


Pittsburgh Supercomputing Center
PROJECTS IN SCIENTIFIC COMPUTING



1 9 9 8

C O N T E N T S

**Foreword from
the Directors 2**

**Pittsburgh
Supercomputing
Center 1998 4**

**Supercomputing in
Pennsylvania 8**

**Partners
with Energy 10**

**Science
Articles 12**

**STRUCTURE OF
PROTEINS AND DNA
..... 14**

**PROTEIN AND
NUCLEIC ACID
SEQUENCE ANALYSIS
..... 18**


**UNDERSTANDING
THE BRAIN
..... 20**

**EARTHQUAKE
MODELING
..... 22**

**MANUFACTURING
PROCESS DESIGN
..... 24**

**UNDERSTANDING
THE EARTH
..... 28**

**EVOLUTION AND
STRUCTURE OF
THE UNIVERSE
..... 30**



*The Pittsburgh Supercomputing Center
is a national,
multi-agency supercomputing center
with major funding from
the Department of Energy,
the National Institutes of Health,
the Commonwealth of Pennsylvania,
the National Science Foundation
and private industry.*

*PSC provides government, academic
and industrial users with access to the country's
most powerful academically-based high-performance
computing resources.*

*PSC advances the state of the art
of high-performance computing and communication
and offers a flexible environment conducive to
solving today's largest and most
challenging computational science problems.*

This year's *Projects in Scientific Computing* continues our tradition of presenting examples, from work done at Pittsburgh Supercomputing Center, of outstanding and award-winning efforts in computational science. The work is in a variety of fields, and highlights manifold ways in which work at PSC spans the gamut of computational science. Besides the intrinsic scientific merit of each of these examples, we point out that:

- Peter Kollman's work, the largest protein molecular dynamics calculation done to date anywhere, and Brooks', both on protein folding are outstanding examples of the use of massive computation in biochemical simulations;
- The Rosenquist and Nicholas work exemplifies the increasingly tight links between high-performance computing and bioinformatics;
- Cohen's work on functional magnetic imaging demonstrates the importance of effective, high-performance networking in biomedical and other distributed applications;
- The video that PSC produced for Bielak's massive earthquake simulations was the only scientific example honored by acceptance in the Siggraph 97 Electronic Theatre show;
- The industrial examples (turbine design at Westinghouse and automotive uses of aluminum at Alcoa) show the importance of high-performance computing to the nation's economic and technical well-being;
- Goodrich's simulation of storms in space demonstrates PSC's flexibility in making massive computing resources available to single researchers on very little notice;
- Lehner's Metropolis award for the best thesis in the world in computational physics in the past year (see Winicour's project on black holes) illustrates how critical high-performance computing resources are to the education of the next generation of computational scientists.

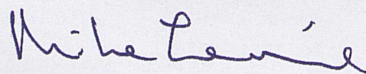
PSC's leading edge facilities, backed by its talented and dedicated staff, have enabled these achievements. Listed as the top-ranking academically-based U.S. machine, PSC's CRAY T3E, as has been the case with the T3D, is routinely focused, in its entirety, on particularly demanding applications. This ability, along with coupled, tuned networking and storage facilities, capable of sustained wide-area transfers of hundreds of gigabytes per day are critical to support applications of this scope and nature.

PSC's expertise and computational capability are proving invaluable to its numerous sponsors, including:

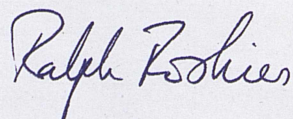
- The Department of Energy, with its interests in working more closely with the university community in accelerating the nation's ability to carry out very large-scale simulations;
- The National Institutes of Health with its concern with aggressively using high-performance computing and communication capabilities in the service of biomedical research;
- The Commonwealth of Pennsylvania with its interest in supporting an outstanding research and educational capability;
- Corporations interested in using advanced technology to improve products and company performance.

The accomplishments reported here provide additional justification for renewed, multi-agency interest in augmenting the national support for high-performance technical computing applied to a wide range of national challenges.

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



Michael J. Levine, scientific director



Ralph Z. Roskies, scientific director



1998

Pittsburgh Supercomputing Center

This Pittsburgh black-and-gold CRAY T3E is known to its users as "Jaromir," in honor of the Pittsburgh Penguins high-scoring forward. It is one of the most powerful computing systems in the world and, by international rating, the *most* powerful academically-based system in the United States. With 512 separate processors that can be harnessed to work simultaneously on the same problem, Jaromir at its peak can do 460 billion computations per second, 80 times more processing power than every person on Earth holding a calculator doing one calculation per second.

"We've had a very positive experience at PSC in both normal and emergency computing. We needed to ask for special treatment to run a simulation of a space weather event while it was still in the news, and we got it. They roll up their sleeves and go to work to produce scientific results."
— Charles Goodrich, research scientist, Space and Plasma Physics Group, University of Maryland

"Our experience computing at PSC has been nothing short of extraordinary. To have a center that dedicates people, time and resources to our problem, storm-scale weather prediction: this has been incredibly productive."

— Kelvin Droegemeier, director, Center for Analysis and Prediction of Storms, University of Oklahoma

"Our project is going extremely well, and we're very pleased with the quality of the computing environment and user support at PSC, which is truly unparalleled. It would be very difficult to maintain this level of production anywhere else. It's great to be cranking out results like this."

— Tomás Arias, Department of Physics, Massachusetts Institute of Technology

problem

The most powerful academically-based

The Pittsburgh Supercomputing Center in 1998 is poised for the future, ready to meet diverse new challenges. Along with entering into a new collaboration with the Department of Energy, PSC supports research goals of the National Institutes of Health, the National Science Foundation, the Commonwealth of Pennsylvania and private corporations.

PSC brings high levels of professionalism, experience and know-how to the task of meeting these research goals. PSC's ability to get scientific work done is well known among computational researchers – minimal red tape, expert consulting and troubleshooting, workshops and training seminars. In Pittsburgh, you get attention to problems, and they get solved.

PSC's ten-year leadership in biomedical supercomputing continues. As part of the National Laboratory for Applied Network Research, the PSC networking group is engineering the network of the future. In 1998, as exemplified by the work described in this publication, PSC continues to make it possible for U.S. researchers to attack the important, difficult problems of engineering and science with the most advanced computational resources available.

Networking the Future


In March 1998, the PSC networking group received \$2.7 million from the National Science Foundation to operate a network engineering center for the National Laboratory for Applied Network Research (NLANR). Through NLANR, PSC staff serve as a nationwide clearinghouse and technical resource to implement the next generation of Internet technologies.

The objective is to provide comprehensive engineering support for American universities and research centers to use vBNS (the very high performance Backbone Network Service), a high-bandwidth, relatively congestion-free network for advanced scientific research. Currently, vBNS links U.S. supercomputing centers and more than 70 universities, including Penn State and Carnegie Mellon, and will eventually link more than 150 research sites.


PSC staff also evaluate and integrate new technologies into existing network infrastructure. "Integrating new technologies is critical to developing advanced research and education applications," said Gwendolyn Huntoon, who manages the PSC networking group. "We work with vendors and researchers to facilitate new technology from research, through engineering and operations, to production use."

Network engineering for American universities and research centers

The PSC managerial team (l to r): Bob Stock, associate director; David Deerfield, biomedical research; Sergiu Sanielevici, parallel applications, Janet Brown, production systems, Jan Asbury, accounting, Elvira Prologo, administrative staff; Rich Raymond, user services; Ray Scott; hardware. Gwendolyn Huntoon, networking, was not present for these photos.



PSC's networking group (l to r): Joe Lappa, Mike Lambert, Peter Berger, Khrys Myrddin, Chris Rapier, Katie Biscan, Steve Petko, Matt Mathis, Steve Cunningham, Ken Goodwin, Jeff Semke, Jamshid Mahdavi.




SOLVING

computing system in the United States.

The 1997 Computerworld Smithsonian Award for Science

Storm Warning

Since 1993, the Center for Analysis and Prediction of Storms (CAPS) at the University of Oklahoma has relied on PSC resources to test its advanced storm-forecasting technology. These springtime experiments successfully predicted location and structure of thunderstorms as much as six hours in advance for the first time, and they have proven the feasibility of storm-scale weather forecasting. For this work, CAPS and PSC won the 1997 Computerworld Smithsonian Award for Science, and CAPS won a 1997 Discover Magazine Award for Technology Innovation.



Solving Biomedical Problems

In 1987, with a grant from the National Institutes of Health, PSC established one of the first high-performance computing programs nationally to focus on biomedical research. In 1998, this program received significant additional NIH support.



PSC's biomedical and scientific visualization group (l to r): David Deerfield, manager, John Burkardt, Alex Ropelewski, Hugh Nicholas, Greg Hood, Nancy Blankenstein, Joel Welling, Greg Foss, Art Wetzel.

The National Center for Research Resources (NCRR) awarded a four-year, multi-million dollar grant for PSC to assess technologies that can improve the ability of biomedical researchers at separate locations to collaborate effectively, and the National Human Genome Research Institute (NHGRI) renewed funding for PSC to train researchers from around the country in DNA and protein sequence analysis.


Through the NCRR program, PSC will implement collaborative tools – shared whiteboard, desktop videoteleconferencing, shared software revision control and others – and evaluate their effectiveness in biomedical research. The program draws on PSC experience with “distance computing” and introduces collaborative technologies into two PSC biomedical projects.

A structural biology collaboration links groups at different locations working on software that simulates structure and movement of proteins and DNA. Research with this software (see pages 14 and 16) has led to new insights in how protein structure relates to its biological function. In the second project, computational pathology, PSC scientists collaborate with pathologists at the University of Pittsburgh Medical Center and a research team at the National Cancer Institute. This collaboration is creating automated techniques to improve the speed and accuracy of tissue diagnosis.

The NHGRI grant recognizes PSC's expertise in sequence analysis. “We're especially qualified to teach these techniques,” said David Deerfield, who leads the PSC biomedical program, “because we developed many of the applications.” Newer sequence-analysis techniques, like those developed at PSC, provide more reliable results than other commonly used methods, and they give researchers a better chance of coping with the explosion of data from sequencing initiatives such as the Human Genome Project.


Pittsburgh Supercomputing Center Workshops (1997-1998)

- Supercomputing Techniques: Parallel Processing
- Supercomputing Techniques: Cray J90 & C90
- Parallel Programming Techniques
- Recommended Programming Practices
- Biomedical Image Analysis and Visualization
- Molecular Mechanics and Dynamics of Biomolecules
- Nucleic Acid and Protein Sequence Analysis



A workshop in progress at the Pittsburgh Supercomputing Center Computer Training Center.

Recognized expertise in sequence analysis



Inaugurating High-Speed Transatlantic Metacomputing


On June 20, 1997, researchers at PSC and the University of Stuttgart, Germany, linked supercomputers on both sides of the Atlantic using high-speed telecommunications – the first time this has been done. The project coupled PSC's CRAY T3E with another at the High Performance Computing Center in Stuttgart (RUS) to run flow simulation software, called URANUS, that predicts aerodynamic forces acting on space vehicles when they reenter Earth's atmosphere.

Joining supercomputers at different locations to work on the same task is called “metacomputing,” and many large-scale problems demand the capability that's potentially available this way. The collaboration with Stuttgart helped to prove viability of this concept and showed that networking capability is in place to collaborate effectively with the European community, a long awaited step.



Looking Deep within the Earth

For realistic modeling of inner-Earth processes that govern Earth's magnetic field, Gary Glatzmaier of Los Alamos National Laboratory won the 1996 Sidney Fernbach Award, the premier annual prize in computational science. Running on PSC's CRAY C90, his model produced the first simulated magnetic-field reversal, and his computational experiments, reported as cover stories in *Nature* and *Science*, offer the first coherent explanation of this phenomenon.



The Great Data Migration

During a three-day span in spring 1998, PSC transferred over a terabyte (a million-million bytes) of data across high-speed networks from Pittsburgh to San Diego. Never before have such large quantities of data moved routinely over a great distance. By careful "tuning" of the software to network capability, PSC staff achieved milestone network performance: data transfer rates reached 120 megabits per second (Mbps) and sustained a 10-hour average of 53 Mbps.

network
performance

Beverly Clayton, PSC executive director, coordinates PSC relations with state government.



A committee representing major research universities and the state System of Higher Education institutions meets to allocate computing resources for Pennsylvania users. Front (l to r): Preston Moore (University of Pennsylvania), Bill Gould (Pennsylvania State System of Higher Education), Jacobo Bielak (Carnegie Mellon University), Steve McMillan (Drexel University), Vijay Agarwala (Pennsylvania State University). Rear: Michael Levine, Ralph Roskies, Elvira Prologo, Beverly Clayton (PSC). Not present: James Gunton, Lehigh University, and Yuan Shi, Temple University.

With Commonwealth of Pennsylvania support, PSC provides advanced computing resources and know-how to Pennsylvania universities and corporations. Reflecting PSC's contribution to both academic and industrial research in the state, six of the nine research projects described in this report are from Pennsylvania. PSC's national presence in high-performance computing and communications gives Pennsylvania a valuable resource for economic development, education and telecommunications.

Economic Development


PSC promotes high-performance computing and networking as tools to boost competitiveness of Pennsylvania business and industry. This includes outreach to Pennsylvania corporations and economic development organizations and participation in conferences, meetings and special events. Along with major corporations like Alcoa, USX and Westinghouse, other Pennsylvania businesses that have used PSC resources include Mine Safety Appliances, Concurrent Technologies Corp., Form Soft and Knowledge Systems.

Education

PSC workshops and seminars train industry scientists and engineers as well as university faculty and students in the most advanced techniques of computational science. PSC training includes extensive hands-on sessions, either in PSC's up-to-date Computer Training Center, or at corporate and academic sites around the state, which this year included Penn State's main campus in State College, Westinghouse Bettis, and Hershey Medical Center. Faculty at Pennsylvania universities incorporate PSC resources into their curriculum.

Pennsylvania universities and corporations have

Through its work with PSC (page 24), Westinghouse Science and Technology Center has significantly improved its ability to design efficient power-generating turbines.



Telecommunications

Having played a leadership role nationally since 1986 in development and evolution of the Internet, PSC's network communications staff provides a valuable resource for network planning at other Commonwealth organizations. PSC staff helped design and deploy local, metropolitan and wide-area networks, and they have helped provide high-quality, cost-effective networking at Pennsylvania universities, including the University of Pittsburgh, Carnegie Mellon and Penn State. They are recognized by the Commonwealth's Department of Administration as a partner to consult on network technologies, and are available to find cost-effective solutions for better connectivity in the Pennsylvania community.

Pennsylvania's Return on Investment

- The Commonwealth's investment over the last 12 years of nearly \$16 million leveraged an additional \$224 million from federal and other sources. Most of this money has been spent in Pennsylvania. PSC provides 100 professional jobs, and 40 internships per year that train students in advanced computing.
- State universities and corporations have ready access to the fastest computers in the world. PSC staff conduct concentrated outreach training at industry research centers.
- While the center's mission is national in scope, a third of PSC users have been affiliated with Pennsylvania institutions, with computer use and training free to educational institutions.

Legislators Support High-Technology Development

"PSC is an economic generator in southwest Pennsylvania, but it also serves the entire state. It is one of few such centers in the United States," says Rep. Joseph F. Markosek (D), 25th District. "Its presence in Pennsylvania helps to move our state forward in high-technology economic development."


"PSC plays an important role as a regional resource and has an impact not only in science and technology," says Rep. John R. Pippy (R), 44th District, "but also a substantial impact on the region's economy."

- PSC educational enrichment programs have included high school teacher-student outreach and introduction of new computer-based curricula in public-school classrooms.
- Because of PSC, Pittsburgh hosted the 1996 national conference on supercomputing, which brought in 5,000 people, including representatives of many computer companies. Economic impact on the region exceeded \$8 million.
- Having one of few such centers in the country reinforces the image of Pittsburgh and Pennsylvania as centers of high technology and helps Pennsylvania attract high-technology employers.

access to the fastest computers in the world.



Through its partnership with PSC, Alcoa realized major cost savings and competitive advantage in designing beverage cans for customers who manufacture aluminum cans. Alcoa also uses supercomputing in the design of aluminum sheet and structural components for the automobile industry (page 26).





Jim Kasdorf, director of special projects, is technical contact for DOE projects at PSC. "The DOE programs we support press the limits of large-scale simulation for scientific purposes. PSC has been a leader in this area for 10 years, and this collaboration appropriately gives us a role in enabling the highest capability academically-based scientific computing in this country."

Through a collaboration launched in February 1998, PSC provides large-scale simulation support on the CRAY T3E for two Department of Energy (DOE) research programs. As part of DOE's Accelerated Strategic Computing Initiative (ASCI), PSC is supporting university-based research, called the Academic Strategic Alliances Program (ASAP), aimed at advancing the state-of-the-art in computational science. Under ASAP, research teams at five universities are addressing complex problems – in fluid dynamics, materials science, astrophysics, combustion and advanced rocketry – to develop highly realistic simulation technologies by early in the next century.

PSC is also supporting DOE's Relativistic Heavy Ion Collider (RHIC), an advanced particle accelerator now under construction at Brookhaven National Laboratory, Long Island. RHIC will smash nuclei together to create high-energy conditions similar to those at the birth of the universe. By simulating these collisions in advance, researchers can develop data-handling and analysis procedures for when RHIC becomes operational in late 1999.


In addition to the projects represented on these two pages, PSC is also working with researchers at the University of Chicago, to simulate thermonuclear burning in stars, and with researchers at the University of Illinois at Urbana-Champaign, who are simulating the dynamic processes of solid-propellant rockets.

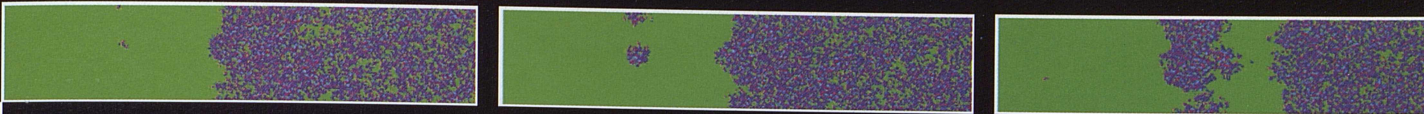
Enabling the highest capability academically-based

Shock Waves in Gas

Researchers at the California Institute of Technology are simulating what happens when a supersonic shock wave traveling through gas of one density strikes an interface with gas of different density. This fluid-dynamics problem has applications in ramjet engines, in fusion reactions, and in astrophysical phenomena such as supernovas.

In this snapshot from a 3D simulation, a Mach 10 shock wave traveling (left to right) through light gas in a square channel rammed into gas three times denser. The colored surface represents the cross section where, because of turbulent mixing between the gases, density is the same. Color indicates vorticity (increasing from blue to red). These simulations help prepare for experiments in Caltech's hypervelocity shock tunnel.





Turbulence in Turbines


At the Center for Integrated Turbulence Simulations (CITS) at Stanford University, researchers are investigating complex flow problems in gas-turbine aircraft engines. Turbulence in the turbine-blade passages affects engine performance. These images from CITS work at PSC give a microscopic view of the transition from smooth to turbulent flow along the blade surface. Such direct numerical simulations, involving 52 million grid points, provide an unprecedented level of detail.

Looking for Quark Soup


The objective of the Relativistic Heavy Ion Collider is to generate heavy ion collisions that liberate the most fundamental particles – quarks – from the nucleus where the “strong force” binds them together. Physicists believe this unbound state of matter, quark-gluon plasma – often called “quark soup” – existed immediately after the big bang, but it has never been observed.

Like the simulated collider events shown here, each collision will produce a chaotic swirl of particle traces, and each such event must be exhaustively analyzed. When RHIC is operational, scientists expect these events to produce a petabyte (a million-billion bytes) of data a year for ten years.

PSC scientists are collaborating with Carnegie Mellon researchers to conduct “mock data challenges.” Physics software generates simulated collider events that prepare the researchers and their processing systems for the real thing. PSC staff played a major role in revising this software to run on the T3E. To date, this work has generated about 250,000 simulated events.



ng
scientific computing in this country.



Metal at Extreme Conditions


Physicist Tomás Arias of MIT, winner of a 1996 DOE Young Scientist Award, calculates the atom-by-atom details of how metals transform when subjected to extremes of temperature and pressure. These visualizations from calculations at PSC show electron density of the two most likely configurations of the “screw dislocation” in molybdenum, a microscopic structure that controls how the metal deforms. A sea of electrons (blue) moves through the metal, except in void regions (purple), while other electrons (gold) remain attached to the atoms.

The computations show the energy difference between these structures to be about three times smaller than expected, critical information for the ASCI project. The excellent inter-processor communications of the T3E made it possible to get these results, which at present would have taken an impractical amount of time on other machines.

1998

projects






STRUCTURE OF PROTEINS AND DNA

14 **Watching a Protein Fold**

Folding of Villin Headpiece Sub-Domain Observed for a Full Microsecond
Peter Kollman, Lu Wang, Yong Duan, University of California, San Francisco

16 **The Bridge from Genes to Proteins**

Exploring Protein-Folding Free-Energy Landscapes
Charles L. Brooks III, Scripps Research Institute



PROTEIN AND NUCLEIC ACID SEQUENCE ANALYSIS

18 **Proteins in Profile**


Profile Analysis: Evidence for a Two-Helix Bundle in Human CCK-58
Grace Rosenquist, University of California, Davis
Hugh Nicholas, Pittsburgh Supercomputing Center



UNDERSTANDING THE BRAIN

20 **Watching the Brain in Action**

Functional Mapping of the Cortex with Conventional MRI Scanners
Jonathan Cohen, University of Pittsburgh and Carnegie Mellon University
Nigel Goddard, Pittsburgh Supercomputing Center



EARTHQUAKE MODELING

22 **Getting Ready for the Big One**

Earthquake Ground Motion Modeling on Parallel Computers
Jacobo Bielak, Carnegie Mellon University

MANUFACTURING PROCESS DESIGN

24 **Turn, Turn, Turn**

Parallel Computation of Rotor-Stator Interaction in Turbines
Paul Cizmas, Westinghouse Science and Technology Center
Ravi Subramanya, Pittsburgh Supercomputing Center

26 **Cars Lite**

Modeling of Aluminum Automobile Components at ALCOA
Edmund Chu, ALCOA Technical Center

UNDERSTANDING THE EARTH

28 **Storms in Space**

Simulation of the January 10-11, 1997 Magnetic Cloud
Charles Goodrich, University of Maryland

EVOLUTION AND STRUCTURE OF THE UNIVERSE

30 **Ripples in Space and Time**

Stable Evolution of 3-Dimensional Single-Black-Hole Spacetimes
Jeffrey Winicour, University of Pittsburgh


Watching a Protein Fold

The protein-folding problem – one of the major challenges of molecular biology in the 1990s – could be thought of as a version of cryptography. Scientists are the code breakers, trying to uncover a set of rules embedded in a protein's sequence of amino acids. This chemical alphabet of 20 characters, strung like beads on a chain along the protein's backbone, carries a blueprint that specifies its folded shape.

Within microseconds after rolling off the protein assembly line (in the cellular ribosome), the stretched-out chain twists itself into a bundle, with turns, helices, sheets and other 3D features. For proteins, function follows from form – the grooves and crevices of its complex folds allow it to latch onto other molecules and carry out its biological role.

But how is it that a particular amino-acid sequence results in a particular folded shape? "The protein-folding problem is still the single most exciting problem in computational biochemistry," says Peter Kollman, a pharmaceutical chemist at the University of California, San Francisco, and a leader in using "molecular dynamics," computational simulations that track the minute shifts of a protein's structure over time. "To be able to predict the structure of the protein from just the amino-acid sequence would have tremendous impact in all of biotechnology and drug design."

In 1997, Kollman and his coworkers Yong Duan and Lu Wang used the CRAY T3D at Pittsburgh to develop software that exploits parallel systems like the T3D and CRAY T3E much more effectively than before. They then simulated the folding of a small protein in water for a full microsecond, 100 times longer than previous simulations. The result is a more complete view of how one protein folds – a look at what hasn't been seen before, offering new insight into the folding process.



Exploiting Massive Parallelism

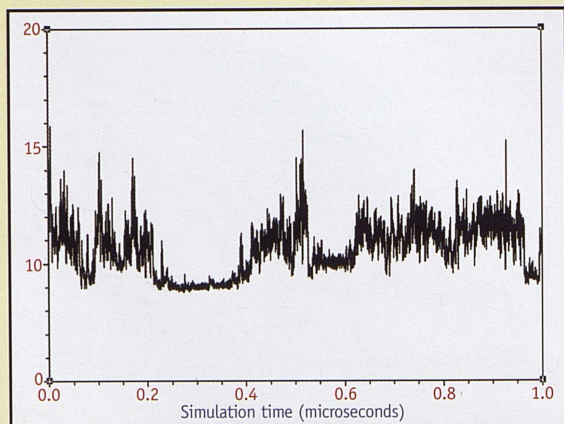
Simulating a millionth of a second of protein movement may sound less than impressive, until you realize that the longest prior simulations of similar proteins extended only 10 to 20 nanoseconds (billionths of a second). The limitation holding back this critical work has been the tremendous computational demand of the simulations, which must account for interactions between each atom in a protein and all the other atoms and surrounding water molecules.

The Kollman team's recent effort focused on a 36 amino-acid protein, the villin headpiece sub-domain. With surrounding water molecules, the computation involved about 12,000 atoms. To make this daunting computation possible, Duan used Pittsburgh's T3D to devise and test changes to the molecular dynamics part of AMBER (a widely used package developed by Kollman's research group). His work boosted single-processor performance about 70% and greatly improved the "load balancing" and communication among processors — resulting overall in 256-processor performance six times faster than before.

"Virtually all other molecular dynamics codes in the literature level off at 40, 50 or 60 processors," says Kollman. "In other words, even if you use all the processors, you don't get any faster because the communication among processors is rate limiting." With Duan's changes, the software now runs 170 times faster on 256 processors than on one alone.

"This was a tour-de-force of parallel programming," says Kollman, "and it wouldn't have been possible except for Pittsburgh making the T3D available to us."


"The dedicated T3D allowed me to conduct extensive tests on a variety of plausible schemes," says Duan, who cites training at a PSC parallel programming workshop and discussions with PSC scientist Michael Crowley as also being instrumental to this work.



Above: A snapshot from protein-folding simulations of the villin headpiece subdomain.

Left: The vertical axis shows the protein's radius of gyration (in angstroms) as it changes over time. The "quiet period" occurs from about 250 to 400 nanoseconds.

Tracking a protein's movement for a full microsecond gives new insight into the protein-folding problem.



Proteins are the action superheroes of the body. As enzymes, they make reactions go a million times faster. As versatile transport vehicles, they carry oxygen and antibodies to fight disease. They do a thousand different jobs. But before a protein can go to work, it must fold into the right shape.

Quiet Time for Protein Folding

What did the researchers learn from a full simulated microsecond of protein folding? Over the first 200 nanoseconds, the protein quickly collapses and then moves back and forth between compact states and more unfolded forms.

The next 800 nanoseconds reveal an intriguing quiet period. From about 250 to 400 nanoseconds the fluctuating movement back and forth between folded globules and unfolding virtually ceases. "For this period in the later part of the trajectory," says Kollman, "everything becomes quiet. And that's where the structure gets closest to the native state."

For Kollman, this behavior suggests that folding may be characterized as a searching process. "It's a tantalizing idea – that the mechanism of protein folding is to bounce around until it finds something close, and stay there for a period, but if it isn't good enough it eventually leaves and keeps searching."

Although only a partial glimpse – even the fastest proteins need 10 to 100 microseconds to fully fold – these results represent a major step forward in protein-folding simulation. With the promise of more powerful computing and higher level parallelism, Kollman sees the goal of simulating the entire folding process as within reach. "We believe our work marks the beginning of a new era of the active participation of full-scale simulations in helping to understand the mechanism of protein folding."

This ribbon representation shows the protein's simulated structure (blue) during the marginally stable period (about 300 nanoseconds) in comparison to its native structure (red) as determined from NMR.

Peter A. Kollman (middle) and Yong Duan, University of California, San Francisco.

The Bridge from Genes to Proteins


Since 1953, when Max Perutz and John Kendrew figured out the structure of hemoglobin and myoglobin, research has steadily added to the knowledge base of protein structures. After 45 years scientists have solved about 7,500 proteins, each representing months or years of work. The human genome has about 100,000 protein-coding regions – DNA sequences that are blueprints for the cell to manufacture proteins – which means there's more than 90,000 proteins left to solve (and that's only the *human* genome).

In the next few years, this sequence data will become available. How can structural biologists keep pace? Solving the protein structures is a major part of realizing the potential offered by genome data for advances against disease, but simple math tells you unless something changes this work could take halfway to the next millenium. "The challenge as a result of the genome mapping projects of the last and present decade," says Charles Brooks of the Scripps Research Institute, "is to deduce protein function starting from genomic sequence."

This emerging field of activity is called "structural genomics," and Brooks has a keen appreciation for what's involved. Over a 20-year period – at Harvard, Carnegie Mellon and Scripps – he's used computational simulations to better understand the complex relations between sequence and structure. His work on the protein-folding problem, as this relation is known, has helped to build theory that explains protein folding in terms of the energy changes of the atom-to-atom interactions. At this juncture, says Brooks, with a tidal wave of sequence data on the way, what's crucial is computational resources: "The complexity of the simulations involved outstrips current computational infrastructure."



Charles L. Brooks III,
Scripps Research Institute



Two partially folded states and the native state (right) of a segment of streptococcal protein G. This segment's 3D structure is comprised of an alpha-helix (purple) between two beta sheets (yellow).

With a tidal wave of sequence data on the way, what's crucial is computational resources.

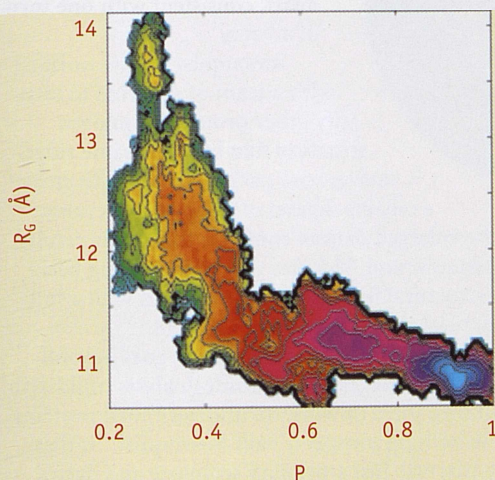
More information:
<http://www.psc.edu/science/brooks98.html>

Structural Genomics and Forcefield Models

“Key to the success of any computational method that aims to provide sequence to structure predictions,” says Brooks, “is that the energy function used to represent the biological system yields the native structure of known proteins as its lowest free-energy state.” Computations can verify the model’s “forcefield” – the mathematical expression of energy relations between the atoms, which the model uses to compute the forces acting on the protein’s structure, moving it from one folded shape to another. Researchers have carried out such verifications for a handful of proteins, but much more must be done.

In structural genomics, these forcefield models – such as CHARMM, widely used protein simulation software Brooks helped develop – could fill the gaps left by other computational methods. In an approach called “inverse folding,” relatively low-resolution models, which run more quickly and cheaply than detailed forcefield models, search the database of known structures, comparing sequences with recurrent folding patterns, to produce structures that are more or less in the ballpark (within three to five angstroms deviation from the native state).

Free Energy Landscape of Protein G



The protein’s radius of gyration (vertical), how far it spreads out from its center, is plotted against the fraction of its interactions that correspond to native state (horizontal). Color contours show free energy, from the lowest (deep blue) of the native state to other (blue to pink to red to orange to yellow to green) less favorable states. The L-shape of this landscape indicates a folding process of early collapse followed by protracted “search” through possibilities to arrive at native structure.

Starting with these low-resolution structures, forcefield models can then refine the data, with the aim of producing structures as accurate as those now produced with x-ray crystallography and NMR methods (precise within less than one angstrom). “Only once we’ve achieved such resolution,” notes Brooks, “can we be confident in using these structures for drug discovery and to assess the protein’s biological function.”

The critical test for a forcefield model is to map how the energy of a protein changes over the course of folding. Brooks’ research group has carried out several such studies. In recent computations at PSC, Brooks and Felix Sheinerman mapped a segment from a small protein, streptococcal protein G, that structurally combines helices and a folded sheet.

The Energy Landscape of Protein G

Explorations of a protein’s energy landscape are very large computational problems, but they represent an alternative to even more costly simulations that track the protein over the complete time span of its folding. Brooks and Sheinerman’s protein G computation involved about 100 gigabytes of data and the equivalent of a month of computing on 512 processors of the CRAY T3E. Experiments suggested that early folding interactions formed a single core that nucleated the rest of the folding reaction, but the simulation suggests otherwise.

“We’re able to look in detail,” says Brooks, “in that region of the energy landscape. What we find is that those points of interaction identified in the experiment as a single nucleus form as several different nucleating points.” By providing deeper insight, the simulation points the way to further experimental studies – underway now – that can more precisely characterize the folding interactions.

“Our understanding in this area is nascent,” says Brooks. “We need to significantly expand the number of proteins we’ve studied computationally, and we need resources to simulate larger proteins. Infrastructure for the investigation of protein-folding free-energy landscapes must be increased or developed immediately if we’re to make a timely impact in the quest to go from sequence to biochemically relevant structure.”

Proteins in Profile



Hugh Nicholas, PSC, and Grace Rosenquist, University of California, Davis. Their collaboration, begun through a 1989 PSC workshop, led to a new approach to sketching the rough outlines of undetermined protein structure.

To live, we digest food. To digest food, we need enzymes. To release digestive enzymes, we need hormones, which go from the small intestine to the pancreas and say, in biochemical language: "Food's here, time for work." One of the hormones that delivers this wakeup call is cholecystokinin, better known to biologists as CCK.

Grace Rosenquist of the University of California, Davis and Hugh Nicholas of Pittsburgh Supercomputing Center recently applied a novel approach to some questions about CCK, with surprising results. CCK is a peptide hormone, a chain of amino-acids, therefore a protein. Biologists for many years thought peptide hormones were a special brand of protein; they seemed to lack the helices, sheets, turns and combinations of these 3-D features that structural biologists call secondary and tertiary structure. They were thought to have primary structure – a linear sequence of amino acids – and that's about it.

The work of Rosenquist and Nicholas suggests that for CCK there's more to it than that. Their results offer credible evidence that a sequence of human CCK, CCK-58, has structure – two closely packed helices. Their research method itself, furthermore, suggests that profile analysis could cut the time and expense involved in determining protein structure, grinding work that is one of the important ongoing tasks of molecular biology.

Protein Carpentry

Helices near each other along the same stretch of peptide often link by folding together like interlaced fingers. Water-avoiding (hydrophobic) sidechains of one helix pack together in the core of the interface with like-minded sidechains from the other, warding off the cell's watery environment. "CCK-58 showed these kinds of signals," says Nicholas, "the right periodicity and frequency of amino-acids, consistent with packed helical structure."

To go beyond these hints, Rosenquist and Nicholas turned to research by Cambridge (England) structural biologist Cyrus Chothia. In the 1970s and early 80s, working with a set of eight packed-helices (from five proteins),

Chothia showed that hydrophobic sidechains on each helix form a diamond-shaped pattern that aligns the sidechains in ridges separated by parallel grooves. The ridges of one helix fit into the grooves of the other, forming a tight joint analogous to tongue-and-groove carpentry.

Profiles of Structure

Profile analysis is often used to analyze amino-acid sequence among proteins related to each other by evolution. Nicholas and Rosenquist, however, applied it in a new way, looking for structural relationships where there's no evolutionary relationship.

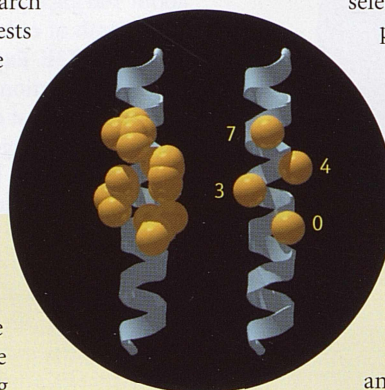
They built sequence-based profiles from Chothia's helix-pairs and used them – applying a classic strategy called "leave-one-out-analysis" – to see if profiles from seven of the helix-pairs could predict helical structure in the remaining sequence. The answer was yes. Chothia's work showed that one pattern of hydrophobic sidechains (called +3 helices) always packs with another (+4 helices). The profiles picked all the +3 helices as +3s and, likewise, all the +4s as +4s.

Encouraged by this high degree of selectivity, the researchers turned to a profile-analysis program, ProfileSS, developed by Nicholas and colleagues at PSC, to apply the helical profiles to CCK-58. The highest scores occur at positions three to four amino-acids apart, consistent with one turn of a helix.

Rosenquist is now working with a team of students to identify other proteins of known structure that have the same ridge-and-groove packed-helices – to create a reliable dataset of 50 to 60 structures.

If the enlarged dataset confirms their early results, Rosenquist and Nicholas will have a strong basis for structural biologists to go forward with a full structural determination of CCK-58.

The real promise of their work goes beyond CCK-58. It suggests that profile analysis could save countless years of work on undetermined proteins. "The exciting prospect," says Rosenquist, "is that you can take just a primary sequence and derive information about structure. You can predict what part of the protein is going to be not just helical, but also packed, and packed in a certain way. And if it works for helix-packing, then we can look at other kinds of structure. This kind of knowledge saves time, and time is money."



Graphical representations of the diamond-shaped pattern of hydrophobic amino-acid sidechains for one helix of a packed helical structure. The diamond shows in the alignment of beta-carbon atoms (right), the atoms closest to the helix. This pattern represents a +4 helix, so called due to the increment of amino acids between those forming two parallel ridges (relative positions 0 & 4 and 3 & 7). The representation (left) shows the full conformation of the sidechains, which form a groove between the parallel ridge surfaces.

A new approach to sketching

the outlines of protein structure.



This graphic, from a PSC educational video, illustrates helix-helix packing, showing a helix-turn-helix motif in the protein lysozyme. Sidechains (green) of one helix (blue) define a ridge that fits into the groove defined by the sidechains (yellow) of the other helix (red). Other sidechains (cyan & orange) participate in the helix-helix contact but not in the ridge-groove interaction.

Watching the Brain in Action

Pittsburgh scientists have created a powerful new technology for viewing the brain at work. By linking an MRI scanner with the Pittsburgh Supercomputing Center's CRAY T3E via high-speed networks, they convert scan data almost instantaneously into an animated 3-D image showing what parts of the brain "light up" during mental activity.

Typically, it takes as much as a day after gathering scan data to produce high-quality 3-D images. The Pittsburgh team – scientists at Carnegie Mellon University, University of Pittsburgh Medical Center and the Pittsburgh Supercomputing Center – cut this delay to a matter of seconds.


"This technology will make it possible to use brain-mapping as a clinical tool in diagnosis and treatment of brain pathology," says Dr. Jonathan Cohen, who codirects the Laboratory for Clinical Cognitive Neuroscience, a joint venture of the University of Pittsburgh and Carnegie Mellon. Real-time capability will aid neurosurgeons in precision surgical planning, and it can be used to test and diagnose cognitive dysfunctions such as schizophrenia, amnesia and epilepsy. With high-speed networking, doctors at locations distant from the MRI scanner can actively consult in patient testing.

Mapping the Brain

For several years, Cohen and his colleagues have used a technique known as functional MRI (fMRI) to do "brain-mapping" experiments that investigate and map the brain regions involved in a particular kind of memory activity known as "working memory." Data from an MRI scanner shows what sites in a subject's brain are active during mental activity. These experiments generate huge amounts of information quickly, and initially it took days to process the data into a high-resolution 3-D image.


To eliminate this bottleneck, the researchers turned to PSC. Carnegie Mellon statistician William Eddy and UPMC physicist Doug Noll worked in collaboration with Goddard and PSC research programmer Greg Hood to exploit the CRAY T3E, a highly parallel system that divides the computing among many processors. In November 1996, the researchers reduced processing time so that a realistic 3-D image of the brain could be viewed live, while the subject was in the scanner, with a delay between mental activity and image availability of about six minutes. The team has now cut this delay to seconds, and they are working to get it under a second, which will allow an improvement in image quality.

Combining MRI, high-speed networks and supercomputing produces a realistic picture of brain activity in seconds.



What parts of the brain are involved in seeing?

This 3-D image shows that the region known as the "primary visual cortex" becomes active when a subject inside an MRI scanner looks at a checkerboard pattern.




In demonstrations of this real-time brain-mapping capability, a test subject – one of the researchers – lies inside an MRI scanner at UPMC and performs a simple mental task. The MRI scanner records data from her brain and transmits it via high-speed network to the CRAY T3E, which converts the raw fMRI data into 3-D images, which compensates for head movement and identifies active areas of the brain. From the T3E, the data travels to a remote location via high-speed network, where observers see the subject's brain as a translucent animation showing what regions "light up" as she does the mental task.

Movies of the Brain at Work

The MRI scanner records data from a region of the brain and transmits it via high-speed network to PSC's CRAY T3E, where a series of computations converts the raw MRI data into a 3-D animated format. First, mathematical data from the scanner is transformed to a representation of physical space. Next, the data is registered to compensate for head movement. Finally, the data is statistically analyzed to determine regions of activity. The 3-D representation is transmitted via high-speed network and "volume rendered" for screen display on the visualization platform.


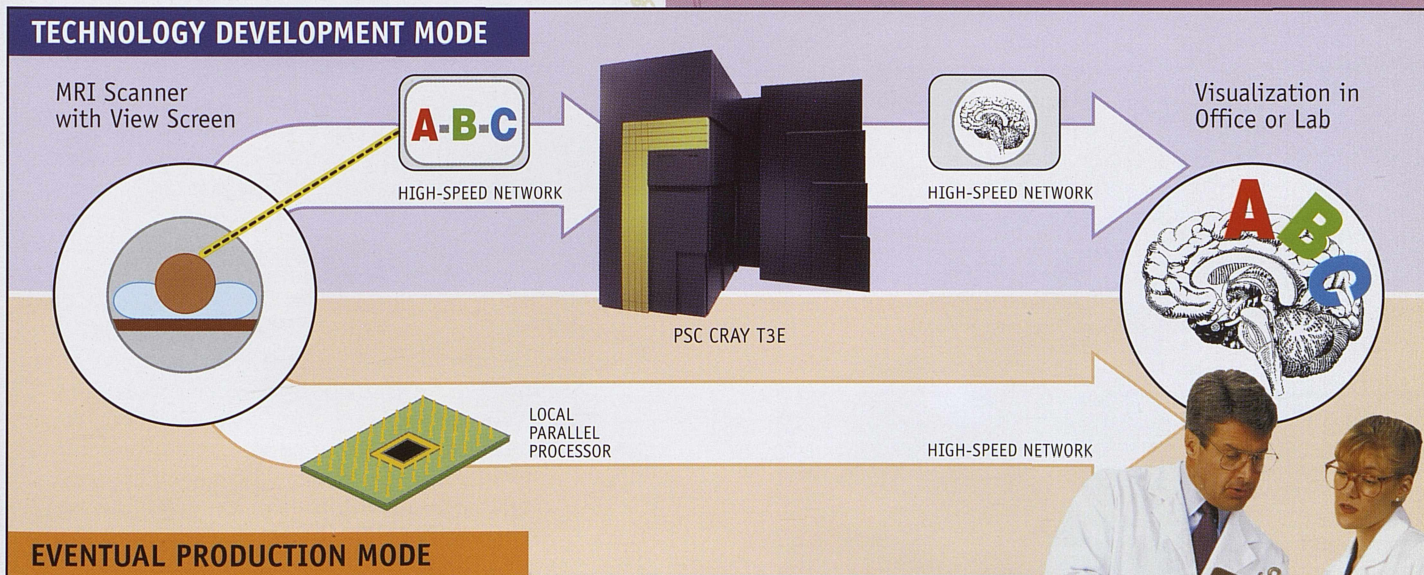
With further improvements in computing and in high-speed networks, the researchers expect to cut the delay between scanning and imaging to under a second. This will make it possible to compensate for movement of the subject's head via a control loop in the scanner, producing better images. Processing now done on the T3E will eventually be incorporated in the MRI scanner, allowing clinical use. New methods, further enhancements, will continue to be tested in the supercomputing environment. By providing a laboratory to develop new software and ideas, PSC opens a window to the future of brain imaging.



From left to right: Nigel Goddard, PSC biomedical scientist, Chris Genovese, Carnegie Mellon, Bill Eddy, Carnegie Mellon, Jonathan D. Cohen, University of Pittsburgh and Carnegie Mellon, Douglas Noll, University of Pittsburgh, and Greg Hood, PSC.

The Pittsburgh Supercomputing Center has taken real-time brain-imaging from idea to prototype demonstration by providing:

- algorithm development
- software optimization
- advanced networking
- scientific visualization
- flexible data archiving
- collaborative focus
- high-performance computing



Getting Ready for the Big One

Mexico City, 1985 — 300 collapsed buildings, up to 10,000 people dead. San Francisco, 1989 — \$5 billion in property loss, 4,000 injured, 62 dead. Calamities like these have led scientists to reevaluate what they know about building structures to withstand major quakes. Research in the wake of these earthquakes and others has shown that local geophysical factors such as basin shape and soil type play a bigger role than previously thought. Because of these factors, the same seismic wave that flattens one block can leave the next unscathed.

“Ground motion can vary significantly over short distances,” says Jacobo Bielak, professor of civil and environmental engineering at Carnegie Mellon University. During the 1989 Loma Prieta quake, for instance, earth movement differed significantly at places only 30 meters apart. How is it possible to account for these factors in building design and construction?

A National Science Foundation “grand challenge” team of CMU engineers and computer scientists, led by Bielak, in collaboration with seismologists from the University of Southern California and the National University of Mexico, is addressing this problem. They are developing 3D computer models that realistically depict how earthquakes shake the Earth.

“Our goal,” says Bielak, “is to develop the capability to predict, by computer simulation, the ground motion of large basins during strong earthquakes. If a large earthquake strikes Los Angeles, which regions will be worst stricken? Which seismic frequencies will most amplify the surface ground motion? Answers will guide development of more rational seismic provisions for building codes, leading to safer, more efficient, economical structures.” Using PSC’s CRAY T3D, Bielak’s team generated the most detailed simulations yet developed on seismic response in the San Fernando Valley.

Archimedes

The Los Angeles Basin is an important problem because of its size and large population and also due to its computational complexity. The model must incorporate information about an earthquake’s source, its propagation path and the site conditions of a large sedimentary basin. “Resolving excitation frequencies of engineering interest,” says Bielak’s colleague Omar Ghattas, “pushes the limits of scientific computing. A big part of our work has been developing numerical algorithms that can harness the power of scalable, parallel systems like the T3D and T3E.”

CMU computer scientist David R. O’Hallaron and colleague Jonathan Shewchuk developed a set of software tools, called Archimedes, that helps automate the modeling process on parallel systems. From an input model of basin geology, Archimedes subdivides the region into computational cells. It partitions the computations for a parallel computer by assigning groups of cells to different processors in a way that minimizes communication.

For its initial effort, the team focused on the San Fernando Valley, a smaller problem than the LA Basin. With input from material models developed at San Diego State University, Archimedes subdivided a region 54 by 33 kilometers and a depth of 15 km into an unstructured 3D mesh of 77 million cells. Archimedes tailors the grid to the local wavelength, 20 times shorter in the softest soil than in the rock that surrounds the valley, so that grid spacing varies from 340 meters in the rock to 21 meters in the softest soil. Even with this wavelength-adaptive methodology, the large scale of the problem requires 13.4 million grid points, resulting in 40 million equations — probably the largest unstructured mesh problem ever solved, says Ghattas. This adaptive approach will facilitate building models for other earthquake-prone regions, such as Mexico City, Salt Lake City and Tokyo.

A Simulated Aftershock

With the San Fernando Valley model, the team simulated an aftershock of the 1994 Northridge earthquake. The model propagates seismic waves through the basin, simulating motion within the basin as well as on its surface. Each simulation requires about seven hours of clock time on 256 T3D processors, using 16 gigabytes of memory, to simulate 40 seconds of shaking.

Their results reveal the tendency of the basin to respond in a complex pattern, with ground motion nine times greater at some sites than others.


“Predicting damage is more complicated than just looking at peak displacements or even the complete history of the ground motion,” says Bielak. “If frequencies at a point on the surface correspond to the fundamental modes of a structure, the potential for damage is greater. Knowing how simple structures with different natural frequencies would respond to the ground motion throughout the valley, one can get a much better idea of the anticipated damage.” The researchers are processing and analyzing these results.

“The CRAY T3D has been unique in providing the level of performance needed for the San Fernando Valley,” adds Bielak. “It allows us for the first time to get to the level of detail we need to realistically simulate the frequencies that do the most damage. Our target application, the entire Greater Los Angeles Basin, is an order of magnitude larger and must be run for multiple scenarios to assess the seismic hazard. The expanded capability of the CRAY T3E gives us hope that this problem is within reach.”

Left to right: Jifeng Xu, David O’Hallaron, Omar Ghattas, Loukas F. Kallivokas, Jacobo Bielak and Hesheng Bao. Jonathan Shewchuk could not be present for this photograph.




The most detailed
simulations yet
developed on
seismic response
in the
San Fernando Valley.



Software called Archimedes partitioned the San Fernando valley into an unstructured mesh, with smaller cells in softer soil, where seismic wavelengths are shorter. This mesh is partitioned into 64 subdomains, indicated by color.

A simulated earthquake in the San Fernando valley. Color depicts the peak magnitude of ground displacement. The simulation covered a 54 x 33 kilometer area, superimposed here on a satellite view of topography, to a depth of 15 km. Los Angeles can be seen in the southeast corner.



PSC video animation of this quake, of which this image is one frame, won the award for Best Database Animation at Siggraph '97, the annual conference of the computer-visualization industry, and it was the only scientific animation selected for showing at the Siggraph '97 Electronic Theater.

Turn, Turn, Turn

Turbine Simulation at Westinghouse Science and Technology Center

Electrical power generation is a multi-billion dollar global business. To gain an edge in this fiercely competitive market, the key for companies like Westinghouse is more efficient turbines. These huge jet-engine like machines do the heavy-duty work of converting raw energy into megawatts of electricity, and even slight improvements in efficiency translate into significant reductions in the cost of generating power.

Can high-performance computing help design more efficient turbines? That question confronted senior scientist Paul Cizmas of the Westinghouse Science and Technology Center in 1996. Existing software was sequential, based on a single-processor computing paradigm. It simulated fluid flow through the blades inside the turbine by calculating the aerodynamics one blade at a time. To take advantage of systems like PSC's 512-processor CRAY T3E, which puts many processors to work simultaneously on the same job, it was necessary to "parallelize" the software.

After attending a PSC parallel-processing workshop in June 1996, Cizmas discussed the problem with PSC consultant Ravi Subramanya. They began collaborating in January 1997 and four months later had working parallel code. "This was a change of paradigm," says Cizmas, "a big step. We needed more realistic simulations and faster turnaround. Jobs that would have taken three months before, now run in under 12 hours."

Part of this impressive performance gain is due to "superlinear" speedup – on test cases, 10 processors together run 15 times faster than one by itself. This seemingly impossible result occurs because parallelizing the turbine aerodynamics, assigning a processor to each turbine blade, improves how data is handled in memory, an advantage that becomes more and more significant with larger simulations. Cizmas has begun tackling these larger problems, which will offer new understanding of how shape and arrangement of the blades affect turbine efficiency.

Democratic Computing: One Processor, One Blade

Modern turbines are comprised of alternating rows of stationary blades (stators) that direct flow into the rotating blades (rotors). Before Cizmas and Subramanya went to work, a test case involving one-and-a-half stator-rotor stages – an eight-blade configuration – would take three months of computing on a single CRAY C90 processor. This made it unrealistic to even think about simulating production turbines, often involving three or more stages with 150 or more blades.

To parallelize this software, Subramanya assigned each blade to a separate processor, thereby avoiding a major problem with the sequential version. “The sequential code ran slowly to begin with,” says Subramanya, “because whenever the calculation shifted from one blade to another, all the data in the processor cache became invalid. It lowered an almost new dataset into memory each time it moved to the next blade.”

Simulations & Design

Larger simulations now underway on the T3E, involving actual turbine blade configurations, should improve turbine design in several ways. Experiments have shown that turbine efficiency varies according to the radial spacing between blades. “A wake sheds from the trailing edge,” says Subramanya, “and interacts with successive rows. This interaction can be constructive or destructive. If it’s destructive, then changing the relative position of a blade can improve efficiency.”

In current and future work, Cizmas expects to extend the new software from 2D to 3D simulations, and to use it to optimize the shapes of turbine blades, a factor that also – along with relative positioning – can affect efficiency. “We are no longer limited,” says Cizmas, “in the number of blades and rows we can simulate. Time is no longer a problem. Our only restriction is the number of processors. From my point of view, what I’m waiting for now is parallel systems with more processors.”



Simulation of a one-and-a-half stage turbine. Color corresponds to temperature. Such simulations pinpoint likely blade hot spots, which helps in the design of cooling conduits and blade coatings.

Jobs that would take three months before now run in under 12 hours, with more accurate results.


This difference, especially significant because of the time-dependent nature of the computation, accounts for the superlinear speedup. A set of mesh-like computational grids for each blade keeps track of fluid properties at that blade, and this data updates with each advance in time. Although one processor doesn’t have enough memory to hold the grid data for 10 blades, it can handle the grids for one blade, greatly reducing the need to swap data back and forth from disk to memory as the computation advances in time.

On an SGI Challenge, for a test case the parallel code runs 15.6 times faster using 10 processors than with one. The CRAY T3E delivers similar speedups, and will perform markedly better than the SGI Challenge, notes Subramanya, for larger simulations that require more processors than available on the SGI machine.

The parallel software is not only faster, it’s also more accurate – due to modifications built into the new code. “In the serial code,” says Subramanya, “you’re using an outdated boundary condition part of the time. This is not physically accurate, but you were forced to do it, because you had to compute sequentially. In parallel, we solve everything at the same time, so we get more realistic results.”

Paul Cizmas (left), senior scientist, Westinghouse Science and Technology Center, and Ravi Subramanya, senior computational science consultant, Pittsburgh Supercomputing Center.

A paper they co-authored, presenting results from their parallel turbine simulation software to an international conference in Stockholm, Sweden, was the Science and Technology Center Technical Paper of the Year.



Cars Lite

As the world's largest aluminum company, Alcoa knows what aluminum can do, and it helped develop the Prowler as well as the Audi A8, a German-produced aluminum car making waves in the upscale market. Both cars have an aluminum skeleton, a "space frame," similar to those used in aircraft, with aluminum panels for their sleek skin. Alcoa automotive engineers like Edmund Chu see these cars as harbingers of the future in automotive engineering.

The key, notes Chu, is computational modeling. Alcoa has sophisticated ability to do computer simulations that predict how aluminum will perform, reducing costly prototyping and trial-and-error processes. "The auto industry tends to look at aluminum's cost per pound," says Chu, "which is substantially higher than steel. But aluminum sheet weighs half as much as steel; cost per pound isn't an appropriate measure of the economics. We emphasize overall cost, and we're working aggressively to make aluminum easier to use."

Alcoa has partnered with PSC since 1987, and Alcoa engineers used PSC resources in designing parts for the Prowler. "Automotive products are a key part of Alcoa's future," says Peter Bridenbaugh, Alcoa executive vice-president of automotive structures. "Scientific modeling on the supercomputer allows us to solve time-critical problems in a competitive manner, which is particularly important in automotive design."

Tool & Die: The Inner Hood

A large part of the development cost for body sheets, such as hood and door panels, is designing the "dies" used to stamp sheet-metal parts in mass production. Chu leads Alcoa's effort in this area. Traditionally, sheet-metal forming relies on the ingenuity of tool-and-die craftsmen, who have sophisticated artisan's know-how for working with steel sheets.

This approach has its limits, however, with new materials like aluminum and with the complex geometries of modern automobiles, that often require many trial-and-error iterations. Using computing, mathematical tools can predict the effectiveness of a design before casting the die and trying it out, potentially saving hundreds of thousands of dollars and weeks of time. And Alcoa's modeling ability with aluminum, says Chu, is more advanced than similar techniques with steel.

As an example, he points to recent work of his engineering group on the underbody of a hood, known as an "inner-hood panel." Traditionally, these panels employ steel with a "beam" design to give strength and rigidity. Cut-outs in the flat part of the sheet reduce weight, but require an extra die and press step, adding production cost. To exploit the unique properties of aluminum, Alcoa developed a "multi-cone" design.

"You don't want to force aluminum to behave like steel," says Chu. "You want to build aluminum characteristics into your design." This meant using a lighter gauge than is possible with steel, leading to the multi-cone design, which gives structural integrity and rigidity equivalent to a steel panel at half the weight, and without cutouts, avoiding potentially millions of dollars in manufacturing cost.

Initial modeling of the panel predicted several locations of high strain. To check the modeling, a die was cast and a Detroit "stamping house" stamped the panel. Analysis confirmed the predicted high-strain regions. Alcoa modified the cone shapes and shifted to a more formable alloy – much easier to do, notes Chu, with simulations than trial-and-error prototypes.

"By integrating material design with process design through computational modeling, we can select the optimum alloy to maximize formability of the part. When we run the model, we can try a number of different alloys." Alcoa's ability to do this relies on its expertise in advanced scientific computing.

Supercomputing and Aluminum Product Design

Alcoa has developed a highly accurate "constitutive model" for aluminum – a mathematical description that relates the microstructure of the metal to how it behaves when formed into a product. Alcoa's partnership with PSC has allowed it to refine this model to its current high degree of accuracy. "This is one of our strengths," says Chu. "We have the ability with simulations to describe all complex loading conditions."


Supercomputing, furthermore, because it gives fast turnaround, makes it feasible to consider multiple possibilities, providing unprecedented design flexibility. "With the Cray," says Chu, "we can look at five or six scenarios all at once. The turnaround is six times faster than on our own workstations, and this is critical in the design stage, where you need to make changes quickly."

In the future of design with aluminum, says Chu, engineers will create new alloys to meet product requirements. "For a particular product, I want to find out what governs the deformation. Is it strain hardening or elongation or something else? With the model, I can put in a design parameter and different material behavior, fictitiously – make up some new alloy, and then we could turn around and, based on the model predictions, say 'Hey, we can develop this alloy.'"



Edmund Chu, senior technical specialist, Alcoa Technical Center.


Below: Aluminum space-frame for the Prowler.





Alcoa engineers used supercomputing to design parts for the Plymouth Prowler.




An engineering tour de force—that's what car buffs are saying about the Plymouth Prowler. This update of a 1950s hot rod is a beauty to look at, and for the discerning its beauty is more than skin deep. The Prowler is the first U.S. car engineered from the ground up to exploit aluminum technology. At 2,800 pounds, 50% less than it would weigh with traditional steel design, the Prowler is quick, fuel efficient and impervious to rust.



This inner-hood panel employs the multi-cone design developed by ALCOA engineers. Computer simulations of thickness distribution predict high-stress areas prior to casting the sheet-metal die, substantially reducing design costs.




Storms in Space



In these visualizations from simulations at PSC, Earth's magnetosphere is "relaxed" prior to the onset of the January 10, 1997 magnetic storm (above).

Three hours later (below), with the arrival of the magnetic cloud, the magnetosphere has transformed. The translucent surface represents the shape and extent of Earth's magnetic field. The colormap indicates plasma density and shows the sunward shock wave around the Earth.

On January 6, 1997, the sun burped, in a manner of speaking, and a billion-ton bubble of hot gas came blasting toward Earth at about a million miles an hour. Four days later, pushing a pileup of solar magnetic field called a magnetic cloud, the bubble rammed the nose of Earth's magnetic field. While most of us were blissfully unaware, the resulting magnetic storm put on a show for space physicists.



"In space science, this was a simulation of

A satellite orbiting 350 miles above the Earth recorded an electrical burst 10,000 times normal intensity. On January 11, the dense trailing edge of the bubble delivered a parting shot to the magnetosphere – the outer boundary of Earth's magnetism, and the energy dissipation in the auroras over the North and South Poles climbed to 1,400 gigawatts, nearly double U.S. power-generating capacity.

This display of space weather also demonstrated how gusts of solar wind – the stream of charged particles from the sun – can play havoc with our information age infrastructure. On January 11, a \$200 million communication satellite, Telstar 401, fell silent, forcing network broadcasts to switch to another satellite.

Severe magnetic storms have occurred before, but January 1997 was the first since a fleet of satellites was in place to capture data. By January 17, Chuck Goodrich and colleagues at the University of Maryland Space and Plasma Physics Group (SPP) were using this data in an unprecedented computational study.


"This event had intense scientific interest," says Goodrich. "NASA asked us to simulate it, and thanks to PSC we were able to provide video animations on very short notice. In space science, this was a simulation of unprecedented proportion." Running on four processors of PSC's CRAY C90, Goodrich and his SPP colleagues John Lyon of Dartmouth, Dennis Papadopoulos and Mike Wiltberger simulated over 45 hours of storm time, from the arrival of the shock front preceding the bubble through the next day and a half. Within a week, they had the first global view of a geomagnetic storm.

The Big Picture

First observed in the early 1970s, these eruptions of solar gas occur as often as once a day. Even the most severe gusts present no known health hazard, since Earth's magnetosphere deflects the mass flow. But some of the flow energy transmits through the magnetosphere and shows up as bursts of electromagnetism in the ionosphere, hence the Northern and Southern lights, and the potential to seriously disrupt power and communications systems.

The advance guard for detecting these blasts of solar wind is satellites that hover about 600,000 miles out on the sunward side of Earth. Earth-orbit satellites closer in gather data on near-Earth effects.

Charles C. Goodrich,
University of Maryland,
Space and Plasma Physics
Group



unprecedented proportion.”

Invaluable as it is, the satellite data represents pinpoints in a huge volume. With supercomputing, the SPP group connects the dots and produces a big picture view.

“The basic question,” says Goodrich, “is how energy from the plasma flow penetrates the magnetosphere and gets absorbed in the inner atmosphere. Until recently, we knew observational effects, but we didn’t know the basic physics. We have theories, but it’s been impossible to check them. Simulations help us understand the detailed physics.”

The SPP numerical model simulates a solar storm transferring energy through the magnetosphere to Earth’s ionosphere. Satellite data initializes the model, and a set of magneto-hydro-dynamic equations calculates effects throughout the magnetospheric volume, represented as a grid in the stretched-out spherical shape of the magnetosphere. The simulation then maps electrical current and velocity through the ionosphere.

The Unforeseen Effect of Plasma Density

The SPP simulations agree well in key details with Earth-orbit satellite data, confirming basic accuracy of the model. Bursts of released energy in the ionosphere happen close in time to when they were recorded at near-Earth satellites and

enter the ionosphere in the same place shown in ultraviolet images from POLAR, a satellite orbiting above the Arctic.


The simulation also shows the first significant energy transfer occurring after a southward turn in orientation of the solar-wind magnetic field. Physicists believe that energy transfer happens predominantly when this occurs, since southward orientation of the solar field is opposite to the northward orientation of Earth’s magnetosphere. The opposed polarities, in theory, allow the magnetic field lines to connect and energy to flow across the connections.

The simulations confirm this view, and offer new insight. Within a prolonged period of southward orientation, the simulation showed ionospheric activity correlating with solar-wind density. The SPP group will investigate this density correlation in future work.

The January storm was not the biggest of its kind; its significance lies in being the first intense magnetic storm followed in entirety, from the sun to the ionosphere. “The animations had a big impact,” says Goodrich. “Space physics deals with invisible fields and particles, and an image carries great weight. Without centers like PSC, we wouldn’t be able to do this.”

This POLAR satellite view of the Arctic shows an ultraviolet image of electro-magnetic radiation from the auroral oval during the January 1997 magnetic storm. Image courtesy of Marshall Space Flight Center.

Ripples in Space and Time



In 1915 Albert Einstein made big waves in physics by propounding a radically new way to think about space, time and gravity. His idea, the general theory of relativity, sent ripples churning through 20th century understanding of the universe that led physicists to the big bang, black holes and brave new worlds of unified theory that reach to embrace all of physics.

Among the still rippling effects of general relativity is Einstein's prediction that moving objects give off gravitational waves. Much like vibrating electrons give off electromagnetic radiation, which allows us to listen to radio and watch TV, the accelerating movements of massive objects in space, such as supernova explosions and black holes, produce gravitational radiation, according to Einstein – ripples moving at the speed of light through the four-dimensional fabric of spacetime. Although indirect evidence suggests Einstein was right, gravity waves are notoriously difficult to detect, and proof of their existence remains a matter of circumstantial evidence.

To clinch the case, scientists are building LIGO, the Laser Interferometer Gravitational-Wave Observatory. LIGO's ability to detect gravity waves early in the next century depends critically on a team of University of Pittsburgh physicists. Astrophysicist Jeffrey Winicour and his coworkers are developing computational tools to simulate how black holes emit gravity waves, and in particular gravity waves from two black holes orbiting each other – a binary black hole. Their 1997 computations on PSC's CRAY C90 marked a milestone. Using an innovative new approach, they carried out the first stable, 3D simulations of a single black hole moving through space over a long period of time – an objective that physicists ten years ago saw as “the Holy Grail of numerical relativity.”

The first stable 3D simulation of a black hole moving through space.



Getting to the Holy Grail

Even with LIGO, finding gravity waves is like the proverbial needle in a haystack, except not only don't you know where to look, but the needle, for all you know, looks like a piece of hay. The job of Winicour's group is to provide reliable pictures of the needle. With accurate knowledge of the frequency and amplitude of gravity waves, LIGO's task will be like tuning in a very weak radio signal. “The more you know about the signal,” says Winicour, “the better chance you have of detecting it.”

Binary black holes are especially important because their gravity waves, according to theory, should be about 100 times stronger than from a single black hole. To simulate a binary black hole, however, is much more complicated, and even for a single black hole with 3D realism, simulation of its spin and motion was out of range until the Winicour team's breakthrough.

The problem is that general relativity presents an extreme challenge to computational capability. For starters, a black hole contains a point of infinite spacetime curvature and infinite density – a singularity, in mathematical terms. Simulating such physical extremes leads to numerical instability – the computer program crashes if it computes near the singularity. Winicour's team has solved this problem by creating a software module that, in effect, explores the black hole, looking for the singularity and cutting it out of the computation.

Riding the Waves

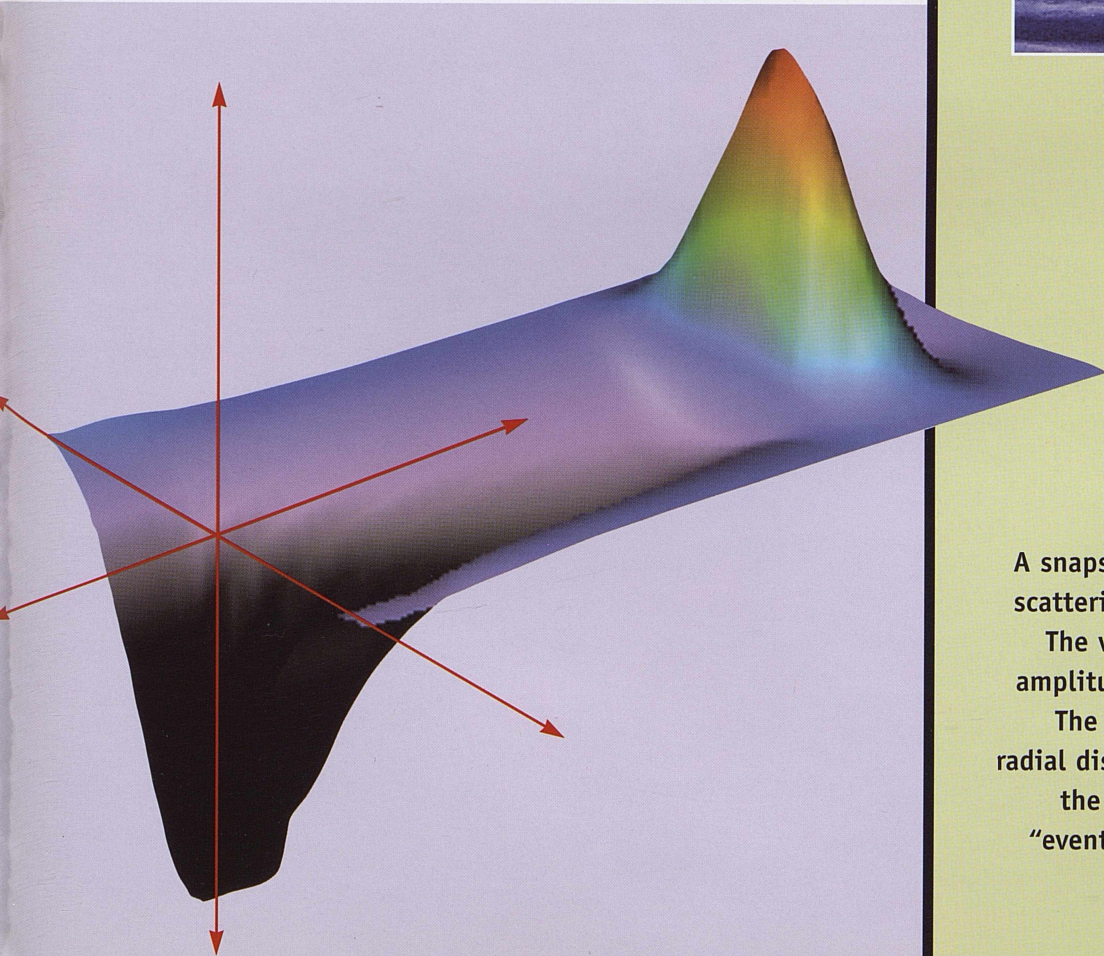
The Winicour group's milestone also reflects a new approach they developed to follow how a black hole evolves. The standard method – the only method until their work – was to recompute properties of the spacetime field going forward in timesteps, much like a computation tracking a particle in space. This method presented serious numerical problems due to the strong, rapid fluctuations of gravity waves, and it had not been possible to simulate even a single moving black hole in 3D over a long time.

The Pittsburgh Relativity Group:
(l to r) Roberto Gómez, Bela Szilagi,
Jeffrey Winicour and Luis Lehner. For his
work on this project, Lehner, who completed
his Ph.D in physics at the University of
Pittsburgh in 1998, received the Nicholas
Metropolis Award. This American Physical
Society award honors the outstanding
computational doctoral thesis of the year.

To overcome this problem, Winicour's team used a different method for labeling points in spacetime, an approach based not on standard time but on how light spreads from a point in space – a concept referred to as a light cone. This approach, called retarded time evolution, smoothes out the irregularities. "You go out along the light cone," explains Winicour, "and it's something like riding with the waves. You smooth out the physical properties."

No one understands better than Winicour the serious work ahead on the path to meeting the goal of LIGO: "The problem of simulating the whole system of a binary black hole has to be broken into regions. Our approach works flawlessly far from the holes and near them, but it breaks down in a region between the two holes." Full simulation of binary black holes will require melding the two types of time evolution, which the group is working on now.

LIGO is a pioneering effort, and it's expected to open a new window on the universe. It should firmly verify that black holes exist. Gravity wave detectors should also tell us more about how neutron stars are born in the explosion of a supernova. They may also give us a look at the instant when space and time came into being. The tools that Winicour's team are developing, along with the resources of centers like PSC, will play a vital role in this work.



More information: <http://www.psc.edu/science/winicour.html>

"Ripple in still water
where there is no pebble tossed
nor wind to blow."


— The Grateful Dead




A snapshot of the gravitational wave scattering from a moving black hole. The vertical axis corresponds to amplitude of the gravitational field. The horizontal axis represents radial distance. Light can't escape from the region (black) inside the "event horizon" of the black hole.

Established in 1986, the Pittsburgh Supercomputing Center is a joint effort of Carnegie Mellon University and the University of Pittsburgh together with Westinghouse Electric Company. It has received significant support from:

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





C R E D I T S

Editor / writer.....

Michael Schneider, PSC

Design.....

John Brady
Design Consultants

Graphics Research.....

Sean Fulton, PSC

Photography.....

Graphic Services at
Mellon Institute,
Jim Schafer,
David Deerfield,
Karen Preuss,
Kevin Roark,
storm graphic
provided by NCAR

Printing.....

Superior Printing



Printed on recycled paper



Waterless Printing

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

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