1996

PROJECTS IN

Pittsburgh Supercomputing Center

SCIENTIFIC COMPUTING

TEN YEARS OF LEADERSHIP

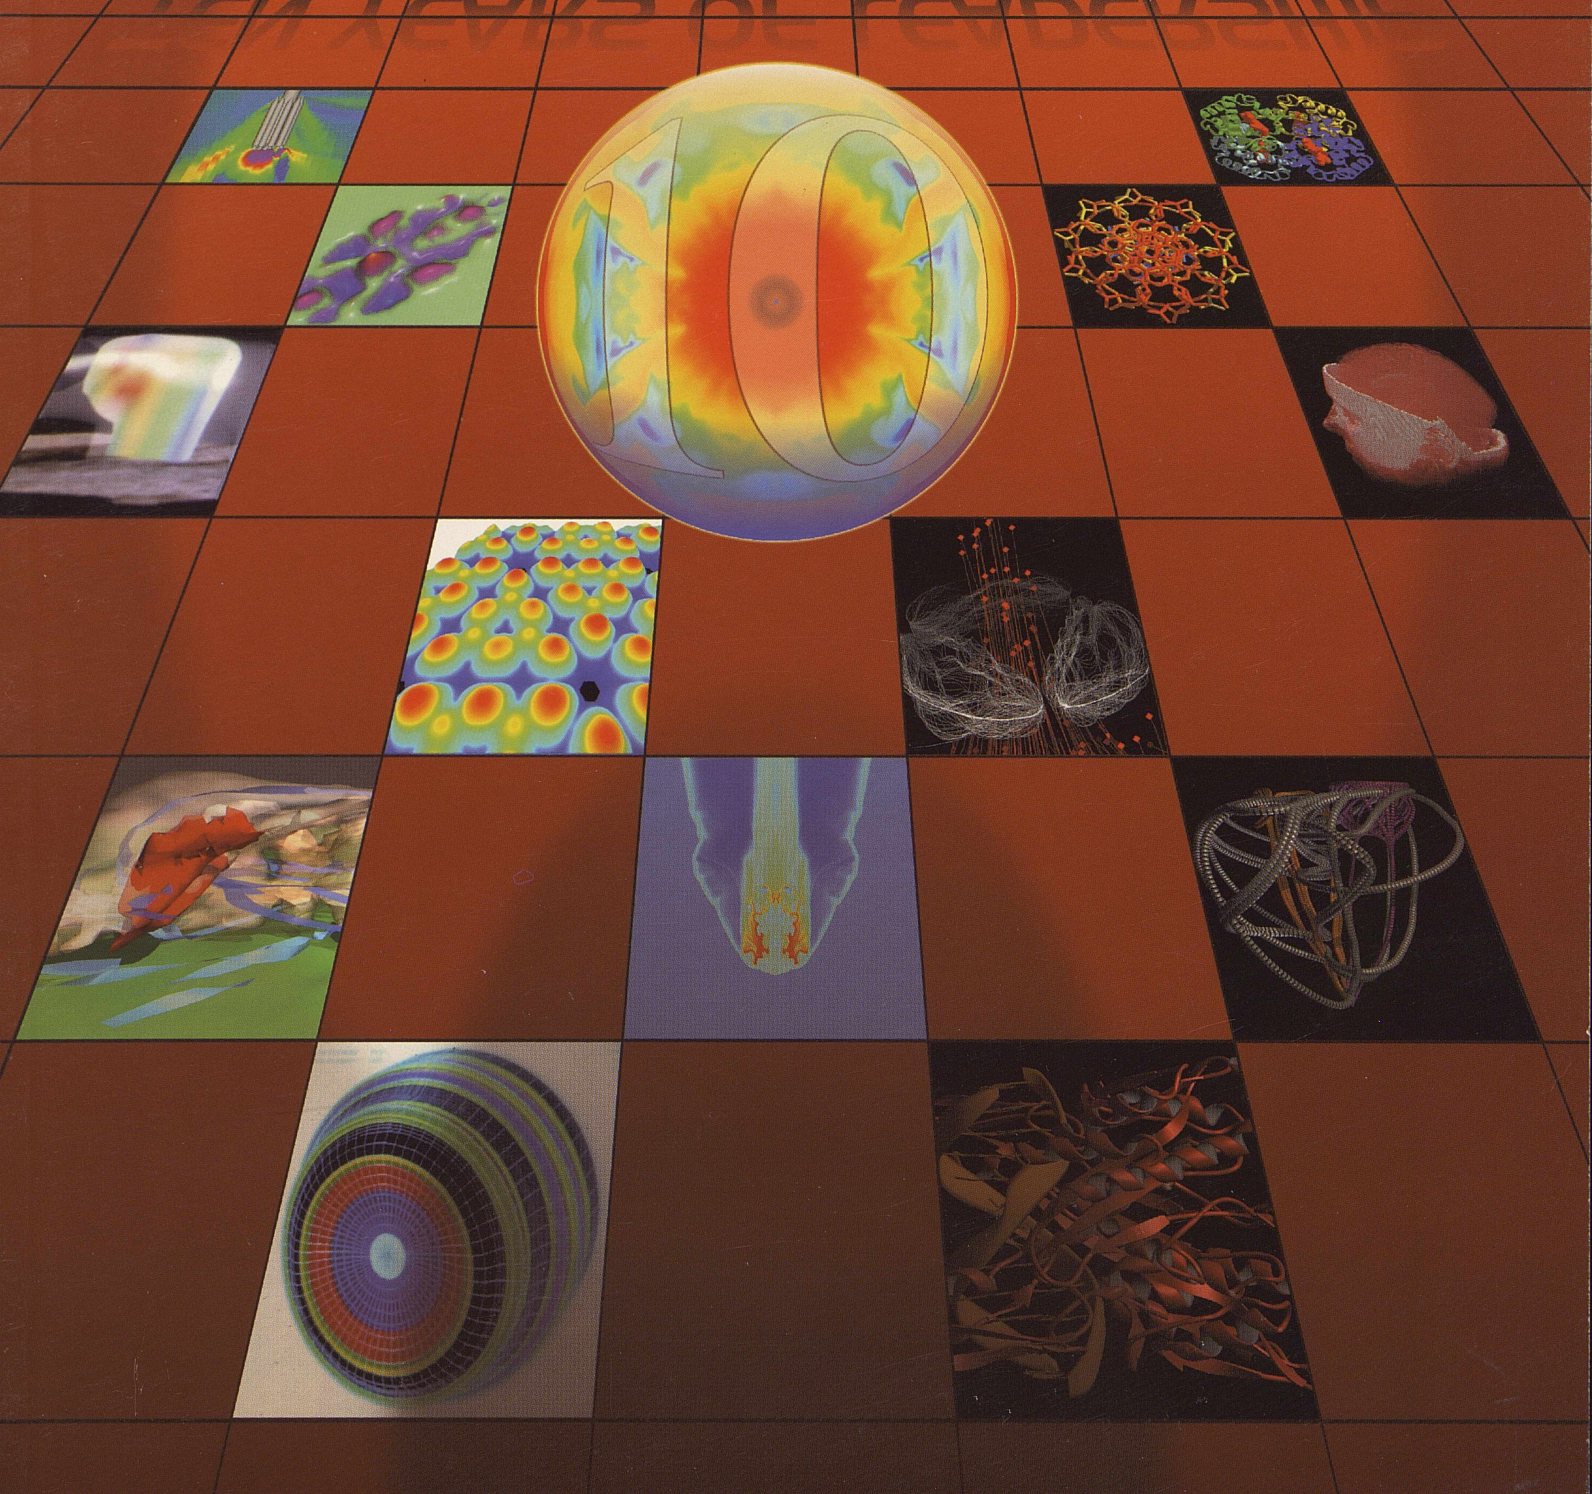

PITTSBURGH
SUPERCOMPUTING
CENTER

Projects in Scientific Computing, 1996

Images on the cover are from computational research at Pittsburgh Supercomputing Center.

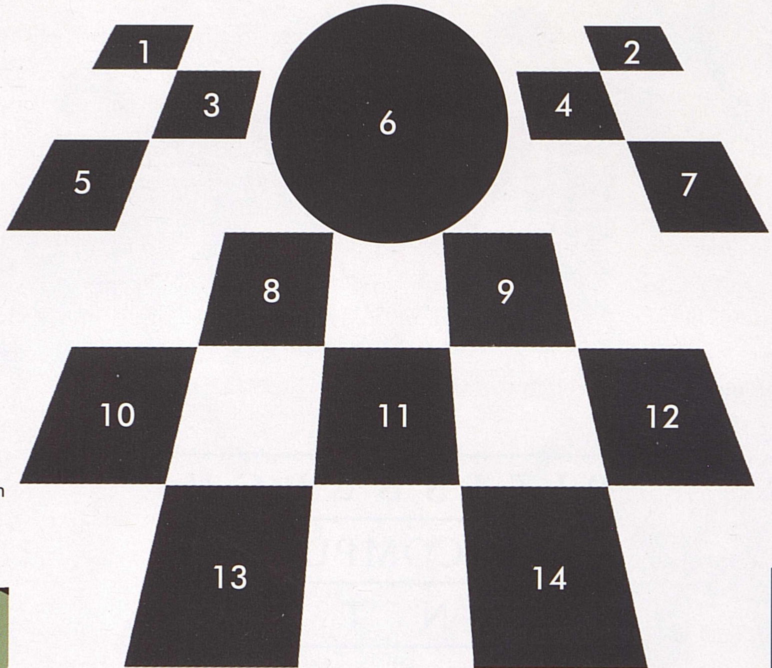

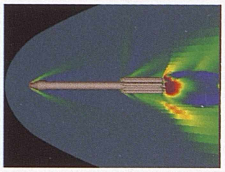

1. Flow field of the Delta II rocket, as simulated by Stephen Taylor, California Institute of Technology (see p. 34).

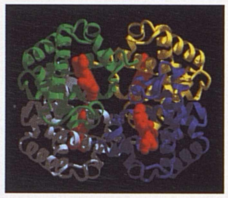

2. Molecular structure of hemoglobin, Marcela Madrid, PSC, and Chien Ho, Carnegie Mellon University.

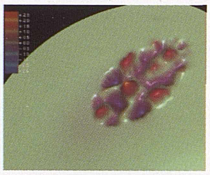

3. Ground motion in a hypothetical earthquake, as simulated by Jacobo Bielak and colleagues, Carnegie Mellon University.

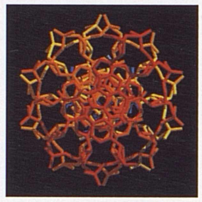

4. Structure of DNA, by David Deerfield, PSC (see p. 7).

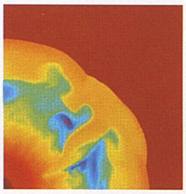

6. Explosion of a supernova, as modeled by Adam Burrows, University of Arizona (see p. 32).

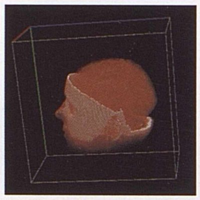

7. Functional MRI (magnetic research imaging) of the brain, from research by Jonathan Cohen, Carnegie Mellon University and the University of Pittsburgh.

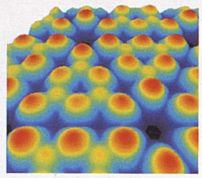

8. Silicon surface reconstruction as computed by John Joannopoulos, MIT.

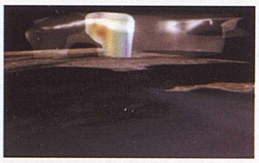

5. Smog over the Los Angeles air basin, as modeled by Gregory McRae and Armistead Russell (see p. 24), from a video produced by PSC visual information systems group.

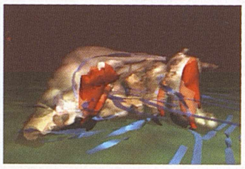

10. A thunderstorm over Oklahoma; image produced by Greg Foss, PSC, from modeling by Kelvin Droegemeier, University of Oklahoma (see p. 22).

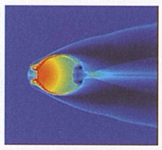

11. Comet Shoemaker-Levy 9 exploding after it crashes into the atmosphere of Jupiter, as modeled by Mordecai-Mark Mac Low, University of Chicago.

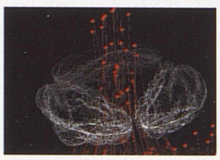

9. Blood flow from the aortic valve, as modeled by Charles Peskin and David McQueen, New York University.

12. Muscle fibers of the heart, from an educational video produced by the PSC visual information systems group, based on computations by Charles Peskin and David McQueen, New York University.

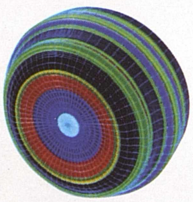

13. The bottom of a beverage can during dynamic snap-through, as simulated by ALCOA Laboratories.

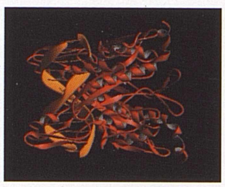

14. DNA interacting with the restriction enzyme Eco-RI, as simulated by John Rosenberg, University of Pittsburgh.

Published August 1996 by the Pittsburgh Supercomputing Center, 4400 Fifth Avenue, Pittsburgh, PA 15213. (412) 268-4960.

Foreword

We have come a long way in a decade. The experiment of harnessing the creative energies in our universities to centers providing support and innovations in computational science has been remarkably successful. The Pittsburgh Supercomputing Center (PSC) is nationally recognized as the place to work on important scientific problems that are computationally challenging. More than half the computing on large problems at the four National Science Foundation supercomputer centers in 1995 was done at PSC.

Coupling the world's best supercomputers and ancillary systems with a willingness to target them as a whole, when appropriate, to single large projects has led to new scientific insights that could not have been predicted ten years ago. We illustrate some of these in this booklet, which describes improved methods for studying the stability of the DNA molecules that contain our genetic material, new milestones in the ability to predict severe storms and consequently save billions of dollars annually, and profound new insights into the origin of the Earth's magnetic field.

Indeed, as we go to press with this booklet, important new findings about the rotational velocity of Earth's inner core are making headlines in newspapers and science journals. Two Earth science projects at PSC, both reported in this booklet (see pp. 26-30), paved the way for this research, which would have been unimaginable without the guidance provided by these supercomputer simulations of Earth's deep interior.

Our work with the CRAY T3D has led us to conclude that massive parallelism offers the most cost-effective approach to most of our very large problems. Just as with the T3D, we demonstrated our world leadership in massively parallel computing by being the first to introduce its successor, the T3E. We worked closely with Cray Research to test the new machine and eradicate its bugs, hardware and software, as well as developing important application software to exploit it.

The past 10 years have taught us several lessons. One is that increasing resolution is often the difference between qualitatively correct and qualitatively inadequate results. Large focused capabilities are essential for this. A second is that interdisciplinary collaborations, especially between computer scientists and computational scientists, are very fruitful. A third is that both the scientific understanding and technological capability required for such projects take time to mature, and the assurance that adequate computational resources will be available is indispensable.

The National Science Foundation is currently reevaluating its supercomputer centers. We look forward to their rapid reaffirmation of PSC's excellence, so that we can focus all our energies on our primary task, enabling breakthrough science using high-performance computing. We gratefully acknowledge NSF's funding, which is critical to our continued vitality, as well as that of major high-performance computing and communications vendors and other government agencies. NIH support has enabled many of the biomedical advances you will read about in this booklet. The Commonwealth of Pennsylvania's support in the face of many other demands also has been critical to our success.

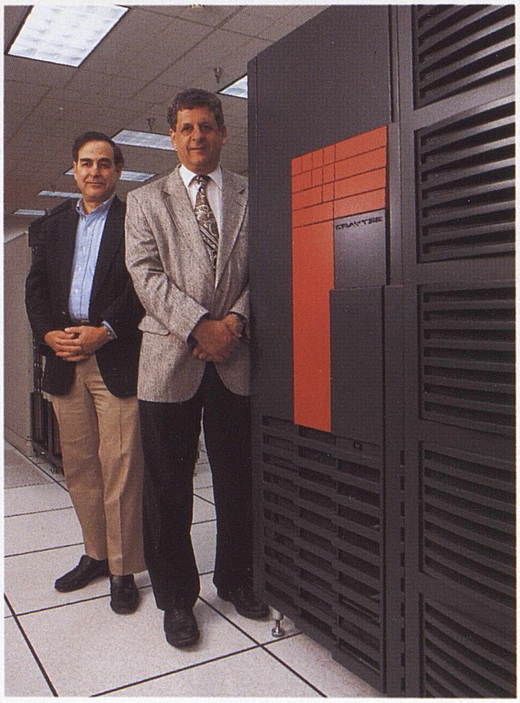

Michael J. Levine, scientific director

Ralph Z. Roskies, scientific director

Projects in Scientific Computing, 1996

Contents

Pittsburgh Supercomputing Center, 1996

10 Years of Making Science Happen

Ten years have passed since June 1986 when the Pittsburgh Supercomputing Center opened its doors. A late-comer among National Science Foundation supercomputing centers, PSC broke fast from the gate and never looked back. A short, invigorating 10 years later, it has become one of the leading institutions worldwide in supercomputing research and technology.

The past year has been a rewarding culmination of this first decade, as a number of PSC-related projects achieved wide recognition. In September, two of the most prestigious scientific journals in the world, *Science* and *Nature*, featured PSC research on their covers. The *Science* cover story (Sept. 8, 1995) reported simulations of solar convection by University of Colorado physicist Juri Toomre and his colleagues. This issue focused on computational fluid dynamics, and three of the four lead articles described results obtained using PSC resources. Earlier in 1995, *Science* also published articles on two of PSC's important "grand challenge" projects: protein folding (see p. 16) and air-quality modeling (see p. 24).

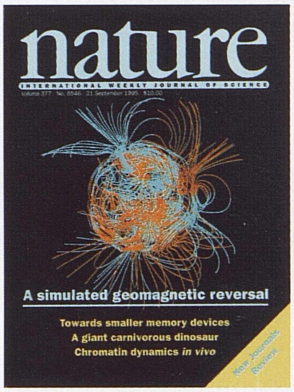

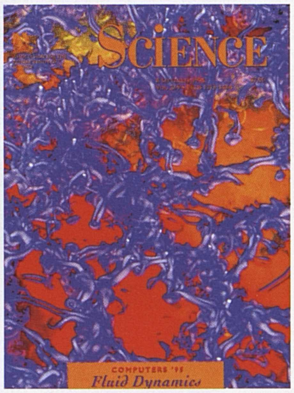

Two weeks later, the cover of *Nature* (Sept. 21, 1995), heralded the first simulated reversal of Earth's magnetic field. This research (see p. 26) used the CRAY C90 to model the electromagnetic fluid-dynamical processes in Earth's core. The results revealed unanticipated geomagnetic phenomena and offer new insight into how Earth's magnetic field is generated.

This year, for the second year in a row, the Sidney Fernbach award, given annually to an individual scientist for innovation in high-performance computing, recognized PSC research. The 1994 Fernbach award honored NYU scientist Charles Peskin, who used PSC's CRAY C90 to model blood-flow in the heart. In December 1995, University of Minnesota physicist Paul Woodward won the 1995 Fernbach award. His computations on PSC's CRAY T3D, believed to be the largest simulation of turbulent convection ever done, revealed features of solar turbulence never before observed or successfully simulated.

The 1996 Computerworld-Smithsonian awards brought further recognition of PSC's leadership in enabling important computational research. The CWSA program is the U.S. computer industry's premier vehicle to recognize information technology that improves the quality of life, and four of the five 1996 finalists in science were PSC-related projects (see pp. 24, 26 & 34). PSC has collaborated with the winner, the Center for Light Microscope Imaging and Biotechnology at Carnegie Mellon, since 1992 on grand challenge research to accelerate display and analysis of three-dimensional image data from cell microscopy.

Four of the five 1996 finalists in science were PSC-related research.

Heart disease. Weather forecasting. Global climate change. Computer visualization is a cornerstone of modern problem solving, and PSC's multimedia, interactive World Wide Web site shows "how scientific visualization is being used to solve the most challenging, socially signifigant problems." (*Iris Universe*, Winter 1996.)

http://www.psc.edu/science/contents.html

Beverly Clayton has been executive director since PSC opened in 1986.

PSC supplied more than half the computing power used in 1995 for large research projects at NSF supercomputing centerS.

Gary Jensen, associate director, who joined PSC in February 1996 after 15 years at the National Center for Atmospheric Research in Boulder, Colorado.

Support for Landmark Research

It is not accidental that research at PSC is achieving significant results. Over the past several years, the center has built one of the most capable computational environments in the world — with two powerful, but radically different systems, the CRAY C90 and T3D. To maximize productivity from these resources, which are in high demand, PSC has emphasized *capability computing*.

Capability computing enables large-scale research that would not be feasible on smaller systems. Scheduling flexibility and access to technical expertise encourage "landmark" projects that can exploit maximum processing over sustained periods. This means using all 16 processors on the C90, for instance, or all 512 processors on the T3D. As a result of this emphasis, more than half of the computing power used in 1995 by large research projects in the NSF supercomputing centers program came from PSC.

While the C90 continues to be the main workhorse for most PSC research, the T3D — entering its third year of operation — has paved the way in making scalable, parallel computing available to scientists in a range of fields. PSC's T3D is now actively used by about 50 research collaborations, and a number of other PSC researchers are adapting their software to exploit the CRAY T3E.

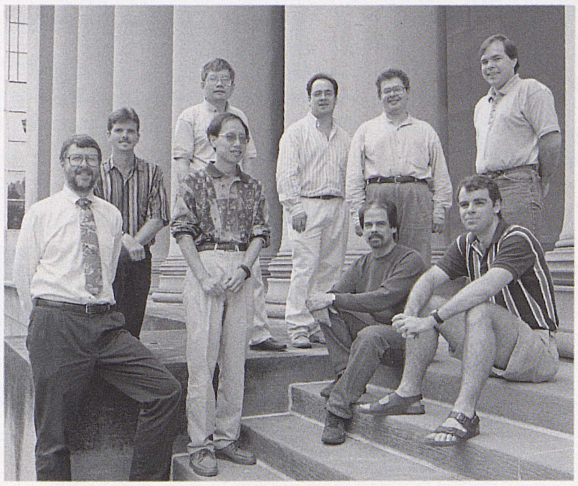

The Parallel Applications Group: (front, l to r) Bruce Loftis, Bin Chen (Cray Research Inc.), Dave O'Neal, Bill Young, (rear, l to r) Nick Nystrom, Qiming Zhang (CRI), Carlos Gonzalez, Sergiu Sanielevici (manager), Rich Graham (CRI). Over the past three years, this group has led PSC development of applications software for the CRAY T3D and T3E, and they are now among the world's leading experts in using scalable, parallel processing to get practical scientific results.

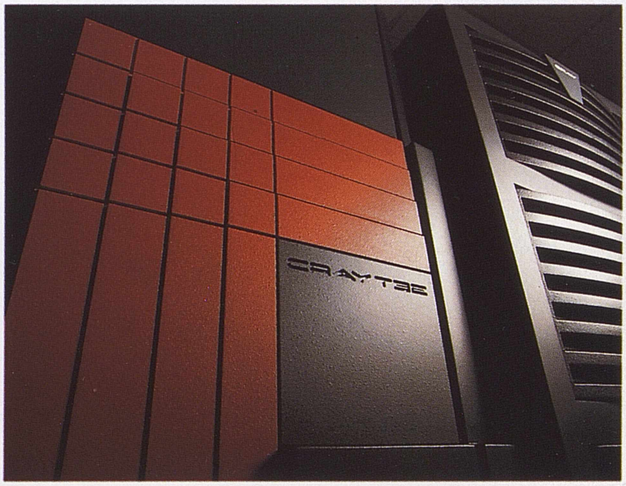

Debut in Pittsburgh: The CRAY T3E

The world's first CRAY T3E was installed in PSC's machine room at Westinghouse Energy Center in March 1996. Applications were running successfully by mid-April. This system will succeed the CRAY T3D as PSC's lead system for scalable, parallel processing, a cost-effective approach to scientific computing that teams tens, hundreds or thousands of processors to work simultaneously on the same task.

Critical to exploiting the potential of scalable, parallel processing is the development of applications software. PSC is one of the leading centers worldwide addressing this need, and the new system will build on PSC's successful work with the T3D. Since August 1993, when the first T3D arrived at Pittsburgh, Cray Research Inc. (CRI) and PSC joined forces in a concerted effort that has made a number of important applications available on the new architecture — in quantum chemistry, materials science, protein sequence-analysis and other areas.

The center will replace its T3D when the T3E expands to 512 processors. The continuity from the T3D to the T3E, which has faster processing and much larger memory, will enhance PSC's ability to attack leading-edge research while maintaining its multi-year investment in highly optimized applications programs. "With this system," said PSC scientific directors Michael Levine and Ralph Roskies, "we expect research productivity to improve by a factor of three to four. This will keep PSC, the NSF supercomputer centers program and American researchers in a world leadership position."

In 1995, PSC added two CRAY J90 systems to its production environment. These systems provide cost-effective, entry-level supercomputing while reserving the center's CRAY C90 for work that requires greater capability. In February 1996, PSC held its first J90 training workshop. PSC also established a J90 Affiliates Group as a forum for communication and support among J90 research sites.

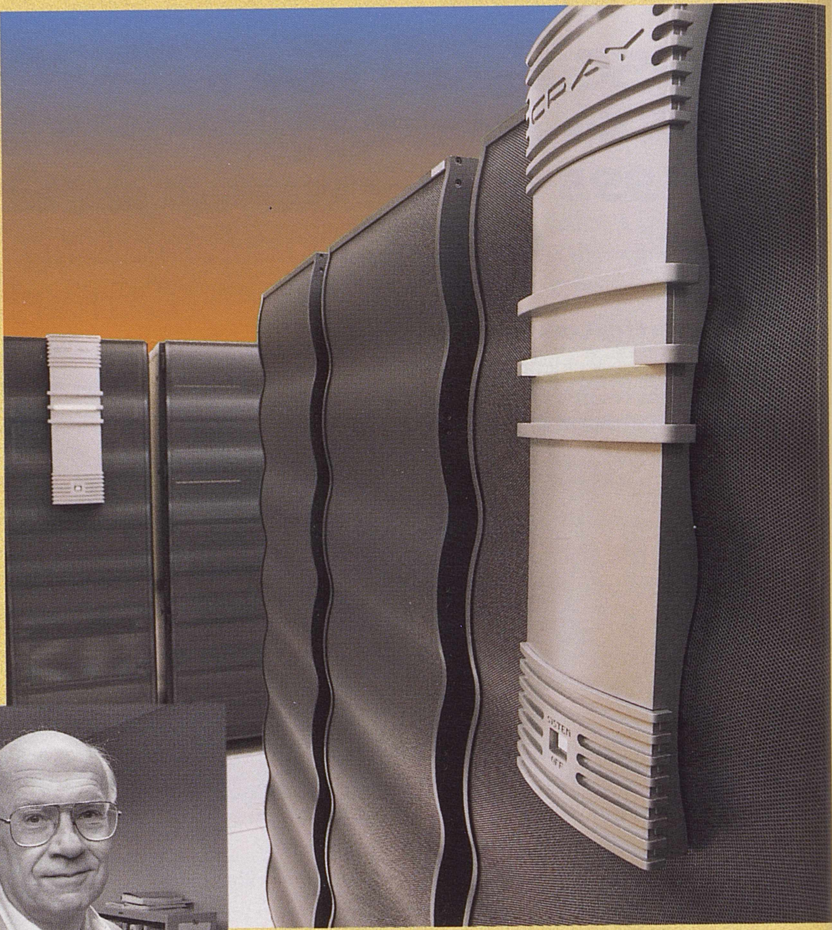

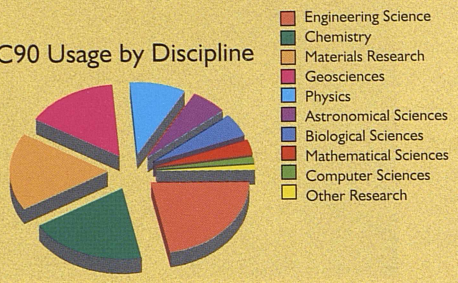

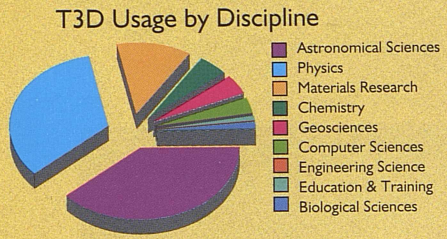

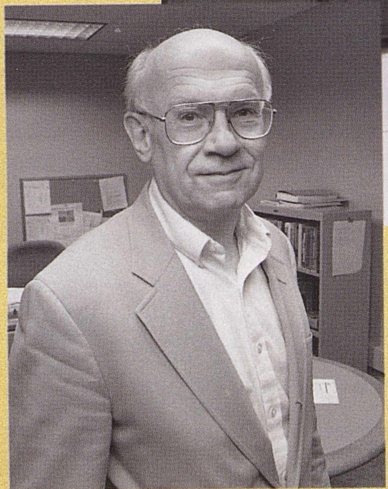

Bob Stock, director of user services. Under Stock's direction, PSC's staff supports nearly 600 active research projects. This includes allocating and administering close to 2,000 user accounts, coordination of training workshops, maintenance of more than 12,000 documents of on-line software documentation, "hotline" troubleshooting that handles 500 to 700 calls a month, and technical consulting on large-scale "landmark" projects.

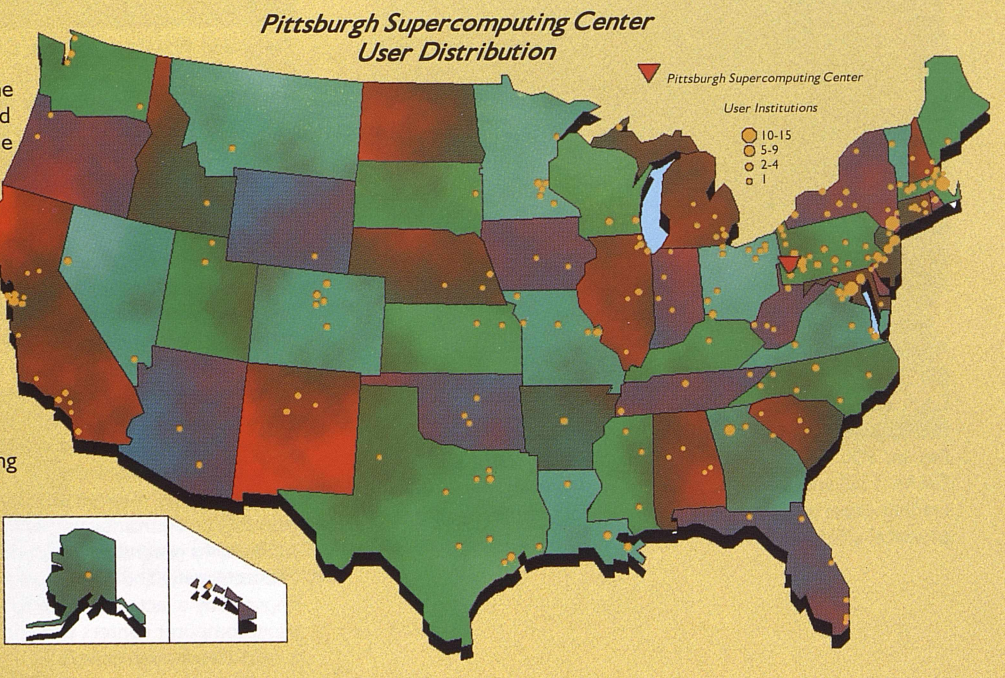

Established in 1986 with a grant from the National Science Foundation supplemented by the Commonwealth of Pennsylvania, the Pittsburgh Supercomputing Center is a joint effort of Carnegie Mellon University and the University of Pittsburgh together with Westinghouse Electric Corp.

To date, more than 11,400 scientists and engineers at more than 670 universities and research centers (color dots) in 49 states and the District of Columbia have used the center's computing resources to advance their research. This work has resulted in more than 3,000 published papers in science and engineering journals. Researchers access the center's resources primarily via the Internet.

Biomedical Supercomputing

In September 1995, the National Institutes of Health awarded a five-year renewal grant of $6 million to PSC. This grant supports new biomedical research by PSC scientists and it continues support for PSC programs to develop software and provide consulting and training in supercomputing techniques to biomedical researchers.

Since its inception, PSC's biomedical program has provided access to computing resources for more than 800 biomedical research projects involving nearly 1,700 researchers in 39 states and the District of Columbia. The center's workshops on computational biology have trained more than 800 researchers in the use of high-performance computing for biomedical research, in such areas as sequence analysis in genome research, the structure of proteins and DNA, and biological fluid dynamics.

"This program fills a crucial need in biomedical research," said Caroline Holloway, director of the Biomedical Research Technology Program of NIH's National Center for Research Resources, which approved the grant. "It bridges the gap between complex biomedical problems and the unique capabilities of supercomputers. PSC provides the national biomedical research community with the world's most advanced high-performance computing resources. It provides excellent user support, and it expands the range of biomedically-relevant software, databases and visualization capability that are available to biomedical researchers."

PSC's biomedical program was initiated in 1987, when the center received a three-year, $2.2 million grant from NIH, the first grant of this kind in the country. In 1990, NIH funded an expansion of PSC's biomedical program with a $6 million, five-year grant establishing the center as an NIH Research Resource. The current grant renews this role, funding five core research projects that advance development of new technologies in computational biology.

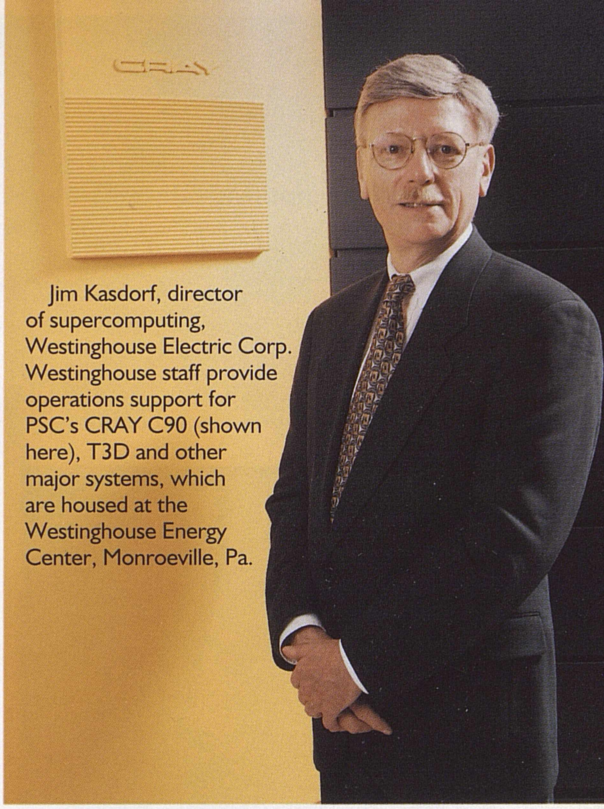

Jim Kasdorf, director of supercomputing, Westinghouse Electric Corp. Westinghouse staff provide operations support for PSC's CRAY C90 (shown here), T3D and other major systems, which are housed at the Westinghouse Energy Center, Monroeville, Pa.

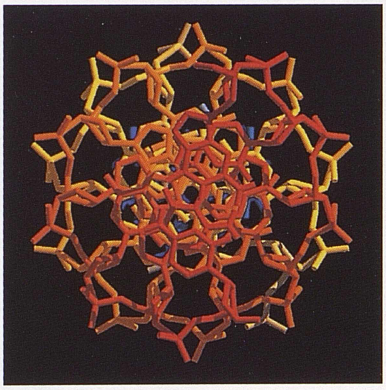

Viewed with ChromaDepth 3D glasses, these images of DNA create a stunning 3D illusion. PSC scientist David Deerfield created them as part of a minute-long video for "Journey into the Living Cell," a multimedia program developed for the Carnegie Science Center by Carnegie Mellon's Center for Light Microscope Imaging in Biotechnology.

"This program fills a *crucial* need in biomedical research."

The Biomedical Group (l to r), Nigel Goddard, Marcela Madrid, Hugh Nicholas, David Deerfield (manager), Michael Crowley, Greg Hood, Alex Ropelewski, Nancy Blankenstein and Joe Geigel.

PSC workshops range from an introduction to supercomputing to advanced research methods.

Supercomputing Education and Training

PSC sponsored 16 workshops during the past 12 months, on a range of topics from introductory methods to advanced research applications. PSC workshops are conducted in the Computer Training Center, equipped with 25 Silicon Graphics INDY workstations, each with 32 megabytes of memory and a high-resolution color monitor.

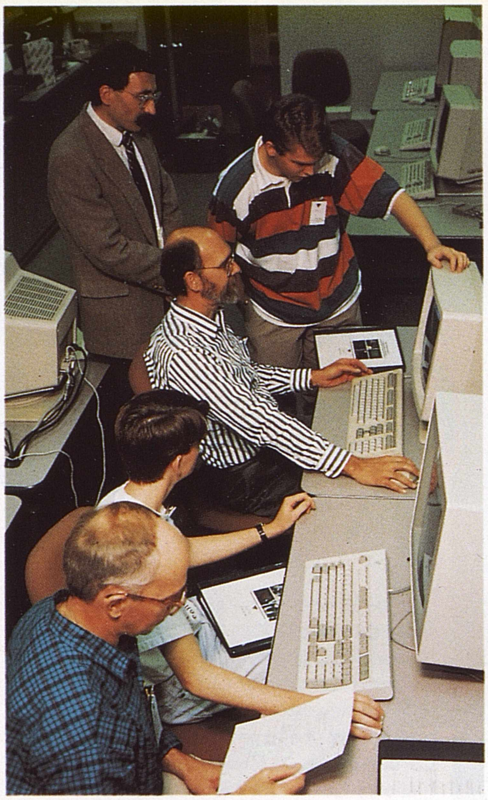

Pittsburgh Supercomputing Center Workshops
(offered at least once July 1995 - June 1996)

Summer Institute on Parallel Computing
Supercomputing Techniques: Introduction to Cray C90
Supercomputing Techniques: Cray J90
PVM and Distributed Computing
Organization of Codes for Cray MPP Systems
Optimized Medium-Scale Parallel Programming
Supercomputing Techniques for Biomedical Researchers
Molecular Mechanics and Dynamics of Biopolymers
Nucleic Acid and Protein Sequence Analysis

Networking and Communications

PSC's networking group this year was actively involved in development of the very high-speed Backbone Network Service (vBNS) that is replacing NSFnet as the main research component of the Internet. Among other activities, the networking staff managed the cutover from NSFnet to a wide-area network connection to MCInet.

Since 1993, PSC has led community networking in the Pittsburgh region through Common Knowledge: Pittsburgh, an NSF-funded national pilot project for integrating networking into the curriculum of an urban public-school system. PSC established local and metropolitan network infrastructure for this project, and PSC staff have worked closely with teachers, administrators and students of Pittsburgh Public Schools to exploit the network environment in teaching and learning.

A national pilot project to integrate networking into the curriculum of an urban public-school SYSTEM.

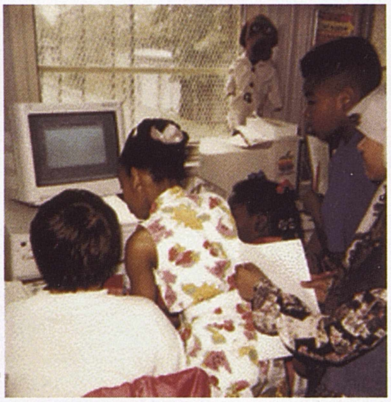

Common Knowledge: Pittsburgh provides Pittsburgh Public School students with valuable tools for problem-solving.

Coordinators of the various areas of PSC activity meet regularly to share information and ideas. Pictured (l to r) are Ken Goodwin (hardware), Rob Pennington (advanced systems), Jamshid Mahdavi (networking), Christina Ricci (allocations), Kathy Benninger (hardware systems), Michael Schneider (public information), Dave Kapcin (financial affairs), Alex Ropelewski (biomedical) and Rich Raymond (user consultants).

ALCOA DUPONT USX WESTINGHOUSE

Corporate Programs

PSC's corporate programs provide outreach to companies of all sizes throughout the United States. The goal is to assist business in gaining access to the knowledge, skilled practitioners and technologies that will allow them to take advantage of high-performance computing. Many corporations maintain a working relationship with PSC, among them ALCOA, Dupont, USX and Westinghouse.

Bill Reight, corporate programs.

Information at the Pittsburgh Supercomputing Center
(412) 268-4960
http://www. psc.edu

Proposals for Computing Time
Christina Ricci
grants@psc.edu

Biomedical Program
Nancy Blankenstein
biomed@psc.edu

Workshops & Summer Institute
Ken McLain
mclain@psc.edu

Corporate Programs
Bill Reight
reight@psc.edu

Newsletter and Documentation
Vivian Benton
benton@psc.edu

Congressman Robert S. Walker (left), chair of the House Committee on Science, toured the PSC machine room at the Westinghouse Energy Center in August 1995. Here, Joel Welling (seated) demonstrates scientific visualization as executive director Beverly Clayton and scientific directors Michael Levine and Ralph Roskies look on.

Ten Years of Leadership

June 9, 1986: With its CRAY X-MP/48 installed and running at Westinghouse Energy Center, PSC officially opens to researchers.

May 1987: ALCOA becomes PSC's first corporate partner.

September 1987: NIH awards PSC a three-year, $2.2 million grant to support supercomputing in biomedical research. PSC is the first supercomputing center to receive an NIH grant for this purpose.

November 1989: PSC user Gregory McRae wins the first Forefronts of Large-Scale Computation award for his computational modeling of large atmospheric systems.

January 1990: The first CRAY Y-MP/832 available outside a government research laboratory is installed at PSC.

April 1990: The Connection Machine CM-2, PSC's first massively parallel system, is installed.

August 1990: PSC receives a five-year $6.1 million grant from NIH to renew its program of support for supercomputing in biomedical research.

December 1990: PSC introduces the Andrew File System for use in high-performance computing.

February 1991: For the first time anywhere, PSC implements direct high-speed data transfer between heterogeneous supercomputing systems, the Y-MP and CM-2, demonstrating "superlinear" speedup on significant scientific applications.

November 1991: University of Pittsburgh biologist John Rosenberg and three colleagues win the 1991 Forefronts of Large-Scale Computation award for their work at PSC on how proteins recognize and bind with DNA.

January 1992: The Connection Machine CM-5 is installed at WEC.

June 1992: Westinghouse wins the 1992 Computerworld-Smithsonian Award in science for its work with PSC.

October 1992: PSC is the first non-government site in the United States to receive a CRAY Y-MP C90.

December 1992: PSC receives "grand challenge" research funding to collaborate with the Center for Light Microscope Imaging and Biotechnology on developing an Automated Interactive Microscope.

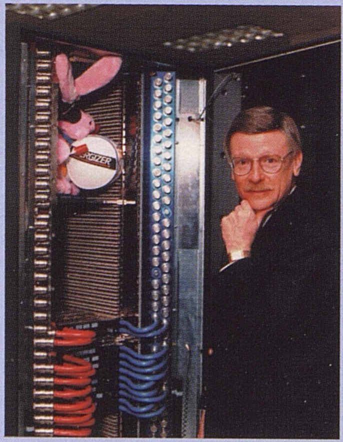

June 1993: PSC wins the 1993 Computerworld-Smithsonian Award for science.

June 1993: PSC, the University of Pittsburgh and Pittsburgh Public Schools receive $2 million from NSF for Common Knowledge: Pittsburgh, a national pilot program to integrate computer networking into the K-12 educational curriculum of a large urban public school system.

September 1993: PSC unveils the world's first CRAY T3D system.

October 1993: PSC receives "grand challenge" funding to collaborate with scientists around the country on astrophysics research.

January 1994: The 1993 Nobel Prize in Physics is awarded to PSC user Joseph H. Taylor for his 1974 discovery of the first binary pulsar. In research at PSC, Taylor has discovered more than 20 pulsars in three years.

June 1994: The 1994 Computerworld-Smithsonian award for Breakthrough Computational Science is awarded to PSC users Charles Peskin and David McQueen for their development of a three-dimensional computational model of blood flow in the heart.

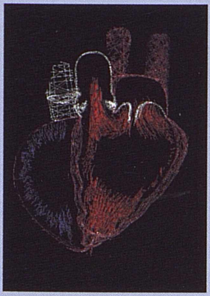

November 1994: PSC user Charles Peskin wins the 1994 Sidney Fernbach award for his modeling of blood flow in the heart.

September 1995: NIH awards $6 million to PSC to continue PSC's program of research and support for biomedical research.

November 1995: PSC user Paul Woodward receives the 1995 Sidney Fernbach award for his investigations of turbulent convection in the sun.

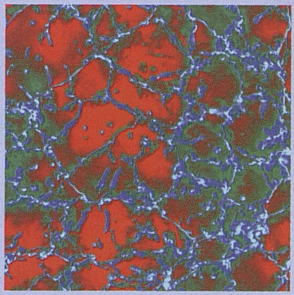

April 1996: The world's first CRAY T3E system is installed at PSC and begins running parallel applications.

May 1996: Four of five finalists for the 1996 Computerworld-Smithsonian award in science are collaborations with PSC, with the award going to the Center for Light Microscope Imaging and Biotechnology.

Scientific Visualization

The Data Intensive Systems Group, (l to r) Anjana Kar, Grace Giras, Joel Welling, Art Wetzel, Greg Foss and Phil Andrews, manager. At left is the "I-Screen," a rear-projection, wide-screen monitor with stereo sound and HDTV capability that provides full-screen display from a workstation. Using stereo images, a head tracker and 3-D wand, it allows creation of a semi-immersive environment.

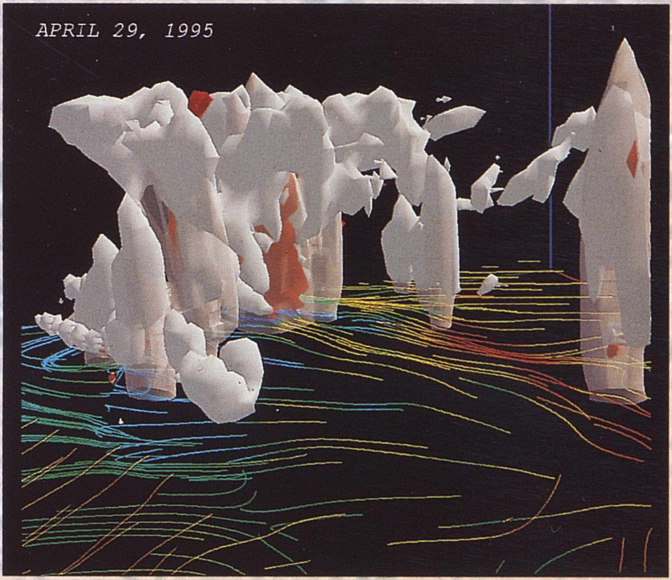

Twister Predictor

During spring 1995 and again in 1996, the Center for Analysis and Prediction of Storms at the University of Oklahoma, used the CRAY T3D to produce a daily weather forecast for parts of Oklahoma (see p. 22). PSC's visualization group processed data from selected forecasts into 3-D animations. This display shows a 152 x 152 kilometer area 35 kilometers in altitude viewed from the southeast. The storm structure shows regions of high water density (white) with columns of high vorticity (red) that potentially spawn tornados. The streamlines, colored according to temperature, show ground-level wind velocity.

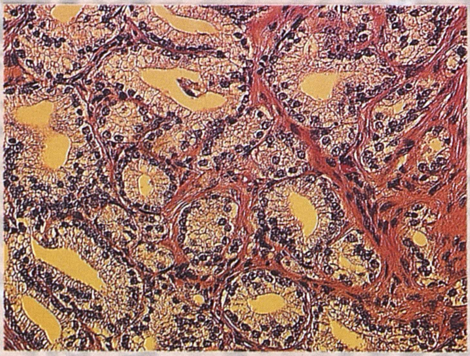

Diagnosing Prostate Cancer

In collaboration with pathologist Michael Becich, University of Pittsburgh Medical Center, PSC's Visual Information Systems Group is developing computerized image classification and pattern-recognition methods to aid in diagnosing prostate cancer. This sequence shows a microscope image of a cancerous prostate section (left), a representation that uses color to differentiate computer-recognized tissue regions (lower left) and additional structural-recognition processing that outlines glandular features (below). From this last representation, statistical analysis of relative cell locations evaluates the degree of malignancy.

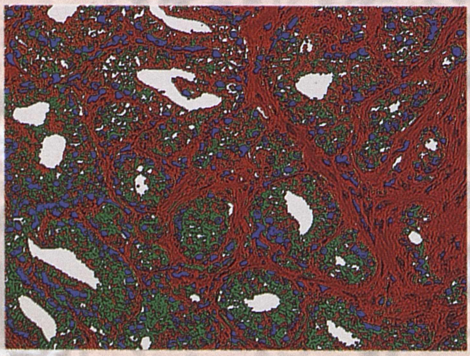

Disks and Jets in Space

Large rotating disks of dust and gas, called "accretion disks," occur in solar systems and in the cores of galaxies. Such disks usually rotate around a central mass, such as a black hole, and often include intense jets of gas and radiation along the rotation axis. PSC visualization specialist Greg Foss created this representation for an animation that introduces research on accretion disks and jets.

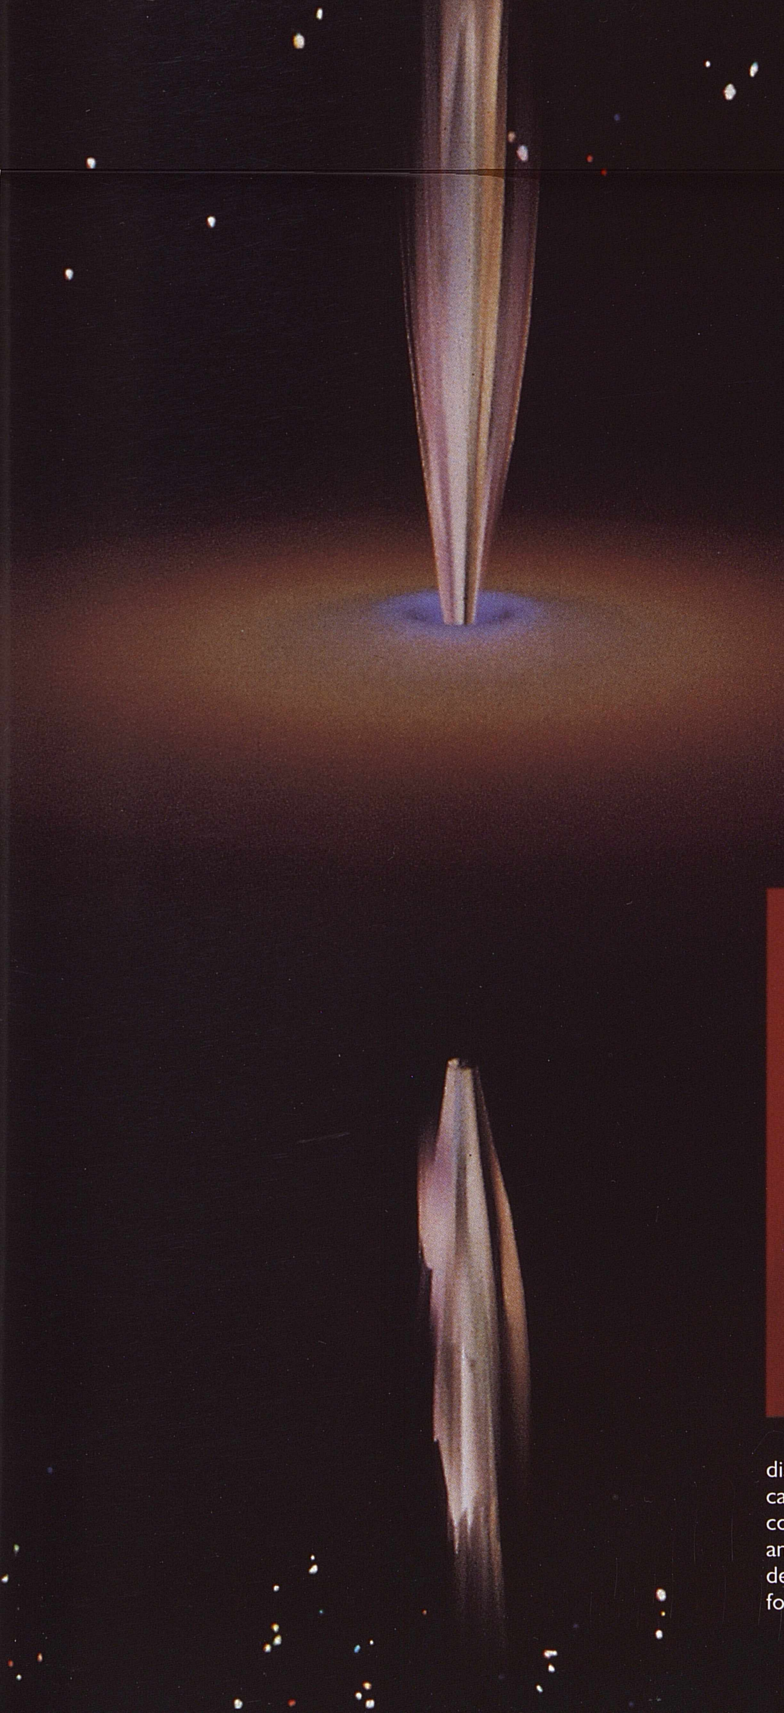

Astrophysicist John Hawley uses the CRAY C90 to examine the forces acting on matter in an accretion disk. These two images represent vertical slabs from an accretion disk and show magnetic field lines (left) in relation to turbulent dynamics (below). In some regions, matter moves toward the disk center (blue) while in others (red) it moves toward the edge.

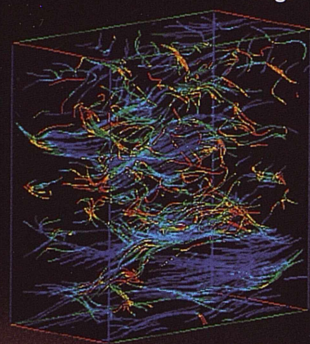

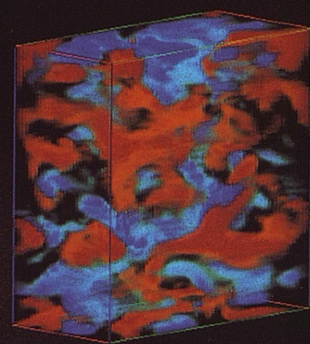

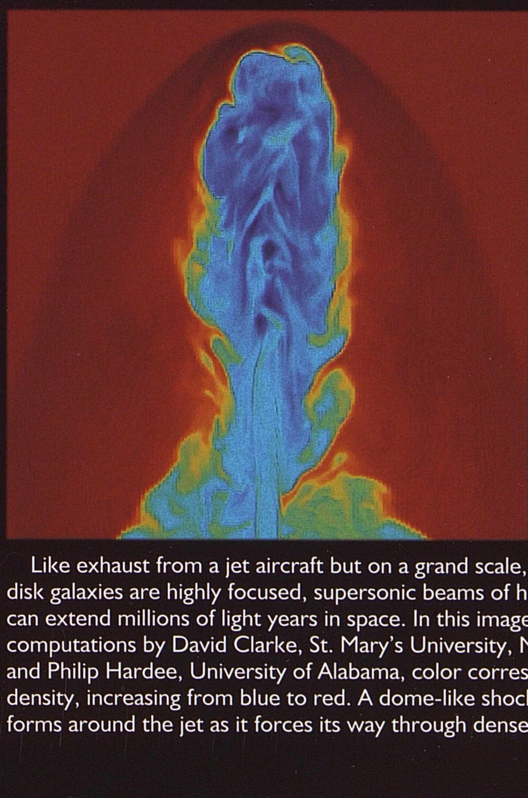

Like exhaust from a jet aircraft but on a grand scale, jets from disk galaxies are highly focused, supersonic beams of hot gas that can extend millions of light years in space. In this image, from computations by David Clarke, St. Mary's University, Nova Scotia and Philip Hardee, University of Alabama, color corresponds to density, increasing from blue to red. A dome-like shock wave forms around the jet as it forces its way through denser space.

Long Distance Charges

One-Two Punch

A software innovation combined with advances in supercomputing hardware has opened a door to better understanding of the fundamental molecules of life. That's the news from pharmaceutical chemist Peter Kollman and his colleague Tom Cheatham.

In computations that simulate the structure of DNA and its dance-like oscillations inside living cells, Kollman and Cheatham used a software innovation called particle-mesh Ewald (PME). Developed by Tom Darden of the National Institute of Environmental Health Science and initially tested at Pittsburgh, PME is an efficient, accurate method to account for the electrical attractions and repulsions between atoms that aren't bonded to each other in a large biomolecule. "This method makes a huge difference," says Kollman. "It provides stability and it leads to dramatic structural improvement over prior methods."

Working closely with Darden and Cheatham, PSC biomedical scientist Mike Crowley implemented PME on the CRAY T3D, bringing unprecedented computing capability to bear on large biomolecule research. Applying this one-two punch — PME and the T3D, Kollman and colleagues have simulated DNA, RNA and proteins with results that herald new possibilities for computational biology.

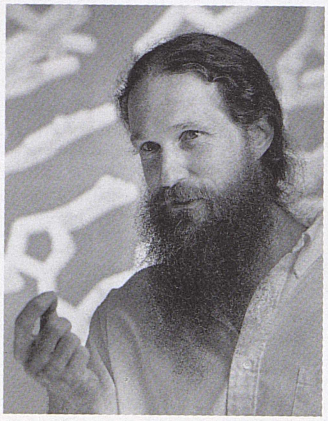

PSC biomedical scientist Mike Crowley, whose efficient, parallelized routine for three-dimensional fast Fourier transforms was instrumental in implementing particle-mesh Ewald on the CRAY T3D. Crowley also has AMBER with PME running on the CRAY T3E.

"Using the T3D," says Cheatham, "is the difference between getting results in less than a week versus months on a workstation. With PME, this opens a new level of what we can investigate in nucleic acids. We can look at larger protein-DNA complexes, because we now have a way to properly represent the DNA. There's a host of mechanisms in the body where these are important, including gene regulation and DNA repair."

Such simulations, which among other things can lead the way to finding new drug therapies for genetic diseases, are now underway in Kollman's research group.

Problem Solved: Electrostatics and Cutoffs

Simulating a large protein or DNA molecule surrounded by water often involves so many atoms (10,000 or more) that before PME it cost too much in computing time to calculate all the forces acting between atoms not bonded to each other. These push-pull interactions, which arise due to positive and negative charges, are called electrostatics. The same basic effect accounts for static electricity. The standard approach has been to assume that these forces, which decrease with distance, don't make much difference when two atoms are farther apart than a certain "cutoff radius" — usually around 10 angstroms.

This assumption made it possible to get worthwhile results, but problems tended to arise as researchers extended their simulations, especially with highly charged molecules such as DNA. Often, it's desirable to track biomolecular motions over a period of nanoseconds (billionths of a second), yet in many cases the molecule would unravel, losing its structural integrity, before the simulation could get that far. Starting in 1993, in a series of computations using PSC's CRAY C90, Darden zeroed in on electrostatics as the source of this problem and devised PME as a way to fix it.

The Spine of Hydration

In recent computations, Cheatham and Kollman simulated a DNA sequence composed of the same base-pair (A-T) repeated ten times. This image shows the average structure of this sequence overlapped with two related snapshots from the final nanosecond of a two nanosecond simulation. Clearly visible, twisting from lower left to upper right, is DNA's "spine of hydration" in the so-called "minor groove" that runs between the double-helical ridges. Density contours for water show the most probable (red) and slightly less probable (yellow) positions for water molecules. Only since implementation of particle-mesh Ewald on the CRAY T3D, say the researchers, is it feasible to carry out simulations like this that reproduce sequence-specific DNA structure and dynamics.

"This opens a new level of what we can investigate in nucleic acids."

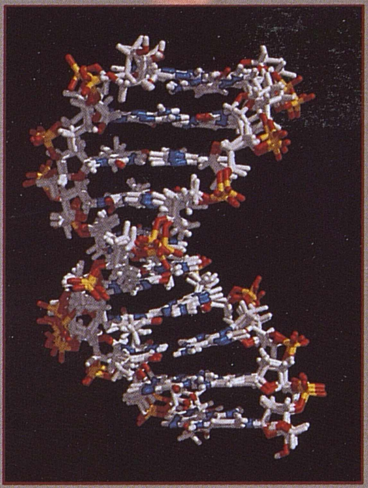

Converged "average" DNA structures obtained from PME simulations starting with A and B form DNA.

A New Wave in Biomolecular Modeling

Darden's central insight was to see how Ewald summation, a method for summing up the charges in a crystal structure, could be combined with a very fast computational method, fast Fourier transforms (FFTs). In early runs on the C90, Darden's innovative synthesis demonstrated remarkable promise: high accuracy — because it computes *all* the electrostatic interactions — with only a modest increase in computing cost over a 10 angstrom cutoff.

The next step was the CRAY T3D. The starting point was a version of AMBER (widely used "molecular dynamics" software) adapted for parallel computing by Jim Vincent and Ken Merz at Penn State and refined by the developers of AMBER. Crowley attacked the task of parallelizing PME and overcame a major obstacle by creating an efficient, parallel three-dimensional FFT routine. With Crowley, Cheatham and Darden working as a team, the pieces came together, and in fall 1995, Cheatham and Kollman ran T3D simulations of DNA and RNA that showed stunning improvement in accuracy and stability over using a cutoff. "Without PME, these molecules just crumpled up in a few hundred picoseconds (trillionths of a second)," says Kollman. "With PME, they remain stable throughout the simulation."

Another study looked at two different forms of DNA. In experiments, DNA occurs in several variant structures depending on its cellular environment, the two most common being A form (high salt, low humidity environment) and B form (low salt, high humidity). Starting with both an A and B sequence, Cheatham ran simulations that converged to the same structure, a close cousin of the B form, within half a nanosecond — exactly what should happen according to theory. "They converged dead on," says Cheatham, "which shows that this method allows you to sample between different conformations and get to the lowest free-energy form — the most representative structure for a given sequence."

Cheatham is exhilarated by the possibilities for new research opened by PME and scalable parallel computing. Likewise, Kollman has a full slate of problems he can barely wait to attack. Some of these studies are already underway, addressing sequence-specific effects, detailed structural variations that depend on the particular base-pairs in a DNA strand. This new level of realism and detail will lead researchers toward solving such mysteries as protein-DNA recognition, a mechanism involved in myriad biological processes, and that depends on a protein's intriguing ability to recognize a specific base-pair sequence along the extremely long helical strands of DNA in our cells. "This is a new wave," says Cheatham, "an exciting time to be in this field." (MS)

References:

T. E. Cheatham and P.A. Kollman, "Observation of the ADNA to BDNA Transition During Unrestrained Molecular Dynamics in Aqueous Solution," *Journal of Molecular Biology* (forthcoming).

T. Cheatham et al., "Molecular Dynamics on Solvated Biomolecular Systems. The Particle Mesh Ewald Method Leads to Stable Trajectories of DNA, RNA and Proteins," *Journal of the American Chemical Society* **11**, 4193-94 (1995).

This research is supported by the National Institutes of Health.

Peter Kollman (left) and Tom Cheatham, University of California, San Francisco.

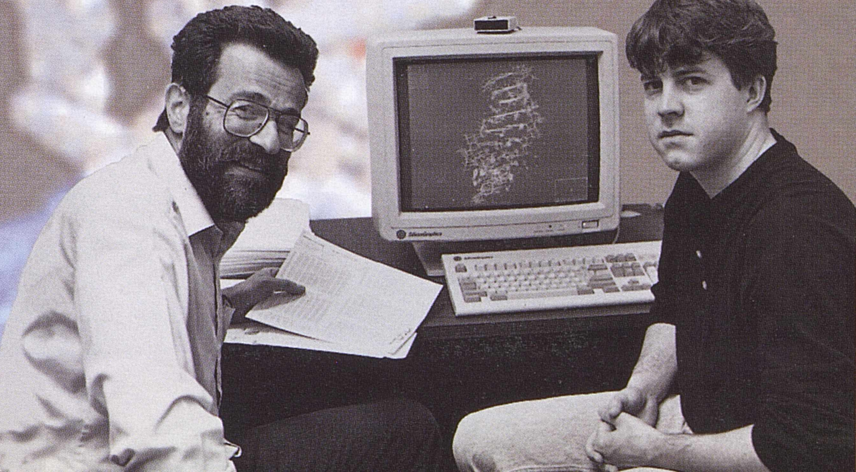

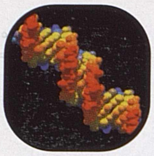

Structure of Proteins & DNA

New Twists in Globs and Zippers

Protein Folding

Anyone who's ever twisted, twirled and tied a strand of shoe-lace licorice knows it can be twisted, twirled and tied only so many times before becoming a tight-knit glob. In biology, the structural equivalent of that glob is the protein, and the twisting process is called protein folding.

Whether serving as enzymes or fundamental units of tissue, proteins are key to successful cellular processes. The progression that takes a loose collection of linked molecular chains and turns it into a snugly-packed, ordered bundle, is an area of keen interest, because only after folding occurs can a protein carry out its assigned task.

Using PSC's C90 and T3D, Charles Brooks, Erik Bozcko and William Young are helping unravel the mystery of how these strands of amino acids ravel in the first place. In addition to aiding future drug design efforts, as well as providing road maps for custom-designed enzymes, deconstructing folding offers the chance to better understand crucial biological processes.

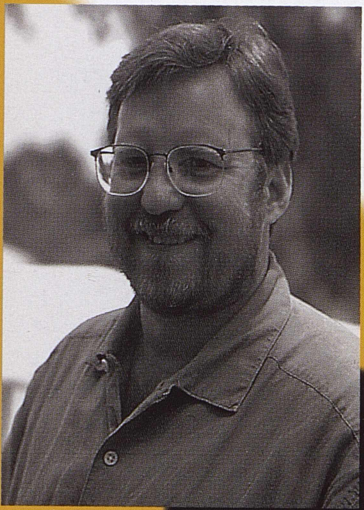

Charles Brooks III,
Scripps Research Institute.

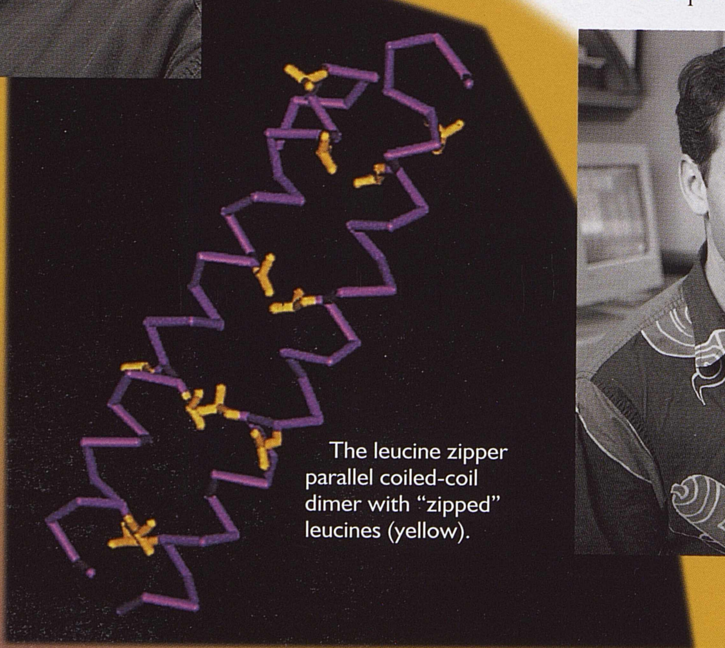

The leucine zipper parallel coiled-coil dimer with "zipped" leucines (yellow).

Filling in the Blanks

Previous experiments have shown that certain proteins, if unfolded then allowed to refold, resume their original shape. But folding occurs so rapidly — a protein will assume millions of different shapes in milliseconds before settling on its final shape, the native state — that only the unfolded and native states have been observed experimentally. For this incomplete photo essay, supercomputing is providing the missing snapshots. "Experiments provide the real information about the system," says Brooks, "and modeling deepens understanding of what the experiments are telling us."

Using the C90, Brooks and Bozcko performed the equivalent of submerging an intact protein in water to examine its folding process. Where past efforts calculated movements for a piece of protein in no medium, Brooks and Bozcko accounted for an entire protein surrounded by water. Their study focused on protein A, which resides on the outer wall of bacterial cells. Their results, published in *Science*, support the energy funnel theory, which explains folding as successive restrictions on a protein's potential to change shape. They also provide the most comprehensive picture yet of protein folding.

The project demanded thousands of multiprocessor C90 hours during a year-and-a-half period. "That was one of the largest problems ever tackled in computational biophysics," says Brooks, "well in excess of 3,000 hours. Erik's work represents a tour de force calculation never before attempted and possible only because of the resources made available by PSC. We wouldn't have done it otherwise."

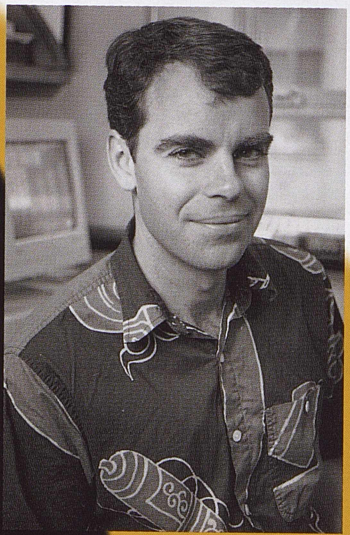

William Young,
Pittsburgh Supercomputing Center
biomedical scientist.

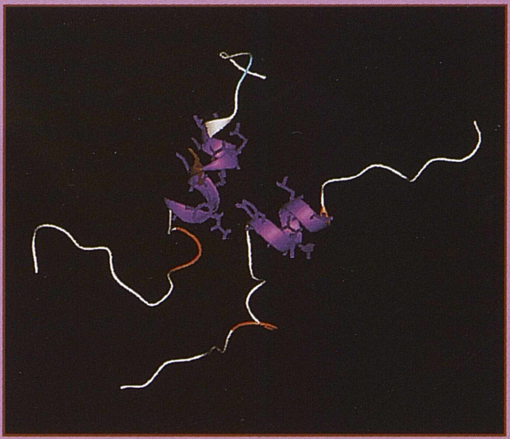

Nucleus of Folding

Initiation of folding (above) occurs at nascent helices, represented as purple ribbon and orange cord. As folding proceeds (right), the coiled-coil dimer starts to form. The orange cord indicates the still unstructured portion of the chain.

Zip-Locked Proteins

In work currently underway, Brooks and Young are using the T3D to study the leucine zipper coiled-coil dimer, a piece of ubiquitous protein architecture. The structure consists of two identical helical strands of protein wrapped around one another and linked by a succession of leucine molecules that "zip" the two units together. Coiled coils can form within the same protein or connect separate proteins and are found in muscle, hair, skin, blood-clotting components and DNA. Though the leucine links keep the helices from touching, the effect of this coupling is akin to two entwined pieces of rotini.

Brooks and Young want to learn whether the helices form before, during or after this molecular marriage. Experiments suggest the coils meet as strands and form helices while wrapping around each other to create the coiled coil. Using supercomputing to heat and hence unfold the protein, the researchers worked backward, capturing snapshots along the way, to examine the progression.

The job totaled about 40,000 T3D processing hours with the molecular dynamics package CHARMM. With "load-balancing algorithms" developed by Young, the parallel version of CHARMM scales well, with only slight performance drop from the communications overhead of using more processors. Going from 64 to 128 processors for this project resulted in an 85 percent speedup, enabling the researchers to meet a self-imposed deadline. T3D parallelism, says Young, was the key to getting results: "We got efficient use of large numbers of processing elements, and that makes a big difference."

Their results provide the most comprehensive picture yet of protein folding.

The computations confirm and expand upon the experimental results, providing more information about the forces involved in protein folding. "Simulations show that these two helical strands are all stretched out, except for one or two turns," says Brooks. "This small amount of helical structure, which falls below the detection limits of lab tools, is the nucleus for folding." By showing that folding depends on the helical interface, the simulations underscore the interplay between experiment and computation, and they put the ball back in the experimentalists' court. If experiments confirm the computational finding, says Brooks, "we will have learned much more about how this protein assembles." (JCW)

References:

Erik M. Bozcko and Charles L. Brooks III, "First-Principles Calculation of the Folding Free Energy of a Three-Helix bundle Protein," *Science* **269**, 393-96 (1995).

Charles L. Brooks III and William S. Young, "Unfolding of the GCN4 Leucine Zipper Dimer Explored Via Molecular Dynamics," (manuscript in progress).

This research is supported by the National Institutes of Health.

Getting to the New World

Waves of Migration

At some point within the last 35,000 years, prehistoric humans walked from the tip of Asia across a land bridge to Alaska, from where they spread through the unpeopled continents of North and South America. What else do we know about the arrival of these earliest Native Americans? Not much. The scant evidence available from archaeological sites suggests it could have been anywhere from 15,000 to 33,000 years ago. One widely held theory, based in part on linguistic analysis, holds that the migration occurred in three distinct waves, related to three well defined groupings of Native American languages.

Until fairly recently, these and similar questions of human origin — are we more like gorillas or chimps? — were the almost exclusive province of archaeologists and fossil-hunter paleontologists like the famous Leakeys. But since about 15 years ago, molecular biologists have been stirring the pot with a trove of new information provided by their ability to deduce evolutionary history from DNA.

University of Pittsburgh geneticist Andrew Merriwether is one among this new breed of genetic detective. Merriwether leads a team of researchers carrying out the most extensive survey yet undertaken of DNA from Native American populations, both living and ancient. Using the CRAY C90, he has analyzed DNA sequences, to identify and map genetic differences, on a scale that would be unthinkable without supercomputing. His recent findings, published in collaboration with colleague Robert Ferrell, are reshaping ideas of how the New World was peopled. "The distribution of genetic patterns," says Merriwether, "best fits a single wave of migration."

Andrew Merriwether and Robert Ferrell,
University of Pittsburgh, Graduate School of Public Health.

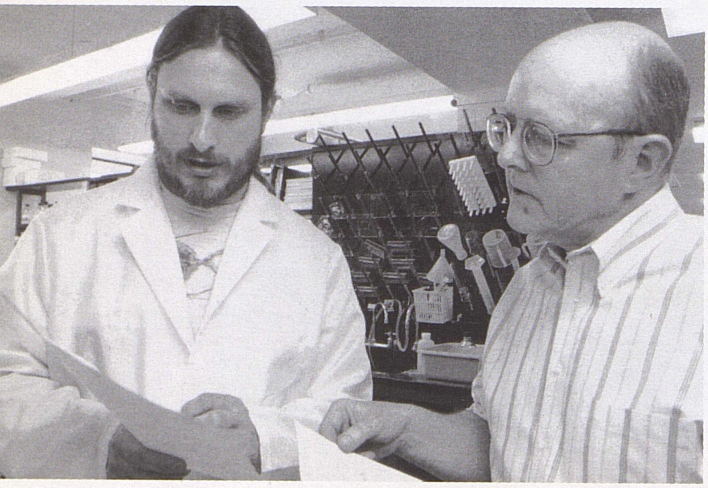

The most reliable results yet *tracing* Evolutionary relationships among Native Americans.

Mitochrondrial Eve

Merriwether's research is based in methods pioneered about 10 years ago by scientists at the University of California, Berkeley, who realized the advantages of using mitochondrial DNA to study evolution. Although most of our DNA resides in cell nuclei, another part of the cell, the mitochondria, also contains relatively short DNA strands. Research has shown that slight changes from generation to generation accumulate steadily in this DNA, like a fast-ticking evolutionary clock — making it possible to extrapolate backward in time.

Mitochondrial DNA, furthermore, is inherited from the mother alone. Nuclear DNA, by contrast, bears the genetic imprint of two parents, four grandparents, eight great-grandparents, etc. Because only one parent contributes to mitochondrial DNA, genetic changes trace backward in a single line, one person in each generation, making it feasible in theory to trace all humans to a single ancestor — so-called "mitochondrial Eve." Several major studies have used these methods to arrive at the still controversial conclusion that all humans are descended from people living in Africa about 200,000 years ago.

Results: A Single Wave

Early research with human mitochondrial DNA showed that genetic patterns cluster by geographical regions and that relationships between populations — evolutionary trees — can be established from these patterns. Some of this research shows genetic relations between Native Americans and Asians, supporting the theory of Asian origin. Merriwether's work is in the same vein, using Native American mitochondrial DNA to shed light on the migration of humans to the New World. "Using these molecular genetic techniques," he says, "we're attempting to trace the ancient migratory routes and determine how long ago contemporary populations separated from each other."

Distribution of Native American Populations

This map shows populations included in Andrew Merriwether's study of evolutionary relationships among Native Americans. Many populations (red dots) were not included in prior studies. Other data (black dots) reported in the literature was also analyzed along with the new populations. The populations in outlined text indicate sites where ancient DNA was collected from teeth, bone fragments and mummified tissue.

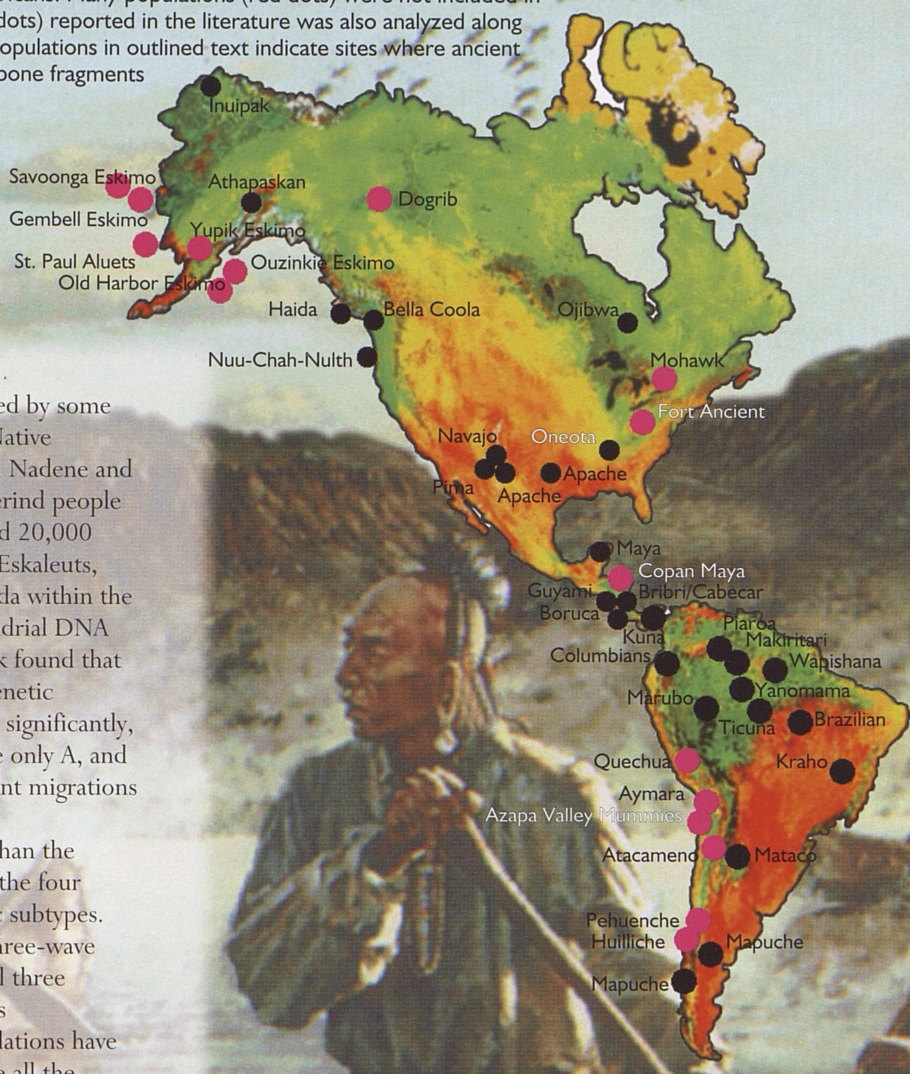

The three waves of migration proposed by some researchers correspond to three main Native American language groups — Amerind, Nadene and Eskaleut. According to this theory, Amerind people arrived about 30,000 years ago, followed 20,000 years later by the Nadene and then the Eskaleuts, who came to Alaska and northern Canada within the last 7,000 years. Some earlier mitochondrial DNA research supports this theory. This work found that all Native Americans come from four genetic lineages, labeled A through D, and that, significantly, Amerinds have all four lineages, Nadene only A, and Eskaleuts A and D — suggesting different migrations at different times.

Working with a much larger dataset than the earlier studies, Merriwether found that the four lineages divide into nine distinct genetic subtypes. Furthermore, directly challenging the three-wave theory, all four lineages showed up in all three language groups. "The key finding," says Merriwether, "is that many of the populations have all four lineages and a number even have all the subtypes. And all types can be found in North, Central and South America. It isn't realistic to believe that the same lineages ended up in all these populations across two continents by separate migrations."

Merriwether's DNA sample includes 1,300 Native Americans representing more than 40 populations throughout the Americas, along with 300 samples from teeth, bone fragments and mummified tissue at three burial sites, one each in North, Central and South America. By including a broader range of populations and large sample sizes, 50 to 100 individuals per population compared to 10 to 20 in other studies, this research offers the most reliable results yet tracing evolutionary relationships among Native Americans.

The large dataset also presents a demanding analytical task, which is where the CRAY C90 comes in. "When you look at thousands of individuals for long DNA sequences," says Merriwether, "you can't do it without the supercomputer. It makes problems feasible that weren't possible to address before. This methodology's been around awhile, but we didn't have the computing power. It's as simple as that."

As the dust settles and researchers in the field contemplate the likely demise of three-wave theory, Merriwether's team has taken up the task of building an evolutionary tree of the Yanomami Indians. This tribe, who live in the tropical rain forest between Brazil and Venezuela, were isolated from modern contact until the 1960s. Results from this unique living population, unmixed with post-Colombian influence, suggest at least three more genetic lineages for the original Native Americans. (MS)

Reference:

D. Andrew Merriwether, Francisco Rothhammer and Robert E. Ferrell, "Distribution of the Four Founding Lineage Haplotypes in Native Americans Suggests a Single Wave of Migration for the New World," *American Journal of Physical Anthropology* **98**, 411-30 (1995).

This research is supported by the W.M. Keck Center for Advanced Training in Computational Biology and the Wenner Gren Foundation.

Illustration: from Departure at Daybreak by Robert Griffing.

Shocking Research

No one is sure why one Waveform outperforms the other, and this Work offers a Compelling answer to the Question.

Cardiac Mayhem

When curious researchers in the 1800s observed a fibrillating heart, they likened it to a bag of slithering earthworms. What they saw was the heart's electrical action, fundamental to its function, gone completely awry.

That breakdown of the heart's exquisitely coordinated electrical activity is the essence of ventricular fibrillation, which causes more than 350,000 deaths a year in the United States alone. Factions of cardiac cells quietly riot against the normal rhythmic motions of the heart, disrupting the pumping process and impeding the flow of blood to vital organs.

Since the early 1980s, the implantable defibrillator has helped manage these unpredictable events in patients with heart disease. Surgically tucked beneath the skin and connected to the heart by electrodes, the battery-powered device continually monitors cardiac performance. If it detects the helter-skelter — and lethal — activity of ventricular fibrillation, it restores calm by delivering an electric shock.

Life-saving benefits, however, come at a price, and research has focused on building more efficient, less energy-hogging implants. "We don't have any detailed knowledge of what cardiac cells are doing in a physiological sense during defibrillation," says Nitish V. Thakor, professor of biomedical engineering at The Johns Hopkins University. "Computer modeling provides the scientific rigor necessary for such an understanding, without having to do a zillion experiments."

Nitish Thakor, Johns Hopkins University.

From the patient's perspective, adds Matthew G. Fishler, "There are psychological, physiological and practical reasons to reduce the energy levels as much as possible." Repeated defibrillations at increasingly higher voltages are painful, can damage heart and surrounding skeletal tissue and hasten battery replacement, an expensive, invasive surgical procedure that takes a further toll on patients.

Wave of the Future

In 1983 animal experiments, researchers replaced the monophasic wave, the wave of choice for implantable defibrillators, with a biphasic wave. By switching electrode polarity, reversing voltage direction, midway through the shock, the biphasic wave significantly reduced the energy and the number of shocks needed to defibrillate, and most implants now incorporate biphasic technology. Yet no one is sure why the biphasic waveform works better. Thakor and Fishler's work offers a compelling answer to the question.

Under Thakor's direction, Fishler designed a computational model of what happens when either mono- or biphasic waves interact with fibrillating heart tissue. This model is the first of its kind to incorporate the gap junctions by which electrical charges travel from cell to cell, thereby simulating cardiac activity both inside and outside cells. "I can't simulate the whole heart," says Fishler, "so I've isolated a fundamental unit of fibrillation, a two-dimensional spiral rotor, sort of like a pinwheel." Rapidly moving and short lived, these roaming, marauding packs of excitation are believed the source of fibrillation.

Modeling these rotors, says Fishler, is computationally demanding. Generating the rotor itself, which represents 350 milliseconds in real time, required more than 100 hours of C90 processing. With the C90, Fishler overcame several problems he encountered with a Connection Machine CM-2. The C90's global memory allowed a significant performance gain in look-up table routines and the algorithms required to simulate defibrillation shocks are better suited to its vector architecture.

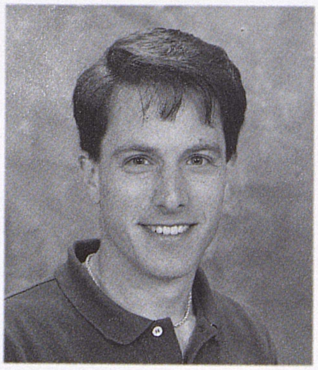

Matthew Fishler

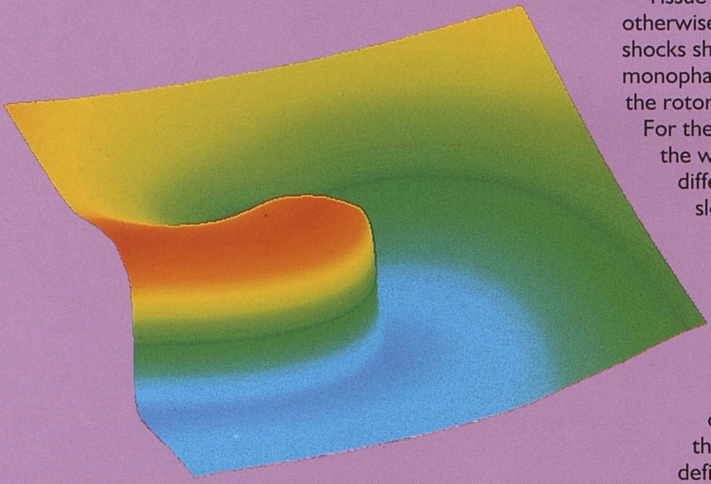

The Shocking Aftermath

Tissue activity two milliseconds after otherwise identical monophasic and biphasic shocks shows contrasting outcomes. For the monophasic shock (left), the leading edge of the rotor (the wavefront) remains steep. For the biphasic shock (below), however, the wavefront is graded, like the difference between a cliff and a gentle slope. The gradedness prevents continued propagation across the tissue by reducing the ability of excited and resting areas to interact, terminating the rotor. In contrast, the monophasic wavefront remains steep enough to maintain its integrity and continue its march through the tissue, foiling the defibrillation attempt.

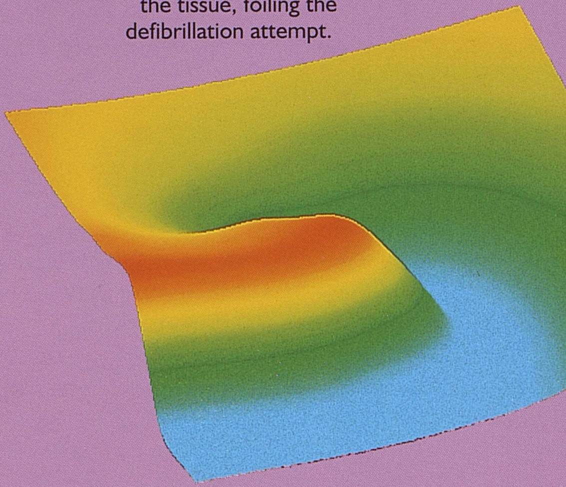

The Right Shock

Defibrillation depends on jolting the majority of heart cells, including fibrillating cells, into a fleeting state of simultaneous excitation. Once the cells collectively return to rest, the sinoatrial node, the heart's pacemaker, resumes its role of rhythm king.

A defibrillation current is weakest in regions furthest from the two implant electrodes. Earlier research showed that a monophasic shock lacked the oomph to excite distant cells and, ironically, could induce further fibrillations, while a biphasic shock had no such shortcomings. To account for this difference, Fishler focused on the distant cells, and the model yielded a new insight.

Normal cardiac cells exist in three states: excited, at rest, or refractory — the interval between excitation and rest. The farther along a cell is in refractory state, the easier it is to excite and, hence, defibrillate. Fishler's simulations showed little difference between mono- and biphasic shocks for cells in early and late refractory: no excitation early, full excitation late. In between early and late refractory, however, biphasics do what monophasics cannot — induce some level of excitation. And the further along in refractory, the greater the excitation, creating a "graded" response. "By switching polarity, we've gained this graded response," says Fishler, emphasizing that "graded amounts of excitation throughout the heart are much better than random pockets of resting and excited regions."

This understanding of the interaction between low-voltage shocks and cardiac cells will aid future research. Just as pacemakers and defibrillators manage heart disease, says Thakor, implants incorporating novel waveforms could help the body's most important muscle do its job more efficiently: "The electrical and mechanical aspects of the heart are linked — good electrical function leads to better mechanical output. If the heart gets damaged from failure, it may be possible to use electrical stimulation to improve mechanical function." (JCW)

References:

M.G. Fishler, E.A. Sobie, L. Tung & N.V. Thakor, "Cardiac Responses to Premature Monophasic and Biphasic Field Stimuli: Results from Cell and Tissue Modeling Studies." *Journal of Electrocardiology* 28 (supp), 174- 79 (1996).

M.G. Fishler, E.A. Sobie, N.V. Thakor & L. Tung, "Mechanisms of Cardiac Cell Excitation with Monophasic and Biphasic Field Stimuli: A Model Study," *Biophysical Journal* (in press).

This research is supported by the National Institutes of Health.

Faster than a Speeding Storm Front

Hailstones and Milestones

April 29, 1995. Around 10 p.m. a vicious thunderstorm rolls into Dallas-Ft. Worth airport. Marble to softball-sized hail pelts planes on the ground. More than 60 commercial jets must be removed from service for repair. The direct insured loss exceeds $20 million, and further loss from canceled flights over the next few weeks runs to $300 million.

What if? The question drives Kelvin Droegemeier's waking days. What if there had been four hours warning? Incoming flights could have been diverted, and with normal flight turnover, about two hours, virtually all the planes would have been removed from harm's way, cutting the airlines' loss and avoiding much of the subsequent inconvenience to travelers. With countless other severe storms as well, including the fierce tornadoes that rampage the Great Plains each spring, better forecasts could greatly reduce property damage and potentially save lives. Since 1986, severe weather has racked the insurance industry with unprecedented losses, and industry studies indicate that more than $14 billion a year could be saved with better forecasts.

Current forecasting gives about 30 minutes warning for storms, with fairly imprecise information about extent and severity. Droegemeier, who directs the Center for Analysis and Prediction of Storms (CAPS) at Oklahoma University, knows it's possible to do better. He and his colleagues are developing new storm-prediction capability that is setting milestones in the field. In 1995, using the CRAY T3D at Pittsburgh, their state-of-the-art computer model, the Advanced Regional Prediction System (ARPS), successfully predicted the location and structure of individual storms six hours in advance, the first time anywhere this has been accomplished.

Zeroing in on Stormy Weather

The weather reports we watch on TV derive from computer models at the National Centers for Environmental Prediction that predict atmospheric structure over the continental United States every 12 hours. The local forecaster extracts from these models to show a regional map, covering perhaps several states, with predictions that weather the next day will be rainy, cloudy, sunny, etc.

In contrast, Droegemeier and his colleagues at CAPS work on storm-scale forecasting, a much tighter focus — a few miles square in space and about 15 minutes in time — that corresponds to the scale on which individual storms evolve. "What we're getting down to," says Droegemeier, "is to say that over Pittsburgh this afternoon from 3:30 to 3:50 there will be a thunderstorm with winds of 30 miles per hour, golfball-sized hail, two-and-a-half inches of rain, and by 3:50 it will be gone. And to give you that forecast six hours in advance."

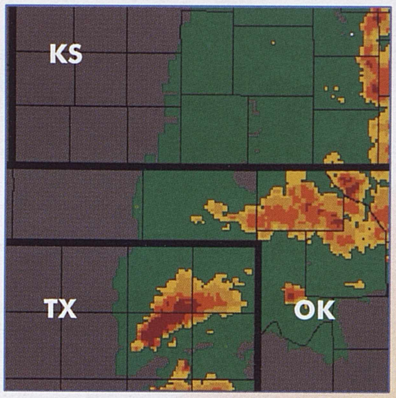

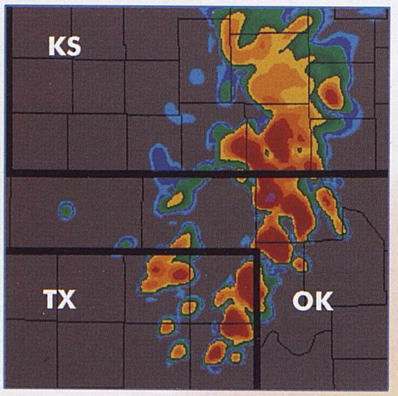

ARPS vs. Reality: June 8, 1995

The ARPS forecast (left) created at 1 p.m. for conditions at 7 p.m. compares well with an actual radar image (above) at 7 p.m. Color indicates rainfall intensity, increasing from light blue to pink. "The model did a remarkably good job," says Kelvin Droegemeier, "getting these storms just about in the right location, especially with the most intensely rotating storms in the northeast Texas panhandle. And it predicted the north-south extent of the storm line up into Kansas. This represents a tremendous success."

Preventable loss from storms in the United States exceeds $14 billion a year.

Miles to Go Before They Sleep

CAPS' goal is to develop storm-scale forecasting to the point where it can be turned over to the National Weather Service early in the next century. Notable among the challenges they have overcome since starting in 1988 is gathering the input data necessary to run a storm-scale model. CAPS innovations now make it possible to deduce all the initializing data — pressure, temperature, wind speed and more — from Doppler radar.

Since 1993, CAPS has carried out experiments during Oklahoma's spring storm season to see how ARPS works in an operational setting. Initializing data each morning feeds the computer model running in Pittsburgh. The output in turn feeds back to forecasters in Norman, Okla. Results have been encouraging. In 1993, running on the CRAY C90 with limited data and a limited version of the model, forecasters used ARPS information in an official National Weather Service forecast.

In spring 1995, CAPS used more ARPS capability, and for the first time exploited parallelism on the CRAY T3D. "The T3D's distributed, globally addressable memory," says Droegemeier, "gave us the ability to run at storm scale over a big area at high resolution." The model did remarkably well (see graphics), especially considering that initializing data was relatively crude — lacking the resolution and accuracy radar can provide. As a result of this success, American Airlines is negotiating a three-year contract with CAPS, investing $1 million to test the new prediction technology as a "smoke alarm" for airports.

CAPS researchers see their goals as within reach, but challenges remain, perhaps none bigger than the extreme computational challenge built into their work. "In meteorology," says Droegemeier, "getting results quickly is essential. If you can't predict the weather significantly faster than it evolves, the prediction has no value. If you're going to create a four-to-six hour forecast, you better do it in half an hour. Parallel computing and high-performance networking are crucial."

In 1995, the T3D took about 75 minutes to run a six-hour forecast. Fine for this stage of development, says Droegemeier, but too long for real-world forecasting. This spring, the model ran with more realistic physics and more complete data, further testing the ability of ARPS to generate forecasts faster than stormy weather changes. "We can do better," says Droegemeier, "and we will." (MS)

Reference:

K. K. Droegemeier et al., "Realtime numerical prediction of storm-scale weather during VORTEX '95, Part I: Goals and methodology," Preprints, *18th Conf. on Severe Local Storms*, American Meteorological Society, San Francisco, 165-68, Jan. 15-20, 1996.

K. K. Droegemeier et al., "Realtime numerical prediction of storm-scale weather during VORTEX '95, Part II: Operations summary and example predictions," Preprints, *18th Conf. on Severe Local Storms*, American Meteorological Society, San Francisco, 178-82, Feb. 19-23, 1996.

This research is supported by the National Science Foundation.

For the first time anywhere, the location and structure of thunderstorms was successfully predicted six hours in advance.

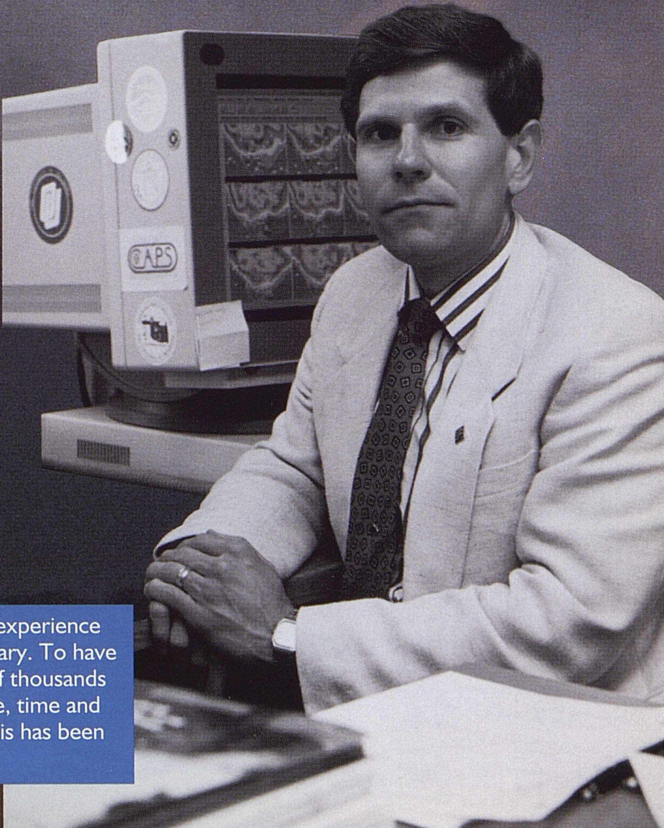

Kelvin Droegemeier, University of Oklahoma. "Our experience computing at PSC has been nothing short of extraordinary. To have a national center with a mission of meeting the needs of thousands of computational scientists and that can dedicate people, time and resources to our problem, storm-scale prediction — this has been incredibly productive."

Zeroing in on Smog

Cleaner Air at Less Cost

Across the United States, cities are working to comply with the National Ambient Air Quality Standard for ozone, a primary constituent of smog. Ozone levels relate directly to the amount of volatile organic compounds (VOCs) dumped into the air by all manner of human activity, from drycleaning to transportation to assaulting underarms with aerosols. Compliance is costly — estimates run as high as $25 billion annually by the year 2000, and the outcome for many control strategies is difficult to anticipate.

VOC Reactivity: A New Approach

Thirty miles above the Earth, the ozone layer protects living things by filtering out the Sun's damaging ultraviolet rays. Earth-bound ozone offers no such benefit. It reduces crop yields by $3 billion annually and contributes to lung disease. The problem occurs when hydrocarbons, a class of VOCs, react in the atmosphere with nitrogen oxides, a combustion byproduct. Mixed by winds and cooked by the sun, they combine to create smog. Automobile exhaust is a primary culprit, simultaneously spewing hydrocarbons and nitrogen oxides.

Their results show how to save tens of millions of dollars while achieving VOC reduction as good or better than current practice.

The current federal effort to manage VOC emissions takes a shotgun approach. Some VOCs contribute more to ozone levels than others, but compliance guidelines target VOCs collectively. This undermines the goal of cleaner, safer air, says air-pollution control engineer Ted Russell. "Since one VOC might be cheaper to manage than another, a city strapped for cash could attain compliance using the most expedient, economical approach while harmful VOCs enter the atmosphere."

Russell and fellow researchers Michelle Bergin and Erik Riedel are using high-performance computing to design scales that can be used on a nationwide basis to help cities target the worst VOCs. Using PSC's DEC Alpha SuperCluster, they have shown such scales are feasible. The new approach could save tens of millions of dollars while attaining VOC reductions equal to or greater than what is achieved currently.

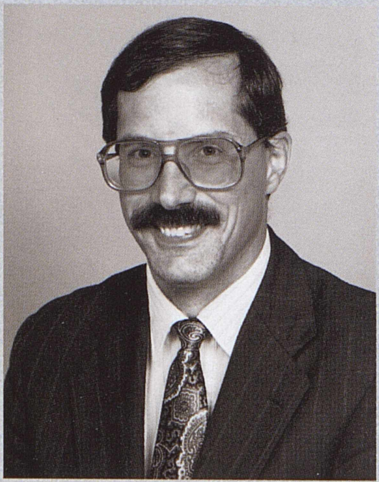

Ted Russell, now at Georgia Tech, carried out this research while at Carnegie Mellon University. For their work using supercomputing to develop strategies for smog reduction, Russell and his colleague Greg McRae of MIT were finalists for the 1996 Computerworld-Smithsonian award in science.

By offering a means to compare VOC reactivity levels in like products, such as gasolines, a reactivity scale can weigh the ozone-causing potential of one VOC source against another. Each VOC is assigned a value that reflects the amount of ozone it will produce. But design depends on the urban setting in question and the level of nitrogen-oxide emissions expected, so a scale for Los Angeles, for instance, might not be useful for New York.

Despite atmospheric differences among cities, says Russell, data suggest differences in VOC reactivities across the nation tend to balance one another out. Using existing scales as a springboard, the Russell team designed a scale that uses a "relative" approach to measuring compounds.

In the case of gasolines, for instance, the most important issue is how an alternative fuel measures up against the status quo fuel. When used for Los Angeles, for example, the relative approach produced results very close to those from the scale specifically tailored to that environment.

"You can use these scales more globally," says Russell, "and you can do that because of the relative relationship between reactivities shown by our calculations." Moreover, results can be used to target VOCs, improving management strategies across the United States. "A relative scale will reveal, for example, that emissions from source x are twice as reactive as source y, so it's twice as beneficial to reduce a pound from source x."

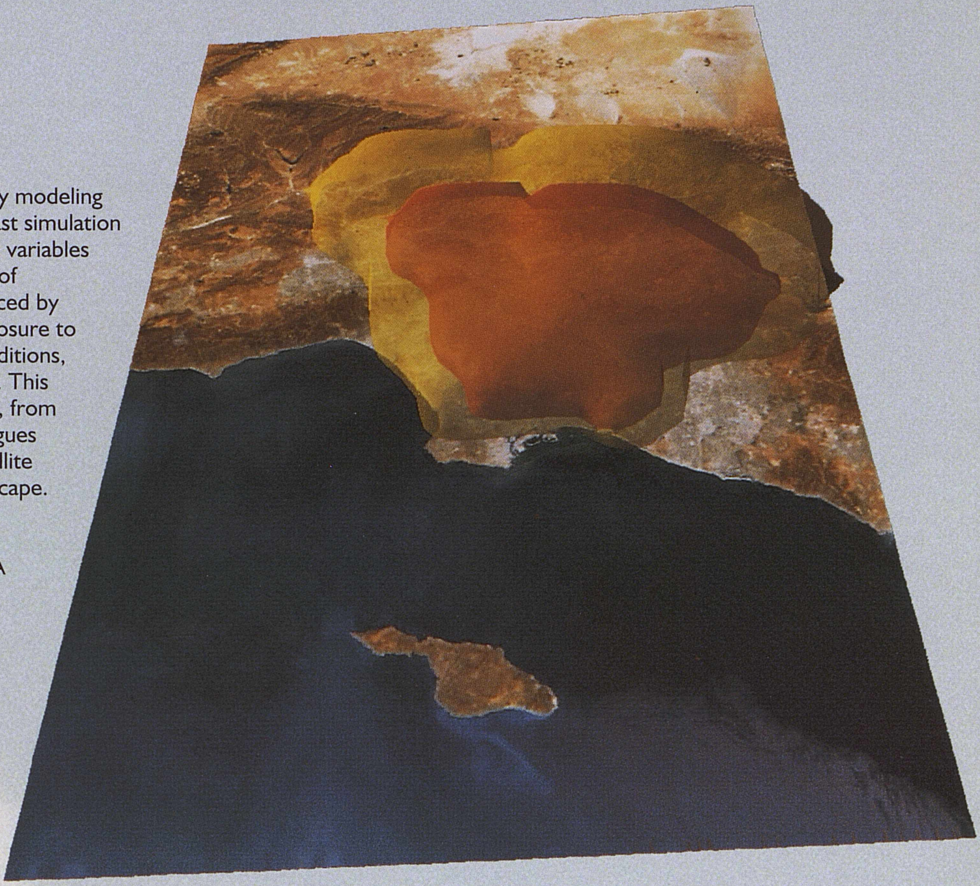

Three-dimensional air-quality modeling provides detail missing from past simulation efforts, accounting for dynamic variables such as movement and mixing of compounds as they are influenced by source location, elevation, exposure to sunlight, time of day, wind conditions, temperature and other factors. This image of the Los Angeles basin, from modeling by Russell and colleagues with data from 1987, uses satellite imagery to represent the landscape. The orange cloud encloses the region where nitrogen oxide concentration exceeds the EPA standard. The yellow cloud encloses the region where concentration is above normal but not exceeding the EPA standard.

A Dynamic Picture

While existing scales were created using atmospheric modeling without spatial variation (zero-dimensional) and based on ozone episodes of short duration, Russell employed 3-D modeling to simulate a three-day episode. The 3-D approach handles movement and mixing of compounds, offering a dynamic rather than static atmospheric representation. "If you envision the study area as a box," says Russell, "the atmosphere in the zero-D model is exactly the same everywhere — at the ground, all the way up, and from one end of the city to the other."

To complete these calculations in time to meet a California Air Resources Board deadline, Russell and his colleagues teamed PSC's SuperCluster with 20 DEC Alpha processors at Carnegie Mellon University. Running continuously on this network of 30 machines, their distributed code finished the modeling in a week. The research produced three new scales: one for predictions of peak ozone scenarios, another to gauge vegetation exposure, and a third to judge exposure to humans. "One of the powers of 3D modeling," says Russell, "is we get a better spatial representation of population exposure, and what you're really interested in is the level of ozone people are exposed to."

Russell looks forward to adapting his models for use on the CRAY T3D. "The T3D will let us develop models faster and link them to other models faster." In particular, his plans include coupling his air pollution models to geographic information systems (GIS), significantly broadening the research possibilities. GIS can hold a huge range of information that makes it possible to better assess exposure to specific populations. "You could look at population by income and gauge their exposure to pollution. Or with the click of a button you could determine the effects of heavy usage of a particular highway system." (JCW)

Reference:
A. Russell, J. Milford, M.S. Bergin, S. McBride, L. McNair, Y. Yang, W.R. Stockwell, and B. Croes, "Urban Ozone Control and Atmospheric Reactivity of Organic Gases," *Science* **269**, 491-95 (1995).

This research is supported by the National Science Foundation, the California Air Resources Board, the Auto/Oil Air Quality Improvement Research Program, the U.S. Department of Renewable Energy Laboratory and the National Aerosol Association.

When North Goes South

Magnetic Flip-Flops

Considering that ships, planes and Boy Scouts steer by it, Earth's magnetic field is less reliable than you'd think. Rocks in an ancient lava flow in Oregon suggest that for a brief erratic span about 16 million years ago magnetic north shifted as much as 6° per day. After little more than a week, a compass needle would have pointed toward Mexico City.

The lava catches Earth's magnetic field in the act of reversing itself. Magnetic north heads south, and — over about 1,000 years — the field does a complete flip-flop. While the Oregon data is controversial, Earth scientists agree that the geological evidence as a whole — the "paleomagnetic" record — proves such reversals happened many times over the past billion years.

"Some reversals occurred within a few 10,000 years of each other," says Los Alamos scientist Gary Glatzmaier, "and there are other periods where no reversals occurred for tens of millions of years." How do these flip-flops happen, and why at such irregular intervals? The geological data, invaluable to show what happened, registers only a mute shrug when it comes to the deeper questions.

For that matter, why is it that instead of quietly fading away, as magnetic fields do when left to their own devices, Earth's magnetic field is still going strong after billions of years? Einstein is said to have considered it one of the most important unsolved problems in physics. With a year of computing on Pittsburgh's CRAY C90, 2,000 hours of processing, Glatzmaier and collaborator Paul Roberts of UCLA took a big step toward some answers.

Their numerical model of the electromagnetic, fluid dynamical processes of Earth's interior reproduced key features of the magnetic field over more than 40,000 years of simulated time. To top it off, the computer-generated field reversed itself.

"We weren't expecting it," says Roberts, "and were delighted. This gives us confidence we've built a credible bridge between theory and the paleomagnetic data." Their surprising results, reported as a cover story in *Nature* (Sept. 21, 1995), provide an inner-Earth view of geomagnetic phenomena that have not been observed or anticipated by theory. Furthermore, the Glatzmaier-Roberts model offers, for the first time, a coherent explanation of magnetic field reversal.

Journey to the Center of the Earth

Roughly speaking, Earth is like a chocolate-covered cherry — layered, with liquid beneath the surface and a solid inner core. Beneath the planet's relatively thin crust is a thick, solid layer called the mantle. Between the mantle and the inner core is a fluid layer, the outer core. According to generally accepted theory — the dynamo theory — interactions between the churning, twisting flow of molten material in the outer core and the magnetic field generate electrical current that, in turn, creates new magnetic energy that sustains the field. "The typical lifetime of a magnetic field like Earth's," says Glatzmaier, "is several tens of thousands of years. The fact that it's existed for billions of years means something must be regenerating it all the time."

How do we know if the dynamo theory is right? To the consternation of our desire to understand what's happening inside the planet we live on, Jules Verne's *Journey to the Center of the Earth* is still fiction. There's no way to penetrate 4,000 miles to Earth's center, nor to monitor fluid motions or magnetism in the outer core.

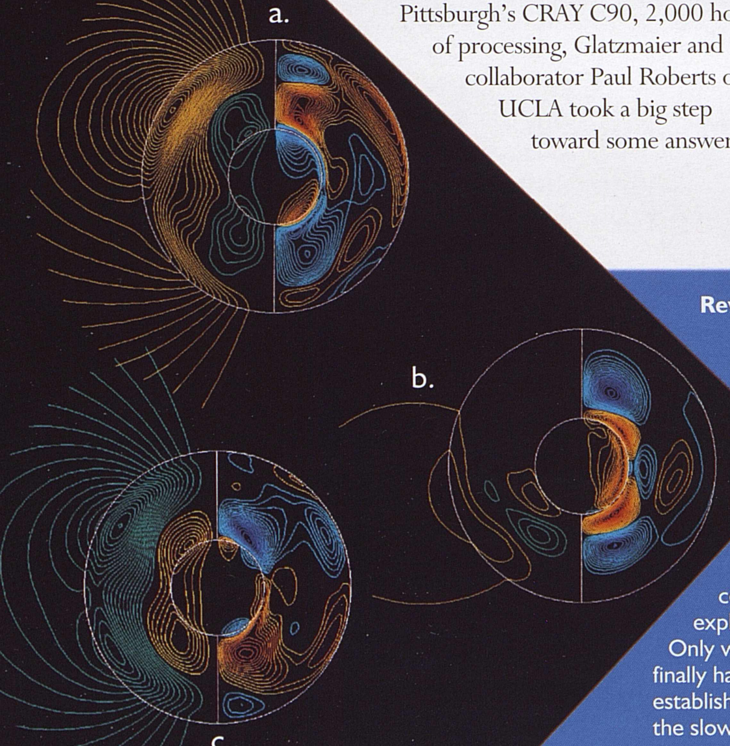

Reversal of Earth's Magnetic Field
Earth's magnetic field evolving for about 9,000 years before (a), during (b) and after (c) the simulated reversal. The outer circle indicates the fluid outer core boundary; the inner circle, the solid inner core. The left hemisphere shows magnetic field contours directed clockwise (green) and counterclockwise (yellow). The right hemisphere shows contours directed westward (blue) and eastward (red),out of and into the plane of the paper. The left hemisphere shows that the field penetrating the inner core is opposed in polarity to the outer core, a feature completely unanticipated by theory. "The outer core polarity," explains Glatzmaier, "is continually trying to invade the inner core. Only when the whole field almost decays away (b), however, does it finally have a chance to diffuse in. Once it does, the opposite polarity gets established. The inner core polarity is the stabilizing force, like an anchor, the slowest thing that can change."

Their model offers the first coherent explanation of magnetic field reversal.

The Glatzmaier-Roberts computational model may be the next best thing to a guided tour of inner Earth. While other models have given good clues that dynamo theory is on track, they have been limited by a two-dimensional approach that required simplifying assumptions. Roberts and Glatzmaier set out to implement a fully three-dimensional model, based on a computer program Glatzmaier developed over many years, that would allow the complex feedbacks between fluid motion and the magnetic field to evolve on their own — in other words, to be solved "self consistently."

Their objectives, in retrospect, were modest. "Mainly," says Roberts, "we wanted to get a geomagnetic field that would maintain itself longer than the decay time. No one's ever done that in a self-consistent manner." After nearly a year running almost daily, as allocated computing time was about to expire, the model produced its Eureka moment.

By itself, the reversal is strong confirmation of the model, and other details — magnitude and structure of the field — also agree well with surface features of Earth's field. The simulation also offers precious insight into the dynamics that sustain the magnetic field and generate reversals. Contrary to what anyone guessed till now, the model shows that in the inner core the magnetic field has an opposite polarity from the outer core, and this stabilizes the field against a tendency to reverse more frequently.

"No one even dreamed about this," says Glatzmaier. "That's the nice thing about a supercomputer. You can just let it do its thing, solve these equations over and over — a large set of variables affecting each other with nonlinear feedback, very hard to figure out. It's a beautiful problem for a supercomputer, and it's really exciting to see this structure and dynamics that no one imagined." (MS)

Simulated three-dimensional structure of Earth's magnetic field, with inward (blue) and outward (yellow) directed field lines. Field lines extend two Earth radii from the core. The location of the core-mantle boundary is evident where the structure becomes complex.

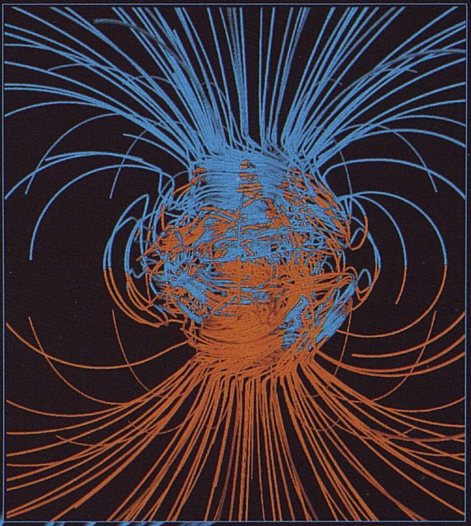

Reference:
Gary A. Glatzmaier & Paul H. Roberts, "A three-dimensional self-consistent computer simulation of a geomagnetic field reversal," *Nature* **377**, 203-209 (1995).

This research is supported by the Institute of Geophysics and Planetary Physics and the LDRD program at Los Alamos.

Illustration: The depiction of Earth with simulated magnetic field lines was created by PSC scientific visualization specialist Greg Foss.

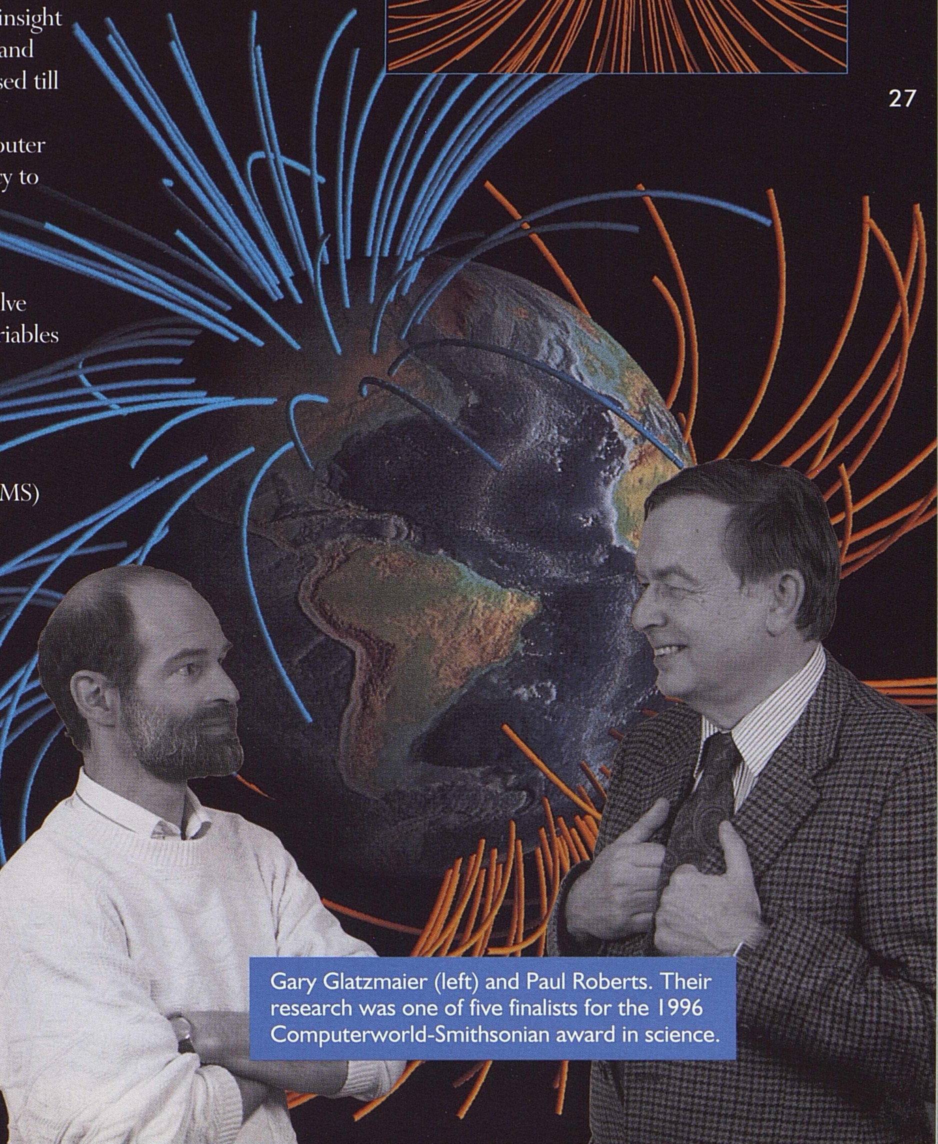

Gary Glatzmaier (left) and Paul Roberts. Their research was one of five finalists for the 1996 Computerworld-Smithsonian award in science.

Crystal at the Center of the Earth

A Seismic Adventure

There's a giant crystal buried deep within the Earth, at the very center, more than 3,000 miles down. It may sound like the latest fantasy adventure game or a new Indiana Jones movie, but it happens to be what scientists discovered in 1995 with a sophisticated computer model of Earth's inner core. This remarkable finding, which offers plausible solutions to some perplexing geophysical puzzles, is transforming what Earth scientists think about the most remote part of our planet.

"To understand what's deep in the Earth is a great challenge," says geophysicist Lars Stixrude. "Drill holes go down only 12 kilometers, about 0.2 percent of the Earth's radius. Most of the planet is totally inaccessible to direct observation." What scientists have pieced together comes primarily from seismic data. When shock waves from earthquakes ripple through the planet, they are detected by sensitive instruments at many locations on the surface. The record of these vibrations reveals variations in their path and speed to scientists who can then draw inferences about the planet's inner structure. This work has added much knowledge over the last ten years, including a puzzling observation: Seismic waves travel faster north-south than east-west, about four seconds faster pole-to-pole than through the equator.

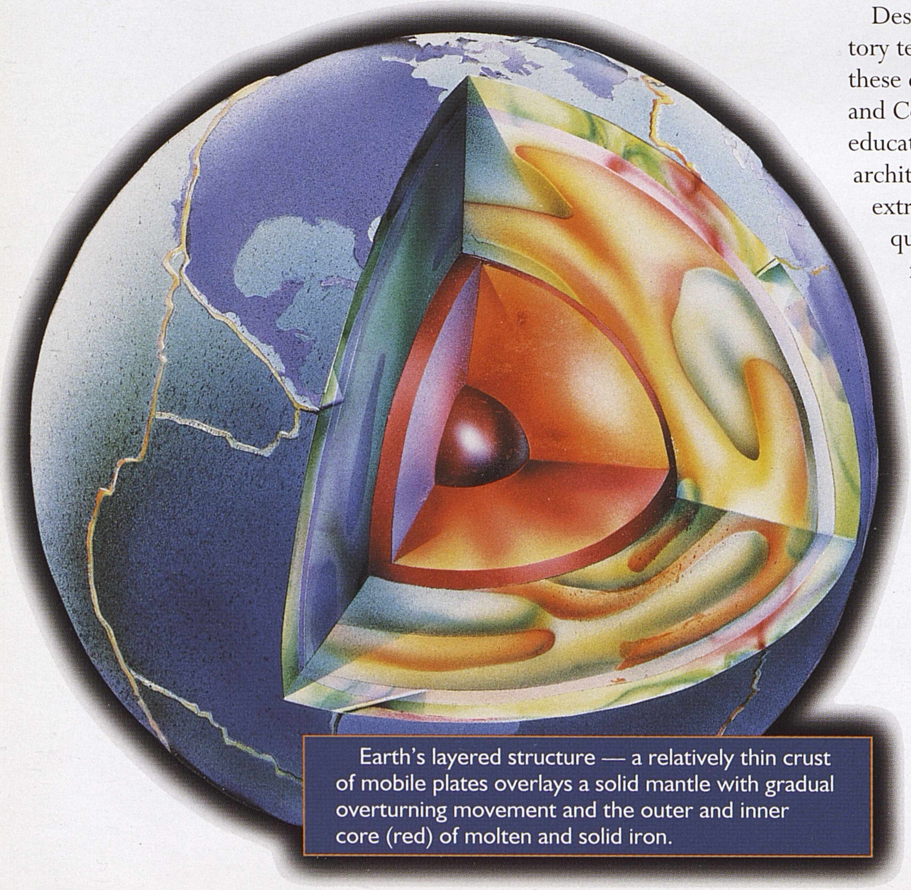

Earth's layered structure — a relatively thin crust of mobile plates overlays a solid mantle with gradual overturning movement and the outer and inner core (red) of molten and solid iron.

This finding, confirmed only within the past two years, quickly led to the conclusion that Earth's solid-iron inner core is "anisotropic" — it has a directional quality, a texture similar to the grain in wood, that allows sound waves to go faster when they travel in a certain direction. What, exactly, is the nature of this inner-core texture? To this question, the seismic data responds with sphinx-like silence. "The problem," says Ronald Cohen of the Carnegie Institution of Washington, "is then we're stymied. We know there's some kind of structure, the data tells us that, but we don't know what it is. If we knew the sound velocities in iron at the pressure and temperature of the inner core, we could get somewhere." To remedy this lack of information, Stixrude and Cohen turned to the CRAY C90.

Getting to the Core

Don't believe Jules Verne. The center of the Earth is not a nice place to visit, unless you like hanging out in a blast furnace. The outer core of the Earth, about two-thirds of the way to the center, is molten iron. Deeper yet, at the inner core, the pressure is so great — 3.5 million times surface pressure — that iron solidifies, even though the temperature is believed to exceed 11,000° F., hotter than the surface of the sun.

Despite rapid advances in high-pressure laboratory techniques, it's not yet possible to duplicate these conditions experimentally, and until Stixrude and Cohen's work, scientists could at best make educated guesses about iron's atom-to-atom architecture — its crystal structure — at the extremes that prevail in the inner core. Using a quantum-based approach called density-functional theory, Stixrude and Cohen set out to do better than an educated guess. With recent improvements in numerical techniques, density-functional theory had predicted iron's properties at low pressure with high accuracy, leading the researchers to believe that with supercomputing they could, in effect, reach 3,000 miles down into the inner core and pull out what they needed.

Could an iron ball 1,500 miles across be a single crystal? This REMARKABLE finding has stirred new thinking about Earth's inner core.

Ronald Cohen (left) and Lars Stixrude, Geophysical Laboratory, Carnegie Institution of Washington.

Rethinking Inner Earth

On Earth's surface, iron comes in three flavors, standard crystalline forms known to scientists as body-centered cubic (bcc), face-centered cubic (fcc) and hexagonal close-packed (hcp). Working with these three structures as their only input, Stixrude and Cohen carried out an extensive study — more than 200 separate calculations over two years — to determine iron's quantum-mechanical properties over a range of high pressures. "Without access to the C90," says Stixrude, "this work would have taken so long it wouldn't have been done."

Prevalent opinion before these calculations held that iron's crystal structure in the inner core was bcc. To the contrary, the calculations showed, bcc iron is unstable at high pressure and not likely to exist in the inner core. For the other two candidates, fcc and hcp, Stixrude and Cohen found that both can exist at high pressure and both would be directional (anisotropic) in how they transmit sound. Hcp iron, however, gives a better fit with the seismic data. All this was new information, but even more surprising was this: To fit the observed anisotropy, the grain-like texture of the inner core had to be much more pronounced than previously thought.

"Hexagonal crystals have a unique directionality," says Stixrude, "which must be aligned and oriented with Earth's spin axis for every crystal in the inner core." This led Stixrude and Cohen to try a computational experiment. If all the crystals must point in the same direction, why not one big crystal? The results, published in *Science,* offer the simplest, most convincing explanation yet put forward for the observed seismic data and have stirred new thinking about the inner core.

Three crystal structures of iron. Yellow lines show bonds between iron atoms.

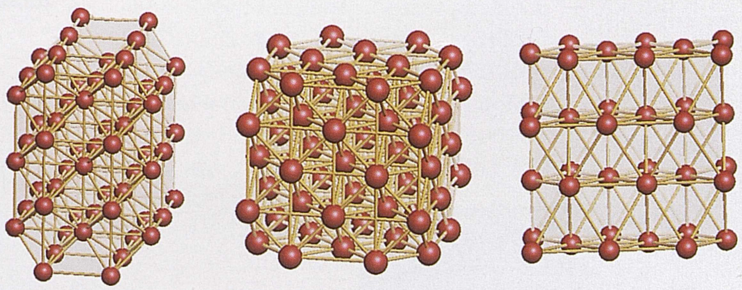

Could an iron ball 1,500 miles across be a single crystal? Unheard of until this work, the idea has prompted realization that the temperature-pressure extremes of the inner core offer ideal conditions for crystal growth. Several high-pressure laboratories have experiments planned to test these results. A strongly oriented inner core could also explain anomalies of Earth's magnetic field, such as tilted field lines near the equator. "To do these esoteric quantum calculations," says Stixrude, "solutions which you can get only with a supercomputer, and get results you can compare directly with messy observations of nature and help explain them — this has been very exciting." (MS)

Reference:

Lars Stixrude and R.E. Cohen, "High-Pressure Elasticity of Iron and Anisotropy of Earth's Inner Core," *Science* **267**, 1972-75 (1995).

Lars Stixrude and R.E. Cohen, "Constraints on the crystalline structure of the inner core: Mechanical instability of BCC iron at high pressure," *Geophysical Research Letters* **22**, 125-28 (1995).

This research is supported by the National Science Foundation and the Alexander von Humboldt Foundation.

Earth illustration: Keelin Murphy, educational outreach program, Center for High Pressure Research.

New Light on Dark Matter

Testing Cosmic Theory

Not many years ago, cosmology was more theology than hard science, says Princeton astrophysicist Jeremiah P. Ostriker. "Supercomputers have helped change that," he adds. Indeed, supercomputing is proving itself one of the most useful means of getting a handle on the nature and structure of the universe.

Ostriker, along with scientists Guohong Xu and Renyue Cen, is using PSC's CRAY T3D to test competing theoretical models of cosmological origin, each of which offers a blueprint for the embryonic universe. One of these, the so-called standard cold dark matter model, had until recently been one of the most widely held theories. While calculations of the other three theories are still being analyzed, the Ostriker team's work has all but pulled the plug on the ailing standard cold dark matter model.

This project is one of the largest supercomputing efforts ever undertaken and occurs under the auspices of the Grand Challenge Cosmology Consortium (GC^3), a collaborative effort among physicists and computer scientists at six universities and research centers. Begun in 1993 with funding from NSF and led by Ostriker, GC^3 aims to exploit scalable parallel systems like the CRAY T3D to investigate the universe.

"What we've done that's new," says Ostriker, "is take some of these speculative scenarios, put them in a computer, and see what they predict in terms of gravitational lensing. Nobody had ever done that before because it's an incredibly hard computational problem."

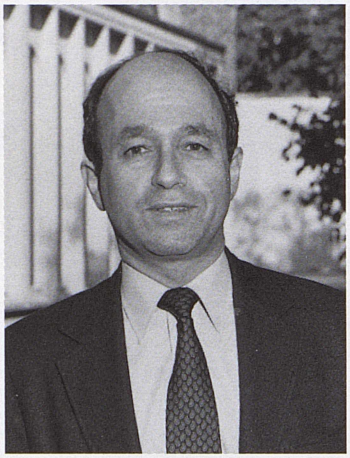

Jeremiah P. Ostriker, Princeton University.

PSC scientific visualization specialist Joel Welling led a GC^3 collaborative effort that tapped the resources of three NSF centers to simulate the Andromeda and Milky Way galaxies colliding, an event predicted to occur some five billion years from now. Presented in a virtual reality medium, the model thrusts observers into the thick of the intergalactic squall. This project was recognized as the Best Integration of Heterogeneous Applications at Supercomputing '95, the annual conference that showcases high-performance computing and communication.

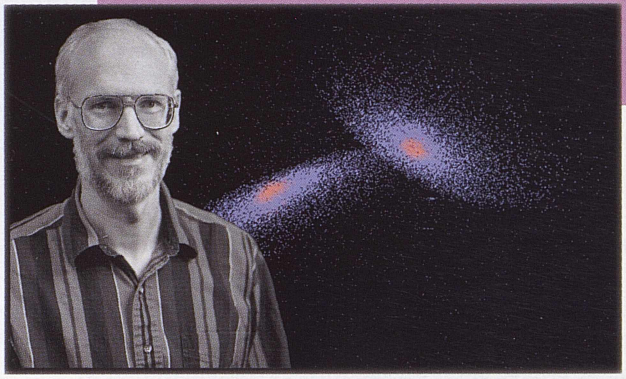

The Original Space-Based Telescope

Gravitational lenses are caused by massive, compact clusters of galactic matter. Much as an optical lens bends light to form an image, the enormous gravitational field of a cluster magnifies and distorts light coming from objects lying far beyond it, such as quasars and galaxies, sometimes producing multiple images that flank the lens itself. It's a valuable, naturally occurring observational tool whose imaging potential also plays a key role in the Princeton simulations.

That lenses exist at all is due in part to dark matter, the invisible glue of the universe. To maintain the structure of the universe, to hold the galaxies together as galaxies, for instance, requires much more gravity than observable matter can generate. Calculations indicate, in fact, that dark matter, the unobservable part of the universe, accounts for as much as 90 percent of all matter. Whether dark matter consists of as yet undetected galactic particles, massive but hidden stellar objects, or some combination of the two, remains undetermined.

But, says Xu, "It doesn't matter what dark matter is made of. The lensing phenomenon detects the gravitational effect of the dark matter. And that's what counts when you're modeling a piece of the universe."

They have all but pulled the plug on the *ailing* standard cold dark matter model.

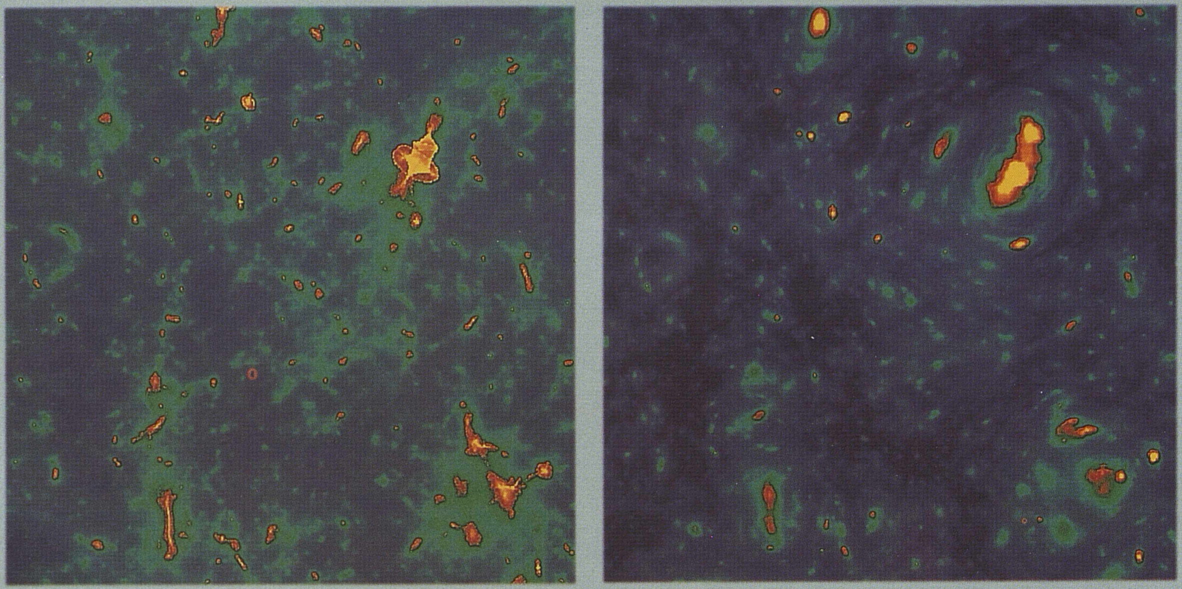

The Cold Dark Matter Model

These two images represent a 3D universe simulated from standard cold dark matter parameters. Density varies from blue (low) through green, red and yellow (high). To gauge the probability of gravitational lensing, the model generates magnification maps (above left). The denser the region, the greater the magnification power. Vibrant yellow indicates the most likely locations that otherwise undetectable objects, such as quasars and galaxies, would be significantly magnified and imaged multiple times.

The second image shows gravitational lensing at work. The large yellow clump with red and green borders (upper right corner) depicts the lens. The two small elongated objects closely flanking it represent lensed images. This model predicted many more gravitational lenses than are actually observed in a given volume of space, in effect ruling out the standard cold dark matter model.

Nailing the Coffin of Cold Dark Matter

A model spreads matter around much the way it currently exists in the cosmos – thick lumpy patches of galaxies and galaxy clusters separated by vast comparatively empty stretches of cosmological range. Density distribution is key, and the number of gravitational lenses occurring in a particular patch of space becomes the test of that distribution. An accurate model will exhibit a distribution of lenses that coincides with observed data. "Different theories predict more or fewer gravitational lenses," says Ostriker, "and that's what we're computing for."

The standard cold dark matter theory, for instance, which calculates just enough mass to prevent eventual collapse or eternal expansion of the universe, produced many more lensed objects of greater intensity than is observed in actuality. Another model, which says the universe will eternally expand because not enough matter exists to prevent it, produced results more in line with today's universe. "The standard cold dark matter model is wrong," says Ostriker. "So the question is, are any of its variants plausible? For now, supercomputers such as the T3D are the primary tools that can help answer such questions."

The Universe and the T3D

Because of limited computing capability, a prior modeling effort by the Ostriker group melded low resolution universe models with high resolution models, producing a portion of the universe that was quite small in cosmological terms (five megaparsecs). Though it offered a good approximation of the interactions between objects near one another, it didn't deal well with remote but interacting galaxies, clusters and other stellar bodies. Such interactions play a key role in the physics of the evolving universe.

The T3D eliminated the need to meld data, boosting the integrity of the results. It generated models 20 times larger (100 megaparsecs) without sacrificing the resolution necessary for small-scale accuracy, and accounted for interactions between remote but gravitationally related objects. "We still get the same resolution achieved with the much smaller model," says Xu. The T3D's memory and data capacities, adds Xu, played a critical role in the project's success. Each of the four simulations generated about 50 gigabytes of data, averaged 40,000 processing hours, and ran mostly using 256 processors, with minimum and maximum runs using 128 and 512. "We ran the biggest simulation of this type ever attempted," says Ostriker. "It could not have been done without the T3D." (JCW)

Reference:
Joachim Wambsganss, Renyue Cen, Jeremiah P. Ostriker and Edwin L. Turner, "Testing Cosmogonic Models with Gravitational Lensing," *Science* **268**, 274-76 (1995).

This research is supported by NASA and the National Science Foundation.

Evolution and Structure of the Universe

Death of a Hot Young Star

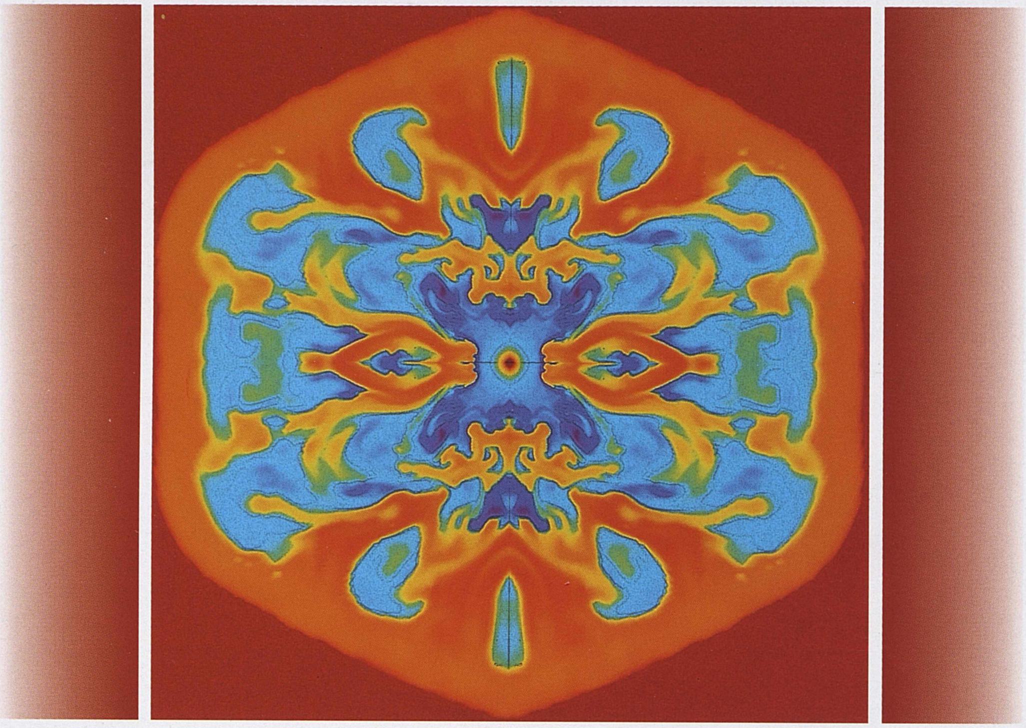

This image shows entropy (red increasing to violet) in an exploding 15 solar mass supernova about 70 milliseconds after explosion. The scale is 1,500 x 1,500 km.

Supernovas, The Real Story

Imagine a flashbulb the size of a planet going off with the brightness of a hundred billion stars. Hard to imagine, but almost everything about a supernova is hard to imagine. A star ten times more massive than the sun, so large and radiant it is destined for a short life — only 10 million years or so, then burnout, with an end stunning in its suddenness, one of the most cataclysmic events in nature.

When this massive star's core is withered to the size of the Earth, but a billion times more dense, the critical instant arrives when the nearly exhausted nuclear furnace no longer radiates enough energy to overcome the fierce inward pull of its own gravity. Click — in less than a second a spherical mass thousands of miles across collapses to the size of a city. Then core bounce, the ultra-dense nuclear material rebounds, sending a shock wave smashing outward through the inrushing outer layers of star material. Bang. The explosion releases as much energy as the sun radiates in 10 billion years, with a flare of light momentarily as bright as an entire galaxy.

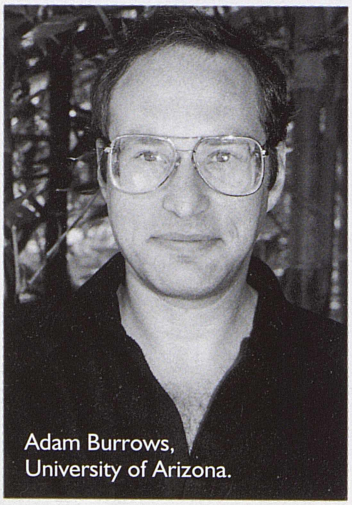

Adam Burrows, University of Arizona.

These spectacular death throes are at the same time the pangs of cosmic rebirth, imparting energy and matter to interstellar space, crucial ingredients for the structure of galaxies, the next generation of stars and the raw material of life itself. The oxygen atoms we breathe, iron in our hemoglobin, calcium in our bones and probably the fluorine in our toothpaste were created by nuclear fusion in these massive stars and thrown into space when they exploded. "It is not an exaggeration," says physicist Adam Burrows, "to say that supernovae affect almost everything astronomical."

Burrows, who chairs the Theoretical Astrophysics Program at the University of Arizona's Steward Observatory, has worked on supernovas for 15 years. With colleagues Bruce Fryxell and John Hayes, he is using the CRAY C90 to model them with greater realism than previously possible, contributing to new understanding of how subatomic particles called neutrinos play a crucial role in driving the supernova shock wave. Researchers in the field have hailed these results as a clear breakthrough in our knowledge of how supernovas explode.

Researchers have hailed these results as a clear breakthrough in understanding how supernovas explode.

Solving the Supernova Problem

For years, the problem in modeling supernova explosions has been the shock wave. Rather than blasting out through the progressively lighter and lighter shells of the dying star, as it does in nature, the shock wave in computer models stalled. It lacked the oomph to blow the lid off.

Researchers believed the missing ingredient was neutrinos. Theoretical studies showed that the high temperature and density of the collapsed stellar core should generate an outward surge of neutrinos behind the shock wave, and that this energy should be enough to explode the star. Until recently, however, the computer models have not cooperated.

Burrows' work demonstrates that the key may be modeling in two or three dimensions. Most prior supernova modeling assumed that the explosion is spherically symmetric, blasting outward the same in all directions. This assumption, in most respects entirely reasonable, made it feasible to do the calculations, which otherwise used too much memory and computing time. But even with neutrinos included, these one-dimensional models were duds; stalled shock, no explosion.

Along with others in the field, Burrows speculated that the problem was not the idea of neutrino heating, but modeling in only one dimension. The neutrino surge, he believed, would introduce convective currents, like in a pot of boiling water, asymmetric bulges and bubbles of energy behind the shock wave that could only be captured in two or three dimensions. Using first the CRAY Y-MP and now the C90, Burrows has shown that indeed this appears to be what happens. As he put it in *The Astrophysical Journal* (September 1995), his two-dimensional model of a 15 solar mass supernova "exploded magnificently."

"The convection and overturn," says Burrows, "actually changes the structure behind the shock." In effect, he explains, it pushes the wave farther out, so that when in one dimension it would stall, in two dimensions it confronts a thinner layer of infalling material. With less pressure holding it back, it is more prone to erupt, and the difference turns out to be critical.

"To do this problem correctly," says Burrows, "you need to do it multi-dimensionally. Only in the last five or so years, with the C90 and similar machines, have we had the computing resources to do this." One of the next steps for Burrows' team is to add a third dimension, which will help to address other crucial supernova questions. They are busy adapting their model to the CRAY T3D and T3E, which Burrows expects will give a tenfold speedup. "That will allow us to do a whole sequence of calculations, explore a range of progenitor stars with different structures. We'll also be able to include more realistic neutrino transfer, which is computationally intensive. We've peeked under the curtain and discovered a new world, but we're still crawling into the sunshine." (MS)

Reference:

Adam Burrows, John Hayes & Bruce A. Fryxell, "On the Nature of Core-Collapse Supernova Explosions," *Astrophysical Journal* **450**, 830-50 (1995).

This research is supported by the National Science Foundation and by NASA.

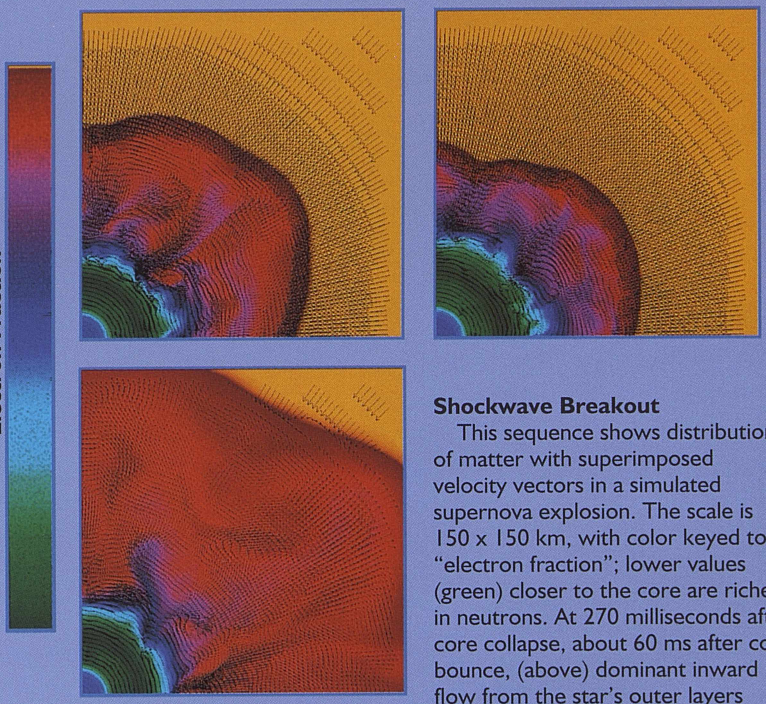

Shockwave Breakout

This sequence shows distribution of matter with superimposed velocity vectors in a simulated supernova explosion. The scale is 150 x 150 km, with color keyed to "electron fraction"; lower values (green) closer to the core are richer in neutrons. At 270 milliseconds after core collapse, about 60 ms after core bounce, (above) dominant inward flow from the star's outer layers holds back the shock wave moving outward from the core. At 307 ms (above left), the shock has moved out slightly but remains essentially stalled. Vigorous and varied convective motions are evident in the region behind the shock. Four ms later (lower left), some material is still falling in, but the star has exploded.

Mystery of the Wrong Orbit

Part 1: A Problem Rears its Ugly Head

To solve some problems, it really does take a rocket scientist. And it doesn't hurt if a computing system like the CRAY T3D is available. Case in point: August 1995. The Delta II rocket puts a commercial satellite into the wrong orbit.

Designed and built to launch satellites for the U.S. Air Force, the Delta II has been a highly successful project. What happened? A committee of engineers from McDonnell Douglas and The Aerospace Corporation, the Air Force contractors on the Delta II, investigate. They find that a mechanical component malfunctioned, causing a lower than intended orbit. What caused the malfunction? Analysis leads them to suspect overheating due to "backflow" from the rocket engines.

Software and hardware ingenuity pushed the envelope of high-performance computing and directly contributed to the success of a high-priority national mission.

How can the engineers confirm this analysis, and in a hurry? No one wants to trust an educated guess, no matter how sophisticated, when millions of dollars are on the line for upcoming satellite launches already scheduled. Wind tunnel testing? This tried-and-true design tool is almost useless for testing rocket performance with the "plume on" — i.e., with the rocket engines firing, and the plume contributes dramatically to the aerodynamics of a rocket in supersonic flight. Furthermore, the Delta II is a clustered or "multi-body" rocket — a core rocket surrounded by nine boosters. This by itself creates problems for wind-tunnel testing, which even in simple cases is costly and time-consuming. Is there a way out of this bind?

Part 2: The Caltech Group

In the late 1980s, Johnson Wang, a senior engineer at The Aerospace Corporation, developed a numerical scheme for simulating the aerodynamics of complex, multi-body rockets like the Delta II with high accuracy. In technical terms, his software is an efficient "flow solver" for the three-dimensional Navier-Stokes equations, the classic mathematical formulation of fluid flow.

To exploit the potential for improved performance on scalable, parallel systems, which team tens, hundreds or even thousands of processors to work simultaneously on a single computing task, Wang in 1991 joined forces with Stephen Taylor, head of the Scalable Concurrent Programming Laboratory at Caltech. Taylor works with a group of scientists who have pioneered basic programming technology to address important, large-scale problems in industrial and defense applications. Their work feeds directly into development of advanced software and the next generation of parallel machines.

Over the past four years, Taylor and Wang worked to implement their flow solver on a range of parallel computers, including work-station networks and shared memory parallel systems. Along with improved performance, the parallel implementation includes features that make it especially useful for rocket design, including the ability to simulate multi-body rockets.

In early 1995, Taylor and Wang put their software through its paces by simulating the Titan IV, another multi-body rocket. These runs proved that the parallel flow solver can provide the kind of data the Delta II team needs. "We can calculate the forces acting on the vehicle," says Taylor. "We can calculate aerodynamic drag and determine the strength of the vehicle's sonic boom. And we can calculate temperature contours that predict heat transfer, which allows engineers to design appropriate shielding and paint."

Stephen Taylor, an instrument-rated pilot, is director of Caltech's Scalable Concurrent Programming Laboratory. For their work developing parallel programming technology in microelectronics, satellite manufacturing and launch vehicle design, SCPL was a finalist in the science category of the 1996 Computerworld Smithsonian Awards.

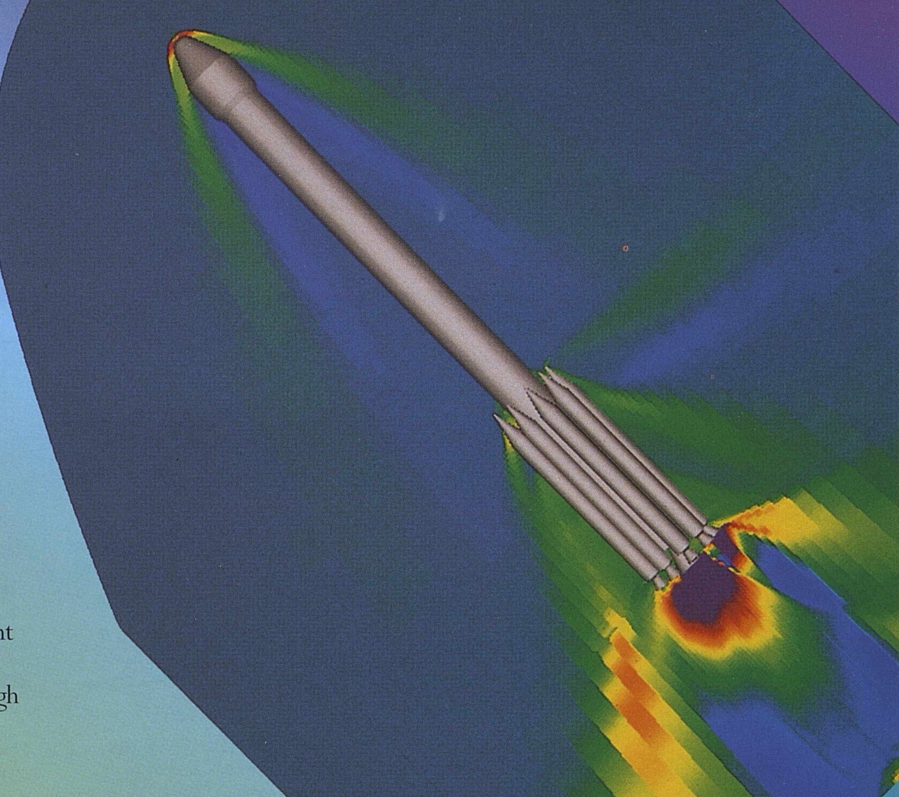

Simulations on the CRAY T3D provided pressure contours around the Delta II with rocket engines firing.

Part 3: A Scalable, Parallel Solution

Simulations like those needed for the Delta II require an enormous amount of computing — months under normal circumstances, but months aren't available. Is there a way to get results more quickly? In work during 1995, Taylor and his student Jerrell Watts and Alan Stagg of Cray Research implemented the flow solver on the CRAY T3D at Pittsburgh Supercomputing Center. Due to the high resolution required for the Delta II computations, the pressing need is memory. The T3D includes 32 billion bytes, which along with its excellent parallel performance should make it possible to get results in a relatively short time span.

PSC responds to the situation by making their T3D available for "dedicated" runs. All 512 T3D processors become a powerful team, working together to compute the complete "plume on" flow field of the Delta II. In a matter of days, the researchers have the results they need. The computations quantify all relevant parameters — velocity, density, pressure — in three dimensions.

The computed results match closely with a pressure gauge reading from the backflow region of the booster rockets during flight, confirming reliability of the simulations. Using the computed flow-field data, the engineers determine that their analysis is essentially correct: Interactions among the booster plumes create a backflow that causes the component to overheat. Adjustments are made, and in November 1995 the Delta II launches a satellite into correct orbit.

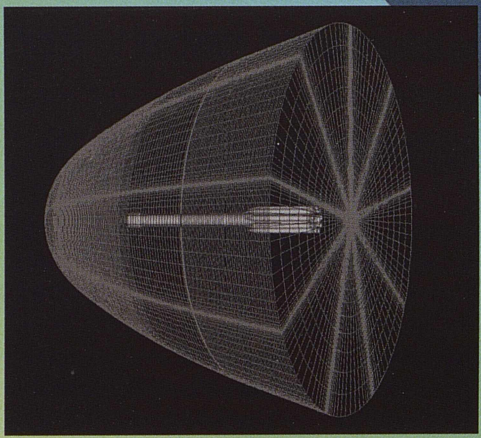

The Delta II simulations used a computational grid composed of about 4.5 million grid points.

Close inspection of the pressure contours revealed that plume interactions create a backflow region between the main engine and a booster.

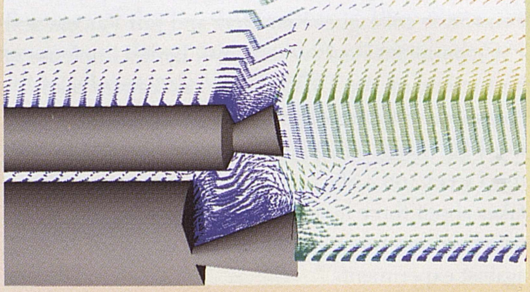

Epilogue

The director of the Delta II program commends the team of researchers and PSC — a combination of software and hardware ingenuity that has pushed the envelope of high-performance computing and directly contributed to the success of a high-priority national mission. "The T3D saved us," says Taylor. "It reduced the turnaround to two weeks, start to finish. Without the T3D, we wouldn't have been able to have an impact on this project." (MS)

Reference:
Taylor and Wang, "Launch Vehicle Simulations using a Concurrent, Implicit Navier-Stokes Solver," *Journal of Spacecraft and Rockets* (forthcoming).

This research is supported by the Advanced Research Projects Agency.

Quantum Chemistry

Wobbly Atoms

NMR and Quadrupolar Nuclei

Since its development in the 1940s, nuclear magnetic resonance (NMR) spectroscopy has become one of the most valuable research tools available to chemists. "It's the single most powerful experimental technique for characterizing the chemistry of any kind of sample," says Alan Benesi, "solid, liquid or gaseous."

In the 1980s, researchers added scanning technology that helped take NMR into hospitals, where as magnetic resonance imaging (MRI) it rapidly proved its worth as a diagnostic tool. With less fanfare, the parent technology, NMR spectroscopy, also expanded its usefulness. It is widely used in industry and since the mid-1980s has become important in biomedical research, where it supplements X-ray crystallography as a method for determining the structure of proteins.

Another NMR revolution is now underway, says Benesi, who directs the NMR facility at Pennsylvania State University. New methods will make it feasible to use NMR spectroscopy for a large class of atoms that are in theory observable, but in practice present difficult problems. These atoms, which happen to comprise most of the periodic table, have "quadrupolar" nuclei — a nonspherical distribution of positive charge in the nucleus — which make the NMR spectrum they produce much more complicated than it otherwise would be. "More than 65 percent of the NMR-observable nuclei in the periodic table are quadrupolar," says Benesi, "and most of the time these nuclei are found in solids such as metals, ceramics and minerals."

Benesi's research involves devising mathematics, numerical methods and computer models to simulate the quantum behavior that underlies NMR experiments, and he focuses on quadrupolar nuclei. "I'm trying to understand what's going on with these nuclei, which will help us find ways, new tricks, to get better information from the experiments." Using the CRAY T3D to great advantage, with unprecedented performance of 39 billion computations a second (Gflops), his simulations have revealed surprising and promising new information.

Ravi Subramanya,
Pittsburgh Supercomputing Center.

Bridging the Gap

The ability to get useful information from NMR spectroscopy goes hand-in-hand with computer simulations that link the experiments with theory. "There's an intimate tie between theory and experiment," says Benesi. "For the more unusual experiments especially, you choose a set of parameters and then, basically, you solve the quantum mechanics to predict what you should observe and compare it with what you see. Then you say 'Oops, that's off.' You go back and change the parameters and say 'Oh, that's better,' back and forth between theory and experiment, until you find the best match."

For quadrupolar nuclei, the mathematics of the theory is complicated enough that the experimental ramifications aren't well understood, and the experiments themselves reveal unusual behavior. A useful technique known as "spinlock," for instance, involves applying a long radio-frequency pulse to the sample to "lock" it at its resonance frequency. For reasons that are not fully understood, explains Benesi, non-quadrupolar nuclei can be spinlocked for much longer than quadrupolar nuclei, where the signal dies away to almost nothing.

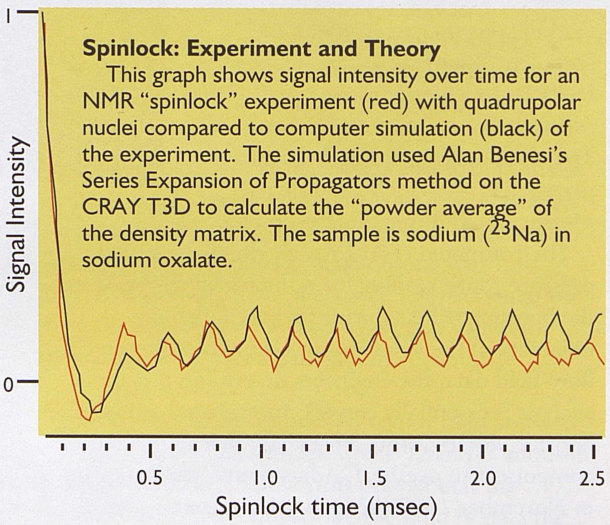

Spinlock: Experiment and Theory
This graph shows signal intensity over time for an NMR "spinlock" experiment (red) with quadrupolar nuclei compared to computer simulation (black) of the experiment. The simulation used Alan Benesi's Series Expansion of Propagators method on the CRAY T3D to calculate the "powder average" of the density matrix. The sample is sodium (^{23}Na) in sodium oxalate.

With his scientific curiosity piqued, in 1993 Benesi devised his own numerical recipe for simulating NMR. Called Series Expansion of Propagators, the method is especially well suited to the quantum traits of quadrupolar nuclei. During 1995, with help from Penn State colleagues Ken Merz and Jim Vincent and PSC consultant Ravi Subramanya, Benesi implemented his method on the T3D. "These calculations take enormous amounts of computing," says Benesi, "and the T3D makes a huge difference. What used to take me a week (on a Silicon Graphics workstation) now takes about three hours. With this turnaround, the data becomes much more useful to compare against experiments."

"If we can devise a way to keep observable signals from dying, that would be *extremely* important in solid-state NMR."

The most time-consuming part of Benesi's computations is what's called the "powder average," so called because samples used in solid-state NMR are almost always in powder form. Simulating the NMR signal from such a powder is, basically, a summing up of many individual simulations, each of which corresponds to one among all the possible directions that the huge quantity of individual crystals in the sample could be oriented. The inherent parallelism of this computation is one of the keys to the outstanding performance (76.5 Mflops per processor scales linearly to 39.2 Gflops on 512 processors) Benesi achieves on the T3D.

New Finding: Decay of the Density Matrix

Benesi's ability to compute the powder average led him to a fascinating discovery. What he calculates is a quantum mathematical quantity called the "reduced density matrix," and what he found — which no one had seen before — was that in simulations of spinlock with quadrupolar nuclei the reduced powder average density matrix decays sharply to almost zero. Sound familiar? The obvious parallel between density matrix decay in simulations and signal loss in actual experiments suggests a connection between theory and experiment that no one had surmised until Benesi's work.

These results are still being absorbed by researchers at other NMR facilities, and the implications have yet to be sorted out. They come as a surprise because no one anticipated that a powder average density matrix would do what Benesi's computations show. "For an individual crystal," explains Benesi, "the density matrix stays constant, no matter what you do to it. Only when you add the matrices together do you get this behavior. In hindsight, it should have been obvious. There's destructive interference among the different frequencies for the differently oriented crystals."

In recent computations, Benesi arrived at another interesting result: With a second radio-frequency pulse, he can revive the near-death density matrix. "I can make it echo back. I can give a pulse, watch the density matrix decay away, then give a refocussing pulse and make the matrix climb back up, in the same amount of time it took to decay. If we can take advantage of that — devise a way to keep observable signals from dying, that would be extremely important in solid-state NMR." (MS)

Reference:
Alan J. Benesi, "Series Expansion of Propagators," *Journal of Magnetic Resonance, Series A* **103**, 230-33 (1993).

Peaks and Valleys: The Inside Story on NMR

A complex quantum phenomenon, NMR works because the nuclei of many atoms — those with an odd number of protons or neutrons — spin like tops, creating a tiny magnetic field. When you put these atoms inside another magnetic field, the spinning nuclei line up in the same direction, as opposed to their normal random orientations, and they wobble or "precess," like a top slowing down. Furthermore — and here's the key — each different kind of atom, carbon versus hydrogen, for instance, has its own distinctive wobble, its own resonance frequency, which makes it possible for skilled researchers like Alan Benesi to read the peaks and valleys of the signals they give off.

"With a solid," explains Benesi, "we put the sample into a little rotor, and we put that down into the sweet spot, the highest magnetic field part of the superconducting magnet, and we spin the rotor and excite it with radio frequency pulses." After the sample is zapped, sometimes in complicated ways that reveal subtle frequency transitions, it gives off signals that can be monitored for the tell-tale clues of a particular atom, including differences that occur according to an atom's surrounding chemical environment. "In a carbon sample, for example," says Benesi, "we can tell the type of carbon, whether it's in a methyl (CH_3-) or carboxyl (-COOH) group, from its frequency. And that's very useful chemical information."

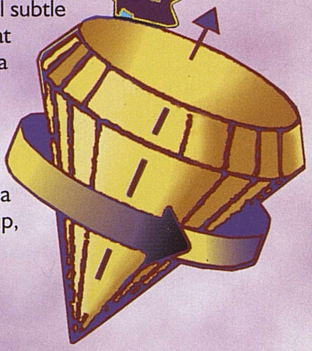

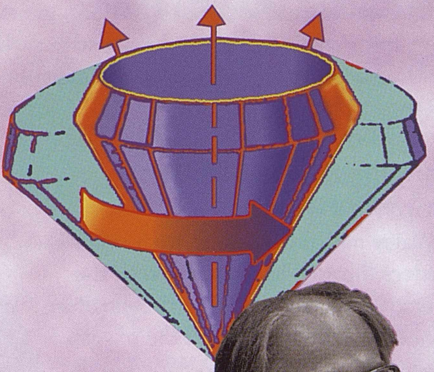

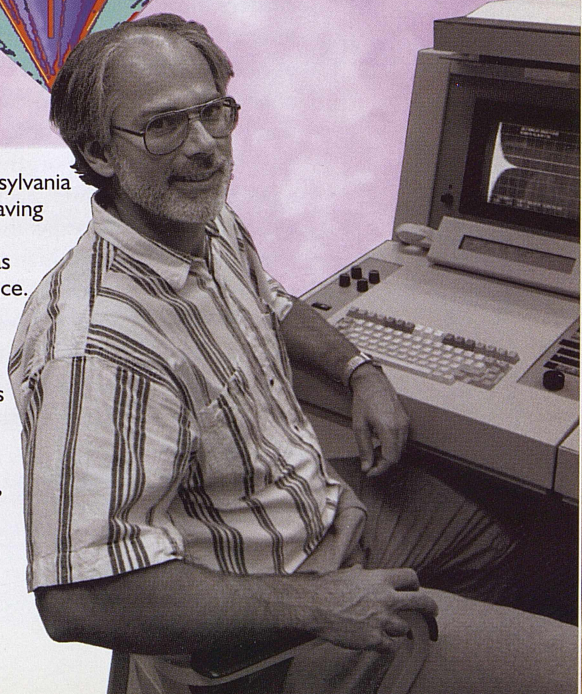

Alan Benesi, Pennsylvania State University. "Having access to the supercomputing center has made all the difference. This is how science should happen. You make resources available so scientists can explore and see what they come up with. That's how advances are made."

Wet Electrons

Electron Catwalks

While electrons are known as the glue that holds matter intact, they also have the ability to flit from bond to bond, molecule to molecule, sparking reactions that in the body produce the energy used for metabolism. This journey, which scientists call electron transport, can be long and winding.

In a protein, for instance, an electron must travel through the molecule's complex chain of atoms, work its way to the surface and onto an adjacent protein. Imagine walking from one end of a street to the other going through all the rooms of each house. A catwalk spanning and connecting the roofs would simplify the stroll. But no such electron catwalk exists in biological systems. Or so scientists thought.

MIT physicist John D. Joannopoulos and research scientist Kyeongjae Cho along with graduate student Ickjin Park have identified an electron pathway, an alternative to the long and winding route of bond-to-bond electron transport, that may be likened to this catwalk. Using PSC's CRAY C90, they simulated what happens when an extra electron slips into a cluster of six water molecules, creating a molecular interaction called a wet electron. What they have learned from this curious entity has implications for electron transport in biological processes in general.

"How different biological molecules interact, store and transfer energy and react with each other," says Joannopoulos, "to a large degree involves electron transport."

Dangling Hydrogens & an Uninvited Guest

Water molecules link to each other in clusters through hydrogen bonds, the interactions between the hydrogen atom of one water molecule and the oxygen of another. These are the same bonds that bind DNA and are ubiquitous in the chemistry of living organisms. In a pure solution of water, each atom will try to satisfy its respective attractions, forming as many bonds as possible in a collective molecular effort to reach a state of lowest possible energy. Like wedding guests searching for dance partners, however, invariably some hydrogen atoms are left dangling.

The researchers simulated what happens when an uninvited guest – an extra electron – shows up at this microscopic reception. Each of the dangling hydrogen atoms – those without dance partners – vies for the companionship of the negatively charged electron, causing the water molecules to form a cage around it, hence the moniker "wet." Although wet electrons have been known to exist since the early 1990s, their atomic interactions have been little understood.

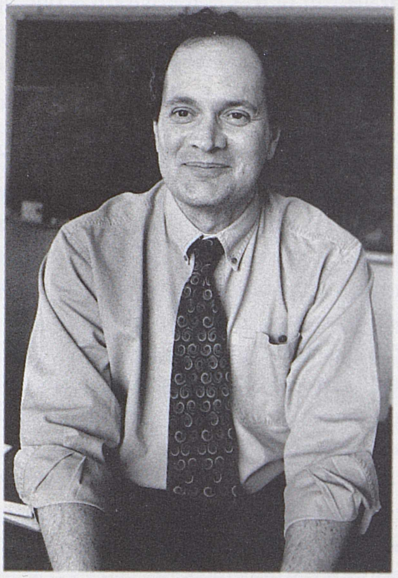

John D. Joannopoulos,
Massachusetts Institute of Technology.

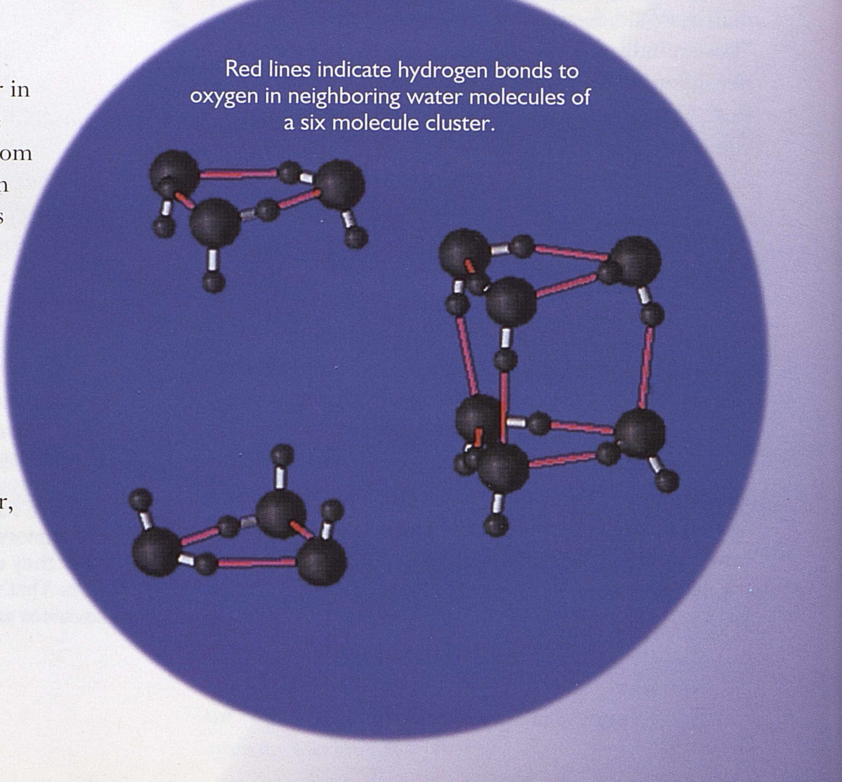

Red lines indicate hydrogen bonds to oxygen in neighboring water molecules of a six molecule cluster.

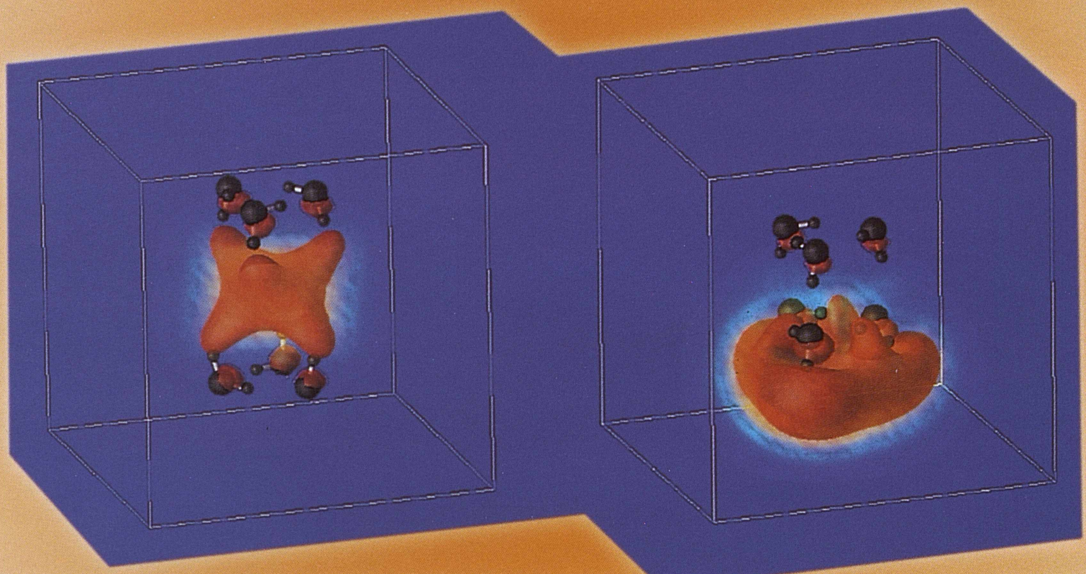

The Wet Electron: Two Possible Structures

Two possible structures occur when an electron is slipped into a cluster of six water molecules. In both structures, hydrogen atoms (small black balls) vie for the attention of the extra electron. The orange cloud represents the potential location of the extra electron. Smaller red clouds depict orbitals between hydrogen and oxygen atoms.

In one structure (left), all six water molecules — three on top, three on the bottom — have dangling hydrogens, non-bonded to oxygen atoms of neighboring water molecules. In the other, all the hydrogen atoms of the three top water molecules are bonded (bonds not shown), and the three lower water molecules, which have three dangling hydrogens (pointing downward), draw the extra electron into their vicinity.

Results: Inside and Outside the Cluster

The computations simulated a process of randomly distorting a wet electron system of six water molecules and allowing it to resettle into its most stable, low energy arrangement. Results showed two possible scenarios. "We found that the system likes to arrange itself in two different ways, both of which bind the electron," says Joannopoulos. "One keeps the electron stuck inside the cluster, and one keeps the electron stuck outside. The typical electron pathway involves moving between atoms, through bonds, but this suggests that the electron could hop through space from one atom to the next. In other words, it doesn't need a bond to relocate."

"If there were a lot of these dangling hydrogens in a line," adds Joannopoulos, "then this extra electron has a means of transport, and that's a completely new idea." With proteins for instance, strategically arranged dangling hydrogens could create a path for an electron to move through space, the equivalent of a hiker crossing a creek by walking over a bridge, as opposed to trudging through the water over slippery rocks. This dangling hydrogen mode of electron transport could be useful in genetic engineering of proteins and in drug design. "It's a possible new pathway and because of that, maybe someday we can engineer it to drive electrons in certain directions along proteins."

They have identified an alternative to the long and winding route of bond-to-bond electron transfer.

To further map this potential new pathway, Joannopoulos plans to expand his wet electron simulations to model what occurs when an extra electron meets up with hundreds of water molecules. This work will require the increased computing capability of a scalable parallel system such as the CRAY T3D. "We'll be studying the dynamics of the system," says Joannopoulos, "and watching how it evolves over time at different temperatures, so the calculations will be much more time and memory intensive. With its huge memory capacity and high speed, the T3D makes these big projects manageable." (JCW)

Reference:
K.S. Kim, I. Park, S. Lee, K. Cho, J.Y. Lee, J. Kim, and J.D. Joannopoulos, "The Nature of a Wet Electron," *Physical Review Letters* **76**, 956-59 (1996).

This research is supported by the Office of Naval Research and United States Joint Services Electronics Program (sponsored by the Army Research Office, Office of Naval Research and Air Force Office of Scientific Research).

Big Mixups

Tackling the Scale Problem

Perfume spreading through a crowded room, an oil spill mixing with ocean currents, Arctic air bringing winter south in the jet stream across the Great Lakes — turbulent diffusion is ubiquitous. At small and large scales, from a teaspoon of cream in hot coffee to the global atmospheric patterns that affect climate change, the effects are the same — random fluctuations, swirls and eddies that accelerate the mixing of one substance with another.

A decades-old question is whether the mathematical laws believed to govern turbulent diffusion hold true. The main problem in testing them has been the immense range of scales involved in some of nature's turbulent systems. The mixing and flow of salinity in the oceans, for example, includes tiny eddies in the tidal pools of a remote cove and currents like the Gulf Stream that span continents. These large systems can't be modeled in a laboratory, and computational models have been limited by algorithms that can deal only with a comparatively small range of scales, on the order of from 1 to 1,000. To test the universal laws over a turbulent diffusion expanse that occurs in nature requires handling a scale range of 1 to 100,000, with the number of variables — and the amount of computation — expanding exponentially with each increase.

Using a creative mathematical approach, New York University researchers Andrew Majda and Frank Elliott have carried out computer modeling that represents a major advance in the ability to study turbulent diffusion across the vast range of scales that comprise such systems in nature. Tapping the parallel processing power of PSC's CRAY T3D and measuring the results against meteorological data, they've shown that the universal laws hold true, and their landmark work offers potential for accurately modeling extremely large, complicated flow systems. "Turbulence is a very difficult problem," says Majda, "but it's a key piece of the atmospheric puzzle. Our model helps put some pieces in place for the first time in idealized circumstances."

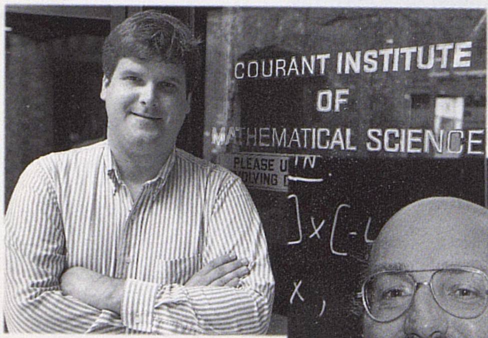

Frank W. Elliott, Jr.,
New York University.

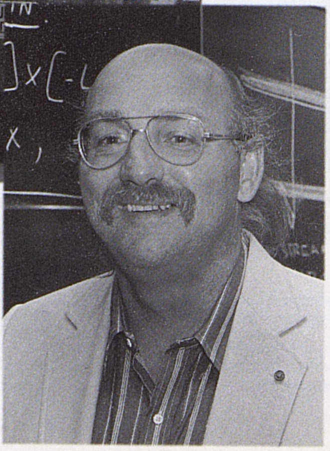

Andrew J. Majda,
New York University.

Sketching to Scale

Turbulent diffusion is driven by an immense collection of interacting, similarly shaped whorls or eddies, ranging in size from millimeters to miles, whose collective effect is measured as a velocity field. The Majda-Elliott approach charts the effects of turbulence on two identical particles released into a modeled velocity field. While the starting point is nearly the same for both, diffusion quickly places the particles on divergent paths, for example sending one east and the other west, as they swirl from one eddy to the next. The system is fractal-like — at every scale, the eddies resemble each other. "Ordinarily it's a very hard thing to model," says Majda. "The idea is to choose an approach that is computationally fast and accurate in that it respects the fractal structure."

The problem with existing models is that the algorithms employed to follow the movements of two particles while simultaneously assessing turbulence data around them can only track the particles for a relatively short time, a few meters as opposed to the thousands of kilometers needed to account for large flow systems. The information loads grow exponentially with each jump to a larger scale, because the calculations try to assess turbulence in all the eddies of each scale. (Each scale represents all the eddies of a particular size.)

Majda and Elliott take a sketcher's approach to cut the data load. Their model scrutinizes a finite amount of information at each scale (several hundred variables as opposed to a possible quintillion), targeting what's happening in a designated range around a particle. Even when the two particles are whirling in large eddies many miles apart, the model measures the energetic activity only in the immediate vicinity of the particle. Yet it still accounts for the direction, speed, correlations in movement, and the distance that separates the two.

"It allows us to use only the necessary data and in a sense, filter out the superfluous noise that limits the number of scales you can examine," says Elliott. "It's like drawing a picture using only lines — it gets the essential information across."

Parallel Power

Majda and Elliott sought out the T3D for its processing power, which allows it to run millions of independent calculations and then pool the results. Using an approach that provides for the random forces at work in diffusion while taking into account the universal laws, their model constructs 1,024 different velocity fields. Each of those runs produces snapshots of every scale, incrementally charting the progress of the particles as they move away from one another. The resulting data are used to build a statistical framework that helps illuminate dispersion paths.

"We're interested in their distance from each other — how rapidly that grows as the particles move away from each other," says Elliott. "That reveals the characteristics of the turbulent system, which relates to the transport of tracers like pollutants, heat or water droplets."

The modeling technique, Monte Carlo simulation, involves independent repetitions of the same experiment. If these repetitions are run simultaneously on different processors, says Elliott, then the code achieves perfect scaling. In other words, doubling the number of processors halves the running time. "Because the algorithm requires small memory and because of the large number of available T3D processors," says Elliott, "we can run many experiments simultaneously, while sampling the particle position intermittently allows the program to run economically."

Majda and Elliott believe their model possibly could be adapted for large-scale endeavors, such as reproducing weather patterns or the mechanisms involved in cloud formation. "In cloud formation modeling for instance," explains Elliott, "we could use our technique to move droplets of liquid water around a cloud, integrate that with simulations of the other forces at work in these scenarios, and use that to understand the mechanisms responsible for producing rain." (JCW)

Reference:
Frank W. Elliott, Jr. & Andrew J. Majda, "Pair Dispersion Over An Inertial Range Spanning Many Decades," *Physics of Fluids* **8**, 1052–60 (1996).

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

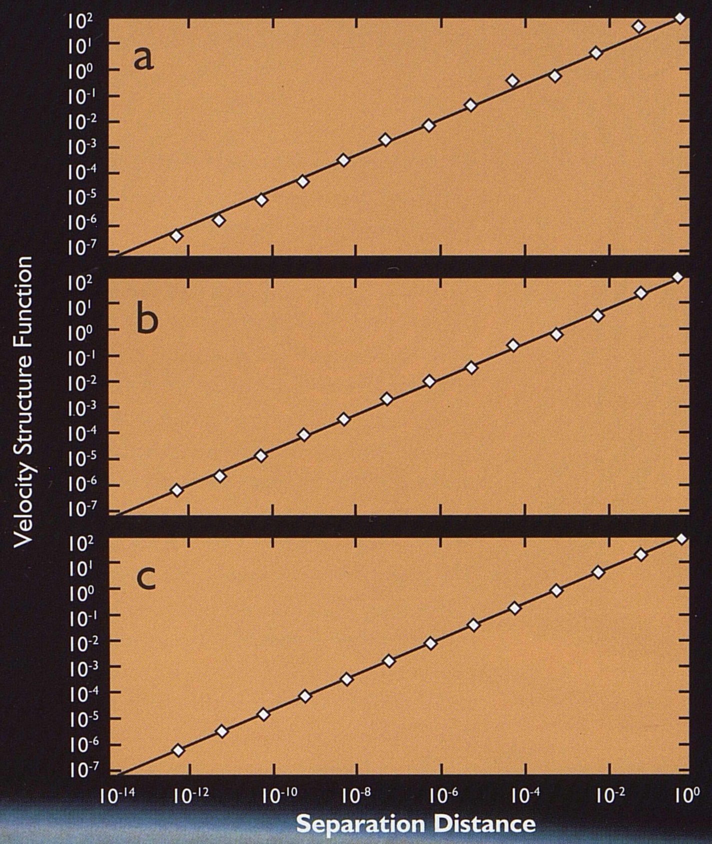

Dispersion of Particle Pairs in a Velocity Field

The fundamental quantity for describing dispersion in a given velocity field is the difference in velocity between any two points. These three graphs show the mean square value of this difference (called the structure function) in relation to distance between the points. In each graph, the solid line shows the theoretical prediction (mean square velocity difference is proportional to separation distance to the 2/3 power). The diamond-shaped points show simulated values of this statistic from averaging over 10 experiments (a), 100 experiments (b) and 1,000 experiments (c). In all cases the measured statistics agree well with theory.

This agreement, moreover, ranges from separations of about 10^{-12} to 1, indicating that this technique can simulate a turbulent velocity field over a range of scales from a millimeter to a million kilometers.

Their work represents a *major advance* in the ability to study **TURBULENT DIFFUSION** across a vast range of ScaleS.

Spiral Chaos

Elusive Heat, Chaotic Patterns

For decades scientists have used experiments replicating a classic form of heat transfer known as Rayleigh-Benard convection to study the effects of a warmer substance displacing a cooler one. These tightly controlled investigations of fluid dynamics allow scientists to scrutinize the fundamentals of behavior that is ubiquitous but otherwise elusive for study. A simmering pot of water, home heating and cooling, a weather system, or the fantastic furnace that is the sun — all possess convection attributes.

For the past several years, James D. Gunton, Haowen Xi and Jorge Vinals have used PSC's CRAY Y-MP and C90 to simulate Rayleigh-Benard convection. The Gunton team's simulations have revealed unpredicted chaotic patterns in R-B convection, unveiled new R-B convection structures and paved the way for mathematically accurate, three-dimensional renditions of chaotic behavior.

"We're trying to understand pattern formation and chaotic behavior in nature," says Gunton. "As a result of using supercomputing to simulate R-B convection, we're discovering previously unknown chaos in fluids."

Dependable Rayleigh-Benard

Rayleigh-Benard convection is significant because for a comparatively easy-to-replicate phenomenon, it continues to provide insight into understanding how heat energy moves through a flow system. Producing R-B convection involves isolating a liquid in a tiny enclosed cylindrical or rectangular cell and creating a temperature difference (gradient) between the bottom and top layers. It's the rough equivalent of heating a covered pan of water, though the experimental arrangement allows for precise control of fluid properties and the temperature gradient between the cylinder top and bottom, in addition to offering a safe view of the heated fluid's surface.

Under long-used experimental parameters, R-B experiments have consistently exhibited the same evolving structures, among them parallel rolls, squares and hexagons (discernible in the convection topography). The initial parameters and the type of fluid govern which kind of pattern will emerge first, but all the patterns are stable and non-chaotic. Given a constant temperature, they maintain their structure over time.

Changing Patterns

In 1992, Gunton's team replicated a laboratory experiment in which the expected transition from hexagonal to a parallel-roll state instead produced an unpredicted global spiral. The Gunton effort confirmed the experimental findings, revealing in a two-dimensional image a mesmerizing, stable rotating global spiral. "Convection often takes place in a very organized fluid motion," says Gunton, "and this has been studied and observed for decades. But supercomputing has helped reveal these novel, unpredicted states."

More recently, the Gunton team used the CRAY C90 to monitor the onset of Rayleigh-Benard from a non-convective state. The resulting two-dimensional images reveal a kaleidoscopic pattern of local spirals. Their findings matched those of an independent though simultaneous experimental effort, thus bolstering the veracity of the unpredicted find.

The spirals not only rotate, says Gunton, they move around in the fluid, annihilate each other, grow at the expense of one another and fluctuate in size and number. Gone is the regularity exhibited in the parallel rolls, with the fluid moving in continuous circular waves. Velocities and temperatures vary throughout the fluid. Some portions of it rise and fall faster than others. Everything about the system changes from one moment to the next. It's a crock pot of symmetry turned bubbly cauldron.

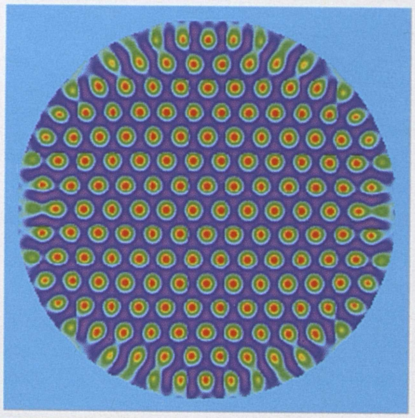

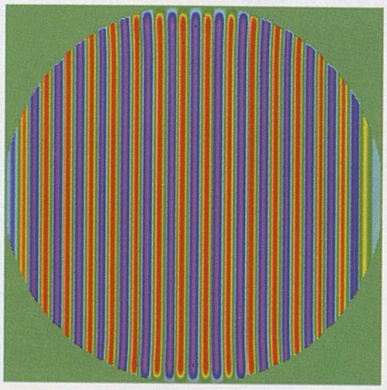

These images show the surface topography of parallel roll (left) and hexagon structures (above) in Rayleigh-Benard convection. Color corresponds to temperature, increasing from violet to red.

"Since these local spirals are erratic in time — changing position and size — there's an irregular behavior in space as well," says Gunton. Thus, he says, a very simple experiment can now be used to study one of nature's most complex occurrences — chaos. Chaos is a realm beyond order, whose disorder has a recognizable pattern over time. Chaotic patterns are seen in everything from weather systems to cardiac activity at the cellular level. "We're now using supercomputing to create three-dimensional models of the chaotic patterns," says Gunton, "which will provide further understanding of both convective behavior and chaos."

The new patterns revealed themselves when experimentalists, and later theorists, devised a means of enlarging the R-B cell, the enclosed setting in which the convection occurs and is observed. Previously, they could examine only a small piece of the convection system, which limited the extent to which pattern activity could evolve. "In order for us to mathematically model bigger cells, we had to solve some very complex equations," says Xi. "It couldn't have been done without supercomputing." The code ran between 300-350 Mflops on the C90 and the researchers have logged approximately 1200 hours on the supercomputers, with 900 more slated for investigations that will produce three-dimensional images of chaos in R-B convection.

"One of the quantities we try to calculate is the upward velocity of the fluid at different points in the system, at different times," says Gunton. "Solving the equations means determining the velocity of the fluid at any given point in the system at any given time. It can't be done without supercomputing." (JCW)

Reference:
Haowen Xi & James D. Gunton, "Spatiotemporal Chaos In a Model of Rayleigh-Benard Convection," *Physical Review E* **52/4**, 4963-75 (1995).

This work is supported by the National Science Foundation.

Chaos in Rayleigh-Benard Convection
The bull's eye-like symmetry of the global spiral (top) compared to the kaleidoscopic nature of a collection of adjacent local spirals (below), as simulated on the CRAY C90 by James Gunton and coworkers. Color corresponds to temperature, increasing from violet to red.

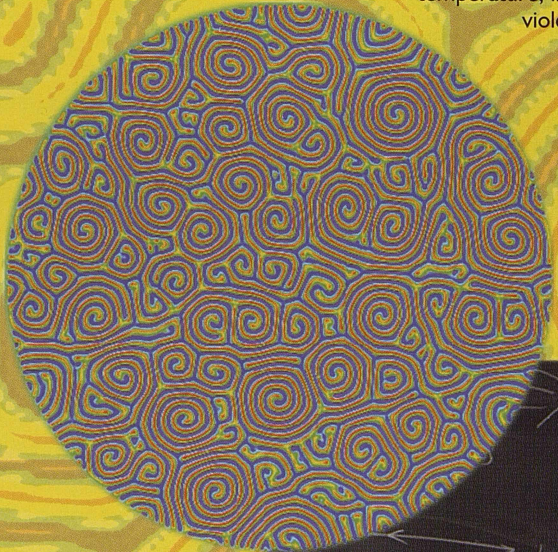

These simulations revealed unpredicted *CHAOTIC PATTERNS* in Rayleigh-Benard convection.

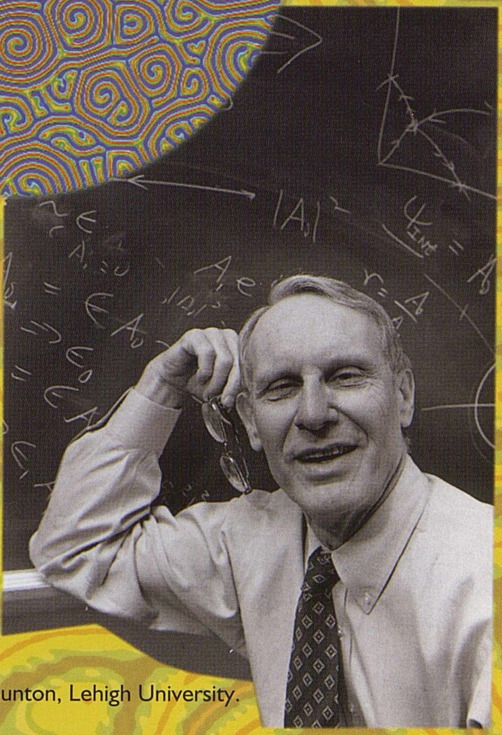

James Gunton, Lehigh University.

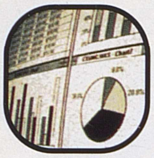

Economics on the Fly

Hubs and Spokes

When the U.S. government deregulated the airlines in 1978, it opened a tumultuous existence for the industry and gave birth to a controversy about what's good and what isn't for American travelers. Three deregulation legacies fuel the debate: a confounding pricing system, hub-and-spoke route structure, and domination of some hub airports by one or two airlines. Critics charge that this environment is the perfect breeding ground for monopoly, inefficient, costly operations and overpriced fares. Some suggest the government should reenter the aeronautical fray to increase competition by limiting the number of gates a carrier can lease at one hub or the number of departures it can schedule.

According to University of California, Berkeley economist Pablo T. Spiller, however, what may look like a quest to monopolize an airport is a natural consequence of operating efficiently. "This system is working," says Spiller, "and consumers are benefiting." Using data from the U.S. Department of Transportation and heavy-duty number crunching on Pittsburgh's CRAY C90, Spiller and his colleagues — Steven Berry of Yale University and graduate student Michael Carnall of the University of Illinois — completed the first study to sort through ticket purchase behavior by modeling supply and demand, and the results document who pays what for airline travel.

Pablo T. Spiller, University of California, Berkeley.

With the freedom brought by deregulation, the airlines quickly embraced hub-and-spoke network structures. As the industry learned on a limited basis prior to 1978, when a direct-route system was the norm, funneling passengers through a central location, or hub, where they can pick up connecting flights, offers the most logical means of moving large numbers of people to many cities many times a day.

Concentration of staff and aircraft at a hub often results in a carrier offering more departures to more destinations than carriers that base their operations elsewhere. Critics view this as behavior calculated to eliminate competition, and they charge that hub carriers have unreasonably high operating costs. Unlike other analyses, however, Spiller's study differentiates costs from markup, and the results show that at a given airport, a hub carrier enjoys 15-20 percent operational savings per passenger over a non-hub carrier at the same site. As its proponents argue, therefore, deregulation has fostered efficiency. "The critics," says Spiller, "are suggesting that the government tinker with the very structure that is allowing these savings to occur."

All's Fare: Modeling Airline Pricing

Part of the debate, says Spiller, represents an unwillingness to view airline prices like other consumer products. "You cannot look at a ticket price and say, *that's exorbitant*, anymore than you can make the same claim of an automobile sticker price without considering the model, make, options offered, and demand for it at a given time. You have to take into account what that fare represents — for instance, coveted departure and return times during peak flying periods, whether or not you book at the last moment, whether or not you'll earn frequent flyer miles."

With these factors in mind, the airlines compete by offering distinct products, presented to consumers as fares. Each fare represents a different market — a route connecting two cities — and departure time. Unlike prior analyses of the industry, Spiller's model captures this labyrinthine system and the buying behavior it breeds. It sorts customers into two groups, business and tourist, comparing the purchasing behavior of both in choosing myriad products within the same market.

Focusing on data from 1985-93, the model shows travelers choosing from among 230,000 combinations of itinerary, fare and carrier in as many as 17,000 markets. Conceptually, says Carnall, the model "takes every product and compares it against the other products in that market, and tries to figure out, given the choices made, what people value. And it looks at these thousands of markets separately in order to make that determination."

The first modeling effort crunched data from the fourth quarter of 1985. During those three months, according to Spiller, tourists using a dominant hub carrier paid anywhere from 1-5 percent above passengers whose flights were booked with non-hub carriers. Business travelers flying hub carriers, however, paid nearly 20 percent more than their counterparts using non-hub carriers. "The problem with previous studies," says Spiller, "is they implied that all travelers who used a hub carrier were paying considerably higher prices. And we find that only the business traveler is paying premiums."

The first study that documents who pays what for airline travel.

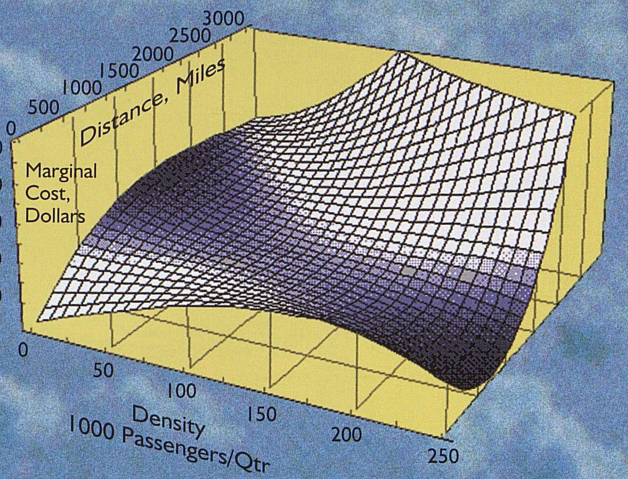

Marginal Cost vs. Distance and Density (4th Qtr, 1985)

The researchers modeled the marginal cost of a leg — a single takeoff to landing — as a function of flight distance and density (the number of passengers flown over the leg in a quarter). Density is a factor because larger aircraft used on dense legs lower per passenger costs. Shaded areas show cost decreasing with increasing density. Aircraft are inefficient for short, low-density legs (takeoffs and landings gobble up fuel). Very long legs are usually covered by large aircraft whenever density justifies a direct flight, so there is little opportunity to reduce costs on these routes.

Good News for Consumers

"As you develop a hub," says Spiller, "your products become more attractive — more direct flights, more frequent flights, more connections — and with that, you gain ability to mark up prices, because those are product qualities that, according to our data, customers are willing to pay for." If airlines are exploiting anything, says Spiller, it is "this peculiar demand for large networks that business travelers have." Consumers are doing well, says Carnall: "For a tourist, the cost of a flight has actually come down a bit." Business travelers are paying more, but are offered more frequent departures and other perks.

Thus far the model has examined individually the fourth quarters for 1985 through 1993. The next objective is to examine consecutive years simultaneously, so that the researchers can get an idea of how the 1990 recession and other economic milestones, such as numerous carrier bankruptcies, affected cost and demand.

The modeling imposes very large memory requirements. Estimation of the smallest quarter (in terms of data) required 96 megabytes of memory and the largest year required 26 megawords. Because no existing software would accommodate the study objectives, initial work involved testing and tweaking, a process, says Carnall, that would have taken years without a supercomputer. The C90 could do three days of workstation processing in an hour. "On the C90, we get results in as quickly as a day."

For complex social policy analysis such as this, says Spiller, supercomputing is essential. "For many issues, the problems we're dealing with are more complex, so you have to start using more sophisticated methods. The supercomputer is there to provide that type of service." (JCW)

References:

Steven Berry, Michael Carnall and Pablo T. Spiller, "Airline Hubs: Costs, Markups and the Implications of Customer Heterogeneity," (working paper).

Jan Brueckner, Nicola Dyer and Pablo T. Spiller, "Fare Determination in Airline Hub-and-spoke Networks," *RAND Journal of Economics* **23**, 309-33 (Fall 1992).

This research is supported by the National Science Foundation.

The Pittsburgh Supercomputing Center is a joint effort of Carnegie Mellon University and the University of Pittsburgh together with Westinghouse Electric Corp. It has received significant support from:

The National Science Foundation
The Commonwealth of Pennsylvania
The National Institutes of Health
The United States Department of Commerce
The Advanced Research Projects Agency
The Ben Franklin Technology Center of Western Pennsylvania
Buhl Foundation
Cray Research, Inc.
Digital Equipment Corp.
International Business Machines Corp.
The W. M. Keck Foundation
Network Systems Corp.
Silicon Graphics Inc.
Storage Technology Corp.
Sun Microsystems, Inc.
TCI of Pennsylvania, Inc.
Thinking Machines Corp.
Apple Computer, Inc.
Bell of Pennsylvania
The Vira L. Heinz Endowment
U. S. Dept. of Energy

credits
editor: *Michael Schneider, PSC.* **design & illustrations:** *Stephen Palmer, Pittsburgh.*
writing & graphics research: *Michael Schneider & John Westropp, PSC.* **copy editing:** *Sean Fulton, PSC.*
photography: *Jeff Macklin, Photography and Graphic Services at Mellon Institute; Jim Schafer; Westinghouse photographic services.*
printing: *Source W, Pittsburgh.*

Thanks to the PSC visual information systems group for their help with computer graphics. Thanks to Gina Paterno, Mary Thomas and Karen Tobkes of the PSC staff for transcribing research interviews.

Printed on recycled paper with vegetable-based inks.

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

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