CRAY CHANNELS

VOLUME 17, NUMBER 1 • A CRAY RESEARCH, INC. PUBLICATION

Response time was about two times faster than other systems tested.

SOLUTIONS IN PLACE

Imagine solving your most challenging computational problems in a small fraction of the time it takes today.

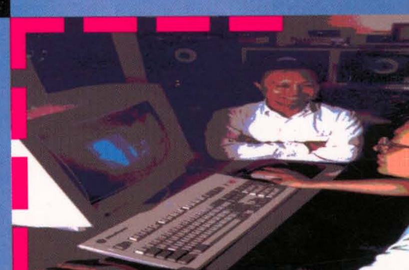

Announcing:

- CRAY T90 series
- Cray Animation Theater
- CMLogic

GET WITH THE PROGRAM. The CraySoft Fortran 90 Programming Environment sets a new standard for programming productivity. More than just an ANSI/ISO compiler, CraySoft Fortran 90 provides a sophisticated parallel programming environment for your SPARC system. With an advanced debugger, program browser, and powerful tools to measure and display parallel processing performance, CraySoft Fortran 90 makes Fortran code easier to write, understand and maintain. It's what you'd expect from the leader in Fortran technology, Cray Research. Isn't it time you got with the program? For more information: craysoft@cray.com, http://www.cray.com/or **1-800-BUY-CRAY.**

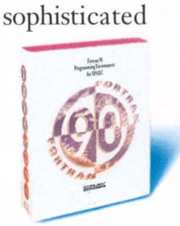

SOLUTIONS IN PLACE

C O N T E N T S

IN THIS ISSUE

In this issue of CRAY CHANNELS, we turn the spotlight on activities within Cray Research. The articles in this issue examine areas in which we have established our leadership—for example, our leap to the forefront of massively parallel processing in the year that our CRAY T3D system has been available. In addition to a discussion of the CRAY T3D system's computing accomplishments, we feature an article on electron-molecule collision simulation done with the CRAY T3D system at the Jet Propulsion Laboratory. We also feature a look at the future of Open Supercomputing, backed by our tradition of excellent software performance.

Three articles present additonal directions: one looks at our recently-formed Business Systems division, one showcases the University of Manchester's CRAY CS6400 system acquisition, and a third illuminates CraySoft's mission to deliver supercomputer software to workstations. Two articles profile divisions with excellent histories of serving customers and helping them develop noteworthy solutions: our Applications and Customer Service divisions.

Finally, we introduce our latest parallel-vector supercomputer, the CRAY T90 system, as well as the Cray Animation Theater (CAT) and CMLogic 2.0 software programs, both designed to enhance supercomputer simulation and visualization capabilities.

All in all, this issue highlights some of the high-performance computing and simulation solutions Cray Research has in place. As always, if you have any comments or questions regarding the content of this magazine, please contact us via email at channels@cray.com.

FEATURES

DEPARTMENTS

solutions

CRAY CHANNELS

CRAY CHANNELS is a publication of the Cray Research, Inc. Marketing Communications Department. Published three times per year, it is intended for users of Cray Research computer systems and others interested in the company and its products. Subscription inquiries and address changes should be sent to Department D, CRAY CHANNELS, Cray Research, Inc., 655D Lone Oak Drive, Eagan, Minnesota 55121. Other comments and queries may be emailed to channels@cray.com. For more information about Cray Research and its products, visit http://www.cray.com/ on the World Wide Web.

Volume 17, Number 1

Editorial staff
Chris Kruell, editor
Molly Sachse, associate editor
Cindy Bielke, Heidi Hagen, Ken Jopp, and Jeff Paulson, contributing editors

Design and production
Barbara Cahlander, electronic production artist
Eric Hanson, art direction and graphic design

The CRAY T90 series of supercomputers

A new chapter from the people who wrote the book on supercomputing

The CRAY T90 series is the newest, most powerful family of supercomputers from Cray Research. The successor to the world's most popular supercomputers, the CRAY C90 series, the CRAY T90 series takes supercomputing to a new level of performance. A combination of Cray Research's world-renowned supercomputing expertise and leading-edge technologies, CRAY T90 systems are the most powerful parallel-vector processing systems on the market today, capable of over 60 GFLOPS.

With the CRAY T90 series, scientists and engineers can cost-effectively simulate chemical reactions, automobile crash tests, ocean currents, financial risk scenarios, and even the delivery of babies. The predictive knowledge that digital science provides helps users design better and safer products, reduce manufacturing costs, and accelerate product design cycles.

The CRAY T932 computer system.

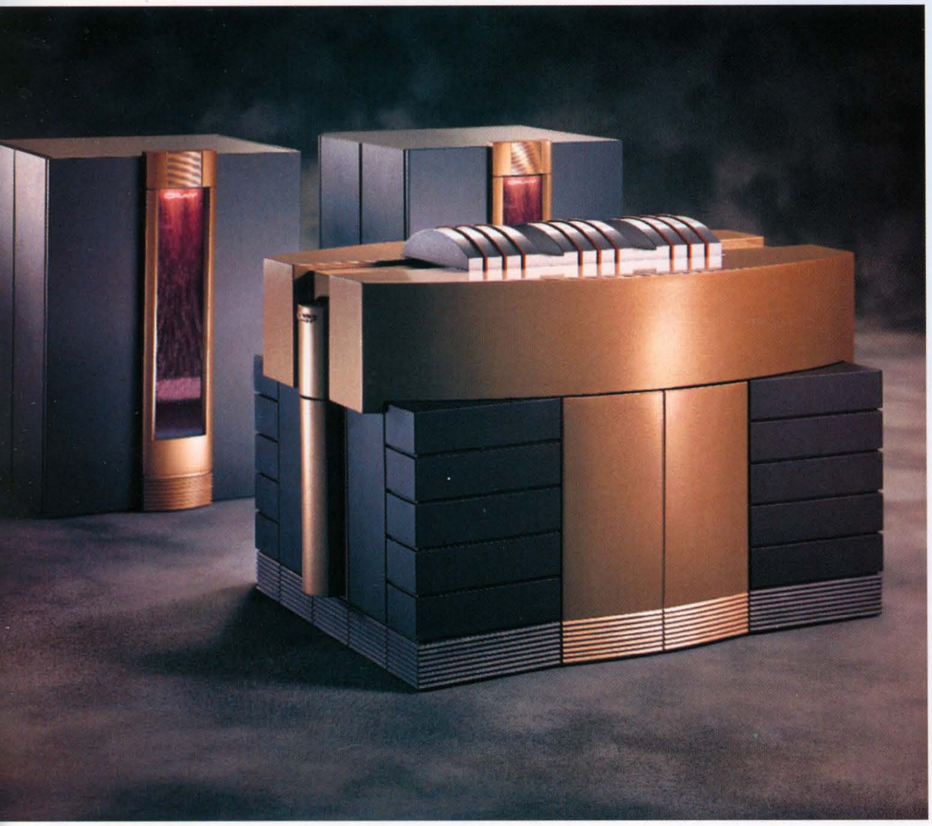

The CRAY T90 series—the most powerful general purpose computers ever built

The general purpose CRAY T90 systems excel in a wide range of applications. Their wide applicability stems not just from their parallel-vector architecture, but also the standards-compliant software and industry-endorsed application packages these systems run.

The CRAY T90 series achieves high throughput through a balanced combination of fast processors, fast memory, and fast I/O. CRAY T90 systems deliver exceptional performance through powerful hardware and software technologies, including the following:

- ☐ Scalable, easy-to-use parallelism
- ☐ Unrivaled memory and I/O bandwidth, removing data transfer limitations found on other server-based architectures
- ☐ Optional SSD solid-state storage device for increased application flexibility and throughput
- ☐ Powerful standards-based system software that reduces the complexities of the system for the user
- ☐ Application software optimized to take advantage of the high-performance system architecture

The demands of high-bandwidth computing go beyond fast processors and fast memory. A balanced simulation server requires an I/O architecture that can feed the computational capabilities of the system. The CRAY T90 series offers the fastest I/O of any computer system available.

I/O performance delivers data

The CRAY T90 series supports today's most advanced industry-standard I/O technologies, including multiple ATM, FDDI, and HIPPI connections. CRAY T90 systems use efficient I/O libraries to maximize system bandwidth. Disk drive technology support includes IPI drives, SCSI, and fiber-channel disks, offering a maximum disk capacity of 64,000 Gbytes (64 Tbytes) of storage. File system extensions include disk striping and disk caching. Large file system support and logical volume support further enhance storage performance. To maximize system performance, the CRAY T90 series also supports SSD solid-state storage devices.

In addition, many third-party tape and tape silo products are also supported through BLOCK MUX and ESCON technology.

Your applications run faster on a CRAY T90 system

The CRAY T90 series is designed to solve problems on your critical path—not just selected benchmarks—problems with hundreds of thousands of variables requiring extraordinary I/O capabilities. Problem jobs that will not run on workstations or even workstation clusters.

The Cray Research applications staff forms partnerships with customers to attack the challenges facing their organizations. Cray Research experts also work closely with application software developers to ensure that the highest level of applications performance is provided to all our customers.

Cray Research's standards-based UNICOS system software transparently delivers scalable application performance. Our industry-leading compilers automatically parallelize, vectorize, and scalar-optimize standard application codes, ensuring that you get the best possible performance from your CRAY T90 system.

We now also offer 64-bit IEEE floating-point compatibility. IEEE floating point enhances file compatibility with workstations and makes it easier to share data in a networked environment.

Reliability moves to a higher level

With its new technologies, the CRAY T90 series improves upon the reliability of past systems. Most of the new CRAY T90 technology minimizes the number of mechanical connections, virtually eliminating the need for wires. The following reliability improvements have also been incorporated:

- ☐ Extra power supplies. Each of the CRAY T90 system voltage busses includes a backup power supply. This lets the system keep running, even if a power supply fails.
- ☐ Tape automated bonding (TAB). Chips are mounted on printed circuit boards using TAB technology, eliminating chip packaging and minimizing connections.
- ☐ Offline processor configuration. A software configuration utility lets you configure a processor offline for repair.
- ☐ Zero-insertion force connections. You can install and remove processor and memory modules easily because of our advanced electrically activated connectors.
- ☐ Hardware error correction. Spare chips in the memory stacks allow memory chip failures to be resolved at the discretion of system users, because hardware error correction prevents system failures if memory chips fail.
- ☐ Single-board functional modules. This technology reduces the number of interconnects, reduces the weight of the machine, and reduces the number of piece parts.
- ☐ Software availability and resiliency capabilities including the following:
  - Checkpoint/restart
  - Interactive session protection
  - Reliable disk storage including RAID technology
  - Alternate disk paths
  - Disk mirroring
  - Guaranteed network file transfers
  - Automatic power on/off

We make sure our computers thrive in your environment

We call it Open Supercomputing—a computing environment built on industry standards to provide familiarity and ensure interoperability across the multivendor systems on your network.

Support for the following networking and distributed computing standards allows the CRAY T90 series to thrive in your existing environment:

- ☐ Communication protocols: TCP/IP and ISO/OSI
- ☐ Distributed client/server computing: ONC and OSF/DCE
- ☐ Distributed data: NFS and OSF/DFS
- ☐ Distributed batch processing: NQS and NQX
- ☐ Distributed parallel programming: PVM parallel virtual machine message passing
- ☐ Media connections: Ethernet, HIPPI, FDDI, and ATM
- ☐ Windowing and visualization: X Window System, MOTIF, Distributed GL

To provide transparent access to distributed data across local and wide area networks, Cray Research offers sophisticated, standards-based data management and hierarchical storage management solutions. From the most sophisticated network of peripherals to networks of PCs and Macintoshes—we talk to them all.

Scalable configurations offer the best fit

CRAY T90 systems come in a wide variety of configurations, from the low-cost single-processor CRAY T94 system to the top-of-the-line CRAY T932 system. Each system offers excellent scalability, which allows your computer resources to grow efficiently with increasing demand.

From zero to 60 billion in one second

The CRAY T94 system offers up to four processors at entry-level pricing and performance of up to almost 8 GFLOPS. The CRAY T916 system offers twice the performance of a CRAY C916 system (over 30 GFLOPS), our previous top-end system, at a fraction of the cost. The CRAY T932 system is capable of over 60 GFLOPS—about four times the performance of a CRAY C916 system. In short, the CRAY T90 series is the most powerful general purpose family of computers on the market today.

The CRAY T94 computer system.

System Configuration

Model	Maximum processors	Central memory (Mbytes)	Approximate peak performance	I/O bandwidth	Cooling
CRAY T94	4	512 to 1024	8 GFLOPS	>8 Gbytes/s	Air or liquid
CRAY T916	16	2048 to 4096	>30 GFLOPS	>17 Gbytes/s	Liquid
CRAY T932	32	2048 to 8192	>60 GFLOPS	>35 Gbytes/s	Liquid

Year one of the CRAY T3D system

An applications progress report

By year-end 1994, more than 25 customers worldwide from the government, university, and commercial sectors had ordered the CRAY T3D massively parallel processing (MPP) system. Phillips Petroleum Company, Chrysler Motor Corporation, the National Center for Atmospheric Research, the European Centre for Medium-Range Weather Forecasts, the Minnesota Supercomputer Center, and the Ohio Supercomputer Center are among the first customers. The average system size acquired is 128 processing elements (PEs).

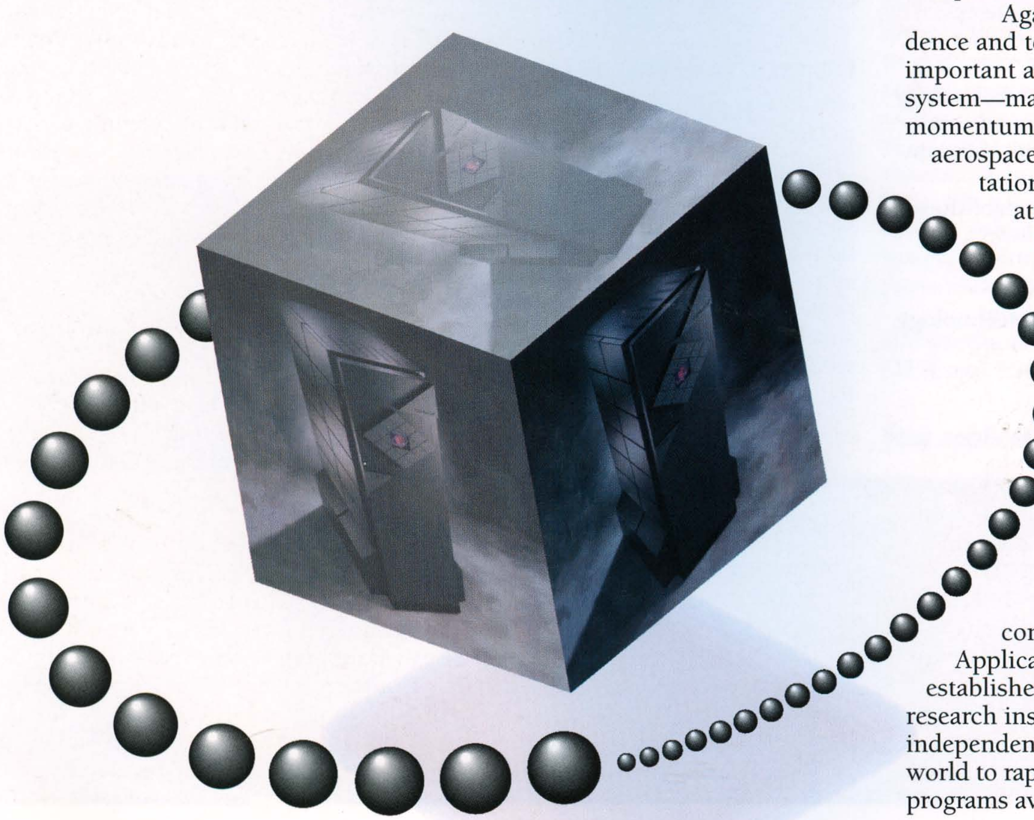

Against the backdrop of customer confidence and technical stability, the initiative to make important applications available on the CRAY T3D system—making it a real, usable tool—is gaining momentum. Chemistry and chemical engineering, aerospace and automotive (including computational fluid dynamics), petroleum, atmospheric modeling, and electromagnetics applications are among the early beneficiaries of the speed and parallelization possible with the CRAY T3D system. This article highlights some of the applications development progress to date. To read a white paper entitled "Applications Update on the CRAY T3D system," access http://www.cray.com/PUBLIC/product-info/mpp/CRAY_T3D.html with your favorite World-Wide Web browser.

Much of the applications work comes under the umbrella of our Parallel Applications Technology Program (PATP), established in 1993 to bring together leading research institutions, industrial firms, and independent software vendors from around the world to rapidly expand the number of applications programs available for the CRAY T3D system in the commercial and industrial arenas. PATP includes the Pittsburgh Supercomputing Center (PSC); NASA's Jet Propulsion Laboratory and the California Institute of Technology; the École Polytechnique Fédérale de Lausanne (Swiss Federal Institute of Technology); and a far-reaching technology transfer agreement struck last June between Cray Research and the U.S. Department of Energy's Los Alamos and Lawrence Livermore National Laboratories. The DOE/Cray Research agreement, under the auspices of the DOE's High Performance Parallel Processor (HPPP) program, includes 15 Cooperative Research and Development Agreements (CRADAs) which involve major industrial firms as well as dynamic small businesses.

Chemistry and chemical engineering

Determining how a drug binds to an enzyme, optimizing the production of a processing plant, developing a new catalyst, and improving the manufacturing processes for silicon devices are among the real-life challenges facing today's biotechnology, chemical, materials, petrochemical, and pharmaceutical companies. Development work on the most advanced algorithms and methods needed to solve these problems occurs largely in universities and government research laboratories. Following are snapshots of several computational chemistry and chemical engineering projects under way on the CRAY T3D system.

Against the backdrop of customer confidence and technical stability, the initiative to make important applications available on the CRAY T3D system is gaining momentum

CAR-PARRINELLO algorithm (quantum molecular dynamics)

Professor Roberto Car and his colleagues at IRRMA (Institut Romand de Recherche Numérique en Physique des Matériaux), an institute for computational materials physics resulting from a joint effort of EPFL and the Universities of Fribourg, Geneva, Lausanne, and Neuchâtel, have developed a new approach to quantum molecular dynamics (QMD) calculations that promises to open new classes of problems to computational study. The new approach, the O(N) formulation, incurs a much smaller computational cost per atom than does the standard CAR-PARRINELLO QMD algorithm.

With the O(N) algorithm, the necessary computations scale linearly with the number of atoms being modeled, in contrast to the cubic scaling of computations required by the standard algorithm. Because of the smaller scaling factor, quantum calculations now can be applied to atomic systems a factor of 10 or more larger than those that can be modeled with the standard algorithm. Instead of being limited to a few hundred atoms, researchers now can model several thousand atoms as well as chemical activity occurring over longer timescales, up to several tens of picoseconds. The new formulation presents these advantages because it relies on a simplified physical description of atomic systems and because the algorithm is highly parallelizable. Car and his associates currently are porting the new algorithm to the CRAY T3D system; a simplified version is already running on the system.

"The new algorithm already is as computationally efficient on the CRAY T3D system as is the standard algorithm, in terms of taking advantage of the computer's hardware features," said Car. "The high-speed communication between processors in the CRAY T3D system has been an important feature in achieving the efficient parallelization of the code. The machine also is very stable and reliable, and now a number of basic mathematical tools, such as the basic linear algebra subroutines, or BLAS, are available to programmers."

By making practical the modeling of several thousand atoms, the O(N) algorithm presents many new prospective QMD applications, including biochemical modeling. Such applications might include the study of membranes and of macromolecules in solution.

Number of processors	Time/step (seconds)	Speedup
16	163.5	1.0
32	68.2	2.4
64	36.2	4.5
128	22.4	7.3

Table 1. Scalability of a Si(100) surface that has about 300 atoms.

CAR-PARRINELLO algorithm (materials science)

Determining the precision needed to manufacture the small-and-getting-smaller silicon devices at the heart of most of today's computer products is a long-term goal of Professor Jerzy Bernholc's work on the CRAY T3D system at the Pittsburgh Supercomputing Center. From a technology standpoint, the Si(100) surface is the preferred substrate for manufacturing such devices, which include silicon semiconductors. Knowing the high-temperature properties of the Si(100) surface—how the atoms move, vibrate, distribute, and reconstruct themselves with each increase in temperature—gives scientists great insight into annealing, growth, and diffusion, which will eventually lead to better growth and manufacturing processes for silicon devices. This investigation, which must run for hundreds of hours to accumulate sufficient statistics about the atoms' behavior, has been impractical on previous computing platforms.

Bernholc, physics professor at North Carolina State University, is currently using a variant of the CAR-PARRINELLO ab initio quantum molecular dynamics method to run in production mode on a heterogeneous CRAY T3D/CRAY C90 system simulations of a Si(100) surface that has about 300 atoms. "We've been extremely pleased with the performance and scalability we've seen on the CRAY T3D system," Bernholc reported (see Table 1). "We're getting 40 to 50 MFLOPS per processor element (PE) for a 16-PE run. In terms of access and availability, this very time-consuming job is most cost-effective on the CRAY T3D system, and we hope to have more access to this machine in the future."

The results of Bernholc's current work will serve as input to a Monte Carlo simulation of growth, which can simulate thousands of atoms and be performed over a much longer time.

The availability of resources on the CRAY T3D system has allowed us to ask questions and run problems we could not do previously.

CHARMM (Chemistry at HARvard Macromolecular Mechanics; Newtonian molecular dynamics)

Scripps Research Institute chemistry professor Charles Brooks and students have developed a heterogeneous version of CHARMM23, in which solvent calculations are performed on the CRAY T3D system and solute calculations are performed on the CRAY C90 supercomputer. The CRAY T3D and CRAY C90 systems complement each other by running portions of applications best suited for each system's specific architecture. A former chemistry professor at Carnegie-Mellon University, Brooks has used this version to conduct computational research into the free energy of protein folding in solution, a problem of major importance to pharmaceutical and polymer applications but currently confined mostly to time-consuming and expensive physical experimentation. "I need to get science done, period," said Brooks. "That need, and the need for more advanced algorithms to exploit the latest hardware technologies, has brought us to confront large-scale computation on MPP machines. The availability of resources on the CRAY T3D system has allowed us to ask questions and run problems we could not do previously. We have carried out an enormous number of peptide folding calculations, and we have used the CRAY T3D to explore the basic physics of the length dependence of helix stability."

A homogeneous (calculations done entirely on the CRAY T3D system) message passing version of CHARMM from Bernie Brooks of the National Institutes of Health also has been ported to the CRAY T3D system.

X-PLOR (x-ray crystallography)

Widely used in both industry and academia, X-PLOR is central to the determination of structures of large molecules such as proteins or nucleic acids (DNA). Knowing the structure of these molecules is essential to studying their chemistry—which drugs might inhibit their activity, for example. X-PLOR can be used for the refinement of structures solved by either x-ray crystallography or nuclear magnetic resonance (NMR). X-ray crystallography requires that an ordered crystal of the molecules be grown. The structure can be solved by observing how x-rays are reflected by this crystal. NMR can measure the distances between specific hydrogen atoms in molecules. With help from computer graphics, each of these methods produces an incomplete experimental description of the molecular structure. X-PLOR is then used to refine the model automatically by using the powerful technique of molecular dynamics combined with simulated annealing to provide the most accurate description feasible of the molecule.

The molecular dynamics portion of X-PLOR currently runs on the CRAY T3D system at the Pittsburgh Supercomputing Center and on CRAY T3D systems at Cray Research. "The interprocessor communication in the CRAY T3D system is very efficient," said Paul Adams, a Yale University postdoctoral associate of X-PLOR developer Professor Axel Brunger. "We've run X-PLOR on up to 128 processors and obtained speedups of factors of 30 to 40 with 64 processors for our initial test cases. The high-speed communications network is giving us better scalability compared to the other systems we have used so far."

AMBER (Assisted Model Building and Energy Refinement; Newtonian molecular dynamics)

Distributed by the Department of Pharmaceutical Chemistry at the University of California San Francisco (UCSF), AMBER is one of the most widely used molecular dynamics codes for simulating the energetics and motions of proteins and nucleic acids. Such simulation is increasingly pertinent in the realm of rational drug design, because proteins and nucleic acids are the key building blocks of viruses, including HIV, the virus

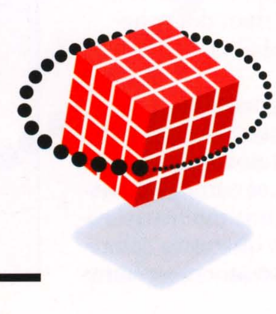

The high-speed communications network is giving us better scalability compared to the other systems we have used so far

responsible for AIDS, and are more generally involved in the etiology of disease. The more chemists can find out about the variety of interactions and bondings of small molecules, proteins, and nucleic acids over time, the better they can understand the structure and function of these molecules and design drugs to "fit" the receptor site on a protein. This additional knowledge requires a longer and better look at larger protein chains and nucleic acids than is practical on current computing platforms.

The CRAY T3D system offers the speed and parallelization necessary for running larger and more complex systems in a reasonable amount of time and at a reasonable computational cost. Although some state-of-the-art simulations in the range of 10,000 to 15,000 atoms are running for nanoseconds on high-performance supercomputers, most are confined to the picosecond timescale possible with workstations.

AMBER is made up of a series of modules. Jim Vincent of Penn State University, assisted by Tom Cheatham of the University of California San Francisco (UCSF) and David Case of Scripps Research Institute, used message passing programming to parallelize the two most compute-intensive modules, Sander (simulated annealing with NMR derived energy restraints) and Gibbs (free energy simulation), on the CRAY T3D system at the Pittsburgh Supercomputing Center.

Structures and fluid dynamics

Planes, trains, and automobiles—and the sophisticated algorithms and methods that drive the high-performance computing applications used in their design—will be major beneficiaries of applications development work being carried out on CRAY T3D systems. Crash simulation, computational fluid dynamics, and the effects of future aircraft on the stratosphere are some of the areas under investigation.

Table 2. Scalability of the MESA software suite's FD-TD module.

Number of PEs	Time (seconds)	Performance versus 64 PEs
64	861.22	1.0
128	423.04	2.04
256	219.37	3.93
512	139.33	6.18
1024	77.93	11.05

PAM-CRASH certification suite successfully run on CRAY T3D system

Cray Research MPP systems will be practical industrial production platforms for car crash simulation. PAM SYSTEM INTERNATIONAL, ESI-Group Software Product Company, developer of the PAM-CRASH industrial software program for crashworthiness analysis, drew this conclusion from a summer study of its code on the CRAY T3D system. A PAM-CRASH certification suite of problems was executed successfully on from 1 to 64 CRAY T3D system processors with consistently accurate results. These problems included the front-barrier-impact of a complete automobile.

Very high parallel efficiencies were achieved with no more than a few hundred finite elements assigned to each CRAY T3D system processor—a direct consequence of the superior interprocessor communication available on the CRAY T3D system. According to PAM SYSTEM INTERNATIONAL, this behavior, which is essential for scalability, has not been achieved on any other MPP system.

Computational fluid dynamics: SAGARMATHA and AVBP

Under the direction of Professors Michel Deville and I. L. Ryhming, researchers at EPFL, Switzerland, are adapting the block-structured SAGARMATHA code and the unstructured-mesh AVBP code to run on CRAY T3D supercomputers. These codes are particularly suited for modeling automotive flows, fluid mixing, and flows in turbines and pumps.

Parallelizing the SAGARMATHA code involves implementing both fine-grain parallelism, using data-parallel techniques for communication among mesh points, and coarse-grain parallelism, using work-sharing and message-passing techniques for communication at the block level. Already researchers plan to use the parallelized version of this code in the design of the proposed Swiss under-

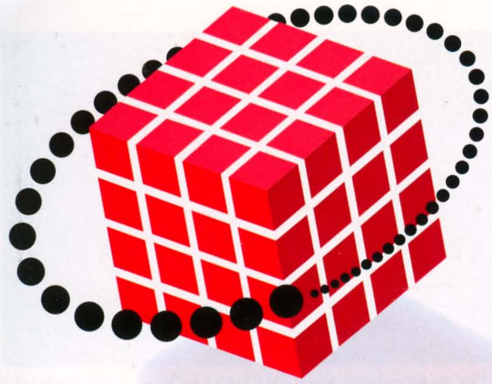

The CRAY T3D system is very stable and easy to use, with a relatively simple but robust programming environment.

ground train, the Swissmetro. A consortium of interests is evaluating the Swissmetro project, with all aerodynamic studies being performed on the CRAY T3D system at EPFL.

"All of the other technological aspects have been considered in the design of the Swissmetro, except the aerodynamics problem," explained Deville. "The shape of the nose of the train for example, will affect the aerodynamic efficiency, and turbulence effects between the tunnel walls and the body may induce instability of the cars inside the tunnel. We expect the CRAY T3D system to allow an order of magnitude improvement in the number of grid points that can be used for CFD studies like the Swissmetro evaluation, leading to deeper insight into the complex physics involved. The Swissmetro train, for example, will be very long, about 200 meters, and because of turbulence problems we must use a fine mesh. The overall problem will require up to several million grid points."

Parallelization of the AVBP code for the CRAY T3D system also is in progress. Work on this unstructured-mesh code is not as straightforward as parallelizing the block-structured code due to the irregular communication pattern among mesh points. Therefore, the parallelization will rely heavily on library software to map communications efficiently onto the hardware. Nonetheless, a coarse-grain parallelism can be implemented through mesh partitioning. Following partitioning, both message-passing and work-sharing programming models can be used, in a manner similar to that used with block-structured meshes. A simplified two-dimensional version of the AVBP code runs today on the CRAY T3D system.

Computational electromagnetics: parallelizing the FD-TD module of the MESA software suite

Cray Research analysts, working in close collaboration with British Aerospace (BAe) computational electromagnetics (CEM) specialists, recently converted the FD-TD module of the BAe electromagnetics software suite, MESA (Modules for Electromagnetic Simulation Applications), to take advantage of the CRAY T3D system. The conversion produced a code that has significant performance improvements and excellent scaling on the CRAY T3D system (see Table 2).

Key parameters in CEM are the electrical size of the problem, the algorithm dependency on the electrical size, and the available computational resources. The electrical size of the problem is proportional to the ratio of the body size to wavelength, or body size times frequency. For the traditional method of moments (MOM) and the FD-TD method, the computer time requirements vary as the sixth and fourth powers of the electrical size of the problem, respectively. Consequently, for a given body size, the FD-TD approach can model higher frequencies than the traditional MOM, moving into the GHz range for some analyses. Capabilities of this type are required for the design of many high-technology products that have to operate at these frequencies. "The FD-TD algorithm has a high degree of parallelism and data locality," explained Cray Research analyst Nick Mayes, "making it ideal for running on an MPP architecture. The shared memory of the CRAY T3D system provided an added advantage; we used it to create a code which runs much faster than an earlier message-passing implementation. The performance on the CRAY T3D system is excellent, with a 1024-processor system providing more than two-and-a-half times the performance of a CRAY C916 system on this highly vectorized FD-TD application. Further improvements will be achieved on the CRAY T3D system by incorporating a more sophisticated load balance algorithm in the MPP version."

British Aerospace welcomes approaches from organizations that are interested in using the CRAY T3D version of the MESA FD-TD module and will consider commercially packaging the module if sufficient interest exists. Applications of MESA include electromagnetic design of aerospace, automotive, and electrical products to shield equipment from potential sources of interference; antenna design; the design of microcircuitry, in which electromagnetic fields can corrupt signal transmissions; and a variety of other electromagnetic design problems.

British Aerospace cited the large memory of the CRAY T3D system as a valuable feature for performing CEM calculations. The FD-TD memory requirement varies with the third power of the electrical size of the problem. Problems that were previously solved on the British Aerospace CRAY Y-MP system with an SSD can now run entirely in memory, with greatly reduced run times. The CRAY T3D system is well-suited for CEM because "its fast interprocessor communications and architecture facilitate the exploitation of the high level of parallelism in the software, providing excellent scalable performance," said British Aerospace CEM analyst Tim Lanfear.

Scientific visualization and image processing: interactive exploration and scientific processing of large science datasets

NASA mission planners at the Jet Propulsion Laboratory (JPL) supercomputer facilities are using a CRAY T3D system to explore the best places for spacecraft to land and for satellites to survey the distant planet. Animated images of Mars created on the CRAY T3D system use a new unified image mosaic of the entire planet surface.

This work is part of JPL's experimentation with the use of a CRAY T3D system to explore and process large scientific databases interactively. Such databases can include satellite optical and radar images of planetary terrain. "The idea is that NASA scientists will be able to sit in front of a terminal and peruse and process large archival databases," explained JPL's principal investigator David Curkendall. "Results of the remote scientific processing can be transmitted over gigabit networks for interactive display and further analysis."

Although scientific data has been rendered on traditional supercomputers for many years, this type of visualization work is relatively new on parallel supercomputers. In the case of Mars, only the CRAY T3D system memory is large enough to accommodate the large images and datasets that have been acquired by NASA missions in the last 30 years. Its primary memory, combined with the computing power of its 256 processors, permits animation to be rendered nearly two orders of magnitude faster than with conventional equipment.

JPL's initial focus has been the software development of a three-dimensional renderer for use on CRAY T3D systems. It was designed to convert information from the fusion of image and terrain data into a synthetic perspective view of Mars. "Getting this to work well in parallel is a subject of considerable intellectual interest because this is a very natural way to view some of these scientific datasets," said Curkendall. "Everybody is used to looking at things in their natural perspective, and scientists are no exception." Curkendall added that "the CRAY T3D system is very stable and easy to use, with a relatively simple but robust programming environment."

A RECORD SETTER FROM THE BEGINNING

Two months after its introduction in September 1993, the CRAY T3D system achieved industry-leading performance on the highly parallel version of LINPACK, a widely recognized independent benchmark test for high-performance computer systems. On this test, a 128-PE CRAY T3D system achieved a sustained performance of 9.4 GFLOPS on a size 12,000 problem. The next-best performance by a second MPP vendor's 128-PE system on a problem of this size was 4.0 GFLOPS. A third MPP vendor's 128-PE system sustained 7.6 GFLOPS but required a much larger problem size to attain this speed. Recently run on a 1024-PE CRAY T3D system, the LINPACK benchmark demonstrated performance greater than 100 GFLOPS.

In March 1994, the first 256-PE CRAY T3D system achieved the highest performance ever reported for an MPP system on the complete set of eight NAS Parallel Benchmark tests, the most widely accepted independent comparison tests for parallel computer systems. Today, parallel applications on a 256-PE CRAY T3D system regularly achieve performance in excess of 10 GFLOPS. In October 1994, Cray Research was the first MPP vendor to report results for an MPP system with more than 1000 processors on the full set of NAS Parallel Benchmarks. A 1024-processor CRAY T3D system substantially beat previous MPP performance and scalability records on all eight benchmark tests. On the three Class A tests considered closest to real-world problems, called simulated applications, the CRAY T3D system was 3.6 to 8.7 times faster than the next-fastest MPP system.

Electron-molecule collisions experiences on the CRAY T3D system

Carl Winstead, Howard P. Pritchard, and Vincent McKoy,
California Institute of Technology

Wherever molecular matter is subject to ionizing energy, collisions between free electrons and molecules are important to the overall kinetics. This is true in natural settings such as interstellar clouds and planetary atmospheres (including Earth's), and in manmade environments ranging from high-voltage switches to the edges of fusion plasmas. Perhaps the most important technology in which electron-molecule collisions play a part is low-temperature plasma processing of materials. Plasma etching, plasma deposition, and several other plasma-based processes are at the foundation of the microelectronics industry.

In low-temperature plasmas, electrons acquire energies from a few electron volts (eV) to a few tens of eV—relatively low as electron energies are classified, but high enough to dissociate and ionize the molecules with which they collide. When the ions and radicals resulting from these collisions reach the surface being processed, they can produce desirable chemical and physical changes.

Efforts are now under way to improve the numerical modeling of plasma reactors. A long-term goal is to develop computer-aided design (CAD) tools for plasma reactors that will replace trial and error and statistical optimization techniques. The success of these efforts will depend on the availability of basic data on many events in the plasma and at the surface,[1] but electron-molecule collisions, because of their role in energy deposition, are especially important.

Determining electron-molecule collision cross sections is experimentally or theoretically difficult. Laboratory scientists must contend with various experimental difficulties, often including—particularly in the case of gases used in industrial plasmas—the hazards of working with pyrophoric, poisonous, or otherwise unpleasant compounds. Theorists face less dangerous but still daunting challenges inherent in the accurate treatment of collisions at low energies. This collision problem, described by many-particle quantum mechanics, is similar in many ways to the problems that arise in studies of bound-state molecular electronic structure—in "normal" computational chemistry. However, electron-molecule collision is different in some key respects. Because it treats free electrons, it is subject to different boundary conditions; for the same reason, electron-molecule collision does not have the structure of an eigenvalue problem that typically arises in computational chemistry.

Accurate ab initio solution of electron-molecule collision problems, in which cross sections are computed without the use of any experimental information, and quantum mechanical effects such as electron exchange are treated without approximation, is numerically intensive for even small molecules and rapidly grows more difficult as the size of the molecule increases. We have been studying these collisions for a number of years, first on sequential processors, including DEC VAX, CRAY X-MP, and CRAY Y-MP systems, and, beginning in 1989, on massively parallel processing (MPP) systems. Our parallel program ran first on the Jet Propulsion Laboratory (JPL)/Caltech Mark IIIfp hypercubes and soon was ported to the Intel iPSC/860 and Touchstone DELTA. (The latter is a unique prototype system with 512 i860 processors, installed at Caltech in 1991 and owned by the Concurrent Supercomputing Consortium.) At present our program runs on Intel's iPSC/860, DELTA, and Paragon computers, on the nCUBE-2, and, since early 1994, on the CRAY T3D system. Here we describe how we structure our computations for efficient solution on MPP systems and how we have used the CRAY T3D system to attack some very large collision problems associated with boron trichloride (BCl_3) plasmas.

Calculating collision cross sections

Various outcomes are possible when an electron collides with a molecule. The electron may scatter elastically, departing with its speed unchanged but possibly in a new direction, or it may scatter inelastically, giving up some of its kinetic energy to

for plasma modeling

the molecule. In the second case, the molecule may be left in an excited state, ionized, or dissociated. The probability that a collision will have a certain outcome is usually given in terms of a cross section, which has the units of area. Cross sections for inelastic processes are typically much smaller than elastic cross sections (meaning that inelastic collisions are much less probable) and also more difficult to calculate. However, inelastic collisions are inherently more interesting than elastic scattering, because they lead to some change in the state of the molecule.

We use a variational method[2] to obtain collision cross sections. Our method, an extension of Schwinger's method,[3] is suited to both elastic and inelastic collisions. Like most variational schemes, our method reduces the computational problem to the solution of a set of linear equations—in matrix form, finding the unknown matrix x in the equation $Ax=b$. All of the physical information about the molecule and the collision process is incorporated in the matrices A and b. From x, we can determine a set of scattering amplitudes, complex numbers whose square moduli yield a set of cross sections.

Constructing, not solving, the linear system is the "hard" part of our calculation. Evaluating the elements of the A matrix, in particular, requires a quadrature approximation to an integral operator. Since the quadrature is carried out over an infinite 3-D domain, tens of thousands of quadrature points may be required to obtain an accurate result. At the quadrature points, we must evaluate all possible six-dimensional integrals of the form

$$\int\int d^3 r_1 \frac{F_a(\vec{r}_1)F_b(\vec{r}_1)F_c(\vec{r}_2)e^{i\vec{k}\cdot\vec{r}_2}}{[\vec{r}_1-\vec{r}_2]} \quad (1)$$

Here, F_a, F_b, and F_c are Cartesian Gaussian functions, which will be familiar to computational chemists, and $\vec{k}$ is the quadrature variable. There are about $G^3/2$ unique ways to choose the indices a, b, and c given a set of G Gaussians, so for G around 200 (typical in our recent work) we may need to evaluate as many as a few million such integrals at each of the quadrature points in $\vec{k}$. Overall, up to 100 billion integrals might be required in a single cross section calculation.

Fortunately, the integrals in Equation 1 may be evaluated analytically. (This is no coincidence: Gaussians are used precisely because the resulting integrals are doable.) Any single integral may therefore be obtained with little computational effort; what makes the problem numerically intensive is the vast number of integrals involved. However, it is not just evaluating all those integrals that is demanding. There are several intermediate steps in the process of building the elements of A and b from the integrals of Equation 1, and the arithmetic associated with these steps often involves more operations than the actual computation of the integrals.

Using parallel computers

From the above discussion, we can see the outline of a numerical procedure: first, evaluate a large number of integrals of the type shown in Equation 1; next, combine those integrals in appropriate ways to construct matrices A and b; finally, solve the resulting linear system $Ax=b$. In fact, this is exactly how we solve the scattering problem, except that we evaluate the integrals in small batches (a few million at a time) rather than all at once, and we do most of the necessary manipulations on each batch before proceeding to the evaluation of the next batch. Approaching the problem this way requires less memory and makes it possible to perform long calculations by combining several shorter runs.

It should be obvious that at least the first step in this procedure—evaluating a batch of integrals—is well suited to MPP platforms such as the CRAY T3D system; we simply assign each processor a different subset of integrals to evaluate. Combining the integrals is trickier, because the rules for building elements of A and b from the integrals are fairly complicated and because the processors must communicate with each other in this phase (each processor has only a small fraction of the current batch of integrals in its own memory). An efficient way of carrying out this step is to use the complicated rules for combining integrals to build a single "transformation matrix" that, when multiplied with the current batch of integrals (arranged as a matrix), gives the necessary combinations as the product matrix. Each processor can build a block of this transformation matrix independently, so that the only communication among processors occurs during the matrix multiplication (see Figure 1). A subsequent step in evaluating A—namely, the angular part of the quadrature over $\vec{k}$—can also be done as a matrix multiplication (see Figure 2). In short, we can formulate both the evaluation and the combination of the integrals in terms of procedures that are either "perfectly parallel" (involving no communication overhead) or highly efficient and relatively easy to program (multiplication of large, dense, distributed matrices). The remaining step, solving $Ax=b$ for x, can also be done on the MPP system; however, since our linear systems tend to be small, it is equally convenient to do this step on a conventional computer.

Porting and running on the CRAY T3D system

JPL's CRAY T3D system was delivered in late December 1993. At that time we had already begun porting our program, using a CRAY T3D system in Chippewa Falls, Wisconsin. Because our program had been running on Intel systems, porting it to the CRAY T3D system was greatly simplified by a library of functions, developed by Phil Campbell of Cray Research, for emulating Intel's NX message passing. By the time the JPL system was dedicated in mid-January 1994, our port was essentially complete, and within a week we were doing full-scale production runs.

We were immediately impressed with the floating point and communication performance of the CRAY T3D system, especially in the distributed matrix multiplications described above, which are sensitive to both factors. Using the LibSci routine SGEMM to multiply local matrix blocks and using direct remote memory access (via shmem_get) rather than message passing, we have seen throughput in this step (including communication and synchronization overhead) around 60 MFLOPS/processor or better. The matrix multiplication also illustrates another of the CRAY T3D system's strong points—low communication latency. This manifests itself in the large fraction of maximum performance that one obtains even for relatively small matrices, where many small data blocks are transferred.

Evaluating the integrals of Equation 1 is reasonably fast in comparison to what we have seen on other RISC processors, though at about 5.5 MFLOPS/processor, it is far from peak. The same factors that inhibit vectorization of this step also inhibit its performance on cache-based microprocessors. However, we believe there is room for improvement as the compiler matures and as we have an opportunity to tune the code. Recently, moreover, we have seen a large improvement in the effective speed of this step (to about 10 MFLOPS) using some nonstandard substitutes for intrinsic functions such as sin, exp, and sqrt.

One area where performance was surprisingly low was in constructing the transformation matrix. Here the problem turned out to be indirect addressing. While we cannot eliminate indirect references in this step, we were able to obtain about a factor of two speedup by increasing the re-use of indices once loaded. A much larger improvement came from an algorithmic change that reduced the number of indirect references substantially (at the expense of storing some auxiliary data).

In most of our work to date, disk I/O has not been a limiting factor. However, when many possible scattering outcomes are considered in the same calculation, I/O requirements grow rapidly; up to a few Gbytes of precomputed intermediate data may be required to construct the A matrix. With very little effort, we have obtained rates of about 6 Mbytes/s overall between the CRAY T3D system and disk with all processors reading a single file simultaneously. We are confident that this rate could be improved significantly, but it is sufficient for our present work.

The overall performance of our program cannot be characterized by a single number. Although evaluating and subsequently combining the integrals are by far the dominant steps in the calculation, the relative importance of these two steps depends strongly on factors such as the size of the molecule and the number of inelastic processes being considered. Therefore, depending on the nature of the calculation at hand, and taking into account load balance, communication, I/O, and other overhead, we typically see throughput speeds in the range between 1 and 10 GFLOPS. To put these numbers in perspective, our original sequential program, which is only partially vectorizable, averages about 30 MFLOPS on a CRAY Y-MP processor—roughly comparable to the speed per processor of our parallel program

Figure 1. Transformation of the one- and two-electron integrals has the form of a distributed multiplication, with each submatrix residing on a different processor. The process is repeated for each G.

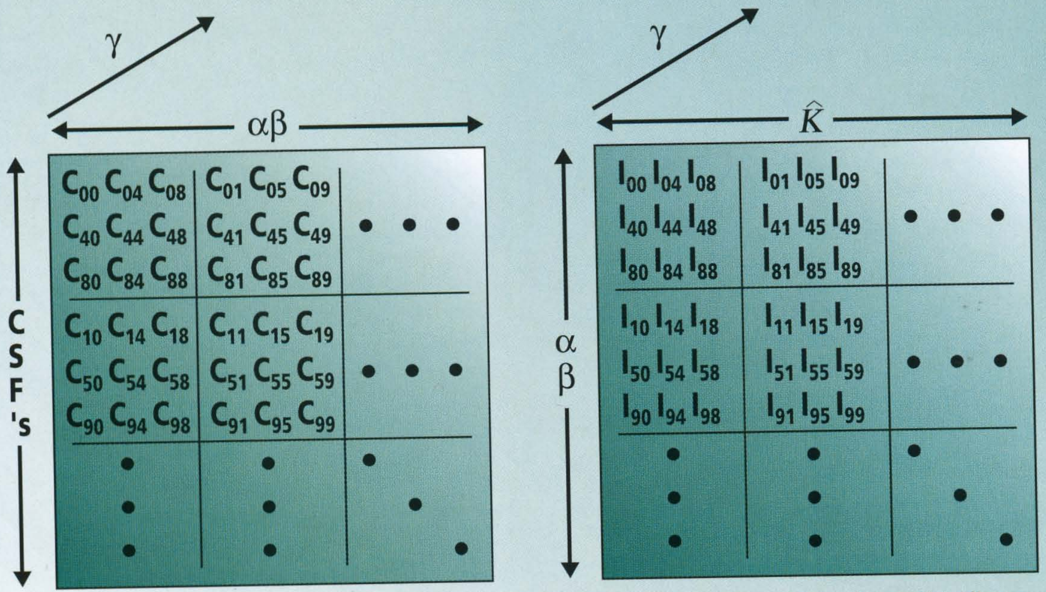

A transformation matrix takes integrals over ao's directly to final matrix elements over CSF's.

A complex matrix of two electron (six dimensional) integrals, I(αβ,*k*;γ) running on the CRAY T3D system.

Applications

We have used the CRAY T3D system at JPL to carry out a number of studies that are part of a project undertaken for the semiconductor industry research consortium, SEMATECH. This project involves calculations of elastic and inelastic cross sections relevant to boron trichloride plasmas used to etch silicon, and thus includes the molecules BCl_3 and $SiCl_4$ together with the radicals that can be formed from them by removing chlorine atoms (BCl_2, $SiCl_3$, and so on).

An example of this work is shown in Figure 3, which represents the cross section for elastic electron collisions with BCl_3 as a function of both the collision energy and the scattering angle. The structure visible in the cross section at low energies is caused by the temporary trapping of the impinging electron in an empty molecular orbital. Analysis of such temporary anions, or "resonances," is, for us, one of the more interesting aspects of studying electron-molecule collisions. To a plasma modeler, of course, the general shape and magnitude of the cross section over a broad energy range may be more important. One of the attractive features of these calculations is that they simultaneously address questions in both basic and applied science.

Molecules as large as $SiCl_4$ have never yet been accessible to ab initio scattering methods like ours. Our study of elastic scattering by $SiCl_4$ would alone consume hundreds of CPU hours on a vector supercomputer; the completion of a comprehensive survey of elastic and inelastic collisions for an entire class of molecules would be prohibitively time-consuming (not to mention expensive). In providing two orders of magnitude in performance improvement, MPP systems have changed the definition of a "feasible" electron-molecule collision calculation and thereby allow us to address problems of practical interest that very recently were beyond our reach.

Continuing rapid advances in the power of MPP systems will make it possible to increase both the scale and the sophistication of theoretical studies of electron-molecule collisions, and to provide increasingly reliable data on an ever broader range of molecules. Equally important, MPP systems will almost certainly provide the "horsepower" for future 3-D plasma models that will allow the scientists and engineers who are the consumers of our cross section data to perform realistic simulations of plasma reactors. It is rewarding to be working in an area where high-performance computing is bridging the gap between basic science and technology.

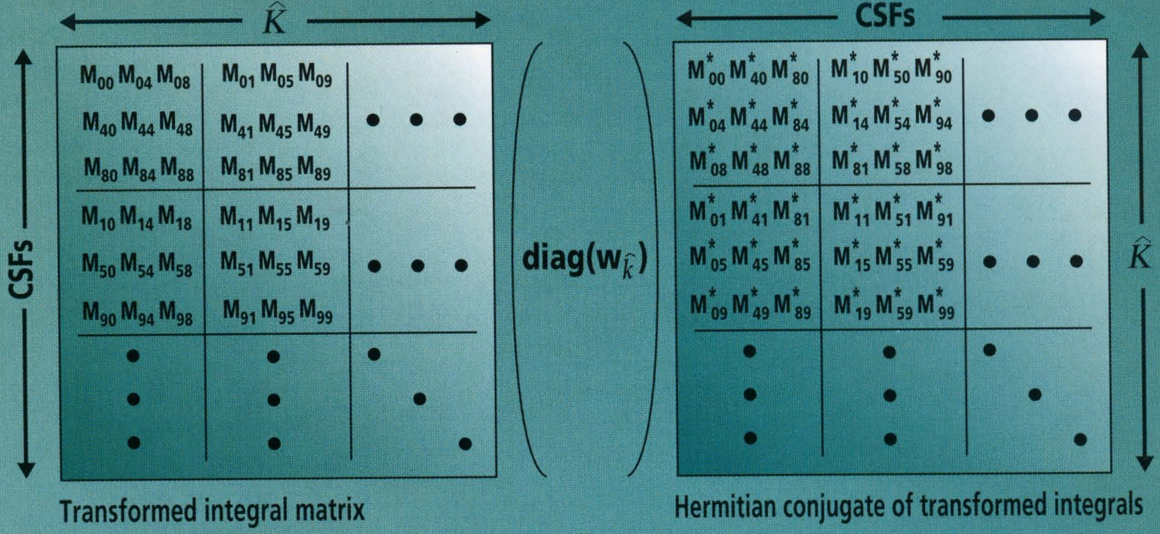

Figure 2. Integration by quadrature can also be organized as multiplication of distributed matrices.

About the authors

Carl Winstead is a senior research fellow at the California Institute of Technology. He received a Ph.D. degree in physical chemistry from Indiana University and a B.A. degree from Rice University.

Howard Pritchard received a Ph.D. degree in chemical physics from the California Institute of Technology and has recently joined the Applications Division at Cray Research. He received a B.A. degree from Rice University.

Vincent McKoy is Professor of Theoretical Chemistry at the California Institute of Technology. He received a Ph.D. degree in chemistry from Yale and a B.S. degree from Nova Scotia Technical University.

References

1. National Research Council, *Plasma Processing of Materials: Scientific Opportunities and Technological Challenges*, National Academy Press, Washington, D.C., 1991.
2. Takatsuka, K. and V. McKoy, *Physical Review A*, Vol. 30, p. 1734, 1984.
3. Schwinger, J., *Physical Review*, Vol. 72, p. 742, 1947.

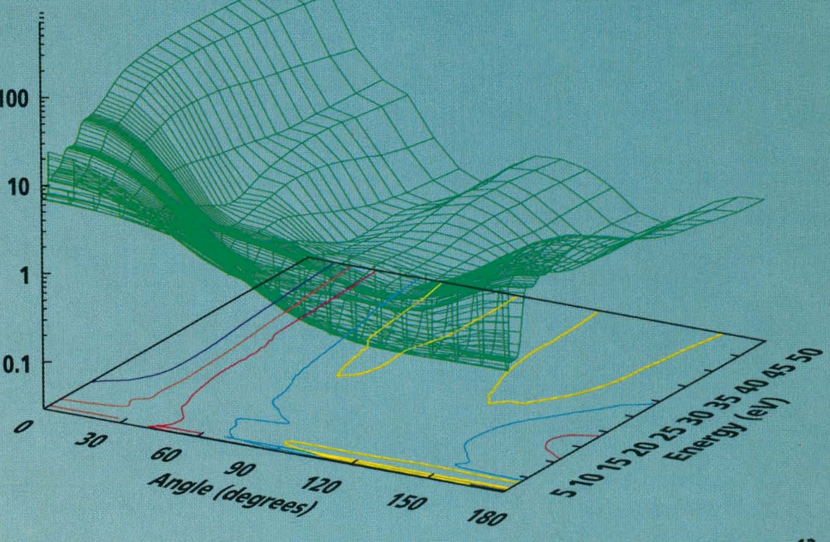

Figure 3. The cross section for elastic collisions of electrons with BCl_3 as a function of impact energy and scattering angle. Structure at low energy is due to resonances.

CRAYSOFT

because no workstation is an island

Judy Smith, Cray Research, Inc.

In almost every computing environment, desktop systems are networked either with other workstations and servers or with mainframes and supercomputers. For years, Cray Research customers have requested that certain Cray Research software be made available to them for use on workstations to increase programmer productivity and more efficient use of their supercomputer resources. We have responded with key products that create a common application development and production environment for the customer's computing complex. CraySoft was created to provide selected Cray Research software on desktop systems and servers in support of Open Supercomputing and to establish de facto industry standards based on Cray Research software technology.

CraySoft brings the power of Cray Research software to desktop and server systems by creating a UNIX-based, enterprise-wide, and distributed scientific and engineering environment for application development and problem solving from the desktop to the supercomputer. This environment includes Cray Research compilers, tools, libraries, networking components, and application products.

CraySoft represents a unique brand name, product style, and distribution capability for Cray Research software products. The brand name is applied to desktop and server software products provided by Cray Research or through technology agreements with third-party vendors. Although Cray Research customers requested these workstation products, you do not have to own a Cray Research system to use them. CraySoft products are available through CraySoft direct sales (800/BUY-CRAY or +1-612/683-3030), Cray Research sales offices, or third-party distributors.

Products sold and distributed by CraySoft are generally adaptations of existing Cray Research software technology for the desktop or server environment. The objective is to leverage much of the time, talent, and energy already invested by Cray Research in the development of UNICOS-based supercomputer products into the production and delivery of CraySoft products for the desktop environment.

CraySoft Network Queuing Environment (NQE)

The CraySoft Network Queuing Environment (NQE) is a suite of products providing enterprise-wide batch service. Released in December 1993, NQE was the first CraySoft product. The user submits a job to queue and NQE selects the most appropriate system on a network, and submits the job for execution on the selected system, providing more efficient use of system resources and better turnaround for user jobs. CraySoft NQE is packaged as a server and a client. The NQE Server consists of the Cray Research-enhanced Network Queuing System, Network Load Balancer, and File Transfer Agent for UNIX workstations. The NQE Server provides a stable network batch queuing environment that is compatible with public domain NQS and supports destination selection, load balancing, and status of tasks across a batch complex.

The Network Queuing Client (NQC) lets users submit jobs to the NQE Server and is a simplified version of the current RQS. NQE is available for systems that run Solaris on SPARC systems and

the Cray Research CS6400 system, SunOS 4.1.3 for SPARC, AIX for IBM RS6000, IRIX for SGI, HP/UX for HP9000, and OSF/1 for DEC Alpha systems.

NQE 2.0, released in the first half of 1995, provides additional capabilities in job dependence and network load balancing.

CraySoft Fortran 90 Programming Environment

The CraySoft Fortran 90 Programming Environment, initially released in August 1994, brings to SPARC platforms a complete development environment built around the Cray Fortran 90 compiler. This programming environment includes an optimizing and parallelizing ANSI/ISO standard Fortran 90 compiler and development tools such as an interactive visual program browser, an advanced symbolic source-level debugger, and a powerful graphical tool for measuring and improving parallel processing performance of application codes. This product provides a Cray Research-like program development environment on SPARC systems. The CraySoft Fortran 90 Programming Environment can be used to develop applications on SPARC systems using the same compiler and programming tools found on Cray Research supercomputers. The CraySoft Fortran 90 Programming Environment supports SPARC-based systems running the Solaris 2.3 or 2.4 operating system, including Sun SPARC systems and Cray Research's CS6400 system.

In the absence of a certified Fortran 90 test suite, Cray Research has invested in the development of an excellent test suite that tests all aspects of the ANSI/ISO Fortran 90 language. The test suite is used extensively in the testing of Cray Research's Fortran 90 compiler products and has been selected by and sold to several other compiler developers and platform vendors to test their own products.

CraySoft LibSci Numerical Library

CraySoft LibSci, also released in August 1994, is a collection of more than 1200 mathematical and scientific analysis subroutines that have been optimized to deliver the highest possible performance on single-processor and parallel-processor SPARC platforms. LibSci is well known as providing high performing mathematical libraries on the Cray Research platforms. With CraySoft LibSci, user programs can take advantage of parallel processing and the underlying SPARC hardware, just as they can on the Cray Research parallel vector systems. CraySoft LibSci includes routines to solve systems of linear equations, FFTs, signal/image processing functions and sparse matrices, eigenvalue problems, and matrix manipulations. CraySoft LibSci routines can be called from Fortran, C, or assembly language programs. CraySoft LibSci is supported for parallel execution with the SunPro F77, SunPro C, Apogee f77, Apogee C, and CraySoft Fortran 90 compilers for SPARC systems.

CraySoft Distributed Programming Environment (DPE)

The CraySoft Distributed Programming Environment (DPE), released in June 1994, brings the interactive portions of the Cray Research CF90 programming environment to workstations. DPE enables users to perform preliminary code cleanup and debugging using the Fortran 90 language parser and interactive program browser on their desktop SPARC systems. Once code has been compiled and executed on the Cray Research system, DPE provides the user with the ability to move performance profiling information to the desktop for further analysis and code improvement using Cray Research performance tools. DPE is available for SPARC systems that run either Solaris 2.3 or SunOS 4.1.3.

CraySoft Open Storage Manager (OSM)

Cray Research and Legent Corporation have entered into a joint licensing and distribution agreement under which CraySoft is marketing, selling, and distributing Legent Corporation's Open Storage Manager (OSM) product to the high-end science, engineering, and industrial marketplaces. OSM improves the efficiency of managing data storage by providing a transparent, scalable method of storing and retrieving large amounts of data and infrequently used files on less costly storage devices. The first release of OSM supports SPARC SunOS and Solaris systems.

OSM is a family of hierarchical storage management technologies that, when combined, provides automated storage management across a network of clients and servers. OSM is packaged by CraySoft as two individual storage management products — CraySoft OSM Client and CraySoft OSM Server. The OSM Client is a file migration product for UNIX environments that extends the capacity of disk-based file systems by routing selected files to alternate storage subsystems. The OSM Server provides a massive capacity storage repository and the application software needed by system administrators to match data and its access characteristics with a hierarchy of network storage resources. It also performs central volume and device management control of access to removable media volumes, drives, and automated library resources within a network.

CraySoft products were the first from Cray Research to be shipped with online documentation readable by using the CrayDoc browser program. All product documentation is distributed online, with only the installation guide and minimal printed documentation being distributed with the product. All CraySoft products and documentation are distributed on CD-ROM.

For additional information about CraySoft or its products, send email to craysoft@cray.com, call 800/BUY-CRAY or +1-612/683-3030, or use your favorite World-Wide Web browser to access http://www.cray.com/.

About the author

Judy Smith has been with Cray Research since 1992 and is the technical project leader within the CraySoft organization. In this position she is responsible for product coordination and packaging. She has had many years of software development and test management, and project management experience at several other computer companies. Smith has a mathematics degree from Texas Tech University.

Applications Division

partner for solutions

W. Derek Robb, Vice President, Applications Division, Cray Research, Inc.

In November 1994, Cray Research elevated its focus on applications solutions by creating the new Applications Division. The mission of the division is to help our customers find superior value solutions to their problems through applications software and expertise partnerships. The most complex problems require or involve independent software vendors (ISVs), one or more Cray Research customers or prospective customers, and a group of experts and products from the Applications Division.

The new Applications Division is already well-positioned to respond to requests for solutions. Over 100 experts with broad industrial experience in engineering and scientific disciplines as well as expertise in using Cray Research supercomputing technology staff the division. We also work as part of a larger "virtual" division of applications experts in our sales support group, software division, and our field offices around the world. In 1995 the Applications Division will enhance its partnerships with ISVs, create applications tools and solutions—internally and at our early CRAY T3D sites—by building on years of product development experience, and find and overcome new computational barriers to industrial productivity by working with customers on the problems they need to solve.

ISV partnerships: the key to industrial solutions

Cray Research collaborates with over 400 ISVs to support the broad spectrum of applications programs available on all of our systems. We work one-on-one with the developers of some of our customers' most important applications to make sure they have access to effective mathematical algorithms, expertise in optimization and programming models, and all of our hardware and system software. The 15-year collaboration between Cray Research and the MacNeal Schwendler Corporation (MSC) is a good example of the payoff of our ISV partnerships.

The goal of our partnership with MSC is to ensure that our mutual customers receive optimal functionality and performance from each new release of MSC/NASTRAN, a popular, general-purpose finite element analysis application. In 1989 Cray Research and MSC began working together to test and optimize new sparse matrix methods on large customer problems. At that time, the biggest model solved had 65,536 degrees of freedom (DOF); with the new methods, models as large as 1.5 million DOF can be solved routinely on Cray Research systems—30 to 40 times faster in CPU time than the older algorithm's sparse matrix allows. In addition, a new set of flexible file I/O (FFIO) methods was developed in the Applications Division as a result of our deep understanding of how applications like MSC/NASTRAN work; this has further cut the elapsed time to solution for large NASTRAN problems.

Dr. Louis Komzsik, chief numerical analyst at MSC, recently commented: "As Cray Research has expanded its product line, from the entry-level systems to the highly parallel end, cooperation between MSC's software developers and members of Cray's Engineering Applications Group (EAG) has become even more crucial. This cooperation produced excellent results on our production system of MSC/NASTRAN version 68, which is the best Cray Research version of NASTRAN that we've ever had.

"Cray's EAG personnel directly suggested some of the performance-related improvements, especially in the I/O area. Other improvements could not have been made without Cray's generous assistance in computing power and computer science support," he said.

Another tangible result of our strong ISV partnerships is the 6 GFLOPS performance obtained on a customer's ANSYS problem. This work was part of the applications suite that won Cray Research honorable mention in the 1993 Gordon Bell competition for computer performance.

In 1995 we will strengthen our ISV program by making more computer resource available to key vendors, communicating more frequently with vendors about business, as well as technical issues of mutual interest, and performing extensive comarketing with videos and brochures.

Partnering with customers

Customer requirements are at the core of each of our vendor partnerships and application development efforts. As the Cray Research MPP architecture emerged a few years ago, it created fertile ground for a new level of customer partnerships.

Rapid acceptance of Cray Research scalable parallel computing platforms depends on the availability of a broad spectrum of MPP applications software. Because there are so few commercially available applications, Cray Research and several of our customers with extensive R&D expertise in MPP applications are aggressively involved in a Parallel Applications Technology Partnership (PATP) program, aimed at providing about 100 new scalable applications on the CRAY T3D system and its follow-on. These partners—L'École Polytechnique Fédérale de Lausanne (EPFL) in Switzerland, the Pittsburgh Supercomputer Center (PSC), Caltech and the Jet Propulsion Laboratory (JPL), Los Alamos National Laboratory (LANL), and Lawrence Livermore National Laboratory (LLNL)—have committed to applications development projects with Cray Research and other commercial firms including ISVs and industrial end-users. PATP involves a total commitment valued at over $100 million on the part of Cray Research and its partners. The projects bring together over 100 staff from the PATP partners.

Each partner brings a history of leadership in high performance computing and computational simulation. Together, they have expertise in a broad range of application and computing domains. At the end of the first year of the three-year partnership, we see a number of successes: 60 projects are active, with a like number of applications running successfully on CRAY T3D systems at the member sites. Several applications are running at over half the peak speed of up to 512-processor systems; and most applications are sustaining 10 to 50 MFLOPS per node in parallel. EPFL is using leading-edge adaptive-meshing technology to solve fluid flow problems in the automotive area. At PSC, huge molecular dynamic simulations for proteins are being carried out to help understand protein folding. At JPL, engineers have carried out complex electromagnetics simulations at over 90 MFLOPS per node; they are also unraveling Synthetic Aperture Radar data and processing 3-D images at unprecedented speeds. Mathematicians and scientists at LLNL are carrying out complex 3-D simulations of groundwater transport in geologically realistic models of soils. This work will be critical in planning environmental remediation for hazardous waste sites. At Los Alamos, in cooperation with Amoco, engineers have developed black oil reservoir models that will enable enhanced management of existing and new oil fields.

In the second year of PATP, we are working with a number of ISVs committed to commercializing several of the PATP applications. Among them are Q-Chem, a new ISV for quantum chemistry modeling; IT Corp., which is commercializing ground water models; Tigresse, which is commercializing the black oil models, and APR, working to provide a high-performance FORTRAN capability for the CRAY T3D system. We will start more projects and work to make many of these applications advances available to the Cray Research customer family as both supported and unsupported products. Finally, we recognize the potential for mutually beneficial partnerships with other leading R&D institutions and will augment the PATP with an international group of Parallel Technology Centers.


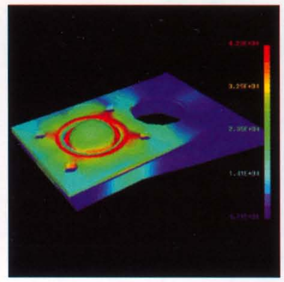


From left: frontal impact simulation with Hybrid III Dummy, courtesy of Kia Motors Corporation; simulation of metalforming produced with the ABAQUS software package; simulation of an automobile instrument cluster bezel.

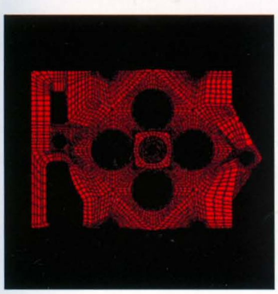
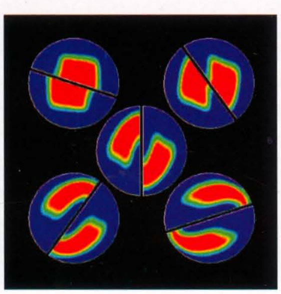

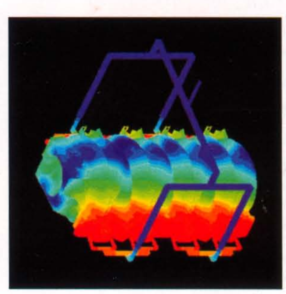
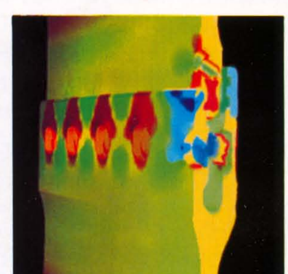

From left: cylinder head model with 192,000 degrees of freedom and 2400 RMS wavefront, analyzed with the ABAQUS software package; concentration fields of the transport of two chemical species, shown at various intersections of an inline static mixer's helical mixing elements as the mechanism has passed through various rotations, courtesy of Chemineer, Inc.; finite element model of a human femur, courtesy of Lamb and Company; a MOLDFLOW filling analysis for an automotive intake manifold; Space Shuttle Solid Rocket Motor factory joint pressure simulation.

Applications consulting services: providing hands-on support

Hyundai Motor Company recently chose a CRAY Y-MP 4E system as its first supercomputing system. Prior to installation of the system in Korea, several Cray Research applications experts were contracted to work with a group of 14 Hyundai engineers in our applications lab to define and refine

Applications Division

how Hyundai could perform an integrated car crash analysis. The group agreed on three objectives for the work: perform a full car integrated frontal crash simulation, teach Hyundai engineers to use and optimize Cray Research's combustion simulation software package, and advance Hyundai's use of production CFD analysis.

Using several of the high-performance computing resources in our corporate computing facility, including a fully dedicated large-memory CRAY M90 system, our applications analysts demonstrated new techniques for building the most comprehensive model possible: how to fold the airbag, how to include the Hybrid III Dummy (the first human-like and most advanced car-crash dummy), and how to integrate the seat belt into the overall model.

Working from an initial model that had only an airbag, toe panel, and seat, and conducting separate analyses for each piece, the team added more detail—a crossbar, a steering column, a steering wheel, and, eventually, the entire car. Consulting with Engineering Systems International, developer of the PAM-CRASH code Hyundai uses, the team determined the optimal numerical model for the airbag, installed it in the full car model, and began to answer questions about how the dummy would interact with each new component.

In 1995 we plan to expand our consulting projects in other industries and are eagerly looking forward to the new challenges this will bring.

Partners for performance

Supporting strategic third-party vendors. Developing and supporting applications solutions to open new markets. Overcoming computational barriers to industrial productivity. It all boils down to providing timely answers to real-world questions that are valuable to our customers and superior to any alternative solution. With our focus on high performance application programs that run on Cray Research platforms—from compact, air-cooled CRAY J90 systems to the top-of-the-line CRAY T90 system and the MPP CRAY T3D system—the mission of the Cray Research Applications Division is to ensure that our customers achieve their technical and business objectives, now and in the future.

About the author

W. Derek Robb is vice president of Cray Research's Applications Division. Since joining Cray Research in 1980, he has also served as vice president of Marketing and director of Sales and Marketing Support. He has a Ph.D. degree in theoretical physics and a B.Sc. degree in mathematics from the Queen's University of Belfast.

Applications software products: developing novel solutions

Our work with customers and vendors yields many opportunities for creating new solutions or enhancing existing ones with software created in the Applications Division. In 1994 we announced two new products and added functionality with new releases of existing products. In 1995 we will continue to leverage our expertise in the same way, making more of our products available beyond the Cray Research hardware family and creating development and marketing partnerships with outside organizations.

The most recently introduced application, **HEXAR**, helps engineers avoid the time and effort of handcrafting spatial meshes used in their analyses. HEXAR uses unique, patented algorithms to convert computer-aided design (CAD) descriptions of objects into 3-D finite element volume meshes. These meshes, or grids, can contain hundreds of thousands of hexahedral (brick-shaped) elements in typical computational fluid dynamics (CFD) or structural analysis problems. With HEXAR, engineers can explore more design options by building meshes in minutes and hours instead of weeks or months.

CRI/*TurboKiva* is a combustion simulation software package developed by Cray Research in collaboration with the Los Alamos National Laboratory to help reduce the average two-year period needed to test efficiency and emissions of new engines. Customers, including the Hyundai Motor Company and Nissan, use CRI/*Turbo*Kiva to study and optimize fuel injection, intake, ignition, and even two-stroke scavenging.

To help customers get better plastic parts to market faster, the Applications Division last year created the Cray Molding Logic software, **CMLogic**. Working with customers in the plastic injection molding industry, the Applications Division combined a new design of experiments (DOE) interface, third-party molding simulation software, and Cray Research hardware to create a "virtual molding machine." CMLogic interfaces to the molding simulation software MOLDFLOW, from Moldflow pty., and C-MOLD from AC Technology. Ultimately, the complete package reduces costly mold trials and optimizes the machine process, thereby improving time to market and profitability. (See the related story on page 28.)

UniChem is our most established application, with 69 licenses in place to date. UniChem integrated an easy-to-use graphical user interface with three quantum chemistry methods—ab initio, semi-empirical, and density functional—through a common user environment. Recently the UniChem development effort has focused on better performance and wider availability of the technology. Within the last year, a new analytic second derivatives for DGauss has provided a fivefold increase in performance, allowing previously intractable calculations to be done. For example, the harmonic vibrational spectrum of C_{60} (Buckminsterfullerene) has been computed with density functional theory (DFT), and the calculated values are in excellent agreement with experiment and suggest that some previous experimental results had been misunderstood. Finally, the standard DGauss Basis Sets developed by Cray Research and optimized specifically for DFT have been made available to the chemical community at large.

One by-product of UniChem is a general **Application Integration Toolkit (AIT)** for creating client-server applications in the job-oriented environment characteristic of supercomputing. It permits remote job initiation, monitoring, and control, a kind of "remote procedure call" for batch jobs. AIT is available for developers of distributed applications on Cray Research systems with clients running on Sun, SGI, IBM, and DEC systems.

Most of the Cray Research application product set is well-established with our customers. However, our continued ability to develop and support these applications depends on finding customer and vendor partners to help determine product directions and new markets outside the traditional Cray Research customer base.

Open supercomputing

A desktop window into the high performance world

Kathy Nottingham, Cray Research, Inc.

Imagine solving your most challenging computational problems in a small fraction of the time it takes today. Imagine using your personal computer (PC) or workstation to transparently access the high-performance features of a supercomputer, including ultra-fast computational power, the fastest available I/O, and the largest capacity storage systems. Your interaction with the supercomputer is simple and familiar to the way you work, your job, your problem. The interface to the supercomputer adapts to you—you do not have to know anything about supercomputing. What's important is that you just got the fastest possible solution to a critical problem with no extra effort on your part.

At Cray Research, we call this vision "Open Supercomputing." Much of this vision is in place today; it will be even more powerful and user-friendly in the future. Open Supercomputing means that computational chemists, Wall Street stock analysts, automotive parts designers, rural physicians, college students—whoever—will interact with their desktop systems and only their desktop systems to do their work. But user access to resources extends far beyond the desktop. With little or no intervention on a user's part, PCs, workstations, clusters, massively parallel processing (MPP) systems, high-end UNIX server systems, and general-purpose parallel-vector processing (PVP) systems, all residing on the "data superhighway," are at the user's disposal. What's more, each of these systems automatically is assigned the specific computing task it does best. Leading-edge Cray Research software does it all.

Removing the roadblocks to high-performance computing

Today, one-half of Cray Research's technical personnel—nearly 450 employees—are devoted to software and our vision for Open Supercomputing. We believe Open Supercomputing is where the world is headed, and we are leading the charge, just as we've led the hardware charge for years.

To fulfill the Open Supercomputing vision, Cray Research is focusing on six areas: standards, scalable performance, resource management, security and availability, high-bandwidth data accessibility, and a cohesive environment.

Standards: the backbone of Open Supercomputing

Standards, the basis for Open Supercomputing, let heterogeneous architectures on a network work together to efficiently solve distributed client/server problems. Standards also simplify the porting of applications from one system to another and let users easily move from one operating environment to another without learning a new interface or changing data formats.

In addition to adhering to software industry standards, Cray Research actively participates in defining and developing the leading national and international standards for operating systems, networking protocols and distributed computing, and software languages with organizations such as IEEE, ANSI, ISO, OSF, COSE, and X/Open. Standards are the basis of Cray Research software: the UNICOS operating system has been certified POSIX 1003.1 System Services standard compliant, Cray Research systems support all major networking communication protocols and media connections, and Cray Research provides industry-leading programming environments to support Fortran 90, C, and C++.

> Today, one-half of Cray Research's Technical personnel—nearly 450 employees—are devoted to software and our vision for Open Supercomputing

Future plans and challenges for Cray Research software personnel include working to address new and emerging standards. For instance, the next major release of UNICOS (Cray Research's 10-year-old UNIX operating system) includes X/Open Base Profile Branding. A challenge for the company includes developing a single microkernel-based (Chorus-based) UNIX for all Cray Research supercomputers. Other challenges include creating interoperability with Windows NT (a desktop operating system) and a future standard API for microkernel technologies.

In networking and distributed computing, Cray Research is currently field testing asynchronous transfer mode (ATM). As part of the SuperCluster project, Cray Research also is developing distributed file system capabilities for transparently accessing remote data across the network by using OSF/DCE Distributed File System (DFS) and ONC+ NFS Version 3.

Future language development calls for a full, native C++ compiler and a continuation of industry Fortran leadership with an early implementation of Fortran 95.

Scalable performance: tapping the speed of supercomputers

Supercomputers provide the fastest solutions possible for large, computationally intensive problems; high performance is the hallmark. Cray Research software will continue to deliver the scalable performance of Cray Research supercomputers, but not at the expense of accessibility, portability, or standards. Open Supercomputing will simply increase the availability and usability of supercomputer performance to the average user.

Today Cray Research provides efficient, scalable parallel processing operating systems with multithreaded UNICOS kernel (for PVP systems) and UNICOS MAX microkernel implementation (for MPP systems). We plan to increase the scalability and efficiency of all Cray Research supercomputers by introducing a common, Chorus microkernel-based serverized UNICOS operating system. We will continue our leadership in the area of automatic optimizing compilers and libraries to ensure the highest application performance.

As computer architectures become increasingly complex, smarter software will be necessary to harness the underlying power of new hardware features. Users will require software that insulates them from hardware differences and masks this complexity. Serverized UNICOS is the first step toward a long-range goal of a distributed operating system that can integrate the variety of architectures on the network to deliver the best possible performance from the network, presenting it as a single system to the user.

Resource management: making the most of a valued resource

In the world of Open Supercomputing all work is distributed across the many heterogeneous architectures on the network. Somebody or something must manage the network resources, their differing capabilities, operating systems, benefits, and costs—and deliver to users the best computing performance at the lowest cost. Multiple levels of resource management are needed, including node, job, and network load management. Cray Research already provides sophisticated resource management facilities that monitor system and network usage and dynamically adjust workload distribution. As future standards emerge, resource management software will expand to automatically distribute jobs to the most appropriate system(s) on the network, account for the distributed resources used, and automatically bill the user online.

Security and availability: protecting your most important asset—data

In any computing environment, data is the most important asset. Open Supercomputing brings into focus issues such as data confidentiality (controlling access to data) and data integrity (protecting data from accidental or malicious corruption). Cray Research is the leader in providing confidentiality, integrity, and availability throughout high-performance, networked computing environments. Integrity and resiliency features such as automatic checkpoint/restart and interactive session protection enhance highly reliable Cray Research hardware to provide even greater system availability and throughput. We will continue to strive for nearly 100 percent processing availability to user applications as we pursue Open Supercomputing.

The dichotomy of "security" and "open" will continue. Finding solutions to the fundamental question "how can you operate in a fully open environment and be assured the data remains confidential and reliable?" is the challenge of the future. Cray Research will continue its leadership in the protection of mission-critical data and processing in a networked environment.

Data accessibility: high bandwidth access to huge amounts of data distributed throughout the network

Solving large-scale, numerically intensive problems requires efficient storage management and data access methods. Today Cray Research provides open systems solutions for distributed data management including ONC/NFS, OSF/DCE DFS, and hierarchical storage management facilities. In the diverse Open Supercomputing environment, users will require smarter hierarchical storage management software to provide automatic, transparent access to data stored on these different layered storage devices throughout the network. The ever-increasing quantities of data generated by today's high-speed processors must be accessible from anywhere on the network. Support for open disk technology with low cost per Mbyte is essential. Because time-to-solution is critical for Cray Research supercomputer users, we must provide transparent access to distributed data at ultrafast speeds. The supercomputer will require connection to networks of the highest possible bandwidth to match its processing speeds and input/output (I/O).

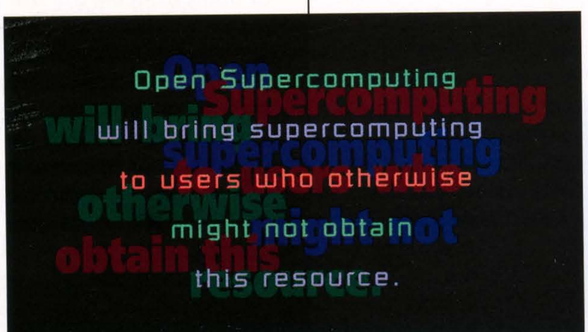

Open Supercomputing will bring supercomputing to users who otherwise might not obtain this resource. Where today 500 workstations may exist on a network, future networks will have thousands or more. In the automotive industry, for example, a car designer, part designer, manufacturing engineer, inventory specialist, purchaser, and pricing team—all those involved in bringing a product to market—will work from the same databases. The integration of technical and commercial computing will result in enormous amounts of data. Advanced technologies to manage and transparently access growing amounts of data distributed throughout the network will be essential. Accessing data by content, rather than location, will be important. Object-oriented databases will be the next level used to find data, requiring a stronger connection between databases and storage management technologies.

Cohesive environment: from desktop to supercomputer

From hardware and user interface perspectives, cohesiveness is an important part of our Open Supercomputing vision. For increased user productivity and better use of resources, we must provide users with common interfaces to all of the computer resources on the network, a common look and feel for all operating systems on the network, and, ultimately, one system view for all hardware resources. This single system image is the desktop system. In the Open Supercomputing world, users will not even know they are using a Cray Research system . . . they'll just know they got an answer, fast.

A cohesive environment is achieved by embracing standards or porting software and tools to other platforms. Cray Research is aggressively doing both. For example, our newest tools are designed to conform to the COSE Common Desktop Environment (CDE) standard; our CraySoft initiative has brought Cray Research software to various vendor platforms. For example, SunSoft recently announced a collaboration with CraySoft to deliver Cray Fortran 90 on Sun Solaris systems. From the microcomputer to the server system and supercomputer, and at all points in between, we intend to provide users with easy-to-use, transparent access to all kinds of systems on a heterogeneous network.

Interaction with computers is rapidly changing. Talking to your computer will be one interface of the future: "Computer show me. Computer tell me. Computer give me the solution." And there will be many more interface techniques to support. Users will interact with computers in their own intuitive ways, requiring the system to adapt to users, rather than users adapting to the system. As a provider of computational servers behind the scenes, Cray Research will interface and interoperate with these diverse interfacing methods.

Summary

Cray Research knows the importance of open computing and has provided Open Supercomputing systems for the past 10 years. Our vision for Open Supercomputing encompasses software technological advancements necessary for continuing our UNIX- and standards-based operating environments and making them even more powerful and more usable for a broader user base. We know some users require the high performance of our products, and we aim to provide these users access to our systems by using the network in the most transparent way. One important user segment is the commercial market, where many technologies traditional to the technical community (such as UNIX) are being embraced by organizations and users. As commercial organizations move to many technologies based in the technical world, Cray Research will be uniquely qualified to bring them high performance solutions.

For an extended version of this article, including information about the current status of Cray Research software, access http://www.cray.com/PUBLIC/product-info/sw/open.sc.vision.html with your favorite World Wide Web browser.

About the author

Kathy Nottingham is a software product manager at Cray Research. She has worked at the company for 10 years in various positions, including benchmarking, software training, and sales support. Her academic background includes a B.S. degree in computer science at St. Cloud State University.

Manchester moves into the future

Cindy Bielke, Cray Research, Inc.

When the Manchester Computer Centre (MCC) chose to replace its Amdahl 5890 mainframe with a 12-processor Cray Research CS6400 system last year, it extended a long tradition of computing excellence and its role as one of the United Kingdom's premier computer centers.

"We always have had leading-edge computer architectures, starting more than 40 years ago with systems developed by a team at the University of Manchester—the Mk 1 in the 1950s, followed by the Mercury and the Atlas," said professor Frank Sumner, the University's director of computing services. Sumner had a personal association with computing legends: noted computer pioneer Alan Turing taught him how to program.

The world's first installed computer went into operation at the University of Manchester in March 1951. It was used as a "national computer machine," although the lack of networking at that time limited its use primarily to University of Manchester researchers and others who traveled to the University. In 1972, the MCC became one of three National Supercomputer Centers in the U.K.;

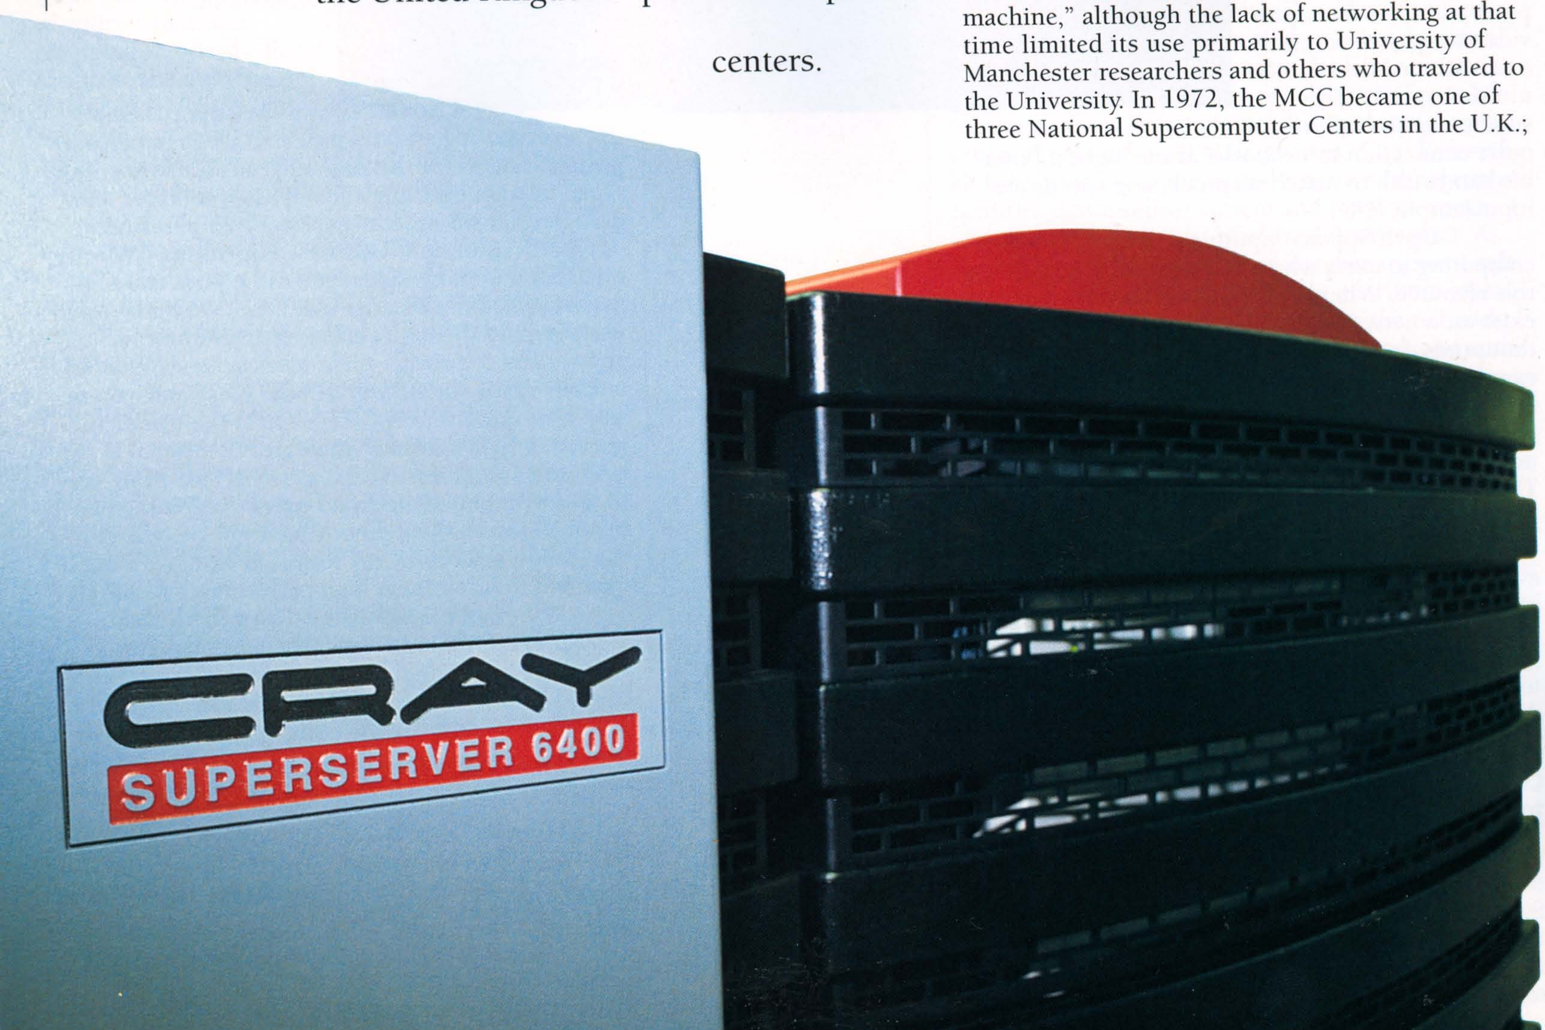

the other two include the London University Center and the Rutherford-Appleton Lab. It provides interactive service to more than 100 institutions and more than 6000 users.

More recently, the UK Joint Information Systems Committee (JISC)—which represents the Higher Education Funding Councils for Wales, Scotland, England, and Northern Ireland—commissioned the MCC to serve as a national database or datasets center available to all U.K. researchers. The distinction represents another important step toward the MCC's goal of becoming the premier database and supercomputing site in the United Kingdom.

Any system Sumner purchased in 1994 had to exceed not only the capabilities of the mainframe it replaced, it had to take the Centre to the next level of computing: scalable parallelism and client-server computing. He chose the CS6400 system after rigorous benchmarking and reliability testing. "Our benchmarks measured combined workload capacity, computational power, and response time. With 175 users accessing eight separate 2-Gbyte datasets, the CS6400's response time was almost instantaneous, a fraction of a second," he said.

Response time was about two times faster than other systems tested. The CS6400 system also met another acceptance criterion—three weeks of uninterrupted operation while benchmarks were running.

The MCC provides users across the country with online access to a wide range of large, complex databases, including International Monetary Fund financial statistics, all UK census data, detailed microcensus data, digital mapping, and scientific and bibliographic data. The average database size is 6 Gbytes and several are loaded concurrently.

The volume and sensitivity of the data, and the number of users, required both mainframe-class security and reliability not usually found in client-server systems. And users are demanding more modern, interactive interfaces to the data, such as the X Window System, and a wider range of analysis tools. Most of the existing tools were written in-house specifically for the mainframe.

The CS6400 system installed at the Manchester Computing Centre meets all these needs. Reliability, availability, and serviceability (RAS) features were designed in from the beginning, with subsystem redundancy to reduce downtime. The Solaris operating system, based on UNIX System V Release 4, provides access to more than 8000 applications. In addition, binary compatibility with Sun Microsystems workstations and servers means applications can be run where best suited—on the CS6400 system or the desktop. Conversion from the old system to the new was very straightforward, and the CS6400 system went into service almost a month ahead of schedule.

JISC set the budget limit and commissioned the Manchester Computing Center to purchase the best system for the database needs. According to Sumner, numerous bids were received, including "reasonable proposals" from IBM, Hewlett-Packard, and Sun Microsystems. Sun has worked closely with the University of Manchester, with which the Manchester Computing Centre is affiliated, and sponsors an annual Sun lecture. The open, scalable architectures of SPARC and Solaris were well known. And the CS6400 system was developed by Cray Research Business Systems Division in close collaboration with Sun Microsystems.

In the end, though, it all came down to trust. "We can trust Cray," said Sumner. "I know a lot about the quality of the designs, the engineering, the professionalism. Cray has very good engineers, attention to detail is great, as is support, but probably the main influence for this procurement has been Cray's continued belief in the highest possible bandwidth."

About the author

Cindy Bielke is a communications specialist in the Marketing Communications department at Cray Research.

Businesses count on Cray Research systems to gain competitive advantage

Q What do Atlantic Portfolio Analytics & Management (APAM) in Orlando, Churchill Insurance in Bromley, England, Citibank in New York, and Fuji Capital Markets Corporation, a subsidiary of Tokyo-based The Fuji Bank, Ltd. in New York have in common?

A They all count on Cray Research systems to meet their business objectives—increasing competitiveness, reducing time to market, and translating massive streams of data into sound business decisions.

Bobbi Hazard
Business Systems Division, Cray Research, Inc.

To maintain a competitive edge—in pricing complex mortgage derivatives at millisecond speed, running an online auto insurance policy application system with 2000 concurrent users, or pricing and hedging derivatives products through Monte Carlo simulations—businesses today require new technologies and infrastructures to keep up with increasing complexity and mounting time-to-market pressures. Faced with an explosion of data, commercial firms are moving away from sole reliance on desktop systems and proprietary mainframes toward a new breed of open systems that are easily integrated into unique and diverse commercial environments. As Shahin Khan, marketing programs manager of Cray Research Business Systems Division, explained, "Open systems are about customer choice and vendor independence. They provide more than one choice to the customer at every level of the computing hierarchy. With open systems, customers depend only on market forces and their ability to identify the mix of vendors who best solve their problems."

In the "data-intensive" market, there are many new variations on the old themes of competitiveness and time to market: real-time transaction processing; decision support solutions (DSS); online transaction processing (OLTP); relational database management systems (RDBMS); and object-oriented database management (OODBM). To respond to these new challenges, Cray Research is adapting its traditional expertise to the business side of supercomputing with an expanded line of products and solution strategies. Today, the company's products include entry level supercomputers, such as the CRAY EL98 system used by APAM, and a SPARC-based superserver, such as Citibank's CS6400 system.

Providing complete solutions to business

In addition to product expansion, Cray Research's involvement in the commercial market has meant a radical departure from the old supercomputing paradigm. Traditional customers—scientists and engineers using powerful vector-parallel supercomputers to solve large-scale problems—

needed to be computational experts and specialists in their fields. On the commercial side of computing, however, customers are looking for the complete solution to their unique business problems, without the steep learning curve. "Giving businesses a bottom-line advantage means that we completely adapt our expertise to provide a custom-made, complete solution to a particular environment that will integrate easily into the existing infrastructure," said Bob Ewald, Cray Research President and COO. "Customers want instant solutions to complex and time-sensitive business problems in real time or near-real time, at a reasonable cost."

Cray Research has a clear advantage in pursuing its strategy of open systems and customized solutions to fit unique business environments. Able to leverage more than 20 years of computing expertise, Cray Research already has the competencies demanded by today's sophisticated business environments: robust parallel hardware and software, network connectivity, a UNIX-based operating system, a large number of optimized applications, tools and standards, and service know-how.

Because of this time-proven expertise, Cray Research has enjoyed great momentum in penetrating the commercial marketplace—particularly in the financial industry. Today, Cray Research's UNIX-based parallel systems and SPARC-based products starting in the $150,000 to $400,000 price range are used by financial organizations such as Citibank, Merrill Lynch, Federal Home Loan Mortgage Association (Freddie Mac), APAM, and Fuji Capital Markets Corporation. Financial applications include collateralized mortgage obligations (CMOs), foreign exchange, option evaluation models, derivatives-based products, real-time analysis of firm-wide trading positions, and data management—at speeds typically 40 to 100 times faster than on desktop workstation systems.

APAM, for example, uses its CRAY EL98 system to conduct complex, timely analyses of all of the dimensions of risk for a particular customer, including how those risks interact. Taking advantage of the parallel and vector processing capabilities of the CRAY EL system, APAM has reduced the time to price complex mortgage derivatives from 2 minutes to between 20 and 30 ms. Wallclock runtimes on routine portfolio structuring and optimization analyses have been reduced from 5 hours to 12 to 13 minutes. APAM's Chief Investment Officer Jon Knight said: "We were not able to contain this kind of 'computational explosion' without employing supercomputing technology."

The complexity and time-sensitivity of today's financial applications were also a factor in Citibank's decision to acquire a Cray Research CS6400 system. Citibank's CS6400 system integrates perfectly with its existing network of Sun Microsystems workstations to support its derivatives business. Existing financial applications can be moved instantly to the Cray Research system; and risks can be monitored and evaluated on an ongoing basis. According to Shahin Khan, Citibank and Churchill Insurance typify the new, open systems customer: "Workstations on a network connected to a supercomputer running analytics make up the computing paradigm for the 1990s."

Initiatives and future strategies

Cray Research intends to be an important player in all areas of the emerging commercial market in which parallel processing systems are used for large-scale database analysis (decision support) and OLTP. This market is expected to grow to a multibillion-dollar level in this decade. Retail, credit card, telecommunications, banking, transportation, entertainment, and cable firms all routinely accumulate massive data volumes from business transactions and need a way not only to manage and store the data efficiently, but also to analyze it to identify buying patterns, fraud, and other information needed to make targeted marketing and business decisions. Running RDBMS on powerful UNIX-based parallel processing systems holds the promise of providing this decision support capability.

As part of its data-intensive market initiatives, Cray Research is forming strategic partnerships, including the recent acquisition of Savant Systems, a privately held consulting firm specializing in database and transaction processing business solutions, to advance into the commercial parallel processing DSS (decision support solutions) market. Renamed Cray Solutions, the former Savant Systems complements Cray Research's traditional competencies with system integration capabilities, particularly an in-depth knowledge of ORACLE and other relational database management systems.

In November 1994, Cray Research concluded an agreement with Informix Software to develop very large database scalability with Informix' Online Dynamic Server 7.10 and the CS6400 system. This will allow Cray Research to support customer databases in the hundreds of gigabytes range for the data warehousing, enterprise DSS, and high-end OLTP markets. Cray Research also is investigating the potential of the scientific data-intensive market, including geographical information systems (GIS), anomaly and fraud detection, and medical imaging.

To continue meeting the demands of increasingly complex and time-sensitive commercial markets, Cray Research is beginning to develop solutions for next-generation data management infrastructures. Among the technologies under consideration are clustered CS6400 systems, and, in the long-term, a follow-on scalable parallel processing (SPP) tool that will evolve out of that architecture. "Our technologies are perfectly adaptive to the business environment," said Bob Ewald. "Because of our core competencies in all aspects of supercomputing, we will lead our commercial customers to great success, just as we did on the science and engineering side of the business."

About the author

Bobbi Hazard is vice president of Cray Research's Business Systems Division, which provides high-performance computing solutions for data-intensive problems. As head of Business Systems, she oversees the operations of the firm's Cray Research Superservers and Cray-Solutions subsidiaries. She joined Cray Research in June 1993 as Project Coordinator, Hardware. Previously, she served as vice president of sales at nCube in Foster City, California.

"Workstations
on a network
connected to a
supercomputer
running
analytics
make up the
computing
paradigm
for the 1990s."

ANNOUNCING

THE CRAY ANIMATION THEATER RENDERING PACKAGE

The Cray Animation Theater (CAT) rendering and animation package brings the power of supercomputing to the world of visualization. CAT software opens worlds of new visual possibilities to engineers, scientists, and commercial animators. Now large, complex datasets that overwhelm even the most capable graphic workstations can be rendered and animated in just minutes using CAT software and Cray Research's industry-leading parallel processing technologies. The CAT package is the first fully supported rendering package developed by Cray Research for use on Cray Research computer systems.

CAT software demonstrates the value of the CRAY T3D massively parallel processing system's large, real memory for delivering high-speed rendering and animation capabilities. The software is scalable to any number of processing elements (PEs) on the CRAY T3D system. This scalability is made possible by the highly parallel design of the CAT renderers and very fast interprocessor communications in the CRAY T3D hardware. By enabling users to render hundred- or thousand-megabyte datasets conveniently in minutes, CAT and the CRAY T3D system bring within reach larger computational projects than you might have imagined possible.

Dynamic simulation of bird strike on fan blade. Each frame contains 573,600 triangles, representing 20 Mbytes of transient data, and was rendered in 1.4 seconds on 32 processing elements of a CRAY T3D system. Data provided by Allison Engine Company.

A new perspective on your bottom line

In engineering disciplines such as structural analysis and computational fluid dynamics, CAT is an ideal tool for animating the output from dynamic modeling projects. The models used in auto crash simulations, for example, can comprise millions of polygons, and the model geometries can change with each timestep. Animating such complex, dynamic models is impractical on workstations due to their lack of processing power, small real memories, and reliance on low-bandwidth network connections.

In workstation environments, users typically must disassemble large output datasets into subsets for rendering, then reassemble the subsets to recreate the original model. CAT software lets you analyze large problems and render the results in their entirety on one hardware platform. Keeping all of your computationally intensive work on the same platform eliminates workstation-performance and network-bandwidth constraints to deliver unprecedented job turnaround. The efficiency and capability gains made possible with CAT let engineers pursue significantly more complex problems with higher mesh resolutions and more degrees of freedom.

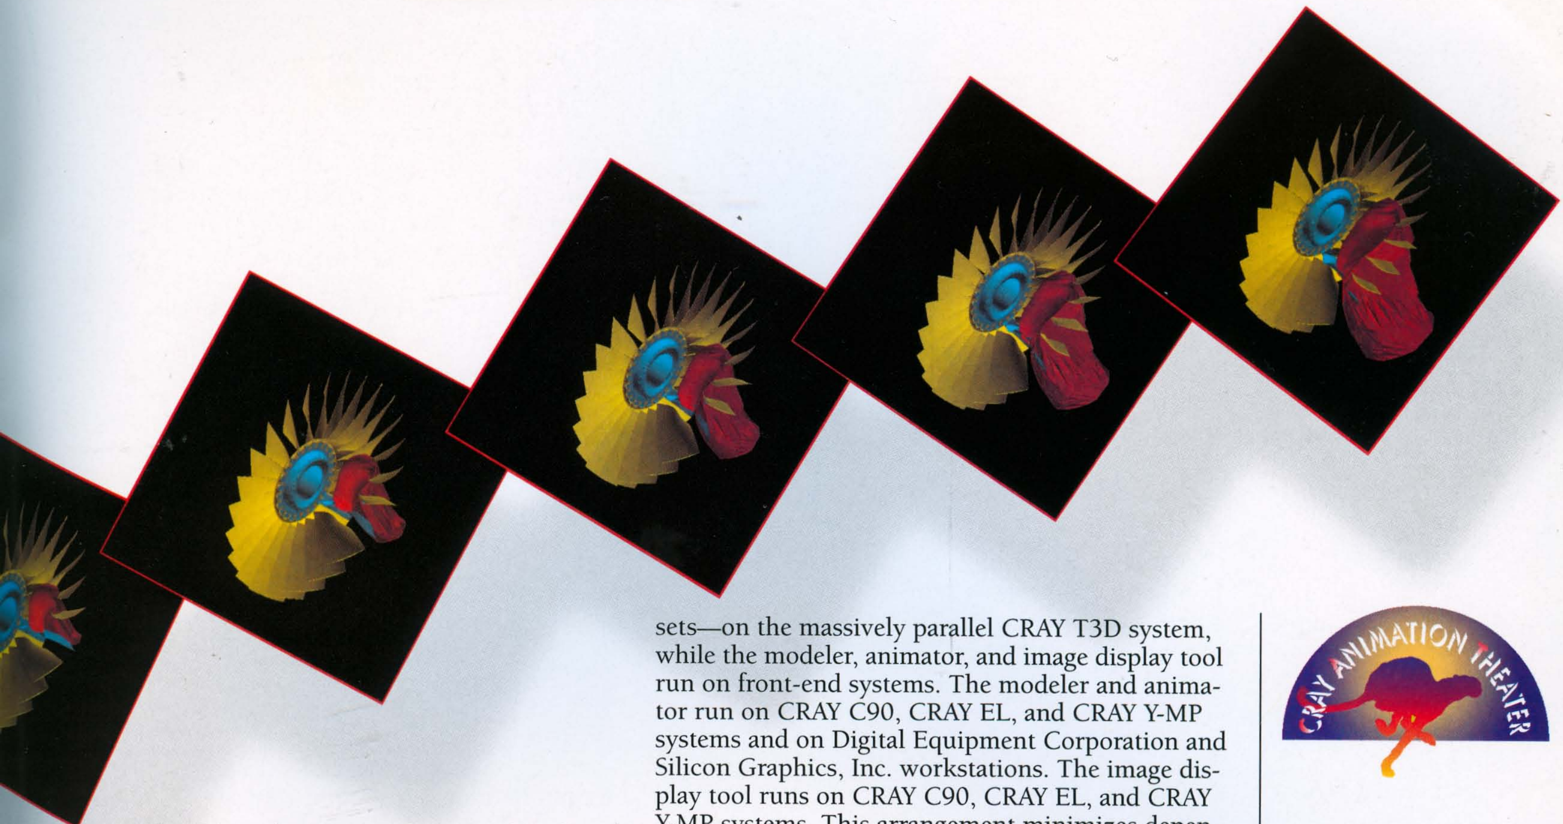

The CAT software's productivity benefits can be applied to an unlimited range of problems, not only in the engineering disciplines, but also in other applications that require fast rendering of large datasets. In chemical research, Earth sciences, astrophysics, and other natural sciences, CAT is the perfect rendering and animation tool for scientists using large computational models. CAT also offers commercial animators unsurpassed rendering capabilities in terms of time and complexity, collapsing production time to a fraction of that required by graphic workstations.

The CAT software package is an integrated toolset that includes the following components:

- ☐ A ray tracing renderer that produces photorealistic images using multiple light sources, texture mapping, and anti-aliasing. This tool is recommended for rendering images that include reflections, refractions, and translucency.
- ☐ A polygon renderer that provides simple and fast rendering capabilities. This tool is recommended for rendering very large datasets, such as those generated by large-scale scientific and engineering simulations. The polygon renderer is available as a library that can be linked to applications running on CRAY T3D systems.
- ☐ A modeler that lets you modify and manipulate objects conveniently through a graphical interface as you prepare animations. The renderers and animator can be invoked from the modeler.
- ☐ An animator that provides interactive capabilities for creating animation scripts.
- ☐ An image display tool that supports output to various media and display devices.

CAT software ensures efficient use of your computing resources by keeping the most computationally intensive work—the rendering of large datasets—on the massively parallel CRAY T3D system, while the modeler, animator, and image display tool run on front-end systems. The modeler and animator run on CRAY C90, CRAY EL, and CRAY Y-MP systems and on Digital Equipment Corporation and Silicon Graphics, Inc. workstations. The image display tool runs on CRAY C90, CRAY EL, and CRAY Y-MP systems. This arrangement minimizes dependence on network data transfers and helps keep desktop resources free for less computationally demanding work.

Since the earliest days of supercomputing, visualization has been an essential tool for taking full advantage of supercomputer power. But the workstations that scientists and engineers typically rely on to provide this capability have not kept pace with the size and complexity of supercomputer analyses. The CAT software package bridges the visualization performance gap by applying Cray Research MPP technology to data rendering and animation. The result is user productivity orders of magnitude beyond that possible with workstations alone.

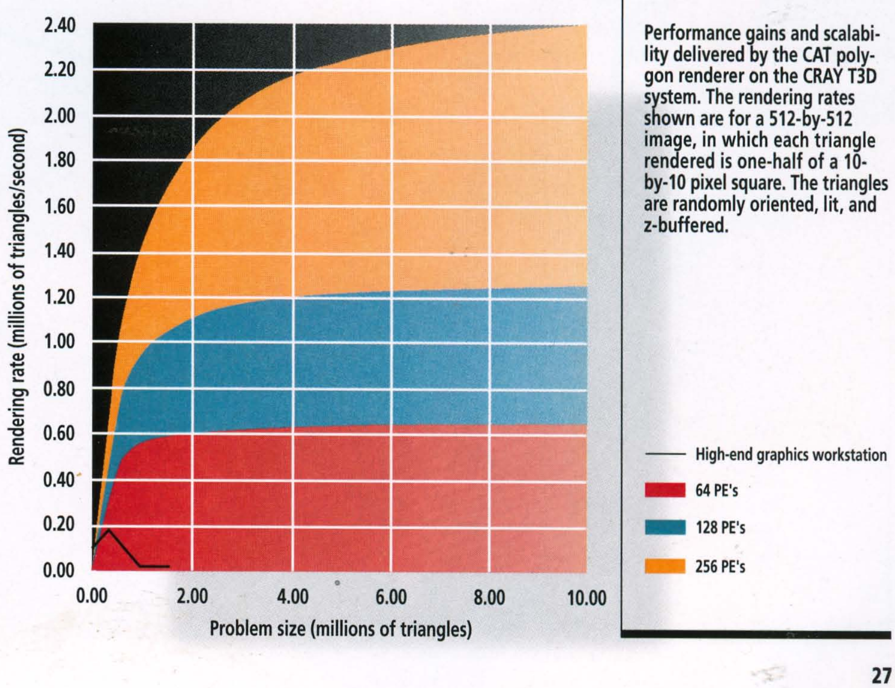

Performance gains and scalability delivered by the CAT polygon renderer on the CRAY T3D system. The rendering rates shown are for a 512-by-512 image, in which each triangle rendered is one-half of a 10-by-10 pixel square. The triangles are randomly oriented, lit, and z-buffered.

REDEFINING THE

Design of experiments, molding simulation software, and supercomputing can each dramatically reduce the cost of developing and testing expensive plastic injection molds. If these technologies were combined into a single cost-effective package, the plastic injection molding industry would have a revolutionary new predictive tool. Cray Molding Logic (CMLogic) from Cray Research is that tool. CMLogic simulates design-of-experiments (DOE) molding trials, using the leading molding applications, on Cray Research systems. It completely predicts mold performance to help you deliver high-quality products, on time, at the lowest possible cost.

According to John Ulicny, research engineer at General Motors, a large automotive body panel mold can cost up to $500,000 and may require up to 10 mold trials. "Considering that a single molding trial may cost as much as $40,000 and take up to six weeks of effort, using CMLogic to eliminate even one of those trials can create substantial savings," he said.

CMLogic completes the simulation solution

CMLogic uses classic DOE principles to help you apply injection molding simulation programs more productively than ever before. Industry standard simulation codes reliably predict a mold's performance for a single molding machine setup. A CMLogic session performs a designed experiment by running multiple setups on a Cray Research supercomputer. Integrating the results from those runs, CMLogic identifies a range of acceptable machine setups (the process window) that can produce your specified part. CMLogic gives you the complete simulation solution.

Try out molding with supercomputers

With CMLogic, you can predict a mold's performance and productivity long before actual molding trials begin. CMLogic gives you a complete understanding of how statistically significant molding parameters interact in an operating plastic injection mold. From filling the mold through final part warpage, a Cray Research system can rapidly analyze all factors to identify a process window for a given mold design. Now you can begin mold tryouts with complete confidence. CMLogic predicts acceptable operating ranges so that you can profit-

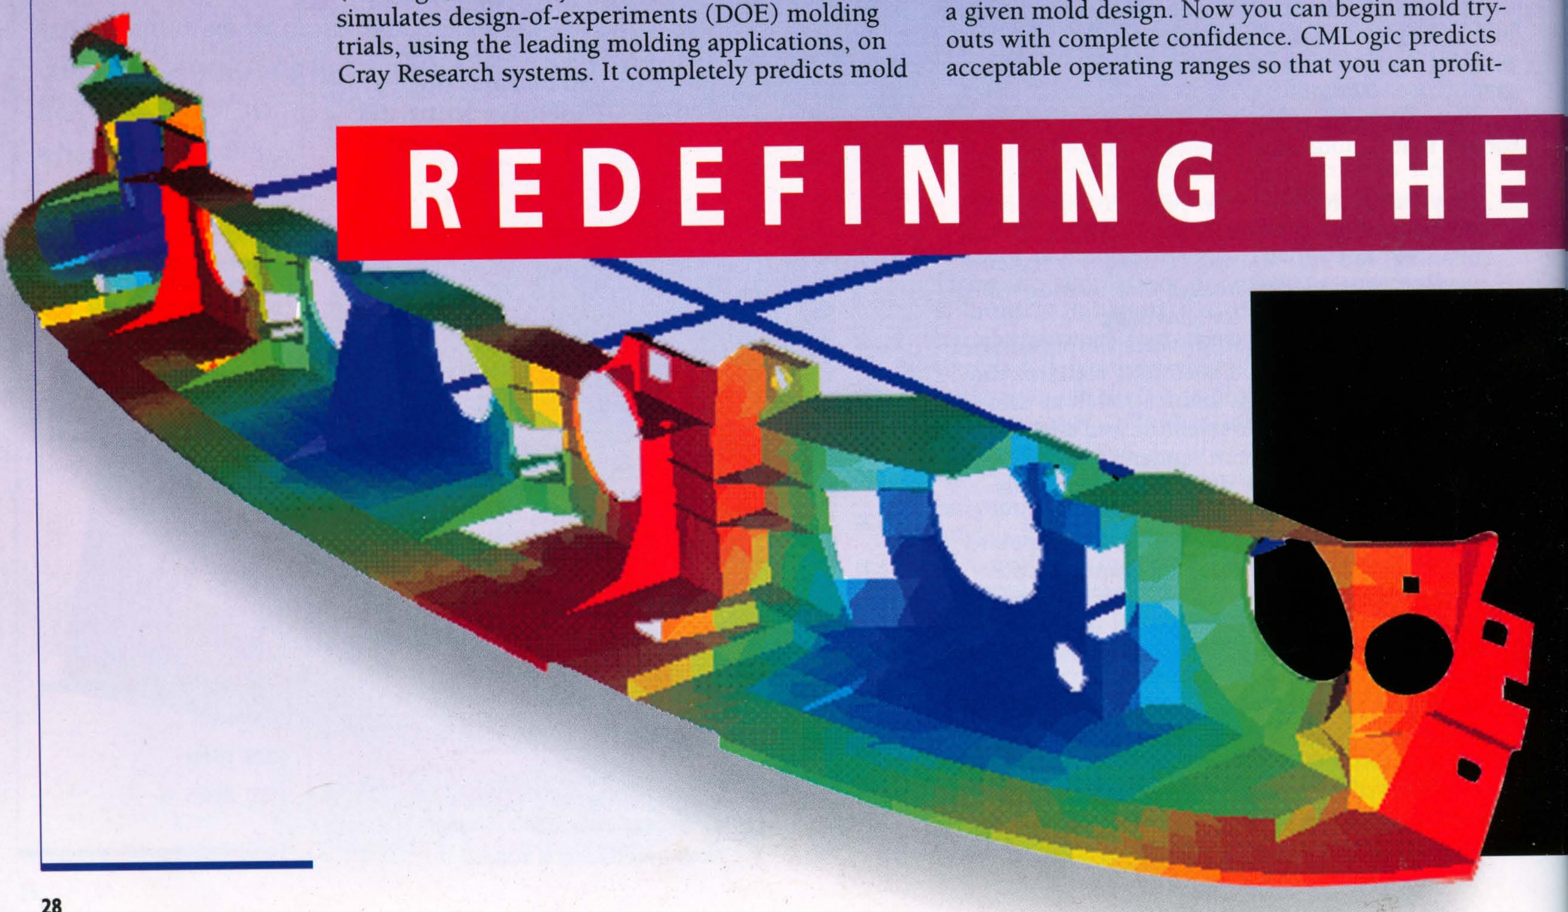

ably optimize machine cycle times or quickly resolve current production problems.

Find the process window

A process window describes the range of machine settings that can produce acceptable parts from a given mold design. CMLogic lets you vary six key injection molding process variables: cavity fill time, plastic melt temperature, mold surface temperature, cavity holding pressure, cavity holding time, and mold closed time.

If you choose to study only two variables, CMLogic will run nine simulations. A complete solution for all six variables requires just 45 jobs. You can rely on the power of Cray Research systems to deliver the results for these simulations quickly and profitably. Once the simulations are run, CMLogic builds a mathematical model that describes how the complex interdependencies between machine settings affect the final part dimensions. If a process window is too small to ensure reliable molding, you will know in time to consider design changes.

Explore the process window . . . interactively!

CMLogic brings you the world's first interactive molding simulator. Using CMLogic session results, you can simply twist knobs on our revolutionary simulator and try out a mold design while sitting at your workstation. CMLogic's unique molding simulator conclusively demonstrates the complex interdependencies of the molding process. How sensitive is shear rate to cavity fill time? How soon can you open the mold without affecting final part dimensions? The simulator uses the mathematical functions found by a CMLogic session to show you—instantly and graphically.

Generate compelling reports quickly

Designed for the real world of industrial manufacturing, CMLogic includes powerful project management and automatic report generation features, making it easy for you to prepare reports in Apple Computer's QuickTime electronic format. QuickTime presents visual information, such as mold filling animations, as a "slide show." The slide show can be viewed using a PC (running Microsoft Windows) or a Macintosh. The electronic format reduces your reports to a single 3.5-inch diskette. CMLogic reports include the following:

- ☐ Mold geometry visualizations and animations
- ☐ A description and visual representation of the process window
- ☐ Your engineering summary
- ☐ Project management information

Real-world management tools

Designed by experienced analysis engineers from the General Motors Research and Development Center, optimization experts, and Cray Research programmers, CMLogic's project management features let you easily manage and track multiple projects. CMLogic easily tracks vital project information such as the following:

- ☐ Contacts (names, phone numbers, and so on)
- ☐ Part names and numbers
- ☐ Critical project dates
- ☐ Engineering notes

For more information about CMLogic, call +1 612/683-3030 or send email to cmlogic@cray.com.

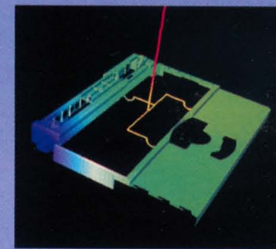

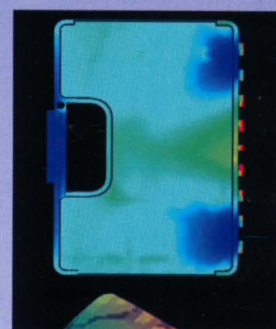

MOLDER'S TOOLKIT

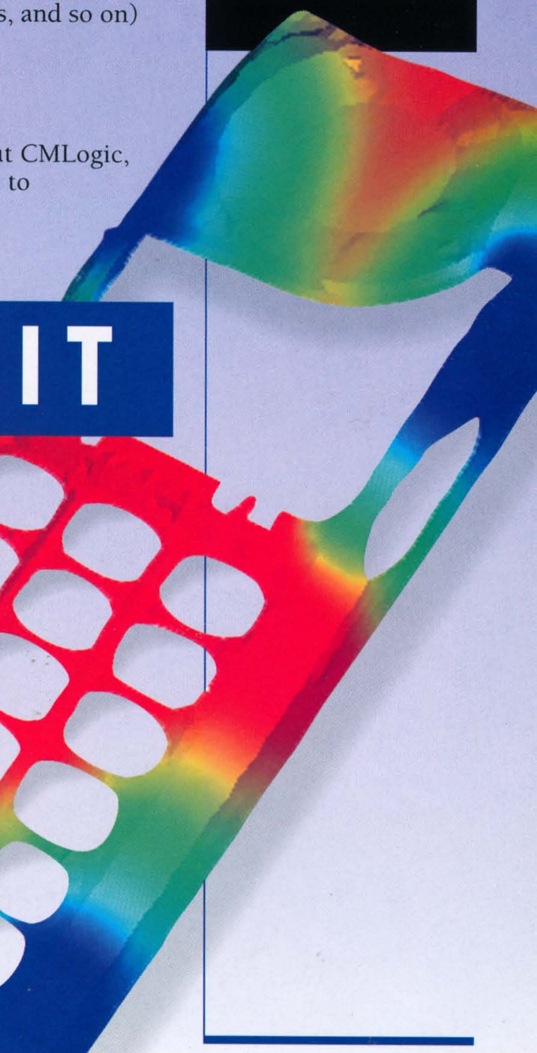

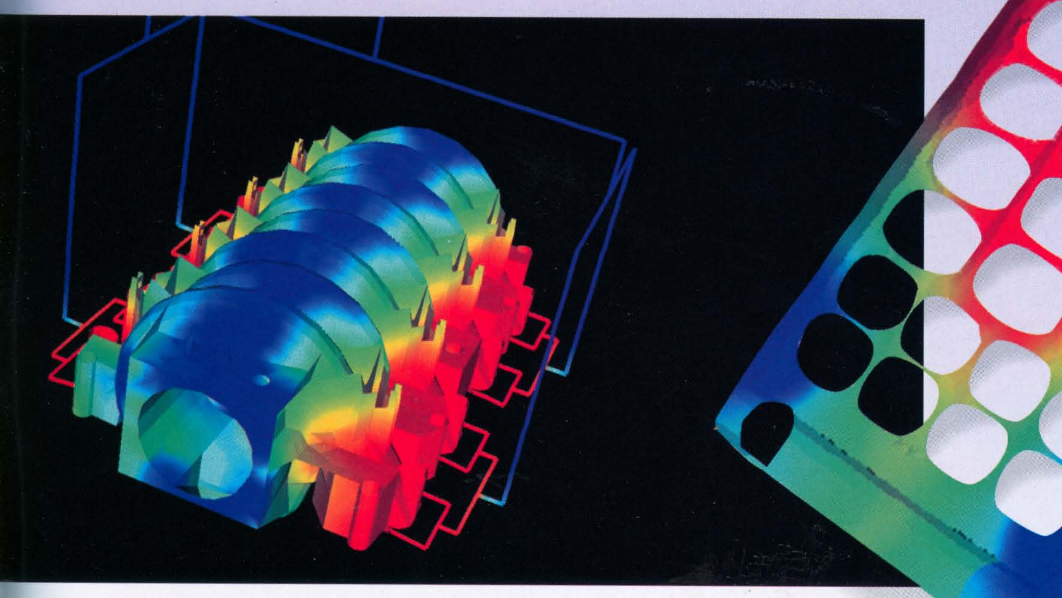

CRAY RESEARCH

Customer Service

MEETING THE CHALLENGES OF A CHANGING ENVIRONMENT

Kelly Jones, Cray Research, Inc.

Information technology (IT) has undergone many changes in the last decade and has become a critical success factor for virtually every major industry worldwide. High performance computing (HPC) continues to be a key segment of the overall IT industry and has seen perhaps some of the most significant changes, especially with the implementation of a variety of architectures. Cray Research has maintained its leadership position in the HPC market by both expanding its parallel-vector processing (PVP) product offerings and by introducing highly competitive symmetric multiprocessing (SMP) and massively parallel processing (MPP) products. This broad range of products, coupled with robust and feature-rich operating system software and an extensive base of applications, has enabled Cray Research to provide complete solutions to its customers.

In addition to offering customers a broad spectrum of HPC solutions, Cray Research provides a comprehensive, worldwide support structure. The Customer Service organization adds value by helping customers maximize the availability of their computing resources, providing efficient and effective information delivery, and offering expert advice and consultation.

Listening and responding to our customers

Strong partnerships with our customers are essential to long-term success. The Cray Research Customer Service division plays a key role in maintaining and growing those partnerships. The day-to-day interaction between our customers and our worldwide field organization provides real-time input for dealing effectively with operational issues and for keeping customers informed.

For the past several years Customer Service has used advisory boards as a means for focused interaction with small groups of customers. In addition to providing feedback on Cray Research products and services, board members offer their perspectives on the overall strategic direction of Customer Service.

"I am really impressed by the sincerity of Cray Research people in listening to their customers and searching for improvement. This separates Cray from the rest of the world."

John Shen
Chrysler Corporation

Customer Service is also active in a variety of customer-sponsored events. The Cray User Group (CUG) is probably the most visible and certainly the largest of these group meetings. We invest significant time and energy in working with CUG to ensure that there is maximum mutual benefit from these meetings.

Listening to our customers is not enough. We are acting on what we have heard. We are working to model our service delivery to match the changing environments of our customers. This includes not only how we respond to customer problems, but also how we develop and deliver information, help customers plan for and install new Cray Research system resources, provide support for users' work, and how we package and price our services.

Rapid response and problem resolution

Our highly trained service professionals maintain their expertise through extensive product training delivered through a variety of mechanisms, from traditional classroom courses to computer-based training (CBT) and state-of-the-art desktop videoconferencing.

The experience and expertise of our field service representatives are augmented by a comprehensive and efficient support structure. Support centers in several locations throughout the world let us quickly channel the resources necessary to solve a customer's problem.

The US-based Customer Service Call Center (CSCC) is the focal point. It provides "one call" access for both customers and Cray Research personnel to our extensive support infrastructure. Field-experienced technical specialists trained on the full range of Cray Research hardware and software products staff the CSCC. These specialists quickly evaluate a situation and provide immediate solutions, or, if necessary, escalate the problem to the appropriate area for further analysis and resolution. This entire network of analysts, engineers, and information specialists lets us resolve problems quickly.

Keeping customers informed

Customers have told us that they want not only rapid response to problems, but also more access to critical support information. We have responded through CRInform and the Cray Research Service Bulletin (CRSB).

CRInform is an online menu-driven information and problem report service available to all Cray Research customers. This tool allows customers to report problems and request technical assistance, learn about problem solutions, order software, receive email notification of service information (including SPR status changes and Software Field Notices), and view service information (such as the CRSB and publications and training catalogs). Customers can obtain a CRInform account by signing a CRInform Program agreement, available from their Cray Research service representative.

The CRSB is a monthly newsletter for system administrators and user support personnel. The CRSB provides information on Cray Research software products and services, including information about new software releases and features, new software tools, usage, tuning and performance enhancement, migrating sites, training and documentation, and symptoms of software problems and how to get solutions.

Installation planning and preparation

Customer Service plays a key role in ensuring successful Cray Research system installations. From preshipment preparations through installation, startup, and the on-site quality assurance check, our support personnel make certain everything goes smoothly.

In 1994, our field service organization in Germany planned and coordinated the entire site preparation and installation process for a CRAY C916 system at Deutsches Klimarechenzentrum (DKRZ) in Hamburg. The new system replaced DKRZ's CRAY-2S and CRAY Y-MP4E systems, and its weight required that a special structure be designed and built to distribute the load safely. The power and cooling support systems also needed extensive modification to accommodate the CRAY C916 system's additional computing power. Knowing from previous experience that such a complex system exchange is not an easy task, Wolfgang Sell, General Manager of DKRZ, asked Cray Research to handle the entire replacement procedure in the framework of the purchase contract. The Cray Research Germany team functioned as prime contractor for the entire operation, hiring all of the necessary subcontractors and craftsmen, including an architect, a stress analyst, a cooling engineer, an electrical firm, a cooling firm, a steel-girder construction firm, and a crane firm.

This installation was a bit more unusual than most because the CRAY C916 system was being installed on the 15th floor of the geophysics department building and had to be hoisted by crane on a specially built platform attached to a window on the 15th floor. All of this (removing the old systems, rebuilding the entire environment, and installing and powering up the new system) was done successfully in only two weeks.

"Cray Service has always been and continues to be the standard which we hold up to other vendors for responsiveness and personnel excellence. Cray is making commendable efforts at bringing the costs for this excellence into line with present-day data center budgets."

Dan Drobnis
San Diego
Supercomputer Center

Cray Research Customer Service coordinated the installation of a CRAY C916 system into DKRZ's 15th-floor computing center.

Sell congratulated Werner Nipken, the Cray Research engineer in charge of the project at DKRZ, and his crew for the careful and cooperative preparation and the excellent work during the short interrupt period, even shorter than had been anticipated in the early planning stages. "This short interrupt time became very important for us, because we decided to replace both the CRAY-2S and the CRAY Y-MP4E computer systems at the same time in order to save money," Sell said.

Performance-centered systems

A challenge of business in the 1990s and beyond centers on how to make people more productive in their jobs in an environment in which technology and the information that it produces expands faster than ever before. The Software Information Services department, based on feedback received through CUG and our advisory boards, initiated the development of an Integrated Performance Support System (IPSS). The objective of such a system is to integrate information, assistance, and learning opportunities and make them available to the user on demand, in the workplace, and at the moment of need.

Several components of this system, such as CrayDoc (Cray Research's online documentation browser) and context-sensitive help (included in tools such as xbrowse, the MPP Apprentice tool, and the TotalView debugger), are available today. Industry acceptance of the Common Desktop Environment (CDE) model for integrated help is expected to expand the number of products that will offer context-sensitive help messages. In addition, Cray Research is developing CBT courses such as CRAY J90 and CRAY EL System Administration Training (ELAD-CBT) using a new CBT environment called Cray*Tutor*.

These components are not the final answer. Cray Research plans to advance along a continuum from loosely integrated elements to a system that is completely transparent to the user. We are committed to making these advances based on collaborative partnerships with our customers.

Delivering solutions—anytime, anywhere

In addition to providing world class maintenance and support for Cray Research systems and our customers, Customer Service has taken two additional proactive steps to help customers find solutions to a variety of strategic and operational challenges: acquiring the Minnesota Supercomputer Center, Inc. (MSCI) and forming a new department called Professional Services.

The acquisition of MSCI in 1994 positions Cray Research to take advantage of the many opportunities in the computer-related network services market. MSCI's commercial offerings give customers the power and expertise to seek solutions to industrial and scientific problems, from small projects to the largest computational problems around. These leading-edge supercomputer resources are readily accessible via a "distance insensitive" network, making them available at the user's desktop. For more information about MSCI, contact your Cray Research account manager or John Sell (+1612/337-3450 or jsell@msc.edu).

"Cray Research provides unmatched service. The competent service people rarely need support beyond their own skills, but when they do, it is available very quickly and easily. They do not have to go through a bureaucratic support organization."

Peter Gray
European Centre for Medium-Range Weather Forecasts

Professional Services offers customers alternate solutions to many operational needs. Customers can outsource system administration and operation procedures, system performance analyses and capacity planning, network management, UNICOS upgrades, help desk services, and more. By tapping into Cray Research experience and expertise, customers can focus on their own core competencies and maximize the value of their investment with Cray Research. For more information about Professional Services, contact your Cray Research account manager or Phil Hernick (+1612/683-5747 or profserv@cray.com).

Value and commitment

As our customers' environments change so do their expectations for service. For the past several years we have used vehicles such as the Cray Research Customer Satisfaction Survey, our advisory boards, CUG, and customer visits to understand how customers' expectations for service are changing. What we have heard is that customers want increased availability, knowledgeable help when they need it, personal attention, more choices, and lower prices.

The Cray Research Customer Service organization has worked hard to meet these expectations. We have structured ourselves to ensure rapid response to problems and have established measurements that let us track the availability of systems from the customer's perspective. We have enhanced the accessibility of critical service information for customers and are working to integrate today's product documentation and help facilities such that the use of Cray Research products and tools becomes intuitive for users. We have strengthened our already strong partnerships with our customers by learning more about their business environments and sharing a sense of responsibility in the accomplishment of their objectives. We continue to explore ways in which we can reduce our maintenance prices and have made significant changes in the kinds of services we offer.

Satisfied customers are the lifeblood of any successful business. Our commitment to ourselves is to improve the level of high customer satisfaction Cray Research Customer Service has enjoyed historically. Our commitment to our customers is to continually meet or exceed expectations by listening to their needs, acting on what we hear, and delivering outstanding value at a competitive price.

About the author

Kelly R. Jones is director of the Professional Services and Decision Support group in the Cray Research Service Division. He joined Cray Research in 1985 as an on-site engineer and has held positions as manager of Hardware Publications and Training as well as director of Quality Management. He received a B.S. degree in engineering technology from Kansas State University. Prior to joining Cray Research, Kelly worked as a technical support engineer for Seismograph Service Corporation in Tulsa, Oklahoma.

CORPORATE REGISTER

Cray Research announces first installation in Thailand

The National Electronics and Computer Technology Center (NECTEC) of Thailand has installed a CRAY EL98 air-cooled, compact supercomputer in Bangkok. The system is the first Cray Research supercomputer installed in Thailand and has four central processing units and 512 Mbytes of main memory (soon to be expanded to 1 Gbyte).

NECTEC will use the CRAY EL98 system for computational chemistry, computational physics, mathematical modeling, and computer science. NECTEC plans to provide access to the system through Thaisarn, a local academic computer network that connects a number of universities in the country and that provides a gateway to the Internet.

NECTEC provides support for research, development, and engineering in universities, government, and the private sector; provides technical services to the private sector in production development; invests in the development and acquisition of technology; and disseminates information regarding electronic and computer technologies and related issues.

UK's Churchill Insurance installs two CRAY CS6400 systems for mission-critical OLTP

Leading United Kingdom insurer **Churchill Insurance** recently installed two 13-processor CS6400 systems at Churchill's Bromley headquarters facility, where they will be used as high-end, mission-critical Oracle online transaction processing (OLTP) servers. Churchill is one of Europe's largest OLTP sites.

Amdahl Corporation served as the prime contractor. The Amdahl competitive win followed an 18-month procurement process in which vendors were asked to run a demanding benchmark simulating the insurer's Oracle-based online auto insurance policy application system with 2,000 concurrent users, according to Churchill chief information officer Carl Ricketts. "We tested a wide range of UNIX multiprocessor systems. The CS6400 system was the only one that passed the test."

Ricketts said added attractions of the CS6400 systems are that each can be expanded to 64 processors and multiple systems could be clustered to handle future growth. "Our business is increasing—we expect to have 750 online users next year. The CS6400 systems can grow with us."

The SPARC-based symmetric multiprocessing (SMP) systems will run Sun's Solaris operating environment with Cray Research and Amdahl extensions. Amdahl is providing systems management software and systems integration.

ARCO (Atlantic Richfield Company) has installed a Cray Research CS6400 system at its Exploration and Production Technology Center, Plano, Texas, for seismic data processing to develop computer models of subsurface terrain. At ARCO, the first customer to acquire a CS6400 system in the petroleum industry, engineers will use the system to process seismic data on the order of hundreds of gigabytes to help the company improve the efficiency of finding and extracting oil. The CS6400 system will be networked to Sun and SGI workstations and an IBM 3090 mainframe.

Fuji Capital Markets Corporation, a subsidiary of Tokyo-based The Fuji Bank, Ltd., ordered a CRAY EL94 minisupercomputer system for installation at Fuji Capital Markets Corporation's U.S. headquarters in the World Trade Center, New York. The CRAY EL94 system will be used to run Monte Carlo-based applications developed in-house by Fuji Capital Markets Corporation's technology group. The system will enable the firm to enhance pricing and hedging of its current and new derivatives products, as well as provide real-time turnaround on pricing requests from Fuji Capital Markets Corporation's customers.

James Kennedy, Fuji Capital Markets Corporation's managing director, said this new Cray Research technology will allow the firm to run hundreds of thousands of "yield curve" scenarios for more accurate statistical analyses of its derivatives products. The system also will be used for "stress testing," a complex analysis used in defining Fuji Capital Markets Corporation's portfolio risks. The new Cray Research system will be integrated with the 125 workstations on Fuji Capital Markets Corporation's wide area network that links the firm's London and Hong Kong offices. Kennedy said traders at these international offices also will use the Cray Research system.

Other Cray Research customers in the financial industry include Atlantic Portfolio Analytics and Management (APAM), Citicorp's Citibank, the Federal Home Loan Mortgage Association ("Freddie Mac"), and Merrill Lynch.

Wayne State University, Detroit, Michigan, ordered a CRAY J916 system for use by its chemistry department for computational chemistry applications such as *Gaussian*, MOPAC, and VENUS to study the electronic structure of molecules; reaction paths and potential energy surfaces for organic reactions; and the molecular dynamics on such surfaces. Installation of this system will make Wayne State the first university in Michigan to acquire its own Cray Research supercomputer.

The CRAY J916 supercomputer will be dedicated specifically to developing new multidisciplinary computing techniques. By combining the processing capabilities of the CRAY J916 system with the improvements in molecular orbital software, it is possible to combine two molecular modeling disciplines—classical mechanics and quantum mechanics—into one integrated solution.

The Department of Defense High Performance Computing Distributed Center at Eglin Air Force Base, Florida, ordered a 128-processor CRAY T3D system. Coupled with the center's existing CRAY Y-MP and CRAY Y-MP EL supercomputers, the new massively parallel processing system will support the research, development, test, and evaluation mission carried out by the center on behalf of the U.S. Air Force and other Department of Defense agencies. The new supercomputer will have 1024 million words of central memory and will be networked to a 1.2 trillion byte StorageTek storage silo and to a visualization laboratory with high-performance engineering and graphics/visualization workstations.

The CRAY T3D system will let Eglin's researchers distribute applications among multiple processors to achieve very fast simulations of highly complex situations. They will use computational fluid dynamics and computational mechanics applications that have been reprogrammed to take advantage of the MPP environment of the CRAY T3D system. In the simulations, which serve the broad goal of delivering weapons in the safest, most effective way possible, aircraft and weapon surfaces are divided into small 3-D grids. This process is repeated at each instant of time along the flight path of the weapon for each grid cell. The highly accurate picture of the effect of aerodynamics on munitions and aircraft made possible with the simulation are vital to effective design.

Cray Solutions provides full services to commercial parallel processing market

In a move that supports Cray Research's pursuit of the growing commercial market for parallel processing computer systems, the company acquired Savant Systems, Dallas, Texas, a privately held consulting firm specializing in database and transaction-processing business solutions.

The former Savant Systems is now a wholly owned Cray Research subsidiary renamed Cray Solutions; it will remain based in Dallas. Cray Solutions will focus on its consulting business, which includes business systems analysis, design and implementation, large-scale project management, and long-term technical support (outsourcing) services. Cray Solutions also will collaborate in Cray Research's overall pursuit of the commercial parallel processing market for decision-support solutions. As part of Cray Research, Cray Solutions will use its industry expertise, along with its knowledge of ORACLE and other relational database management systems (RDBMS), application products, and development tools in crafting leading-edge business solutions.

Clients includes BP America, Compaq Computer Corp., Martin Marietta, Metropolitan Transit Authority of Houston, Oracle Corp., Pennzoil Exploration and Production Co., and the North Carolina Supercomputer Center.

BioNumerik, Cray Research partner to speed development of computer designed drugs to combat cancer, heart disease

Pharmaceutical research and development firm BioNumerik Pharmaceuticals, Inc. (BioNumerik) will enter a beta site agreement for new supercomputer systems from Cray Research, which are designed to speed the development of powerful new drugs to combat cancer and heart disease. Under this unique agreement—intended to use parallel supercomputer power to design and develop new substances with optimal, highly specific disease-fighting capabilities—BioNumerik will operate Cray Research's newest systems in a collaboration to optimize proprietary BioNumerik software for use on these supercomputers.

The goals of the agreement are to test the operation of BioNumerik's advanced software and Cray Research's new parallel supercomputer systems in the development of an innovative, time-saving approach to drug development. The Cray Research systems are expected to help accelerate BioNumerik's ongoing research in this area. One of the most important goals will be the introduction of new pharmaceutical agents into clinical practice.

BioNumerik currently is engaged in preclinical development of several anticancer drugs aimed at treating the most common and deadly adult cancers, including tumors of the lung, breast, colon, prostate, pancreas, and melanoma. The company also is developing two new cardiovascular drugs: a new agent to elevate HDL, or "good cholesterol," and a second product that reduces LDL, or "bad cholesterol." These therapies are being developed under an exclusive licensing agreement with the Southwest Foundation for Biomedical Research.

BioNumerik chairman and CEO Fred Hausheer is one of the world's foremost pharmaceutical researchers. He has served on the faculty of the Johns Hopkins Oncology Center and the University of Texas Health Science Center, is coholder of many patents that involve new therapeutic agents, was the 1992 recipient of the IEEE Forefronts in Large-Scale Computation Award, and was a finalist for the 1993 Computerworld Smithsonian award for Breakthrough Computational Science.

BioNumerik Pharmaceuticals was formed in 1992 to develop novel, breakthrough drugs for the treatment of cancer and other diseases, using a unique methodology that involves supercomputer technology and proprietary molecular simulation software.

Cray Research and Informix Software announce joint effort to support very large databases (VLDB)

Cray Research and Informix Software will work together to provide very large database (VLDB) scalability with Informix-OnLine Dynamic Server 7.10 and Cray Research's CS6400 system.

Informix quality assurance staff will use a CS6400 system installed at Informix headquarters in Menlo Park, California, to test OnLine Dynamic Server 7.10 and its scalability enhancements. Planned testing will encompass single and multiple instances of OnLine Server querying very large databases. The goal of the joint effort is to support customer databases in the hundreds-of-gigabytes range.

SCIENTIFIC's Linda software is now available for the CRAY T3D system

SCIENTIFIC Computing Associates' Linda software, a product for parallelizing and running applications on parallel computers, is now available for the CRAY T3D MPP system. Linda has been optimized for use on the CRAY T3D system in both C and Fortran—existing Linda-based applications can now be run on the CRAY T3D system with little or no modification, and new highly parallel applications can be generated efficiently for the CRAY T3D system from users' sequential codes.

Linda is an important addition to the programming and applications support tools available for the CRAY T3D system. Many production applications have been parallelized using SCIENTIFIC's Linda, especially for workstation networks and clusters. These applications can be moved to the CRAY T3D system when customers need faster solution times or want to keep the workstation systems free for less demanding work. For more information contact SCIENTIFIC Computing Associates at 203/777-7442 or email software@sca.com.

Cray Research and ANSYS, Inc. team to expand parallel processing in ANSYS

Cray Research and ANSYS, Inc. are jointly expanding parallel processing capabilities of the ANSYS finite element analysis code on Cray Research's parallel vector supercomputers. (The ANSYS program is used for structural analysis simulations in product design, testing, and manufacturability. The software package is popular among industrial users that compute large-scale modeling and simulation problems on products such as automobiles and components, aircraft, and other manufactured goods.)

ANSYS software has been available on Cray Research supercomputers for more than 15 years. The improved parallel features are anticipated to be released in ANSYS, revision 5.2, for the entire line of Cray Research supercomputers.

Cray Research and ANSYS officials said that the parallel processing performance of the ANSYS program on Cray Research supercomputers, especially for very large problems, has been internationally recognized in recent years. Last year the ANSYS/Cray Research solution was nominated as a finalist for the prestigious Gordon Bell Prize and also cited in Europe at SuParCup '93 in an international parallel processing performance competition. This recognition was given for sustained performance of 6 GFLOPS on a CRAY C916 system for an analysis of a new design for NASA's Space Shuttle Liquid Oxygen Pump.

The parallel enhancements in the ANSYS program are based on a newly designed database and code restructuring. The changes allow simultaneous computation of individual elements—which may number in the hundreds of thousands for detailed model simulations—on multiple processors such as those in Cray Research supercomputers.

CraySoft delivers new storage management tools for SPARC platforms

CraySoft, Cray Research's open systems software unit, announced the early access availability of the Open Storage Manager (OSM) product for SPARC systems under the terms of a joint licensing and distribution agreement with Legent Corporation. CraySoft will market, sell, and distribute Legent Corporation's OSM product to the high-end science, engineering, and industrial marketplace. OSM, a family of open hierarchical storage management (HSM) solutions for UNIX systems, provides automated storage management across networks of SPARC-based clients and servers. Hierarchical storage management improves the efficiency of managing data storage by providing a transparent, scalable method of storing and retrieving large amounts of data and infrequently used files on less costly storage devices.

The CraySoft OSM product consists of the OSM Client and OSM Server. The OSM product complements Cray Research's own HSM solution, called Data Migration Facility (DMF) software, which is available only on the company's supercomputing systems that run the UNICOS operating system. OSM software provides the same scalability on SPARC that DMF provides for Cray Research's multiprocessing supercomputers.

CraySoft OSM is available for early access release on all Solaris SPARC platforms, including the Cray Research CS6400 system. General availability is expected in early 1995. CraySoft products are available through Cray Research sales offices worldwide, and Cray Research's high-performance hotline, 1-800-BUY-CRAY, as well as email at craysoft@cray.com. Information about CraySoft and other Cray Research products is available on the World-Wide Web at http://www.cray.com/.

Cray Research releases HEXAR 1.1 mesh-generation software

HEXAR 1.1, an enhanced version of the mesh-generation software product, is now available. The HEXAR software can save product developers months of time and hundreds of thousands of dollars by automatically creating 3-D computer models from raw computer-aided design (CAD) data in minutes. Traditionally, it has taken months to handcraft models, or meshes, needed for computer analysis for product designs, medical phenomena, and other complex research.

HEXAR 1.1 has a 70 percent success rate in creating complete hexahedral (six-sided "brick") models on the first attempt, without user intervention, considerably better than the 40 percent automatic mesh-generation rate many people hoped for in a product of this type. The new version also features improved mesh-element quality, a 50 percent reduction in computer memory requirements, and a 50 percent improvement in speed.

As with the original version, HEXAR 1.1 is suited for the automotive, aerospace, chemical, civil engineering, casting and molding, materials and plastics, electronics, semiconductor, and biomedical industries. HEXAR software produces meshes for the most widely used computer-aided engineering (CAE) formats, including PATRAN, IDEAS, and EnSight. Because of this, the HEXAR product works with nearly all third-party engineering analysis software packages used today, including ABAQUS, ANSYS, NASTRAN, FIDAP, STAR-CD, FLOW-3D, PROCAST, and MSC/EMAS.

SCIENTIFIC's Linda software is now available for the CRAY T3D system

SCIENTIFIC Computing Associates' Linda software, a product for parallelizing and running applications on parallel computers, is now available for the CRAY T3D MPP system. Linda has been optimized for use on the CRAY T3D system in both C and Fortran—existing Linda-based applications can now be run on the CRAY T3D system with little or no modification, and new highly parallel applications can be generated efficiently for the CRAY T3D system from users' sequential codes.

Linda is an important addition to the programming and applications support tools available for the CRAY T3D system. Many production applications have been parallelized using SCIENTIFIC's Linda, especially for workstation networks and clusters. These applications can be moved to the CRAY T3D system when customers need faster solution times or want to keep the workstation systems free for less demanding work. For more information contact SCIENTIFIC Computing Associates at 203/777-7442 or email software@sca.com.

Cray Research and ANSYS, Inc. team to expand parallel processing in ANSYS

Cray Research and ANSYS, Inc. are jointly expanding parallel processing capabilities of the ANSYS finite element analysis code on Cray Research's parallel vector supercomputers. (The ANSYS program is used for structural analysis simulations in product design, testing, and manufacturability. The software package is popular among industrial users that compute large-scale modeling and simulation problems on products such as automobiles and components, aircraft, and other manufactured goods.)

ANSYS software has been available on Cray Research supercomputers for more than 15 years. The improved parallel features are anticipated to be released in ANSYS, revision 5.2, for the entire line of Cray Research supercomputers.

Cray Research and ANSYS officials said that the parallel processing performance of the ANSYS program on Cray Research supercomputers, especially for very large problems, has been internationally recognized in recent years. Last year the ANSYS/Cray Research solution was nominated as a finalist for the prestigious Gordon Bell Prize and also cited in Europe at SuParCup '93 in an international parallel processing performance competition. This recognition was given for sustained performance of 6 GFLOPS on a CRAY C916 system for an analysis of a new design for NASA's Space Shuttle Liquid Oxygen Pump.

The parallel enhancements in the ANSYS program are based on a newly designed database and code restructuring. The changes allow simultaneous computation of individual elements—which may number in the hundreds of thousands for detailed model simulations—on multiple processors such as those in Cray Research supercomputers.

CraySoft delivers new storage management tools for SPARC platforms

CraySoft, Cray Research's open systems software unit, announced the early access availability of the Open Storage Manager (OSM) product for SPARC systems under the terms of a joint licensing and distribution agreement with Legent Corporation. CraySoft will market, sell, and distribute Legent Corporation's OSM product to the high-end science, engineering, and industrial marketplace. OSM, a family of open hierarchical storage management (HSM) solutions for UNIX systems, provides automated storage management across networks of SPARC-based clients and servers. Hierarchical storage management improves the efficiency of managing data storage by providing a transparent, scalable method of storing and retrieving large amounts of data and infrequently used files on less costly storage devices.

The CraySoft OSM product consists of the OSM Client and OSM Server. The OSM product complements Cray Research's own HSM solution, called Data Migration Facility (DMF) software, which is available only on the company's supercomputing systems that run the UNICOS operating system. OSM software provides the same scalability on SPARC that DMF provides for Cray Research's multiprocessing supercomputers.

CraySoft OSM is available for early access release on all Solaris SPARC platforms, including the Cray Research CS6400 system. General availability is expected in early 1995. CraySoft products are available through Cray Research sales offices worldwide, and Cray Research's high-performance hotline, 1-800-BUY-CRAY, as well as email at craysoft@cray.com. Information about CraySoft and other Cray Research products is available on the World-Wide Web at http://www.cray.com/.

Cray Research releases HEXAR 1.1 mesh-generation software

HEXAR 1.1, an enhanced version of the mesh-generation software product, is now available. The HEXAR software can save product developers months of time and hundreds of thousands of dollars by automatically creating 3-D computer models from raw computer-aided design (CAD) data in minutes. Traditionally, it has taken months to hand-craft models, or meshes, needed for computer analysis for product designs, medical phenomena, and other complex research.

HEXAR 1.1 has a 70 percent success rate in creating complete hexahedral (six-sided "brick") models on the first attempt, without user intervention, considerably better than the 40 percent automatic mesh-generation rate many people hoped for in a product of this type. The new version also features improved mesh-element quality, a 50 percent reduction in computer memory requirements, and a 50 percent improvement in speed.

As with the original version, HEXAR 1.1 is suited for the automotive, aerospace, chemical, civil engineering, casting and molding, materials and plastics, electronics, semiconductor, and biomedical industries. HEXAR software produces meshes for the most widely used computer-aided engineering (CAE) formats, including PATRAN, IDEAS, and EnSight. Because of this, the HEXAR product works with nearly all third-party engineering analysis software packages used today, including ABAQUS, ANSYS, NASTRAN, FIDAP, STAR-CD, FLOW-3D, PROCAST, and MSC/EMAS.

Doomsday at a distance

Contributed by Michael Schneider, Pittsburgh Supercomputing Center

Fireworks on Jupiter

All telescopes pointed to Jupiter last summer as comet Shoemaker-Levy 9 crashed into the jumbo planet. While people day-by-day followed the news of Shoemaker-Levy 9's fiery death, a basic scientific question entered public consciousness: Could a similar large comet or asteroid someday come hurtling toward Earth? The short answer is yes. It apparently happened 65 million years ago, with disastrous consequences for the dinosaurs, and the odds are, say astronomers, it will happen again within the next 100 million years. What are the chances human life could survive such a calamity? The event gave scientists an unprecedented chance to observe doomsday from a safe distance and learn from it.

"It was a unique event," said astrophysicist Mordecai-Mark Mac Low."It's the first time we've been able to predict a large planetary impact and then observe it, and it will probably be the only time in our lifetimes that an impact this large occurs." In preparation for the comet crash, Mac Low, a postdoctoral researcher at the University of Chicago, used the CRAY C90 system at Pittsburgh to forecast what would happen. "Basically, I was trying to predict the results of the impact so that observers could plan for the event. I looked at things like how bright the flash of the explosion would be and how much material from beneath the Jovian clouds would be lifted above them where it can be observed."

Results from analysis so far suggest that Mac Low's simulations were consistent with what actually happened on Jupiter. Using the CRAY C90 system and ZEUS, a program developed at the Laboratory for Computational Astrophysics of the National Center for Supercomputing Applications, Mac Low ran simulations that showed the comet would penetrate less deeply and explode more violently than other models predicted. Perhaps the most important result, at least in terms of forecasting the effect on Earth impacts, is that the simulations agree with a mathematical theory called "the pancake model." This model assumes that after the aerodynamic force from the comet's impact into the atmosphere overcomes its material strength, the comet flattens like a pancake, which greatly increases drag—essentially stopping the comet in its tracks.

Forecasting a big splash

In collaboration with Kevin Zahnle of NASA Ames Research Center, Mac Low ran three types of simulations: (1) a comet fragment entering Jupiter's atmosphere until it exploded, (2) the initial fireball from the explosion, and (3) what happened in the atmosphere after the initial fireball. The researchers used ZEUS for modeling the gas dynamics of astrophysical phenomena such as the violent shock waves from supernovae.

For the last six years, Mac Low has used ZEUS to study interstellar gas dynamics, and he realized he could apply the same method with relatively minor changes to simulate a comet crashing into Jupiter's atmosphere—basically by shifting the scale from light years to kilometers. "The physics," said Mac Low, "is virtually identical. It's only the details of the composition of the atmosphere that change, and of course the length scales, time scales, and density changed—by 20 orders of magnitude in some cases."

Their simulations predicted that the flash from the explosion would last about one minute, with about as much brightness as the sunlit side of Jupiter. Unfortunately for Earth observers, the comet crashed into the back side of Jupiter. Mac Low's calculations suggested, nevertheless, that the fireball would be bright enough to be seen from the NASA spacecraft Galileo or with Earth telescopes as a reflection off one of Jupiter's moons.

The strongest prediction from the simulations, the one of which Mac Low was most confident, had to do with how deep the comet would dive into Jupiter's atmosphere before exploding. Other models showed it going in hundreds of kilometers, so that its energy is absorbed relatively slowly—what one researcher called a "soft catch." Mac Low's results showed, however, that impact with Jupiter's atmosphere would rip the comet apart more quickly and violently, with a fierce explosion, after penetrating only about 110 km below the cloud tops. The resulting plume of superheated debris, according to Mac Low's study, would shoot hundreds of kilometers above Jupiter's layered clouds, giving astronomers a good chance to observe the aftereffects and, in the process, learn more than we know now about the composition of Jupiter's atmosphere. Inferences from observational data indicate that this prediction was essentially accurate.

Mac Low's simulations used a computational grid finer than the other models, suggesting that his results more closely approximated reality. As a check on this surmise, Mac Low ran his code at much lower resolution and got a result similar to the other models. "At low resolution, we got one result," said Mac Low, "and at high resolution, we got another, and as we continued increasing resolution, the result stayed the same." The high degree of detail in Mac Low's study was made possible by the CRAY C90 system.

Agreeing with the pancake model

The most important result of Mac Low's study is his finding that the numerical simulations agree well with the pancake model. The pancake model can be used to predict what will happen from comet and asteroid impacts on Earth and other planets.

"One of the scientific issues we're hoping to get a handle on in terms of Earth impactors," said Mac Low, "is how big a rock do we need to worry about? One of the motives for modeling this impact is to see if we can do reasonably accurate predictions. If we can, we can start talking about how well the Earth's atmosphere protects us."

This 640 x 640 pixel image was rendered using the ray tracer from the Cray Animation Theater (CAT) rendering and animation software package. The image took 10 minutes and 12 seconds to render on 64 PEs of a 256-PE CRAY T3D system. Final rendering and modeling was by Keith Fredericks; Rob Krawiec modeled the dart board, Nancy Rowe modeled the darts, Jim Johnson modeled the numbers, and Jim Morgan created the CAT logo.