CRAY CHANNELS

SPRING 1991 · A CRAY RESEARCH, INC., PUBLICATION

Aerospace Applications

Announcing new hardware and software products

CRAY CHANNELS

In this issue

On December 17, 1903, the Wright brothers flew their airplane into history. Their longest flight that day lasted 59 seconds and covered 852 feet. Sixty-six years later, within the lifetimes of some observers, astronauts walked on the moon and returned safely to Earth. The technological advancements in air and space travel during the past century have been astonishing. But as the vehicles for such travel have become more complex, so have the challenges associated with further advancements.

This issue of CRAY CHANNELS features several articles on computational challenges in the aerospace industry — challenges that require the power of Cray Research computer systems. These complex problems are diverse, ranging from the development of new algorithms for compressor design to predicting the effects of bird strikes. This issue also introduces new Cray Research hardware and software products: the CRAY Y-MP8E and CRAY Y-MP8I supercomputers; version 6.0 of the UNICOS operating system; and UniChem, a supercomputing environment for computational chemistry.

Advances in aerospace technology are becoming increasingly impressive as we near the end of the century. The shape of aerospace travel during the next century will depend on the availability of tools, such as Cray Research computer systems, that reduce the time and cost of new systems development. Large-scale numerical modeling is an essential methodology for aerospace design and engineering, and Cray Research computer systems are the aerospace industry's most cost-effective large-scale modeling tools.

Features

28

Departments

CRAY CHANNELS is a quarterly publication of the Cray Research, Inc., Marketing Communications Department, Tina M. Bonetti, Director. It is intended for users of Cray Research computer systems and others interested in the company and its products. Please mail feature story ideas, news items, and Gallery submissions to CRAY CHANNELS at Cray Research, Inc., 1440 Northland Drive, Mendota Heights, Minnesota 55120. Subscription inquiries and address changes should be sent Attention Dept. D.

Volume 13, Number 1

Editorial staff
Ken Jopp, editor
Mary Miller, associate editor
Vicky Frank, contributing writer
Rich Brueckner, contributing writer

Design and production
Barbara Cahlander
Eric Hanson
James Morgan
Cynthia Rykken

On the cover: A variety of jet engines from Rolls-Royce plc represents the fruits of engine design and engineering using Cray Research computer systems.

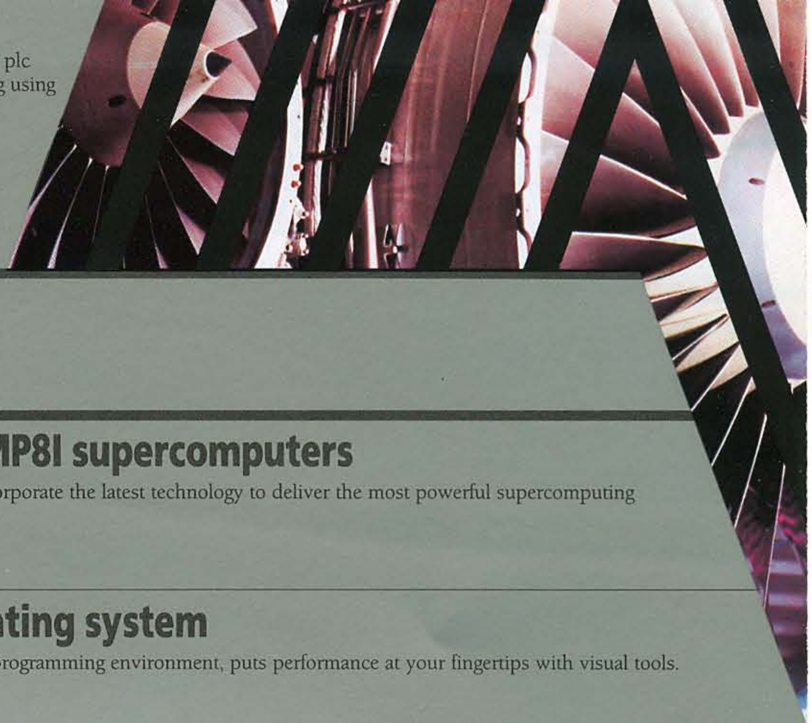

You are challenged to solve problems that stand in the way of a better future. Your success depends on having the best possible tool to bridge the gap between inspiration and innovation.

The CRAY Y-MP8 Supercomputer Systems

The best gets better

To help scientists and engineers bridge the gap between inspiration and innovation, Cray Research has introduced two powerful additions to the CRAY Y-MP product line: the CRAY Y-MP8E system, the most powerful supercomputer to date, and the CRAY Y-MP8I system, which provides comparable performance to previous CRAY Y-MP8 systems at a significantly lower price. These new supercomputers offer the highest sustained performance available for a variety of applications.

"The CRAY Y-MP series has been our most successful product series yet," said John A. Rollwagen, chairman and chief executive officer of Cray Research. "The products we are announcing today combine the strength of this series with new technologies, innovative packaging, and attractive new prices. The result is a series of systems that will solve larger problems faster, and provide users with increased reliability and productivity."

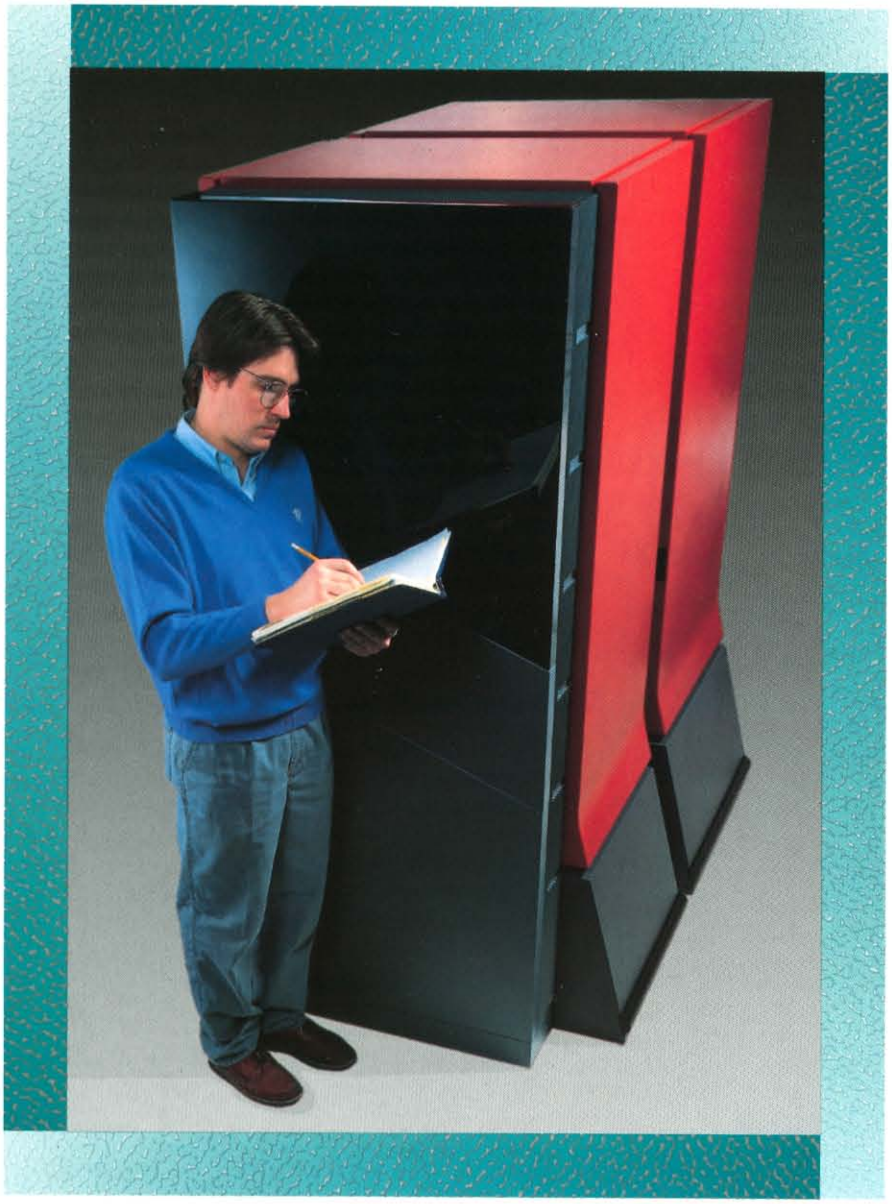

Rollwagen added that these new Cray Research systems are the heart of a total supercomputing solution. "We offer more than just fast hardware," he said. "Working closely with end users, we have developed a system of hardware, software, and networking capabilities that work in concert. This lets users focus on their research instead of the supercomputer's requirements."

The best gets better

The new CRAY Y-MP8 supercomputers set new standards for high-performance computing with outstanding functionality, reliability, and I/O capabilities, all at a lower cost than previous CRAY Y-MP8 systems. With a wide range of options and new technologies, the new CRAY Y-MP8 systems provide an outstanding level of price, performance, and productivity.

With all new I/O and SSD technology, the CRAY Y-MP8 systems offer unmatched performance for real applications. The CRAY Y-MP8 models use a balanced system architecture with large, fast memories matched and tuned with up to eight central processing units working in parallel. This architecture provides sustained performance of over two GFLOPS on a variety of applications. In fact, more codes and applications run at GFLOPS speed on CRAY Y-MP systems than on any other systems. With this performance at their fingertips, users can tackle problems previously thought intractable.

A tradition of performance and reliability

The original CRAY Y-MP8 computer system, with its multiple processors and vectorization capabilities, shattered all performance records. The new CRAY Y-MP8 models build on these strengths with expanded I/O and secondary storage capacity, innovative packaging, and an overall increase in integration at the component level.

The most powerful Cray Research system ever — the CRAY Y-MP8E

The CRAY Y-MP8E system is the most powerful, general-purpose scientific supercomputer available. With up to eight processors, the CRAY Y-MP8E system offers the largest central memory, SSD capacity, and I/O capacity ever available on CRAY Y-MP supercomputers.

The CRAY Y-MP8E system offers an unprecedented level of computing power and throughput. With up to 256 million 64-bit words (256 Mwords) of directly addressable central memory and 2 billion words (2 Gwords) of SSD capacity, the CRAY Y-MP8E system allows users to solve previously intractable problems.

To use its computational capacity efficiently, a high-speed supercomputer requires a vast amount of input/output data. To fulfill this requirement, the CRAY Y-MP8E system offers the most powerful and versatile I/O capabilities in the industry. The IOS comprises one to eight I/O clusters (IOCs), that support high-speed DD-60 drives. This allows more data to be accessed at faster rates than previous systems. Each IOC supports up to 16 channel adapters for disk, tape, and communications connections.

The CRAY Y-MP8E system offers an optional integrated SSD solid-state storage device. The SSD provides very high speed secondary memory of up to 2,048 Mwords (2 Gwords) at a lower cost per Mword than with previous SSDs. With support for up to four 1000-Mbyte/sec channels to the SSD, the CRAY Y-MP8E system provides rapid access to this massive storage capacity, which allows users to tackle larger problems.

High-end supercomputing made more affordable — the CRAY Y-MP8I

The CRAY Y-MP8I system reduces the cost of full-scale supercomputing by providing up to eight CPUs, an I/O subsystem (IOS), and an optional SSD in the same cabinet. By offering the performance capability of previous CRAY Y-MP8 systems at a significantly lower price, the CRAY Y-MP8I system gives customers more computing for their money. This outstanding price/performance ratio makes high-end supercomputing available to a broader range of users.

The CRAY Y-MP8I system offers powerful supercomputing capabilities with up to 128 Mwords of central memory. The CRAY Y-MP8I requires less floor space, less support equipment, and costs less to operate and maintain than previous CRAY Y-MP8 models. Its cost effectiveness is further enhanced with support for low cost/Mbyte DD-60 and DD-61 disk drives.

The CRAY Y-MP8I integrated IOS provides improved performance and reliability. The IOS includes up to four I/O clusters (IOCs), which allow more data to be accessed at faster rates than with previous CRAY Y-MP systems. Each IOC supports up to 16 channel adapters for connection to disk storage units, tape units, and communications products.

The CRAY Y-MP8I system can be configured with an optional integrated SSD solid-state storage device. The SSD provides very-high-speed secondary memory of 256 or 512 Mwords at a lower cost per Mword. This large storage capacity allows users to address larger problems that require extensive I/O and out-of-memory solution techniques.

The CRAY Y-MP8I computer system offers practical upgradability, allowing it to grow with users' needs. The standard CRAY Y-MP8I system includes four CPUs, 64 Mwords of central memory, and one IOC. Upgrade options include additional central processing units, IOCs, central memory, and an SSD.

All new I/O technology

The CRAY Y-MP8 models use an all new IOS that provides improved performance, versatility, and reliability. The IOS is an integral part of the CRAY Y-MP8 design, acting as the mainframe's data distribution point. The IOS allows high-speed communication between the central memory of the CRAY Y-MP8 system and peripherals such as disk drives.

To increase the CRAY Y-MP8 production workload capacity, the new I/O technology offers increased bandwidth that allows users to connect to more peripheral devices and perform more simultaneous activities. The IOS consists of up to four (CRAY Y-MP8I system) or eight (CRAY Y-MP8E system) I/O clusters, each with up to 16 channel adapters.

The new I/O technology provides customers with a flexible framework that can grow with their I/O and peripheral needs. The standard configuration of the CRAY Y-MP8 models includes one I/O cluster with eight channel adapters. Additional I/O clusters and channel adapters can be configured easily in the field.

All new SSD technology

The CRAY Y-MP8 models offer all new SSD technology that offers reliable storage capacity at a lower cost per Mword. Using VLSI chips and increased system integration, the SSD is available with up to 512 Mwords (CRAY Y-MP8I) or 2 Gwords (CRAY Y-MP8E). These large storage capacities allow users to solve larger problems.

The optional SSD solid-state storage device is a very fast random-access device used for large prestaged or intermediate files that are generated and manipulated repeatedly by user programs. The SSD also can be used for swapping programs and for holding commonly accessed libraries and other frequently accessed programs, thereby improving overall system performance.

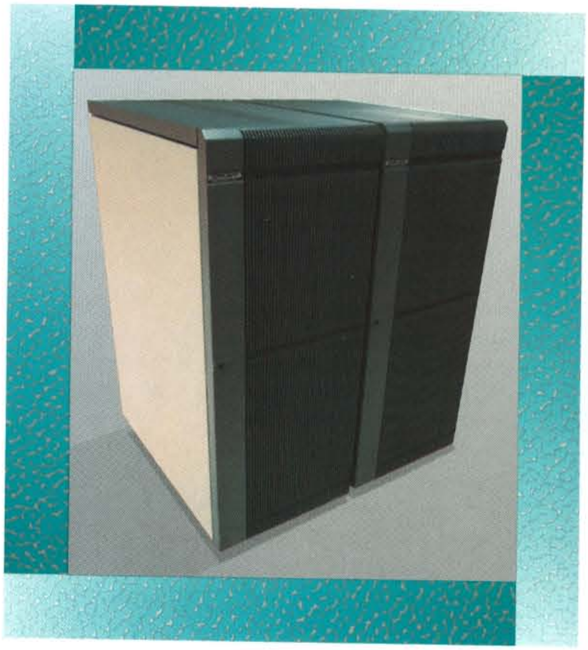

Two DE-60 disk cabinets, each containing up to eight DD-60 or DD-61 disk drives.

CRAY Y-MP8E system configurations

Model	CPUs	Central memory (Mwords)	IOCs	Optional SSD (Mwords)
CRAY Y-MP8E/4128	4	128	1-8	512, 1024, or 2048
CRAY Y-MP8E/4256	4	256	1-8	512, 1024, or 2048
CRAY Y-MP8E/8128	8	128	1-8	512, 1024, or 2048
CRAY Y-MP8E/8256	8	256	1-8	512, 1024, or 2048

CRAY Y-MP8I system configurations

Model	CPUs	Central memory (Mwords)	IOCs	Optional SSD (Mwords)
CRAY Y-MP8I/464	4	64	1-4	256 or 512
CRAY Y-MP8I/4128	4	128	1-4	256 or 512
CRAY Y-MP8I/864	8	64	1-4	256 or 512
CRAY Y-MP8I/8128	8	128	1-4	256 or 512

CRAY Y-MP8 computer system configuration options.

The CRAY Y-MP8 systems communicate with the SSD through one to four 1000-Mbyte/sec channels. The SSD is connected to the IOS through one to eight 200-Mbyte/sec channels. These connections enable data to be transferred directly between an IOS and the SSD without passing through central memory, thereby increasing overall performance.

Disk Drives

Cray Research offers fast, reliable mass storage devices that provide large storage capacities in a small space. The CRAY Y-MP8 models support all current Cray Research disk storage devices including the DD-60 and DD-61 disk storage units.

The DD-60 disk drive offers outstanding performance and large storage capacities when matched with the I/O capability of the CRAY Y-MP8 models. With the capability to support over a terabyte of disk storage, the CRAY Y-MP8E system gives users high-speed access to over six times more data than was possible with previous CRAY Y-MP8 systems. The DD-60 is a 24-Mbyte/sec disk drive with a sustained transfer rate of 20 Mbytes/sec. Up to eight DD-60 disk drives can be connected to each DCA-2 channel adapter in the IOS.

The DD-61 disk drive delivers large storage capacities at a low cost. The DD-61 provides users with access to large amounts of data using highly reliable, 8-inch disk technology that gives the DD-61 a lower cost per Mbyte, a small footprint, and low power consumption. The DD-61 is a 3-Mbyte/sec disk drive with a sustained transfer rate of 2.6 Mbytes/sec. Up to eight DD-61 disk drives can be connected to each DCA-2 channel adapter in the integrated CRAY Y-MP8 IOS.

Performance-oriented, feature-rich software

Because hardware performance depends heavily on software performance, Cray Research emphasizes performance-oriented, feature-rich software products that enhance the hardware capabilities of the CRAY Y-MP8 models. Cray Research provides the most complete body of system software available on any supercomputer system. In addition, Cray Research systems are unsurpassed in their ability to connect to computer hardware from other vendors.

Cray Research also supports leading-edge applications for nearly every scientific and engineering discipline including over 600 of the most widely used third-party application programs. Applications are available for industries such as aerospace, automotive, chemistry, electronics, energy, and petroleum. These applications accelerate product development, increase productivity, and solve basic research problems.

The CRAY Y-MP8 supercomputers — the best gets better

The new CRAY Y-MP8 systems are part of a total supercomputing solution that bridges the gap between inspiration and innovation. For more information on the CRAY Y-MP8E and CRAY Y-MP8I supercomputer systems, contact your nearest Cray Research office.

Announcing version 6.0 of the UNICOS operating system

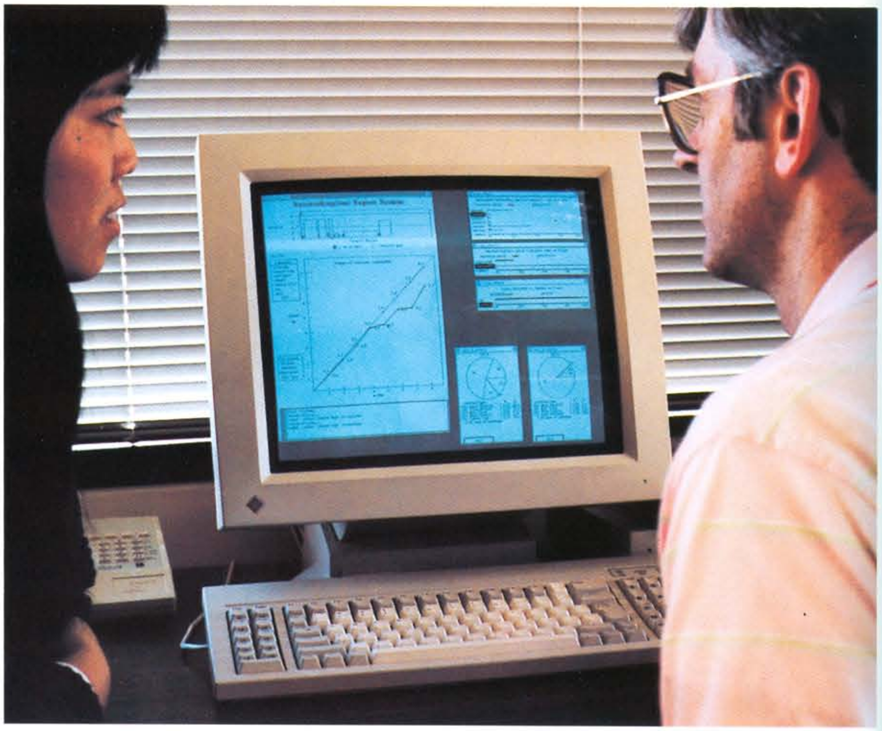

Cray Research's tradition of innovative hardware design is matched by its leadership in software development. Now, to enhance the performance of the world's fastest computer systems, Cray Research offers release 6.0 of the UNICOS operating system, the most advanced and feature-rich UNIX programming environment available on any computer system.

UNICOS provides a powerful and portable software environment for both batch and interactive processing. Introduced in 1986, UNICOS was the first UNIX-based operating system for general-purpose supercomputing. UNICOS includes Berkeley extensions to the UNIX System V standard as well as numerous enhancements critical to the high-performance production environment. Designed to conform to the POSIX 1003.1 standard, release 6.0 of UNICOS sets new standards for performance, usability, maintainability, and connectivity.

The ultimate performance operating system

UNICOS combines all the inherent strengths of UNIX, such as portability and a powerful user interface, with production-oriented features such as high-performance I/O, multiprocessing support, ANSI/IBM tape support, resource control, job scheduling, and batch processing. These features and others have added the robustness necessary to optimize performance in a production setting.

UNICOS 6.0 includes the following performance features:

- ☐ The UNICOS Storage System — the world's first high-performance, UNIX-based file server for supercomputers. It provides transparent data access, file access control, system administration, and automated storage management capabilities.
- ☐ Multitasked support for Level 2 and Level 3 BLAS kernels, the new LAPACK algebra package, and new one-, two-, and three-dimensional fast Fourier transform routines.
- ☐ Flexible File I/O (FFIO) that now includes memory resident and secondary data segment layers. These achieve tremendous performance gains by allowing the user to declare all or part of a file to reside in memory or on an SSD solid-state storage device as a secondary data segment. In addition, the ability to perform automatic conversion of IEEE numeric formats has been added. FFIO has been enhanced to permit users to process non-native files through standard Fortran I/O by identifying them with the *assign* command. FFIO now can now process files from a wide variety of vendors in this manner. It also is designed to be extended easily by Cray Research or a customer to handle new formats.
- ☐ Many enhancements that optimize parallel performance in batch production without user intervention. Improved scheduling allows multitasked jobs to run with the same or better throughput than non-multitasked jobs.

A new level of usability

In addition to the X Window System tools, UNICOS 6.0 includes several features that enhance usability:

- ☐ A standardized message system allows users to call up detailed explanations of many error messages and enables them to control the format and language of received messages.
- ☐ Docview online documentation viewing utility allows users to view information on line, write it to a file, and print it.
- ☐ The Korn Shell programming environment offers C shell funtionality with Bourne Shell compatibility.
- ☐ The GNU EMACS text editor provides much greater functionality than the standard UNIX editors.

Robust, maintainable code

A major portion of the time spent developing the UNICOS 6.0 release was devoted to achieving the most stable, robust UNICOS ever released. To reach this goal, UNICOS 6.0 was tested aggressively in demanding production environments with multiple daily tape mounts, round-the-clock production jobs, and wide-ranging production networks. UNICOS 6.0 also is easier to install and maintain than previous releases due to the following enhancements:

- ☐ UNICOS 6.0 reduces the amount of non-production time required to support configurations or reconfigurations. Installability has been improved with the UNICOS installation manager, a highly automated, menu-driven system that provides a common look and feel for all major components of the installation process. In addition, new pre-compiled, relocatable UNICOS modules greatly reduce system compile/build time.
- ☐ (SMARTE) An online maintenance environment automatically reports errors and provides an X Window System and command line interface to all online diagnostics and concurrent maintenance tools.
- ☐ Improved online diagnostics support asynchronous mode execution, NSC adapter testing, and HSX high-speed connection testing.
- ☐ A UNICOS source management tool adds binary identification and level information to modules.
- ☐ The file system now includes long file names, dual-residence for migrated files, and high-speed file system maintenance operations.

New standards for connectivity

Cray Research is dedicated to making supercomputing more accessible through adherence to industry standards. UNICOS 6.0 supports this goal by meeting the International Standards Organization Open Systems Interconnection (OSI) standard. Release 6.0 also supports X Window System version 11 release 4 from MIT, as well as enhanced NFS and NQS. High-speed networking media also are supported, including Ethernet, NSC HYPERchannel, Fiber Distributed Data Interface (FDDI) and High Performance Parallel Interface (HIPPI). Transmission speeds can be as fast as 100 Mbytes/sec, allowing quick movement of vast amounts of data.

Hardware support

The UNICOS 6.0 release supports CRAY Y-MP, CRAY X-MP EA, and CRAY X-MP systems, with Model B, Model C, and Model D I/O Subsystems, and CRAY-2 computer systems. A broad set of peripheral and communications hardware products is supported as well, as are the Cray Research Operator Workstation (OWS) and many third party communications hardware products.

Rich programming environments

UNICOS 6.0 provides the platform for Cray Research's programming environments, including the CF77 compiling system release 4.0.1 or later, Cray Research Standard C 2.0, Pascal 4.1, CFT2, Cray Ada 1.1, and CAL version 2, release 6.0.

Cray Research's programming environments are much more than just compiling systems. They include capabilities for optimizing, vectorizing, and multitasking code to take the best advantage of a Cray Research system's architecture. These reliable optimization features are turned on by default, assuring that user code runs as fast as possible. The programming environments also contain program development tools, debuggers, and subroutine libraries. Many tools have an optional X Window System interface (see sidebar).

Some of these compiler products are released separately and require separate licensing. Contact your Cray Research representative for more information.

Easy-to-use tools put peak performance at your fingertips

Cray Research is the first in the supercomputer industry to harness the power of visualization for its flexible, user-level performance analysis tools. The following tools are available with an X Window System interface that makes application performance data much more accessible to the user. The output of some of these tools can be shown and manipulated graphically, which makes it easy for users to interpret performance parameter relationships and optimize their codes.

Performance analysis tools

profview reveals which areas of a program have the highest CPU activity and thus are likely candidates for optimization.

flowview reveals the time spent in each routine.

perfview provides interpretations of performance data with suggestions for improvement.

Autotasking aids

atexpert provides interpretation of autotasking performance data (predicts dedicated performance without a dedicated system).

atscope provides information to help users parallelize loops that are not automatically detected by Autotasking software.

atchop automatically isolates code responsible for numerical differences in the loops of subroutines using the Autotasking software.

Debugging aid

CDBX aids in debugging codes, including multitasked codes (functionally equivalent to the UNIX dbx and the CTSS ddt debuggers).

System administration aid

The Cray-based Network Monitor is a visual tool for examining network connections in real time and pinpointing problem areas.

Training and publications

Cray Research offers a full range of training in the UNICOS operating system. Courses are available for first time users, programmers and analysts, system administrators, scientists, and engineers. Special focus courses also are offered.

The documentation set includes manuals specific to Cray Research computer systems, products, libraries, and utilities, as well as a UNICOS primer and UNICOS overview for beginning system users. Many of the UNICOS manuals have been updated for the 6.0 release.

Cray Research computer systems and UNICOS 6.0 — the world's most powerful team

With UNICOS 6.0, the performance of the world's fastest Cray Research computer systems is combined with the world's most powerful UNIX-based software to enable you to solve your most complex computational problems faster and more efficiently.

UniChem

a supercomputing environment for computational chemistry

The four main goals of UniChem

Unification of a variety of techniques within one framework

Visualization — the ability to manipulate molecules directly in a graphical fashion

Enabling chemistry that is possible only with a supercomputer
- ☐ Sophisticated scientific visualizations, allowing researchers to see things that cannot be seen otherwise
- ☐ Highly accurate simulations of large systems

Novel methods such as density functional theory

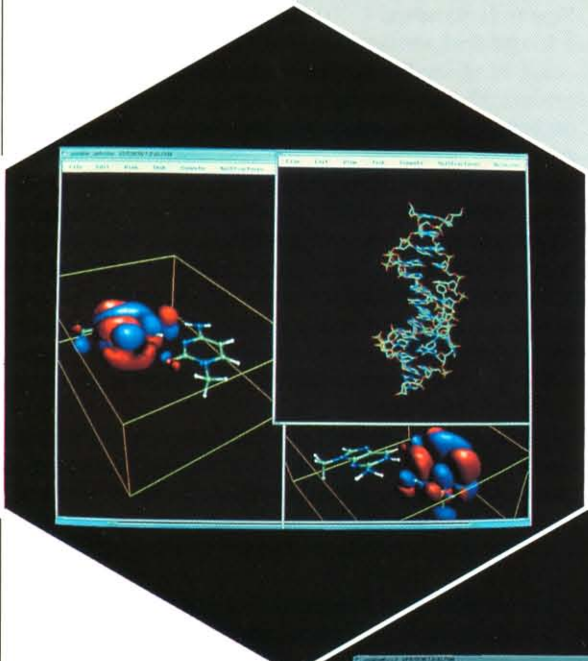

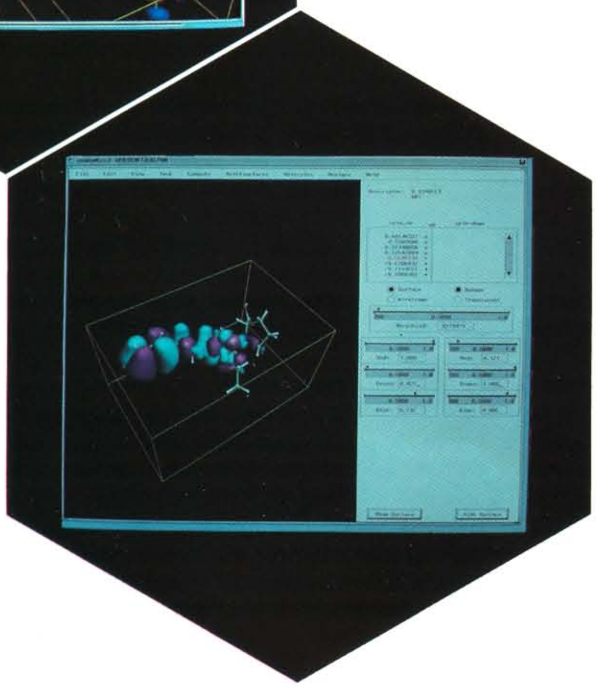

Above: X-ray structure of a DNA strand, along with the highest occupied molecular orbital (HOMO) and the lowest unoccupied molecular orbital (LUMO) of a guanine-cytosine base pair contained within the strand.

Right: Display of the highest occupied molecular orbital (HOMO) of the ACE inhibitor enalopril. Also shown are the menu of orbital energies and the color selectors used for setting up visualization of molecular orbitals.

Cray Research's UniChem computational chemistry environment is an easy-to-use, integrated system for product design and development in the chemical, pharmaceutical, and petrochemical industries. The UniChem environment links the power and networking capabilities of Cray Research computer systems to the convenience and visualization capabilities of Silicon Graphics IRIS workstations. It provides a common interface through which users can manage data, import molecular structures or build them with atoms or fragments in the UniChem fragment libraries, choose a computational method, select parameters, control computations, and visualize results. The use of a common interface to perform these tasks enables users to run many simulations in a short time.

At the core of the UniChem package is an integrated set of programs for performing numerically intensive calculations. These high-performance quantum mechanics programs, which run on Cray Research systems, include

- ☐ MNDO90, a semi-empirical program (including AM1, PM3, and MNDO) applicable to molecules that comprise up to about 500 atoms
- ☐ DGauss, a density functional program applicable to molecules and clusters that comprise up to about 150 atoms, including elements from hydrogen to xenon
- ☐ CADPAC, an ab initio program applicable to systems that comprise up to about 50 atoms

These programs cover the most widely used quantum chemistry methods and were developed on, and optimized for, Cray Research systems. In the UniChem environment, the MNDO90, DGauss, and CADPAC programs are accessed through the common user interface. UniChem is an open system that also allows users to attach other chemistry programs, run them through the user interface, and analyze and visualize the output on a graphics workstation.

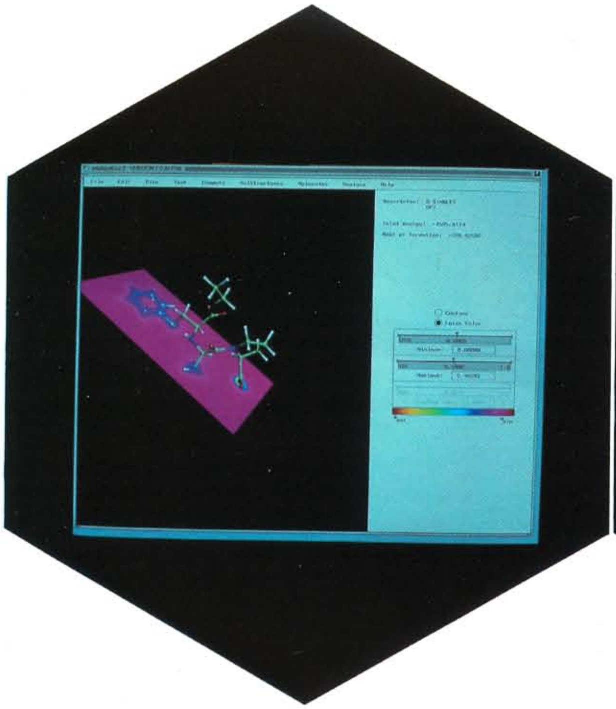

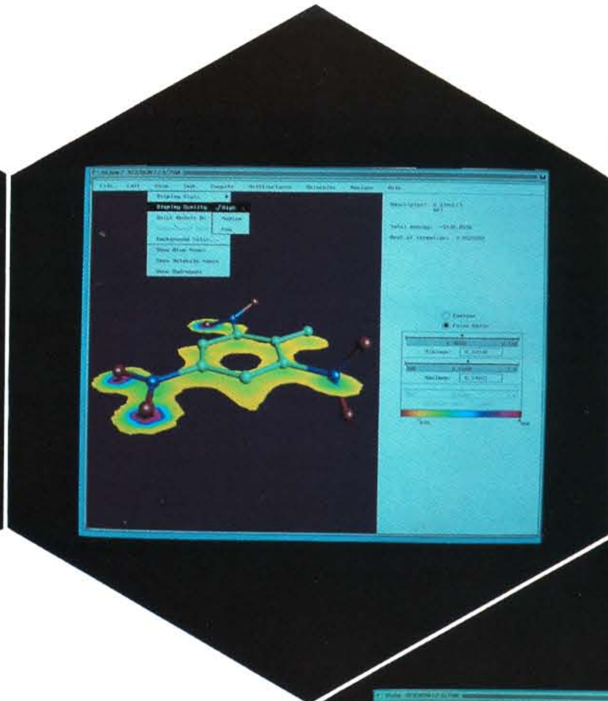

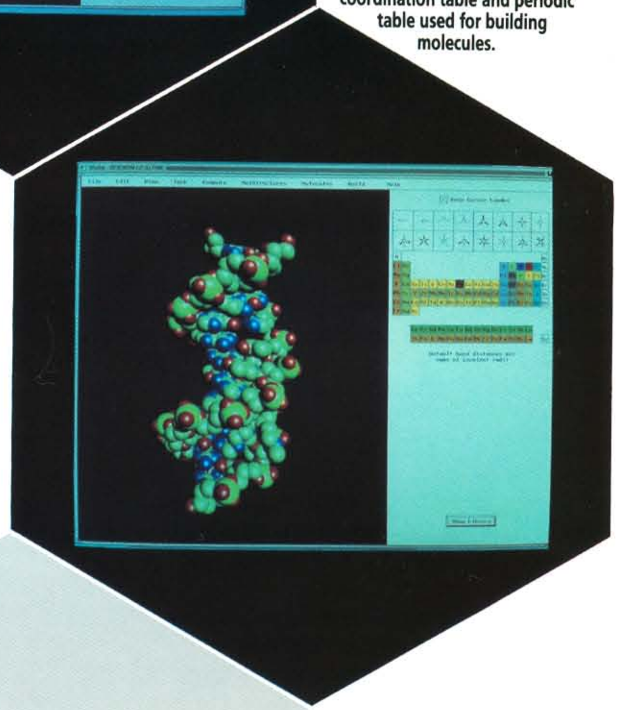

Far left: Two-dimensional display of the electron density of the ACE inhibitor enalopril. Also shown is the color map used for displaying the electron density.

Left: Two-dimensional display of the electron density in the plane of the aromatic ring of TNT. (Hydrogen atoms are hidden in this view.) Also shown are pull-down menus revealing alternatives for the quality of display, along with the color map used for displaying the electron density.

Below: CPK rendering of a 500-atom fragment from a strand of DNA. Also shown are the coordination table and periodic table used for building molecules.

UniChem thus complements and enhances the value of existing software.

The UniChem environment provides chemists with an invaluable computational tool for the design and development of

- ☐ Pharmaceuticals
- ☐ Agrochemicals
- ☐ Polymers
- ☐ Catalysts
- ☐ Specialty chemicals
- ☐ Advanced structural, optical, electronic, and magnetic materials

The UniChem environment's distributed architecture allocates each task to the appropriate computing resource, thereby maximizing computing efficiency. The Cray Research supercomputer solves the complex equations of quantum mechanics and can be linked to workstations, mainframe computers, and file servers. The Silicon Graphics IRIS workstation manages the user interface.

The UniChem software environment was developed by Cray Research in collaboration with researchers from major chemical and pharmaceutical companies, including 3M, Du Pont, Eli Lilly & Company, Exxon Research and Engineering, and Monsanto and its subsidiaries G. D. Searle and Nutra Sweet. UniChem 1.0 runs on all Cray Research computer systems under release 5.1 or later of the UNICOS operating system, which is based on the UNIX System V operating system from UNIX System Laboratories, Inc. The user interface component runs on Silicon Graphics IRIS workstations under version 3.3 or later of the IRIX operating system. Workstations must be connected to the supercomputer through TCP/IP links.

Through UniChem, Cray Research provides chemists with the most powerful tool available for building and interacting with computational chemistry models. This unique environment brings the full power of supercomputing to the desktops of product designers and developers in the chemical industries.

UniChem summary

UniChem delivers multiple molecular modeling and simulation techniques in a single, integrated software environment. The package features

- ☐ **Accurate, large-scale molecular simulation on a Cray Research supercomputer**
- ☐ **State-of-the-art applications software from research groups worldwide**
- ☐ **Integrated quantum chemistry programs**
- ☐ **Import and export of molecular structures**
- ☐ **Three-dimensional model building**
- ☐ **Editing, scientific visualization, and direct manipulation on powerful graphics workstations**
- ☐ **Easy set-up, launch, control, visualization, and data management.**

For more information about the UniChem software environment, contact Elaine Frankowski, Cray Research, Inc., Industry, Science & Technology Department, 655-E Lone Oak Drive, Eagan, MN, 55121; telephone: (612) 683-3683.

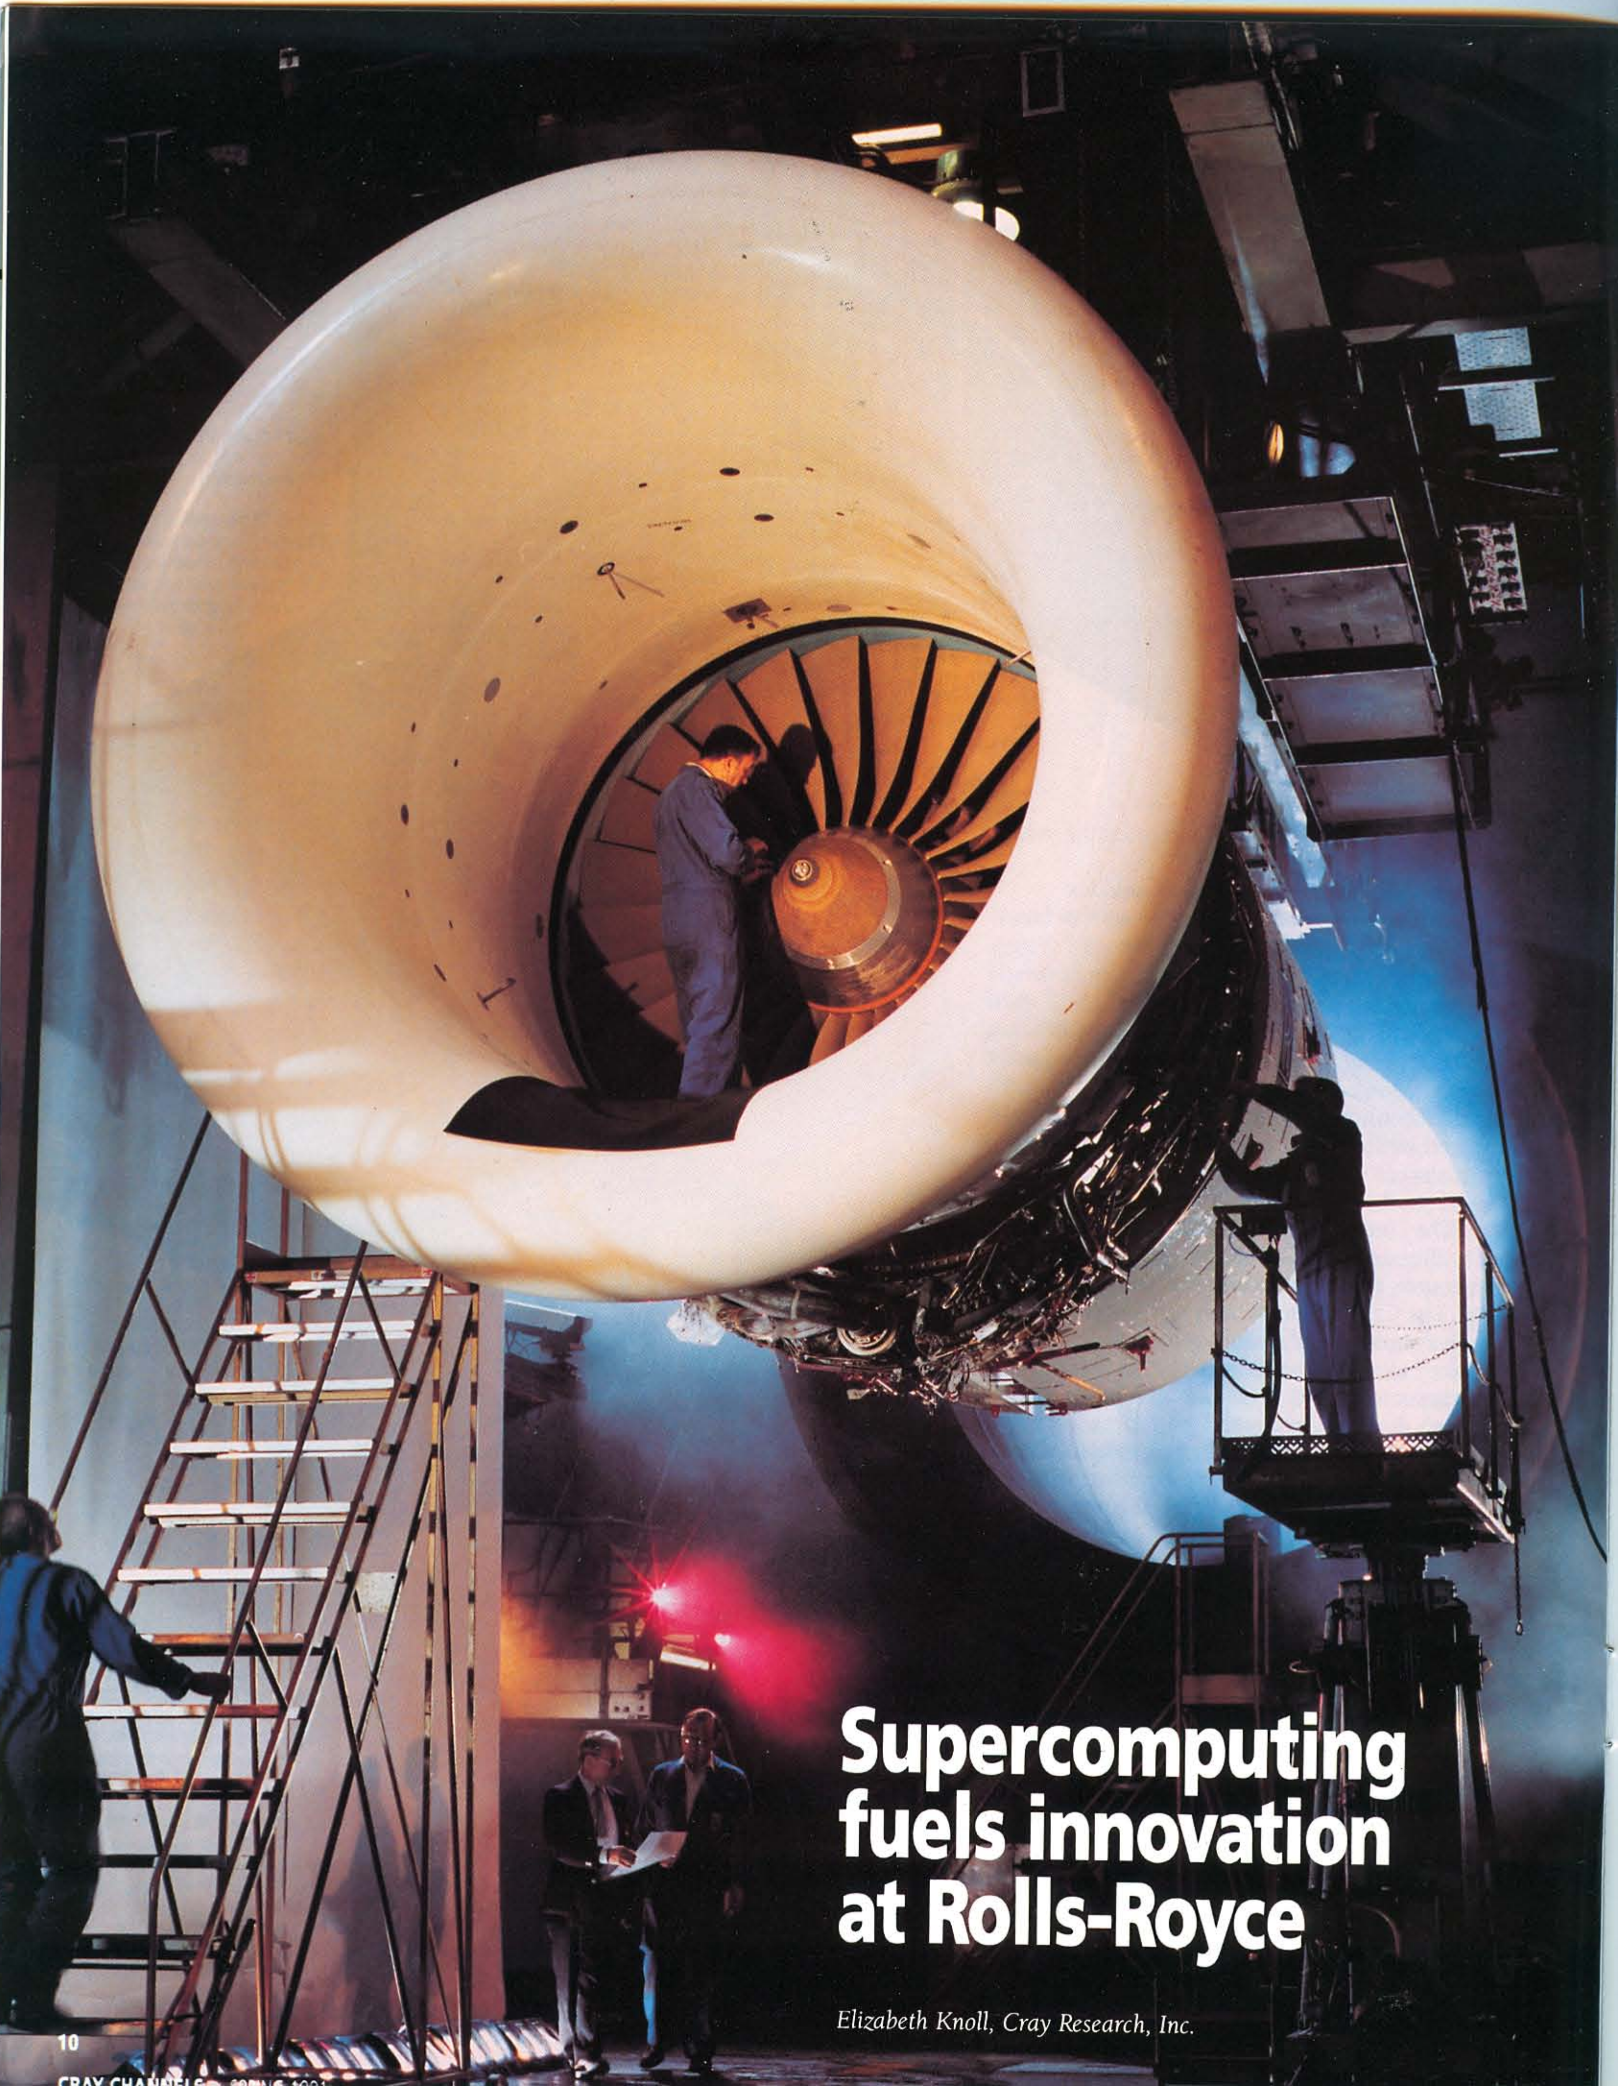

Supercomputing fuels innovation at Rolls-Royce

Elizabeth Knoll, Cray Research, Inc.

Although simple in concept, a working jet engine is one of the most highly stressed and complex engineering structures. A large commercial turbofan engine comprises thousands of moving parts and moves as much as 50 tons of air per minute. A compressor forces the air to 35 times atmospheric pressure, and the high-pressure turbine works at the highest temperature for maximum efficiency. About 90 high-pressure turbine blades are connected to a disc that spins about 10,000 times per minute. Each of these blades is subjected to a centrifugal force about equal to the weight of a London bus.

Because each of these components must work in perfect balance from takeoff to touchdown, Rolls-Royce engineers use supercomputers to design and evaluate the components of gas turbine engines before they reach the skies. Since 1985, Rolls-Royce has been using Cray Research supercomputers to reduce the time and cost of designing and testing its engines. In October 1990, the company installed a two-processor CRAY Y-MP system at its Derby, U.K., facility.

The research and development cycle for a new engine design may span 10 years and cost about $1 billion. Up to six years are spent researching materials and manufacturing methods, and up to five years are spent developing the engine. Once a project is approved for development, the risks must be low, because backtracking is extremely expensive. By using supercomputer analysis, the company can invest its resources in "advanced engineering" — research and development done before resources are committed to new engine designs. As a result, Rolls-Royce has streamlined its design and testing cycle. Thirty years ago the company might have built 20 development engines to validate a new design; today it builds only 6 to 8, with significant cost savings.

Rolls-Royce uses its Cray Research supercomputer in three major research areas: structural analysis, computational fluid dynamics (CFD), and process modeling. About 45 percent of the company's supercomputing resources is devoted to structural analysis, including static, dynamic, and thermomechanical analysis of fan blades, turbine blades, and other engine components. Another 45 percent is devoted to CFD for the design of various stages of the engines, compressors, turbines, combustion, and engine installations. An additional 10 percent of the computer time is allocated to process modeling, or simulating the behavior of materials during the manufacturing processes of forging, casting, and heat treatment.

Structural analysis of components

Ensuring the mechanical integrity of engines is a critical step in the design process. A wide range of finite element programs is used to study the static and dynamic response of engine components, from turbine blades to engine casings. Supercomputer power enables engineers to model large, integrated structures, increasing the accuracy of findings and creating new design approaches.

"In the 'good old days,' engines were relatively simple things made from components that could be analyzed independently," said Rob Tuley, a stress computing technologist at Rolls-Royce. "The trend over the past 10 years is for engines to get larger, for thrusts to increase, but also for reduced weight and increased performance. As a result, the structure must be analyzed as an integrated whole. As you model a component, maybe one-third of the model is what you are really interested in — but you have to add the other two-thirds to make the model accurate."

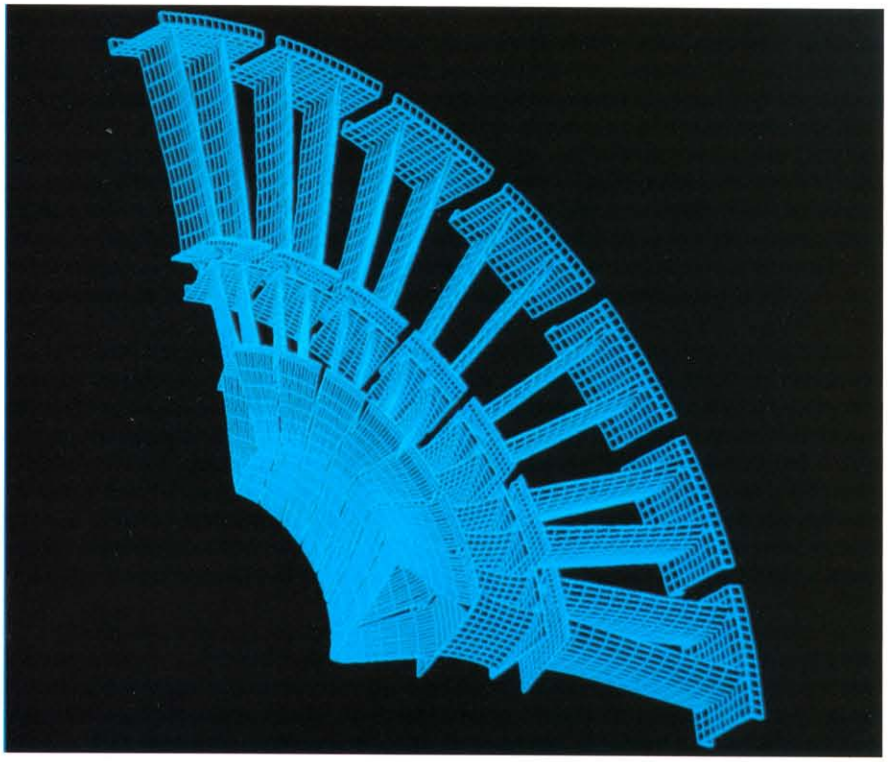

Figure 1. Compressor outlet guide vane assembly.

For example, Figure 1 shows an "exploded" model of a set of compressor outlet guide vanes. In addition to their aerodynamic function, the vanes act like spokes in a wheel, transferring loads from hub to rim. In this lightweight design, only alternate vanes carry the load. To understand the modes of vibration of one of the outer vanes, it is necessary to model a sector comprising two inner and two outer vanes. This 2300-element model runs in 30 minutes on the CRAY Y-MP system.

Tuley would like to take the idea of integrated analysis to its limits. His group is developing algorithms that will improve analysis times by a factor of 400. "These speedups, and efficient network communications between the Cray Research supercomputer and Silicon Graphics workstations used for visualization, can pave the way for modeling whole-engine dynamics," said Tuley. "We are looking at running a dynamics model of a complete engine on the supercomputer, calculating results on the fly, and displaying them on the workstation. It is like a numerical testbed. The user can play with the throttle, change the model parameters, and observe the results in real time."

Predicting the outcome of a bird strike

It is impossible to eliminate the hazard of a jet engine ingesting birds, especially at low altitude during takeoff and landing. The engine must remain under control during and after a bird strike. It must either continue to deliver power or safely shut down.

Opposite page: Rolls-Royce has carried out the first run of its new Trent engine, the latest member of the RB211 family, at the company's Derby, U.K., test facility. The Trent powers the McDonnell Douglas MD-11, Airbus A330, and is offered on the Boeing 777 Airplanes.

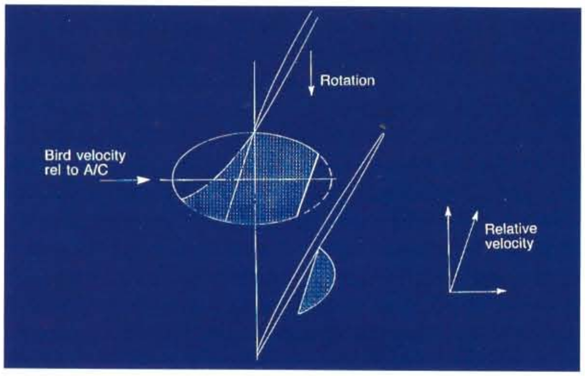

Figure 2. Slicing action of a rotor blade.

At Rolls-Royce, all fan blades go through a series of rig tests before they are subjected to an engine certification test. For certification, up to eight birds are fired into the engine within two seconds. But before any test is performed, engineers optimize the design by computer methods and predict the outcome.

Using DYNA3D, a finite element program optimized for use with Cray Research systems, Mike Lawson, a mechanical technology specialist, has simulated worst-case bird strikes. He explained that when a bird strikes a fan blade tip, the bird and blade meet with a relative velocity of about 400 meters per second, giving a normal velocity of around 130 meters per second onto the blade pressure surface. This causes the blade to slice through the bird. The mass of this slice striking a blade is determined by the slice taken by the previous blade and the path of the leading edge of the impacted blade (Figure 2).

The simulation involves many complex phenomena. First, the bird behaves like a fluid. Second, the impact energy on a single blade due to a typical mass of 0.35 kg will deform the blade at high stress rates. Third, the slice of bird impacting a blade plastically is dependent on the developing deformation of both bird and blade. Finally, the yield strength of the blade material is highly dependent on the strain rate of the plastic deformation.

A full finite element model of a single Rolls-Royce wide-chord fan blade may consist of 20,000 eight-node brick elements, and the bird model may consist of 5000 additional elements, resulting in a model of 80,000 degrees of freedom.

Lawson uses the company's own finite element code to calculate the initial stresses and deflections. During the DYNA3D simulation, the impact forces are calculated using the contact-impact algorithm. From this, the response of the blade is calculated in terms of deflections, and from these results, the stresses, strains, and plastic deformation are calculated. New forces, including internal forces from the stresses, and new contact forces, then are calculated from the displaced shape (Figure 3).

The dynamic analysis stops when the plastic deformation is finished. The displaced shape and elastic stress state of the blade are used as initial conditions for a final nonlinear stress analysis to calculate the static deformed shape of the blade after the bird strike. This then is used for air-flow calculations and for comparison with the subsequent rig test.

Postprocessing initially is performed using TAURUS, a postprocessor for DYNA3D. Still pictures of displaced shapes and stress contours are transmitted from the Cray Research supercomputer and displayed on a local terminal. Improved networking capability now enables researchers to produce animated displays of the simulation by using Cray Research's Multipurpose Graphic System (MPGS) software package run locally on the CRAY Y-MP system and Silicon Graphics workstations. The speedup of the postprocessing phase greatly increases user productivity and improves the design timescale.

Analyzing stresses in an engine casing

Figure 4 shows a finite element model of one section of the compressor casing of the Rolls-Royce Tay engine. The model was used to analyze stress in the casing. The inner conical ring houses the compressor

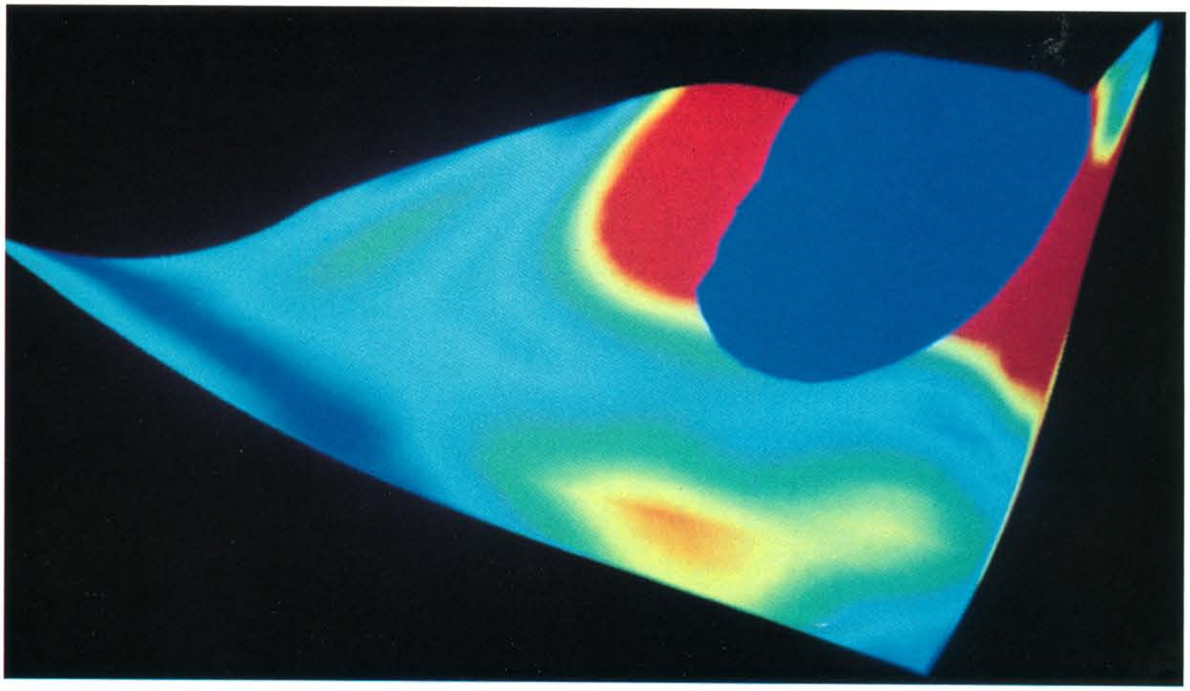

Figure 3. Color fringe plot of effective stresses at 240 microseconds on a fan blade.

shaft bearings, the intermediate ring contains axial air bleed slots visible on the end face of the casing, and the outer ring includes engine-to-airframe mounting points and other local features. The rings are connected by hollow spokes.

The model comprises 5200 elements and has 110,000 degrees of freedom. A stress analysis was performed with a differential axial load applied between the inner front and outer rear flanges of the casing. This analysis was performed in 69 minutes of wall-clock time on the two-processor CRAY Y-MP system. The heaviest computational task in the analysis was factorization of the 1.5 Gbyte assembled stiffness matrix. "The out-of-core matrix factorization was run using five Mwords of memory and ran in 30 minutes of CPU time and 35 minutes of wall-clock time, demonstrating the excellent I/O performance of the Cray Research system," said Tuley.

The resulting deflections are color-coded in Figure 4, with the red regions showing the largest movement. The corresponding stress field is shown in Figure 5, which plots the Von Mises stress function. The need for detailed three-dimensional modeling of the component is evident from the complex pattern of local stress concentrations in the structure.

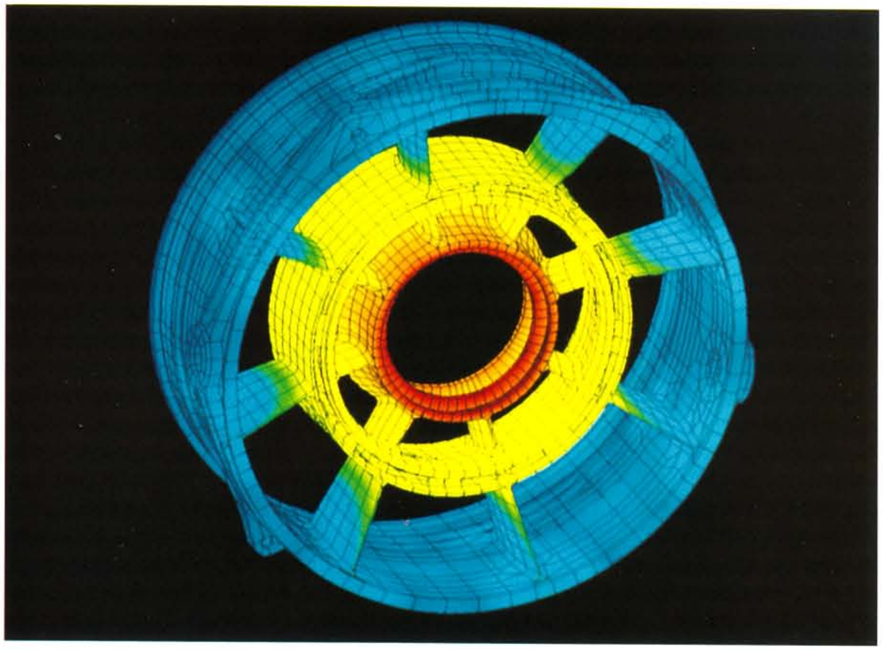

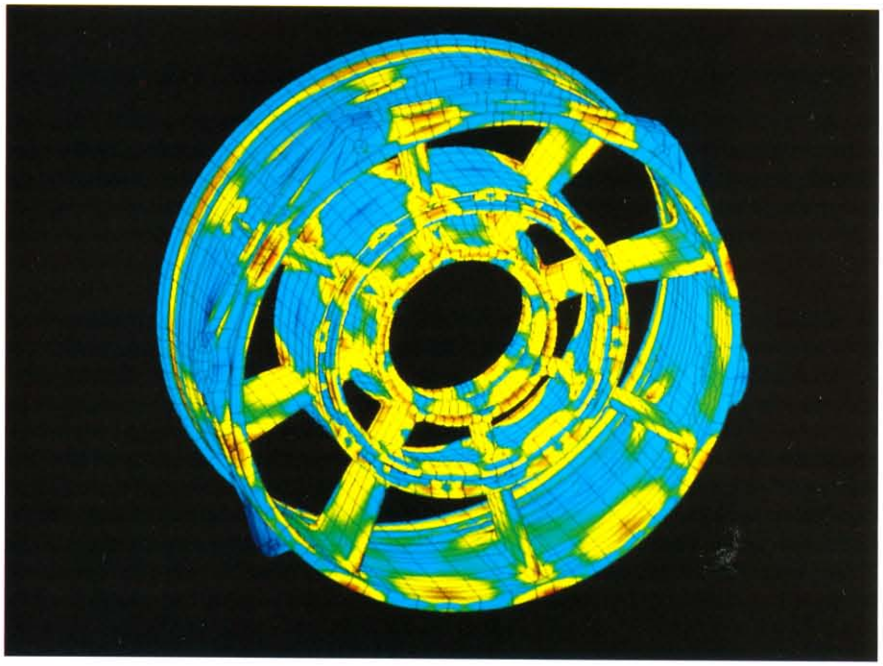

Figure 4 (top). Color-coded deflections of a section of the compressor casing of a Rolls-Royce Tay engine. Red shows areas of greatest movement.

Figure 5 (bottom). Von Mises stress field corresponding to deflections shown in Figure 4.

Modeling fluid flow

"Computational fluid dynamics enables us to perform numerical experiments relatively cheaply and efficiently. The challenge is to use the information produced to generate deeper understanding of complex flow features to improve component design," explained Peter Stow, chief of Theoretical Science at Rolls-Royce.

A typical three-dimensional viscous flow calculation for a single blade row can require about 50,000 grid points, with a number of flow variables at each point, totaling about 300,000 unknowns. About 2000 iterations are required to solve the nonlinear equations; a complete solution may involve 300 billion calculations. According to Stow, "The availability of supercomputers such as the CRAY X-MP and CRAY Y-MP systems has greatly reduced the time needed to perform such calculations. These large three-dimensional calculations now can be performed routinely on a production basis."

"In fact, numerical solutions provide far more information than could be acquired from experimental rigs or engine tests," said Stow. Computer modeling can provide quantitative data such as pressure loads on the components, heat transfer rates, and loss levels. Stow added that CFD is used to analyze properties of all major turbomachinery components at Rolls-Royce, such as compressor aerodynamics, turbine aerodynamics and heat transfer, combustion, installations (such as intakes and nozzles), disc cavity flows, and unsteady aerodynamics of flutter and forced response.

Improving the design of transonic fan rotors

One application of three-dimensional CFD at Rolls-Royce has been the simulation of flow in transonic fan rotors. The fan is a particularly crucial engine component because it supplies about 70 percent of engine thrust. An inefficient fan can increase a fan engine's fuel consumption greatly.

The company's development of a new wide-chord fan has been called one of the most significant advances in aerodynamic engineering in years. The new wide-chord fan has demonstrated the ability to pass a greater flow per unit frontal area with an increased pressure ratio and 2 to 3 percent more efficiency than traditional blades. The Tay, V2500, and both the -535E4 and -524G/H variants of the RB211 engine are in airline service with these fans; the wide-chord fans currently are being specified for the Trent series. Compared to conventional fans on earlier engines, the new fans have fewer blades, longer chord, and no part-span vibration dampers.

The latest analysis technique used by Rolls-Royce to assess fan rotor designs solves the three-dimensional Reynolds-averaged Navier Stokes equations for the time-averaged turbulent flow in individual blade rows with a semi-implicit, time-marching algorithm. The code is capable of operating over the entire fan flow range between choke and stall. It is inherently superior to the coupled, viscous-inviscid interaction methods that are used routinely, because it automatically incorporates key phenomena such as tip clearance flows, end-wall boundary layers, and secondary flows. The goal is to match engine requirements for high thrust at top speed with efficiency at cruising conditions.

Three-dimensional viscous modeling programs have been particularly useful for diagnosing poorly understood phenomena, such as the behavior of the fan when it is subjected to inlet flow pressure distortion. This may occur in air-frame applications with intakes far removed from the engine face, or when a conventional nacelle is subjected to crosswind.

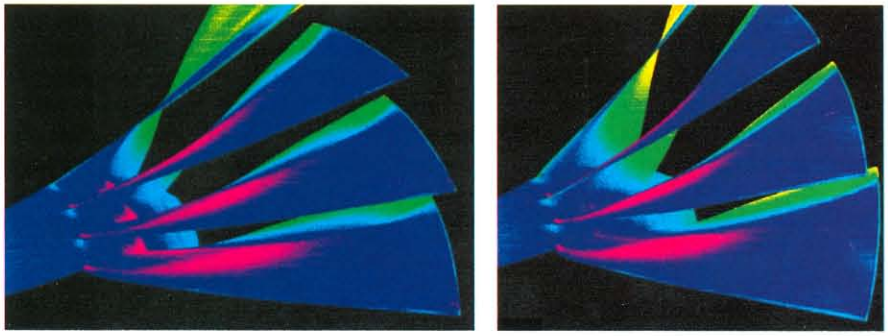

Figure 6. A first-generation wide-chord fan blade modeled in three-dimensions with a viscous code; inlet tip total pressure (left), with clear inflow (right). The tip distortion results in distress for the fan hub.

Figure 6 shows a three-dimensional viscous analysis of a first-generation wide-chord fan blade both with clean inflow and when subjected to an inlet total pressure distortion at the casing. The distortion leads to a radial redistribution of the flow, which, in this design, results in high Mach numbers near the hub. Figure 7 shows a second-generation design, run with the same inlet conditions. This analysis shows no evidence of the high hub Mach number caused by tip distortion in Figure 6.

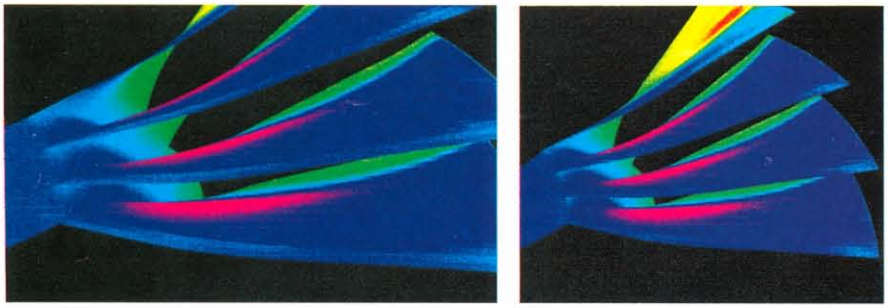

Figure 7. A second-generation wide-chord fan blade modeled in three-dimensions with a viscous code; inlet tip total pressure (left), with clear inflow (right). This analysis shows no evidence of hub distress caused by distortion.

According to Christopher Robinson, research group leader in Compressor Technology, "Designers have been able to improve new designs, such as that of the Trent wide-chord fan to make them more tolerant to inlet distortion than the first-generation wide-chord fans." To further improve engine efficiency and performance, Rolls-Royce is researching and installing three-dimensionally designed blading in the fan, compressor, and turbine stages of its engines.

"The use of Cray Research supercomputers has greatly improved the turnaround of jobs, such that three-dimensional analysis techniques are now an integral part of the fan and turbine product definition iteration," said Robinson. "A new workstation-based output visualization program has been introduced for the detailed interrogation of results from the three-dimensional calculation performed on the Cray Research system. This has greatly improved the speed of output processing and depth of understanding."

These techniques enable designers to understand the important flow mechanisms contributing to loss reduction and to optimize the blade airfoils and flow path to achieve maximum efficiency.

Calculating combustion

Principal variables in combustion modeling include three velocity components: fuel fraction, turbulence, and energy. "We are trying to calculate the flow and temperature distribution inside the combustion chamber because those conditions are distributed into the turbine," explained Christopher Priddin, manager of Combustion Methods. "We try to locate where any hot spots occur in the temperature distribution and, more importantly, where they originate."

Supercomputer simulations also enable engineers to identify the hottest areas in the combustor. As a result, they can modify the combustion process by adding air to, or cooling the sides of, the combustion chamber. The goal is to achieve ideal mixing, reduce emissions of nitric oxide, and eventually tackle smoke and carbon dioxide emissions.

"With the Cray Research system we are able to create finer meshes and do longer runs than we are able to do on other computers," said Priddin. "It's really a matter of being able to include finer detail. The processing speed and storage space let us address larger problems."

Modeling manufacturing processes

Rolls-Royce uses process modeling techniques to improve the quality of the materials in its products, as well as its response to the marketplace, according to Gordon Higginbotham, chief of Metal Manufacturing Technology. "That, of course, equates to an overall reduction of lead time and cost," he said. "To achieve that, we want to ensure that when we accept something for manufacture, we can actually make it. We want to understand the processes that are involved in manufacturing components and model them accurately."

The process modeling group uses computers to simulate the casting, forging, and heat-treating of metals. These complex processes have varied attributes and requirements. For example, forging simulations generally require a large-scale deformation analysis of a component, thermal analysis, and a small-scale deformation simulation for the die. Three-dimensional modeling of forging requires powerful computing resources. Important engine components that are forged include compressor and turbine discs and blades.

Casting models involve non-isothermal fluid-flow simulation of mold-filling, thermal analysis of solidification, and stress and deformation analysis. Further analysis may be required to study the convection-driven currents in the liquid metal once the mold is filled. The data produced from the analysis need to be interpreted further to predict the features of interest. Components produced by this method include turbine blades and jet engine casings.

Heat treatment is used to ensure that forged discs and turbine blades have the microstructures and properties required to work correctly. Castings also are heat treated to acquire the required properties, and to relieve stress in welded assemblies. For example, nickel-based superalloy discs for jet engines often are oil quenched following heat treatment at temperatures of about 1050° C. This induces rapid cooling, which strengthens the alloy. Heat treatment modeling is used to predict disc features such as residual stresses, cracking, microstructural defects, and distortion (Figure 8).

Computers enable researchers to see inside these processes — some of which cannot be examined physically. "You could never see inside a forging as it is being forged, but process modeling is enabling us to do just that," said Ian Williams, a research analyst. "This is providing good insight into the physics of the situation."

Ian Craighead, a technologist in the group, models the heat-treatment process in two dimensions using the ABAQUS software package on the Cray Research system. "The Cray Research system provides rapid computation, which allows us to check the model after a few minutes," he said. "There's a world of difference between three minutes on this system and three minutes on a workstation. This difference allows me to run large, highly nonlinear models."

Process modeling also enables Rolls-Royce engineers to go through the trial-and-error process via computer before parts are made and tested. For example, process modeling was used to simulate the forging of a turbine disc for a jet engine in a new turbine disc alloy. As a result, the lead time was reduced from two years to four months largely because the manufacturing route developed on the computer resulted in the first disc being forged correctly the first time. And once researchers understand the process, they can build in controls; they can inspect the process rather than the product and increase product integrity.

"By using the Cray Research system, the quality of research is improving significantly and rapidly," said Higginbotham. "Introducing a Cray Research supercomputer into the loop enables us to tackle things that we couldn't have tackled before. We expect to see a very good payback from that based upon the experience we have had to date."

Modeling heat treatment processes and predicting residual stresses

Turbine and compressor disc chambers for engines are oil quenched to achieve the required material properties. Due to the plastic deformation that occurs during quenching, residual stresses are "locked in" to the components. The magnitude and distribution of these stresses will determine the life of the component in service.

Traditionally, these stresses have been determined by physical methods that require destroying the components. These methods are expensive and time consuming. Craighead and his colleagues have found that, by predicting these stresses using finite element methods (FEM) on the Cray Research system, they can achieve results that equal the accuracy of experimental measurements — without destroying components. More specifically, they used FEM to predict the development of stresses as a result of heat quenching in simple model discs made of Waspaloy, a nickel-based superalloy.

To develop a model, an approximation to the disc shape was machined from nickel alloy and the thermal history was used to calculate the thermal boundary conditions. Those conditions then were used as the basis for a finite element analysis to simulate quenching. This stress analysis, based on the calculated thermal history, was run using the appropriate material properties. Finally, a disc was quenched and its stresses were measured to verify the calculated stress distribution. The temperatures calculated closely matched those measured, and the predicted and measured stresses also showed a good match.

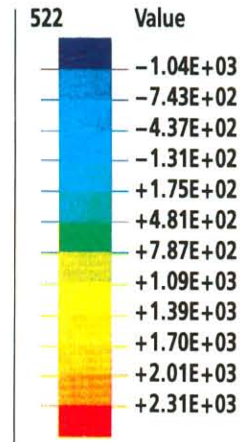

Figure 8. Axial stress contours in a forged disc after quenching. This information helps researchers predict the likelihood of cracking.

Maintaining the technical edge

About 100 changes were made during the design of blades that went into service in 1965. Today, three to four changes may be made during the design of a new blade. "Engineers now are able to predict well in advance how an engine will perform — not only its thrust, but also internal temperatures of components, fluid flow around components, stress, and vibration levels. We have to do less and less engineering design validation," explained Derek Gale, head of Engineering Computing.

Competition in the aero-engine field is fierce, and Rolls-Royce is determined to continue developing products at the leading edge of technology. "We need to maintain state-of-the-art knowledge — part of that is having the right computing facilities," Gale said.

"Maintaining a technical edge is achieved in two ways at Rolls-Royce," added Gale. "One is through technical methods, so researchers understand how components behave in a physical way. But there is also a technical edge in terms of how rapidly you can bring your product to market. The Cray Research system plays a key role in both of these."

Acknowledgments

Special thanks to Ian Craighead, Derek Gale, Gordon Higginbotham, Mike Lawson, Christopher Priddin, Christopher Robinson, Peter Stow, Rob Tuley, and Ian Williams of Rolls-Royce plc.

About the author

Elizabeth Knoll is a communications specialist in Cray Research's Marketing Communications department.

Large-scale methods in computational electromagnetics

Daniel S. Katz and Allen Taflove, Northwestern University, Evanston, Illinois
Jeffrey P. Brooks and Evans Harrigan, Cray Research, Inc.

The numerical modeling of electromagnetic wave phenomena can be a computationally intensive task. To date, the design and engineering of aerospace vehicles has been the primary application driving the development of large-scale methods in computational electromagnetics (CEM). Efforts in this area have been aimed primarily at minimizing the radar cross section (RCS) of aerospace vehicles. RCS minimization enhances the survivability of vehicles that are subjected to precision-targeted ordnance. The physics of RCS is determined by Maxwell's equations and the constitutive properties of a vehicle's materials. As a result, the interesting situation arises in which the effectiveness and cost of state-of-the-art aerospace systems in part depends on the ability to develop an efficient engineering understanding of 120-year-old equations that describe the propagation and scattering of electromagnetic waves.

Two algorithms are of primary interest in this field: the robust, traditional, full-matrix, frequency-domain integral equation method of moments (MoM); and emerging time-domain, grid-based direct solutions of Maxwell's curl equations. Both types of algorithm make efficient use of Cray Research hardware and software capabilities.

Full-matrix MoM field computations at 2 GFLOPS

In the MoM area, one group of important codes originated with the Rao-Wilton-Glisson triangular surface patch technique for RCS analysis of arbitrarily shaped three-dimensional conducting structures.[1] Cray Research analysts determined that the primary task here involves the solution of very large, dense, complex-valued matrices (10K by 10K and larger) that exceed the available central memory. A strategy evolved to develop a complex-valued lower-upper matrix decomposition program that utilizes an efficient out-of-memory scheme and is adaptable to multiple CPU usage. The result was CLUD — Complex Lower-Upper Decomposition, with versions developed for the CRAY-2, CRAY X-MP, and CRAY Y-MP computer systems. This work rapidly gained popularity among MoM users, and Cray Research scientists have provided assistance to members of the CEM community in adapting these matrix solvers to many MoM codes.

A second group of important MoM codes originated with the Newman ESP-3 rectangular surface patch technique for RCS analysis of arbitrarily shaped three-dimensional conducting structures.[2] Cray Research analysts adapted their out-of-memory scheme for this code and subsequently developed an out-of-memory solver suitable for simultaneous solution of the monostatic RCS at a large number of illumination angles (right-hand sides). In fact, the number of right-hand sides could be in the thousands, approximating the order, *N*, of the MoM matrix. Subsequently, a parallel-processing version of the "*N* right-hand-sides" code was developed.

Although CLUD works well, two drawbacks had to be addressed for very large problems:

- ☐ The input/output (I/O) for CLUD is either synchronous to disks or synchronously staged from disk to Cray Research's SSD solid-state storage device. If the matrix is scaled to fit entirely in an SSD, this is not troublesome, and near-peak performance is achieved on the CRAY X-MP and CRAY Y-MP systems. However, a 20K-by-20K complex-valued MoM matrix requires an 800 Mword SSD, which is not currently available. In CLUD, very large problems of this size require synchronous I/O between disks and SSD, which reduces overall performance.
- ☐ The CLUD algorithm is based on a SAXPY type kernel that works on individual columns. This kernel runs at peak performance on the CRAY X-MP and CRAY Y-MP systems, but not on the CRAY-2 system because of a high ratio of memory operations to computation.

Because the Cray Research mathematical software group had optimized the BLAS-3 (Basic Linear Algebra Subroutines) to run at near-peak performance on all Cray Research computer systems, an improved algorithm was developed that was based on these kernels. A block-oriented method was adapted from LAPACK to run out-of-memory by Jeffrey Brooks of Cray Research's benchmarking department. The routine, CGETRF, made use of two BLAS-3 kernels, CGEMM (complex matrix multiply) and CTRSM (complex triangular backsolve).

To adapt CGETRF to run out-of-memory, the matrix is divided into slabs. (A slab is a matrix block consisting of a large number of adjacent columns of the matrix.) The matrix is decomposed from left to right, one slab at a time. Computation works on pairs of slabs. To compute a new leading slab, all preceding slabs need to be brought into memory, one at a time, for computation. This is an I/O pattern similar to that used in the existing CLUD code. However, three slab-sized memory buffers are used in the new code to allow for asynchronous I/O. The partial-pivoting scheme used in CGETRF is preserved in the new out-of-memory version.

A routine called CMXMA was written to take advantage of Golub's identity, which reduces the multiplication count for complex-number products. CMXMA converts complex matrix products to three real matrix multiplies and several matrix additions. (This routine is available as CGEMMS in SCILIB 6.0.) The Strassen's real matrix multiply (SGEMMS) was used to save further on operations. SGEMMS is a Strassen's algorithm extension to the standard BLAS-3 matrix multiply routine, SGEMM. SGEMMS was written by Cray Research's mathematical software group and is included in version 6.0 of Cray Research's UNICOS operating system.

When automatically multitasked to run on all eight processors on a CRAY Y-MP system, the new out-of-memory code ran at average computation rates exceeding 2.1 GFLOPS. Only 1.99 hours were required to process a 20K-by-20K matrix. During this run, 138 Gbytes of I/O were discharged to and from seven DD-40 disk drives. Yet, only 228 seconds (3.8 minutes) represented I/O wait time. In fact, 90 percent of the actual I/O operations were performed concurrently with the floating-point arithmetic by virtue of the asynchronous I/O scheme and therefore did not contribute to the observed wait time. As matrix size increased, the relative efficiency of the asynchronous scheme improved, with the I/O concurrency factor rising to 95 percent for a 40K matrix. Thus, the massive I/O associated with solving huge, dense, complex-valued MoM matrices could be buried almost completely.

When automatically multitasked to run on all eight processors on a CRAY Y-MP system, the new out-of-memory code ran at average computation rates exceeding 2.1 GFLOPS.

Multiprocessing space-grid time-domain codes

Although the LU decomposition strategy described here is highly efficient, the fundamental [order (N^3)] computational burden of LU decomposition remains dimensionally large. In fact, it is so large that there is virtually no prospect for using the traditional, full-matrix MoM to computationally model entire aerospace structures, such as fighter planes, at radar frequencies much above 150 MHz. Yet, radar frequencies of interest can greatly exceed 150 MHz, climbing to 10 GHz and higher. Much research effort, therefore, has been invested in the development of alternative iterative frequency-domain approaches, including conjugate gradient and spectral methods, that preserve the rigorous boundary-integral formulation of MoM while realizing dimensionally reduced [order (N^2) or less] computational burdens. Such methods would permit, in principle, the modeling of entire aircraft at radar frequencies above 1 GHz. However, these alternatives may not be as robust as the full-matrix MoM, insofar as they may not provide results of engineering value for a wide class of structures without the user having to wonder if the iterative algorithm has converged.

Problems involved in applying frequency-domain, full-matrix MoM technology to large scale RCS modeling have prompted much new interest in an alternative class of non-matrix approaches: direct space-grid, time-domain solvers for Maxwell's time-dependent curl equations. These approaches appear to be as robust and accurate as MoM, but have dimensionally-reduced computational burdens [approaching order (N)] such that whole-aircraft modeling for RCS can be considered in the near future. Currently, the primary approaches in this class are the finite-difference time-domain (FD-TD) and finite-volume time-domain (FV-TD) techniques.[3,4] These are analogous to existing mesh-based solutions of fluid-flow problems in that the numerical model is based upon a direct, time-domain solution of the governing partial differential equation. Yet, FD-TD and FV-TD are very nontraditional approaches to CEM for detailed engineering applications, where frequency-domain methods (primarily full-matrix MoM) have dominated.

FD-TD and FV-TD methods for Maxwell's equations are based on volumetric sampling of the unknown near-field distribution within and surrounding the structure of interest. The sampling is at sub-wavelength (λ_0) resolution to avoid aliasing of the field magnitude and phase information. Overall, the goal is to provide a self-consistent model of the mutual coupling of all the electrically-small volume cells that comprise the structure and its near field, even if the structure spans tens of λ_0 in three dimensions and there are tens of millions of space cells.

The primary FD-TD and FV-TD algorithms used today are fully explicit, second-order-accurate grid-based solvers that use highly vectorizable schemes for time-marching the six vector components of the electromagnetic near field at each of the volume cells. The explicit nature of the solvers is maintained either by leapfrog or predictor-corrector time-integration schemes. Present methods differ primarily in the set up of the space grid (almost-completely structured for FD-TD, body-fitted or unstructured for FV-TD) and the enforcement of EM field continuity at the interfaces of adjacent cells. As a result, the number of floating point operations needed to update a field vector component over one time step can vary by about 20 to 1 from one algorithm to another.

However, the choice of algorithm is not straightforward, despite this wide range of computational burdens. There is an important tradeoff decision to be made. Namely, a faster, simpler solver such as FD-TD uses meshes that may not be compatible with those used in other aerospace engineering studies, computational fluid dynamics (CFD) studies in particular. As a result, there is much "homework" to be done as researchers learn to generate a new class of three-dimensional meshes specific to Maxwell's curl equations. On the other hand, the more complex FV-TD solvers can utilize existing CFD mesh generators, but require substantially more algorithmic computer arithmetic and storage. Both FD-TD and FV-TD algorithms are highly vectorized, having been benchmarked at over 200 MFLOPS on one processor of a CRAY Y-MP system for real models. However, the attainment of even higher MFLOPS rates may be hampered by the fact that the the space grids have an unavoidable number of non-standard cells that require either scalar or odd-lot vector operations. These nonstandard cells result from the need to program a near-field radiation condition at the outermost grid boundary (simulating the grid continuing to infinity), and the need to stitch together varying types of meshes to accommodate complex structure shapes. Despite this, it has been found possible

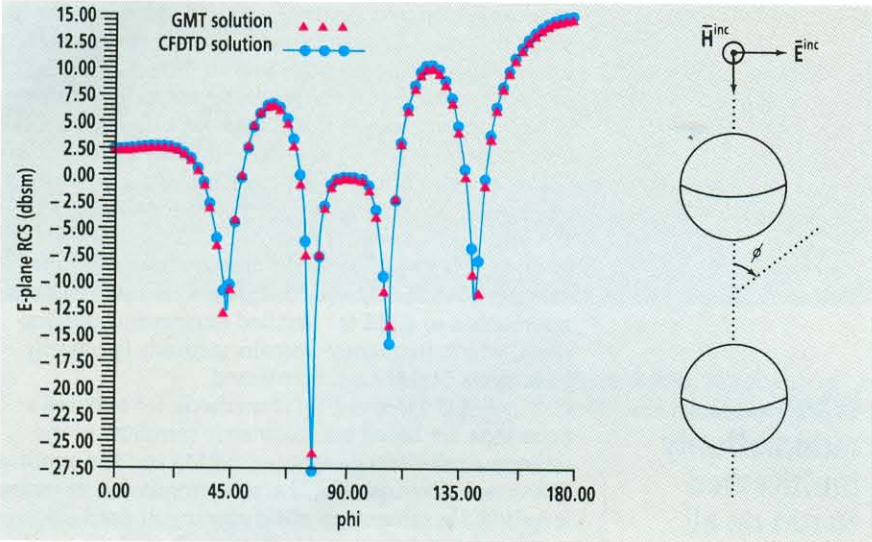

Figure 1. Agreement of FD-TD and generalized multipole technique (GMT)[6] bistatic RCS within 1 dB over a 42-dB range for a pair of 1-λ_0 diameter conducting spheres separated by a 1-λ_0 air gap.

to achieve nearly 100 percent concurrent utilization of all eight processors on a CRAY Y-MP system using Cray Research's Autotasking automatic multitasking software feature for three-dimensional FD-TD and FV-TD codes. Only relatively minor modifications were required to the original single-processor Fortran code.

Three-dimensional FD-TD validation example

Excellent validations of FD-TD have been obtained for three-dimensional problems that involve some of the key electromagnetic wave physics involved in RCS phenomena: near fields, monostatic RCS pattern, and bistatic RCS pattern. Here, we detail the results of a canonical, but difficult, bistatic RCS pattern validation.

Figure 1 shows the bistatic (side-scatter) RCS of a pair of 1-λ_0 diameter conducting spheres separated by a 1-λ_0 air gap.[5] The spheres are illuminated by a plane wave that propagates along a line connecting the centers of the spheres, and the bistatic pattern is observed in the plane of the incident electric field. (Note: when $\phi = 0^0$, the response is in the backscatter direction; that is, it is the monostatic RCS.) Here, the comparison is between FD-TD (using a mostly Cartesian, partially unstructured mesh to model conformally the spheres' surface curvatures) and an analytical approach well-suited for this problem, the generalized multipole technique (GMT).[6] Agreement between the two methods is excellent: within ± 1 dB (approximately ± 25 percent) over a wide 42-dB (16,000 to 1) range of RCS. This modeling accuracy occurs despite some tough electromagnetic field physics: the spheres interchange energy across the air gap in a tightly-coupled manner. For this problem, alternative FV-TD approaches using body-fitted meshes may introduce artifacts due to refraction and reflection of numerical waves propagating across global mesh distortions in the air-gap region. In fact, the two-sphere problem is a canonical example of difficult three-dimensional structures having substantial EM coupling between disjoint regions.

Electrically-large FD-TD application: jet engine inlet

The multiprocessing in-memory FD-TD code was used to model the RCS properties of an electrically large three-dimensional structure of engineering significance: a serpentine jet engine inlet (Figure 2). The overall system design problem involved sizing and shaping the engine inlet to meet specifications for both aerodynamics (thrust) and monostatic RCS at 10 GHz. The inlet was assumed embedded within a simple rectangular metal box coated with commercially available radar-absorbing material that provides approximately 30 dB (1000 to 1) suppression of electromagnetic wave reflections at 10 GHz. Thus, the FD-TD computed near-field and far-field electromagnetic response was primarily a function of the inside wall shaping of the inlet and not any exterior embedding.

As shown in Figure 2, the incident wave was assumed to propagate from right to left and be polarized with its electric field pointing across the narrow gap dimension (y direction) of the inlet. In this figure, the aperture of the inlet is located at the right, and the inlet is shorted by a conducting wall that represents the turbine assembly at the far left. With the box dimensions set at 30" x 10.5" x 10", the overall inlet and box target configuration spanned $25.4\lambda_0$ x $8.89\lambda_0$ x $8.47\lambda_0$ at 10 GHz. For this target, the FD-TD space cell size was 1/8" ($\lambda_0/9.43$); and the overall lattice had 270 x 122 x 118 cells that spanned $4608\lambda_0^3$ and contained 23,321,520 unknown vector field components. Starting with zero-field initial conditions, 1800 time steps were used (95.25 cycles of the incident wave) to march the field components to the sinusoidal steady state. The computer running time was only 3 minutes and 40 seconds per monostatic RCS calculation on the CRAY Y-MP system using automatic multiprocessing across eight processors (7.97/8 processor concurrency), yielding an average computation rate of 1.6 GFLOPS.

In addition to simple data for the RCS pattern, the FD-TD modeling provided details of the complex near field. Figure 3 shows the instantaneous distribution (positive and negative values) of the total gap (E_y) electric field component in a two-dimensional observation plane that cuts through the center of the three-dimensional engine inlet. This photograph was derived from a color videotape display of the propagating electric field penetrating the inlet, generated directly by the FD-TD time-stepping. The display was taken late in the time-stepping when the field had settled into a repetitive sinusoidal oscillation (standing wave). It may be possible to use such highly detailed near-field information (very difficult to obtain from measurements) to improve future RCS designs. Comparatively, if MoM were applied to model the same engine inlet,

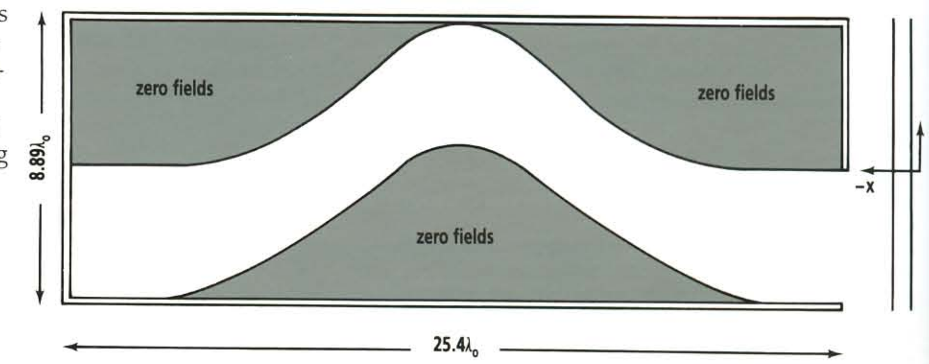

Figure 2. Geometry of engine inlet embedded in a rectangular metal box coated with commercial radar-absorbing material. The 10-GHz incident wave propagates from right to left.

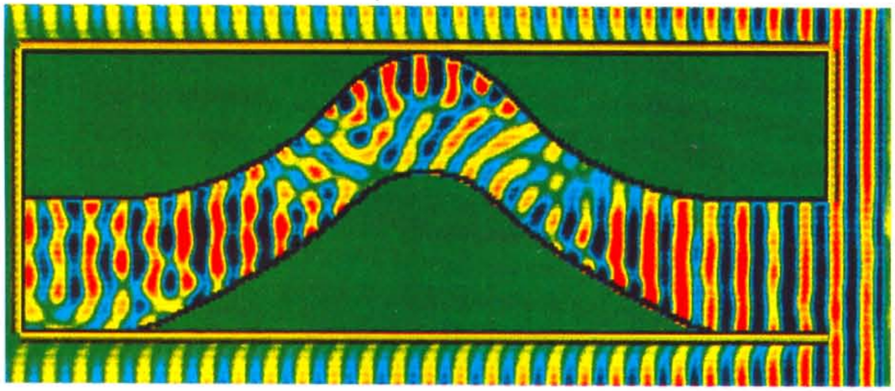

a complex-valued linear system involving approximately 450,000 equations would have to be set up and solved. This assumes a standard triangular surface patching implementation of the electric field integral equation[1], with the $1500\lambda_0^2$ surface area of the engine inlet discretized at 10 divisions per λ_0. Using the 2.1-GFLOPS out-of-memory subroutine for LU decomposition discussed earlier, the CRAY Y-MP system running time for this matrix would be about 2.6 years for 5000 monostatic angles. This compares to only about 12.7 days for FD-TD for the same number of monostatic angles, a speedup factor of 75. Additional problems involved in error accumulation in the LU decomposition and reliability of the computer system over the multiyear solution time probably would combine to render a traditional MoM solution useless for this target and those of similar or larger electrical sizes. We note also that MoM does not directly provide details of the penetrating near-field distribution.

Figure 3. FD-TD computed map of the instantaneous distribution of the total E_y vector field component in a two-dimensional cut through the z-center of the engine inlet geometry (at the sinusoidal steady state).

Present work and future directions

At present, grid-based time-domain CEM models of three-dimensional structures that span more than $30\lambda_0$ are being developed for the eight-processor CRAY Y-MP system. Work at this time addresses several areas:

- ☐ Automated mesh generation
- ☐ Multiprocessing out-of-memory software
- ☐ Subcell models for fine-grained structural features such as coatings
- ☐ Higher-order algorithms
- ☐ Application to nontraditional CEM areas, including design of ultra-high-speed electronic computer circuits, electro-optic components, and all-optical switches.

Extrapolating from benchmarks with the eight-processor CRAY Y-MP system, the next-generation CRAY Y-MP/16 system should provide a steady 10 to 13 GFLOPS computation rate for grid-based time-domain CEM codes when using automatic multitasking across 16 processors. The proverbial "billion-unknown" CEM problem (a three-dimensional computational volume of about $150{,}000\lambda_0^3$) could be completed in as little as 40 minutes per monostatic RCS observation. Multiprocessing out-of-memory software should enable even larger volumes to be modeled in their entirety. Using such software, the era of the "entire airplane in the grid" would be opened for a number of important aerospace systems for radar frequencies of 1 to 10 GHz. Automated geometry generation would permit CEM modelers to use structure databases developed by non-electromagnetics engineers, leading to lower design costs and the possibility of innovative design optimizations.

Acknowledgments

Daniel S. Katz and Allen Taflove were supported in part for this work by Cray Research, General Dynamics P.O. 4059045, National Science Foundation Grant ASC-8811273, and Office of Naval Research Contract N00014-88-K-0475. Jeffrey P. Brooks and Evans Harrigan were primarily responsible for the MoM, FD-TD, and FV-TD code optimizations and benchmarks. Thomas G. Jurgens of Fermilab, Batavia, Illinois, implemented the double-sphere FD-TD model.

About the authors

Daniel S. Katz is a graduate student pursuing his Ph.D. degree in the Electrical Engineering and Computer Science Department, McCormick School of Engineering, Northwestern University. He served an internship in Cray Research's Industry, Science & Technology Department during the summers of 1989 and 1990.

Allen Taflove is a professor of electrical engineering and computer science at Northwestern University. He recently was named an IEEE Fellow for his pioneering of FD-TD methods in CEM since 1973. He also is a national lecturer for the IEEE Antennas and Propagation Society and technical program chairman of the 1992 IEEE Antennas and Propagation Society International Symposium in Chicago.

Jeffrey P. Brooks is a manager of benchmark analysts in Cray Research's benchmarking department. He received an M.S. degree in mathematics and an M.B.A. degree from the University of Minnesota in 1985 and 1991, respectively. Before joining Cray Research in 1989, he was an instructor of mathematics at St. Olaf College in Minnesota and a technical representative at the Minnesota Supercomputer Center.

Evans Harrigan is a senior systems engineer at Cray Research, with 23 years of experience in supercomputing. He received a degree in mathematics from Michigan State University.

References

1. Rao, S. M., D. R. Wilton, and A. W. Glisson, "Electromagnetic Scattering by Surfaces of Arbitrary Shape," *IEEE Transactions on Antennas and Propagation,* Vol. 30, pp. 409-418, May 1982.
2. Newman, E. H. and D. M. Pozar, "Electromagnetic Modeling of Composite Wire and Surface Geometries," *IEEE Transactions on Antennas and Propagation,* Vol. 26, pp. 784-789, November 1978.
3. Taflove, A., "Review of the Formulation and Applications of the Finite-Difference Time-Domain Method for Numerical Modeling of Electromagnetic Wave Interactions with Arbitrary Structures," *Wave Motion,* Vol. 10, pp. 547-582, December 1988.
4. Shankar, V., A. H. Mohammadian, and W. F. Hall, "A Time-Domain Finite-Volume Treatment for Maxwell's Equations," *Electromagnetics,* Vol. 10, pp. 127-145, January-June 1990.
5. Jurgens, T. G., "Conformal Finite-Difference Time-Domain Modeling," Ph.D. dissertation, Northwestern University, EECS Dept., 1990.
6. Ludwig, A. C., General Research Corp., Santa Barbara, California, private communication with T. G. Jurgens.

ASTROS: the Automated Structural Optimization System

Raymond M. Kolonay and V. B. Venkayya, Wright-Patterson Air Force Base, Ohio
Mike Long, Cray Research, Inc.

The Automated Structural Optimization System (ASTROS) is a multidisciplinary computer program for the preliminary design of aircraft and spacecraft structures. It was developed by the Aeronautical System Division's Wright Laboratory at Wright-Patterson Air Force Base. Northrop Aircraft Division, Universal Analytics, Inc., and Kaman Avidyne Corporation developed the system under contract with Wright Laboratory's Flight Dynamics Directorate.

The ASTROS program is unique in many respects. It integrates structural and aerodynamic analysis, controls, and optimization for aircraft and spacecraft design. Because the program integrates formal mathematical optimization and sensitivity analysis, it can reduce design time and costs significantly. It also facilitates effective communication between engineers working in different disciplines. Although originally designed for aerospace applications, ASTROS also can be used in related areas, including the automotive, marine, civil engineering, and transportation industries. When run on supercomputers, such as those from Cray Research, ASTROS solves integrated large-scale modeling problems quickly enough to be a practical tool for engineers and designers.

The heart of the ASTROS program is the structural analysis module, which is based on the finite element method and performs statics, statics with inertial relief, normal modes, transient response, and frequency response analyses. The module incorporates commonly used finite elements, such as rod, beam, membrane, bending, and solid elements. The membrane and bending elements support full composite modeling capability. The aeroelastic modules calculate steady and unsteady aerodynamic loads and perform static aeroelastic, flutter, gust, and nuclear blast response analyses. The aerodynamic loads are based on paneling methods.

Figure 1. Finite element model of ASTOVL fighter plane.

The design optimization module makes ASTROS useful for design as well as for engineering analysis. Using mathematical optimization techniques, ASTROS can find the lightest-weight design that meets a given design criterion. To accomplish this, ASTROS uses the well-known method of modified feasible directions found in the MDOT engineering program from Vanderplaats, Miura, and Associates, Inc. Optimality criteria methods are used to to carry out the optimization once the problem is properly posed for efficient solution. Strength, deflection, frequency, flutter and divergence constraints, aeroelastic trim parameters, and aeroelastic stability derivatives can be satisfied while designing the finite elements. The design variables in ASTROS consist of element thicknesses, cross-sectional areas, and concentrated mass values. The elements can be designed individually, in groups, or by using manufacturable shape functions.

The optimization techniques used in ASTROS require the calculation of the derivatives of the response functions and the objective function with respect to the design variables. This sensitivity analysis is based on analytical derivatives. Both direct and adjoint variable methods are used and are applied appropriately to ensure computational efficiency.

Along with its multidisciplinary analysis and design capabilities, ASTROS has a state-of-the-art system architecture that gives users broad flexibility in applying the program to their needs. This includes the ability to add engineering modules, bulk data entries, database entities, and system messages with relative ease.

Large-scale application

The ASTROS program recently was run on a CRAY Y-MP computer system to model an Advanced Short Take Off and Vertical Landing (ASTOVL) fighter plane. The ASTOVL concept describes a single-engine fighter that has a short conventional take off and can land vertically. The finite element model (Figure 1) was a right-half model of the complete vehicle with centerline plane-of-symmetry boundary conditions. The model had 870 grid points and 1389 elements with the predominant elements being quadrilateral and triangular membranes. It included 1362 local physical variables, including cross-sectional areas, membrane thicknesses, ply thicknesses, and concentrated masses. These variables were linked to 173 global design variables that were submitted to the mathematical optimizer. Composite materials were used to model the majority of the wing and fuselage super- and substructures.

The objective of the design study was to minimize the weight of the structure and satisfy the mission performance requirements of the aircraft. The study concentrated on two critical areas of the mission requirements: the elimination of flutter in the supersonic flight envelope (Mach 1.5 at 15,000 feet) and the satisfaction of strength requirements for two landing gear static load cases (8.51 g impact ultimate load in the nose gear, 6.9 g impact ultimate load in the main landing gear). Figure 2 is a velocity vs. damping (v-g) diagram for the initial design, and Figure 3 shows the initial thickness/cross-sectional areas of the physical variables. The initial design satisfies the strength requirements, but Figure 2 shows that the flutter constraint is severely violated. Using the ASTROS program, engineers were able

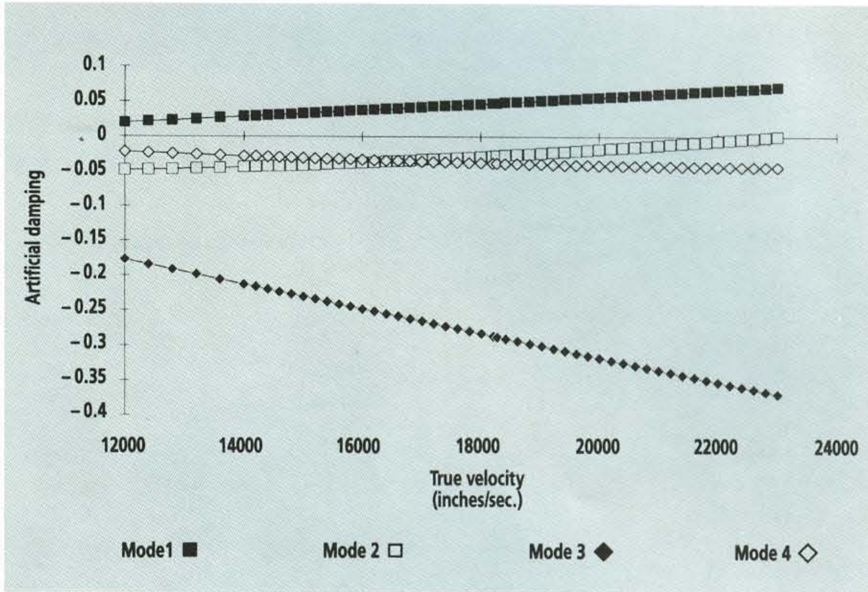

Figure 2. Velocity vs. damping results for the initial ASTOVL design.

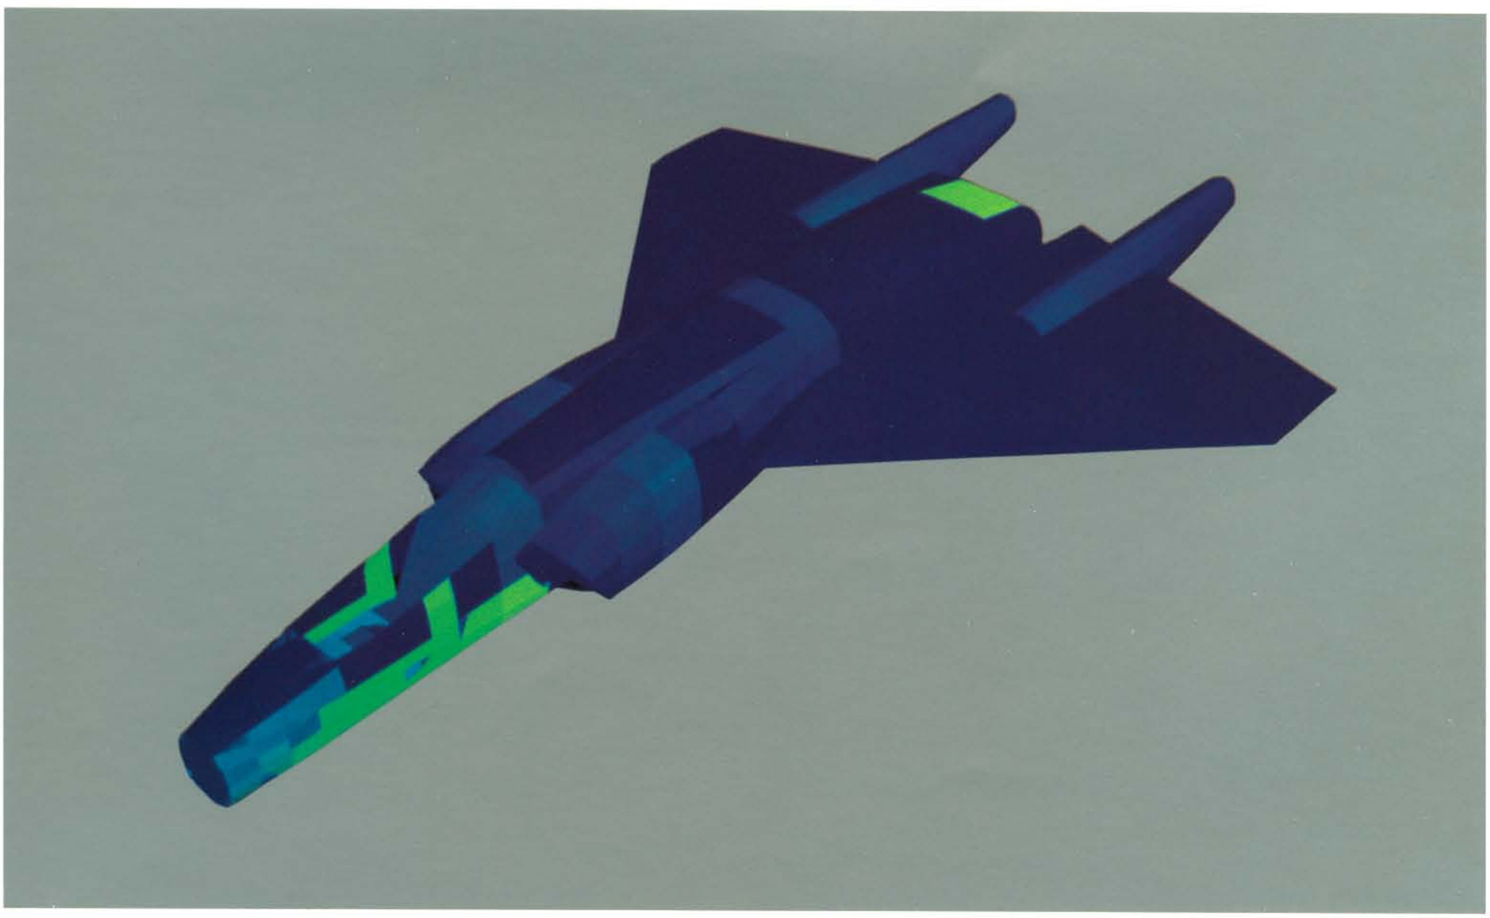

Figure 3. Initial element thickness of the first ply layer of the composite aircraft skins. Thickness increases from blue to green to red.

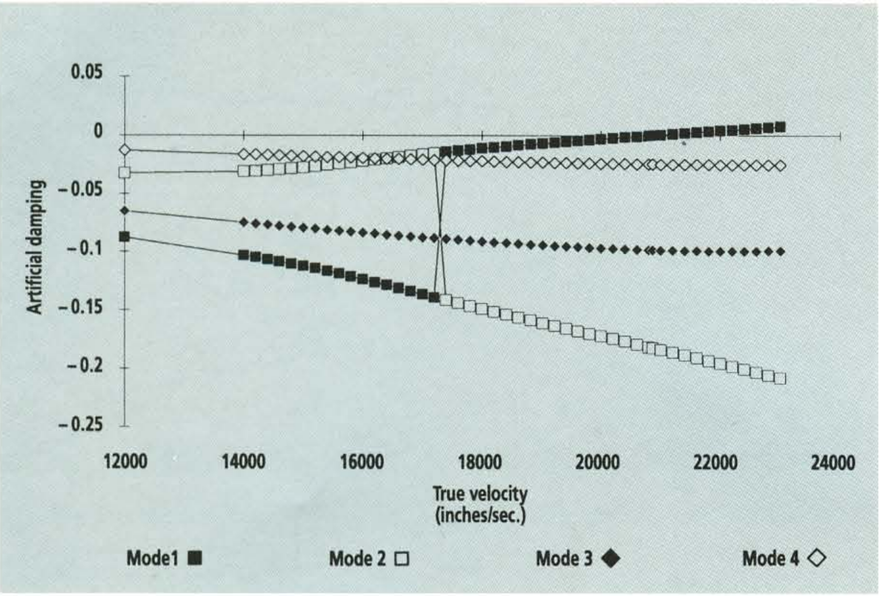

Figure 4. True velocity vs. damping results for the optimized ASTOVL design.

to modify the design to satisfy both the flutter and strength constraints. Figure 4 illustrates that the new flutter speed occurs outside the flight envelope. Figure 5 shows the final distribution of materials on the aircraft. The ASTROS program required 12 iterations and only 2 hours and 13 minutes of CPU time on the CRAY Y-MP system to complete the design.

Although the finite element and aerodynamic models of the ASTOVL are of modest size, the coupling of multiple engineering disciplines and the iterative optimization solution generates a computationally intensive problem. Solving realistic large-scale problems with ASTROS in a reasonable amount of time requires the computational power of supercomputers such as those from Cray Research.

Acknowledgments

The ASTOVL model development and original design studies were performed by J. A. Hagen and D. J. Neill of the Aircraft Division of Northrop Corporation, under contract to Wright Laboratory.

About the authors

Raymond M. Kolonay is an aerospace engineer in the analysis and optimization branch, Structures Division, Flight Dynamics Directorate at Wright Laboratory, Wright-Patterson Air Force Base, Ohio. He received a B.S. degree in civil engineering in 1985 and an M.S. degree in civil engineering in 1987 from The Ohio State University.

V. B. Venkayya is the principal scientist in the optimization branch, Structures Division, Flight Dynamics Directorate at Wright Laboratory, Wright-Patterson Air Force Base, Ohio. He received his undergraduate education in India and M.S. and Ph.D. degrees from the University of Missouri at Columbia and the University of Illinois at Urbana-Champaign, respectively. He is a Fellow of Wright Laboratory, an Associate Fellow of the AIAA and Senior Member of the ASME and ASCE professional associations.

Mike Long is manager of the structural applications group in Cray Research's Industry, Science & Technology department. He received B.S. and M.S. degrees in civil engineering from Brigham Young University.

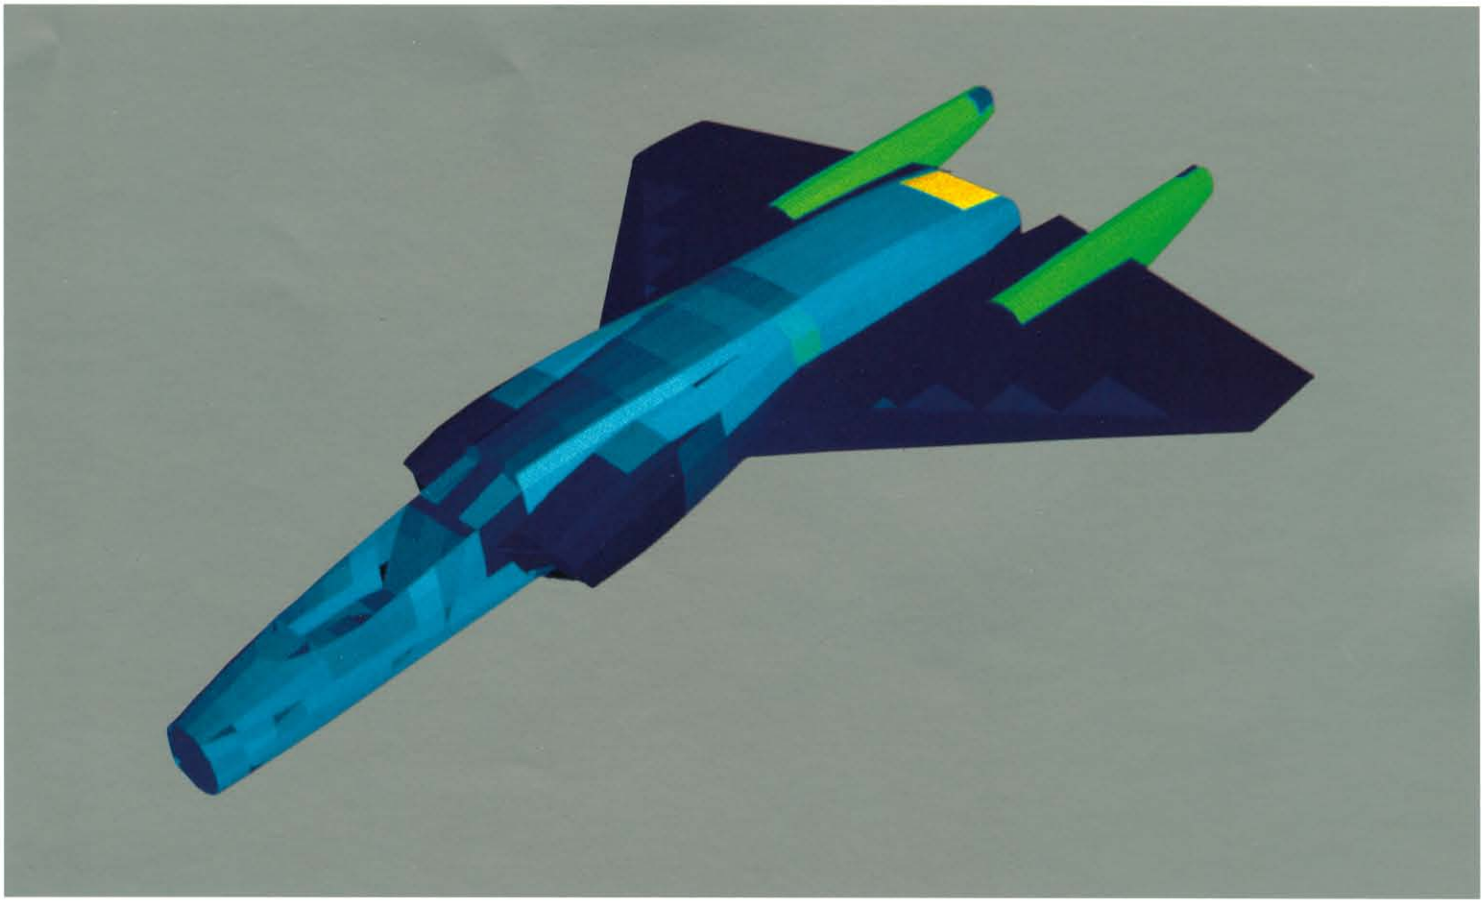

Figure 5. Final element thickness of the first ply layer of the composite aircraft skins. Thickness increases from blue to green to red.

Supercomputer modeling for centrifugal compressor aerodynamics

Jon S. Mounts, United Technologies Research Center, East Hartford, Connecticut
Joost J. Brasz, Carrier Corporation, Syracuse, New York

United Technologies Research Center is developing a three-dimensional, Navier-Stokes algorithm for the unsteady analysis of centrifugal compressor aerodynamics. The algorithm will be used to enhance the design of more efficient, stabler, and quieter centrifugal compressors. This development effort depends on the power of supercomputers, such as the CRAY Y-MP system, to provide the necessary computational resources.

The centrifugal compressor has a wide range of uses, from large industrial chillers for commercial air conditioning systems, to fuel and oxidizing pumps in the space shuttle's main engine. It has long been characterized as one of the most complicated fluid-dynamic components ever designed, but its versatility is due to its ability to attain high pressure ratios within a relatively short distance.

Past design work

In the past, industry design methodology for centrifugal compressor applications in commercial air conditioning systems was based on analyses ranging from the use of largely empirical techniques to quasi-three-dimensional, steady codes. These essentially inviscid programs were used to calculate design parameters, such as aerodynamic blade loading and internal flow diffusion. Then, the centrifugal impeller design was modified until the design parameters reached certain acceptable values as determined by experience. While this methodology yielded some impressive results, it gave little insight into three-dimensional, viscous effects and no information about the flowfield interaction between stationary and rotating components within the compressor.

The problem

If the effects of viscosity and component interactions are not properly accounted for during the design phase, the finished compressor may be noisy, inefficient, and have a limited range of operation. To design more efficient, stabler, and quieter compressors, an advanced analysis tool is needed. United Technologies Research Center is working on such a tool: a three-dimensional Navier-Stokes algorithm for the unsteady analysis of centrifugal compressor aerodynamics.

At the core of this new tool is the three-dimensional Navier-Stokes algorithm. This algorithm was developed at United Technologies Research Center, based on a technique by M. M. Rai for unsteady rotor/stator interaction for applications in multicomponent flowfield interactions within a centrifugal compressor.[1,2,3]

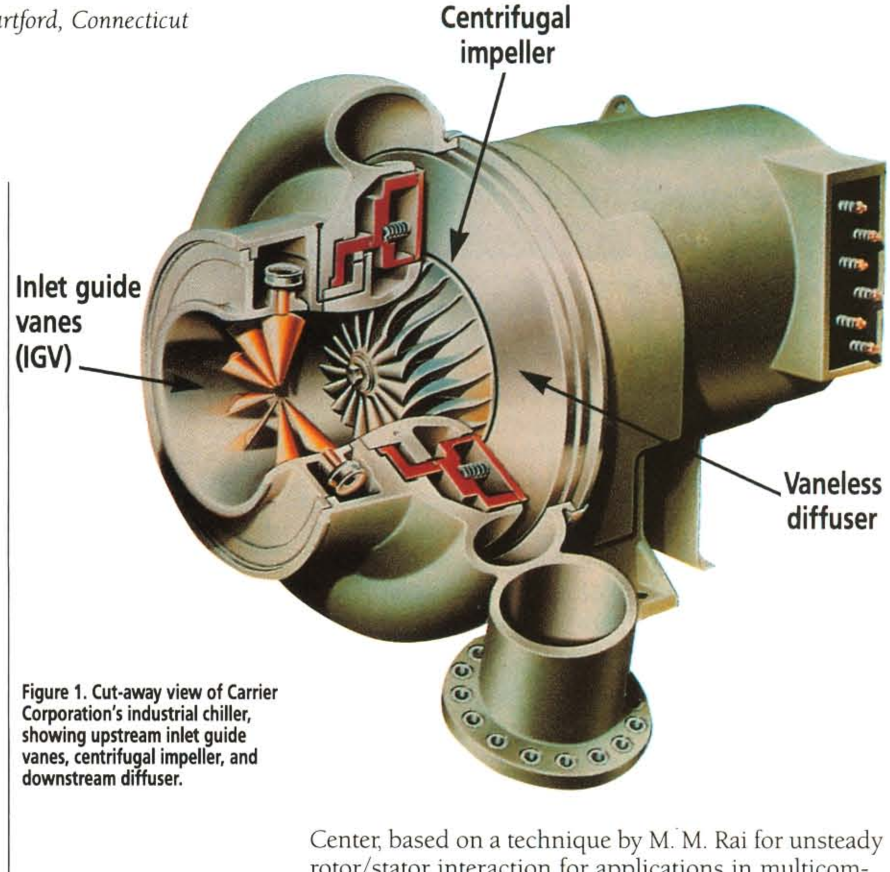

Figure 1. Cut-away view of Carrier Corporation's industrial chiller, showing upstream inlet guide vanes, centrifugal impeller, and downstream diffuser.

This unsteady analysis is needed because distortions in the centrifugal impeller exit velocity profile, interacting with the diffuser, cause highly unsteady flow phenomena with large flow angle variations. Both the efficiency and operating range of the centrifugal compressor appear to be directly linked to the behavior of the flow as the centrifugal impeller interacts with other components of the machine, particularly the upstream inlet guide vane and the downstream diffuser (Figure 1). The goal, then, is to design centrifugal impellers that produce more nearly uniform exit profiles, thereby reducing the flow angle fluctuations to the diffuser.

Critical interactions between compressor components must be incorporated into the analysis to produce realistic numerical simulations of the flowfield. Such modeling requires extremely fine mesh resolution to capture the fluid physics as the flow develops through the impeller passage, from the inducer, through the radial bend, and on past the exit plane to interact with the downstream diffuser. Because the flowfield is unsteady, the CFD analysis of the centrifugal compressor presents a computational challenge.

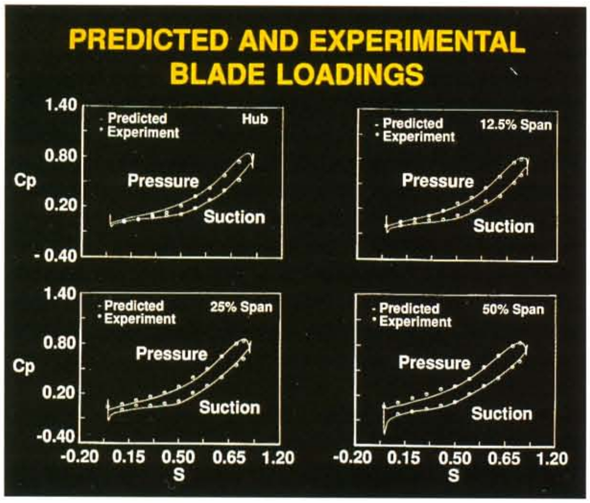

Figure 2. Comparison of blade loading in surface pressure coefficients for the unshrouded impeller (with tip leakage) at design conditions with previously obtained experimental data.

This study has focused on the flowfield physics of a backswept centrifugal impeller using a time-dependent, three-dimensional, Navier-Stokes algorithm. Advanced scientific visualization techniques have been used to analyze the development of the flow structures in the impeller passage due to tip leakage at both on- and off-design conditions. For validation, blade loading comparisons have been made between solutions obtained for an unshrouded impeller (with tip leakage) and available experimental data.[4] Off-design analyses also have been performed to investigate the overall flowfield patterns within the impeller passage.[3]

The algorithm

The governing equations in the present analysis are the time-dependent, three-dimensional, Navier-Stokes equations incorporating the thin-layer assumption. The numerical procedure employs an implicit, upwind, finite-difference scheme. In this scheme, inviscid fluxes are discretized using a third-order, spatially accurate scheme, and viscous fluxes are discretized using standard central differences. Closure is obtained through an algebraic turbulence model with a modified mixing length assumption to account for corner flow through the passage and tip leakage effects near the shroud. No-slip boundary conditions are applied at solid surfaces, and the theory of characteristics is used to determine both inlet and exit flow conditions. An alternating direction-approximate factorization method is used to compute the rate of change of the dependent variables. Newton subiterations are employed to eliminate factorization and linearization errors.

To obtain inlet guide vane-impeller-diffuser interaction, the computational domain was broken into three distinct blocks: 1) the upstream block of mesh, containing the inlet guide vane flowfield; 2) the middle block of mesh, containing the centrifugal impeller; and 3) the downstream mesh, containing the diffuser flowfield (vaneless, in this case). In this model, the upstream and downstream blocks of mesh remain stationary, and the middle block rotates relative to the two stationary blocks, thus allowing for the proper flowfield interaction between centrifugal compresser components.

Analysis

The centrifugal impeller used in this study is Carrier Corporation's standard 17DK industrial chiller centrifugal compressor. It consists of 17 equally spaced blades with 30 degrees of backsweep. Figure 2 shows blade loadings for the unshrouded (with tip clearance) impeller versus experimental data at several spanwise locations. The distorted exit velocity profile of the classic "jet-wake" causes disturbances downstream from the impeller, reducing efficiency and operating range, and increasing noise.[5] Understanding this phenomenon is crucial for creating quiet and efficient designs.

Figure 3 shows the highly three-dimensional nature of the impeller flowfield in terms of rotary static pressure. Figure 4 shows the the development of the flowfield in terms of entropy contours for the unshrouded (with tip clearance) impeller. The solutions show that a pocket of low velocity, low momentum fluid has been driven through the tip clearance gap and distributed along the shroud surface. This results from pressure to the suction side of the blade.

Further study

Further analyses are underway to investigate the development of the flow structures through the impeller exit plane to their interaction with downstream, vaneless diffusers. In addition, upstream interactions with pre-rotation and swirl imposed by various inlet guide vane settings also are being investigated. Because the number of inlet guide vanes and diffusers typically is not equal to the number of impeller passages for a given centrifugal compressor design, multiple impeller passages have to be modeled to solve the equations on a periodic structure.

Figure 3. Three-dimensional nature of centrifugal impeller flowfield (rotary static pressure contours).

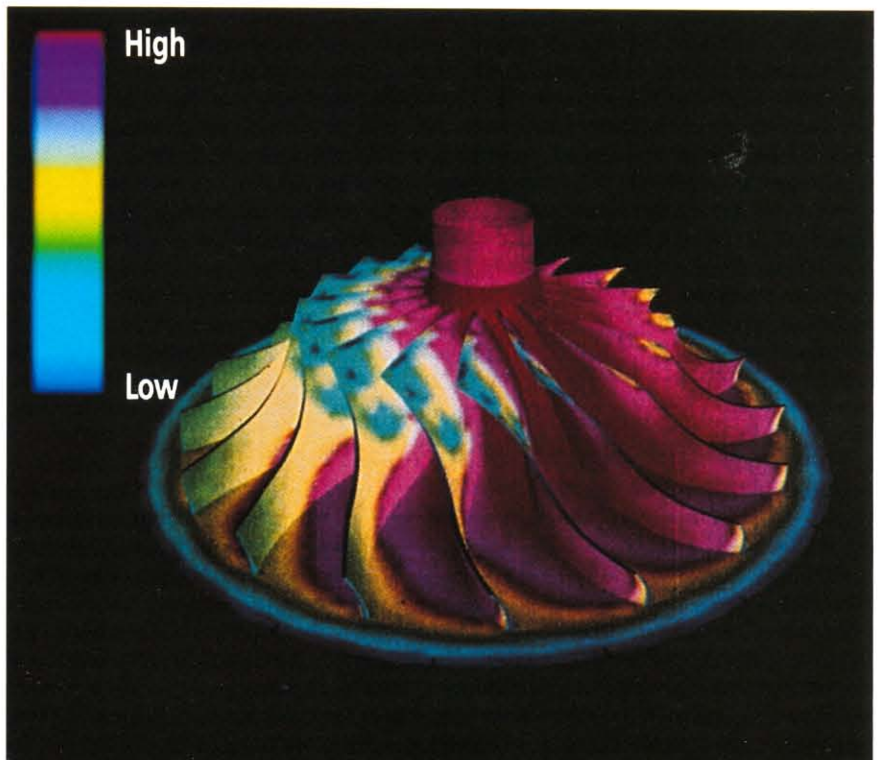

The global features of the impeller have been modeled on a CRAY Y-MP system using approximately 65,000 mesh points (75 by 25 by 35) for a single impeller passage, requiring about 10 million words of memory. Calculations typically require 18 to 20 CPU hours on a single processor to reach an acceptable time-accurate solution.

The additional mesh resolution required to define more complicated geometries both upstream and downstream could potentially increase the memory requirements to put the total number of grid points at 750,000 to 1,000,000 and the memory at close to 100 million words. In addition, the increased mesh and complexity of the solution is likely to raise the CPU requirement to 180 to 200 CPU hours on a single processor.

The ability to accurately resolve the unsteady, three-dimensional flowfields associated with the interaction of centrifugal compressor components can have a dramatic impact on compressor design. Resources such as the CRAY Y-MP system and efficient use of its advanced architectures, including solid-state storage and parallel processing, are critical to the successful completion of this and many other projects on the leading edge of technology.

About the authors

Jon S. Mounts is an associate research engineer at United Technologies Research Center investigating centrifugal compressor aerodynamics for Carrier Corporation. He received his B.Sc. degree in aeronautical engineering from the U.S. Air Force Academy and his M.S. degree from the University of Florida. He is currently pursuing his Ph.D. degree in applied mathematics at Yale University.

Joost J. Brasz is a senior staff engineer at Carrier Corporation and is responsible for centrifugal chiller development engineering and the Carrier research programs at United Technologies Research Center. He received both his B.Sc. and M.S. degrees in physics from the Free University of Amsterdam, and his Ph.D. degree in technical physics at Eindhoven University of Technology, the Netherlands.

Acknowledgments

The authors thank Charlie Finan of Cray Research for contributions regarding the efficient use of the CRAY Y-MP system architecture, Joe Caspar and Tom Barber of United Technologies Research Center and Meredith Spears of Cray Research for their support of the REACH program, and Daniel J. Dorney and H. David Joslin of United Technologies for their assistance. Scientific visualization was provided by David Edwards of the United Technologies Research Center.

References

1. Rai, M. M., "Navier-Stokes Simulations of Rotor/Stator Interaction Using Patched and Overlaid Grids," *Journal of Propulsion and Power*, Vol. 3, No. 5, September-October, 1987.

2. Dorney, D. J., and R. L. Davis, "Centrifugal Compressor Impeller Aerodynamics: A Numerical Investigation," ASME Paper 90-GT-213, June 1990.

3. Mounts, J. S., D. J. Dorney, and J. J. Brasz, "Flowfield Analysis of a Backswept Centrifugal Impeller," to be published by AIAA, June 1991.

4. Joslyn, H. D., J. Brasz, and R. P. Dring, "Centrifugal Compressor Impeller Aerodynamics: An Experimental Investigation," ASME Paper 90-GT-214, June 1990.

5. Eckhardt, E., "Detailed Flow Investigations within a High-Speed Centrifugal Compressor Impeller," ASME Paper 76-FE-13, January 1976.

Figure 4. Development of the flowfield in the unshrouded (with tip leakage) impeller passage at design conditions (entropy contours).

Supercomputer applications in explosives research

David A. Jones, David L. Smith, and Rodney A. J. Borg
Materials Research Laboratory, Defence Science and Technology Organisation, Melbourne, Australia

The numerical modeling of explosions requires the solution of highly nonlinear equations through the use of sophisticated finite difference or finite element continuum mechanics computer codes. Many of these codes are now either fully or partly vectorized and are ideally suited for use on supercomputers. Scientists at the Materials Research Laboratory (MRL) in Melbourne, Australia, have recently been using a CRAY X-MP/14 system at Leading Edge Technologies, Ltd., to understand aspects of explosive cutting devices, detonation transfer, and the molecular dynamics of sedimentation in explosive formulations.

Explosive cutting devices

The Explosives Division of MRL has investigated explosively driven cutting devices for several years. One device that has received a lot of attention is known as ladder fracture tape, which consists of a channel-sectioned molding made from a neoprene rubber compound filled with plastic or moldable explosive. Figure 1 shows a typical design. Designed originally to cut pipe casings, fracture tape has been optimized for cutting thin metal plates. It also has been used successfully for military purposes such as ordnance disposal and demolition.

Figure 1. Schematic diagram of a section of ladder fracture tape.

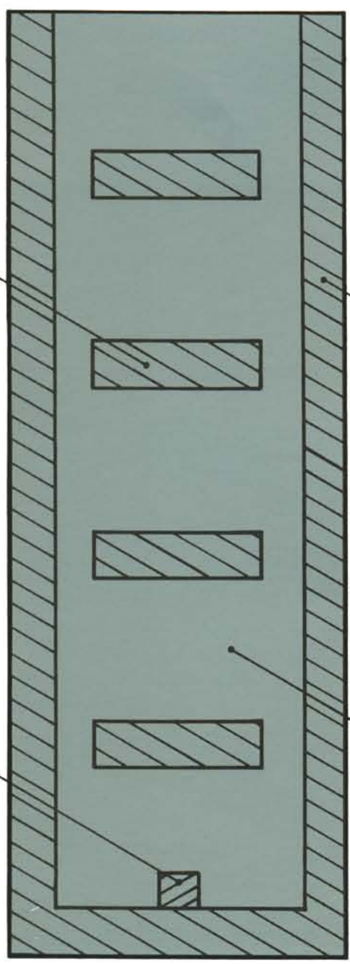

Ideally, ladder fracture tape works through a sequence of events: the explosive is detonated at the center of one end, and a complex pattern of interacting shock and detonation waves travels along the tape. Each successive barrier absorbs the shock caused by the impact of the detonation wave but also allows the detonation to proceed through the gaps between it and the side walls of the tape. If the gap is larger than the minimum thickness for maintenance of stable detonation in the explosive, the detonation passes the corner of the barrier and expands towards the center of the tape. When the two fronts from each side of the barrier collide at the centerline of the tape, the resultant pressure is much higher than the normal detonation pressure of the explosive. If the tape is placed on a metal surface and detonated, the detonation front generates shock waves within the metal as it propagates along the tape. These waves are also reflected from the rear surface of the metal as tension waves, and the interaction between the incident and reflected waves creates a spall surface within the metal where the metal's tensile strength has been exceeded. Figure 2 shows the effect of the detonation of a section of ladder fracture tape placed in contact with a 10-cm-thick lead witness plate. The enhanced damage pattern along the centerline of the tape is clearly visible.

Two design constraints are important for ladder fracture tape to function optimally: as the point of collision of the two detonation waves moves along the centerline, the collision angle is reduced, thereby reducing the enhanced collision pressure; and the next barrier in the tape must be placed so that this process is repeated before the enhanced pressure decays too much. Numerical simulations of the behavior of the propagating detonation in a section of ladder fracture tape previously have been run at MRL using the reactive hydrocode 2DL.[1] These simulations indicated that for some geometries the undesired second mechanism described above can operate. These simulations recently have been repeated and extended at MRL using the HULL software package on the CRAY X-MP/14 system.

The HULL package is a self-contained system of programs for solving problems of hydrodynamic flow in two and three dimensions.[2] HULL allows calculations in Eulerian or Lagrangian space or a combination of the two. Originally developed under contract to the U.S. Air Force, the HULL system generates Fortran code appropriate to the problem under study and optimizes the code for many computers. The size and architecture of the HULL system has prevented full vectorization (the system contains 125,000 lines of Fortran), but "hot" routines contain specific vector instructions for Cray Research's CF77 Fortran compiling system, enhancing performance on Cray Research computer systems. Figure 3 shows pressure profiles from a HULL simulation of fracture tape, where the "ladder" runs from the bottom left to the top right of the image. The calculation was performed in axisymmetric two-dimensional Eulerian space and used 80,000 computational cells.

Detonation transfer

The transmission and reflection of shock and detonation waves through layered explosives is important in several areas of military interest. When solid explosives are used to study detonation transfer between layers, the complex interactions that develop are hard to observe. However, in gaseous layers the problem is much simpler because pulsed laser Schlieren photography can be used to study complex structures in gaseous explosives. At the University of Michigan the technique has been combined with a specially designed shock tube to visualize transmitted shock and detonation structures in layered detonations for an extensive combination of explosives.[3] MRL is collaborating with U.S. researchers at the University of Michigan and the Naval Research Laboratory (NRL) to understand the basic physics of these interactions through numerical simulation techniques. The computer code used for these simulations was developed at NRL and uses the Flux-Corrected Transport (FCT) algorithm for the solution of the compressible flow equations.[4]

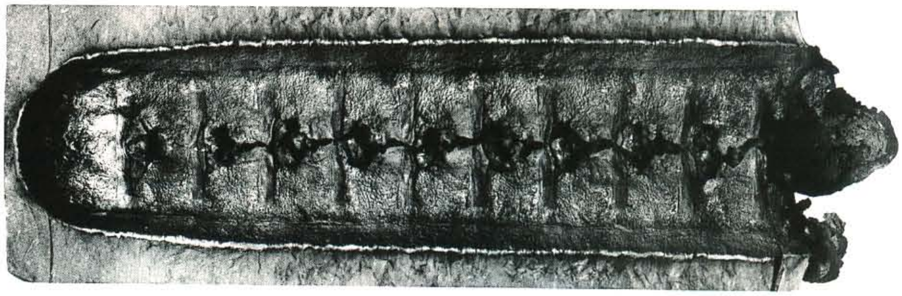

Figure 2. Damage to a 1-cm-thick metal plate after detonation of a section of ladder fracture tape placed along the surface. The imprint of the explosive ladder is evident, as is the enhanced damage pattern along the centerline.

Study of the decay, transfer, or re-ignition of detonation as a reactive shock travels from one explosive material into another requires careful resolution of the reaction zone and a detailed understanding of the chemistry of the processes that convert reactants into products. The experiments conducted at the University of Michigan have been simulated without including the full details of the chemical reactions in the code. Instead, a two-stage induction parameter model is used that is calibrated from accurate solutions of the full set of chemical rate equations and reproduces the essential features of the induction period and energy release processes. The gaseous reactants and products are modeled using a perfect gas equation of state. The induction parameter model is then coupled to the time-dependent equations that describe the conservation of mass, momentum, and energy for the reacting flow. The model is implemented on a two-dimensional rectangular finite difference mesh, and the equations are solved using the FCT algorithm and direction- and timestep-splitting. The code is fully vectorized and runs approximately 60 times faster on the CRAY X-MP/14 system than on the MRL VAX 8700 computer.

In the experiments, a detonation is initiated in one channel of a three-meter-long, two-channel steel detonation tube, and pressure switches along the length of the channel ensure that the detonation velocity quickly reaches a steady value. The other end of the detonation tube leads to a test cell that contains trans-

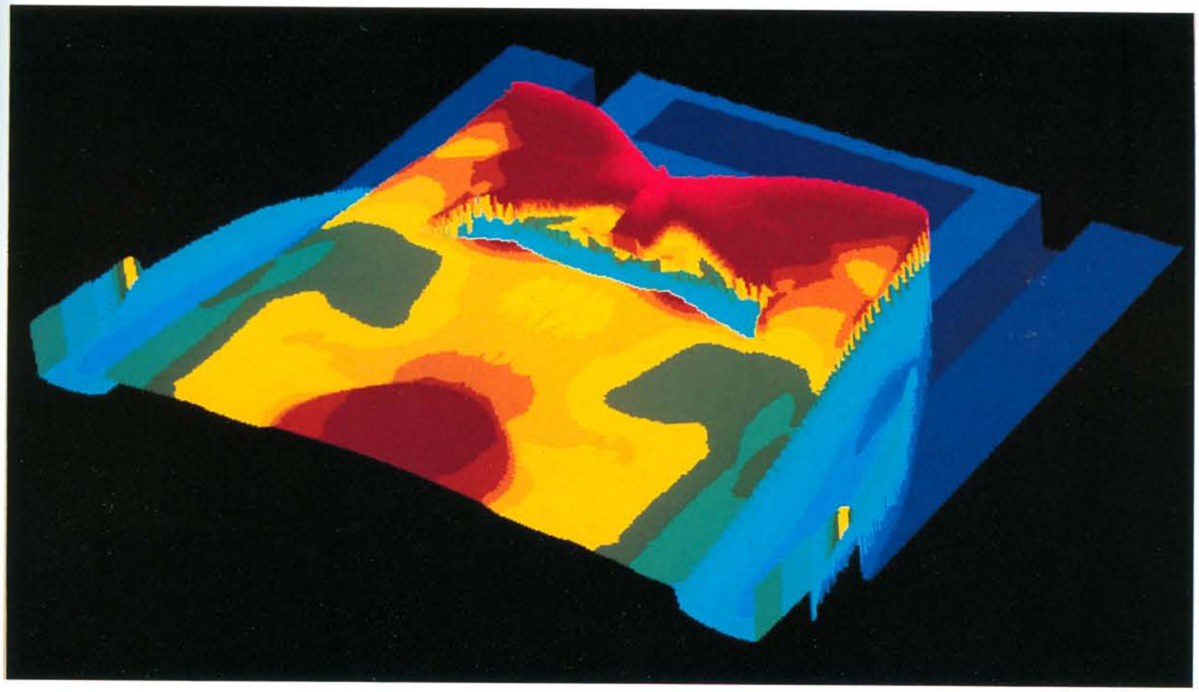

Figure 3. Pressure profiles from a HULL simulation of ladder fracture tape. The "ladder" runs from bottom left to top right.

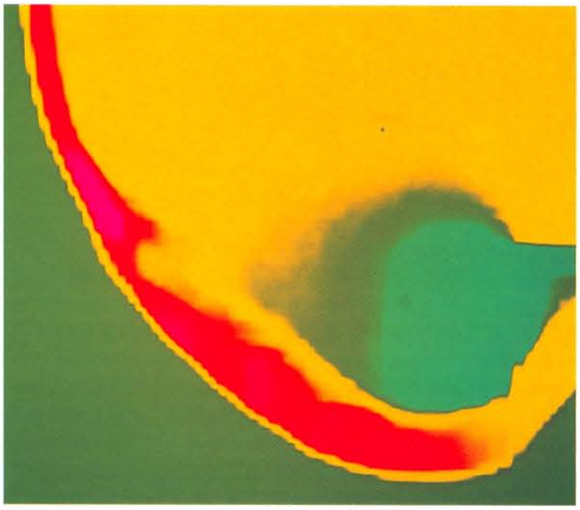

Figure 4. Simulated shock density pattern formed by a detonation in an argon-diluted stoichiometric mixture of hydrogen and oxygen as it emerges from the end of a detonation tube.

parent side walls so that laser Schlieren photographs can be taken. In the test cell, the dividing wall between the two channels has been removed to allow the detonation to interact with another material in the second channel. Figure 4 shows the simulated shock density pattern in the test cell produced by the code when both channels contain an argon-diluted stoichiometric mixture of hydrogen and oxygen. The green at the left and bottom of the picture represents undisturbed gas; the yellow, the detonation products; and the varying shades of red, the reaction zone. The end of the dividing wall is visible as the short horizontal line segment between the green and yellow areas halfway up the right side of the picture. The reaction depicted here is very transient, and if expansion were allowed to continue into an unconfined space the shock front would separate further from the reaction zone and the detonation would fade out. In the simulation sequence from which Figure 4 was taken, the reactive shock was allowed to interact with a rigid boundary at the bottom of the grid. This interaction resulted in an irregular reflection, and the subsequent high pressure reinitiated detonation behind the shock. The numerical simulations so far have reproduced many of the features seen experimentally and have assisted in the interpretation of the experimental results.[5]

Molecular dynamics of sedimentation

As these examples have shown, the major applications of supercomputers in explosives research lie in the area of computational fluid dynamics and large continuum mechanics hydrocodes. However, other areas of explosives research also can benefit from the power of supercomputers. Military explosives typically contain a mixture of explosives of varying particle sizes, various inert materials to enhance the specific impulse, and plasticisers to decrease the sensitivity of the explosive. The viscosity of the resulting mixture depends heavily on the relative proportions of components of different particle sizes and must be kept within certain limits so that the explosive can be poured easily. Consequently, much time and effort are invested in formulating explosives with low viscosity but high explosive output.

Researchers at MRL are investigating new approaches to the formulation problem by examining sedimentation during the solidification of the explosive Composition B, which is a 60/40 mixture of hexahydro trinitro-s-triazine (RDX) and trinitrotoluene (TNT). Ordnance is filled by a casting process in which Composition B is heated above the melting point of TNT (82° C), poured into the shell, and allowed to cool and solidify. During this process the heated mixture is essentially a suspension of solid RDX particles in liquid TNT, and sedimentation of RDX can occur in the shell during cooling and solidification of the TNT. The sedimentation will lead to localized regions of higher RDX content and higher sensitivity, because RDX is more sensitive to initiation than TNT.

A molecular dynamics computer simulation technique was used at MRL to study this sedimentation process. In molecular dynamics calculations, the forces on all particles are calculated, and the particles are moved according to Newton's equations of motion.

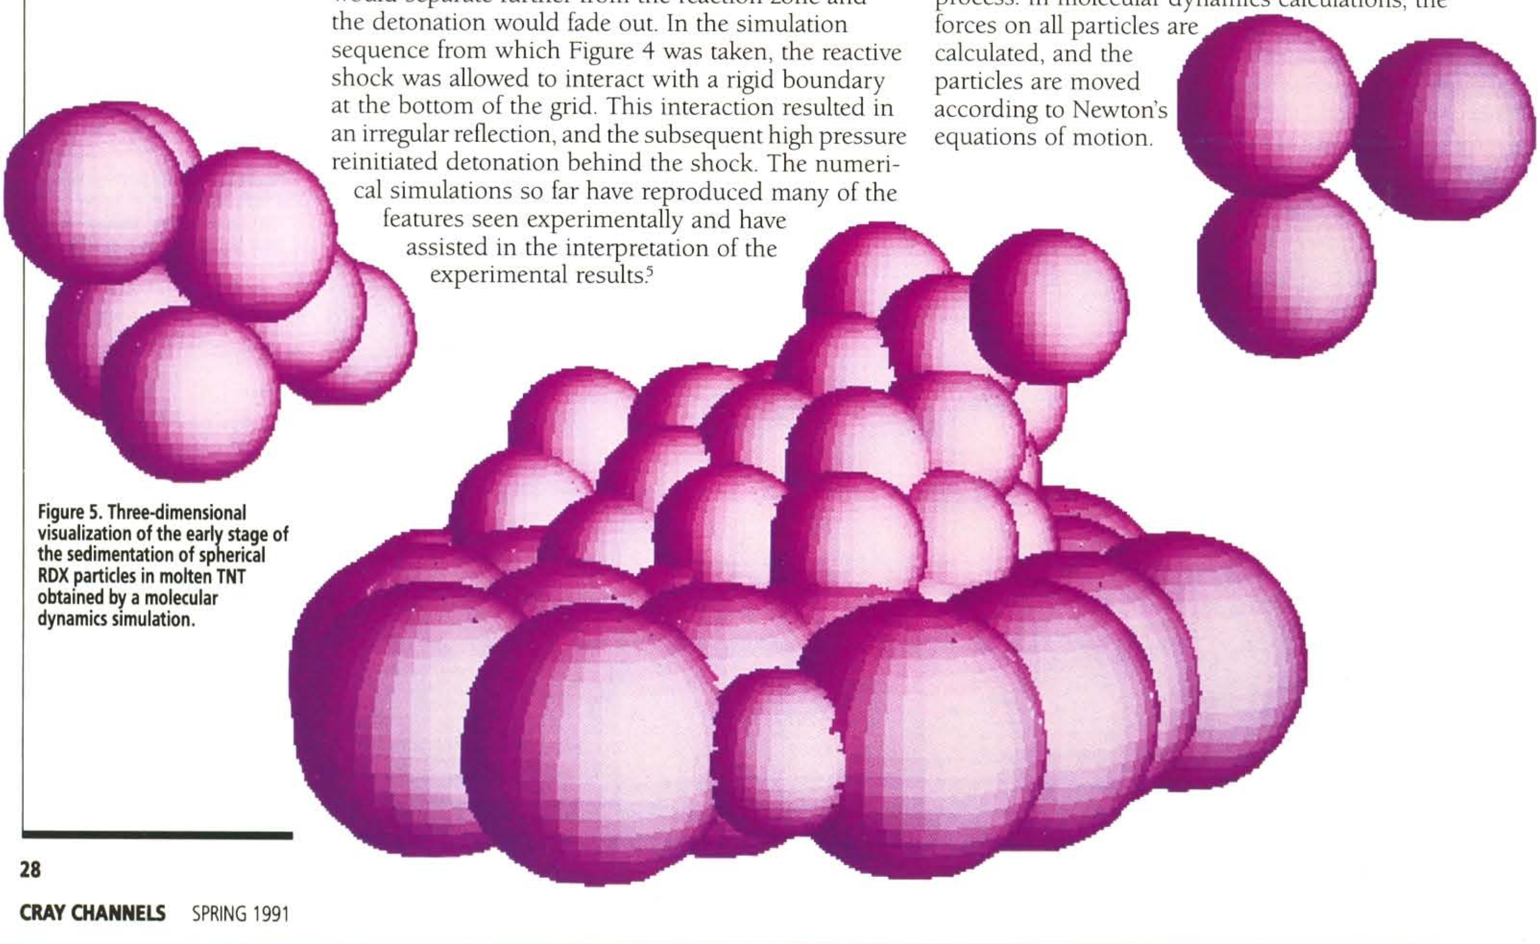

Figure 5. Three-dimensional visualization of the early stage of the sedimentation of spherical RDX particles in molten TNT obtained by a molecular dynamics simulation.

The initial state for the sedimentation system consists of N particles dispersed homogeneously in a viscous medium and subject to gravity and interparticle forces. The final state consists of a completely sedimented system in which none of the particles has any significant motion. Figure 5 shows a three-dimensional visualization of spherical particles of RDX in molten TNT during the early stages of sedimentation.

The initial program was designed to run on a VAX 8700 computer and gave useful information on sediment density as a function of particle size distribution, but the program was impractical to run on systems large enough to give reliable statistics because the CPU time varied as N^2. This is a common problem in molecular dynamics simulations and arises because in a system of N particles randomly distributed in space, the motion of any one given particle is in principle determined by the remaining N-1 particles. In practice, however, the particle will be influenced significantly by only a relatively small number of nearby particles. These are defined to be the particle's "nearest neighbors." Finding the nearest neighbors in an efficient manner is a central part of any molecular dynamics simulation. A natural choice is to calculate the distance between the particle and the N-1 other particles in the system. A cutoff radius, RC, is then defined such that particles separated by a distance greater than RC have negligible effect on one another. The time taken to calculate the force on the given particle is then reduced because only those particles within the cutoff radius RC need to be considered. The difficulty with this approach is that calculating the distance to each of the N-1 other particles is almost as expensive as including all N-1 particles in the force calculation.

At MRL we have redesigned our molecular dynamics code to overcome the nearest neighbor problem by implementing a new algorithm that keeps track of near-neighbor relationships. The algorithm is based on constructing a "monotonic Lagrangian grid" (MLG), a data structure in which adjacent particles in space have close grid indices.[6] The data structure is arranged so that a particle's nearest neighbors are readily identified but without the need to continually calculate the distance to each of the other N-1 particles. The computational cost of the scheme scales as N, and the algorithm is ideally suited to vectorization because it uses data from contiguous memory locations. The scheme first was implemented on a VAX 8700 computer to check the coding and verify the linear scaling of CPU time with N. This improvement enabled us to increase our sample size, N, conveniently from 200 to 1000. To further increase computational efficiency, the vectorization capabilities of the algorithm must be used. In an MLG, the N particles in the system are specified by assigning values to three numbers: NX, NY, and NZ, where $N = NX \times NY \times NZ$. In a "regular MLG" structure, all of the computationally intensive parts of the code are contained within a nested set of DO loops over indices i, j, and k, which vary between 1 and NX, NY, and NZ, respectively. In a sample of 1000 particles we would have $NX = NY = NZ = 10$, which means that the innermost loop has a count of only 10. This is far too short for the increase in speed due to the pipeline architecture to be effective, and when our regular MLG code was run on the CRAY X-MP system with $N = 1000$ only a modest increase in speed was obtained.[7] One way of overcoming this problem is to use a "skew-periodic" MLG structure.[8] In essence, this maps a plane of points onto one vector, so that the innermost loop now has a count of, say, $NX \times NY$. For $N = 512$ this would result in vectors of length 64 and would lead to significant savings in CPU time over a scalar computer. We currently are considering implementation of a skew-periodic MLG in our molecular dynamics code for application to the sedimentation problem and other calculations that involve the viscosity of solid suspensions of varying particle sizes.

Acknowledgments

The authors thank Ross Kummer, formerly of the MRL Explosives Division, for producing the color graphics used in this article. The work on detonation transfer was done in collaboration with Elaine Oran of the Naval Research Laboratory and Martin Sichel of the University of Michigan, and the molecular dynamics simulations were done in collaboration with Sam Lambrakos of the Naval Research Laboratory. The authors also thank Mark Watson and Tony Poriazis of Leading Edge Technologies, Ltd.

About the authors

David Jones is a research scientist in the Explosives Division at MRL. During 1987 and 1988 he was a visiting scientist in the Laboratory for Computational Physics and Fluid Dynamics at the Naval Research Laboratory, Washington D.C. He was awarded a Ph.D. degree in theoretical physics from Monash University in Australia in 1977.

David Smith works in the Explosives Division at MRL and has used hydrocode modeling to study a variety of explosives-related phenomena. He has a B.Sc. degree from the University of New South Wales.

Rodney Borg is a cadet research scientist in the Explosives Division at MRL and is currently studying for a Ph.D. degree at the Flinders University of South Australia in the area of molecular photodissociation of explosive molecules. He has a B.Sc. degree in chemistry from the University of Melbourne.

References

1. Masinskas, J. J. and E. H. van Leeuwen, "Theoretical Investigation of the Wave Shaping Process in a Section of Ladder Fracture Tape," Report MRL-R-1081, September 1987.
2. Matuska, D. A. and J. J. Osborne, "HULL Documentation," Orlando Technology, Inc., Florida, 1985.
3. Liu, J. C., M. Sichel, and C. W. Kauffman, "Dynamics of Explosions," *Progress in Astronautics and Aeronautics,* Vol. 114, p. 264, 1988.
4. Jones, D. A., M. Sichel, R. Guirguis, and E. S. Oran, "Numerical Simulation of Layered Detonations," *Progress in Astronautics and Aeronautics,* in press.
5. Jones, D. A., M. Sichel, E. S. Oran, and R. Guirguis, "Detonation Transmission in Layered Explosives," The 23rd Symposium (International) on Combustion, Orleans, France, July 1990.
6. Boris, J., "A Vectorised Near Neighbors Algorithm of Order N Using a Monotonic Logical Grid," *Journal of Computational Physics,* Vol. 66, pp. 1-20, 1986.
7. Borg, R. A. J., D. A. Jones, and S. G. Lambrakos, "Molecular Dynamics Simulation of Sedimentation Using the Monotonic Lagrangian Grid," Report MRL-TR-89-30, July 1989.
8. Lambrakos, S. G. and J. P. Boris, "Geometric Properties of the Monotonic Lagrangian Grid Algorithm for Near Neighbor Calculations," *Journal of Computational Physics,* Vol. 73, pp. 183-202, 1987.

CORPORATE REGISTER

Cray Research gains new customers in Australia, England, the United States, and France

The Australian Bureau of Meteorology, Australia's weather agency, a new customer for Cray Research, has installed a CRAY X-MP14/se system. The supercomputer will be used to run global models of the atmosphere to improve the accuracy of weather and climate forecasts.

Ros Kelly, Australia's environment minister, said that the CRAY X-MP system has enabled the Australian weather bureau to run operational weather forecasts for the first time. Operational forecasting provides timely and accurate weather predictions that are broadcast to the public.

The United Kingdom (UK) Meteorological Office in Bracknell, Berkshire, one of the world's leading weather and climate research centers, has ordered a CRAY Y-MP8 system. The new system, which includes a solid-state storage device with 128 million words of memory, is scheduled to be installed in the second quarter of 1991. The supercomputer will be used by the UK Meteorological Office's newly established Hadley Centre for Climate Prediction and Research, which was developed to conduct research on climate change using numerical models. The UK Meteorological Office issues five-day weather forecasts and provides high-level wind predictions to many of the world's airline companies.

The University of Texas System's Center for High Performance Computing has installed a CRAY Y-MP8 system and a solid state storage device. The new system replaces a CRAY X-MP24 system and a CRAY X-MP EA/14se system. Each of the 14 universities in the system is able to access the new, more powerful CRAY Y-MP system. According to Charles Warlick, executive director of academic information systems for the University of Texas System, the CRAY Y-MP system is a central tool in fulfilling the university system's research mission. The supercomputer will be used to perform leading-edge research in chemistry, aerospace, petroleum engineering, advanced materials science, biomedicine, electronics, advanced physics, and rational drug design in cancer research.

The National Institute of Standards and Technology (NIST) has ordered a CRAY Y-MP2 system, including a solid state storage device with 32 million words of memory. The new system will be installed early in 1991 at the institute's facility in Gaithersburg, Maryland, and will run Cray Research's UNICOS operating system. NIST is a civil agency that conducts research in chemistry, materials, electronics, manufacturing, building technology, and fire safety. Through an interagency computer sharing arrangement, NIST's CRAY Y-MP system also will be used for environmental research and radio spectrum analysis.

The French National Organization of Aerospace Research (ONERA) has ordered a CRAY Y-MP system with 128 million words of memory. This system, which replaces a CRAY X-MP system, is scheduled to be installed in the third quarter of 1991 at ONERA's facility near Paris. The supercomputer will be used for aerospace research and design applications, including computational fluid dynamics and structural analysis.

Peugeot S.A. (PSA), an organization consisting of two French automobile manufacturers, Automobiles Peugeot and Automobiles Citroen, has ordered a CRAY Y-MP2E system. The new system, scheduled to be installed in the second quarter of 1991 at PSA's computer facility near Paris, replaces a CRAY X-MP system and will be the first CRAY Y-MP2E system to be installed in France. PSA's system will be used for general automotive research and design including structural analysis, crash simulation, acoustic design, and car behavior simulation.

Scripps Clinic and Research Foundation of La Jolla, California, has ordered a CRAY Y-MP2E system and the UniChem interactive chemistry environment. Using this powerful combination, researchers will investigate various molecules and their reactions within the human body, which is crucial in understanding and working toward a cure for diseases ranging from cancer and AIDS to the common cold.

Conoco is the first petroleum company to order a CRAY Y-MP2E system, which will be installed in the second quarter of 1991 at the company's facilities at Ponca City, Oklahoma. The new system will be used for oil exploration and production, and will allow the company to conduct advanced oil exploration and reservoir modeling techniques.

Grumman Corporation has ordered a CRAY Y-MP2E system to be installed in the second quarter of 1991. The new system will be used for aerospace research and design, and will provide increased power and memory capabilities to perform larger computational fluid dynamics (CFD) and structural analysis simulations.

Cray Research and Yokogawa Electric form joint venture company

Cray Research, Inc., through its Entry Level Systems division, has finalized an agreement with Yokogawa Electric Corporation of Tokyo to establish a new joint venture company, Yokogawa Cray ELS, Ltd. The new company will have exclusive rights to sell and support the CRAY XMS minisupercomputer and its follow-on products in Japan. The joint venture company will be owned equally by Cray Research and Yokogawa. A board member of Yokogawa will be named president of the company.

"Establishing the joint venture with Cray Research and participating in the minisupercomputing business has a great significance for us," said Tak Yamanaka, president of Yokogawa. "Many of our current customers who are engaged in process control and factory automation have recognized the need for more efficient research and development, as well as shorter product development cycles."

Cray Research and Canon Sales form agreement to market CRAY Y-MP2E system in Japan

Under an agreement signed February 1, 1991, Canon Sales Company will market the CRAY Y-MP2E supercomputer in Japan on behalf of Cray Research. Canon Sales will focus its efforts on industries such as finance and construction, where it has extensive experience. In accordance with the agreement, Canon ordered two CRAY Y-MP2E systems for delivery in 1991 and 1992.

Cray Research's existing organization in Japan, Cray Research Japan, Ltd., will continue to sell all of the company's supercomputers, including the CRAY Y-MP2E system, throughout the country.

APPLICATIONS UPDATE

Cray Research announces CRI/TurboKiva engine design tool

CRI/TurboKiva is a comprehensive simulation and analysis tool for modeling engine flow, fuel injection, combustion, and pollutant emission. The CRI/TurboKiva program was developed by Cray Research in cooperation with Group T3 t the Los Alamos National Laboratory and with participation from researchers at the University of Wisconsin. CRI/TurboKiva serves as a numerical test bench for the design of homogeneous-charge or fuel-injected reciprocating engines. It combines experimental analysis and high-performance supercomputing to minimize engine design time.

CRI/TurboKiva is based on the kiva2 program from the Los Alamos National Laboratory. It incorporates numerous enhancements that improve the program in terms of

- Ease of use. CRI/TurboKiva includes a pre- and a postprocessor based on the Massachusetts Institute of Technology's X Window System.
- Computational efficiency. CRI/TurboKiva is highly optimized to take advantage of Cray Research hardware and software systems.
- Versatility. CRI/TurboKiva can accommodate detailed geometries and can be integrated with other computer-aided engineering tools. It also can incorporate user-developed combustion and chemical reaction models.
- Accuracy. CRI/TurboKiva includes detailed physics.

The CRI/TurboKiva input processor uses Open Software Foundation's MOTIF graphic user interface and facilitates the input and management of hundreds of numerical, geometrical, physical, and chemical quantities. The flow simulation program has been greatly enhanced to take full advantage of the vector and parallel processing capabilities of Cray Research computer systems. Complex computational meshes generated on a variety of industrial mesh generators can be input, and experimental data, swirl and/or tumble can be read as initial velocity conditions. Inlet and exhaust ports and moving valves also are included.

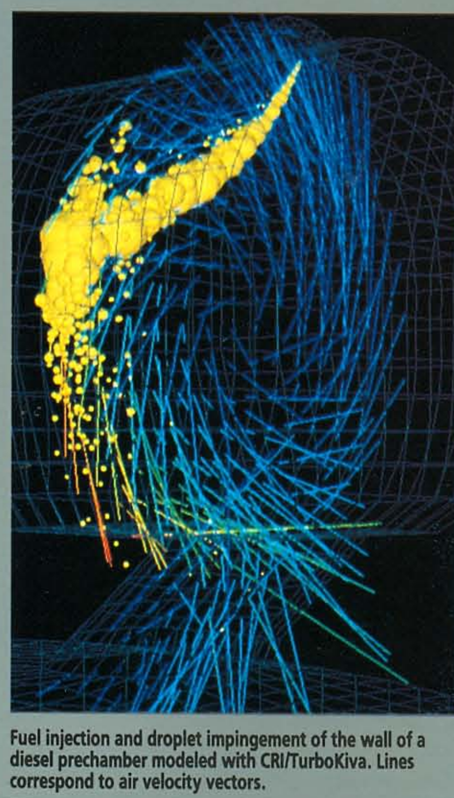

Fuel injection and droplet impingement of the wall of a diesel prechamber modeled with CRI/TurboKiva. Lines correspond to air velocity vectors.

CRI/TurboKiva includes two postprocessing options:

- An interactive postprocessor based on the MOTIF interface. This postprocessor is an enhanced version of the postprocessor included in the original kiva2 program.
- Cray Research's Multipurpose Graphic System (MPGS). This graphics package combines the high-speed computing capabilities of Cray Research systems with the three-dimensional graphics capabilities of advanced workstations. MPGS can be used for three-dimensional postprocessing, animation, and video generation.

For more information about using CRI/TurboKiva on Cray Research systems, contact Reza Taghavi, Cray Research, Inc., 655-E Lone Oak Drive, Eagan, MN, 55121; telephone: (612) 683-3643.

Gigaflop-winning CFD program available on Cray Research systems

The FEAT computational fluid dynamics (CFD) program from the Engineering Analysis Centre of Berkeley Nuclear Laboratories, based at Nuclear Electric in the United Kingdom, now is available commercially for use on Cray Research computer systems. The program uses finite element methods to solve the equations of incompressible turbulent flow and general heat transfer in general-purpose flow configurations. The code has been validated against a wide range of power-plant applications, including electrostatic precipitators, environmental flows, gas flows in boilers, and fast reactor liquid-metal thermohydraulics. Researchers from the Engineering Analysis Centre received a 1990 Cray Research Gigaflop Performance Award for running the FEAT program on an eight-processor CRAY Y-MP computer system at a sustained rate exceeding 1.5 GFLOPS.

When modeling fluid flow, engineers often cannot determine whether the underlying equations have been solved to the required level of accuracy, due to flow complexity. As a result, engineers are unable to guarantee that essential embedded flow structures have been suitably resolved. In addition, numerical errors sometimes cannot be disentangled from turbulence model inadequacy, and as a result, models can be difficult to refine. In many cases, these factors have limited the impact of industrial CFD relative to that of structural engineering analysis. The FEAT program makes CFD more useful to the engineer through the use of quadratic elements to represent

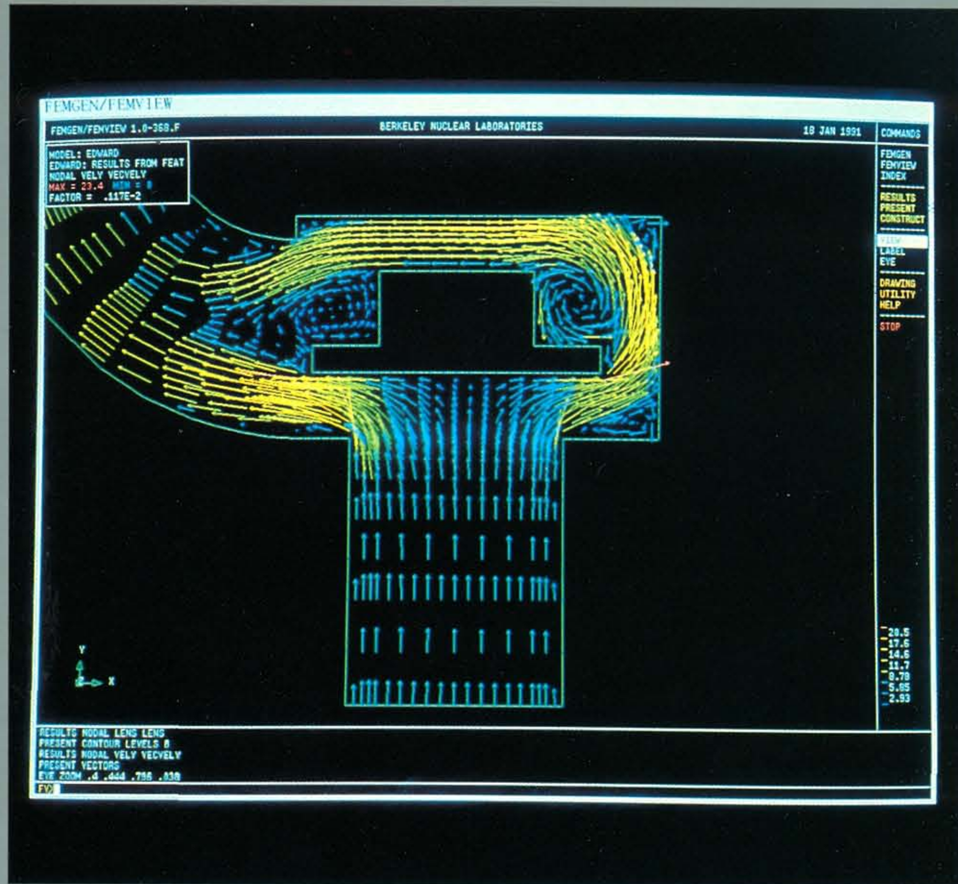

Fluid flow through a fire hydrant modeled with the FEAT software package. Courtesy Stanton, plc.

convected flow variables and by generally eliminating the need to use upwinding. This enables high spatial accuracy (fourth order compared to the normal first order) to be achieved for mesh spacing. In addition, the flexibility of finite element methods allows meshing to be concentrated into subregions in which errors have been unacceptably high. FEAT also provides a grid-error analyzer that allows the engineer easily to assess the degree of grid convergence achieved and therefore to solve problems with a guaranteed level of grid convergence.

The GLOPS performance was achieved during a simulation of turbulent flow in a cubical cavity stirred by a body force. The simulation required many of the FEAT program's features and represented a problem that the program has been used to solve in many contexts, including buoyancy-driven flows in liquid pools and electromagnetically stirred induction furnaces.

The mesh used for the problem comprised 3136 quadratic elements, and the fluid stresses were modeled using the k-ε turbulence model, the standard option in FEAT. To attain GFLOPS performance, 5 of the 563 subroutines within the standard commercial version of the code were modified to use preprocessor compiler directives included in the Autotasking feature of Cray Research's CF77 Fortran compiling system. The remaining routines were run through the Autotasking feature without modification.

The problem also required the transfer of a large amount of data. To minimize the overhead of the data transfer, all I/O was handled through the SSD as secondary data segments, which reduced the I/O wait time from several thousand seconds to one second. The original problem would have required more than 15 hours of computing time on one CPU of a CRAY Y-MP system. The final elapsed time for the job was 2.2 hours, running at a sustained rate of 1.514 GFLOPS on an eight-processor CRAY Y-MP system, using 40 Mwords of memory and 128 Mwords configured as secondary data segments on an SSD solid-state storage device. Work is ongoing to duplicate this performance on a wider range of fluid dynamic problems, particularly vehicle aerodynamics.

For more information about using FEAT on Cray Research systems, contact Peter Whitton, Berkeley Nuclear Laboratory, Berkeley, Gloucestershire, England, GL13 9BP; telephone: (0453) 810451; or Dave Harrison, Cray Research (UK), Ltd., Oldbury, Bracknell, Berkshire, England, RG12 4TQ, telephone: (0344) 485971.

USER NEWS

Gigaflop winners set the standard

By setting new standards in supercomputer performance, Cray Research computer systems help engineers and scientists set new standards in efficiency and productivity. To recognize researchers who are achieving the highest levels of performance on Cray Research systems while solving practical problems in engineering and science, Cray Research instituted the Gigaflop Performance Award Program in 1989. In its first year, the program acknowledged 20 individuals and teams who had solved problems using Cray Research computer systems at a sustained processing rate exceeding 1.0 Gigaflops (GFLOPS), or one billion floating-point operations per second. In 1990, the qualifying performance level was raised to 1.5 GFLOPS.

The following descriptions highlight the work of three 1990 Gigaflop Performance Award winners. For a brochure that describes the work of each of the 30 winners, contact Diane Ciardelli, Marketing Communications, Cray Research, Inc., 1440 Northland Drive, Mendota Heights, MN, 55120; telephone: (612) 683-7215.

Optimizing space structures

Many structural engineering problems involve structures tightly coupled with nonstructural subsystems. Three-dimensional simulation of these fluid-structure, thermal-structure, magneto-structure, and control-structure interactions requires extensive numerical computations, which are only possible using high-performance algorithms and hardware. To make such analyses easier, researchers Charbel Farhat, Luis Crivelli, and Eddy Pramono of the University of Colorado, Boulder, developed the PVFENS program to conduct nonlinear finite element analysis of space structures. They achieved a processing rate exceeding 1.6 GFLOPS on a CRAY Y-MP system using the program to validate the proposed design of a space shuttle cabin compartment. "The key aspects we studied were the external shape, the thickness of the vehicle's outer shell, and the sizing of the stiffeners used to stiffen the shell," Farhat explained.

Their analysis included 8,880 four-noded shell elements and generated 72,654 equations. PVFENS made those calculations using a computational strategy that parallelizes and vectorizes the finite element analysis from beginning to end, rather than focusing on specific computational stages. "This is mostly a nonlinear static analysis to evaluate the stresses and strains on the craft," Farhat said. "Our goal is to calculate the minimum weight possible under a given maximum allowable level of stresses and strains."

The researchers used the finite element method to solve a set of three-dimensional, nonlinear, partial differential equations, which involved solving several nonlinear algebraic sparse equations. GFLOPS performance was attained by a complete redesign of the numerical algorithms to exploit parallelism at the outer loops, vectorization at the inner loops, and minimization of processor-to-memory traffic requirements.

The program features several algorithms that support parallel and vector computational strategies. Embedded in these strategies are versions of well-known iterative and direct solution algorithms for both static and dynamic finite element analyses, as well as innovative parallel algorithms that have been spurred by the advent of parallel and vector hardware. "We found this variety of algorithms necessary because there is no unique parallel or vector finite element algorithm that is 'ideal' for any given problem," Farhat explained. As a result, PVFENS can be used to evaluate a wide range of structural engineering problems.

To further improve parallelism and adjust the code to run on eight processors of the CRAY Y-MP computer system, PVFENS features two processor-mapping techniques. These mapping schemes are triggered by a finite element preprocessor that includes several mesh decomposition algorithms. PVFENS is implemented using the Force program, developed at the University of Colorado, Boulder, by Harry Jordan. Force allows programmers to write their programs and run their code on any of several shared-memory multiprocessors. "The ability to perform very rapidly repeated analyses of realistic complex structural models is essential for the evaluation, sizing, and optimization of structural designs," Farhat said.

Structural design and analysis of a cabin for a launch vehicle modeled on a CRAY Y-MP system using the PFVENS software program.

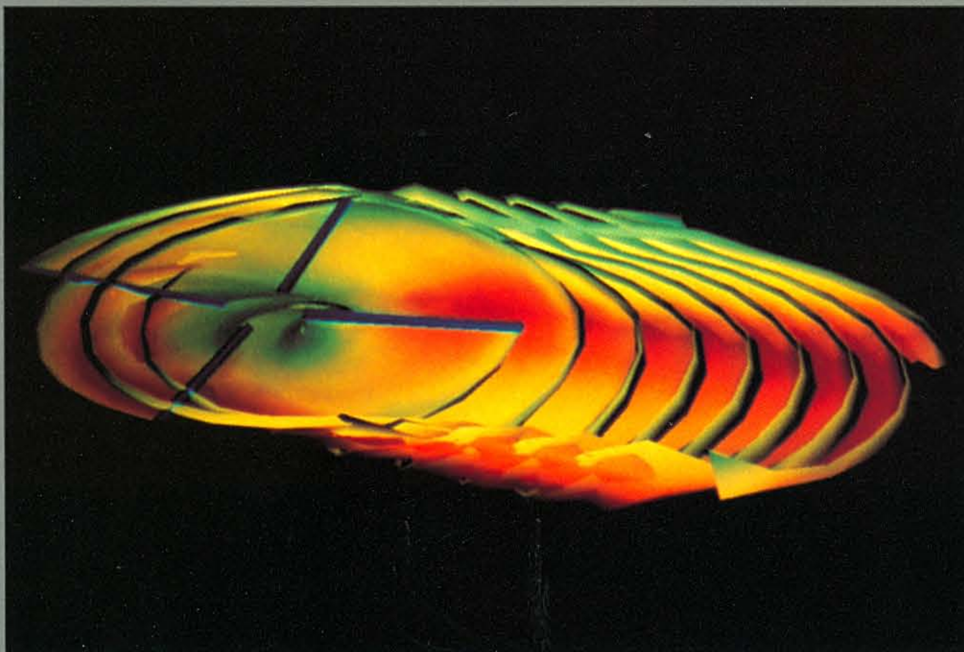

Predicted vortex lattice wake structure of the complete wake surface of an H-34 helicopter rotor as modeled on a CRAY Y-MP system using the FREEWAKE software program.

Model describes helicopter wakes

The limitations of early computer architectures prevented helicopter design engineers from cost-effectively using computer simulation to predict real-world helicopter rotor blade performance. The calculations needed to determine the complete wake structure of helicopter rotor blades in forward flight grow as n^2, where n is the number of wake elements. These calculations must be performed for many time steps. In the past, time and cost prohibited the computing of such large problems. Therefore, engineers could evaluate rotor blade performance only by conducting costly wind tunnel tests.

This situation is changing, however. T. Alan Egolf, senior principal engineer at the United Technologies Research Center in East Hartford, Connecticut, has written a program that, when run on a multiprocessor supercomputer, can calculate the wake structure of helicopter rotor blades in forward flight. The program, called FREEWAKE, uses a vortex lattice structure to model the complete wake surface. The code first calculates the mutually induced velocity of all wake elements, then combines the results with calculations of free-stream velocity to predict the wake displacement for a given moment in time. FREEWAKE then calculates the influence of this distorted wake geometry on the rotor blade.

Minor changes to the FREEWAKE code improved its vector performance on a single-processor CRAY Y-MP system. Measured single-processor performance was 206 MFLOPS. Egolf obtained 1.535 GFLOPS performance on an eight-processor CRAY Y-MP system after Andrew Zachary of Cray Research made an additional change to a primary DO-loop structure in the code and added a compiler directive. The case run for GFLOPS timing evaluated four rotor blades, each with 16 wake filaments with 256 segments per filament, and defined a wake lattice structure of over 16,000 wake elements.

"FREEWAKE should improve researchers' ability to predict more accurately the distorted wake geometry of helicopter blades," Egolf said. "It allows faster turnaround and accommodates larger problems. Ultimately, this will help us better define the influence of helicopter wake geometry on helicopter rotor blade air loads. In time we hope to produce a numerical rotor test facility using this FREEWAKE code and others where we can do more tests and refine the rotor blade design to a greater degree before going into the wind tunnel. This should bring our design costs down and allow the wind tunnel to become used more as a verification tool than as a part of the design process."

Supercomputing the super collider

When the Superconducting Super Collider (SSC) is completed near Dallas, Texas, the main collider ring will extend more than 50 miles in circumference and will contain approximately 12,000 magnet elements. Physicists will use the collider to explore the fundamental structure of matter by accelerating and colliding subatomic particles. Successful completion of the project will require judicious tradeoffs between the quality and cost of the main magnets and error compensation systems. Throughout the project, researchers will rely on computer simulations to predict the effects of machine parameters and will use the information to make many important engineering decisions. Physicists George Bourianoff and Yiton Yan, of the Superconducting Super Collider Laboratory, and analysts Gerald Kirschner and Michael Talian, of Cray Research, are simulating the movements of particles through the colliders to study design alternatives. While running the simulations on a CRAY Y-MP system, the researchers achieved a processing rate of 1.590 GFLOPS.

A key technical issue for the SSC is the control of beam loss during the critical periods when proton particles are being loaded and during the initial phases of acceleration. Researchers can calculate a reasonable statistical estimate of particle loss that is likely to occur in the first 40 minutes of operation by following an ensemble of 100 particles for 10 million turns around a lattice with 12,000 magnet elements. However, tracking a single particle through a single magnet element requires approximately 200 floating-point operations, and the entire simulation requires 1015 floating-point operations. About 2,000 hours of CPU time on a 100 MFLOPS system would be required to complete the simulation. The researchers are working to reduce that CPU time to more manageable levels and increase the predictive capabilities of the simulation. "We tracked 512 particles simultaneously through this situation 100,000 times," Kirschner explained. "After several thousand turns you can determine whether the design needs to be changed because, as the particles move closer to the speed of light, it becomes almost impossible for them to deflect very much off their line. During the earlier stages that evaluation is more critical."

To achieve GFLOPS performance, the researchers first modified their code to achieve optimal performance on a single-processor CRAY Y-MP system. "We made sure that every crucial calculation area was highly vectorized and then attained as much overlap of functional units as possible," Kirschner said. The team then used the Autotasking feature of Cray Research's CF77 Fortran compiling system to spread the improved code across the eight-processor CRAY Y-MP system to track an ensemble of 512 noninteracting proton particles, with 64 particles running on each of the processors. The particles were tracked for 100,000 turns around a 53-mile lattice that comprised the 12,000 electromagnetic elements that make up the collider ring.

GALLERY

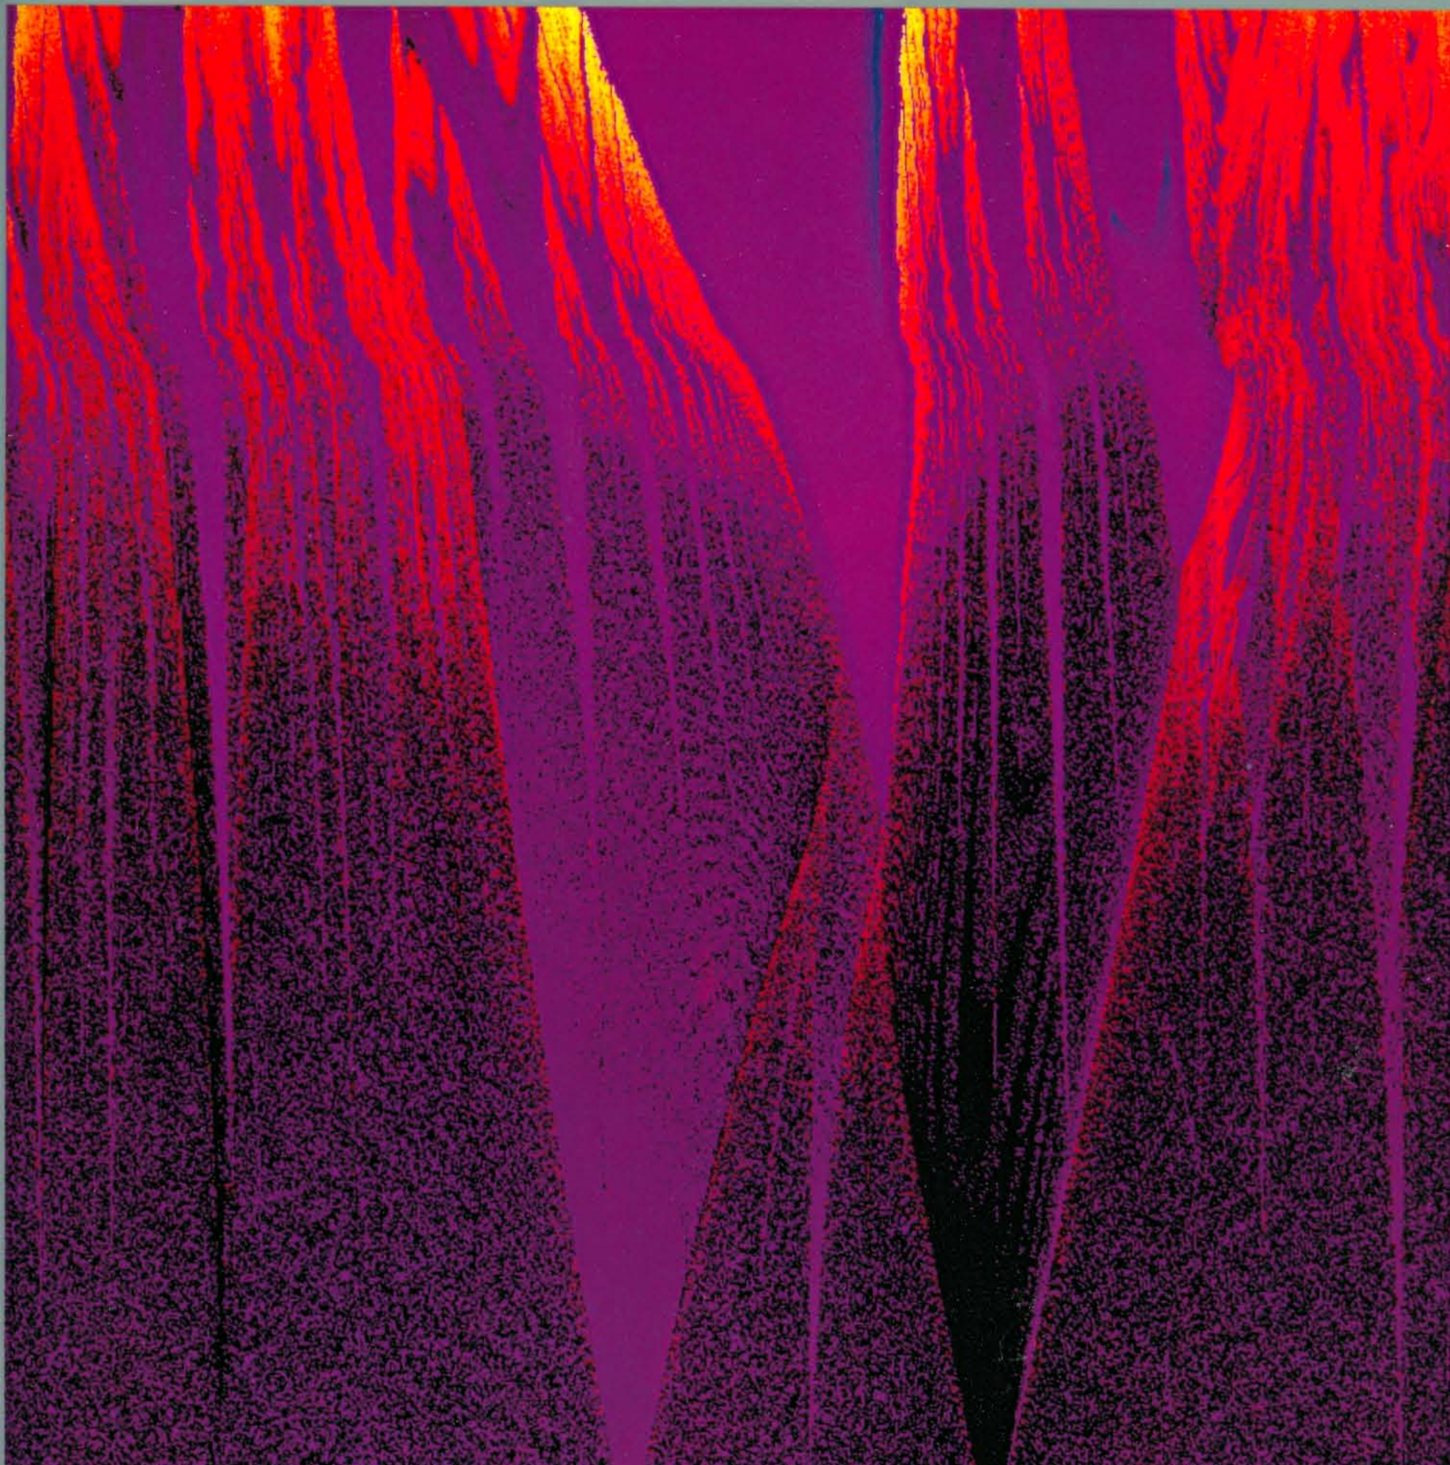

This output from a neural network model run on a CRAY Y-MP supercomputer shows the complexity that simple nonlinear systems can produce. The image was created by University of California at San Diego student Michael Casey on the CRAY Y-MP8/864 system at the San Diego Supercomputer Center.

CRAY CHANNELS welcomes Gallery submissions. Please send submissions to the address inside the front cover.