A CRAY RESEARCH, INC. PUBLICATION

CRAY CHANNELS

Fall 1985

ANNOUNCEMENT!
CRAY X-MP Series expanded and enhanced

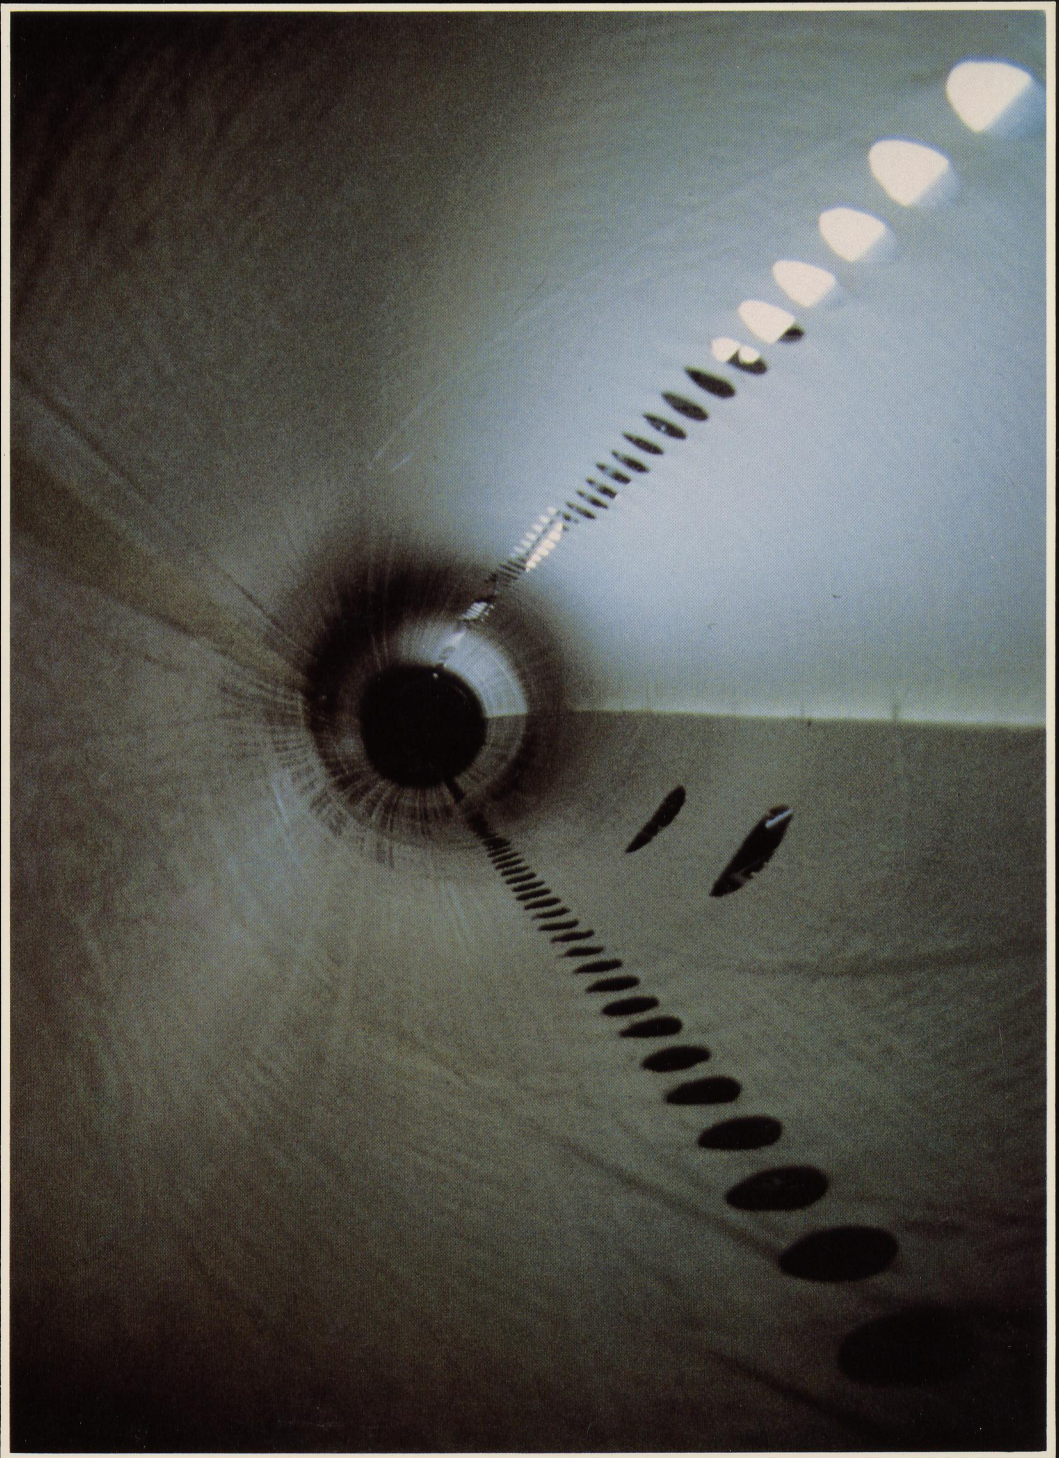

IN THIS ISSUE

Automotive manufacturing, once clearly a U.S.- dominated industry, has drifted into international waters in recent years. European and Japanese manufacturers now compete successfully in the United States and on their home grounds. Among the effects of this global competition is an accelerated drive by the industry to modernize. Increasingly, auto makers in the United States, Europe and Japan are exploring supercomputing as part of this effort. Applied to engineering and research problems, supercomputers can dramatically reduce the time and resources needed for design, testing and refinement.

In this issue of CRAY CHANNELS we see how automotive engineering and basic research have benefited from computational analyses and numerical simulation made possible by supercomputers. In addition, our regular departments report on the latest engineering software available for CRAY systems and on diverse applications from astrophysics to semiconductor manufacturing.

If you've given little thought to the science behind the car or truck you drive, we hope you'll appreciate the behind-the-scenes look we take here. The evolution of the automobile has been a carefully guided process — one that will become increasingly refined as engineers and researchers gain access to the best computational tools available.

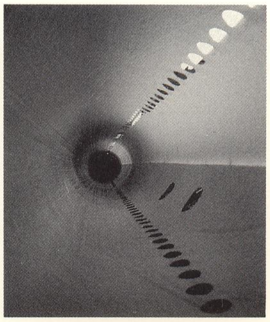

On the cover is an air diffuser used at Cray's silicon integrated circuit manufacturing facility. The 36-foot long diffuser is one of four that distribute fresh air evenly throughout the facility's enclosed airspace. The fresh air is brought in to replace air exhausted from the facility. Diffusing the fresh air evenly with the room air before circulating it minimizes temperature and humidity gradients and helps to regulate air pressure and flow in the clean room.

CRAY CHANNELS

A Cray Research, Inc. publication

Fall 1985

FEATURES

DEPARTMENTS

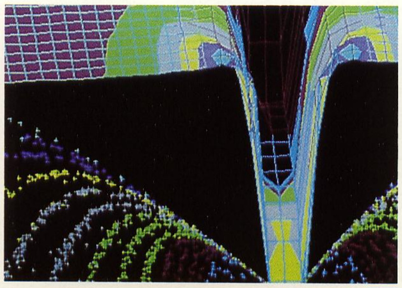

6

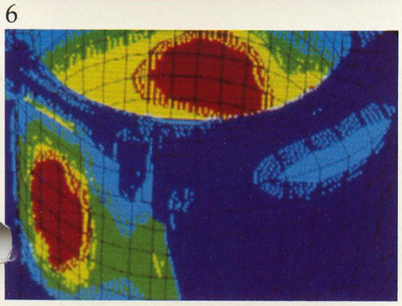

12

CRAY CHANNELS is a quarterly publication of Cray Research, Inc. intended for users of Cray computer systems and others interested in the company and its products. Please mail subscription requests, feature story ideas and news items to CRAY CHANNELS at Cray Research, Inc., 608 Second Avenue South, Minneapolis, MN 55402.

Volume 7, Number 3
Fall 1985

Editorial Staff
Tina Bonetti
Anne Hazelroth
Ken Jopp
Carol Smith
Jack Thompson

Design
Cindy Erickson
Eric Hanson

Expanding capabilities: the CRAY X-MP Series of Computer Systems

In the mid-1970s, marketing studies predicted that the future of supercomputing would be limited by the relatively small demand for such systems. In the decade since that forecast, however, the marketplace for supercomputers has exploded, and with this broad expansion has come the demand for ever-faster, ever-larger systems. In keeping with its tradition of leadership, Cray Research is once again meeting the growing needs of its customers by introducing further enhancements to the CRAY X-MP Series of Computer Systems.

Based on the field-proven CRAY X-MP central processing unit (CPU), the new X-MP family incorporates nine models with one, two or four CPUs. Central memory has been expanded across the line, with up to 8 million words available on uniprocessor systems and up to 16 million words available on the two- and four-processor systems. Options for field upgrade of memory are available on all models.

Gather/scatter and compressed index hardware has been extended from the CRAY X-MP/48 to all models in the series, providing significant performance improvements for a variety of applications. Gather/scatter allows the vectorized loading and storing of randomly organized data anywhere in memory. With compressed index, gather/scatter permits the vectorization of previously scalar-bound code, such as loops which contain conditional instructions.

In addition to the new mainframes, Cray Research has also announced a revision to the Solid-state Storage Device (SSD). A new model, featuring 512 megabytes (Mbytes) of very fast secondary MOS memory, has been added to the 256- and 1024-Mbyte SSDs. Due to technological changes, the smaller, 64- and 128-Mbyte models have been discontinued. When connected to a four-processor CRAY X-MP through two very-high-speed channels, the SSD provides a maximum aggregate transfer rate of 2000 Mbyte/sec. Outstanding I/O performance is also provided by state-of-the-art DD-39 and DD-49 disk drives.

Les Davis, Executive Vice President for Engineering and Manufacturing, commented, "The new large memories, when coupled with the power of an SSD and our new I/O devices, offer users much greater problem solving capabilities. Applications in aerospace, petroleum and weather modeling will be among the early beneficiaries, and other application areas will see marked improvements as well."

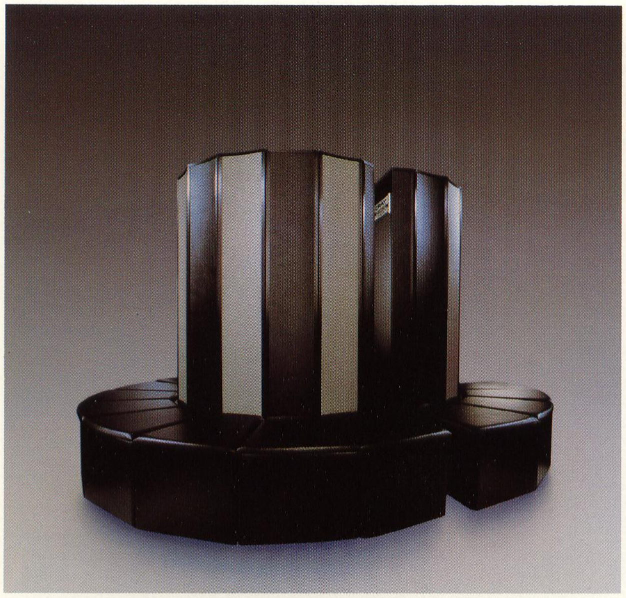

The new memory options of the CRAY X-MP family, coupled with multiprocessor architectures, provide improved performance which goes beyond the ability to handle large single jobs. In multitasking and multiprogramming applications, more and bigger jobs are able to remain in memory at once. This minimizes the need for roll-out to I/O devices, thus improving throughput for each task and user.

The introduction of the new CRAY X-MP models is consistent with the company's philosophy of providing the newest technology for its customers as soon as it becomes available. As Steve Chen, Vice President of Development explained, "This announcement is another example of Cray's commitment to its customers. It demonstrates the company's unique ability to keep customers at the forefront of technology."

CRAY X-MP/4 models

The top-of-the-line CRAY X-MP/4 systems offer an order of magnitude greater performance than the original CRAY-1. They are configured with 8 or 16 million 64-bit words of fast ECL bipolar memory, providing a maximum memory bandwidth 16 times that of the CRAY-1. Central memory has a bank cycle time of 38 nanoseconds (nsec) and is shared by four identical CPUs. All four CPUs and all 16 million words of memory may be used by a single program. The mainframe is arranged in the familiar 12-column, 270° arc.

CRAY X-MP/2 models

The proven CRAY X-MP/2 models have become the established price and performance leaders in the supercomputer industry. The new CRAY X-MP dual-processor systems offer up to four times the memory and require only half the electrical power of the original CRAY X-MP/2 systems.

The CRAY X-MP/2 models are arranged in an 8-column, 180° arc and are available with 4, 8 or 16 million words of static MOS central memory, providing a maximum memory bandwidth four times that of the CRAY-1. The additional memory means that even larger problems can be solved with the CRAY X-MP/2 systems.

CRAY X-MP/1 models

The CRAY X-MP/1 systems combine a single X-MP CPU with 1, 2, 4 or 8 million words of static MOS memory and provide memory bandwidth four times that of the CRAY-1. The enhanced CRAY X-MP/1 models offer up to twice the memory of the old CRAY X-MP/1 along with a new 1000-Mbyte/sec channel connection to the SSD. They are arranged in a 6-column, 135° arc.

I/O Subsystem

The I/O Subsystem (IOS) is an integral part of all CRAY X-MP computers and contributes to the X-MP models' outstanding performance. The IOS acts as a data concentrator and data distribution point for the CRAY X-MP mainframe. It offers parallel disk drive capabilities, I/O buffering for disk-resident datasets, high performance on-line tape handling and efficient front-end system communication. Up to 8 million words of buffer memory can be configured on the IOS, enabling faster and more efficient data access.

Solid-state Storage Device

The optional Solid-state Storage Device (SSD) is a very fast random-access device especially designed to complement the CRAY X-MP. The SSD, in conjunction with multiprocessor architectures, allows the development of algorithms to solve larger and more sophisticated science and engineering problems. New developments in MOS technology have allowed significant price reductions across the SSD line, thus increasing the cost-effectiveness of adding an SSD to the CRAY X-MP.

System performance is significantly enhanced by the SSD's exceptionally high transfer rates and short data access times. The SSD may be configured with 256, 512 or 1024 Mbytes of rapid-access MOS memory. Two SSD channels on X-MP/4 systems and one SSD channel on X-MP/2 and X-MP/1 systems provide transfer rates of 1000 Mbyte/sec per channel. Access times of less than 25 microseconds are achievable between the SSD and an X-MP mainframe. The SSD offers marked improvement on I/O-bound applications, and thus allows users to develop new algorithms that would not have been practical with traditional disk I/O.

DD-39 and DD-49 disk drives

Complementing and balancing the CRAY X-MP computing speeds are the DD-39 and DD-49 disk drives. Both are high-density (1200-Mbyte) storage devices offering fast data access and retrieval. The DD-39 can sustain a data transfer rate of 5.9 Mbyte/sec with an average access time of 18 milliseconds (msec); the DD-49 can sustain a rate of 9.8 Mbyte/sec with an average access time of 16 msec. These disks are the fastest available, and when combined with the data handling and buffering capability of the IOS, they provide unsurpassed I/O performance. Typically, DD-49 disks are configured on CRAY X-MP/4 and CRAY X-MP/2 systems, while DD-39 disks are configured on CRAY X-MP/1 systems. Up to 32 disk drives may be configured in any combination.

CRAY software

The Cray Operating System (COS) and Cray FORTRAN compiler (CFT) will support 16-million-word memories and the new size SSD. COS and CFT have been modified to provide extended memory addressing, and CFT is able to generate vector code to take advantage of the gather/scatter and compressed index hardware. Cray's other software products have also grown to provide users ready access to the power of the new CRAY X-MP systems. □

CRAY X-MP mainframe highlights:

- □ Four processors sharing 8 or 16 million words of ECL bipolar memory on the X-MP/4, or
- □ Two processors sharing 4, 8 or 16 million words of MOS memory on the X-MP/2, or
- □ One processor with 1, 2, 4 or 8 million words of MOS memory on the X-MP/1
- □ 9.5 nsec clock time
- □ 38 nsec (on X-MP/4) or 76 nsec (on X-MP/1 and X-MP/2) memory bank cycle time
- □ 6-Mbyte, 100-Mbyte and 1000-Mbyte channels
- □ SECDED memory protection
- □ Four parallel memory ports per processor
- □ Flexible hardware chaining for vector operations
- □ Second vector logical unit per processor
- □ Gather/scatter and compressed index vector support on all models
- □ Flexible processor clustering for multitasking applications
- □ Dedicated registers for efficient interprocessor communications and control

I/O Subsystem highlights:

- □ Parallel disk streaming capabilities, one controller per disk cabinet
- □ I/O buffering for disk- and tape-resident datasets
- □ Software support for parallel disk striping
- □ Buffer memory-resident datasets
- □ High-performance disk drives
- □ Front-end system communication with IBM, CDC, DEC, Honeywell, Data General and Sperry computer systems
- □ Linkage to workstations such as Apollo and Sun via Network Systems Corporation (NSC) network adapters

SSD highlights:

- □ Memory sizes of 256, 512 or 1024 Mbytes
- □ Support for:
  - Two 1000-Mbyte channels for linkage to CRAY X-MP/4 systems
  - One 1000-Mbyte channel for linkage to CRAY X-MP/1 or X-MP/2 systems
- □ SECDED memory protection
- □ Software support to allow existing programs to use the SSD without program modification
- □ Direct data path (100-Mbyte channel) between SSD and IOS

Computational research at General Motors Research Laboratories

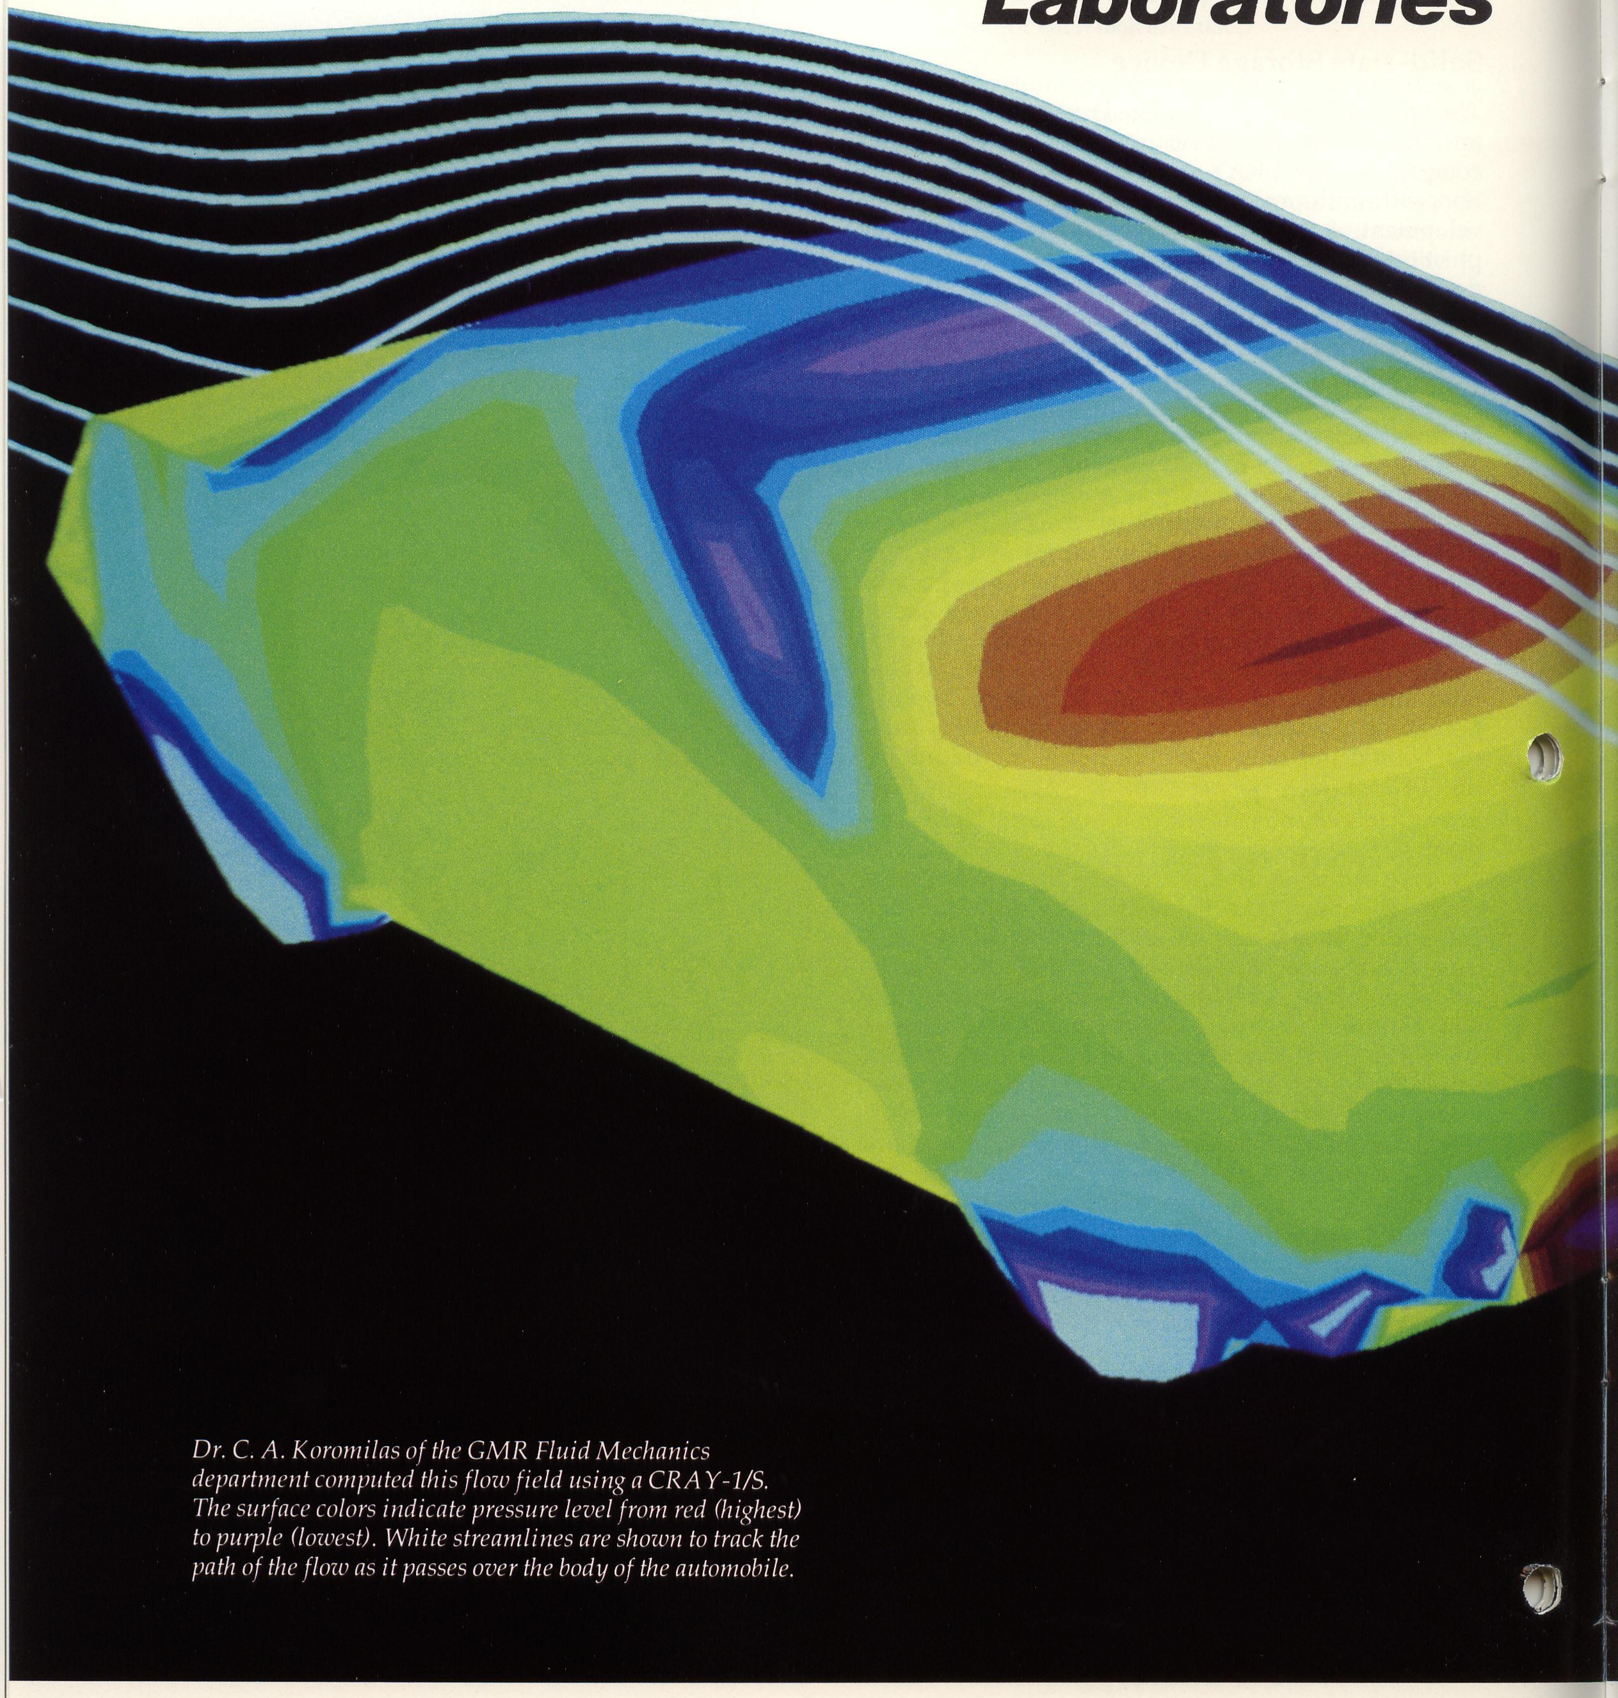

Dr. C. A. Koromilas of the GMR Fluid Mechanics department computed this flow field using a CRAY-1/S. The surface colors indicate pressure level from red (highest) to purple (lowest). White streamlines are shown to track the path of the flow as it passes over the body of the automobile.

Few who saw the first "horseless carriage" could have imagined the impact automobiles would have on our world, but today many of us take them for granted. Cars and trucks have fostered the growth of suburbs and helped turn a former frontier into a vacationland. Increased consumer demand has driven the automotive industry into fierce competition. To meet consumer demand effectively and to stay competitive, automotive manufacturers today must develop safer, more economical and more comfortable vehicles at an increasingly rapid pace.

Computational analysis is a tool that helps manufacturers meet this need by dramatically accelerating vehicle design cycles. Design proposals now can be analyzed via computer, thanks to the high-speed capabilities of supercomputers such as the CRAY. Applied early in the development cycle, computational methods enable designers and engineers to predict the impacts of design changes before any resources are spent building and testing prototypes. By enabling manufacturers to bypass many cycles of time-consuming and expensive prototype tests, computational analysis provides a cost-effective competitive edge.

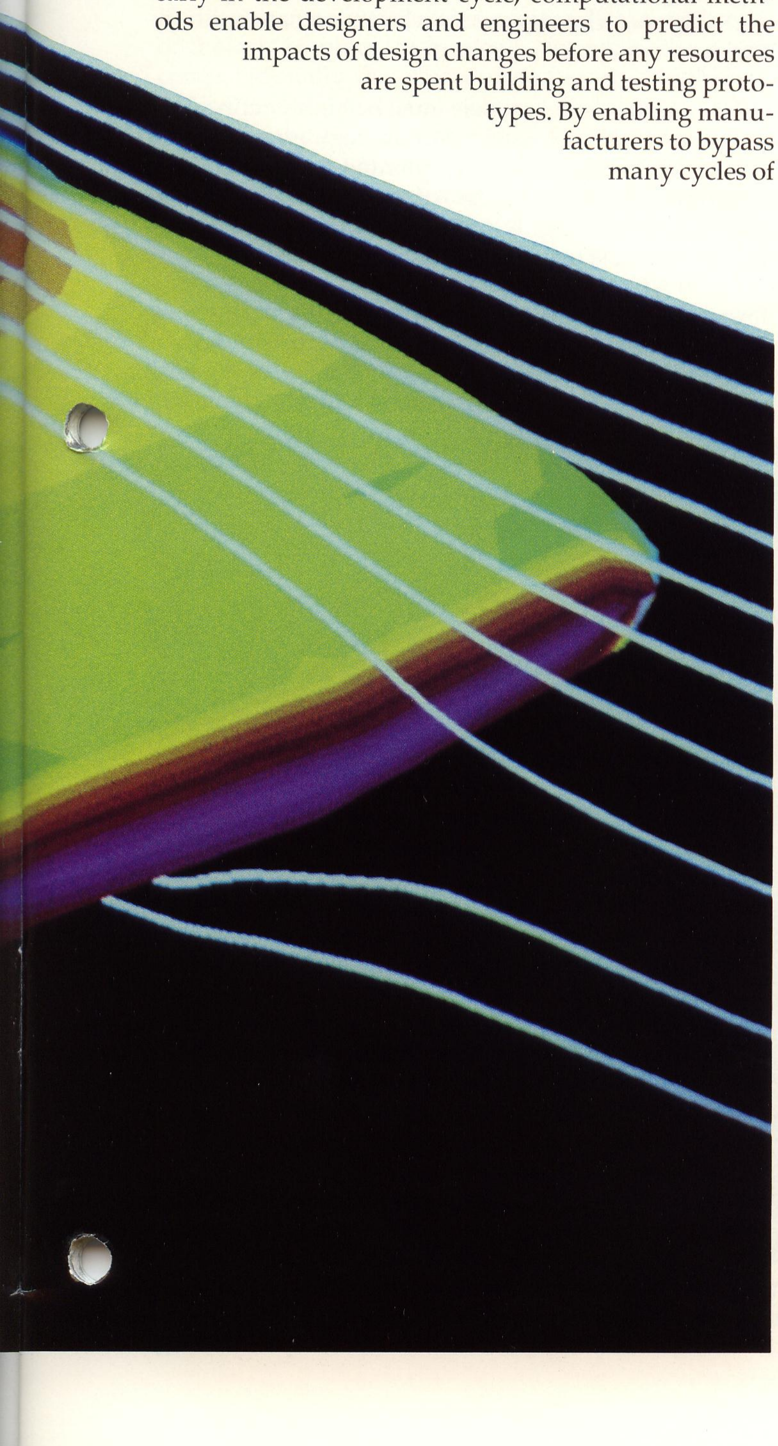

Scientists at the General Motors Research Laboratories (GMR) in Warren, Michigan, are developing computational tools for GM designers and engineers. At the heart of the lab's computational efforts is its CRAY-1/S supercomputer. Several ongoing research projects rely on the CRAY to solve problems too large for conventional computers. Computational efforts currently focus on vehicle aerodynamics, fuel efficiency, emission levels and passenger compartment noise.

Aerodynamics

Automakers study aerodynamics because it affects two primary aspects of vehicle performance: 1) fuel economy, which is directly affected by aerodynamic drag, and 2) handling, which is directly affected by crosswind-generated aerodynamic forces and indirectly affected by aerodynamic lift.

Designers traditionally evaluated vehicle aerodynamics in wind tunnels. While such testing is still extensive, interest is growing in using computers to simulate wind tunnel conditions. Computational research, however, does not stand alone. "We see it as a kind of symbiotic relationship, with each half of the problem feeding the other," explained Dean Hammond, senior staff research engineer in GMR's Fluid Mechanics Department. "At this point, we're not setting out to take a paper car and predict the aerodynamics. We're trying to gain an understanding of flow fields by combining computation and experiment." By learning to predict computationally the flow fields associated with various body shapes, designers can save time and money that would otherwise go into trial-and-error testing of prototypes.

GMR's Fluid Mechanics Department is working to refine aerodynamic simulation programs so they can be applied during the initial stages of vehicle development. Codes similar to those used at GMR have been developed by and for the aerospace industry and are commercially available. But GMR has contracted for its own customized software because the standard aerodynamics codes are not well suited to automotive needs. Unlike aircraft aerodynamicists, auto aerodynamicists must be able to predict strongly separated flows and handle ground proximity effects.

Results from a sample potential flow calculation performed on the CRAY at GMR are shown at left. This problem is typical of the design exercises for which the CRAY is used. These computations were performed using panel methods, a variant of the boundary-integral technique, that involve solving a linear system of equations of order equal to the number of surface elements — typically 1000 to 3000.

"In such a case, the computational speed of the CRAY and its large real memory become mandatory," Hammond said. "The CRAY ran this problem 14 times faster than the mainframes we were using before. Now we can consider doing some of our problem-solving interactively."

Ironically, work intended to develop computational tools has actually increased wind tunnel use. Whereas wind tunnel tests were previously run on virtually every proposed vehicle shape, access to the CRAY inspired GMR's aerodynamicists to attempt more systematic evaluations. However, refining and verifying the computational models requires additional wind tunnel data, more detailed data and different kinds of data.

For example, turbulence is a dominant factor shaping flow fields, particularly with separated flows, which are important to automotive aerodynamics. However, modeling turbulence is very difficult because it must be done statistically. "We can vary turbulence in the computational model and track it against experimental results," Hammond explained. "But to make the comparison you need very detailed data from the wind tunnel. We used to measure just forces. Now we measure forces, pressures all over the surface, three components of velocity, pressure and some turbulent properties, such as kinetic energy, in the exterior field. The wind tunnels are still being used, but in a new way."

"We are trying to combine this modeling with parallel experimental efforts to accelerate our learning. Ultimately, we're interested in developing computational tools to speed the design process..."

Devising ways to optimize vehicle aerodynamics during early design is a major goal of this research. The CRAY not only makes it practical to apply aerodynamic analysis, it also speeds up the development of the tools and methodology to do so. The sooner aerodynamic codes are refined for use by designers and engineers, the sooner their impact will be felt on vehicle development. This point is not unappreciated by the GMR aerodynamicists.

"Before we had the CRAY, we would have to wait two weeks just to see if we'd made the right algorithm modification," Hammond noted. "Now we can get the same answer in a couple of hours. The CRAY has really made our computational aerodynamics program possible. It gives us answers in a reasonable amount of time and with enough geometric complexity that we can address realistic problems."

Engine flow and combustion

Along with aerodynamic considerations, auto designers and engineers would like to predict analytically the fuel efficiency and emission levels of various engine types. This will require accurate models of flow and combustion processes within the engine. Engineers at GMR are using the CRAY to develop such models. Roger Krieger and Keith Meintjes of the GMR Fluid Mechanics Department are members of a group devoting considerable effort to achieving this goal.

GMR's engine interests include conventional spark ignition engines and direct-injection (DI) engines. DI engines, such as stratified charge and diesel engines, have potential fuel efficiency advantages over conventional engines. In DI engines, fuel and air mix in the combustion chamber rather than in the intake manifold or port. The stratified charge engine appears to avoid some of the soot emission problems associated with diesels.

"We are trying to combine this modeling with parallel experimental efforts to accelerate our learning," Krieger explained. "Ultimately, we're interested in developing computational tools to speed the design process and improve the degree to which we can optimize engines."

"The computational models must be highly refined because conventional spark-ignition engines are already highly optimized," Meintjes pointed out. "But direct-injection engines, because of their complexity, are generally not as highly optimized, so they will probably benefit most from the modeling."

Engine flows are extremely complex. Typically they are three-dimensional, compressible, transient, turbulent, multiphase (when sprays used, as in DI engines) and reactive (during combustion) and are characterized by flow separations. The interactions of these characteristics are modeled with KIVA, a three-dimensional code generated at Los Alamos National Laboratory (LANL) on a CRAY and further developed at GMR.

"The two-dimensional forerunner of KIVA was also developed on a CRAY at Los Alamos, but could easily be run on other scalar machines," Meintjes explained. "When work began on KIVA, LANL realized very quickly that a supercomputer was required. So the final version is vectorized for the CRAY."

The equations used to model the flows are the conservation (transport) equations applied to mass, momentum in each coordinate direction (Navier-Stokes equations) and energy. In addition, they account for each chemical species (perhaps ten) and for turbulence, droplet spray motion and combustion.

Depending on the particular submodels used to treat turbulence, spray and combustion, on the order of 20 equations require solution. These are nonlinear partial differential equations solved in finite-difference form over a computational grid representing the geometry of the combustion chamber (Figure 2).

"We've found that detailed solutions for these processes require 45 or more cells in each of the three coordinate directions, which amounts to about 91,000 cells," Krieger explained. "With 20 equations to solve simultaneously in each cell, we need computer storage of roughly 100 million words to make fully resolved three-dimensional calculations. Our current CRAY memory restricts calculations to grids of about 25 cells in each direction. But even for these coarse grid calculations, a supercomputer is required for practical computation times."

Because of the massive amounts of data the code produces, GMR researchers are also developing graphics programs to help them interpret the data. (A calculation of 20 equations for 91,000 cells and 2000 time steps will produce about 3.6 billion data points.) A sample graphic representation of an injection spray in a stratified charge combustion chamber is shown in Figure 3. Such a figure can reveal at a glance any anomalies and help identify correlations among the several variables plotted.

Using supercomputer codes such as KIVA and state-of-the-art graphics technology, scientists can study computationally the effects of numerous variables on engine performance. For example, the size and shape of the combustion chamber and the angles at which fuel is injected into the chamber can affect performance profoundly. Researchers are using the codes to improve understanding of the relationships among these factors. Attempting to do this by purely experimental methods is impractical because of the time and expense required and the difficulty of obtaining measurements as detailed as those computationally derivable. Applying computational methods early in the design cycle reduces costs and development time.

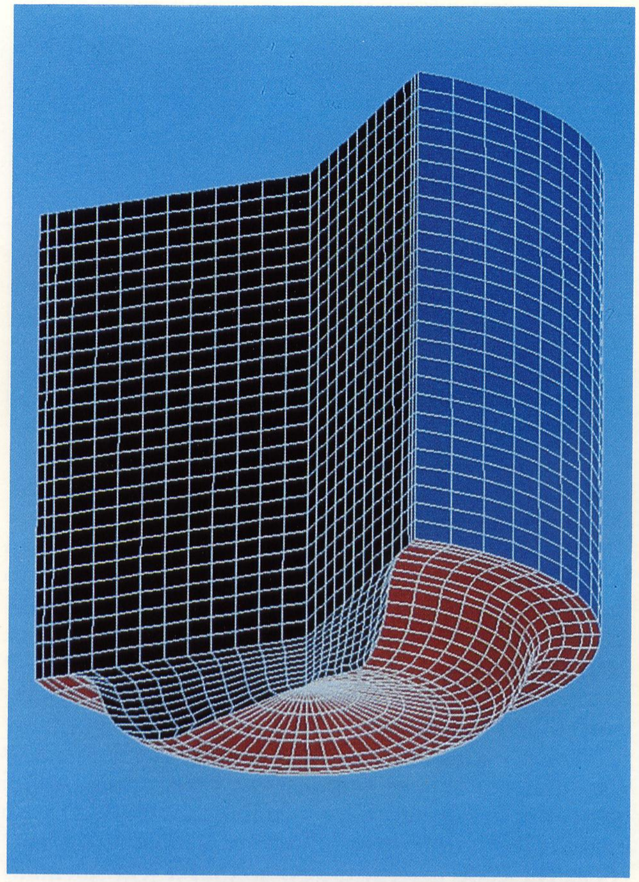

Figure 2. Computational grid representing combustion chamber geometry.

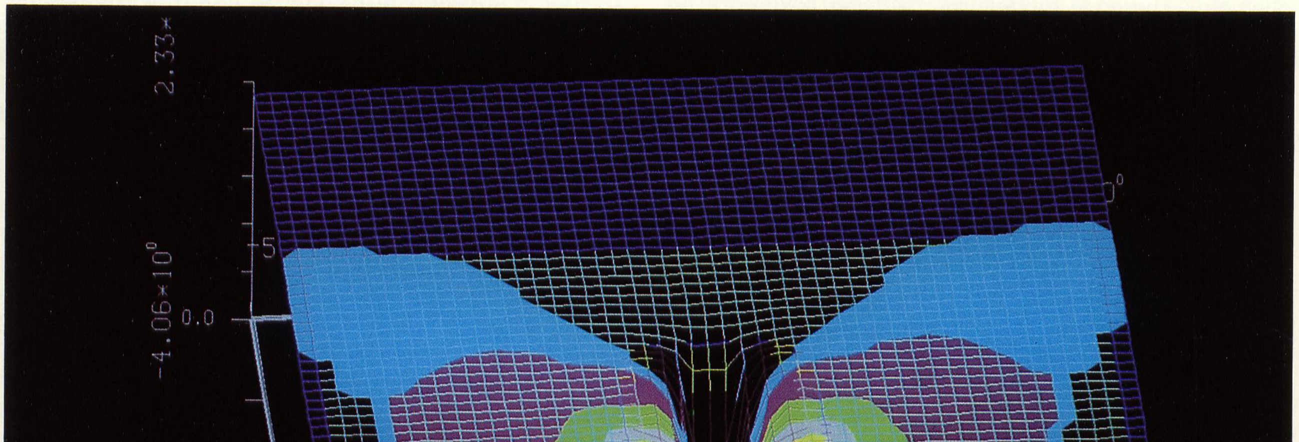

Figure 3. The interaction of an axisymmetric hollow-cone spray with an oncoming airstream. The colored symbols show drop size and location. Surface height shows the axial gas velocity component; grid color shows the turbulence length scale and surface color shows drop number density.

Emission modeling

An automobile or truck is powered by a carefully timed sequence of explosions in its engine's combustion chambers. Unfortunately, this process produces more than horsepower. It also generates several types of air pollutants. Increasing public concern over these pollutants, as well as increased government regulation, has prompted manufacturers to step up their emission control research. Along with devising various types of converters and valves to trap and control pollutants, the industry is studying combustion chemistry and physics in an attempt to combat pollution at its source. Steve Harris, a chemist in GMR's Physical Chemistry Department, and Cleve Ashcraft of the Computer Science Department have been modeling soot formation to better understand how engine design affects this type of emission.

When fuel burns in an automobile or truck engine, most of it is transformed chemically into carbon dioxide and water. However, some of it can end up as bits of carbon, or soot. Harris and Ashcraft have developed a code that models the formation of soot particles during combustion. As soot particles collide, they tend to stick together and grow in size, a process called coagulation. Particles can also grow by collision with gas particles. This is called surface growth. Both processes are incorporated into Harris and Ashcraft's code. The code is refined now so that it will calculate the number of soot particles and their size distributions for a given set of experimental results. To accomplish this, the code models the collision frequency of several thousand particle sizes. A differential equation associated with each particle size describes a collision coefficient that quantifies the frequency with which that particle size will collide with each of the other particle sizes.

. . . computational analysis is a new tool with infinite promise for accelerating product development cycles and conserving resources that traditionally have gone into building and testing development models.

"The number of 'size 10' particles can change depending on how often they collide with 'size 1', 'size 2', up to 'size 3000', particles," Harris explained. "So for every given size particle there are thousands of other size particles with which it can collide. Each equation will have several thousand terms, and there are several thousand equations. Depending on the initial number of particles and their size distribution, the code typically must calculate several million collision coefficients."

Running such a complex code on conventional mainframes proved impractical, so Harris transferred the problem to the CRAY. "The CRAY turned out to be an ideal machine for what we're doing because we're solving thousands of simultaneous differential equations," he said. "The solutions aren't dependent on each other, so everything vectorizes. We've seen well over a factor of ten speedup using the CRAY."

The ultimate goal of the research is to gain enough understanding so that engine design or operating conditions can be modified to minimize soot production. Using this and other codes during the development cycle may minimize the need for pollution control devices once a vehicle is built.

Ignition chemistry

The chemistry of fuel-air ignition is another factor engine designers would like to take into account early in the design process. Ignition is the complex sequence of chemical events imediately preceding combustion. Researchers of ignition chemistry attempt to gather information to resolve problems such as how to ignite the fuel-air mixture using the least amount of energy and how to ignite it to best facilitate combustion.

"Ideally, we'd like to run engines lean, using lots of oxygen to completely burn the fuel," explained Tom Sloane, a GMR chemist. "Running lean provides some thermal efficiency advantages and possibly pollution advantages as well. But these lean mixtures are difficult to ignite and burn."

To study the chemistry of ignition, Sloane performs one-dimensional simulations of various flame shapes. Modeling enables him to determine changes in the temperature and concentration of several chemical species in a combustion chamber as a function of time. The calculations typically involve 24 chemical species and 150 chemical reactions and account for physical phenomena such as molecular diffusion, convection and conduction.

"Using the code, I can put energy into the system in the form of heat or chemical radicals, or a combination of heat and radicals. This way I can investigate the relative effects these different energy forms have on the ignition process," Sloane explained.

Sloane uses commercially available combustion modeling software and a combustion modeling code written at Lawrence Livermore National Laboratory. He said the commercial code was converted to the CRAY, resulting in a speedup factor of about nine.

"The complete flame calculation took about ten hours on the mainframe we had been using," he explained, "so one job would get spread over a few days. The speedup on the CRAY brings it down to a little over an hour, which is practical to run in one shot. Calculations that took me a week now take about a day."

By combining ignition chemistry and other codes, designers and engineers may someday be able to calculate all the important performance characteristics of an engine during its first development stages. Such an ability would signal a revolution in development methodology, perhaps as significant as the revolutionizing effect of the assembly line on manufacturing.

Structural-acoustic design

A very different supercomputing application from those discussed so far is the analysis of passenger compartment noise. A structural-acoustic analysis code developed at GMR has already been used to minimize noise in a production vehicle. Once again, the application is intended to provide designers and engineers with a means to identify and correct potential problems early in the development phase.

Low-frequency compartment noise caused by vibration of the surrounding body panels is a potential vehicle problem. Bumpy roads, aerodynamic inputs and a vehicle's powertrain can initiate low-frequency vibrations in panels. The vibrating panels then act as drum heads or loudspeaker diaphragms, producing sound waves — and creating an annoyance to the vehicle's occupants.

Several years ago, GMR engineers developed a finite element acoustic model. However, it did not provide enough information to identify either the panel source of a noise or the structural modifications that would reduce it. To identify the source, it was essential to couple a body vibrations model with the acoustic model. Don Nefske and Shung Sung, GMR staff research engineers, set out to do so and succeeded in coupling the acoustic model with a structural system model that incorporates MSC/NASTRAN, a widely-used structural analysis program.

The structural model provides the vibration response of body surfaces during actual vehicle operating conditions and, with a coupling methodology devised by Nefske and Sung, links it with the acoustic model to translate vibration data into acoustic data. The resulting structural-acoustic model accurately predicts quantitative measures of low frequency noise at any location in the vehicle interior for simulated road or operating conditions.

If the model predicts an unacceptable noise level, a separate post-processor program is used to determine the contribution of the various panels and thus to locate the primary source of the noise (Figure 4). "By defining and quantifying the noise paths," Sung explained, "the model can suggest structural modifications — perhaps the redesign of a panel — early enough in the design stage to avoid more expensive add-on treatments that can add weight to the vehicle."

"The body model itself is a very large and complex code," Nefske added. "Using the CRAY makes it practical to run. Before we had the CRAY we had to limit considerably the number of degrees of freedom in the body model, and that compromised the model's accuracy. Being able to increase the number of degrees of freedom gives us higher fidelity in the structural-acoustic model — it increases the range of frequencies we can examine."

The structural-acoustic analysis code is being applied to problems in several ongoing vehicle programs. The model's first application throughout a program came early in the design phase of GM's new-sized van, the 1985 Chevrolet Astro/GMC Safari. By using prescribed forcing inputs that occur during on-the-road operation, a potential boom noise problem was found in a pre-prototype design of the van, well before any hardware was built. The post-processing program was used to identify the noise paths and to pinpoint the body panels that were causing the problem. Subsequent structural modifications were made to the prototype design. The model indicated that the modifications reduced vehicle noise in the 45 to 60 Hz range by approximately ten decibels. It also verified that the changes reduced the panel contributions to the interior noise over the entire 20 to 60 Hz range. "The overall noise level was reduced by an estimated seven decibels, which is enough for people to notice the difference," Nefske said.

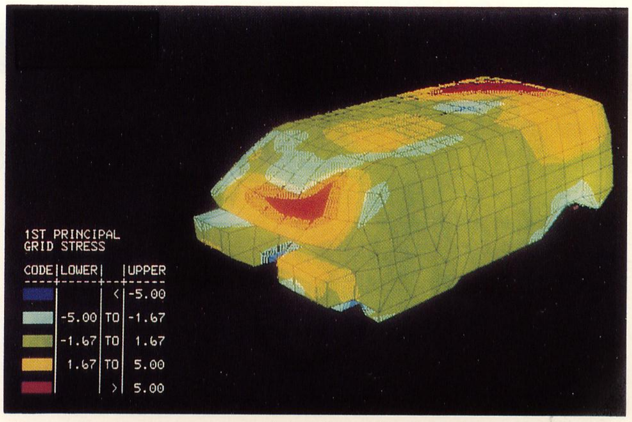

Figure 4. An energy participation display showing the contribution of various interior panels to interior noise at a given highway speed.

Conclusion

The newest cost and time saving technologies have to be exploited for any business to remain competitive. In the case of the automotive industry, computational analysis is a new tool with infinite promise for accelerating product development cycles and conserving resources that traditionally have gone into building and testing development models.

The balanced design of CRAY computers makes them ideal tools for addressing research problems in many disciplines. With its interests in both basic and applied research, the automotive industry is a microcosm of the world of supercomputer applications. However, despite knowledge already gained at GMR and elsewhere, the potential for applying computational analysis to vehicle development is just beginning to be realized. □

Acknowledgement

Special thanks to the following people at GMR for their help in preparing this article: Lee Buzan, George Dodd, Myron Ginsberg, Dean Hammond, Steve Harris, Roger Krieger, Keith Meintjes, Don Nefske, Tom Sloane, Shung Sung and Cray site analyst Bob Sabourin.

Computational analysis in automotive design

Dean Grubbs
General Motors Corporation

Introduction

Like most modern industries, automobile manufacturing is turning to state-of-the-art technology to meet the demands of the 1980s and 1990s. Automotive design in particular is being revolutionized by new technologies, primarily through improved engineering analysis software and supercomputers. At first, such applications seem radically different from traditional design engineering practices. However, a closer look shows that only the approach is different; the problems, the principles of design and testing and the process of redesign remain the same.

The design process

Historically, automobile designs were developed from an empirical process. Engineers relied on design experience and prototype tests in designing components. If a test part failed, it was redesigned until it was acceptable. This process required moving designs from engineering to experimental manufacturing to lab test or proving grounds, and then back to engineering. Although this process was very time consuming and costly, it was practical because many component designs from previous production models could be reused or modified. Engineering refinements seemed

Figure 1. Stiffness and buckling analysis of an automobile hood. The color plot shows the areas where buckling would occur in a crash. Barrier crash loads were simulated in the model to produce the effect. Advanced hardware and software subroutines will enable design engineers to conduct complete nonlinear crash analyses of similar and more complex components.

Figure 2. Section view of one quarter of a piston. Color shows the stress due to side loads for the prediction of piston configuration durability. Heat transfer analysis has been used in predicting operating temperature distribution combining cylinder pressure and inertia effects.

Figure 3. This color stress plot represents a single cylinder of an engine block. The color pattern maps out the various degrees of stress. In this analysis, complex geometric loads were analyzed for minute distortions. Loads transmitted to the engine block walls under combustion pressures are major contributors to bore distortions. Analyses of such loads and distortions are important to achieve fuel economy, longer product life and lighter engine blocks.

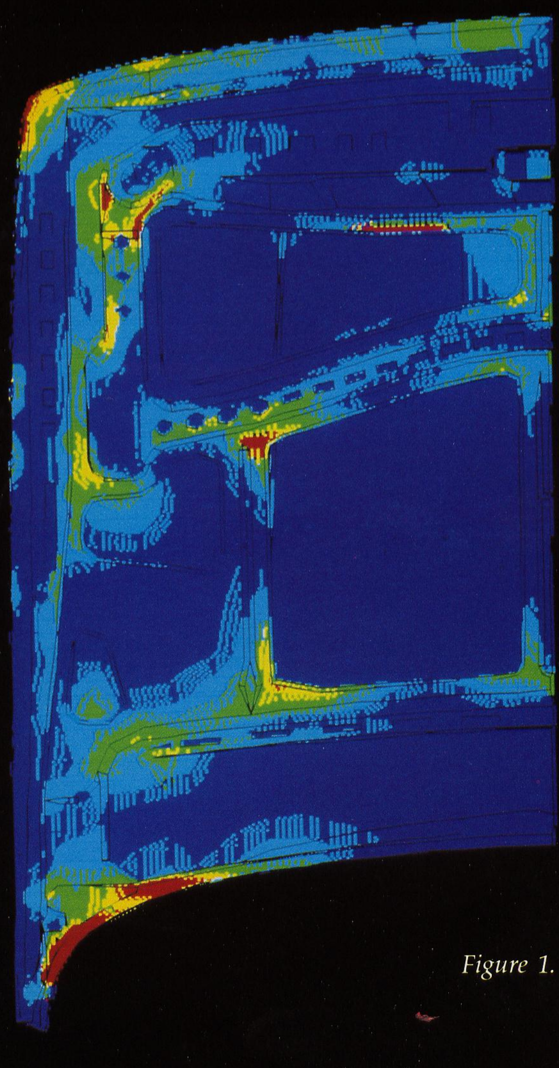

Figure 1.

to be a matter of keeping or modifying time-tested designs and discarding those that did not work.

Today, finite element analysis (FEA) and finite element modeling (FEM) allow engineering staffs to execute the traditional design process mathematically via computer. Instead of drafting, constructing and testing various prototype design stages, engineers use simulations to design and evaluate components. Once constructed mathematically, the models are subjected to structural evaluations. The evaluation results guide design engineering staffs in the redesign of components to achieve optimum strength and mass before actual experimental hardware is built and tested.

The design challenge

Through necessity, automotive engineers have embraced FEA and FEM, once used exclusively by the aerospace industry. This is largely due to design challenges first encountered in the 1970s when the U.S. federal government required automotive manufacturers to adhere to its new Corporate Average Fuel Economy, emission and safety regulations. The new regulatory demands, rising fuel prices and changing consumer demands combined to strain the traditional empirical approach by contributing to the already growing complexity of auto designs. To be competitive, new designs had to provide weight reduction and improved fuel economy while adhering to many new regulations.

Development time also became a critical factor in creating competitive and marketable designs. Initially, engineers responded by accelerating traditional design practices. However, it soon became clear that creating lighter and more fuel efficient designs demanded innovative procedures. To meet this need, computerized analytical design was introduced in the mid-1970s and is becoming increasingly necessary to produce efficient, strong and lightweight components.

Design applications

Analytical design methods were not incorporated overnight. At first, they were restricted by the lack of efficient software, interactive graphics capabilities and sufficient computer capacity. But the revolution in computer-aided design (CAD) and computer-aided manufacturing (CAM) supported a number of computer graphics developments. As CAD/CAM-oriented hardware and software were applied to FEA and other design analysis tools, CAD graphics routines significantly accelerated the modeling process.

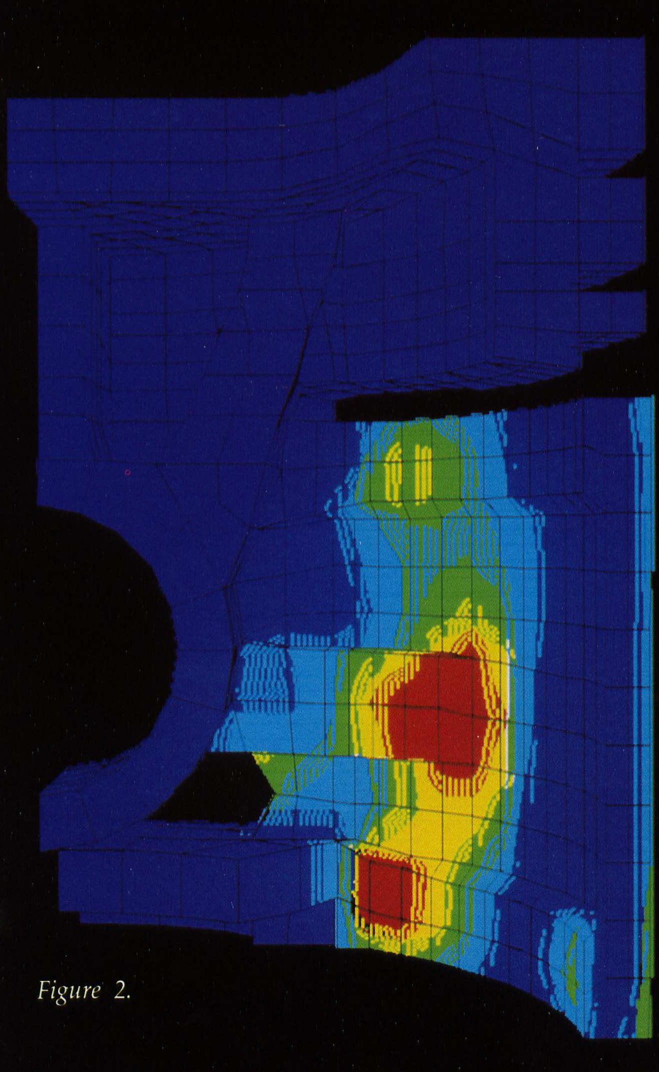

Figure 2.

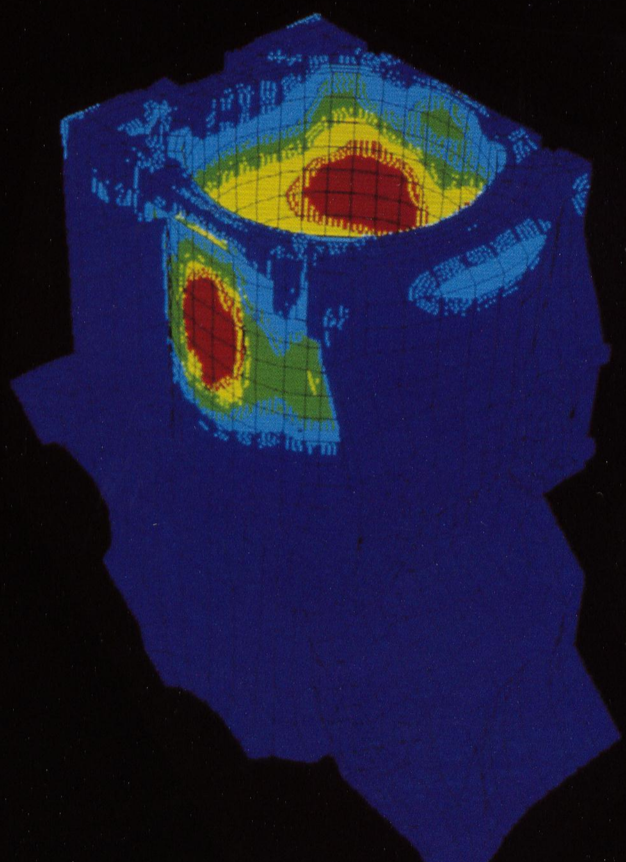

Figure 3.

The CAD revolution moved designs from paper to a computer database. Initially, FEA was severely limited by the labor-intensive task of mapping large models onto layout drawings and revising them manually. The process was time consuming, inflexible and subject to human error. Engineers had to summarize data from stacks of output to analyze the problem as a whole. Today, supercomputers, advanced CAD software and graphics capabilities have liberated engineers from early FEA restrictions.

Shorter analysis run times and interactive graphics have similarly revolutionized automotive design. Graphic displays show exactly where component structural deflection from stress occurs. These results can be presented in "time slices" of component motion. Color graphics enable designers to view contour lines mapping out varying degrees of stress. While early use of FEA was confined primarily to mechanical engineering, software has also been developed for electrical engineering.

A growing concern of automotive manufacturers is the electromagnetic compatibility (EMC) of their products. As automotive electronics have become more complex with several on-board microprocessors, digital displays and cathode ray tubes, EMC and interaction of system components have become crucial design factors. Because vehicle electronics are not developed "under one roof" and cannot be empirically tested as a system until shortly before production deadlines, traditional design methods limit systems testing to late in the design process. The implementation of electrical engineering analytical design methodology will have the potential for predicting automotive EMC during the design concept phase and vehicle prototype construction.

Future applications

The combination of supercomputers and improved convergence routines for nonlinear analysis may allow engineers to "crush" entire automobiles in crashworthiness simulations before initiating construction and testing of prototypes. Such crash simulations were virtually impossible using earlier FEA methods and computer systems. Analytical crash tests will enable engineers to "see" behind surfaces such as body sheet metal to obtain accurate insights of component interaction.

New capabilities are emerging as software designers create faster programs with improved model generation capabilities. New algorithms to generate FEM

Figure 4. Magnetic flux distribution of a horn. Such analyses are used to obtain device impedance. The multicolored contours represent the magnetic flux density within the design structure and surrounding air.

Figure 5. Connecting rod analyses such as this take shear stresses and compression loading into account. Design weight reductions with increased structural integrity have been achieved in this way. Finite element analysis has helped tremendously in predicting rod failures by evaluating the effects of the bearing journal and distribution of the oil film.

Figure 6. Instrument panel and vibration analysis of component interaction aids in producing structurally sound designs. Such analyses are becoming more important as instrument panels have to support the weight of radios, computers and other mounted systems. Information obtained from these analyses enables design engineers to provide sufficient panel structure to avoid unwanted vibrations.

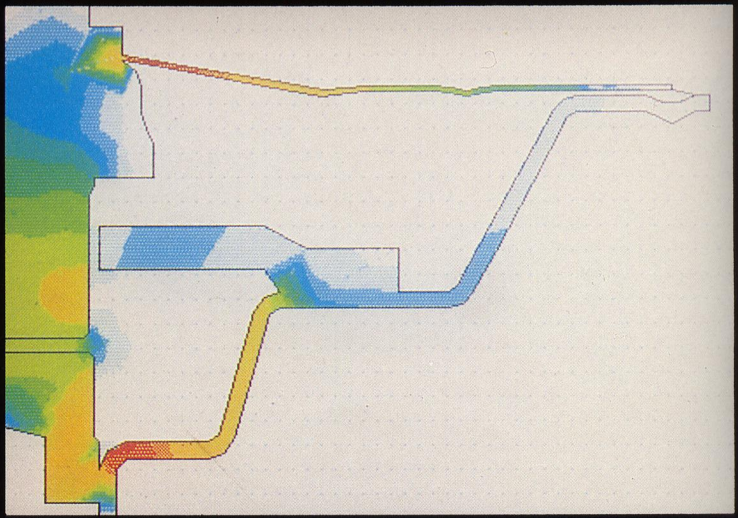

Figure 4.

meshes from wire frame and geometric CAD models are also being developed. These capabilities, along with engineering workstations that continue to improve by providing more computational power at less cost, will further improve the cost-efficiency of analytical methods. Analytical design analysis in the auto industry has already proved savings that have increased geometrically each year.

Eventually, it can be expected that automobile manufacturers will predict every aspect of design performance analytically before building any parts. The only limitations to analysis lie in the availability, speed and memory capacity of computers and in the availability of appropriate software.

In addition, manufacturing capital investment decisions will be aided by mathematical models that calculate the effects of production variation, design tolerance in relation to performance and vehicle aesthetics. Such analyses will be conducted before significant resources are committed for production tooling and will help determine if manufacturing systems will yield acceptable and consistent production variations.

Shorter design lead times, stringent emissions standards and demands for safety and overall quality improvements continue to challenge automotive engineers. Analytical design is becoming crucial to meeting these challenges.

Some advances are still needed to fully exploit the potential of analytical methods. Dynamic simulation of individual component characteristics, design criteria and full vehicle performance during collision will require a major increase in computing power. Mathematical evaluation of manufacturing processes before final procurement will require consolidating engineering talents and resources with state-of-the-art software and hardware.

It is becoming clear that supercomputers will soon be mandatory for engineering staffs to remain competitive. Though expensive, they will pay for themselves by enabling engineering groups to produce cost-effective automobile designs that are lighter, more fuel efficient and more durable than current models. □

Acknowledgement

Dean Grubbs is Manager, Analytical Design Analysis, Flint Engineering Center, Buick Oldsmobile Cadillac Group, General Motors Corporation, Flint, Michigan.

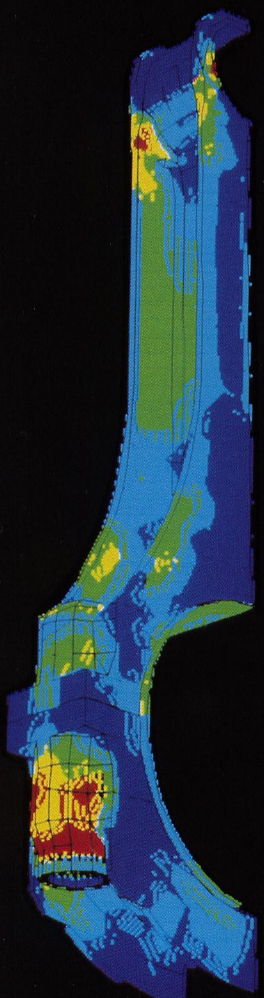

Figure 5.

Figure 6.

Examples of computational methods in automotive engineering

Werner Assmann and Dr. Erich Schelkle
Porsche Development Center

As computer hardware and software become increasingly sophisticated, they enable the automotive industry to exploit computational analysis for design and engineering work. The benefits of computational analysis complement the strengths of other development tools. Today, the combination of physical testing and computer simulation provide unique insights into automotive design and engineering while helping to reduce cost and development time.

The time saved during development — although not an easily quantifiable advantage — is a major justification for the use of computational methods. Computational analysis is valuable for optimally laying out a component, an aggregate or a system, while minimizing the number of experiment cycles required. This not only reduces design and test effort, but permits development of higher-grade products.

Along with the time they save, computational methods are valued for the accuracy of their results. Boundary conditions and model parameters can be defined uniquely for computational models. In this sense, results are more accurate than those of physical tests, where parameters are subject to statistical variation. Computational methods can save considerable time and cost since physical tests require several trials to guard against statistical error.

But perhaps the greatest benefit of computational analysis is derived from its use early in the development process. In contrast to physical tests, theoretical analysis of a problem enables designers and engineers to study an arbitrary number of variable effects, such as degree and location of stress or deformation. In some cases where no suitable measurement methods are available (for example, nonstationary temperature computations for a brake disc), only computational methods can provide insight into a problem. Furthermore, such methods not only reveal weak points of a design but also localize overdimensioned areas.

Ultimately, only the integrated use of all development tools — design, testing and computation — provides optimum insight into problems and ensures the smooth flow of development operations. The complexity of automotive design and development problems makes it necessary to balance the relationship between simulation and physical tests. Feedback from tests provides data for improving computational models and the accuracy of boundary condition definition. Nonetheless, only very fast vector processors, such as the CRAY X-MP, enable researchers to find economically viable solutions to the kinds of large projects currently being discussed in the aerodynamics and finite element sectors of automotive engineering.

Following are examples of several interesting design problems requiring high-performance computation.

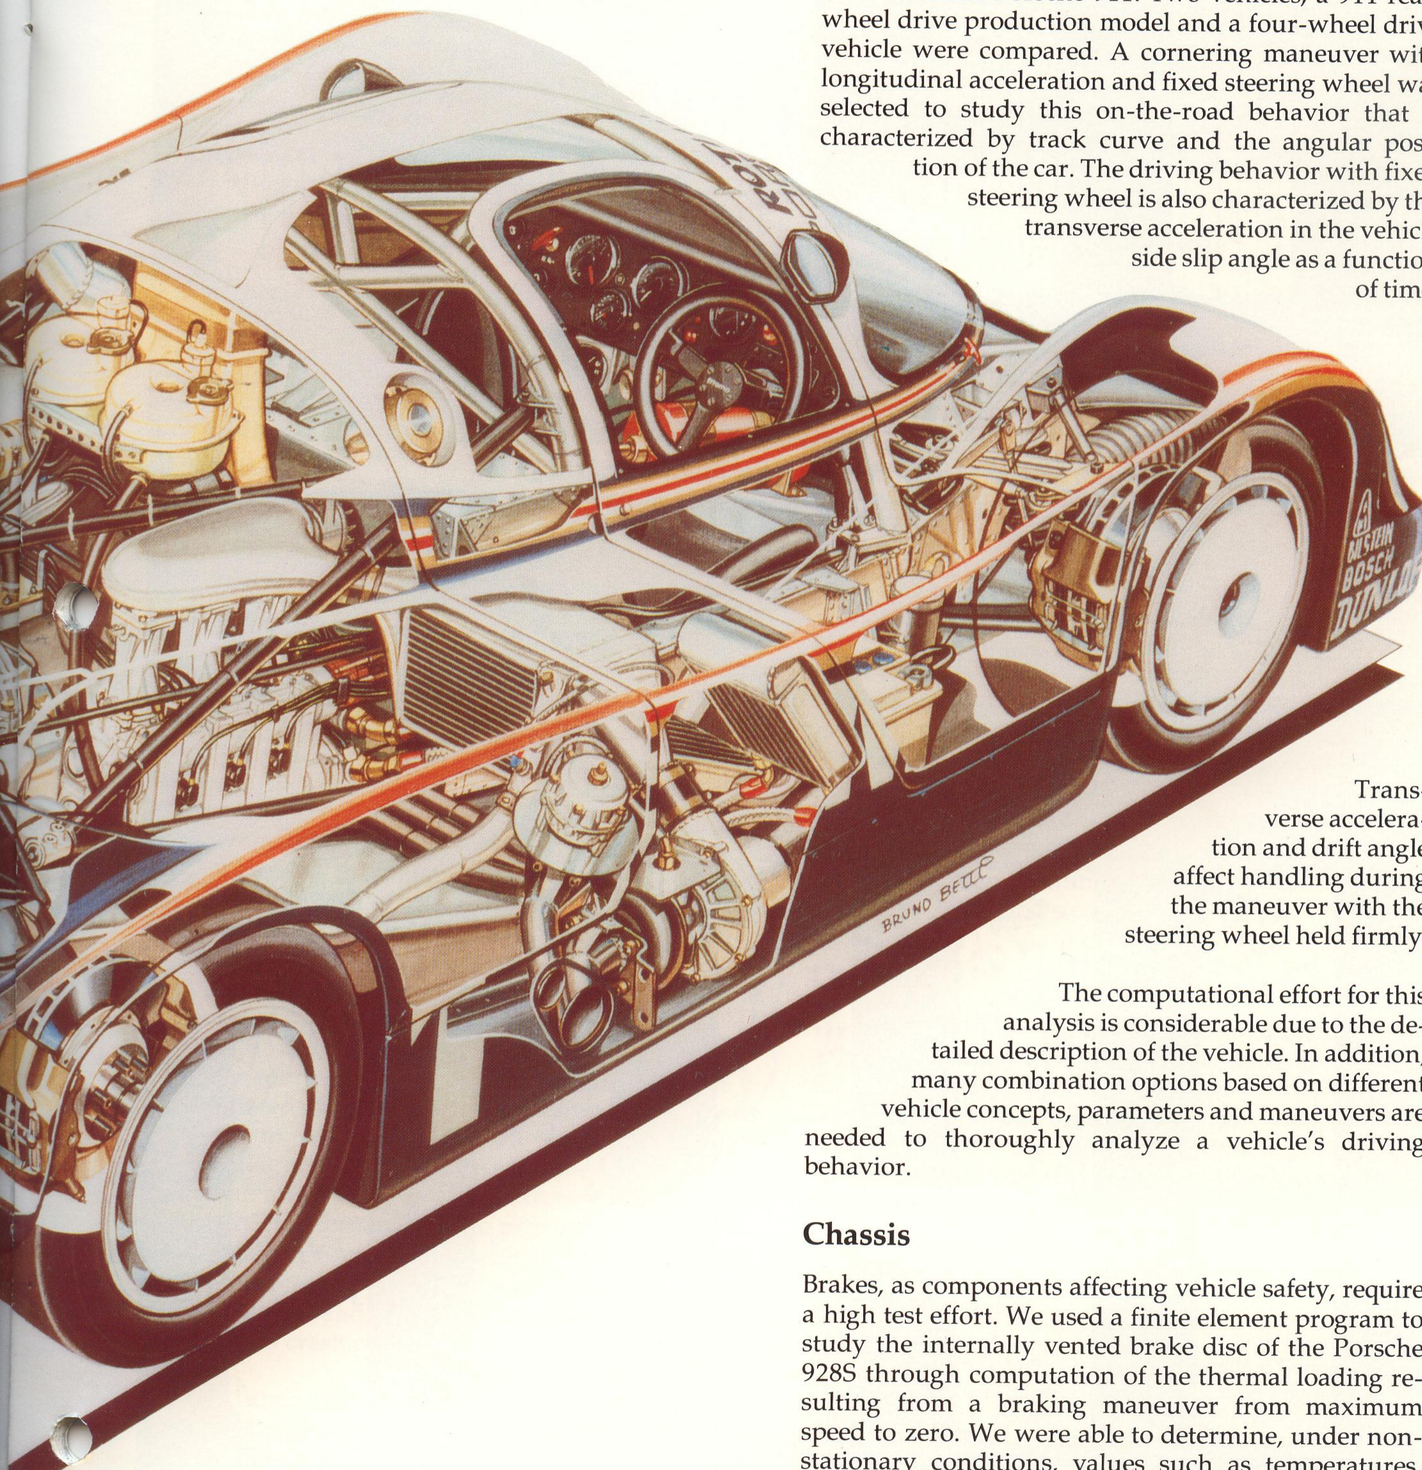

Porsche 956, Group C

Driving simulation

Handling characteristics can be determined by simulating performance of an entire vehicle. Driving simulation programs can also be used to rerun certain test cycles to obtain information on load collectives and to study vehicle concepts during different maneuvers such as steady-state cornering, accelerating or braking in a turn or driving with a cross wind.

This was done, for instance, in a performance study on a modified Porsche 911. Two vehicles, a 911 rear-wheel drive production model and a four-wheel drive vehicle were compared. A cornering maneuver with longitudinal acceleration and fixed steering wheel was selected to study this on-the-road behavior that is characterized by track curve and the angular position of the car. The driving behavior with fixed steering wheel is also characterized by the transverse acceleration in the vehicle side slip angle as a function of time.

Transverse acceleration and drift angle affect handling during the maneuver with the steering wheel held firmly.

The computational effort for this analysis is considerable due to the detailed description of the vehicle. In addition, many combination options based on different vehicle concepts, parameters and maneuvers are needed to thoroughly analyze a vehicle's driving behavior.

Chassis

Brakes, as components affecting vehicle safety, require a high test effort. We used a finite element program to study the internally vented brake disc of the Porsche 928S through computation of the thermal loading resulting from a braking maneuver from maximum speed to zero. We were able to determine, under non-stationary conditions, values such as temperatures, stress and deformations — values that are extremely difficult to measure in physical tests.

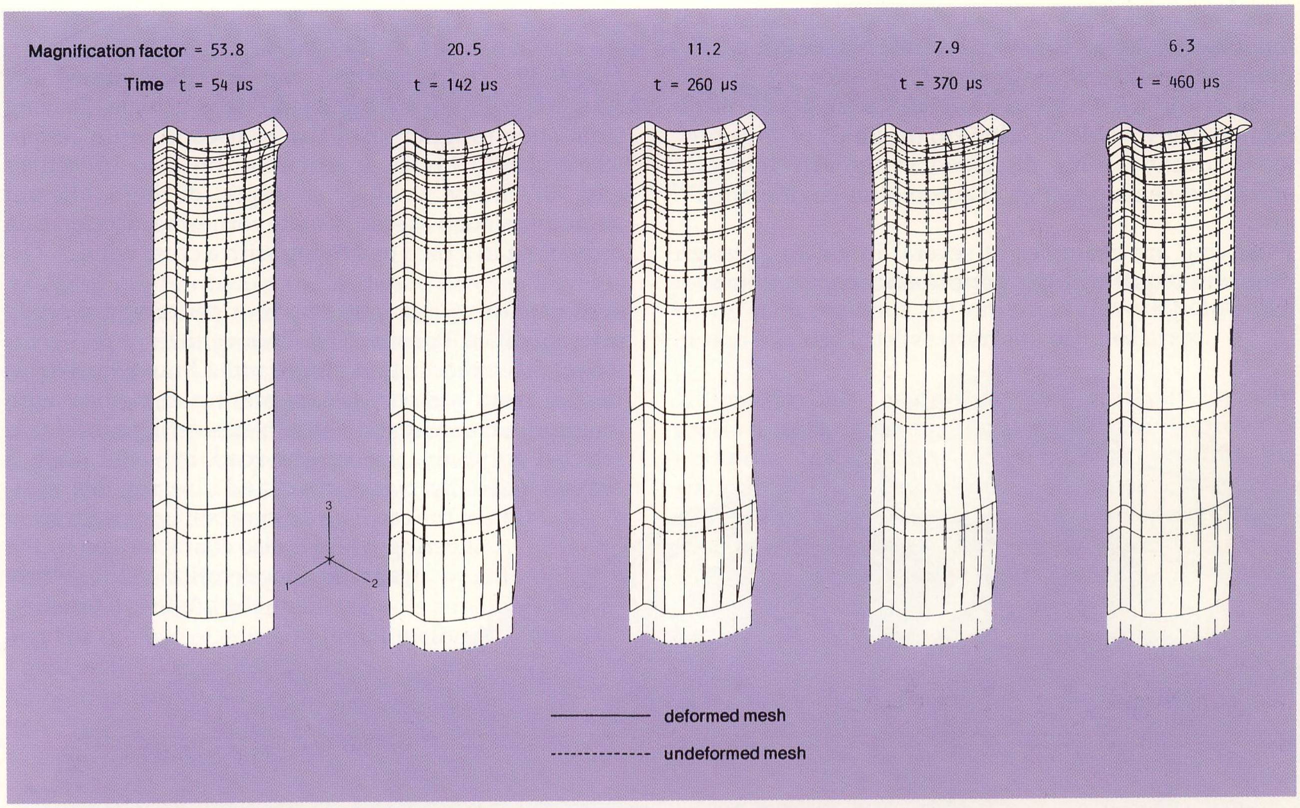

Figure 1. Deformations of longitudinal front member during crash simulation for Porsche 928

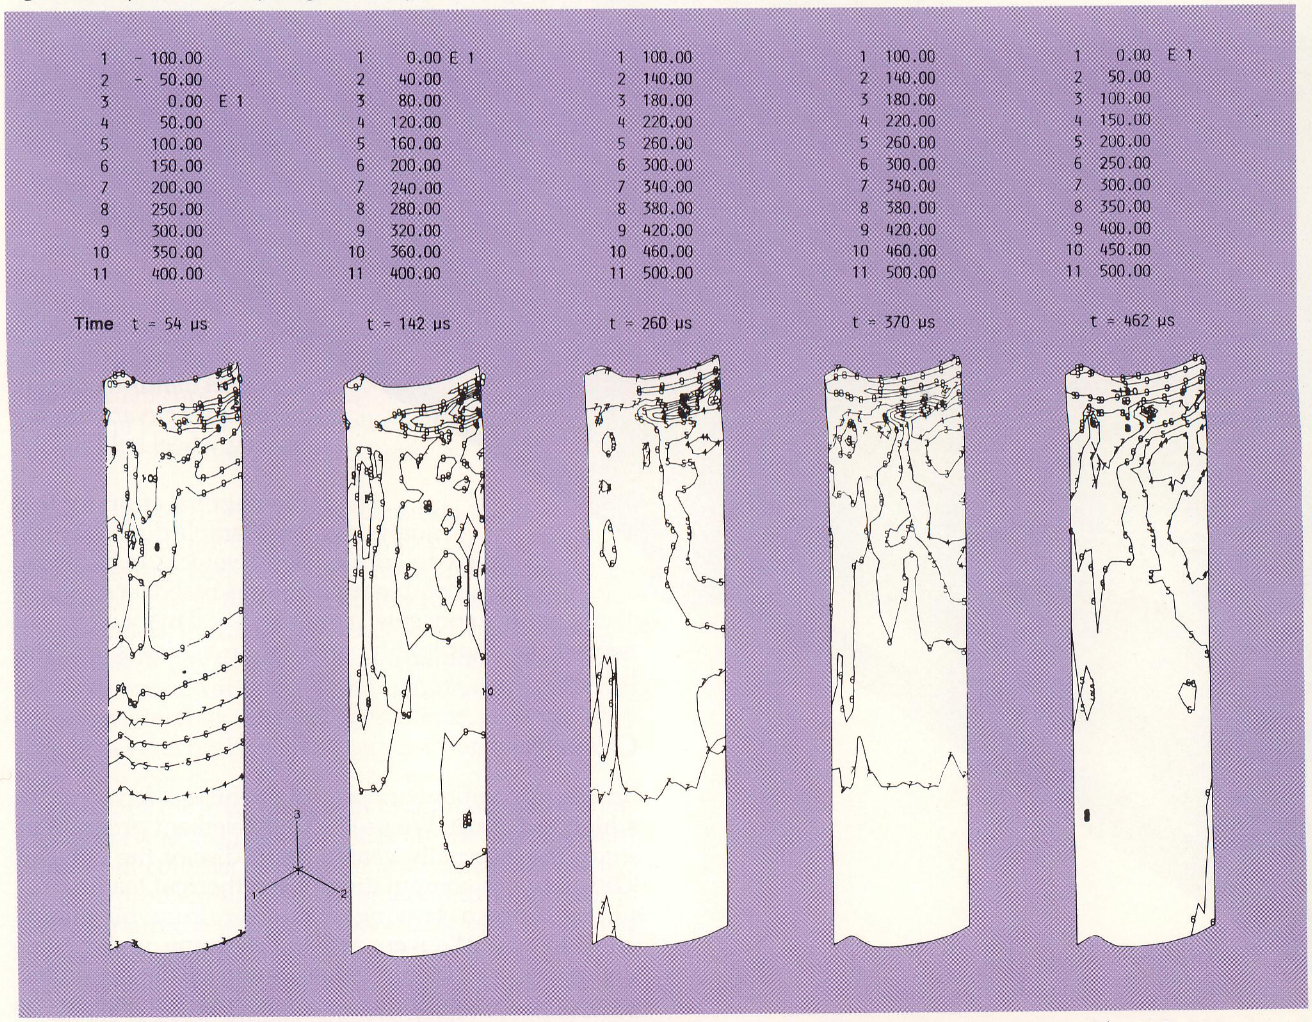

Figure 2. Equivalent stresses (N/mm²) of the longitudinal front member during crash simulation for Porsche 928

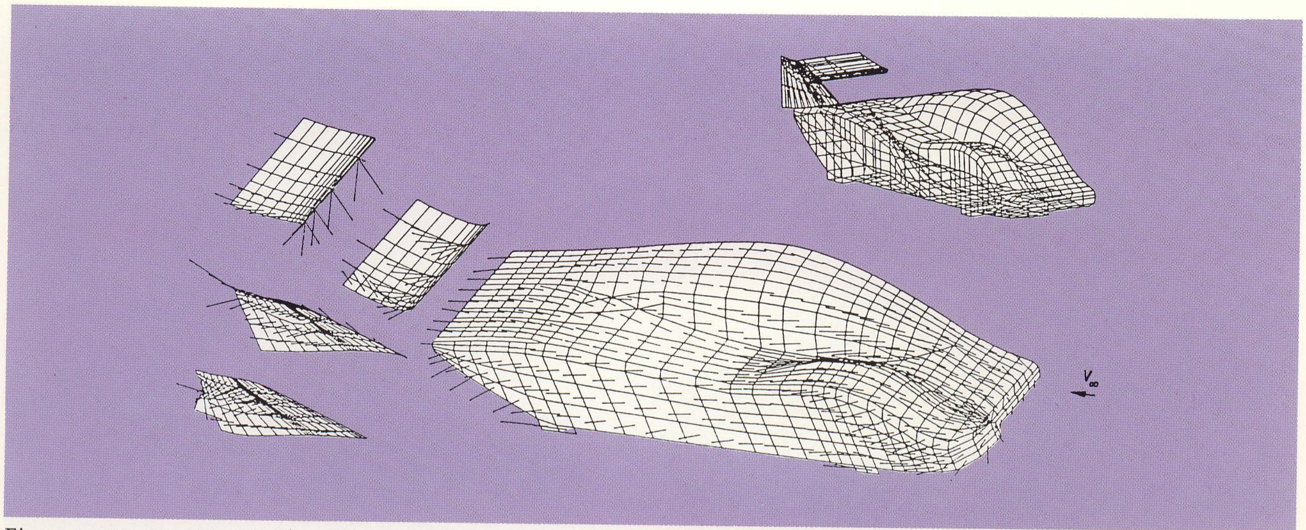

Figure 3. Velocity distribution after panel-boundary iteration for Porsche 956 (Group C)

Body

Sophisticated linear finite element programs provide detailed data on deformations and stresses for bodies and body components. The details one can obtain depend on the refinement of the mesh structure in critical areas. Computations are aimed at attaining optimum overall rigidity while minimizing vehicle weight.

In crash computations, the objective is to achieve optimum conversion of the kinetic energy existing prior to a crash into deformation energy by means of suitable structural supports. Because front-end crash computations with the aid of the finite element method require considerable CPU time, in the past a suspended mass simulation model was used that relied on experimental data of spring stiffness of structural sections. These values were obtained mainly through static compression tests.

However, an inherent disadvantage of suspended mass systems is that they do not allow simulation of the supporting effect of the sheet metal adjacent to the structural supports. For this reason, nonlinear finite element programs are increasingly used for crash computations of body components and complete front-end or rear-end sections.

A highly complex and compute-intensive example is demonstrated by a nonlinear dynamic elasto-plastic computation of a Porsche 928 longitudinal front member (Figures 1 and 2). This example was simulated for a crash into a solid wall (frontal crash) at a speed of 30 km/hr. The objective was to prove that a nonlinear finite element program can correctly evaluate a highly nonlinear process, particularly with regard to the folding mechanism and von Mises equivalent stresses for five different crash phases.

The CPU time requirement for some of the examples shows quite clearly that only large-scale computers such as the CRAY-1 or CRAY X-MP are capable of providing fast and economical solutions for problems of this size.

Aerodynamic considerations

Because wind tunnel testing for body design is time consuming and expensive, aerodynamic computation methods will be used increasingly in the predevelopment stage. With the current panel/boundary layer methods, reasonable values can be obtained for velocity distribution (Figure 3) and pressure distribution up to a flow separation point. With wind tunnels, only very incomplete flow simulation is possible. The difficulty of simulating the appropriate boundary conditions in an experiment is also inherent in wind tunnel testing.

Flow velocities in wind tunnels differ from those encountered on roads because of the absence of relative movement betweeen vehicle and floor. The programs used in the analysis illustrated in Figure 3 can take into account some effects of this relative movement. However, satisfactorily computing viscous flow with separations and wake flow is possible only with the aid of a computer program that can solve the Navier-Stokes equations running on a very powerful computer such as a CRAY system.

Outlook

We believe that the use of computational methods has a favorable impact on the cost/performance tradeoff for product development. Using computational methods early in the design development stage and integrating them into the overall vehicle concept can reduce design cycle time significantly, slow rising development costs and improve vehicle quality. □

Acknowledgement

Werner Assmann is Manager of the Computation Department and Dr. Erich Schelkle is Supervisor of the Structural Engineering Section, Porsche Development Center, Weissach, West Germany.

This article appeared in an expanded form in **CRAYationen**, *a publication of Cray Research GmbH, Munich, West Germany. Special thanks are extended to the authors for reprint permission.*

A hard look at fast wares
(a perspective from the designers)

It's hard to believe that three years have passed since the introduction of the CRAY X-MP multiprocessor computers. In the past three years, customers have benchmarked on the CRAY X-MP and have experimented with the hardware and software. The time and effort taken has been worthwhile. Our multiprocessor and multitasking approach to supercomputing has proven itself on large-scale problems. In fact, as testimony to its success, the CRAY X-MP Series has been expanded twice since its original introduction. Today, customer interest is strong and the CRAY X-MP is in wide demand.

Many features on the CRAY X-MP contribute to its exceptional performance. Some are examined in this article, excerpted from a longer article on the X-MP written by several system designers. This excerpt discusses certain aspects of the X-MP's unique architecture and describes how it evolved from the CRAY-1. Those interested in reading the complete version of the designers' paper or in obtaining more information about the X-MP system are encouraged to review the note and references listed at the end of this article.

Historical note

In 1979, while Seymour Cray was leading the development of the CRAY-2, a separate effort in the company was initiated under the direction of Les Davis and led by Steve Chen. The task was to design a machine more powerful than, but instruction compatible with, the CRAY-1. Three years later, in April 1982, checkout of the CRAY X-MP/2 prototype was completed.

Soon thereafter, design work on a four-processor model began. A parallel effort was undertaken by another Cray engineering group to use the central processing unit (CPU) of the X-MP with the relatively inexpensive MOS memory technology in developing a single-processor model, the CRAY X-MP/1. This project was completed in 1983. Today the expanded and enhanced CRAY X-MP Series includes one-, two- and four-processor models (see related article in this issue).

Design strategy

Several important design strategy decisions were made at the onset of the X-MP project. To shorten the time for circuit design, the designers decided to use the same 16-gate ECL gate array chips used in the CRAY-2. A distinctive characteristic of the X-MP lies in its packaging technology. It is rather remarkable that Cray has been able to use chips with such a low level of integration. The use of 16-gate gate arrays (eight times the integration of the 2-gate chips used in the CRAY-1) had great implications for packaging and a faster machine cycle time. The 16-gate gate array may consume three to four times the power of a 2-gate chip and can create heat dissipation problems if high wattage chips are concentrated in one place. But in order to shorten the signal travel time on the metal foil that connects the chips on the printed circuit board (PCB), chips could not be placed too far apart. Furthermore, to achieve a faster clock rate, the design rules for fan-outs and loading were also strengthened over those of the CRAY-1 design.

These difficult factors made chip placement and PCB routing very challenging. With 200 to 300 chips per module and thousands of latch-to-latch paths to check, the enforcement of design rules became impossible without some kind of computer aids. As a result, for the first time at Cray Research, a CRAY-1 computer was used to enforce and checkout the design rules throughout the design process. This approach not only sped development, but also assured that the design goals were met on target.

There was another packaging difference between the CRAY X-MP and the CRAY-1. On the CRAY-1, standard 6 x 8 inch circuit boards are mounted on both sides of a copper plate. The unit is called a module. Inter-module communication is done through twisted wires three to four feet long. In order to shorten the signal traveling time, a tighter packaging technique was necessary. On the X-MP, where a double module is used, two CRAY-1-like modules were fastened together. All four circuit boards communicate directly through fixed locations via jumpers. Hence, in the same sense that a CRAY-2 module is viewed as being three-dimensional, the CRAY X-MP module can be viewed as being two-and-one-half-dimensional.

The designers had the option of using a new cooling technology to increase packaging density and shorten clock cycle time. But since they were experimenting with new architecture and component technologies, a conservative decision was made to use an enhanced version of the CRAY-1 cooling technique.

The designers also had a choice between super vector speed and faster scalar performance. It was conceivable that adding more vector units to the X-MP would double or even quadruple vector speed for long vectors. Although it was necessary to improve vector speed, it was decided that it was more important to improve the scalar speed and thus system throughput for a wider range of applications.

The company's belief that parallel processors would be the trend for future machine architectures also came into play. Therefore, a deliberate decision was made to pursue a multiprocessor design on the X-MP instead of multiple vector units. Vector performance was increased through other means as described later in this article.

Performance improvements

As indicated earlier, the CRAY X-MP incorporates many design innovations that set it apart from other supercomputers. I/O demands have been addressed with the development of the Solid-state Storage Device. Disk technology has been advanced with powerful new disk drives, and the I/O Subsystem remains a powerful integral complement.

In this section we look at two other important design features that are unique to the CRAY X-MP — the system's handling of scalar and vector operations and its multiprocessing communications mechanisms.

Scalar and vector operations

The scalar performance of each processor is improved through a faster machine clock, shorter memory access time, larger instruction buffers (twice that of the CRAY-1 per processor), multiple data paths and above all, multiple processors.

The vector performance of each processor is improved through a faster machine clock, parallel memory ports and a hardware-automatic flexible chaining feature. The X-MP design allows simultaneous memory fetches, a sequence of computations and memory stores in a series of related vector operations.

Figure 1 illustrates the benefits of these features for a common vector computation in linear algebra. On the CRAY-1, this computation takes three chained operation times, also known as chimes. The chimes consist of (load), (load,*,+) and (store). The single port to memory on the CRAY-1 prevented any further overlap of operations. The compiler had to ensure that the multiply and add instructions were issued at the proper time to catch the fixed chain. Each operation proceeded at a rate dictated by the CRAY-1 clock period of 12.5 nanoseconds (nsec).

On the CRAY X-MP, this computation takes only one chime. All operations are pipelined and proceed at a faster clock period rate (9.5 nsec). Additionally, the compiler has greater freedom in scheduling these and other supporting instructions since there is no fixed chain slot time. The elimination of fixed chain slot time and concurrent bi-directional memory access

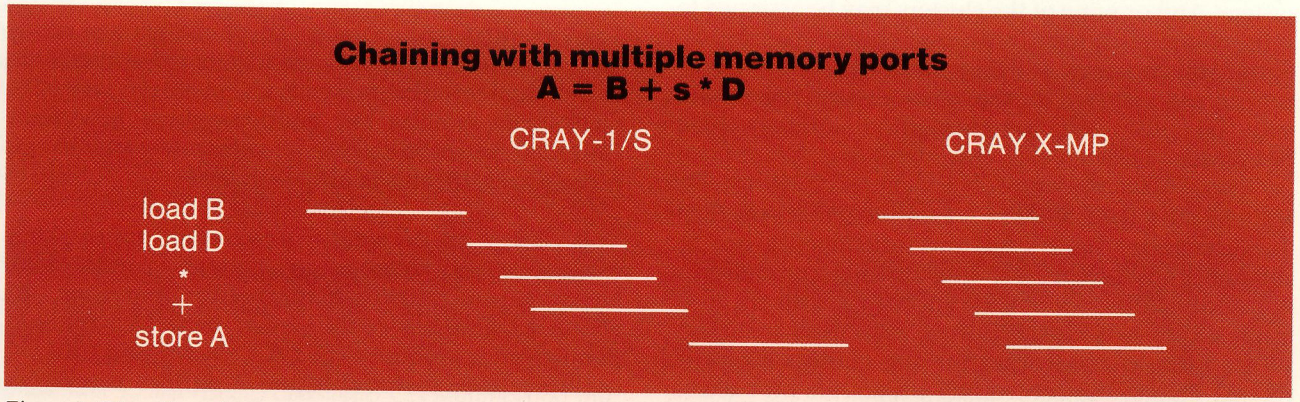

Figure 1. Time lines for vector computations

make the CRAY X-MP more amenable to the FORTRAN environment than its predecessor. As a result, the processor design provides a higher deliverable speed when processing in a general purpose environment.

Additional new features support vector indirect addressing and vector conditional executions. The hardware gather/scatter unit and hardware support for compress/expand enable vectorization of sparse matrices and certain FORTRAN loops.

Another distinction of the CRAY machine family is that the vector and scalar units and controls are so intimately integrated that, physically, there are no clear scalar or vector processor sections. This design philosophy requires tighter packaging and allows for shorter vector pipe startup time and faster data flow between scalar and vector units, as required in the execution of typical user codes with interspersed scalar and vector code segments.

Processor communication and multiprocessing

Orchestrating multiple processors to work in concert was a major design challenge. Without a mechanism enabling communication between processors, the potential of multitasking would be lost.

On the CRAY X-MP multiprocessors, processors operate independently of one another and may execute different jobs simultaneously. For the processors to communicate with one another efficiently in a shared central memory environment, new hardware mechanisms called clusters are provided for interprocessor communication and control. The cluster system enables very low communication overhead for communication among processors without disrupting operations and computations executing in other processors.

Each cluster consists of eight 32-bit shared B registers, eight 64-bit shared T registers and thirty-two 1-bit semaphore registers. The two-processor X-MP has three clusters, the four-processor X-MP has five. One cluster is typically reserved for operating system use and the others are available for user jobs.

Assigning clusters is a function of the operating system. Any or none of the clusters may be assigned to a processor. However, only one cluster may be assigned to a processor at a time. When two or more processors are assigned to the same cluster, the cluster may be used for communication among those processors.

Because the CRAY X-MP/1 is a uniprocessor system, it does not have clusters but the functionality can be simulated in memory for multitasking communication. On multiprocessor CRAY X-MPs, hardware arbitrates access to a cluster during each machine cycle. When there is contention for the cluster, access is rotated among the contending processors on successive cycles.

Clusters can be involved in several types of operations as follows:

- □ The clustering of a processor may be interrogated.
- □ 24-bit values may be transferred between A registers and shared B registers.
- □ 64-bit values may be transferred between S registers and shared T registers.
- □ A transfer may be made between the high-order 32 bits of an S register and all 32 semaphore bits simultaneously.
- □ Single semaphores may be set or cleared.
- □ A processor may wait for a single semaphore bit to clear and then set the bit.

This last operation, the wait and set function on semaphore registers, is a mechanism that includes hardware interlock. This interlock prevents a simultaneous wait and set operation on the same semaphore register by more than one processor. This basic control operation can be used to implement any of the other common software synchronization mechanisms.

This 'busy wait' operation has two significant advantages over most other hardware interprocessor exclusive access controls, such as test and set or compare and swap. First, it does not reference memory. Spin wait sampling of the same location in memory is avoided and hence disruption of memory accesses by other processors is eliminated. Second, the hardware can detect whether a processor is waiting. As a consequence, the waiting processor can be selected for interruption to do other useful work such as processing I/O interrupts.

Furthermore, a primary deadlock condition is detectable by the hardware. When all processors assigned to a particular cluster are waiting on semaphores, that group of processors is deadlocked and the hardware issues a deadlock interrupt to all of them. The deadlock can then be resolved through the governing system software.

Deadlock interrupt is useful not only in resolving a deadlock situation, but also in scheduling work. For example, when there are more schedulable units (tasks) to be executed for a user code than there are physical processors, ready tasks can be put into execution when the deadlock interrupt is detected.

The interprocessor communication mechanism allows the processors to send and acknowledge messages, and to synchronize activities in a timely way. The mechanism enables the multiprocessors to execute simultaneously the tasks of a single user program in a coordinated manner (multitasking).

This dimension of parallel processing, multitasking, extends beyond vector processing and is considered to be higher level parallelism than vector processing. The ability to simultaneously exploit these two levels of parallelism is unique to the CRAY X-MP and CRAY-2 multiprocessor systems. Figure 2 clarifies this multilevel parallelism.

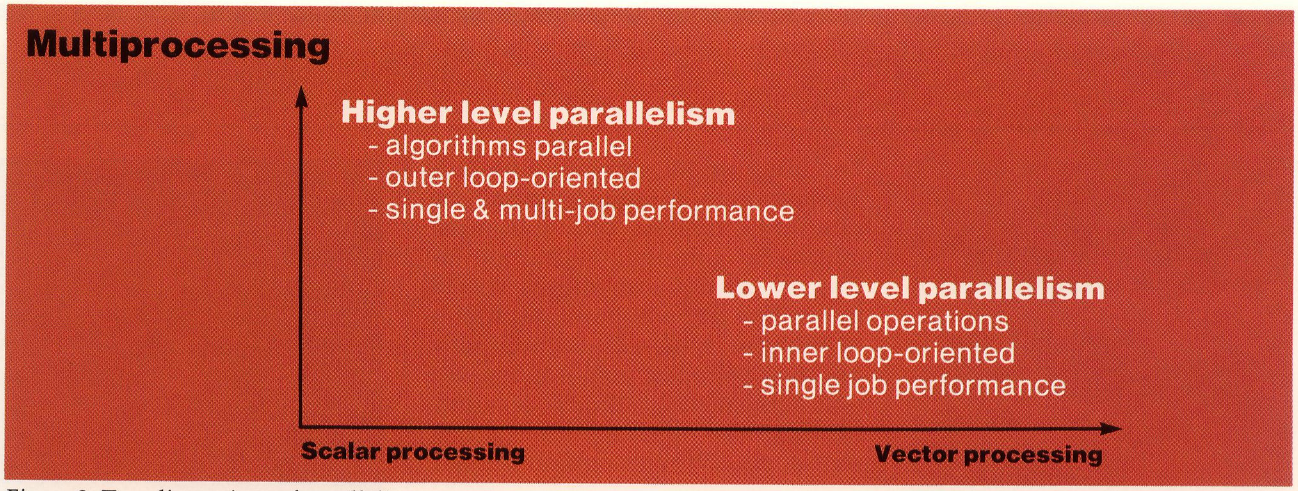

Figure 2. Two dimensions of parallelism

The clustering concept makes the CRAY X-MP a flexible and powerful computer that allows for very small granularity (in the order of microseconds) and coarse granularity multiprocessing environments. The multitasking software designed with built-in queuing capabilities provides a very flexible environment for multiprocessing with coarse granularity. At the same time, fine multiprocessing granularity can be realized by using hardware features directly.

A new era

The hardware and software features described above, along with many others, make CRAY X-MP computers innovative high-performers practical for the real-world user. Software is a key element in delivering the hardware capabilities to the user. A variety of operating systems have been adapted to the hardware and the X-MP accepts user codes written in several languages — FORTRAN, C, PASCAL. A variety of CRAY-developed and private vendor application software are available for use on the X-MP. Recent successes in using the multiprocessors in solving real-life large-scale scientific and engineering problems have ushered in a new era of high-speed computing. □

Note

This article was excerpted from an original article entitled "CRAY X-MP: A Multiprocessor Supercomputer," written by Steve S. Chen, Christopher C. Hsiung, John L. Larson and Eugene R. Somdahl. The article has been submitted for publication in the book **Vector and Parallel Processors: Architecture, Applications, and Performance Evaluation**, *Myron Ginsberg, Editor, to be published by North Holland.*

Those interested in obtaining the original version of this article can write to: Christopher C. Hsiung, Cray Research, Inc., 900 Lowater Road, Chippewa Falls, WI 54729.

References

The CRAY-1 Computer Systems, Cray Research, Inc., Pub. No. 2240008B.

Barton, John T., and Hsiung, Christopher C., "Multitasking the Code ARC3D," *Proceedings of the GAMM Workshop on the Efficient Use of Vector Computers with Emphasis on Computational Fluid Dynamics, Karlsruhe, West Germany, March 1985.*

Bucher, Ingrid Y., "The Computational Speed of Supercomputers," *Proceedings of SIGMETRICS, 1983.*

Chen, Steve S., "Large-scale and High-speed Multiprocessor System for Scientific Applications — CRAY X-MP Series," *Proceedings of NATO Advanced Research Workshop on High Speed Computation, J. Kowalik, ed., Springer Verlag, Julich, West Germany, June 1983.*

Chen, Steve, S., Dongarra, Jack J., Hsiung, Christopher, C., "Multiprocessing Linear Algebra Algorithms on the CRAY X-MP/2: Experiences with Small Granularity," *Journal of Parallel and Distributed Computing, Vol. 1, No. 1, August 1984.*

Hockney, R.W., and Jesshope, C.R., **Parallel Computers**, *Adam Hilger Ltd., Bristol, 1981, pp. 69-95.*

Hsiung, Christopher C., and Butscher, Werner, "A New Numerical Seismic 3-D Migration Model for Vector Multiprocessors", *Parallel Computing, Vol. 1, No. 2, December 1984, pp. 113-120.*

Johnson, Paul M., "An Introduction to Vector Processing," *Computer Design, February 1978, pp. 89-97.*

Kneis, Wilford, Industrial Real Time FORTRAN Standard, SIGPLAN Notices, July 1981, pp. 45-60.

Kozdrowichi, Edward, W., and Theis, Douglas J., "Second Generation of Vector Supercomputers," *IEEE Computer, Vol 14, No. 11, November 1980, pp. 71-83.*

Larson, John L., "Multitasking on the CRAY X-MP/2 Multiprocessor," *IEEE Computer, Vol. 17, No. 7, July 1984.*

Srini, Vason P., and Asenjo, Jorge F., "Analysis of CRAY-1/S Architecture," *Proc. of the 10th Annual International Symp. on Computer Architecture, IEEE & ACM, 1983, pp. 194-206.*

Worlton, Jack, "The Philosophy Behind the Machines," *Conference on High-Speed Computing, Glenden Beach, Ore., 1981, sponsored by Los Alamos and Lawrence Livermore National Laboratories.*

Worlton, Jack, "Understanding Supercomputer Benchmarks," *Los Alamos Internal Report, March 1984.*

CORPORATE REGISTER

Cray scholarships, grants address science education needs

High-quality education for computer scientists, electrical engineers and other science professionals is crucial to computer industry growth. Unfortunately, the computer industry's own appeal is so great that few scientists in computer-related fields consider becoming professional educators. As a result, science and engineering faculties throughout the United States remain understaffed. This problem is one of several that Cray Research is addressing through its selective grant and scholarship programs.

Twice each year, Cray Research commits funds to several efforts aimed at improving the quality and accessability of math and science education. "We're concerned about the quantity and quality of both students and faculty at all educational levels," said Bill Linder-Scholer, coordinator of Cray's grant and scholarship giving programs. "Our programs address three priority areas: faculty development, undergraduate scholarships and career awareness, and innovation and excellence in math, science and engineering."

In addressing the first priority area, Cray earlier this year awarded a fellowship of $52,000 to the American Electronics Association (AEA) faculty development program. The fellowship will be used to fund one electrical engineering or computer science graduate student for four years at the University of Minnesota. The money will be divided into two equal parts, with one half awarded as a grant, the other as a loan that will convert to a grant if the student, after attaining a doctorate degree, teaches at the university level for at least three years.

"This is one way we hope to address the shortage of new doctorate holders going into teaching," explained Linder-Scholer. "Last year in the United States the number of faculty vacancies in electrical engineering was between three and four times the number of electrical engineering doctorates produced." He went on to say that the situation is similar in computer science and computer engineering. "We want to encourage new doctorate holders to get a taste of teaching, and we hope some of them will make a career of it," Linder-Scholer said.

Another program addressing teacher shortages is aimed at the needs of secondary education. Cray Research and Breck school, a private high school in Minneapolis, have joined in a program to encourage science graduates to teach at the school. Cray is offering a stipend equal to about one third of a starting teacher's annual salary and summer employment at Cray as incentives. The teaching fellowship requires a three-year commitment to begin in the fall of 1985.

One awardee was chosen earlier this year by Breck. Stephanie Knutson, a recent graduate in mathematics from Saint Olaf College in Northfield, Minnesota, will teach geometry to freshman and sophomore classes. "I applied for work at several computer companies and also considered teaching," she said. "The Cray fellowship made the teaching job at Breck much more attractive. I'm also looking forward to working with software development at Cray next summer."

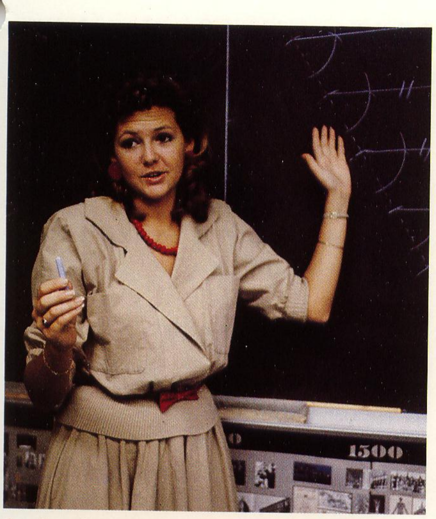

Stephanie Knutson

A discouragement to college science graduates who might otherwise consider teaching is the need for additional education coursework to receive state certification. "Since Breck is a private school, teachers needn't be state certified," explained Linder-Scholer. "If the program works well there, we might try to design a similar one for public schools."

A second priority guiding Cray's grant-making decisions is the need for increased career awareness in the community. To address this need, Cray presented a $10,000 grant to the Minnesota Alliance for Science, an education, business and government coalition promoting science education in Minnesota. The Alliance had previously received a grant from a local foundation to draft a long-term science and math primary and secondary education plan. The money from Cray was used to fund a series of conferences in outlying regions of the state to discuss the plan with parents and educators.

"Half of the public school students in Minnesota live outside the Twin Cities metropolitan area," explained Fran Jackson, assistant director of the Alliance. "We felt it was important to inform families and schools about science education issues in the state. Because decisions are made at the state Department of Education, people in outlying areas often feel left out."

A third priority for Cray's grant and scholarship programs is encouraging innovation and excellence in math, science and engineering education. Efforts in this area focused on two programs during the first half of this year: the Presidential Young Investigators Awards program and the Talented Youth Mathematics Programs at the University of Minnesota and the University of Wisconsin, Eau Claire.

For its part in the Presidential Young Investigators (PYI) program, Cray has taken up a National Science Foundation (NSF) offer to match grants given to selected scientists. Eligible scientists are pre-tenure, junior faculty members in science and engineering already awarded $25,000 PYI grants by NSF. Cray is giving $5000 to each of ten PYI awardees to be matched dollar-for-dollar by NSF, in addition to the original $25,000 PYI grant.

"Ten candidates were chosen whose work had some relationship to supercomputing," explained Linder-Scholer. "This is an important national program and we're glad to participate. Research support helps professors get tenure and encourages them to continue teaching."

The Talented Youth Mathematics Programs are accelerated mathematics training grounds for gifted junior high school students. "Cray is a major contributor," said Linder-Scholer. "The program packs four years of high school math into two. By the time participating students are sophomores, they're ready for college calculus."

Cray Research's financial support of mathematics, science and engineering education will approach $600,000 this year, with plans to expand further in the years to come. Cray understands the need and feels a responsibility to promote education in the sciences. Ensuring the growth of the computer industry begins with promoting science and math at all educational levels.

1985 Cray grants to Presidential Young Investigators awardees

Brian A. Barsky
Department of Electrical Engineering and Computer Sciences/University of California, Berkeley

Joan M. Centrella
Department of Physics and Atmospheric Science/Drexel University

Catherine French
Department of Civil and Mineral Engineering/University of Minnesota

Jorge E. Hirsch
Department of Physics/University of California, San Diego

Mark A. Horowitz
Department of Electrical Engineering/Stanford University

Paul Hudak
Department of Computer Science/Yale University

John A. Pearce
Bio-Medical Engineering Laboratory/University of Texas at Austin

Daniel Rehak
Department of Civil Engineering/Carnegie-Mellon University

M.K. Vernon
Computer Science Department/University of Wisconsin at Madison

Cray announces new orders

A flurry of orders for CRAY supercomputers has been received from Cray's European subsidiaries. The company has announced orders from England, France and, for the first time, Switzerland. The United States is represented with an order from a familiar domestic customer as well.

CORPORATE REGISTER

The European Centre for Medium Range Weather Forecasts (ECMWF) recently ordered a CRAY X-MP/48 computer system with a Solid-state Storage Device. Located near Reading, England, ECMWF is a cooperative venture of 17 European member states established in 1975 to provide medium-range (7-to 10-day) weather forecasts. The system is to be installed in the first quarter of 1986, subject to export license approval.

In July, Cray Research announced that it has been awarded a contract by British Petroleum (BP) for the installation of a CRAY X-MP/12 computer. The system will be installed at the London headquarters of British Petroleum Exploration Co. Ltd. during the fourth quarter of 1985. It will be used to support the oil and gas exploration and production activities of BP Exploration. The installation is subject to export license approval.

Cray Research announced in June that the Swiss federal government has ordered a CRAY-1 S/2000 computer system to be installed at l'Ecole Polytechnique Federale de Lausanne (EPFL) in Lausanne, Switzerland. The system will be the first CRAY installation in Switzerland and is subject to export license approval. EPFL will provide CRAY computing services for basic scientific research to professors, researchers and students at EPFL and at l'Ecole Polytechnique Federale de Zurich, another Swiss technological institute.

Electricite de France (EDF) recently ordered a CRAY X-MP/24 computer system to be installed in the third quarter of 1985, pending approval of an export license. The contract calls for an upgrade of the system in the second quarter of 1986. EDF is the national utilities company in charge of nuclear plant design and construction, electricity production and distribution. The system will be installed at EDF's research center facility in Clamart, France. The first CRAY system at EDF was installed in 1980 and was a joint venture between EDF and the Compagnie Internationale de Services en Informatiques (CISI) of Paris.

Cray Research also announced in May that Sandia National Laboratories of Albuquerque, New Mexico, ordered a CRAY X-MP/24 computer system with a Solid-state Storage Device. Sandia is a multiprogram research and development organization that operates under a no-fee, no-profit contract for the U.S. Department of Energy.

Cray ranks as a top R&D spender

It seems to happen every year, and last year was no exception. *Business Week* reported in its July 8th issue that Cray Research's expenditures for R&D were noteworthy. As a percent of sales, the company ranks number 13 among top companies in the United States, spending 16.4% of revenue on development. In real dollars that works out to $37.5 million in 1984 compared with 1983 R&D expenditures of $25.5 million. The company was listed number 15 in R&D dollars spent per employee at $17,038. As a technology-driven company, Cray Research is committed to maintaining R&D as a priority and to developing and implementing the newest computer technologies.

First foreign CRAY-2 order received

Cray Research announced in July that it has been awarded a contract by the University of Stuttgart, West Germany for the installation of a CRAY-2 computer system. This order is the first for a CRAY-2 outside of the United States.

The decision to obtain a CRAY-2 system was announced by Lothar Spaeth, head of the state of Baden-Wuerttemberg, who said the contract signed by the University of Stuttgart was subject to the approval of his state's parliament. The order is also subject to U.S. export license approval.

Ministerpresident Spaeth said that a CRAY-1/M installed at the University of Stuttgart in 1983 has also been used by the Universities of Freiburg, Constance and Tuebingen, and a number of commercial companies, and that industry usage of the computer facility is expected to increase in the future.

Cray forms new subsidiaries

In reponse to the worldwide demand for supercomputers, Cray Research, Inc. has announced the formation of two new European subsidiaries.

Cray Research S.R.L., incorporated in late 1984, recently established headquarters in Milan, Italy and is responsible for all marketing and support activities in Italy. This subsidiary was responsible for the installation of the CRAY system in Bologna, Italy for CINECA, a consortium of Italian universities.

The new Swiss subsidiary, Cray Research (Suisse) S.A. is based in Lausanne, Switzerland and was incorporated in May 1985. Its first installation will be at l'Ecole Polytechnique Federale de Lausanne, a Swiss technological institute.

Corrections

In the summer issue of CRAY CHANNELS, an article on page 27 provided a telephone number for readers interested in ordering the proceedings of the 1985 Cray Science and Engineering Symposium. The number listed was incorrect. Readers wishing to order a copy of the proceedings should contact Dennis Abraham at (612) 681-3091.

On page 14 of the same issue, we directed readers interested in the Cray Multitasking User Guide to the Cray distribution center in Mendota Heights, Minnesota. Because the Multitasking User Guide is intended only for prospective and current CRAY users, those interested in it should contact the nearest Cray region office.

APPLICATIONS IN DEPTH

Department members share knowledge

Cray Research takes a vital interest in the advancement of supercomputing applications. Recently, several members of Cray's Applications Department have been involved in organizing or otherwise contributing to professional conferences around the world. Many of the department's members work in specific application areas and are committed to their field of expertise. As a result, efforts to integrate Cray systems into specific application areas are guided by application experts rather than just computer experts. This orientation is a strength greatly appreciated by Cray customers. To give you a quick round-up of who has been doing what and where, some of the group's activities are described here.

Dr. Henry Makowitz is the session organizer and chairman of an American Nuclear Society Winter Meeting special session. The conference will be held November 10-14, 1985 in San Francisco, California. Makowitz is also contributing an invited paper on "Tightly Coupled Simulation of Nuclear Reactor Transients."

Dr. Makowitz also presented a talk in August at the Workshop/Conference of Applications of Supercomputers at Michigan Technological University. The purpose of the workshop was to expose members of the academic community to the concepts involved in designing and implementing codes on parallel processors.

Gray Lorig was asked to present a tutorial at Eurograph, the European Association for Computer Graphics this summer in Nice, France. His workshop entitled "Advanced Image Synthesis" offered information on advanced techniques and technologies in raytracing, texture mapping, shading, depth and topology.

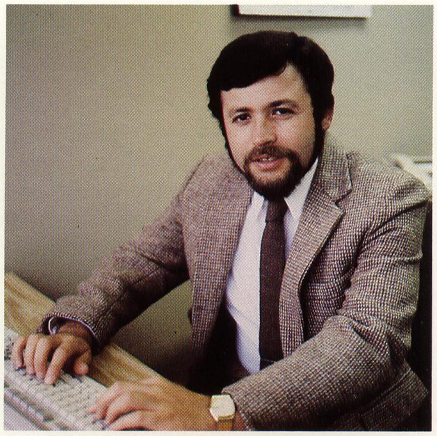

Dr. Henry Makowitz

APPLICATIONS IN DEPTH

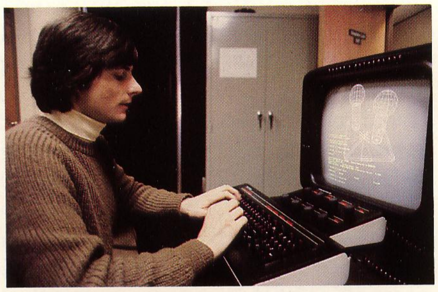

Gray Lorig

During the AIAA 7th Computational Fluid Dynamics Conference held in July in Cincinnati, Ohio, Kent Misegades of Cray Research made a special presentation on "Optimized Supercomputing." Because most speakers discussed their research, it was gratifying for Cray to have an opportunity to relate some of its contributions to the field. At the same conference Steve Chen, lead designer of the CRAY X-MP computer, led a meeting entitled, "Advanced Large-scale and High-speed Multiprocessor System for Scientific Applications."

Erich Wimmer presented three papers at a meeting of the American Physical Society earlier this year. They were "Evidence for Internal Crystallographic Distortion in GIC's (graphite intercalation compounds): C_6Li," with Swiss researchers M. Posternak and A. Baldereschi, and H.J.F. Jansen and A.J. Freeman of Northwestern University; "Energetics of Tungsten (001) and Vanadium (001): Multilayer Relaxation and Surface Energy," with C.L. Fu, A.J. Freeman, S. Ohnishi and H.J.F. Freeman, all of Northwestern University; and "Total Energy Structural and Electronic and Magnetic Structure Study of Chemisorption: c(2 x 2) Oxygen on Nickel (001)," with S.R. Chubb and A.J. Freeman of Northwestern University.

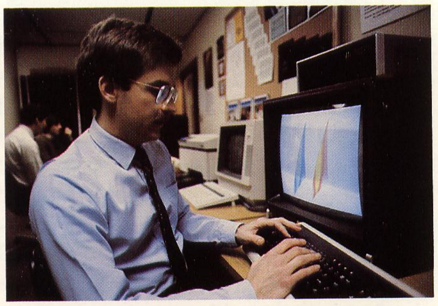

Kent Misegades

Later this fall, Wimmer will present a paper at "Supercomputer Simulation in Chemistry," a satellite symposium to the Fifth International Conference of Quantum Chemistry. His paper is entitled "Catalytic Promotion and Poisoning: All-electron Local Density Functional Calculations of Carbon Monoxide on Nickel (001) coadsorbed with Potassium and Sulfur," with C.L. Fu and A.J. Freeman.

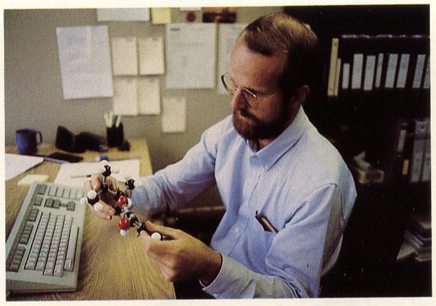

Erich Wimmer

"The topics of these papers relate to our theoretical understanding of catalysis, which is still in its infancy," Wimmer explained. "Supercomputer simulations give us, for the first time, the ability to realistically model metallic surfaces with adsorbed atoms and molecules."

KIVA models combustion cylinder dynamics

KIVA is a comprehensive code for modeling internal combustion engines. The code was written at Los Alamos National Laboratory (LANL) for the CRAY computer and makes extensive use of vectorization. The numerical model represents the spray dynamics, fluid flow, species transport, mixing, chemical reactions and accompanying heat release that occur inside the cylinder of an internal combustion engine.

The latest of several generations of engine in-cylinder dynamics modeling codes developed at LANL, KIVA embodies a number of improvements over its predecessors. The program is essentially a three-dimensional version of the earlier two-dimensional program CONCHAS-SPRAY. Aside from the dimensionality, the main differences in features between the two programs are:

- □ KIVA uses an acoustic-subcycling method for computer efficiency at low Mach number, instead of the pressure iteration used in CONCHAS-SPRAY.
- □ The simple subgrid scale turbulence model used in CONCHAS-SPRAY has been generalized to include a transport equation for the turbulent kinetic energy associated with the subgrid scale motions.
- □ The spray model used in CONCHAS-SPRAY has been augmented by a model for droplet collisions and coalescence, which have been found important in many typical applications.
- □ KIVA has been written for the CRAY computer with extensive use of vectorization.

Spatial differences are formed with respect to a generalized three-dimensional mesh of arbitrary hexahedrons whose corner locations are specified functions of time. This feature allows a Lagrangian, Eulerian, or mixed description, providing a great deal of geometric flexibility and is particularly useful for representing curved or moving boundaries. Mesh options include both planar and cylindrical coordinates for two-dimensional calculations and an unwrapped three-dimensional planar mesh. In the latter case, in general, the cells need not be square or rectangular, so many geometries can be represented by distorting the cells and deactivating other cells to form obstacles.

The capability to represent an evaporating fuel spray has been included because several engine concepts of current interest, in particular diesel and direct-injection stratified charge engines, are of the fuel-injection type.

For more information on KIVA, contact: Group T-3, Mailstop B216, Los

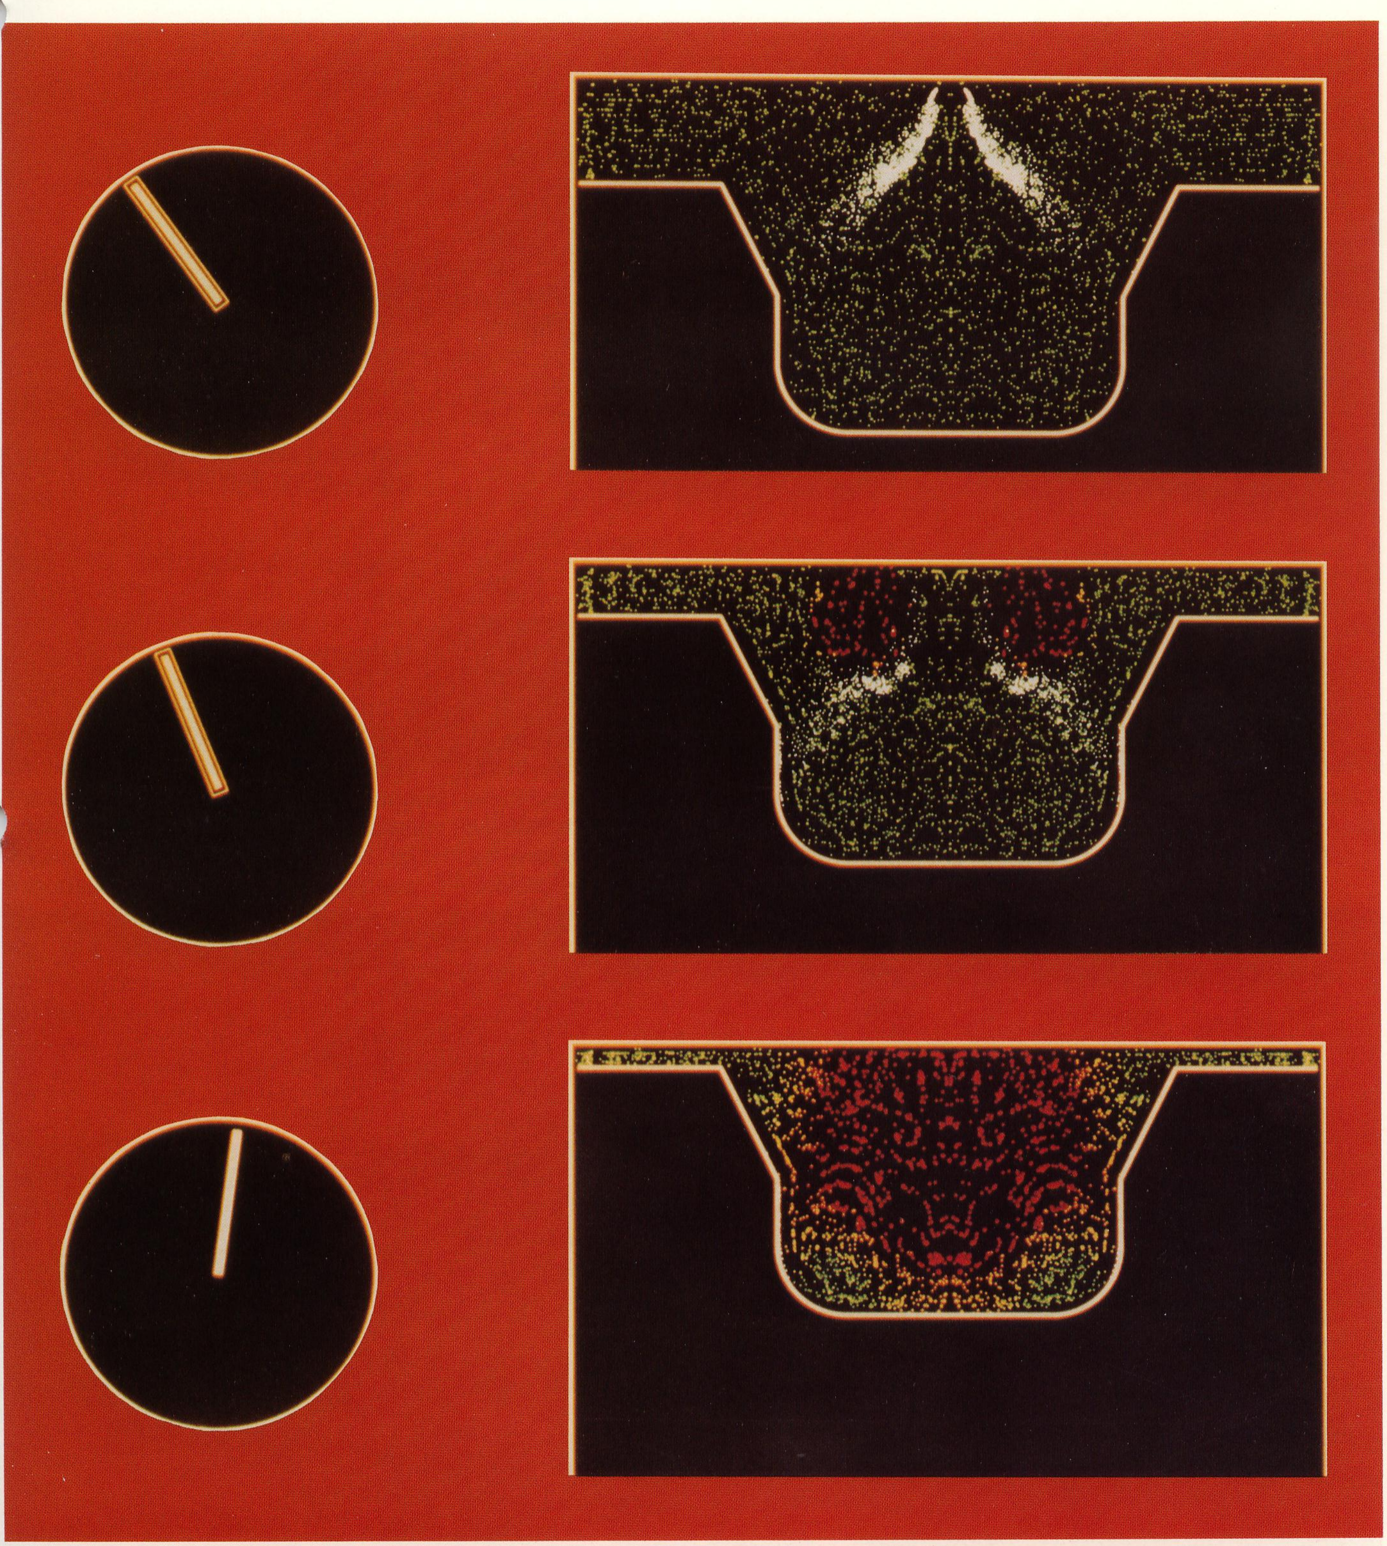

Selected frames from a CRAY-generated movie of a simulated direct-injection stratified charge combustion chamber. The top frame shows the fuel spray prior to ignition; crank angle is at 35.75 degrees BTDC (before top dead center). The second frame is early in the burn, at 22 degrees BTDC, five degrees after ignition, and shows fuel spray still unevaporated. The bottom frame, at 7.5 degrees after TDC, shows burning essentially complete in the bowl. The color of the background marker particles is keyed to temperature, ranging from pure green at temperatures below 1500 K, through yellow, to pure red at temperatures above 2000 K. Although KIVA is primarily a three-dimensional code, it provides an axisymmetric two-dimensional option for simulations such as this.

Alamos National Laboratory, Los Alamos, NM 87545.

SKS on X-MP in Mendota Heights

In late spring, Merlin Profilers' Seismic Kernel System (SKS) was installed on the CRAY X-MP/48 at Mendota Heights. SKS is a large seismic data processing package that includes a number of features particularly interesting to the seismic community. The system includes:

- □ Modern structured design techniques.
- □ Correctness proofs in key places based on the propositional and predicate calculi.
- □ True device-independent graphics based on the GKS standard, but including three-dimensionality via a polysurface primitive and a seismic raster sub-language.
- □ Modern information hiding concepts.
- □ A Problem Oriented Language (POL) based user interface. It is simple, but complete, via the Bohm and Jacopini structure theorem.
- □ Portability in excess of 99.8% using an unambiguous subset of the FORTRAN 77 standard.
- □ An intrinsic software development environment which reduces the software cycle time from months to days.

Established jointly by Merlin Profilers and Cray Research, the SKS installation at Mendota Heights provides a facility for demonstrating seismic data processing. Working together, the two staffs recently ran a wide range of seismic benchmarks on the X-MP with good results. For example, during the prestack stage, a single processor of the X-MP was performing virtually the entire prestack processing sequence (some 25 processes, including geometry, F-K and normal filtering, deconvolutions, CDP gather, miscellaneous gains, NMO, stack and four types of electrostatic raster plotting). This was done at the rate of about seven minutes per input tape. Similar gains were observed for more exotic time and depth migration processes.

Merlin Profilers (Research) Ltd., which is located in the United Kingdom, is involved in geophysical exploration and processing. Those interested in additional information about SKS and Merlin Profilers should contact: Les Hatton, Merlin Profilers (Research) Ltd., Duke House, 1 Duke Street, Woking, Surrey GU21 5BA, England; Telephone: Woking (04862) 23321; Telex: 859348 Merlin-G.

Performance improved for latest MSC/NASTRAN

MSC/NASTRAN is a large-scale general purpose program that solves a variety of engineering analysis problems by the finite element method. Newly released Version 64 was built using the latest Cray FORTRAN Compiler and runs under all versions of the Cray Operating System. Enhancements proposed by Cray's Engineering Applications Group and incorporated into this version improve performance on CRAY computers approximately 15 percent over the previous version.

Program capabilities include static and dynamic structural analysis, material and geometric nonlinearity, heat transfer, aeroelasticity, acoustics, electromagnetism and other types of field problems. MSC/NASTRAN is used extensively in diverse industries such as automotive, aerospace, civil engineering and petroleum, and in government research.

The program includes several capabilities for large problems, such as sparse matrix routines, multilevel superelements, component mode synthesis and nonlinear analysis. MSC/NASTRAN's nonlinear analysis capability is designed for efficiency in large solutions and compatibility with existing linear models. Users have the choice of several yield functions and hardening rules for modeling structural materials and may select from a number of popular solution methods or allow MSC/NASTRAN to automatically select the optimum path.

For more information on MSC/NASTRAN, contact MacNeal-Schwendler Corporation, 815 Colorado Blvd., Los Angeles, CA 90041; telephone: (213) 258-9111; telex 4720462 MSC; or the Engineering Applications Group, Cray Research, Inc., telephone (612) 681-3652.

Stress analysis module added to PATRAN

PATRAN, a finite element modeling program for engineering analysis, generates finite element meshes for use in thermal and stress analysis and aerodynamics. Version 1.6 includes a finite element analysis module that integrates directly with PATRAN, enabling engineers to model and analyze structures interactively without the need for translators and input files. Creation and submittal of an analysis using the stress analysis module and evaluation of the results are accomplished with simple menu selections.

The module's architecture allows it to run in parallel with PATRAN. Users can combine high-priority interactive modeling with analysis execution in the background, under direct user control. The stress analysis module performs linear static and normal modes analysis for isotropic and orthotropic materials. Loads include forces, temperature, displacement and pressure.

PATRAN is available on CRAY computers and offers performance improvements of 50 to 200 times over conventional superminicomputers. For more information on the Stress Analysis Module for PATRAN contact PDA Engineering, 1560 Brookhollow Drive, Santa Ana, CA 92705-5475; telephone: (714) 556-2800; telex: 683392; or the Engineering Applications Group, Cray Research, Inc., telephone (612) 681-3652.

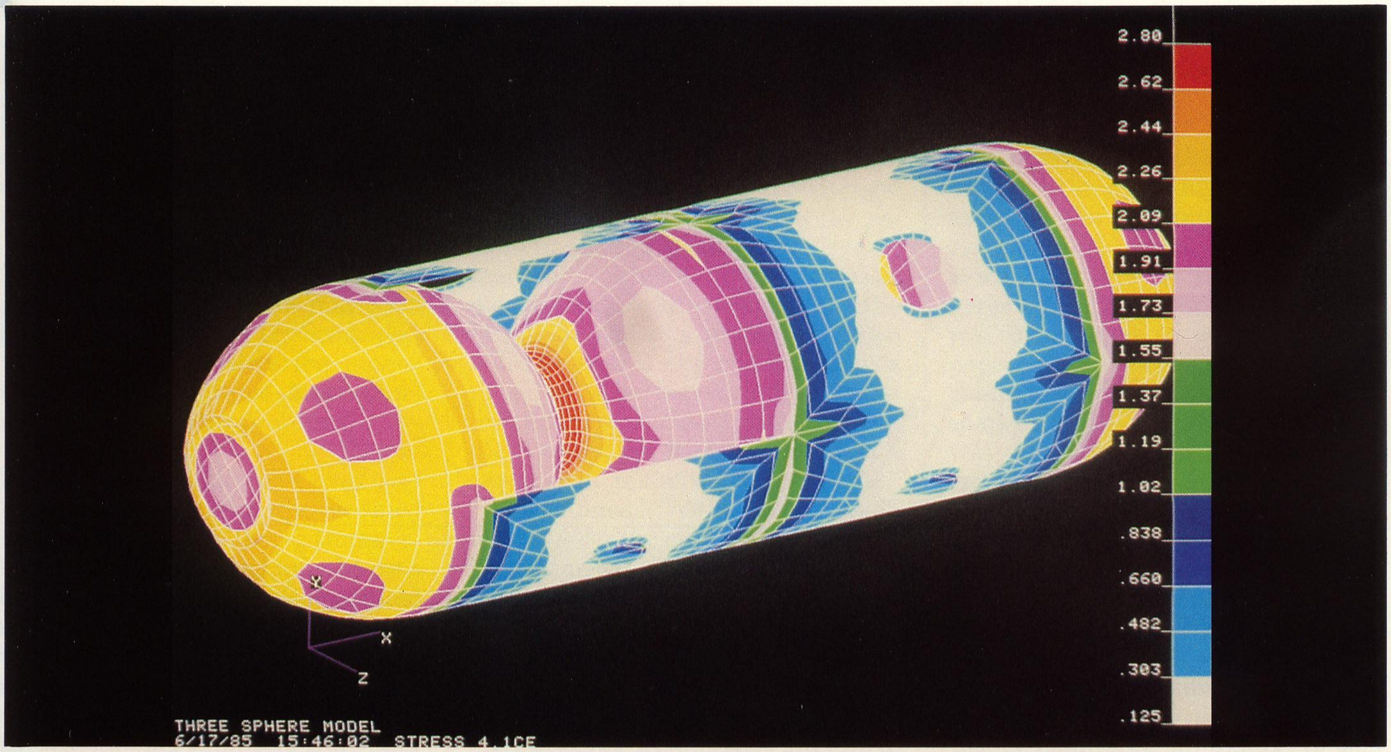

Three-sphere pressure vessel model generated with PATRAN 1.6A on a CRAY X-MP/48. The 3200-element model had 18000 degrees of freedom. Analysis was accomplished with the PATRAN STRESS module in a few minutes.

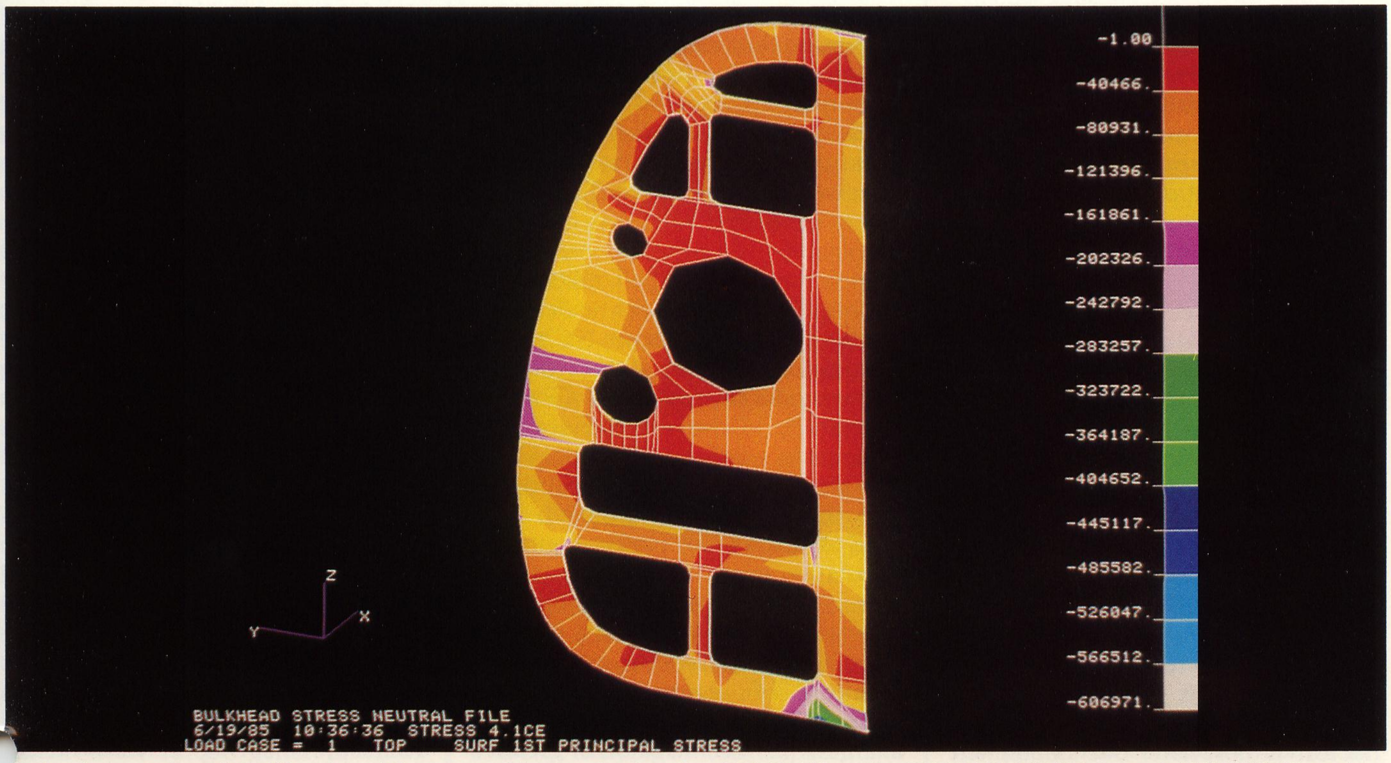

Finite element aircraft bulkhead model generated interactively using PATRAN on the CRAY X-MP/48 through VAX/VMS station. A finite element analysis using the PATRAN STRESS module was spawned from the interactive PATRAN session.

USER NEWS

Nuclear winter debate continues

Using computer simulation on high-speed machines like the CRAY, climatologists and air chemists have been able to explore the long-range global atmospheric effects of a nuclear war. Although results vary, there seems to be one general truth accepted by the scientific community: the Earth's surface temperature would be reduced if the atmosphere were laden with massive amounts of high altitude smoke and aerosols generated by the fires following a nuclear exchange. Clouds of sooty smoke could be transported long distances by the atmosphere and prevent sunlight from reaching much of the Earth's surface, scientists say.

Researchers Turco, Toon, Ackerman, Pollack and Sagan (TTAPS) garnered extensive media attention with their 1983 report and coined the phrase "nuclear winter" to describe a post-nuclear environment. Their findings indicated that continental-scale sub-freezing temperatures would ensue following a large-scale nuclear war and that this freeze might destroy any remaining human, animal and plant life. Many researchers now feel that the TTAPS report may have overstated the global effects of a post-nuclear atmosphere. But as Stephen Schneider of the National Center for Atmospheric Research (NCAR) explained, "Although their initial claims may have been overstated, the fundamental issue is real."

The TTAPS calculations involved a simple model of the energy balance within the atmosphere. The researchers used a one-dimensional, annually-averaged, radiative-convective model and raised issues that needed to be addressed by a more complex model. A three-dimensional model that included important radiative processes and also accounted for regional and seasonal variations in the weather could make more complete predictions. New research was begun where TTAPS left off.

Curt Covey, now at the University of Miami, along with Schneider and Starley Thompson, also an NCAR climatologist, modified a three-dimensional model for use on the CRAY-1. Specifically, Covey, Schneider and Thompson wanted to investigate the effects of massive smoke injections into the atmosphere. In 1984, they used a CRAY to perform general circulation model (GCM) simulations of nuclear smoke effects.

The GCM is a three-dimensional calculation of regional and global climatic effects of smoke generated by a large-scale (100 megaton or greater) nuclear exchange. The results suggested that circulation changes caused by aerosol-induced atmospheric radiative heating could spread the aerosols well beyond the altitude and latitude zones in which the smoke was initially generated.

The GCM used at NCAR is based on the Community Climate Model, a model used for other atmospheric research. It uses the spectral transform technique for horizontal discretization with rhomboidal truncation, corresponding to a horizontal resolution of 4.5 degrees in latitude and 7.5 degrees in longitude. Atmospheric winds, temperature and moisture are computed based on algorithms describing standard conservation laws. Essential to the model are semi-empirical parameterizations that govern phenomena too small in scale to be resolved and computed by the model. The parameterizations include interactive clouds that form and disperse as determined by relative humidity and convective activity. Also included are the absorption of sunlight within the atmosphere by ozone, water vapor and smoke.

Covey and his colleagues found that the zonally-averaged meridional circulation of the atmosphere is greatly affected by the smoke perturbation, in both summer and winter Northern Hemisphere scenarios. In a Northern Hemisphere spring scenario, normal rising motion near the equator and sinking in the sub-tropics is replaced by a massive circulation upwards and away from the smoke area. Similar results were obtained in the summer scenario. The modified circulation would probably transport smoke and aerosol particles upwards and southwards, spreading them to non-combatant regions.

"Moving the smoke around has proved to be bad news for non-combatant countries," said Dr. Schneider. "This result emphasizes the importance of using three-dimensional models to address the problem."

Although the studies of Covey and his associates significantly modified the TTAPS findings, they qualitatively agreed with the fundamental conclusion of the one-dimensional model: smoke generated by a nuclear war could lead to dramatic reductions in land surface temperature. Transient patches of sooty smoke would be lifted into the stratosphere by their own heat and transported out of the war zone. More specifically, the three-dimensional model suggested the possibility of the rapid freezing of land surfaces in the Northern Hemisphere.

Nuclear winter modeling continues to be refined by Drs. Covey, Schneider and Thompson, and also by researchers at the Lawrence Livermore and Los Alamos National Laboratories using CRAY X-MPs. Their efforts have stimulated genuine scientific and political concern about this possible threat to humanity. Further improvements to the model include the results of a study by Dr. Thompson involving smoke movement and scattering of sunlight. Researchers at Los Alamos National Laboratory are refining their model to include rain removal of smoke and smoke motion. Improvement of the three-dimensional model means improved studies will be possible in areas other than nuclear winter modeling. Studies of the greenhouse effect theory, ice ages and prehistoric climates will benefit from future generations of three-dimensional global atmospheric models.

Still, many questions and concerns remain unanswered. Due to the nature of the research, the scientists cannot verify their results. "The largest uncertainties surrounding the climatic effects of nuclear smoke clouds don't arise from the climatic models themselves, but from the assumptions used to drive them," commented Schneider. "No climatic model can explicitly calculate the height to which a plume of smoke will rise over a city consumed by a firestorm. The biggest questions in the theory of nuclear winter occur at the local scale: that is, how much smoke will initially get how high." He concluded, "We hope the inside of a CRAY is the only place where we ever have to observe the results of these models."

Twists and scrolls

They form elegant shapes that would appear to be of primary interest to mathematicians, yet they have been found in a surprising and potentially dangerous place in the human body. Called "scroll rings," these three-dimensional waves are being depicted in great detail with the help of a CRAY computer.

Two-dimensional scroll rings can be observed in chemistry. On a thin film of Belousov-Zhabotinsky reagent, oxidation-reduction reactions progress from centers in a richly colored dye. As the oxidation reaction spirals outward, the orange medium is turned a light blue. The oxidized regions gradually recover their orange hue, only to be turned blue by the next extending wave of the spiral. This gives the perception of rotation about the center point, with a complete revolution requiring about one minute.

The spiral pattern is not limited solely to a thin layer of chemical reagent. Analogous patterns arise in heart muscle tissue, where they can be particularly dangerous if they override the heart's normal pacemaker system. In this physiological application, the study of scroll rings in their two-dimensional form is inappropriate; three-dimensional representations are needed to better understand their behavior.

One possibility would be for three-dimensional scroll rings to be simple extensions of the spirals seen in the reagent film, where waves would extend outward in the shape of a paper scroll. Instead, their shape was determined to be more like a paper scroll whose ends had been connected, forming a hollow, ring-shaped scroll. This representation was then further complicated by the fact that scroll rings behave as a grass fire would — two colliding waves annihilate each other, as though unable to penetrate and burn through each other's ashes. This phenomenon makes the transition from two-dimensional to three-dimensional representation more challenging.

Despite the complications, scroll rings do have a high degree of symmetry and could thus be drawn without the assistance of computer graphics. But what of more complex structures? Beyond simple scroll rings are *twisted* scroll rings, where the wave front is slightly twisted as it progresses through one full turn about the central axis. It is also possible to link two originating rings together, producing complicated interactions as the wave forms collide.

To visualize these structures, the aid of a computer graphics package was enlisted. At Los Alamos National Laboratories, a package called GRAFIC was used on a CRAY-1 to plot both the simple ring and its twisted counterpart. A series of logical meshes, each tying an (x,y,z) coordinate to a node in a rectangular mesh, is the input to GRAFIC.

The coordinates are used to manipulate what would otherwise be a flat, rectangular screen mesh. The mesh can be distorted into any shape by specifying new coordinates for the intersections of each of the "wires" contained on it. GRAFIC removes any

Three manifestations of scroll rings. Oxidation-reduction reactions form two-dimensional spirals in a thin layer of the Belousov-Zhabotinsky reagent (top). A three-dimensional representation of a scroll ring as plotted by GRAFIC (bottom left); the window into the center of the structure is created by coloring selected cells "invisible." A twisted scroll ring plot (bottom right) illustrates the intricacies caused by the added twist.

hidden surfaces and allows the user to specify the location of the light source and the viewing angle.

The coloring capabilities of GRAFIC are used to peer inside the structures. Each cell in the mesh can be colored independently; to remove a slice from the wave mesh, selected cells are colored "invisible." This makes it possible to view almost any cross-section or window of the structure.

A great deal of research remains to be done. As Dr. Arthur Winfree, the researcher who used the CRAY at Los Alamos, explained, "We have thus far used the CRAY for drawing the pictures. The next step is to use the machine to solve the complex mathematical equations which characterize these waves, and ultimately to solve the electrophysiological equations that need to be included. When this can be accomplished, the research can have significant impact on the study of scroll waves in the human body."

Note: Dr. Winfree is interested in contacting others who have produced three-dimensional graphic output from partial differential equations that model chemical reactions and local diffusion. Dr. Winfree can be contacted at: The Project in Nonlinear Science, B003, University of California, San Diego, La Jolla, CA 92093; telephone: (619) 452-6753.

Evolution of Tycho's remnant modeled

A mere 45 thousand trillion miles from Earth, in the constellation Casseopeia, floats the corpse of a dwarf. The "corpse" is the remnant of a white dwarf star that exploded, lighting the sky for several months in 1572 and 1573. The explosion was faithfully reported by Danish astronomer Tycho Brahe and today the remnant intrigues scientists who hope to explain an unusual shell of material that surrounds it.

Researchers at the Los Alamos National Laboratory and the University of Illinois at Urbana-Champaign are refining a computer model that simulates the behavior of exploding stars, or supernovae. The model is being run on a Los Alamos CRAY computer, among others. Preliminary results suggest that the shell of material surrounding Tycho's supernova remnant is debris shed by the dwarf star during its 100,000-year death. Periodic explosions threw matter into space, creating a ring of material called a planetary nebula.

About 400 years ago, matter dumped on the dwarf by a companion star triggered the final, fatal eruption. Based on their simulations, the scientists believe that the observed shell around the supernova remnant is the shocked planetary nebula.

"Our results contribute to a large circumstantial case for a binary system," said Eric Jones, an astrophysicist developing the code at Los Alamos. "We played parameter games, simulating shells of various sizes and densities, until we got a circumstance that looks like what's really out there. It's just circumstantial evidence — no one can yet prove the shell originated in a binary system — but the model looks promising."

The model reveals that the shock wave from the white dwarf's final explosion rolled into interstellar space and, by the 1700s, hit the shell of earlier dust and gas. "Much of the shock wave passed through the shell, but some of it recoiled and is spreading back toward the point of the explosion," Jones said. "More refined computer models should yield a lot more information, and we believe the models will confirm these findings."

The current one-dimensional model solves hydrodynamic equations used to describe nuclear explosions in the star's atmosphere. Refinements will include a more detailed and realistic stellar model as well as radiation effects and ionization.

The CRAY was needed to get a useful resolution with the current model, and will definitely be needed to run the code as it becomes more complex," Jones said. "To get better checks on the model, we will be predicting x-ray spectra for comparison with observations from satellites. The spectra calculations require integration of extensive ionization networks that will require use of a CRAY."

Along with helping to explain an esoteric interest of astrophysicists, the research has generated improved data for atomic theorists to compare with their previous approximations. "A colleague has already been stimulated by our results to generate new atomic physics data for an iron ion," Jones said.

Astrophysical objects such as stars, nebulae and galaxies may forever escape laboratory confinement and manipulation, but that doesn't mean they can't be subjected to experimentation. Supercomputers as laboratories are not limited in their applicability by the size, distance or ferocity of cosmic phenomena. Using them, today's astrophysicists can go beyond observation and conjecture to test their theories.

CRAY-laser combination reveals importance of CVD gas-phase chemistry

Chemical vapor deposition (CVD) is a process essential to the microelectronics fabrication industry. The industry uses it to produce coatings in the manufacture of integrated circuits. During the CVD process, a source gas introduced into a chamber deposits a thin film of solid material onto a heated substrate. Other applications include electro-optics, superconductors and the making of wear- and corrosion-resistant surfaces.

Although a sophisticated technological art, CVD's underlying physics and chemistry have not been well understood. To gain a better understanding of the physics and chemistry involved, scientists at Sandia National Laboratories have been modeling CVD reaction chambers on the Lab's CRAY-1/S and using laser technology to verify their results.

"The modeling was initiated in support of experimental work," said Michael Coltrin, a scientist at Sandia's Albuquerque, New Mexico facility. "Measurements of chemical species and gas temperatures in CVD chambers could be explained in several ways. Using the model, we were able to predict the chemical mechanisms responsible for the measured results."

A significant finding of the studies is that gas-phase chemistry strongly influences the deposition process. It had been generally believed that complex surface reactions were of primary importance. Coltrin's model, developed in collaboration with Robert Kee and James Miller at Sandia's Livermore, California facility, treats gas-phase chemistry and fluid mechanics in detail. "Such a treatment of the gas phase is unique in CVD models, which often concentrate on the surface chemistry and treat the gas phase very simply," Coltrin said.

The model includes fluid velocity and energy equations, species mass fractions and chemical reaction rates for 27 chemical reactions. It predicts gas-phase temperature and velocity profiles and concentration profiles above the heated substrate for 17 chemical species. The model also predicts deposition rates and deposition uniformity for varieties of temperature, pressure, carrier gas composition, flow rate and cell size.

"We originally ran the code on a conventional mainframe," said Coltrin. "But that required going to out-of-core memory. We have to load about 250,000 words. One array contains 200,000 elements. The primary advantage of running this code on the CRAY is that I can keep everything in core memory."

In a particular case, the model predicted a significant density of diatomic silicon, Si_2, a gas phase intermediate in a common CVD process. The

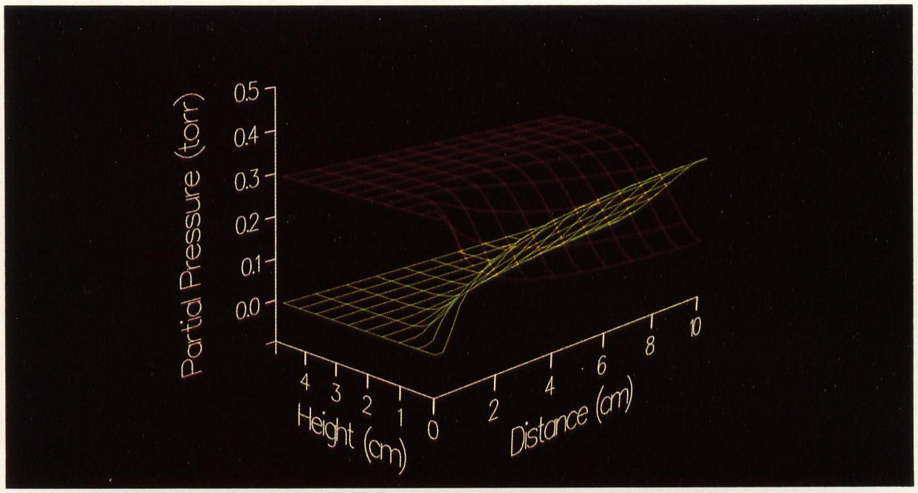

SiH_4 and H_2 profiles - predicted concentration profiles of the CVD reactant silane (red grid) and a product, hydrogen gas (green grid), in a reactor as a function of height above the heated substrate and distance along the substrate. The silane is depleted as it flows over the substrate and the loss is accompanied by the production of hydrogen.

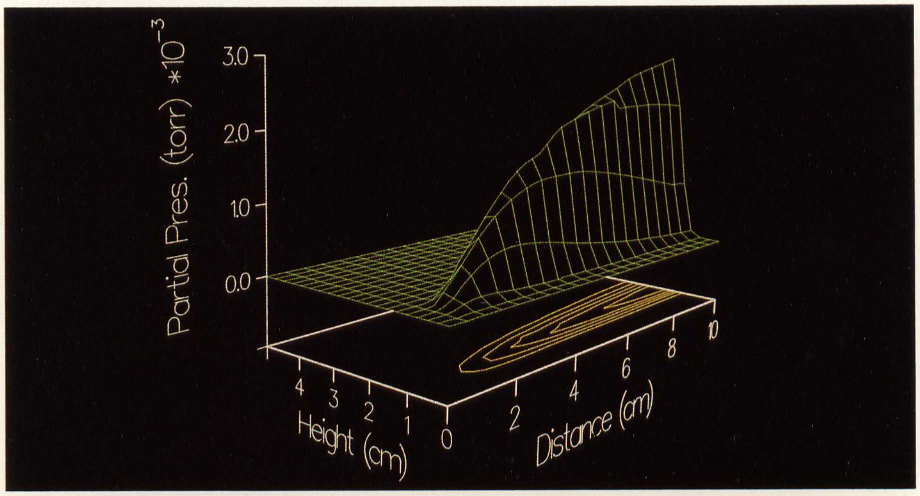

Si_2H_6 profiles - predicted concentration profile of disilane, Si_2H_6, an intermediate produced during the silane decomposition. A contour map is projected below the green grid.

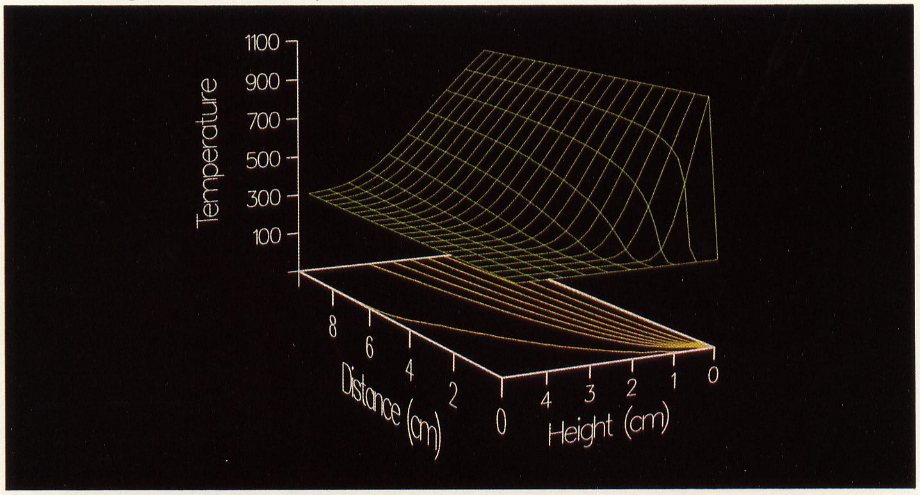

Temperature profiles - predicted gas-phase temperatures

process uses silane gas, SiH_4, to deposit a layer of solid silicon on a substrate. The hot silane gas thermally decomposes, leaving silicon behind.

The model's prediction was verified experimentally by laser-excited fluorescence, which involved focusing light from a dye laser above the substrate inside a CVD reactor, causing the gas to fluoresce. The light emitted perpendicular to the laser beam passed through a quartz window where its spectra were recorded, processed and analyzed. This technique provides a means for unequivocally identifying a given chemical species.

"Our experiments constitute the first direct observation of a reactive intermediate species in the silane CVD system. They also support strongly the importance of gas-phase chemical reactions," said William Breiland, a Sandia scientist who, along with Pauline Ho, conducted the laser-excited fluorescence experiments.

"This observation is inconsistent with many previous theoretical models for silicon deposition, which emphasized either thermodynamic equilibrium or surface kinetics," Ho said. Ho and Breiland point out that their observations are predicted by Coltrin's new theoretical model.

"In addition to the specific discoveries," Ho added, "our work shows that recent advances in computers and laser technology make it possible to investigate and understand complex systems such as CVD on a very fundamental level."

Discovering the importance of gas-phase chemistry to CVD presents the exciting possibility of custom-tailoring the process. Gas-phase chemistry, as complex as it may seem, is more readily understood and altered than the phenomena of surface chemistry and physics. By explaining the chemical "recipes" used in CVD reactors, the studies open up new possibilities for improving recipes in the future.

WORLDWIDE BUSINESS CENTERS

Domestic

Corporate Headquarters
608 Second Avenue South
Minneapolis, MN 55402
(612)333-5889

**Government Marketing/
Government Relations Office**
655 15th Street NW, Suite 210
Washington, DC 20005
(202)638-6000

Central Region Office
5330 Manhattan Circle, Suite F
Boulder, CO 80303
(303)499-3055
Albuquerque
Chicago
Detroit
Minneapolis/St. Paul

Eastern Region Office
11710 Beltsville Drive, Suite 500
Beltsville, MD 20705
(301) 595-5100
Atlanta
Boston
Laurel
Pittsburgh
Rochester
Tampa

Petroleum Region Office
5847 San Felipe, Suite 3000
Houston, TX 77057
(713) 975-8998
Dallas
Tulsa

Western Region Office
5776 Stoneridge Mall Road
The Atrium, Suite 350
Pleasanton, CA 94566
(415) 463-2800
Los Angeles
Seattle

International

Cray Canada Inc.
4141 Yonge Street, Suite 302
Toronto, Ontario
Canada M2P 2A8
(416) 229-2729

Cray Research France S.A.
7 Rue de Tilsitt
75017 Paris, France
(01)-766-01-55

Cray Research GmbH
Perhamerstrasse 31
8000 Munich 21, West Germany
(089)-56014-0

Cray Research Japan, Limited
W 1154, Shin Aoyama Building
1-1 Minami-Aoyama 1-chome
Minato-Ku, Tokyo 107, Japan
(03)-403-3471

Cray Research S.R.L.
Via Vivaio, 15
20122 Milano, Italy
(02)70.98.27

Cray Research (UK) Ltd.
Cray House
London Road, Bracknell
Berkshire RG12 2SY
United Kingdom
(0344)-485971

CRAY RESEARCH, INC.

Corporate Headquarters
Cray Research, Inc.
608 Second Avenue South
Minneapolis, MN 55402
Tel: (612) 333-5889