PITTSBURGH
SUPERCOMPUTING
CENTER

BIOMEDICAL APPLICATIONS

PITTSBURGH
SUPERCOMPUTING
CENTER

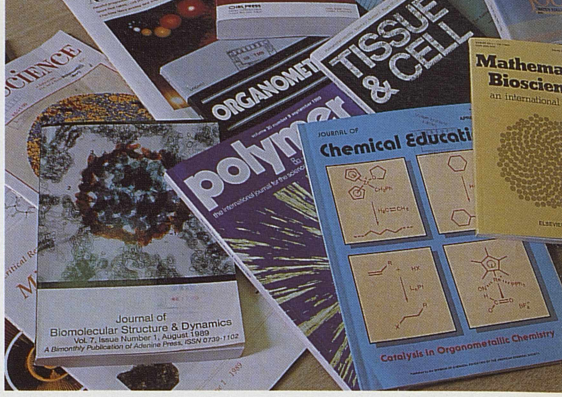

In 1986 a grant from the National Science Foundation established the Pittsburgh Supercomputing Center. The mission: Make high-performance computing available to scientists and engineers nationwide. To date over 1600 researchers at more than 180 universities and research centers in 46 states have used the Pittsburgh Supercomputing Center to advance their projects. This research spans a diverse range of disciplines, including air quality studies, the microeconomics of Wall Street, bone tissue healing, protein structure and galaxy evolution. It has resulted in publication of over 600 papers in science and engineering journals.

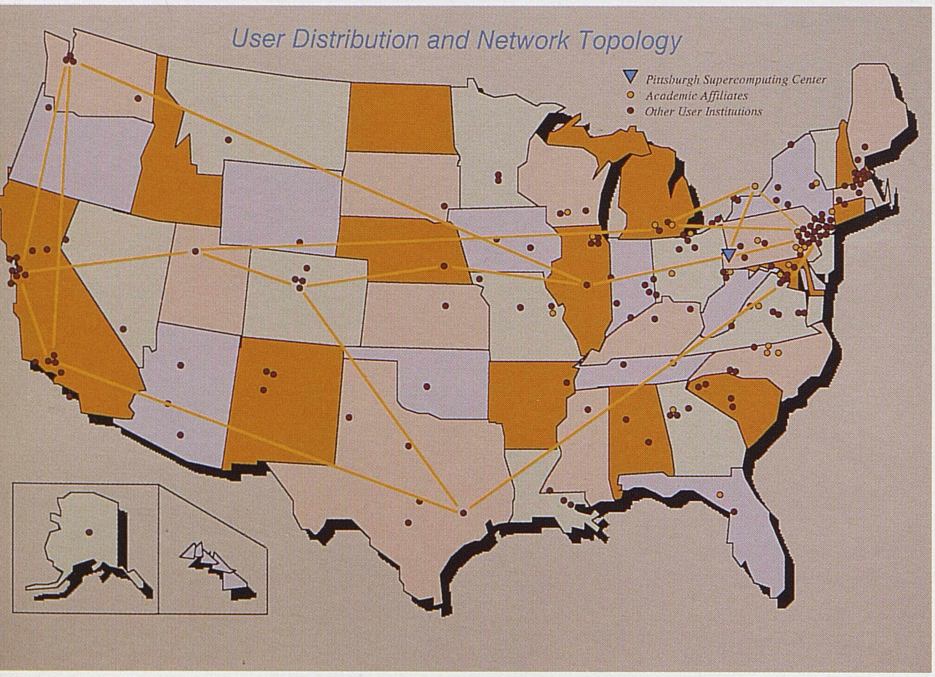

Researchers throughout the United States use the computing power of the Pittsburgh Supercomputing Center via electronic networks. Regional communication networks connect to a National Science Foundation high-speed pathway, "the NSF backbone."

Twenty-seven universities have become Pittsburgh Supercomputing Center Academic Affiliates. Representatives from these campuses form the Center's main advisory body.

University of Pittsburgh

As a joint project of Carnegie Mellon, the University of Pittsburgh and Westinghouse Electric Corporation, the Pittsburgh Supercomputing Center has an institutional base uniquely well-suited to its task. It can draw on the resources of one of the outstanding computer science programs in the country, a large state-related university known for research in science and medicine and a major corporation with a strong track record in high-performance computing.

Carnegie Mellon

"On the global level, it's fundamental: whoever has the fastest supercomputers has the best chance to find answers to the tough questions facing all societies."

—Richard M. Cyert, president
Carnegie Mellon University

Mike Levine & Ralph Roskies

Under the leadership of scientific directors Michael Levine and Ralph Roskies, physicists with wide experience in computational research, the Center has built a comprehensive program of service to users, on-line documentation and software development.

▼ Training—workshops on supercomputing techniques and applications reached over 1500 participants in 1988 alone.

▼ Software—a comprehensive library, including a variety of application packages, meets the software requirements of most research projects. Center staff have also developed programs that enhance the user's computing environment.

▼ Documentation—extensive on-line help files, including many selected examples, complement hardcopy documentation.

▼ Consulting—from "hotline" troubleshooting of immediate problems to Ph.D. scientific specialists who advise on research strategy, Center staff assist users in various stages of their work.

▼ Scientific Visualization—a variety of graphics packages, coordinated through GPLOT, a Center developed translator program, allows users to produce video animations with relative ease.

Through an agreement with Cray Research, Inc., the Center in December 1988 received the first Cray Y-MP installed at a non-government site. The agreement also provides that the Center will receive the first non-governmental CRAY-3 as soon as it is available. The CRAY-3 will be five times more powerful than the Y, with 16 times as much memory.

"This contract creates confidence within the research community that National Science Foundation supercomputing is an ongoing national priority," says Michael Levine. "Scientists can attack the most difficult questions in their field because they know that the most advanced computing resources will be at hand."

Beverly Clayton, executive director, Pittsburgh Supercomputing Center.

"Through our workshops and on-line documentation, and especially through the commitment and problem-solving skill of our consulting staff, we've gained a reputation for excellent service to users."

—Beverly Clayton

The Pittsburgh Supercomputing Center's CRAY Y-MP, generally regarded as the most powerful computer in the world.

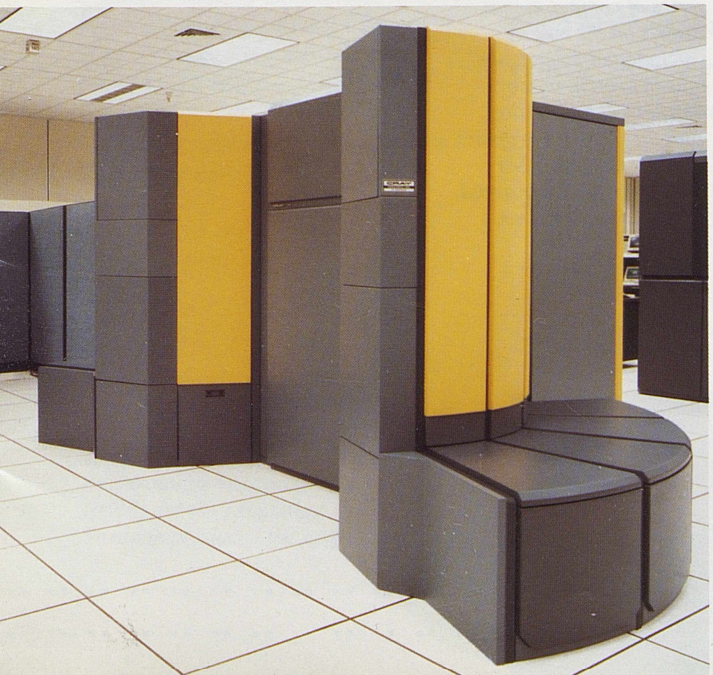

Jim Kasdorf, manager of engineering computing services at Westinghouse Energy Center, site of the Center's CRAY Y-MP/832. Opening the Y's front panels reveals the laced piping of its fluid cooling system.

THE PROGRAM IN BIOMEDICAL SUPERCOMPUTING

Pioneering efforts like Salk polio vaccine and Thomas Starzl's organ transplant team have established the University of Pittsburgh as an important center of medical science. Since its inception, the Pittsburgh Supercomputing Center has been a strong partner to this health research capability, with noteworthy projects in drug design, macromolecular structure and prosthetic design among the research it has supported.

In September 1987 the National Institutes of Health, through the Biomedical Research Technology Program in the Division of Research Resources, awarded the Center a grant to expand its effort in biomedical applications. This grant, the only one of its kind to a supercomputing facility, marked a major initiative to promote supercomputing for biomedical research.

Two chemists with strong backgrounds in biology and computing, Hugh Nicholas and David Deerfield, soon joined the Center's staff and quickly moved to implement training, support and development activities.

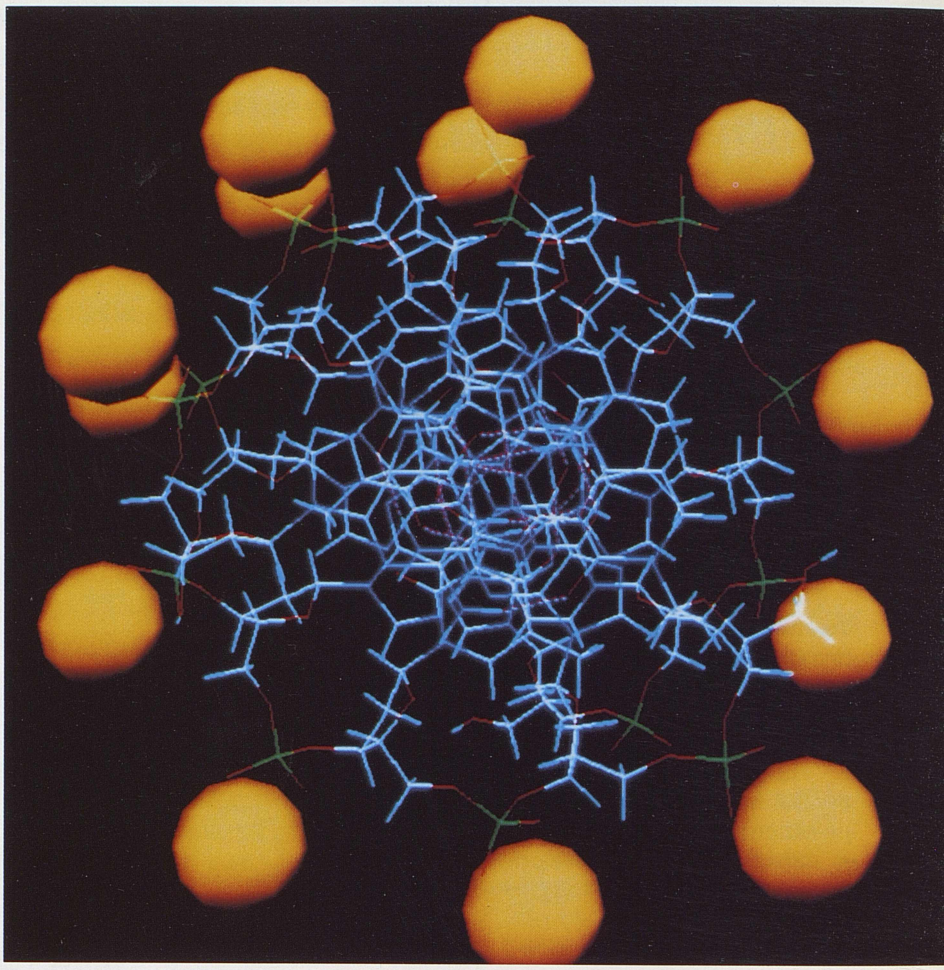

David Deerfield has applied quantum and molecular mechanics in a series of computational studies on the role of calcium and magnesium ions in protein interactions. In collaboration with the Center's graphics staff, he has produced several three-dimensional molecular animations.

A series of workshops have introduced over a hundred biomedical researchers to supercomputing.

▼ Twice a year the Center holds a five-day workshop on supercomputing techniques for biomedical researchers. These sessions teach efficient use of the Cray computing environment, including special features of the CRAY Y-MP. Half the time goes to hands-on sessions in which participants work on their own problems under expert guidance.

▼ Two to three-day workshops address particular biomedical applications:

—Molecular Mechanics and Dynamics

—Macromolecular Structure Refinement

—Nucleic Acid and Protein Sequence Analysis

—Fluid Dynamics in Biological Systems

—Advanced Image Processing in Electron Microscopy

"You can't crunch numbers on a CRAY without learning how," notes Hugh Nicholas, "and to be productive you have to learn quickly. The workshops accomplish that."

These workshops are led by authorities in their field, including among others: Charles Brooks, Carnegie Mellon University; Joachim Frank, National Center for High Voltage Electron Microscopy; Jacob Maizel, National Cancer Institute; Charles Peskin, Courant Institute and Gary Stormo, University of Colorado.

▼ Structure and Function of Biological Molecules

Determining the structure of proteins and other biological molecules and examining their functional parameters has been an active area of usage on the Center's CRAY. With the NIH grant, the Center has expanded its resources in this area.

The biomedical software library includes two molecular structure databases: Brookhaven (macromolecules) and Cambridge (small molecules). The Center has several programs for structure refinement by least-square methods, and it supports Axel Brünger's X-PLOR program.

"X-PLOR uses molecular dynamics to significantly speed the search for an optimized structure," notes David Deerfield. Deerfield supports packages in molecular mechanics and dynamics—including AMBER, CHARMM and GROMOS—along with semi-empirical packages and GAUSSIAN and a set of utilities he coded to transfer data among these programs.

"As predictions become more accurate due to improved methodologies and the increased availability of supercomputers, simulations of protein dynamics will be applied to practical problems in genetic engineering and the search for inhibitors to cure a wide range of diseases."

—Martin Karplus

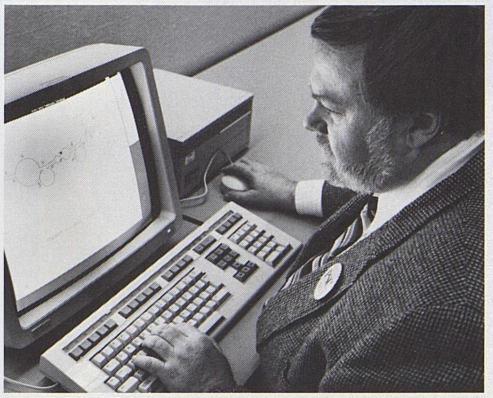

Hugh Nicholas has worked extensively in X-ray crystallography and biological sequence analysis. He developed computer-aided sequence comparisons that successfully predict the regions of transfer RNA recognized by amino-acid attaching enzymes.

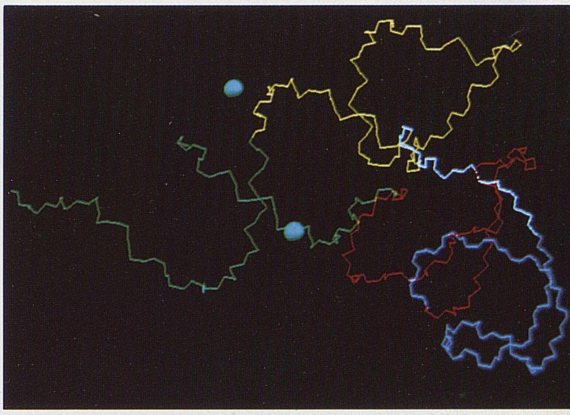

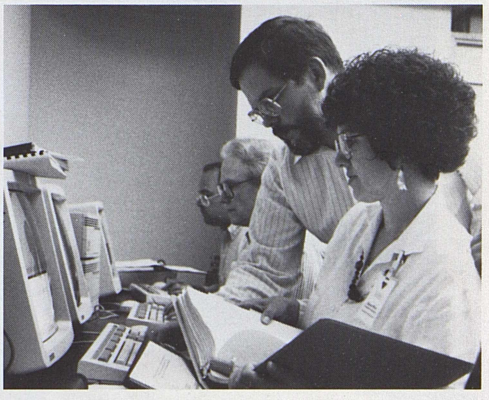

David Deerfield consults with Susan West, University of Rochester, at a workshop hands-on session.

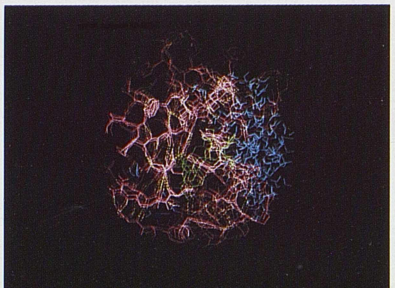

A molecular dynamics simulation of the drug trimethoprim "mutating" to a related compound at the active site of the enzyme dihydrofolate reductase. Charles Brooks (p. 10) did the computing, and Scott Sneddon produced the graphic.

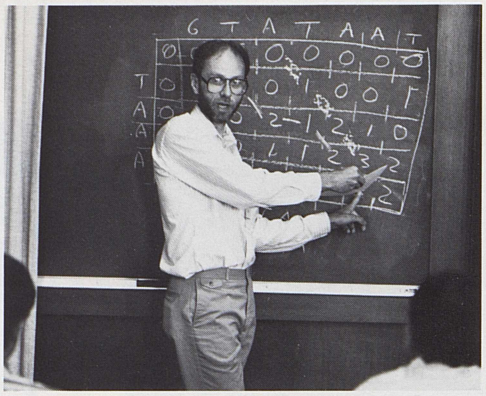

Michael Waterman, University of Southern California, explains a sequence alignment algorithm at a recent workshop.

"Examples arise daily of newly discovered relationships between cancerous and normal genes, of protein classes with diverse origins but common structural motifs, of insights into molecular evolution, and of mechanisms to generate the seemingly enormous diversity that organisms show."

—Jacob Maizel
National Cancer Institute

UUU Phe	UCU Ser	UAU Tyr	UGU Cys
UUC Phe	UCC Ser	UAC Tyr	UGC Cys
UUA Leu	UCA Ser	UAA stop	UGA stop
UUG Leu	UCG Ser	UAG stop	UGG Trp
CUU Leu	CCU Pro	CAU His	CGU Arg
CUC Leu	CCC Pro	CAC His	CGC Arg
CUA Leu	CCA Pro	CAA Gln	CGA Arg
CUG Leu	CCG Pro	CAG Gln	CGG Arg
AUU Ile	ACU Thr	AAU Asn	AGU Ser
AUC Ile	ACC Thr	AAC Asn	AGC Ser
AUA Ile	ACA Thr	AAA Lys	AGA Arg
AUG Met	ACG Thr	AAG Lys	AGG Arg
GUU Val	GCU Ala	GAU Asp	GGU Gly
GUC Val	GCC Ala	GAC Asp	GGC Gly
GUA Val	GCA Ala	GAA Glu	GGA Gly
GUG Val	GCG Ala	GAG Glu	GGG Gly

▼ Sequence Analysis

Analyzing the human genome is one of the major challenges in biological science, and the quantity of data alone pushes biologists toward computing. The Center has compiled the best tools available to assist researchers in this task.

"We provide a comprehensive computing environment for molecular biologists."

"We have three major nucleic acid sequence data bases and the facilities to search them," explains Hugh Nicholas. "Two of them, GenBank and EMBL, currently have over 30 million base-pairs. We support the IDEAS analysis programs from the National Cancer Institute and several alignment programs—including code we've implemented, ALIGN, for very long sequences. We also use the University of Wisconsin Genetics Computer Group programs; with this package and our more computer-intensive programs, we provide a comprehensive computing environment for molecular biologists."

Cheri Brooks, coordinator of the biomedical program, organizes the biomedical workshops and administers proposals for biomedical research. For first-time users of supercomputing, the Center awards starter grants of a few hours. Proposals for production grants go to the Center's Biomedical Review Committee.

PROJECTS IN BIOMEDICAL SCIENCE

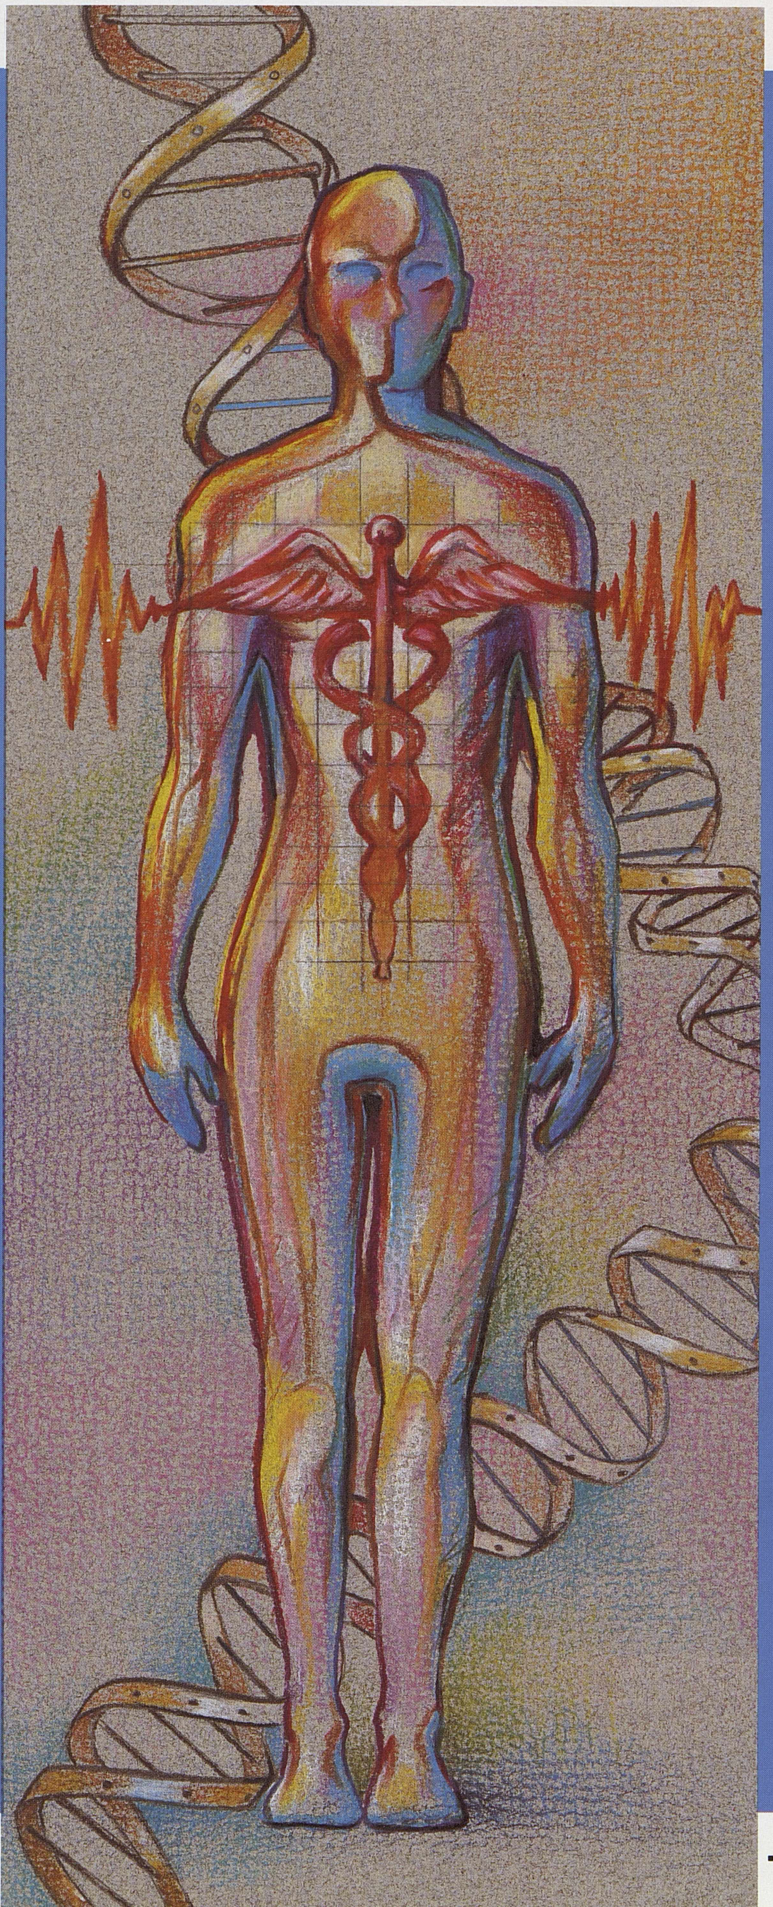

Health-related research at the Pittsburgh Supercomputing Center has been ambitious, innovative and concerned with a wide variety of problems. Molecular biology, especially molecular mechanics and dynamics studies of structure-function relationships, has been an active focus of the Center's work. At the same time, studies in cell growth, blood flow, cardiac arrhythmia and electrocardiography among others have also contributed to our knowledge of human biology and our hopes for an improved quality of life.

The common denominator of these projects is supercomputing. Most of the results from biomedical work at the Center to date relied on the capability of its first supercomputer, the CRAY X-MP. Even with this computing power, many projects—especially protein and other large molecule studies—had to be scaled to its limits of speed and memory. With the advanced capability of the Center's CRAY Y-MP/832, we can expect knowledge to progress at an even faster pace than we have seen in recent years.

"With these advances in supercomputing, Pittsburgh stands on the verge of exciting progress in biomedicine."

—Wesley Posvar, president
University of Pittsburgh

The Charm of Biomolecular Motion

Accustomed to seeing protein structures as fixed images from X-ray crystallography, we can lose sight of how in living systems these molecules never stop moving. At picosecond tempos, they vibrate as if dancing a sub-microscopic jig, fluctuating constantly around their average structure.

Martin Karplus and coworkers at Harvard have developed molecular dynamics methods to study these motions and have implemented them with a software package called CHARMM—Chemistry at Harvard Macromolecular Mechanics. Working with Karplus, Jiali Gao recently used CHARMM at the Pittsburgh Supercomputing Center to simulate the difference in the free energy of hemoglobin when one amino acid changes into another, creating a mutant protein.

This "computer alchemy" made it possible to decompose the free energy change into separate contributions—the mutating amino acid interacting with the surrounding water and with other amino acids of the protein. The magnitude of some of these individual changes greatly exceeded the overall difference. In *Science* (June 2, 1989), Karplus and his colleagues conclude that this simulation "provides new insights into the origin of thermodynamic changes in mutant proteins."

▼ Automated Structure Refinement?

Hours turning rapidly to weeks, the crystallographer works at his graphics display, shifting a bond angle, rotating a side-chain. . . . New techniques make it easier than ever to collect X-ray diffraction data on biomolecules, and the laborious refinement process, even with interactive graphics, is one of the main bottlenecks to determining the structure of crystallized proteins.

Axel Brünger of Yale, who collaborated with Martin Karplus during post-doctoral work at Harvard, applied Karplus' molecular dynamics approach to the problem and found a computational method, "simulated annealing," that automates the structure refinement process with considerable success. Brünger has incorporated simulated-annealing in his program X-PLOR, supported at the Center with documentation and an on-line help file.

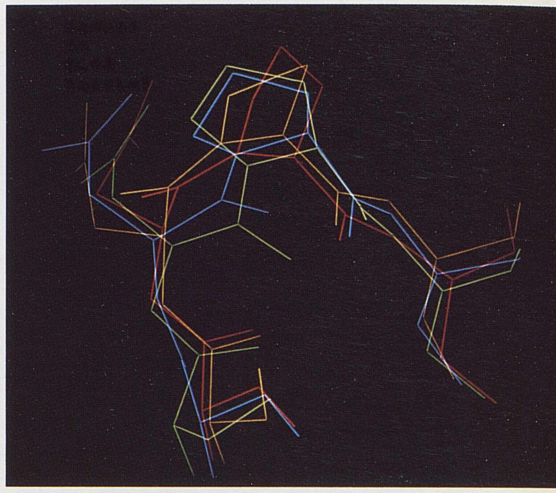

In tests with aspartate aminotransferase, simulated annealing made structural changes computationally that are not possible with traditional least-squares optimizing routines. The initial crystallographic structure (green) is overlaid with structures obtained by restrained least-squares minimization (blue), simulated-annealing (red) and manual correction (yellow).

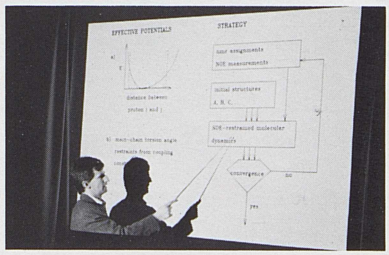

At a Pittsburgh Supercomputing Center seminar, Axel Brünger of the Howard Hughes Medical Institute, Yale University, explains a molecular dynamics strategy for refining protein structure using nuclear magnetic resonance data.

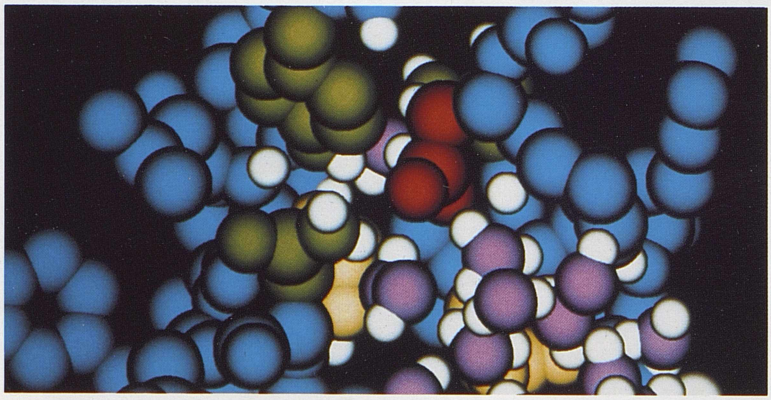

A representation of hemoglobin, the subject of the mutant protein study by Gao, Karplus and colleagues. The side-chain of the mutating amino acid, aspartic acid, is shown in red. Interacting water molecules are purple and white and some of the interacting side-chains are shown in green and yellow.

"Spray and pray"—that's one description of the traditional approach to finding a new drug. From millions of compounds in nature, drug companies identify thousands with potential medicinal value, and it's generally estimated that one in 10,000 actually tested becomes a drug. On the average, it takes 7 to 10 years and $50 to $70 million to bring a drug to market.

Supercomputing offers an alternative: biomolecular simulations that specify the parameters, in effect a blueprint, to synthesize drugs that have little or no side effects—drugs molecularly "aimed" at the desired result. For expert biochemists armed with sophisticated mixed quantum mechanical and empirical approaches and computational know-how, this is the foreseeable future, not science fiction. Though still short of full-scale computer-aided design, biomolecular simulations have proven capability to eliminate false trails and identify promising areas for further research.

William Lipscomb directs much of his research toward understanding structure-function relations in complex enzyme systems. Since 1967 he has studied an enzyme, aspartate transcarbamylase (ATCase), positioned like a porter at the gate of pyrimidine biosynthesis, a metabolic pathway crucial to DNA replication and normal cell division. Lipscomb's work on ATCase led to the release of phosphonacetyl-aspartate (PALA), an inhibitor of ATCase, as a cancer treatment.

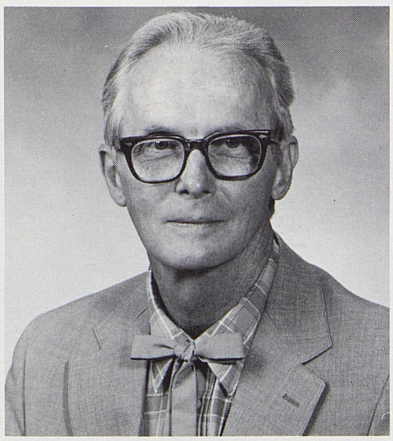

William Lipscomb, Harvard University, 1976 Nobel laureate in chemistry.

▼ Supercomputing and Enzyme Research

Lipscomb used Axel Brünger's computationally demanding X-PLOR program to great advantage in refining ATCase structures. "We'd have enormous difficulty doing these calculations without it," says Lipscomb. He emphasizes the point by relating ATCase to work he recently began on another complex enzyme, fructose 1, 6-bisphosphatase, which could lead to an effective diabetes drug:

"Using the supercomputer, it took us about seven months to get the structure. We didn't have all this computing power when we started working on ATCase, and it was about twelve years before we established either of the big three-dimensional structures. But even that's not as impressive as the fact that we can work on much more complex problems when we know the supercomputer is available, problems we never would have tried."

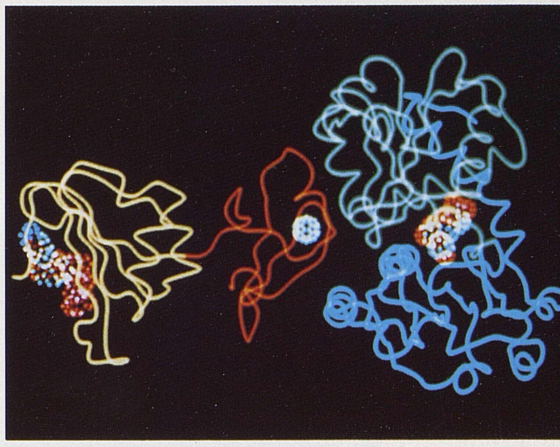

Bound configuration of ATCase, showing one of the six chain pairs that constitute the enzyme. A CTP molecule binds the R-chain (regulatory) active site (left), initiating a cooperative reaction that deactivates a paired C-chain (catalytic) active site 60 angstroms away, halfway across the protein. This structure shows PALA bound to the C-chain active site.

The metabolic pathway catalyzed by ATCase. The reaction is feedback inhibited by CTP.

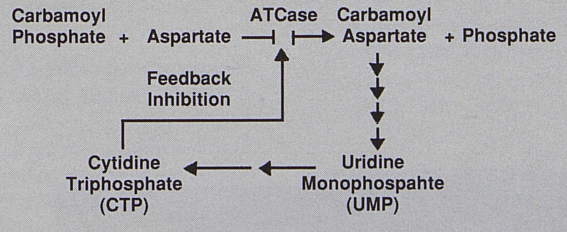

"We can work on much more complex problems when we know the supercomputer is available, problems we never would have tried."

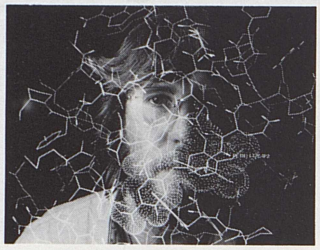

Charles Brooks of Carnegie Mellon University's chemistry department studies a computer-generated image of trimethoprim. As a post-doctoral researcher at Harvard, Brooks collaborated with Martin Karplus. Brooks has added thermodynamic simulation methods to the capabilities of CHARMM.

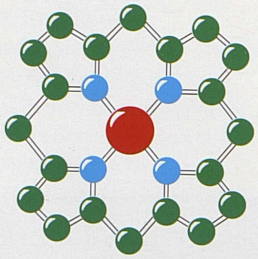

The heme unit, active site of heme proteins. A central iron atom bonds to four nitrogens, each of which in turn forms a ring with four carbons.

▼ Looking for the Best Inhibitor

Nearly all drugs work because they inhibit an enzyme, and the goal of drug design is to decipher the complex interactions that control enzyme-inhibitor binding. Charles Brooks has conducted a series of studies on the antibacterial drug trimethoprim. Trimethoprim inhibits an enzyme, dihydrofolate reductase (DHFR), important in cell division.

A series of compounds closely related to trimethoprim also inhibit DHFR, and this allows Brooks—with the help of supercomputing—to examine how slight differences among closely related compounds can affect binding. These features of the trimethoprim/DHFR system allow Brooks to treat it as a development prototype for a general drug design strategy. "The supercomputer's not a magic solution to drug design," notes Brooks, "but it allows chemists to skip steps one through five and go directly to step six, and that represents an enormous saving of person hours."

▼ The Heme Protein Puzzle

Gilda Loew directs the Molecular Research Institute in Palo Alto, California and has studied heme proteins for twenty years. Recently she applied molecular dynamics simulations to a particularly interesting segment of the heme protein family tree. Myoglobin is virtually identical at its active site, chemically and geometrically, to its cousin cytochrome C peroxidase (CCP). Myoglobin transports oxygen, and CCP is an enzyme. "One of the enigmas of heme proteins," says Loew, "is that they have essentially the same active site, the heme unit, and the biological function centers there, usually on the iron itself, yet the biological functions differ."

Loew successfully characterized an unstable intermediate form of CCP, the only form in which CCP performs its enzyme function. Hydrogen peroxide (H^2O^2) displaces the water molecule normally bound to CCP's heme unit. "The peroxide state is so transient," says Loew, "that it's never been isolated experimentally." The molecular dynamics image of this enzyme state corresponds to twenty picoseconds of simulation, a snapshot each picosecond.

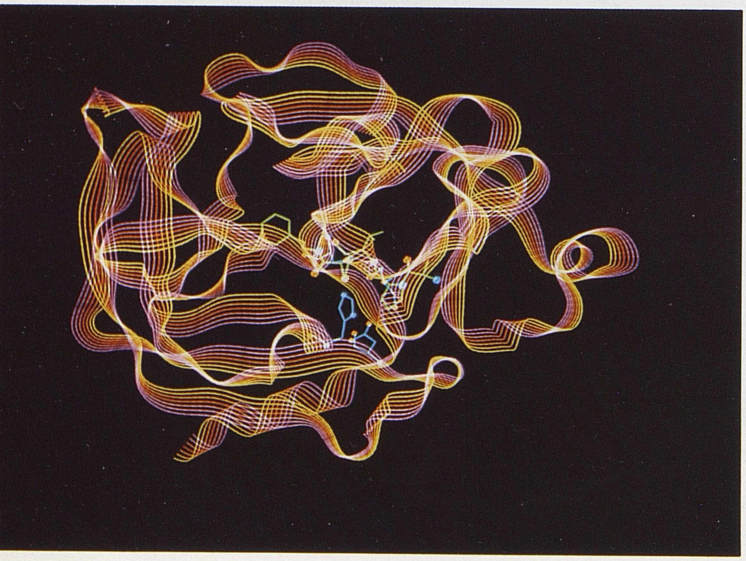

Using AMBER, a set of molecular dynamics programs developed by Peter Kollman of the University of California at San Francisco, Edgar F. Meyer, Bogdan Lesyng and Maciej Geller of Texas A&M University conducted a thermodynamic free-energy difference calculation for the pancreatic elastase (PPE) enzyme-inhibitor system. The polychromic "ribbon" represents the backbone conformation of PPE with an inhibitor, petidyl-difluoroketone, bound at the active site.

▼ How Calcium Gets Around

Calmodulin is a ubiquitous calcium-sensing regulatory protein. To understand how calcium (Ca^{2+}) gets around in the human body, where it is involved in stress response, heart-rhythm, muscle contraction and bone strength, you need to understand calmodulin. Harel Weinstein and colleagues at Mount Sinai medical school have used Pittsburgh's CRAY to simulate calmodulin interacting in the watery world it inhabits.

They learned that when the calcium-bound crystal form goes swimming it loosens its erect helical "backbone," curls over and pulls the two Ca^{2+} binding domains closer together. "This structure explains the activity of the molecule," says Weinstein. He notes that these results agree with experimental data that had not clearly indicated the aqueous structure. "This shows what theoretical simulations can do—people have been using the crystal structure to explain the activity of this very important molecule."

▼ Unraveling the Double Helix

The Y-MP may be a special boon to DNA research. The sheer size of DNA molecules has limited computational approaches and challenged the ingenuity of researchers.

David Beveridge of Wesleyan University has a series of DNA-related projects underway. Using GROMOS, he has simulated DNA-water interactions for several DNA configurations. In recent work Beveridge performed Monte Carlo simulations for the "spine of hydration" on the minor groove. By adjusting intermolecular potential functions, he obtained "nearly perfect" congruity with observed crystallographic results.

Peter Kollman is using Pittsburgh's CRAY to extend his research on basepairs in water. In this work, he uses AMBER to simulate basepair-water interactions. He also draws on thermodynamical theory to evaluate free energies of association as a function of distance. Kollman expects to deepen understanding of the factors affecting stability of the double helix structure from these basepair studies.

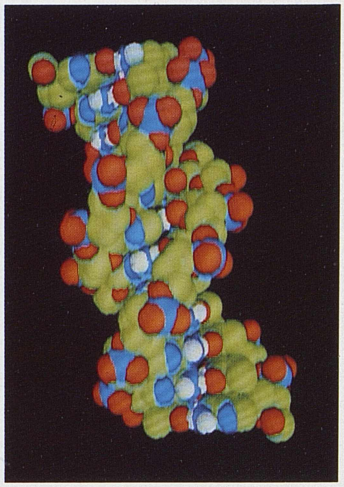

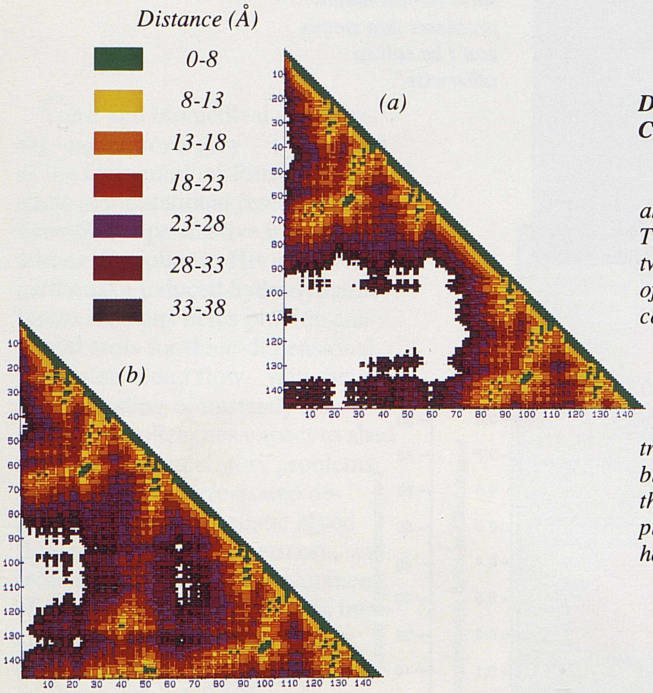

Distance Matrices for Calmodulin, Crystallized (a) and in Solution (b).

The x & y axes correspond to the amino acid sequence of calmodulin. The color-coding shows distances between residue positions, an indication of folding (tertiary structure). Non-colored areas indicate distances greater than 38 angstroms (Å), the limit of the color scale.

"You see that each matrix has two colored regions, both roughly triangular. These represent the two binding domains. The hole between them is much smaller in the second plot, showing that the binding domains have moved closer together."

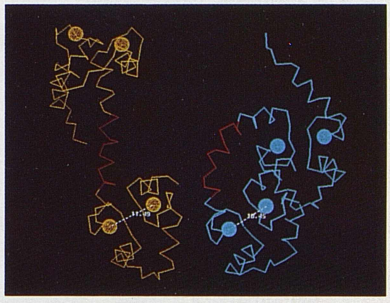

Crystallized calmodulin, left, and in solution, right—after 246 picoseconds of molecular dynamics simulation (using CHARMM). The red lines represent the helical backbone connecting the two Ca^{2+} binding domains (2 active sites in each domain) of the protein.

In Ovid's classical masterpiece, people become deer, trees, donkeys and goats; in Kafka, a man becomes a cockroach. The one-celled bacterium *Bacillus subtilis* undergoes a metamorphosis no less extreme. When it's too hot, too poisonous or there's no food, *B. subtilis* pulls its nuclear material together and covers it with resistant, protective layers—spore coats. The "mother cell" then self destructs as it releases this survival-adapted form of itself.

Sporulation is a basic study in cell differentiation, a process involved in the growth of all higher life forms. "If we understand differentiation in its simplest form," says Mohammed Ataai, "we should be able to understand higher cellular systems." A biochemically-oriented chemical engineer at the University of Pittsburgh, Ataai has built a comprehensive model of *B. subtilis*, with differential equations and rate constants for 39 separate metabolic reactants. His aim is to isolate the origin of the sporulation response.

In initial work, he has successfully simulated the transition from normal cell division—the "exponential phase"—to the "stationary phase," when growth stops to prepare for sporulation. Computation per time-step increases dramatically during this transition, reflecting the radical changes in the cell's biochemistry. "This is very difficult to simulate," says Ataai, "and our model does it with no numerical instability. This means we should be able to see which changes in cell processes cause differentiation."

When increased strain on a bone persists over time, the bone "remodels"; new tissue reshapes it and adds strength. Though remodeling is a healing process, it poses a problem for orthopedists. A high-percentage of prosthetic joint replacements fail after five to ten years; for knee replacements, the failure rate is 36%. Joint replacement redistributes strains on the bone, and consequent remodeling eventually loosens the steel pin anchor.

Biomedical engineer Richard Hart has attacked this problem with modeling more sophisticated than previously used in the field. The objective is prosthesis design that accurately accounts for remodeling. Hart developed a three-dimensional finite-element code that updates shape and stiffness after each iteration. He has achieved good agreement with experimental results on sheep bone and is extending his work to include remodeling of tooth sockets in the human jaw.

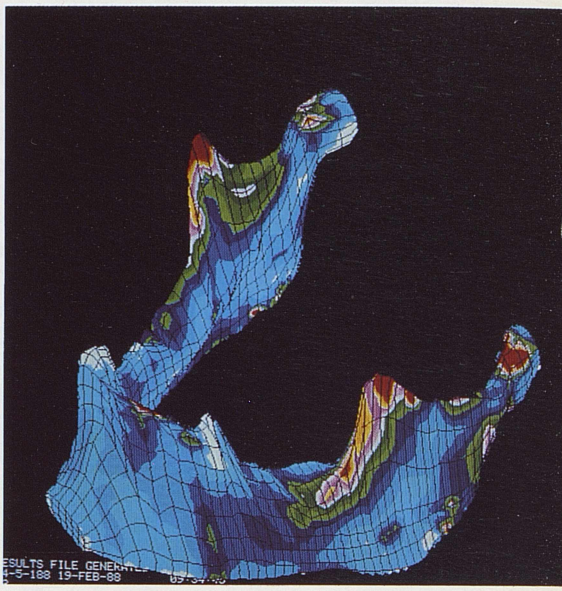

Richard Hart's 4000 node three-dimensional finite-element model of the human jaw.

Richard T. Hart, Tulane University.

"The supercomputer makes it possible to solve biomechanics problems that simply can't be solved otherwise."

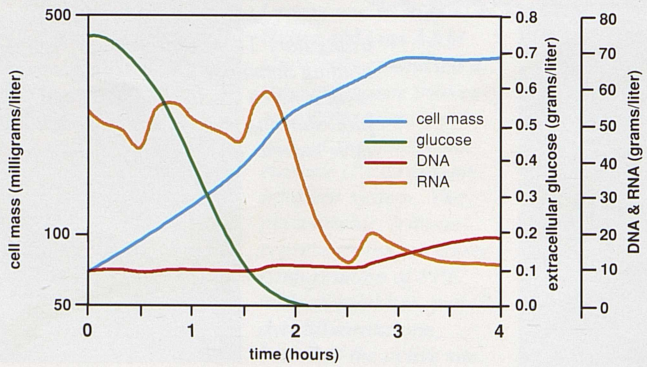

The first two hours correspond to the last two cycles of cell division. At two hours, when glucose is gone, growth rate—indicated by cell mass—levels off. During the transition (2 to 3.5 hours), the model predicts intracellular concentrations for all the components modeled.

Inverse electrocardiography—the idea is to construct a map of electrical events on the surface of the heart noninvasively. "Improved electronics now make it practical to record potential distribution over the entire torso, from 100 sites or more," says Yoram Rudy. "With this information, we can in principle reconstruct epicardial potentials, and this has great clinical importance. It would show the location of ventricular arrhythmias and various conduction abnormalities."

The classical mathematical approach to the "inverse problem" is "ill posed," says Rudy a biomedical engineer at Case Western Reserve University: A certain amount of irregularity, "noise," is inherent in surface potential data, and small variations can lead to large variations in the epicardial solution. To overcome this problem, Rudy has adopted a mathematical method known as Tikhonov regularization. His computations show good qualitative agreement with canine experimental data. He concludes that the inverse procedure is a feasible, practical and valuable tool in cardiac diagnosis.

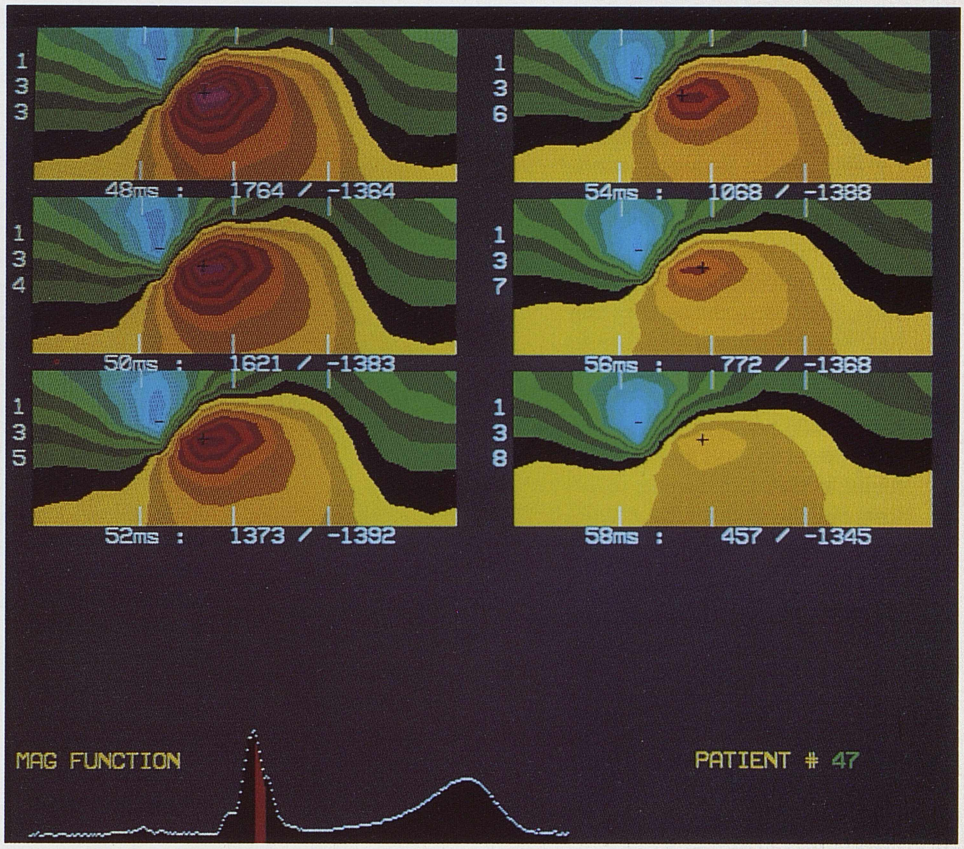

Civil and biomedical engineering and neurosurgery—T. K. Hung of the University of Pittsburgh conducts computational research that exemplifies productive interaction among disciplines. His work on earthquake induced hydrodynamic forces on dams helps provide analytical tools for three-dimensional studies on blood flow. With better understanding of unsteady flows, Hung and colleagues expect to shed new light on circulatory problems such as atherosclerosis and deformed arteries. Realistic blood flow simulations, furthermore, will reduce live testing in development of surgical techniques such as intra-aortic balloon pumping, another area of Hung's computational work on Pittsburgh's CRAY.

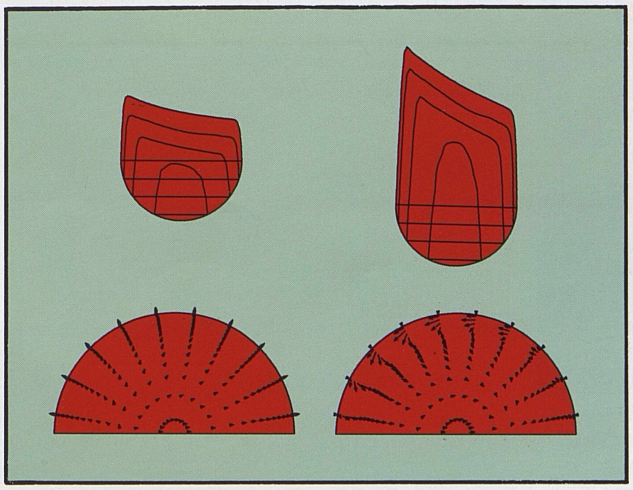

Half-section velocity profiles from the same point along a curved artery at two different times during pulsed-flow simulation. Upper plots represent the longitudinal component. Lower plots indicate secondary flow (velocity components in the cross-sectional plane). Coupling the longitudinal and secondary flow gives a three-dimensional picture of spiral processes.

Los Angeles smog has been so bad for so long that it begins to seem part of nature's design. Greg McRae believes there's a workable solution. McRae together with Ted Russell developed a detailed, comprehensive model of Los Angeles air basin emissions, meteorology and air pollution chemistry: 500,000 coupled, non-linear equations, fifty chemical species and 10,000 mesh points, tested and verified against field measurements.

With 200 hours of CRAY supercomputing, Russell and McRae found interactions among pollutants indicating that EPA's mandated strategy, requiring expensive controls on hydrocarbon sources, might not work. Furthermore, by changing emission parameters, they found that converting from gasoline to methanol as a vehicle fuel could reduce smog roughly by half. They made a video animation and went on the road, presenting their results to environmental officials in Washington and California. Some of them listened. Fuel conversion is one of the recommended strategies in proposed federal clean air act legislation.

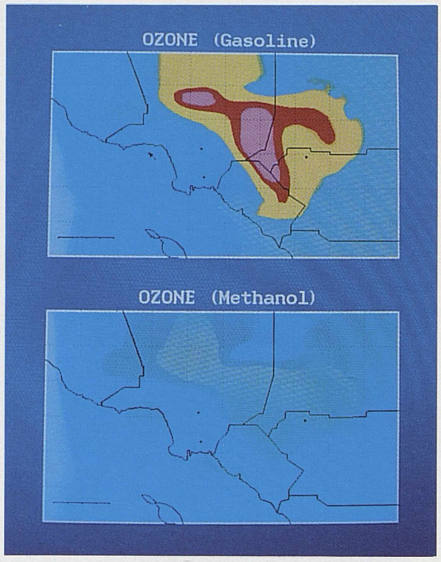

Predicted ozone distribution in the Los Angeles air basin at 1:00 p.m. on a weekday with gasoline and methanol-fueled vehicles. Methanol reduces peak ozone almost 50%, and reduces the area exposed to levels above the EPA standard (green).

At the University of Pittsburgh Graduate School of Public Health, coke oven emissions are not merely data. When the wind blows north, coke works in Clairton and East Pittsburgh contribute to the air Sati Mazumdar breathes. Biostatistician Mazumdar and her colleagues Carol Redmond and Joseph Costantino use the Center's CRAY to improve methods for assessing the carcinogenic risk of air pollutants. Their particular focus is multistage modeling, an approach to cancer risk that accounts for time dependent factors such as latency periods. "Multistage theory," says Mazumdar, "has been successful in describing many experimental and epidemiological observations."

The project uses two sets of lung cancer mortality data: arsenic-exposed copper smelter workers and a study that followed 1100 Pittsburgh coke oven workers from 1953-1970. "Extrapolating from occupational data to general environmental exposure is often the only way to quantify the health effects of low exposure levels."

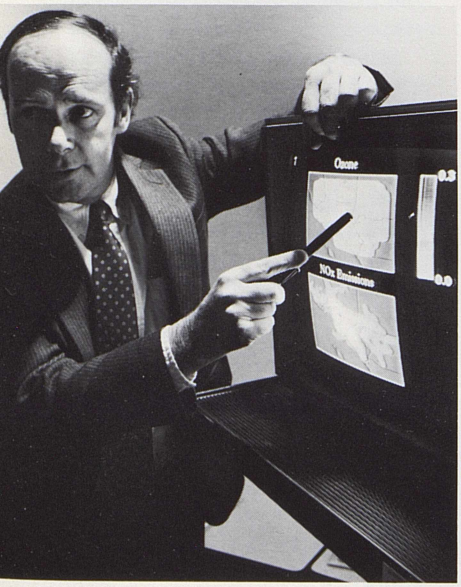

Gregory J. McRae, Carnegie Mellon University, departments of chemical engineering and engineering and public policy.

"When you think about $30 billion a year spent on air pollution control, supercomputing is an absolute bargain."

A "likelihood surface" for a four-stage multistage model, first stage and third stage dose-related, fitted to Pittsburgh male coke oven workers lung cancer mortality (1953-1970). "This surface shows a distinct peak, giving the maximum likelihood for the model's two parameters."

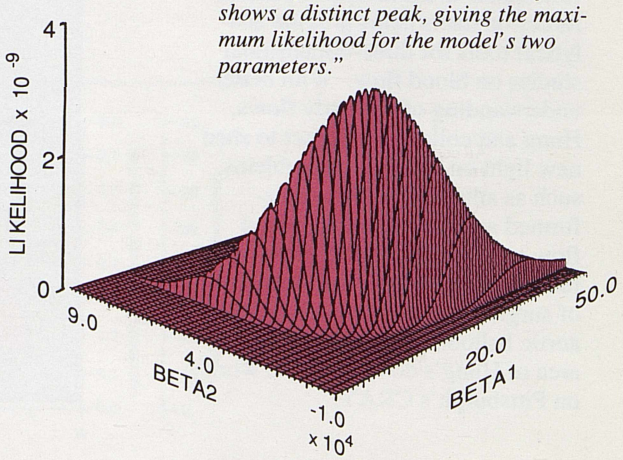

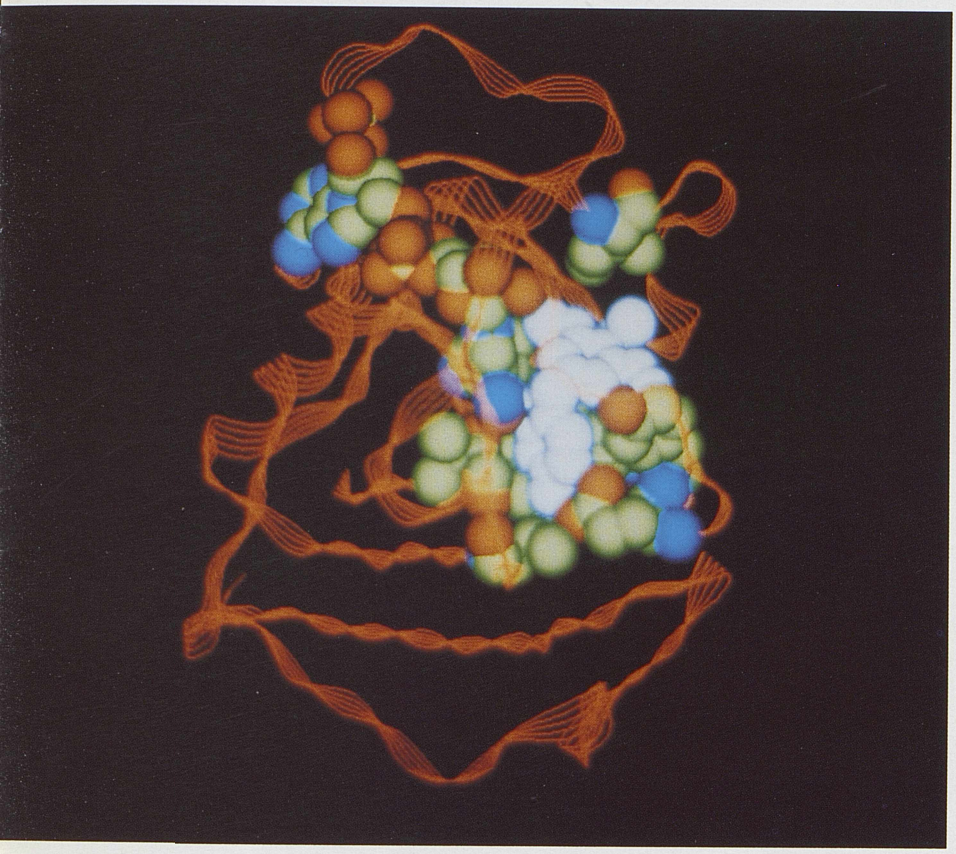

Pittsburgh Supercomputing Center
BIOMEDICAL ADVISORY COMMITTEE

BIOMEDICAL PROGRAM PRINCIPAL INVESTIGATORS

Bruce G. Buchanan
University of Pittsburgh
Department of Computer Science &
School of Medicine

Michael Levine
Scientific Director
Pittsburgh Supercomputing Center

Robert F. Murphy
Carnegie Mellon University
Department of Biological Sciences

Ralph Roskies
Scientific Director
Pittsburgh Supercomputing Center

For more information about the Pittsburgh Supercomputing Center's program in biomedical research or for workshop applications and research proposal forms, please call or write:

Cherolyn Brooks
Biomedical Coordinator
Pittsburgh Supercomputing Center
4400 Fifth Avenue
Pittsburgh, PA 15213

All computer graphics derive from computations at the Pittsburgh Supercomputing Center. Dr. Yoram Rudy provided the elelectrocardiograph map (p. 13). The graphic on this page, from the research of Charles Brooks (p. 10), produced by Scott Sneddon, Carnegie Mellon University, shows the drug trimethoprim (white) bound to the enzyme dihydrofolate reductase (red ribbon) with the co-enzyme NADPH and other protein residues. The DNA image on the cover was rendered by Hugh Nicholas and David Deerfield on an Iris workstation and photographed by Jim Schafer.

Back cover quotation: Ralph Roskies, invited talk, "Supercomputing and Biomedical Science," Conference on Grand Challenges in Computational Science, Molokai, Jan. 3, 1989.

The Pittsburgh Supercomputing Center biomedical program is supported by a grant from the National Institutes of Health, Division of Research Resources, Biomedical Research Technology Program, Cooperative Agreement #U41 RR04154.

Editor: Mike Schneider Design: Steve Palmer

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

"Improving public health,
understanding and eradicating disease,
understanding the fundamental processes of life,
developing deeper insights into evolution
and a deeper appreciation of our place in the universe —
these are truly grand challenges."