CRAY CHANNELS

VOLUME 18, NUMBER 1 • A CRAY RESEARCH, INC. PUBLICATION

computational

structural

Computational Fluid Dynamics and Structural Analysis

Modern computational fluid dynamics tools offer a vast array of physical modeling features, making them applicable to a wider variety of practical engineering problems than ever before.

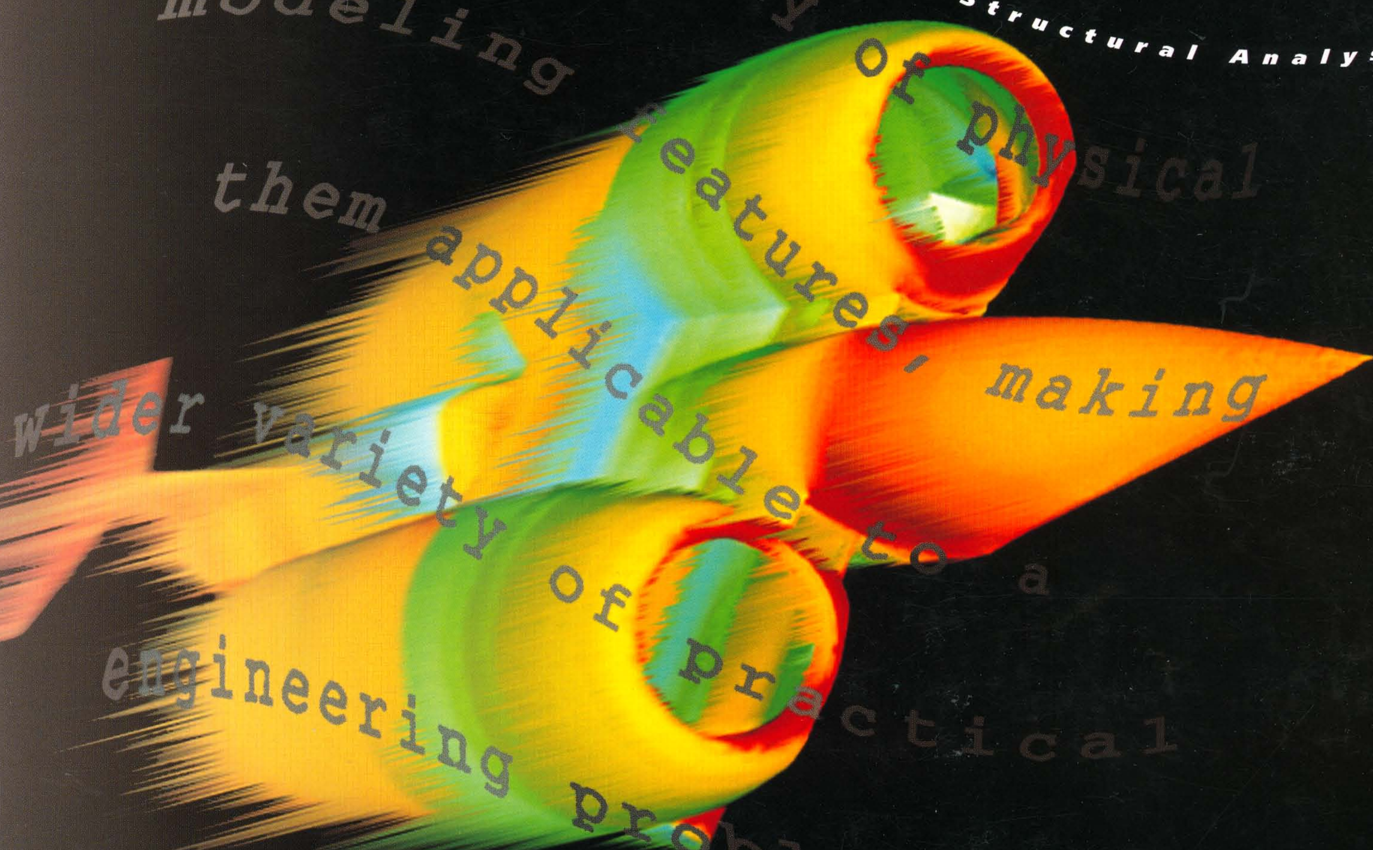

Introducing the CRAY T3E series of highly scalable computer systems

cray

Redefines Advanced Computing

Again

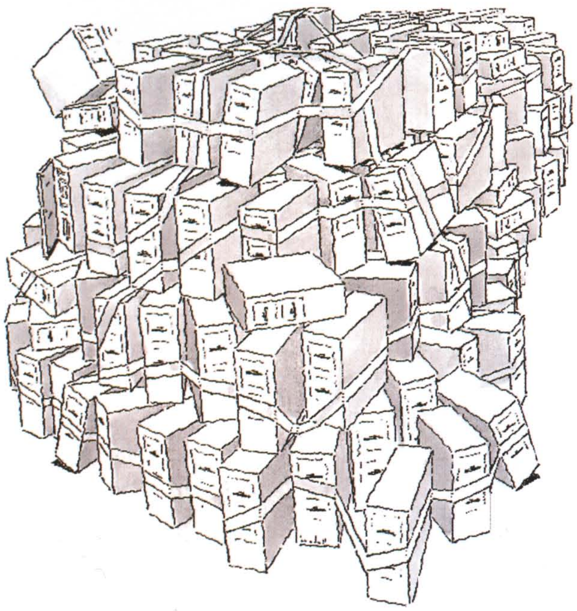

Theirs

Hundreds of computers, plenty of duct tape.

Ours

A State-of-the-Art Computer System

Introducing

The CRAY T3E system, the first truly scalable supercomputer

For those familiar with advanced computing, it will come as no surprise that the company that just redefined the category is the same company that defined it in the first place – Cray Research. In a world where price/performance is everything, Cray delivers like never before. But don't just take our word for it, here's what some industry experts have to say about the CRAY T3E system.

- "(The CRAY T3E) will make scalable parallel supercomputing affordable for lots of companies that used to just dream about it."
  Michael P. Burwen, President, Palo Alto Management Group
- "IDC believes the T3E provides Cray with a powerful product family that will compete across a broad range of markets. This product is the centerpiece of Cray's continuing commitment to high-performance computing."
  Debra Goldfarb, Director of Workstations and High-Performance Computing, International Data Corp
- "With the T3E and its across-the-board true scalability, Cray has redefined what it means to have scalable technology.
  Joe Clabby, Director of Distributed Systems, Aberdeen Group
- "...Cray engineers have built an industrial-strength SPP system capable of tackling the most challenging problems of its most demanding customers..."
  Gary Smaby, President, the Smaby Group
- "With the T3E, Cray has addressed all key user concerns, including outstanding performance and price/performance, scalability, affordable entry price...and ease of use."
  Omri Serlin, President, ITOM International Company

CRAY RESEARCH, INC.

C O N T E N T S

IN THIS ISSUE

Cut costs. Save time. Improve product quality. These objectives are shared by engineers in all manufacturing industries. To help them achieve their objectives, Cray Research provides engineers and researchers with a range of high-performance computational tools.

In this issue of CRAY CHANNELS we survey newer uses of structural analysis and computational fluid dynamics to meet today's research and engineering challenges. Through computer modeling, engineers can spot problems early in the design cycle of a product or manufacturing process, information that can contribute directly to an organization's efficiency and bottom line. In this issue we also introduce the CRAY T3E series of highly scalable systems. Innovative technologies allow the major components of CRAY T3E systems—processors, memory, I/O, and operating system—to scale efficiently from 16 to 2048 processors.

Companies that invest in leading-edge computer modeling technology can reduce waste, minimize expensive prototyping, and collapse their product design cycles to beat their competitors to market. From fluid mixing to metal stamping to mesh generation, Cray Research delivers a competitive advantage through its high-performance computing solutions.

FEATURES

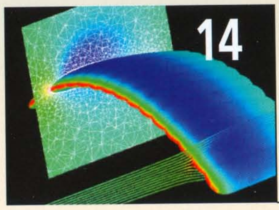

DEPARTMENTS

CRAY CHANNELS

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

Volume 18, Number 1

Editorial staff
Ken Jopp, editor
Gerri Dyrek, Heidi Hagen, Jeff Paulson, contributing editors

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

Scalable power that works

Scalable computing has reached a new level

A new level where a single, powerful system grows with your applications. Where performance and price/performance scale efficiently and cost-effectively from 16 to 2048 processors. Where each feature and function is designed for larger problems, faster solutions, bigger workloads, and a greater return on your investment.

A new level where true scalability is available from under $1 million. Where price no longer keeps budget-conscious organizations and departments from investigating scalable algorithms and developing production applications. Where university-based R&D centers, government laboratories, and automotive, petroleum, aerospace, weather, and environmental research organizations of all sizes have access to the most powerful computer on the planet.

A new level where flexibility, configurability, and reliability reign

Introducing the CRAY T3E series, the world's most powerful and flexible scalable parallel computing systems. Across a range of configurations from 16 to 2048 processors with corresponding peak performance levels from gigaflops to teraflops, CRAY T3E systems deliver superior scalable performance and price/performance on technical applications in scientific and engineering R&D and industrial production computing.

The power and flexibility of this new generation come from global system scalability. Every CRAY T3E system delivers scalable performance through a combination of the world's fastest microprocessor and the industry's best interprocessor communications, the world's first scalable operating system, and scalable I/O. In any configuration, every function and feature contributes to balanced scalable performance. Every price point offers tremendous room for growth with a corresponding payback in performance, throughput, and productivity. This new generation redefines scalable performance and price/performance to deliver scalable power that works.

The CRAY T3E series of scalable parallel systems

Liquid-cooled CRAY T3E system.

Scalable performance

For top performance, we integrated the world's fastest microprocessor into a refined architecture that reflects our unmatched experience in designing the world's most powerful computer systems. CRAY T3E scalable parallel systems use the DECchip 21164 (DEC Alpha EV5) from Digital Equipment Corporation, capable of 600 MFLOPS peak performance. This reduced instruction set computing (RISC) microprocessor is cache-based, has pipelined functional units, issues multiple instructions per cycle, and supports IEEE standard 32-bit and 64-bit floating-point arithmetic. CRAY T3E processing elements (PEs) include the DEC Alpha microprocessor, local memory, and performance-accelerating control logic designed by Cray Research engineers and fabricated in low-cost CMOS.

Each PE has its own local DRAM memory with a capacity of from 64 Mbytes to 2 Gbytes. A shared, high-performance, globally addressable memory subsystem makes these memories accessible to every PE in a CRAY T3E system.

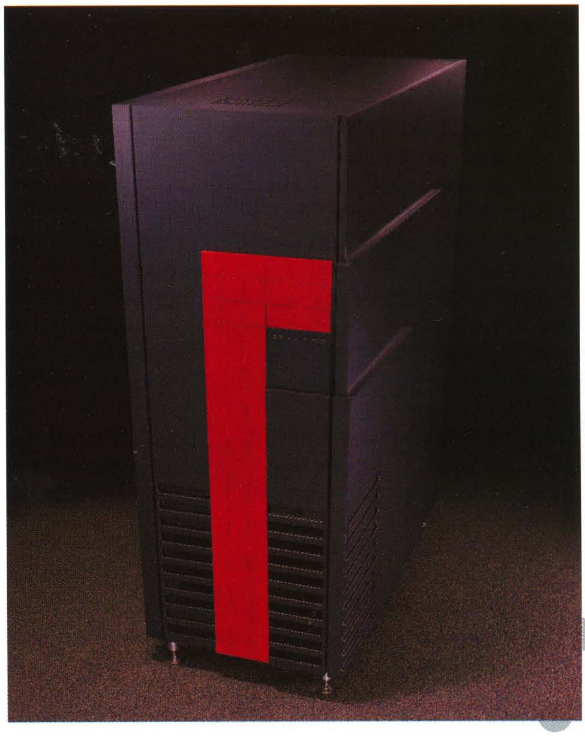

Air-cooled CRAY T3E system.

Scalable interprocessor communications

PEs in the CRAY T3E system are connected by a high-bandwidth, low-latency bidirectional 3-D torus system interconnect network six times faster per PE than that used in our first-generation scalable system. We added adaptive routing to allow messages on the interconnect network to be rerouted around temporary "hot spots." Interprocessor data payload communication rates are 480 Mbytes per second in every direction through the torus; in a 512-PE CRAY T3E system, bisection bandwidth exceeds 122 Gbytes per second.

I/O channels are integrated into the 3-D torus network and increase in number with system size. As PEs are added to a CRAY T3E system, interprocessor and I/O bandwidth scale accordingly. This means that scalable applications run as efficiently on larger configurations as they do on smaller ones.

Scalable operating system

To support the scalability of the CRAY T3E system, we created UNICOS/mk, a scalable, serverized version of our UNICOS operating system. UNICOS uses the UNIX System V operating system as a foundation for high-performance capabilities that meet the demands of our supercomputer customers. UNICOS/mk expands these capabilities to provide unsurpassed parallel efficiency and scalability.

UNICOS/mk is distributed among the CRAY T3E system's PEs, not replicated on each. This distribution of operating system functions provides a global view of the computing environment—a single-system image—that allows administrators to manage a systemwide suite of resources as a single entity.

UNICOS/mk is divided into "servers," which are distributed among the processors of a CRAY T3E system. Local servers process operating system requests specific to each user PE. Global servers provide systemwide operating system capabilities such as process management and file allocation. The result is the world's first scalable operating system.

In addition to the user PEs that run applications and commands, CRAY T3E systems include dedicated system PEs which run the global UNICOS/mk servers. Because global services are provided by system PEs and not replicated throughout the CRAY T3E system, UNICOS/mk efficiently scales, with full functionality, to service from tens to thousands of PEs with minimal overhead.

Through standard UNIX operating system utilities and commands and the same interface as our UNICOS operating system, UNICOS/mk provides a familiar operating environment for users and administrators alike. Scalability is enhanced through 64-bit addressing, which allows the operating system to support large data sets and file systems and thousands of users.

UNICOS/mk complies with the POSIX 1003.1 (system services) and POSIX 1003.2 (commands and utilities) standards. But UNICOS/mk does more than ensure a familiar, open operating environment through compliance with industry standards. It also incorporates our more than 10 years of experience in high-performance parallel UNIX operating systems and includes significant enhancements for scalability; system availability; resource management; and scalable, high-bandwidth I/O.

UNICOS/mk supports our scalable I/O architecture, which provides I/O capability that scales with system size and ensures superior I/O capacity for high-performance, scalable applications. UNICOS/mk delivers this performance through software innovations including distribution of file management services; distio (I/O centrifuge); and multiple global file servers.

Cray plans new configurations to meet demand for scalable systems

When the CRAY T3E series was announced in November 1995, Cray Research already had an order backlog of $92 million for the new systems. By year-end 1995, the backlog had grown to more than $160 million, and it has increased since then. To meet demand for CRAY T3E technology from an even broader range of customers, Cray plans to introduce several smaller configurations later this year. These configurations will include new memory options, as well as options for smaller numbers of PEs. The smaller systems will provide unique capabilities for engineers and research scientists in the petroleum, automotive, aerospace, chemical and other industries. In addition, Cray will offer the smaller systems to qualified university customers at an attractive price. As part of the CRAY T3E Academic Program, these systems will be packaged with a service contract, system software, and programming environments.

Programming environments

To simplify the porting and development of high-performance applications, Cray Research programming environments provide an integrated set of the most powerful application development tools available: optimizing compilers, high-performance parallel libraries, advanced debugging and performance analysis tools, and the most flexible and complete parallel programming models in the industry.

For more than a decade, Cray Research has been the recognized leader in optimizing compilers. This tradition continues with the Fortran 90, C, and C++ compilers for the CRAY T3E system. These compilers go beyond code optimization to provide access to comprehensive parallel programming models. Each compiler supports explicit parallel programming via the Cray Message Passing Toolkit. The toolkit contains high-performance implementations of the standard PVM and MPI libraries, for fast point-to-point and global message passing communication.

Explicit parallel programming capabilities are further extended through the innovative Cray Shared Memory Library, which provides data passing or one-sided communications. This library delivers the fastest interprocessor communication in the industry.

Fortran programmers also can benefit from the simplicity of implicit parallel programming, with system software based on the CRAFT and HPF parallel models. CRAFT is an implicit parallel model developed by Cray Research for the CRAY T3D system to provide users with flexible work distribution among processors and high-level control over data distribution. The HPF implicit parallel model addresses many of these same issues and is an open model that provides increased application portability.

Our programming environments also include visual programming and analysis tools that provide detailed insight into code performance and structure. They support the development of applications written in any combination of the CRAY T3E system's programming languages (Fortran 90, C, and C++) and programming models (MPI, PVM, Shared Memory Library, and CRAFT/HPF). These tools include the Cray TotalView Debugger, an advanced debugger designed for symbolic, source-level debugging of parallel applications; the Cray Program Browser, an interactive environment for browsing and editing files and applications; and MPP Apprentice, a state-of-the-art parallel performance analysis tool that not only displays performance data but interprets the information and provides recommendations to improve performance.

Scalable I/O architecture

Breakthrough I/O technology available for the first time on the CRAY T3E system sets new standards for I/O scalability, performance, capacity, and configurability. Designed to work with all other components of the system for peak efficiency, this new architecture provides a balance between CPU, memory, and I/O demands.

The CRAY T3E system performs I/O through multiple ports onto one or more scalable GigaRing channels. Each dual-ring I/O channel, with data in the two rings traveling in opposite directions, delivers high I/O data bandwidth and enhances reliability.

This high-performance I/O architecture scales to meet your needs by growing in step with your computing system resources. Regardless of size, all CRAY T3E systems may be configured with an I/O channel for every eight PEs (in air-cooled models) or 16 PEs (in liquid-cooled models). All I/O channels are accessible and controllable from all PEs. On the CRAY T3E system, each GigaRing channel has a maximum data payload bandwidth of 1 Gbyte/s and provides high-speed access to peripherals, networks, and other Cray Research systems.

The scalable I/O architecture is so flexible that most configuration designers will find they are nearly free from configurability constraints. For example, a single disk node on an I/O channel can provide access to up to 1.6 Tbytes of disk storage. On a CRAY T3E system, multiple I/O channels and I/O nodes can easily be configured to support up to quadrillions of bytes (petabytes) of disk capacity.

Each GigaRing I/O channel can be configured with multiple nodes of various types: a Multipurpose Node that accepts standard Sbus type controller cards to support FDDI, Ethernet, ATM, SCSI disk, and SCSI tape access; disk drive nodes that provide connectivity to disk arrays and single disks over industry standard channels such as IPI, SCSI and SCSI/Fibre Channel; tape drive nodes that provide access to Block Mux and ESCON tape devices; and network nodes providing ANSI 32-bit or 64-bit HIPPI channels (100 or 200 Mbytes/s).

Superior price/performance from competitively priced entry-level systems to affordable high-end solutions for the world's greatest challenges. Breakthrough hardware, software, and I/O designs configured to meet each customer's needs and balanced to deliver the same efficient performance in every configuration. The CRAY T3E system from Cray Research—scalable power that works.

Computational fluid dynamics in the pharmaceutical industry

Jef Dawson
and Richard D. LaRoche
Cray Research, Inc.

Modern computational fluid dynamics (CFD) tools offer a vast array of physical modeling features, making them applicable to a wider variety of practical engineering problems than ever before. This article highlights two CFD simulations of interest to the pharmaceutical industry.

Clean room ventilation

Pharmaceutical products must be kept sterile and free of foreign material during processing. For this reason, much of the processing is done in clean rooms in which contaminants are eliminated by very high efficiency air filtration. Other industries, including microelectronics and food processing, require similar environmental controls. There are many questions that designers and operators of clean rooms must answer, including

- Is the material being processed safe from contamination?
- Is the filtered air being used most efficiently?
- What happens if a blower fails or a door is opened during processing?

In the pharmaceutical industry, these questions are especially important because when the U.S. Food and Drug Administration approves a drug, it also approves the process used for its production. If any kind of variance occurs and conformance to the approved process conditions cannot be guaranteed, expensive batches of drugs may have to be discarded. In this article, we use one simple example to show how computational fluid dynamics is being used in the design of clean rooms and similar environments.

In this simulation, a 15′ x 12.5′ x 10′ clean room (Figure 1) is analyzed. The room has three work areas along one wall, each 5 feet wide and 3 feet deep with its own ventilation hood. Each hood draws air from the room through a return vent, and passes it through a High Efficiency Particulate Air filter before discharging it above the work table. In addition, there is a single air inlet in the ceiling of the room and an outlet under the far left end of the work table. This analysis was performed using FLOW-3D, a general-purpose CFD code developed by Flow Science, Inc. The geometry input capabilities of FLOW-3D allowed this simulation to be set up and run on a CRAY J90 system in less than two working days.

With no other equipment, furniture, or technicians in the model, this clean room has a fairly simple geometry. Even so, the flow in the room is quite complex (Figure 2). The most prominent feature of the flow is the thick sheet of air that emerges from the gap between the ventilation hood and the work table. This sheet of air creates two counter-rotating, horizontally oriented vortices that extend the length of the room. The larger vortex fills the bulk of the room, and the smaller one is under the table. This would be a relatively simple flow pattern, except for the effect of the column of air that issues from the ceiling inlet.

Figure 1 (below). Layout of the clean room.

Figure 2 (bottom). Velocity vectors in a clipping plane through the center of the clean room.

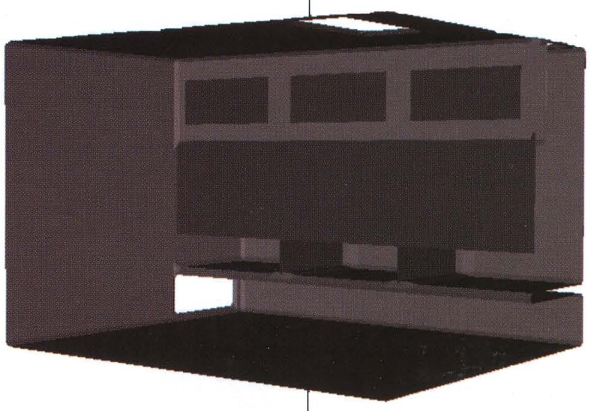

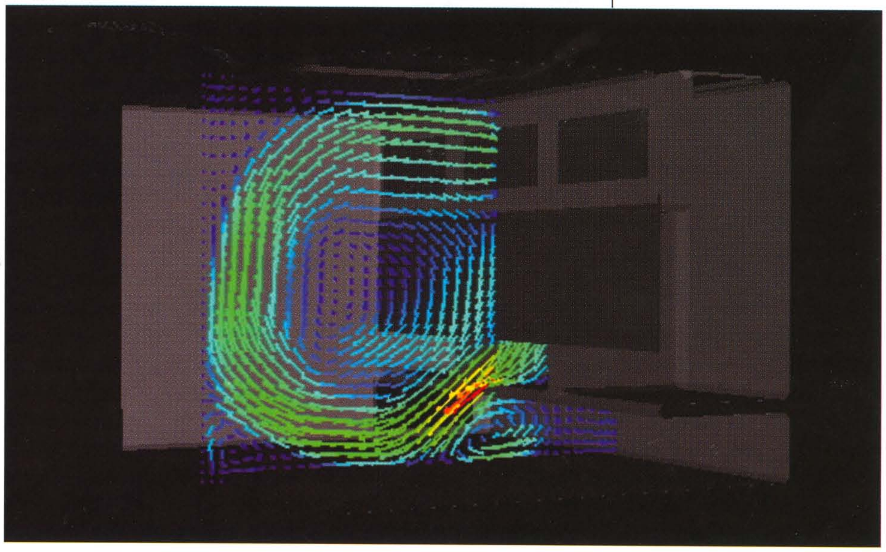

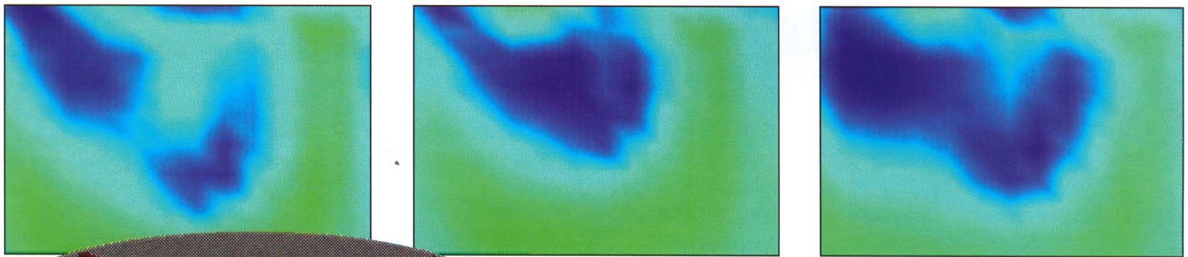

Figure 3. Velocity magnitude on a clipping plane approximately two feet below the clean room ceiling at three different times.

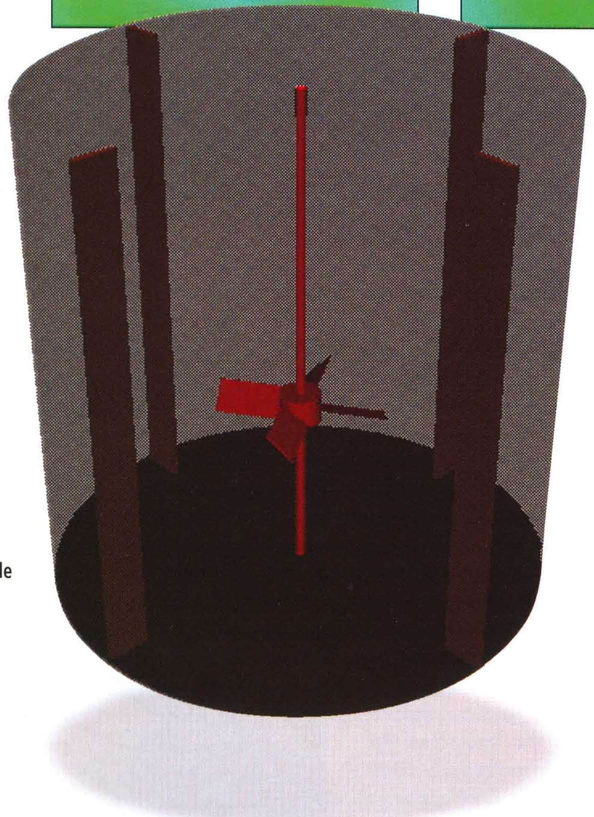

Figure 4. Baffled stirred tank with a four-blade pitched blade impeller.

The interaction of this inlet stream and the large vortex creates smaller, time-dependent vortices which are carried downstream by the two main vortices, resulting in complex, unsteady flow in all regions of the room. Figure 3 shows the velocity magnitude on a clipping plane about two feet below the ceiling, between the main inlet and the work table. Red and blue regions indicate higher and lower velocities, respectively. The three images show the variation of the velocity in time.

At any given point in the room, the time-dependent nature of the flow manifests itself as periodic "puffs" of air. This raises a question of fundamental importance in a clean room: are any of these puffs of air strong enough to carry contaminants back into one of the work areas? We address this issue by considering the velocity variation in the region just outside the gap between the hood and the table. The relatively high-speed air emerging from the gap prevents any contaminants from entering the hood. A similar issue arises if a door into the clean room is opened. To prevent contaminants from entering through an open door, clean rooms are operated at a slightly higher pressure than adjoining rooms and corridors so that air will only flow out when a door is opened. This simulation could easily be modified to determine whether any of the small, fluctuating vortices could overcome a moderate pressure gradient and bring contaminants into the room through an open door.

Even though the mesh resolution in this simulation is relatively coarse, the time-dependent nature of the flow makes it a computationally intensive problem.

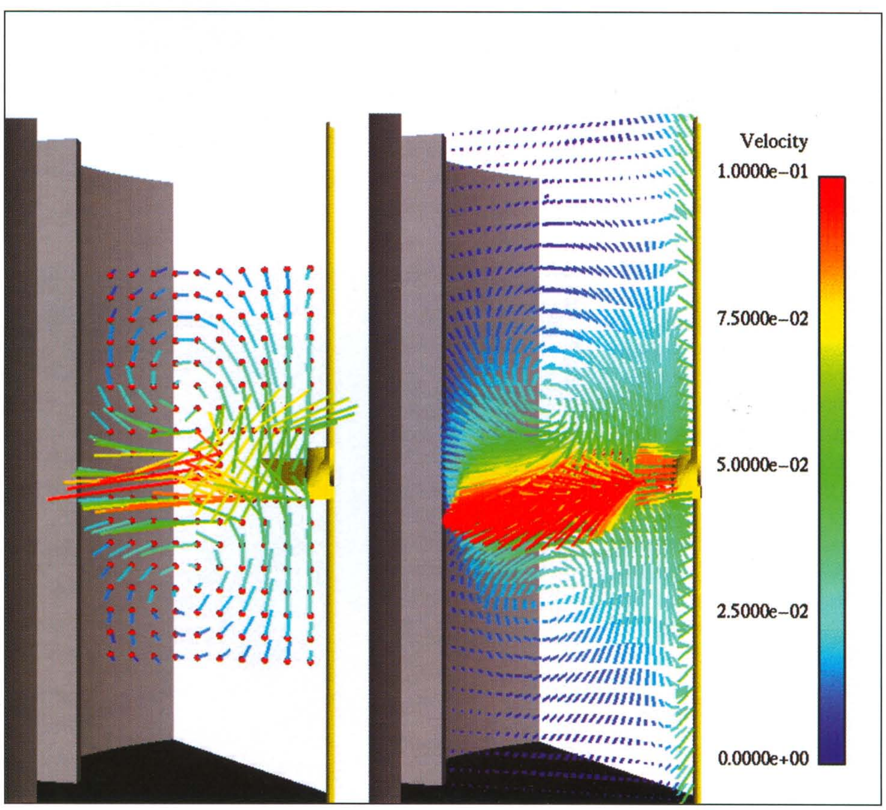

Figure 5. Experimental laser doppler velocimetry (LDV) vs. FLUENT sliding mesh model (right) of a Dow Chemical pilot-scale stirred tank, RE = 20.4.

Fluid mixing in stirred tank reactors

The stirred tank is a common piece of equipment in the pharmaceutical and specialty chemical industries. Figure 4 shows a typical configuration with a four-blade pitched blade impeller in a tank with four baffles. Product quality and yield are often highly dependent on the nature of the fluid mixing inside a stirred tank reactor. CFD has become a valuable engineering tool which can be used with traditional pilot-plant studies to understand the fluid mixing processes within stirred tanks.

Time-averaged impeller CFD modeling in stirred tanks has been established as an effective engineering tool used in conjunction with an experimental program to answer large-scale fluid mixing questions. One major limitation of this technique is its heavy reliance on experimental laser Doppler velocimetry (LDV) data, which may be very difficult to obtain for many real industrial situations. Another limitation stems from the fact that this technique simplifies the fluid flow details in the near-impeller region by imposing time-averaged velocity or momentum boundary conditions.

Effectively modeling the smaller-scale flow structures in the near-impeller region will be essential to addressing many engineering questions such as enhanced micromixing for chemical reaction and mass transfer in gas-liquid systems. One step toward more detailed impeller modeling is sliding-mesh technology.

In baffled stirred tanks, the impeller blades continuously sweep past stationary baffles, resulting in time-dependent geometry configurations and flow patterns. Commercial CFD software packages such as CFX, FIDAP, FLUENT, and STAR-CD, are able to model explicitly the time-dependent geometry in a baffled stirred tank. Figure 5 shows a comparison of the laminar flow field in the baffle plane of a pilot-scale stirred tank obtained from Dow Chemical experiments[1] and FLUENT sliding-mesh calculations.

In Figure 6 we have plotted the axial, radial, and tangential velocity components along a line in the baffle just above the impeller-swept region and compared them with experimental LDV data at laminar flow conditions. Figure 7 shows a similar comparison along a line immediately under the impeller-swept region at turbulent flow conditions. These figures show that sliding-mesh CFD is able to model fluid velocities in stirred tanks well within the established accuracy of LDV measurement techniques. Such verification enables the engineer to use the sliding-mesh CFD model with confidence in scale-up to production stirred tanks where LDV measurements are not practical.

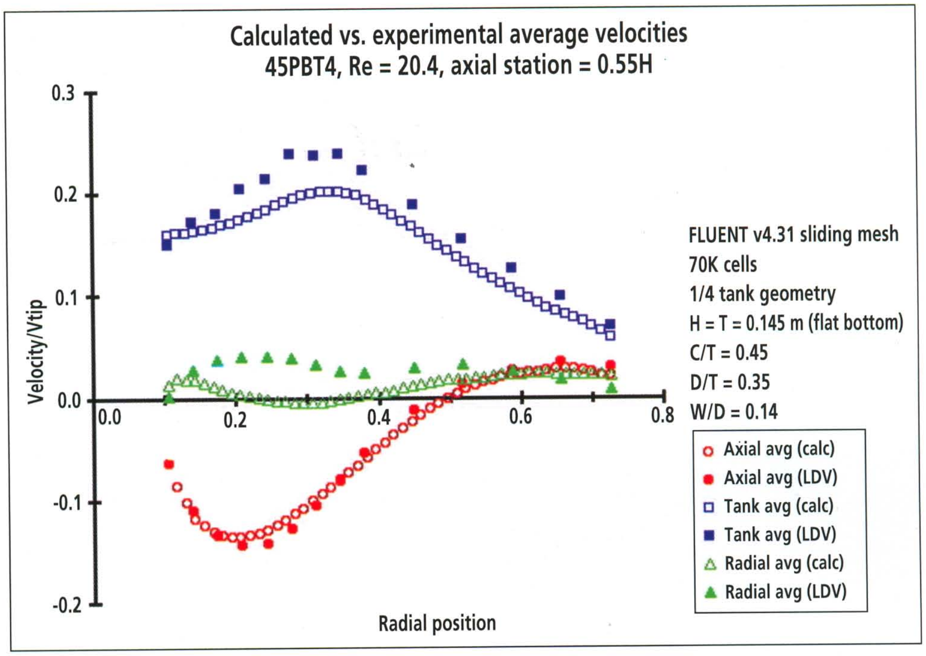

Figure 6 (left). Experimental and computed velocity components just above the impeller region in a laminar stirred tank.

Figure 7 (below). Experimental and computed velocity components in a baffled stirred tank (turbulent flow) with a four-blade pitched-blade impeller.

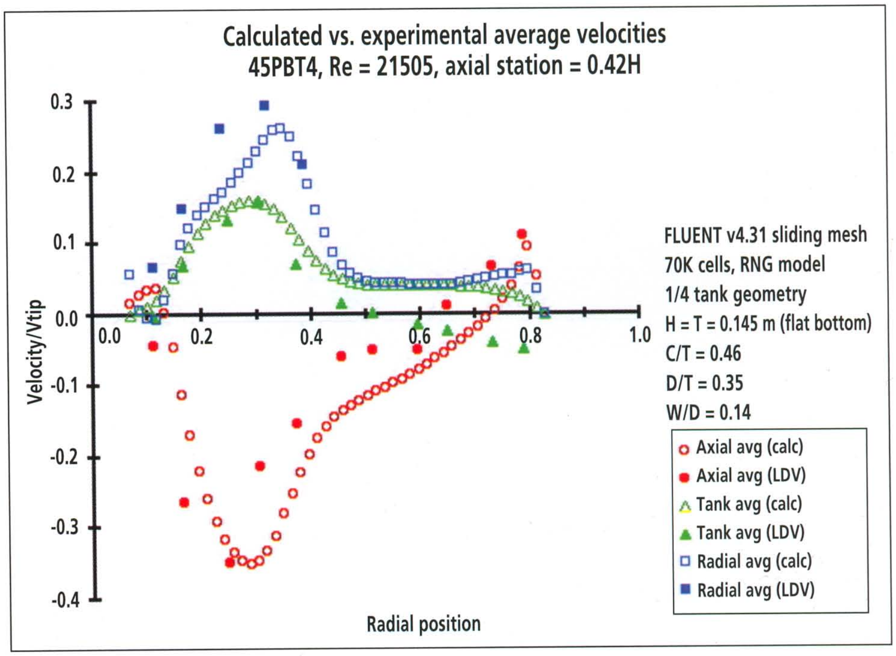

About the authors

Jef Dawson is the manager of the CFD group within the Applications Division at Cray Research. He has been in the CFD group for 13 years. He received a B.S. degree in Mathematics at the University of Minnesota.

Richard D. LaRoche is a senior chemical engineering consultant for the Applications Division at Cray Research, where he focuses on the application of CFD to process industry problems. He received a B.S. degree in chemical engineering from Montana State University and M.S. and Ph.D. degrees in chemical engineering from the University of Illinois at Urbana-Champaign. He previously held positions at Dow Chemical and Penn State University.

Reference

1. Sen, S., C. K. Lee, and D. E. Leng, "Modeling of Turbulence in Stirred Tanks with Experimental Validation," Paper No. 189a, 1994 Annual AIChE Meeting, San Francisco, November 1994.

SUPERSONIC CAR
prepares for race against the sound barrier

Kevin Fox, Cray Research, Inc.

A brand new supersonic car—the Thrust SSC—is being designed to make a record-breaking attempt to surpass the world land-speed record and exceed the sound barrier at speeds that will approximate 750 mph at ground level. Today, the Thrust SSC and its jet pilot RAF Flight Lieutenant Andy Green are being prepared for the challenge at hand.

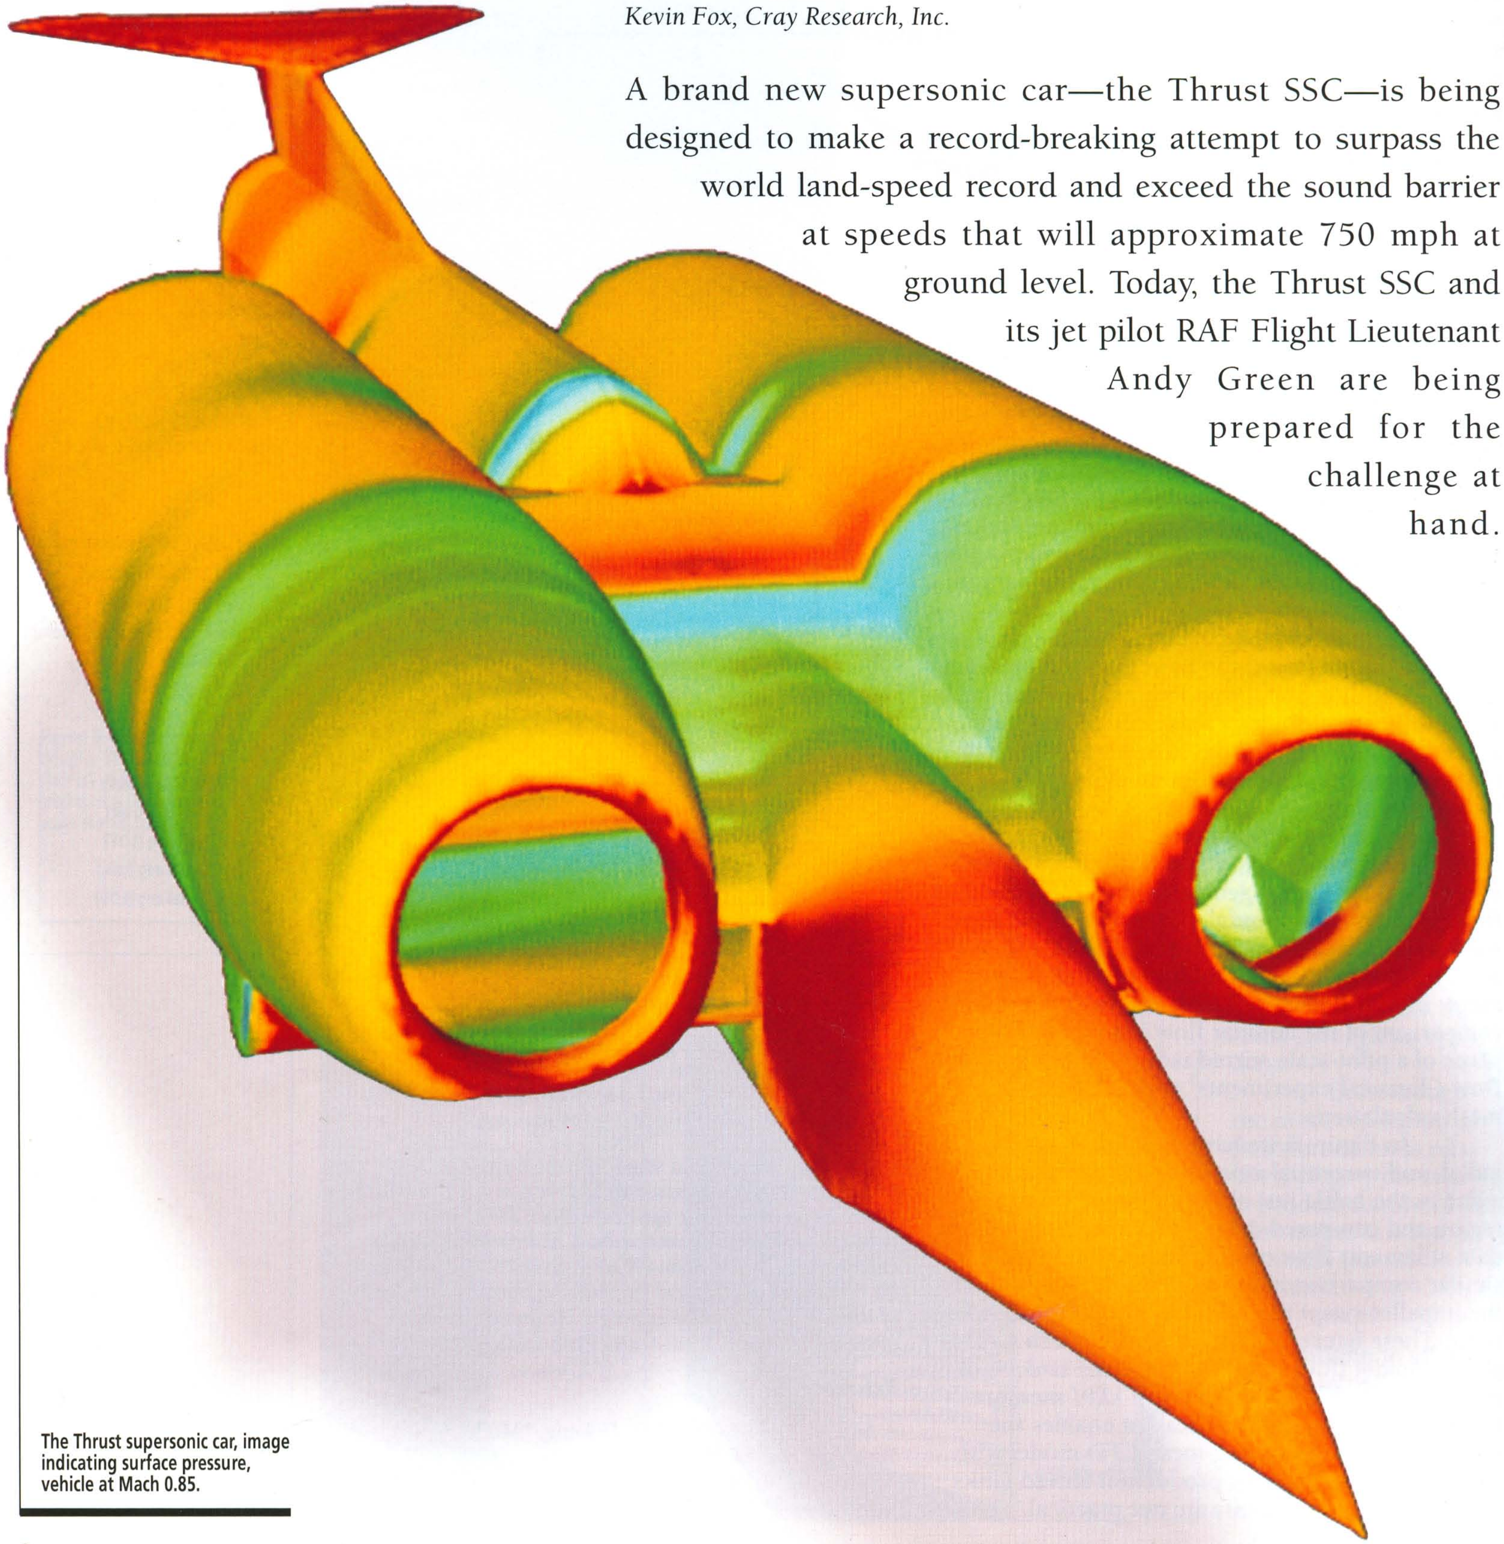

The Thrust supersonic car, image indicating surface pressure, vehicle at Mach 0.85.

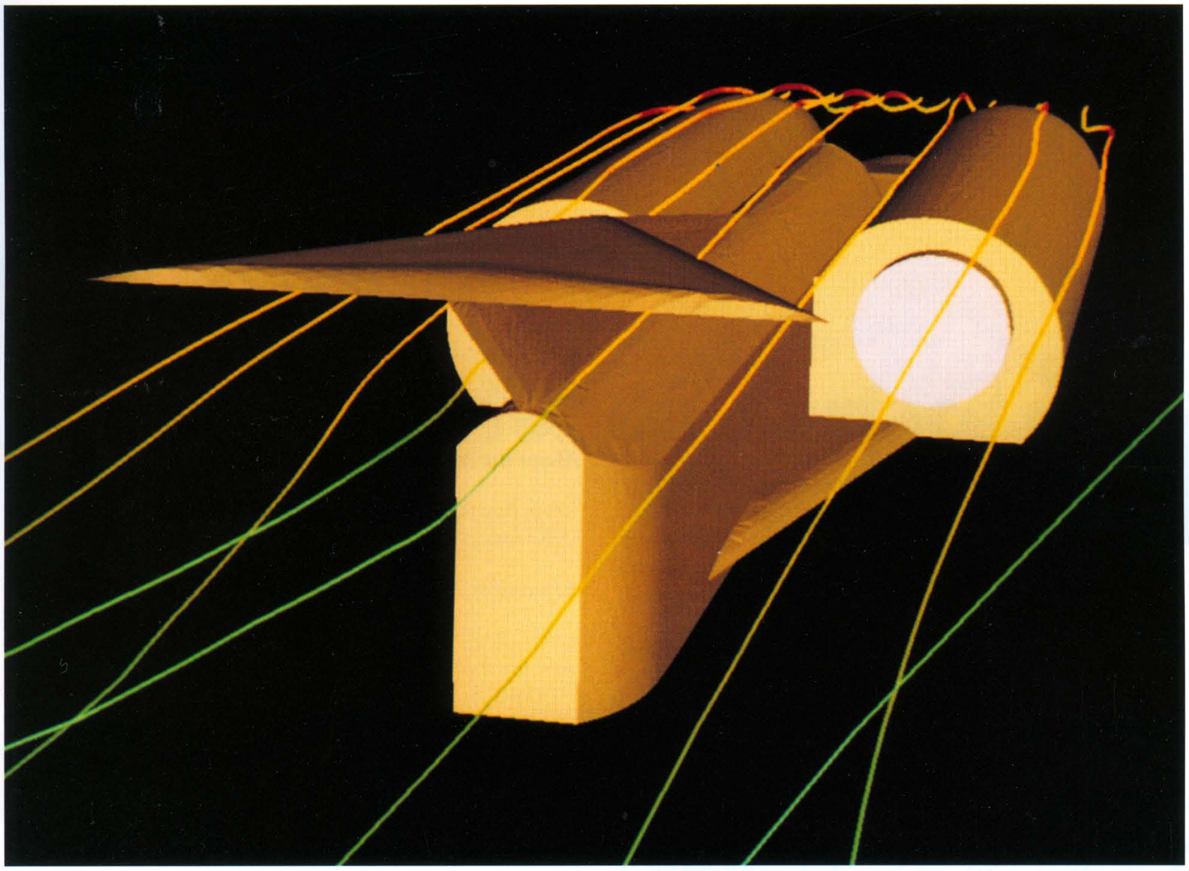

Airflow streamlines, vehicle at Mach 0.85.

The Thrust SSC project is being spearheaded by Richard Noble. This is not the first time that Noble has been involved in a record-setting adventure of this kind. Noble—"the fastest man on earth"—set the land-speed record twelve years ago in the Thrust 2 at a speed of 633.5 mph.

With the help of a retired aerodynamicist, Ron Ayers, and several companies and engineers, Noble's idea for the new Thrust SSC is well on its way to becoming a reality. It was a chance meeting between Noble and Ayers, the lead designer of the Thrust SSC, that set the project in motion over two years ago. Since then a variety of UK companies have provided the sponsorship needed to complete the project either by offering services-in-kind or direct financial funding. Now, Noble and Ayers are busy building a car that could very likely earn them a page in history.

Clearly, if the attempts to break the world land-speed record and the sound barrier are going to be successful, the Thrust SSC is going to need raw power—enough power, for example, to launch the car from 0 to 800 mph in less than a minute. Achieving the sheer power for a project such as the Thrust SSC requires a radical design. The Thrust SSC features a tubular steel and carbon composite structure with a totally unique rear wheel steering system. Power is provided by two side-mounted Rolls Royce Spey 205 jet engines. This type of engine is normally found on jet fighter aircraft.

Approximately 54 feet long, this seven-ton, 100,000-horsepower land vehicle is the most powerful car ever built. The challenge in designing it lies in ensuring stability and control. At speeds in excess of 850 mph, even a small flaw can spell disaster. For instance, at supersonic speed, even the slightest twitch upwards (0.5 degrees) could cause the vehicle to flip over onto its back. Keeping the nose level is absolutely crucial to the safety of the Thrust SSC and a secret part of its design.

Evaluating the design of a land vehicle that can charge at supersonic speeds cross country is another challenge. To accomplish this task, the design of the Thrust SSC has been analyzed using a combination of state-of-the-art simulation and experimental testing. The state-of-the-art technology includes Cray Research supercomputers. As a sponsor of the Thrust SSC project, Cray Research offered services-in-kind to provide the computing time to evaluate the aerodynamics of the vehicle.

Using FLITE3D—an application system developed by Computational Dynamics Research and the Department of Civil Engineering at the University of Wales Swansea—Dr. Obey Hassan modeled the entire vehicle and studied the airflow around the Thrust SSC at various speeds on a CRAY C92 system. The boundaries of the computational domain are represented by an unstructured assembly of triangles, and the computational domain is represented by an unstructured assembly of tetrahedra. The computations assumed steady and inviscid flow and took account of the engine effects by defining appropriate engine inflow and outflow conditions. The steady state is achieved by means of an approach based upon explicit timestepping. Each simulation took approximately three hours on a CRAY C92 system.

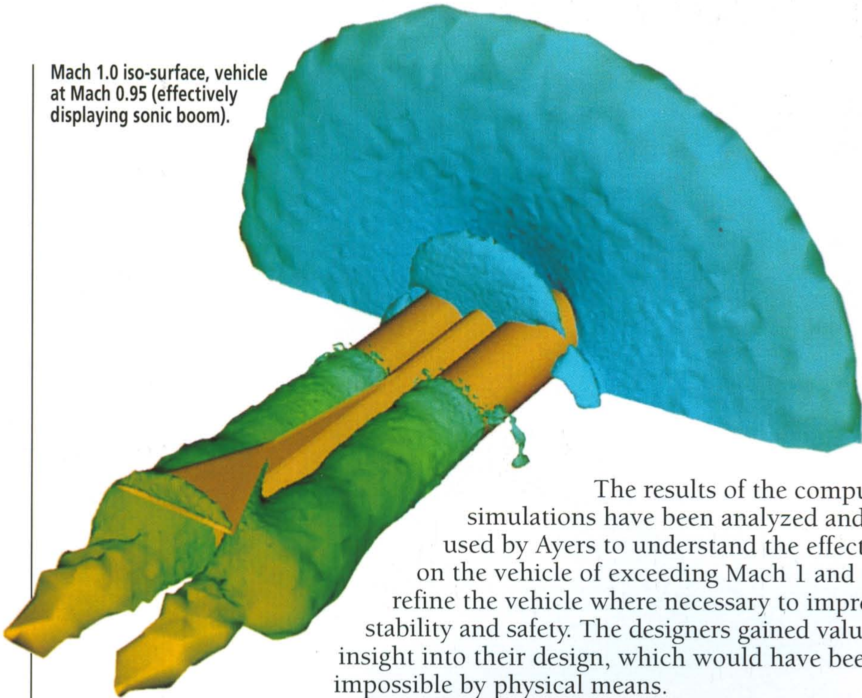

Mach 1.0 iso-surface, vehicle at Mach 0.95 (effectively displaying sonic boom).

The results of the computer simulations have been analyzed and used by Ayers to understand the effects on the vehicle of exceeding Mach 1 and to refine the vehicle where necessary to improve stability and safety. The designers gained valuable insight into their design, which would have been impossible by physical means.

To further evaluate the design, experimental testing was done to study ground effects. A rocket-powered model of the Thrust SSC was created and launched down a track traveling at speeds as high as 850 mph. The results? The simulated ground effects closely matched the data gathered from the computer simulations. According to Hassan, "It was quite exciting to have the results of the simulation validated through the experiments." The data gathered through these and other tests are crucial to the continued progress of the project.

Despite these careful preparations, there is undeniably inherent risk in traveling at the speeds required to break the world land-speed record. Consequently, safety is a primary concern. Indeed, considerable design effort has gone into providing the driver, Andy Green, as much protection as possible against physical injury. The results of any test runs will be used to continue to improve upon or enhance the vehicle during the lead-up phase. The record breaking attempts will take place within the next two years.

When test runs and the supersonic car are completed, the plan is to take the Thrust SSC to the Nevada desert for its race against the record. The Nevada desert hosted Noble 12 years ago when he became "the fastest man on earth." Now Green is planning to go to the same place in man's historic challenge against the forces of nature.

About the author

Kevin Fox is an applications consultant at Cray Research (UK) Ltd.

References

Bailey, Eric, "Jet pilot volunteers at speed of sound for land record ride," *The Daily Telegraph*, London, England, February 4, 1995.

Farish, Mike, "Building a Super Car," *Engineering*, Vol. 236, No. 3, pp. 16-18, March 1995.

Fexall, James, "Professor out to cut hit record," *The Auto Express*, No. 328, pp. 28-29, January 13, 1995.

Acknowledgment

All images were produced using EnSight from Computational Engineering International, Inc. EnSight is a distributed, interactive results postprocessor for various applications including stress and computational fluid dynamics analyses.

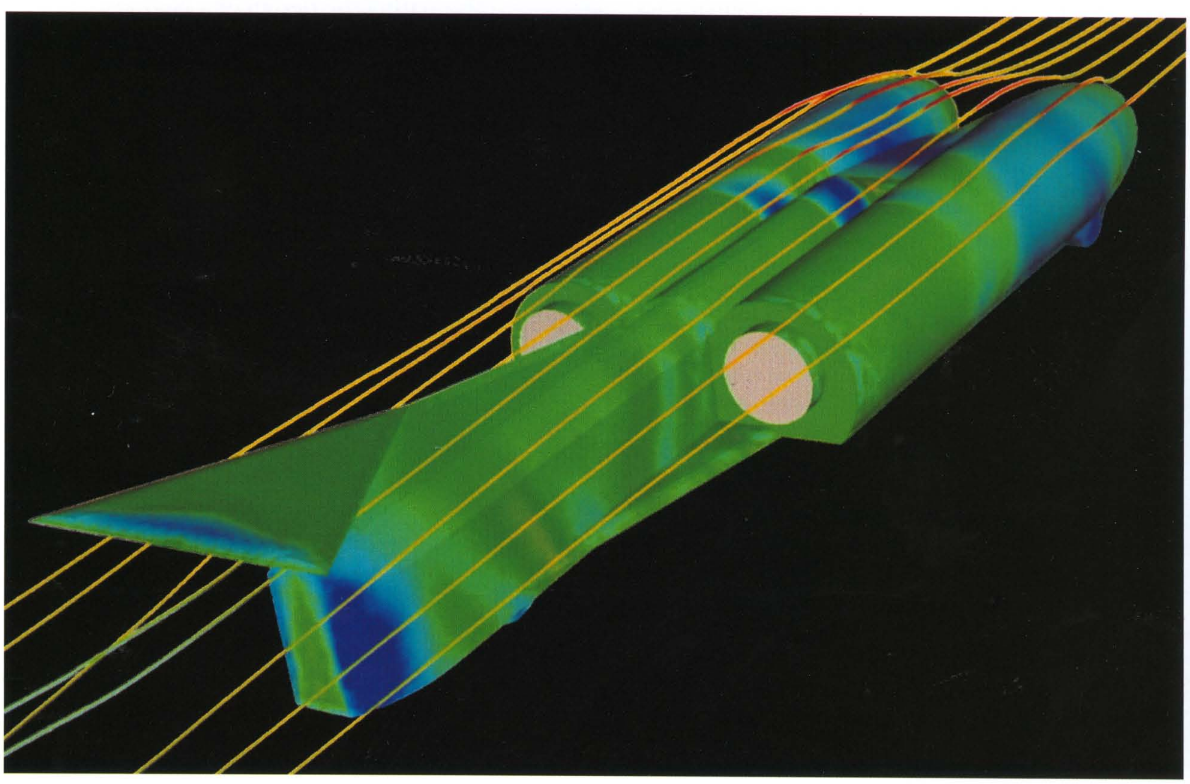

Surface pressure with airflow streamlines, vehicle at Mach 1.15.

A new algorithm for die surface development in sheet metal forming

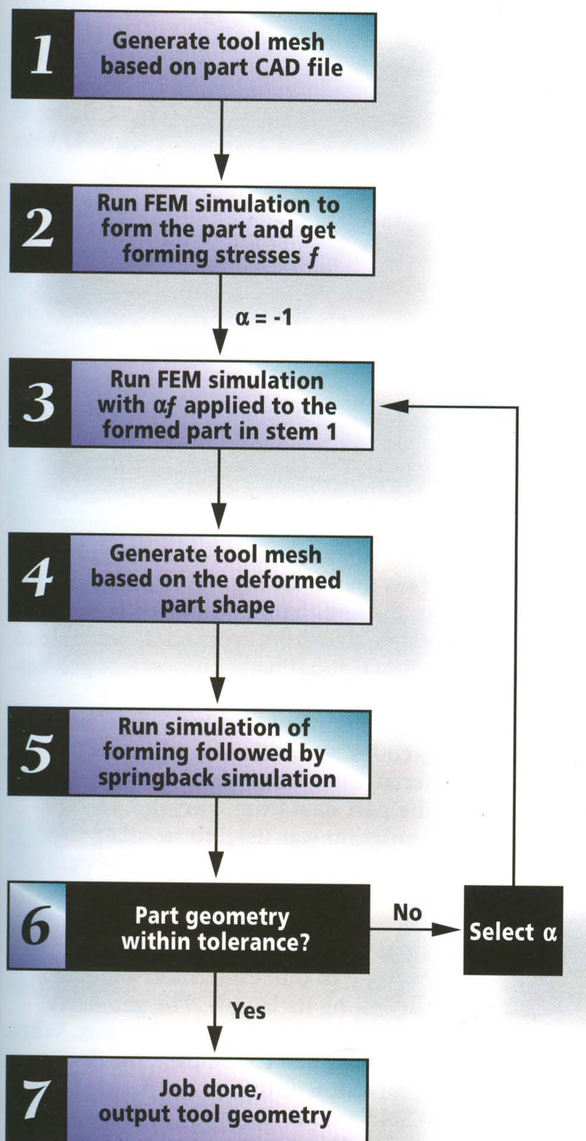

Figure 1. Flow diagram of the die-surface modeling algorithm.

Changquing Du
and Li Zhang
Chrysler Corporation

Longwu Wu
Cray Research, Inc.

Cray Research recently developed a design algorithm to help sheet-metal stampers produce three-dimensional sheet-metal parts faster and more cost-effectively. This new algorithm reduces the amount of guesswork involved in compensating for springback in die surface development. Based on recent advances in finite element method (FEM) analysis, the iterative FEM die design algorithm generates tooling geometry that produces three-dimensional sheet metal parts of the desired shape after springback.

FEM analysis has been employed increasingly by die surface development engineers to simulate sheet-metal-forming processes and produce information on splitting and wrinkling. Researchers now report success in predicting springback in sheet-metal forming through FEM simulations. The result of springback analysis, usually shown as the amount of over-bending or over-crowning, can only provide limited clues as to how to modify the die surfaces to compensate for springback. It requires much guess work, building new tooling models, and conducting new analyses to check out those alternatives. Even though this approach is already far better and cheaper than traditional hardware tryouts, it is still a trial-and-error methodology.

In collaboration with Chrysler, Cray Research has developed an algorithm to systematically update the die surface while compensating for springback using iterative forming and springback simulations. The end result is a die surface definition which can produce the desired net part shape.

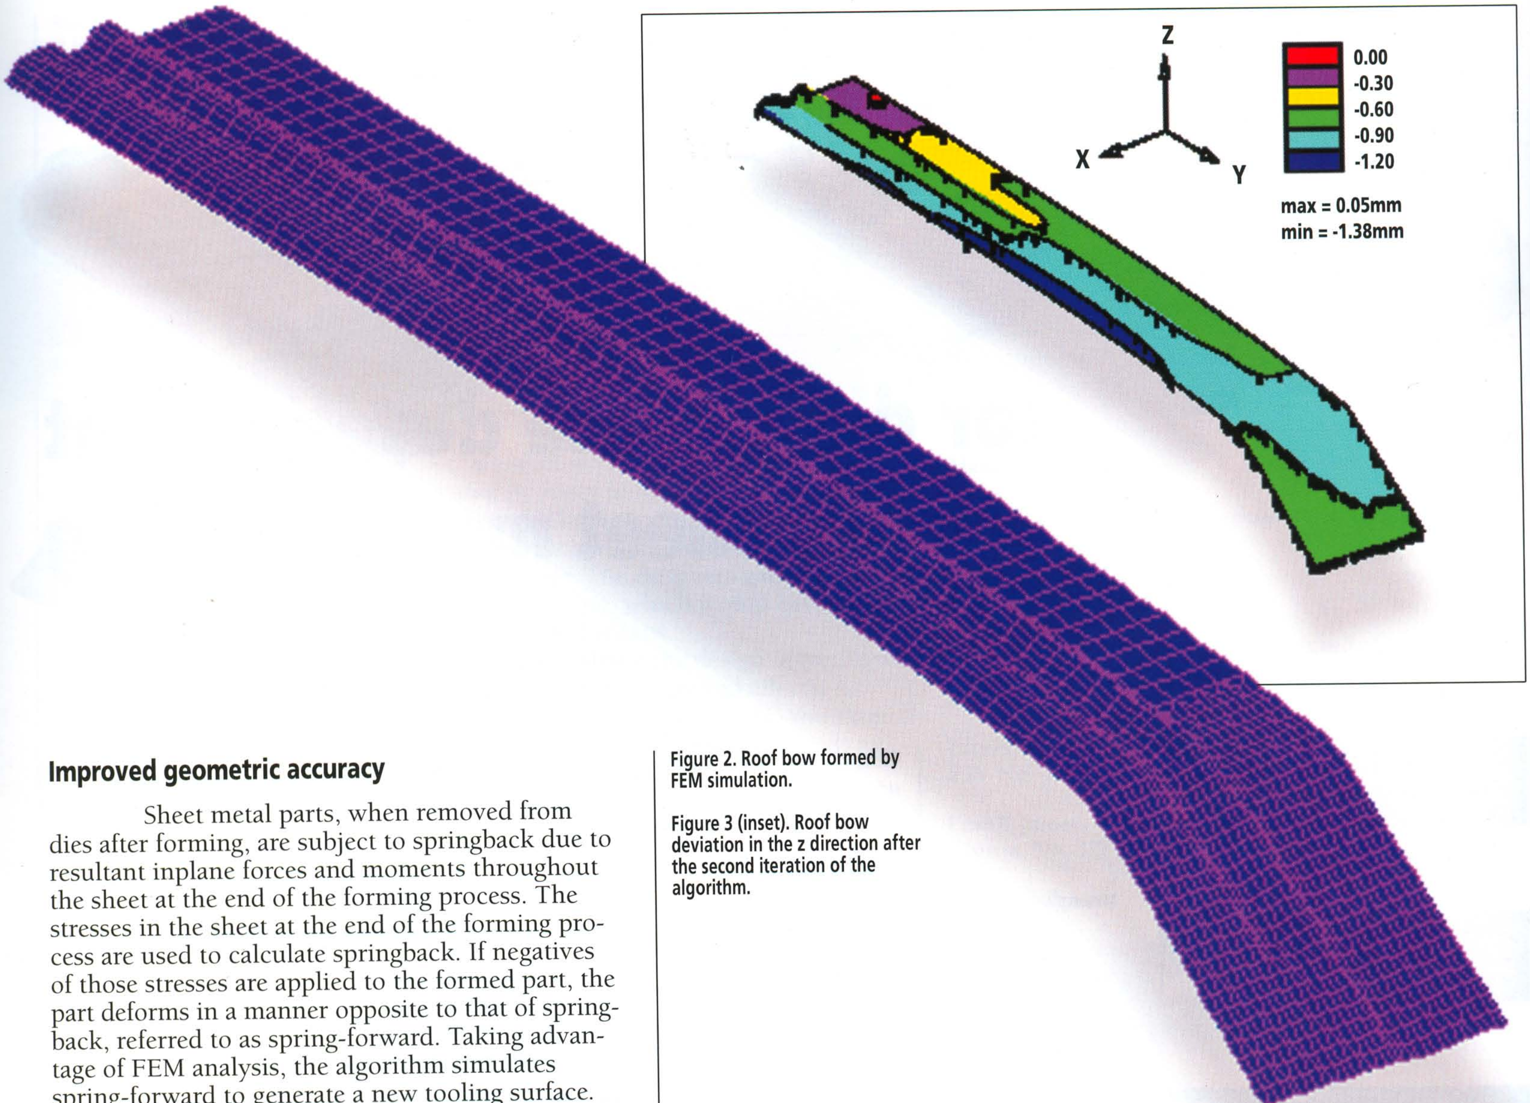

Figure 2. Roof bow formed by FEM simulation.

Figure 3 (inset). Roof bow deviation in the z direction after the second iteration of the algorithm.

Improved geometric accuracy

Sheet metal parts, when removed from dies after forming, are subject to springback due to resultant inplane forces and moments throughout the sheet at the end of the forming process. The stresses in the sheet at the end of the forming process are used to calculate springback. If negatives of those stresses are applied to the formed part, the part deforms in a manner opposite to that of springback, referred to as spring-forward. Taking advantage of FEM analysis, the algorithm simulates spring-forward to generate a new tooling surface. Die surface designs based on the shape of the part after spring-forward will produce more geometrically accurate parts.

The objective of the algorithm is to obtain a geometric description of tooling that will form parts of desired shape after springback. As in any iteration scheme, iteration zero starts with an initial guess, which in this case is the designed part shape. Therefore, the FEM mesh is first generated from the CAD file of the part design, and a standard simulation of forming is performed by FEM analysts.

After the simulation of forming, analysts conduct two additional simulations. The first is the springback simulation, during which the deviation of a critical dimension from the design value is measured. The second simulation is the full spring-forward. It results in a deviation that is usually very close to the negative of the initial deviation in the critical design dimension. Based on the part shape that results in spring-forward simulation, the FEM mesh of tooling for the next iteration is prepared. The algorithm iterates until the geometry of the part after springback is within tolerance. Once the geometry of the part after springback is within tolerance, the die geometry data used in this iteration can be output and fed back to a CAD/NC system. Figure 1 presents a flow diagram of the algorithm.

Partnership with Chrysler

With the help of Chrysler Corporation, this new algorithm was applied to two sheet metal parts—a roof bow and a rail cap. The tests performed on these parts showed that this new iterative FEM die surface design algorithm worked. It was able to handle three-dimensional parts and correct geometrical deviations to an acceptable level.

Livermore Software Technology Corporation's (LSTC) codes were used on a CRAY C90 supercomputer at Cray Research to test the algorithm using the part data from Chrysler. LSTC's LS-DYNA3D program was used to simulate the forming of the part, and the LS-NIKE3D program was used for the springback and spring-forward calculations.

The first test was run on a roof bow (Figure 2) that was long in one dimension and had a gentle curvature along its length. The geometry features caused severe springback in this part. When the algorithm was applied to the part, the deviation in the normal direction at the end of the part was reduced from 9.73 mm to 1.38 mm, a reduction of 80 percent (Figure 3).

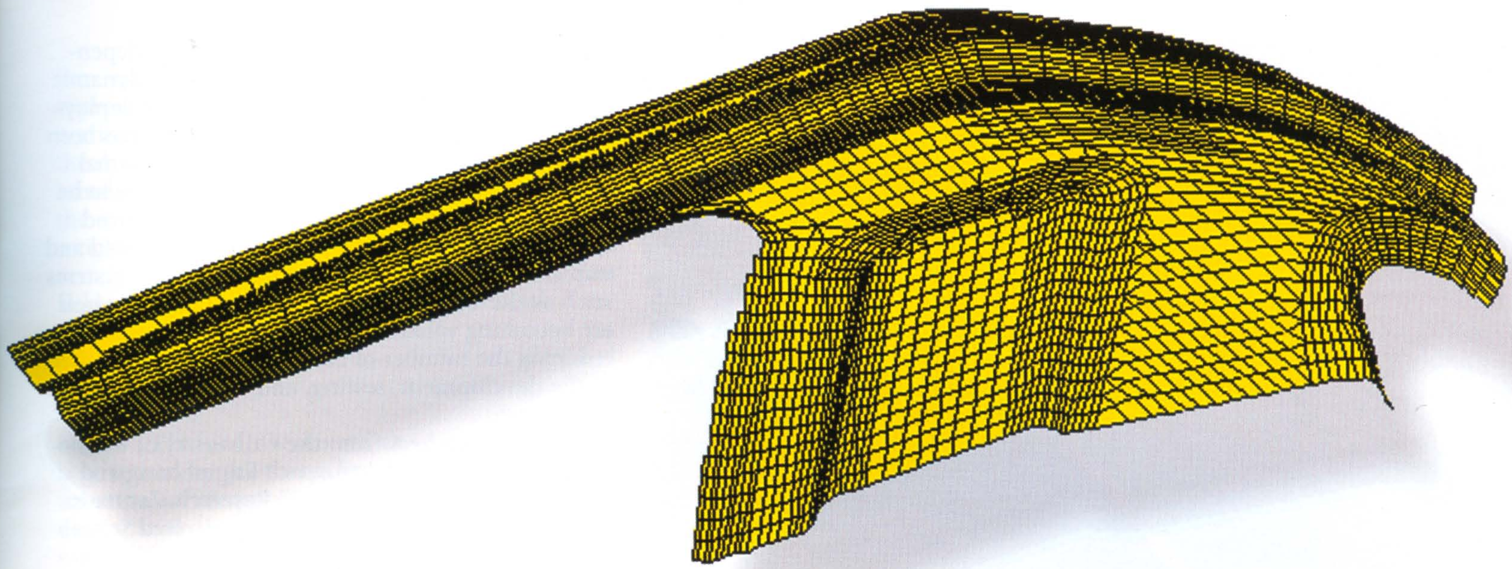

Figure 4. Rail cap formed by FEM simulation.

In the second set of tests, the algorithm was applied to a Chrysler rail cap. The rail cap had a more complicated geometry and forming process. The tooling was composed of several pieces, including a punch, die, pad, and flanging punch. The formed part is shown in Figure 4. After springback and spring-forward simulations, the geometric inaccuracy at the left end of the part was reduced from 10.3 mm to 0.88 mm, a reduction of 90 percent (Figure 5).

Successes to date warrant further verification with more difficult parts, such as pillar and rail sections where springback is more pronounced. A semiautomated algorithm is also being developed to help guide the engineer through the iterative processes.

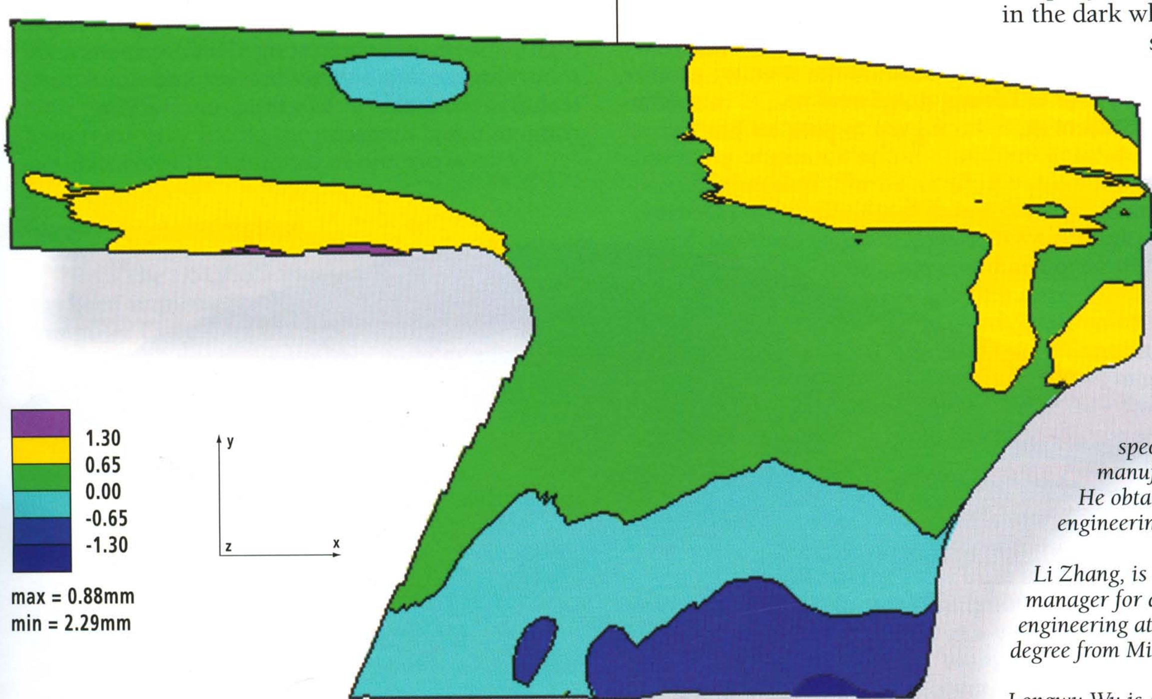

Figure 5. Rail cap deviation in the z direction after the second iteration of the algorithm.

A valuable algorithm

The new Cray Research algorithm directly addresses problems that die surface design engineers face every day when determining die geometries from the shapes of parts. It is the only systematic approach to date to compensate for springback in three-dimensional parts of complicated geometry. When the algorithm is applied properly, it reduces the number of physical tooling tryouts needed to produce parts within geometrical tolerances.

The Cray Research algorithm is applicable to sheet-metal-forming processes in any manufacturing industry, and automakers have begun showing interest in it. "The algorithm will be a very valuable tool for improving our stamping processes," said Dr. Sing Tang, a senior staff scientist at Ford Motor Company. "As stampers, we are otherwise in the dark when trying to compensate for springback." The potential for manufacturers is significant. Use of the algorithm can directly impact the bottom line through increased production efficiencies and significant cost savings.

About the authors

Changquing Du is a stamping specialist in advanced stamping manufacturing engineering at Chrysler. He obtained a Ph.D. degree in structural engineering from the University of Maine.

Li Zhang, is a product and development manager for advanced stamping manufacturing engineering at Chrysler. He received an MSME degree from Michigan Technological University.

Longwu Wu is a senior applications analyst at Cray Research. He received a Ph.D. degree in mechanical engineering from Drexel University in Philadelphia.

Modeling *large ram air* parachutes *on the CRAY T3D system*

Tayfun E. Tezduyar
Aerospace Engineering and Mechanics
Army HPC Research Center
University of Minnesota

The ram air parachute ("parafoil") has replaced the more familiar hemispherically shaped parachute in amateur skydiving. Because of their gliding capability, reliability, and aerodynamic efficiency, these parachutes are used increasingly for the delivery of large payloads. Larger parachutes, with canopy sizes of 7000 square feet, are now being developed by the Army to deliver payloads of up to 21 tons from altitudes as high as 25,000 feet.

One of the features of ram air parachutes of particular interest to the Army is their ability to make a soft landing by means of a flare maneuver. During this maneuver the trailing edges of the parachute are retracted to increase the effective camber. The increase in lift and drag forces results in a reduction of velocity, reducing the impact on landing.

Realistic three-dimensional, time-dependent computer-based modeling of the aerodynamic forces and stresses encountered during the deployment and gliding stages of parachute flight has been limited in the past by insufficient computational resources. While these simulations continue to be a major computational challenge, new advanced finite element simulation techniques developed and implemented on very large memory, scalable systems such as the CRAY T3D system from Cray Research are becoming valuable and cost-effective tools for reducing the number of costly airdrop tests required in the development, testing, and evaluation of these parachutes.

Researchers from the University of Minnesota's Department of Aerospace Engineering and Mechanics and the Army HPC Research Center in collaboration with researchers at the United States Army Natick Research, Development, and Engineering Center are developing the advanced finite element flow simulation techniques to model a complete parachute configuration in three-dimensional space, from its initial deployment and unfolding through its gliding phase and impact with the ground. The goal of this research is to develop techniques that will allow the rapid and cost-effective design and fielding of new airdrop systems.

The aerodynamics of ram air parachutes requires the analysis of a large number of complex phenomena. Deployment, expansion, and transition to the gliding stage involve rapidly changing geometries, unsteady and turbulent flow behavior, and nonlinear interactions between the parachute structure, aerodynamics forces, and the payload. Major deformations of the canopy, changes in the orientation of the parachute, and relative motion between the canopy and the payload occur even during the gliding phase and need to be modeled carefully. These time-dependent conditions with changing geometries require very large memory, scalable high-performance computing (HPC) systems and robust and accurate three-dimensional simulation techniques capable of handling time-varying computational domains.

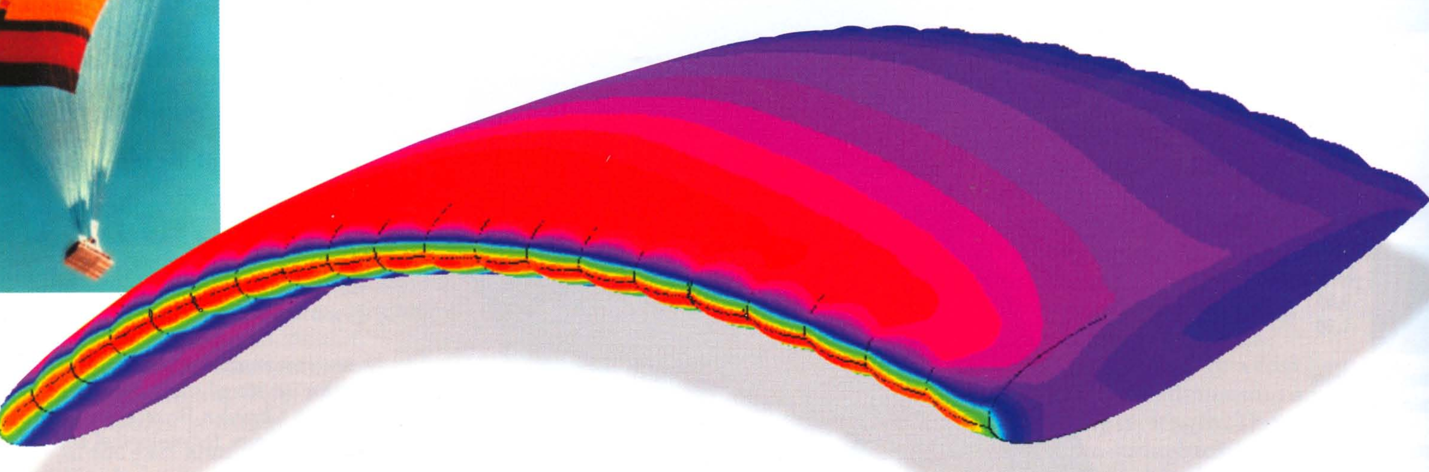

Figure 1. Steady-state flow around a parachute; pressure distribution on the parachute surface. The mesh consists of 9,741,930 nodes and 9,595,520 hexahedral elements. To reach the steady-state solution, 38,391,715 coupled nonlinear equations are solved simultaneously at every pseudo-timestep. The insert shows an actual drop test of a medium-sized parachute.

At the present phase of the research, the time variation of the geometry of the canopy is approximated based on the initial and final configurations and dimensions of the canopy. However, the dynamics of the canopy, i.e., its motion with translation and rotation, is determined as part of the overall solution. This motion depends on the weight and motion of the payload as well as on the aerodynamic forces generated by the unsteady flow field which is an even greater computational challenge.

Formulation

The air flow around the parachute canopy is governed by the three-dimensional Navier-Stokes equations of incompressible flows with time-dependent spatial domains. The three-dimensional dynamics of the canopy is governed by Newton's laws of motion with the forces acting on the canopy calculated from the simulated flow field. The deforming-spatial-domain/stabilized-space-time (DSD/SST) finite element formulation is used to handle the time-dependent spatial domains.[1,2] In this formulation, the finite element interpolation polynomials are functions of both space and time, and the stabilized variational formulation of the problem is written over the associated space-time domains. Globally, the interpolation functions are continuous in space but discontinuous in time. The solution is obtained one space-time slab at a time.

A major algorithmic issue is the methodology for updating the finite element mesh as the spatial domain changes its shape over time. There are several ways to manage this.[3] In the work described in this article, the motion of each finite element grid point is prescribed explicitly. This is accomplished without remeshing (i.e., without generating a new set of nodes and elements). With this approach, the connectivity of the mesh remains the same throughout the simulation. As a result, the cost of mesh generation and that associated with parallelization setup are minimized. This is important in reducing the computational time of simulations with hundreds of timesteps.

Parallel implementation

The DSD/SST finite element formulation has been implemented on the Minnesota Supercomputer Center's 512-node CRAY T3D supercomputer. The code is written in Fortran 77 optimized for the CRAY T3D system's Alpha chip nodal architecture and Parallel Virtual Machine (PVM) library for interprocessor communication. The finite element mesh is partitioned into load-balanced contiguous subdomains, with each subdomain assigned to one processing node. The need for PVM-based communication between elements is therefore reduced to the subdomain boundaries. The coupled nonlinear equations encountered at every timestep or pseudo-timestep are solved iteratively using matrix-free iterations. The scalable performance currently observed on the CRAY T3D system is of the order of 20 MFLOPS per processing node. A typical simulation executes at about 10 GFLOPS on the 512-node CRAY T3D system.

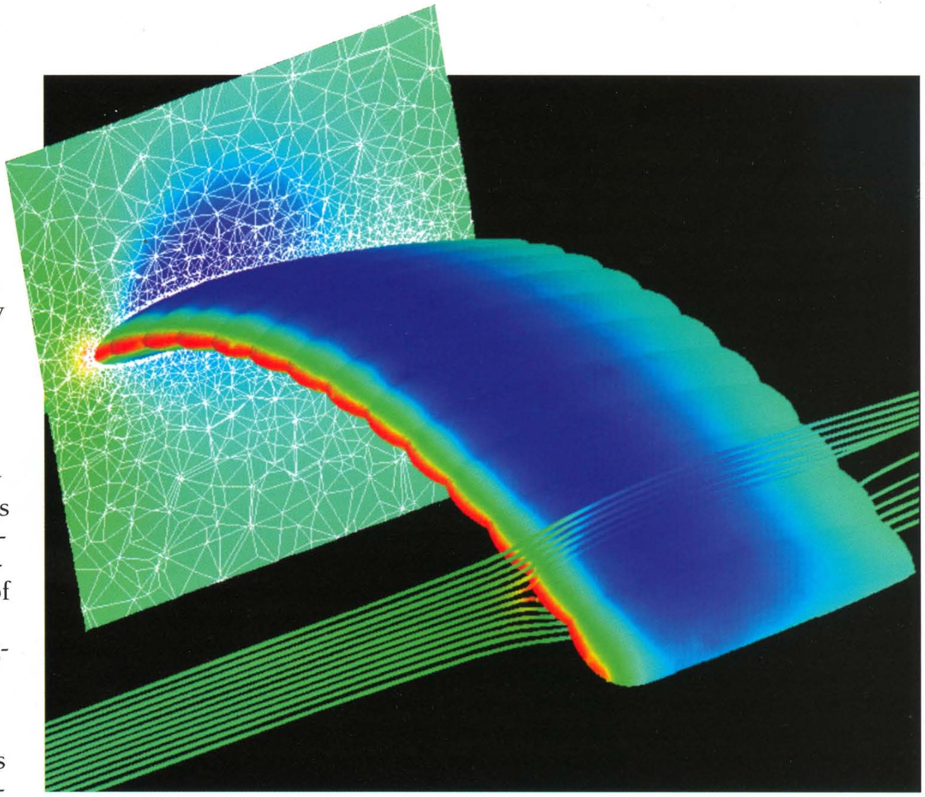

Figure 2. Steady-state flow around a parachute using an unstructured mesh. Pressure distribution on the parachute surface, mesh, and pressure distribution at a cross section, and stream tubes color-coded with the pressure. The mesh consists of 144,649 nodes and 905,410 tetrahedral elements. To reach the steady-state solution, 506,119 coupled nonlinear equations are solved simultaneously at every pseudo-timestep.

Simulations

The air flow around a parachute is very complex, requiring very fine grids for adequate resolution of flow features. The availability of cutting-edge computational resources such as the CRAY T3D system takes us one step further toward obtaining more accurate solutions. Additional examples are included in Reference 4.

Two simulations are presented here. The first simulation depicts the steady-state performance of the parachute, and the second addresses its pitch stability. The steady-state performance is simulated using two finite element meshes. Figure 1 shows the result from the finite element simulations with over 9.7 million grid points. To reach the steady-state solution, over 38 million coupled nonlinear equations are solved simultaneously at every pseudo-timestep. The input file for this simulation consisted of 1.2 Gbytes of data representing the nodal coordinates, mesh connectivity, and boundary conditions. The output file contained over 300 megabytes of data for a single timestep.

Figure 2 shows the steady-state solution obtained from a simulation using an unstructured finite element mesh. Both the modeler and mesh generator were developed by the Aerospace Engineering and Mechanics Finite Element Group at the Army HPC Research Center.

In the second simulation, the dynamic pitch stability of the parachute is investigated (Figure 3). Dynamic analysis is important in evaluating the ability of the parachute to rapidly settle into steady glide. In this simulation the parachute is released almost vertically at a very low speed. It exhibits an initial tendency to rapidly pitch forward, go into a dive, and then recover into steady glide. During the next phase of our research we will extend this simulation to begin to model the behavior of the parachute during a flare maneuver.

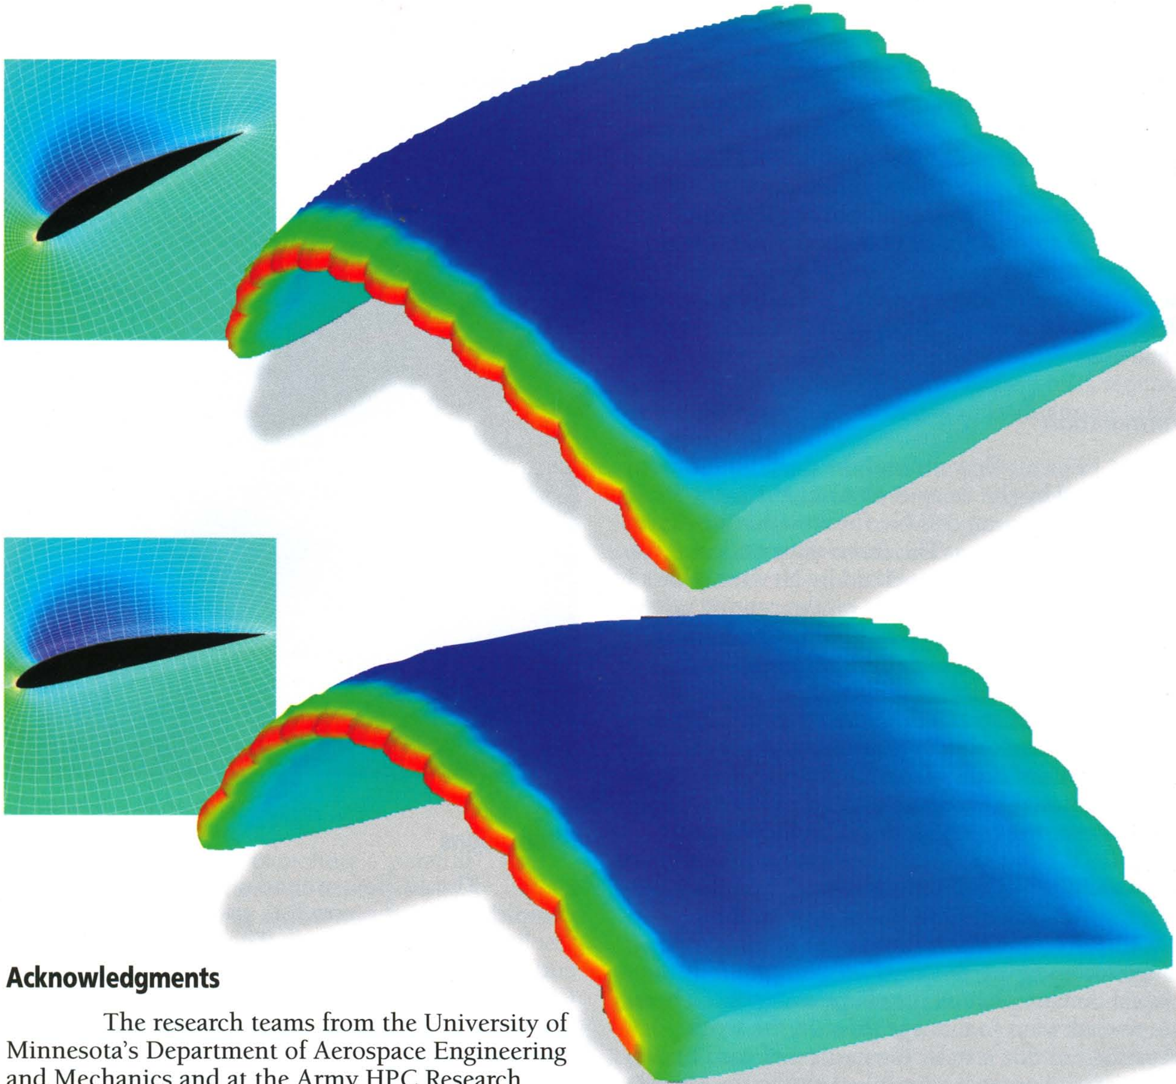

Figure 3. Parachute in gliding descent. Pressure distribution on the parachute surface at two instants during descent. The mesh consists of 291,347 nodes and 279,888 hexahedral elements. At every timestep 2,225,496 coupled nonlinear equations are solved simultaneously.

Acknowledgments

The research teams from the University of Minnesota's Department of Aerospace Engineering and Mechanics and at the Army HPC Research Center include Shahrouz Aliabadi, William Garrard, Vinay Kalro, Sanjay Mittal, Tayfun Tezduyar, and Chris Waters. The research team from the US Army Natick Research, Development, and Engineering Center includes Earl Steeves and Keith Stein. Marek Behr and Andrew Johnson of the Army HPC Research Center provided parallel programming and automatic mesh generation support. The project is supported by the Army Research Office (Program Officer Jagdish Chandra) and by Army Research Laboratory, Army HPC Research Center (Program Officer Walter Sturek). No official endorsement should be inferred. CRAY C90 system time was provided in part by the University of Minnesota Supercomputer Institute. CRAY T3D and CRAY C90 system time was provided in part through arrangements with the Minnesota Supercomputer Center, Inc.

About the author

Tayfun Tezduyar received an M.S. degree in mechanical engineering from The California Institute of Technology in 1978 and a Ph.D. degree in 1982. After postdoctoral work at Stanford he joined the University of Houston as assistant professor in 1983 and the University of Minnesota as associate professor in 1987. He has been a professor of aerospace engineering and mechanics since 1991 and the director of the Army HPC Research Center since January 1994. Dr. Tezduyar holds a 1986 Presidential Young Investigator Award and was elected a Fellow of the American Society of Mechanical Engineers in 1994. He has published over 120 papers and is the Computational Fluid Dynamics editor for IEEE Computational Science & Engineering.

References

1. Tezduyar, T. E., M. Behr, and J. Liou, "A New Strategy for Finite Element Computations Involving Moving Boundaries and Interfaces—The DSD/ST Procedure: I. The Concept and the Preliminary Numerical Tests," *Computer Methods in Applied Mechanics and Engineering*, 94, pp. 339-351, 1992.
2. Tezduyar, T. E., M. Behr, S. Mittal, and J. Liou, "A New Strategy for Finite Element Computations Involving Moving Boundaries and Interfaces—The DSD/ST Procedure: II. Computation of Free-surface Flows, Two-liquid Flows, and Flows with Drifting Cylinders," *Computer Methods in Applied Mechanics and Engineering*, 94, pp. 353-371, 1992.
3. Johnson, A. A. and T. E. Tezduyar, "Mesh Update Strategies in Parallel Finite Element Computations of Flow Problems with Moving Boundaries and Interfaces," *Computer Methods in Applied Mechanics and Engineering*, 119, pp. 73-94, 1994.
4. Tezduyar, T. E., S. Aliabadi, M. Behr, A. Johnson, V. Kalro, and C. Waters, "3D Simulation of Flow Problems with Parallel Finite Element Computations on the CRAY T3D," *Computational Mechanics '95, Proceedings of International Conference on Computational Engineering Science*, Mauna Lani, Hawaii, 1995.

Mesh Generation the Easy Way

HEXAR

The Automatic Mesh Generator

Jean Cabello, Yoshihiko Mochizuki, and Thomas Raeuchle, Cray Research, Inc.

Defining mesh generation's role in the design process

Today, mesh generation is the single most time-consuming part of the computer-aided engineering (CAE) process. Using traditional manual mesh generation techniques, it can take an engineer months to hand-craft the mesh needed to analyze product designs through computer simulation. Even though an engineer might spend months building a mesh, the quality of the resulting mesh is not only limited by time, but also by the skills and experience of the engineer. And since the quality of the mesh directly impacts the accuracy of the simulation, mesh generation is critical to the effectiveness of the entire design process.

Mesh generation faces other issues as part of the design process. First, design problems are becoming bigger and more complicated. Second, engineers are making remarkable progress in improving the accuracy and efficiency of the numerical schemes that represent the physics of the simulation. As a result of these two issues, engineers are demanding even more complex meshes. The use of more complex geometries is making it even more challenging to accelerate time-to-market and improve the quality of new products through CAE.

Regardless of the area of industrial application, mesh generation is setting the pace of the design process, limiting the accuracy of simulation, and dictating the overall length of the design cycle (Figure 1). And in a highly competitive marketplace, quality designs and time-to-market can determine a product's success or failure.

Eliminating the mesh generation bottleneck

To remove the mesh generation bottleneck in the design process and improve the efficiency with which sophisticated computer meshes are created, Cray Research developed HEXAR, a pioneering software application targeted for use by the automotive, aerospace, chemical, civil engineering, casting and molding, materials and plastics, electronics, semiconductor, and biomedical industries. HEXAR is designed to save product developers months of time and hundreds of thousands of dollars by automatically converting raw computer-aided design (CAD) data for any complex object into the sophisticated three-dimensional hexahedral meshes that engineers use to simulate and analyze new product or component designs. A million-element, hexahedral mesh can take six months and $150,000 or more to develop. Now, using HEXAR software, this process typically can take less than 30 minutes on a Cray supercomputer.

Surpassing the competition

No product available on the market today is equivalent to HEXAR. HEXAR is fully automatic. It automatically generates hexahedral meshes which in some analyses are preferred for more accurate modeling and simulation. Other mesh-generation software products available today either require substantial manual intervention or fail to produce meshes that contain only hexahedrals. Moreover, competitive products are not as robust and often contain degenerated elements that detract from the

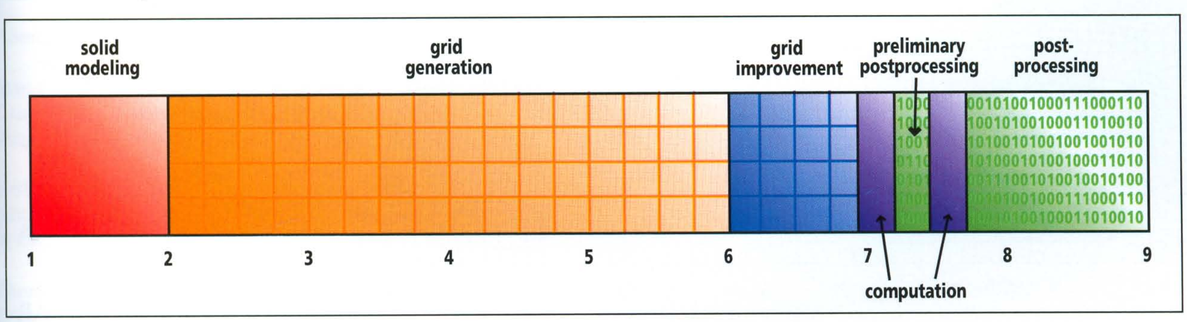

Figure 1. Relative duration of each task in a large-scale CFD project.

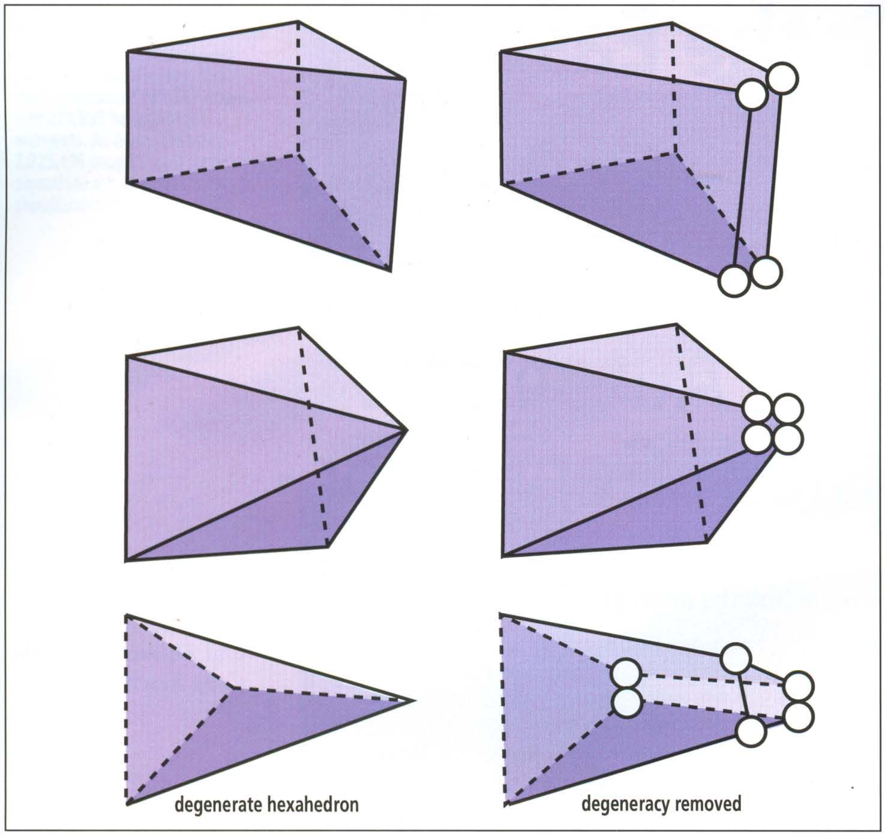

Figure 2 (above). Removal of hexahedral degeneracies.

Figure 3 (below). Hexahedral mesh of a soil formation that was analyzed for static stability.

quality of the mesh and ultimately the results of the analysis.

Generating meshes the easy way

HEXAR is based on a patented approach that mimics manual mesh-generation methods. HEXAR software begins by automatically converting CAD data to a generic triangulated surface definition (TSD) format using one of the utility programs provided in the HEXAR package. The TSD provides HEXAR the information it needs about the outer geometry of the object to successfully generate the mesh inside. As it generates the mesh, HEXAR automatically handles gaps and overlaps that are often produced by CAD packages. To initiate this process, all the user has to do to get set up is input the TSD name, the required mesh size, and, if needed, a gap tolerance. Once the TSD is given, the set-up time and requirements are constant regardless of the complexity of the geometry.

After the TSD file is read, HEXAR extracts various features of the surface and generates a surface map that represents the outer skin of an entire family of all hexahedral volume meshes. Once the map file is created, the first stage of mesh conditioning takes place, first at the surface level and then at the interior level. To provide the user flexibility in managing the size of the resulting mesh, Cray has included a grid coarsening feature. Grid coarsening gives the user the ability to reduce the number of elements in the mesh in order to reduce the CPU time during an analysis.

HEXAR applies complex conditioning methods to ensure the best-quality mesh elements. However, in some cases, this mesh generation procedure results in the creation of degenerate hexahedra in certain areas. HEXAR automatically detects flaws and replaces defective hexahedral elements (Figure 2) with hexahedra that are not degenerate at the expense of incurring some departure from the CAD geometry.

HEXAR software produces meshes for the most widely used CAE formats, including PATRAN, IDEAS, and EnSight. Consequently, the HEXAR product works with nearly all third-party engineering analysis software packages used today including ABAQUS, ANSYS, NASTRAN, FIDAP, STAR-CD, FLOW-3D, PROCAST, and MSC/EMAS. These formats are readable by most commercial CFD and structural simulation software.

Using an automatic mesh generator for real-world application

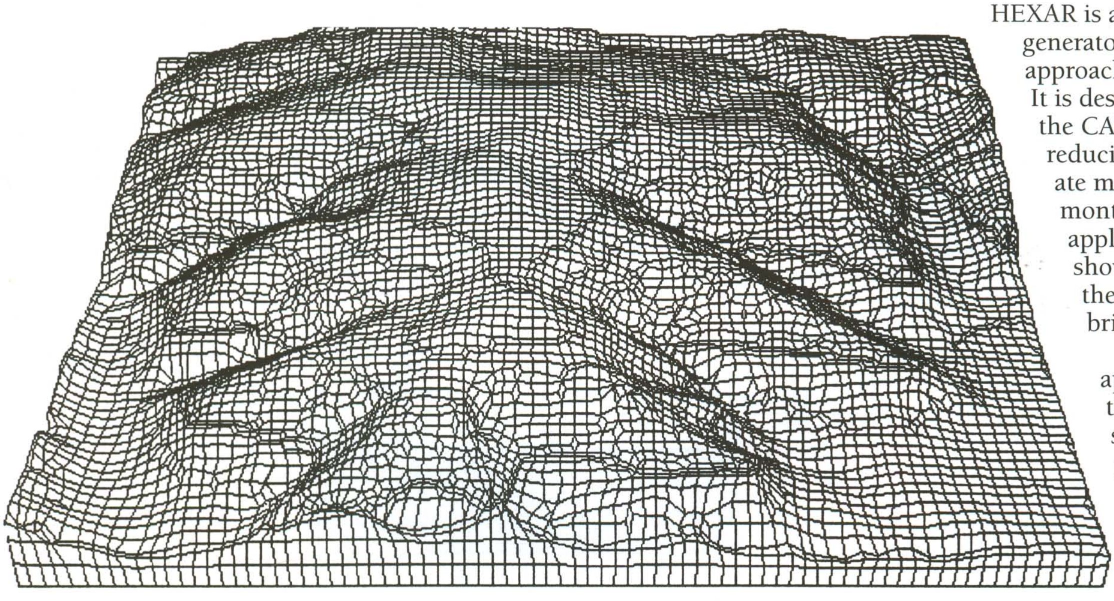

HEXAR is a high-quality automatic mesh generator that represents a novel approach to mesh generation for CAE. It is designed to minimize the length of the CAE process by significantly reducing the time required to generate meshes—in some cases from months to minutes. The real-world applications described below clearly show the speed and effectiveness of the software and the value it brings to the design process.

In a civil engineering application, a hilly soil formation was analyzed for static stability. Running on a single processor of a Cray supercomputer, HEXAR generated a 30,000-element hexahedral boundary-fitted mesh in less than one minute (Fig. 3).

In a biomedical research application, HEXAR

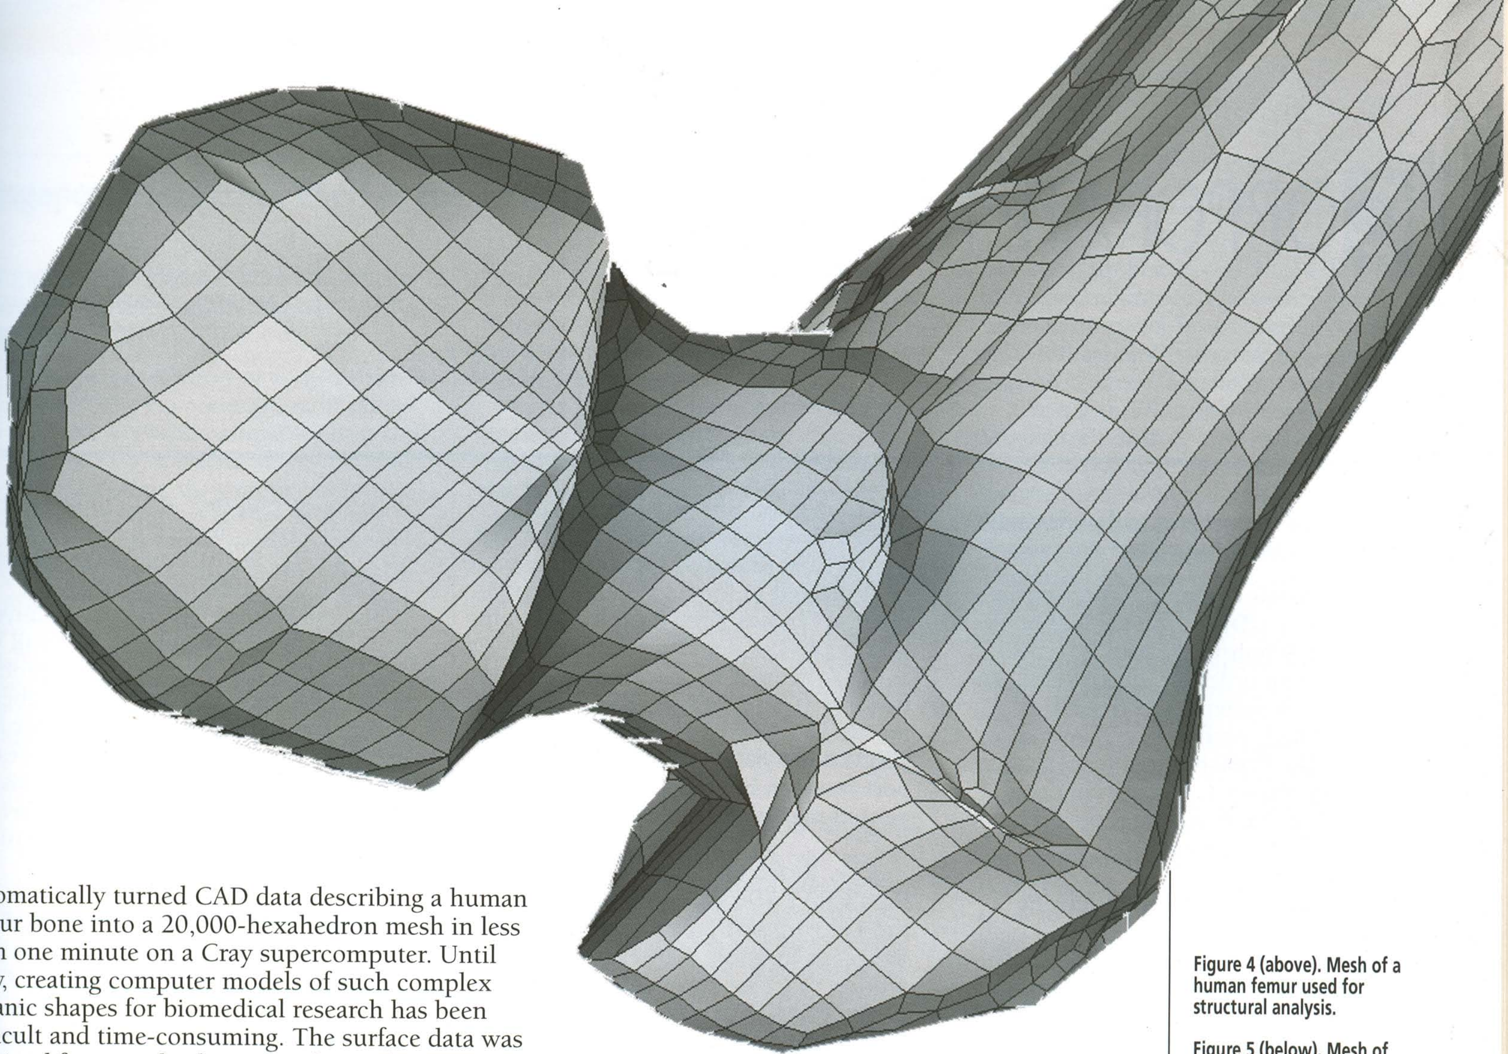

automatically turned CAD data describing a human femur bone into a 20,000-hexahedron mesh in less than one minute on a Cray supercomputer. Until now, creating computer models of such complex organic shapes for biomedical research has been difficult and time-consuming. The surface data was obtained from medical imaging data. The mesh was used to compute structural load distribution using a commercial finite element package (Figure 4).

In a CFD application, HEXAR automatically generated a 400,000-hexahedron mesh for an underhood cab compartment of a Ford Taurus engine. The analysis was completed on a Cray supercomputer in just over two hours (Figure 5).

HEXAR is well suited for use in a wide range of industries and areas of application, including casting and metal forming, biomedical engineering, petroleum reservoir simulation, soil mechanics, structural analysis, and electromagnetics. Combined with the speed of Cray supercomputing systems, HEXAR makes it easier for industries worldwide to move from CAD to CAE. HEXAR is a genuine breakthrough product for engineers and researchers. It's automatic. It's easy-to-use. It's affordable. And it provides an effective solution to the most time-consuming aspect of the design process—mesh generation.

Figure 4 (above). Mesh of a human femur used for structural analysis.

Figure 5 (below). Mesh of underhood cab compartment of a Ford Taurus engine, used for airflow and thermal analyses.

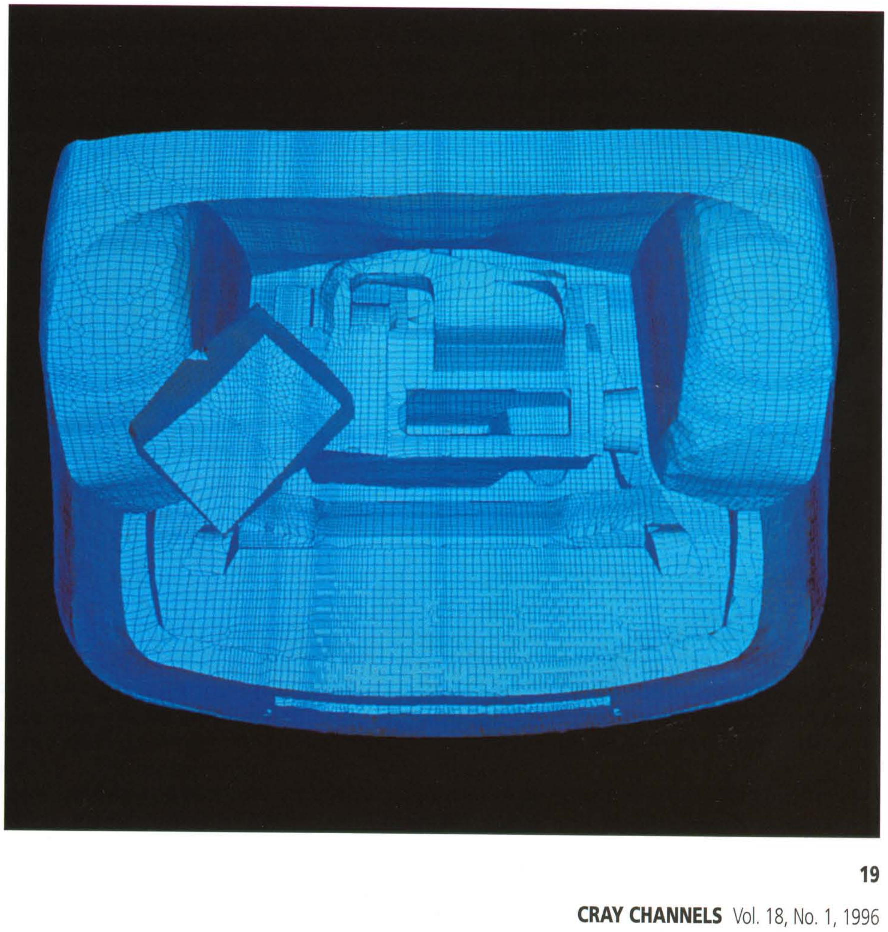

About the authors

Jean Cabello is the technical leader of the HEXAR project at Cray Research. He earned a Ph.D. degree in numerical analysis from the University of Paris Sud in 1990.

Yoshihiko Mochizuki developed the prototype of HEXAR. He joined Cray Research in 1994 after obtaining a Ph.D. degree in nuclear engineering from the University of Tokyo.

Thomas Raeuchle is the Director of Applications Products at Cray Research.

CORPORATE REGISTER

Cray Research receives orders for new scalable CRAY T3E system

The French National Supercomputing Facility (IDRIS) for the National Center of Scientific Research (CNRS) has placed an order for a CRAY T3D and a CRAY T3E system. The systems will serve the high-performance computing needs of the French scientific community and are targeted at grand challenge problems in environmental research, chemistry, and physics.

The **Pittsburgh Supercomputing Center (PSC)** will receive the first high-end version of the CRAY T3E scalable parallel computer system. Due to start shipping in early 1996, the new systems are compatible with the current CRAY T3D series, which in 1994 made Cray Research the revenue leader in the technical market for enterprise-level scalable parallel systems. The Pittsburgh Supercomputing Center is upgrading from a CRAY T3D system. The CRAY T3E system will be faster and have more memory than the CRAY T3D system, said Ralph Roskies and Michael Levine, PSC scientific directors, who expect that performance on applications running at PSC will improve three- to fourfold. This will benefit the scientific community nationwide, they said, by providing a powerful new tool that will make a difference on important problems, including turbulence, cosmology, drug design, weather forecasting and global climate change.

Japanese telecommunications giant **Nippon Telegraph and Telephone Corporation** (NTT) has installed a top-of-the-line, 32-processor CRAY T90 supercomputer at the company's Musashino R&D center in Musashino City, Japan. The system will be used for new materials development and advanced multimedia research, including image processing technologies. The system demonstrated sustained performance of 30 gigaflops for this procurement, which was based on an open solicitation published by NTT.

Providence, Rhode Island-based **Brown University** has installed the first of two Cray Research supercomputers to be dedicated to the university's Department of Physics. Brown's is the first university physics department in the United States to acquire a Cray system: a large-memory (four gigabyte) CRAY EL98 system, to be followed by a CRAY J916 system. These systems will be used to support research in high-energy and condensed matter theoretical physics.

Cray Research has received an order for a CRAY T94 supercomputer from the Tokyo-based **Japan Atomic Energy Research Institute Center for the Promotion of Computational Science and Engineering**. The order was officially received from Fujitsu Ltd., which will act as the systems integrator for this site, combining the Cray Research system with equipment from other computer vendors.

Akzo Nobel has ordered its second Cray Research supercomputer system. A CRAY J90 supercomputer system will be installed in 1995 at Akzo Nobel's Central Research Facilities in Arhem, The Netherlands. Akzo Nobel's computational chemists, among others, will use the system with the UniChem client/server computational chemistry software package for molecular modeling applications related to new Akzo Nobel products. A CRAY EL system also installed at Akzo Nobel will be used as a file server, providing storage facilities for Akzo Nobel's range of workstations. Cray's Network Queuing Environment (NQE) software will also network the two supercomputers to existing workstations.

3M installed a CRAY J916 compact supercomputer at its Electronic Products Division in Austin, Texas. The supercomputer is used for applications related to electronics packaging.

The **University of Los Andes**, Santafé de Bogotá, Colombia, installed a CRAY J916 supercomputer at Colombia's new national supercomputing center, the MOX Center for Advanced Computing in Engineering, marking the first Cray supercomputer in that country. University officials said the new center will help solve Colombia's computational problems in the areas of advanced product design and manufacturing; complex phenomena such as earthquakes and weather prediction; and financial calculations under uncertain scenarios. Data warehousing will also be a target application at the new MOX center.

MascoTech Engineering installed a CRAY J916 compact supercomputer system at its United Kingdom facility. The CRAY J90 supercomputer will be used by MascoTech engineers for a wide range of computer-aided engineering and design work using such popular commercially available software as CSA/NASTRAN, PAM-CRASH, OASYS-DYNA, RADIOSS, and ABAQUS. According to MascoTech officials, the CRAY J90 supercomputer was chosen over competing solutions because of its leading performance on the customer's benchmark, which included strenuous tests of the system's memory and input/output (I/O) bandwidth; the system's ability to support multiple concurrent users operating a wide range of application software; Cray's expertise in the customer's application set and in the worldwide automotive industry; and the company's long-standing relationship with the commercial software vendors for these applications.

The U.S. Department of Energy (DoE) ordered two CRAY J916 supercomputers for installation at AlliedSignal Kansas City Division. AlliedSignal Kansas City Division operates a defense manufacturing plant for the DoE. One CRAY J90 system will be dedicated to finite element applications using the popular ABAQUS software package from Hibbitt, Karlsson & Sorensen, Inc. The second system will serve as backup for these applications and will support other plant applications. Both systems will be equip-

Singapore establishes itself as the premier high-performance computing facility in South East Asia

Singapore's National Supercomputing Research Centre installs top-of-the-line supercomputer

With the installation of a top-of-the-line CRAY T94 supercomputer at Singapore's National Supercomputing Research Centre (NSRC) in December 1995, Singapore has become the premier high-performance computing facility in South East Asia. Two other Cray supercomputers, a CRAY J916 system at the Computer Centre of the National University of Singapore (NUS) and a CRAY EL98 system at the NUS Centre for Computational Mechanics, were installed in June and November of 1995, respectively.

NSRC, which is funded by the National Science and Technology Board and is affiliated with NUS, will use the high-end CRAY T94 supercomputer as a general-purpose, national resource to support industry, academia, research organizations, and government agencies. In industry, NSRC will advance the use of applications in chemical process engineering, fluid dynamics, structural analysis, physics, chemistry, electronics, and precision engineering.

According to Dr. William Hale, director of the National Supercomputing Research Centre, "The investment of this new and powerful system will help us fuel industrial development, encourage further R&D work, and drive the use of high-performance computing for a wide variety of industrial and scientific endeavors in Singapore."

Peter Boek, general manager of Cray Research in Australia, South East Asia and India, said "Cray appreciates that the Singapore Government—through the National Science and Technology Board, National University of Singapore, and NSRC—placed faith in Cray's technology for this important new phase in supercomputing in Singapore. With the CRAY T94, Singapore has established itself as the premier high-performance computing facility in South East Asia."

Pioneering high-performance computing in Singapore

The National University of Singapore (NUS) placed the country's first-ever order for a Cray supercomputer in November 1994. The CRAY J916 system, installed at the NUS Computer Centre in June 1995, has since been upgraded from 8 to 12 CPUs to meet user demands. The compact air-cooled supercomputer is used to train undergraduate and graduate students and researchers in supercomputing simulation technologies.

"The CRAY J916 supercomputer and its wide range of application software will open up new research and development opportunities, encourage new users to come on board, and lift our research capability to new heights," said NUS Vice-Chancellor Lim Pin.

During the National University of Singapore's 90th Anniversary celebration in September 1995, the University's Centre for Computational Mechanics (CCM) received a CRAY EL98 system—Singapore's third Cray supercomputer. Through this entry-level supercomputer local industries have access to the academic and technical expertise of CCM and to affordable computational power for industrial research and development. The research, conducted under the Computational Mechanics Industrial Programme (CMIP), recently established by CCM with a grant of $2.89 million from the National Science and Technology Board, targets areas of computational mechanics relevant to local industries. These areas include the application of composite-technology; the modeling of crystal growth; and the performance of welded structures for aerospace applications. CCM is also promoting industry use of high-performance computing in the areas of finite element analysis, computational fluid dynamics, fluid flows and advanced materials technology.

The CRAYQuest Awards: Recognizing the role of high-performance computing in R & D excellence

Cray Research and the National Supercomputing Research Centre (NSRC) in Singapore will jointly organize and launch CRAYQuest Singapore '96, the inaugural competition of the CRAYQuest Awards. The CRAYQuest Awards are an initiative of Cray Research to encourage research and development projects at NSRC that will advance the application of supercomputing in industry throughout Singapore.

CRAYQuest Singapore '96, Singapore's premier high-performance computing contest, is an eight-month contest beginning in March 1996 and ending in October 1996. It is open to all researchers in Singapore and will be an annual event. Winning projects are those that best demonstrate the use of NSRC's CRAY T94 system in solving problems in computational science or engineering with actual or potential application to industrial problems. For more information on CRAYQuest, send e-mail to crayquest@cray.com.sg

ped with Cray's Data Migration Facility (DMF), a high-performance hierarchical storage management software package. According to AlliedSignal officials, DMF is used to manage more than one terabyte (trillion byte) of engineering data associated with ABAQUS simulation problems running on the current Cray system.

Cray Research has delivered a CRAY J916 low-cost compact supercomputer and Cray's UniChem client/server computational chemistry software package to **Solvay Germany GmbH**. Solvay is a diversified chemical company involved in the manufacturing of alkali products, peroxides, plastics, catalysts, and pharmaceuticals. It is the first European chemical customer to order a Cray supercomputer for computational chemistry.

Macklanburg-Duncan, a leading home improvement products manufacturer headquartered in Oklahoma City, installed two CRAY CS6400 enterprise database servers for commercial client/server applications. Macklanburg-Duncan installed the CS6400 database servers for production computing, software development, and disaster recovery to support the consolidation of its Oracle databases. The consolidation is designed to improve data access for company employees, according to Michael Mack, manager of IT technical services at the company. The new systems will replace Sequent and Honeywell Bull equipment and will house and manage the company's financial, manufacturing, order processing and manufacturing equipment maintenance databases. In addition to database consolidation, the CS6400 systems will support Macklanburg-Duncan's plans to advance its automated ("paperless") order processing, as well as the effectiveness of its manufacturing database.

MSC/NASTRAN version 68.2 provides improved performance on Cray Research supercomputers

The MacNeal-Schwendler Corporation has released Version 68.2 of the MSC/NASTRAN structural analysis package. This new version performs better for certain types of jobs typically run on Cray systems. In particular, the CPU time and disk space required by the Lanczos normal modes extraction method has been substantially reduced for large models with many modes. MSC/NASTRAN version 68.2 is available for CRAY J90 systems, in addition to CRAY Y-MP and CRAY C90 systems. The CRAY J90 version of MSC/NASTRAN runs up to 30 percent faster than the CRAY Y-MP version when run on a CRAY J90 system.

Each of these Cray versions of MSC/NASTRAN has additional parallel processing capabilities that are available only on Cray systems.

The figure below shows the disk space, I/O, and time savings for an average sized car body analysis performed on a nondedicated CRAY C90 system.

In addition to performance improvements, MSC/NASTRAN Version 68.2 includes other user-requested capabilities:

- An improved shell element formulation automatically results in better and more accurate answers, due to its unique shell-normals calculations. FEA users spend a great amount of time in modeling and tuning their models for accurate results. This effort is significantly reduced in version 68.2.
- A string-based DMAP ALTER capability enables users to write DMAP ALTERs independent of line numbers and subDMAP names. This capability is particularly helpful when a user upgrades to a new version of MSC/NASTRAN.
- An inertia relief sensitivity and optimization capability was supported as a DMAP ALTER in Version 68. It is part of the standard optimization solution sequence in Version 68.2.
- Design optimization improvements have been made. The optimizer and several of its defaults have been modified for better and faster results.
- More than 170 errors have been corrected in Version 68.2.

For a comprehensive list of Version 68.2 enhancements, please contact your representative at the MacNeal-Schwendler Corporation; or Doug Petesch, Cray Research, Inc., 655E Lone Oak Drive, Eagan, MN 55121; telephone: 612/683-3654; email: djp@cray.com.

MSC/NASTRAN Performance Improvements on CRAY C90 Systems

Version	Disk Space (GB)	I/O (GB)	1 Processor CPU	1 Processor Elapsed	3 Processors CPU	3 Processors Elapsed
68.1	10.8	280	5370	7110	5390	6620
68.2	8.5	240	4290	5900	4410	4740

(70,000 shell elements, 384,000 DOF, 777 modes)

A new parallel version of *Gaussian 94* is now available for Cray systems

Gaussian, a popular integrated system of programs for carrying out ab initio, density functional, and semiempirical calculations, is now available in a parallel version for Cray Research parallel-vector supercomputers. *Gaussian* is currently installed at over one hundred of Cray Research's customer sites and is one of the top three applications at approximately 45 of these sites. Cray Research has been working on this new version in collaboration with Gaussian, Inc. since February 1995.

Ultimately, as a result of the combined efforts of both corporations, the end user will benefit from the greatly increased performance of the newly released parallel version. Users doing large Self Consistent Field (SCF) energy and gradient calculations will typically see a speedup of five times on eight processors, which will greatly reduce the time it takes to optimize molecular structures. Users doing large MP4 energy calculations will see speedups ranging as high as seven times on eight processors making highly accurate G2 calculations tractable for larger molecules.

To facilitate the development of the code, Cray Research placed a fulltime computational quantum chemist and a CRAY EL92 system on-site at Lorentzian, Inc., in North Haven, Connecticut, the primary development organization for *Gaussian*. The shared location benefitted both corporations. It allowed Cray Research's chemist to work side-by-side with the original developers of the code making the optimization process much more efficient. In addition, it provided Lorentzian with convenient access to a Cray Research supercomputer and applications expertise, which helped lower the cost of maintaining the Cray Research version of *Gaussian*.

The shared location not only provided an environment for close collaboration on the development of the new parallel version but also created a natural opportunity for Cray Research to aid Gaussian in new code development. Although specific details about future plans have not been released, work is being done to enhance the performance and functionality of the pseudopotential parts of the program. This will enable future versions of the program to more efficiently carry out calculations involving elements in the lower part of the periodic table such as platinum or zirconium which are metals that often appear in catalytic systems.

Michael Booth, vice president of Cray Research's applications division said, "The decision to place a Cray Research scientist on-site proved to be advantageous for all parties involved. The newly released version of *Gaussian* is going to enable our customers to complete *Gaussian* calculations in a matter of hours instead of days and will make the CRAY J90 a very effective machine for running *Gaussian*."

For more information on the latest version of *Gaussian*, email info@gaussian.com, or contact Carlos Sosa, Cray Research, Inc., 655E Lone Oak Drive, Eagan, MN 55121; telephone: 612 683-3602; email: cpsosa@appsdiv.cray.com

New version of ANSYS available

ANSYS Inc. has released version 5.2 of the ANSYS structural analysis package. This new release promises expanded parallel processing capabilities and an improved Cray Research version of the ANSYS PowerSolver.

For the past five years, ANSYS has been available on Cray Research vector SMP systems using a highly optimized parallel-vector implementation of the standard frontal solvers. Large models run in the past have achieved a sustained performance rate as high as 6 GFLOPS on a 16 CPU CRAY C90 system. These high sustained elapsed-time rates were achieved using parallel processing and a special I/O library that incorporates asynchronous I/O and an intelligent cache to reduce I/O wait time.

Now, with new changes in version 5.2 of ANSYS, parallel processing will be expanded to perform thousands of computations at the individual element level on multiple processors. This increased use of parallel processing in ANSYS version 5.2 will greatly reduce time to solution for more sophisticated nonlinear analyses. For these analyses, the individual element computations are often more expensive than the time required to solve the global system of equations.

In addition, new solver technology introduced in ANSYS version 5.1 has now been optimized for Cray Research systems in version 5.2. This new solver, known as the ANSYS PowerSolver, is a state-of-the-art iterative equation solver licensed from Computational Applications and System Integration, Inc. The new solver technology has greatly reduced the computational requirements of the equation solution phase of ANSYS analyses compared to the direct frontal solvers. Although the performance of this solver varies from problem to problem, the new optimized version of the PowerSolver in ANSYS version 5.2 is substantially faster than the previous Cray implementation used in version 5.1.

The new parallel processing capabilities in ANSYS version 5.2 and the optimization of the Powersolver are the results of an ongoing relationship between Cray Research's Applications Division, ANSYS, Inc., and Computational Applications and System Integration, Inc. The chart shows sustained improvements in performance on Cray Research supercomputers with each new release of ANSYS (including version 5.2) for a representative ANSYS analysis.

In collaboration with ANSYS, Cray Research is experimenting with other direct solver options which may be made available in the near future in special release versions of ANSYS version 5.2. In the meantime, Cray Research will continue to work with ANSYS to improve code performance. For more information about ANSYS, please contact Gene Poole, Cray Research, Inc., 655E Lone Oak Drive, Eagan, MN 55121; telephone: 612/683-3684; email: gene.poole@cray.com.

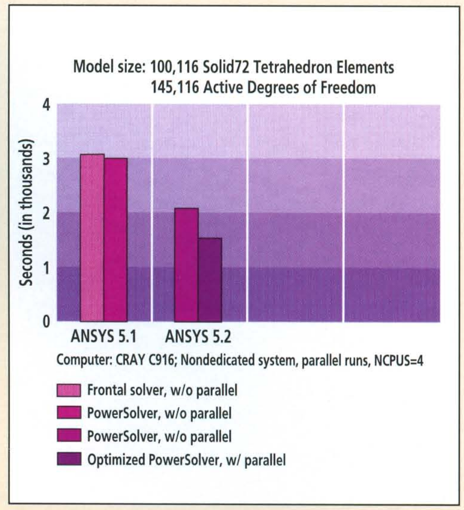

Kumho purchases Cray applications consulting services

Kumho, a Korean-based conglomerate, used Cray Research's applications consulting services to coordinate several projects: a civil engineering analysis, a tire analysis, and a chemical analysis for newly developed materials. Projects were determined by mutual agreement and carried out by Kumho and Cray Research technical staff with the assistance of AC Engineering from West Layfayette, Indiana. Key engineers from Kumho received training on the company's CRAY C90 system as well as the targeted applications needed to complete the analyses. For the Kumho Construction Division, ABAQUS/Explicit and ABAQUS/Standard from Hibbitt, Karlsson and Sorensen, Inc. were employed to simulate vibration dampers on bridges to prevent earthquake damage. In addition, ABAQUS/Standard was used to define tread patterns in tires for the Kumho Tire Division. For the Kumho Chemical Division, DMol and DISCOVER, chemical codes from BIOSYM Technologies, Inc. were used to gain a better understanding of the function of the enzyme DHFR at the atomic level. The CRAY C90 is installed at the Kumho Computer Center in Kwangju.

Raytheon redefines light jet design

Raytheon Aircraft has developed a breakthrough plane: The Raytheon Premier I, the first all-new light jet in three decades. The Premier I is bigger, faster, and more affordable than any light jet flying today. Its wide passenger cabin, swept-wing design, and Mach cruise speed resulted from leveraging not only Raytheon's extensive knowledge of aerodynamics and construction techniques, but also the high performance of Cray Research supercomputers.

Using a CRAY C90 system at the NASA Ames Research Center, as well as their own new CRAY J916 system to design the Premier I, Raytheon engineers designed a light jet that provides performance, comfort, and value to the customer. The Raytheon Premier I is in a class by itself, redefining the light jet arena. It is expected to take the market by storm.

Raytheon chose to design this aircraft after extensive market research identified comfort, performance, and value as the features most desired by customers. The composite fuselage allows for a luxuriously appointed cabin sized seven inches higher and eight inches wider than competitive aircraft. The jet is powered by two next-generation Williams-Rolls FJ44-2A fan jet engines, each with 2300 pounds of takeoff thrust.

Raytheon used Integrated Product Teams (IPT), made up of representatives from each major discipline, to develop the Premier I. The IPT approach ensures that all of the necessary people are involved early on in the project to minimize later rework. Assisted by Computer Aided Three-dimensional Interactive Analysis (CATIA) software, Raytheon Aircraft's team simulated assembly sequencing and created electronic aircraft mock-ups, which allowed them to analyze structures and optimize system design. This procedure reduced part counts, product flow times, and operating expenses.

The aerodynamic design of an entirely new light jet wing was carried out on an aggressive schedule through a cooperative agreement with NASA Ames Research Center. A combination of 3-D transonic computational tools was used to do the actual design work and validate the results.

Neal Pfeiffer, Senior Technical Specialist at Raytheon Aircraft Company, explains that this design method was an order of magnitude faster than anything else available: "The design iteration took less than 8 hours of computer time. A total of 18 computational wing models were made during a three-month period to arrive at the desired wing design. This three-month design time is less than half of that required for similar, previous NASA design efforts." The swept-wing design contributes to the high cruise speed of 530 mph.

The TRANAIR code, which was developed by Boeing for NASA, was used because of its ability to analyze complicated configurations. This was especially important for this project because of the effects of the engine nacelles and pylons on the distribution of lift across the span of the wing. Unfortunately, the NASA TRANAIR code lacks the ability to do design work directly.

To reduce the wing's design time, a new, very efficient design code, SYN87, was used. SYN87 was developed jointly by James Reuther, a researcher at NASA Ames, and Antony Jameson at Princeton University. The code is highly vectorized and even without special parallel programming effort, shows significant speed-up on parallel computers. Further development is underway to improve the parallel efficiency of the code.

There was only one problem with using the SYN87 code; it was written to handle only simple wing and fuselage configurations. This limitation was overcome by using SYN87 to design a bump on the aft fuselage that simulated the blockage effect on the wing that was apparent in a TRANAIR simulation with the nacelles and pylons present on a baseline wing.

A single, wing-design cycle at the primary flight condition can be made in less than eight hours on either the NASA Ames CRAY C90 or the Raytheon CRAY J916 system. Additional analysis runs for other flight conditions can be made in less than 30 minutes on the Raytheon system with SYN87. Additional TRANAIR solutions, with viscous modeling, can be made in two to eight hours for medium to very fine grid resolutions on the CRAY J916 system.

"We are very happy with our Cray system, and we are exceedingly happy with the level of support that we receive. The support is as important as the system itself," Pfeiffer said.

The first flight of the Raytheon Premier I will take off in the summer of 1997, with deliveries planned for the fall of 1998. It will be certified for single-pilot operation.

Paramics software helps motorists go with the flow

Poorly managed traffic takes a toll not just on the unfortunate commuters whose time and money go up in car exhaust, but also on the environment and, indirectly, on all motorists. Road congestion squanders fossil fuels and concentrates toxic emissions. It contributes to more frequent accidents—fueling insurance rates—and accelerates road deterioration—fueling tax rates. As traffic piles up in urban centers around the world, city planners and public departments of transportation are reaching an inescapable conclusion: There must be a better way.

A better way to orchestrate traffic flows is exactly what the software devel-

Visual interface to the Paramics traffic modeling program. In combination with highly parallel computing systems, the program can give traffic managers and city planners valuable insight into traffic flow patterns and routing options.

opers at Quadstone Ltd. had in mind when they designed the Paramics traffic simulation package. Located in Edinburgh, Scotland, Quadstone is demonstrating the benefits of modeling traffic flow computationally as an aid to traffic management and urban planning. Using microscopic modeling techniques (in which vehicles are represented individually) rather than traditional macroscopic methods (in which vehicles are modeled collectively as a continuous fluid), the Quadstone developers have demonstrated the value of their approach convincingly enough to receive development contracts from the British government. The computing power needed to make microscopic modeling practical for large, complex traffic flows is provided by the CRAY T3D system installed at the Edinburgh Parallel Computing Centre and the CRAY EL system installed at Quadstone.

Among the projects commissioned by the United Kingdom Department of Transport is a study of traffic on the London ring highway, the M25, now being widened. Quadstone analysts are using the Paramics package to research the practicality of using programmable speed-limit signs to help manage traffic. Another project involves predicting the effects on traffic of proposed new road layouts around Edinburgh. Quadstone is also using Paramics to anticipate the effects of improved public transport schemes and extended accommodations for pedestrians in other Scottish cities. Using the Paramics package, planners hope to identify likely bottlenecks in traffic flow and to run experiments on the computer to assess alternative routing schemes.

> *As traffic piles up in urban centers around the world, city planners and public departments of transportation are reaching an inescapable conclusion: There must be a better way.*

Detailed traffic modeling has applications not just for future roadways, but for existing ones as well. A current project at Quadstone involves studying the feasibility of reserved "bus only" lanes in Edinburgh and the effects they would have on overall traffic flow. In addition, the spread of monitoring technology, such as pavement sensors, online video, and even satellite detectors, creates an opportunity for real-time analysis of traffic flow, which could predict congestion on a day-to-day, or even hour-to-hour, basis. Predictions made early enough could be used to alleviate or prevent congestion through signal timings or other traffic control mechanisms.

"Using the highly parallel CRAY T3D supercomputer, we can model complex traffic flows considerably faster than realtime," explained Mark Smith, Quadstone marketing director. "Twelve times realtime is a realistic possibility, which would give us the likely congestion an hour from now in just five minutes, given current ground positions from remote sensors. Alternative routes then could be suggested to motorists via programmable message signs along the highway." The microscopic methodology that makes these applications practical gets its power from its simplicity and accuracy. More precise and more flexible than macroscopic methods, microscopic methods are much better able to scale up to model very high traffic volumes and large numbers of intersections (nodes or links). Traditional methods also lack the precision needed to model some features common to complex traffic flows. "Backup waves," for example, cannot be anticipated by macroscopic methods. These are a common occurrence in which large numbers of vehicles slow down simultaneously and the traffic flow breaks down, or becomes nonlinear.

"Microscopic methods were held back by the lack of computing power," said Smith. "The CRAY T3D not only lets us tackle more complex problems, but Cray's standard software environment lets us do development work on workstations and move the code directly to the supercomputer for computationally intensive production runs. At this point we can model 120,000 vehicles on 30,000 kilometers of road at three times realtime rates on 32 nodes of the CRAY T3D. The Paramics-supercomputer combination opens many possibilities for addressing practical traffic planning and management problems."

The precision made possible with microscopic modeling and parallel supercomputers also has the potential to help inform environmental policy. By identifying areas prone to heavy traffic idling or a high proportion of stop-and-start activity, microscopic modeling can be used to anticipate areas of high pollution. This information then can be used for both transportation planning and environmental policy evaluations.

Interest in large-scale microscopic traffic modeling continues to grow. This growth is fueled by simple economics. For example, the traffic-management consultancy SIAS Ltd., also in Edinburgh, estimates that a reduction by 20 percent in congestion delays through the Scottish trunk road network would produce about a £250 million ($400 million) annual savings for the Scottish economy. Moreover, now that parallel supercomputers have made large-scale microscopic modeling practical, the methods can be applied to decision support in many areas beyond traffic modeling. The methods could be used to model the demands placed by individuals on a network of service or business outlets. They also may have applications in predicting the performance of a pension or savings fund based on individual contributors' investments and withdrawals.

(For more information on the Paramics software package, contact Mark Smith, Marketing Director, Quadstone Ltd, 16 Chester Street, Edinburgh, EH3 7RA, UK; tel: +44 131 220 4491; email: mark@quadstone.co.uk)

Quantum molecular dynamics on the CRAY T3D system at École Polytechnique Fédérale de Lausanne.

Until recently, companies and researchers using computers to design new chemicals and materials were unable to perform molecular dynamics with the full accuracy of quantum mechanical theory on large systems (a few hundred atoms). Methods developed originally by Professor Roberto Car (Geneva University and the Swiss Federal Institute of Technology) and Professor Michele Parrinello (Max Planck Institute) in conjunction with modern supercomputers such as the CRAY T3D system have made these kinds of simulations possible.

Car and Parrinello developed the CAR-PARRINELLO algorithm, which is now the established technique for performing quantum molecular dynamics simulations and is included in several commercial application codes. CAR-PARRINELLO is important to industrial and research organizations studying industrial processes, chemistry, solid state physics, and advanced materials science because these research areas typically require accurate modeling of systems containing hundreds of atoms. These types of simulations, while extremely parallel, require global communications due to the extended nature of the functions describing the electrons. The CRAY T3D system, with its balance of a fast communication network and high-speed RISC microprocessors, is very well suited to these types of calculations.

"With new algorithms and powerful parallel computers like the CRAY T3D, it is possible to realistically study large systems that were previously inaccessible or studied with less exact classical models," said Car. "This opens the way for accurate analysis via computer simulations of new pharmaceuticals, chemical processes, and materials, reducing the amount of costly or difficult new laboratory experiments and complementing existing experimental knowledge."

Under Car's direction, representatives from a group of universities in Western Switzerland called IRRMA (Institut Romand de Recherche Numérique en Physique des Matériaux) are working to solve important—and highly parallel—problems in materials science and associated technologies and understand fundamental issues in physics and chemistry.

In collaboration with Cray Research, IRRMA has written two large quantum molecular dynamics codes for the CRAY T3D supercomputer: a first-principles molecular dynamics code using the CAR-PARRINELLO method, and a so-called O(N) tight-binding code which is based on a simplified quantum mechanical approach and incurs a much smaller computational cost per atom than does the standard CAR-PARRINELLO algorithm. The O(N) algorithm uses localized functions to represent the electrons, allowing quantum calculations to be applied to atomic systems a factor of ten or more larger than those that can be modeled with standard algorithms. Instead of being limited to a few hundred atoms, researchers now can model several thousand atoms, as well as chemical activity occurring over longer timescales, up to several hundreds of picoseconds. "Both of these codes run extremely efficiently on the CRAY T3D," said Car. "The reliability of the system and the familiar operating environment have made it easy to write and run these codes, which have now been in production for many thousands of hours on different physical problems."

The first-principles code has been used to study graphitization of a 288-carbon-atom system. The O(N) code has been used to study the deposition of fullerenes on a semiconducting substrate in a 4000-atom system. Important new scientific results have been obtained on the CRAY T3D supercomputer for these systems, both of which are extremely important from an industrial point of view for the production of synthetic diamonds and hard coatings

Medtronic accelerates product development with Cray technology

The quest to alleviate pain, extend life, and restore health drives product design and development at Medtronic, Inc., one of the world's leading biomedical device manufacturers.

Minnesota-based Medtronic supplies half of the world's pacemakers. The company also produces heart valves, angioplasty catheters, drug administration and pain management implants, along with an array of other products used in open heart surgery. Most of the products, such as pacemakers and drug pumps, are implanted directly into the body. Others, such as angioplasty catheters, blood pumps, and oxygenators, are used during invasive procedures or for life-sustaining functions during surgery.

Because biomedical implant devices are used in conjunction with high-risk procedures, achieving 100 percent reliability and optimal quality control are Medtronic's primary objectives in product design and manufacturing. "Physicians and their patients depend on our products. We at Medtronic seek to continually improve product capabilities and to make our designs as robust as possible," says Paul Citron, Medtronic vice president for Science and Technology.

Ensuring biocompatibility, reliability, and functionality is the core responsibility of the Metals and Structural Mechanics Research Group within Medtronic's Center for Biomaterials Research. The group's world-class scientists use computer-based simulation to optimize the designs of Medtronic's biomedical devices, applying interdisciplinary expertise in fluid dynamics, materials sciences, structural mechanics, electromagnetics, and thermodynamics. Key design challenges include

- ❑ Making certain that pacemaker leads are highly resistant to tension, torsion, bending, and compression. The leads, manufactured from extruded polymer sheaths, contain signal wires connecting the implanted pacemaker to the heart muscle and are in constant contact with body tissue, artery walls, and heart valves.
- ❑ Ensuring that oxygenators and blood pumps provide a smooth flow of oxygenated blood, with minimal shear and turbulence, while the devices temporarily replace the patient's lungs and heart during open heart surgery.
- ❑ Designing lithium batteries that provide optimal reliability over the full seven-to-ten-year implant life of pacemakers.

While several steps removed from the final-approved, marketable device, being able to provide numerical simulations that reliably predict product or subcomponent design performance is crucial to Medtronic's ability to bring higher-reliability designs to the marketplace in a shorter period of time," explained Svenn E. Borgersen, the group's senior research scientist.

To help meet the challenge of designing products with continually enhanced capabilities, Medtronic's Center for Biomaterials Research installed a CRAY J916 supercomputer.

Running complex simulations on the CRAY J916 is significantly speeding up the development and optimization of new Medtronic biomedical devices. Take pacemaker leads, for example. "Simulation studies of leads performed on workstations often required days, weeks, or months to complete, and in some cases could not be solved due to memory limitations," Borgersen said. "These same complex simulations are now being solved on the CRAY J916 in a matter of hours."

Medtronic's goal here—optimizing the lead's ability to withstand the in vivo environment of the human body—is addressed using the CRAY J916 system and the MARC K6.2 finite element analysis (FEA) software product from MARC Analysis Research Corporation. The numerical simulation studies allow researchers to evaluate the mechanical and structural performance of the lead design based on tension, compression, bending, and torsion loads due to in vivo conditions. The computations are very complex. The lead assembly contains multiple components that interact with each other during the loading process. Because in vivo loads applied to the leads often exceed normal limits for the materials used, the numerical analysis becomes even more complicated, necessitating the use of nonlinear material properties.

Another key Medtronic use of supercomputer simulation is to optimize the design of blood oxygenator products that act as lung substitutes during open heart surgery. The primary challenges associated with designing blood oxygenators are optimizing the efficiency of the devices, and minimizing or eliminating

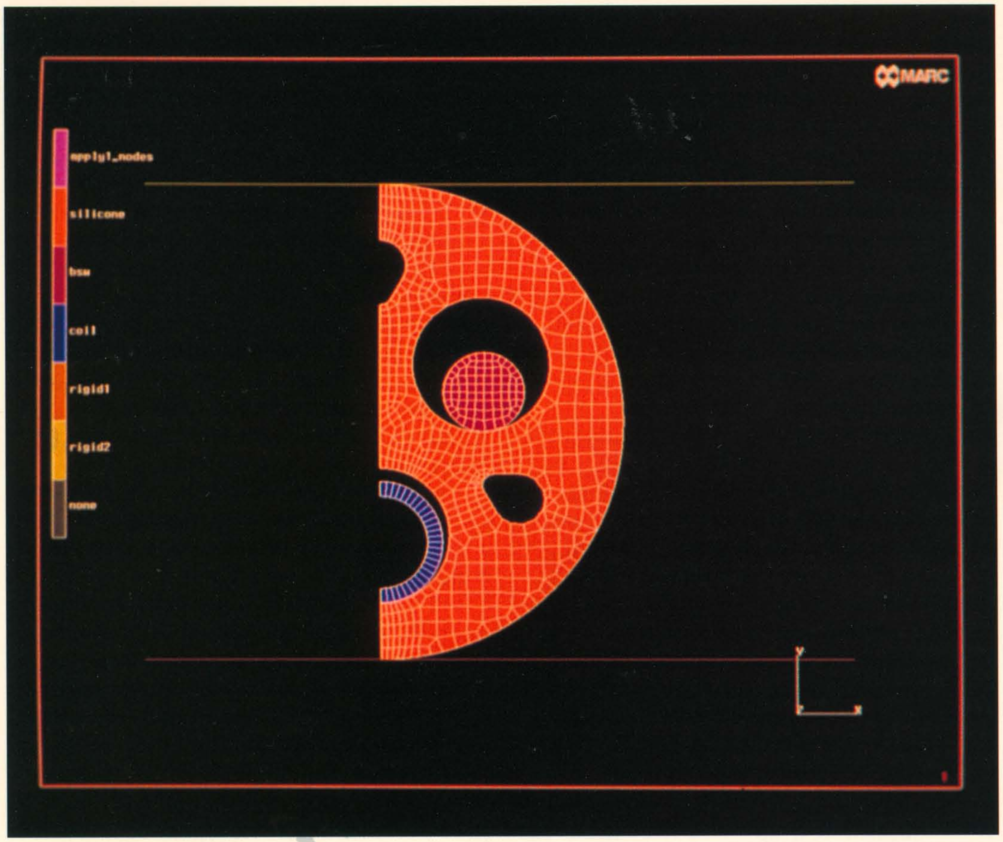

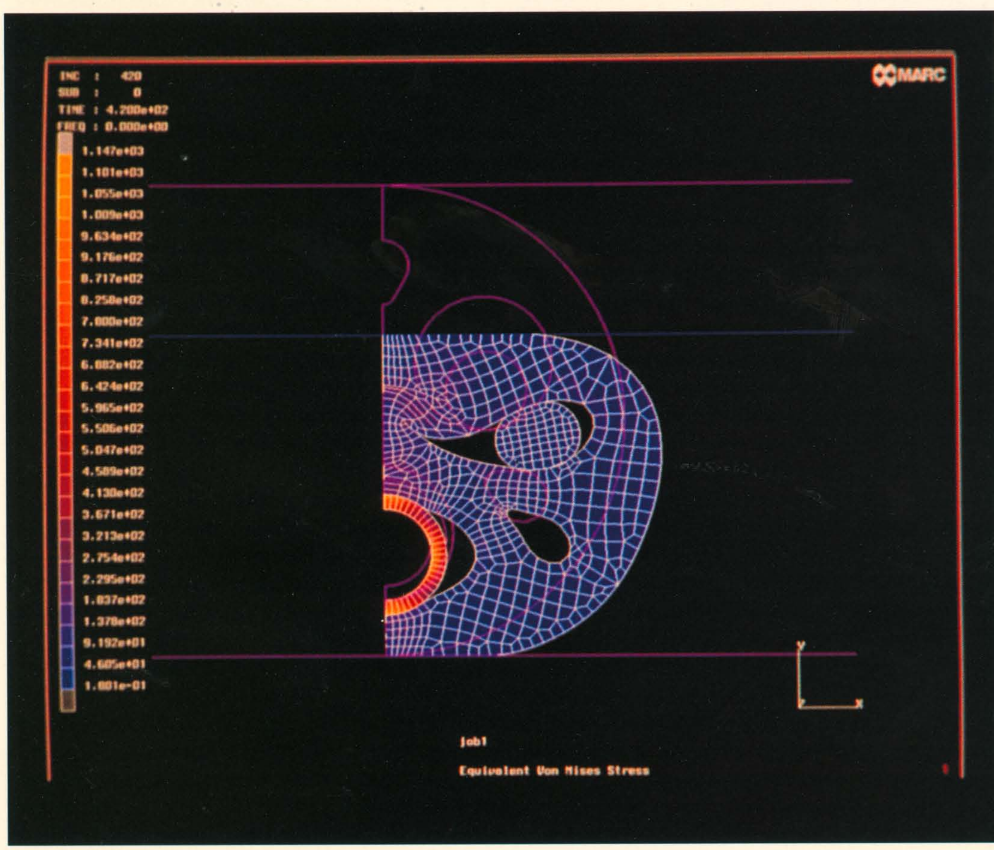

Medtronic biomedical engineers evaluate the mechanical and structural performance of pacemaker lead designs computationally by "subjecting" the leads to tension, compression, bending, and torsion loads on the CRAY J916 system using the MARC finite element software product from MARC Analysis Research Corporation. The images show the multi-lumen, extruded polymer lead body with braided stranded wire and helically coiled multifilar electrical leads within the lumens being compressed between two rigid surfaces.

damage to the blood (hemolysis). In addition, to prevent air from being trapped in the oxygenator during manufacturing, the blood flow path within the device is filled with a fluid that must be purged before surgery—minimizing this "priming volume" is another important design goal. Medtronic addresses these design challenges by running the popular FIDAP computational fluid dynamics software package on the CRAY J916 system.

Creating the detailed three-dimensional grid-like models ("meshes") needed for computer simulation of complex phenomena such as the oxygenator's blood flow path is typically a time-consuming manual task. By using Cray's HEXAR automeshing software, Medtronic is now able to create these models in approximately one-fifth of the previous time.

Medtronic also uses the CRAY J916 supercomputer system to design the optimal preform for the lithium battery used to power pacemakers. Highly moisture-sensitive and very reactive with other materials, lithium is difficult to handle in laboratory and high-volume production environments. It has a strong tendency to adhere to other materials, making production tooling design challenging. It also exhibits highly nonlinear mechanical properties, approximating a fully plastic behavior when subjected to tensile loads.

Because the mechanical properties of lithium are not fully described in the scientific literature, Medtronic's Metals and Structural Mechanics Research Group launched a pioneering materials testing program to determine the compression behavior of lithium preforms. This research effort was coupled with the concurrent supercomputer simulation of the material test specimen, using the MARC K6.2 software. Thanks to very fast turnaround times on the CRAY J916 system, Medtronic's researchers were able to develop a materials model that showed excellent correlation between the load-deflection results obtained in physical tests and those produced in the computer simulations. As a result, Medtronic successfully optimized the preform design for high-volume battery production.

Given the growing demand for extremely reliable biomedical devices, Citron foresees many new development opportunities for the company's sophisticated tools and processes: "The Cray solution is exactly what we needed to solve our complex computational problems associated with new product designs."

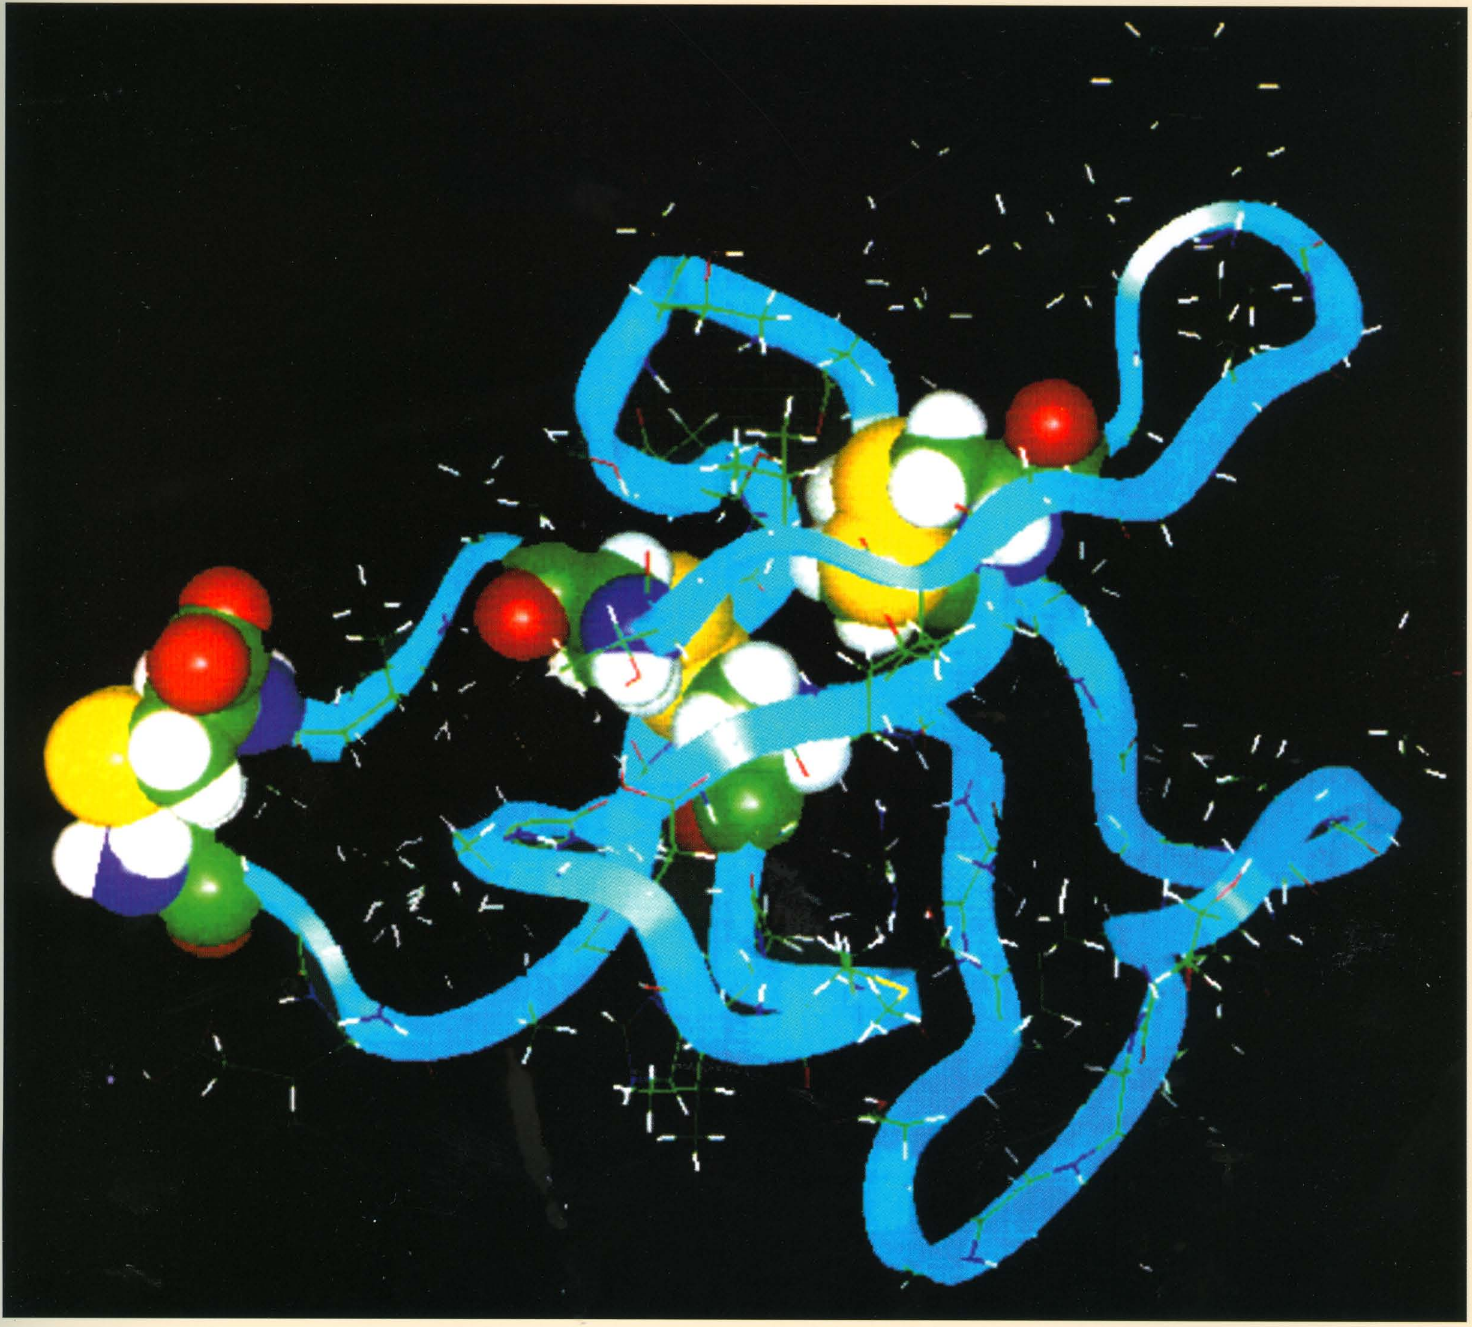

Three-dimensional conformation of kringle-like (K) domain of Dengue-2 virus envelope protein. K domains are found in the proteins associated with the human blood clotting system. The calculations were carried out on a Cray supercomputer at the Institute for Biomedical Research at the University of Mexico by researchers Blanca Ruiz and Gustavo Ortega.