided by Computer Graphic Metafiles (CGM). We've now introduced a similar system for 3-D." P3D essentially accomplishes for 3-D what CGM does for 2-D, and it can use the same device drivers as GPLOT, the Center's widely-used CGM translator.

Joel Welling (standing) and Chris Nuuja of the Center's graphics group in the expanded new graphics lab. New hardware includes a Silicon Graphics Personal Iris, two Sun Sparcstations, an IBM System/6000 workstation and a rack of animation equipment.

A new release of GPLOT this year added several new features, including point-and-click capability for most workstations. Developed by Scientific Specialist Phil Andrews, GPLOT is designed to interpret any valid CGM file and make it available to a wide variety of different operating systems and hardware devices. Probably the most widespread CGM interpreter available, GPLOT is used at hundreds of sites in several countries, including other NSF centers.

New Corporate Affiliates

The CRAY Y-MP is widely acknowledged to be the most powerful general-purpose scientific computer available, and this year two major U.S. corporations, USX and Pfizer Inc, added this capability to their research programs by becoming Corporate Affiliates of the Center. "It's a creative and valuable approach to industrial research in this country," says Levine. "Through public-private partnerships, we can accelerate the transition between new insights in basic science and improved products in the marketplace."

"Affiliation with the Pittsburgh Supercomputing Center opens new doors for our engineers," said Keith K. Kappmeyer, USS Vice President of Engineering and Research, in an April news release. "They can now solve problems beyond the scope of conventional computer-aided engineering." USX's steel-making R & D group, the USS Technical Center, is using the Y-MP to model steel rolling (see p. 32) and continuous casting.

In May, Pfizer, a New York corporation with worldwide pharmaceutical sales of $2.7 billion, moved to enhance its use of molecular-modeling techniques in the development of new drugs. Pfizer chemists and biologists will work with the Center's staff to implement their research objectives, notes Ira Hochman, the Center's Corporate Liaison. Tests indicate that the Y-MP will run typical problems 45-60 times faster than Pfizer's on-site system.

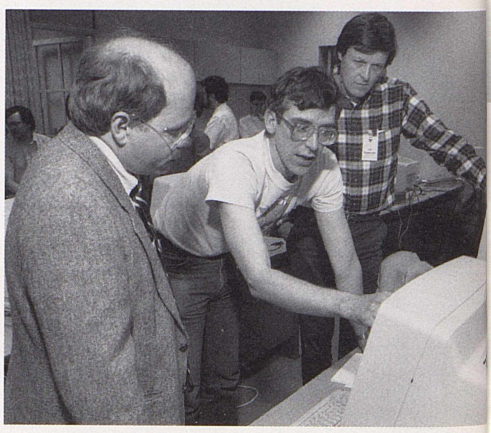

Leaning over a workstation in the Center's new Computer Training Center, John Burkardt of the Scientific Support staff coaches two participants in the Biomedical Fluid Dynamics workshop, April 18-20, 1990. This workshop introduced a group of biomedical researchers to the immersed-boundary technique, an efficient numerical method for simulating biological fluid flows.

The Biomedical Initiative

In September 1987, the National Institutes of Health awarded a grant to the Pittsburgh Supercomputing Center to advance the use of supercomputing in biomedical research. Now in its third year, the Biomedical Initiative has introduced more than 200 biomedical scientists to the techniques of supercomputing. These researchers have come from over 100 different institutions in more than 40 states to workshops at the Center, and more than 60 of them have begun to use the Center's CRAY Y-MP in their research.

Along with coordinating this very active training program, the Center's two Biomedical Scientific Specialists, Hugh Nicholas and David Deerfield, have compiled a comprehensive set of biomedical-related programs and databases, including several they have developed. "We provide a comprehensive computing environment for molecular biologists," says Nicholas, who has focused on tools for nucleic acid and protein sequence analysis. This year he guided work on an improved implementation of MaxSegs, a state-of-the-art program for high-speed searching of sequence databases to find similar regions of different biomolecules.

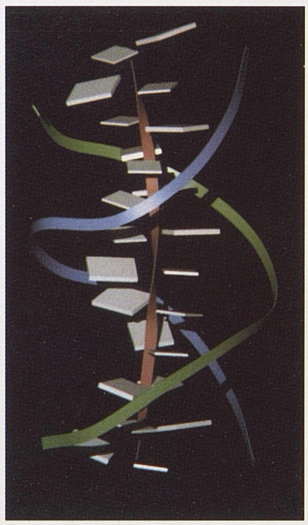

This three-dimensional representation of DNA developed by David Deerfield provides a precise way to visualize how basepairs align along the central axis of the DNA structure.

Expanding the Base of Knowledge

In another active year teaching the techniques of supercomputing, the Center organized 50 workshops and seminars with more than 1300 participants.

Who to Contact at the Pittsburgh Supercomputing Center

Requests for Computing Resources:
Wendy Janocha
412-268-5005

Information on Biomedical Workshops:
Nancy Kiser
412-268-5206

Other Workshops & the Summer Institute:
Casey Porto
412-268-7808

Industrial Affiliates Program:
Ira Hochman
412-268-2776

Newsletter and Documentation:
Vivian Benton
412-268-6355

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

Preparing for an on-site workshop at the University of Alaska Fairbanks, Education Coordinator Chuck Maiden (left) met with fellow travelers Jamshid Mahdavi and Casey Porto and Alaska Fairbanks' liaison Joe Hawkins (right) at the igloo in Polar World. The exhibit is a popular feature at the Carnegie Museum of Natural History, across the street from the Center's offices in the Mellon Institute building.

"Users can't always come to Pittsburgh for our workshops," says Maiden, "so we're making an effort to go to them." In addition to the University of Alaska Fairbanks, Maiden and company have conducted workshops at the University of Pennsylvania and Shippensburg State University.

Twenty scientists and engineers spent July 17-28, 1989 at the Pittsburgh Supercomputing Center's third Summer Institute. This highlight of the Center's training program is an intense introduction to supercomputing. As usual, it elicited many approving comments: "Thrilling. I was truly amazed by the staff availability, helpfulness and interest in the project work." "Great job making us feel at home." "Information that will be put to use not only by me, but by anyone at my institute that I can help to get on PSC's supercomputer."

Deb Nigra (left) and Vivian Benton handle the imposing task of keeping the Center's documentation up-to-date with its ever changing array of software and hardware. Nigra works on the extensive set of on-line help and example files. Benton, who edits PSC NEWS, the Center's newsletter, also edits and distributes hardcopy documentation.

Jan Asbury, Business Manager, works with Janette Hanchak (left) and Cheri Brooks (right) to keep tabs on financial matters.

Scientific Visualization

While supercomputing has been a tremendous boon, providing unprecedented ability to generate enormous amounts of data, it also presents researchers with an embarrassment of riches. What do you do with all the numbers? A computational model of Los Angeles air pollution, for instance, yields about 2000 paperback books worth of data. How can anyone analyze this amount of information and draw scientific conclusions in time to do any good?

The solution has been scientific visualization—images rather than numbers. Translate the data into a concrete form and then carefully look to see what you can see. In this way, relying on the brain's natural facility with visual information, researchers can in a few seconds extract insights that would take days, weeks or longer to get at digging through piles of numbers.

It's fascinating that this very high-technology science is evolving in a way that makes fresh demands on the basic, low-technology roots of science—careful observation. The images shown here were produced using the computer graphics facilities at the Pittsburgh Supercomputing Center.

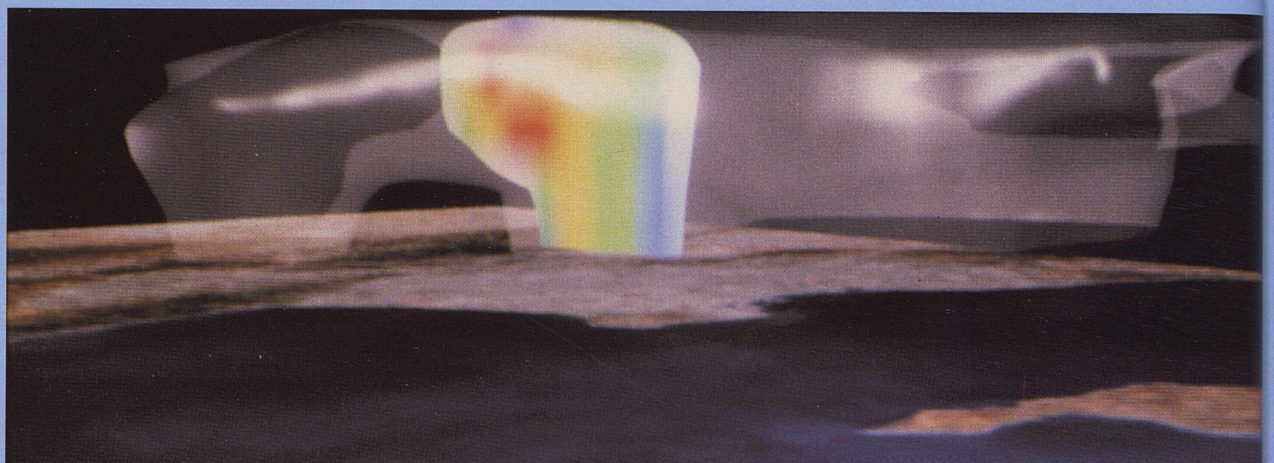

Smog over Los Angeles, Catalina Island in foreground. Gregory J. McRae, Gray Lorig, Chris Nuuja & Margaret Maschnei. The fine mesh indicates a volumetric contour of ozone, the primary constituent of smog, at concentration above EPA standard (.12 parts per million). Color shows gradations of ozone concentration above .12, increasing from blue to red, in a vertical slice through this ozone "cloud."

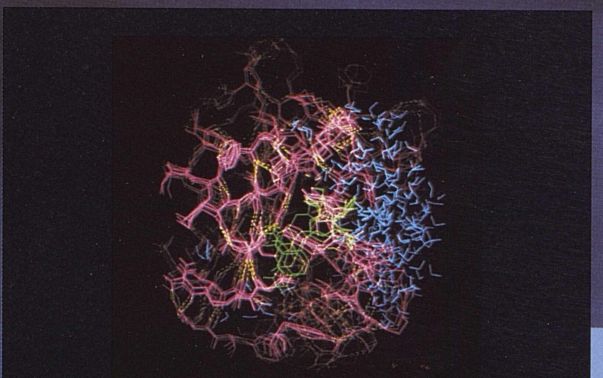

Dynamics of mutation for the drug trimethoprim at the active site of the enzyme dihydrofolate reductase. Charles Brooks & Scott Sneddon. Red indicates the active site and associated residues, green is the drug and blue is surrounding water.

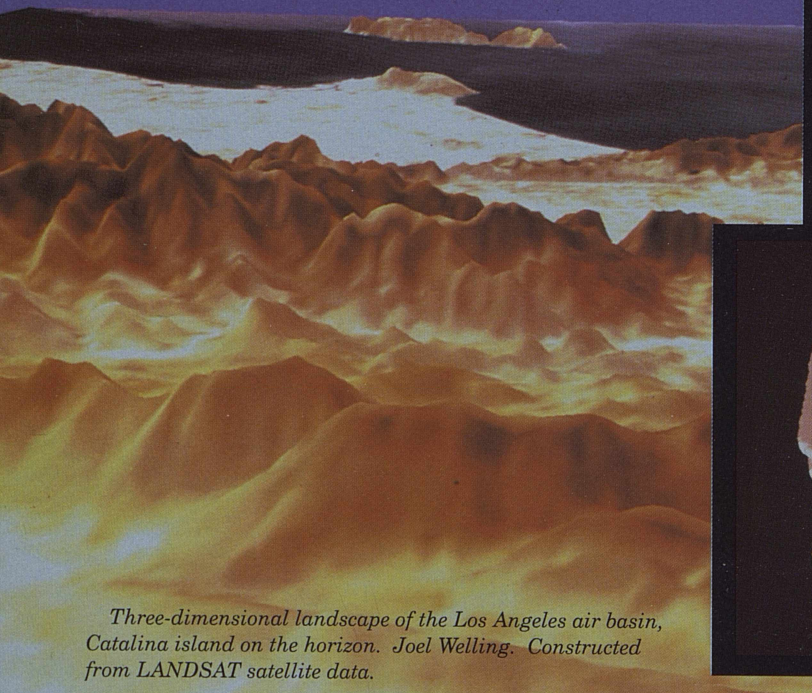

Three-dimensional landscape of the Los Angeles air basin, Catalina island on the horizon. Joel Welling. Constructed from LANDSAT satellite data.

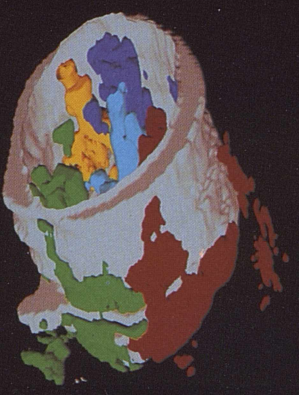

Three-dimensional reconstruction of lymph-node cancer. Rick Siderits, Cy Evans & Joel Welling. White indicates the lymph node, and the five colors indicate five separate branching growths of cancer tissue. This research was featured in the June 1990 issue of BioTechniques.

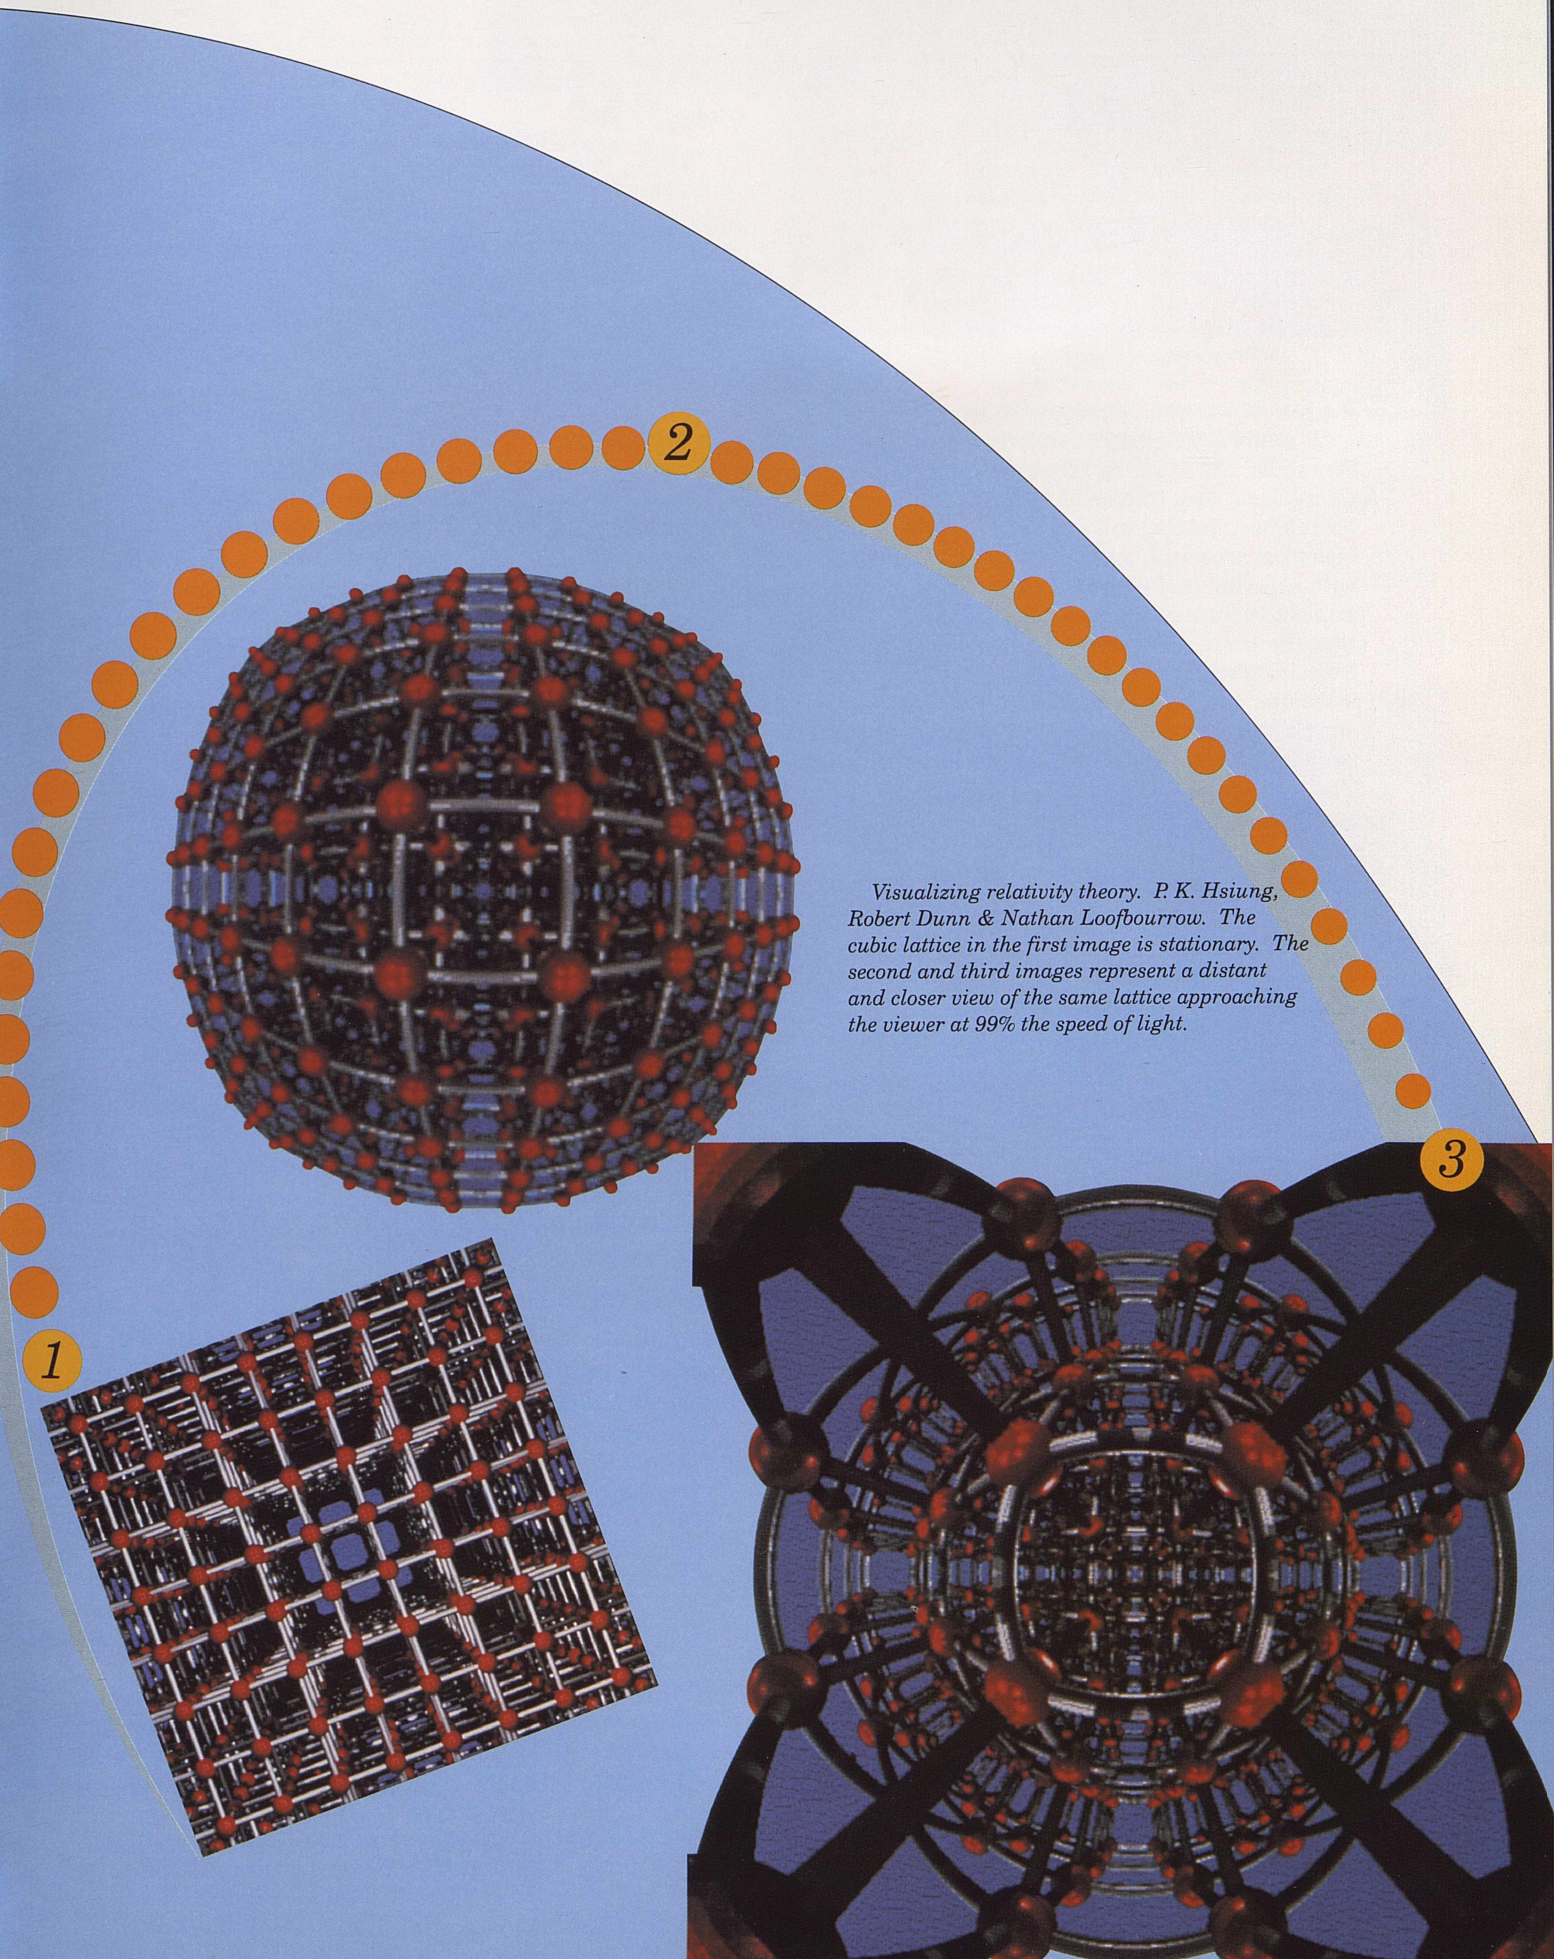

Visualizing relativity theory. P. K. Hsiung, Robert Dunn & Nathan Loofbourrow. The cubic lattice in the first image is stationary. The second and third images represent a distant and closer view of the same lattice approaching the viewer at 99% the speed of light.

Getting a Grip on Calcium

Structure, Function and Selectivity of Calcium-Binding Proteins

Harel Weinstein, Mount Sinai School of Medicine

"His hands were small and prehensile, with fingers knotted like a cord."

—Robert Louis Stevenson

The New World of Molecular Biology

"It really is an entire new world out there," says Harel Weinstein. He's talking about how improved experimental methods have combined with advances in computing techniques, clearing the way to new understanding of biological processes, and he's definitely excited. Armed with experimentally-derived data on the three-dimensional structure of biomolecules, it's becoming possible, says Weinstein, to do computer simulations that give fresh insights into the complex relationships between molecular structure and biological function. Weinstein and his colleagues at New York City's Mount Sinai School of Medicine have stepped into this new world. Their computer studies on calcium-binding proteins ask and answer questions not posed before about this interesting protein family.

Harel Weinstein chairs the department of physiology and biophysics at New York City's Mount Sinai School of Medicine. Here he explains some of the structural features of calcium-binding proteins during a seminar at the Pittsburgh Supercomputing Center this past March.

Binding Calcium with the EF Hand

Calcium-binding proteins, which include calmodulin, troponin-C and calbindin, play an essential mediary role in muscle contraction, heart rhythm, stress response and many other cellular functions triggered by calcium ion (Ca^{2+}) concentration. Calmodulin in particular is ubiquitous in animal cells, where—once it is activated by binding with Ca^{2+}—it in turn activates over 30 different target proteins. By opening and closing channels in the cell membrane, the calcium-calmodulin complex directly controls the calcium ion's vital ability to act as an intracellular messenger, setting off responses to various external stimuli.

What's especially interesting about this very diverse family is that all of its members have the same structural motif, a helix-loop-helix sequence called the "EF hand," at the calcium binding-site. Like a primordial expression of the prehensile reflex, a cluster of eight oxygen atoms in the EF-hand loop wrap around and hold the ion in place. Robert Kretsinger first isolated this motif in parvalbumin, a fish-muscle protein, and named it for parvalbumin's E and F helices, which can be visualized as the thumb and forefinger of a right hand.

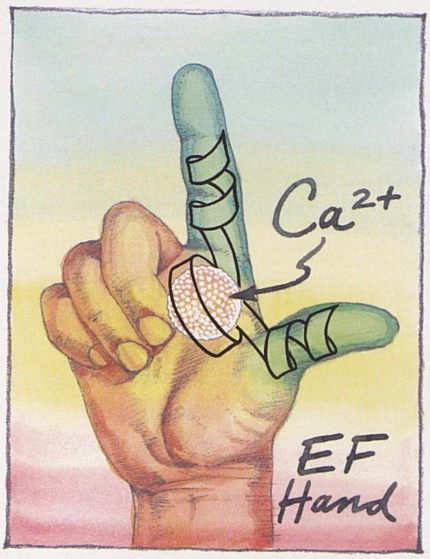

Probing the EF Hand with Molecular Dynamics

This structural similarity of the EF-hand binding-loop from protein to protein is extraordinary, notes Weinstein, because even the amino-acid side-chains in the loop region are the same. To examine the structural and electronic characteristics of this binding-site, to get a grip on why this precise configuration is so exclusive, favored over the multitude of other possibilities, Weinstein turned to computations based in the theory of molecular dynamics (MD). "Molecular dynamics is the only approach that allows us to treat molecular systems of this magnitude."

In MD methods, a molecular structure is thought of as a collection of atoms in space. Their changing behavior over time (dynamics) is computed according to Newton's equations of motion, which express the basic relationship that force equals mass times acceleration. The mass of each atom is known, and the force is derived

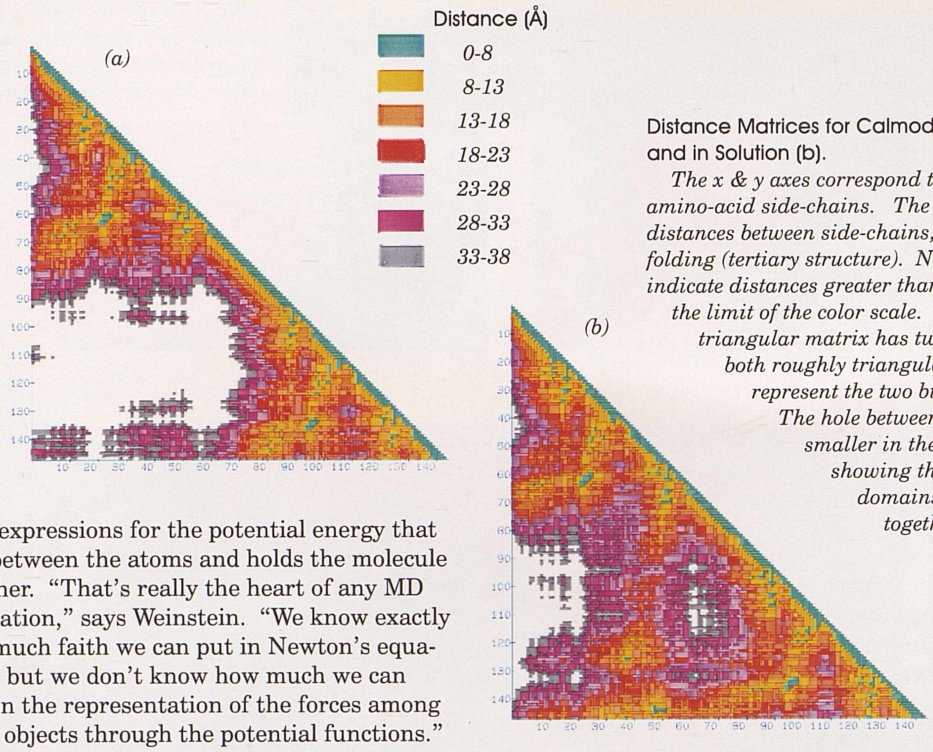

Distance Matrices for Calmodulin: Crystallized (a) and in Solution (b).

The x & y axes correspond to the sequence of amino-acid side-chains. The color-coding shows distances between side-chains, an indication of folding (tertiary structure). Non-colored areas indicate distances greater than 38 angstroms (Å), the limit of the color scale. "You see that each triangular matrix has two colored regions, both roughly triangular, within it. These represent the two binding domains. The hole between them is much smaller in the second plot, showing that the binding domains have moved closer together."

from expressions for the potential energy that acts between the atoms and holds the molecule together. "That's really the heart of any MD simulation," says Weinstein. "We know exactly how much faith we can put in Newton's equations, but we don't know how much we can rely on the representation of the forces among these objects through the potential functions."

Weinstein and his colleagues did quantum-theory calculations to develop potential functions for molecules involving Ca^{2+}. They tested these parameters by running MD calculations on two calcium-binding molecules for which experimental data is available—cyclo (-L-Pro-Gly-)$_3$ peptide and calbindin-D. In both instances, they obtained "very satisfactory agreement" between the MD-calculated structure and experiment. These simulations also disclosed that water molecules in the protein's natural environment are essential for stabilizing its structure. "Solvent," says Weinstein, "is a very important component of this story."

Results: Calmodulin is No Dumbbell

Having established the viability of his basic method, Weinstein went on to other MD simulations. Computations investigating the competition between calcium and magnesium ions for the EF-hand binding-site shed new light on why Ca^{2+} is preferred. Water again plays a key role. Computed binding-energies indicate that the protein can bind equally well with Ca^{2+} or Mg^{2+} but that Mg^{2+} binds with water more strongly than Ca^{2+}. The wraparound cluster of oxygen atoms in the EF-hand binding loop is in effect an island, a sequestered site high and dry from the surrounding water, and Weinstein's computations suggest that Ca^{2+} is preferred because it's harder for Mg^{2+} to leave the water.

Weinstein's most exciting finding to date came from computations on calmodulin. Once more, being able to account for the water, possible only with high performance computing systems like the CRAY Y-MP, revealed new information. "We were enormously surprised." The calmodulin structure from X-ray crystallographic studies is a "dumbbell"—two binding-domains, each with two Ca^{2+} binding-sites, joined by an erect helical backbone. Weinstein used MD to simulate calcium-bound calmodulin in solution, in effect throwing the dumbbell into a swimming pool. He learned that when calmodulin goes swimming it curls over in the middle and pulls the two binding domains much closer together.

Other studies have suggested that when the calcium-calmodulin complex binds to its target proteins it

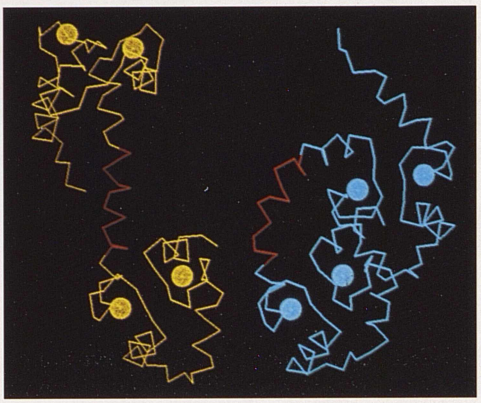

The "dumbbell" structure of crystallized calmodulin (left) and compacted in solution (right) after 246 picoseconds (.000000000246 sec.) of molecular dynamics simulation. Red represents the helical backbone connecting the two Ca^{2+} binding domains (2 binding-sites in each domain).

assumes this "compacted" form. Weinstein's supercomputer studies, however, are the first to indicate that compaction occurs not simply as a result of protein-binding but as a direct consequence of changes in the helical backbone helped along by Ca^{2+} binding. Calcium binding actually enables calmodulin's interaction with other proteins. "This shows the great value of theoretical simulations," notes Weinstein. "The crystal structure is what people have been using to explain the activity of this very important molecule, and we now know that could be misleading."

References:
F. Sussman & H. Weinstein, "On the ion selectivity in Ca-binding proteins: The cyclo (-L-Pro-Gly-)$_3$ peptide as a model," *Proceedings of the National Academy of Science, Biophysics* **86**, 7880 (1989).
K. Hori, J. N. Kushick & H. Weinstein, "Structural and energetic parameters of Ca^{2+} binding to peptides and proteins," *Biopolymers* **27**, 1865 (1988).

This research was supported by the National Institutes of Health, the National Institute on Drug Abuse, the United States Environmental Protection Agency and the National Science Foundation.

The Tell-Tale Heart

Epicardial Potentials Solutions to the Inverse Problem in Electrocardiography

Yoram Rudy, Barbara Messinger-Rapport
& Howard Oster, Case Western Reserve University

"a low, dull, quick sound–much such a sound as a watch makes when enveloped in cotton."

Inverse Electrocardiography

In Poe's classic tale, the sound of a beating heart becomes the inescapable signal of a diseased soul. For cardiologists also, signals from the heart tell a great deal about health, electrocardiography being one of the most valuable diagnostic tools developed by modern medicine. Still, conventional electrocardiography is inherently limited; using a few electrodes on the body surface, it reports what amounts to a muffled echo—a smoothed-out, low-resolution imprint of the electrical impulses acting on the muscle fibers of the heart.

If an accurate map of electrical potentials on the heart surface—the epicardium—were available, it could save lives. The immediate cause of death in most fatal heart attacks, for instance, is arrhythmia, a breakdown in the normal pattern of cardiac electrical activity. Many arrhythmias occur because an abnormal region inside the heart begins generating impulses that overwhelm the heart's natural pacemaker, the sinus node. Surgical removal of the problem tissue can correct the arrhythmia, but it must be located, which requires electrically mapping the heart with a hand-held probe during open-heart surgery.

Researchers at Case Western Reserve are developing a technique to obtain epicardial potentials noninvasively. Along with reducing the risk of open-heart surgery, this could make it possible to use minimally-invasive, precise techniques like catheter ablation, in which electric current is applied through a fine plastic tube to destroy the abnormal tissue. "With the inverse solution, we will be able to noninvasively localize abnormal arrhythmogenic activity within a centimeter," says Dr. Yoram Rudy. Inverse electrocardiography, as Rudy terms his approach, is feasible because improved electronics make it practical to use hundreds of electrodes, giving a complete map of electric potentials on the body-surface. From this information, a corresponding map of epicardial potentials can be constructed by way of a mathematical transformation.

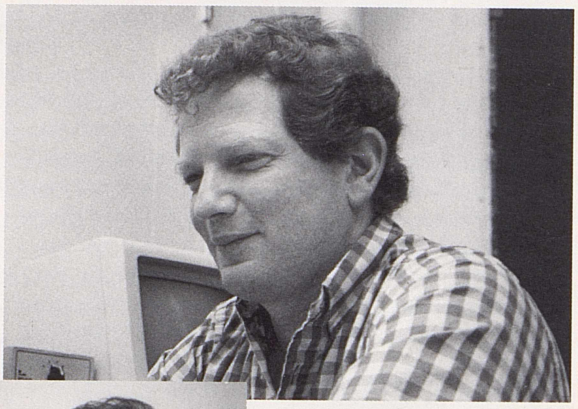

Yoram Rudy. In another project at the Pittsburgh Supercomputing Center, Rudy and his colleague at Case Western, Weilun Quan, are using the CRAY Y-MP to conduct theory-based simulations at the cellular level of how electrical impulses propagate in heart-muscle .

Captain Barbara Messinger-Rapport is currently serving her residency at Wilford Hall Medical Center, San Antonio, Texas.

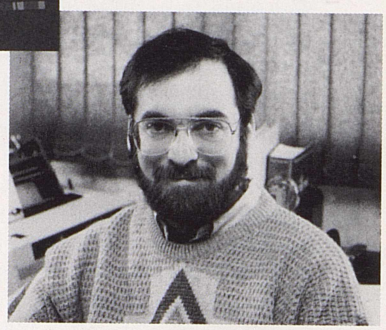

Howard Oster

The Eccentric-Spheres Model

The specific problem that has occupied Rudy, his colleague Barbara Messinger-Rapport and Ph.D. student Howard Oster is that the mathematical transformation is "ill-posed"—slight inaccuracies in the body-surface data magnify to enormous errors in the epicardial map. "Basically, there are two ways to go at it," says Rudy.

This cross-section of the eccentric spheres model shows the epicardium (red) in relation to the torso surface, which incorporates concentric surrounding layers of lung (blue), muscle tissue (light brown) and fat (beige).

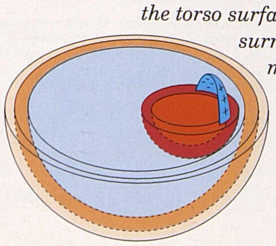

Measured and Computed Epicardial Potentials

These contour maps show electric potentials on the epicardial surface of the beating dog-heart model at an early stage of an advancing wave of activation. Each contour indicates an interval of 50 microvolts. The regularized inverse solution (right) located the maximum and minimum values within one centimeter of the measured values (left).

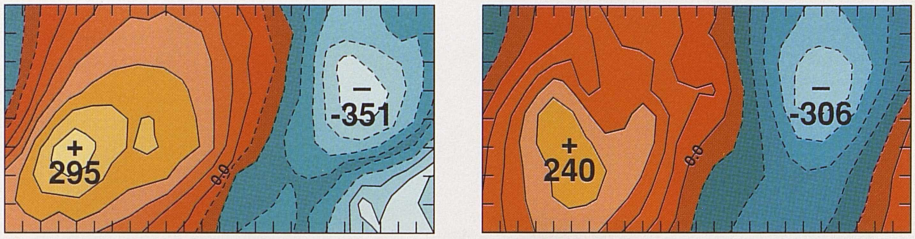

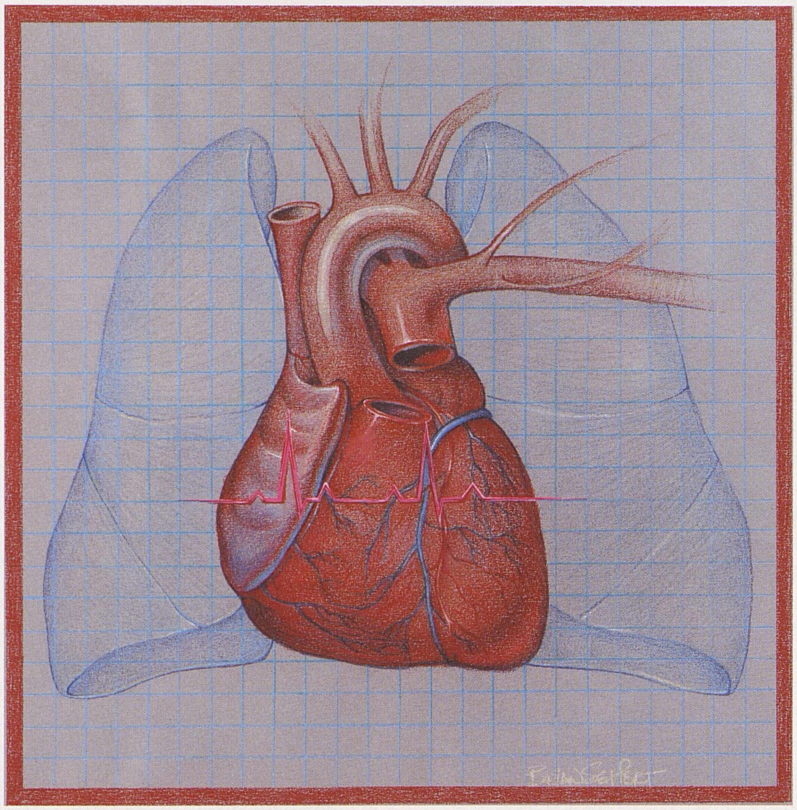

"You can try to reduce measurement inaccuracy—noise, or you can incorporate information about epicardial potentials into the solution procedure, in effect constrain it within realistic bounds." Noise is inherent in the process, and analyses by Rudy and others show that even a ten-fold improvement in measurement accuracy only slightly improves the epicardial solution.

To investigate the second approach, "regularizing" the solution by incorporating physiological information into the equations, the researchers used an eccentric-spheres model of the heart-torso system. "The model," notes Messinger-Rapport, "gave us an idea of the relative contributions of the lungs, skeletal muscle and fat layers surrounding the heart to the measured signal on the body surface." Though a simplified approximation of reality, the model's regular geometry made it possible to solve the inverse equations analytically, thereby providing a practical way to see how varying physical parameters affects the solution. The model showed, for instance, that fat conductivity had negligible effect and that slight variations in heart size and placement could be very significant. "We learned that you can improve the quality of the solution tremendously," says Rudy, "by introducing intelligent physiologically-based constraints."

Computing with a Realistic Model

Rudy and Messinger-Rapport incorporated physical constraints into the inverse solution through a regularization technique known as Tikhonov regularization, and they computationally verified it using the eccentric-spheres model. They then applied their method to a realistic heart-torso system developed by Dr. Bruno Taccardi of the University of Utah. The beating heart of a dog is inserted at the proper position in a torso-shaped tank filled with a conductive solution; 400 electrodes record surface potentials. Inverse computations from this surface data showed good agreement with measurements from 200 electrodes on the dog-heart epicardium. "Our results indicate," says Messinger-Rapport, "that it would be possible to locate the source of an arrhythmia on the epicardium within a one centimeter radius."

"With 400 measured points on the surface and over 200 on the heart," says Messinger-Rapport, "the supercomputer was indispensable. The speed of the CRAY along with the vectorizability of our code, allowed us to experiment with different regularization methods on the entire set of data." "The geometry is irregular," adds Rudy. "You have to segment the heart surface and body surface into a large number of discrete elements, and that results in very large-scale matrices. Without supercomputing, this project would be dead."

References:

Yoram Rudy & Barbara J. Messinger-Rapport, "The Inverse Problem in Electrocardiography: Solutions in Terms of Epicardial Potentials," *CRC Critical Reviews in Biomedical Engineering* **16**, 215 (1988).

Barbara J. Messinger-Rapport & Yoram Rudy, "Computational issues of importance to the inverse recovery of epicardial potentials in a realistic heart-torso geometry," *Mathematical Biosciences* **97**, 85 (1989).

This work was supported by the National Institutes of Health, the American Heart Association Northeast Ohio Affiliate and the National Science Foundation.

Looking into a Clouded Globe

Computational Aspects of Moisture Transport in Global Models of the Atmosphere

David L. Williamson, The National Center for Atmospheric Research

"If we hope to live a long time, we must begin to think like a geological force. That is, we must become the first geological force to learn to think. In geological terms, time is not money. Time is everything. In geological terms, what can happen will happen."

Global Climate: Bad News & Good News

How can we rationally cope with the difficult (and difficult to face) questions of global climate change? How can we assess what will happen in ten years or 100 years if worldwide emissions of CO_2 (and other greenhouse gases) continue rising at the exponential rate of the past ten years? Except for small pieces of the puzzle, laboratory methods aren't much help; the land-ocean-atmosphere system is far too large and complex. The Earth itself is, in one sense, the laboratory and human activity is the variable factor, but you can't do a controlled experiment. By the time we know the results, it could be too late.

Twentieth-century technology has created grave problems, and it has brought forth sophisticated new tools for dealing with them. General-circulation models (GCMs) combine our most advanced knowledge of climate processes with our most advanced computing capability. Scientists devise mathematical expressions for the governing physical processes—hot air rises, cold air sinks, the earth absorbs and gives off solar heat, water vapor condenses into clouds, gravity plays its role, energy and mass are conserved and so on. The equations themselves would be unsolvable—fruitless theory—without supercomputing. In the last five to ten years, advances in computing have brought GCMs into the mainstream of Earth science.

David Williamson heads the Climate Modeling Section at the National Center for Atmospheric Research (NCAR), which is supported by the National Science Foundation. The group of scientists he works with develop and maintain one of the most sophisticated GCMs in the world. NCAR's Community Climate Model (CCM), as the name implies, is used by scientists in a variety of disciplines. "It's an attempt to provide a well-documented, reasonably state-of-the-art model," says Williamson, "so that everybody doesn't have to develop their own." In effect, the CCM provides a locus for constructive dialogue across disciplines directed at improving climate modeling. "Many diverse computational experiments can be seen to relate to each other when performed with the same model."

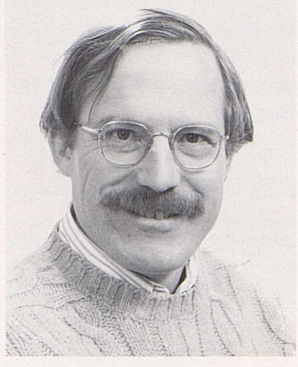

David L. Williamson

Moisture Transport: Fixing the Fixer

Williamson's group at NCAR has worked primarily on the atmospheric component of the CCM, and recently he has focused on the CCM's mechanisms for "moisture transport." Using the CRAY Y-MP at the Pittsburgh Supercomputing Center, Williamson and colleagues Philip Rasch and Jeffrey Kiehl have worked on a "serious error" in the way the CCM calculates the movement of water vapor in the atmosphere. The problem is related to the mathematical characteristics of the CCM: It is a "spectral" model and as such calculates water vapor movement in terms of a wave-like spectrum of frequencies. The mathematical transformation to "spectral

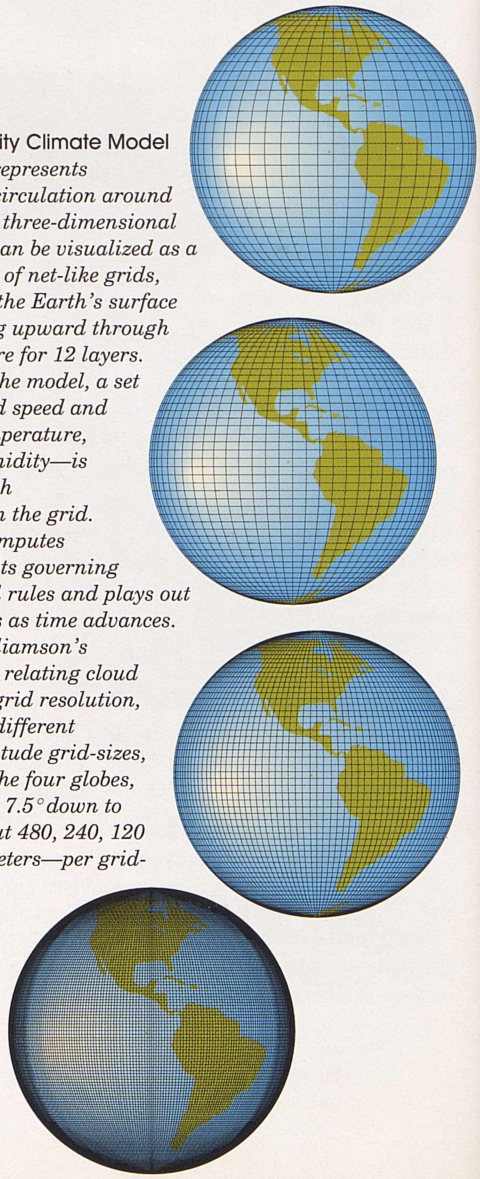

The Community Climate Model

The CCM represents atmospheric circulation around the globe in a three-dimensional space which can be visualized as a layered series of net-like grids, beginning at the Earth's surface and extending upward through the atmosphere for 12 layers. To initialize the model, a set of data—wind speed and direction, temperature, pressure, humidity—is entered at each intersection in the grid. The model computes according to its governing mathematical rules and plays out what happens as time advances. In David Williamson's computations relating cloud amount and grid resolution, he used four different latitude-longitude grid-sizes, as shown in the four globes, ranging from 7.5° down to 1.1°—or about 480, 240, 120 and 60 kilometers—per grid-box side.

The Jagiellonian Globe
The original of this globe was crafted about 1510 at Jagiellonian University, Krakow, Poland. It is believed to be the first model of the Earth to include America as a separate continent.

space" is an efficient way to do the calculation, but it also requires an approximation, known as "truncation," that can introduce error.

The most noticeable result of this error is "undershooting"—some areas of the model, certain grid "boxes," can end up with negative moisture. The model compensates for undershooting with a "fixer." "Because of basic properties of the scheme itself, the truncation," says Williamson, "can give these negative values, and something has to be done. People have tended to just bring the negative values up to zero or some small number above zero." Williamson and Rasch did a series of computations in which they isolated the water vapor produced by the fixer and found that in certain areas it was as much as 40% of the total.

These studies also revealed an anomaly: running the CCM with higher resolution of longitude-latitude reduced the amount of clouds. Increased resolution improves the model's fluid dynamics and should give a better result, but the spectral fixer to some degree compensates for the dynamics at low resolution. "We knew we had some missing processes," notes Williamson, "but this pointed out that those processes are crucial. When we improve the dynamics by going to higher resolutions, the other processes must be included to get the cloudiness right—a very important aspect of the climate problem."

To avoid the spectral fixer, Williamson and Rasch have developed and tested a non-spectral method of calculating the moisture. This approach incorporates "shape-preservation"—it won't cause any negatives or excess positives. "We developed a scheme which is not perfect yet, but it eliminates many of the problems we had with the spectral method we had been using." Williamson plans to implement the shape-preserving technique in the next version of the CCM.

GCMs and Supercomputing

GCMs are a sure sign that we've come a long way from relying on omens, oracles and crystal balls to allay our anxiety about the uncertainty of nature. Though GCMs aren't perfect, they're the best tool we have—nothing else comes close—for assessing the impact of the greenhouse effect, ozone depletion, nuclear winter and other climatic phenomena. Their biggest limitation, even with today's supercomputers, is computing capability. Williamson and Kiehl's study on the relation between model resolution and clouds concludes by observing that finer resolution, both horizontally and vertically, is necessary to get reasonably accurate predictions of global cloudiness—key information for assessing the greenhouse effect. These improvements await the next generation of supercomputing.

References:
Philip J. Rasch & David L. Williamson, "Computational Aspects of Moisture Transport in Global Models of the Atmosphere," *Quarterly Journal of the Royal Meteorological Society* (July 1990), in press.
J. T. Kiehl & D. L. Williamson, "Dependence of Cloud Amount on Horizontal Resolution in the NCAR Community Climate Model," *Journal of Geophysical Research*, submitted.

Epigraph: Jonathan Weiner, *The Next One Hundred Years* (New York: Bantam, 1990), pp. 232-33.

The Jagiellonian Globe reproduction is in the Polish nationality classroom at the University of Pittsburgh's Cathedral of Learning.

This work is supported by the National Science Foundation and the U.S. Navy.

Feeding the Monster

Tidal Triggering of Starbursts and Active Galactic Nuclei

James E. Gunn & Lars Hernquist, Princeton University

"In what distant deeps or skies burnt the fire of thine eyes."

Cosmology and Galaxy Diversity

Try to imagine a heat wave a hundred-thousand times hotter than the core of the sun. Then forget about the sun because when the "big bang" happened, about ten billion years ago, it was too hot for stars to form. Galaxies and stars began to appear after about a billion years of expansion and cooling. Hydrogen and helium started collecting in isolated clouds, gravity pulled in more gas, and the clouds became proto-galaxies, spinning clumps of gaseous matter. As the gas cooled further, much of its matter and energy converted to points of intense radiance—stars. A hundred billion of them speckle the Milky Way, one galaxy among billions in the cosmos, and off in a corner orbiting an unsingular star, we live.

More or less, that's the prevailing story as 20th century astronomers and physicists have discovered it, though vast uncertainties remain. With respect to galaxies, for instance, what accounts for the wide variations in size, mass, density, temperature, illumination and other properties? In shape alone, there are the huge star footballs (ellipticals) and the spirals—two-armed, four-armed and crab-like, some flat and round as frisbees, others humpbacked and distorted. What coherent theory of galaxy formation can encompass this diversity? What can the differences tell us about pre-history?

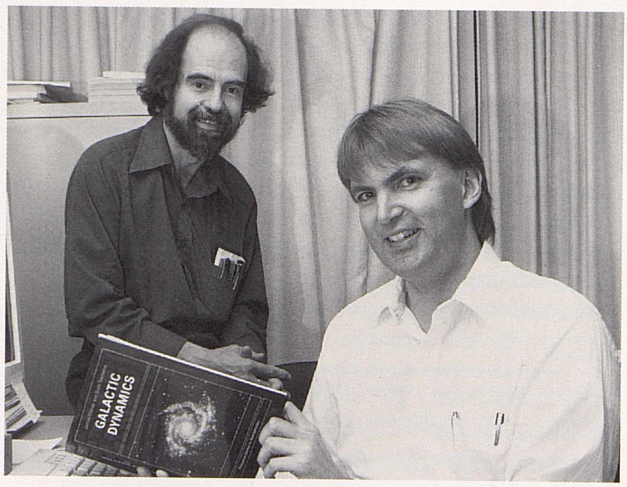

James E. Gunn and Lars Hernquist

"Problems in science are never really solved," says Gunn, "but Lars' calculations go a long way toward giving us some real understanding of what happens, and in that sense they solve a long-standing problem."

Starbursts and Black Holes Rimmed with Gas

In research at the Pittsburgh Supercomputing Center, Lars Hernquist has been collaborating with James Gunn and graduate student Neal Katz to probe these questions. Recently, Hernquist has followed up on some earlier work by Gunn investigating unusual phenomena in which large masses of gas concentrate in the center of a galaxy. In one of these phenomena, called starbursting, galaxies with very bright infrared radiance convert their central gas to stars much more rapidly than the relatively steady rate of star formation in most spiral galaxies.

Seyfert galaxies are a related phenomenon. Though there is no anomalous burst of star formation, they display intense radiance in a central region even more confined than starbursts. From this small nuclear region, Seyfert galaxies give off as much light as from all the rest of the galaxy. It has become generally well-accepted that a black hole could account for the intense centralization of this gas activity. Nevertheless, as Gunn pointed out in his 1979 article "Feeding the Monster," the existence of a black hole doesn't resolve some more basic questions that also apply to starbursts and quasars, the brightest objects in the sky: Where do the large amounts of gas come from and how do they get transported into a dense central region?

Gunn demonstrated that angular momentum transfer alone couldn't account for these phenomena. "Nature has presented us with a *fait accompli*. Matter finds its way from the outer regions of a galaxy into a very small volume in the middle, and back-of-the-envelope calculations tell you that simply can't happen." He postulated collisions between galaxies as a possible explanation, an insight that could not be tested adequately until Hernquist attacked the problem with supercomputing. "Lars has solved a problem that's been around since quasars were discovered. How do you get this material in? It requires a 'full-up' treatment of the gas dynamics along with gravity."

Smooth Particle Hydrodynamics with Tree Structures

Due to the scope of the processes involved in galaxy formation—light year distances and complex interactions involving gas and billions of stars—computing capability is an inherent limitation, one that Hernquist has handled with some innovative algorithms. In earlier work at Pittsburgh, Hernquist took the "tree-structure" approach developed by his colleagues Josh Barnes and

Tree Structures and TREESPH

Though the mathematics is quite complex, the basic concept of tree-structured data follows from its name. With a direct calculation (without tree structures), the computer must calculate each star's interaction with every other star. Thus with 40,000 stars in a galaxy, the computer does 40,000 calculations for each star, 1.6 billion in total (40,000^2). Tree structuring, on the other hand, is hierarchical, like a corporate chain of authority, rather than star by star. The algorithm calculates each star's interactions with clusters of stars grouped by proximity. This reduces the total calculations to about NlogN rather than N^2—about 160,000 calculations rather than 1.6 billion for 40,000 stars, which makes an enormous difference in the size of the galaxies which can be modeled.

Hernquist's innovation with TREESPH is applying tree structure to the gas dynamics. "Numerically, it's one of the most difficult things. The gas is thermally unstable; it can't exist at a uniform temperature, and it fragments into small, very dense regions. In these very dense regions, things change much more rapidly, which means the force needs to be computed much more often. And that's the complication, because if you compute for the entire system at this very small time scale, it would be too expensive to do anything. My algorithm allows fast evolving regions to evolve separately and still interact with the rest of the system. You compute forces at different times in different regions, depending on the local properties of the fluid."

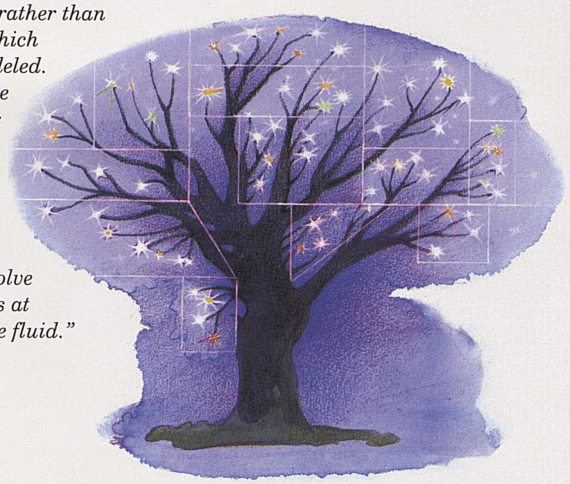

Piet Hut on VAX computers and adapted it to a CRAY X-MP, boosting the speed by a factor of 150-200 in the process. By providing an accurate approximation of the interactions among billions of stars, tree structures greatly expand the size of a galaxy that can be modeled.

To simulate starburst galaxies, Hernquist combined tree structures with a technique called smoothed-particle hydrodynamics (SPH). As the name implies, SPH uses fluid dynamics equations to model galactic gas. It divides the gas into parcels, defined by a "smoothing length," which interact with the particles used to represent stars. Hernquist's new code, TREESPH, developed in collaboration with Katz, joins the two methods in a synthesis that can simulate larger galaxies than was possible before with both gas and stars and, says Hernquist, with greater accuracy. "This code allows you to treat the gas dynamics rigorously and include other effects, like the self-gravity of the gas and cooling due to radiation."

When Worlds Collide

With the necessary tools in place, Hernquist set out to investigate Gunn's earlier suggestions about the relation between intense gas activity and galaxy collisions. The prime example is a small satellite galaxy orbiting and eventually merging with a larger galaxy. The gravitational field of the satellite creates tidal effects in the large galaxy's gas similar to the lunar influence on the sea. Hernquist modeled gas and stars of a galaxy merging with a satellite having 10% as much mass. By the end of the simulation, 40% of the gas initially distributed throughout the large galaxy collapsed into a dense central core. The time scales and gas densities of this process agree well with observations of starbursts.

Will subsequent evolution of this dense region transport the gas further, until it spirals onto the black hole and dumps into the very center, like a Seyfert galaxy? Hernquist believes it may and that tidal processes may also explain quasars. But to test this requires more modeling ingenuity. "You need to include how gas converts into stars," explains Hernquist, "a process we don't understand very well. I think I might focus on that dense, centrally concentrated region and ignore the rest of the galaxy. Then include effects of the conversion to stars and see how that evolves. Starbursting and nuclear activity may be intrinsic to galaxy formation itself."

References:

James E. Gunn, "Feeding the monster: gas discs in elliptical galaxies," in *Active galactic nuclei*, ed. C. Hazard & S. Mitton (Cambridge: Cambridge Univ. Press, 1979).

Lars Hernquist & Neal Katz, "TREESPH: A Unification of SPH with the Hierarchical Tree Method," *Astrophysical Journal Supp.* **70**, 419 (1989).

Lars Hernquist, "Tidal triggering of starbursts and nuclear activity in galaxies," *Nature* **340**, 687 (1989).

Epigraph: from "The Tyger," William Blake.

In addition to allocations of CRAY computing time at the Pittsburgh Supercomputing Center, this work was supported by a New Jersey High Technology Grant.

Gas Distribution in the Large Galaxy

The gas component of the large disk, shown isolated from stars, 425 million years after it begins to merge with the satellite indicates the disparity between "hot" and "cold" gas regions. The coldest gas (purple/violet) is around 100°Kelvin; the hottest (red) about 10,000°K. Because of thermal instability, relatively small amounts of gas are at intermediate temperatures (green and blue). The spiral arms of the gas structure develop as a result of satellite tidal effects.

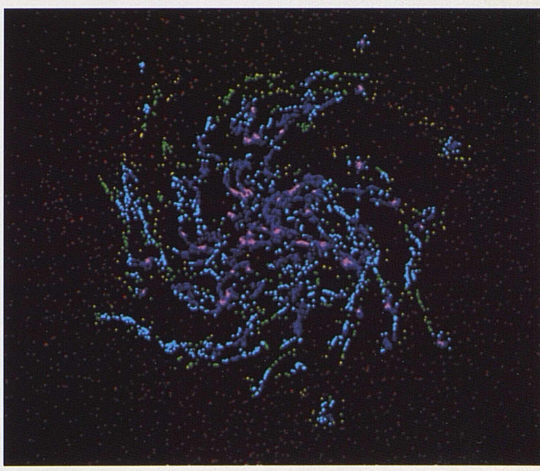

Collision of a Galaxy Disk and a Smaller Satellite Companion

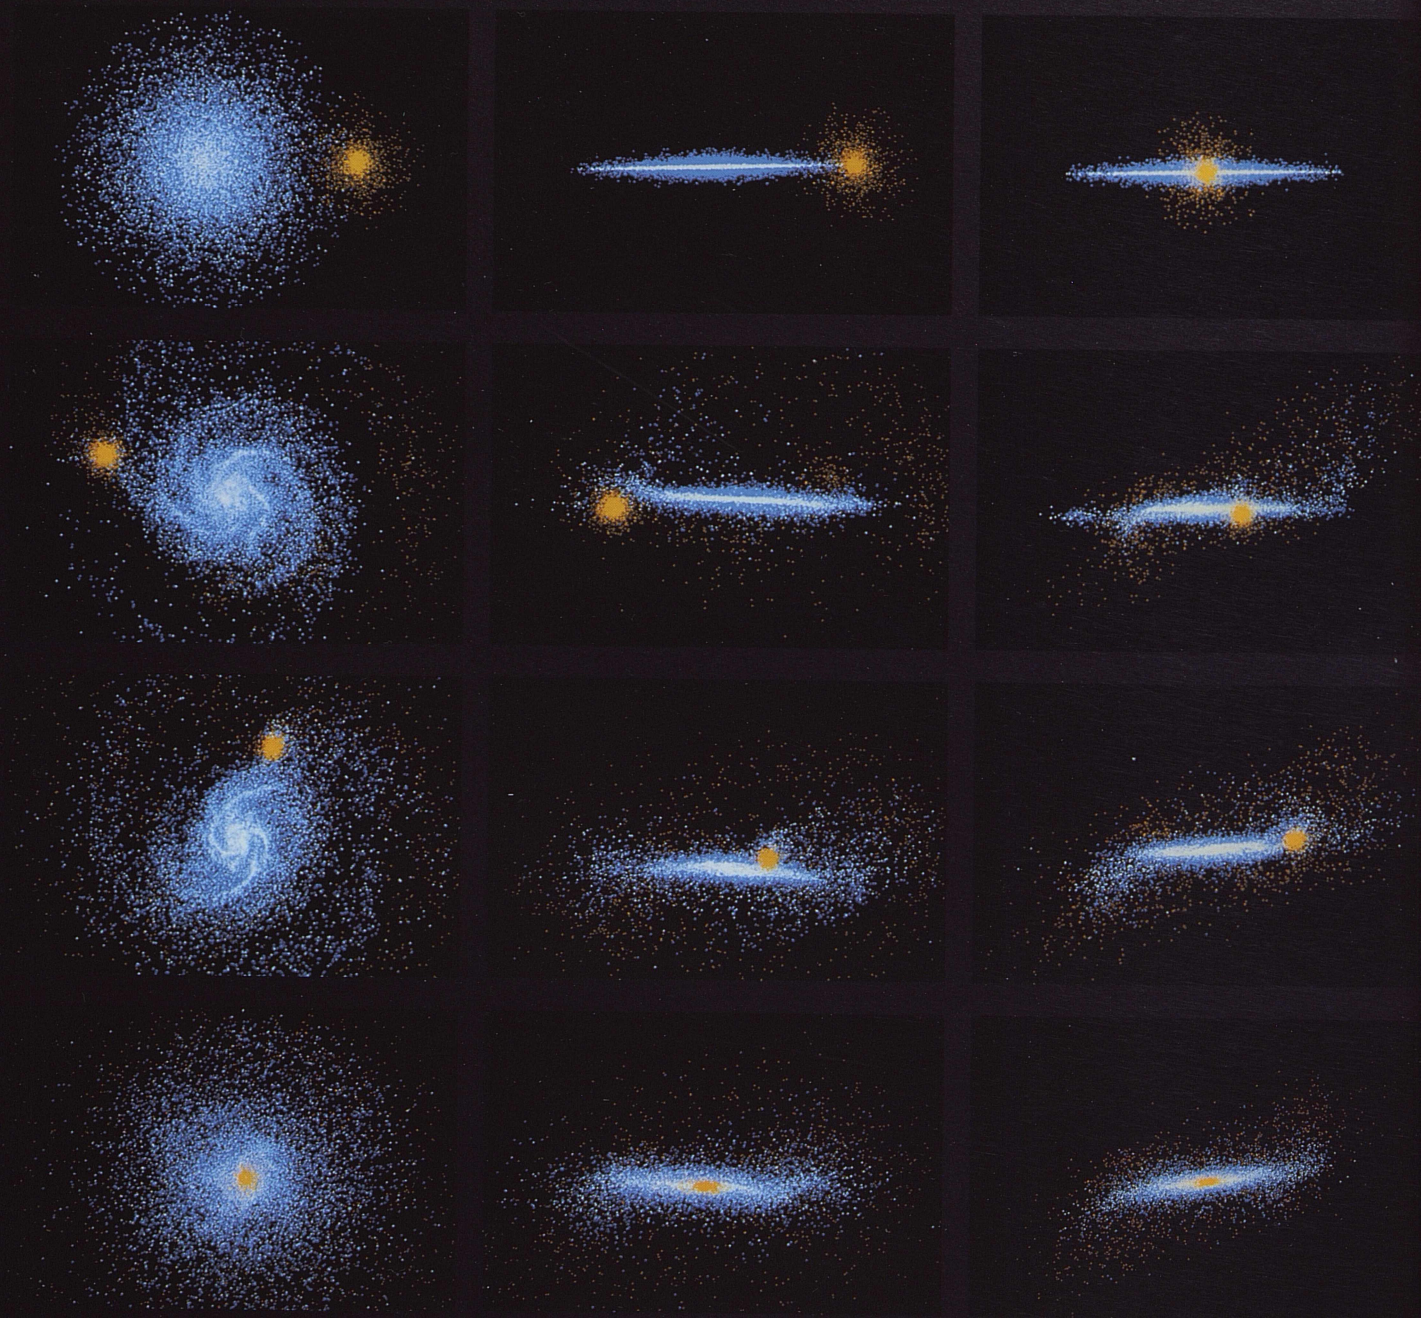

Each row of three frames depicts the disk-satellite interaction from three perspectives—face-on, edge-on and end-on projections. The top row is at time = 0 and each succeeding row represents an evolution of 800 million years. The disk mass is 90% stars (blue) and 10% gas (white). The satellite has ten times less mass than the disk, and this mass is entirely stars (orange); the model gives it no gas of its own. As gravity exerts itself, the satellite orbit gradually decays, and the larger galaxy consumes its smaller companion. "Note the strong response of the disk at intermediate times," says Hernquist, "especially in the gas phase, and the damage done to the vertical structure of the disk. Of special interest is the large mass of gas in the disk center at late stages of the encounter."

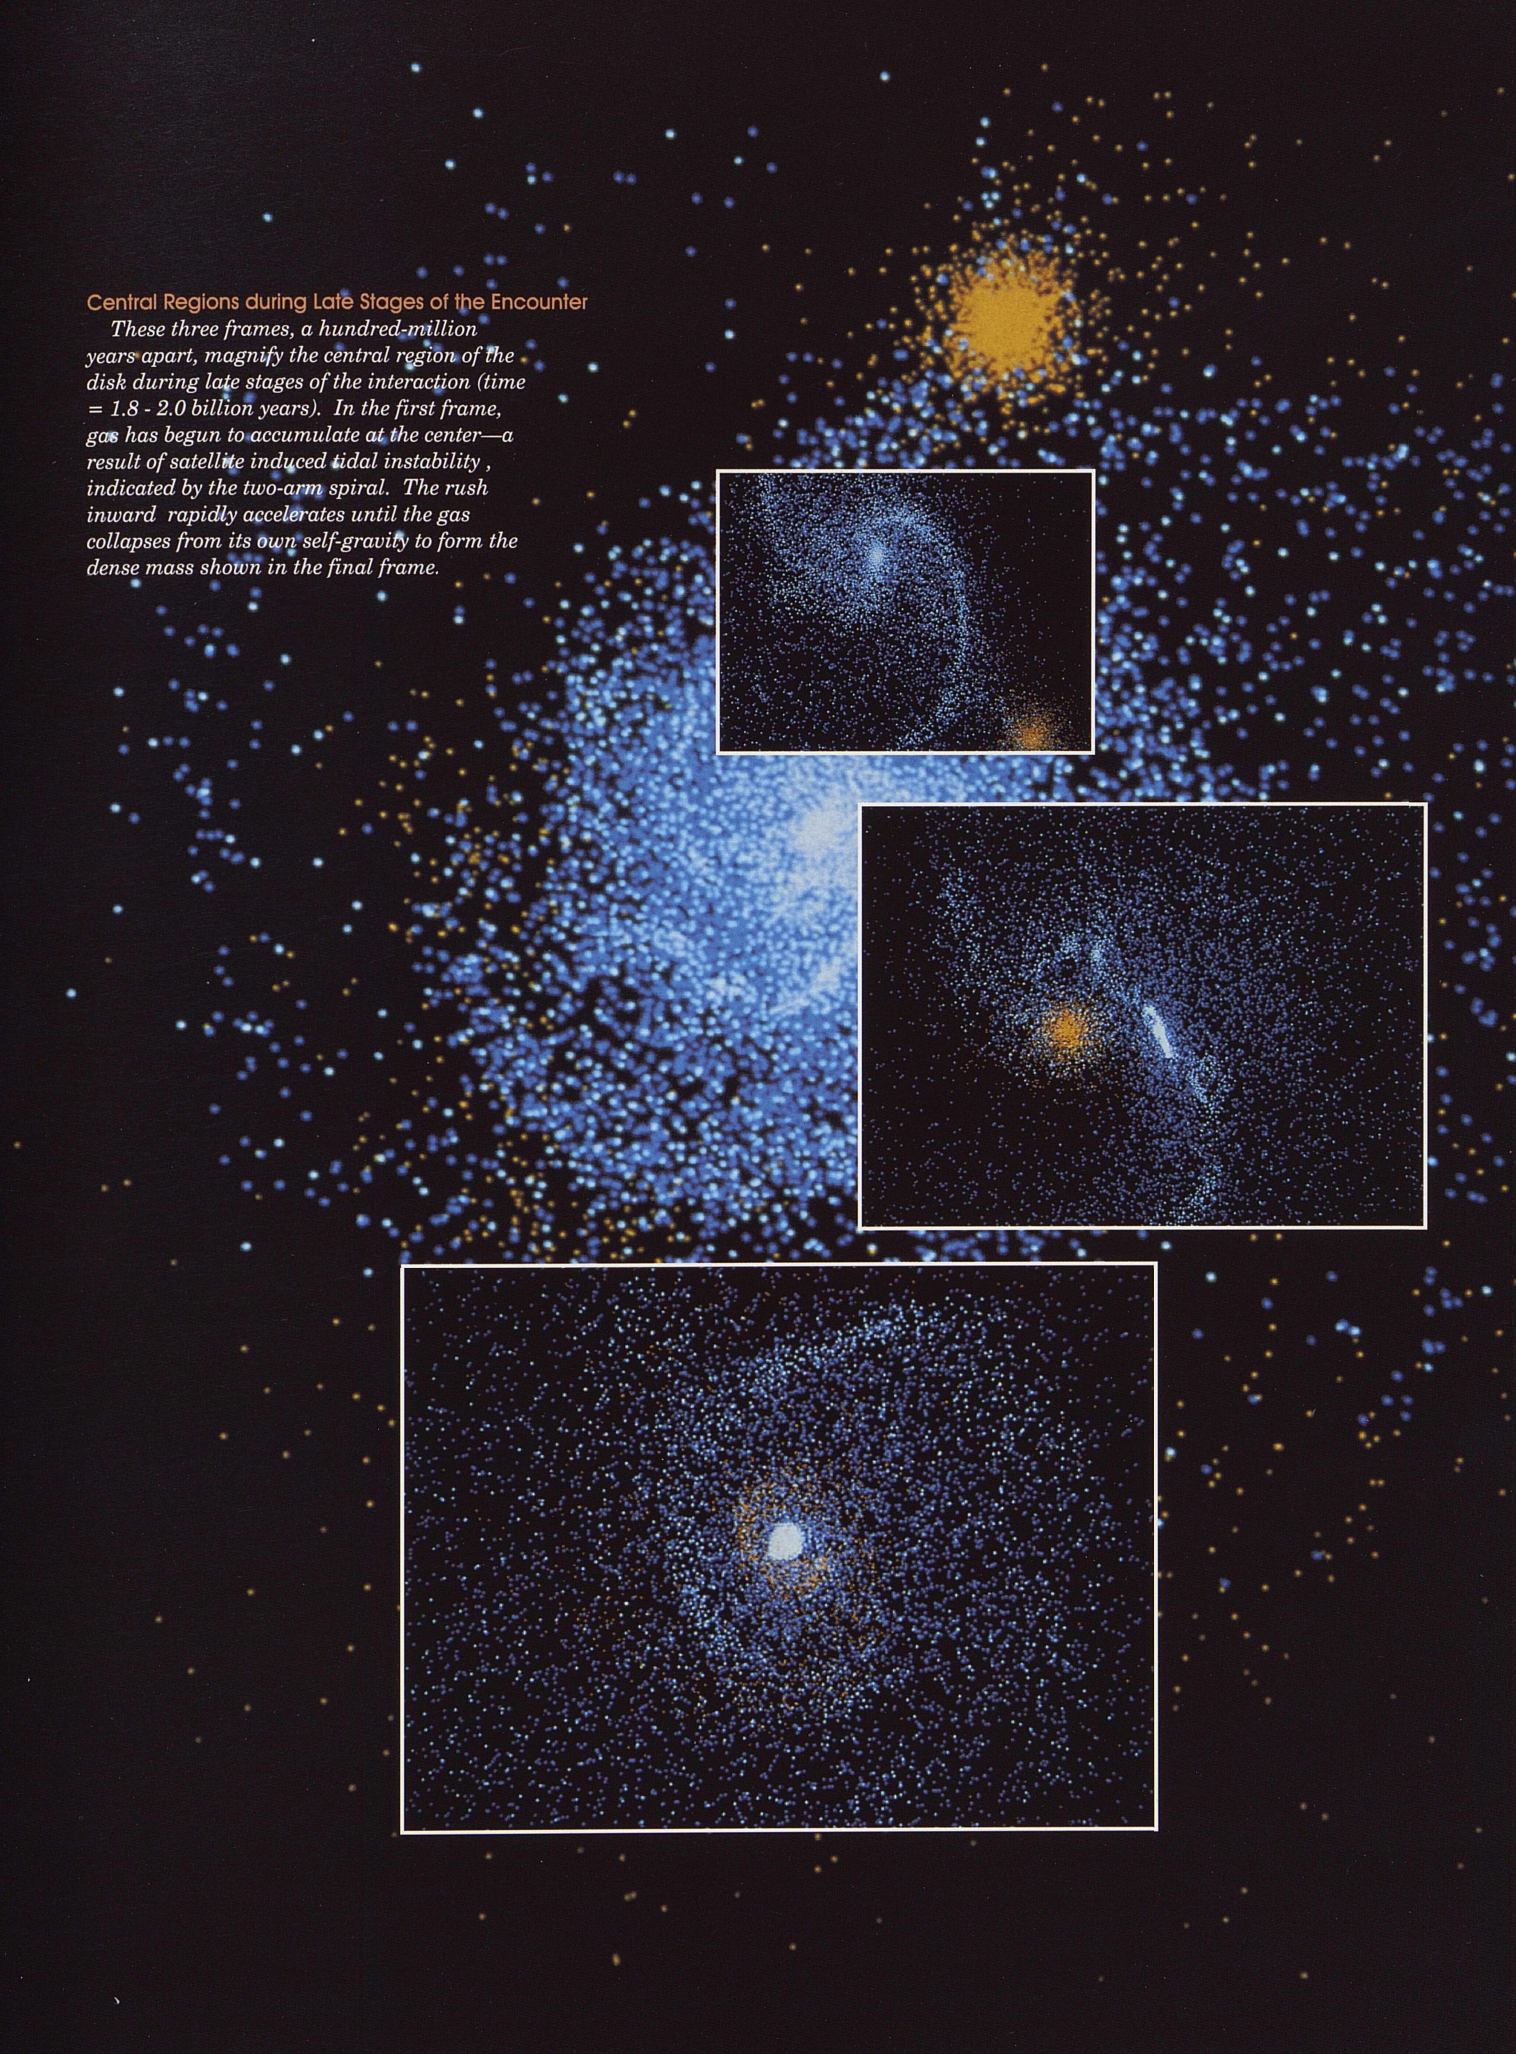

Central Regions during Late Stages of the Encounter

These three frames, a hundred-million years apart, magnify the central region of the disk during late stages of the interaction (time = 1.8 - 2.0 billion years). In the first frame, gas has begun to accumulate at the center—a result of satellite induced tidal instability , indicated by the two-arm spiral. The rush inward rapidly accelerates until the gas collapses from its own self-gravity to form the dense mass shown in the final frame.

Restless Ions

Molecular Dynamics Simulation of Multivalent β"-Aluminas

Gregory C. Farrington, University of Pennsylvania

Superionic Conductors

Though it didn't make the 6:00 news, a revolution got underway in 1967 when Yung-Fan Yao and James Kummer of Ford Motors reported the conductivity of sodium beta (β)-alumina. What they found was astonishing—a crystalline solid with very low electronic conductivity, virtually an insulator, had room-temperature sodium ion (Na^+) conductivity nearly as high as liquid electrolytes. Unlike metals, in which electricity flows by electron transport, electrolytes—battery acid for instance—conduct electricity through the movement of ions. In the vast majority of solids, however, the atoms are tightly bound in the crystalline molecular structure, not free to roam as ions.

Chemists had identified solid electrolytes before 1967, but they were relatively soft, unstable compounds like lead fluoride and silver iodide. Sodium β-alumina is an extremely hard ceramic, and its unique properties opened a range of potential new technologies. Gregory Farrington, who directs the University of Pennsylvania's Laboratory for Research on the Structure of Matter, describes it as "the liberation of electrochemical technology from the liquid state." The most immediate prospect, high-energy-density batteries, has been realized in a solid-state battery for heart pacemakers; and this technology still holds the promise of a viable electric car.

A Rich Family of Inorganics

In collaboration with coworkers at Penn and at the University of Uppsala in Sweden, Farrington has carried out an extensive program of research in the β-alumina family, focusing especially on beta double-prime (β")-alumina, a composition closely related to β-alumina with a slightly different structure. "These are fascinating materials for studying how changes in composition affect structure and properties," says Farrington. "The β-and β"- aluminas are not just one or two compounds; dozens and dozens of compounds can be synthesized from the same structure." In work with John Thomas of the University of Uppsala, Farrington demonstrated that the β"-aluminas have a rich chemistry, arguably unique among inorganic materials. The Na^+ ion in sodium β"-alumina can be exchanged for virtually any ion of +1, +2 or +3 valence in the periodic table. Some of these substitutions, furthermore, lead to luminescent properties with potential for use as lasers and opto-electronic devices.

Structural studies based on various diffraction techniques suggest that vacant sites in the conduction plane (see figure) give rise to the remarkable ion mobility of these materials. Diffraction techniques, however, are inherently limited in what they can tell about the structure of a mobile-ion system. "It's as if you wanted to know which squares on a checkerboard have checkers on them. X-ray diffraction could tell you, for example, that on the average two-thirds of the checkers are on red and a third on black. But you still don't know which squares are occupied. It doesn't tell you the specific local arrangements that lead to this average arrangement, and as with everything else, the specific arrangements are critical."

Molecular Dynamics with FUNGUS

To compensate for these limitations, Farrington has turned to molecular dynamics (MD) simulation—implemented on the CRAY Y-MP with FUNGUS, a program initially developed in England and adapted by Farrington and his coworkers to their purposes. With advances in computing power, MD has become especially useful in elucidating complex structure/function relationships in proteins and other biomolecules. In applying MD simulations to solid ionic conductors, Farrington exploits their capability to track the evolution of structure at the atomic level.

Initial work centered on validating the parameters required to initialize the MD simulation, the interatomic potential energies and the shape of the "simulation box"—the fixed number of molecular units chosen to approximate the continuity of the solid-state. The basic test is how well the simulations reproduce experiment. Results with sodium β"-alumina gave a persuasive indication of the power of the MD approach. Tracing the movement of

Gregory C. Farrington & Cathy Lane, a Ph.D. candidate at the University of Pennsylvania, with their model of sodium β"-alumina.

Na^+ ions over 1000 five-femtosecond (10^{-15} sec.) timesteps closely reproduced the ion distribution observed in experiments.

Encouraged by these results, Farrington's group applied MD to a more complex problem. Sodium-barium β"-alumina is a mixed-ion system, and it exhibits a peculiar phenomenon called the mixed-alkali effect. Its ion conductivity when high proportions of barium are mixed with sodium is much lower than either single-ion system. The simulation predicted very low conductivity when Ba^{++} replaces 80% of the Na^+, and experimental studies with the same material confirmed this finding.

Crystals of copper (Cu^+) β"-alumina under ultraviolet excitation. The combination of interesting optical properties and the ability to control and modify the chemistry suggests important potential applications for β"-aluminas as lasers, phosphors and optical components.

Predicting New Materials

Farrington's MD simulations have demonstrated the value of this approach, and they have also pointed up what could be accomplished with further refinements to the software and, especially, with more computing power. One of Farrington's objectives for MD with the β"-aluminas is to take advantage of this family's rich chemical versatility by substituting from the periodic table to simulate the behavior of new compounds. "Can we predict from the computer screen which compositions are most interesting to make in the real world? It's a significant challenge." The foremost obstacle, says Farrington, is the large amount of computer time required.

References:
B. Dunn, G. C. Farrington & J. O. Thomas, "Frontiers in ß"-Alumina Research," *Materials Research Society Bulletin* **XIV**, 22 (1989).
C. Lane, G. C. Farrington, J. O. Thomas & M. A. Zendejas, "Molecular Dynamics Simulation of Ion Transport in Na^+-Ba^{++}-β"-Alumina," *Solid State Ionics* **40/41** (1990), in press.

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

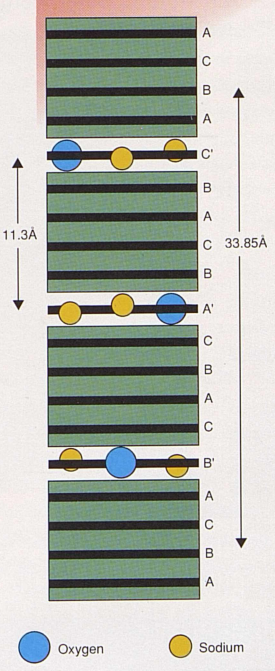

Unit-Cell and Conduction Plane of Sodium β"-Alumina

The remarkable ion transport properties of β"-*alumina are the result of its unusual structure. The solid lines at A, B, and C indicate close-packed blocks of aluminum, magnesium and oxygen, called "spinel blocks" from their resemblance to the mineral spinel ($MgAl_2O_4$). The three more open layers—about four angstroms (Å) thick (roughly 16 billionth of an inch)—are the "conduction planes." Ion motion occurs in honeycomb-like passageways around the oxygen atoms in these layers.*

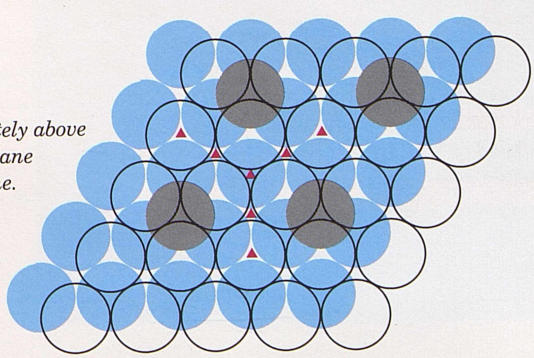

Open circles represent oxygens in the layer immediately above the conduction plane, dark circles are oxygens in the plane and light circles are oxygens below the conduction plane. The small triangles represent sites in the conduction plane that can be occupied by mobile ions.

The World According to GAUSSIAN

Developments in the Application of Molecular Orbital Theory

John A. Pople, Martin Head-Gordon & Douglas J. Fox
Carnegie Mellon University

"The vast majority of chemistry remains to be discovered."

Hungry Chemists

"As computers have increased in capability so have chemists' appetite for studying ever larger molecules," says Doug Fox. This appetite may well be insatiable, but that hasn't stopped Fox and members of Professor John Pople's group of theoretical chemists from creating tools to feed it. The GAUSSIAN system of programs, developed under Pople's guiding hand at Carnegie Mellon, has been a mainstay of quantum chemistry since the early 1970s, and it is one of the most-used applications running on the Pittsburgh Supercomputing Center's CRAY Y-MP. The two latest versions, GAUSSIAN 88 and, most recently, GAUSSIAN 90, incorporate major advances in numerical technique that, with a vigorous push from computers like the Y-MP, open the door on a new realm of unsolved problems in molecular chemistry.

Setting the Table: Direct SCF & QCI

One of the major improvements in GAUSSIAN 88 is direct computing of the Hartree-Fock self-consistent field (SCF). The SCF method of calculating interactions among the electrons in a molecule, the electron-correlation energy, involves an extremely large "list" of integrals. In the traditional approach, this list is computed one time and stored on disk. Direct SCF recomputes the list, or portions of it, as needed, avoiding the performance bottleneck of disk input/output. While the concept is not new, Martin Head-Gordon, who did his dissertation with Pople's group, developed a highly-vectorizable algorithm to implement it. As a result, says Pople, GAUSSIAN can be used to unravel the complexities of much larger systems. "The program is no longer limited by disk storage."

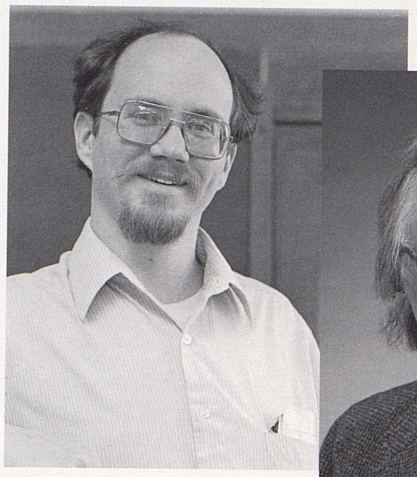

Quantum chemist Doug Fox occupies two separate but interacting orbits in the GAUSSIAN world. He works hands-on with GAUSSIAN code, in collaboration with John Pople's research group, exercising his background in configuration interaction theory to help maintain the program's track-record as a reliable performer in a variety of computing environments. As a scientific specialist at the Pittsburgh Supercomputing Center, he consults with chemists across the country who use the CRAY Y-MP, often drawing on his inner knowledge of GAUSSIAN to suggest research strategies.

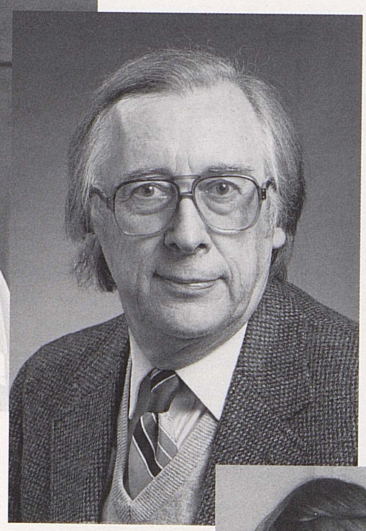

John A. Pople, John C. Warner University Professor of Natural Science, Carnegie Mellon University. "Extension to big systems is the interest now. The availability of the CRAY has enabled many people to do a wide range of computations, and it will permit application of the methods implemented in GAUSSIAN 88 [G1 Theory] to larger systems."

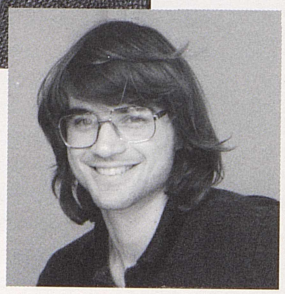

Martin Head-Gordon currently works at Bell Labs in New Jersey.

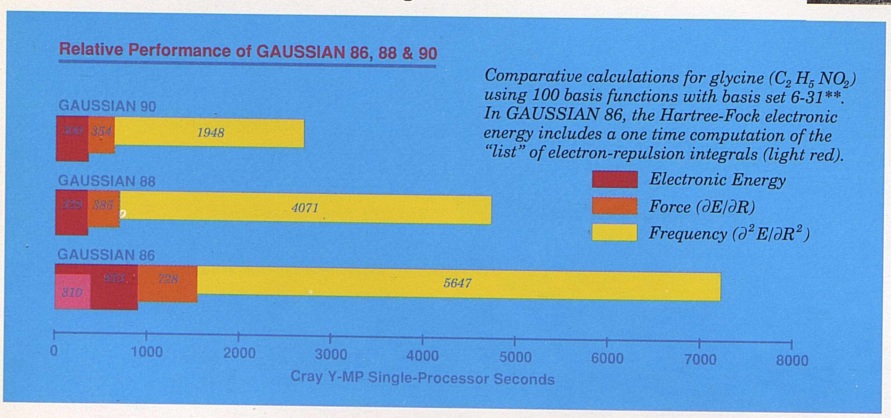

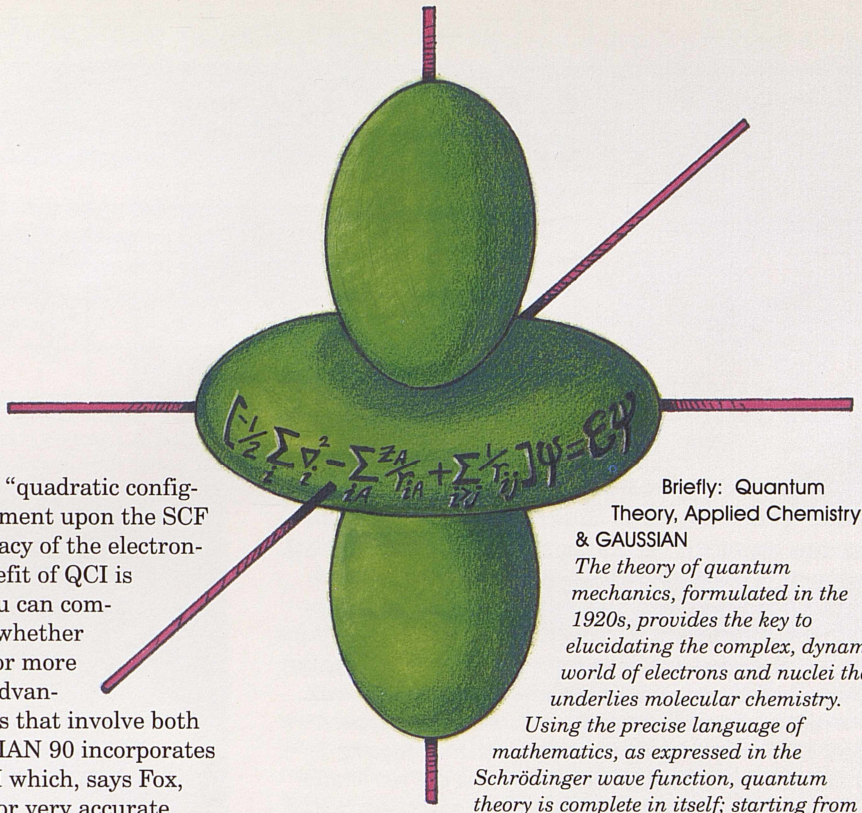

GAUSSIAN 88 also implemented "quadratic configuration interaction" (QCI), a refinement upon the SCF calculation that improves the accuracy of the electron-correlation energy. A principle benefit of QCI is size-consistency; how accurately you can compute the energy doesn't depend on whether there's only two as opposed to five or more atoms in the molecule. This is an advantage in investigating many reactions that involve both large and small molecules. GAUSSIAN 90 incorporates performance improvements for QCI which, says Fox, "should make this a powerful tool for very accurate structure and energy determinations."

A New Recipe for Good Molecular Cooking

The Pople group has used the recent enhancements to GAUSSIAN in developing a generalized methodology—a computational recipe—called "Gaussian-1 (G1) Theory." G1 is a composite method for accurately computing the total electronic and vibrational energy for a molecular system using (1) the Hartree-Fock ground-state geometry, (2) vibrational frequencies and (3) corrections to the electron-correlation energy. High-level corrections, such as QCI, are computed separately from other steps in the method, such as Hartree-Fock frequencies and Møller-Plesset perturbation theory, making G1 adaptable to a wide range of molecules.

In recent work, the Pople group applied G1 to a set of 31 experimentally-measured molecules and achieved their objective: energy results as accurate as experiment within two kilocalories per mole across the board. The CRAY Y-MP facilitated QCI calculations that met the target accuracy for several large-molecule multiple-bond systems, such as HCN and H_2CO. G1, says Pople, "achieves an accuracy of two kcals per mole systematically through all the cases where experimental data is well known. I think that's the first time that's been demonstrated. The technique reproduces essentially all the known experimental facts to that accuracy, and it can be used in a predictive manner for molecules for which experimental energy is not available."

References:
Martin Head-Gordon & John A Pople, "A method for two-electron Gaussian integral and integral derivative evaluation using recurrence relations," *Journal of Chemical Physics* **89**, 5777 (1988).
John A. Pople, Martin Head-Gordon, Douglas J. Fox, Krishnan Raghavachari & Larry A. Curtiss, "Gaussian-1 Theory: A General Procedure for Prediction of Molecular Energies," *Journal of Chemical Physics* **90**, 5622(1989).
Epigraph: from Warren J. Hehre, Leo Radom, Paul v.R. Schleyer & John A. Pople, *Ab Initio* Molecular Orbital Theory (N.Y.: Wiley, 1986), p. 3.

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

Briefly: Quantum Theory, Applied Chemistry & GAUSSIAN

The theory of quantum mechanics, formulated in the 1920s, provides the key to elucidating the complex, dynamic world of electrons and nuclei that underlies molecular chemistry. Using the precise language of mathematics, as expressed in the Schrödinger wave function, quantum theory is complete in itself; starting from only the atomic number of the atoms that comprise a molecule, using no experimental data, it can in principle predict most, if not all, chemical phenomena. Nature, nevertheless, has in large part resisted this form of interrogation, usually called ab initio *or "from first principles."*

The problem is that it isn't feasible to exactly solve Schrödinger's equation. There are simply too many variables, too many unknowns and—even with supercomputers—not enough computing power. By the 1960s, with the advent of computers, chemists found some useful shortcuts to full quantum theory. With "semi-empirical" approaches, they rely on experimentally-derived values to reduce the complexity of certain terms. And with various theoretical refinements, reasonably accurate ab initio *calculations began to become feasible—in limited cases.*

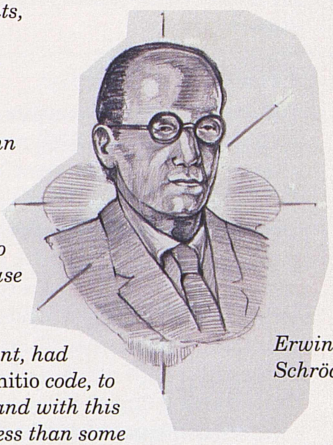

Erwin Schrödinger

GAUSSIAN originated, explains Doug Fox, when John Pople's research group in the late 1960s wanted a way to evaluate the accuracy of its semi-empirical research. "No one was doing ab initio *because it was felt to be too computationally expensive." Warren Hehre, then a graduate student, had the task of producing an* ab initio *code, to use largely as a benchmark, and with this code,* ab initio *began to cost less than some semi-empirical approaches. A fundamental theoretical refinement underlying GAUSSIAN is the Hartree-Fock "self-consistent field." Instead of a separate calculation for each electron's interaction with every other electron, as an exact solution would require, GAUSSIAN treats each electron as moving in a field that is, in effect, an average of the rest of the molecule.*

Path of Least Resistance

A Possible New Class of Metal-Oxide Superconductors

Thomas A. Albright, University of Houston

"They sought it with thimbles, they sought it with care;
They pursued it with forks and hope;
They threatened its life with a railway share;
They charmed it with smiles and soap."

The Philosopher's Stone Revisited?

Superconductivity is an amazing, almost magical phenomenon because it breaks the rules; electricity without resistance is having your cake and eating it too. And compared to the payoffs that await a commercially viable, room-temperature superconductor, transmuting a base metal to gold would be a cheap trick. Beginning in 1986 with Bednorz and Müller's Nobel Prize winning discovery of superconductivity in lanthanum copper-oxide, new laboratory-synthesized superconducting materials have excited scientists around the world and spurred a diligent effort to understand how these materials work and to solve the problems of making them practical.

There have been discouraging words, mainly "critical current density"—superconductivity in the new materials breaks down when current rises above relatively low levels. Some research suggests this problem can be surmounted, and the overall superconductivity effort shows no serious signs of flagging. In this country, four major corporate research labs and numerous smaller teams at universities are advancing the quest, accruing knowledge step-by-step. Still, it remains an open question: Does the enormous potential represent real possibility? Or a twentieth-century version of the philosopher's stone?

A Possible New Direction

Current research centers almost exclusively on a related group of copper-oxide materials, the only materials to display "high-temperature" superconductivity—above 77° K, the temperature of liquid nitrogen. Time-consuming "shake and bake" laboratory methods substitute and mix elements based on the recipe given by these proven high-T_c compounds. "There are probably about 2000 groups making copper-oxide materials," says Thomas Albright, "and sooner or later a finite limit will be reached for that domain of elemental composition. The issue's still going to be *In what parts of the periodic table other than copper should one look for superconductors*?" With this in mind Albright has identified another class of metal-oxides, another recipe for experimentalists to cook up, which he believes could lead to new high-T_c compounds.

Since 1988 Albright has collaborated with experimentalist Paul Chu and his colleagues at the Texas Center for Superconductivity. Chu pushed T_c to an astonishing new high of 93° K in 1987 when he discovered yttrium-barium-copper-oxide ($YBa_2Cu_3O_7$), the "1-2-3" superconductor. With detailed calculations of electron band structure, Albright has been puzzling out some of the tough questions associated with 1-2-3 and the related bismuth and thallium-based copper-oxides that have an even higher T_c. What structural mechanisms explain their unique properties? With these problems on his mind, Albright went browsing one day in a book of solid-state structures.

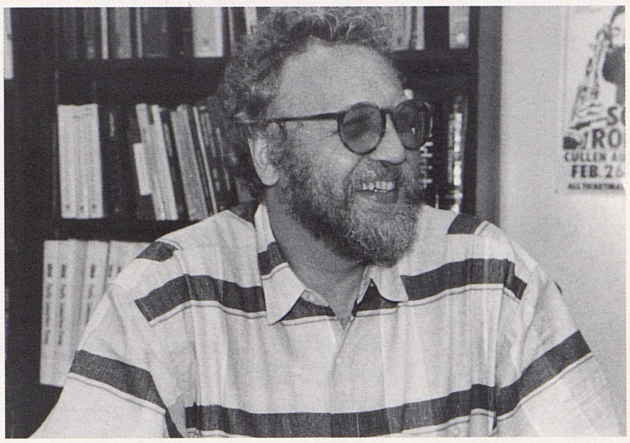

Thomas Albright and his coworkers carry on an extensive project at the Pittsburgh Supercomputing Center examining some of the formidable problems posed by the chemistry of transition metal compounds. They use John Pople's GAUSSIAN quantum chemistry software (see previous article) to examine structure at the molecular level, and they do electronic bandstructure calculations to predict potential new solid-state compounds.

"This particular group of structures jumped off the page." He had come upon Aurivillius compounds, named for Danish solid-state scientist Bengt Aurivillius, who discovered them in 1921. Though Aurivillius compounds were extensively studied in the 50s and 60s, Albright's work on high-T_c copper-oxides led him to see them in a fresh light. "The crystal structures of Aurivillius phases are strikingly similar to the bismuth and thallium-based superconductors." Both have vacancies that result in layered planes in the molecular structure, a feature that appears to play a key role in all the high-T_c superconductors.

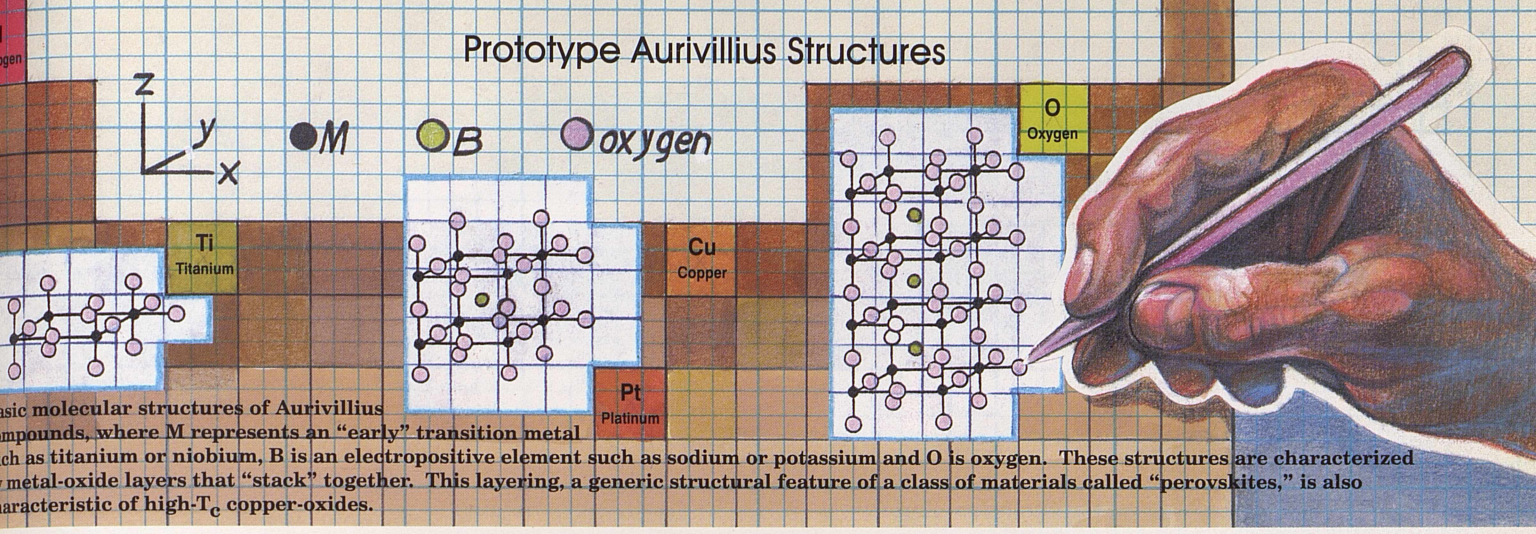

asic molecular structures of Aurivillius mpounds, where M represents an "early" transition metal ch as titanium or niobium, B is an electropositive element such as sodium or potassium and O is oxygen. These structures are characterized metal-oxide layers that "stack" together. This layering, a generic structural feature of a class of materials called "perovskites," is also haracteristic of high-T_c copper-oxides.

Calculating the Electronic Bands

For a metal-oxide to show superconductivity, explains Albright, it must have d-orbital electrons. Because of their relative high energy, d-electrons have freedom to move around in the molecular structure. Aurivillius compounds synthesized to date have no d-electrons, thus no superconductivity, and they are not even normal conductors of electricity. Albright suggests, nevertheless, that the Aurivillius "recipe" has a wide tolerance for substituting elements, and it could be possible to synthesize Aurivillius forms with d-electrons.

Using Pittsburgh's CRAY, he did "tight-binding" calculations on several titanium-oxide (TiO_4) layered Aurivillius structures with 1 to 5 d-orbital electrons. These calculations, which give detailed indications of how electrons interact within the molecular structure, confirm Albright's initial insight. "The Fermi surfaces with three d-electrons and, with slight structural modifications, one d-electron, are very close in shape to those of the copper-oxide superconductors." Albright concludes that a layered Aurivillius structure with a d-electron count close to one or three "might give rise to a new class of superconductors."

Albright's quantum theory-based calculations point the way toward potentially fruitful laboratory work and stand to reduce the costly "needle in a haystack" process involved in searching for superconducting materials. "We're proposing a structure that we're reasonably sure can exist based on what's known about the materials." For practical purposes, such calculations require supercomputing. Albright estimates that the CRAY does band-structure computations 20 times faster than the University of Houston's VAX 8650. "If you're looking at electronic-structure, you need to do a gigantic number of calculations. It may be technically feasible on a smaller computer, but the clock-time involved and the need to share the computing resource would make this kind of study untenable."

References:
Kyeong Ae Yee, Thomas A. Albright, Dongwoon Jung & Myung-Hwan Whangbo, "A Possible New Class of Metal Oxide Superconductors," *Angewandte Chemie, Intl. Ed.* **28,** 750 (1989).

Epigraph: from "The Hunting of the Snark," Lewis Carroll.

This work is supported by the Robert A. Welch Foundation, the Petroleum Research Fund as administered by the American Chemical Society, the Texas Center for Superconductivity at the University of Houston, the Defense Advanced Research Projects Agency and the State of Texas.

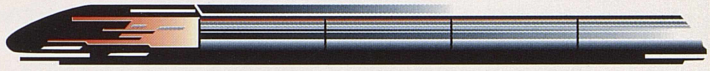

Superconductivity—What is it, and why is it super?

Normally, materials resist electrical current. "Flowing" electrons induce vibrations in the microscopic structure, raising the temperature and in effect dissipating some of the electrical energy as heat. Lowering the temperature reduces the tendency to vibrate; and in a normal metal, resistance falls with lower temperatures but never disappears. If a material is superconducting, however, below a critical temperature, T_c, the vibrations in the molecular structure create an attractive force that links electrons to each other in pairs and resistance vanishes—current can flow forever.

H. Kamerlingh Onnes discovered superconductivity in 1911. He found that mercury became superconducting at about 4° Kelvin, 4° above absolute zero (-273° Centrigrade). Recently discovered high-T_c materials raise the prospect of a room-temperature superconductor, a material with potential limited only by imagination. Lightweight, powerful magnets could drive electric motors that would make batterypowered cars practical, ending our environmentally destructive dependence on fossil fuel. Trains that "levitate" on magnetic rails could travel hundreds of miles an hour, providing an alternative to inter-city air travel. Computers built with superconducting circuitry (see next article) could run 100 times faster than today's fastest supercomputer, and communications lines could carry 100 times more data than the most advanced optical fibers.

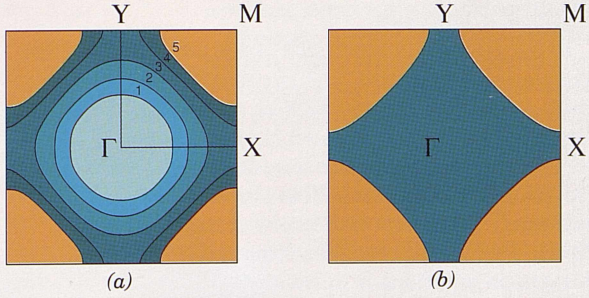

(a) (b)

Fermi Surfaces for Titanium Oxide Aurivillius Phases and the "1-2-3" Superconductor

The Fermi surface gives a geometric representation of the relation between molecular structure and electron energy levels. In particular, it indicates whether the energy-level is high enough to allow the "flow" of electric current. Here, (a) shows Fermi surfaces of the xy band for titanium oxide (TiO_4) with one to five electrons in the d-orbital. The three-electron surface is nearly indistinguishable from the Fermi surface for the x^2-y^2 d-orbital of yttrium-barium-copper-oxide, the 1-2-3 superconductor, shown in (b).

A New Phase

New Coherent States in Periodic Arrays of Ultrasmall Josephson Junctions

Jorge V. José, Northeastern University

A New Kind of Superconductivity

"We had no clue before we did the calculations. What we found was completely unexpected." This is how Jorge José describes his computations applying quantum theory to simulate the low-temperature behavior of an array of Josephson junctions. His work has identified a potential new state of matter, a superconducting phase with different properties than "normal" superconductivity. The transition to this phase occurs at a temperature below T_c, the critical temperature at which paired electrons—the mechanism of superconductivity—first occur.

Taking the lead from José's work, three laboratories are doing experiments aimed at physically verifying the new phase, an instance—like the discovery of the Josephson effect itself and many elementary particles—where theory tells experimentalists where to look. But in this case, theory needed a boost from computational science. "Supercomputing is crucial," says José. "The transition occurs outside the high and very-low temperature regions where asymptotic perturbative calculations can be done." The equations that describe quantum behavior in the temperature range in question can be solved only with iterative methods—calculations that successively narrow the range of error to converge on a solution. Applying iterative methods to this problem would be completely impractical without supercomputing capability.

Josephson Junctions

Brian Josephson received a Nobel prize for his 1962 calculations predicting that superconducting electron pairs could "tunnel" through an insulating junction. Experiments soon confirmed his predictions, and important practical applications have followed, especially in metrology; a Josephson-effect alternating current allows extremely precise measurement of electromagnetic fields. Perhaps the most exciting potential application of Josephson junctions is in computer microcircuitry. Research suggests the possibility of switching times 100 to 1000 times faster than the fastest silicon semiconductors. IBM sponsored a 15-year program, dropped in 1983, that overcame a number of technical obstacles. Japanese firms have continued work in this area and have reported a switching time of 4.2 picoseconds (.0000000000042 seconds), probably the fastest yet achieved.

José's research focuses on arrays of Josephson junctions, and he examines effects that have become feasible only recently. It is now possible to fabricate arrays of "ultrasmall" junctions, junctions in which the "feature size"—the facing area of the superconducting electrodes—may be as small as .02 microns (a millionth of an inch) square. A Josephson junction, explains José, is like a capacitor: the junction area and width determine its capacitance. Low capacitance facilitates fast Josephson circuitry but also increases the "charging energy," energy required to overcome the quantum fluctuations that occur when a negatively-charged electron pair crosses the junction. As the charging energy increases, the critical temperature decreases and the superconducting phase tends toward instability.

A Quantum-Induced Phase Transition

José examined this effect in 1984, and his analytic calculations, based in renormalization group theory, suggested that superconductivity broke down, even at temperatures below T_c, when the capacitance was low enough for charging energy to be a factor. These calculations didn't allow a direct look at the temperature range where this transition appeared. Working in collaboration with Alan Goldman of the University of Minnesota, Laurence Jacobs of MIT and Mark Novotny of IBM (now on the faculty at Florida State University), José developed a Monte Carlo algorithm tailored to the problem, and in 1988 an extensive series of computations completed at the Pittsburgh Supercomputing Center revealed the previously undetected phase.

"We found a transition from the known superconducting phase to a different superconducting phase, one completely dominated by these quantum fluctuations." The new phase differs from known superconductivity in several ways. The transition at T_c, explains José, is a second-order transition; in thermodynamic terms, it doesn't have latent heat. When you heat water to the boiling point, for instance, there is a time lag before it starts evaporating. The quantum-induced phase transition, T_{QUIT}, revealed by José's calculations is—like boiling water—a first-order transition. "That was very unexpected and very significant."

References:

Laurence Jacobs, Jorge V. José, M. A. Novotny & A. M. Goldman, "New coherent states in periodic arrays of ultrasmall Josephson junctions," *Physical Review B* **38**, 4562 (1988).

Jungzae Choi & Jorge V. José, "Quasiparticle Effects in Quantum-Induced Transitions in Superconductors," *Physical Review Letters* **62**, 1904 (1989).

This research is supported by the National Science Foundation.

Quantum Tunneling and the Josephson Effect

Current can flow across an insulator separating two conductors even if—by classical physics—the electrons don't have enough energy to cross the barrier. This phenomenon, predicted by quantum theory, is called "tunneling." It can be understood in relation to the wave-like character of electrons.

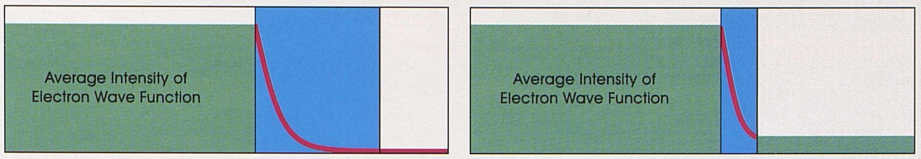

In classical physics, an electron encountering an insulating barrier is like a marble rolling uphill; it can get to the other side only if it has enough kinetic energy to roll over the top. Quantum mechanics, on the other hand, says the marble might get to the other side even if it can't get over the top. The probability of this happening depends on the intensity of the electron wave, which at any point in space is simply the wave amplitude squared. At a barrier, the intensity declines as a smooth exponential curve, and if the barrier is thin enough, a current will appear on the other side. The Josephson effect, basically, is tunneling between superconductors—paired rather than single electrons.

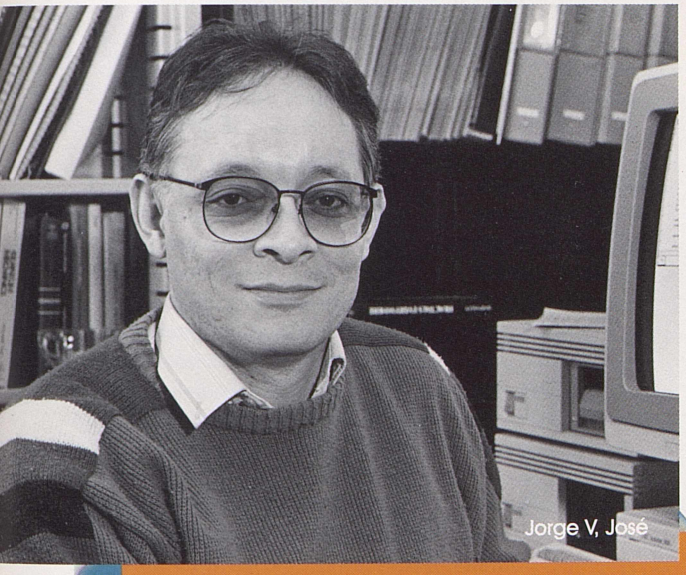
Jorge V. José

José is continuing his investigation of the new superconducting phase, in collaboration with Jungzae Choi of Oxford University, with simulations that include the energy dissipation that occurs when normal electron tunneling mixes with Josephson tunneling—usually the case in reality. He anticipates that increases in computing power will also allow him to model disordered arrays such as occur from film-deposition techniques, where the superconducting grain-sizes are not evenly distributed.

Dipole Vortex-Pairs v. Temperature

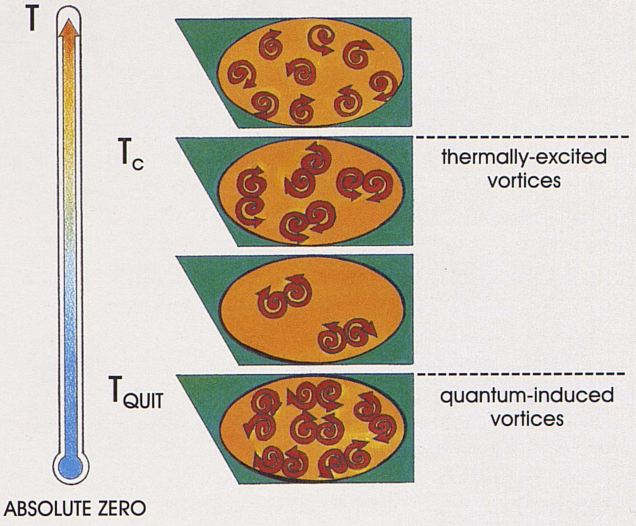

José's simulations are based in "BKT theory"—a two-dimensional model of superconductivity that is particularly useful in describing thin superconducting layers like Josephson junction arrays. José has extended BKT theory—named for solid-state theorists Berezinskii, Kosterlitz and Thouless—by accounting for charging-energy effects. Basically, this theory describes the superconducting transition in terms of opposite-spin vortices of electromagnetic energy—a "dipole vortex pair." Vortex interaction increases as temperature decreases, until at critical temperature, T_c, the vortices bind in pairs.

While the T_c transition is dominated by thermal fluctuations, José's simulations showed a transition below T_c to a new superconducting phase dominated by quantum fluctuations. "As you start cooling the system, the number of vortex pairs decreases until an intermediate temperature below which the number of vortex pairs begins to increase." José calls this second critical temperature T_{QUIT} for quantum-induced transition.

Segment of a million-element Josephson junction array fabricated using photolithography by Sperry Corporation, St. Paul, Minnesota. A thin layer of silicon (light blue) insulates between two layers of niobium etched to create a cross-pattern of electrodes–dark blue in the top layer and white in the underlayer. The small squares where the dark blue and white crosses overlap are the junctions. The "feature size" is 8 microns (.0003 in.) on a side, larger than the "ultrasmall" arrays involved in José's current research.

Wood Vibrations

Compliances of Wood for Violin Top Plates

Robert T. Schumacher, Carnegie Mellon University

Isn't it good, Norwegian wood?

Centuries of craftsmanship and the human ear have settled the issue: the best material for the top plate of a violin is *picea abies*—Norway spruce. This wood is distinguished by its low density and its high "mechanical anisotropy," explains physicist Robert Schumacher. It's about 20 times stronger across the grain than with the grain.

Many of the oldest stands of Norway spruce in Europe have been obliterated by the *Waldsterben* (forest death) phenomenon, a combination of the effect from air pollution—which significantly modifies plant nutrition and soil chemistry—and natural stresses. During the 1980s, 20 to 25% of European forests have been moderately to severely damaged as a result. The potential scarcity of high-quality Norway spruce has prompted some violin makers to experiment with other materials, such as epoxied carbon fibers. Can these materials duplicate the acoustical characteristics of a top plate carved from Norway spruce? Schumacher, a life-long violinist who has done extensive research in acoustics, enlisted CRAY supercomputing to investigate this question.

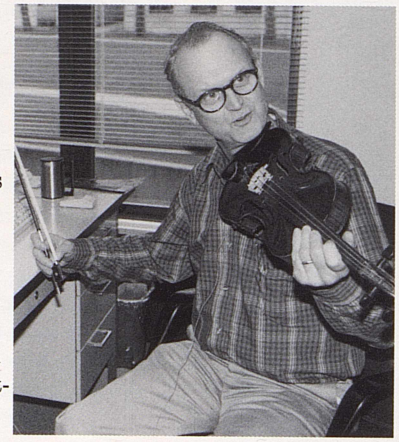

Robert T. Schumacher

Elastic Compliance

For centuries, violin makers have tested top-plate wood by bending it both ways and twisting it. In effect, they judge the wood's vibrational characteristics according to three of its "elastic-compliance elements"— experimentally derived constants that give the ratio of elastic deformation to applied force. Because materials like Norway spruce are stronger in one direction than another, their compliance differs depending on the direction of the deforming force. A violin top plate, furthermore, is arched, not planar, and therefore—according to well-developed theory—it takes nine, not three, elastic-compliance elements to fully describe it.

"One purpose of my research," says Schumacher, "was to see how important these six other elastic constants are. In other words, do violin makers— including modern scientific ones—fool themselves by relying on only three?" Some recent finite-element modeling of top plates and whole violins gives reasonable agreement with measured vibrational characteristics using only three or four elastic constants, suggesting that the neglected compliances don't matter. CRAY computing capability allowed Schumacher to include all nine constants and to do extensive calculations analyzing to what extent top plate normal mode frequencies depend on each one of them.

Computing with ABAQUS

Schumacher used the ABAQUS finite-element package to compute normal mode frequencies for a flat top plate and several different arched plates. He approximated a top plate's complex arched shape with a three-dimensional "stadium model"— basically a rectangle capped with semi-circles. Though much simplified from a real violin, this model made it relatively easy to vary the arch, the characteristic which in theory necessitates additional constants, and allowed Schumacher to investigate the primary question with a practical amount of computing.

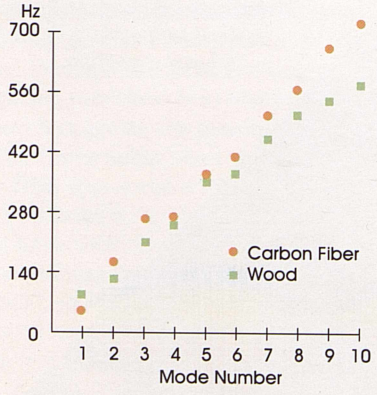

Normal Mode Frequencies for Norway Spruce and Epoxied Carbon Fiber

The "Stadium" Model

This model used over 200 elements and over 1000 nodes and required about five minutes of CRAY time to compute normal mode frequencies. Analyzing the sensitivity of the normal mode frequencies to the nine compliance elements tested required 19 separate runs per top-plate configuration. "Within the practical limits of available funding, I could not have done this work without a CRAY or comparably powerful computing system."

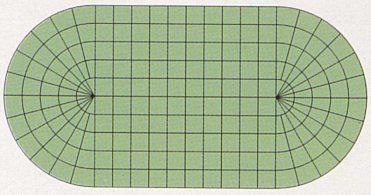

It appears that violin makers know what they're doing. "When you cut a piece of Norway spruce into the configuration of a top plate, the three elastic constants relevant to a flat or slightly-curved plate are sufficient within a few percent." This implies that man-made composites, which tend to have higher symmetry properties (hence fewer compliance elements) than wood, are feasible substitutes. To further test this question, Schumacher used measured elasticity parameters for a carbon fiber-epoxy top plate to calculate the normal mode frequencies of his stadium model. The results compare well with Norway spruce.

References:
Robert T. Schumacher, "Compliances of wood for violin top plates," *Journal of the Acoustical Society of America* **84**, 1223 (1988).

This research was supported by the National Science Foundation .

"The violin is the most astounding acoustical phenomenon ever perpetrated by and on the human nervous system. I doubt we will ever solve its mysteries completely."
— Carleen Maley Hutchins, master violin maker, recipient of the first Silver Medal (1981) awarded by the Acoustical Society of America.

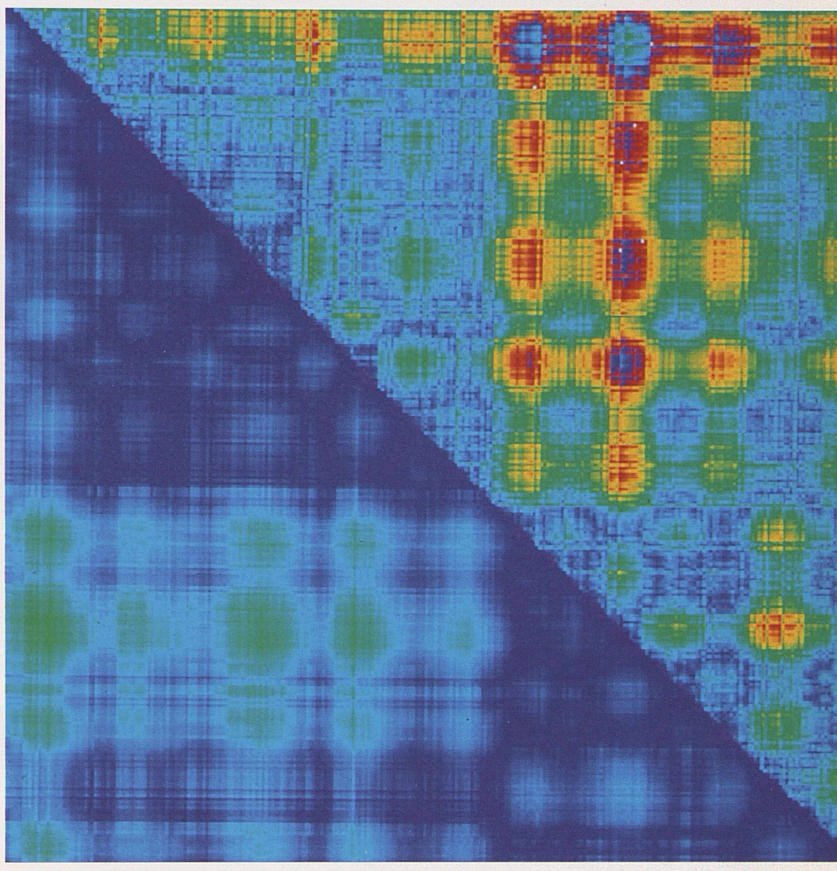

Violin D String Bowed Upward: Norm-Difference Plot

In current research, Schumacher is examining subtle irregularities in the waveform produced when a violin string is bowed. Such irregularities differentiate the sound of a real violin from computer-generated simulations. To visualize these phenomena, slight shifts in frequency and amplitude from period to period, Schumacher has adapted a plotting technique, suggested by research in dynamical systems, which he calls a "norm-difference plot."

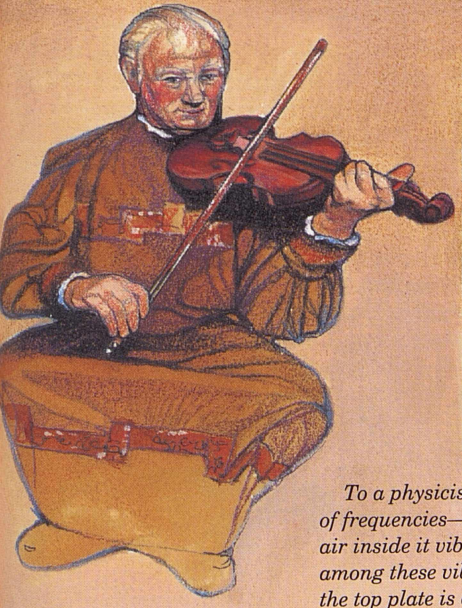

The first three normal modes of a vibrating string

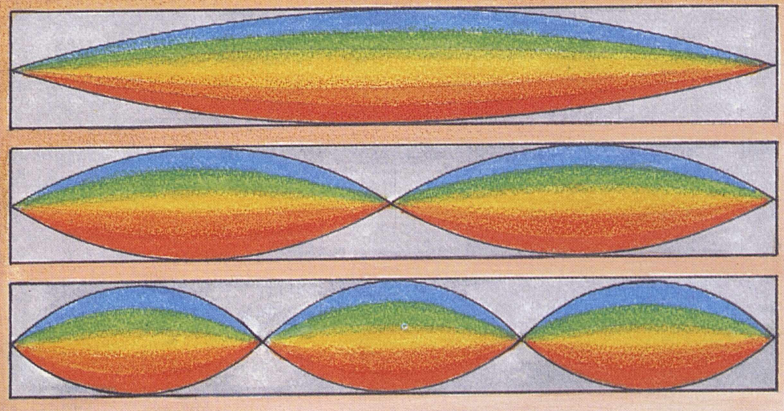

To a physicist, a violin is a vibrating system. Drawing a bow across the string makes it vibrate at a number of frequencies—its "normal modes." The string alone would be nearly inaudible, but the violin body and the air inside it vibrate in response to the string; and the violin's distinctive rich sound results from the resonance among these vibrating elements. To create a high-quality violin, every element in this system must be right, but the top plate is crucial. If it's too thick or thin or stiff, it won't vibrate properly, and the sound will be flawed.

Stress Memory in Heavy Fluids

Numerically Stable Finite-Element Methods for Viscoelastic Flow

Robert C. Armstrong & Robert A. Brown, Massachusetts Institute of Technology

"What! You think that Isaac Newton told a lie?
Where do you hope to go when you die?"

What Newton Didn't Know

Polymers weren't around in 1687, so it's hard to fault Isaac Newton for not thinking about them when he formulated classic fluid mechanics. His basic equation defines a fluid's viscosity, the molecular adhesion that exerts a friction-like resistance to stress, as a simple linear proportion between the stress and the rate of the resulting strain. The equation works well for most liquids and gases, but along about 1930 chemists found ways to build long chains of repeating molecular units—polymers. The wondrous new materials created this way, rayon, nylon and many others, have molecular weights ranging from 10,000 to 100 million—extremely large molecules. Unlike "thin" liquids and gases, polymers respond elastically to stress, recoiling like a rubberband when let go, and in many circumstances they exhibit this "memory"and a non-linear relation between stress and strain rate. In short, they totally defy Newton's equation.

By about 1950, polymers spawned a new field of engineering science—non-Newtonian fluid mechanics. "Polymers, in all shapes and forms, including slurries and suspensions," says Robert Brown, "are non-Newtonian. If you want to predict polymer flows—for injection molding of parts, printed circuit boards, whatever—you must be able to model viscoelastic flows." Brown and his colleague Robert Armstrong conduct a broad research program in non-Newtonian fluids, combining laboratory work and theoretical studies. One of their main objectives is to develop computer models for the complex flows involved in polymer processing. In computations at the Pittsburgh Supercomputing Center, they have achieved a fundamental breakthrough.

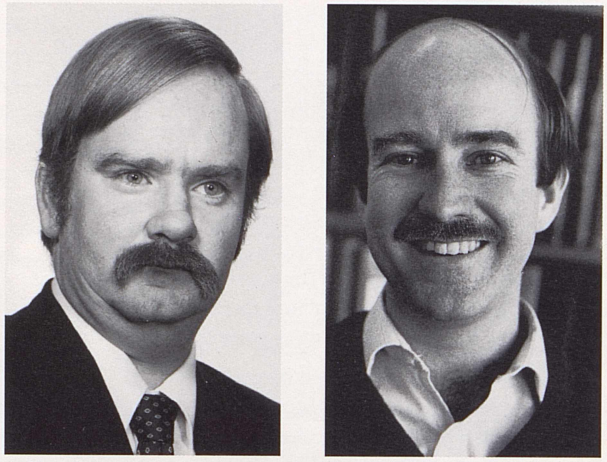

Robert A. Brown (left) & Robert C. Armstrong are professors in chemical engineering at MIT. Brown heads the department. "You can't do these problems without a supercomputer," says Armstrong. "Either you have the supercomputer and you do it, or you don't do it." Brown estimates that their extensively vectorized and chained EEME code runs at about 310 megaflops, near maximum single-processor performance on the CRAY Y-MP.

The Future of Polymer Process Design

The properties of polymer products depend on the shaping and forming process as well as intrinsic molecular structure. "The place we'd like to get 10-15 years down the road," says Armstrong, "is to have an engineer work with a chemist to design a molecule and a process that are matched to give a product with the desired

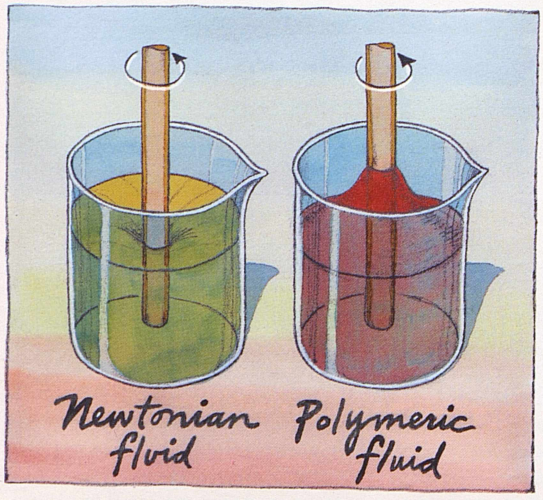

Newtonian & Non-Newtonian Fluids

Fluids come in two flavors. Newtonian liquids and gases obey Newton's law of viscosity, which mathematically expresses a fairly intuitive principle: Some fluids resist force more than others—water, for instance, is easier to stir than cream; and for a given applied force (stress), the slower a fluid flows in response (strain rate) the more resistance (viscosity) it has. This proportional (linear) relation between stress and strain rate, by definition, makes a fluid Newtonian.

Polymers don't obey Newton. They're more complex, more non-Newtonian. "For example," says Robert Brown, "if you pull a spoon across a bucket of Newtonian fluid, when you stop the spoon the fluid stops, and it doesn't exert any stress on the spoon. But a non-Newtonian fluid remembers where it came from, and it wants to go back. If you hold the spoon, you'll feel the fluid pull."

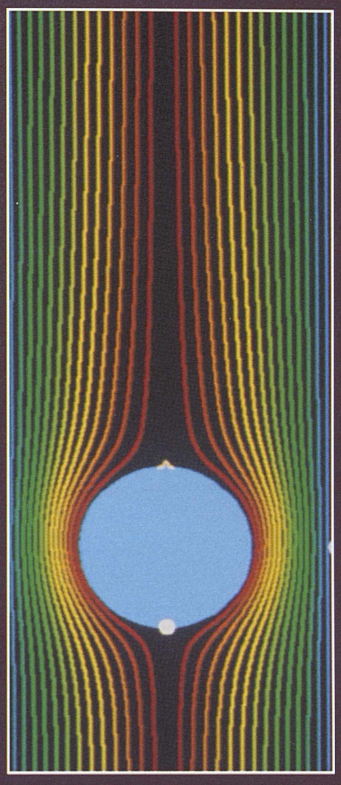

Streamlines (left) and stress contours (right) computed using EEME for a sphere falling through a tube filled with non-Newtonian fluid. For the streamlines, color indicates proximity to the central axis, red closest and blue at the walls. Similarly for stress, red is maximum, going toward blue at the minimum.

The dramatic downstream wake shows how polymer fluids "remember" stress, like a stretched rubberband, for a considerable distance. Steep stress gradients form along the surface of the sphere (closely packed contours). "This boundary layer structure," notes Armstrong, "is what demands supercomputing in the end. Even in smooth geometry, unless you put in enough elements to resolve that structure, you can't get converged solutions."

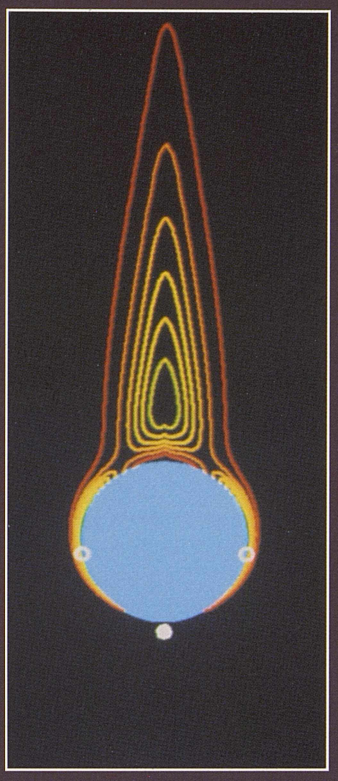

properties." Accurate design tools can also boost cost-effectiveness. "For a given shaping die, for example, you want to pump material through as fast as you can. But if the flow begins to oscillate—become time-dependent—the product will look bad, so you turn the flow-rate down. If you can design a flow geometry that allows you to extrude faster and still get the same shape, you get a better return on your capital investment."

The Explicitly Elliptic Momentum Equation

Until recently, there had been scant progress in numerical modeling of these flows. Armstrong and Brown's breakthrough was to understand that the governing set of coupled differential equations is partly "elliptic" and partly "hyperbolic." Proven methods exist to solve either of these two generic classes of equations, but they didn't work for viscoelastic flows. Once Brown and Armstrong realized the mixed character of the equations, they reformulated them to split the elliptic part of the mathematics from the hyperbolic part. With this new formulation, the Explicitly Elliptic Momentum Equation (EEME "eemie"), both parts can be accurately solved.

They have applied EEME with finite-element methods in a series of test problems, including flow between eccentric rotating cylinders and around a sphere in a tube. In previous viscoelastic calculations, reducing the size of the finite elements threw off the results. EEME, however, is consistent and accurate with increasing mesh refinement. "What we've shown," says Armstrong, "is that once you split the problem properly you can use well-developed methods and get converged solutions." With good solutions for the first time in this field, EEME is becoming the standard of comparison for non-Newtonian flow calculations. "It sets a tenor," says Brown, "for how one can design numerical algorithms for complex fluids."

Toward Realism: Car Bumpers and Supercomputing

Having solved a problem that plagued non-Newtonian fluid mechanics for over 20 years, Brown and Armstrong have set their sights on modeling realistic polymer flows. The obstacle: computing power. "Now the problem is efficiency," notes Armstrong. "Can we go from two-dimensional steady-state flows to three-dimensional flows and time-dependency? Given the time these calculations require, 3-D is going to be a difficult problem, yet the experiments show that real flows are 3-D. That's where instabilities at high flow-rates take you." Take a car bumper, for example, says Brown. Most modern bumpers are made from suspensions of glass fibers in a polymer. "You want to be able to model that and understand where the fibers end up, because the regions with no fiber will be very weak. That's where the game will be played out. Those are the exciting applications."

References:

R. K. King, M. R. Apelian, R. C. Armstrong & R. A. Brown, "Numerically stable finite-element methods for solution of steady viscoelastic flows," *Journal of Non-Newtonian Fluid Mechanics* **29**, 147 (1988).

P. Northey, R. C. Armstrong & R. A. Brown, "Finite element calculations of time-dependent two-dimensional flows of an Upper-Convected Maxwell fluid using the EEME formulation," *Journal of Non-Newtonian Fluid Mechanics* (1990), in press.

P. J. Coates, W. Lunsman, R. C. Armstrong & R. A. Brown, "Viscoelastic flows computed with the EEME/finite-element method," *Journal of Scientific Computing* (1990), in press.

Epigraph: anonymous, quoted from Morris Kline, *Mathematics and the Physical World* (N.Y.: Crowell, 1959), p. 224.

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

A New Wrinkle in Rolled Steel

Finite-Element Modeling of Hot-Steel Rolling

Jerry Kahrs, USS Technical Center

Computer Modeling & the Steel Market

Steel is gearing up for the 90s. In the 1980s, foreign steel and new materials cut deeply into the markets served by American steel, and the industry responded with a strong effort to modernize its production facilities. Pittsburgh-based USS, a division of USX Corporation, has turned to supercomputing as one component of this drive toward the technological leading-edge. Through affiliation with the Pittsburgh Supercomputing Center, researchers at the USS Technical Center, the R & D arm of USX's steel-making group, have access to computational capability that gives a substantial boost to USX's steel development program.

Jerry Kahrs, an engineer at the USS Technical Center, spends his workday investigating ways to improve the process control systems used in steel manufacturing. "There's been a virtual explosion of intelligent devices, sensors and actuators directly connected to the processes and networked to higher-level process-control computers," says Kahrs. "These devices let us use mathematical models not only for off-line simulation studies but also, even more importantly, we can imbed them in the core of our process control software." Kahrs believes that developing more accurate, versatile computer models is one of the keys to producing quality steel products at competitive cost.

Complexity of the Hot-Rolling Process

Kahrs has worked extensively on hot-rolling processes, a major step in the production of a wide variety of raw steel forms—structural plates, beams, bars, pipes and steel sheet for products from automobile bodies to food containers. Process control models for hot-rolling use real-time data from electronic sensors along with specifications for a particular product to determine how much width-reduction can be accomplished on each pass through the mill. The model computes how much force will be exerted on the mills and aims to maximize productivity while avoiding damage to the millions of dollars of capital investment that the mills represent.

To achieve better fine-tuning of this process, the mathematical models have evolved toward greater complexity, incorporating more information to more precisely mirror reality. Current state-of-the-art modeling, says Kahrs, includes not only temperature and chemical composition of the steel; it accounts for stretch and bending in the mill and complex heat-transfer processes—cooling of the steel by convection and conduction and the heating that occurs as the mechanical work of rolling converts to thermal energy.

The ultimate result of these improved models, says Kahrs, is more consistent control of product dimensions and finer control of the steel microstructure. As an example of how improved process control responds to a market need, Kahrs cites the heavy steel plates and beams used in high-rise buildings and offshore oil platforms. "Tougher building standards and safety regulations have created a need for steel alloys with better fracture toughness and weldability. With common structural steel alloys, this can be accomplished by producing steel with a relatively fine grain-size, and the most economical way to do this is through proper temperature control during rolling."

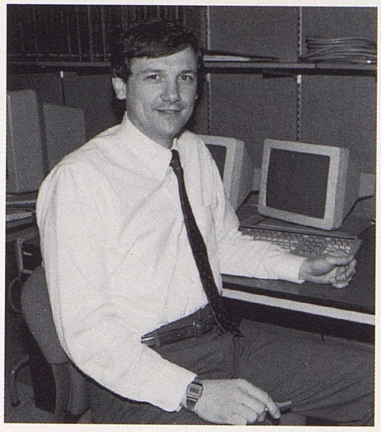

Jerry Kahrs, Associate Research Consultant at USS Technical Center.

Advantages of the Y-MP

In Kahrs' work on the hot-rolling of steel slabs, he has turned to three-dimensional finite-element modeling, using the DYNA 3-D package developed by Lawrence Livermore Laboratories and the processing power of the CRAY Y-MP. "You have to handle large displacement, large rotation and a highly non-linear material response, not to mention the complexity of a sometimes sliding, sometimes sticking interface between the roll and the steel. When you add the closely coupled thermal effects, the computational task is daunting." Kahrs' simulations take about an hour of Y-MP time per run. Without access to the CRAY, turn-around time for modeling of this complexity would be unmanageable, says Kahrs. "I wouldn't even attempt it; it just wouldn't be worth it."

Steel Microstructure in Relation to Temperature and Deformation History

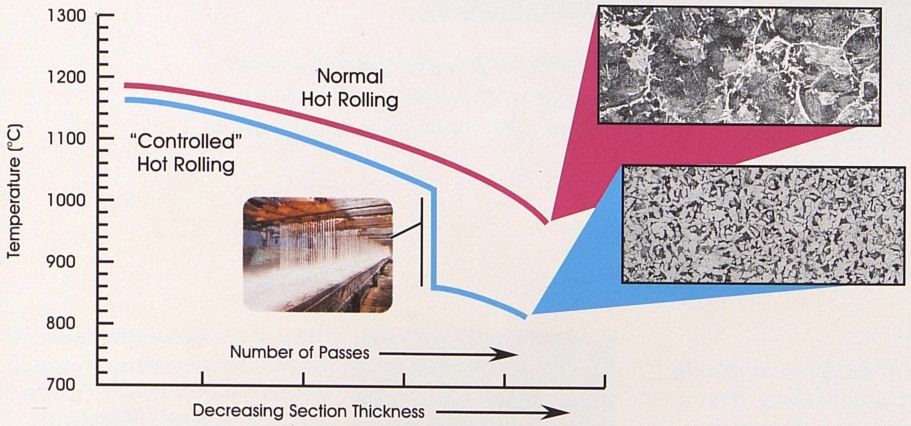

"The drop in temperature indicates a stage in the rolling sequence where the steel is rapidly cooled using water sprays. This is a predetermined strategy for a particular product, and the computer model keeps track of the process. The model can determine, for instance, the correct flow rate of water and how fast the steel should move through the cooling sprays, which will vary from run to run according to the particular product and thickness of the steel."

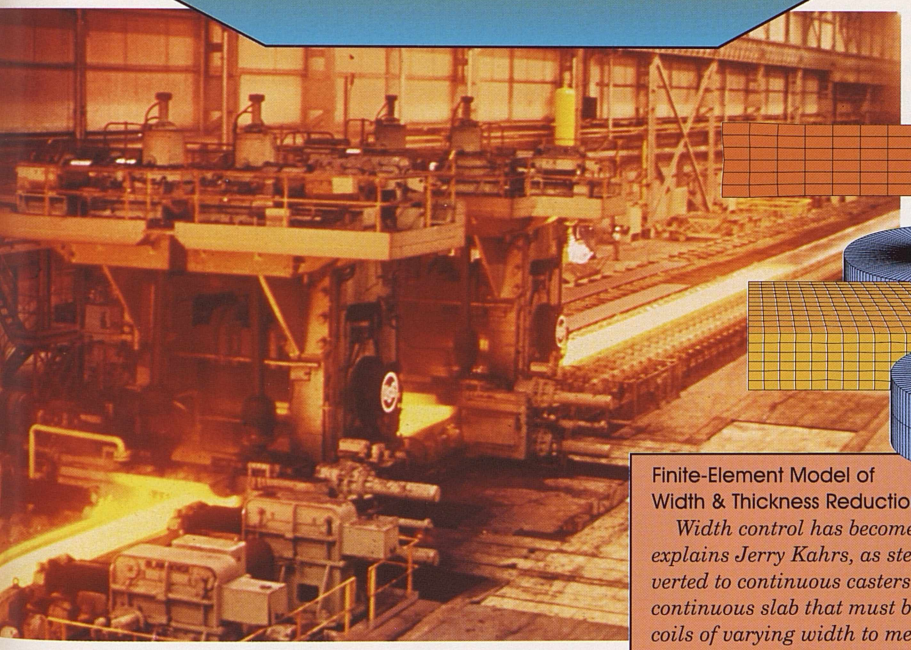

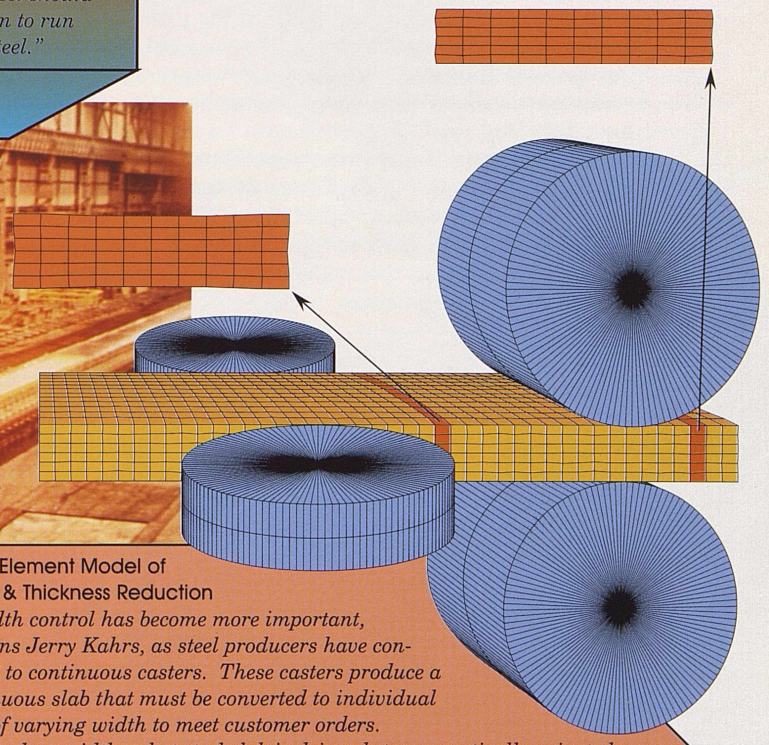

Finite-Element Model of Width & Thickness Reduction

Width control has become more important, explains Jerry Kahrs, as steel producers have converted to continuous casters. These casters produce a continuous slab that must be converted to individual coils of varying width to meet customer orders.

To reduce width, a hot steel slab is driven between vertically-oriented rolls. This squeezing tends to bulge the face of the slab, a deformity eliminated by the ensuing thickness-reduction stand. The thickness reduction, however, tends to restore some of the slab's width—a process that varies according to the properties of a particular slab and rolling stand.

Modeling on the supercomputer, says Kahrs, has made it possible to investigate the limits of width-reduction without risking the equipment. "It's extremely expensive to shut down production to run tests on a new product. We really hate to do that. With the supercomputer, we can get a very good idea without affecting production at all."

Widening the Width-Acceptance Window

Kahrs' modeling of hot-steel rolling has opened the "width-acceptance window" at the Hot Strip Mill of the USS Gary, Indiana Works. If the dimensions of a slab don't fit within the width-acceptance window for a particular product, it must be diverted to a different order, placed in inventory or down-graded and sold as lower-quality material. "The wider the acceptance window, the less inventory they have to keep around, and that becomes very important financially."

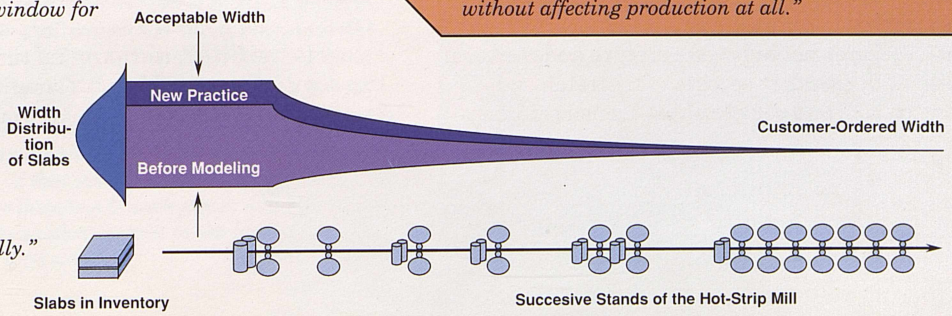

The Underworld of Ice

Three-Dimensional Microanalysis of Temperature-Gradient Metamorphism

Pat Burns & Mark Christon, Colorado State University

"One must have a mind of winter
To regard the frost and the boughs
Of the pine-trees crusted with snow."

Avalanches and Acid Pulse

Snow and mountains—for many of us, these words signal skis, a lodge, a log fire to drive the chill away. For Pat Burns and Mark Christon, snow in the mountains also means depth hoar, temperature-gradient metamorphism, avalanches and acid pulse, terms with a menacing edge. Avalanches are deadly, and science has not yet been able to muster a reliable predictive capability. Christon and Burns have turned the computing power of the CRAY X-MP and Y-MP at the Pittsburgh Supercomputing Center toward a fifty-year-old question: Why does depth hoar, the ice-crystal structure that develops under deep snow, cause avalanches in some cases and not others?

Nature makes avalanches by herself but needs help from civilization to bring about an acid pulse. Sulfates, nitrates, sulfur dioxide and other urban-industrial air pollutants deposit in snow, and the same water vapor transport mechanisms that cause depth hoar tend to concentrate these impurities. Spring—Alpine meadows start to bloom, snow begins to melt. "The snowpack will release 40 to 50% of its impurities in the first meltwater," explains Christon. "That rapid flush of impurities is usually realized as pH, and in some Alpine regions, where the hydrological systems are fairly dilute, a pH of about 5 can kill anything in the water—trout, frogs, turtles. For the plantlife, it can boil down to *how long will it survive?* These are fragile ecosystems."

A Crystal Clear Look at TGM

The winter air in Alpine regions is usually colder than the ground, which retains heat and holds a relatively constant temperature of about 0° C. As a result, ice at lower levels of the snowpack sublimates to water vapor, which diffuses upward through the air space between snow crystals. This phenomenon—temperature-gradient metamorphism (TGM)—recrystallizes the ice, changing its structure in ways that can weaken underlying layers and create avalanche conditions. Movement of water molecules across the vapor space in snow also offers a key to understanding how pollutants deposited on the snow surface distribute and concentrate at deeper layers.

Most research on TGM has developed complex models of water vapor transport that treat snow as a homogeneous substance. These efforts have produced large variations in the diffusion coefficient—"anywhere from 0.7 to 24," notes Burns—a number that relates the rate at which water diffuses in snow to the dry air rate; accurate determination of this parameter is fundamental to developing an accurate description of TGM. Burns and Christon have taken a different approach. "As opposed to modeling nonequilibrium thermodynamics," says Burns, "and applying nonlinear equations at the macroscopic level, we use a relatively simple physical model and look at the microscale. We emphasize the geometrical complexity of snow crystals, something we could do because we're using the CRAY."

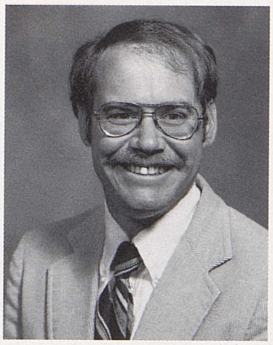

Pat Burns

"We couldn't have begun to do this in three dimensions without a supercomputer." The finite-element deforming meshes track 9000 nodes with eight variables at each node. Relying on vectorization and gather/scatter features of the CRAY, the model runs at a sustained X-MP single-processor execution rate better than 100 million arithmetic operations a second.

Modeling Branch-Grain Theory

Burns and Christon base their model on the "branch grain theory" of TGM posed by Forest Service geologist Richard Sommerfeld. Relying on extensive experimental observations, Sommerfeld found that the geometry of individual crystals, in particular their "branch grains," is the major determinant in TGM. Christon first developed a two-dimensional computer model of snow crystals that tended to confirm this approach. With access to the CRAY at Pittsburgh, he worked with Burns to develop a three-dimensional finite element model. Using a coupled pair of deforming meshes, one for the entire volume and one for the vapor space, the model tracks both heat transfer and mass transfer as it changes over time with various crystal geometries.

"In snow there's a distribution of geometries," says Christon, "and it can be different for every location in the snow cover. We're trying to determine a representative range of geometries and crystal sizes, which we've arrived at from measurements of disaggregated snow. Given a particular configuration, we can then take our model and determine the effective rate that water vapor moves across the lattices." Christon checked his model against analytical solutions for several idealized geome-

tries and also confirmed that it produced results for thermal conductivity that matched well with field observations.

He applied the model to over 150 separate three-dimensional geometries. Compared to most prior research, the results substantially narrow the limits of the diffusion coefficient. "We've nailed it between 1 and about 2.5," says Christon, who notes that this agrees with Sommerfeld's experimental work. Burns and Christon believe these results provide a way to meaningfully reinterpret earlier studies. "Looking at geometrical complexity allows us to synthesize a whole new theory," says Burns. "People in the past argued about whether the experiments were correct. What we're seeing is that perhaps they were all correct; they were just dealing with different types of snow."

References:
Richard A. Sommerfeld, "A Branch Grain Theory of Temperature Gradient Metamorphism in Snow," *Journal of Geophysical Research* **88,** 1484 (1983).
Mark Christon & Pat Burns, "3-D Transient Heat and Mass Transport in Microscopic Ice Lattices," to be presented at AIAA/ASME Thermophysics and Heat Transfer Conference, Seattle, June 18-20, 1990.
Salamander photo courtesy of Bill Allen, Pittsburgh.

Epigraph: from "The Snowman," Wallace Stevens.

This research is supported by the Rocky Mountain Forest and Range Experiment Station and the National Science Foundation.

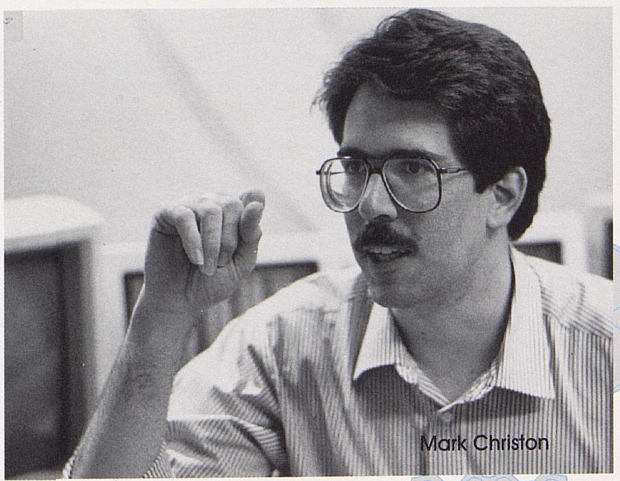
Mark Christon

Using the Center's graphics facilities, Mark Christon has produced several 3-D animations showing thermal gradients in ice crystals.

"With the CRAY, we can basically go in and stand inside one of these lattices," says Christon, "and watch the process as it happens. That hasn't been done before. You can see that basically the horizontal planes are isothermal and the connective members carry the temperature gradient. What we've found, essentially, is that the rate of diffusion is going to be enhanced anywhere there is a branch grain, and the closer the branches are to another surface in the lattice the greater the potential for diffusion."

Researchers from the University of California, Berkeley have recently found a 65% decline in the population of tiger salamanders at a subalpine watershed in the Colorado Rockies. Tiger salamander eggs are sensitive to pH in the range produced by an acid pulse, and the study hypothesizes that this may be the first quantifiable indication of "biological damage from acid deposition in the western United States." Conservation Biology ***3,*** *149 (1989).*

Geometrical Evolution of an Ice Crystal

This three-dimensional idealized representation of aggregated ice-lattice cells approximates the geometrical complexity of ice in a snowpack. This lattice has four vertical "branch grains" and a density of 266 kilograms per cubic meter. Applying a temperature gradient for 15 hours results in the transfer of ice mass from the lower, horizontal surface to the branch grains and the upper surface. The cup-shaped branches, says Mark Christon, correspond to an observed characteristic of temperature-gradient metamorphism. Color corresponds to temperature, with red warm (0°C) and blue cold.

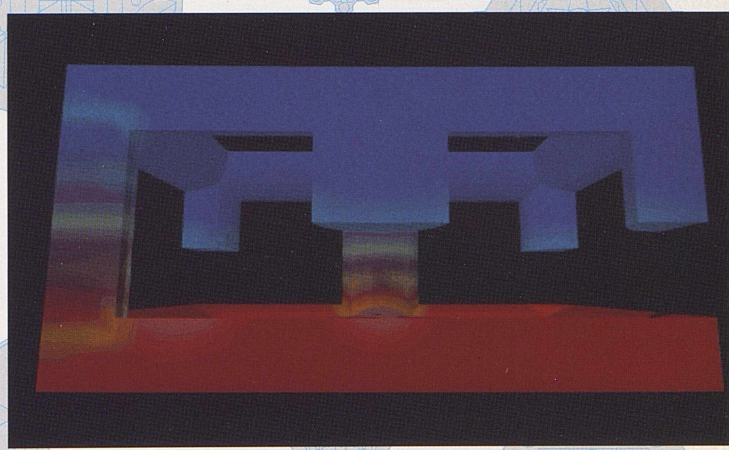

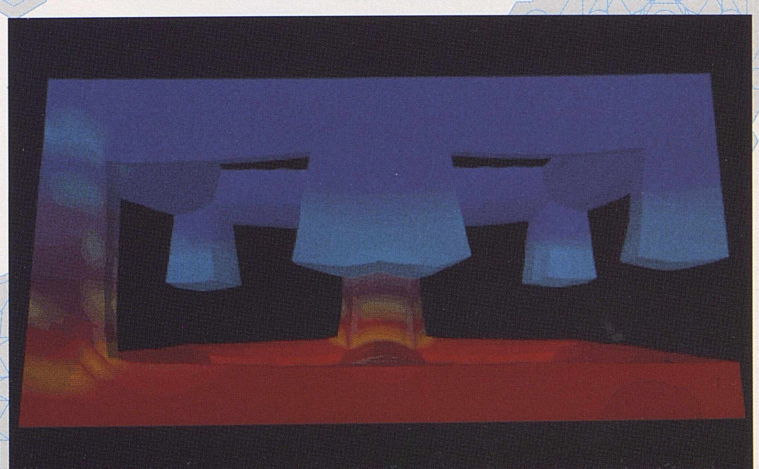

Risky Breathing

Multistage Modeling of Cohort Data for Carcinogenic Risk Assessment

Sati Mazumdar & Rita N. Patwardhan
University of Pittsburgh

"I durst not laugh for fear of opening my lips and receiving the bad air."

Assessing the Risk of Airborne Cancer

"Growing evidence demonstrates that pervasive contamination of air, water, soil, and food with a wide range of industrial carcinogens, generally without public knowledge and consent, is important in causation of modern preventable cancer." So wrote Samuel Epstein of the University of Illinois Medical Center and 16 cosigners to a May 10, 1988 open letter in *Science*. Though it's well accepted that some industrial air pollutants can cause cancer, regulators at the Environmental Protection Agency face the imponderably difficult task of balancing these risks against the cost of reducing them. By executive order, to regulate these substances requires a showing that "the potential benefits to society . . . outweigh the potential costs."

To deal with this problem, EPA in the 80s turned to quantitative risk assessment. Mathematical techniques model or "fit" data from human and animal studies on whether exposure to a particular substance causes cancer. The models measure the probability that cancer incidence may be due to chance or to factors other than exposure to the substance, and they quantify the "excess risk"—cigarette smokers, for instance, are ten times more likely to get lung cancer than non-smokers. As a general rule, EPA considers that the risk of getting cancer from lifetime (70 years) exposure to an industrial carcinogen in the ambient air must be better than one-in-a-million to require an industry to reduce emissions.

This acceptable risk "bright line" implicitly acknowledges the inexactness inherent in risk assessment. Most cancers have a latency period, often as long as 20 to 30 years, and cancer deaths must be correlated to earlier exposures. Risk assessors must account for variations in concentration and duration of the exposure, and they must consider what incidental factors—diet, smoking, age—may have influenced results. Information is often incomplete, and many of the variables can only be estimated or, more realistically, treated as a range of probability. To accommodate these uncertainties, risk assessment is evolving toward complex "stochastic" approaches that require extensive computing.

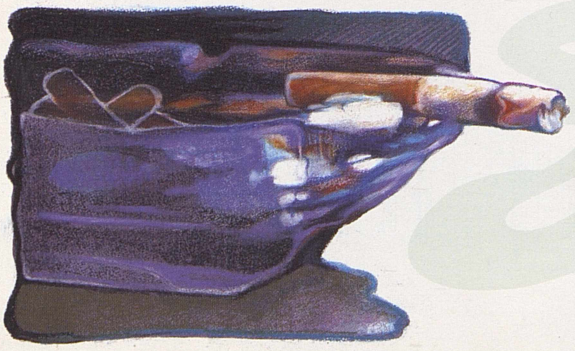

The Multistage Model

Rita N. Patwardhan

One such approach, multistage modeling, is based on studies of the way cancer develops. "Multistage theory assumes that a single cell goes through distinct stages of alteration in order to generate a malignant tumor," explains Professor Sati Mazumdar of the University of Pittsburgh's Graduate School of Public Health (GSPH). Multistage modeling distinguishes between a carcinogen's effect as an "initiator" of cancer in healthy cells and as a "promoter" in cells already affected. It provides a way to analyze whether different carcinogens affect different stages and if the effect is dose-related in either or both stages.

In a project at the Pittsburgh Supercomputing Center, Mazumdar and Rita Patwardhan implemented multistage modeling of large datasets on the CRAY. They developed algorithms and software packages to apply multistage modeling to epidemiological data. Using these tools and CRAY number-crunching power, they modeled data from the Steelworkers Mortality Study, an important epidemiological study on the cancer risk of working in a steel mill, conducted by GSPH biostatistician Carol Redmond and others.

Coke-Oven Emissions

By tracking the case histories of 59,000 Pittsburgh-area steelworkers on the payroll in 1953, the steelworkers study indicated that one subset or "cohort" of this data, nonwhite males who worked at the top of coke-ovens, were seven times more likely to die of lung cancer than other workers. In baking coal to produce coke, a key ingredient in steelmaking, coke ovens pro-

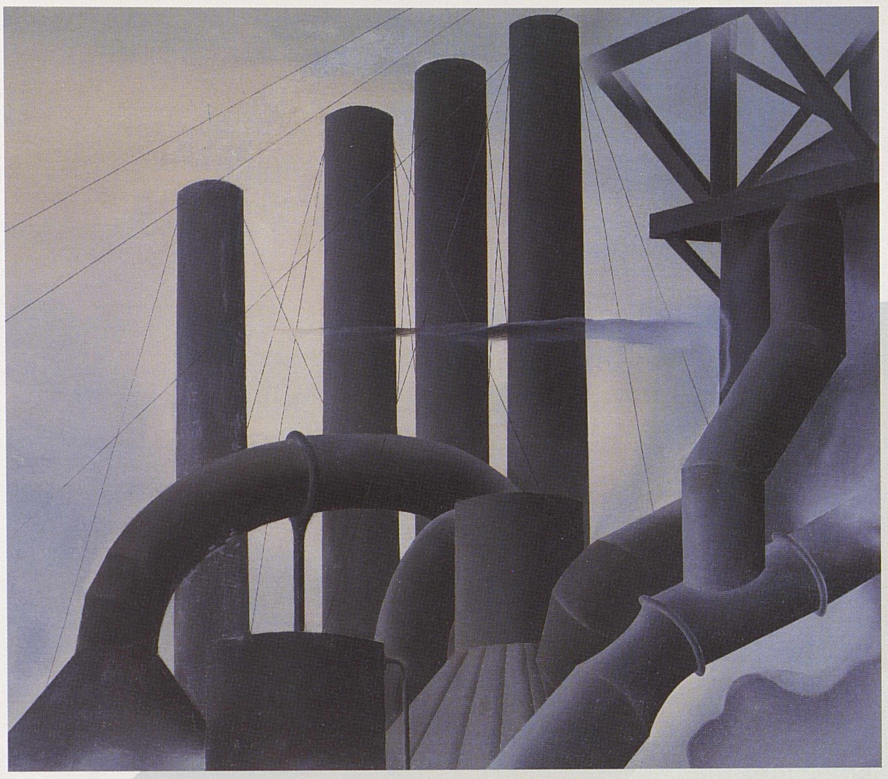

duce a hydrocarbon-laden exhaust that includes several known carcinogens; one of them—benzo(a)pyrene—is one of the most potent carcinogens in cigarette smoke. In subsequent research, Mazumdar, Redmond and others traced the case histories of 1100 coke-oven workers through 1970 and found that lung- cancer risk increased with exposure. This work established for the first time a dose-response relation between coke-oven emissions and lung cancer.

Using the CRAY, Rita Patwardhan and Susan Zhou incorporated time-dependent factors into the dose-response analysis of this data. The effect of a carcinogen can vary with age at initial exposure and duration of the exposure. Also, though many occupational studies assume constant exposure levels, exposure often varies significantly with different job assignments. To more realistically account for these factors, Mazumdar and Patwardhan divided the work-periods for each worker into one-month intervals and assigned an exposure level to each interval. These carefully constructed "exposure profiles" became the raw material for their computations. The multistage model indicates that coke-oven emissions both initiate and promote lung cancer and that the risk of lifetime occupational exposure is even greater than previously calculated—eight times the normal population.

The complex modeling undertaken in this project can be done effectively only with very high-performance computing. A "maximum likelihood" analysis, used to validate an essential constant, the potency parameter, requires an iterative numerical approach that, basically, replicates the calculation 1000 times for the 1100 workers in the coke-oven cohort. "This project required extensive calculations on a large amount of data," says Mazumdar, "exposure patterns and vital statistics for each case history, and calculations of age-specific death-rate, lifetime risk and potency parameters." None of the standard statistical packages could efficiently handle the complex evaluations that had to be done. "It seems easy after the fact," she says, "but multistage modeling of cohort data with time-dependent exposure was feasible only because we had access to the CRAY."

References:
Sati Mazumdar, Carol K. Redmond, Joseph P. Costantino, Rita N. Patwardhan & Susan Y. J. Zhou, "Recent Developments in the Multistage Modeling of Cohort Data for Carcinogenic Risk Assessment," *Environmental Health Perspectives* (1990), in press.
Rita N. Patwardhan, *Inferential Procedures for Multistage Models for Carcinogenic Risk Assessment with Applications*, Doctoral Dissertation, Department of Biostatistics, Graduate School of Public Health, University of Pittsburgh (1989).

Epigraph: *Julius Caesar*, I, ii, 249-50.

Steelmill Image: Elsie Driggs, *Pittsburgh*, 1927, oil on canvas, Whitney Museum of American Art, New York City.

This research was supported by the Environmental Protection Agency.

Likelihood Surface

In multistage theory, cancer risk depends on exposure, represented as an index defined by the model, times a constant—the potency parameter (β). The potency parameters, independent for each stage, are estimated using "maximum likelihood" analysis to obtain the best possible fit with observed data. This "likelihood surface" for the multistage model fitted to coke-oven workers cancer mortality shows a distinct peak, indicating maximum likelihood estimates for the two potency parameters.

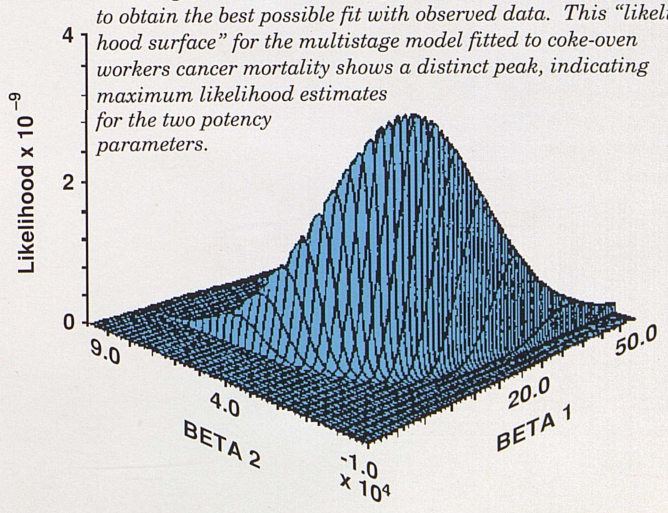

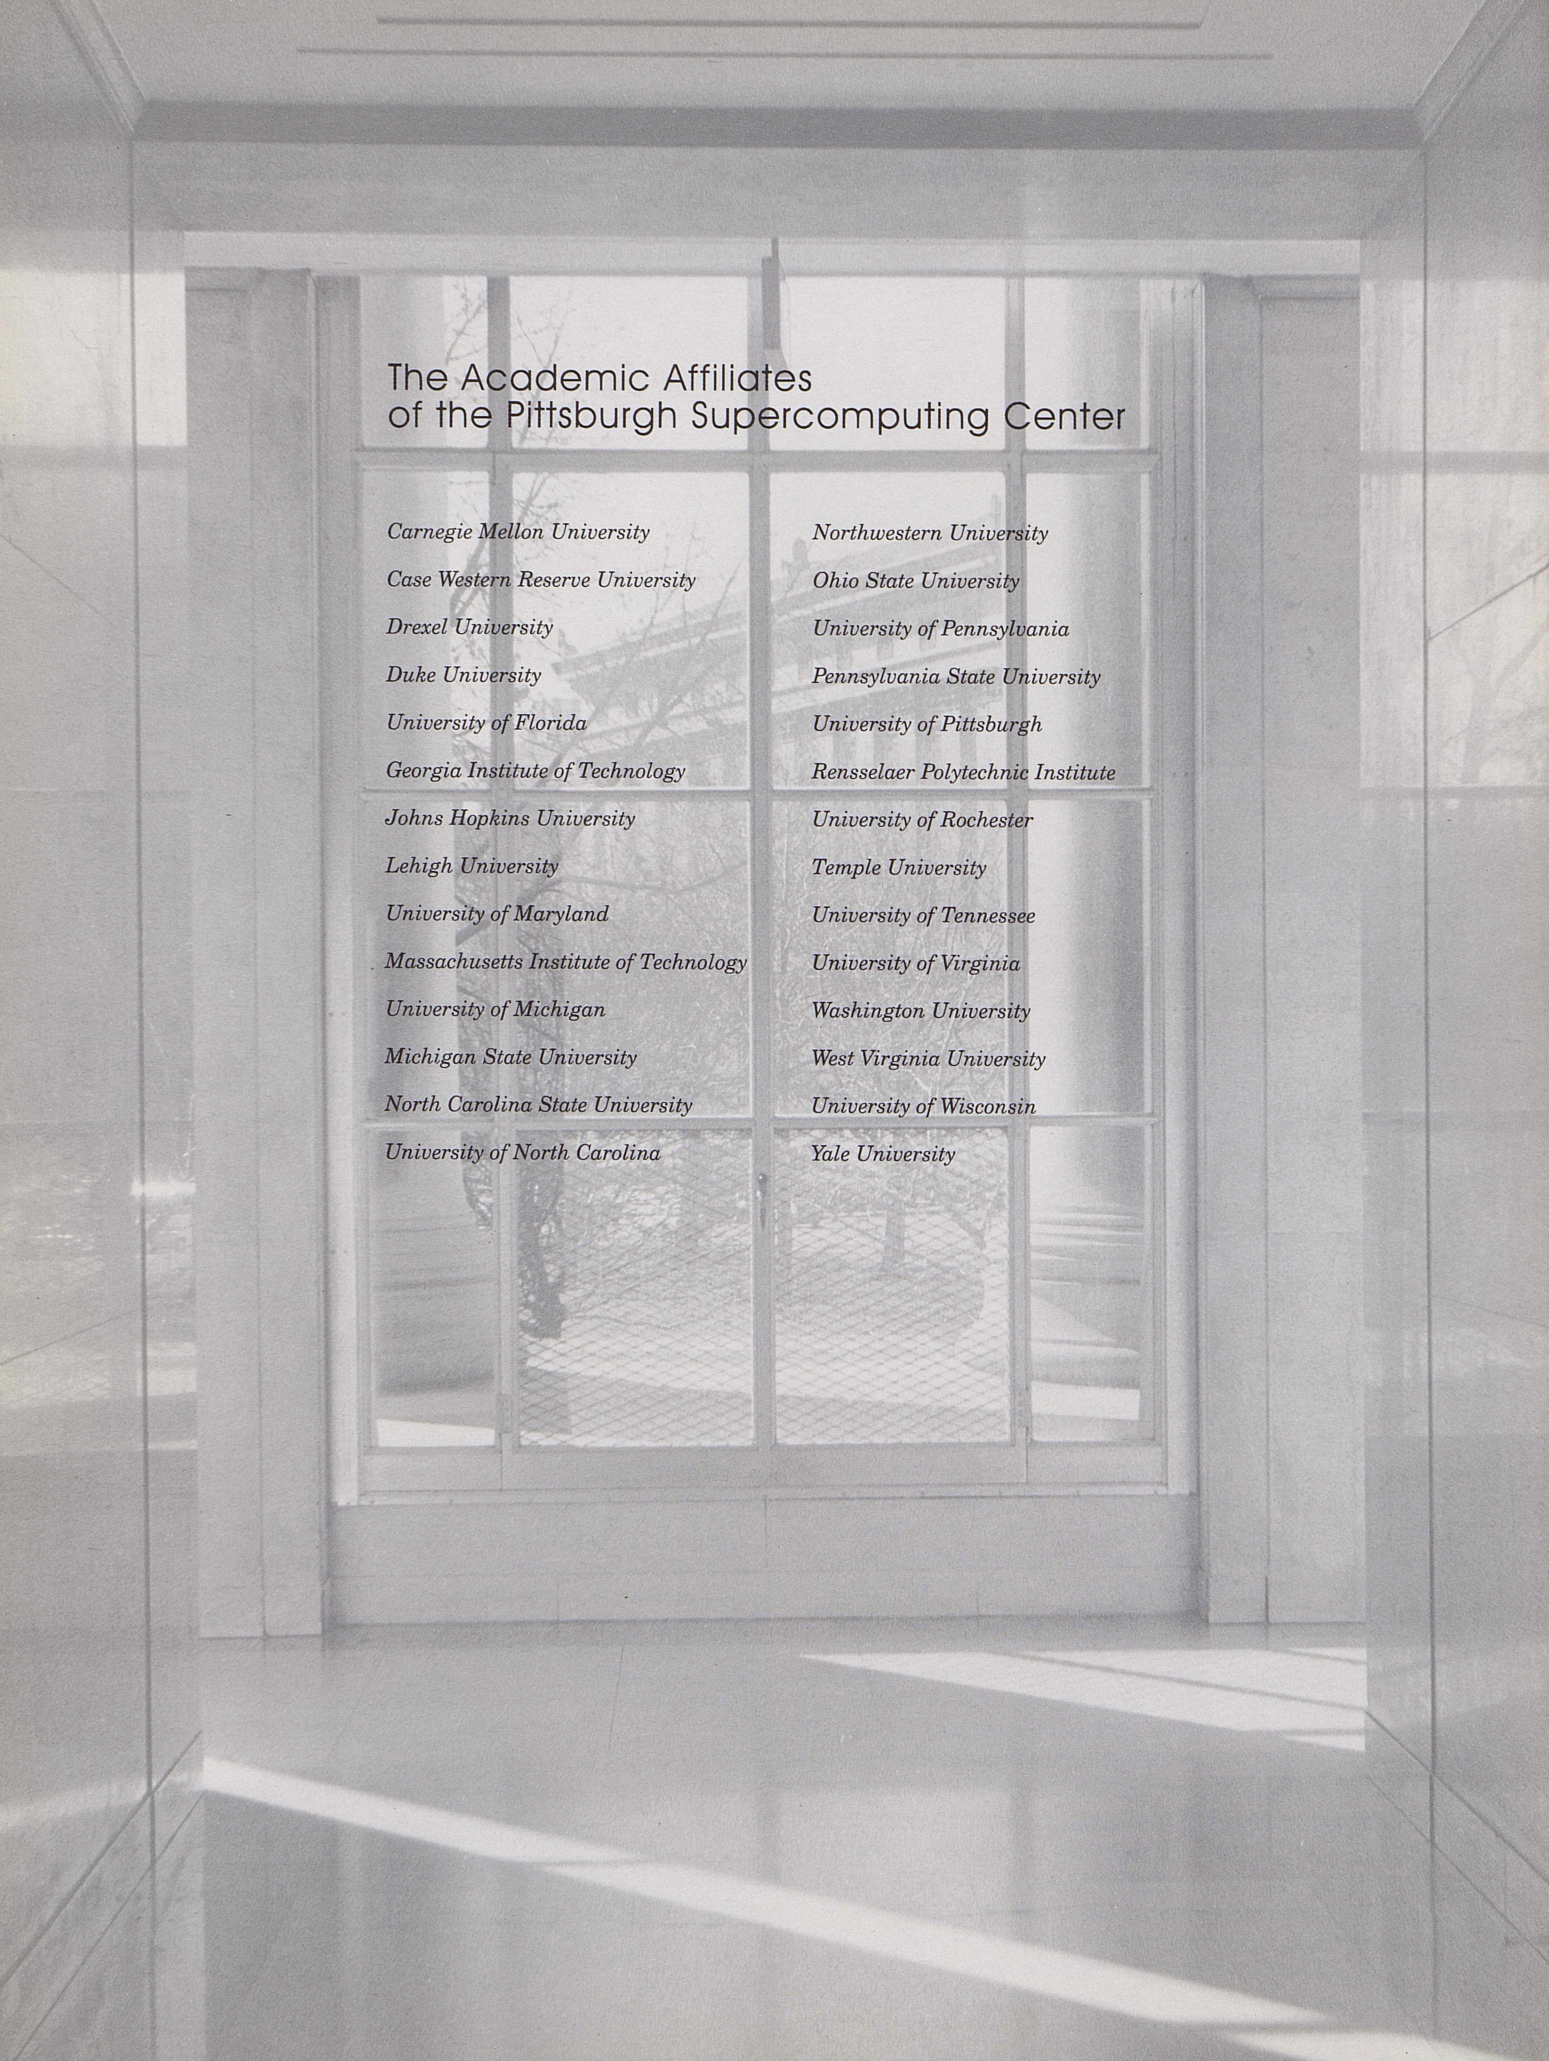

The Academic Affiliates of the Pittsburgh Supercomputing Center

Carnegie Mellon University

Case Western Reserve University

Drexel University

Duke University

University of Florida

Georgia Institute of Technology

Johns Hopkins University

Lehigh University

University of Maryland

Massachusetts Institute of Technology

University of Michigan

Michigan State University

North Carolina State University

University of North Carolina

Northwestern University

Ohio State University

University of Pennsylvania

Pennsylvania State University

University of Pittsburgh

Rensselaer Polytechnic Institute

University of Rochester

Temple University

University of Tennessee

University of Virginia

Washington University

West Virginia University

University of Wisconsin

Yale University

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

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