









1995

# P R O J E C T S I N


Pittsburgh Supercomputing Center

# SCIENTIFIC COMPUTING





# Projects in Scientific Computing 1995



---

P I T T S B U R G H

---

S U P E R C O M P U T I N G

---

C E N T E R

---

Published July 1995 by the Pittsburgh Supercomputing Center,  
4400 Fifth Avenue, Pittsburgh, PA 15213. (412) 268-4960.

Copyright © 1995, Pittsburgh Supercomputing Center.  
All rights reserved.

# Foreword

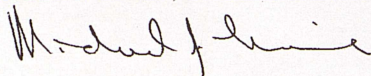
We are in a time of national political reassessment, which includes reexamining our scientific priorities as well. As the Pittsburgh Supercomputing Center enters its tenth year, it is worthwhile to look at how far we have come as well as what lies ahead. We, and our sister National Science Foundation centers, have changed dramatically in the past decade. Our computers are 100 times more powerful than they were. This has led to new scientific insights into problems that were intractable a decade ago. We are driving important new developments in computational science, which is enhancing users' productivity, and ensuring continued U.S. predominance in this vital area of international competitiveness.

What has remained constant throughout this period of change is that high-end computational capability has contributed to advancing scientific research. The experiment of harnessing the creative energies in our universities to centers providing support and innovations in computational science has been remarkably successful.

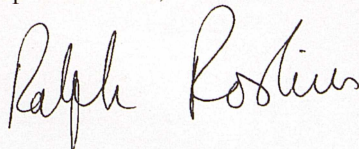
We at PSC are very proud of what our users have accomplished. With our support and using the world's best computing equipment, they have pushed forward scientific understanding in many areas. By stressing capability computing, encouraging single users to exploit all 16 processors on the CRAY C90, or half or even all 512 processors on our CRAY T3D for sustained periods, we have enabled insights in fields as diverse as ocean currents in the North Atlantic and the origin of galaxies.

At this time, while high-performance computing technology is increasingly realizing the advantages of parallelism, both vector supercomputers and massively parallel machines are essential. The choice is and must be driven by scientific considerations. Today's balance is shown in the projects highlighted in this booklet — six have been carried out on the C90, and six on the T3D, while two have used both. Our major effort in making easily used parallel software available is contributing substantially to increased scientific productivity. Where such codes are available, we have found that they are more cost-effective.

Michael J. Levine, Scientific Director




Ralph Z. Roskies, Scientific Director



This past year has seen an explosion in the National Information Infrastructure and our ability to access, retrieve, store and transmit information. The NSF centers have often led the way in developing and deploying new technology. New communications and data-handling capabilities are transforming several fields of science and also radically changing how we disseminate information and how we teach.

The past 10 years has 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.

It is with profound appreciation that we acknowledge our major funding sources. The centers were founded by NSF and nurtured by major high-performance computing and communications vendors and other government agencies. While NSF has continued to be the largest single source of funding, we would not have nearly as powerful massively parallel machines were it not for Advanced Research Projects Agency support. National Institutes of Health support has enabled many of the biomedical advances you will read about in this booklet. And the Commonwealth of Pennsylvania's support in very demanding times has been critical to our success.



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


# Projects in Scientific Computing, 1995

## Contents

<b>Pittsburgh Supercomputing Center, 1995</b>	<b>4</b>
<b>Pushing Technology Forward</b>	<b>10</b>
<b>Scientific Visualization</b>	<b>12</b>

---


### Modeling Global Climate



---

<b>High Tide in Ocean Modeling</b>	<b>14</b>
------------------------------------	-----------

High Resolution Simulation of Global Ocean Circulation  
Matt O'Keefe & Aaron Sawdey, University of Minnesota  
Rainer Bleck, University of Miami



---

<b>Unlocking the Sun's Secrets</b>	<b>16</b>
------------------------------------	-----------

Simulation of Turbulent Compressible Convection  
Paul Woodward & David Porter, University of Minnesota

---


### Structure of Proteins & DNA



---

<b>Melting Proteins</b>	<b>18</b>
-------------------------	-----------

The Effect of Long-Range Forces in Molecular Dynamics Simulations  
Tom A. Darden, National Institute of Environmental Health Science  
Lee Pedersen, University of North Carolina, Chapel Hill




---

<b>Shifty Biomolecules</b>	<b>20</b>
----------------------------	-----------

Macromolecular Structures in Solution via Analysis of 2D NMR Experiments  
Thomas James, University of California, San Francisco

---

### Understanding the Brain




---

<b>A New Look at Seeing</b>	<b>22</b>
-----------------------------	-----------

A Unified Self-Organizing Model of the Primary Visual Cortex  
Risto Miikkulainen & Joseph Sirosh, University of Texas at Austin

---

### Respiration and the Lungs




---

<b>Breathing Lessons</b>	<b>24</b>
--------------------------	-----------

Computational Analysis of Small Airway Fluid Dynamics  
Jeffrey R. Hammersley, Medical College of Ohio

---

## Evolution and Structure of the Universe



---

### Groping in the Dark

26

The Formation of Galaxies and Large-Scale Structure of the Universe  
Edmund Bertsehringer, Massachusetts Institute of Technology

---

## Manufacturing Process Design



---

### Molecular Spaghetti

28

Studies of Concentrated Polymer Solutions in Strong Flows  
Gary Leal, University of California, Santa Barbara

---

## Semiconductor Design and Production



---

### Travels with Silicon

30

Temperature-Dependent Properties, Diffusion and Growth on the Si(100) Surface  
Jerzy Bernholc, North Carolina State University

---

## Design of New Materials



---

### Dancing with Electrons

32

Magnetic Properties of Thin Films  
Arthur J. Freeman, Northwestern University



---

### The Amazing Metal Sponge

34

Constant Chemical Potential Simulations of Palladium-Hydride  
Khalid A. Mansour, Cray Research, Inc.  
Ralph J. Wolf, Westinghouse Savannah River Company

---

## Quantum Chemistry



---

### Swimming Ions & Jumpy Electrons

36

Solvent Free Energies for Electron Transfer at the Water-Metal Interface  
Gregory Voth, University of Pennsylvania

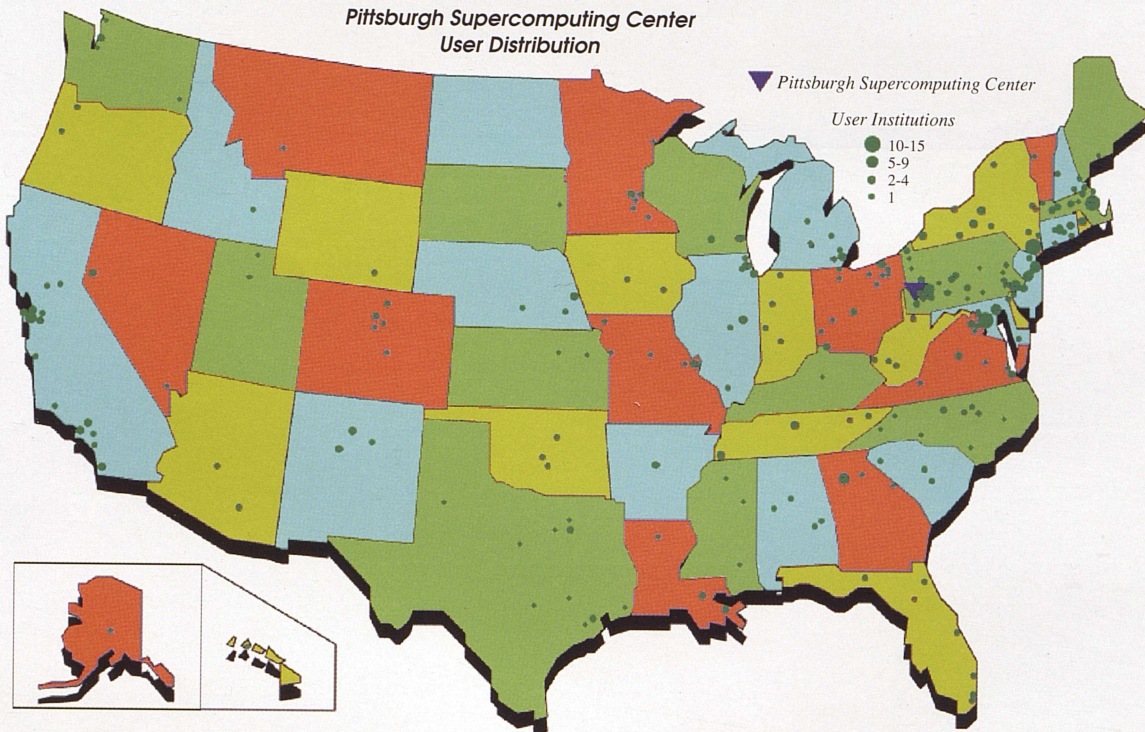
## Pittsburgh Supercomputing Center, 1995

The Pittsburgh Supercomputing Center's primary mission is to provide hardware, software and know-how in high-performance computing that facilitate science and engineering research. It is gratifying, therefore, that the past year brought recognition for several research achievements made possible by PSC resources. The 1995 Computerworld Smithsonian Awards recognized simulations of Comet Shoemaker-Levy 9 crashing into Jupiter carried out at PSC by Mordecai-Mark Mac Low of the University of Chicago. These computations showed that the comet would penetrate less deeply and explode more violently than other models predicted and helped observers worldwide plan for the crash. They continue to play a role in interpreting observational data from this remarkable event.

For his research on blood flow in the heart using PSC's CRAY C90, New York University scientist Charles Peskin received the 1994 Sidney Fernbach Award, which recognizes "outstanding contribution in the application of high-performance computers using innovative approaches." Peskin's three-dimensional model of blood flow in the heart, its nearby valves and major vessels will make it possible to address many questions difficult or impossible to get at in animal research and clinical studies.

For his PSC research on the interaction between proteins and DNA, University of Pittsburgh biologist John Rosenberg was a finalist in the 1994 *Discover* magazine Technology Awards. Rosenberg's simulations led to fundamental insights into the process by which proteins recognize and attach to just the right spot on the long, twisted strands of a DNA molecule.

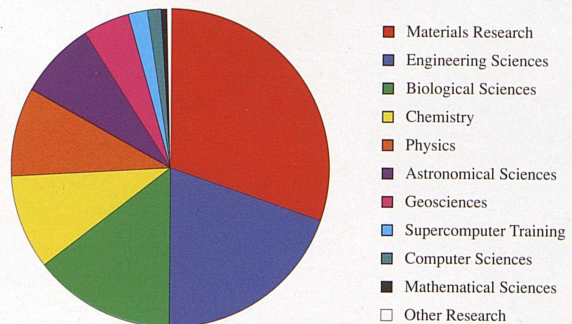
*Each of these research projects is described in previous versions of this publication, available in interactive, multimedia format on World Wide Web: <http://pscinfo.psc.edu/MetaCenter/MetaScience/Articles/Contents.html>*



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

To date, more than 10,000 scientists and engineers at more than 1,300 universities and research centers (green dots) in 49 states and the District of Columbia have used the center's computing resources to advance their research. This work has resulted in more than 3,000 published papers in professional science and engineering journals. Researchers access the center's resources primarily via the Internet.

Twenty-six universities are Pittsburgh Supercomputing Center Academic Affiliates (see back page). Representatives from these campuses form the center's main advisory body.



**CRAY C90 Usage by Discipline**

### Capability Computing in a Heterogeneous Environment


A key to enabling these breakthroughs, as well as many other significant research efforts at PSC, is a focus on “capability” computing — enabling large-scale computations that would not be feasible on smaller systems such as workstations. This is made possible by PSC technologies that unite the center’s computing platforms into an integrated, heterogeneous system. The CRAY T3D scalable massively parallel processing (MPP) system is tightly coupled to the CRAY C90, a parallel vector supercomputing system, via a high-speed data transfer link. The result is the world’s first single-vendor heterogeneous system.

As described elsewhere in this booklet, researchers during the past year have used the T3D to great advantage. With significant results in diverse research areas — cosmology, materials science, neural simulation, turbulence and ocean circulation — the T3D is proving to be a highly capable system. It has been remarkably stable, even through this year’s upgrade from 128 to 512 processors, and it has been used for production research at a high level of efficiency as development work has proceeded. With PSC’s success at converting heavily used applications software to the T3D (p. 10), large-scale parallelism has stepped to the fore of general-purpose scientific computing.

Through the center’s focus on capability computing, a number of “grand challenge” projects — large-scale computational undertakings with important scientific and social impact — receive consulting support from PSC staff. Among these are an NSF grand challenge project on geophysical and astrophysical turbulence (p. 16) and the Grand Challenge Cosmology Consortium (p. 26).

*Support from PSC scientists and systems engineers for large scale “landmark” projects is part of the center’s focus on capability computing. PSC scientists Gary Beckwith (left)*

*and Nigel Goddard worked with University of Texas researchers Risto Mäikkulainen and Joseph Sirosh to get their neural model of the primary visual cortex (p. 22) up and running on the CRAY T3D.*



*Beverly Clayton, PSC executive director.*



*Jim Kasdorf, director of supercomputing, Westinghouse Electric Corp. and members of the Westinghouse Energy Center (WEC) operations staff (l to r): Ernest L. Harvey, Kasdorf, Ron J. Flex and Kim Wilt. Westinghouse staff provide operations support for the CRAY T3D (left) and C90, which are housed at WEC in Monroeville, Pa.*

The C90 continues to serve as the workhorse of computational productivity for most PSC researchers. A number of projects (e.g., p. 30) employ the heterogeneous capability of the C90/T3D linked system, and other researchers use the two systems for different parts of their work. PSC staff actively encourage and provide consulting to support multi-processor parallelism. As a result, the portion of C90 usage devoted to parallel computing has steadily increased. It now averages around 30 percent, and the C90 overall is used at a high level of efficiency (about 3.4 billion calculations per second for all 16 processors).

The T3D/C90 is closely linked to the 14-processor Digital Equipment Corp. Alpha workstation cluster. All three platforms are also linked to archival and file servers through high-speed interconnects. The DEC Alpha cluster complements the C90 and T3D, providing a versatile alternative that meets the needs of many research problems. This year, the center's consulting staff has helped to transfer several projects that did not effectively exploit the C90's vector capability to the Alpha cluster, where 14-processor parallelism results in faster turnaround for these essentially scalar jobs.

The center's work on creating software tools includes fruitful collaborations with computer scientists. CMU computer scientist Thomas Gross developed a library for transposes (a frequently used mathematical operation) on the T3D that is 30 times faster than the vendor supplied library. This leads to much improved FFT (Fast Fourier Transform) performance. Adam Beguelin, who holds a joint appointment at PSC and CMU's School of Computer Science, extended his PVM (Parallel Virtual Machine) development work to new products that address load balancing in distributed environments. He also organized the national PVM workshop at PSC in May 1995.

### Networking for Research and in the Community

This year saw a major changeover in network connections among the NSF-supported supercomputing centers. Effective in April, NSFnet, which has served as the "backbone" of the Internet, was replaced by the very high-speed Backbone Network Services (vBNS) network. Based in fiber-optic technology and operated by MCI, vBNS is three times as fast (155 million bits per second) as NSFnet and is expected to be sixteen times faster yet within five years.

PSC's networking group has been actively involved for several years in planning for this changeover, which began with a connection between PSC and the San Diego Supercomputer Center. Through participation in the Internet Engineering Task Force, they helped develop standards upon which vBNS and other networks are based, and locally they coordinated a smooth transition to the new network.

The network group also plays an active role supporting networks in the Pittsburgh-area community. Common Knowledge: Pittsburgh is an NSF-funded pilot program for integrating networking into a large urban public school. Now in its third year, CK:P involves the center staff in technical support and planning to effectively use the Internet in the classroom. PSC staff helped design the local area network (LAN) for each of four initial sites, two elementary and two high schools, and seven more added this year. Each LAN connects to a router at PSC that lets the schools communicate with each other.

PSC staff also provide technical support for BRIDGES, a project with funding from the U.S. Department of Commerce to connect the Carnegie Public Library in the "Hill District," a Pittsburgh inner-city community, to the Internet. The networking group is investigating ways to provide cost-effective connectivity for these community projects. One option being explored is innovative use of the local cable-TV infrastructure.




*Members of the PSC communications group (standing, l to r): Mike Lambert, Gene Hastings, Steve Cunningham, Shirl Grant, Kevin Sullivan, Andy Adams. (seated): Jamshid Mahdavi, Wendy Huntoon (manager), Tina Sabeh, John Studarus.*

### Biomedical Supercomputing

One part of PSC's very active biomedical program, funded through the National Center for Research Resources' Biomedical Research Technology Program, is a series of workshops that offer training on the use of supercomputing in biomedical research. The importance of these workshops as a means of knowledge sharing was exemplified this past year by the July 1994 workshop in "Computational X-Ray Crystallography." Nobel Prize-winning biophysicist Herbert Hauptman led the workshop, which offered the first formal training in the use of Shake-and-Bake, a computational method developed by Hauptman and colleagues at the Medical Foundation of Buffalo. Shake-and-Bake is an "automated" method for determining the structure of biomolecules from x-ray diffraction data. From over 90 applicants, 20 researchers from around the country were selected to participate.

Another of this year's biomedical training highlights was "Applications of High Performance Computing in Bioengineering" held in October 1994. This two-day workshop focused on the current status of computational methods in bioengineering and how high-performance computing can help to overcome barriers to progress in the field. It addressed topics such as modeling of bone-healing processes and computer-aided prostheses design. The workshop brought together 23 experts from a range of computational and bioengineering disciplines to present papers and participate in roundtable discussions.



*Chad Vizino (left) and Prithvi Rao of the production systems group, managed by Janet Brown. Along with maintaining the center's mass storage system, this group keeps the CRAY C90 in tiptop running condition and helps in providing software support for the Andrew File System. AFS was adapted by PSC for use in a supercomputing environment and is now used as a shared system among all NSF supercomputing centers.*




*Members of the PSC administrative staff (clockwise from left): Tracy Mills, Pat Tomko-Buck, Anne Marie Zellner, Elvira Prologo (office manager), D. Cook and Tina Sabeh.*

### Learning About MPP and "Distance Learning"

This year the center sponsored 22 workshops in a range of topics, from introductory to advanced applications, with nearly 400 participants. Hand-in-hand with the center's development effort making scientific application packages available on the CRAY T3D (p. 10), the education staff has increased training opportunities in massively parallel processing (MPP). "Parallel Processing/CRAY T3D" is a new five-day workshop (offered four times this year) that introduces users to MPP in general and, in particular, to the 512-processor CRAY T3D.

A new wrinkle in PSC's workshops, instituted this year, is "distance learning." The education staff provides lecture materials on World Wide Web, so that workshop participants can preview the materials and follow lectures on their workstations during the workshop. PSC will refine these materials, adding detail to the notes along with links to PSC documentation, text, audio and video. As this work proceeds, these WWW documents will become a source for researchers to learn independently from workshop attendance.



*PSC conducts its workshops in the Computer Training Center, upgraded this year with 25 new INDY Silicon Graphics workstations, each with 32 megabytes of memory and a high-resolution color monitor.*

### Pittsburgh Supercomputing Center Workshops (July 1994 - June 1995)

Summer Institute on Supercomputing  
 Supercomputing Techniques: Parallel Processing/CRAY T3D  
 Supercomputing Techniques: Cray C90  
 PVM and Distributed Computing  
 Computational X-Ray Crystallography  
 Heterogeneous Supercomputing for Biomedical Research  
 High-Performance Computing in Bioengineering  
 Nucleic Acid and Protein Sequence Analysis  
 Structure Determination Using Nuclear Magnetic Resonance



### The PSC Hotline

*Casey Porto, staff consultant, and Tom Maiden, student consultant, on duty at the PSC hotline. By responding directly to trouble calls, hotline staff are the frontline problem solvers for PSC researchers.*

*Ken Hackworth, user consultant, trains and supervises a group of 10 undergraduate students who staff the hotline in rotating shifts, backed up by a fulltime PSC consultant. More than 50 students have served as hotline consultants since PSC opened, and several have "graduated" to fulltime staff positions.*

*The hotline operates Monday to Friday 9 a.m. to 8 p.m. and Saturday 9 a.m. to 4 p.m. Hotline consultants field 500 to 700 calls a month on questions that run the gamut, from forgotten passwords to code debugging. The hotline goal is that each call results in a more educated user.*


## USX • ALCOA • WESTINGHOUSE

### Corporate Partners

PSC corporate partners have found that attacking research problems with computational science offers a core competitive advantage in their industry. Availability of the CRAY C90 and T3D for commercial research and product development has played a key role in establishing these partnerships.



*Martin Straut, corporate program officer.  
"Global competition and the need to improve time-to-market for new products," says Straut, "are producing a growing convergence between corporations and high-performance computing."*



*J. Ray Scott joined PSC this year as manager of technical facilities. Terry Sebben coordinates PSC support systems, which include the allocations data base and Macintosh support.*



*Members of the advanced systems group work on various projects in distributed data management, heterogeneous computing and the Alpha workstation cluster. (l to r): Rob Pennington (coordinator), Juan Leon and Kathy Benninger.*

### Information at the Pittsburgh Supercomputing Center

#### Proposals for Computing Time

*Susan Lambl*

#### Biomedical Initiative

*Nancy Blankenstein*

#### Workshops & Summer Institute

*Cathy Milligan*

#### Corporate Programs

*Martin Straut*

#### Newsletter and Documentation

*Vivian Benton*

These Pittsburgh Supercomputing Center staff members can be contacted by telephone (412-268-4960) or electronic mail: [lastname@psc.edu](mailto:lastname@psc.edu). Information about PSC is also available on World Wide Web: <http://pscinfo.psc.edu>

Research articles from *Projects in Scientific Computing* are available in an interactive, multimedia format on World Wide Web: <http://pscinfo.psc.edu/MetaCenter/MetaScience/Articles/Contents.html>

## Pushing Technology Forward

For years, massively parallel processing — teaming hundreds or thousands of processors to work simultaneously on one computing job — has been a good idea. Success at exploiting the potential of MPP, nevertheless, has suffered from a lack of applications software. In 1995, after two years of work with the CRAY T3D at Pittsburgh Supercomputing Center, this is changing.

---

*PSC scientists and consultants are now among the world's leading experts on developing scalable parallel applications for Cray systems.*

---

The T3D at Pittsburgh is taking MPP to a new level of scientific productivity, and it is proving itself as a cost-effective way to gain computing power. As evidenced by projects described in this booklet, researchers using this system have obtained important results in ocean modeling, cosmology, materials science, neural modeling and turbulence simulation. For many of these applications, the T3D outperforms the CRAY C90.

Along with facilitating user projects, PSC has for the last two years carried out a concerted effort in T3D software development. A working group of PSC scientists and systems engineers joined forces with on-site Cray developers. For this joint effort, the Parallel Applications Technology Program (PATP), the objective has been to make effective applications software available on the T3D, especially packages frequently used on PSC's CRAY C90.

PATP has borne fruit. In June 1995, PSC announced T3D availability for a number of widely used scientific software packages (see sidebar) — GAMESS, CHARMM, MultiSearch, AMBER and Shake-n-Bake. Other PATP results include an efficient parallel finite-element method, much improved Fast Fourier Transform capability, and a parallel implementation of simulated annealing, a method for locating the global minimum of a potential function.

The widespread use of these third-party packages — 30 to 40 percent of PSC's CRAY C90 cycles — indicates their importance. With implementation of this software on the T3D, MPP is no longer confined to a narrow range of applications. Researchers in many fields can exploit the raw potential of MPP hardware.

Through PATP, PSC staff have worked closely with Cray developers, learning many things about the soul of this new machine. An invaluable byproduct is that PSC scientists and consultants are now among the world's leading experts on scalable parallel applications for Cray systems. PSC personnel adapt and write utilities for T3D users and provide consulting and instruction on parallel code development. They also advise on optimum approaches to research problems on the T3D and other scalable parallel systems.



*Members of the PSC Parallel Applications Technology Program development team gathered around PSC's CRAY T3D at Westinghouse Energy Center, Monroeville (l to r): William Young, Qiming Zhang (Cray) David O'Neal, Asiri Nanayakkara (Cray), Nick Nystrom, Frank Wimberly, Sergiu Sanielevici (PSC parallel applications manager), Alex Ropelewski, John Kyle (Cray), Michael Crowley, Raghurama Reddy and Carlos Gonzalez.*

## **Parallel Applications at PSC**



*The following applications packages are now available for use on PSC's CRAY T3D:*

- *AMBER is used to model and simulate biological macromolecular systems. An empirical potential energy function describes the interaction between atoms. AMBER scales well on the T3D, with performance on 16 T3D processors roughly equivalent to one C90 processor.*
- *CHARMM is another program that uses empirical energy functions to model macromolecular systems. It is widely used to study protein and peptide systems in aqueous and gas phase and can also be used to refine structures from NMR and x-ray crystallography.*
- *GAMESS is a widely used, general purpose, ab initio quantum chemistry program best known for its ability to perform multiconfiguration self-consistent field calculations. Performance on the T3D is extremely cost-effective relative to the CRAY C90.*
- *MultiSearch, an extension of MaxSegs, is a portable parallel application for sequence database searching and alignment. Developed at PSC, it now runs on most MIMD machines, including the T3D. MultiSearch can use sequence databases in a number of file formats, and it incorporates four different searching and alignment algorithms.*
- *Shake-n-Bake implements a direct methods phasing algorithm for determining crystal structure from raw crystallographic reflection data. The remarkable speed and scaling of this program on the T3D makes possible very large refinements of structures.*

*More information on these packages and their T3D implementation can be located on the World Wide Web: <http://pscinfo.psc.edu/general/software/alphabetical/alphabetical.html>*

### Helix-Helix Packing

An educational video being produced by PSC's biomedical group, shows how two protein helices pack together. This structure, from the protein lysozyme, is a helix-turn-helix motif, with residues (purple) defining the turn. Sidechains (green) of one helix (blue) define a ridge that fits into the groove defined by sidechains (yellow) of the other helix (red). Other sidechains (cyan and orange) participate in the helix-helix contact but not in the ridge-groove interaction. PSC biomedical scientist David Deerfield and visualization specialist Greg Foss produced the image.



### Airway of a Child


What toys are likely to cause choking in the airway of a small child? In collaboration with PSC scientist Joel Welling, researchers from Incheape Testing Labs and Children's Hospital in Pittsburgh are working to build models that will help make better predictions. The 3-D image is reconstructed from 2-D scans (red planes) of the airway of a young child. An interpolation algorithm plots airway location (white) between planes. The image shows the change produced (yellow) by slightly altering the algorithm.

# Scientific Visualization

### Storms in Oklahoma


Throughout May 1995, the Center for Analysis and Prediction of Storms, a National Science Foundation Science and Technology Center at the University of Oklahoma, used the CRAY T3D at Pittsburgh to produce a daily weather forecast for parts of Oklahoma. PSC's visualization group processed data from selected forecasts into 3-D animations.

Shown here is visualized data for an area 67 kilometers square extending upward 17 kilometers. The white area represents high concentrations of combined cloud water and rain water. Red shows columns of high vorticity that potentially spawn tornadoes. The blue ribbons indicate wind velocity at two planes, approximately 600 meters and 4800 meters above ground.



### A Dwarf Galaxy and its Parent

A "dwarf" galaxy, shown as yellow, green and blue points, is formed from streams of stars and gas ripped from a large spiral galaxy (red) by its interaction with another galaxy. This image represents one time-frame from simulations of galaxy interaction carried out at PSC by astrophysicists Lars Hernquist and Josh Barnes.



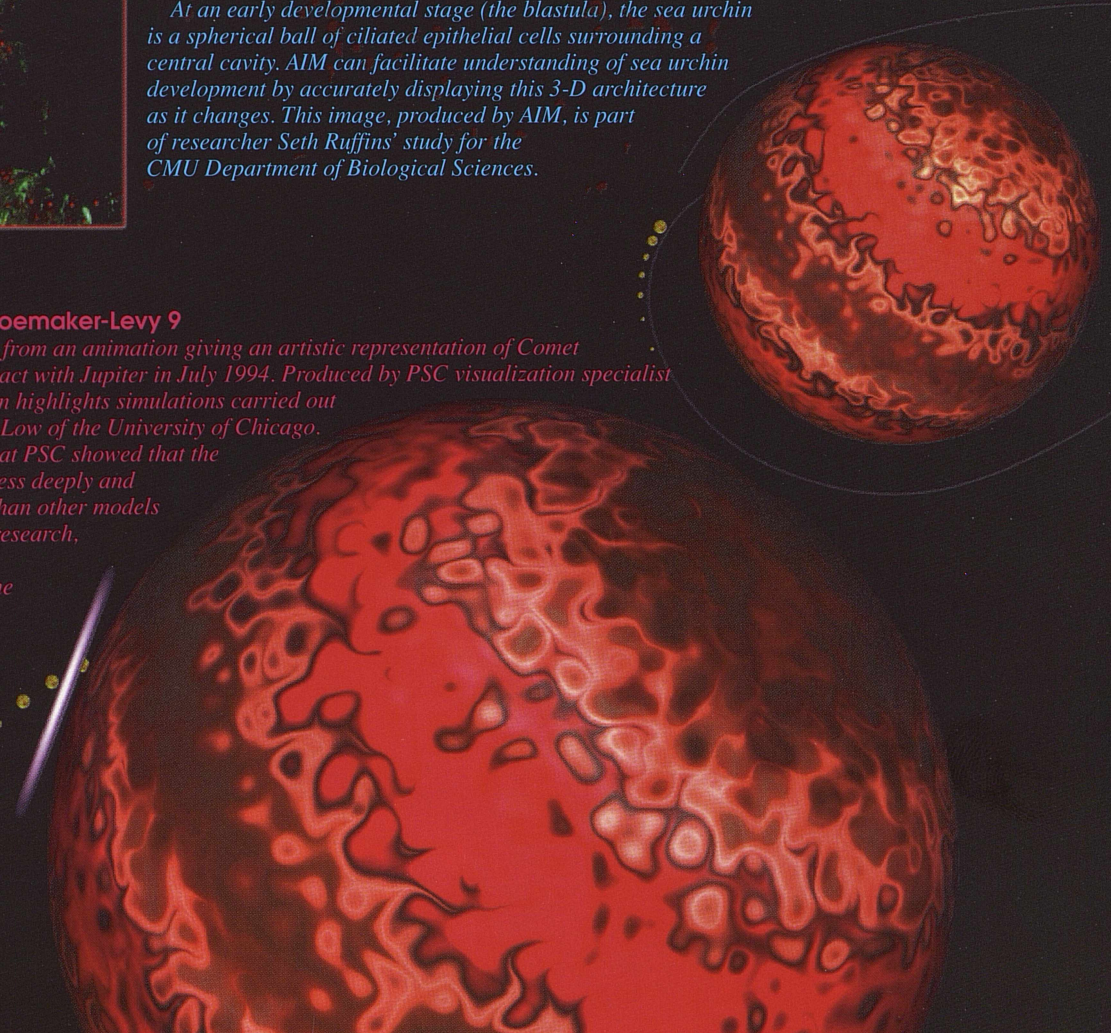
### Sea Urchin Development


The Automated Interactive Microscope (AIM), a collaboration between PSC and the Center for Light Microscope Imaging and Biotechnology at Carnegie Mellon University, develops high-speed network links between supercomputers and advanced microscopes. AIM will greatly accelerate the acquisition, analysis and display of 3-D image data.

At an early developmental stage (the blastula), the sea urchin is a spherical ball of ciliated epithelial cells surrounding a central cavity. AIM can facilitate understanding of sea urchin development by accurately displaying this 3-D architecture as it changes. This image, produced by AIM, is part of researcher Seth Ruffins' study for the CMU Department of Biological Sciences.

### Impact of Comet Shoemaker-Levy 9

These two images are from an animation giving an artistic representation of Comet Shoemaker-Levy 9's impact with Jupiter in July 1994. Produced by PSC visualization specialist Greg Foss, the animation highlights simulations carried out by Mordecai-Mark Mac Low of the University of Chicago. Mac Low's calculations at PSC showed that the comet would penetrate less deeply and explode more violently than other models had predicted. For this research, PSC and Mac Low were honored as finalists in the science category of the 1995 Computerworld Smithsonian awards.





*“A two-foot rise in sea level could have precisely the same effect on ocean waves that a 2° C. rise in temperature would have on heat waves. It would make the freak disaster commonplace.”*

—J. Weiner, *The Next One Hundred Years*

### Toward a Realistic Computerized Ocean

The ocean is rising! The ocean is rising! It's a cry reminiscent of Chicken Little, who thought the sky was falling when an acorn hit him on the head. The important difference is the ocean *is* rising.

Scientists estimate sea level is roughly four inches higher than it was 100 years ago, and satellite measurements show it going up one to three millimeters a year. Do the rising levels reflect climate change associated with alarms about global warming, as many scientists believe? Or, as others argue, are they part of normal fluctuations in weather cycles that will even out over time? Either way, should we be worried? What are the potential effects for coastal population centers?

The need for answers runs smack up against the vast uncertainties inherent in a system as complex as Earth's climate. The best tool — the only tool — we have for assimilating the multitude of variables and trying to make rational predictions is computer modeling, and it's not yet good enough. “In order for climate models to be reliable,” says University of Miami ocean scientist Rainer Bleck, “you want them to be able to accurately predict climate for stationary conditions, without the added variabilities of increasing carbon dioxide in the atmosphere. We're not at that stage yet.”

Recent computations at the Pittsburgh Supercomputing Center are a notable step in the right direction. Taking advantage of the parallel-processing ability of the CRAY T3D, Bleck and his collaborators, University of Minnesota computer engineers Matt O'Keefe and

Aaron Sawdey, proved the feasibility of a revised approach to ocean modeling.


In a simulation that ran for 10 days on 256 T3D processors, half the machine, their model of circulation in the Atlantic Ocean correctly predicted the course of the Gulf Stream. No other circulation model of the entire Atlantic Ocean has done this. “This degree of realism,” says Bleck, “is very much a step up from previous simulations.”

### Putting the Gulf Stream on Course with MPP


For years, the sticking point for modeling the North Atlantic has been Cape Hatteras, the bump on the U.S. east coast between North and South Carolina. The Gulf Stream hugs the coast as it comes up from Florida until at Cape Hatteras it veers off on a line, east-to-northeast, toward the open Atlantic. In computerized oceans, however, the Gulf Stream has preferred turning left past Cape Hatteras, clinging to the shoreline rather than taking to the open sea.

“People have suspected grid-size is the problem,” explains Bleck, referring to the mesh-like grid that divides the ocean into segments to compute the variables of interest — temperature, salinity, current direction and speed, etc. With a finer grid, smaller segments, chances improve that the results mirror reality. Until now, however, modeling the entire Atlantic at high enough resolution to get the currents right required a prohibitive amount of computing.

Rainer Bleck (left), University of Miami,  
Matt O'Keefe and Aaron Sawdey,  
University of Minnesota.



*“We've crossed a threshold. This kind of code actually runs faster on MPP than the current vector machines.”*



*"This degree of realism is very much a step up from previous simulations."*

Availability of the CRAY T3D at Pittsburgh provided an opportunity. O'Keefe and Sawdey adapted Bleck's model, the Miami Isopycnic Coordinate Model (MICOM), to massively parallel processing (MPP). In October 1994, PSC made computing time available for a dedicated run on 256 of the T3D's 512 processors. With grid-size at 0.08° longitude (roughly six kilometers at mid-latitude), and with the model covering the North Atlantic and extending south of the equator, MICOM ran for 10 days, carrying out 10 months of simulation. The results, captured in a video animation, show that the Gulf Stream immediately begins establishing itself on the right path.

"This model was an attractive test for MPP," says O'Keefe, and the T3D has done well. Performance data suggests MICOM will hum along at about 9.6 billion calculations a second (Gflops) on a fully configured T3D (1024 processors) as compared to 4.4 Gflops on a 16-processor CRAY C90, the current top-of-the-line vector system. The availability of a large amount of memory (16 gigabytes) on the T3D, says O'Keefe, is a big factor in achieving this performance: "We've crossed a threshold. This kind of code actually runs faster on MPP than on current vector machines."

#### **MICOM: An "Isopycnic" Ocean Model**

These results place MICOM at the forefront of ocean modeling, not only because the Gulf Stream goes where it should, but also because MICOM — which has a 15-year history — has now proven that a significantly revised approach to modeling the ocean's vertical dimension works.

Ocean modeling began in the 1960s, but progress has lagged behind atmospheric models. The relative sparsity of observational data is one

factor, but just as limiting, says Bleck, is the greater complexity of closed-basin systems like the ocean and the highly nonlinear equations required to describe buoyancy effects in sea water.

In particular, conventional ocean modeling has been vexed by distortions in the interaction between the ocean surface, which captures heat from the sun, and deeper, colder regions. "The ocean is heated from above," explains Bleck. "Warm water sits on top of cold water, and you need to make sure that this temperature contrast doesn't get diminished over time." The problem has been that it is inherent in the numerical scheme of conventional models for heat to trickle between warm and cold regions.

"We have designed a set of equations," says Bleck, "where this type of heat diffusion is removed." The ocean is divided into 11 layers, each of which maintains its own density — hence the term *isopycnic*, meaning constant density. Each layer is separate, as if divided by plastic sheets from the layers above and below it — greatly reducing the spurious effects of heat diffusion.

The T3D results at Pittsburgh show that this approach works. "It was not clear 10 years ago," says Bleck, "that this would be practical. We've proven that the equations from this more sophisticated scheme are numerically solvable. This is important for climate research." (MS)

#### **Reference:**

Aaron Sawdey, Matthew O'Keefe & Rainer Bleck, "The Design, Implementation, and Performance of a Parallel Ocean Circulation Model," (preprint, 1995).

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

Photo of researchers: Eric P. Chassignet, University of Miami.

**Beneath the Surface**


Some things are easy to take for granted. Like the Sun showing up every day to cast light and heat in our direction. The ancient Egyptians, who worshipped the Sun as the source of all being, intuitively understood that life on Earth depends absolutely on that daily dose of energy.

In some ways, scientists like University of Minnesota physicists Paul Woodward and David Porter are the modern version of ancient Egyptians. They devote their intellectual energy to unlocking secrets of how the Sun works. To be sure, modern science knows more than Pharaoh's priests, yet with the Sun — as with many things in Nature — the more we know, the more new puzzles reveal themselves. How long will it be, for instance, till the Sun burns out? Current estimates range between five and six billion years, and the range reflects uncertainty in our knowledge of even this most familiar star.

Woodward and Porter direct their research, part of a National Science Foundation "grand challenge" project in geophysical and astrophysical turbulence, at what happens in a layer of gas — roughly the outer third of the Sun — called "the convection zone." In the fall of 1994, using the CRAY T3D at Pittsburgh, they carried out the largest simulation of compressible convection ever done. "We're trying to find out in some detail," says Woodward, "what's going on beneath the surface, where you can't see. This is the first calculation in three dimensions where we had enough resolution to treat the small-scale structure of turbulence in this region with a reasonable degree of accuracy."

Visualizations from this computation show features of solar turbulence the researchers didn't expect to see, and they show other features predicted by theory that simulations until now lacked sufficient resolution to test. "This is the first time," says Woodward, "that we've had results we could really sit down and compare with the predictions of simplified models for this kind of convection in stars."


Paul Woodward (right)  
and David Porter,  
University of Minnesota.

**The Sun's Turbulent Outer Layer**

An average star in every respect, the Sun is one of 100 billion in our galaxy. Astronomers classify it as a "yellow dwarf," and no dramatic changes are expected for another five billion years or so, when in the last stage of its life, the Sun will transform to a "red giant," expanding into space and gobbling up Mercury, Venus and probably Earth.

At its dense core — which at about 14 million degrees Kelvin is the furnace for every living room on Earth — hydrogen atoms fuse to form helium, converting four million tons of matter into energy every second, which radiates outward as light rays. The dense, ionized gas becomes progressively cooler closer to the surface, until at about 30,000 degrees, in the convection zone, electrons join up with protons to form atoms. Here the gas begins to move with a convective motion — essentially the same as heat circulating in a radiator-heated room. Rising streams of hot gas carry heat toward the Sun's surface where, at about 5,000 degrees, cooler gas sinks back toward lower reaches of the convective layer. This roiling motion is highly turbulent, and it affects the rate at which energy leaves the Sun, radiating outward into space.

Understanding these flows is important for understanding the evolution of stars, and it has a direct bearing on predictions of solar lifetime. Violent turbulence in the convection zone also affects the periodic flares of gas that shoot out into space, releasing great bursts of energy that can disrupt electronic systems on Earth.


**Vorticity in the Convection Zone**

*This rectangular slab is a volume rendering showing a side view of the solar convection zone, roughly the outer third of the sun. It is idealized, with hard walls at the top and bottom. Energy is added from the bottom, to model radiation from nuclear fusion in the Sun's core. Colored fields represent "vorticity," how strongly the gas is spinning. Black areas are weak; green is slightly stronger, and white the strongest.*

*The knotted, densely packed vortex tubes, explains David Porter, show vigorously turbulent regions, especially along the top boundary and in the large downflow lane near the right edge. Vertical vortex tubes at the lower center resemble Earth tornados. Here the flow converges in both horizontal directions and expands upward.*

"Give me the splendid silent sun with  
all his beams full-dazzling!"

— Walt Whitman




### Results: Smoke Rings in the Sun

Because these processes happen beneath the Sun's visible surface — the photosphere — they can't be seen from Earth, and it's very difficult to conduct experiments. Computer simulations provide a way to see what's otherwise unseeable. The difficulty has to do with the complexity of the flow. Along with general features that occur over relatively large scales, such as the granulation pattern on the Sun's surface, the researchers need to observe fine details.

"You want to see the small vortices and how they interact," says Woodward. "Very fast time scales and small length scales are involved in this problem." To approach realistic results, the researchers must allow the gas to be *compressible* — its density changes as pressure and temperature change, which increases the computing requirements. The modeling must also replicate large pressure differentials between the bottom and top of the convection zone. With the CRAY T3D, it became possible to more realistically model this depth differential.

In October 1994, PSC made two weeks of dedicated time available. Woodward and Porter are still analyzing the pile of data (250 gigabytes) that resulted, a process that could take a year or longer. Nevertheless, scientific visualization has already revealed phenomena not seen before.

In general, the researchers see more turbulence than was apparent in previous simulations. Representative of this are vortex rings, which in the visualizations look like smoke rings, moving up and down. The visualizations also reveal a hierarchy of dynamics. "You have small convection cells at the surface," explains Porter, "and they're embedded in a large convective cell that spans the depth of the layer. So we're seeing the



### Temperature in the Convection Zone

This view looking down at the surface of the convection zone (above) shows temperature according to color — aqua (cold) through blue, green and red to yellow (hot). A cellular network of cool downflow regions surrounds isolated, warm upflows in a pattern that resembles "solar granules" observed on the Sun's surface.

The second slice (left) is from deeper in the convection zone (one pressure scale height below the surface). Here, the cool downflow lanes have merged into fewer, larger cells. Effects of turbulence can be seen in the knotted downflow structure. Some downflow lanes, crushed into isolate downflowing plumes, are seen as blue dots.

hierarchy of convection theorists had predicted, but only now is numerical experiment capable of resolving and testing these theories."

"This is the first three-dimensional calculation of compressible convection we've done," says Woodward, "where we have real confidence in the results." From here, the researchers plan to add more realism — a stable layer beneath the convection zone and a free surface on top, like taking the lid off a boiling pot. "It's a scientific philosophy of our grand challenge team that we don't want to take a step towards a more complex formulation of the problem, a more realistic model, until we've understood the simplified one. We're ready now to take that step." (MS)

#### References:

- D. H. Porter & P. R. Woodward, "Effects of Turbulence on a Compressible Convection Zone," (in preparation, 1995).
- D. H. Porter & P. R. Woodward, "High Resolution Simulations of Compressible Convection using the Piecewise-Parabolic Method," *Astrophysics Journal Supp.* 93, 309-49 (July 1994).

This research is supported by the National Science Foundation.

**Simulated Glue for DNA & Proteins**

Sometimes proteins are like the Wicked Witch. Add water, wait a nanosecond and you can almost hear that piteous whimper. "I'm mel-l-l-ting!"

Well, not really. But there may be an audible groan from biochemists as they observe simulations of a large protein or DNA in water. For a few hundred picoseconds of simulation time it looks fine, then the molecule appears to unravel. It's been a troubling problem, but Lee Pedersen and Tom Darden now seem to have it fixed.


The problem has been molecular dynamics (MD). As the name suggests, MD is a way to simulate and predict how molecular structure changes over time. Over the past 15 years, MD has evolved — along with increasing computing power — to become an important part of the molecular biology toolkit. The colorful protein structures in biology textbooks are analogous to butterflies mounted in a museum showcase. They represent molecules painstakingly removed from cells and crystallized so that X-ray crystallography can reveal their three-dimensional structure. With these structures as starting data, MD puts proteins and DNA back into a cell-like

environment. The most realistic studies include surrounding water molecules and ions, thus replicating many of the structural forces acting in the living environment.

Nevertheless, trying to account for the dynamic interactions among not only the atoms of the molecule itself but also thousands of water molecules is an extremely complicated computational task. As computing power has increased, it has become feasible and desirable to track how the structure changes over as long as several nanoseconds — less than a fast eyeblink in human time,

but as good as a lifetime in protein biochemistry.

Such simulations can take hundreds of computing hours, and the outcome can be disappointing.



Lee Pedersen (left), University of North Carolina, and Tom Darden, National Institute of Environmental Health Science.

A few years ago, Pedersen and Darden encountered protein melting face-to-face. "The protein we were simulating would literally shake itself apart in a couple hundred picoseconds," says Darden, a biomathematician at the National Institute of Environmental Health Science. "This melting behavior is even more pronounced for DNA. Generally, the longer you run, the worse the situation."

With simulations at the Pittsburgh Supercomputing Center, the researchers diagnosed the malady — simulating *electrostatics*: the attraction-repulsion forces between atoms that aren't bonded to each other, and Darden devised a cure. He came up with a new method — "particle mesh Ewald" — that is fast and therefore has the added advantage of making it feasible to study larger structures. "There's two significant outcomes," says Pedersen, a physical chemist at the University of North Carolina. "This provides the accuracy missing in previous calculations, and it makes bigger systems possible."

**Zeroing in on the Cutoff Radius**


Treating electrostatic forces accurately has become a problem for MD because of the sheer size of the problems researchers want to attack. Simulations of large proteins or DNA with water involve 10,000 or more atoms, making it infeasible to explicitly calculate each atom's interaction with every other. As a reasonable approximation that makes the computation possible, researchers traditionally use a "cutoff radius" — usually around 10 angstroms — beyond which electrostatic forces (which decrease with distance) have been considered insignificant.

---

*"This provides the accuracy missing in previous calculations, and it makes bigger systems possible."*



---

This works fairly well in some cases. In others it doesn't. Pedersen and Darden started running into the melting problem when they simulated DNA and a large protein involved in cancer (*H-ras* p21). Both are highly charged molecules, which suggested electrostatics could be the problem. "Tom and I," says Pedersen, "were doing simulations that five years ago people would have said were all right, but we were extending them to longer times. We really wanted the answer, yet we could see before our eyes that something was fundamentally wrong." Other researchers were running into the same problem, and many began to conclude it was an inherent weakness of MD.



Graphical depiction of bovine pancreatic trypsin inhibitor, the protein used for testing the particle mesh Ewald algorithm developed by Tom Darden and Lee Pedersen. Colors indicate oxygen atoms (red), nitrogen (blue), hydrogen (cyan), carbon (white) and sulfur (yellow).

The "backbone" of bovine pancreatic trypsin inhibitor coded according to secondary structure: Random coil (wheat) is tube shaped; helices (aqua) are flat; and sheets (magenta) are arrows. Dots show the solvent accessible molecular surface.



Mike Crowley, PSC biomedical scientist. Working with Darden, Crowley has developed an efficient, parallelized routine for three-dimensional FFT, and PME simulations are now underway on the CRAY T3D.

Pedersen and Darden began to zero in on a solution when one of their students, Darrin York, spent a summer working at the Pittsburgh Supercomputing Center. York and PSC biomedical scientist David Deerfield were running simulations that included magnesium and sodium ions, called "counterions" because their positive charges neutralize the negative charge of DNA. Pedersen visited to check on progress: "Darrin kept saying look at this. It was pretty apparent what was happening."

The ions separated from each other at exactly the distance of the cutoff radius. Because they aren't bonded to other atoms, counterions have greater freedom to move around in the water: "They respond to long range electrostatic forces," explains Darden, "and with a cutoff radius they never reach natural equilibrium. When they come closer than the cutoff, they repel each other, and as soon as they're outside, they don't feel a thing. The end result is they hang out right at the cutoff." With this strong clue that cutoff radius was the culprit, Darden set about finding a better way.

### Particle Mesh Ewald

To eliminate the inaccuracy introduced by a cutoff radius, Darden turned back to a much cited 1921 paper by Paul Ewald, who helped pioneer X-ray crystallography. Darden's central insight was to see how "Ewald summation" — a method for summing up the collection of charges in a crystalline lattice — could be combined with a very fast computational method — fast Fourier transforms (FFT). "We're the first people to do macromolecular Ewald summation," says Darden, "because no one had a way to do it fast."

To make Ewald summation workable for the irregular charge distribution typical of MD, Darden turned to plasma physics for the "particle mesh method." As a check on the accuracy of this innovative synthesis, dubbed "particle mesh Ewald" (PME), Darden, York and Pedersen ran a one nanosecond simulation of a large protein (bovine pancreatic trypsin inhibitor) and compared results to the crystal structure. The deviation was lower than observed between different crystallized forms of the protein.

"We can now run long simulations, and we don't get the melting," says Darden. "We get high accuracy. Furthermore, simulations with PME are less expensive than with a 10 angstrom cutoff." Using PME, Darden and Pedersen have now completed the simulations of DNA and the H-ras p21 protein that previously stymied them. Because PME combines accuracy and fast performance, Darden and Pedersen see it as bridging the gap between crystal structures and MD's ability to simulate structural changes inside the cell. (MS)

### References:

- Darrin M. York, Alexander Wlodawer, Lee G. Pedersen & Tom A. Darden, "Atomic-level accuracy in simulations of large protein crystals," *Proceedings of the National Academy of Sciences* 91, 8715-18 (1994).
- Darrin M. York, Tom A. Darden & Lee G. Pedersen, "The effect of long-range electrostatic interactions in simulations of macromolecular crystals: A comparison of the Ewald and truncated list methods," *Journal of Chemical Physics* 99, 8345-48 (1993).

This research was supported by the National Cancer Institute, Cray Research, Inc. and the National Science Foundation.

Photo of Lee Pedersen & Tom Darden: David Deerfield, PSC.

Computer graphics: David Deerfield & Amanda Holland-Minkley, PSC.

**Proteins at the Movies**

When you open a biochemistry textbook, you'll see dozens of molecular structures frozen in the pages. These images have immensely advanced our understanding of proteins. Nevertheless, for most biological molecules, static poses give an incomplete, potentially misleading picture, because the molecules within our liquid innards are constantly changing shape.

*"A driving force for this research is understanding the structural features that could make it possible to impede viral replication."*

These changes in shape, says Thomas James, a pharmaceutical chemist at the University of California at San Francisco, may drive certain chemical liaisons. A twisting, turning protein may bind only with another protein that recognizes similar shape-changing dynamics. In the case of environmental toxins, for instance, this insight may help researchers develop artificial proteins that trap them. And in the billion-dollar pharmaceutical industry, researchers always are striving to design drugs that bind selectively with proteins involved in disease processes.

"If we're going to put in this kind of effort," James says, "it would be nice to target a realistic picture of what the molecule is doing and what it looks like. My goal is that in a few years if I want to show you a molecule's structure, instead of a snapshot or slide, I'll show you a video."

*The structure of this 13-basepair sequence of DNA from the HIV virus was determined by*

*Thomas James and colleagues using NMR methods and computing at Pittsburgh Supercomputing Center.*

*The spiral staircase-like double-helix structure forms from basepairs (red) that link the two DNA "backbones" (yellow).*


*Aligning DNA sequences from several HIV strains showed that this sequence, from the "long terminal repeat region" of the HIV-1 genome, is "highly conserved" — i.e., it remains the same in different strains. This suggests that the sequence plays a role in HIV's ability to reproduce itself in human cells, and the structure is a potential target for antiviral drugs.*

Thomas L. James,  
University of California,  
San Francisco.

James is a leader in using nuclear magnetic resonance (NMR) methods to determine molecular structure. In recent work at the Pittsburgh Supercomputing Center, he has examined the dynamic structure of several proteins, including a regulatory protein that activates HIV, the virus that causes AIDS. James and his colleagues have also developed a promising new computational approach to the analysis of NMR structural data.

**NMR and Supercomputing**

For the last 15 years, James and other researchers who wanted to understand how molecules move have been able to glean some understanding using NMR, which bombards atoms with radio waves when they're in magnetic fields, creating a spectrum that acts like a fingerprint. These spectra, explains James, make it possible to decipher molecular structures in solution, and they can potentially give insight into dynamic aspects of those structures. "We were able to talk to a certain extent about frequencies and amplitudes of motion," James says, "but we have not been capable of saying anything very meaningful. For example, I may know that this particular bond moves 60 degrees and that it fluctuates back and forth in one nanosecond. But, in fact, by my saying that, I still didn't know exactly what the motion was."



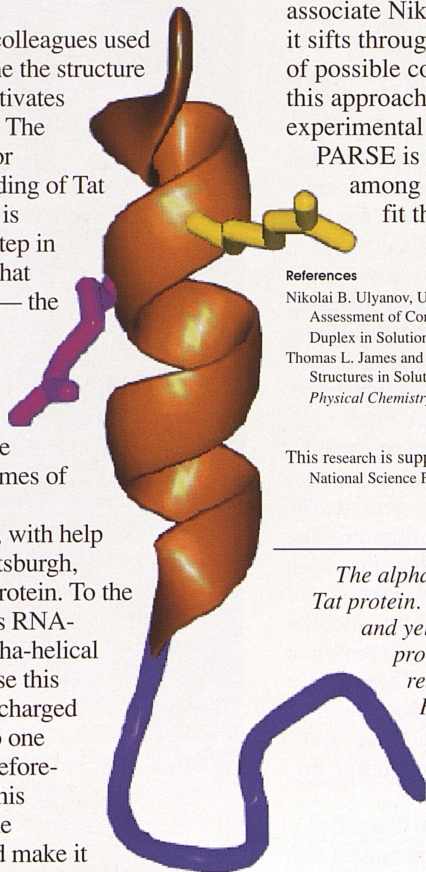
Nowadays, James is developing computational techniques that incorporate experimental structure and dynamics results from NMR, which essentially supplies details about distances between hydrogen atoms and torsion angles between bonds, as well as previously available information such as length and angles between bonds. "If we use that available information, plus experimental information, we should be able to conduct an intelligent search of the millions of possible conformations to find out which ones are consistent with experimental data. So what we'll end up with is not just a single structure, but a family of structures that details the molecule's movement in solution."

Using Pittsburgh's C90, James and his colleagues have performed their calculations with software packages they developed. Known as MARDIGRAS and CORMA, they determine distances between atoms based on NMR data. And to determine how the atoms move and also their energy levels, the researchers use the AMBER software package developed by colleague Peter Kollman's group at the university.

### Attacking AIDS

Recently, James and his colleagues used their techniques to determine the structure of a protein segment that activates HIV to produce more virus. The protein is known as "Tat" for transactivating protein. Binding of Tat protein to RNA in the virus is believed to be a necessary step in recruiting cellular proteins that promote viral transcription — the process by which the viral DNA is "transcribed" to create matching strands of RNA that, in turn, guide the process that produces the structural proteins and enzymes of a new virus.

James and his colleagues, with help from supercomputing at Pittsburgh, derived a structure for Tat protein. To the surprise of researchers, Tat's RNA-binding region forms an alpha-helical structure in solution. Because this region has many positively charged amino acids, says James, no one would have predicted this beforehand. "A driving force for this research is understanding the structural features that could make it possible to block Tat function and consequently impede viral replication. We're now using the Tat protein structure to try to find a molecule that will bind specifically with this structure."



### A New Procedure: PARSE

"The ability to determine dynamic structures is not a solved problem yet," says James. "The obstacles, in some sense, are that if it's so difficult to determine the average structure for a molecule in solution, then how could I figure out the many interconverting structures for that same molecule in solution. If a molecule assumes one structure part of the time and another structure part of the time, how do I figure out which ones these are? And the spectral parameters used to determine structure are also influenced by how fast the exchange between structures occurs."

Conventional computational techniques, for instance, do not specify the minimum number of conformations — different structural configurations — possible and the probability that each one exists. Rather, these techniques average together the many conformations and develop an "average" structure that may never exist.


A solution, says James, is a computational approach developed by James' research associate Nikolai Ulyanov. Known as PARSE, it sifts through NMR data, creating hundreds of possible conformations. The viability of this approach depends on high quality experimental data, but with quality input PARSE is capable of determining which among many possible structures best fit the data. (SE)

#### References

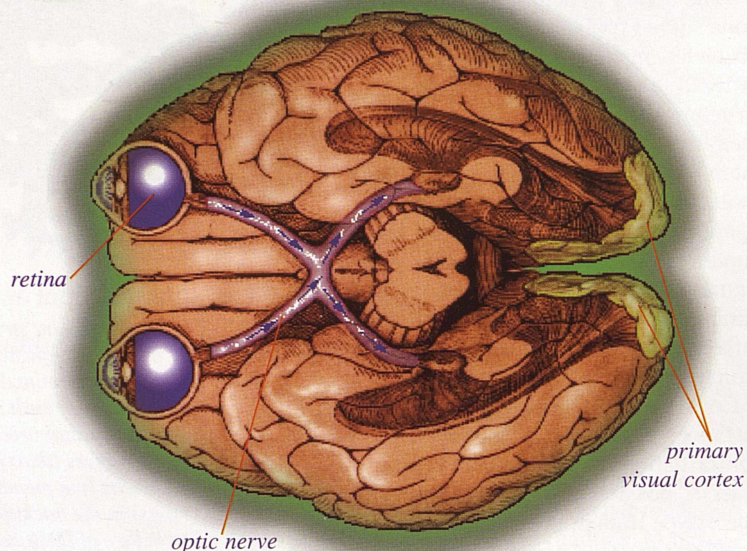
- Nikolai B. Ulyanov, Uli Schmitz, Anil Kumar and Thomas L. James, "Probability Assessment of Conformational Ensembles: Sugar Repuckering in a DNA Duplex in Solution," *Biophysical Journal* **68**, 13-24 (1995).  
 Thomas L. James and Vladimir J. Basus, "Generation of High-Resolution Protein Structures in Solution from Multidimensional NMR," *Annual Review of Physical Chemistry* **42**, 501-42 (1991).

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

*The alpha-helical backbone segment (gold) of Tat protein. Two amino-acid sidechains (magenta and yellow) indicate the site at which the protein interacts with the transactivation response (TAR) sequence of RNA from HIV-1. This interaction plays a vital role in transcription of the HIV virus.*



*Nine different structures of the same sequence of DNA as determined by PARSE. This result shows that the spatial orientation of the sugars (deoxyribose, shown as five-membered rings) in the backbone structure of DNA can shift under physiological conditions from the typical S-conformation (blue) to N-conformation (green). This could affect interactions between DNA and proteins such as RNA-polymerase.*



### The Shifting View of Vision

There's more to seeing than meets the eye. Quite a lot more actually. What meets the eye is light waves. Their impact on the retina triggers impulses that travel a pathway along the optic nerve through the thalamus to the back of the brain, where they register on the *primary visual cortex*. From there, via myriad connections, neuron to neuron, visual impulses reach out to other parts of the brain, and we get an instantaneous mental image of a colorful, three-dimensional world. How does it happen?

The human brain, say scientists who study it, may be the most complex structure in the universe. A convoluted mass of tissue made of a 100 billion interconnected neurons spawns a mental world — all the processes we call thought, consciousness, creativity, emotion, memory and vision.

Probably the most studied of these is vision, and contrary to long accepted views that brain structure is genetically determined, research since the mid-70s has shown that seeing is a learning process. When a kitten sees only horizontal lines for the first few months of its life, for example, it loses the ability to see vertical lines. These and other experiments suggest that the interconnections between neurons aren't fixed at birth, but evolve depending on visual experience.

"Experiments have shown that much of the structure thought to be innate or hard-wired actually develops during the early days or weeks of life," says Risto Miikkulainen, a computer scientist at the University of Texas, "and even in the adult these structures can reorganize." Using the CRAY T3D at Pittsburgh Supercomputing Center,

Miikkulainen and Ph.D. candidate Joseph Sirosh have developed the first computational model to comprehensively mimic this "self-organizing" process.

"We couldn't have done this 10 years ago," says Sirosh, "because the technological capability to simulate these large neural networks didn't exist. Machines like the T3D are absolutely crucial to this work."

### The Visual Cortex and Self-Organization

Nobel prize winning studies by Hubel and Wiesel in the 1960s showed that the primary visual cortex is organized in relation to features in the visual field of the retinas. Neurons that respond to vertical lines — trees or telephone poles in the visual field, for instance — are grouped together in patches called *orientation columns*. Other orientation columns respond to horizontal lines, 45° lines, etc. Other neuron groupings form according to whether they respond predominantly to the right or left eye.



*Orientation columns in a patch of monkey visual cortex. Colors from red to violet indicate the orientation preference of the neurons, from 0 to 180 degrees. Overlaid white lines show the orientation preference of each area. (From research by Blasdel, 1992.)*

Since the mid-70s, experiments have demonstrated that these patterns aren't genetically fixed, but develop from visual experience. "An organism," explains Miikkulainen, "seems to store a set of learning rules, genetically programmed procedures expressed in the behavior of each neuron, so that information from the outside world organizes the connections from neuron to neuron. Because the organizing mechanism is unsupervised and local to each neuron, it is referred to as self-organization."


Experiments have shown that connections among neurons are initially more or less randomly scattered. With visual experience, the neurons appear to divide their labor, in effect saying, "You take this kind of input. I'll take that kind of input." "Neurons that reinforce each other," says Miikkulainen, "wire together, and unnecessary connections are eliminated. Such filtering eliminates redundancies and allows the brain to efficiently process the massive amounts of visual information impinging on the retinas. This trimming down process is what we're modeling."

Risto Miikkulainen  
and Joseph Sirosh,  
University of Texas.



## Receptive Field: Inputs and Lateral Interaction

Simulated spots of light on the retinas (two connected panels) activate the model. Initially, the neurons respond in a diffuse pattern (left). The activation spreads through the lateral connections, and the neurons settle into a focused patch with sharp contrast between areas of high and low activity.



## Fire Together, Wire Together

Miikkulainen and Sirosh's essential insight has been to adapt a learning mechanism first proposed in 1949, called Hebbian learning, to self-organization in the visual cortex. The Hebbian mechanism adjusts connections according to the idea that neurons active at the same time increase the strength of their connection. "You can think of it as 'Neurons that fire together, wire together,'" says Sirosh. "As a result of this process, connections that are required are retained and others are eliminated."

---

*"The model allows us to predict how lateral connections organize in the visual cortex and what visual knowledge is stored in them."*

They implemented their Hebbian algorithm in a network model they call LISSOM (Laterally Interconnected Synergetically Self-Organizing Map). LISSOM models a sheet of neurons in the primary visual cortex — typically 65,000 neurons with 700 million interconnections. Each neuron connects to "receptors" in two fields, representing the right and left retina, and to other neurons in the network. LISSOM is the first model of the visual cortex to simulate self-organization of these neuron-to-neuron or "lateral" connections.


To date, LISSOM — which performs at a sustained speed of eight billion calculations a second on 256 T3D processors — has shown a promising ability to reproduce observed patterns of neural organization. What is most interesting is that a single learning mechanism can reproduce various features of self-organization.

"The model allows us to predict how lateral connections organize in the visual cortex," says Miikkulainen, "and what visual knowledge is stored in them." In the brain, this information is preliminary to processes that organize distinct visual features into coherent wholes. LISSOM's success at modeling self-organization suggests that neural modeling will play a vital role as vision research presses on to explore these higher level processes. (MS,SE)

### References:

- Joseph Sirosh & Risto Miikkulainen, "Ocular Dominance and Activation Dynamics in a Unified Self-Organizing Model of the Visual Cortex," *Advances in Neural Information Processing* 7, 1995 (in press).
- Joseph Sirosh & Risto Miikkulainen, "Topographic Receptive Fields and Patterned Lateral Interaction in a Self-Organizing Model of The Primary Visual Cortex," *Neural Computation*, 1995 (in press).

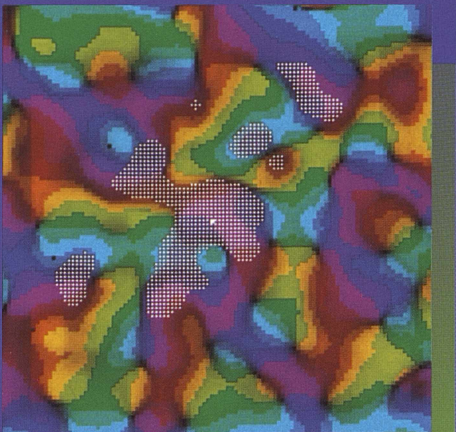
This research is supported by the National Science Foundation.




## Self-Organization of the Orientation Map

Using oriented light spots as input, LISSOM modeled how orientation preference and neuron-to-neuron connection patterns develop in the cortex. The color of each neuron in the network (192 x 192 neurons), from red to magenta to blue to green, corresponds to its orientation preference from zero to 180 degrees. The small white dots show lateral connections of the neuron marked with a big white dot.

Initially (above) the orientation preferences were random and lateral connections covered almost the entire map. After a series of several thousand inputs, the neurons organize (below) into orientation columns, and lateral connections link areas of similar orientation preference. These patterns agree in key features with maps obtained by experimental imaging.





## Respiration and the Lungs

# Breathing Lessons


### A Better Way to Take Your Medicine

A cigarette smoker takes a deep drag and holds it, waiting to feel the tingle of nicotine hitting the bloodstream. It's not a pretty picture if we care about our health. Nevertheless, says Jeffrey Hammersley, M.D., the smoker is on to something: "The airway is an excellent way to administer drugs. You get rapid transmission of vapors and small particles deep into the lung and almost instant absorption into the bloodstream."

Take cardiac arrhythmia, for instance. Many people with dangerously irregular heartbeats carry an "autoinjector" of antiarrhythmic medication, typically lidocaine, for jabbing into their thigh when their heart acts up. In three to four minutes, the heart usually calms down. For some, however, three to four minutes may be too late. In the future, Hammersley and his colleagues expect that inhalers will speed lidocaine from the lungs to the heart in 10-15 seconds.

Other potential inhaled medications include insulin for diabetics and antibiotics to treat chronic infections that now require injections. When the lung is the target of medication, inhaling requires less total medication and less frequent doses, reducing side effects and cost, and it should be possible to extend these benefits to other drugs. "If you could also deliver antibiotics or insulin through the airways," says Hammersley, "you have the potential to eliminate the risk and discomfort of needles and immensely reduce treatment costs for a wide range of diseases."

Hammersley is a lung specialist at the Medical College of Ohio, and he combines his medical knowledge with engineering training in fluid dynamics. Using the CRAY C90 at PSC, he has begun to create a computational model of airflow in the lungs. He collaborates with colleagues in engineering and computer science at the University of Arkansas, the NSF Engineering Research Center at Mississippi State and the University of Toledo. Together they have converted computational fluid dynamics (CFD) techniques used in the design of cars, airplanes and aerospace vehicles for use in the complex branching geometry of the lung's small airways.



Jeffrey Hammersley, M.D.,  
Medical College of Ohio,  
Toledo, Ohio.


### How Little We Know

What we know about how air swirls and weaves its way through the lung's intricate passages is very limited. The approach to delivering drugs to the lung for asthma and cystic fibrosis, for example, says Hammersley, has been largely trial and error.

"We know drugs reach the main airways beyond the voicebox, but most particles are knocked out by turbulence and mixing in the central airways." Even in normal subjects, he adds, drugs often fail to reach the outlying, small airways where the airflow obstruction in asthma is seated. To get around these limitations, lab researchers try various particle sizes, and they alter particle properties and evaluate different propellants. "Nevertheless," he says, "all you can really do is toss it in and see what happens."


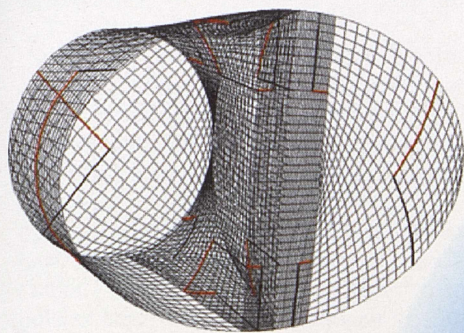
The problem is the complexity of the lung's structure and how this affects airflow. Most pulmonary medicine textbooks describe airflow as turbulent down through the windpipe into the two main bronchial tubes, and as a smooth, symmetrical flow through the smaller airways beyond. Research on scale models (much of it carried out by Hammersley's mentor and MCO colleague Dan Olson), shows that this simplistic view is deeply flawed.

Flows in the small airways are strongly affected by the 16 or more generations of branching that extend out to the air sacs. This structure is highly asymmetric; branches go off from the parent airway at different angles and extend for different lengths as the airways narrow from several centimeters to less than a millimeter at the end. As a result, airflows are irregular and skewed with difficult to predict effects on the delivery of inhaled aerosols.



These graphical results from 3-D lung airflow modeling on the CRAY C90 depict flow velocity at selected cross-sections in a single bifurcation. Velocity (in dimensionless units) ranges from zero (black) through dark blue to red to maximum (white).

*In lung airways, the region where a parent branch splits into two daughter passageways, called the "carina," is the most difficult region to model. It requires generating grids that accurately reproduce the "central flow splitter" and wall geometry.*



---

*"You have the potential to eliminate the risk and discomfort of needles and immensely reduce treatment costs for a wide range of diseases."*

#### **Computational Fluid Dynamics & the Lungs**

Because of the complexity of this branching structure, physical scale-model research on lung airflows is limited to about three generations of branching. The small airways deep within the lung, furthermore, are inaccessible to most experimental techniques. The group's goal is to use CFD tools to simulate flows in these inaccessible regions. With accurate computer modeling, it should be possible to develop the detailed understanding of lung flow patterns and particle transport that can make new aerosol medications a reality.

"If we want to deliver medications to the spot where the lung doesn't want them to go," says Hammersley, "we have to be more sophisticated. We must understand the transport process and learn to redirect the particles — by modifying their speed, density or their electrical charge. With computer modeling, we can provide information that will avoid inappropriate clinical trials and wasted effort."

Using a grid generation program, EAGLE, developed by the Department of Defense, the group has produced grids that mimic the complexly curved walls and branching of the small airways. This has been a daunting task in that EAGLE was designed to simulate flows over high speed aircraft, which have a much more regular geometry. "With computations at PSC, we've shown not only that we can accurately reproduce airway geometry, but that it's possible to extend these grids and flow solutions to virtually any tubular biological structure."



The researchers used two different flow solver codes to simulate flows over this three-dimensional grid — NASA's INS3D and a new code developed by collaborators at NSF's Engineering Research Center for Computational Fluid Simulations. Because each flow-splitting branch requires between 125,000 and 450,000 grid points, these are very demanding calculations. The results to date, carried out to three generations of branching, compare well with flow measurements made in scale models.

Future objectives are to include flexible airway walls and to simulate how particles of different sizes are deposited. "We want to be sure the simulations match up well with experiment," says Hammersley, "then we can take the next step, to give us answers we can't get experimentally." (MS, SE)

## A Universe in Ferment

Cosmology. The branch of science that takes as its subject, well, everything — the universe. The seemingly infinite reaches of space, the expanse of time — past, present, future. Cosmology asks how and when did all this come into being? What structure does it have? How is it evolving?

Cosmologists have plenty to think about lately. Although understanding of the visible universe has progressed by leaps and bounds in the last 20 to 30 years, the pace of new information — from sensitive electronics harnessed to Earth telescopes, sophisticated satellites and, most recently, the Hubble Space Telescope — shows no sign of slowing down. Cosmology is in the news, with theorists and observers contending to make sense of it all.


To evaluate the new information, to see how it fits into the general framework of the “big bang” model of cosmic origin, cosmologists increasingly turn to supercomputers. Edmund Bertschinger, a theoretical astrophysicist at MIT, is one among a group of physicists and computer scientists who in 1993 formed a nationwide collaboration, the Grand Challenge Cosmology Consortium (GC<sup>3</sup>). Funded by NSF through the

federal High-Performance Computing and Communications program, GC<sup>3</sup> is a concerted effort to attack cosmological questions with computational modeling — in particular, to exploit the capability of massively parallel systems such as the CRAY T3D at Pittsburgh.

Bertschinger has focused on one of the most fundamental questions: the origin of large-scale structure in the universe. Detailed observations since the late 1980s show that galaxies are not evenly distributed in space, as previously believed, but clustered, with areas of high and low density. Through a two-pronged

approach, modeling microwave radiation from space and modeling the evolution of matter itself,

Bertschinger looks at how small



Edmund Bertschinger,  
Massachusetts Institute  
of Technology.

irregularities at the start of the big bang could have led to the universe we see around us now. “Somehow the universe started with this tremendous, hot fireball,” says Bertschinger. “The early universe was quite smooth. How did it get to be so lumpy today?”

## Afterglow of the Big Bang

One of Bertschinger’s two approaches involves modeling the “glow” of the early universe. The cosmic microwave background, discovered in 1965 by Arno Penzias and Robert Wilson, who subsequently received a Nobel prize, is thermal radiation from the big bang, still traveling through space 15 or so billion years later. “Imagine that our eyes could detect microwave radiation,” explains Bertschinger, “like that produced by microwave ovens. The night sky would be luminous.” For cosmologists, this afterglow is a fossil of the universe in its infancy.

This microwave radiation is almost perfectly uniform from all directions, as predicted by the big bang theory. In 1992, however, using NASA’s Cosmic Background Explorer satellite, scientists found minute fluctuations. This was the first direct evidence of irregularities in early stages of the universe that could account for how matter clumped into planets, stars, galaxies and the clustered chains of galaxies found in the past decade.


With the CRAY C90 and, more recently, the CRAY T3D, Bertschinger has carried out what he believes are the most accurate calculations yet that track the consequences of these microwave fluctuations. Using code he developed, called LINGER (for linear general relativity), his basic approach is to apply a random fluctuation to the first second of the big bang and evolve forward in time. The results are then compared with observed data.

## LINGER Zooms on the C90

In May 1994, Bertschinger’s cosmic microwave background evolution code, LINGER, using all 16 processors of the CRAY C90, ran at a sustained speed of 8.6 billion calculations a second. This is the best performance achieved to date on PSC’s C90.

The numerical approach Bertschinger uses in this modeling is analogous to an electronic synthesizer sampling music from the universe. “Basically, what we’re doing is decomposing acoustic waves into their different frequency components.” Each harmonic can then be computed independently, an approach that fits extremely well with parallel processing. The speedup from T3D scalability, going from 32 to 64 or 128 processors, says Bertschinger, is nearly linear, better than 99.9 percent. “The T3D has been a wonderful machine for this project. We get turnaround in a couple of days that would take months on a single Alpha chip workstation.”

“We had the sky, up there, all speckled with stars, and we used to lay on our backs and look up at them, and discuss about whether they was made, or only just happened.”  
— Mark Twain, *Huckleberry Finn*



To evaluate new information, to see how it fits with the "big bang," cosmologists increasingly turn to supercomputers.

### Dark Matter Evolution

Both arms of Bertschinger's GC<sup>3</sup> work bear on what is probably the key question in 1990s cosmology: dark matter. Galaxy mapping and other measurements over the last two decades have raised a perplexing problem. Luminous matter — the visible stars and planets — doesn't have enough mass, hence enough gravity, to prevent galaxies from flying apart. This leads cosmologists of many stripes to conclude that most of the universe, probably more than 90 percent, is unseen — dark matter.

Various theories have been proposed, suggesting different kinds of not yet discovered particles pervading the universe. Until these particles are detected in laboratories, the only way to test the theories is cosmological modeling. The basic idea is to posit an initial configuration of matter and let the big bang happen until, billions of years later in the computer, galaxies form. One can then compare the structure produced by one theory against others and with reality.


The computational challenge is the incredible range of scales between galaxies, about 10,000 light years across, and the visible universe, at least 10 billion light years. One approach is to model a volume smaller than the observable universe, which can still give potentially useful results. Taking this path, Bertschinger's dark matter simulations on the Connection Machine, CM-5, at the National Center for Supercomputing Applications evolved 16 million particles starting less than a hundred million years after the big bang. The results show that the character of structure formed — whether dense clumps or filament-like chains — depends on the nature of the initial fluctuations. This modeling also suggests that a purely "cold dark matter" scenario — one of various dark matter theories proposed — forms too much structure in comparison with the real universe.

"The T3D opens up new possibilities," says Bertschinger, who is working to port his dark matter evolution code to Pittsburgh's T3D. "The long-term goal of our project, besides trying to answer these questions, is developing a portable, scalable parallel cosmology evolution code that includes everything from soup to nuts, from microwave background to galaxy formation and particle and gas dynamics. We're making substantial progress." (MS)


#### References:

[http://zeus.ncsa.uiuc.edu:8080/GC3\\_Home\\_Page.html](http://zeus.ncsa.uiuc.edu:8080/GC3_Home_Page.html)  
 E. Bertschinger, "Cosmic Structure Formation," *Physica D* 77, 354-79 (1994).  
 C. P. Ma & E. Bertschinger, "Do Galactic Systems Form Too Late in Cold+Hot Dark Matter Models?" *Astrophysics Journal Letters* 434, L5-L9 (1994).

This research is supported by the National Science Foundation.



No, they're not Easter Eggs. They're the infant universe hatching from its structureless shell. The COBE science working group produced the blue and pink all-sky map from data collected by the Cosmic Background Explorer (COBE) satellite. It depicts minute temperature fluctuations, ranging from 0.00015° Kelvin colder (blue) than the background to 0.00015° K. warmer (pink). The second map represents Edmund Bertschinger's simulations on the CRAY T3D at Pittsburgh. This map, at higher angular resolution (0.5 degrees) than COBE could detect (7 degrees), shows negative (blue) and positive (red) fluctuations of 0.0002° K. The simulation assumed a mixed hot and cold dark matter model with 5 eV neutrino mass.



### The Grand Challenge Cosmology Consortium

Jeremiah P. Ostriker,  
Princeton University

Edmund Bertschinger,  
MIT


Dennis Gannon,  
Indiana University

Lars Hernquist,  
U. of California, Santa Cruz

Michael Norman,  
U. of Illinois at Urbana-Champaign

Ralph Roskies,  
Pittsburgh Supercomputing Center

David Spergel,  
Princeton University



PSC scientist Joel Welling (left) and University of Pittsburgh computer scientist Rami Melhem are working with the Grand Challenge Cosmology Consortium to develop cosmology modeling and visualization software.



*"Materials are going to be a major force in the American economy, and materials processing research is important to lay the groundwork."*

### The Deborah Number Dilemma

Research groups around the world are working to determine which theoretical models describing polymer flows most closely match experimental findings. Over the past 15 to 20 years, dozens of models have been proposed, but they could be tested only under limited conditions. "As a consequence," says Leal, "nobody really had a good handle on which, if any, of these equations describe the behavior of real materials in more complex flows typical of manufacturing systems."

A major stumbling block has been a computational limitation associated with the Biblical prophet Deborah: She said the mountains flow according to the Lord's time scale and not humans. In the case of polymers, a high Deborah number means the flow is strong enough that the polymer becomes highly oriented in one direction and stretched. For most flows, this occurs when the time it takes for the polymer to relax is long compared to the rate at which the flow is deforming it. Fluids that act this way are called "non-Newtonian." Unlike Newtonian fluids such as water, after they're stretched they don't immediately return to their unstressed state.

"When computational studies of polymer flow began 15 to 20 years ago," explains Leal, "researchers found that numerical methods used for Newtonian fluids wouldn't work for any flow beyond a Deborah number of approximately one. Not only are the equations more complicated, with more unknowns, but they couldn't be solved for any flow in a Deborah number range of technological significance. Gradually, driven by the work of researchers worldwide and a series of international workshops, the so-called high Deborah number problem has been whittled away. Most of the issues are resolved, and we're at a point where we can compute nontrivial flows."

### The Dumbbell Model

To begin understanding the behavior of polymers in flow, Leal and his colleagues have initially adopted a simple mathematical model: the elastic dumbbell model, two spheres with a spring between them. "The dumbbell is oriented and stretched by the flow," says Leal, "and the resulting tension in the spring makes it want to return to its non-stretched state." Though the dumbbell is a simple shape, it provides a reasonable model of the polymer chain dynamics, and the resulting mathematical model for flow is more complicated than previous models that lack a description of molecular structure.

For the last three years, Leal and colleagues have been running and modifying their computer code at Pittsburgh, complementing this theoretical research with laboratory work. "What we're finding," he says, "is that the relatively simple computational models help to explain the connection between what you see macroscopically in the fluid mechanics and what's actually happening at the microscopic level. In most of the problems we're looking at computationally, there's a one-to-one counterpart with a flow geometry we're looking at in the laboratory."

An example is how polymers flow between two rotating side-by-side cylinders in an experimental device called a two-roll mill. Leal's simulations of flow characteristics such as velocity and the polymer's shape and how much it is stretched show good agreement with observed data. "We find encouraging qualitative and even quantitative comparisons," says Leal, "between the model calculations and what we see experimentally." (SE)

#### Reference:


P. Singh and L.G. Leal, "Computational Studies of the FENE Dumbbell Model in a Co-Rotating Two-Roll Mill," *Journal of Rheology* 38, 485-517 (1994).

This research is supported by the National Science Foundation.

### Simulating a Stretched Polymer

*These graphics represent simulations of polymer flow in a planar cross-slot device. Dilute polymer solution enters from both sides of the input channel (horizontal axis). The two inflow streams splash into one another, creating a stagnation flow pattern at the center, where they merge and flow out both directions of the cross channel (vertical axis). The red lines indicate flow streamlines. The vertical colored stripe (above) is a closeup of the stagnation region.*


*Colors emanating from the stagnation point show how much the polymers stretch with reference to equilibrium state. The polymers enter unstretched, and near the stagnation point stretch out. Since the relaxation process takes time, the polymers remain stretched a distance downstream. With this model (a "dumbbell model," adapted from Hinch and DeGennes), maximum stretch is 300 (dimensionless length squared), compared to two at equilibrium. (MS)*



### Chips and Steps

Where would we be without silicon? This abundant element — it comprises 27 percent of the Earth's crust — provides the starting material for nearly all the microelectronic "chips," from VCRs to the Space Shuttle, that have transformed the way we live. In computing and communications, 40 years of continuing innovation in smaller, faster circuitry, with silicon as the base material, is driving us headlong into the Information Age.

*The first starting-from-scratch calculations to give the atomic and electronic structure of silicon "steps."*


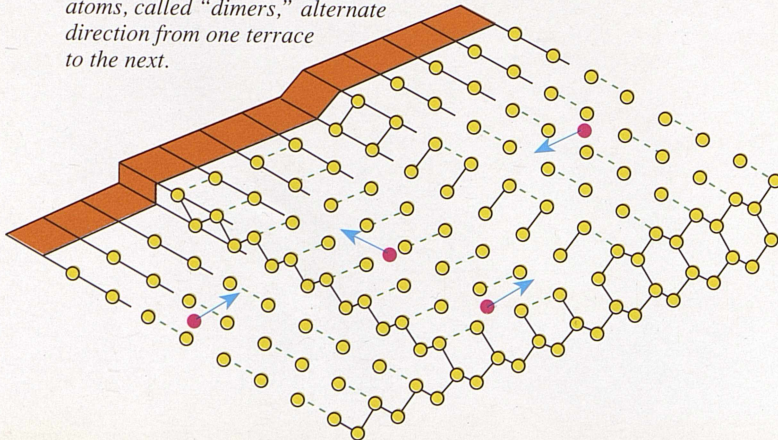


Jerzy Bernholc, North Carolina State University. At a 1992 PSC seminar, Bernholc described quantum molecular dynamics and its applications in his research.

Silicon remains the choice for many applications, despite other materials that promise faster electronics, because it is durable and reliable, and over the years many problems associated with producing chip-quality crystals in commercial quantity have been worked out. That hardly means, however, that no challenges remain in silicon chip technology. Whether it is to solve important problems in science or for the rapidly evolving new world of communications and entertainment, we need faster circuitry, smaller, more perfect chips.

With these needs in mind, Jerzy Bernholc and his colleagues are working to understand the atom-by-atom details of how silicon grows. The process of producing the relatively pure, defect-free material required for chips is called "growing" because it involves starting with a small "seed" crystal of the desired structure and adding to it one atom at a time. Bernholc, a solid-state physicist, has used the CRAY C90 and, recently, the CRAY T3D at Pittsburgh to investigate "step-flow" growth, a prevalent silicon growth process by which layers build up, an atom at a time, into a staircase-like formation of terraces.

*This schematic shows three terraces separated by one-atom layer steps. Adatoms (magenta) arrive at a surface and travel to the step edges, where they fill out the crystal by extending the terrace. The bonds (green dashed lines) between paired surface atoms, called "dimers," alternate direction from one terrace to the next.*




*The Digital Equipment Corporation Alpha microprocessor, the core silicon "chip" used in the CRAY T3D, can do 150 million calculations a second.*

Using the C90, Bernholc's research group completed the first starting-from-scratch calculations to give the atomic and electronic structure of silicon "steps." Extending this work, further C90 calculations explored the pathways by which a deposited atom, an "adatom," moves across a silicon surface and fills in an empty space at the growing edge of a step. With availability of the CRAY T3D, Bernholc further expanded the project, adding complexity and realism to the calculations. The results are the beginnings of a detailed road map of the travels of silicon adatoms.

### MBE and QMD on the C90

"The goal of our project," says Bernholc, "is to understand in atomic detail the growth modes of semiconductors and to identify the conditions for high quality layer-by-layer growth." Step-flow growth is the most common growth process in molecular beam epitaxy (MBE), a relatively new method of producing chip-quality silicon crystals. In MBE, a gaseous beam deposits a one or two atom layer of semiconductor material on a base layer or "substrate," with the inter-atomic spaces of the substrate forming an egg-carton-like template for the adatoms.

To simulate these processes, Bernholc uses a powerful computational approach called quantum molecular dynamics (QMD). His code, adapted from a method developed by European physicists Roberto Car and Michele Parrinello, is a truly quantum approach that allows the electrons and atoms to move freely with time, so that the atom-atom interactions are computed *ab initio* ("from the beginning"), directly from the electronic forces. Bernholc's C90 adaptation of QMD runs at 693 million calculations a second (Mflops) on a single C90 processor, one of the fastest codes running on this machine.



This sideview shows Jerzy Bernholc' CRAY T3D simulations of adatom diffusion at high temperature on a seven-layer silicon slab with a one-atom layer step. Three adatoms (red) are deposited on the surface, two on the lower terrace (violet), one on the upper terrace (aqua). Dimer tilting is somewhat disordered due to the high temperature. Silicon "bulk" atoms (blue) are sandwiched between the stepped surface and a bottom layer (purple). The slab bottom is saturated with hydrogen atoms (white). The slab includes 232 silicon and 72 hydrogen atoms.

Taking advantage of "friendly user" testing on Pittsburgh's new C90 in early 1993, Bernholc carried out a series of four-processor runs (2.3 Gflops) that resulted in the first *ab initio* calculations giving the structure of steps. "We compared the calculated atomic and electronic structure," says Bernholc, "to atomic resolution images from scanning tunneling microscopy (STM). We verified that we did indeed reproduce what was observed experimentally, and our results provide a better interpretation of the experiments than others that had been proposed. So we now know what the step structures are with some precision."

#### On the T3D: Adatom Jumps & Dimer Buckling

Further *ab initio* calculations on the C90 explored how adatoms move across the silicon surface, jumping from one inter-atomic space to the next, until they become incorporated in the step edges. To make this calculation feasible, these simulations were done at zero temperature and investigated only certain possible pathways. The results showed the energy barriers that must be overcome for the adatom to "jump" and the binding sites where they end up.

A surprising finding was that the jumps depend to a degree on "dimer buckling." The silicon surface forms bonded pairs of atoms, called "dimers," that tilt, with one atom higher than the other, resulting in an uneven, buckled surface. Bernholc' calculations indicate that it is easier for an adatom to move past a "down" than an "up" atom and that the dimers can shift their tilt in response to the adatom. "This information," notes Bernholc, "would have been difficult or impossible to obtain with experimental methods."

With availability of the CRAY T3D, it became feasible to do these calculations at high temperature, which more realistically simulates the silicon

growth process as it occurs during MBE. Qiming Zhang, a materials scientist working at PSC as part of Cray Research's Parallel Applications Technology Program, translated Bernholc' QMD code to run on the T3D while also taking advantage of high speed communication between the T3D and C90. A relatively small part of the computing, the atomic movements, runs on the C90 while the electron calculations run on the T3D. The result — a substantial performance improvement: 6 Gflops using 128 T3D processors, compared to 2.3 Gflops on four C90 processors, one-fourth of the machine in both cases.


Using the T3D code, Bernholc carried out a series of heating up experiments. The results reproduce the high temperature structure observed with STM. More recent simulations focus on diffusion across a heated surface. Further improving on the C90 studies, these calculations allow the adatoms to follow any pathway available to them. So far, despite frequent oscillations in the tilting of the dimers at high temperatures, the jumps correspond to the zero temperature pathways.

"That's the first question," says Bernholc. "Do the atoms follow a different pathway? For the flat surface, the pathway is consistent with low temperature, exactly what we'd expect. We will do the same thing across the step edges. In this high temperature regime, we can see how the jumps occur in real time and probe for new and unforeseen diffusion mechanisms. The T3D has made these simulations possible." (MS)

#### References:

- Q.-M. Zhang, C. Roland, P. Boguslawski & J. Bernholc, "Ab Initio Studies of the Diffusion Barriers at Single-Height Si(100) Steps," *Physical Review Letters* 1995 (in press).
- P. Boguslawski, Q.-M. Zhang, Z. Zhang & J. Bernholc, "Structure of Monatomic Steps on the Si(001) Surface," *Physical Review Letters* 72, 3694 (1994).
- M. G. Wensell, Q.-M. Zhang, C. Roland & J. Bernholc, to be published.

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



### Unpaired Spins and Magnetic Moments

When we think of magnetism, most of us remember iron filings on a piece of thin cardboard. Moving a bar magnet under the paper makes the normally invisible lines of the magnetic field visible. At this point, the inquisitive among us may ask questions: What is a magnetic field? Where does magnetism come from? Think of an attractive dancer circling the outskirts of a crowded dance floor, spinning alone to the music. At a fundamental level, says Arthur Freeman, that's magnetism.

In the three-dimensional dance space occupied by an atom, electrons pair up as they revolve around the nucleus, each spinning on its axis in an opposite direction to its partner — in effect, canceling each other's spin. Among the outermost electrons of some metal atoms, however — those most weakly bound to the nucleus, a few unpaired electrons spin in the same direction. It's this net unpaired spin, explains physicist Freeman of Northwestern University, that in these magnetic metals gives rise to a "magnetic moment," creating the pull of a magnetic field.

Relying on powerful computational methods to simulate the bewildering complexity of this atomic dance floor, Freeman has revolutionized ideas about magnetism. During the last decade he has shown, contrary to what physicists believed prior to his work, that a surface atomic layer of a metal can have more magnetic moment than the bulk form of the same metal. Research stimulated by this finding has led to increasing magnetic data storage on compact discs by more than 40 times.

Freeman's pioneering work with magnetism along with important computational studies in superconductivity and other solid-state properties has helped to herald a new branch of science, computational materials science. "We are now making materials," says Freeman, "with exotic properties that nature doesn't give us, new materials from old elements. What happens is once you have a new tool — like Galileo had the telescope — you make discoveries. We have a new tool — the computational capability of the computer. So we make new discoveries."


*"O body swayed to music, O brightening glance,  
How can we know the dancer from the dance?"*

— W. B. Yeats, "Among School Children"

### Thin-Layer Magnetism Comes Alive

Freeman's interest in thin-layer magnetism began in the mid-1970s when he didn't believe experiments showing that the surfaces of bulk metals — such as nickel, iron and cobalt — have no magnetic moment: "We did computational investigations, and we found that at least theoretically, they weren't magnetically dead. They were more alive than the bulk." With iron, Freeman and collaborators discovered that magnetism at the surface actually increased 40 percent.

Freeman next modeled a single-atom layer, a "free monolayer," and discovered that the magnetic moment is even stronger than on a bulk metal surface. Experimentalists couldn't make free monolayers in the laboratory to test these predictions, so Freeman began simulating a metal monolayer on a surface of silver or gold, which because they have no unpaired electrons are magnetically neutral. He found that the magnetic moment dropped only slightly from the free monolayer value, less than 1 percent for iron.



*Calculated electron spin density for the surface of iron (left) compared to the free monolayer (right). Dark blue indicates negative spin, and other colors are positive, increasing from light blue through pink. A larger area of positive spin in the free monolayer accounts for a magnetic moment of 3.18 bohrs, compared to 2.98 for the surface and 2.20 for bulk iron.*

The increased magnetic moment that occurs in a monolayer, explains Freeman, happens because each atom has fewer atomic neighbors than in the bulk: "You have a reduced dimensionality. In a single layer, an iron atom, for instance, has four atoms surrounding it, whereas in the solid it has eight. The interactions with other atoms in the bulk tend to reduce magnetism."

In recent work on Pittsburgh's C90, Freeman and former post-doctoral fellow Ruqian Wu, now at the California State University-Northridge,

determined that magnesium oxide is even more neutral than silver or gold. They also found that sandwiching layers of metal between neutral materials helps protect the metal and doesn't significantly reduce its magnetic properties. "You've covered it so it won't rust or oxidize," says Freeman, "so for practical purposes, multilayers are much better."

### Improved Magnetic Storage

When laboratory researchers first tried to duplicate Freeman's simulations of a monolayer on a neutral surface, the experiments failed. "One of the crucial experiments for monolayer magnetism was iron on silver," says Freeman. "They did photoemission, looking at what happens to electrons emitted from the surface. They found no magnetism."

The researchers, it turned out, were looking in the wrong direction. In most cases, magnetic moments align parallel to the plane, but in this case, they align perpendicularly, vertical instead of horizontal. With corrected technique, the researchers found the magnetism. "In an attempt to confirm or refute the predictions of enhanced magnetism," Freeman says, "they discovered something very important for applications in magnetic recording."

The perpendicular alignment permits packing more information. This factor coupled with the stronger magnetic moment of a single-atom layer is transforming the magnetic recording industry of compact discs and lasers. Freeman, along with Wu and visiting professor Dingsheng Wang of the Academy of Sciences of China in Beijing, developed a new computational approach — the state-tracking scheme — that predicts which materials have magnetic moments oriented perpendicularly. Along with enhanced storage, the new magnetism revealed by Freeman's work broadens the choice of materials beyond conventional oxide discs.

### Taking Supercomputing to the Limit

Because detailing electronic structure requires complicated quantum-mechanical calculations, Freeman's research demands supercomputing. About 18 person-years — by Freeman and three post-doctoral fellows — went into development of a very precise computational approach known as the full-potential linearized augmented plane wave (FLAPW) method. With FLAPW, Freeman and his collaborators are able to study 75 atoms in a unit cell on the C90. Fifteen years ago, even using less precise methods, the number was only 10. But he's still not satisfied.

"I'm still overwhelmed that we can deal with so much complexity. The ability to calculate spin and orbital magnetic moments layer-by-layer and atom-by-atom is remarkable. But we want to answer questions that are more complex. We need every increase in computing power. We never get enough for our simulations of real materials."

Currently Freeman and his colleagues are further developing their ability to calculate perpendicular magnetism in overlayer and multilayer systems. They also are studying the single-layer magnetism of metals from the next row in the periodic table — the so-called 4d elements, such as rhodium and ruthenium. While these elements have no magnetism in the bulk, a magnetic moment arises when a monolayer is deposited on a substrate.

"We're creating magnetism," says Freeman, "where none existed before." (SE, MS)

#### References:



- Ruqian Wu, Dingsheng Wang and A.J. Freeman, "Magnetism in Low Dimensional Systems: Magnetic Properties of Thin Films," in *Handbook of Surface Imaging and Visualization*, A.T. Hubbard ed. (CRC Press: in press).
- Ruqian Wu, Dingsheng Wang and A.J. Freeman, "First Principles Investigations of MCD Spectra and Sum Rules for 3d Transition Metal Surfaces," *Journal of Magnetism and Magnetic Materials* (in press).

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

Arthur J. Freeman,  
Northwestern University.

*Calculated spin density for a ruthenium monolayer deposited on magnesium oxide. The large positive spin area (light blue through pink contours) of the ruthenium atom corresponds to a magnetic moment of 2.1 bohrs. Since ruthenium is not magnetic as a bulk metal, this is an important prediction.*

*"We're creating magnetism where none existed before."*



### Soaking Up Hydrogen

A silvery-white metal named after the goddess of wisdom. It doesn't tarnish in air. Mixing it with gold produces the white gold used in jewelry. Closely related to platinum, it's used in dentistry, watchmaking, surgical instruments and electrical contacts.

The metal is palladium, number 46 on the periodic chart, and its most remarkable property, actually, is none of the above. Palladium soaks up hydrogen like a sponge — that is, if you can imagine a sponge that soaks up hundreds of buckets of water. At room temperature and atmospheric pressure, palladium can absorb up to 900 times its own volume of hydrogen. "That means," says Khalid Mansour, "that if you were to pump hydrogen into a bottle, it would take enormous pressure to store the same amount easily absorbed in a palladium bed of the same volume."

This makes palladium an efficient, safe storage medium for hydrogen and hydrogen isotopes, such as tritium, a byproduct of nuclear reactions. Nevertheless, the microscopic details of how this absorption process works are poorly understood. Better understanding of what happens at the molecular level, such as phase changes that occur as hydrogen fills the metal — forming a metal hydride — and the effect of defects in the material, could give clues to designing metal hydrides that perform better. Such research also bears directly on technologies that use metal hydrides in fuel cells and batteries.

At Westinghouse Savannah River Company (WSRC) in South Carolina, Mansour (now at Cray Research), Ralph Wolf and their colleagues Myung Lee and Clemson University physicist John Ray are working to fill the gaps of knowledge. "Basically," says Mansour, "the experimentalists did things by alchemy. We know that palladium has this remarkable ability to soak up hydrogen, but there's a lot we don't understand. What we need is

physical understanding of the underlying behavior at a microscopic level. And that's where we come in."

The researchers are using the CRAY T3D to simulate palladium filling up with hydrogen, forming palladium hydride. Essentially, they're doing experiments with the computer to find out how the number of hydrogen atoms in palladium changes relative to the external pressure of hydrogen gas. To do this, they've developed a new approach to modeling the key thermodynamical property, chemical potential. Applying this method on the T3D allows them to model a much larger number of atoms, and they're getting results that, for the first time, make it possible to directly compare simulations with observed experimental data.



### Extending the Life of Palladium Beds

A principal objective of the WSRC research is to extend the useful life of the palladium storage medium. The palladium-hydride storage process in batteries and other storage technologies is cyclic. "Hydrogen is loaded into the metal hydride," explains Wolf, "and then desorbed back off, through as many as 10,000 cycles. You keep absorbing and desorbing hydrogen, or charging and discharging, and like a sponge getting wet and drying out over and over again, it expands and contracts. Over time the material begins to break down."

By understanding this breakdown at the level of atoms, the researchers hope to find solutions. Palladium in reality, for instance, is seldom a perfect crystal. Like the grain in wood, it has grain boundaries and dislocations. As the metal expands and contracts over time, it's apt to fracture at grain boundaries. An additional wrinkle is that tritium stored in a metal decays radioactively to helium, which becomes trapped in the material and forms new defect sites. "For some reason," explains Mansour, "the helium stays in, creating a permanent deformation in the palladium crystal structure that alters the material's performance."


The researchers want to simulate this complex process. Mansour sees it as a multi-step objective: "We want to understand how dislocated palladium behaves. We then want to understand how the dislocated material acts in the presence of hydrogen. Then, we'd like to understand how it behaves with helium. The next step would be to put hydrogen and helium both in there, which is what happens in real life. These studies are underway, and this is what we're doing on the T3D."

Ralph Wolf,  
Westinghouse Savannah  
River Company.



Khalid Mansour, Cray Research Inc.

Beta phase in palladium hydride near the phase transition at 300° Kelvin. In the alpha phase, hydrogen atoms (yellow) populate the interstices in the lattice structure of palladium (purple) randomly. At the phase transition, the lattice expands allowing hydrogen to cluster at higher density, as visualized here. This image shows the lattice from the (001) direction.



### Modeling Palladium Hydride on the T3D

To describe the interactions among the atoms in palladium hydride, the WSRC team uses a method, the embedded atom method (EAM), originally developed at Sandia National Laboratory. The WSRC group, however, made an important revision. Unlike prior approaches, they allow the volume of palladium atoms to expand as the number of hydrogen atoms increases, thereby tracking a unique physical reality that occurs in the palladium-hydride system.

“Volume expansivity is critical,” says Wolf. “What we needed was a method to fix the chemical potential of hydrogen atoms, the number of palladium atoms, temperature and pressure. The talents of Professor Ray were invaluable.” The final result is a simulation that allows hydrogen atoms to be created and destroyed in an expandable metal lattice.

In recent work, the group applied their method on the T3D at Pittsburgh (and at the Arctic Research Supercomputing Center). Vance Shaffer of Cray Research helped convert the code to run on the T3D. The results have been gratifying. “The T3D is a very stable machine,” says Mansour, “and the performance has been wonderful. We’ve been able to simulate systems roughly 20 times larger than we’ve been working with till now.” Their simulations of palladium with dislocations and grain boundaries show increases in the hydrogen concentration in the metal that correspond to expansion of the palladium atoms as hydrogen pressure goes up.

These results, and the potential for even faster turnaround with better optimized code (which they’re working on), have encouraged the WSRC group. They can now model systems of a size that approach reality, a factor especially important, says Wolf, to understanding grain boundaries and dislocations. “The systems we’re running now are 3,000 to 5,000 atoms, and this range will allow us to investigate the role of defects in the materials.”


A related goal that now appears in sight is to simulate alloys that could prevent or retard fracturing. One possibility, says Wolf, would be to alloy palladium with nickel at grain boundaries: “Hydrogen, as far as we know, doesn’t like nickel, so it wouldn’t go near the grain boundary. You can map out scenarios like that in your head, but you’d like to be able to prove them with simulations. Much of the experimental strategies in metal hydrides research are trial and error. A pinch of this, a tad of that. It’s like the lottery — you could get lucky. What we really need is better knowledge, and that’s what simulations can give us.” (MS)

#### References:

- Ralph J. Wolf, Myung W. Lee, Ricardo C. Davis, Patrick J. Fay & John R. Ray, “Pressure-composition isotherms for palladium hydride,” *Physical Review B* 48, 12415-418 (1993).  
Ralph J. Wolf, Khalid A. Mansour, Myung W. Lee & John R. Ray, “Temperature dependence of elastic constants of embedded-atom models of palladium,” *Physical Review B* 46, 8027-35 (1992).

This research is supported by the U.S. Department of Energy.

*“What we really need is better knowledge, and that’s what simulations can give us.”*



## Electrochemistry, Batteries and Rust

When Gregory Voth and his colleagues say “ET,” chances are they’re not talking about the cute alien in the Spielberg movie. To Voth, a theoretical chemist at the University of Pennsylvania, ET means electron transfer, one of the most fundamental processes in chemistry.

“There’s this whole field of electrochemistry,” he says, “that goes back to Faraday.”

Michael Faraday’s 19th century experiments passing electricity through water and other compounds laid the groundwork for modern understanding that matter has discrete particles of electricity — electrons — that transfer their allegiance during chemical reactions.

In daily life, ET affects us in many ways. The batteries in our cars and portable stereos depend on it. Perhaps its most familiar and unwelcome guise is reddish-brown oxidized metal, AKA rust. In the United States alone, more than \$10 billion a year goes to replace corroded equipment and to protect existing structures. “Imagine how much money we could save,” says

Voth, “in infrastructure and repair, and in our own personal lives, if we could find better ways to prevent corrosion.”

Voth’s research is aimed in that direction. He and his colleagues have mounted the first major foray of computer simulation into this important field. Using the CRAY C90 at Pittsburgh, they recently completed a massive computation that, for the first time, simulates a typical electrochemical reaction — an electron transferring from a platinum electrode to an iron ion immersed in water. His surprising, interesting results lead to the conclusion that quantum methods are required to accurately simulate ET processes in water.


## Fluctuating Solvent

The difficulties of simulating ET processes relate to the three-way interaction among an electrode, either metal or a semiconductor, an electrolytic solution — often water — and ions in the electrolyte. “You have a very complex interface,” says Voth. “There’s water interacting with the metal, water interacting with the ion, and the ion interaction with the metal.”


Still, the most perplexing theoretical problem has to do with how an electron from the sea of electrons in the metal jumps to a swimming ion. With ferric iron, for instance, the +3 charge of the ion will be in happy equilibrium with the water around it. It violates laws of energy conservation, explains Voth, for the ion to instantaneously gain an electron, shifting to ferrous iron (+2), without first going out of equilibrium with water.


The anomaly has been accounted for by Voth’s former graduate advisor at Caltech, Rudolph Marcus, who received a Nobel prize for his theoretical picture of electron transfer. “The solvent needs to fluctuate,” says Voth. “Of course, it’s a liquid so it’s constantly fluctuating, and instantaneous fluctuations every now and then allow that electron to transfer from the metal to the ion without this energy cost. Calculating the free energy associated with that state tells us about the rate of electron transfer.”

Taking off from Marcus, Voth and his graduate students Jay Straus and August Calhoun developed a set of equations and put them into a framework that makes it possible, with the help of supercomputing, to simulate electrochemical ET processes. To “see” the transition states of the fluctuating solvent computationally, they adapted a method called “umbrella sampling” that samples different energy states of the ion-water system.



Gregory A. Voth,  
University of Pennsylvania.





Voth and his colleagues are beginning to simulate more complex electrochemical systems such as multiple ions of different species and a first-principles treatment of the electrode surface. This visualization shows a snapshot from a simulation that includes a chlorine counterion (green) and an iron ion in water with a platinum electrode. Including the counterion more realistically represents the chemical environment of an ET reaction in electrolytic solution. Ultimately, Voth and his colleagues want to include many ions and counterions in a single simulation.

### Results: The Influence of Water

Using the CRAY C90, Voth and his co-workers simulated an iron ion in water at 300° Kelvin interacting with the surface of a platinum electrode. To realistically represent the ion's interactions with the water around it, the simulation included 671 water molecules. The first part of the work applied an essentially Newtonian scheme for calculating the forces between atoms. The results show an energy barrier (about 10 kilocalories per mole) that must be overcome for the ion to shift between the ferrous and ferric state.

"There's been debate for years," says Voth, "about whether the solvent really influences an electrochemical electron transfer reaction. We showed fairly conclusively that under these conditions it does. This information can be very difficult, if not impossible, to get from experiment."

The next phase was to treat the water quantum mechanically, which meant quantizing the hydrogen nuclei of each water molecule. Voth and his co-workers applied the Feynman path integral technique, an approach to quantum computations that Voth adapted to predict the rate of ET processes. "There are in principle very large quantum effects in water from a variety of sources," says Voth, "and we wanted to investigate what difference this could make."

With 400 hours on a CRAY C90 processor and an additional eight months on five workstations, Voth believes this is the largest liquid-state Feynman path integral simulation carried out to date. "What we found was this striking effect that, given the same parameters, the quantum modes in water are very significant in determining the free energy of the reaction. You see that the barrier is lowered going from left to right in the free energy curve, which means the reaction will happen faster. More than that there's now a driving force. You go downhill from left to right, which means the reaction will occur spontaneously in that direction, whereas in the classical case, it's just as happy to go in either direction."

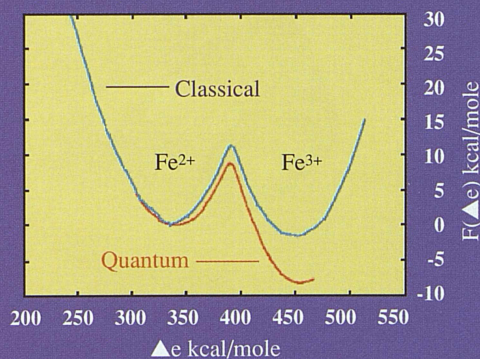
Voth and his colleagues are now working out a mathematical theory to describe this quantum effect. "That's the beauty of computer simulation. This study exemplifies how it adds a third tier to science, on top of theory and experiment. You can find novel effects that you wouldn't have expected. And now we're going to develop the theories to explain their origin." (MS)

#### Reference:

Jay B. Straus, August Calhoun & Gregory A. Voth, "Calculation of solvent free energies for heterogeneous electron transfer at the water-metal interface: Classical versus quantum behavior," *Journal of Chemical Physics* **102**, 529-39 (1995).

This research is supported by the Office of Naval Research, the National Science Foundation and a David and Lucile Packard Fellowship in Science and Engineering.

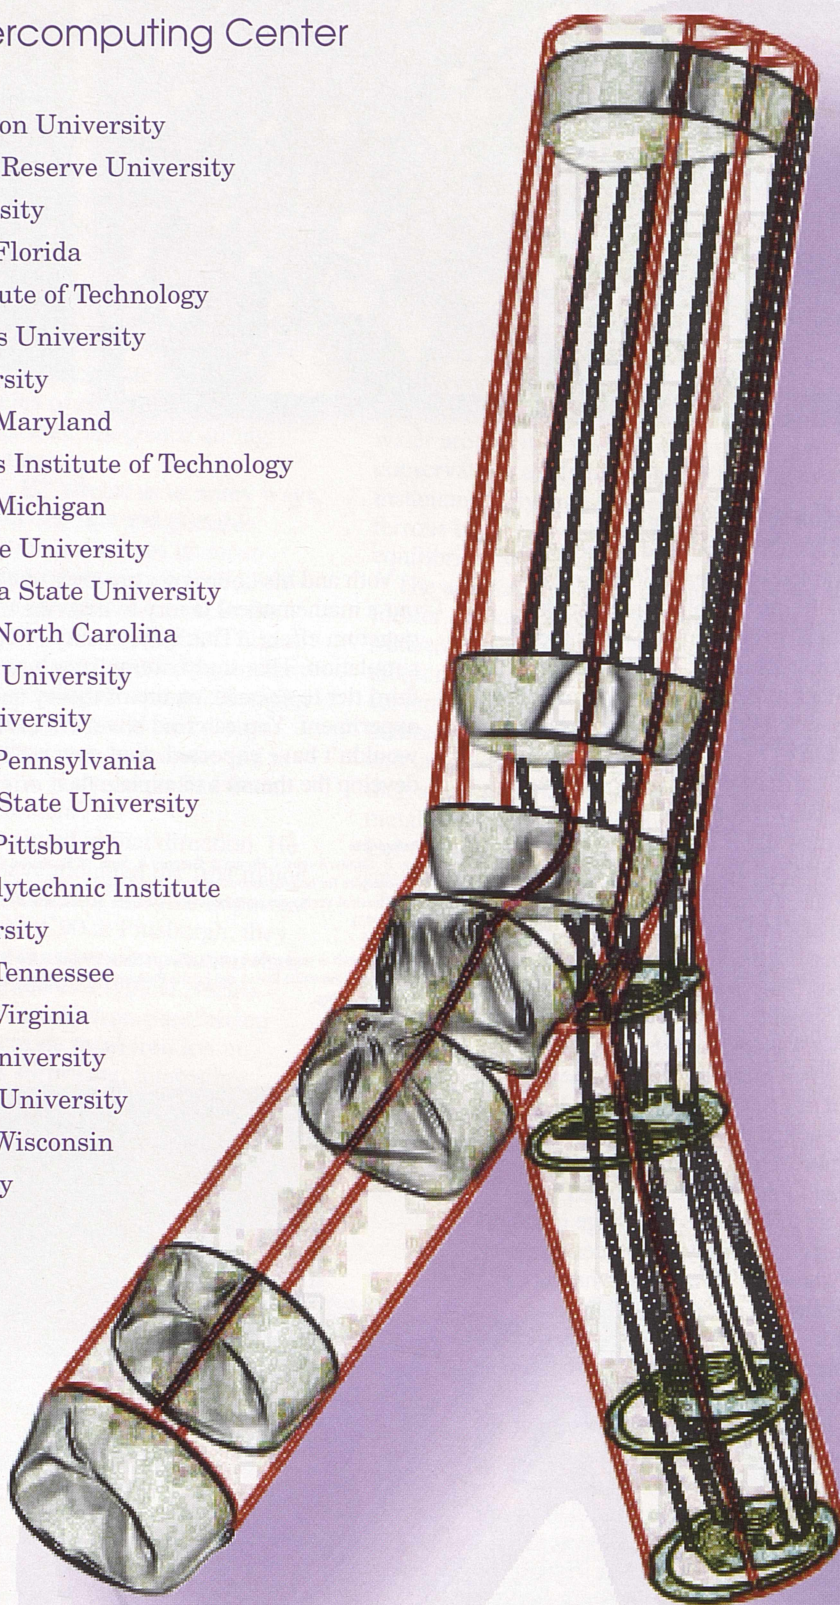
"We found this striking effect. The quantum modes in water are very significant in determining the free energy of the reaction."



Free energy for electron transfer at the water-metal interface. The two wells in the curve correspond to the  $Fe^{2+}$  (left) and  $Fe^{3+}$  (right) state of the ion. The hump between the wells indicates the "energy barrier," approximately 10 kilocalories per mole, between the two states.

## The Academic Affiliates of the Pittsburgh Supercomputing Center

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



The Pittsburgh Supercomputing Center is a joint project 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 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.  
Storage Technology Corp.  
Sun Microsystems, Inc.  
Thinking Machines Corp.  
Apple Computer, Inc.  
Bell of Pennsylvania  
The Vira L. Heinz Endowment  
The National Institute of Standards & Technology  
U. S. Dept. of Energy

#### **credits**

**editor:** *Michael Schneider, PSC.* **design & illustrations:** *Stephen Palmer, Pittsburgh.*

**writing & graphics research:** *Steve Eisenberg & Michael Schneider, PSC.* **copy editing:** *Ryan Little, PSC.*

**photography:** *Jeff Macklin, Photography and Graphic Services at Mellon Institute, Jim Schafer.*

**printing:** *Westinghouse Printing, Pittsburgh.*

*Page layout & design were done by Steve Palmer on a Macintosh Quadra 700 at PSC using PageMaker 5.0 desktop publishing software and Adobe Photoshop. These digital layouts included spot color, encapsulated postscript graphics, including computer-created illustrations, and tag image file format (TIFF) scientific graphics from calculations by researchers at PSC.*

*Thanks to the PSC Sci-Vis group for their help with computer graphics. Thanks to Lisa Andra and Natalie Noble of the PSC staff for transcribing research interviews.*



Printed on recycled paper with vegetable-based inks.

