A CRAY RESEARCH, INC. PUBLICATION

CRAY CHANNELS

Spring 1987

Announcing new CRAY X-MP and CRAY-2 computer systems

FEATURE ARTICLES:

DEPARTMENTS:

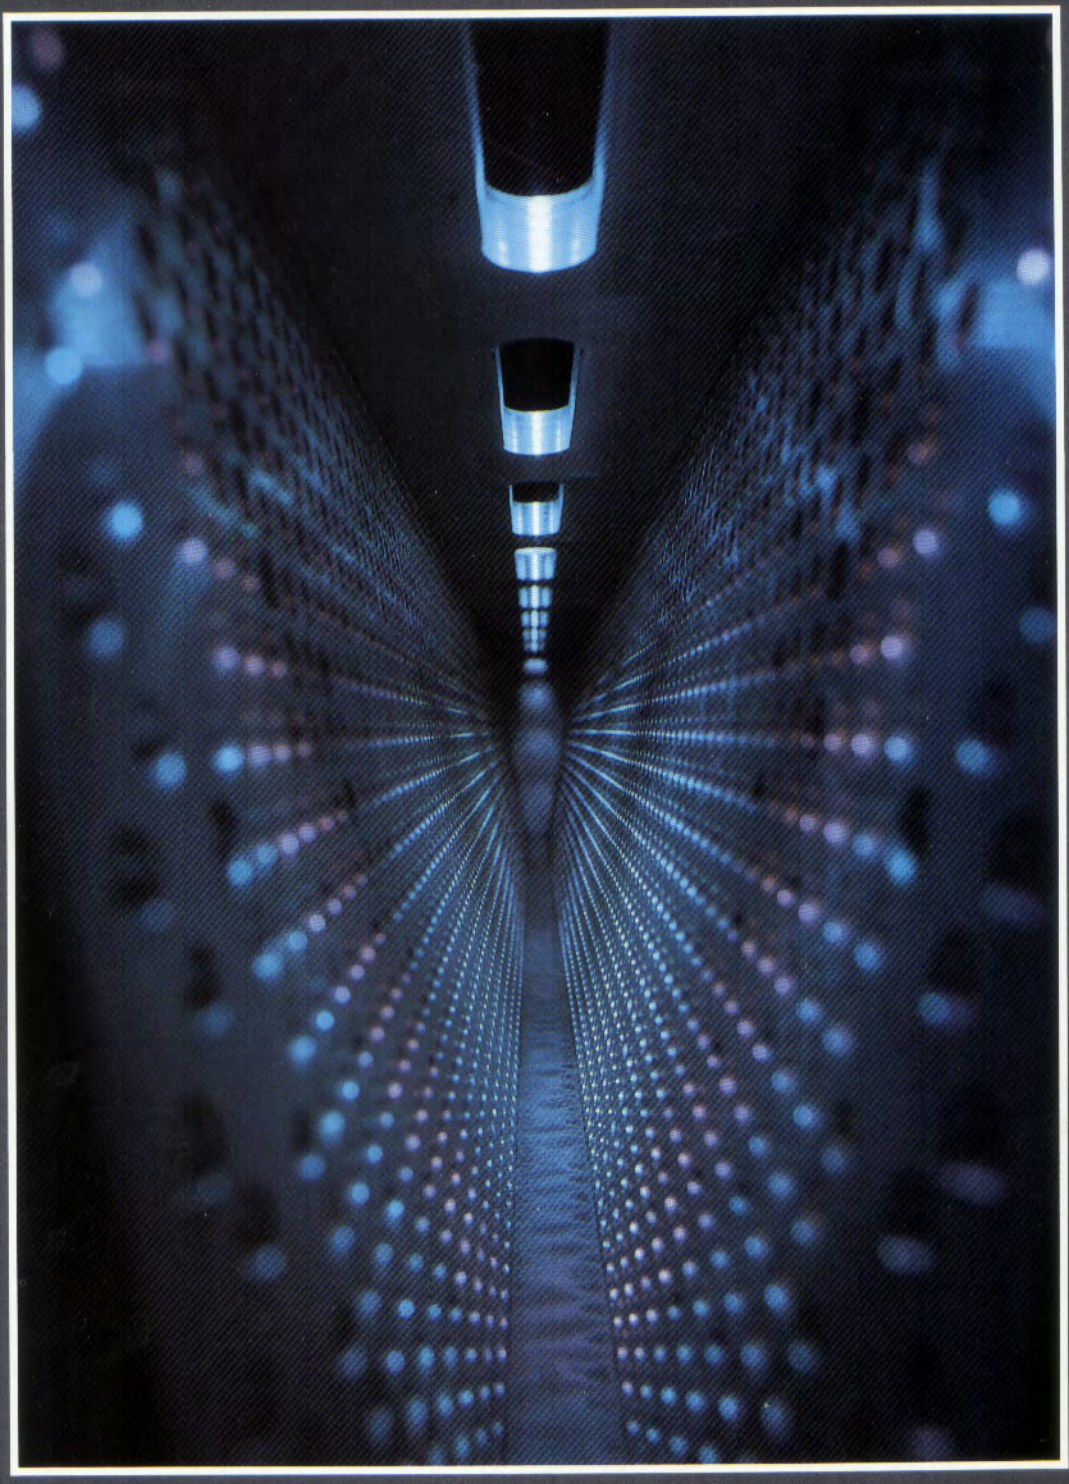

IN THIS ISSUE

Supercomputer demand has grown steadily since the first Cray system was shipped to a customer in 1976. Today, the electronics industry is becoming an important new market for Cray computer systems. The microelectronics revolution has spawned a need for powerful computing tools to handle the industry's design challenges. Cray systems currently installed in the semiconductor and computer industries give customers a competitive edge by accelerating design turnaround and making large-scale modeling practical.

This issue of CRAY CHANNELS provides examples of the advantages supercomputing brings to electronic design and engineering. In this issue we also announce the latest models of the CRAY X-MP and CRAY-2 computer systems. Both lines have been expanded to serve the broadening needs of supercomputer users. We also announce the latest releases of Cray Research's UNICOS operating system and C language compiler, and report on research in the academic community and on the high seas.

The term *cybernetics* was coined for the study of feedback among elements in a system. The electronics and computer industries have such a relationship, each feeding back its most recent advances into the design of the other. As this cycle of mutual enhancement intensifies, it benefits every area of electronic information processing. For its part, Cray Research benefits from the development of faster electronic components, just as the people who design those components benefit from the advanced computing solutions Cray systems provide.

On the cover is a look down a chassis column that will be used in a CRAY-2 computer system. The metal bar with oblong holes is a spacer inserted in the column prior to its shipment to the assembly facility. During assembly, the spacer will be removed and replaced by a bus bar. The liquid coolant that maintains a constant temperature in the CRAY-2 system will circulate through the many small holes in the column walls.

CRAY CHANNELS

A Cray Research, Inc. publication

Spring 1987

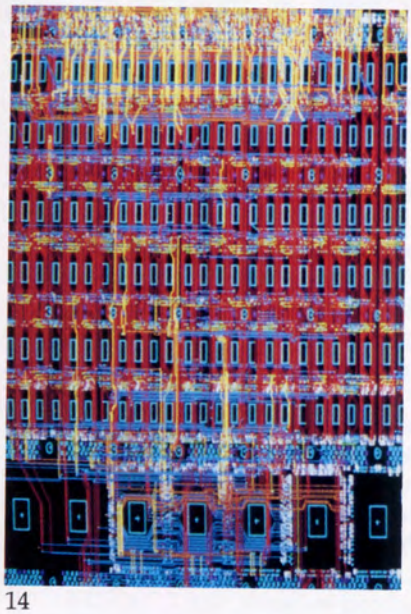

14

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 9, Number 1

Editorial Staff
Tina Bonetti
Ken Jopp
Carol Smith
Darrell Stone
John Swenson

Design
Eric Hanson
Cynthia Rykken

FEATURES

DEPARTMENTS

Cray Research product lines expand to meet user needs

Product development at Cray Research is an ongoing effort to provide better value and more options to new and current Cray supercomputer customers. Because Cray systems are called upon to solve diverse problems, adaptability and upgradability are primary design considerations.

Recent additions to Cray Research's product lines enhance their flexibility and price/performance. The CRAY-2 computer system has been extended to a line of three supercomputers. The CRAY X-MP series now offers new models and enhanced upgradability. And a new HSX high-speed external channel can now be added to CRAY-2 or CRAY X-MP systems.

The new CRAY-2 computer systems

Since its introduction, the CRAY-2 computer system has combined four processors with common memory consisting of 256 million words. The price of this original system has been reduced, and now two new CRAY-2 models make the advanced CRAY-2 technology available at a significantly lower cost. Each new version of the CRAY-2 system offers a common memory of 128 million words. One new model features four background processors, the other model features two background processors.

The same standard features are offered on all three models. The compact CRAY-2 components are immersed in a fluorocarbon liquid that dissipates the heat generated by the components. The logic and memory circuits are contained in eight-layer, three-dimensional modules. The large common memory is constructed of the most dense memory chips available, and the logic circuits are made with the fastest silicon chips available. Common memory is randomly accessible from any of the background processors and from any of the high-speed and common data channels.

Common memory is one of the more important features of the CRAY-2 systems. In conventional memory-limited computer systems, I/O wait times for large problems that use out-of-memory storage can run into hours. With the large common memory of the CRAY-2 computer systems, many of these problems can fit entirely in memory, reducing I/O wait times.

Control of network access equipment and the high-speed disk drives is integral to the CRAY-2 mainframe hardware. A single foreground processor coordinates the data flow between common memory and all external devices across four high-speed I/O channels. The synchronous operation of the foreground processor with the background processors and the external devices provides a significant increase in data throughput.

The background processors used in the CRAY-2 systems are each more powerful than a CRAY-1 computer system. The extremely fast 4.1-nanosecond clock cycle gives all three CRAY-2 models exceptional scalar and vector processing capabilities. The multiple background processors have the ability to operate independently on separate jobs or concurrently on a single problem.

All three CRAY-2 models look identical and use the existing 270°, 14-column chassis. Each of the three models occupy a mere 16 square feet of floor space with a total height of 45 inches. The new models support the same software as the original CRAY-2 system. CRAY-2 systems are not field upgradable.

For computationally intensive, large-scale applications the CRAY-2 computer systems continue to provide balanced, cost-effective performance. The new configuration options now make CRAY-2 technology available to a wider spectrum of supercomputer users.

CRAY X-MP series enhanced, new models available

The CRAY X-MP series of computer systems represents the evolution of field-proven technologies into the widest selection of supercomputers available today. CRAY X-MP systems offer powerful, cost-effective computing solutions for advanced scientific applications — for experienced supercomputer users with the most demanding computing requirements and for newer users whose research needs now require supercomputer power.

Cray Research has discontinued production of CRAY X-MP/11 and CRAY X-MP/12 systems, and has added two new models to the CRAY X-MP series. The new CRAY X-MP/116 system combines a single CPU with 16 million words of MOS memory. A new entry-level system, the CRAY X-MP/14se, combines a single CPU with four million words of central memory. In addition, all CRAY X-MP/1 systems except the CRAY X-MP/14se are now field upgradable to dual-processor systems.

The CRAY X-MP/14se computer system

The new CRAY X-MP/14se system provides users true supercomputer performance for an entry-level price. This




new model has most of the same standard features of other models in the CRAY X-MP series, and is an excellent solution for users who are just entering the world of supercomputing or are in need of a powerful processor to perform compartmentalized or distributed processing.

The new CRAY X-MP/14se system consists of a six-column chassis containing a single CRAY X-MP CPU, the I/O Subsystem, and four million 64-bit words of static MOS memory. The system occupies about 20 square feet of floor space and requires less electrical power than an upgradable CRAY X-MP/1 system. The CRAY X-MP/14se supports up to eight DD-39 disk drives and up to 16 IBM-compatible tape drives, but does not support an SSD connection.

The performance of the CRAY X-MP/14se system is approximately 80 percent that of the CRAY X-MP/14 system. All CRAY X-MP software products, including the Cray operating systems COS and UNICOS, are supported on the CRAY X-MP/14se. As with all CRAY X-MP systems, two Fortran compilers also are offered. Cray Research's CFT and CFT77 Fortran compilers fully meet ANSI 78 standards while offering a high degree of automatic scalar and vector optimization.

Installation and maintenance of the CRAY X-MP/14se system is simpler than is required for the larger CRAY X-MP systems. A prewired power distribution unit and flexible cooling system hose make installation much faster and easier. A new environmental monitor also ensures low maintenance costs for the CRAY X-MP/14se system. This feature will monitor the environment and automatically shut down the system in the event of power interruption or cooling system failure, enabling unattended operation.

Upward growth potential

The enhanced CRAY X-MP/1 systems are now upgradable to dual-processor systems in the field, significantly reducing system downtime. Because a CRAY X-MP/1 system can be easily converted to a CRAY X-MP/2 system, Cray Research customers can begin supercomputing with one processor and then add the dimensions of multiprocessing and multitasking as their needs and resources grow.

All CRAY X-MP systems are carefully balanced to deliver optimum overall performance. Fast long and short vector processing is balanced with high-speed scalar processing, and both are supported by powerful input/output capabilities. Each CPU also offers gather/scatter and compressed index vector instructions.

The Cray I/O Subsystem (IOS) is an integral part of the CRAY X-MP design and acts as a data distribution point for the mainframe. The IOS handles I/O for a variety of front-end computer systems and peripherals such as disk units and plug-compatible tape subsystems.

The DD-39 and DD-49 disk drives are high-density magnetic storage devices that complement and balance the power of CRAY X-MP systems. These disks are the fastest available, and when combined with the data handling and buffering capability of the I/O Subsystem, they provide superior I/O performance.

The optional SSD solid-state storage device allows the development of algorithms to solve larger and more sophisticated problems. It enhances CRAY X-MP performance by providing fast access to large datasets and temporary storage for system programs. The SSD functions as a very high-speed secondary memory, significantly reducing I/O wait time when compared to systems configured with conventional disk storage. The complete line of SSDs includes five models with memory sizes ranging from 32 million to over 512 million words.

HSX high-speed external channel

The new HSX high-speed channel is available for interconnection of Cray systems or for connecting external equipment to either a CRAY X-MP or CRAY-2 system. Few devices are capable of achieving the high transfer rates offered by the HSX. The channel provides full duplex point-to-point communication at rates of up to 100 Mbytes/sec over a distance of up to 50 feet. The HSX channel is ideal for high-speed graphics support and supercomputer-to-supercomputer connection.

Diversity and balance

"Cray Research continues to address the needs of an increasingly diverse supercomputer market," said John A. Rollwagen, chairman of Cray Research. "New models of the CRAY-2 computer system will make this unique technology available to a wider range of users. The upgradability of new CRAY X-MP systems will allow our customers to expand their computing solution easily as their needs grow. And the CRAY X-MP/14se system provides true supercomputer capability at an entry-level price unprecedented in our company's history."□

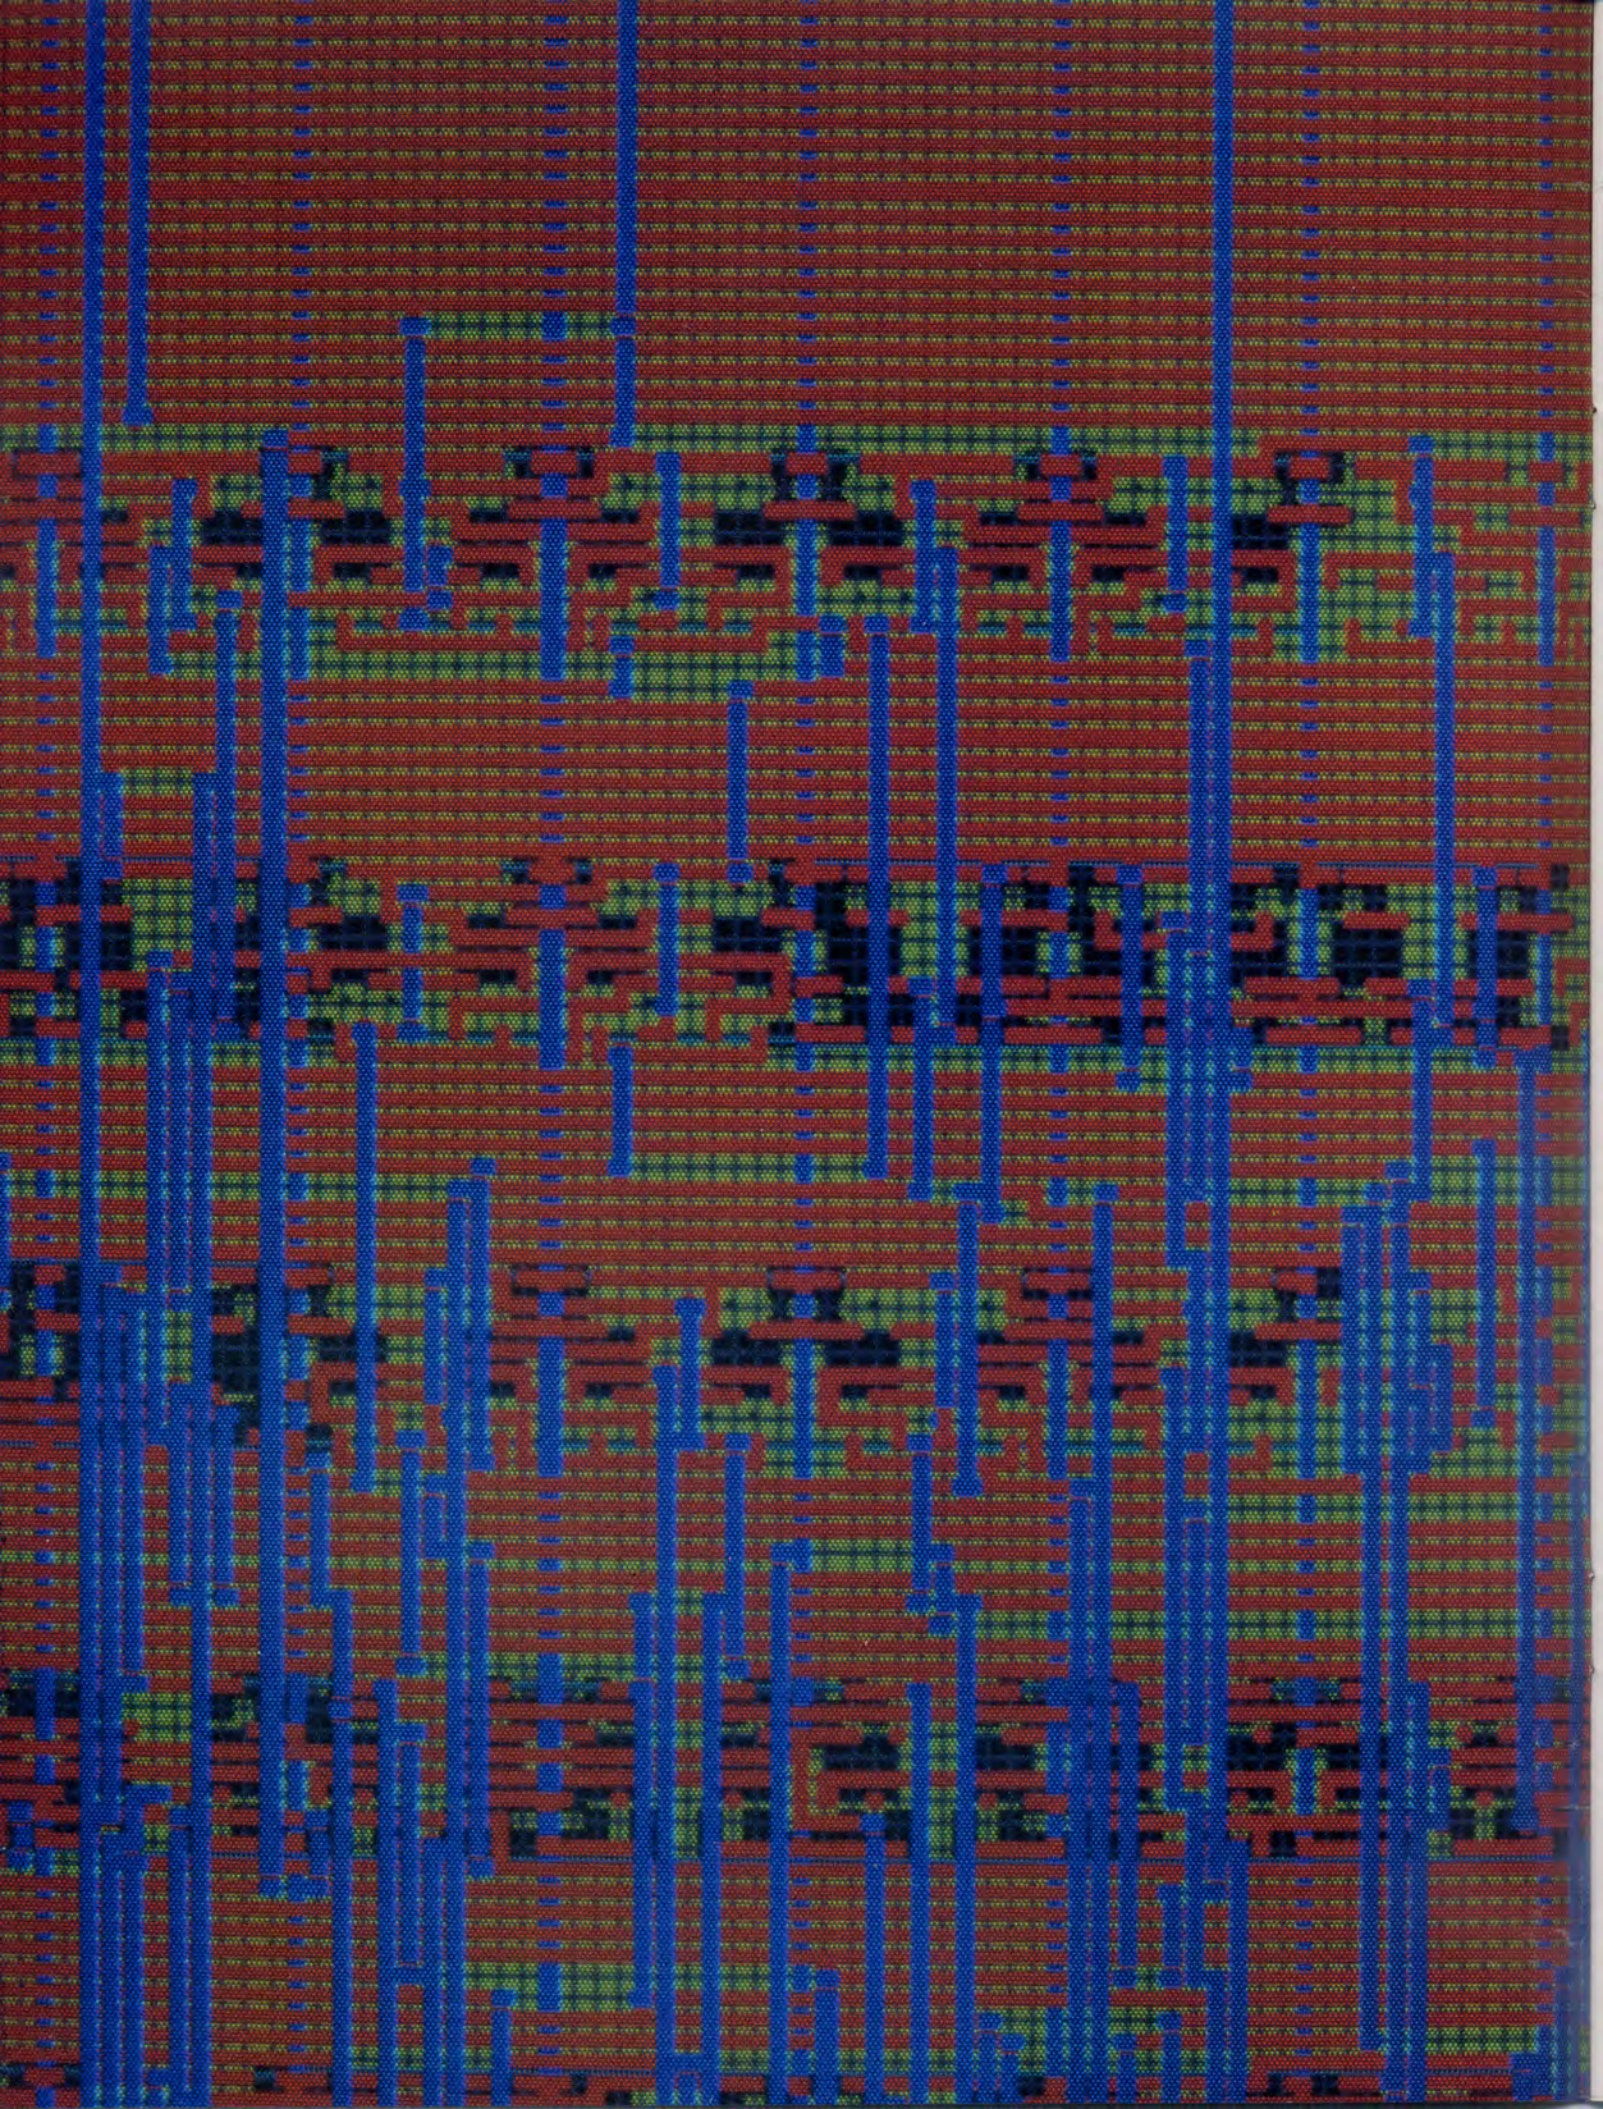

ASIC design and supercomputing at Fairchild

Rodolfo Betancourt, Carlos Dangelo, and Daniel Fabre,
Fairchild Semiconductor Corporation, Milpitas, California

Pioneering new territory is a risky venture that businesses often must undertake to stay competitive. New technologies can leapfrog tried-and-true methods of production, giving technological pioneers leverage in the marketplace. Past successes, however, offer no guarantees when deciding which new technological avenues to follow.

At Fairchild Semiconductor Corporation, a subsidiary of the Schlumberger Company, we leapt to a new design technology in 1985 with the purchase of a CRAY-1S/2000 supercomputer. The Cray system is used primarily by the gate-array division, where customers use it to prototype application-specific integrated circuits (ASICs). As the size and complexity of these devices have grown, a need has evolved for fast, large-memory computer hardware to model designs and simulate performance. Once our engineering staff had gained experience with Cray computer systems through service bureaus, we came to believe that we could turn the supercomputer into a decisive tool for semiconductor design.

Application-specific integrated circuits

At Fairchild's gate-array division, ASICs are designed using one of two design methodologies: gate-array, and more recently, standard-cells. In gate-array design, patterns of unconnected transistors are laid out in multiple arrays on a semiconductor wafer. These arrays of transistors are then connected with a user-specified pattern of conductive metal that dictates the function that the silicon will perform. In a standard-cell design methodology, small blocks that perform certain functions are designed first by Fairchild engineers. Customers then combine these blocks to create a larger block that performs a specified function.

ASIC design at Fairchild

Customers design their application-specific integrated circuits using a workstation-based schematic entry tool that allows a design to be entered into a software database. The customer then runs a proprietary logic simulator (FAIRLOGS) to verify that the design entered will in fact perform the desired function. A proprietary fault simulator (PFAULT) is run to verify design testability; that is, to ensure that the design, once manufactured, can be tested, therefore ensuring that the customer will get a product that works. The engineer then runs a placement and route program (GARDS for gate-array designs, or

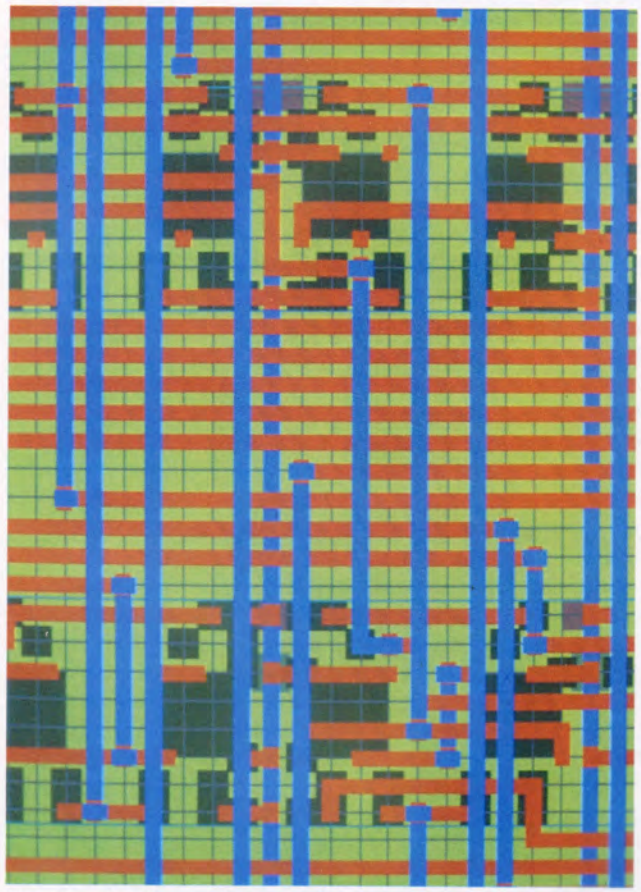

Graphic displays, above and facing page, showing placement and routing of 2-micron CMOS gate arrays designed on Fairchild's Cray computer system.

CAL-MP/TIMBERWOLF for standard-cell designs) to lay out the design onto silicon. Finally, a collection of tools are run for design-rule verification, mask data generation, and other functions. These programs are written in Fortran, C, and Pascal.

The Faircad network ties Fairchild's own design software, which resides in a VAX/VMS environment and in the on-site CRAY-1/S system, to a host of remote design centers and customer engineering workstations. A Cambridge Instruments electron-beam direct write-on-silicon system is also linked to the network, and is used for prototype manufacturing. The electron-beam system is another high-performance tool purchased by Fairchild to provide the fastest turnaround possible for prototyping customer products. CPU-intensive tasks, such as simulations, design expansions, and placement and routing, are all done on the Cray system in a way that is transparent to the designers.

We use the noninteractive Cray operating system COS. However, the ability to receive job turnaround in hours instead of days, and minutes instead of hours, makes the design process more interactive. A lack of electronic CAD software to run on COS could be a major deficiency. Currently in this field, third-party vendors offer far fewer programs that run on a Cray system than on VAX computers. This problem, however, is not crucial because most commercial electronic CAD software is designed to serve a broad market, whereas actual production work typically involves special considerations. We often have found it effective to write application-specific software for design work on the Cray system rather than rely on off-the-shelf programs. In such cases the productivity gains provided by the supercomputer justify the resources spent developing the job-specific software.

Turnaround for ASIC designs on average has been cut to one-third the time previously experienced.

Cray system meets computing needs

Fault simulation was the application that most dramatically demonstrated the need for supercomputing at Fairchild. In 1981 an example involving a relatively small circuit of 750 gates required upwards of 100 hours to run on a superminicomputer. This slow process was unacceptable, but our organization had determined that it needed a fault simulation capability. It became clear that worthwhile fault simulation was practical, however, only after we had ported our code to a Cray system at a service bureau.

To benchmark the supercomputer, we compared two simulation runs on the Cray system and on a VAX 11/780 system. In the design of a 2500-gate ECL gate array, some 5000 test vectors were applied over 525,000 timesteps. Process time on the run took about 73 seconds on the Cray system but over an hour on the VAX system.

In the second logic simulation, performed on a 6000-gate CMOS design, the 14,000 vectors used with over 700,000 timesteps ran in just 3 minutes and 46 seconds on the Cray system but required more than 2.5 hours on the VAX 11/780 system. We have observed that logic simulation programs typically will run 30 to 60 times faster on the Cray system than on a VAX 11/780 system.

The number of designs we have completed during the past year is more than twice that of the previous year.

Turnaround for ASIC designs on average has been cut to one-third the time previously experienced. At any given time we may have two or more fault simulations running on the Cray system that require two to four hours each to run. This productivity would not be possible even on hardware accelerators. Logic simulation of our most complex circuits would not have been possible without the Cray system. The development of some circuits would have taken far too long, specifically the development of an 8000-gate array and our large standard cells. Other areas of positive impact include applications of layout verification tools and, more recently, placement and routing.

With the Cray system we now routinely design digital logic circuits that comprise approximately 40,000 transistors and produce a product in about two calendar months, or four designer man-months.

The large speedup in turnaround is significant. It allows gate-array designers to simulate a design many times during a design session, thereby improving the design while reducing the total design time. Most products designed by customers are subjected to a narrow "window" of profitability; that is, a narrow timeframe during which the product is relatively free of competition. For this reason speed is essential, and the speed of the supercomputer allows us to deliver prototypes to customers faster than we could by any other means. In addition, we can solve much larger problems using the Cray system. This factor will

While we have been using the supercomputer, our revenues have tripled.

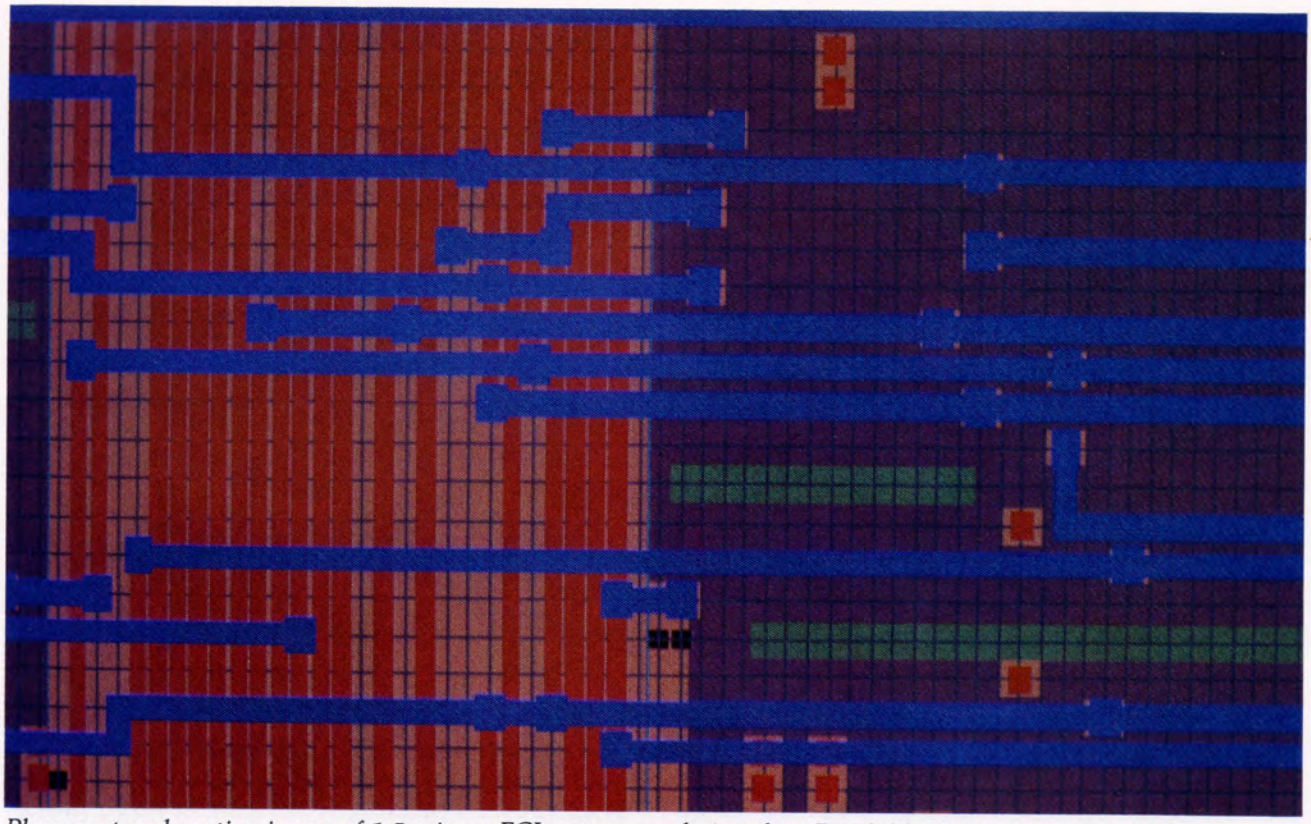

Placement and routing image of 1.5-micron ECL gate array designed on Fairchild's Cray computer system.

become increasingly important as we confront larger scales of integration in future circuits.

Achieving our goals

In 1982, when Fairchild's gate-array division was formed, our production rate was sufficient to produce a 1000-gate design in three to six months. Our goal at the time was to increase productivity by a factor of 10, which we accomplished in three years. Today we are able to design standard cells of 15,000 gates or gate arrays of 8000 gates in less than two months. The number of designs we have completed during the past year is more than twice that of the previous year. Our goal now is to increase productivity again by a factor of 10 by 1988 because we will have to enable customers to design parts using 100,000 gates in the same timeframe as before: three to six months. Our very fast computer hardware will help, though software will have to improve as well. But here, too, the power of the supercomputer has shown itself to be extremely valuable. It has helped us develop new software and test new algorithms.

The speedup that we have achieved has not yet taken full advantage of the power of the Cray system. We have converted all of the more time-consuming programs to the system, but have not yet fully optimized the majority of programs to exploit all of the Cray system's features.

During the past year, while we have been using the supercomputer, our revenues have tripled. Previously, revenues had doubled annually on average. We do not attribute these improvements exclusively to the use of the Cray computer system, but we believe we would not have been able to achieve them without having the supercomputer as an integral part of our CAD system.

Pioneering

Our decision to acquire a supercomputer at Fairchild included an element of uncertainty, but we also knew that the endeavor offered a potentially enormous payoff. As with any industry, even one as new as the semiconductor industry, certain research and production methods become established, and eventually entrenched. Fear of the unknown can be a potent deterrent to exploring new methodologies, and fear exists in some quarters that supercomputers are monsters that are impossible to understand. This attitude recalls an attitude common in the 1950s that computers themselves were mysterious machines impossible to understand.

A supercomputer requires additional hardware such as front-end systems, workstations, and, ideally, graphics terminals. But our experiences have demonstrated that, given the proper configuration, a supercomputer can be a powerful, cost-effective design tool. Ultimately, our customers are the beneficiaries of this technology, receiving a dramatically accelerated turnaround on their circuit designs. □

About the authors

The authors are employed at Fairchild Semiconductor Corporation's Gate Array Division. Rodolfo Betancourt, Ph.D., is the division's CAD software development manager. Carlos Dangelo, Ph.D., is division design automation manager. Daniel Fabre, M.S., is senior staff engineer and project manager.

At Apple a Cray keeps the doldrums away

Kent Koeninger, Apple Computer, Inc., Cupertino, California

During the past decade the design and engineering of personal computers has evolved from a garage-based, hand-tool enterprise to one that taps the computational might of supercomputers. This leap in design and engineering technology has occurred within one company, Apple Computer, Inc. In 1986 Apple purchased a CRAY X-MP/48 computer system for advanced design projects and to serve as a resource for all Apple engineers. Apple's "Cray Evangelist" Kent Koeninger here explains Apple's reasons — practical and philosophical — for purchasing a Cray supercomputer.

Our decision to acquire a Cray computer system at Apple grew out of a central driving force, which was the need to design new user interfaces. An efficient way to evaluate new designs was needed as an alternative to actually building and testing a repertoire of hardware prototypes. Modeling interfaces, however, requires a powerful computer system on which to run the models. Apple's advanced computer development group also wanted to run complete simulations of multiple-layer printed circuit (PC) boards and complex VLSI chips. Although minicomputers could have handled partial simulations, a major computing resource — a supercomputer — was needed for the exhaustive hardware simulations the group wanted to run.

The manager of advanced computer development at Apple, Sam Holland, had previously completed a survey of U.S. and Japanese supercomputers. With the results of his research in hand, management determined that a Cray system would best meet the company's needs.

The computing environment

Our use of the Cray system so far has been strictly interactive, making response time for jobs critical. Apple was the first site to run Cray Research's interactive operating system, UNICOS, on a CRAY X-MP system. Originally we had planned to run the Cray operating system COS,

and switch to UNICOS at a later time. But a last-minute decision was made to install UNICOS from the outset. This decision turned out to be a good one. UNICOS has been very stable. Cray Research has done an excellent job maintaining the flavor of AT&T's popular UNIX operating system, on which UNICOS is based.

To further enhance the interactivity of our computing environment, we currently are involved in a cooperative effort with several universities to produce software for a workstation environment that runs on a standard Apple Macintosh personal computer and interacts with the UNIX operating system. The Macintosh workstation will provide distributed editing and network communications from itself to the Cray system using the AppleTalk, Ethernet, and HYPERchannel networks. The workstation will use the TCP/IP protocol.

Our experience with the Cray system during the past year convinces us that it will significantly shorten the design cycles of future products. Already it has made possible some very creative and technically demanding progress. Two examples follow.

Significant bits

The value of the Cray system's speed was demonstrated recently by the solution of a large numerical study. Steve Perlman of Apple's advanced development group needed to determine the minimum number of significant bits required to give a type of integrated circuit (IC) acceptable precision. Preliminary numerical analyses generated four answers to the problem. Consequently, the group members chose not to trust these results. The group could have turned the problem over to expert mathematicians and waited for an eventual answer, but instead they chose to compute all possible cases and found all of the problem's boundaries by sheer brute-force computing using the Cray system.

Using dedicated time on a VAX 11/780 computer, the program would have required two weeks to run to completion. But when other users of the VAX computer were unable to use the system, Perlman lowered his program's priority status. At the lower status the program would have required six months to run to completion, so Perlman moved the program to the Cray system.

The code conversion to the Cray system took only half an hour, requiring only minor changes. An interesting modification was the removal of a "printf" statement in an inner loop that was included to monitor the progress of the program. On the VAX computer it printed about once every minute. The loop was executed so quickly on the CRAY X-MP system that the printer was unable to keep up. The print statement in the inner loop and one in the next outer loop were removed, completing the conversion to the Cray system.

As a result of moving the problem to the supercomputer, the solution was obtained in two and one-half hours. That is, a problem that would have required six months to run on a loaded VAX computer became one afternoon's work on the Cray system, including code conversion time. The practical result of the conversion was that the problem was solved 1700 times faster on the Cray system than would

have been possible on the VAX computer (180 days at 24 hours per day divided by 2.5 hours = 1728).

When using only one of the Cray system's four CPUs, the program ran 25 times faster than on an unloaded VAX 11/780 computer. When the problem was divided into sections that ran simultaneously on the Cray system's four CPUs, the run required one 1/100th the time needed on the unloaded VAX computer. These results were obtained using the first release of the Cray C compiler, which does not optimize or vectorize. Vectorization, included in the latest version of the Cray C compiler (see this issue's Corporate Register), would have generated a significantly faster running time. Perlman now plans to execute several more similar runs as a result of his first experience with the Cray system.

Disk head

A recent study of disk head configurations also benefited from the capabilities of the Cray system. Jim White and Conrad Chen in Apple's peripheral development group have been able to evaluate head configurations that they could not have evaluated without the simulation capabilities provided by the supercomputer.

The disk-head model simulates a three-dimensional, compressible, viscous, rarefied, time-dependent air flow between a disk head and a flexible medium. It solves a system of partial differential equations that describes the physics of the elastic medium, fluid flow, and rigid body dynamics. The Cray Fortran compiler CFT compiled the program in 10 seconds on the CRAY X-MP/48 system. The same compile using the F77 UNIX Fortran compiler on a VAX 11/780 computer required 30 minutes. On the Cray system, an average run of the program takes 30 minutes, whereas on a VAX computer the same run would take about one day. This reduction in turnaround time enabled White to try 20 different configurations. The fast turnaround has allowed him to be creative and exhaustive in searching for superior disk heads. Results such as these are only preliminary indications of the effect we expect the Cray system to have on productivity and turnaround.

Advanced development

General engineering needs serve as one justification for acquiring the Cray supercomputer. The primary motivation for looking into supercomputing, however, came from the needs of the advanced computer development project. This group was organized to model future Apple products in their entirety prior to building new hardware. The models will simulate all product features including IC functions, PC board layouts, the system-level architecture, heat flow, and the user interface including graphics. The models also will be used to perform "crashworthiness" studies of physical frames. As is clear from the sample of applications, our need was for a general-purpose computing system to handle design work, not hardware optimized for a particular application. If this methodology works as planned, we will know what future products look and feel like before any of their parts actually exist.

As it is for all Apple products, the user interface will be our primary concern and will act as a springboard for the products' other features. Our observation is that our competitors seem to think very little about the user interface, and add it on almost as an afterthought. We consider the user interface to be primary and we plan to design the new products around it.

Ideas for new product interfaces will be dry run by modeling a display screen in real time using the HSX channel on the Cray system. This is the fastest way we have to determine the optimal screen display for a user interface. Once that is done, the challenge for us will be to see if a machine capable of providing that interface can be built inexpensively enough to be marketable; that is, through simula-

The Mandelbrot set shown at left was computed in less than two minutes on Apple's CRAY X-MP computer system. By linking Ultra graphics terminals via Ultra Corporation's UltraBus to the Cray system through the HSX high-speed external channel, Apple engineers have achieved data display rates of more than 80 Mbytes/sec. This has allowed them to display dense color images (1280 x 1024 pixels x 24 bits of color per pixel) at a rate exceeding 16 screens per second. At right are Newt Perdue (center), vice president of Ultra Corporation, with Sam Holland (left) and Kent Koeninger (right) of Apple Computer.

The essence of speed

Apple's rapid acquisition of a supercomputer may seem miraculous to anyone familiar with the usually lengthy process. But our need for supercomputer capabilities was clear, and the lack of bureaucratic obstacles within the company hastened the process considerably. As a result, Apple boasts one of the fastest Cray system installations to date, having prepared its installation site in only six weeks.

For most of these six weeks, crews worked three shifts, 24 hours a day, seven days a week preparing a machine room for the system. The crews gutted the building, installed the cooling tower, the chilled water and high-pressure freon plumbing, the wiring, the air conditioning, the walls, and the false floor. Nothing could be tested, however, because the 480-volt transformer did not arrive until the Friday before the Cray system was shipped. Everything worked perfectly except the transformer, which had to be exchanged Friday night.

The Cray system was installed and ready to use two weeks after its arrival. The first thirty days of use constituted an acceptance period, during which the system had excellent uptime (98.6 percent). Since then, it has shown itself to be a reliable performer. The speed with which the Cray System was installed was a good omen, considering the use to which the system would be put: giving engineers and designers rapid turnaround.

tion we want to design the optimal user interface that we can sell to the public.

The emphasis in using the Cray system for advanced development is not on solving problems of physics, but rather it is on trying out many ideas to solve problems creatively. We do not want our creative people to feel compromised or inhibited by a lack of computer resources. They should be able to try something and get turnaround fast enough so that they are encouraged to try more ideas. This approach in essence redefines the whole design cycle. If 30 minutes is all the time needed to try out an idea, you don't have to think about it much — just try it. If it doesn't work, try something else.

As far as we have been able to determine, the Cray system will be adequate for simulating new products completely. If we are able to compute user interface graphic displays in real time, the HSX channel will enable us to display frames at 600 to 800 Mbits/sec. If slow time simulations are necessary, we plan to run the simulation and store the information on striped disks. In this case, the playback speed will be determined by the speed of the disk drive. Using more conventional techniques, a 5-Mbit interface for example, it would take an hour or so to produce a segment of tape a few minutes long. With the Cray system we will be able to play the tape back virtually in real time and get a real-time display of the graphics. Of course, this strategy is somewhat speculative; we really can't know for sure if it will work this way until we try it.

If all goes as currently envisioned, engineers throughout Apple will move their most demanding problems to the Cray system, motivated by faster turnaround and by the opportunity to tackle problems unsuitable for conventional mainframe computers. Also, thanks to the Cray system, members of the advanced project group will be able to use the new products they are designing before the products exist. These "products in software" should function as if they were physical objects, and any bugs in their design should be evident from diagnostic checks. To the extent that we can accurately model the products' components, we will be able to bypass hardware-prototype testing in favor of software-model testing for these products. Achieving this end will mark a milestone in the evolution of CAD applications and could set the standard for similar future design projects. □

About the author

Kent Koeninger received his B.S. degree in mathematics from California State University in 1977. Prior to joining Apple Computer in March 1986, he worked for ZeroOne Systems, Inc. as manager of the systems group at NASA Ames Research Center's Advanced Computational Facility. As "Cray Evangelist" at Apple, Koeninger is responsible for "making the Cray system the most effective tool possible for Apple." His activities include coordinating the computer network through which Apple personnel access the Cray system, assisting in converting electronic design packages, and teaching courses to Apple personnel on using the Cray system.

Experiences with SPICE-GT at Toshiba

As one of the leading semiconductor manufacturers in the world, Toshiba is greatly concerned with the quality and accuracy of its semiconductor products. Therefore, Toshiba makes extensive use of its circuit simulation program, SPICE-GT, which is based on the SPICE program from the University of California, Berkeley. SPICE-GT provides Toshiba's engineers with the ability to model and simulate a circuit, and thus determine very accurately how the circuit will behave when it is built. Because of the program's importance, as well as the tremendous amounts of CPU time required to run the application, Toshiba has optimized the program and converted the code to run on the company's CRAY X-MP/22 computer system. The optimized program offers significantly faster performance, improved convergence, more accurate physical models, and several useful interface capabilities.

Faster performance

By enhancing the basic algorithms in SPICE, Toshiba has achieved significantly faster simulation performance. For example, one of the major algorithms in SPICE uses iteration to converge to the correct result. By modifying the algorithm, SPICE-GT is able to find the result with fewer iterations. Figure 1 compares the number of iterations required by the original SPICE and Toshiba's SPICE-GT.

Improved convergence

A major problem with SPICE is the inability of the iterative algorithm to reliably converge on the correct results for complicated circuits. Toshiba solved this problem in SPICE-GT by enhancing the timestep control portions of the algorithm. Of course, if the circuit is designed improperly, convergence will not occur. In this case calculations will be stopped, and an error message will be given.

More accurate physical models

The program's physical models were improved, particularly for MOS circuits. The physical models describe how

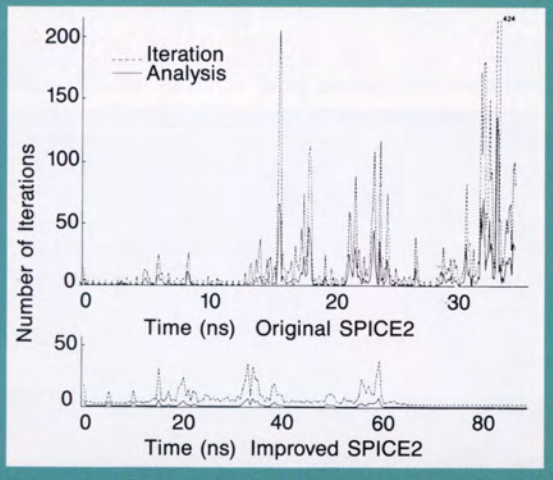

Figure 1. Comparison of iterations to calculate 1-Mbit DRAM circuits.

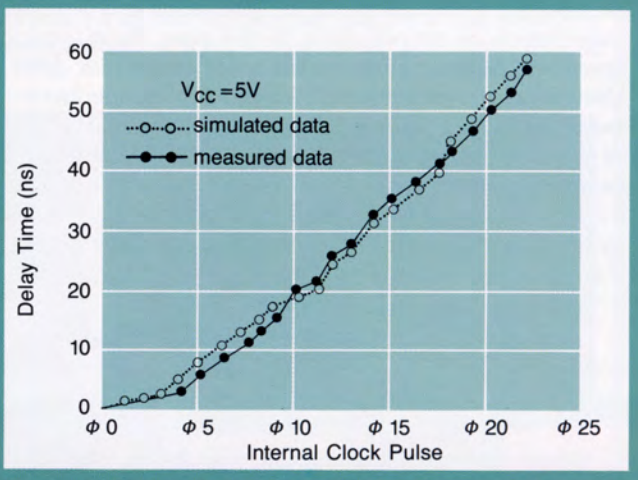

Figure 2. Comparison of simulated and measured internal clock pulses for a 1-Mbit DRAM.

each device will behave when built. SPICE-GT uses analytical models as well as table-look-up models. Toshiba was the first company in the electronics industry to put into practical use table-look-up models for MOS transistors. Figure 2 compares simulated data for the internal clock pulse for a 1-Mbit DRAM circuit and data measured by electron beam testing. The excellent agreement between simulated and measured data indicates the accuracy of the physical models in SPICE-GT.

Enhanced interfaces

To add flexibility, SPICE-GT can be used with preprocessors and postprocessors. For example, a preprocessor can be used to extract parameters and create data for table-look-up models. In addition, a postprocessor can be used to provide added information about the circuit being simulated, such as power consumption and current gradient, and to display calculation results graphically.

Supercomputing

Toshiba engineers have optimized SPICE-GT to best take advantage of their CRAY X-MP computer system. Originally, the program was run without any optimizations. Table 1 shows a typical distribution of CPU time for our most important simulation model. Examination of this data shows that the circuit matrix solver consumes the majority of CPU time and, therefore, should be the first part optimized. The circuit matrix solver consists of two parts: the first part rearranges a sparse matrix (called decomposition), and the second part solves the sparse matrix.

The algorithms in SPICE do not easily lend themselves to vectorization. In addition, when the program was written, the data structures and coding techniques used did not consider vector processing. Toshiba's CRAY X-MP/22 computer system does not have gather/scatter hardware, which would be required if vectorization was used. Toshiba engineers, therefore, chose a code-generation technique that would improve the performance of scalar operations for the solver portion.

The code-generation technique consists of a special solver in memory written in machine language. This solver uses the data structure characteristics of a SPICE-GT model, compared to the original Fortran solver, which did not. Furthermore, the code-generation solver does not use the loop control step. Therefore, the instruction buffer facility of the CRAY X-MP system works very well. In addition, the optimized code uses the instruction shuffling technique to delete unnecessary load/store instructions. Table 1 also shows the new distribution of CPU time for SPICE-GT.

Multitasking

Once the solver was optimized, engineers began to look at tuning the decomposition portion. To solve this problem, multitasking was applied to exploit natural parallelism. Table 2 shows the distribution of elapsed time for several cases.

Once gather/scatter hardware is added to the system, Toshiba engineers will begin vectorizing both the code generation and multitasking areas to further improve the simulator's performance.□

About the authors

Yoshinari Fukui received his B.S. degree in electrical engineering in 1972. His research at Toshiba has concentrated on round-off error analysis, supercomputer algorithms, and symbolic processing.

Hiroyuki Kinoshita received his B.S. degree in electronics engineering in 1971 and joined Toshiba later that year. From 1972 to 1982 Kinoshita worked as a chief design engineer on MOS dynamic memory circuits. Since 1983 he has supervised circuit simulation technology and is currently a manager in Toshiba's integrated circuit division.

solver	circuit matrix construction	circuit matrix solution	other
Fortran	21.8%	75.9%	2.3%
code-generation (SPICE-GT)	78.3%	15.8%	5.9%

Table 1. Comparison of CPU time distribution for matrix solvers.

case	elapsed time (sec)		CPU time (sec)	
	single-task	multi-task	single-task	multi-task
1	373	253	371	418
2	55	37	54	65
3	55	40	48	70
4	47	31	45	57
5	47	29	45	53
6	47	30	47	54
7	169	107	168	191
8	354	231	354	398

Table 2. Total elapsed and CPU timings for multitasked and single-CPU runs.

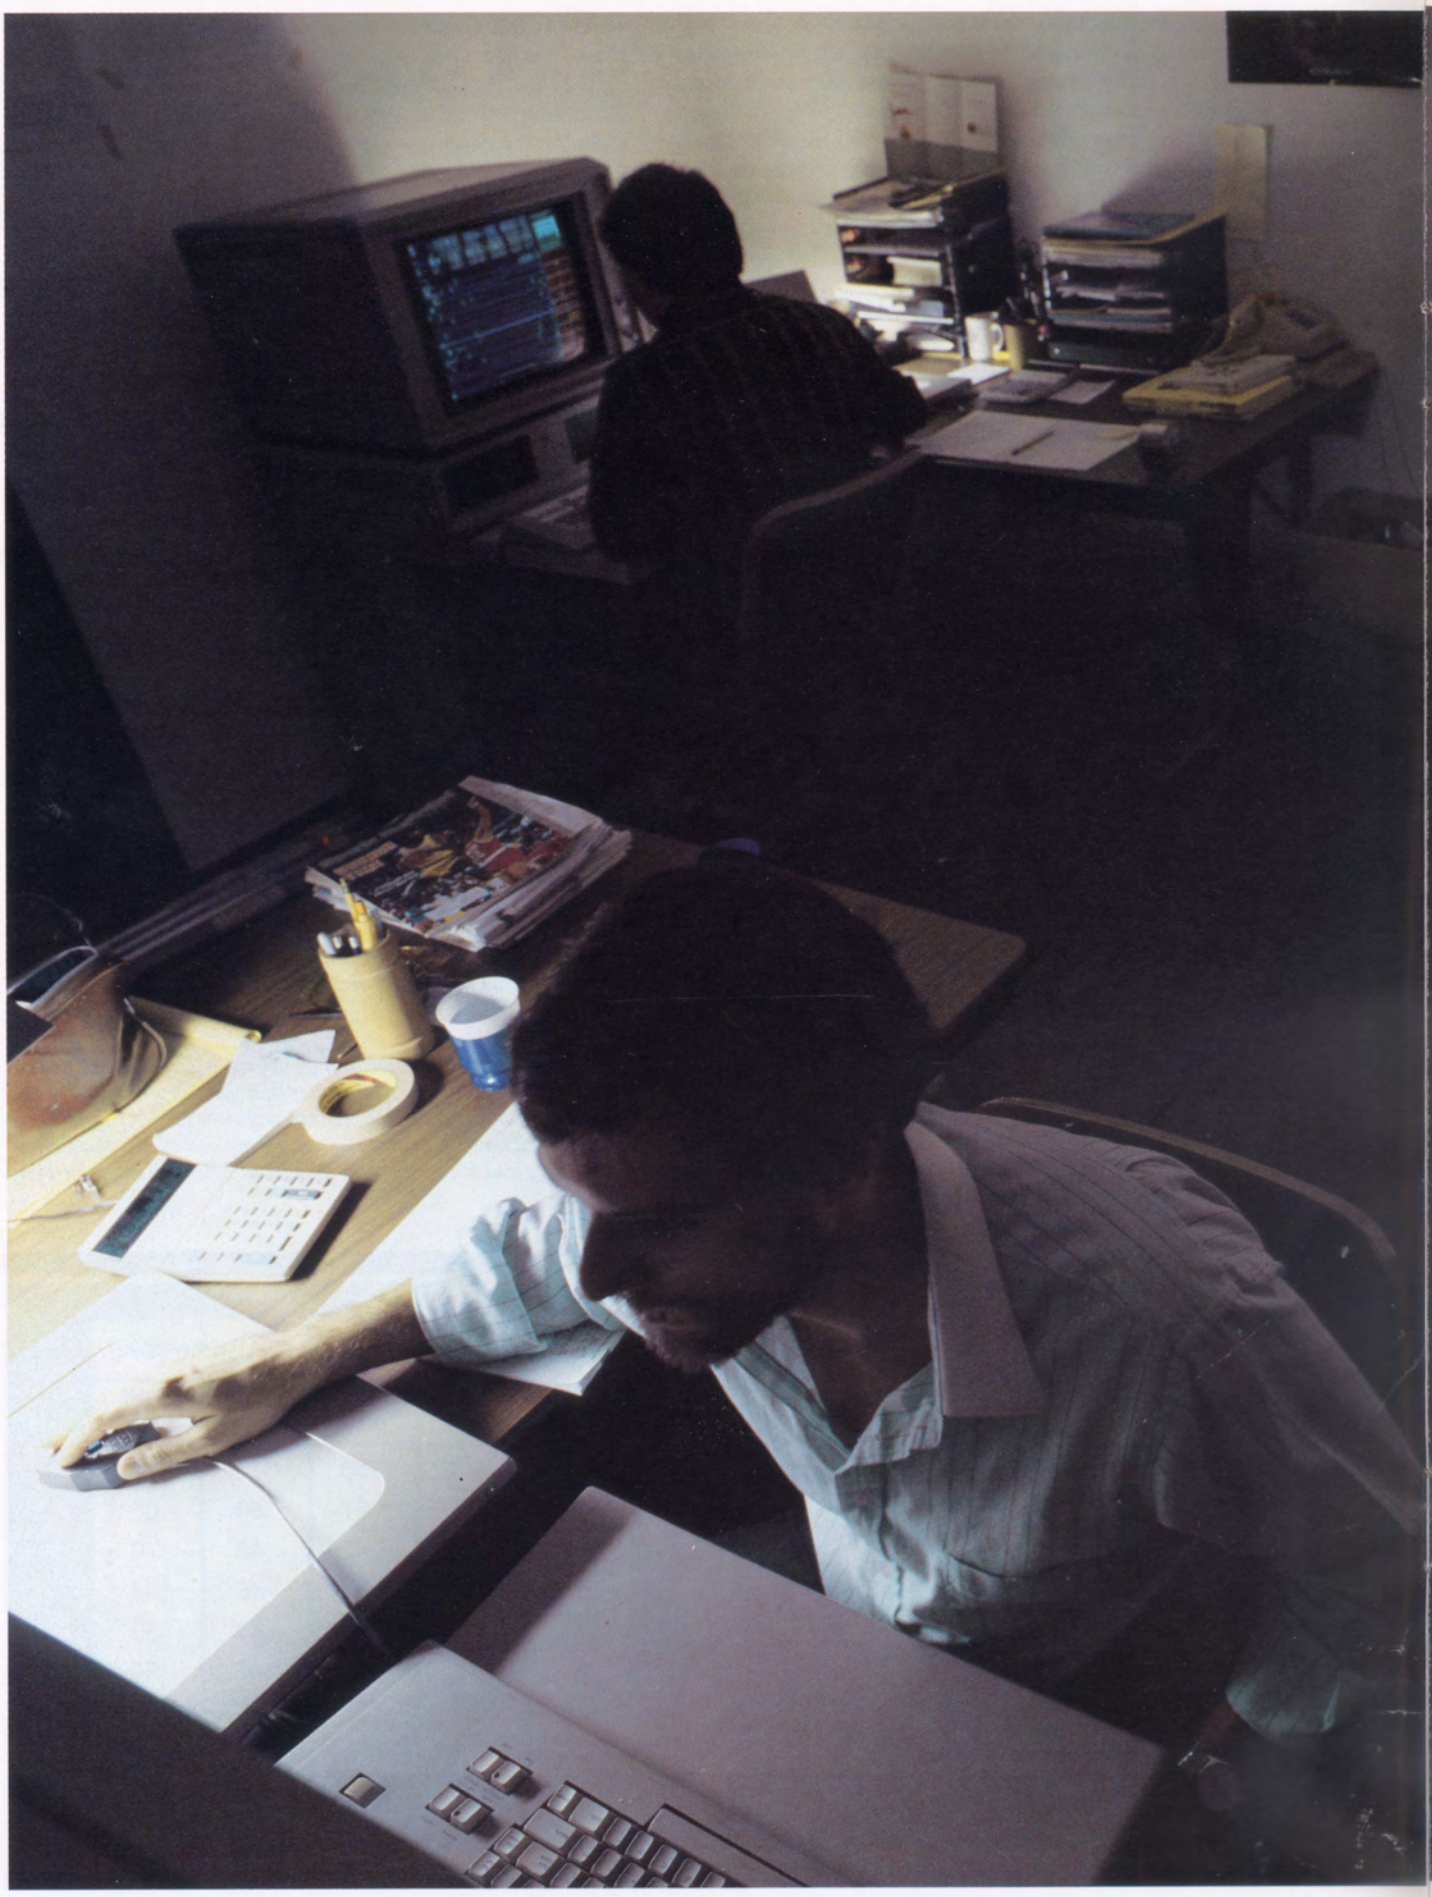

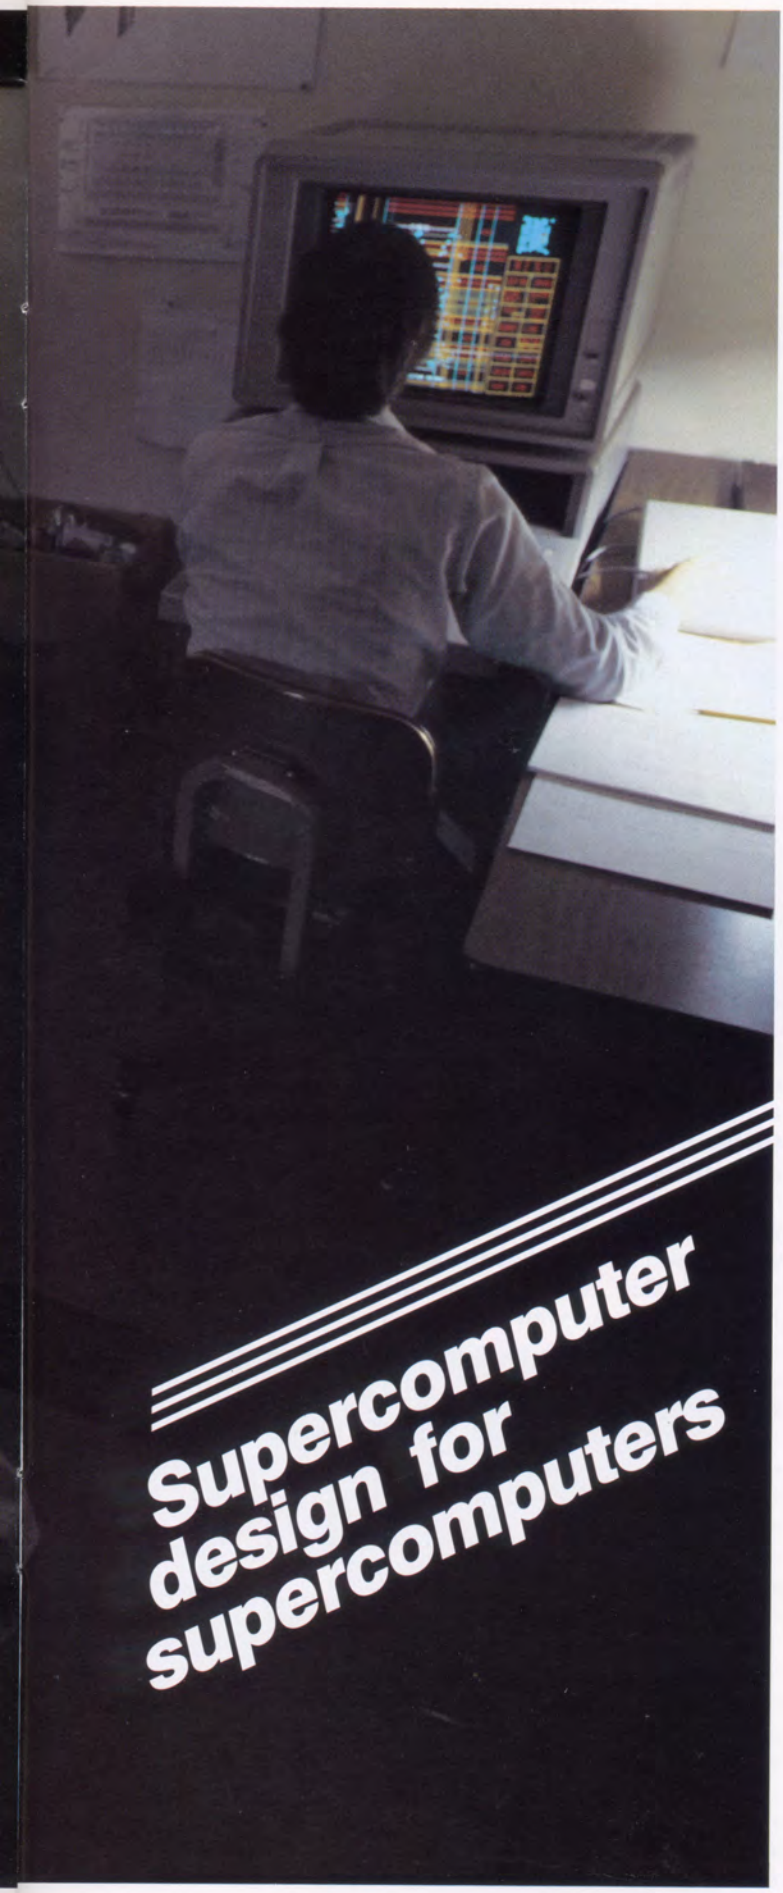

Supercomputer design for supercomputers

Creating tomorrow's computers is a challenge being met with the help of today's computers. Computer-aided design and engineering, CAD and CAE, are now used extensively in the computer industry. In this way computer technology has become self-augmenting, with computers themselves playing a vital role in the creation of faster, more powerful computing systems.

At Cray Research's product development facilities in Chippewa Falls, Wisconsin, designers and engineers are developing future Cray products using a CRAY-1/M system dedicated to electronic CAD and CAE, as well as other in-house Cray systems that serve many of the company's engineering needs.

Cray product development personnel rely on the in-house systems to reduce design turnaround time and solve large-scale computational problems. The supercomputers give designers and engineers rapid feedback on their ideas, shortening the design cycles that generate new products. The systems are used extensively for integrated circuit and printed circuit board routing, design-rule checking, and logic and fault simulation. The Cray systems provide time savings and large-scale simulation capabilities that would not be possible using any other type of computer system.

Production software

The Cray systems serve Cray's product development personnel as practical industrial design and engineering tools. "We use the CRAY-1/M system to accelerate the development and production of supercomputers," stresses Ed O'Neill, who maintains the commercial and public-domain software used in Cray Research's electronic CAD efforts. "Much of the CAD software we run on the system has been implemented on other hardware, but using the CRAY-1 system enables us to accomplish more with these programs, and do so faster. This is in addition to solving problems too large for conventional computers."

The speed and large memories of Cray systems are particularly well-suited for large-scale circuit simulations. The CRAY-1/M system typically can be found running three IC simulations at a given time, each containing from 2000 to more than 20,000 elements. A 23,000-element simulation with a 90-nanosecond transient time recently was run in just over seven hours. By contrast, 600- to 800-element runs are typical of electronic CAD work done on conventional mainframe computers.

A recent incident demonstrates a unique advantage in using Cray systems for large-scale simulation. O'Neill explains, "We had a situation with a circuit containing many identical subsets where simulating one subset did not indicate a problem, nor did simulating two identical subsets. But simulating seven identical subsets did reveal a problem, and the problem was present in the silicon. In this instance only a very large-scale simulation could have detected the problem."

"This is not to say that if you can't simulate 20,000 elements you are handicapped," O'Neill emphasizes. "Someone simulating small-scale integration probably doesn't need a Cray system to accommodate their component size, but

a design methodology for VLSI simulations that is based on small machines may not emulate faithfully the manufactured IC. The machines should be able to simulate enough circuitry to re-create any problem in the design. Often the ability to turnaround 20,000-element circuit simulations quickly is needed to detect flaws that only manifest themselves at scales of thousands of elements."

Another incident illustrates the value and sometimes unexpected benefits of fast turnaround. A designer was experiencing a two-week turnaround on a mainframe computer running a circuit simulation model. When he transferred the problem to the CRAY-1/M system, turnaround was reduced to a few hours. "As it turned out, in this case the long run on the mainframe was due to a defect in the program," O'Neill notes. "An overnight run on the Cray system made it obvious that something was wrong. Discovering the problem at the outset would have saved this designer some time and frustration."

Specialized software

Cray Research's electronic design and engineering efforts also take advantage of software designed in-house that is customized to take full advantage of the Cray systems' capabilities. An example is a recently developed design program for integrated circuits (ICs) that automates the placement of macros on a chip. Macros are groups of semiconductor gates that can be internally connected to perform a variety of logic functions. Arranging macros on a chip in a satisfactory manner constitutes a difficult optimization problem.

The automatic placement program begins with a random arrangement of macros and proposes successive random swaps. The program always accepts changes in the positions of macros if they improve the chip's overall arrangement. Bad changes may also be accepted, with a probability that decreases slowly to zero. This allows the system to jump out of local minima and reach near-optimal configurations. The algorithm used in the program differs from other widely-used algorithms in its overall goal, however. Generally, the goal of placement is to minimize the total metal length of the connections inside a chip and to minimize the chip area occupied. "Our primary goal is to make the chip function at the highest possible speed, so placement is driven by the propagation time of signals through the macros and along the connections," explains Philippe de Forcrand, a physicist at Cray Research who helped design the placement program. "Placement using this algorithm is computationally intensive, but access to the CRAY-1 system makes it a practical option. As a result, we know we're designing the fastest logic circuits we can."

Once placement of macros within an IC is complete, the circuit's interconnections must be routed. Routing is performed using the program DRAW, which enables engineers to route printed circuit (PC) boards and ICs interactively on the CRAY-1/M system using a graphics terminal. "DRAW was written for the CRAY-1 system because design-rule checking was already running on the system and we wanted to demonstrate interactive graphics directly on the Cray system," explains Joel Garcia, designer of DRAW. "Having a computing system powerful enough to handle both processes also spares us the inconvenience of sending files around among smaller systems."

Design-rule checking is performed using another program developed in-house, DELAY. DELAY converts the designer's logic into a working design file that inserts default logic, checks syntax, and calculates the time needed for the circuit. Within an IC are logic groups called macros, with their own design rules that must be checked. When DELAY verifies that an IC layout meets the design rules, the design goes to mask processing. Another program does design-rule verification for PC boards.

These programs were written specifically to check components for a particular Cray product under development. Writing such programs exclusively for a single project is a calculated strategy. "The prevailing philosophy in the electronics industry promotes large all-purpose software packages for CAD," explains Harro Zimmermann, author of Cray's design-rule checking programs. "But such programs strive to be so flexible that they burden themselves and end up very slow, whereas our programs are application-specific. They are very fast because they are only solving problems relevant to one system's design needs."

General-purpose software packages also are limited by their applicability to certain technologies, and so become compromised when new technologies appear. Fiber optics may be an example. "By the time programs appear to handle fiber optics, even newer technologies will be around. So I'm convinced that fast, application-specific programs are most efficent for long-range productivity," says Zimmermann.

New levels of miniaturization exert another limit on existing packages. Proximity effects, for example, appear at submicron scales where circuit lines become smaller than the wavelength of light used to make masks. As a result, two such lines near each other can blur together on a mask. "In that event, the routings would have to be modified to get around the problem. Existing commercial CAD packages and dedicated hardware would take an incredibly long time to do this," notes Zimmermann. He adds, "By combining our software with the power of a Cray system, we're

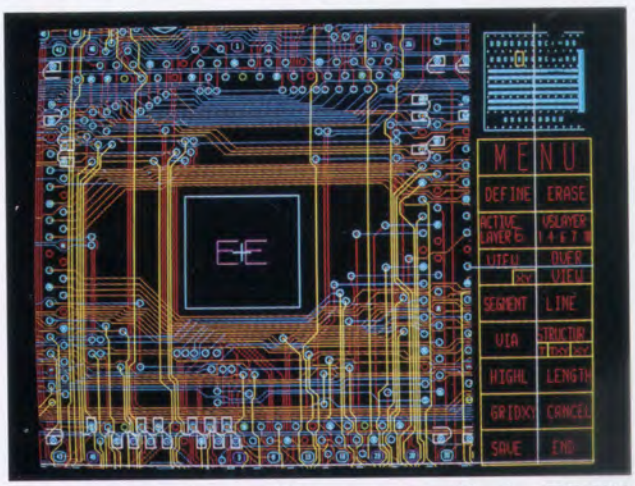

IC being routed on a Cray system with the DRAW program.

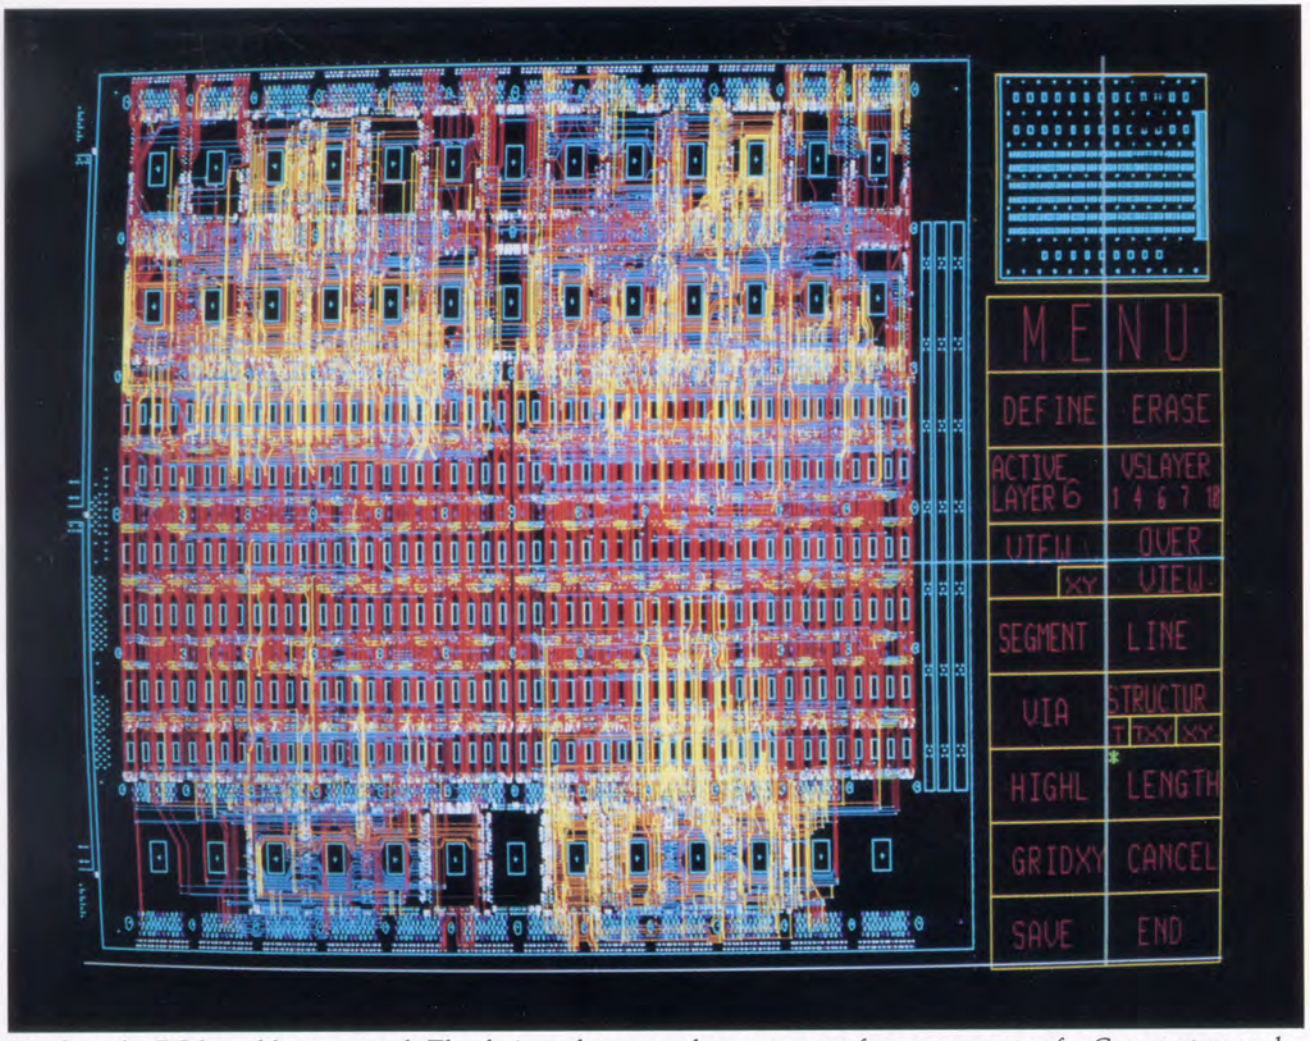

Display of a PC board being routed. The designs shown on these pages are for components of a Cray system under development.

able to tackle extraordinarily large problems, as well. For example, we are considering using the CRAY-1/M system to perform design-rule checking for the CPU of the follow-on to the CRAY X-MP series — the entire CPU in one run. We couldn't consider such a large run on a computer smaller than the CRAY-1 system."

Fault simulation is another application for which development personnel rely on the CRAY-1 system. A fault simulator has been written specifically to run on the Cray system, although engineers using programs developed on other systems also benefit from the Cray system's capabilities. "A vendor helping us build chips has a fault simulator on which a single run costs $500," notes programmer Tim Sirianni. "The same simulation runs in one or two minutes on the CRAY-1 system, at a cost significantly less."

A comprehensive tool

Speed is the essential ingredient that makes Cray systems superior electronic design tools. This speed is complemented by the systems' large memories. Both factors are essential to meet the CAD and CAE requirements of Cray's new product development groups. "The Cray systems are absolutely necessary for our work," says Ken Seufzer, who oversees PC board layout, routing, and design-rule checking. "For the follow-on to the CRAY X-MP system, we are working on PC boards that measure 11 by 22 inches. Because of the large size and complexity of these boards, a single verification can take up to eight hours on a CRAY-1 system. Runs like that would be unthinkable if we were restricted to anything less than a supercomputer."

Cray systems demonstrate their superior performance every working day in Cray Research's demanding electronic design and engineering environment. Tapping the power of its own products has played a significant role in making Cray Research the world's foremost supercomputer producer. By no means other than through the use of Cray systems could Cray Research have so successfully turned design innovations into deliverable, reliable, high-performance products.

But the electronic CAD and CAE capabilities of Cray systems extend beyond the demands of supercomputer design and engineering. Cray systems provide a competitive advantage in any electronic design and engineering environment. The rapid turnaround they provide can dramatically shorten design cycles, enabling manufacturers to introduce their components to the market in the shortest possible time.□

The greening of computational science

supercomputing at the University of Illinois

Many things have changed since 1982 when the need for supercomputing capability in academia was just beginning to be appreciated in the United States. Thanks to the outpouring of financial and political support from all levels of government and the National Science Foundation, U.S. universities are being girded with computing power the likes of which have never been seen before. A total of 13 supercomputers have been installed at U.S. universities since 1985. Prior to that time, only three U.S. universities had supercomputers.

In 1985, the University of Illinois became the second U.S. university to acquire a Cray supercomputer and the first university to acquire a Cray system under the National Science Foundation (NSF) supercomputing initiative. (Ultimately, three of the five NSF supercomputing centers installed Cray supercomputers.) Initially, a CRAY X-MP/24 system was installed at the National Center for Supercomputing Applications (NCSA) on the Urbana-Champaign campus. At the close of 1986, the center doubled its supercomputing capacity with an upgrade to a CRAY X-MP/48 supercomputer. As an NSF supercomputing center, one of NCSA's primary responsibilities is to provide computing cycles to the research community, awarded based on peer review. Along with the other NSF centers, the NCSA at Illinois has been very successful in providing time to qualified researchers.

Under the direction of Dr. Larry Smarr, professor of astronomy and physics at the University of Illinois, the endeavors of the NCSA go far beyond just providing supercomputing time. In fact, with additional funding provided by the state of Illinois, the University of Illinois at Urbana-Champaign, and donations from computer vendors, the workings of the NCSA are designed to spawn Smarr's proposed "revolution in computational science."

The NCSA is dedicated to effecting real change in the way computational research is approached. Dr. Robert Wilhelmson, associate director of NCSA and atmospheric scientist at the University of Illinois, explained, "We are committed to being responsive to user needs, to offering an environment and the computing tools that allow scientists to maximize their research efforts. That includes providing tools that help researchers execute the computations and study the results more productively. In addition, NCSA provides a local environment in which researchers from different geographic and intellectual areas can work side-by-side on problems of common interest, such as computing techniques or methods of visualizing data. Coupled with good internal and external networking capability, NCSA acts as a catalyst and partner in new discoveries by researchers throughout the country."

The evolution of computational science

Smarr contends that fundamental changes in the way computational science is approached and conducted will develop rapidly over the next five years. Today, changes in technology are most recognized and accepted. However, as advances in all facets of computing technology — including computer graphics, software, and networking — are enhanced and mainstreamed, changes will occur in the human enterprise or sociology of doing science.

Changes in technology

The maturation of many facets of computing technology, including mass storage, front-end systems, software, networking, and computer graphics, are having a real impact in the evolution of computational research. While advances on these various technologies are changing the approach to computational research, a balanced computational system linking various components in an effective way also is critical. Smarr summarized, "The real question we face is: How should the entire supercomputer environment be designed to turn raw computing power into advances in human knowledge? That is the question we believe NCSA answers."

Several key technology components dramatically affect the computing environment at NCSA. Those reviewed in this article are the supercomputer, personal computers and workstations, computer graphics, and networking.

The great partnership — supercomputer and workstation

Two great technologies, the supercomputer and the personal computer, fundamentally have changed scientific research and the ease with which all computing can be done. The supercomputer has allowed researchers to address questions and find solutions that could not previously be obtained. Meanwhile, the personal computer makes computing accessible to millions of people.

Smarr observes that today we are seeing the unification of these two great technologies, the supercomputer and per-

sonal computer, in a national distributed network system of computing. User workstations provide user friendliness, software, and productivity tools that are not practical for the supercomputer. At the same time, supercomputers offer unequaled computational capabilities. By coupling the supercomputer with personal computers, the entire scientific production cycle, not just the computation, can be optimized.

The National Center for Supercomputer Applications is dedicated to building a comprehensive computing environment based on these emerging technologies. The computing system at the NCSA consists of three basic components: the supercomputer, comprising the CRAY X-MP/48 computer system with SSD solid-state storage device, the workstation, and the scientist.

The NCSA brings supercomputer power closer to the user, but in the user's own environment, that of the workstation. The key to maintaining that environment is that the system at large should retain as much of this familiar character as possible as the user interacts with the various system components. Users should not be badgered with inconveniences that can hamper productivity, yet they need access to different machines, both locally and remotely, with differing capabilities to do what they do best: science.

Wilhelmson explained, "NCSA is putting computational tools before researchers so they can work as efficiently as possible. We are preserving the workstation environment that researchers are comfortable with, while delivering the power of the supercomputer through appropriate commercial software and development of new software."

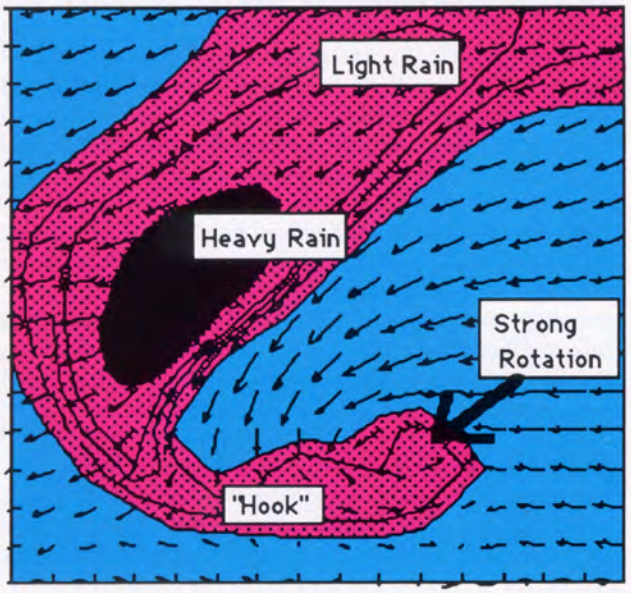

Figure 1. Results from a portion of a thunderstorm model simulation on the Cray system at NCSA. Notice the classic "hook" pattern in the simulated radar echo, a feature often associated with observed tornadic thunderstorms. The strong rotation in the wind field at the end of the hook is the parent circulation from which most tornadoes are born. (Courtesy Wicker, Wilhelmson, Shaw, and Brooks, University of Illinois)

An important component of the NCSA program is providing a workstation environment in which transfer of data between the workstation and the Cray system is straightforward and easy and in which effective software is available for manipulating test and graphics data on the workstation. Figure 1 illustrates the results of sending a graphics file to an Apple Macintosh Plus from the Cray system and the addition of text, annotations, and other manipulations performed on the Mac with standard commercial software.

Consultants, research scientists, and the workstation group staff at NCSA are all involved in meeting NCSA's workstation objectives. In addition, NCSA has sought cooperation and received support from major computer vendors including Apple, DEC, Hewlett Packard, IBM, Silicon Graphics, and Sun Microsystems.

A key challenge facing the development team is mastery of the multiple user interfaces resulting from the linking of different computers, operating systems, software packages, and network interfaces. For example, in a currently experimental IBM PC environment, each computer is activated by typing a disk command whereby each appears as a different IBM PC disk drive. One command automatically issues a series of commands that logs a remote user on to the CRAY X-MP system. From then on, the normal PC-DOS commands send out their corresponding commands to the Cray supercomputer. With this software, users only need to know the operating system of their computer to carry out many activities.

Computer graphics

Another area of technology having a profound effect on computational science is computer graphics. Because computational science has the ability to find solutions of "realistic complexity," it relies heavily on the visualization of data to represent solutions. Thousands of numbers are not processed by the human brain very effectively or efficiently. Upon entering the eye-brain system, however, images are converted into recognizable structures that can be analyzed.

Smarr observes that visualization allows for a direct coupling of the supercomputing results into the eye-brain system, which further couples with our intuitive and associative memory capabilities. Spatial and time complexities generated by supercomputer simulations can be illustrated via computer graphics. With the help of video or movie presentations, for example, the human eye can couple the keen senses of time variation, causality, and intuition from a computation.

An important aspect of the NCSA's effort is the development of computer graphics capabilities — both graphics for easy daily use and high-performance graphics. In the former case, producing color graphics or images from large multidimensional computations or from raw observational or laboratory data involves interaction with four basic elements of the NCSA computer environment: the CRAY X-MP computer system, the VAX front-end, a network, and a workstation. Means for streamlining passage of data through this system are being developed.

For high-performance graphics NCSA is establishing an Advanced Visualization Facility. This facility will provide scientists with the capability to interact with and explore their data in real-time using graphical animations of their simulations and to produce videotapes and films of their research results.

Summarizing the need for good visualization tools, Wilhelmson said, "It is important to remember that understanding and knowledge are often hidden in masses of data and that our job as researchers is not only to collect and present the data, but to interrogate and learn from it."

Networks

The advancement of networking technology is crucial to the greening of computational science. Access to supercomputer power by users, both on-site and remote, is a key issue among computer scientists and users around the country. Major questions being asked today are: How do we handle high-volume data? How do we move information effectively and efficiently from the scientist's workstation to the supercomputer so productivity is not negatively affected? Meaningful answers to these questions will dramatically change the ease with which scientists can conduct their research.

It is expected that thousands of remote users will access the supercomputer centers over NSFnet, a national network supported by NSF. The network will be an overlapping set of regional networks hooked together by a high-speed backbone connecting the NSF centers. The backbone currently transmits data at a rate of 56 Kbits/sec, and NCSA plans to upgrade that performance in the near future. The current goal is to use a common protocol across all these networks, namely TCP/IP. The University of Illinois Computing Services Office, as a part of its involvement with NCSA, has played a major role in the implementation and management of the backbone and will continue for the next year to manage the backbone through a grant from the NSF.

It is not only important to provide remote users with good access to supercomputers, but also to provide local improvements in data communication between the supercomputer, mass storage, and visualization tools. As supercomputers become more powerful, more data is produced. The balance between what the scientist can store and see, and how much data is produced must be continually assessed as new technology becomes available. As part of NCSA's visualization program, future plans include the development of high-bandwidth communications (100 Mbytes/sec) between its components and the CRAY X-MP/48 computer system.

Changes in the sociology of science

It generally is accepted that computation becomes an integral part of scientific discovery by linking theory and experiment. Theory and analysis precede computation, and the results of a complex code are validated against experimental data. The expectation is that solutions obtained in this manner are more accurate than if computation were omitted. The computational link between theory and experiment impacts the intellectual environment in which science is conducted in addition to affecting the qualitative results of science.

Smarr explained, "In many areas of research today, the ability to computationally simulate real phenomena pulls theorists, experimentalists, and computationalists into a natural team in an effort to understand a complex problem. Traditionally, analytic science has been done in the 'single warrior combat mode.' In computational science, there are so many new technologies in hardware and software that one needs experts in the use of such tools working with the scientist.".

Smarr went on to say that this level of interdisciplinary teamwork relates to the universal character of the computer. "By that I mean each of us, regardless of our field of study, must edit and transfer files, learn vectorization, use workstations, etc. Until now, we had to spend much of our time learning how to use this technology on our own. We had to investigate the literature to discover what products were available, then we had to try them out and wire them up ourselves; we had to wade through manuals to figure out how to use the software, and so forth."

Smarr believes there is a major opportunity to alter this process with the NSF supercomputing centers. The national centers should serve as a reservoir of expertise on all these issues, and they should provide national training centers where individuals can go to learn the technological side of the field in a short period of time. The centers have already done a great deal of investigation to determine which commercial pieces of hardware and software are most effective. If the centers provide this training function, individual researchers can save an enormous amount of time.

Beyond the technology considerations, computational communities that differ from the traditional departmental structure of a university are emerging. For example, any investigator who is solving a quantum mechanical problem will be led to inherently probabilistic algorithms such as Monte Carlo. These investigators may come from the departments of physics, chemistry, material properties, biochemistry, etc. However, they have a common algorithmic interest, which if shared, would increase their productivity. Right now there is not a good forum for this shared interest because most journals are defined along content lines rather than algorithmic lines.

To facilitate the germination of ideas across disciplinary boundaries, the NCSA established the Interdisciplinary Research Center (IRC). Essentially, the IRC is a think tank for researchers fostering new interdisciplinary approaches in computational science. At the IRC researchers confer with other visiting scientists, computing consultants, staff research scientists, and computer development professionals.

Visitors are sponsored by the NCSA Visitors Program. A user can combine both training and research in one visit to the IRC. Each new visitor has the advantage of being able to use the assets left behind by previous visitors.

The center's scientific program is organized around new interdisciplinary groupings based on the common underlying mathematics of research areas. One such community studies continuum systems based on classical field theories. These researchers use, for example, computational algorithms based on finite element, finite differencing, and spectral methods. Another computational community concerns itself with quantum and statistical systems for which the underlying mathematics are probability theory and discrete math. These and other computational communities are composed of University of Illinois faculty; in the future, they will expand to include national users from academia and industry.

Still another example of NCSA's commitment to interdisciplinary teams is evidenced by the center's Industrial Supercomputing Program. "Corporate partners are invited to participate by sending teams of corporate researchers to work shoulder-to-shoulder with the best minds from universities," said Smarr.

In 1986, Eastman Kodak Company became the first partner in the program, committing $1 million a year for a three-year period. In exchange, Kodak researchers are situated at the IRC and are working on basic research and product development. The company is using supercomputing to learn new techniques that can be applied to its research. In fact, Kodak is discovering that some of its old assumptions about product design no longer apply. From experiences with supercomputing at NCSA, Kodak researchers can now consider a broader solution space and then focus on the best alternatives.

Interdisciplinary cooperation at NCSA may prove to be of value to agriculture as well. David Onstad, a scientist, and Donna Cox, an artist, recently teamed with a computer programmer to devise new ways to present data obtained from modeling the population of the corn borer, an insect that wreaks havoc with Midwestern corn crops. Color graphics were used to illustrate the subtle changes in infected and healthy larvae of the corn borer over the 140 days of the corn growing season.

A recent initiative for computational electronics illustrates the impact a common tool such as the supercomputer can have on a shared goal. The initiative is an outgrowth of Karl Hess's use of the Cray system as a Cray Research Affiliate and his long-term involvement with NCSA. Hess and Umberto Ravaioli, both professors of electrical and computer engineering at the University of Illinois, are leading this national initiative with the goal of improving communication between a large number of groups, which have devoted their efforts to the large-scale simulation of semiconductor electronics and optoelectronics.

Recent work by Hess and coworkers shows the kind of impact that research in this area can have. They have simulated details of nonlinear transport in high electron mobility transistors as well as in optoelectronic devices such as avalanche photodiodes. Of particular interest has been the study of fast transients and the ultimate speed limitations of these devices. They have found from two-dimensional simulations and visualization that the switching on and off of transistors is highly asymmetric, a fact that had not been known previously. The discovery was made possible through the making of a movie, a frame of which is shown in Figure 2. Significant progress has also been made in the understanding of so-called "dead spaces" in avalanche photodiodes and other devices based on impact ionization using ensemble Monte Carlo techniques and extensive simulations.

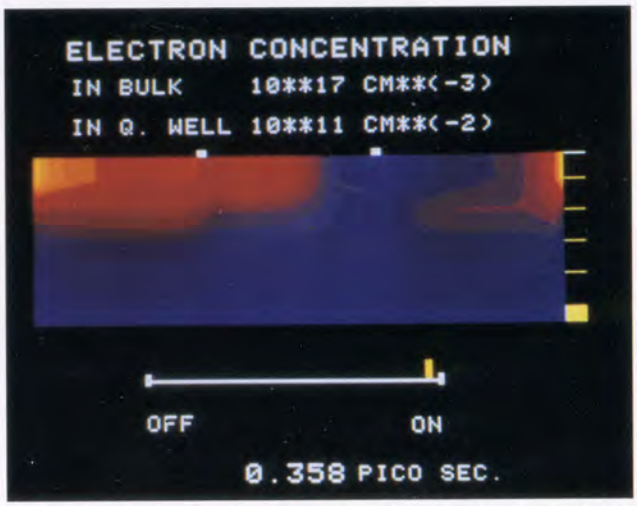

Figure 2. Graphic output showing the switching on of the novel high electron mobility transistor. Yellow denotes high electron concentration; blue denotes low electron concentration. The output is a result of a cooperative effort between Cray Research and the University of Illinois (Karl Hess, et al).

Summary

The National Center for Supercomputing Applications, founded on the premise that we are in the midst of a revolution in computational science, is ambitious and rich in potential. Smarr and his colleagues have generated the energy and momentum that will make a real difference in the manner in which science is conducted. The support of the funding organizations, vendors, and corporate partners contributes to the current accomplishments and future successes of the center.

Along with other university supercomputing centers, the NCSA is breaking new ground in the quantity and quality of academic scientific computing. The national center has many development projects that are helping to make computational research easier for scientists in all areas. Projects at NCSA such as the Visitors' Program, computer graphics work, developing the personal computer/supercomputer link, and many others not mentioned here, are allowing scientists to get back to being scientists.□

Acknowledgments

Thanks are extended to Larry Smarr and Robert Wilhelmson for their assistance in the preparation of this article. Much of the information presented here is taken from Smarr's review article entitled "The Computational Science Revolution: Technology, Methodology, and Sociology," which appears in High-Speed Computing: Scientific Applications and Algorithm Design, *University of Illinois Press, 1987.*

COMPUTATION OF FAST FOURIER TRANSFORMS

Multidimensional transforms come of age

Mickey Edwards, Cray Research, Inc.

Fourier transforms are used in virtually every area of physics and applied mathematics computing. There are many reasons for using transforms. Numerous mathematical operations, such as convolution or filtering and differentiation, can be performed more efficiently in the Fourier domain. In areas such as wave propagation theory, electrical engineering, and signal processing, fundamental theoretical development is accomplished in the Fourier domain. Laplace and other transforms found in many theoretical disciplines can be expressed as Fourier transforms with a change of variable or substitution.

The birth of FFTs

In 1965, J.W. Cooley and J.W. Tukey published the paper "An Algorithm for the Machine Computation of the Complex Fourier Series" in *Mathematics of Computation*. This was the birth of fast Fourier transforms (FFTs) and launched an important field of algorithm research. The value of Fourier domain algorithms as an alternative to spatial and spatial-time domain formulations was well understood prior to 1965, but Fourier transform computation was generally not feasible because high-speed digital computers did not exist.

The discrete Fourier transform (DFT) of an input sequence of N real or complex numbers $f_0, f_1, f_2, \ldots, f_{N-1}$ is defined as

$$F_k = \sum_{n=0}^{N-1} f_n e^{-i(2\pi nk/N)}$$

for $k = 0, 1, 2, \ldots, N-1$. The inverse discrete Fourier transform (IDFT) of the complex sequence $F_0, F_1, F_2, \ldots, F_{N-1}$ is then

$$f_n = \frac{1}{N} \sum_{k=0}^{N-1} F_k e^{i(2\pi nk/N)}$$

for $n = 0, 1, 2, \ldots, N-1$ the symbol i denotes the square root of -1 and e denotes the natural logarithmic base 2.71828.... Computation of the above DFT or IDFT requires on the order of N^2 complex multiply-and-add operations.

An FFT is a computational algorithm that orders DFT computations in a manner requiring on the order of $N \text{ Log}_2 N$ or fewer operations. In this article, all DFT algorithms are identified as FFTs. For brevity, this article will not attempt to acknowledge the many of researchers who have worked in FFT algorithm development nor will it differentiate among algorithms.

One-dimensional FFTs

When FFT computation was in its infancy, computations of base 2 or binary radix FFTs were the first to be implemented. To compute such FFTs, N is required to be a power of 2 ($N = 2^M$). If necessary, zeros are appended to the input sequence to obtain a power of 2 length. Implemented in the "simplest" vectorized form, an outer loop of M passes successively performs multiplication and addition operations with inner loops (or vector lengths) of $N/2, N/4, N/8, \ldots, 4, 2$, and 1 element. The short vector lengths of the last outer loop passes can be modified by inverting or interchanging inner and outer loops for those passes. For example, inner loop vector lengths of $N/2, N/4, N/8, \ldots, 2^{M/2+1}, 2^{M/2}, 2^{M/2}, 2^{M/2+1}, \ldots, N/8, N/4, N/2$ can be used for even M.

Assuming a complex valued input sequence, the so-called butterfly, or twiddle, operation requires four real multiplies and six real adds. Assuming pre-computation of complex exponential multipliers, $5MN$ floating point operations are required. If it is recognized that the real or imaginary component of some complex exponential multipliers is zero, the floating-point operations count can be reduced slightly. Input and output arrays can coincide or use the same memory locations.

The only drawback of the "simplest" FFT is that the final transformed result is stored with bit-reversed indices. The kth result ($k = b_0 + b_1 2 + b_2 2^2 + \ldots + b_{N-1} 2^{N-1}$ where each b_i is equal to 0 or 1) actually belongs in the k'th position ($k' = b_{N-1} + b_{N-2} 2 + b_{N-3} 2^3 + \ldots + b_0 2^{N-1}$). Reordering final bit-reversed results is known as unscrambling. The "simplest" inverse FFT is computed in the same manner and the resulting complex valued time or spatial sequence is then unscrambled.

The unscrambling problem was eliminated with new algorithms that reordered computations such that after the Mth outer-loop pass, final results were stored in natural order. These newer algorithms require a work array and the results of intermediate outer-loop passes ping-pong between input and work arrays. They are computationally comparable and maintain the same number of outer- and inner-loop passes.

Some of these natural-ordered or unscrambled algorithms are less suitable for vector processing due to memory addressing. The algorithms typically load or store vectors with a multiple of two memory strides or increments. In this article, however, a discussion of parallel FFT computation shows that all of these algorithms can be implemented on vector processors.

Implementation on Cray systems

The work discussed here has been done primarily on CRAY X-MP supercomputers, although much work was initally done on CRAY-1 systems. Memory in CRAY X-MP systems is interleaved in 64, 32, or 16 memory banks. Each memory bank can be accessed every four or eight clock periods, depending on whether the system has ECL bipolar or MOS memory technology. Older CRAY X-MP systems feature a 9.5-nanosecond (nsec) clock period; those

installed after mid-1986 have an 8.5-nsec clock; CRAY-1 supercomputers have a 12- or 12.5-nsec clock period.

The criterion for selecting FFT algorithms for efficient computation on Cray supercomputers is rather simple. Memory bank performance must be assessed. To perform efficient vector operations, users must be able to load and store vectors without memory bank conflicts. In memory, vector elements can be stored in consecutive locations, with a uniform stride, or with a nonuniform increment. The last case requires the random gather/scatter hardware instruction.

For most natural-ordered base 2 FFT algorithms, it is impossible to avoid some undesirable memory addressing. Problems can be eased by introducing a suitable storage scheme for intermediate outer-loop passes. Because CRAY X-MP systems have multiple numbers of memory banks and because the multiple CPUs share the same common memory, which makes memory behavior less predictable, it is difficult to arrive at one simple memory bank conflict formula. The common practice is to avoid strides that are powers of 2.

On CRAY X-MP computer systems, implementation of the "simplest" FFT uses a memory stride or increment of two for all vector loads and stores. Butterfly computation requires four vector loads and four vector stores. The memory stride of two for conventional real-and-imaginary complex storage is accommodated by CRAY X-MP and CRAY-1 architecture.

Because butterfly computation requires four vector loads and four vector stores, implementation on CRAY-1 computer systems was memory-bound due to the single memory path on those systems. Optimal implementation required eight chimes (a chained operation time).

The large memory bandwidth of the CRAY X-MP supercomputer almost doubled CRAY-1 computer system performance. With the CRAY X-MP system, the "simplest" FFT becomes add-functional-unit-bound. The six butterfly adds require six chimes and the vector loads, stores, and multiplications can be performed within the six chimes. The reduction in the number of chimes needed with CRAY X-MP computers and the difference in clock periods between CRAY X-MP and CRAY-1 supercomputers yields the performance improvement realized with the CRAY X-MP computer system.

The base 2 or binary radix FFT routines CFFT2, RCFFT2, and CRFFT2 found in the Cray Research SCILIB mathematical library support complex-to-complex, real-to-complex, and complex-to-real transforms. Table 1 shows CRAY X-MP computational times for various transform lengths.

Prime radix FFTs

The desire to improve base 2 or binary radix FFT performance led to the implementation of prime radix FFTs. That is, N is a power of p ($N = p^M$), where p is a prime integer (3, 5, 7, 11, ...). Many natural-order prime radix FFTs can be implemented with an outer loop of M passes and inner-loop vector lengths that are powers of p. Memory addressing when a stride is a multiple of p is suitable for all Cray supercomputers.

The implementation of efficient prime radix FFTs is difficult because of the more complicated butterfly computation. On Cray supercomputers these algorithms become vector-register bound requiring storage and reloading of vector registers. However, the large memory bandwidth of the CRAY X-MP system is designed to support this kind of memory traffic without much delay.

The tightest known CRAY X-MP base 3 butterfly requires 26 chimes. Base 5 butterflies require almost twice as many chimes. Prime radix FFTs are desirable because the number of floating-point operations required is reduced. For example, consider $5^3 = 125$ and $2^7 = 128$. Due to computational complexity, a base 5 butterfly will achieve fewer floating-point operations per second than a base 2 butterfly. However, if sufficiently optimized, three outer-loop passes in the base 5 algorithm (compared to seven in the base 2 algorithm) can yield a reduced computation time.

Mixed radix FFTs

Mixed radix FFTs are the most desirable algorithms. Here N is equal to the product of factors, each of which is a power of a prime integer base. A new outer loop processes each base. For each base, outer- and inner-loop computations are similar to the binary and prime radix FFT algorithms. For implementation on Cray computer systems, minimization of memory bank conflicts typically requires that any base 2 computations be performed after computations in any other bases. Programming difficulties typically limit the number of bases actually em-

FFT sizes	CRAY X-MP times (ms) (1 CPU, 8.5 nsec clock)		
length	CFFT2	RCFFT2	CRFT2
64	0.0275	0.0209	0.0207
256	0.0915	0.0564	0.0559
1024	0.4037	0.2255	0.2239

Table 1. Calculation times in milliseconds for one-dimensional FFTs on a single CPU of a CRAY X-MP with an 8.5-nsec clock.

FFT sizes		CRAY X-MP times (ms) (1 CPU, 8.5 nsec clock)		CRAY-2 times (ms)	
length	number computed in parallel	CPU time to compute all FFTs	CPU time/FFT	CPU time to compute all FFTs	CPU time/FFT
64	64	0.582	0.0091	0.42	0.0066
64	128	1.14	0.0089	0.77	0.0060
64	256	2.25	0.0088	1.48	0.0058
1024	64	15.7	0.245	10.8	0.169
1024	128	31.2	0.244	20.7	0.162
1024	256	62.3	0.243	40.4	0.158

Table 2. Single-CPU calculation times for parallel complex-to-complex FFTs.

ployed. Mixed radix FFTs are desirable because the number of necessary appended zeros is minimized and the exponent or number of outer passes for each base can be kept small.

Parallel one-dimensional FFTs

If computation of more than one FFT of the same length is necessary, programming difficulties can be greatly diminished. Consider, for example, an array of LN real or complex elements with L time or spatial sequences of N samples stored across the L rows. Parallel computation of L FFTs can be readily achieved by vectorization of down columns. Any FFT algorithm can be easily extended by replacing each operation with a single element f_i by L operations on the L f_i's. Each vector load or store uses a real memory increment of one or two. Alternatively, the L sequences could be stored down columns and vectorization can be achieved across rows. Memory conflicts will occur if N is divisible by numbers that are high powers of two. In this event, memory conflicts can be avoided by storage of a dummy row.

Typically, if the final transformed results are unscrambled, a work array equal in size to the input is required, which may be expensive. Alternatively, the size of the work array can be decreased by segmenting the computations and by performing 64 or less FFTs with each subroutine call. Subroutine arguments must include the word increment between successive elements of each sequence and the word increment between sequence-first elements.

On the other hand, if final transformed results are scrambled, a single array of N elements can be used to locate scrambled rows or columns. Typically, the user program computes or modifies the transformed results with no program degradation. The advantage of having to locate scrambled rows or columns is not having a work array requirement. Typically, parallel FFT computation is used with large problems that are amenable to parallel processing.

Cray users have made extensive use of parallel mixed radix FFT routines XFFT991 and XCFFT99 developed by Clive Temperton of the Canadian Meteorological Centre. New BENCHLIB routines CFFT999 and FFT999, derived from Temperton's work, have been developed by Dennis Kuba of Cray Research. Table 2 summarizes CRAY X-MP and CRAY-2 system performance of CFFT999.

Two-dimensional FFTs

For large problems, two-dimensional FFT computation can be achieved by two successive calls to a parallel FFT routine. When the matrix cannot reside in central memory, out-of-memory solutions using the SSD solid-state storage device and distributive disk storage can be employed successfully. The matrix is partitioned into submatrices small enough that a strip parallel to the x or y axis can reside in central memory. If disk storage is used, submatrices for each strip are stored on K disks. By staggering disk assignments for successive strips, users can issue K I/O requests and simultaneously read or write a strip. Figure 1 gives an example using four disks. The strip data motion solution is a transpose operation. Before or after retrieval or storage, the program must perform a sequence of move operations.

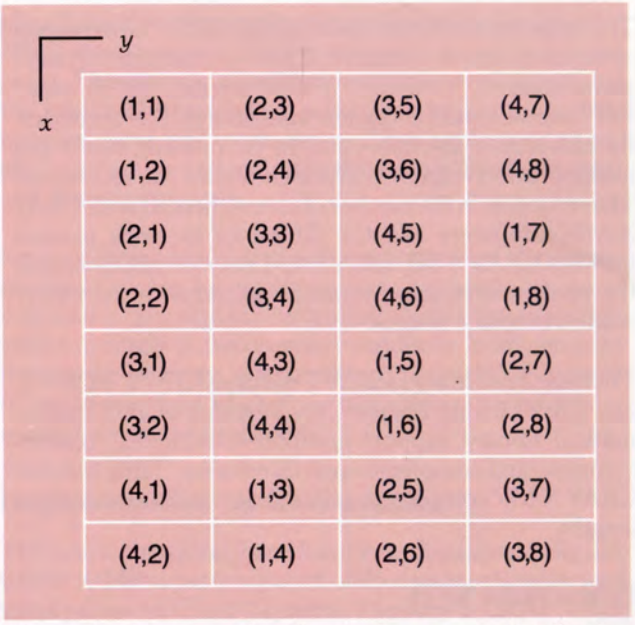

Figure 1. An example of out-of-memory distributive storage. The (3,8) at bottom right, for example denotes the eighth logical record of the third logical unit. Column or row strips parallel to the x or y axis can be retrieved or stored with four I/O operations.

FFT computation can be performed by parallel FFT routines or a series of one-dimensional FFT executions. For smaller problems, special purpose two-dimensional FFT routines may be in order.

Three-dimensional FFTs

Currently, Cray Research does not offer a generalized out-of-memory three-dimensional FFT capability. However, computation of large three-dimensional FFTs has been demonstrated. A forward FFT computation followed by inverse FFT computation using a 1024 x 1024 x 1024 real-valued dataset was completed in an elapsed or wallclock time of less than eight minutes.

The system configuration consisted of a dedicated CRAY X-MP/48 supercomputer, 40 million words of SSD storage, and eight DD-49 disk units (90 percent capacity on each unit). Using BENCHLIB routines PACK21 and EXPAND21, which pack and unpack 32-bit samples at the approximate rate of 2.5 clock periods per element, 32-bit samples were stored on disks. A distributive storage scheme was employed such that constant-x or constant-y planes could be simultaneously retrieved or stored by issuing four I/O requests. The SSD and queued I/O routines were used to effect a transpose of the second and third (x and y) storage indices. Triple buffering was used for concurrent retrieval of four disk streams, storage of four disk streams, and SSD transfers. Three data passes were required to perform forward and inverse FFT computation. The first pass computed x and z transforms, the second pass computed y and k_y transforms, and the third pass computed k_x and k_z transforms.

Total disk and SSD data movement of 3.2 billion words and 4.1 billion words, respectively, was performed. Dividing the total disk and SSD data movement by the elapsed time of 7.7 minutes yielded sustained transfer rates of 6.9 and 8.7 Mwords/sec on disk and SSD, respectively. In other words, a total sustained data motion rate of 15.6 million words per second was achieved. The 7.7 minute elapsed time did not include data staging or initialization of the input dataset.

The FFT computation used the aforementioned Temperton and SCILIB one-dimensional FFT routines. Parallel processing was achieved by Cray Research microtasking. Essentially, program execution was CPU-bound. Removal of the SSD would have necessitated two additional data passes to perform transpose before and after y and k_y transform steps.

Recently, three-dimensional FFT computation was demonstrated using a 512 x 512 x 512 real-valued dataset. Using 32-bit samples and the PACK21 and EXPAND21 routines, the dataset resided in 64 million words of SSD storage; disk storage was not used. Except for FFT computation by SCILIB of one-dimensional FFTs, the demonstration was identical to the above 1024 x 1024 x 1024 problem. An elapsed or wallclock time of approximately 1.33 minutes was required to perform forward or inverse FFTs using the CRAY X-MP/48 supercomputer in dedicated mode. Use of parallel FFT routines would reduce the wallclock time even further.

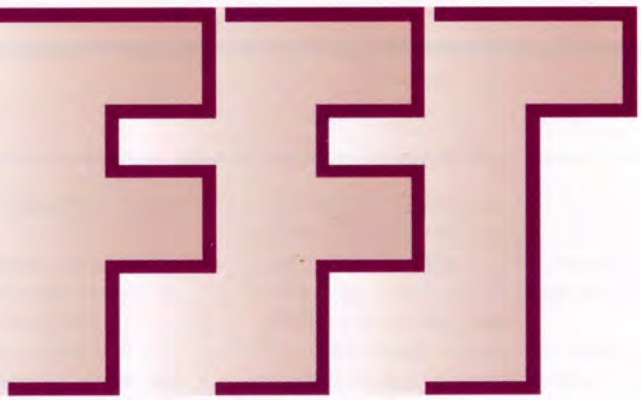

Four-dimensional FFTs

Successful demonstration of the three-dimensional FFTs discussed earlier suggests that out-of-memory four-dimensional FFT computation may be feasible. Obviously, 512 x 512 x 512 x 512 or 1024 x 1024 x 1024 x 1024 datasets are out of the question. It is conceivable, however, that meaningfully-sized four-dimensional problems, for example, 512 x 512 x 512 x 32, could reside on the disk farm and the resulting computational time could be bounded by some reasonable number of CRAY X-MP compute hours. Certainly, we must hope that, at some time in the not-so-distant future and with a future system, meaningfully-sized four-dimensional FFT computation will be feasible.

Conclusion

It can be argued that many FFT algorithms currently being used by the computing community are obsolete. This is probably due to programming difficulties that traditionally have led to the implementation of relatively simple algorithms. Considering current ongoing algorithm research and development, more sophisticated algorithms will emerge.

The importance of FFT computation has been greatly enhanced by the computational speed and I/O capabilities of supercomputers. The use of multidimensional FFT computation is rapidly becoming commonplace in many application areas. The capability to compute multidimensional FFTs in a timely manner is fueling the development of increasingly complex and sophisticated algorithms. Looking to the future, the feasibility of computing large multidimensional FFTs will provide real-time capabilities and revolutionary advances in some application areas.

Those interested in additional information may contact Mickey Edwards at Cray Research, Inc., 5847 San Felipe, Suite 3000, Houston, TX 77057; telephone: (713) 975-8998.□

About the author

Mickey Edwards manages an applications support group in Cray Research's Petroleum Region located in Houston, Texas. Since joining Cray Research in 1982, his activities have included the development of large three-dimensional algorithms. He has worked in the fields of signal processing and geophysics since 1964 after having earned B.S. and M.A. degrees in mathematics from the University of Oklahoma.

CORPORATE REGISTER

Cray Asia/Pacific Inc. serves the Far East

In January Cray Research announced the formation of Cray Asia/Pacific Inc. to support marketing efforts in many of the industrial countries of the Pacific Rim. The new business center is managed by Barry Utting, formerly an account manager at Cray Research (UK) Ltd. In his former position, Utting opened the Scandinavian market to Cray Research, establishing Cray system installations in Sweden and Norway and serving as general manager for those countries.

"Several East Asian countries have developed substantial industrial capacity in automobile design and manufacturing, electronics, petroleum, and energy production — all market segments with a demonstrated need for supercomputers," said Marcelo Gumucio, Cray's executive vice president of marketing. "Weather forecasting is also an area in which there is much interest. In addition, the Asian academic community has emerged as a strong potential user of Cray systems. We expect that the Far East will develop into a meaningful market for Cray Research in the coming years."

Cray Asia/Pacific Inc. is headquartered at Suite 1300, Sutherland House, 3 Chater Road Central, Hong Kong; telephone: (852)-5-214669.

Cray opens new U.S. sales offices

Cray Research recently opened a new U.S. sales office to serve the Los Angeles area. The new office is located at 222 North Sepulveda Blvd., Suite 1422, El Segundo, CA 90245; telephone: (213) 640-8402.

Cray Research also has opened a new sales office in Darien, Connecticut. The new office can be contacted at P.O. Box 3423, 397 Post Road, Darien, CT 06820-0854; telephone: (203) 655-0854.

A third new U.S. sales office was opened in New Jersey. The New Jersey Cray Research sales office is at 150 Morristown Road, Suite 100, Bernardsville, NJ 07924; telephone: (201) 221-1611.

New orders reflect diverse market

In January Cray Research announced that the Naval Research Laboratory (NRL) of the U.S. Department of the Navy had exercised its contract option to upgrade from a CRAY X-MP/12 computer system to a CRAY X-MP/24 system. The new system was installed in the first quarter of 1987 at the Naval Research Laboratory in Washington D.C., and will be used for a range of research projects.

The Unisys Corporation, acting as prime contractor for the U.S. Navy, ordered a CRAY X-MP/22 computer system in January. The system is scheduled for installation during the second quarter of 1987 at the U.S. Department of Navy's David Taylor Naval Ship Research and Development Center in Bethesda, Maryland.

Also in January, Cray Research announced that the General Electric Aircraft Engine Business Group had acquired a CRAY X-MP/28 computer system with an SSD solid-state storage device. GE is using the system in the design and engineering of advanced technology products within their diversified business structure. The system was installed during the fourth quarter of 1986 at the customer's headquarters in Cincinnati, Ohio.

Cray Research confirmed in February that Honda Motor Company, Ltd., of Tokyo, Japan, had ordered a CRAY X-MP/12 supercomputer. The system will be installed in the second quarter of 1987 at Honda's research and development facility in Tochigi, Japan. Honda will use the system for vehicle research, development, and design, primarily in the areas of structural analysis, aerodynamics, combustion modeling, and crash analysis. This system will be the seventh Cray system installed in Japan.

In February Cray Research also announced that the Commissariat a l'Energie Atomique (CEA) had ordered a CRAY X-MP/416 computer system. The system will be installed in the second quarter of 1987 at CEA's computer facility in Limeil, France, pending export license approval. This system is CEA's third Cray supercomputer and will be used for scientific research.

Also in February, Cray Research announced that Grumman Aerospace Corporation installed a CRAY X-MP/14 supercomputer with an SSD solid-state storage device. The system was installed in the first quarter of 1987 at Grumman's scientific computer facility in Bethpage, New York. The system replaces a CRAY-1 M/2300 computer system installed in 1983 and will be used for engineering and scientific applications needed in the design of complex aircraft systems. These applications include computational fluid dynamics, structural analysis, image processing, and aircraft mission analysis.

Cray releases vectorizing C compiler

Release 2.0 of the Cray C compiler is now available. This compiler is based on the Portable C Compiler from AT&T. This release is the initial implementation of vectorization in the Cray C compiler.

New features added in this release include

- Vectorization of FOR loops. Benchmarks run for an electronics industry customer yielded performance improvements of up to 35 times that of the previous, nonvectorizing C compiler release.
- Common subexpression elimination. A compile-time option allows users to specify an optimization level.
- CPU targeting. A compile-time option allows generation of code for hardware with extended memory addressing capability, allowing memory references up to addresses requiring a full 24 bits.
- Long variable names. Implemented for the UNICOS operating system, this option allows users to use variable names of up to 255 characters, in 2 cases.

The C preprocessor, cpp, is included as a part of the Cray C compiler. The preprocessor allows macro substitution, conditional compilation, and the inclusion of name files in the compilation process.

Release 2.0 of the Cray C compiler will execute on any CRAY X-MP or CRAY-1 computer system. The compiler translates C language into assembler instructions that make effective use of the target Cray computer system. The Cray C compiler that executes on CRAY-2 computer systems is released as part of the UNICOS operating system.

UNICOS release 2.0 now available

Release 2.0 of the Cray operating system UNICOS enhances this powerful, interactive operating system designed to run on Cray computer systems. UNICOS 2.0 adds enhanced capabilites to maximize the performance of Cray computer systems and reflects Cray Research's commitment to provide its users with software technology that takes full advantage of Cray hardware.

UNICOS 2.0 includes the following enhancements for all Cray systems.

- Multitasking is now supported.
- Many enhanced or new user and administrator commands have been added to UNICOS.
- Enhancements have been made to online diagnostics maintenance-level tests.
- File system allocation has been improved to allocate contiguous disk space, to allow placing files on specific devices, and to allow placing a single file on several devices.
- The critical path for system calls and interrupts has been shortened, increasing system performance.
- The SCP protocol, linking UNICOS to front-end stations, is provided.

This release includes the following enhancements to UNICOS running on CRAY-2 computer systems.

- Batch job queues can be defined by maximum memory size or time limit.
- Error recovery and error logging have been improved on the DD-49 disk drives.
- Floating-point tables include improved reciprocal and square root functions.
- System accounting and related administrator commands have been improved to provide better information and an account ID for billing.

The following enhancements have been made for UNICOS running on CRAY X-MP and CRAY-1 computer systems.

- Users can now specify an alternate target architecture.
- The deadstart parameter file can be used during UNICOS system initialization to specify the number of CPUs, memory size, and the number of buffers in the buffer pool.
- Recovery from checkpoint or "drop" files is provided.
- A system activity report utility is provided, which can be used to monitor system performance.

For current Cray customers running COS, Cray Research offers an environment and tools that are helpful when migrating from COS to UNICOS. The COS Guest Operating System (GOS) feature (available with COS 1.15 and later releases) enables CRAY X-MP multiprocessor users to run COS and UNICOS concurrently. The GOS feature, the migration tools, and the special migration support team help users make a smooth and gradual transition to UNICOS.

APPLICATIONS IN DEPTH

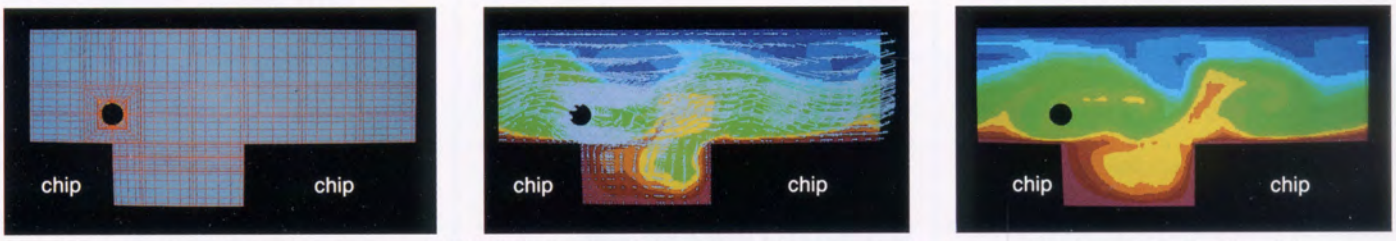

Graphic examples illustrating the use of NEKTON in the design of cooling systems for arrays of high-power-density chips: spectral element mesh for one periodic length of the array (left), flow patterns and temperature color fills for a Reynolds number of Re = 525 (center), temperature field at a slightly later time, demonstrating the unsteady nature of the flow (right). The cylindrical eddy promoters are placed to destabilize the flow, creating highly efficient unsteady heat transfer from the chip surfaces to the coolant stream. These NEKTON-generated results have been experimentally verified.

NEKTON models heat transfer for electronic design

Nektonics, Inc., has optimized its NEKTON fluid flow/heat transfer package, NEKTON, for Cray computer systems. NEKTON is applicable to coupled solid-fluid problems associated with the cooling and thermal control of high-power-density microelectronic components and systems.

NEKTON simulates incompressible and steady and unsteady natural convection flows in internal, external, and periodic two- and three-dimensional domains. User interface capabilities include menu-driven interactive mesh generation and data input; streamline, pressure, and temperature contour plots on arbitrarily defined slices of the three-dimensional domain; velocity, pressure, heat flux, and temperature surface maps and profiles; and time histories of velocity, pressure, and temperature. All pre- and post-processing features are implemented in full-color three-dimensional graphics.

NEKTON solves the full Navier-Stokes and energy equations by the spectral element method, a high-order technique that combines the geometric flexibility of the finite element methods with the rapid convergence rate of spectral techniques. The spectral element method

- Obtains high accuracy with relatively few grid points due to high-order interpolation, minimal numerical diffusion and dispersion, and automatic clustering of points near domain boundaries
- Accurately describes curved boundaries using high-order isoparametric elemental mappings
- Uses advanced iterative solution procedures that economize on storage and CPU use, and exploit a high degree of vectorization and parallelism

For more information on using NEKTON with Cray computer systems, contact Brian M. McCay at Nektonics, Inc., P.O. Box 22, Bedford, MA 01730-0001; telephone: (617) 275-4011.

HSPICE models analog circuits

The HSPICE analog circuit simulator is a widely used industrial-grade version of the Berkeley SPICE program. HSPICE is distributed and supported by Meta-Software, Inc., of Campbell, California, which has made the program available for use on Cray computer systems.

HSPICE simulates the electrical behavior of circuits using elements such as transistors, diodes, resistors, capacitors, inductors, and transmission lines. For sources, the program allows pulse, sinusoidal, exponential, piecewise linear, and single-frequency FM. Controlled sources include linear and nonlinear voltage and current-controlled voltages and currents. Using these it can perform various circuit analyses, including worst-case analysis, transient analysis, DC operating point, DC sweep, DC sensitivity, DC small-signal transfer function, AC small-signal analysis, Fourier analysis, and analyses of temperature, noise, and distortion. The program uses the Newton-Raphson solution technique to provide maximum modeling accuracy and flexibility.

HSPICE provides improved convergence and stability, additional analysis modes, engineering and software support, and extended modeling capabilities. HSPICE offers a rich modeling technology. It includes 25 MOS IDS model equations; improved subthreshold and noise modeling; a large choice of mobility, channel length modulation, and capacitance equations; and three choices of parasitics. Models include the Aspec models, Huang and Taylor depletion model, BSIM, Casmos, special HSPICE MOS models, vertical and lateral BJT models, soft and hard Zener diode breakdown, two gallium arsenide models, and others. Meta-Software provides model parameter sets based on lab measurements and optimized by the SUXES-10 program. Lab measurements are derived from the special MetaTest-Chip, providing complete and highly accurate models.

For more information on using the HSPICE analog circuit simulator with Cray computer systems, contact Mike Smith, Meta-Software, Inc., 50 Curtner Ave., Suite 16, Campbell, CA 95008; telephone: (408) 371-5100.

I-G SPICE available on Cray systems

Interactive Graphics SPICE (I-G SPICE), a version of the widely used SPICE 2 circuit simulation program, is offered for Cray systems running COS or CTSS. I-G SPICE offers all the standard SPICE capabilities including DC and AC analysis; transient, noise, sensitivity, Fourier, and distortion analysis; and analysis of temperature sweep, transmission lines, transfer characteristics, and driving point impedance.

I-G SPICE gives engineers the flexibility and speed of designing on-line. Updated and enhanced since 1978, version IV includes the following features

- ☐ User-defined equations and tabular functions
- ☐ Multi-plot capability that allows the user to plot multiple independent variables on the same axis
- ☐ User-defined Fortran subprograms for complete modeling flexibility
- ☐ Built-in digital gates and functions for true mixed-mode analysis
- ☐ User-defined parameters that allow input of any variable or combination of variables. Any function can be plotted against any other function
- ☐ Automatic reruns to allow any circuit and model parameters to be changed for as many reruns as desired
- ☐ Worst-case calculations based on the ± range of component values to evaluate the worst deviation from specified values of all combinations
- ☐ Monte Carlo capability to allow complete statistical analysis of DC, AC, and transient responses
- ☐ Optimization to help designers assign ideal component values to optimize component value combinations for specific circuits

I-G SPICE can be used as a stand-alone package or may be integrated into a workstation. Output is available on graphics terminals or can be channeled to all types of output devices. Documentation is furnished, including a 500-page user's guide and a 350-page government report illustrating in detail a comprehensive I-G SPICE example. The fully illustrated user guide is supplemented by a 22-hour videocassette seminar. AB Associates furnishes the complete Fortran source code to help companies conveniently add their own special features. I-G SPICE has been structured to accom-

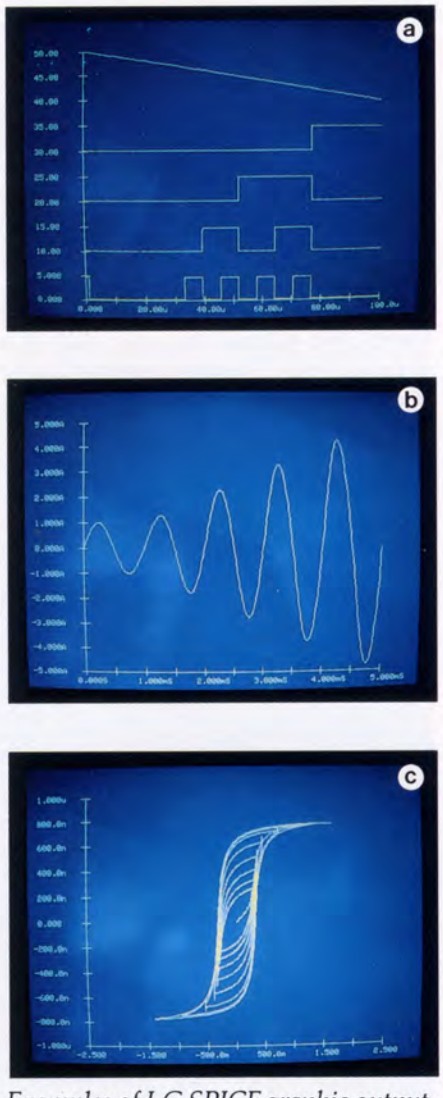

Examples of I-G SPICE graphic output: (a) hybrid analog-to-digital converter using built-in digital functions, (b) variable inductor test circuit, (c) switching regulator analysis.

modate updates and enhancements as they become available.

For more information on using I-G SPICE with Cray computer systems, contact AB Associates Inc., P.O. Box 82215, Tampa, FL 33682; telephone: (813) 932-9853.

PRECISE circuit simulation package offered for Cray systems

PRECISE is a proprietary software product of Electrical Engineering Software, Inc., for circuit simulation. Now available for use on the CRAY-1, CRAY-2, and CRAY X-MP systems, PRECISE is a comprehensive tool for simulating circuits and performing tests and analyses that cannot be performed efficiently by breadboarding or prototyping. Aerospace and semiconductor companies currently use PRECISE in analog circuit design.

PRECISE allows users to make adjustments, test alternatives, and work at a detailed level. The user creates a circuit description and a simulation is performed in the engineer-oriented language. System controls and commands can be utilized in an interactive or background mode. Analysis and results can be examined immediately, allowing the user to modify conditions and continue.

PRECISE features include

- Interactive and background execution
- Correct numerics
- Improved convergence
- Sweep variables across ranges
- Device parameter alterations to vary process
- Output graphics
- SPICE compatibility
- Full-function capability
- External model and subcircuit files
- Discrete library options
- User-defined device modeling option

PRECISE is designed to work with active devices such as diodes, bipolar transistors, JFETs, MOSFETs, and GaAs models and allows the standard types of simulation and analysis to be performed. Compatible networks allow saved versions of simulations to be transferred between machines saving engineering and processor time. This combines with the rest of the PRECISE features to shorten design cycles and improve circuit performance.

For more information on using the PRECISE circuit simulation package on Cray systems, contact Electrical Engineering Software, Inc., 4675 Stevens Creek Blvd., Suite 101, Santa Clara, CA 95051; telephone: (408) 296-8151.

The CAD Group Inc. offers SALT logic simulation and timing analysis on Cray systems

SALT™ (Simulation Analysis of Logic and Timing) is a design automation tool that provides engineers with the ability to simulate concurrently all levels of design hierarchy in a mixed mode (analog/digital) environment. SALT runs under the Cray operating systems COS and UNICOS. Support for the CTSS operating system is currently in development.

Two other modules, SHDL and PFG, interface with SALT to provide additional capabilities. SHDL (SALT Hardware Description Language) provides the ability to quickly develop behavioral models which may then be added to SALT. SHDL also can translate behavioral models into gate structures. PFG (Probabilistic Fault Grading) provides a fast fault grading capability.

SALT quickly analyzes the behavior of both digital and analog designs of any combination of a multitude of hierarchical levels. A Cray computer system allows SALT to perform tasks in a fraction of the time required on other systems. Current SALT development will further optimize the code to take advantage of vectorization on Cray systems.

SALT can be run either in a batch or interactive environment. The multiple checkpoint restart feature allows simulation to continue from a restart point after changes have been made, saving users from having to restart simulations from the beginning. Other important features include timing analysis, minimum/maximum worst case timing analysis, and tester mode (output assertions).

For more information on using SALT with Cray computer systems, contact Vinnie Apicella, The CAD Group Inc., 3911 Portola Drive, Santa Cruz, CA 95062; telephone: (408) 475-5800

Information/data analysis packages available on Cray systems

SPSS Inc. of Chicago, Illinois now offers two products, SPSS-X and its table-producing option, SPSS-X Tables, for use on CRAY-1 and CRAY X-MP computer systems.

The SPSS-X Information Analysis System is a file and data management package that offers more than 40 statistical procedures and report writing functions. The SPSS-X package for statistical and data analysis serves as a

- Data management tool to form and select information variables for analysis
- File management tool for flexible and efficient file handling
- Statistical tool to produce simple tables or plot complex multivariate functions
- Modeling and forecasting tool to build models, identify trends, and predict future events
- Multiple application tool for communications, marketing, research, and administration

The SPSS-X Tables option produces publication-quality tables that combine large amounts of information and obey simple commands in flexible formats and functions. Users have complete control over table display and a wide range of statistical variables can be incorporated to communicate results to a variety of audiences.

For more information on using SPSS-X or SPSS-X Tables with Cray computer systems, contact SPSS Inc., 444 North Michigan Ave., Chicago, IL 60611; telephone: (312) 329-3500.

USER NEWS

Cray systems are a winner on the water

Stars & Stripes '87, the U.S. yacht that won the right to challenge Australia for the America's Cup, sailed to victory on February 4 in races off Fremantle, Australia. The 12-meter yacht was designed with the help of supercomputer simulations that were run during the past two years on CRAY X-MP/48 and CRAY X-MP/216 systems at Cray Research's Mendota Heights, Minnesota computer center.

Skippered by Dennis Conner, the man identified early on as the favorite to challenge the Australians, *Stars & Stripes* won the best-of-seven series in four straight races over *Kookabura III*, the Australian defender. Conner's yacht defeated the yacht *New Zealand* in the challenger finals by a margin of 4-1, and defeated Tom Blackaller's *USA* 4-0 in the challenger semifinals.

Conner and his crew won only about 80 percent of their races in the round robin series that selected the four challenger semifinalists. In fact, *New Zealand* bested *Stars & Stripes* in two out of three meetings during the round robin series. But the early round robin races were in light air (moderate winds), while the designers of *Stars & Stripes* had chosen to optimize the boat for the much stronger winds and heavy seas anticipated to begin during the late round robins and semifinals. The doctor came as expected — that is, the *Fremantle Doctor*, the name locals have given to the strong, cooling breeze that blows off the Indian Ocean during the hot months of the Western Australian summer. *Stars & Stripes*, predicted to be at her best in winds above twenty knots, went on to defeat her competition handily.

The boat's new-found speed also owed to some modifications made before the challenger semifinals. In the early fall of 1986, *Stars & Stripes'* designers used the Cray system to model new winglets (small horizontal stabilizing fins mounted to the keel that also help to reduce drag) and a new rudder for the yacht. The new components were shipped to Fremantle and installed on the boat without any experimental testing. During the design of the full boat, the designers had become so confident in their computer models that they trusted the computational results over tow tank testing, the usual method used to validate boat designs. *Stars & Stripes'* improvement also was attributed to an experimental drag-reducing film that was applied to the hull of the yacht. Because the tiny grooves in the film had to be aligned with the direction of local water flow around the boat hull, large prints were made of some of the graphics generated by the Cray system to serve as a guide for applying the film. The film was developed by 3M Co. and was based on research conducted at NASA for reducing drag on aircraft.

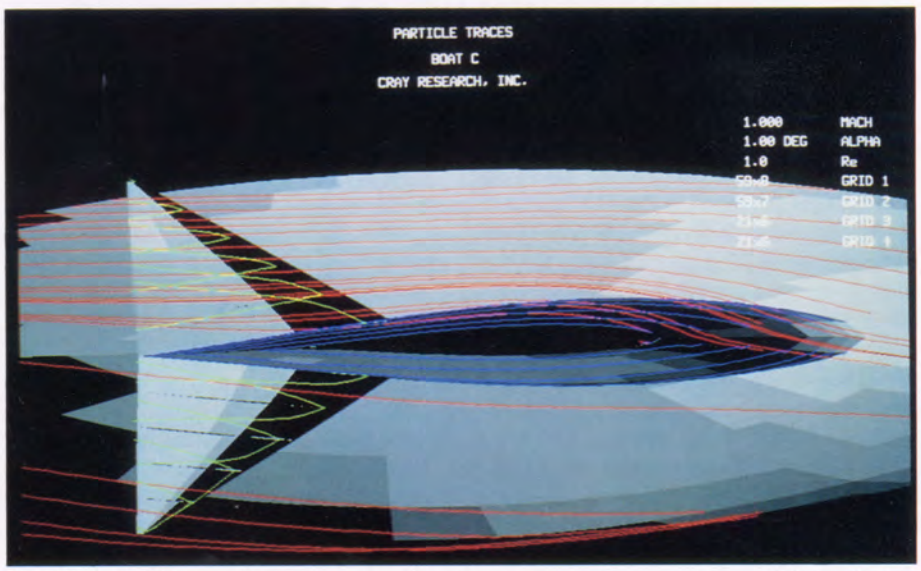

The Stars & Stripes *design team made extensive use of computer models to design the trophy-winning yacht. This image shows a result of calculations run on a Cray computer system to predict flow around the boat's keel and winglets.*

Experienced observers of the sport will serve notice that a fast boat is only part of what it takes to be a winner. Dennis Conner is believed to have logged more hours at the helm of a 12-meter than any other living skipper. He also had a crew that had sailed with him almost daily for the last 18 months. The yacht *USA*, also designed with the help of a Cray system, lost to *Stars & Stripes* in the semifinals, but actually matched or outpaced *Stars & Stripes* on 13 of the 16 legs of the first two races. *USA* sported a radical forward rudder and eliminated the traditional keel, making it very tricky to sail. Heiner Meldner, chief computer scientist of the *USA* design effort, felt that *USA* could have defeated *Stars & Stripes* if the crew had had more time to learn how to handle the unusual boat. *USA* entered the water in June 1986, a full two years after the *Stars & Stripes* crew began shakedowns and training on trial yachts.

CRAY CHANNELS featured an article about the design of *USA* and *Stars & Stripes* in its Summer 1986 issue. Details of the designs were limited, however, because of concerns over highly secret aspects of the new yachts. But, because the winning syndicate must release the design of its boat for public scrutiny, many of the secrets of *Stars & Stripes* may be revealed in an upcoming issue of CRAY CHANNELS.

The circuitous design ornamenting AT&T Bell Laboratories' CRAY X-MP/24 computer system comes from a microprocessor designed on the labs' earlier CRAY-1 system.

High-tech attire makes this system unique

A computer system's appearance may not be the main concern of its users, but there is something to be said for adding that personal touch. The CRAY X-MP/24 computer system at AT&T Bell Laboratories in Murray Hill, New Jersey, was given its fancy window dressing by circuit designers "who wanted to individualize the system with a decorative motif relevant to work being done at the lab," explains John Drobny, a senior designer in Bell Labs' art department. "The pattern is from part of an AT&T WE 32000 Microprocessor, a chip that was designed on our CRAY-1 system."

Drobney said the mural generally draws favorable comments from passersby. "A lot of people mistake it for a city skyline," he notes. "I've seen people point to it and say 'There's the World Trade Center' or 'There's the Empire State Building.' Even technical people who come by scrutinize it pretty closely trying to figure out what it is." The Bell Labs system may be the most ornate to date, but if custom paint jobs for supercomputers catch on, who knows what we'll see next. Racing stripes anyone?

Plasma-assisted CVD research

Plasma-assisted chemical vapor deposition (CVD) is a process widely used in the microelectronics industry to deposit thin solid films. Plasmas are also used to remove films, a process called etching. Plasmas used for CVD and etching are low-pressure gases that are weakly ionized by radio waves. These plasmas resemble those used in neon signs and fluorescent lamps; the major difference being that gases used in microelectronics manufacturing processes are specialized and highly purified. The chemical action of the plasma results from the action of highly energetic electrons, freed during the ionization process. These energetic electrons collide with gas molecules and break them into smaller molecular fragments. The molecular fragments, usually very active chemically, then find their way to surfaces where they react chemically to form a film.

David Graves, assistant professor in the department of chemical engineering at the University of California, Berkeley, is simulating plasma-assisted CVD to better understand the physics and chemistry involved. Graves is using a CRAY X-MP/48 computer system in-

stalled at the San Diego Supercomputer Center with a satellite hook-up to the Remote Users Access Center (RUAC) on the Berkeley campus.

Graves is focusing on the movement of electrons and ions and the electrical fields associated with them that make up the plasma. The electrons and ions are treated as continuous fluids in the mathematical model Graves is using, a strategy which resulted in the most complete description of the plasma behavior to date. To demonstrate that the mathematical model is accurate, Graves and the graduate students working with him are simultaneously making measurements on laboratory plasmas and comparing the measurements to the model predictions.

Graves currently uses a simulation in which the chemistry is greatly simplified. This makes the calculations much simpler, but yields a less realistic representation of plasma-assisted CVD. In addition, the current model assumes that plasma properties depend only on a single dimension in space. Most industrial reactors depend on two dimensions. As a result, Graves and coworkers are changing the model to incorporate realistic chemistry and two spatial dimensions.

The one-dimensional simulation itself involves vast calculations because of variations in time as well as space. Computational speed and large on-line memory capacity are the main advantages of using the CRAY X-MP/48 system in these simulations.

Model equations are solved numerically on the CRAY X-MP/48 using a technique known as the spectral-Galerkin finite element method. The radio waves applied to the plasma are periodic in time, and this forces all plasma properties to be time periodic as well. The spectral-Galerkin method takes advantage of this time periodicity and is capable of resolving rapid variations in space by virtue of its ability to refine the finite element mesh locally.

Researchers are still trying to understand the intricacies of plasma-assisted CVD, but Graves' research into its physics has provided some insight. Recent simulations on the CRAY X-MP/48 have predicted a pattern of light emission that has been observed in laboratory plasmas by several other research groups. This work is significant because regions of light emission correspond to regions where the molecular dissociation leading to film deposition takes place. Graves' findings will be reported in the *Journal of Applied Physics*.

"The Cray system is an absolutely necessary component. It is the central part of the research," Graves said. If more detailed simulations can be successfully compared with experimental results, then engineers using CVD will be able to design new and better processes based on a solid understanding of the principles underlying plasma-assisted CVD. This is particularly important in light of the intense international competition in the semiconductor industry.

Bell Labs seeks a better IC process

In the quest for high-performance integrated circuits, electronic designers pack transistors so densely that unexpected problems sometimes result. Researchers at AT&T Bell Laboratories are using the labs' CRAY X-MP/24 computer system to study the performance limits of extreme miniaturization. Mark Pinto, a member of the labs' technical staff, uses the Cray system to model a phenomenon called latch-up, which electrically grounds a circuit, destroying not only its logic state, but also sometimes the physical chip itself. Latch-up occurs when transistors are placed so near each other that an electric pulse, from a power or radiation source for example, can initiate and sustain a high current through the structure separating the transistors. The charged structure grounds the circuit, sometimes literally burning it up. By simulating processing methods and device performance, Pinto and his coworkers have been able to suggest changes in semiconductor processing that will eliminate latch-up.

"The problem is usually avoided in the industry by simply providing enough separation between transistors," Pinto says. "But doing so compromises a circuit's performance by increasing interconnection lengths. Our primary interest is in modeling very high density digital CMOS circuits, in which two transistors form a single gate, so we want to keep the transistors as near to each other as possible. If we can devise processing methods that prevent latch-up, we can

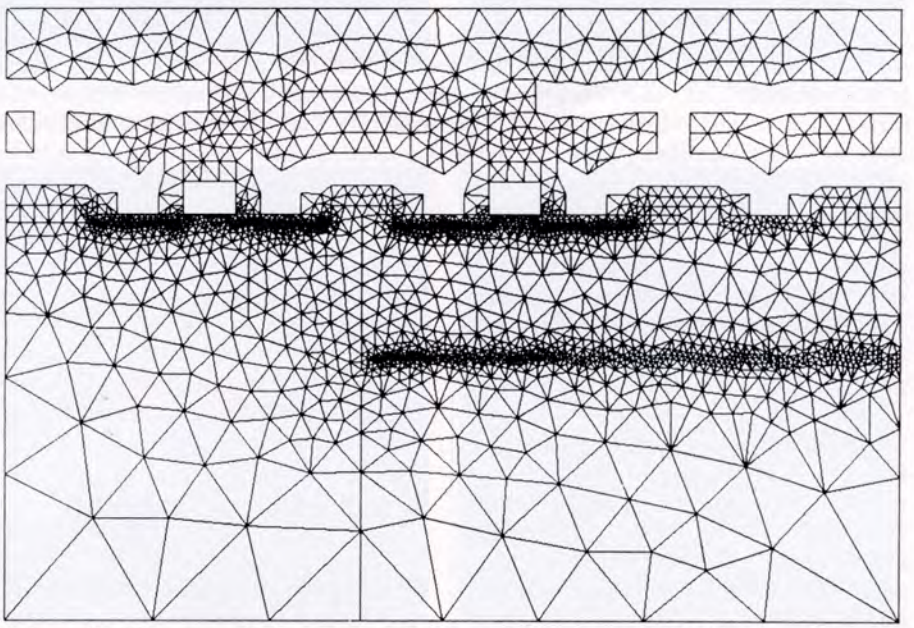

A large, highly nonuniform computational grid describing a CMOS device. This grid is typical of those used in the latch-up studies at AT&T Bell Laboratories.

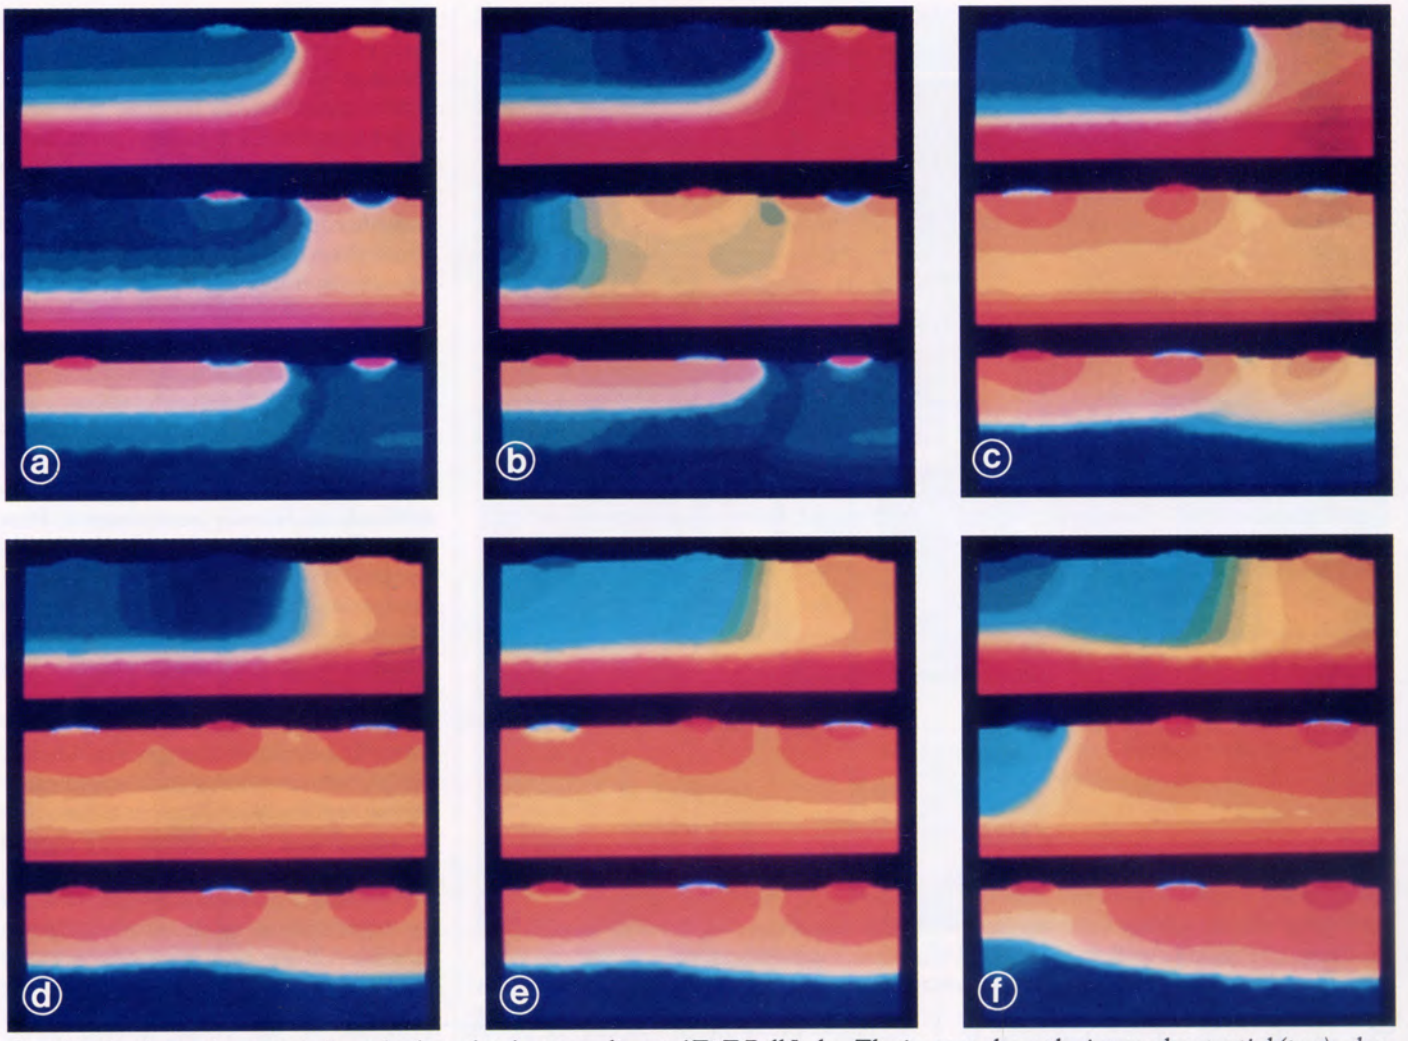

Graphic displays of computed results from latch-up studies at AT&T Bell Labs. The images show the internal potential (top), electron distribution (middle), and hole distribution (bottom) in a CMOS device responding to a voltage pulse. Red represents the largest potential and most concentrated electron and hole distribution; blue represents the smallest potential and least concentrated electron and hole distribution. The sequence shows the device (a) in its normal state, before the pulse is applied, (b) just after the pulse is applied, (c) later, with the pulse still applied, (d) as the pulse is terminated, (e) some time after the pulse is terminated, and (f) later still, in its steady state. The final image shows the device latched; the two terminals are shorted together. A device that has reached this state will no longer function.

proceed to further levels of miniaturization. But we can study this phenomenon only through computer modeling. And only the Cray system has the power we need to run the models."

The research involves first modeling the fabrication steps used to make an integrated circuit. Then the electrical characteristics of the transistors (devices) and the structures between them are modeled according to the results of the fabrication model. The size and complexity of a device model are determined by the device's operating conditions. Modeling MOSFETs, for example, requires considering only electrons moving along very restricted paths. Modeling phenomena such as latch-up is more complicated, requiring that the model include electrons and holes moving in many directions.

"Our models are primarily two-dimensional, but to optimize a structure to avoid latch-up, even in two dimensions, requires using the Cray system," says Pinto. "At this point, our three-dimensional modeling is restricted to basic research, not production work. But we ultimately have to consider three-dimensional effects, so this work is very important. The Cray system is even more essential for these models."

Pinto's process and device simulations make extensive use of the Cray system's vectorization and gather/scatter capabilities. His work also relies on a sparse matrix solver written at Bell Labs. "The solver is not improved much by running in vector mode," Pinto says. "But the optimization for the Cray system, the special solver, and the fast scalar processing give us a performance on the

Cray system that is at least 200 times that of our VAX mainframes for this application."

Varying the amount of dopant applied to different parts of a wafer might be a solution generally applicable to the latch-up problem, Pinto suggests. If doing so turns out to be a useful approach, modeling can provide the necessary quantitative information. At this point the researchers are analyzing many processing methods to discover which parameters affect circuit performance and how. Once this information is applied, latch-up may no longer be an obstacle to miniaturization. By using Cray systems to model phenomena that could not be studied in any other way, Pinto and his colleagues may enable circuit designers to exploit fully the high-performance potential of very small circuitry.

Computer-aided forging helps cut corners

Computer modeling can reduce the cost of trial-and-error testing in virtually any industrial design or engineering environment. The cost savings result from replacing burdensome physical experiments with mathematical models that a computer can test quickly and repeatedly until the best procedure or design is found. Not only is this method usually cheaper and faster than conventional trial-and-error methods, but it also can reveal important information unobtainable from physical experiments.

Innovative studies of computer modeling to solve problems in metal forging were recently carried out at Lawrence Livermore National Laboratory in Livermore, California. Laboratory researchers Elane Flower and John Hallquist used a CRAY-1 computer system to simulate metal stock, dies, and forging processes, achieving results that solved practical forging problems.

Using software developed at the Livermore lab, Flower and Hallquist tackled a problem encountered by Precision Forge of Oxnard, California. The company had had problems with the final dimensions of a cup forging produced in six stages from a stainless steel alloy. Several attempts were made to obtain proper die fill during the final stage without creating a "kink" on the inner free surface. Design changes were evaluated by machining the fifth-stage forging and changing the sixth-stage die to alter its closure. But modeling the same components using the finite element program NIKE2D on the Cray system turned out to be more efficient than traditional trial-and-error tests.

To model the sixth forging stage, a computational mesh was generated for the die and the fifth-stage forging. Material properties were assigned to both, boundary conditions were established between the die and forging, and a displacement history was given. Values for the part dimensions, stress/strain, strain rate, coefficient of friction, and die closure were handled by an input generator, and values for each element of the mesh were computed at every 0.01 inches of displacement. The problem required approximately 20 minutes of CPU time on the CRAY-1 system. Analysis of the computed stresses and strains showed a lack of die fill in the upper outer corners of the die. The analysis also showed good correlation between regions of large plastic strain and high shear stress where kinking occurred.

From the modeling, the researchers also obtained interface pressures that correspond to load requirements for deformation. This information generally is calculated less accurately from estimates of flow stress and geometry. But com-

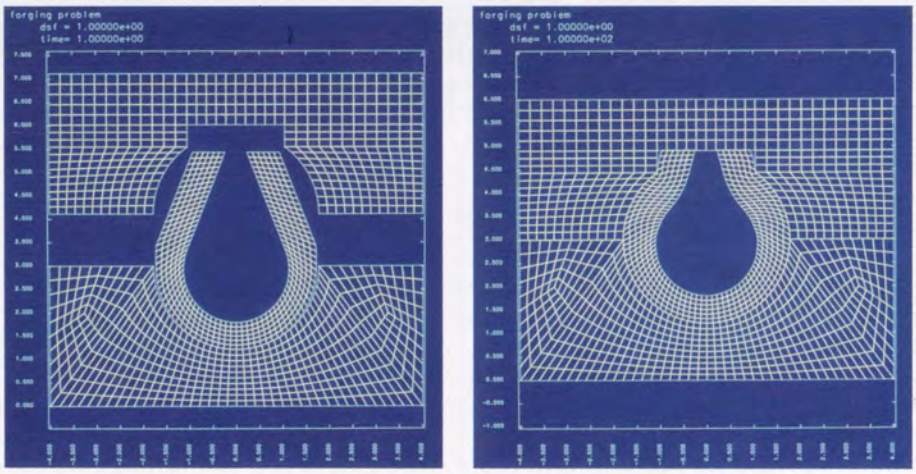

Computational mesh of fifth-stage forging and sixth-stage die (left) used to model cup-forging process and deformed mesh after modeling (right).

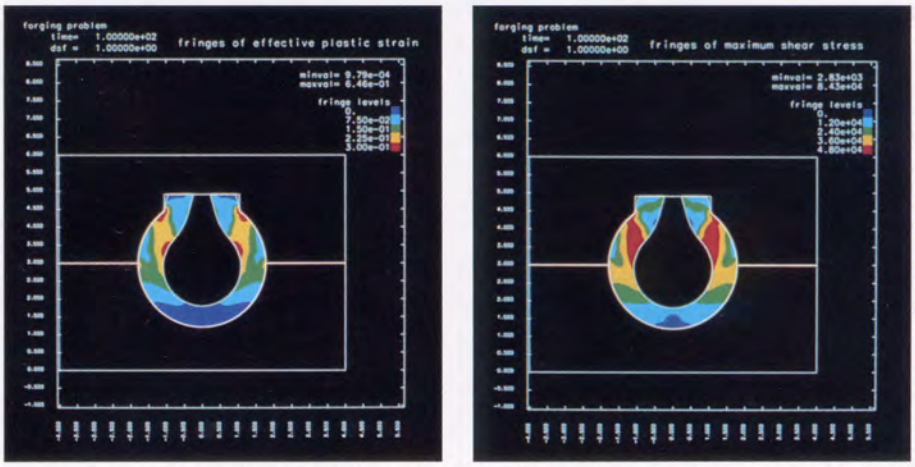

Computed regions of strain (left) and stress (right) from sixth-stage cup forging model.

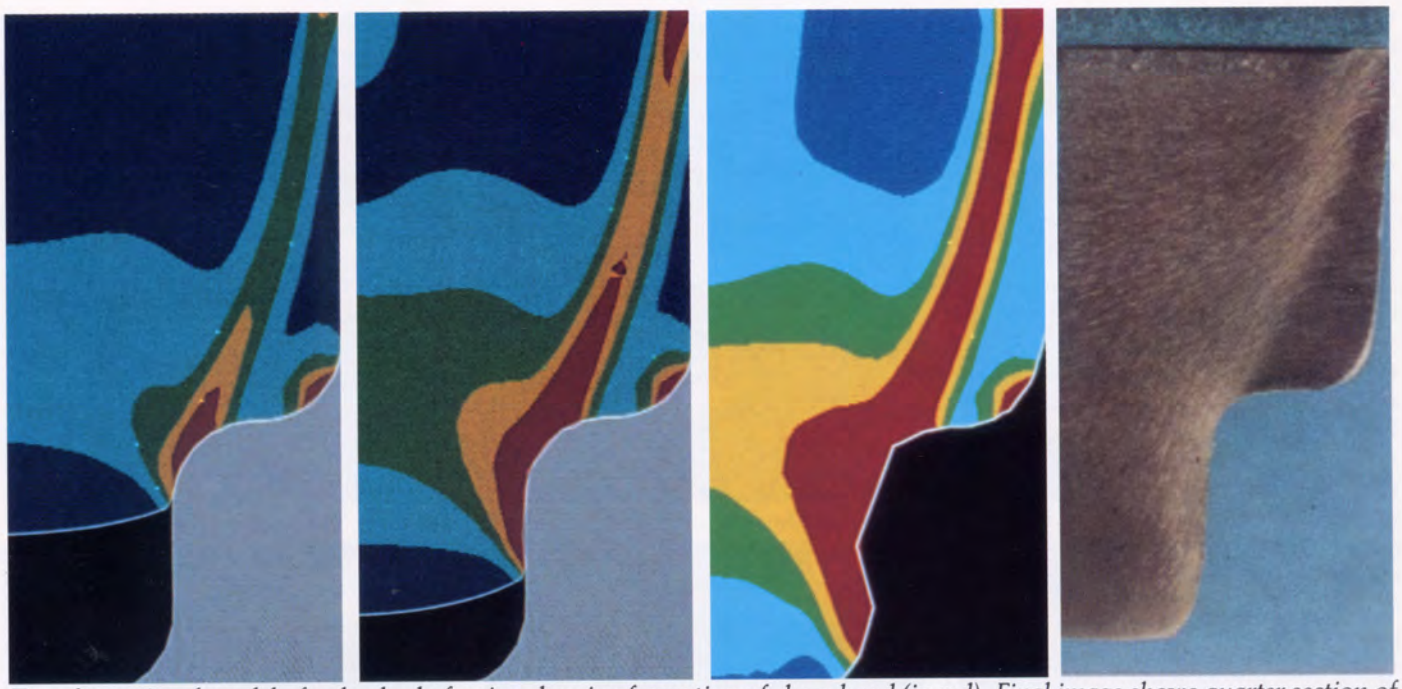

Two-dimensional model of valve body forging showing formation of shear band (in red). Final image shows quarter section of actual forged part. Computer modeling could have been used to predict accurately the formation of shear bands in these parts.

puter modeling provided more accurate values based on exact geometry, stress as a function of strain, and friction.

Another problem tackled by the Livermore lab researchers involved finding the cause of shear bands in a valve body forging. The forged part required a horizontal grain flow, which was achieved during die development using a plate for starting stock. However, the plate stock was found to be of insufficient purity, so for production a bar of the required purity was placed horizontally in the die.

But the change in starting stock geometry changed the flow characteristics of the metal during forging and produced parts that contained shear bands. *Shear bands* are regions of inhomogeneity within the forging, which form at an interface where the metal is flowing at different strain rates. Because of these defects, the forgings were rejected.

Special forgings were then run to find the causes of the shear bands. From these runs, bar length was determined to be the primary factor correlated with the presence of shear bands, though the reason for this was unknown. Computer modeling using NIKE2D on the Cray system revealed the reason bar length affected flow localization, resulting in shear bands, and why bar stock created a problem where plate stock did not.

The first analysis did not account for friction and did not reveal any localized strain. When friction was included in the calculations, however, a region of high strain developed that corresponded to the shear-band region in the actual forging. This led to the realization that friction was an increasingly dominant force that accentuated localized strain as bar length increased. The analysis with friction took about three CPU minutes on the CRAY-1 computer system.

Additional three-dimensional analyses of the problem using the code DYNA3D were assigned a ram velocity of 600 cm/sec. These studies corroborated the results of the two-dimensional studies. The DYNA3D analysis took approximately 5 hours on the CRAY-1 computer system. Higher ram velocities would require proportionately shorter CPU time; lower ram velocities, longer time. As a result of the computational analyses, the shear band problem was solved by minimizing bar length.

"These sorts of production problems are handled much easier and faster on a computer than they are with the usual methods," said Flower, a metallurgist at the Livermore lab. "Although the metal forming industry is moving toward computer modeling, there is still a lot of trial-and-error physical testing. But it is slow and expensive to cut a die, try it, shave it, try it again, shave it again, and so on. For repetitive work like this, computers are much more efficient; there is no doubt that in time most metal parts will be computer-designed."

In addition to the cup forging and the shear band problems, other problems were solved computationally that would have required considerable time, labor, material, and destructive testing had they been studied in the usual way. As NIKE2D, DYNA3D, and other engineering analysis codes are refined and developed for use on supercomputers, computer modeling will be applied to increasingly large and complex engineering and design problems. As these metal forging examples indicate, industries that augment their traditional experimental methods with computer modeling stand to gain substantial time and cost savings advantages.

WORLDWIDE BUSINESS CENTERS

Domestic

Corporate Headquarters
608 Second Avenue South
Minneapolis, MN 55402
(612) 333-5889, telex 4991729

Government Marketing/
Government Relations Office
National Place
1331 Pennsylvania Avenue, Suite 1331 N
Washington, DC 20004
(202) 638-6000

Central Region Office
5350 Manhattan Circle
Boulder, CO 80303
(303) 499-3055
- Albuquerque
- Chicago
- Colorado Springs
- Detroit
- Minneapolis/St. Paul
- St. Louis

Eastern Region Office
4041 Powder Mill Road, Suite 600
Calverton, MD 20705
(301) 595-5100
- Atlanta
- Bernardsville
- Boston
- Cincinnati
- Darien
- Huntsville
- Laurel
- Pittsburgh
- Tampa

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

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

International

Cray Asia/Pacific Inc.
Suite 1003
Sutherland House
3 Chater Road Central
Hong Kong
(5)-214669, telex 89765

Cray Canada Inc.
4211 Yonge Street, Suite 210
Willowdale, Ontario
Canada M2P 2A9
(416) 229-2729

Cray Research France S.A.
7 Rue de Tilsitt
75017 Paris, France
(01)-766-01-55, telex 660568

Cray Research GmbH
Perhamerstrasse 31
8000 Munich 21, West Germany
(089)-56014-0, telex 5213211

Cray Research Japan, Limited
W 1154, Shin Aoyama Building
1-1 Minami-Aoyama 1-chome
Minato-Ku, Tokyo 107, Japan
(03)-403-3471, telex 28893

Cray Research S.R.L.
Via Vivaio, 15
20122 Milan, Italy
(02)-70-98-27, telex 315403

Cray Research (UK) Ltd.
Cray House
London Road, Bracknell
Berkshire RG12 2SY
United Kingdom
(0344)-485971, telex 848841

CRAY RESEARCH, INC.