# CRAY CHANNELS

SUMMER 1990 · A CRAY RESEARCH, INC., PUBLICATION


# COLLABORATION

## Introducing the CRAY Y-MP2E supercomputer

CRAY CHANNELS

# In this issue

Sometimes cooperation increases competitiveness, as when corporations become more competitive in the marketplace by cooperating on research projects with academic researchers. Such cooperation can provide industrial engineers with technical expertise and resources to help them solve problems quickly and cost-effectively. The availability of Cray Research supercomputers at universities around the world is encouraging cooperative efforts that benefit industrial engineers as well as university researchers.

In this issue of CRAY CHANNELS, we profile collaborative university-industry efforts at the San Diego Supercomputing Center, the National Center for Supercomputing Applications, and the University of London Computing Center. We also look at academic collaborations with the public sector to develop a Great Lakes forecasting system and with a private clinic and research foundation to design new anticancer and anti-AIDS drugs. In addition, we look at new ways to visualize program performance on Cray Research systems. Our regular departments describe Cray Research's university R&D grant program, image processing applications of Cray Research systems to study brain anatomy, and Cray Research's proposed acquisition of Supertek Computers, Inc., marking Cray Research's entry into the minisupercomputer market.

In this issue we also announce the CRAY Y-MP2E computer system, which makes true supercomputer power affordable to a broad range of potential users. The CRAY Y-MP2E system costs less to own and operate than previous one- and two-processor CRAY Y-MP systems, but it does not compromise performance. And at the high end of the CRAY Y-MP product line, a new 256-Mword memory option is available. With these announcements Cray Research reaffirms its commitment to provide complete supercomputing solutions.

## Features

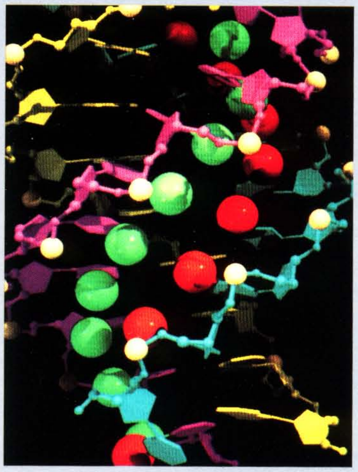

18



## Departments




CRAY CHANNELS is a quarterly publication of Cray Research, Inc., intended for users of Cray Research computer systems and others interested in the company and its products. Please mail feature story ideas, news items, and Gallery submissions to CRAY CHANNELS at Cray Research, Inc., 1440 Northland Drive, Mendota Heights, Minnesota 55120.

Volume 12, Number 2

**Editorial staff**
Ken Jopp, editor
Elizabeth Knoll, associate editor

**Design and production**
Barbara Cahlander
Eric Hanson
James Morgan
Cynthia Rykken

**On the cover:** When university and industry researchers collaborate, their shared knowledge enriches both the commercial and academic environments. Supercomputers are helping to catalyze such efforts. (Computer graphic produced with Cray Research's MPGS graphics package.)

INTRODUCING THE

# CRAY Y-MP2E

SUPERCOMPUTER

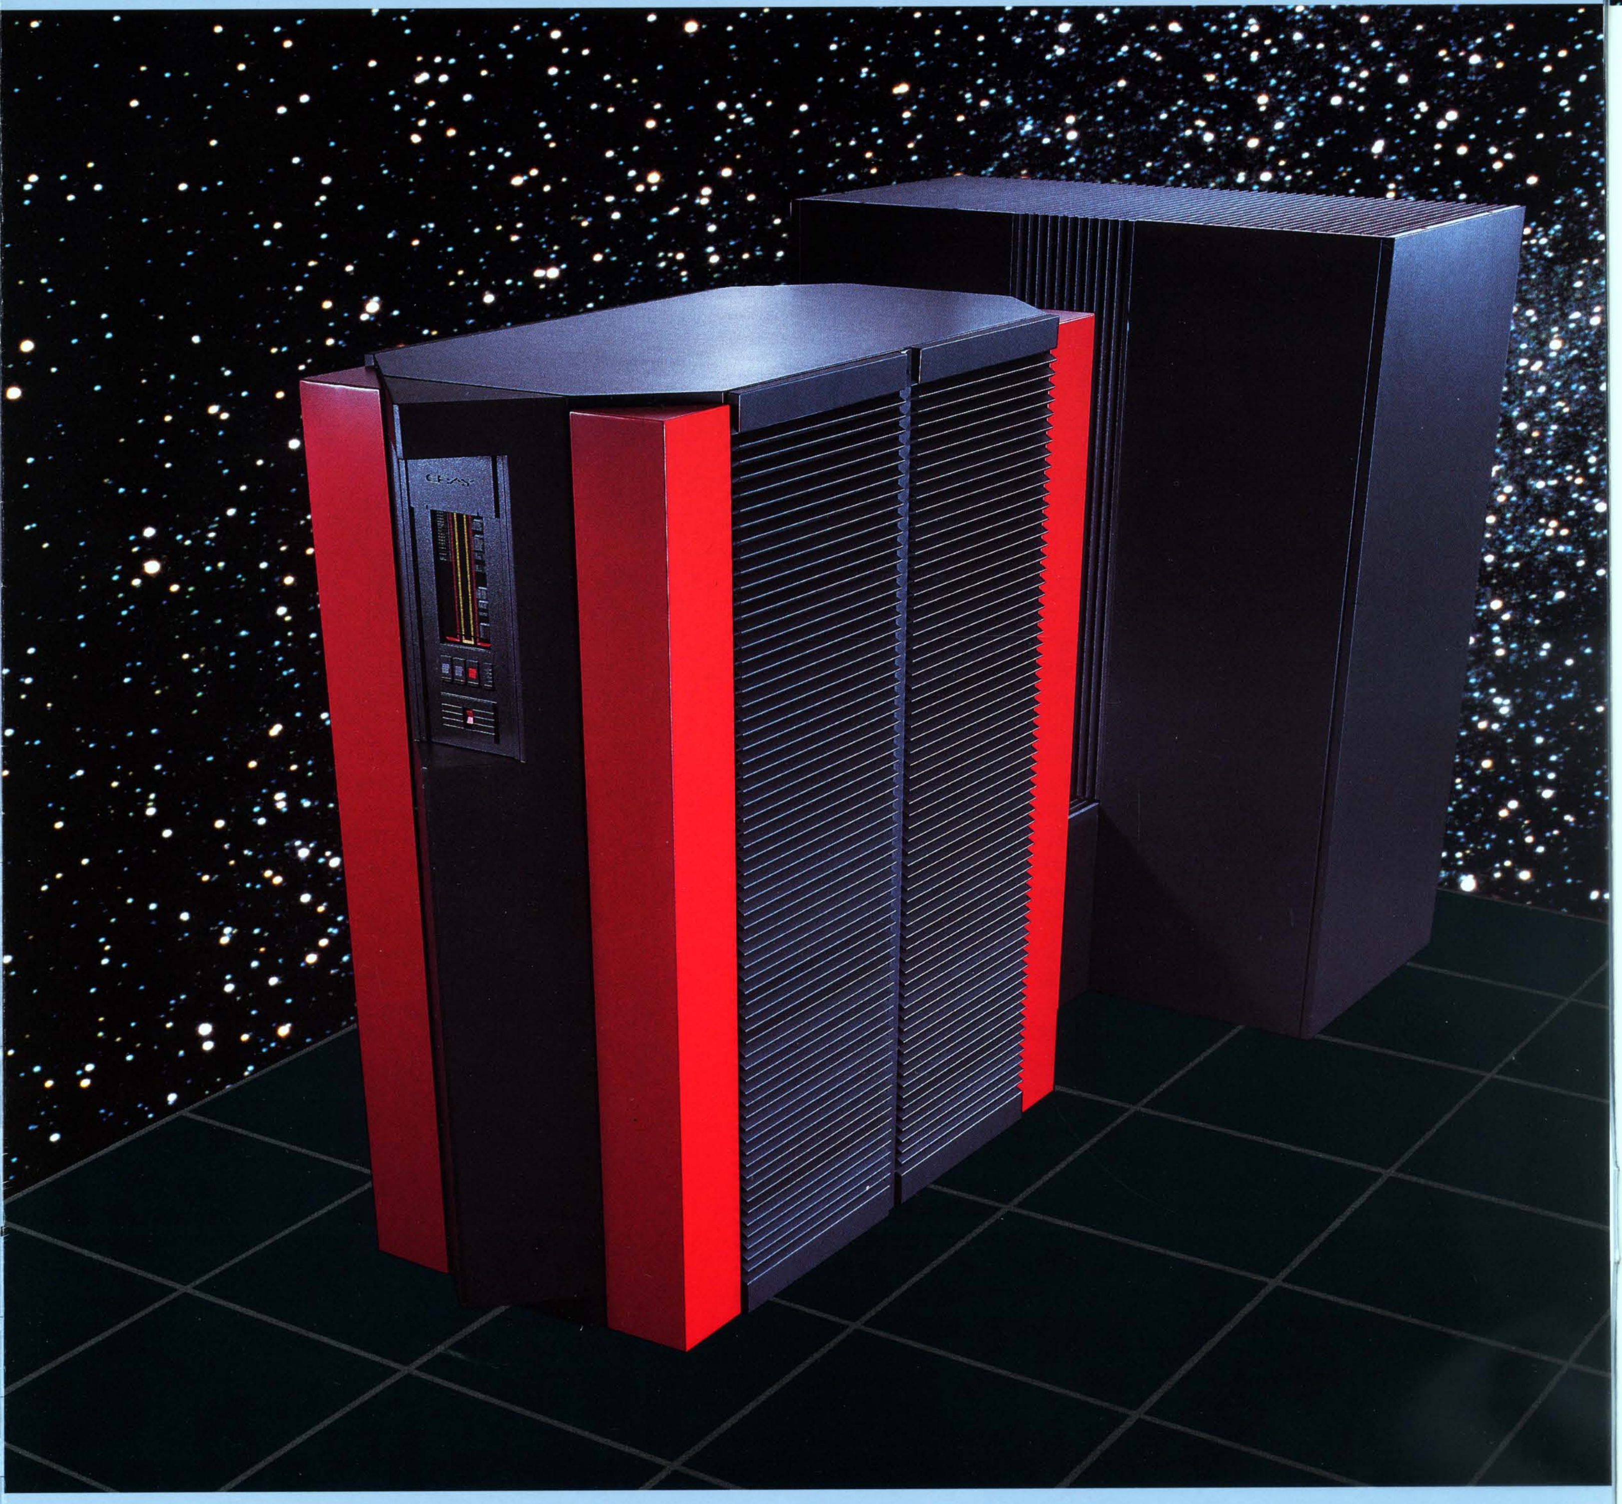

# A new standard in supercomputing

For 15 years, Cray Research has set the standard for high-speed, large-scale computing, revolutionizing problem solving in science and engineering. Now, in response to customer requests for more supercomputing options, Cray Research announces the CRAY Y-MP2E supercomputer — setting yet another standard in supercomputing. The CRAY Y-MP2E system establishes the lower boundary of real supercomputing, delivering CRAY Y-MP-class performance in a range of configurations. It leads all other systems in its class in performance and price/performance.

## A complete supercomputing solution

The CRAY Y-MP2E system is available in one- and two-processor configurations with from 16 to 64 million words of central memory. The CPU — the same as that used in the high-end CRAY Y-MP systems — achieves the highest sustained supercomputer performance in its class. For both vector and scalar code, the CRAY Y-MP2E system has a significant advantage over minisupercomputers, due to its considerably faster clock period of 6.0 nanoseconds. Its multiple-port memory also sets it apart from minisupercomputers.

## Real supercomputing within reach

The CRAY Y-MP2E computer system makes supercomputing affordable to a broad range of users. Its low cost of ownership is due to several factors:

| Model | CPUs (Processors) | Memory in Mwords |
|---|---|---|
| CRAY Y-MP2E/264 | 2 | 64* |
| CRAY Y-MP2E/232 | 2 | 32 |
| CRAY Y-MP2E/216 | 2 | 16 |
| CRAY Y-MP2E/164 | 1 | 64* |
| CRAY Y-MP2E/132 | 1 | 32 |
| CRAY Y-MP2E/116 | 1 | 16 |

* 64-Mword models and upgrades will be available in 1991.

**CRAY Y-MP2E computer system configuration options.**

- ☐ It can be air cooled and thus does not require any special site plumbing; it requires only a standard computer room environment with commercial air conditioning capacity. The CRAY Y-MP2E system also can be water cooled, at the customer's option, without extensive plumbing requirements.
- ☐ It runs on 50 or 60 Hertz electrical power; the system does not require a motor generator set.
- ☐ Customer personnel can prepare the site before Cray Research personnel arrive to install the system.
- ☐ It has a limited number of connections, making installation quick and easy. And it requires minimal maintenance due to its high level of integration and modular technology.
- ☐ It consumes only 29-50 kilowatts of electrical power.
- ☐ Its high reliability and enhanced maintenance features make possible a low-cost, on-call hardware maintenance option.
- ☐ Upgrading from one to two processors and from a smaller to a larger memory can be done quickly and easily on site.

The CRAY Y-MP2E system is a cost-effective step up for users of traditional mainframe equipment and departmental minisupercomputers; it provides a significant performance and price/performance advantage over these other systems. The CRAY Y-MP2E also offers an attractive upgrade or addition for current Cray Research customers.

The CRAY Y-MP2E system is fully compatible with other members of the CRAY Y-MP supercomputer product line. It runs Cray Research's UNICOS operating system (Release 6.0 and later) and the same compilers and applications software as other CRAY Y-MP systems. Cray Research's UNICOS operating system is based on the AT&T UNIX System V operating system. As with other Cray Research computer systems, a customer can network the CRAY Y-MP2E into almost any existing computer environment. Customers who upgrade to a larger Cray Research system can be assured of full hardware and software compatibility.

CRAY Y-MP2E systems come with the following standard characteristics

- ☐ Each CRAY Y-MP2E processor has a 6.0-nsec clock speed and four ports to main memory, the same as the larger CRAY Y-MP systems.

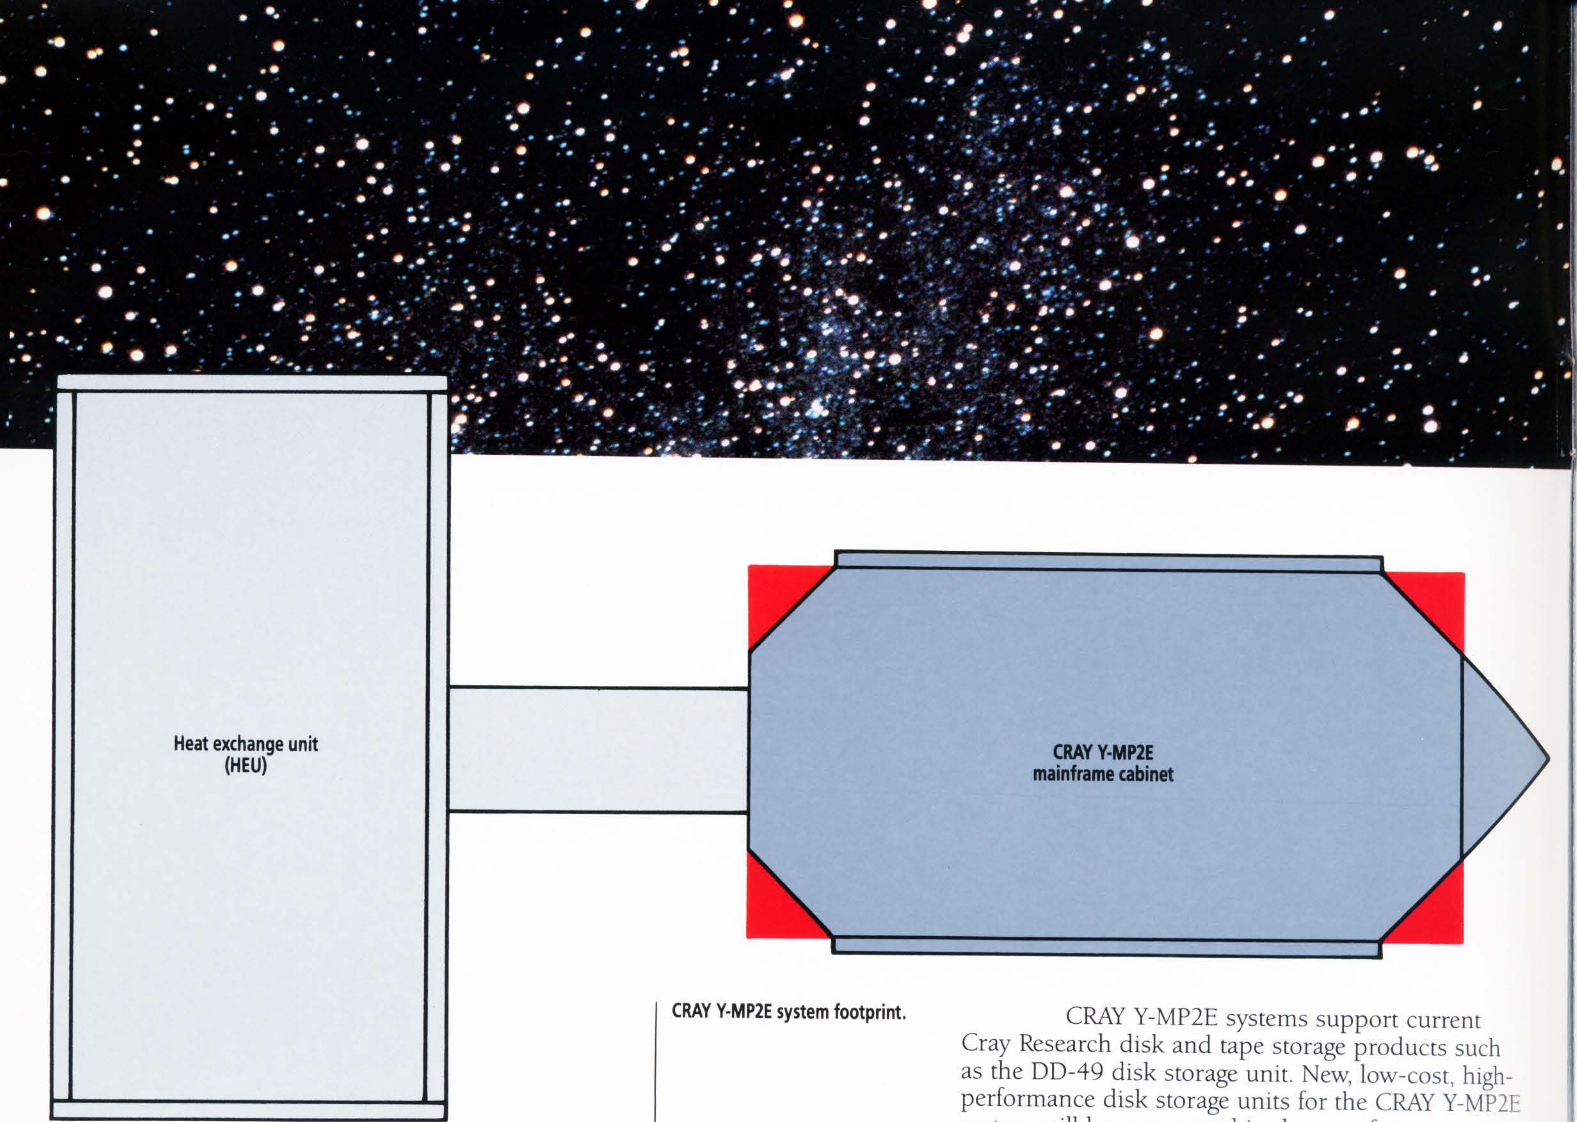


CRAY Y-MP2E system footprint.

- ☐ The CRAY Y-MP2E system uses 15-nsec access memory devices, the same as in larger CRAY Y-MP systems.
- ☐ The new integrated Input/Output Subsystem (IOS) comprises one or two I/O clusters (IOCs), which are housed in the mainframe cabinet along with the CPUs and memory. Each IOC supports up to 16 channel adapters for connection to disk units, tape units, and communications products.
- ☐ The mainframe cabinet is connected to the heat exchange unit (HEU), which contains the fully integrated cooling apparatus.
- ☐ The optional new-technology, air-cooled SSD solid-state storage device is housed in a separate cabinet connected to a separate HEU.

The smallest-configuration CRAY Y-MP2E systems are upgradeable by adding a second CPU, a second IOC, or additional main memory. Initially, main memory can be upgraded to 32 or 64 Mwords, with a future option to upgrade to 128 Mwords. Very high-speed secondary memory also can be added via an optional SSD solid-state storage device, available in 128-Mword, 256-Mword, 512-Mword, and 1024-Mword (1-Gword) sizes. As a customer's computing needs grow, a smooth upgrade path exists through the CRAY Y-MP product line.

CRAY Y-MP2E systems support current Cray Research disk and tape storage products such as the DD-49 disk storage unit. New, low-cost, high-performance disk storage units for the CRAY Y-MP2E system will be announced in the near future.

## Power and cooling simplified

The CRAY Y-MP2E system runs on common sources of commercial electrical power: 208 or 480 Vac, 3 phase, 60 Hz; or 380 to 415 Vac, 3 phase, 50 Hz. The systems do not require a motor generator set.

Cooling the CRAY Y-MP2E system also is simple. No refrigerant connections need to be made directly to the system; the system's modules are cooled by captive Fluorinert liquid coolant, which then transfers the heat to air-cooled Freon in the HEU. A CRAY Y-MP2E system also can be cooled with closed loop (treated) building water (40-85° F; 4.4-29.4° C) at the customer's option.

## A stand-alone unit or a specialist

A CRAY Y-MP2E system can function as the sole supercomputer at a site or work alongside larger Cray Research systems as a complementary machine. It is well-suited, for example, for use as a developmental system or as a separate system for secure applications. Its low ownership costs enable users to access Cray Research high-performance I/O to address file-storage needs. A file server based on the CRAY Y-MP2E system can satisfy requests from multiple supercomputers over gigabit networks while providing service to small systems, workstations, and personal computers. The diagram at right indicates potential roles the CRAY Y-MP2E

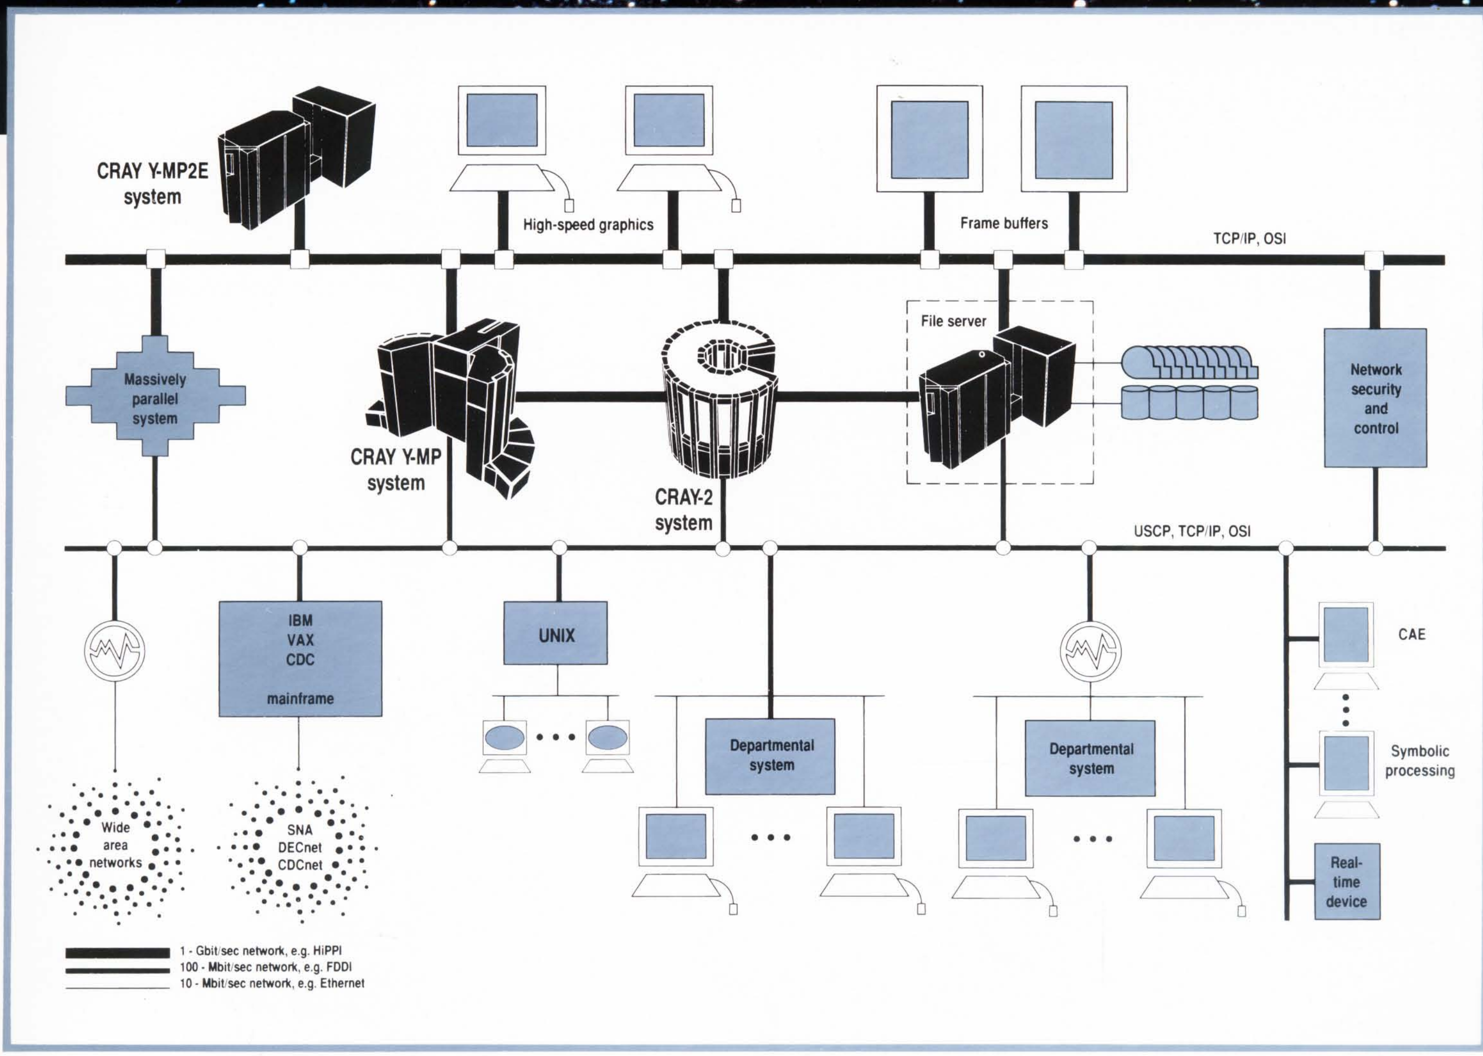


Network supercomputing environment with integrated CRAY Y-MP2E systems.

system can play in a typical heterogeneous networking environment.

## Low maintenance costs

New maintenance features and service strategies for the CRAY Y-MP2E supercomputer allow Cray Research to offer a full range of maintenance options, including an on-call service option for hardware. This option can reduce maintenance costs significantly; customers that allow remote access for service receive an additional discount.

CRAY Y-MP computer systems have established themselves as indispensable research and engineering tools in diverse industries. The latest addition to the product line, the CRAY Y-MP2E system, delivers the advantages of real supecomputing to a broader market through reduced costs of ownership. The benefit to customers is enhanced creativity and productivity and a resulting strategic advantage.

## Cray Research announces new high-end memory option

A 256-Mword memory option now is available for CRAY Y-MP8 computer systems. The larger memory is double the size of the largest previously available CRAY Y-MP memory. "The addition of a 256-million-word system to our CRAY Y-MP line reaffirms our commitment to high-performance supercomputing," said John Rollwagen, chairman and CEO of Cray Research. The larger memory improves CRAY Y-MP8 system performance particularly for memory-intensive applications such as petroleum exploration, weather prediction, and computational fluid dynamics. The new memory option incorporates 1024K BiCMOS memory technology.

# Supercomputing in the U.K.

## Initiatives for academic and industrial collaboration

*Malcolm Keech and Christopher Lazou*
*University of London Computer Centre, London, England*

One of the goals of the University of London Computer Centre (ULCC) is to encourage collaborative projects between United Kingdom industry and academia. The computing center has developed a practical program to demonstrate the benefits of supercomputing to industry. The program, which is funded partially by the U.K. Department of Trade and Industry, aims to demonstrate the feasibility and practice of applying high-performance computing to a wide range of commercial and industrial applications, and to promote wider academic and industrial awareness of the routes to collaborative use of advanced research computing facilities, such as those at ULCC.

### The development of the center

In the mid-1960s, numerically intensive computation was becoming an essential tool for advanced scientific research and engineering design in the U.K. A working party of the Department of Education and Science found that the U.K. academic community needed access to powerful computational facilities.[1] These would be funded centrally and provided from select university locations on a regional basis. The University of London Computer Centre became operational in 1968 for this purpose, and 10 years later it had graduated to a full national service. Another working party in the mid-1980s endorsed the national supercomputer center policy and recommended the collaborative use of supercomputing resources by academic projects with industrial support.[2]

### Facilities at ULCC

Today, as a national center for advanced research computing, the ULCC provides large-scale, high-performance computing services to the U.K. academic and industrial communities. Resources include a CRAY X-MP/28 supercomputer, an Amdahl 5890/300 system, a VAXcluster, and a StorageTek 4400 automated cartridge system of 1200 Gbytes capacity. ULCC is one of the network operations centers for JANET, the U.K. Joint Academic NETwork, and provides specialized gateway services connecting to international academic communities via Internet. Interactive access is available over JANET to the Amdahl system and VAXcluster where Cray Research station software provides job submission, status, and control facilities. The standard protocol for batch traffic adopted by the U.K. academic community is known as "Red Book" JTMP, Job Transfer and Manipulation Protocol, which allows remote job entry and output re-direction. Figure 1 shows network communications between JANET sites and hubs.

Applications software supported on the Cray Research system includes civil engineering, computational chemistry, and crystallography packages alongside graphics, mathematics, and statistics libraries and programs. The civil engineering codes available are DYNA3D, LUSAS, and MSC/NASTRAN, with assorted pre- and postprocessors, and the NAG finite element library. GAMESS and GAUSSIAN 86 are the premier quantum chemistry codes, but a large number of additional programs also are provided through the National Working Party on Computational Chemistry. Crystallography traditionally has been a major area of

**Figure 1. U.K. Joint Academic NETwork.**

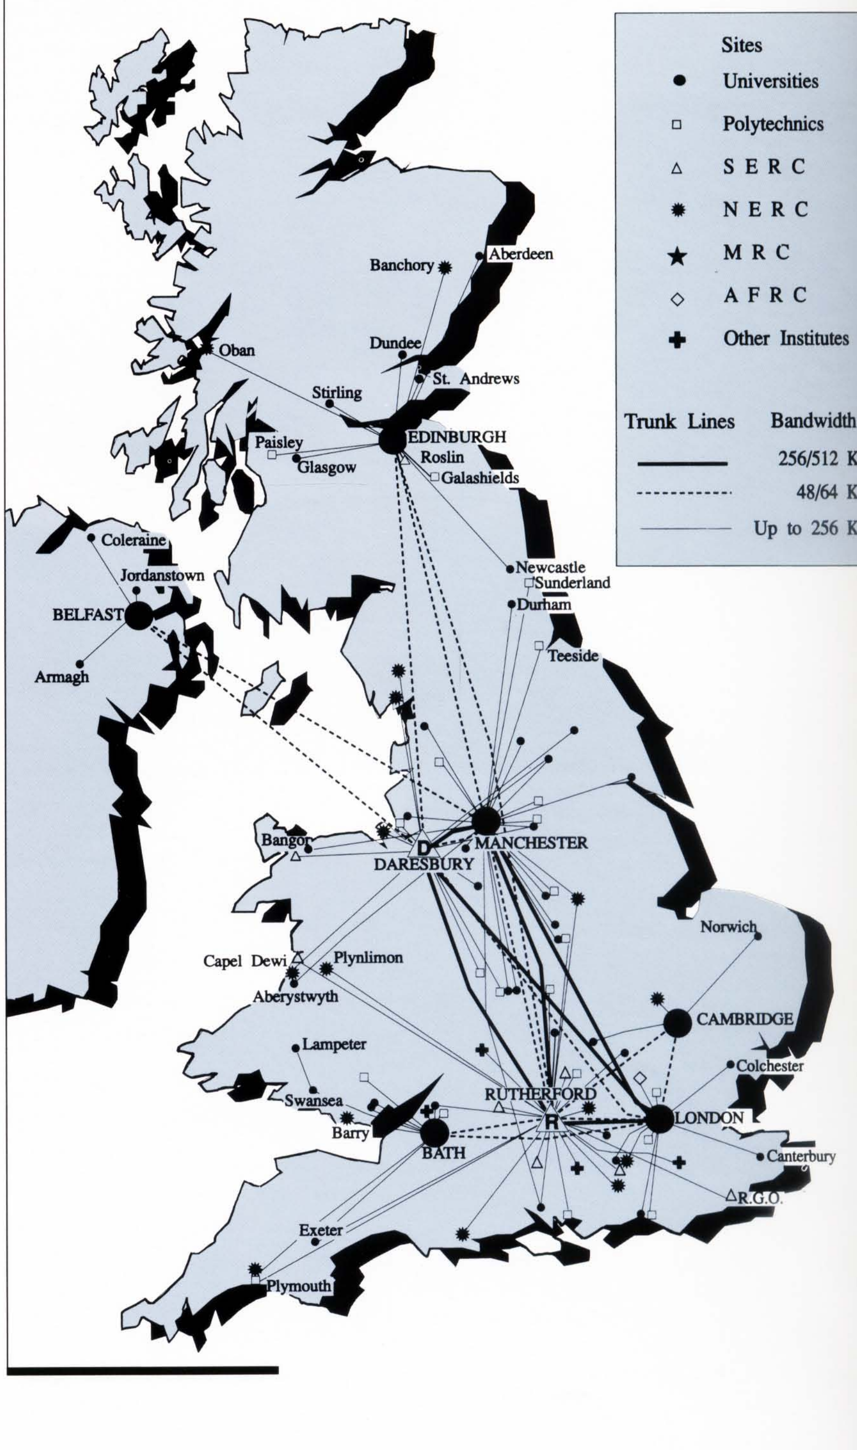

# Case study: The retroviral proteinase from HIV-1

Andrew Hemmings of the Crystallography Department at Birkbeck College is pursuing several experimental, theoretical, and computational approaches to research into the HIV-1 retrovirus. His work, partially supported by Cray Research, is one example of the many varied research projects at ULCC.

The discovery that aspartic proteinases play an active role in the process of retroviral maturation has opened an exciting new area of research for those interested in their internal mechanism. These enzymes are novel because they are active catalytically as homodimers, placing important restraints on our ideas of the role of protein environments in catalysis. The HIV-1 retrovirus is one of a large family of enveloped viruses containing a single-stranded RNA genome, which replicates via an obligatory DNA intermediate. Its notoriety stems from its role as the etiological agent of AIDS. The retroviral genome consists of three major genetic elements that are arranged in the order 5'-gag-pol-env-3'. The pol gene codes for the enzymes involved in viral replication, including a reverse transcriptase, an integrase, and a proteinase. This aspartic proteinase is responsible for cleavage of its own precursor, a gag-pol fusion polyprotein.

Attempts to develop antiviral drugs so far have centered on the viral reverse transcriptase, but the proteinase also may prove to be a suitable target, as the enzyme has been shown to be essential for maturation. Thus, inhibitors of viral proteolytic processing, which do not affect the metabolism of the host cell, could prove to be of use as drugs. The techniques of x-ray crystallography have been used to solve the structure of the native proteinase at 2.7Å resolution[1], correcting an earlier published structure containing a misinterpreted inter-subunit strand packing arrangement. In this process, Hemmings extensively used the CRAY X-MP/28 system at ULCC for refinement of the structure against observed structure factor amplitudes. Both conventional restrained least squares and simulated annealing procedures were used to give the final model shown in Figures 1 and 2.[2]

The details of the inter-subunit region obviously are important if inhibitors of the dimerization process are to be designed rationally. From the refined coordinates it is observed that the structure of this dimeric enzyme provides an active site region with very similar disposition of hydrogen bonding functions to that of the monomeric examples of this family from fungi and mammals. Knowledge gained from work on these systems can be used in retroviral proteinase inhibitor design studies. On the other hand, because the structure of the viral enzyme is symmetric, binding of an asymmetric inhibitor will lead to a loss of symmetry. One potentially useful avenue of research should then be through the development of symmetric inhibitors, hopefully maximizing binding affinity through exploitation of entropic considerations.

***References***


1. Lapatto, R., et al., *Nature,* 342, pp. 299-302, 1989.
2. Brünger, A. T., "Crystallographic Refinement by Simulated Annealing on Supercomputers," *CRAY CHANNELS,* Fall 1988.


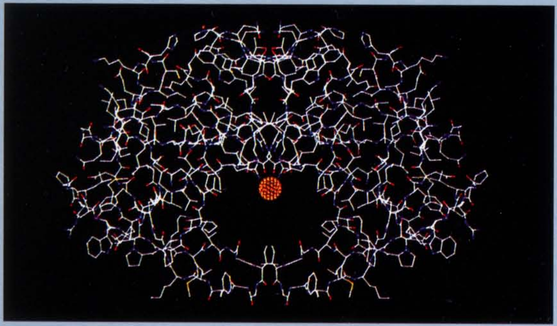

**Figure 1 (left). Structure of HIV-1 proteinase as refined at 2.7Å resolution.**

**Figure 2 (right). Space-filling representation showing the two subunits of the HIV-1 proteinase.**

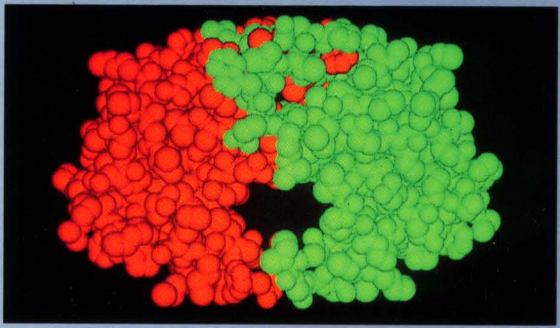

interest at ULCC. The range of codes available includes EXPLORE, MULTAN80, SHELX, and XTAL. Other software is coordinated through collaborative computational projects by the Science and Engineering Research Council in areas such as aerodynamics, fusion plasmas, and lattice gauge theory.

## Apportioning ULCC's resources

About 4500 users access the computer center.[3] Resources are allocated in four sectors: a fully peer-reviewed sector, a pump-priming sector, the university sector, and the externally-supported sector. The majority of the center's resources — between 55 and 65 percent — is granted to users in the fully peer-reviewed sector. Individual allocations are made by research councils and the British Academy after a full peer review, and at no cost to universities or other higher-education institutions. The pump-priming sector is apportioned up to 5 percent of the resources for users' pilot projects. Small allocations of up to 0.1 percent for periods of up to three months are made by the ULCC director, again at no charge to the academic community. For small or beginning users, the university sector has access to up to 20 percent of the center's capacity. Individual allocations may not exceed 0.1 percent. Universities and other institutions are invited to take part in an annual round of bids for resources. Up to 10 percent of the center's resources may be used by academic projects with external support, such as commercial or industrial collaborations. A charge is normally levied in this sector.

## The Association of Supercomputer Users

ULCC launched the Association of Supercomputer Users (ASU) in 1989 to provide a forum for academia, industry, and government research establishments to discuss issues of common concern. At a practical level, ASU membership provides facilities to explore numerically intensive computation with minimal commitment of expenditure on equipment and staff.

The project, partially funded by the U.K. Department of Trade and Industry, began in 1989 with four seminars on the use of supercomputers, two of which were management presentations on systems architecture initiatives, academic initiatives, and academic-industrial collaboration. The two technical presentations covered drug design in the field of pharmaceutics and the use of MSC/NASTRAN, a large-scale structural analysis package.

ASU membership includes the use of ten CPU hours on the CRAY X-MP/28 system at ULCC, including front-end access, training courses for two people on how to use the facilities, one-day seminars on high-performance computing and its benefits, and a place at an annual seminar organized by ULCC with hardware vendors and applications software presentations. Members also receive two copies of the monthly ULCC News to keep up to date with the service, a copy of a report describing current work on the Cray Research system at ULCC, and a copy of the ULCC Annual Report. Affiliates pay £5000 in advance and receive computer services including production running of codes and advanced consultation in areas including numerical methods, vectorizing techniques, and parallel processing.

## Externally supported research

The collaborative projects at ULCC represent fields from computational chemistry to mechanical engineering and bring together academic and industrial organizations throughout the U.K. For example, researchers from the Crystallography Department of Birkbeck College are using ULCC resources to study the HIV-1 retrovirus (see sidebar). Researchers from the Mechanical Engineering Department of the University College London (UCL) are developing a model that will predict the motion of a marine structure in a hostile ocean environment, an important topic in marine hydrodynamics.

The total response of the structure to the ocean wave is determined by superposition of responses to individual regular waves. A coupled finite element method has been employed at UCL that represents the fluid flow surrounding the structure by finite elements, matched by other forms of the solution in the remaining fluid region. Because at least 10,000 nodes will be needed for sufficient accuracy in some industrial cases, the power of a Cray Research system is needed to calculate the 10,000-by-10,000 solution matrix. The researchers also will be using the Cray Research system to investigate the response to large waves in the time domain, a pressing and more complicated problem.

Here are some examples of defense-related collaborative research projects currently being conducted at ULCC:

- ☐ Researchers at the University of Exeter's Mathematics Department are examining the structure and stability of oceanic fronts using mathematical and computer models.
- ☐ At the Marine Science Laboratories of the University College of North Wales, researchers are simulating the creation of oceanic eddies around seamounts with high aspect ratios in stratified waters to investigate acoustic propagation.
- ☐ At the Department of Applied Mathematics and Theoretical Physics at the University of Cambridge, researchers are using the Cray Research system to determine the onset of instabilities in the motion of a rotating body caused by fluid motion within a contained cavity.
- ☐ Researchers in the Chemistry Department at Royal Holloway and Bedford New College are performing ab initio molecular orbital calculations on magnesium-containing systems.

**The collaborative projects bring together academic and industrial organizations throughout the U.K.**

## Conclusion

The ULCC program to foster links between U.K. industry and academia aims to demonstrate the benefits of supercomputing to industry and in the process enhance the competitiveness of the United Kingdom. In anticipation of the European single-market community in 1992, appropriate initiatives are being explored, and ULCC is building close ties with other European centers of excellence.

*For more information about ULCC's Association for Supercomputer Users, please contact Malcolm Keech or Christopher Lazou at the University of London Computer Centre, 20 Guilford Street, London WC1N 1DZ, England.*


### *Acknowledgments*

*The authors thank Andrew Hemmings of Birkbeck College for the case-study material, Guoxiong Wu of University College, London, for his input, and ULCC Workstation Support.*



### *About the authors*

*Malcolm Keech is manager of Marketing and National Services at the University of London Computer Centre. He received a B.A. degree in mathematics from the University of Oxford in 1973 and a Ph.D. degree in numerical analysis from the University of Manchester in 1977. He has worked for several years in academic and industrial computing, notably at the University of Manchester Regional Computer Centre and the Florida State University Supercomputer Center, before joining ULCC at the start of this year.*

*Christopher Lazou is manager of Supercomputer Futures Development at the University of London Computer Centre. He received a B.Sc. degree in mathematics and physics from the University of London in 1966. After three years in industry, he joined ULCC as a senior systems programmer in 1969 and subsequently specialized in compilers and graphics software. He is vice president of the Cray User Group and author of the book* Supercomputers and Their Use, *published by Oxford University Press in 1986 and 1988 (revised edition).*


### *References*


1. Flowers, B., "Report of a Joint Working Group on Computers for Research," Her Majesty's Stationary Office Command 2883, High Holborn, London, 1965.
2. Forty, A. J., "Future Facilities for Advanced Research Computing," Joint Working Party Report for Advisory Board for the Research Councils, Computer Board for Universities and Research Councils, Universities Grants Committee, Department of Education and Science, Elizabeth House, London, 1985.
3. Altmann, J. A., ed., *Supercomputer Assisted Research*, University of London Computer Centre, London, 1989.

# Industrial engineering-university collaboration

## A new mold

*Merry Maisel, San Diego Supercomputer Center, University of California at San Diego*

To remain economically competitive, countries that have long been industrialized must find new ways to coordinate the research and production segments of their economies. The more efficient technology transfer that results can help get products to market sooner. Efforts to bring research and production together are being made by the four national supercomputer centers established in the United States by the National Science Foundation. The center at San Diego provides an example of industry-university cooperation for technology transfer.

The San Diego Supercomputing Center (SDSC) is located on the campus of the University of California, San Diego. Like the other centers, it has a mission to lead in expanding the use of supercomputers among U.S. academic and industrial researchers. The main computer at SDSC is a CRAY Y-MP8/864 system, which includes a 128-Mword SSD storage device, integrated into a hierarchy of powerful minisupercomputers, graphics engines, and workstations. SDSC maintains a team of researchers in various disciplines and specialists in scientific visualization and research optimization.

Among SDSC's industrial participants is GenCorp Research, a technology-based company with strong positions in aerospace, automotive, and polymer products. GenCorp has three operating segments: Aerojet, GenCorp Automotive (which includes General Tire), and GenCorp Polymer Products. The product line ranges from solid rocket propellants to reinforced plastic auto components to tennis balls.

Jose Castro and Dick Griffith, scientists at GenCorp Research in Akron, Ohio, are specialists in polymer technology, a technology that lies behind virtually all of GenCorp's products. An important component of GenCorp's polymer business is the production of auto body panels from compression-molded sheet-molding compound (SMC), a fiber-reinforced plastic. A research partnership between the GenCorp scientists and university researcher Charles Tucker led to an application of the SDSC Cray Research system that helped solve a compression-molding problem. As a result, the quality of the panels was assured during the design of the manufacturing process and academic researchers saw their code validated in a commercial application.

### Why plastic?

Reinforced plastic has several advantages for use in automobile "skins." Although gas lines are not as long as they were in the 1970s, the need still exists for lightweight, fuel-efficient cars. Moreover, interest in conserving fossil fuels has grown in step with environmental concerns. Reinforced plastic parts are not only lighter than competitive materials, but also can be stronger and more flexible on impact. High glass content and long fibers can give SMC parts an impressive combination of stiffness and shatter resistance relative to competitive materials, yet SMC can be molded into virtually any shape.

Castro and Griffith, along with associates C. C. Lee and Elliott Straus, wanted to find the optimum molding conditions for compression-molding of SMC into automotive body panels.[1,2] Fiber-reinforced plastic had been injection-molded and compression-molded to produce the body panels for the Chevrolet Corvette for the past decade, but the Corvette is produced in low volume. To manufacture panels for a higher production volume, the researchers needed to consider the speed with which the compression-molding process could be accomplished, which is partly a function of the press. GenCorp's new plant in Shelbyville, Indiana, has press capacities that range from 500 to 2000 tons, all highly instrumented and automated. But production speed also is a function of the way in which the material flows through the mold, which is what the scientists set out to simulate.

Fortunately, they didn't have to start their modeling effort from scratch. GenCorp had been supporting computational research by Professor Charles Tucker of the Mechanical and Industrial Engineering Department of the University of Illinois at Urbana-Champaign, where Tucker and his students were studying ways to model the molding of fiber-reinforced polymers. Castro and his colleagues were able to modify and use the codes developed by Tucker and his group, but they needed to run many models to explore the parameter space and design the optimum system, a task for which GenCorp's own computers were too slow.

### Finding the right computer

The researchers then turned to Alonzo Church and Dan Fleming, scientists in the Applied Mathematics and Computer Applications Department at GenCorp. Fleming had gained experience with supercomputing over the past several years working in oil reservoir simulation while at Phillips Petroleum Company in Bartlesville, Oklahoma, and later at GenCorp. In 1986, GenCorp became an industrial participant of SDSC through work conducted by the company's Aerojet division. Church, in charge of the department, had helped engineers from Aeroject ElectroSystems use SDSC's CRAY X-MP/48 system, which preceded the CRAY Y-MP system, to calculate the wear on bearings for a space-borne earth-scanning antenna.

"It was logical for Castro's group to do the work at SDSC," Church explains. "As an industrial participant in the SDSC program, we had access to a

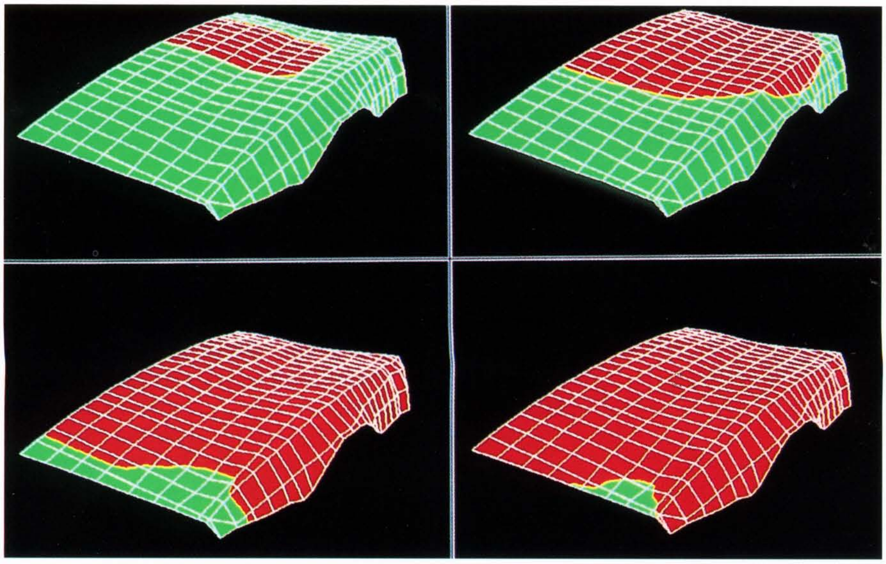

into the practical arena of materials design. "That's where computational power is an advantage," he adds.

Castro adapted and used codes developed by Tucker in a computational environment supplied by Fleming and Church, along with SDSC, to design the manufacturing process for a specific product: body panels for three vehicles. These were a new Chevrolet Lumina minivan, the Pontiac Trans Sport, and the Oldsmobile Silhouette. The researchers verified that TIMS correctly simulates the shapes of SMC flow fronts by taking measurements in an automotive hood mold.[1]

During the SMC compression molding process, a premeasured SMC charge is placed between the heated halves of a mold. The halves then are brought together to squeeze the SMC and fill the mold, after which pressure is maintained while the SMC cures. A complete molding cycle has four stages: SMC charging, that is, putting the compound in the mold; mold closing; curing; and mold opening so the part can be removed. The GenCorp body panels are given

supercomputer, but we didn't have to staff or maintain it. We could build our own expertise for the future at an appropriate level of investment. And we had access to the expertise of the SDSC's User Services group, some of whose consultants are themselves engineers with training in finite-element methods."

The code used by Castro and his group was called TIMS: Thin Injection Molding Simulation. It had been developed by Tucker and two of his students, Tim Osswald,[3] who is now an assistant professor at the University of Wisconsin, and Suresh Advani,[4] now an assistant professor at the University of Delaware. The code was developed specifically to model the flow of fiber-reinforced polymers in compression molding. Professor Tucker notes that similar codes are used to model a somewhat simpler problem, unreinforced polymers shaped by injection molding. "But our code is unique," he adds, "because it simulates compression molding and predicts the flow of the material and the small-scale structure of the finished product — the overall orientation of the fibers."

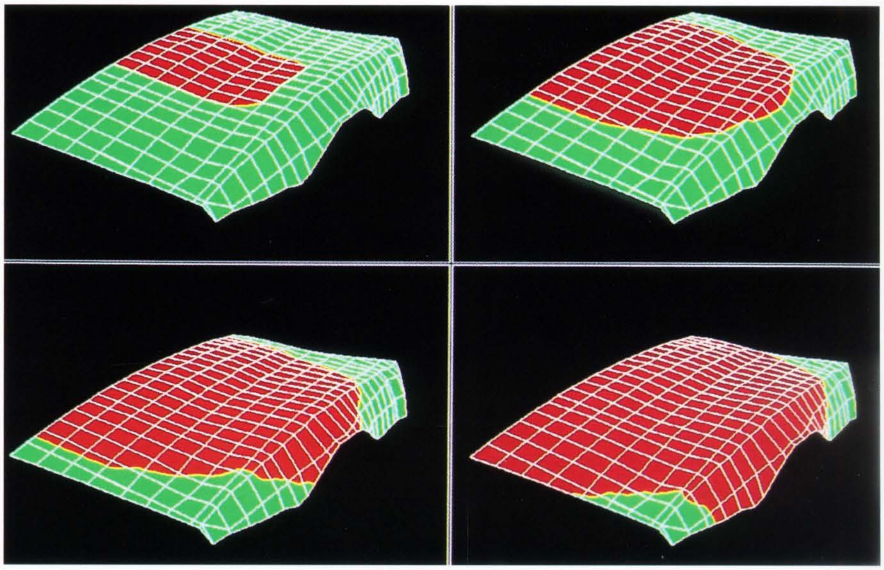

**Filling pattern for automotive hood outer panel with charge placed at rear (top left sequence). Simulation predicted formation of a knit line at the corner nearest the viewer. As simulations were run with the charge placed successively nearer the front, the knit line continued to form (lower left and lower right sequences), until the charge was placed very near the front, over the headlamp. Simulation predicted that the knit line would not form with the charge placed here (top right sequence), and in practice it did not. Simulation by Dan Fleming, GenCorp Research, Akron, Ohio.**

## Control-volume filling

The usual finite-element mesh, Tucker points out, represents the entire domain. "But in our problem, as the material fills the mold, the shape of the material domain is constantly changing." The solution was to begin with a cavity of known shape, charged with a blob of the polymer, also of known shape, and then to follow its spread through a "control volume" scheme. The code partitions the compression-molding cavity into control volumes and tracks the filling of each volume as the compound reaches it. Tucker says, "The method is a bit crude at the edges, but it is very easy to do complex geometries. So it's the only code we know of that will model compression molding for realistic part geometries — the complex moldings that make up auto bodies, for example."

Tucker also points out that many people who model materials processing or manufacturing routinely calculate heat transfer, fluid flow, and stress. By calculating the fiber orientation, however, his group was adding the capacity to predict the microstructure of a material, bringing the dreams of materials science

a protective coating while in the mold, a process that also had to be modeled. "It's a little like making waffles," Church says. "You have to put the batter — the mold charge — in the right place, and you have to start with the right amount."

Fleming installed TIMS on the SDSC CRAY X-MP system in 1988, accessing it from a MicroVAX II computer over a DECnet dial-up link.[5] To preprocess the input and display the predicted fill patterns as a function of time, senior research engineer Doug Millward implemented a package from Structural Dynamics Research Corporation called I-DEAS. Millward wrote interfaces between I-DEAS on the GenCorp workstations and TIMS on the Cray system. Castro's group ran the overall system, and Straus later went to Shelbyville to oversee the implementation of the molding process.

## Knit picking

"One of the problems shown by the modeling," Castro says, "was the appearance of knit lines, which form where two flow fronts meet or when inserts or obstructions are present in the mold cavity. These will be weak spots in the molded part, because when the flow is nearly parallel to the knit lines, very few fibers bridge them. We discovered that we could use the simulation to eliminate knit lines or move them to less critical areas by relocating or redesigning the shape of the SMC charge."

Today, the manufacturing process is working at GenCorp's giant presses in Shelbyville. The vans are out on the streets, and GenCorp is continuing to design and manufacture similar parts for other vehicles, to be made by General Motors and other automobile manufacturers. "I consider this a model industry-university collaboration," says Church.

Tucker agrees. "Industry benefits from this kind of cooperation with academic engineering, and vice versa. That's how it should be done." Tucker exemplifies the method: his graduate work had industrial sponsorship and he has obtained such sponsorship for 10 years. Current sponsors include GenCorp, Premix (which also produces auto body parts), and Norsolor, a French firm in a similar business. Tucker's

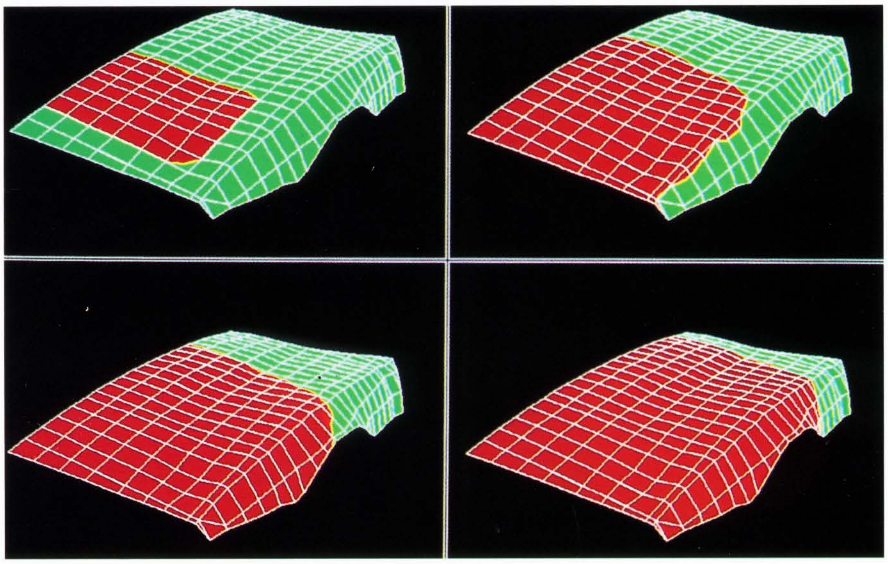

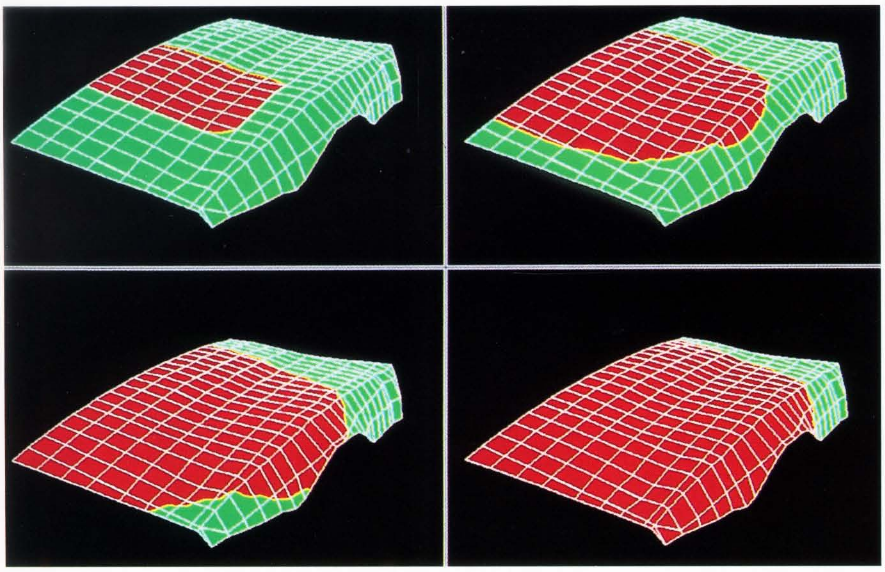

group continues to work on the simulation of mold-filling, developing theories and numerical methods to handle cases beyond the current capabilities of TIMS. Tucker has just edited *Fundamentals of Computer Modeling for Polymer Processing* (Hanser Publishers/Oxford University Press, 1989), for researchers who want to write or use such simulations.

Tucker regards the collaboration with GenCorp and SDSC as particularly successful for several reasons. "GenCorp's people had the wisdom to see that the code we had written was not a black box. Jose Castro and the members of his group took the trouble to understand the code. They knew it well enough to tell us about the bugs. And Al Church and his group took time to work with the code and integrate it into their own computational environment, making it work with their other computational tools.

"For our part," he adds, "we took the trouble to write clean code and document it. But no academic research group is meant to be a commercial software packager, so it's great to work with an industrial group that can build a bridge between university research software and something that can be used as a reliable engineering tool." For Tucker, the products are ideas, techniques, and well-trained students. The software carries the ideas and techniques to industry, and it functions as a tool for the industrial collaborators. The efficiency of the tool is increased by using it in a supercomputer environment, and the San Diego center was crucial. The center benefited as well, learning to provide greater expertise and service to university and industrial researchers in collaboration. For Castro, Griffith, Fleming, and Church, the ultimate products are the auto body parts, which demonstrate the advantages of polymer technology.

### About the author

*Merry Maisel is a science writer at the San Diego Supercomputer Center and recently became a graduate student in a new program at the University of California at San Diego called Science Studies. Her primary research interest is in the impact of large-scale computing on the sciences.*

### References

1. Castro, J. M., and R. M. Griffith, "Kinematics and Dynamics of SMC Compression Molding," 43rd Annual Conference, Composites Institute, The Society of the Plastics Industry, Inc., February 1-5, 1988, session 17-E.
2. Castro, J. M., and C. C. Lee, "Thermal and Cure Analysis in Sheet-molding Compound Compression Models," *Polymer Engineering and Science*, Vol. 27, No. 3, p. 218, 1987.
3. Osswald, T. A., and C. L. Tucker, "Compression Mold Filling Simulation for Nonplanar Parts," *International Polymer Processing*, in press, 1990; Osswald, T. A., and C. L. Tucker, "An Automated Simulation of Compression Mold Filling for Complex Parts," SPE Technical Papers, ANTEC, p. 169, 1985; Osswald, T. A., "Numerical Methods for Compression Mold Filling Simulation," Ph.D. dissertation, Department of Mechanical and Industrial Engineering, University of Illinois, 1984.
4. Advani, S. G., and C. L. Tucker, "A Numerical Simulation of Short Fiber Orientation in Compression Molding," Polymer Composites, Vol. 11, No. 3, 1990.
5. Fleming, P. D., J. M. Castro, and E. J. Straus, "Compression Mold Filling on a Remote Supercomputer," presented at the First Ohio Supercomputer Center Symposium on Supercomputing, Columbus, Ohio, September 8-9, 1988; Fleming, P. D., and R. L. Riedinger, internal GenCorp report (unpublished).

# The impact of synergy
## An example from the NCSA-FMC partnership

*Ilhan Dilber, National Center for Supercomputing Applications, Urbana, Illinois*

*Wing Cheng, FMC Corporation, Santa Clara, California*

The National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign was established in February 1985 with a National Science Foundation grant and was opened to the national research community on January 15, 1986. The state of Illinois, the university, and private industry, through industrial partnerships, provided additional funding. The center's mission statement calls for the center "...to use the tools and techniques of leading-edge computational science to contribute to new advances in science and engineering, strengthen the industrial competitiveness of American industry, and create new advances in computational science." To accomplish this goal, NCSA offers industrial partners access to a CRAY-2 computer system and a CRAY X-MP/48 system.

The mission of NCSA's industrial program is to form partnerships between the center and major corporations to strengthen corporate research and support new areas of research and service within NCSA. Full partners make major commitments over a period of years in exchange for participation in research at NCSA. Participants receive supercomputing training individually tailored to their needs along with access to the center's technology. In 1986, the Eastman Kodak Company became the first corporation to join NCSA. Today the program includes six additional partners: Amoco Corporation, Caterpillar Inc., The Dow Chemical Company, Eli Lilly and Company, FMC Corporation, and Motorola Inc.

FMC Corporation joined NCSA in February 1989 as the fifth industrial partner. The Chicago-based company is one of the world's leading producers of machinery and chemicals for industry, government, and agriculture. The company's businesses fall into five major segments: industrial chemicals, precious metals, defense systems, performance chemicals, and machinery and equipment.

One of the first projects to use the resources at NCSA originated at FMC's Corporate Technology Center (CTC) in Santa Clara, California. The CTC is the corporate research and development center for machinery, manufacturing automation, and computer sciences. Wing Cheng, a senior technologist at CTC, working with Scott Langlie, a project leader in FMC's Advanced Systems Center in Minneapolis, Minnesota, was researching the development of insensitive munition systems for use onboard naval ships. The project was aimed at improving the design of the munition casing material to prevent the explosive or propellant inside from being detonated by an accidental fragment or bullet impact. The armor had to be designed not only to defeat or adequately withstand such threats, but also to produce only low shocks, which would not detonate the contents. (To attain this design goal, new advanced materials, such as composites and sandwiched structures, also are being tested.)

Experimental approaches to this problem would be not only very costly, but also destructive, and therefore would not reveal details of the process, such as shock propagation, gradual deformation, and the rupture of the materials. Hence, successful numerical simulations were an invaluable part of the design process. For the numerical analysis, researchers used a hydrodynamic code based on the DYNA finite element code from Lawrence Livermore National Laboratory (LLNL). The finite element method provided the best approach to the analysis due to the complexity of the geometry. The models had to be discretized with fine meshes because the problems being modeled were very sensitive to the details of the shocks. Supercomputing was required because even the two-dimensional problems required substantial core memory, and three-dimensional problems would not have been solvable on most mainframes. To run the models, a special FMC-enhanced version of the DYNA code was ported to the NCSA's CRAY-2 system, which has 128 Mwords of static memory.

The researchers modeled an 80-gram fragment-simulating projectile hitting a 1.75-inch woven roving Kevlar-49 reinforced polyester layered composite panel at a velocity of 5699 feet per second. A 1.5-inch steel witness plate was placed behind the composite panel to stop the projectile after it perforated the panel. This problem corresponded to a configuration that previously had been subjected to full-scale ballistic tests. The results of the analysis were saved in a series of binary files, which then were postprocessed by the ORION program, also from LLNL. The ORION program allows the generation of contour plots for various parameters at given time steps. Static frames, however, do not provide enough information about the problem dynamics, such as shock propagation and deformation during puncturing. Soon after the first results were obtained, the need for real-time animation became obvious.

The visualization goal became the generation of color raster images in NCSA's Hierarchical Data Format (HDF) which then could be animated on desktop computers, such as the Macintosh II computer, or workstations such as those from Sun Microsystems. The HDF was developed at NCSA to store different types of data, such as floating-point numbers, raster images, and annotations, in a single file and to transfer the file among machines and operating systems. NCSA has developed various software tools to manipulate HDF files, making possible animation, annotation of images, floating point data manipulation, and conversion to raster images. Once the color raster images were stored in the HDF files on the Cray Research system, they could be transferred easily to a Macintosh computer and animated in real time.

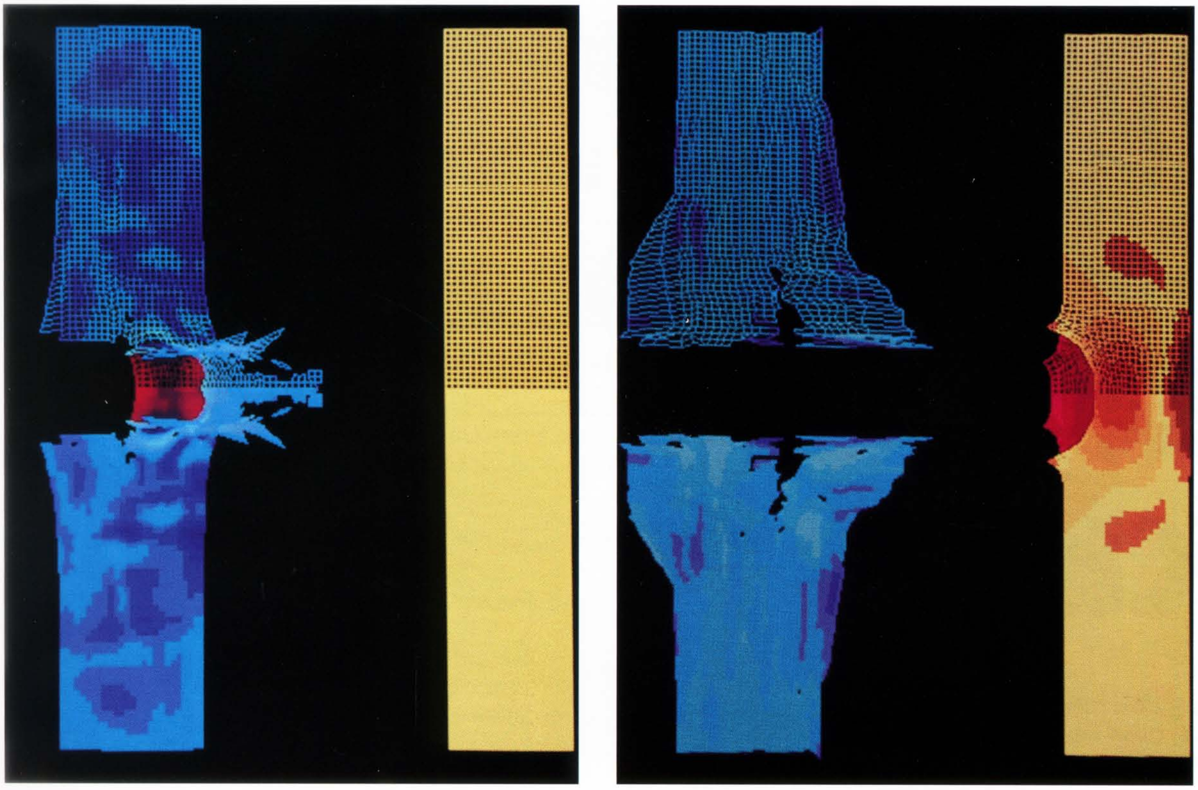

**Figure 1. Selected frames from the animation of the projectile simulation.**

The first step was to modify the ORION package to write out the nodal coordinates, the element connectivity, and the nodal values of a given parameter to a binary file for each time step. Next, a program was developed to read that information and convert it to color raster images and create the frames to be animated. The frames then were transferred to a Macintosh computer, annotated using the NCSA Layout program, and animated with the NCSA Image program. Once the annotations and colors were found to be satisfactory, the frames were recorded on videotape for presentation. Selected frames from the animation are shown in Figure 1.

FMC engineers completed their first scientific visualization project within three months. The presentation included the two-dimensional animations generated as described above, as well as more sophisticated three-dimensional animations created by NCSA's Visualization Services and Development Group. The visualization was prepared for presentations to the company's upper management. FMC video production specialists added live interviews and narrations to the animation sequence and made a 25-minute promotional tape that describes the value of visualization. The presentation later was made available to NCSA and since has been distributed widely in the academic and industrial communities.

After demonstrating the importance of scientific visualization, FMC's Corporate Technology Center acquired a graphics workstation from Stellar Computer, Inc., to produce similar video presentations in-house. As part of that effort, technologist Philip Ho has ported the HDF libraries to the workstation and has incorporated the HDF graphics into diglib, the graphics library for various LLNL software, including DYNA and ORION. As a result of this work, HDF is one of the device-driver options in those programs, and the images can be displayed as color raster images and saved in HDF files directly from within the analysis or postprocessing program. The revised graphics library then was ported back to the CRAY-2 system, and linked to other LLNL programs.

This research project is one example of academia and industry helping each other. Throughout the first year of its partnership with NCSA, FMC has enjoyed substantial benefits from its exposure to the technological resources and technical knowledge available at NCSA. FMC also has purchased a second Stellar workstation, which has been installed at the company's NCSA offices to develop visualization tools based on the Application Visualization System software available through the center. Hence, another opportunity exists for technology transfer, as each organization benefits from the software developed by the other.


***About the authors***

*Ilhan Dilber is an industrial consultant at the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign. He also teaches training courses at the center, including program optimization for Cray Research computer systems, and supports applications software for use in solid mechanics and other engineering areas. He received a B.S. degree in mechanical engineering from Bosphorus University, Istanbul, and M.S. and Ph.D. degrees in mechanical engineering from Northwestern University, Evanston, Illinois.*

*Wing Cheng is a senior member of the technical staff at FMC's Corporate Technology Center. He has extensive experience with large-scale finite element programs such as NASTRAN, MARC, NISA, and DYNA. Cheng received a B.S. degree in mechanical engineering from Washington University, St. Louis, Missouri, and an M.S. degree in engineering mechanics from Santa Clara University, California.*

# Preparation of real-time Great Lakes forecasts

*Keith W. Bedford, The Ohio State University, Columbus, Ohio*
*David J. Schwab, NOAA-Great Lakes Environmental Research Laboratory, Ann Arbor, Michigan*

Over the past 25 years, a partnership of Canadian and U.S. federal, state, and local governments in the Great Lakes region has achieved something rarely seen in regard to any other large body of water — the reversal of the effects of pollution and a consequent reinvigoration of the resource. As a result, commercial and recreational organizations are competing intensely with each other and with environmental preservation organizations for use of the lakes.

Consequently, public officials must make responsible planning decisions, including resource allocations. However, the lack of a centralized database with information about the lakes impedes these decisions and gives rise to strong and continuing pressures for accurate information. Intensifying these demands are the recent concerns about toxic substances and the need to determine the relevance of global warming scenarios. Private groups also need organized information to make optimal use of the limited resources.

To meet these needs, researchers from The Ohio State University (OSU) and the National Oceanic and Atmospheric Administration (NOAA) are using OSU's CRAY Y-MP8/864 supercomputer to develop the Great Lakes Forecasting System, which will provide data on lake circulation and thermal structure systematically and cost-efficiently. The model will provide six-hour forecasts of the physical characteristics of the lakes.

The most fundamental information required for responsible decisions involves the physical behavior of the lakes. The Great Lakes (Figure 1), particularly the shallower ones (Erie and Ontario), respond rapidly to surface heating and cooling and surface shear stress from wind storms. These fluctuations cause not only complex and sometimes dangerous surface wave con-

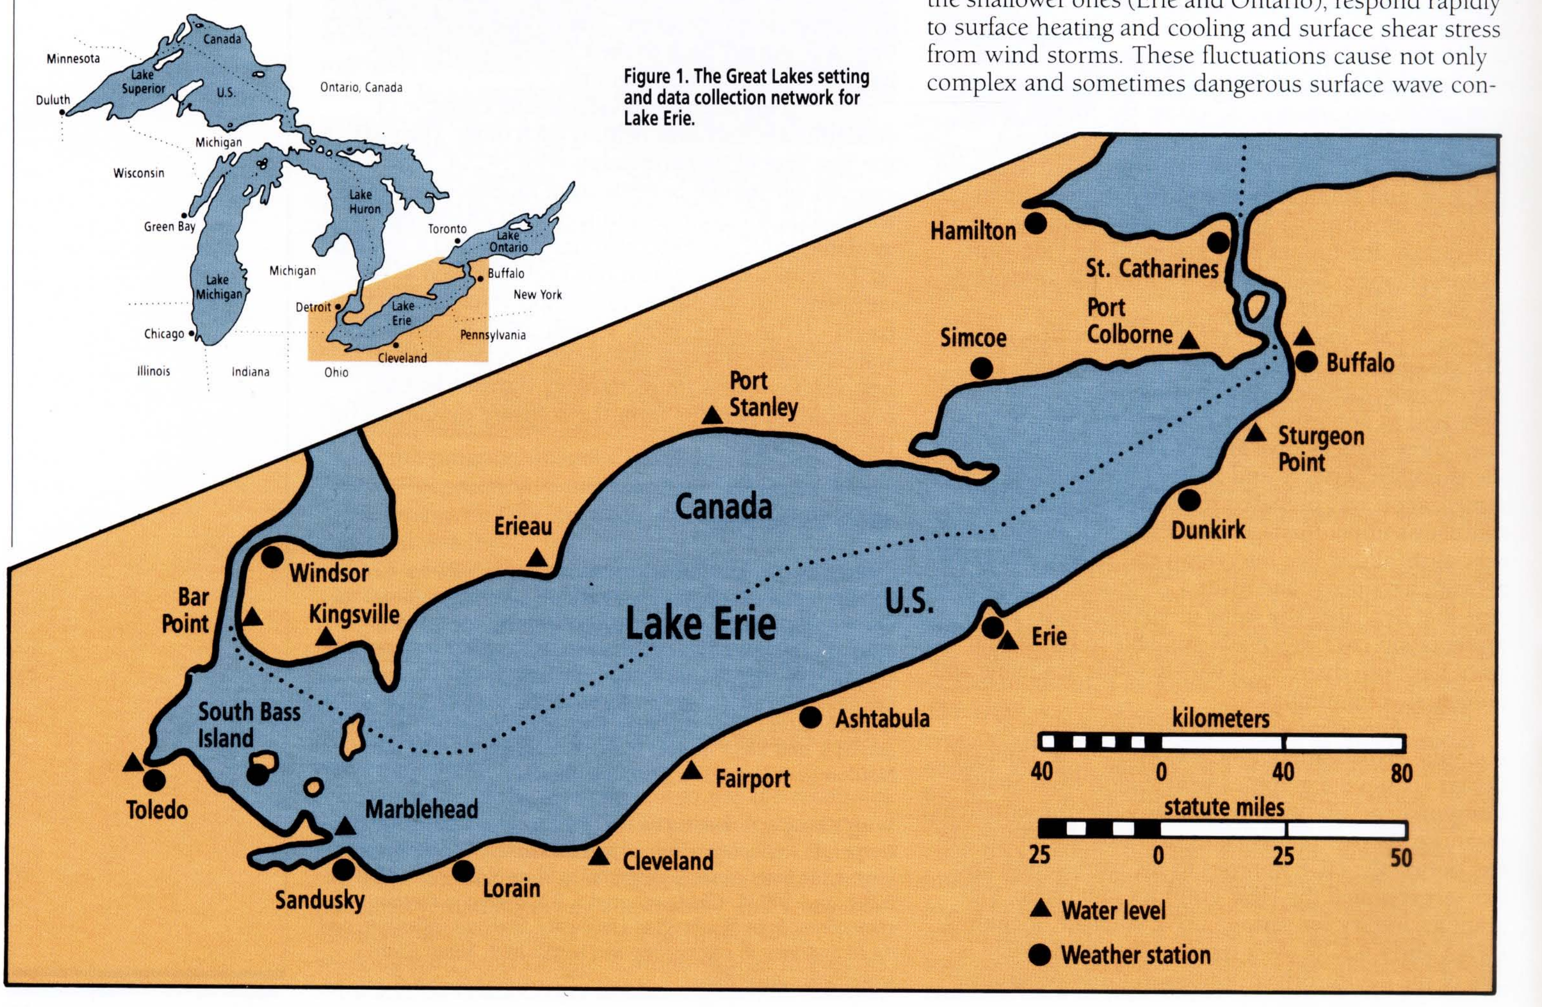


**Figure 1. The Great Lakes setting and data collection network for Lake Erie.**

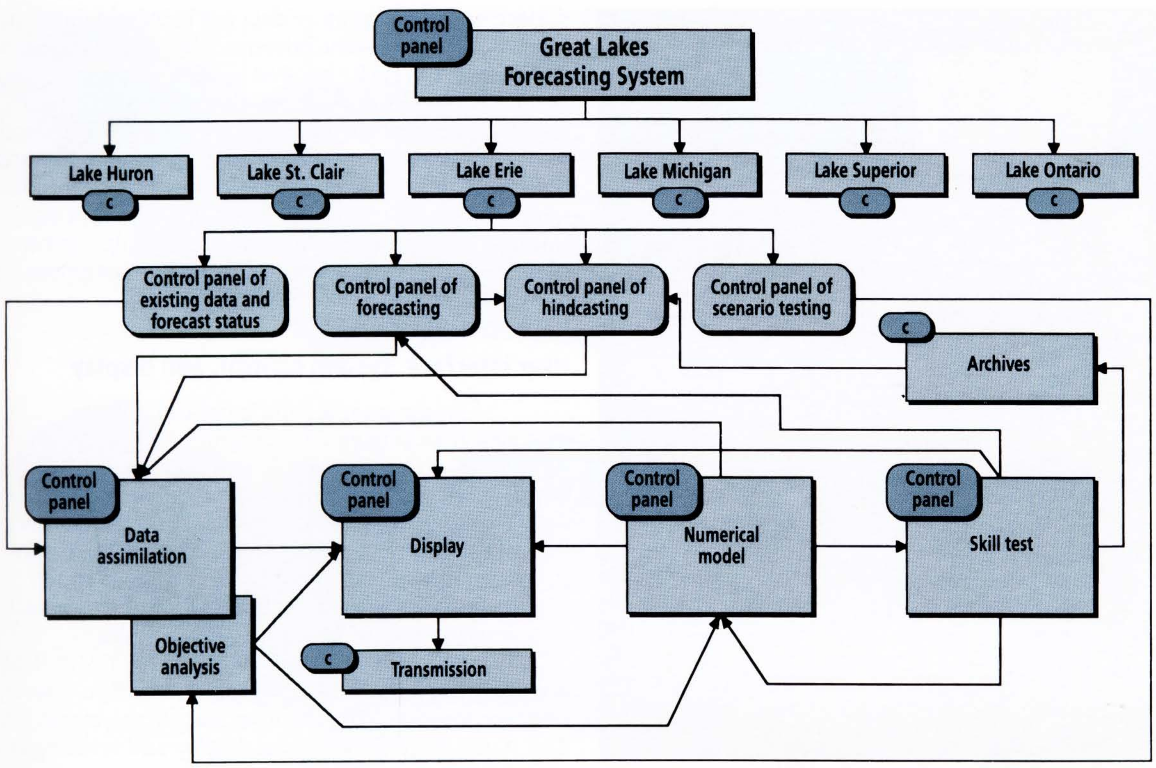


**Figure 2. The system block diagram for the Great Lakes Forecasting System.**

ditions, but also near-shore coastal jets and temperature gradients. Summer thermocline formation and internal wave activity compound this already complex picture. Clearly, the lakes exhibit no constant circulation and temperature characteristics upon which to base the required decisions.

The first stage of the forecasting system concentrates on predictions of three-dimensional current and temperature fields as well as the distribution of water levels. Additional forecasts on wind waves and shore erosion potential will be developed in the second phase. The output from this system addresses four areas of use: hazard warning and avoidance, enhancement of commercial and recreational activity, scenario testing and risk assessment, and natural resource preservation.

The system prototype is being developed for Lake Erie, due to its shallowness and intermittent, often extreme response to atmospheric events. This model, called the Lake Erie Information System (LEIFS), is believed to be the first such forecast system for large bodies of water. The complexity of the model required the power of the Ohio Supercomputer Center's CRAY Y-MP8 computer system. Earlier test versions were studied on the CRAY X-MP system previously at the center, but it was unable to provide the desired performance and forecast resolution, which are possible with the CRAY Y-MP system.

## Forecasting system structure

Because the components required to make the predictions already were available, LEIFS was developed to link all the components efficiently, to perform and evaluate the predictions, and to disseminate the information to the user groups. The system (Figure 2) comprises four control modules (or panels) for data assimilation and analysis, forecasting runs, hindcasting runs, and scenario testing; and comprises six functional modules for data assimilation and objective analysis, numerical modeling, evaluation (skill testing) of model results, three-dimensional stereo-graphics display and tracking of model predictions, data transmission, and data archiving.

The heart of the modeling module is a three-dimensional model based on a turbulence closure scheme[1] and adapted for use in closed basins such as the Great Lakes by the NOAA Great Lakes Environmental Research Laboratory. Higher-order Bayesian filtering procedures[2] are being implemented in the model to improve aspects of the advection and turbulence schemes. Real-time forecasts are to be made with a 0.5-km square horizontal grid covering 23 vertical intervals in the depth normalized vertical coordinate. Two- and five-kilometer grids (containing 178,000 and 29,160 grid points respectively) with 15 vertical intervals have been used for testing purposes. The barotropic/baroclinic mode-splitting solution procedure results in an outer time step of 20 minutes and an inner time step of 25 seconds for the 2-km grid. A 10-day calculation with the 2-km horizontal grid with 15 vertical increments consumes 14 minutes of CPU time on the CRAY Y-MP system.

When used in the hindcast mode, the evaluation module provides a rational system for comparing a time series, or maps, of predicted variables to the observed field data. Skill tests derived from nonparametric statistics[3] form the basis of the comparison. While tunable coefficients in the models are few, the skill test score is used as a basis for coefficient tuning. If the operator chooses, a calculation can be stored for future design calculations; however, the archive module is configured to store all the data necessary to recreate the calculation rather than the calculated results. These data include the measured input data, the tuned model coefficients, and the skill test scores.

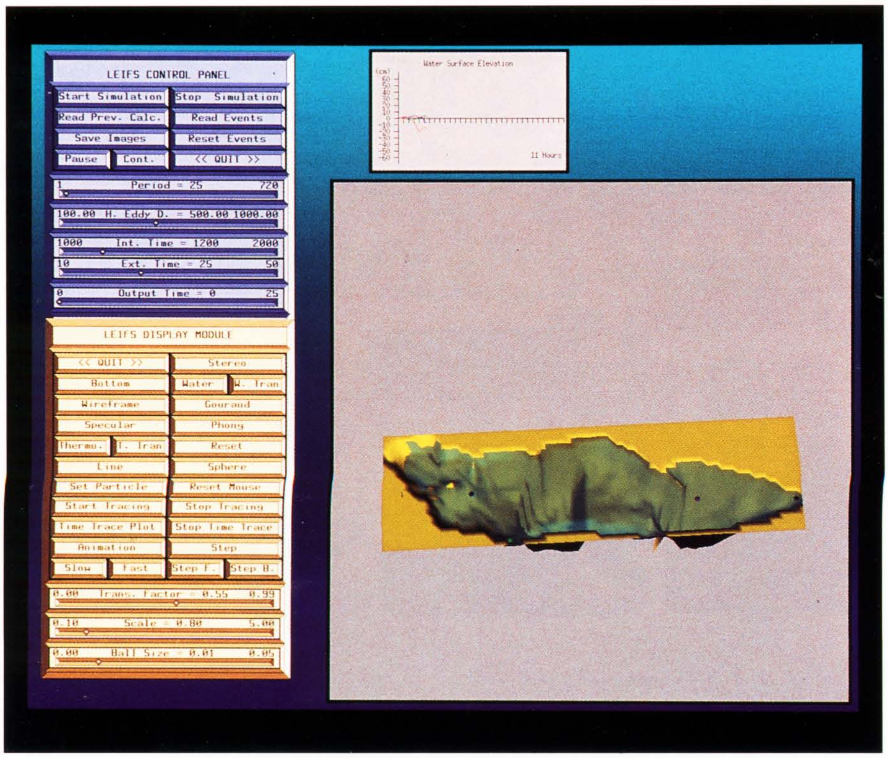


**Figure 3. Example LEIFS Display and Control Module Panels. The water level is rendered in real time as the calculation proceeds and the balls are placed by the user at points where time series of data are to be collected and displayed at the top of the graphics screen.**

In this way, the three-dimensional circulation and temperature fields can be recalculated instead of being stored on expensive on-line storage media or bulky, slow, tape storage.

## Data requirements and assimilation

At first glance, the problem of obtaining the required input data might seem intractable; however, due to governmental concern about pollution effects in the Great Lakes, an extensive base of operationally collected data now exists. Therefore, the problem, with one exception, is again one of configuring existing components. The principal data required for model initialization are water levels, tributary inflows, wind speed and direction, air temperature, and surface heat flux and temperature distributions. Water levels are measured every 15 minutes at a network of 13 gauges around Lake Erie maintained by the NOAA-National Ocean Service and the Canadian Hydrographic Service. The tributary flows are measured hourly by the United States Geological Survey. Four National Weather Service stations surround the lake (at Detroit, Toledo, Cleveland, and Buffalo; see Figure 1). Also, an automated meterological station is operated by the National Weather Service on South Bass Island, Lake Erie. The data are collected operationally and transmitted to centralized databases.

The surface temperature and heat flux estimates will require some additional work, but the conceptual processes are complete.[4] The heat fluxes are estimated by the procedure developed by Michael McCormick of the NOAA-Great Lakes Environmental Research Laboratory. The conceptual model is based on standard light transmission and scattering models with the operational input data consisting of the wind and air temperature data, collected as described above, plus estimates of the water surface temperature. The surface water temperature data are being obtained in two ways. First, real-time Advanced Very High Resolution Radiometer (AVHRR) infrared images, mapped to the surface temperature with approximately 1-km$^2$ resolution, are being obtained several times per week. Second, to provide in-situ data, we have identified over 40 water intakes around the lake that are owned by private companies or municipalities who regularly monitor water intake temperature. We now are configuring a subset of 20 of these deep-water intakes into an automated temperature-data collection network.

## User interface, system control, and display

As the number and complexity of the modules grow (Figure 2) it becomes necessary to provide an extremely reliable and understandable interface between the various components. Therefore, we have integrated interactive visual displays at every step of the control interface.[5,6] For this reason, a considerable amount of development is being directed toward the display module (Figure 2) and its interaction with the control functions and modules.

The interface runs on the CRAY Y-MP8/864 system and is accessed through Sun-3 and Sun-4 series color graphics workstations from Sun Microsystems, Inc. The stereo three-dimensional display and tracking is done with a Stellar GS2000 graphics supercomputer workstation. Full use is made of the UNIX and UNICOS system capabilities, including full use of version 2 of the X Window System. Future versions will provide an icon window interface.

Graphics are embedded in the interface at three levels: first, as a mechanism for interacting with the control panels; second, as a means of tracking the calculations during execution; and third, as a display and analysis system for intermediate and final calculations and for display of derived results. Three types of controls are implemented on the control panel for each module (Figure 3): one shot "buttons" to activate or deactivate procedures, a file browser to select files for the various procedures, and "dials" or "sliders" to set the values of various coefficients in the modules.

Tracking and displaying functions are handled in the display modules, the heart of which is a series of two- and three-dimensional graphics procedures composed as a toolbox.[6,7] This toolbox system, the *animation production environment*, is being developed by the Ohio Supercomputer Graphics Project of the Ohio Supercomputer Center. The complex physics in Lake Erie requires three-dimensional rendering of features; therefore, algorithms for smooth surface shading, transparency rendering, volumetric rendering (Figure 4), stereo viewing, particle tracing, and vector display are selectable from the control panel.

We envision a transmission and distribution system that provides prediction maps in three forms, tailored to the needs of major users. For recreational fishers and boaters, the charts would be obtainable at local marinas and other locations by a fax system. We also plan to test a system for delivering charts to boaters via satellite or radio-telephone. Finally, NOAA is developing a regional coastal information network for the Great Lakes that could be used to transmit the forecasts to the various regulatory and planning agencies.

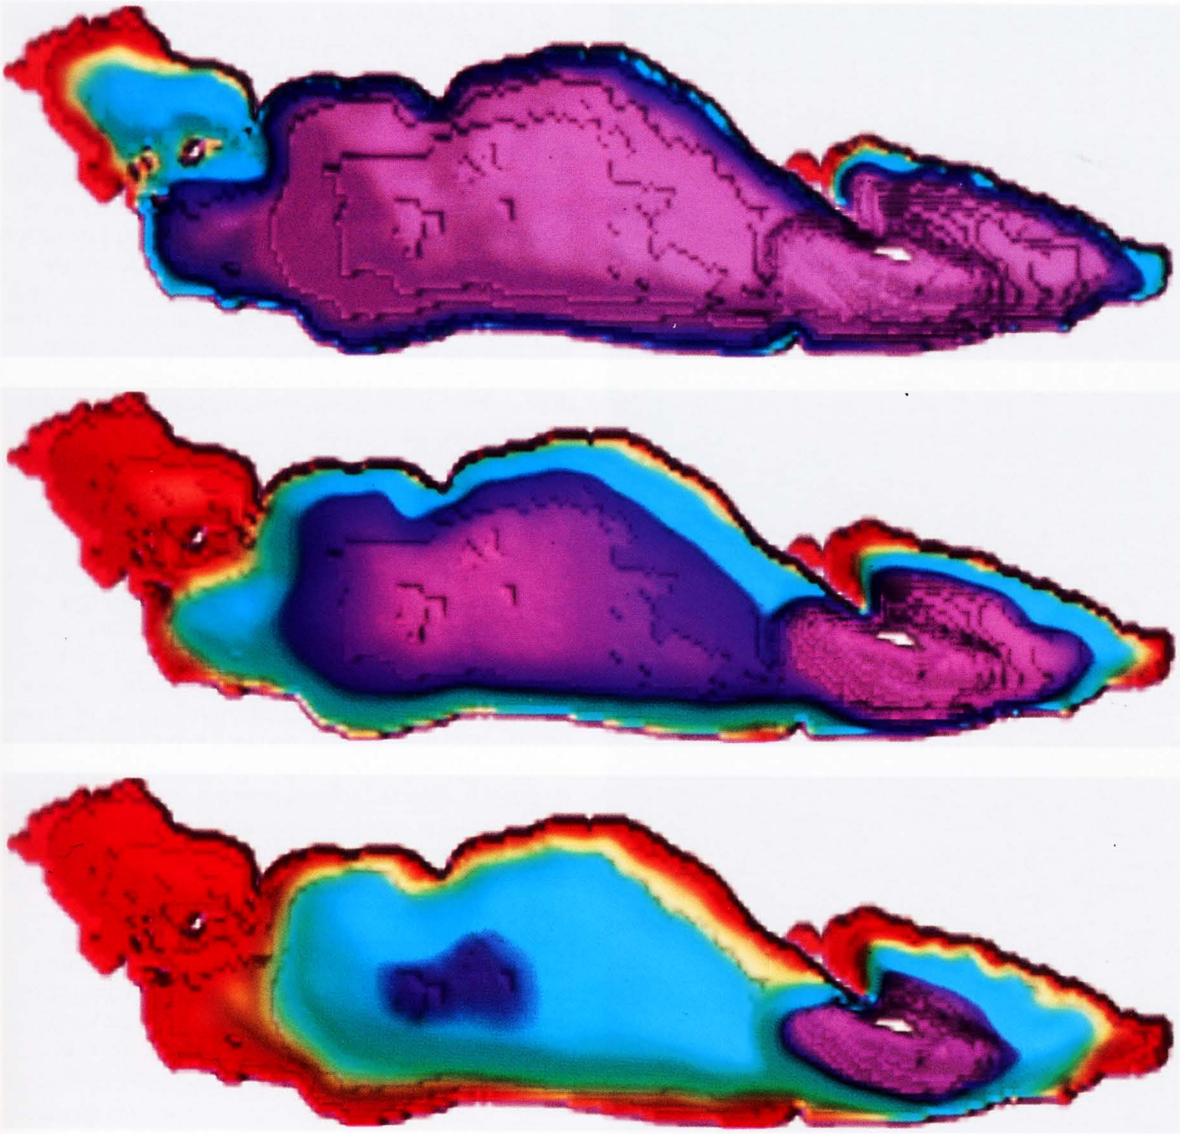

**Figure 4. A series of volumetric renderings of the Lake Erie water temperature from early to late spring (top to bottom). Red represents the warmest water, while purple represents the coldest water (4° C). Opaqueness is correlated with the density gradient; therefore, the top and center images show the bottom contours, while the bottom image shows a thermocline in early development.**


### *Acknowledgments*

*The authors thank the research team and sponsors for their support. Other project personnel include Carolyn Merry of The Ohio State University and Michael McCormick of NOAA-GLERL; R. Marshall and J. Kempf of the Ohio Supercomputer Graphics Project; and research assistants C. C. Yen, D. Welsh, G. Lockwood, C. F. Kuan, P. Chu, and D. Podber, all Ph.D. students in the Coastal Engineering Program at The Ohio State University. Financial support for this project includes grants from Cray Research, Inc., NASA (-973), NOAA-Ohio Sea Grant, and the OSU Office of Research and Graduate Studies. Support from the Ohio Supercomputer Center, in the form of generous allotments of CRAY Y-MP8 system time, is also appreciated, as is the continuing support of C. Bender and S. Dyer of the Ohio Supercomputer Center and R. Stoddard and T. Sweeney of The Ohio State University.*


### *About the Authors*


*Keith W. Bedford and David J. Schwab are project directors of the Great Lakes Forecasting System project. Bedford is a professor of civil engineering and atmospheric sciences at The Ohio State University. He received a Ph.D. degree from Cornell University in 1974. David Schwab is an oceanographer at the NOAA-Great Lakes Environmental Research Laboratory, specializing in physical limnology. He received a Ph.D. degree in 1981 from the University of Michigan.*


### *References*


1. Oey, L. Y., G. Mellor, and R. Hires, *Journal of Physical Oceanography,* Parts I-III, Vol. 15, pp. 1676-1720, 1985.
2. Bedford, K., J. Dingman, and W. Yeo, *Three-Dimensional Models of Marine and Estuarine Dynamics,* ed. J. Nihoul and W. Jamart, Elsevier Publishing Company, pp. 113-126, 1987.
3. Dingman, J. and K. Bedford, *Journal of Hydraulic Engineering,* Vol. 112, No. HY2, pp. 124-141, 1986.
4. Bedford, K., C. Merry, and D. Schwab, 1989 Fall Conference, proceedings, American Society of Photogrammetry and Remote Sensing, Falls Church, Virginia, pp. 77-78, 1989.
5. Bedford, K., C. C. Yen, J. Kempf, D. Schwab, R. Marshall, and C. Kuan, *Estuarine and Coastal Modeling,* ed. M. Spalding, American Society of Civil Engineers, New York, in press, 1990.
6. Marshall, R., J. Kempf, S. Dyer, and C. C. Yen, 1990 Symposium on Interactive 3D Graphics, Association for Computing Machinery, Snowbird, Utah, March 25-28, 1990.
7. Kempf, J., R. Marshall, and C. C. Yen, *Proceedings of the International Society for Optical Engineering,* Santa Clara, California, February 14-16, 1990.

# Computational design of pharmacologic agents for cancer and AIDS therapy

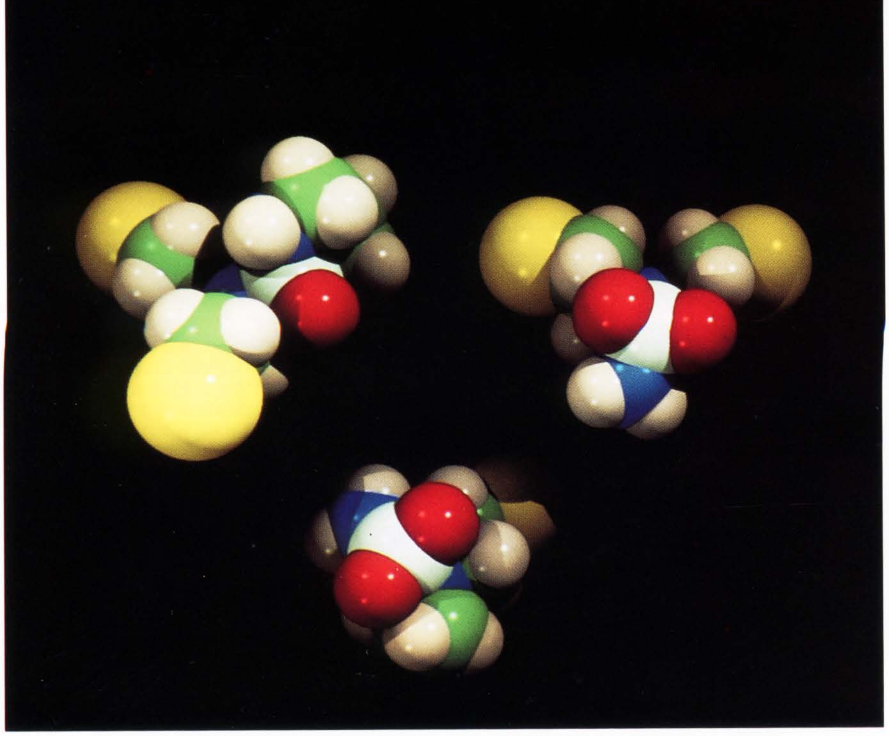

Ab initio quantum mechanical (RHF 3-21G*) optimized geometries for cyclophosphamide (upper left), phosphoramide mustard (upper right), and phosphoramide aziridine (bottom).

*Frederick H. Hausheer*
*Cancer Therapy and Research Center, San Antonio, Texas and the University of Texas, San Antonio*

*U. Chandra Singh*
*Scripps Clinic and Research Foundation, La Jolla, California*

The treatment of cancer has advanced significantly during the past four decades as a result of dedicated clinical and experimental efforts. Many of the therapeutic gains related to more effective medicines are the result of technological advancements and an enhanced understanding of the fundamental chemical and biological interactions in the pathogenesis and pathophysiology of cancer. Several types of malignancy now can be cured and significant palliation achieved in a variety of other tumors. Unfortunately, many of the more common types of neoplasms, such as carcinomas of the lung, breast, gastrointestinal tract, and melanoma, resist therapy with available agents. Multiple drug resistance is a significant obstacle in which the cytotoxic action of pharmacologic agents is rendered ineffectual by tumor cells.[1]

During the past decade, the human immunodeficiency virus (HIV) has created a major health problem for several reasons: it is lethal, a significant latent interval exists between infection and disease manifestation, its epidemiologic patterns fluctuate, and no effective therapy exists. Important considerations in the development of new therapeutic agents for AIDS and neoplastic disorders are the immediate and long term clinical toxicities that frequently affect the patient's quality of life. The untoward effects of these agents are monitored and managed expectantly, because some toxicities can be lethal to the patient.[1] Our primary research is directed toward generating new classes of pharmacologic agents possessing highly specific cytotoxicity for neoplastic cells and cells that have incorporated the HIV genome. The major goals of developing such agents will be to cure these diseases, which currently resist therapies that are minimally toxic to the patient. The work described here is ongoing in conjunction with Sterling Drug, Inc. Supercomputers used during the course of this work include Cray Research systems at the National Cancer Institute, Scripps Research Institute, and the University of Texas.

## Current approaches

The discovery that carcinogenic and viral alterations in DNA play a central role in the pathogenesis and pathophysiology of malignancy and AIDS has been a significant accomplishment.[1] Evidence suggests that critical molecular targets in neoplastic cells and cells with the HIV genome differ distinctly from normal host cells. The partial redundancy of DNA includes a regulatory molecule that could be selectively inactivated by pharmacologic agents. Other molecular targets for therapy include various functional proteins and polysaccharides that are accessible to inactivation by pharmacophores and monoclonal antibodies. The most important marker of cytotoxicity is the loss of the cell's capacity to undergo replication, and many independent cellular processes are being identified that will lead to cell death when subjected to pharmacologic agents. Researchers identify and characterize molecular targets such as DNA and cellular proteins increasingly through molecular biology, NMR, pharmacology, cytogenetics, and x-ray crystallography.

Biomedical research is entering an exciting new era with the advent of advanced molecular computational methods and supercomputers. The problems of complex molecular processes can be resolved by the combined efforts of medicinal chemists, NMR spectroscopists, molecular biologists, bio-organic chemists, physical chemists, and x-ray crystallographers. The multidisciplinary efforts of these teams should be tightly coupled to computational analysis to predict the molecular interactions of interest, and in turn guide and complement experimental results. Supercomputers are critical to this approach because rapid feedback and design considerations must guide and complement experimental efforts to maximize productivity. The operating environment of these machines must be adaptable to many algorithms and permit frequent program enhancements to simulate specific problems of varying magnitude.

New anticancer agents traditionally have been developed largely by trial and error, more recently incorporating in vitro preclinical screening using a large panel of various human malignant cell lines.[1] If an agent significantly inhibits the growth of certain types of cell lines, it is tested in vivo (murine) against the same cell lines to establish its ability to selectively inhibit tumor growth. This level of testing is followed by toxicological studies in larger animals, using tumor models exhibiting well-defined characteristics. If a drug demonstrates an acceptable profile at this stage, production and formulation are carried out, and final consideration is given to human clinical trials. This approach has several limitations, including incomplete correlation between screening and clinical activity. Despite such limitations, the in vitro screening approach has identified several clinically useful agents and probably will continue to do so. The projected rate of screening at the National Cancer Institute is estimated at 10,000 to 15,000 compounds per year, which corresponds to one potentially useful anticancer agent being identified every two to four years. The development of a clinically useful anticancer drug currently costs between $125 million and $150 million. About 40 anticancer agents currently are approved for clinical use, and perhaps 10 of these have an established curative role.[1]

## Computational approaches

We believe that the speed and capacity of the developing generation of supercomputers will help to establish a tightly coupled theoretical and experimental effort that will supersede traditional approaches to drug discovery and development. Supercomputers and numerical experimentation can provide rapid feedback between various laboratory experimental disciplines and thereby help to better focus research efforts. Coupled theoretical/experimental efforts are beginning to provide qualitative and quantitative insights into molecular interactions of interest, instead of experimentation within a "black box" realm. The results of such computations are used in most instances to test ongoing hypotheses and to optimize experimental designs by reducing (focusing) the number of possible considerations to the most relevant ones. Use of these computational approaches will depend increasingly on supercomputers to provide both rapid feedback to experimental research teams and an environment for the rapid development and enhancement of molecular computational algorithms.[2]

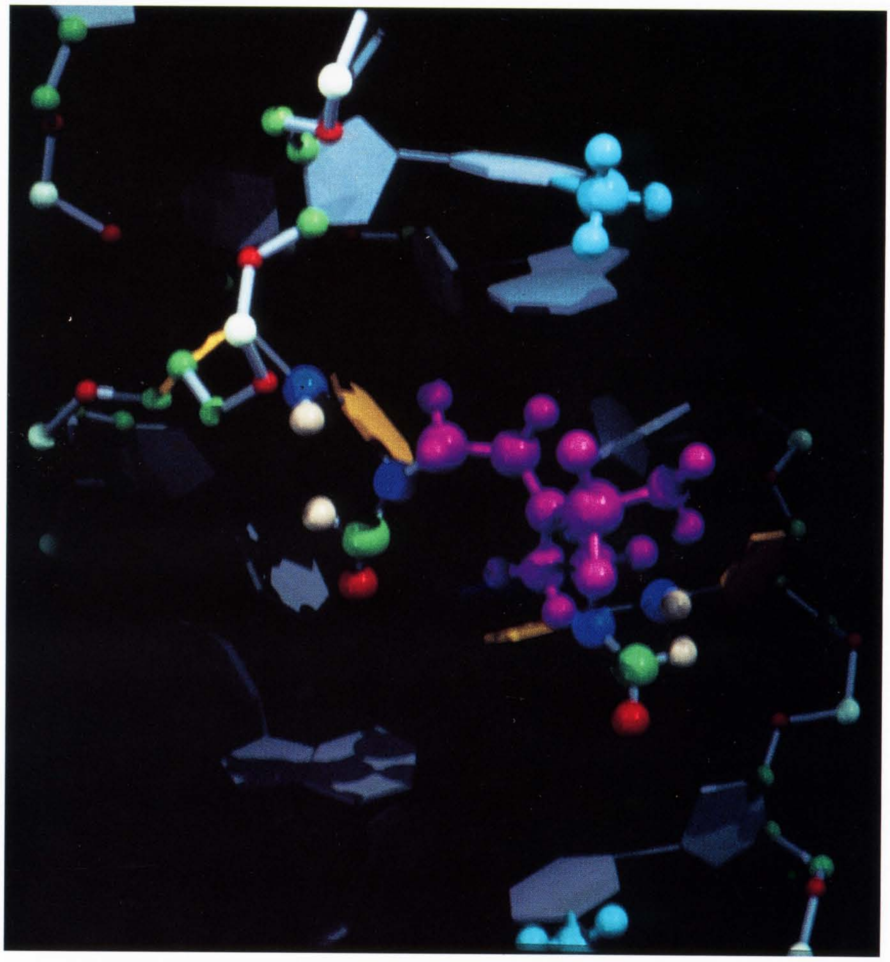

(Above) An interstrand DNA crosslink with phosphoramide mustard (the active metabolite of cyclophosphamide). This ray-traced image represents the time-averaged structure of the drug-DNA crosslink during 50 picoseconds of molecular dynamics at 300 K in water.

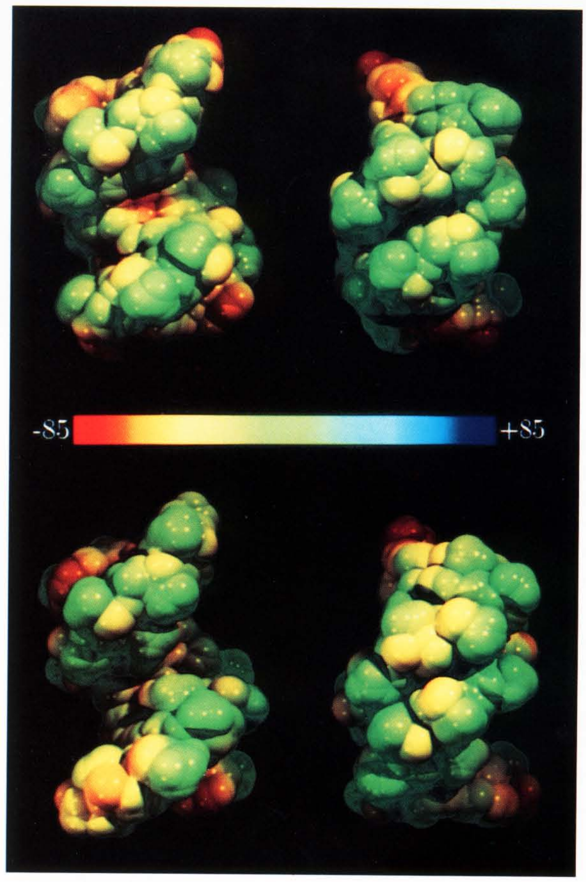


(Right) Electrostatic potential surface of neutral DNA heteropolymers. These ray-traced images show the electrostatic potential surface of two DNA heteropolymers containing 10 base pairs. This approach is useful in identifying sequence-specific regions in DNA for nucleophilic reactions with a variety of anticancer drugs. The top molecules represent the major and minor groove views of the surface of a guanine-cytosine heteroduplex, and the bottom molecules are of an adenine-thymine heteroduplex. The charges fitted to the molecular electrostatic potential were calculated using ab initio quantum mechanical methods (6-31G*), and the molecular dynamics simulations of the DNA duplexes were carried out in water (approximately 9200 atoms per simulation) at 300 K for 50 picoseconds.

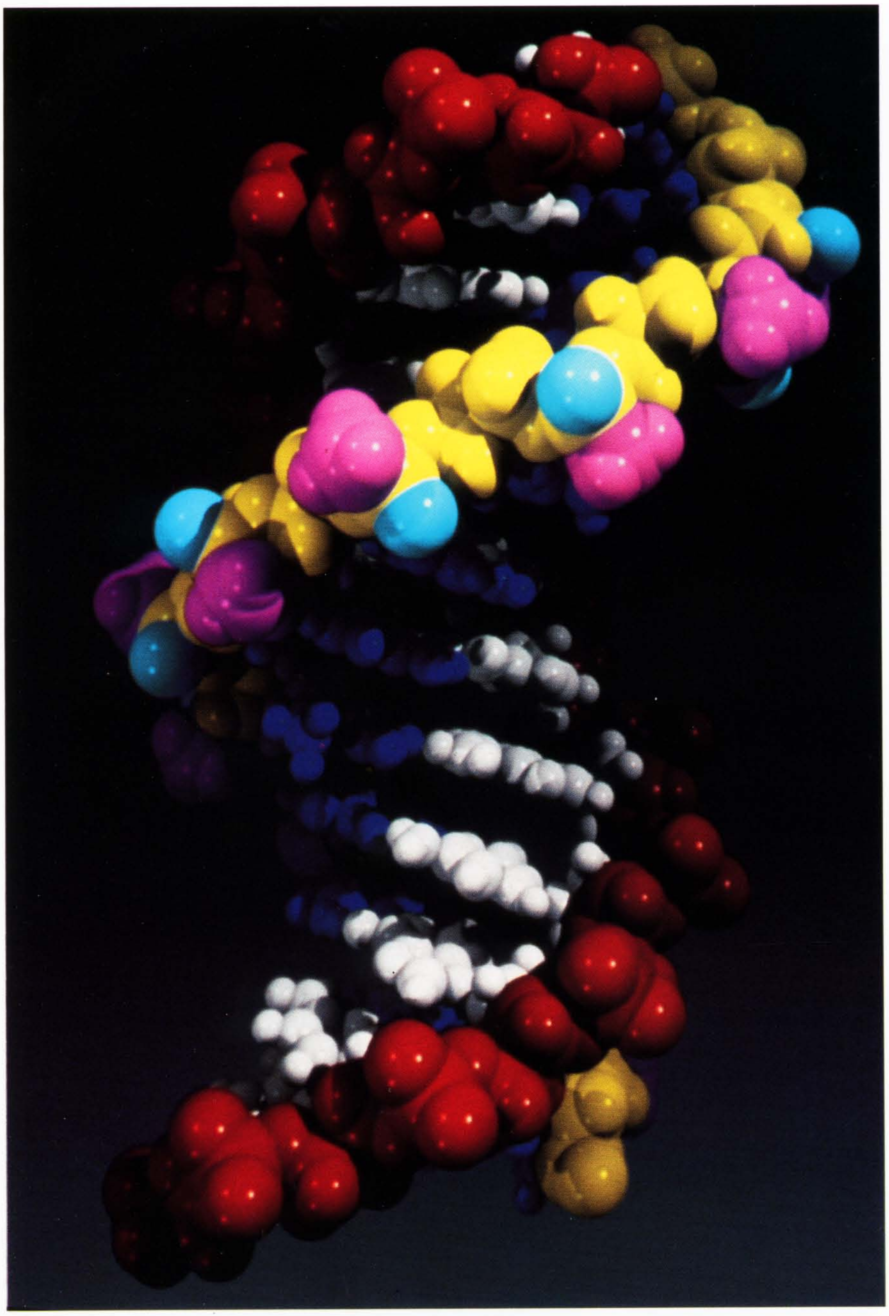

Methylphosphonate substituted antisense hybridization to an AIDS (HIV) DNA target. The antisense sequence (yellow, magenta, and blue) is strongly bound to the target strand (shown in white and blue) coding for a viral protein. This is a ray-traced image of a molecular dynamics simulation with counterions and water at 300 K (water removed for clarity).

## Applications

One application of large-scale computation in pharmacological research involves the use of ab initio methods to study chemical, conformational, electronic, and energetic properties of various clinically active drug molecules that are difficult to study experimentally. Cyclophosphamide, one of the most widely used and effective anticancer and immunosuppressive agents, provides an example. This drug has a broad activity profile, including solid tumors, leukemias, and lymphomas. Synthesis of activated metabolites of this drug, 4-hydroperoxy-cyclophosphamide, for example, has facilitated ex vivo purging of autologous bone marrow for transplantation. With the supercomputer and a hierarchy of molecular computational methods coupled to experimental efforts, the local chemical determinants of cytotoxicity for this class of agent have been identified for the first time.[3] Certain metabolites of cyclophosphamide interact with DNA and proteins of cells producing covalent crosslinks.

DNA is the major target of this class of drugs, which mediates cytotoxic events, and most laboratory approaches in this area dictate the use of highly complex experimental methods, which in most instances can characterize only partially the interactions at each step. Ab initio quantum mechanical calculations have provided structural, stereochemical, and chemical data for active metabolites of the parent compound and have prompted several previously unconsidered experimental efforts that are yielding greater insight into the chemical properties of the drug. Using molecular mechanics and molecular dynamics, we have simulated the conformational and energetic interactions of six crosslink complex ensemble trajectories in water at 300 K each for 50 picoseconds.[3] These calculations helped to identify several previously unrecognized molecular interactions that could affect the antitumor activity of this drug significantly, and have suggested new leads for the design of agents that bind efficiently to specific regions of DNA. In the past, numerous attempts to describe specific structural/stereochemical relationships have eluded formidable efforts by NMR or x-ray crystallography for several molecular species in this class of antitumor agents. For example, ab initio calculations of several electronic states of one metabolite (phosphoramide mustard) led us to design and use several previously unconsidered experimental approaches, such as solid-state NMR, and have provided new structural, chemical, and stereochemical data that are being used to design more effective agents. Computer simulations carried out over the past year and a half provide new evidence that agents in this class consistently bind DNA selectively. Such binding is determined largely by specific factors, such as nucleotide sequence, conformation of the drug and DNA helix, electrostatic attraction and repulsion, and electron interactions involved in covalent adducts. During this time we have used these data to design a series of multidisciplinary experimental approaches to delineate the magnitude of the effects indicated by these calculations.

Efforts to design molecules rationally that bind specifically to DNA regions that regulate gene expression, or to the mRNA product of such genes, are aimed ultimately at the development of effective disease-specific therapies with minimal toxicity to the patient.[4] Such efforts involve the study of specific chemical alterations in the base, sugar, and phosphate groups of DNA-like molecules that can be coupled with reactive fragments of certain anticancer groups. In this way, certain desirable chemical and biological properties are retained in a "homing" molecule that inactivates specific cellular targets unique to tumor or AIDS-infected cells. Such properties include regional DNA-mRNA sequence binding with a high degree of specificity, capacity to assume a similar conformation with native DNA and mRNA, and resistance to cleavage by cellular enzymes. When these molecules are near their targets, reversible binding is favored initially, followed by irreversible covalent modification, which inactivates DNA and mRNA targets. These fragments could be targeted to oncogenes or regions of DNA that code for drug resistance. The use of similar agents is likely to benefit AIDS patients because the

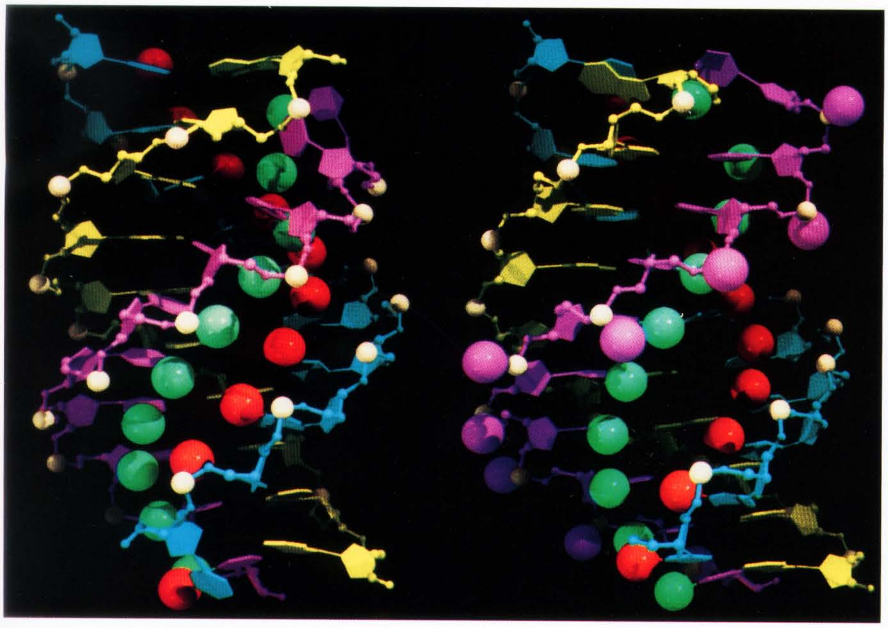

pathogenesis of the disease involves direct alterations in the genome of host cells and is accompanied by abnormal protein production.[4]

Chemical modification of the phosphate backbone is a challenging experimental problem in this area. Chemical modifications by sulfur, methyl, or ethyl groups, for example, of the phosphorus atom result in two diastereomers. It is important to identify whether one conformation can enhance or inhibit the binding of the parent molecule to the target. When a molecule contains 15 phosphorus atoms, $2^{15}$ (32,768) possible diastereomeric conformations are possible. To reduce the number of experimental considerations, we are using a supercomputer and advanced molecular computational methods to identify these relationships for a given molecular sequence. Computational methods have identified a significant chemical interaction between the backbone and methyl groups on the modified base, which could decrease the binding affinity.[5,6,7] Several specific binding agents have demonstrated cytotoxicity in vitro and have supported the feasibility of developing this as a new class of therapeutic agents. Numerical simulation by supercomputer has enabled us to overcome several experimental obstacles and has identified molecular interactions of potential significance.

**Triple helix formation with native (left) and methylphosphonate-substituted (right) oligodeoxynucleotides. The methylphosphonate antisense third strand (magenta in the right image) is binding more tightly to the target DNA duplex because of decreased electrostatic repulsion and hydrophobic effects (from the methyl groups). The methyl groups (shown as spheres) of thymine (red on the first strand, green on the third strand) are more closely associated in the native triple helix and are pulled apart with the methylphosphonate substitution. Both molecular dynamics simulations were carried out with explicit water and counterions (the average number of atoms per simulation was 10,400) with periodic boundary conditions for 50 picoseconds at 300 K.**

## Combined efforts

Advancement in molecular computational methods coupled with the power and flexible environment of general-purpose supercomputers provides a new approach to the development of more effective pharmacologic therapy for malignancy and AIDS. By integrating this modern approach with the strengths of a coupled multidisciplinary experimental effort, it will become possible to understand the pathogenesis and pathophysiology of these diseases and to develop new agents of greater clinical efficacy to inactivate certain specific molecular targets. Although this field is still in its infancy, the dramatic progress in computational algorithms and hardware will reduce these problems commensurate with that of new experimental approaches. The integration of theoretical and experimental efforts will be essential to identify and resolve scientific problems expeditiously in the future. The major result of these developments will be a dramatic reduction in mortality, morbidity, and lost time and resources. Future developments in this field will affect not only cancer and AIDS, but all areas of medicine involving pharmacologic management of disease.


### *Acknowledgments*

*The authors thank Tom Palmer of Cray Research, Inc., for computer graphics assistance and Cray Research, Inc., and Sterling Drug, Inc., for grant support.*


### *About the authors*


*Fred Hausheer received his M.D. degree from the University of Missouri-Columbia School of Medicine in 1982. He is board certified in medical oncology following completion of postdoctoral fellowship training at the Johns Hopkins Oncology Center and Hospital. He has held staff appointments at the Johns Hopkins Oncology Center and at the Advanced Scientific Computing Laboratory at the National Cancer Institute. Currently, he is involved in clinical development of new therapeutic agents for patients with cancer in the Drug Development Section of Medical Oncology at the Cancer Therapy and Research Foundation of South Texas and the University of Texas in San Antonio. His research interests include anticancer and anti-AIDS drug design by the application of advanced molecular computational methods and molecular pharmacology of anticancer and other therapeutic medicinal agents.*

*U. Chandra Singh received his Ph.D. degree from Madurai University, Madurai, India. He is an associate member in the Molecular Biology Department at Scripps Research Institute. His research interests include rational drug design by the application of advanced molecular computational methods and supercomputers, and he is widely recognized for his expertise and contributions in scientific programming developments in molecular mechanics, molecular dynamics, free energy perturbation, and quantum mechanical programs.*


### *References*


1. DeVita, V. T., S. Hellman, and S. A. Rosenberg, eds., *Cancer: The Principles and Practice of Oncology,* Saunders, 1985.
2. Singh, U. C., F. K. Brown, P. Bash, and P. A. Kollman, "An Approach to the Application of Free Energy Perturbation Methods Using Molecular Dynamics," *Journal of the American Chemical Society,* Vol. 109, pp. 1607-1614, 1987.
3. Hausheer, F. H., U. C. Singh, and O. M. Colvin, "Identifiation of Local Determinants of DNA Interstrand Crosslink Formation by Cyclophosphamide Metabolites," *Anti-Cancer Drug Design,* Vol. 4, pp. 281-294, 1989.
4. Miller, P. and P. O. P. T'so, "A New Approach to Chemotherapy Based on Molecular and Nucleic Acid Chemistry: MATAGEN," *Anti-Cancer Drug Design,* Vol. 2, pp. 117-128, 1987.
5. Palmer, T. C., F. H. Hausheer, and J. D. Saxe, "Applications of Ray Tracing in Molecular Graphics," *Journal of Molecular Graphics,* Vol. 7, pp. 160-164, 1989.
6. Hausheer, F. H., U. C. Singh, J. D. Saxe, O. M. Colvin, and P. O. P. T'so, "Enhanced Binding of Methylphosphonate Modified Oligodeoxynucleotides in Triple-Stranded Poly [dT] [dA] [dT] Demonstrated by Molecular Dynamics," *Proceedings of the American Association of Cancer Research,* Vol. 30, p. 629, 1989.
7. Hausheer, F. H., U. C. Singh, J. D. Saxe, O. M. Colvin, P. O. P. T'so, "Can Oligonucleotide Methyphosphonates Form a Stable Triplet with a DNA Duplex?" *Anti-Cancer Drug Design* (in press), 1990.

# Visualizing program performance

*Bob Swanson, Cray Research, Inc.*

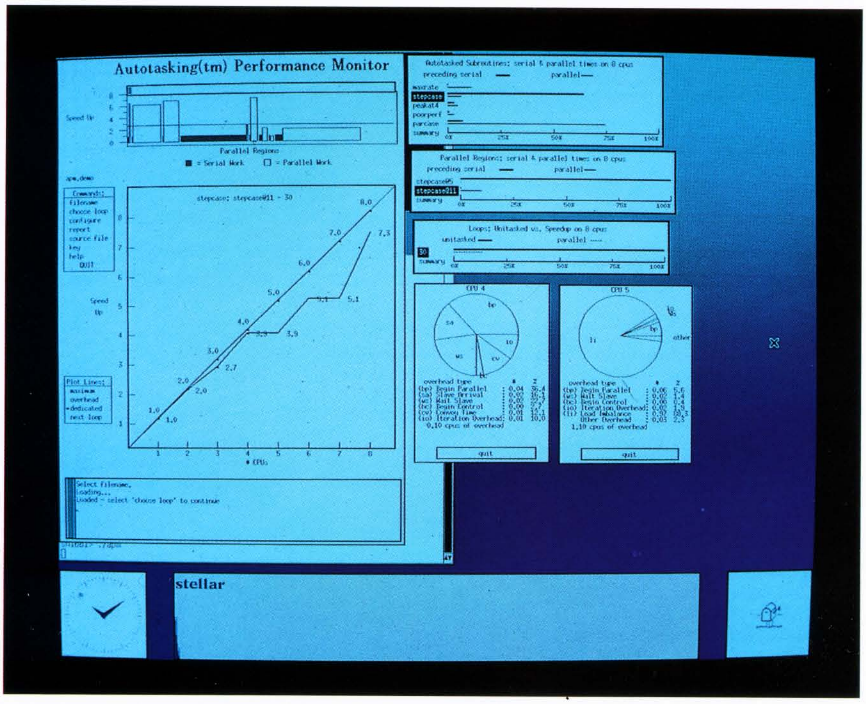


**Figure 1. Typical apm screen setup. Users can view several aspects of a program's performance simultaneously through multiple windows.**

Time. Many people would say it is today's most precious commodity. Users of Cray Research supercomputers want their programs to save them time by providing solutions as quickly as possible, but they don't want to spend much time improving program performance.

However, programmers are finding that even small optimization efforts pay off. An excellent example is provided by the Perfect Benchmarks, an industry-wide set of application codes that accurately measures supercomputer performance.[1] These benchmarks demonstrate performance for baseline and optimized versions of 13 application programs. When optimized, most programs performed two to four times faster. In one example, programmers tripled the performance of a code by investing only one hour in optimization. These performance improvements were made possible with performance tools supported by Cray Research's UNICOS operating system. The tools enable programmers to identify portions of their programs where minimal improvements can reap the highest performance paybacks. After all, spending hours optimizing the performance of a function that accounts for only 1 percent of total program time would be wasted effort.

In recent years, a number of performance analysis tools have become available for use with Cray Research supercomputers, including Flowtrace, Perftrace, SPY/prof, and PERFMON/hpm. However, a major difficulty in using these older tools is the inflexibility of their data displays. In most cases, a different view of performance data either was unavailable, or required recompilation or re-execution of the user's program.

The availability of the X Window System on the UNICOS operating system has presented Cray Research with an excellent opportunity to help programmers visualize performance information. In addition to offering graphical displays of the performance data, the X Window System can provide a noticeably improved user interface to the tools' features and func-

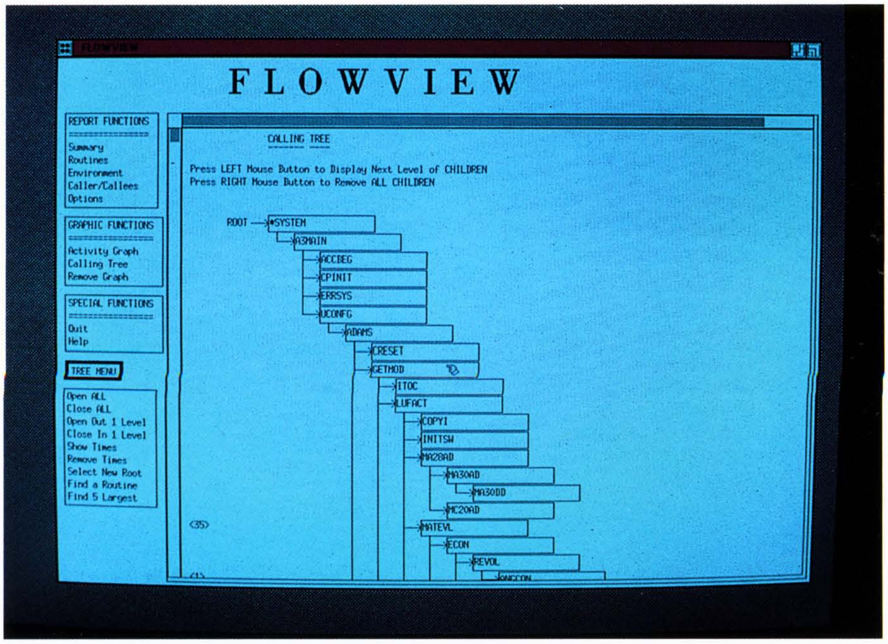


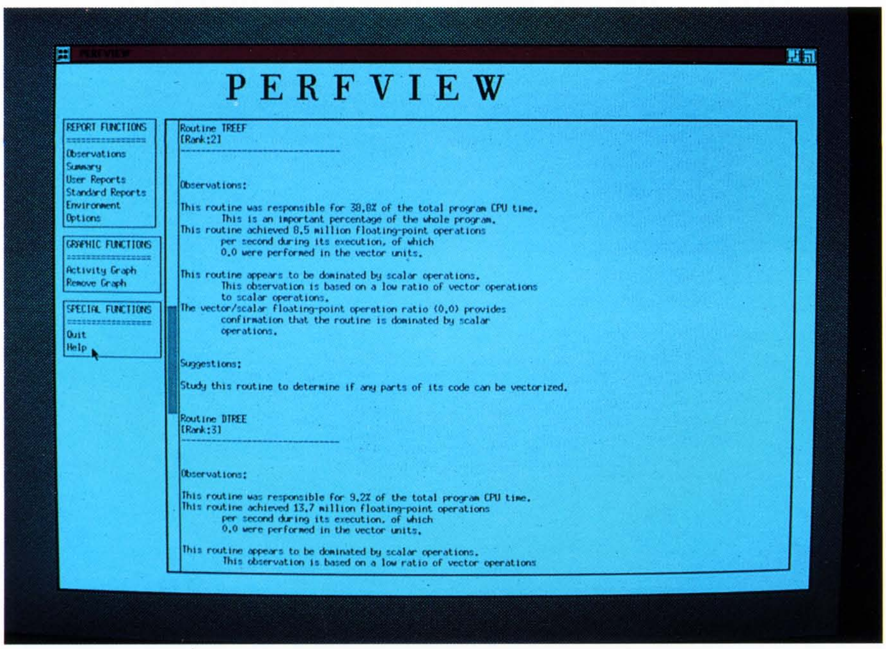


**Figure 3 (top). A flowview screen, showing its graphical calling tree display. Users can navigate through this display, exploring the dynamic structure of the program.**

**Figure 4 (bottom). A perfview screen, showing observations about routines within the user's program. This report presents complex HPM data with simple explanations of the program's hardware performance.**

the program is running. Users of flowview can view the performance data in a variety of ways. For instance, users of the X Window System interface can traverse a graphical calling tree of their program. This tree shows the relationships of different routines and allows study of portions of programs (sub-trees) with a simple point-and-click mouse interface. Figure 3 illustrates a flowview screen, showing its graphical calling tree display. Using the mouse buttons, users can navigate through this display, exploring the dynamic structure of the program.

## Perfview

The perfview command processes the data generated by the Perftrace library. Perftrace times user routines, using the Hardware Performance Monitor (available only on the CRAY Y-MP and CRAY X-MP systems). A number of hardware statistics, such as the MFLOPS rating of a subroutine, are gathered by this device. The perfview command can process up to four executions of the same user program, where each was run under a different hpm counter group. This tool also can process the output of the hpm command, which gathers HPM statistics for an entire program.

In addition to reporting the Perftrace data in several forms and arrangements, this tool can offer observations about the program and routine performance. These observations are based on general rules governing the expected behavior of the HPM counters for certain types of codes. For instance, the tool can determine and report whether or not a routine appears to be well-vectorized.

Because HPM data can be very complex (up to 34 statistics for each routine), The perfview command enables users to design their own reports. Users can define special statistics that use simple arithmetic to combine several HPM counters. They then can define precisely which fields will be listed on a line of report output. Figure 4 illustrates a perfview screen, showing its report of observations about routines within the user's program. This report is designed to present complex HPM data by giving simple explanations of the program's hardware performance.

## Summary

One of Cray Research's goals is to provide users with a wide variety of performance information. Each tool uses a different method of gathering performance data, varies in the type and scope of information recorded, and varies in its impact on program performance. For example, Perftrace provides very detailed, low-level information, but requires recompilation and causes significant routine CPU overhead. The prof command provides less detailed information with minimal overhead. While some users have tool preferences, we encourage users to try several, and to discover the tools that provide the most appropriate types of information for their needs.

*All performance analysis tools currently available from Cray Research are described in detail in the UNICOS Performance Utilities Reference Manual, publication SR-2040. Note, however, that the new tools described in this article will be described in documents supporting UNICOS 6.0.*


### Acknowledgments

*The author thanks Al Stipek and Charles Grassl of Cray Research for providing the material for several key sections of this article.*



### About the author

*Bob Swanson is a programmer in the Tools, Libraries, and Commands group in the Software Division of Cray Research, Inc. He has worked at Cray Research for seven years, mainly as a field analyst. For the past two years, his primary focus has been to develop and enhance performance analysis tools.*


### Reference


1. Grassl, C. M., and J. L. Schwarzmeier, "A New Measure of Supercomputer Performance: Results from the Perfect Benchmarks," *CRAY CHANNELS*, Spring 1990, pp. 18-21.

# CORPORATE REGISTER

## Cray Research gains new customers in Australia, Japan, U.K., and U.S.

**The Commonwealth Scientific and Industrial Research Organisation** (CSIRO) in Canberra, Australia, has installed a two-processor CRAY Y-MP supercomputer. CSIRO is an organization funded by the Australian government and other sources to carry out research to contribute to Australia's quest for enhanced economic performance, living conditions, environmental quality, and community understanding of science and technology.

**Daihatsu Motor Corporation** in Japan, a new customer for Cray Research, has ordered a two-processor CRAY Y-MP supercomputer system. The purchased system will be installed in the third quarter of 1990 at Daihatsu's headquarters outside Osaka and will be used for computer-aided engineering applications such as structural analysis, crash simulation, computational fluid dynamics, combustion simulation, vibration analysis, and design optimization.

**The Royal Aerospace Establishment** (RAE), a British government defense research laboratory, has installed a four-processor CRAY-2 supercomputer at its main computer facilities in Farnborough, Hants, England. The CRAY-2 supercomputer will be used by both the RAE and the Admiralty Research Establishment for computational fluid dynamics and structural analysis research.

**Rolls-Royce plc,** a leading international aero engine manufacturer, has ordered a four-processor CRAY Y-MP system to be installed in October in Derby, England. The company will use the new system for engine design processes including computational fluid dynamics, turbulance thermodynamics, and stress analysis.

**E. I. du Pont de Nemours & Company** has ordered a four-processor CRAY Y-MP supercomputer to be installed this fall at its Experimental Station in Wilmington, Delaware. The new CRAY Y-MP supercomputer is Du Pont's third Cray Research system. Du Pont uses supercomputing for design and analysis of its manufacturing facilities and processes, new product development, and environmental studies.

**The University of Nevada-Las Vegas** (UNLV) has purchased a two-processor CRAY Y-MP supercomputer to be installed this summer in the university's engineering complex. UNLV is a new customer for Cray Research. The new system will use the UNICOS operating system. UNLV intends to use the CRAY Y-MP system to create a world-class center for the study of radioactive waste management and other scientific problems of regional and national interest.

## Cray Research to enter minisupercomputer market

In March, Cray Research announced its intent to enter the minisupercomputer market by signing an agreement-in-principle to acquire Supertek Computers, Inc., a privately held Santa Clara, California, designer and manufacturer of a Cray Research-compatible minisupercomputer system. Under terms of the proposed agreement, which is subject to negotiation of a definitive agreement and a Hart-Scott-Rodino review, Supertek would become a wholly owned subsidiary of Cray Research, Inc.

"Supertek was the only firm that met our criteria for acquiring compatible technology for this market segment," said John A. Rollwagen, chairman and chief executive officer of Cray Research. "Their product is Cray compatible and is available now. We believe the minisupercomputer system will stimulate demand for Cray Research supercomputers and will allow us to meet our customers' needs for a Cray-compatible product of this kind."

Supertek's system, called the S-1, is compatible with Cray Research's CRAY X-MP supercomputer series. By the end of 1990, the company expects to port its UNICOS operating system to this technology, resulting in an interim Cray Research minisupercomputer system. Supertek also has done developmental work on a follow-on product that Cray Research expects to turn into a CRAY Y-MP-compatible system available for worldwide marketing in the second half of 1991.

## Cray Research receives Bell-Perfect Award for overall supercomputer performance

In March, Cray Research won the overall performance category of the first Bell-Perfect Award. The company was recognized for the performance of its CRAY Y-MP8/832 computer system, which achieved the fastest execution time on the 13 Perfect Benchmark codes, a collection of application programs that solve basic scientific and engineering problems. The benchmarks are administered by the Center for Supercomputing Research and Development (CSRD) at the University of Illinois at Urbana-Champaign.

Charles Grassl and Jim Schwarzmeier of Cray Research reported an execution time of 53 seconds using a CRAY Y-MP8/832 computer system — about 3500 times faster than the execution time achieved on a DEC VAX 11/780 computer.

Timings indicate that the CRAY Y-MP-8/832 system significantly outperforms Fujitsu, Hitachi, and NEC supercomputers. (See "A New Measure of Supercomputer Performance: Results from the Perfect Benchmarks," *CRAY CHANNELS,* Spring 1990.) The CRAY Y-MP system also achieved the best overall price/performance for four mathematical algorithms, and Jeff Brooks and Tom Hewitt of Cray Research will receive an award for that category. "This award demonstrates the CRAY Y-MP system's unmatched performance," said John Rollwagen, chairman and chief executive officer of Cray Research. "The CRAY Y-MP system clearly delivers the best performance of all the computer systems tested over the cross-section of application codes."

## Cray Research to host science and engineering symposium

Cray Research is holding its Fifth International Symposium for Science and Engineering on Cray Research Supercomputers, in London, England, on October 22-24, 1990. The symposium provides an international forum for discussing recent progress in advanced scientific and engineering computing. Participants will explore opportunities and requirements for new applications of advanced computing systems, under the theme "Supercomputer Challenges of the 1990s."

More than 50 speakers from research organizations and corporations throughout the world will present papers in their areas of expertise. The symposium will be held at the Royal Lancaster Hotel, Lancaster Terrace, London. For registration information, please contact the Cray Research sales office nearest you.

# APPLICATIONS UPDATE

## Call for entries: Cray Research 1990 Gigaflop Performance Awards

In 1989, Cray Research initiated the Gigaflop Performance Award Program to highlight achievements in computational science. In its first year, the program recognized 20 individuals and teams who had solved problems using Cray Research computer systems running in excess of one gigaflop — that is, over one billion (Giga-) FLoating Point Operations Per Second.

This year the program award target has been raised to 1.5 GFLOPS. If you are conducting important work and achieving performance at this level, Cray Research encourages you to submit an entry for evaluation. Entries for this year's program are being accepted through September 1, 1990. A panel of Cray Research scientists and engineers will review each entry to confirm that it meets the award citeria. Awards will be presented in New York City at a reception and dinner the week of November 12, 1990, during the IEEE/ACM Supercomputing '90 conference.

The Cray Research Gigaflop Performance Award Program is conducted to recognize scientists and engineers who are working at the leading edge of their respective disciplines while furthering the science of supercomputing. For an entry form or more information, contact Vicky Frank, Gigaflop Performance Award Program administrator, Cray Research, Inc., 1333 Northland Drive, Mendota Heights, MN, 55120.

## Grant program fuels research, technology transfer

In 1985, at about the time that the National Science Foundation was establishing supercomputer centers at several U.S. universities, Cray Research introduced its Research and Development Grant Program. The initial purpose of the program was to assist universities in the funding of Cray Research computer systems and to promote computational research in science and engineering. The program since has been expanded in funding and in scope. Cray Research donated $10 million in 1990 to universities that operate Cray Research computer systems. The program has several objectives:

- ☐ To complement Cray Research's internal R&D efforts in vector and parallel algorithms and applications development
- ☐ To educate students at all levels in the use of supercomputers
- ☐ To transfer university research results to industrial users of supercomputers
- ☐ To provide financial support to academic centers

### How the program works

Proposals from investigators undergo a peer review at Cray Research. Considerations made during the review include the quality of science and engineering involved, requirements for a supercomputer (including a plan for vector and parallel processing), and the potential usefulness of the results to industrial scientists and engineers. Awards are made annually and are renewable. Awards can take several forms, including support for graduate students and postdoctoral fellows, travel to conferences, and computer time on Cray Research systems. The program funds projects distributed across a wide spectrum of scientific and engineering disciplines, as shown in the chart.

### Technology transfer

In 1990, the program is providing funding for over 200 research projects in computational science. The results of these projects, which can include annual reports, journal publications, videotapes, algorithms, and application programs, are disseminated by Cray Research to a wide audience within the company and externally to customers and prospects. This dissemination facilitates the transfer of technology from universities to industrial users of supercomputers.

### Results

An example of the effectiveness of the program can be found within the University of Texas (UT). The university operates two Cray Research systems installed at the Center for High Performance Computing in Austin. University administrators want the supercomputers to function as resources for the entire 14-campus wide Texas university system. The R&D grant program is helping address this concern by funding projects at 13 of the 14 campuses. At the UT Health Sciences Center in San Antonio, researcher Fred Hausheer uses the Cray Research systems to design pharmacological treatments for cancer and AIDS (see his article, page 18). At the UT Tyler campus, researchers George Whitson and Arun Kulkarni have developed an artificial neural network model that runs on the Cray Research systems; the model is being used for image analysis of satellite data and for gene identification and classification. At the Austin campus, researcher Graham Carey has set up a computational fluid dynamics laboratory that makes extensive use of supercomputing. One of his students, Alfred Lorber, received a Cray Research Gigaflop Performance Award in 1989 for his work in two-dimensional flow simulation.

The grant program recently has become international, and in 1990, nine researchers at the University of Toronto received funding. Three of the funded projects involve advanced computational fluid dynamics at the Institute for Aerospace Studies. Two projects in atmospheric physics involve modeling the behavior of mesoscale storms.

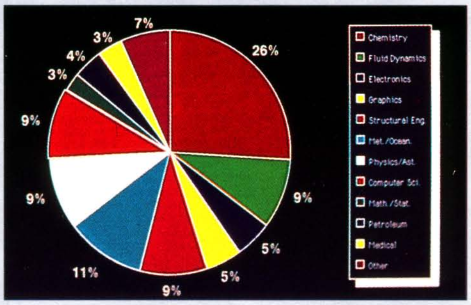


**Cray Research R&D Grant Program funding by discipline.**

Medical research receiving funding includes blood flow modeling, drug design, and radiation dosimetry simulation. Grant recipients at the university also are optimizing an image processing system to process vast amounts of radioastronomy data.

### Mutual benefit

The R&D grant program benefits not only academic researchers, but also Cray Research. It establishes communication between Cray Research personnel and hundreds of researchers working on supercomputer algorithms and applications, including students who eventually will enter industry. The grant program also allows Cray Research to establish formal collaborations between academic researchers and researchers at Cray Research and industry sites. For example, researcher Stanley Ahalt at The Ohio State University is developing an artificial neural network program called the Neural Shell, which runs on Cray Research supercomputers, and on workstations, using the X Windows graphics interface. Ahalt collaborates with Cray Research's C compiler development group to improve compiler performance and with the company's applications group to research artificial intelligence applications.

Much of the innovative software that runs on today's supercomputers originated in academic environments. Cray Research sponsorship of diverse research projects at universities demonstrates the company's continuing commitment to foster computational science and to support computational scientists. In most cases, these researchers are developing applications or processes that can be valuable to other users in academia, government, and industry, enabling Cray Research to expand the range of computational disciplines that benefit from supercomputing.

## INGRES RDBMS now available on Cray systems

The INGRES relational database management system offers an easy-to-use, time-saving method for storing, retrieving, accessing, and updating information. Engineers and scientists now can use INGRES 5.06 on Cray Research computer systems as a stand-alone data management tool and can embed structured query language (SQL) statements in Fortran and C applications to take advantage of the data organization INGRES provides. The availability of INGRES on CRAY X-MP and CRAY Y-MP systems gives Cray Research customers industry standards for data management and the potential for large-scale applications integration.

Data management has become a substantial requirement at many Cray Research sites. As INGRES becomes optimized for Cray Research systems in future releases, users increasingly will be able to apply the large memories, high I/O bandwidth, and computational performance of Cray Research systems to data management problems. This capability will save time and enhance productivity in traditional application areas as well as in emerging markets such as operational systems, chemistry, and finance.

To the user, INGRES functions on Cray Research computer systems exactly as it does on other computer systems. In addition, the INGRES/NET product will allow INGRES users on one computer system to manipulate a database residing on another system. Because INGRES/NET works in the TCP/IP environment supported by UNICOS, a Cray Research system can be both the database client and server on a network with other machines running INGRES.

The Ingres products available in release 5.06 are INGRES RDBMS, INGRES/MENU, INGRES/QUERY, INGRES/REPORTS, INGRES/FORMS, INGRES Interactive QUEL and SQL, EMBEDDED QUEL and SQL for C and Fortran, INGRES/NET for the TCP/IP protocol, and INGRES/APPLICATIONS.

Ingres Corporation is a leading supplier of quality information integration products and solutions for mission-critical applications. The company's primary product is the INGRES RDBMS, which enables organizations to quickly access data that span a variety of computer systems in a network. The INGRES family of products includes tools for developing applications and gateways for accessing data across heterogeneous computer operating environments. For more information about using INGRES 5.06 on Cray Research computer systems, contact Steve Mattos, Ingres Corporation, 1080 Marina Village Parkway, Alameda, CA, 94501; telephone: (415) 748-3686; or Sara Graffunder, Cray Research, Inc., 655-E Lone Oak Drive, Eagan, MN, 55121; telephone: (612) 683-3671.

## Motorola teams up with NCSA to improve electronic products

Electric currents traveling between components on the surface of a circuit board generate electric and magnetic fields that can interfere with each other and cause the circuit to malfunction. This is especially true in today's compact, high-frequency equipment. Motorola Inc. has found an effective way to analyze this problem through its working relationship with the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign (UIUC).

Motorola Inc. became NCSA's fourth industrial partner in January 1988. As one of seven industrial partners at NCSA, Motorola has gained extensive access to the center's CRAY X-MP/48 and CRAY-2 supercomputers. Shortly after joining NCSA, Motorola began working with the UIUC College of Engineering to develop a finite-difference program that simulates in three dimensions the propagation of electric and magnetic fields on a multilayered circuit board.

Motorola's goals in this venture were to use the speed and memory of the CRAY-2 system to solve Maxwell's equations, which describe the behavior of electric and magnetic fields, and to transform the massive amounts of solution information into an animation that engineers could access from a desktop system. Motorola's circuit engineers then would be able to see how the frequency and strength of a signal relate to the electric and magnetic fields.

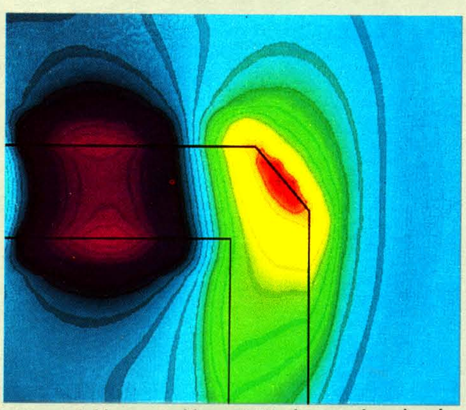

An electric field generated from a 3 GHz sine wave introduced onto the circuit board from the left side. The circuit board trace is outlined in black. The field strength ranges from red to violet showing positive field strength and negative field strength respectively.

Problems can occur when the electric field generated from one trace (microstrip or stripline transmission line) induces a signal on an adjacent trace. This is called cross talk, and, if serious enough, it can cause the circuit to malfunction. The cost of moving a product from design to production can be very high if these problems are not resolved during the design phase. By using computer simulations to model the circuit boards and their electric and magnetic fields, problem areas can be identified and corrected during design, before an expensive prototype is produced. This greatly reduces the time between design and production.

The massive amounts of memory required by this simulation make NCSA's CRAY-2 system, with 128 Mwords of static memory, a natural choice for the computations. The simulation saves the transient solution information in an ASCII output file. This file then is converted into eight-bit color raster images and saved in an NCSA HDF (Hierarchical Data Format) file using a public-domain NCSA utility. These HDF files are transmitted over the network to the engineer's local workstation. Using NCSA Image for the Macintosh and NCSA X Image for X Windows workstations, researchers interactively can animate the solution data on their desktops, changing the color palette to identify problem areas.

One study looked at the transmission coefficient of two electrical signals traveling around a 90° turn as a function of the input signal. The study showed that the pulse input breaks down as it enters the corner and radiates away from the trace. This radiation contributes to cross talk. By redesigning the bends, this radiation can be minimized, reducing the problems related to cross talk.

Through this industry-university relationship, Motorola Inc. has been developing tools that can be used to improve the quality of its electronic products and bring them to market quickly. Combining access to cutting-edge supercomputing technologies and experience in numerical simulation programs, they have found solutions to some of today's engineering problems. Similar NCSA/UIUC relationships with other corporations have produced fruitful results and can serve as models for future industry-university partnerships.

## Cray Research system carries brainstorm to reality

At The Ohio State University, medical students soon will be able to dissect a human brain on a computer screen. By slicing sections, rotating the brain, and making structures transparent, they will be able to explore brain changes associated with Alzheimer's disease, schizophrenia, AIDS-related dementia, and depression.

The power behind the screen will be a CRAY Y-MP supercomputer, whose speed enables researchers to efficiently construct three-dimensional images from magnetic resonance imaging (MRI) data. The images will help researchers visualize and quantify relationships among various brain structures in ways previously not possible.

"This venture began as a basic science project," says Michael Torello, an assistant professor in OSU's Department of Psychiatry and Psychology. "We wanted to pinpoint the relationship between brain activity measured by an electroencephalogram and the underlying anatomy of the brain."

Torello's co-investigators are Don Stredney, an assistant professor in the College of the Arts and a medical illustrator; and Wunjang Shieh, a postdoctoral researcher who developed a renderer to create three-dimensional images from MRI data.

"The project led us to the MRI facility, looking into slice data, which are two-dimensional images," says Torello. "Some of our research questions could be answered by two-dimensional MRI images, but this method was clumsy because we had to imagine a third dimension to comprehend the relationship between structures. Then, one day I realized that with the computer power available today, we easily could construct a solid image from MRI data."

MRI reveals mental illness through biological markers. However, one MRI slice may show the same pathology for many abnormalities, such as Alzheimer's disease, depression, anorexia, and schizophrenia. "Brain pathology on one slice doesn't say much, but when you can visualize that pathology as a three-dimensional entity, you can perform shape analysis and express the abnormality quantitatively," says Torello.

The team first created three-dimensional views using MRI data and a Sun Workstation. Each image was reconstructed from 55 MRI slices, 256-by-256 in resolution and 3 mm in thickness. On a Sun 3/160 Workstation each view could be calculated in 2.5 hours. "When you are rotating the brain through a 360 degree axis, showing 36 images — one every 10 degrees — that's a hefty amount of time," says Torello. By transferring the process to OSU's CRAY Y-MP8 system, the researchers pared the time to reconstruct one image to 80 seconds, which is about 113 times faster than the previous method. Optimizing the program will yield additional time savings, according to Torello.

"Using the CRAY Y-MP system enabled us to fine tune our images quite rapidly. This shortens our development time significantly and facilitates the creative process. Now we can sit by the monitor, examine images, and brainstorm much easier." He adds that a prototype of the on-screen brain dissection device should be ready by the end of this summer, and an optimized version should be ready by the end of the year. Next, Shieh will build a user interface to help students manipulate the images more easily.

Torello says that this project will revolutionize methods for teaching brain anatomy and detecting and defining brain abnormalities in the mentally ill. "Mental illness used to be thought of as an emotional disorder, but it's a brain disease. MRI can help us make quantitative statements about the abnormalities, but this supercomputer project will give us another view — another window to what the brain really looks like."

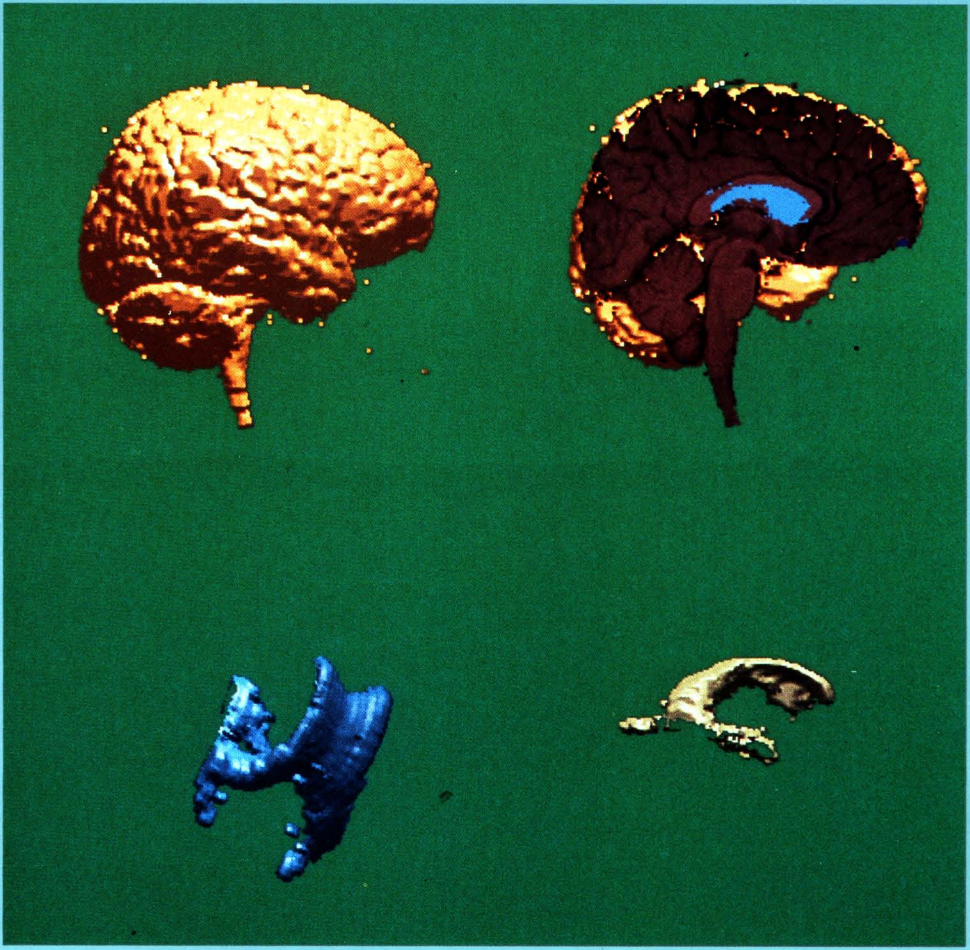

**Four three-dimensional views of brain components: (top left) brain exterior surfaces; (top right) side view of brain interior, in which cerebral spinal fluid has been color-coded; (bottom left) oblique view of extracted brain fluid compartments; and (bottom right) lateral view of brain fluid departments.**

## Cray Research system investigates bird song dialects

How do birds learn to sing? This question is at the heart of a research project being conducted by professors and Ph.D. students at the University of Toronto's Department of Zoology. With assistance from the Cray Research R&D grant program, the four-person research team is exploring how territorial birds form dialect, or shared-song, groups distributed over specific geographical neighborhoods, and how individual birds learn their song repertoires.

Beginning with field data analyzed by sonograph, the researchers designed a large theoretical model which, when run on the University of Toronto's CRAY X-MP/24 supercomputer, sheds new light on the question of how birds learn to sing. Sonographs are used to convert tape-recorded bird songs into digital information. Sonograph data were not used directly in the supercomputer model, but were used to guide the design of the computational simulation.

Birds learn to sing from each other. Among chaffinches (*Fringilla coelebs*), a small reddish-brown songbird that was the subject of the study, only males sing. Observation indicates that infant male chaffinches learn to sing (albeit poorly) from their fathers. Around age one year, a bird takes up a territory and learns very rapidly to sing an adult repertoire of songs, which he copies from a small cluster of one to eight adult male neighbor birds.

If all birds in an area sing a particular song or a particular set of songs, this is referred to as a dialect, which may differ from area to area. While these patterns have been observed in the wild, the larger questions of how dialect groups form and how birds learn their repertoires are harder to investigate, and require a more theoretical approach.

With the power of the university's Cray Research computer system, the team constructed a theoretical model of a large neighborhood and projected it through many biological generations, varying the influence of three important control parameters.

tions. Release 6.0 of the UNICOS operating system, scheduled for release by the end of 1990, will include several new commands that will provide users with an X Window System interface to performance utilities. In addition, these new tools will provide other capabilities for users who require hard copy, or who are unable to access workstations. These new tools are described in the following sections.

## The Autotasking Performance Monitor

Users of multiprocessor computer systems have found it difficult to determine how effectively their programs run in parallel. Multitasking speedups are best measured on a dedicated system, but the majority of users have access only to busy production systems.

The Autotasking Performance Monitor (APM) is a tool that can predict parallel performance speedups of a program on a dedicated system, based on statistics gathered during execution of the program on a production system. It shows where a program is spending its time, and whether those areas are executed serially or with multiple processors. APM focuses on timing each parallel region identified by Cray Research's automatic parallel processing software, called Autotasking. The tool gathers measurements of the times required to start, stop, and perform the parallel work in those regions. In addition, it measures the time required to perform the work with a single processor (unitasked time), and the serial time between parallel regions.

The overhead for APM timings is kept low (10 to 20 percent) by carefully locating the timing points to be sampled and by avoiding large numbers of subroutine calls. This strategy enables APM to make accurate projections of speedups associated with using any number of processors, up to the maximum available on the system. The use of APM requires instrumentation by the Fortran multitasking translator (FMP) phase of the Autotasking system. The user compiles, loads, and executes the program with this instrumentation. During normal program termination, the collected performance data are written to a file. This file may be processed as a standard ASCII file, but much more can be learned about the parallel performance by using the apm command to view the data.

The apm postprocessing tool uses an X Window System interface to allow graphical viewing of the performance data. This tool also provides a non-X Window System interactive and batch user interface. Figure 1 shows a typical apm screen setup. Multiple windows enable users to view simultaneously several aspects of a program's performance.

## The view tools

The three view tools planned for UNICOS 6.0, called profview, flowview, and perfview, will help provide a better way for users to study the data generated by the existing performance tools: prof, Flowtrace, Perftrace, and hpm. Each of the view tools first displays an overview of the program's performance. Additional options and reports then allow the user to focus on increasingly greater details of performance information. These tools are helpful to users at all experience levels. For the novice user, the more general overview may

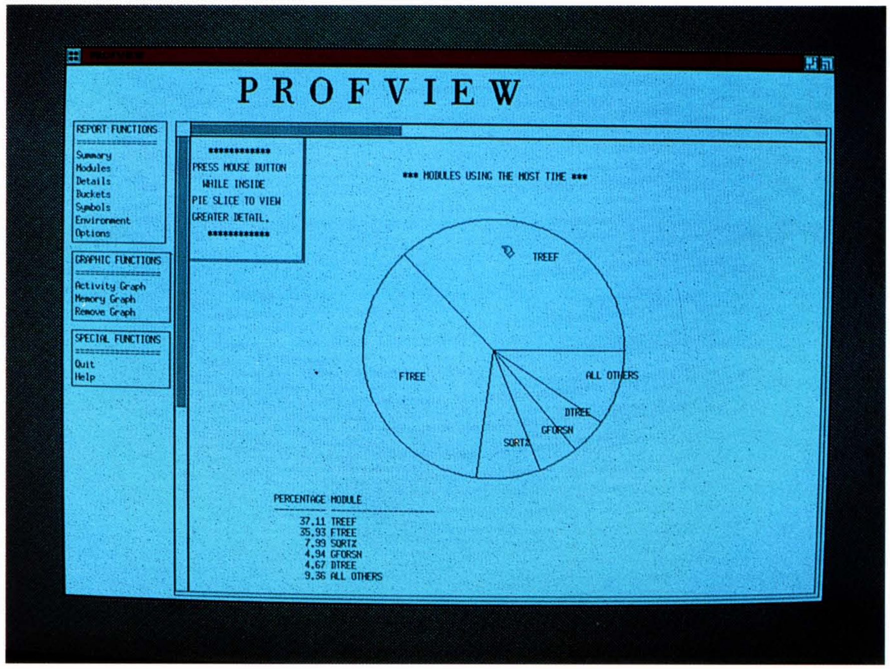


**Figure 2. A profview screen. The pie chart shows only the top five routines in the user's program, providing visual evidence of the proportion of the total time measured for each routine. By pressing the left mouse button with the cursor placed within a pie slice, users can view higher-level details for that selected routine.**

be enough to begin optimization efforts. For sophisticated users, the level of detail available is limited only by the granularity of the original raw data. Each of the new view tools is discussed in the following sections.

### Profview

The profview command processes the data generated from the profiling system (prof, and the profiling library). This system gathers performance statistics based on a regular sampling of the program's current executing instruction (P) address. Program areas that use larger amounts of CPU time show higher statistical samples ("hits"). A map of these hits, combined with the known memory locations of routines, can provide users with highly detailed views showing precisely where CPU time is spent in programs and routines.

Users of profview can view the statistical information in a variety of ways and at several levels of detail. For instance, if the debugging symbols were enabled during the original program's compilation, this tool can show performance percentages down to the level of one line or loop of Fortran code. For users of the X Window System interface, this tool also provides a memory-oriented view of performance. Users can use the mouse to point to "spikes" of high CPU activity, and then view detailed information about the routine at that memory location.

Figure 2 illustrates a profview screen, showing the first graphic display shown to a user. The pie chart shows only the top five routines in the user's program, providing visual evidence of the proportion of the total time measured for each routine. If the user presses the left mouse button with the cursor placed within a pie slice, the program will display higher level details for that selected routine.

### Flowview

The flowview command processes the data generated by the Flowtrace library. Flowtrace gathers execution timing information about user routines while

## CUG reports

*Users of Cray Research computer systems established the Cray User Group (CUG) in 1977 to provide a forum for the exchange of ideas related to Cray Research systems and their applications. The group holds two general meetings each year. Its first meeting of 1990 was held April 9-13 in Toronto, with the theme, "Network Supercomputing." Below, CUG president Mary Zosel of the Lawrence Livermore National Laboratory offers her comments on the meeting and other CUG-related business.*

The 25th meeting of the Cray User Group is now a successful part of our history. Thanks go to Edmund West and Anna Pezacki of the Ontario Centre for Large Scale Computation for coordinating and hosting the meeting. Thanks also to Christopher Lazou, ULCC, and the program committee for arranging a fine technical program with several distinguished guests, including Fred Weingarten from the Office of Technology Assessment of the U.S. Congress, John Rollwagen, CEO of Cray Research, and Professor Kahan, this year's Turing Award winner.

The Cray User Group welcomed several new sites during the past six months and now has 159 member sites, about 85 percent of Cray Research's current customer base. This growing membership provides CUG treasurer Howard Weinberger with an amusing challenge. He has been receiving notices of electronic fund transfers for membership dues that are anonymous except for a transaction number. If you think your site has paid the membership fee, and CUG does not seem to have a record of payment, you might want to investigate to see if you are one of our mystery patrons.

In response to requests from the Cray User Group, Cray Research formally established the Customer Accessible Software Problem Report database in 1988 to enable customers to access the SPRs that have been filed with the company. Sites can sign up as members and examine the SPRs of any of the other sites that have joined the system. Cray Research and the CUG Board of Directors are concerned, however, because use has been too low to justify the expense of the system. We are soliciting your participation or feedback about lack of participation.

We are working to adapt our organization as Cray Research and our user community evolve. To this end, there are a few changes in our CUG structure to announce. Michel Jaunin of EPFL will be the new Latin representative to CUG. The Mass Storage Systems SIC had its organizational meeting and judging from the number of interested people, it will be an active group. We accepted the recommendation of the chairs of the CRAY-2 MIG and the CTSS SIC that their groups be dissolved. In the process of reviewing our organization of technical committees and other special groups, we have noticed that some groups overlap while others, such as accounting, lack discussion entirely. If you are aware of a topic that seems to be missing from our agenda, please bring it to our attention.

It is reasonable to step back occasionally and ask how well CUG is achieving its goals. The organization's formal statement of purpose is to "provide an open forum to promote the free interchange of information and ideas which are of mutual interest and value to users of Cray computers, and to provide a formal communications channel between members of CUG and Cray Research." Informally, I would summarize the purpose with three "I"s: information, interaction, and influence. Cray Research is doing an excellent job of providing CUG members with up-to-date information and education about company products. We also seem to be prospering in our purpose as a forum for interactions between sites. I am concerned, however, that we are missing some opportunities to exercise organized influence on Cray Research. Some time ago, CUG had a formal User Requirements Committee. The interaction with Cray Research was less than satisfactory, so it was discontinued in favor of coordinating requirements through the individual SICs. This has not been happening in any formal manner. In fact, a question recently has arisen about what kind of customer response/review is required before a requirement is seen as a valid customer requirement from CUG. The CUG Board of Directors will address this issue internally and with Cray Research during the coming months.

Our next CUG meeting will be hosted by the University of Texas Center for High Performance Computing in Austin Texas, Oct. 1-5. The theme of the meeting will be "Benefits of Supercomputing Research," and several interesting speakers already have been proposed. I will look for you in the great state of Texas.

The project team includes Ph.D. candidates D. Bartholomew Brown and Alejandro Lynch; Professor Allan J. Baker, who also is the head of the Ornithology Department of the Royal Ontario Museum; and Professor Bruce Falls. Brown, the group's programmer, is a theoretical biologist interested in spatial population dynamics. The other three team members are field ornithologists with extensive experience in field observation.

The model is a Fortran program that simulates the interactions of a 100-by-100 matrix of territories projected in time through at least 1000 iterations, as influenced by mutation, or copying error; mortality, or the rate at which territories are vacated; and biased copying, or the tendency of birds in small clusters to copy common songs rather than copying at random.

For each model iteration, a proportion of randomly selected males was removed according to the mortality parameter and replaced by untutored birds. Copying was either exact or inexact (thus creating a new unique song type), according to the mutation parameter.

According to the model, birds learned their songs by either unbiased copying at random from one of their neighbors, or by biased copying based on the squared frequencies of song types sung by these neighbors. For example, if a bird has eight neighbors, six singing song A and two singing song B, and a mutation rate of zero is assumed, the probability of the new bird learning song A is 0.75 for the unbiased model and 0.9 for the biased model.

The unbiased copying model produced no dialect groupings. The biased model, on the other hand, showed a sharp transition between small group and large group behavior at rates of about 40 percent mortality and five percent mutation, and it indicated dialect groups of approximately 100-500 birds. All these results are consistent with field observation. "One of the most important things we learned," says Professor Baker, "is that the element of bias in dialect formation is particularly relevant."

Earlier in the study, the investigators tried to run a smaller simulation on a Silicon Graphics IRIS workstation and a Motorola 68020 coprocessor in an IBM personal computer. The model was too small, and it could not be run through enough iterations to do more than hint at the dynamics of large bird populations.

By moving the problem to the Cray Research system and running a much larger matrix through 1000 iterations of the biased copying model, the researchers found a dramatic change from small group behavior to the formation of dialect groups.

Having identified promising new directions in dialect formation in their first year, the researchers are beginning the second year of the project, which will be devoted to how individual birds learn their song repertoires.

CRAY RESEARCH, INC.
Marketing Communications
1400 Northland Drive
Mendota Heights, MN 55120
(612) 452-6650