A CRAY RESEARCH, INC. PUBLICATION

CRAY CHANNELS

Summer 1985

ANNOUNCEMENT!
CRAY-2 Computer System introduced

FEATURE ARTICLES:

DEPARTMENTS:

IN THIS ISSUE

When the CRAY-1 supercomputer was unveiled nearly a decade ago, it set new standards for scientific computer performance. Now those standards have been exceeded by an order of magnitude. This year we are pleased and proud to announce the CRAY-2 supercomputer, the newest computer to emerge from Seymour Cray's research laboratory in Chippewa Falls, Wisconsin. The CRAY-2's unique immersion cooling and three-dimensional circuit technologies mark a radical break from previous computer designs. These breakthroughs make possible the CRAY-2's unprecedented 4.1-nanosecond clock cycle and 256-million word memory. Important as these advances are, they are all the more impressive when one sees the CRAY-2's four powerful CPUs operating in a see-through cabinet merely 45 inches tall and 53 inches in diameter!

In this issue of CRAY CHANNELS we also explain multitasking, how to do it, when and why. A look at the roles supercomputers play in seismic data processsing and nuclear reactor safety analysis is also offered, as well as a profile of CUG, the Cray User Group. Our regular departments will take you to the second Cray science and engineering symposium, our expanding operations in Chippewa Falls and to Jupiter, for a computational look at the jovian red spot.

The development of the CRAY-2 supercomputer demonstrates Cray Research's commitment to excellence in scientific computing. When Keats wrote, "Beauty is truth, truth beauty," it's unlikely he was contemplating the arrival of the CRAY-2, but he might as well have. We hope you'll join in our excitement as we introduce this new research tool to the scientific community.

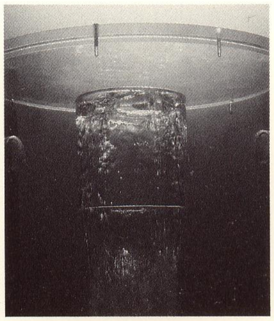

On the cover is the top of a coolant reservoir for the CRAY-2 supercomputer. The cascading liquid is the CRAY-2's inert fluorocarbon coolant being pumped through a standpipe. Th standpipe is used to maintain a coolant pressure close to atmospheric. The reservoir itself serves as a coolant holding tank when the mainframe must be serviced.

CRAY CHANNELS

A Cray Research, Inc. publication

Summer 1985

6

FEATURES

DEPARTMENTS

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

Volume 7, Number 2
Summer 1985

Editor
Carol Smith

Contributing Writer
Ken Jopp

Design
Cindy Erickson
Eric Hanson

INTRODUCING

The CRAY-2 Computer System

Curiosity about the CRAY-2 has been high ever since late 1981, when Seymour Cray announced a technological breakthrough enabling him to complete the system's design stage. On November 19th of that year, Cray told a news conference gathering about the key features of his new design and explained some of the unusual technologies chosen. Now, Cray Research is proud to announce the CRAY-2 Computer System — the first of a new generation of supercomputers developed by Seymour Cray.

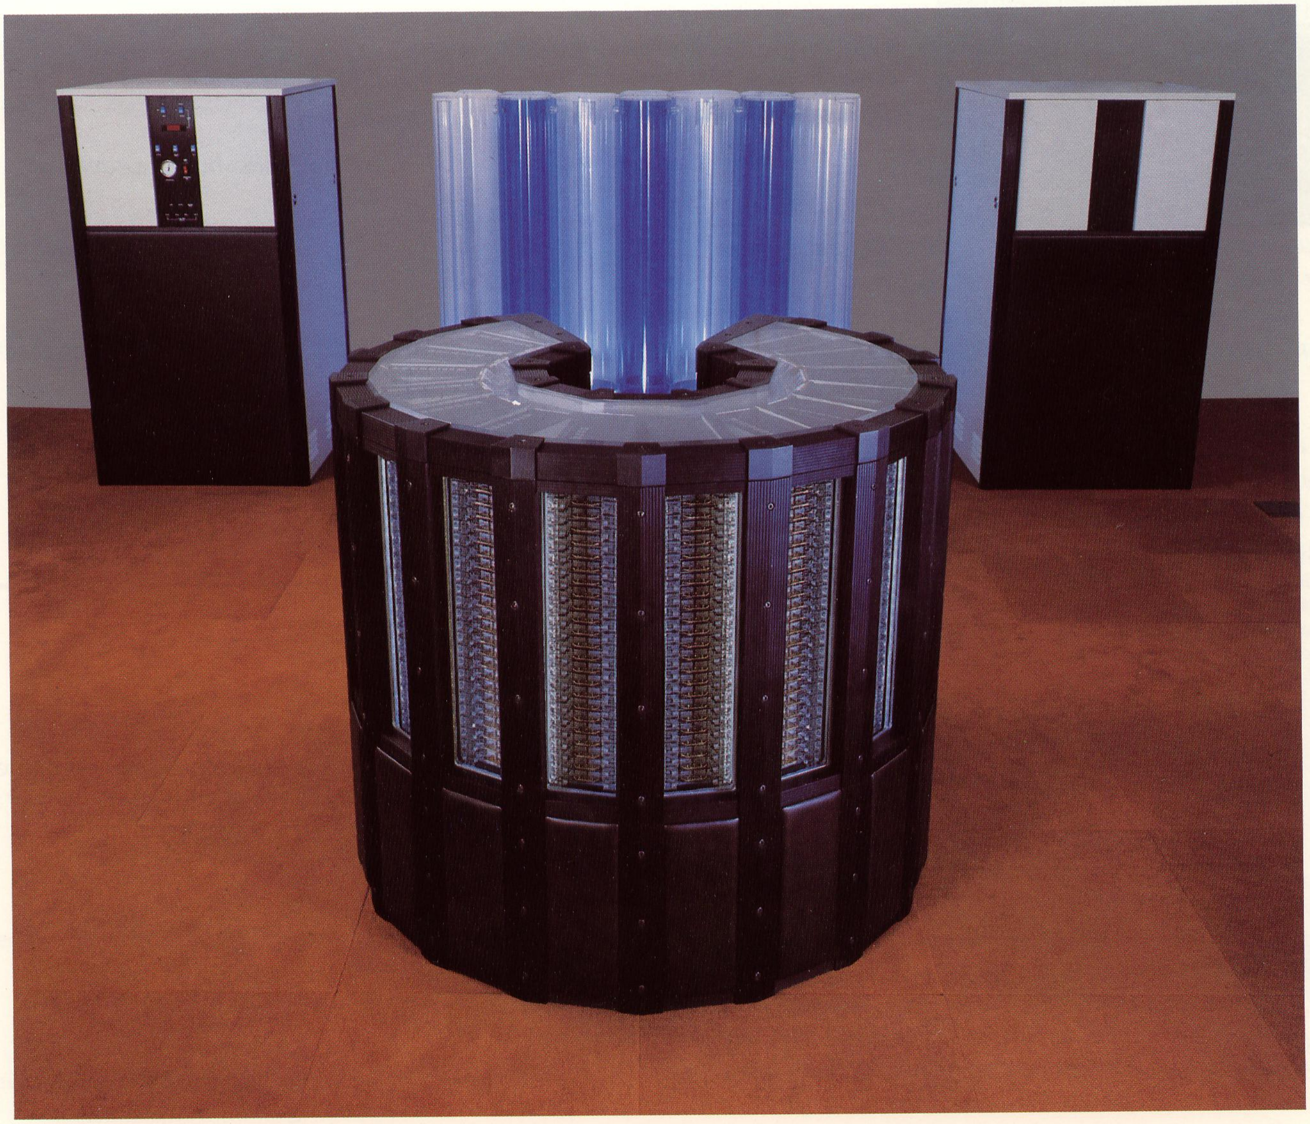

CRAY-2 architecture and design overview

A pioneer in the use of liquid immersion cooling technology, the CRAY-2 features a 256-million-word Common Memory, four independent processors and a 4.1 nanosecond clock cycle — all in a package just 45 inches tall and 53 inches in diameter. The CRAY-2 delivers effective CPU speed six to twelve times that of the CRAY-1 and runs an operating system based on AT&T's widely accepted UNIX™ System V.

The CRAY-2's extremely high processing rates result from its compact size, its balanced integration of scalar and vector capabilities and its large Common Memory in a multiprocessing environment. The significant architectural components of the CRAY-2 Computer System (Figure 1) include four Background Processors, each with a high-speed Local Memory, a Foreground Processor, 256 million 64-bit words of Common Memory and a maintenance control console.

Each of the four Background Processors contains registers and functional units to perform both vector and scalar operations. The high-speed Local Memory integral to each Background Processor is available for temporary storage of vector and scalar data. The single Foreground Processor supervises the four Background Processors, while the large Common Memory complements the processors and provides architectural balance, thus assuring extremely high throughput rates. The maintenance control console enables routine onsite maintenance.

Background Processors

Each of the four identical Background Processors is more powerful than a CRAY-1 computer, offering exceptional scalar and vector processing capabilities. The four processors can operate independently on separate jobs or concurrently on a single program. Clock cycle time on each is 4.1 nanoseconds — faster than any other computer system available.

Each Background Processor consists of a computation section, a control section and a high-speed Local Memory of 16,384 64-bit words. The computation section performs arithmetic and logical calculations. These operations and the other functions of a Background Processor are coordinated through the control section. Local Memory is used for temporary storage of scalar and vector data during computations. It replaces the B and T registers on the CRAY-1 and is readily available to user jobs.

Common Memory

The CRAY-2's large Common Memory is one of its primary technological advantages. Remarkable but true: one CRAY-2 system has more memory than the total of all CRAY-1 and CRAY X-MP systems installed to date! Common Memory on the CRAY-2 consists of 256 million 64-bit words randomly accessible from any of the four Background Processors and from any of the high-speed and common data channels.

Common Memory is arranged in four quadrants of 32 banks each, for a total of 128 interleaved banks. It is shared by the Foreground Processor, the four Background Processors and the peripheral equipment controllers. Each bank of memory has an independent data path to each of the four Common Memory ports. Each bi-directional Common Memory port connects to a Background Processor and a foreground communications channel. Total memory bandwidth is 64 gigabits (1 billion) words per second.

All memory access is performed automatically by the hardware. Any user may use all or part of this memory. Significantly larger than that offered on any other commercially available computer system, the CRAY-2's Common Memory allows the individual user to run programs impossible to run on any other system. It also enhances multiprogramming by allowing an exponential increase in the number of jobs that can reside concurrently in memory.

Foreground Processor and I/O Section

The Foreground Processor supervises overall system activity among itself, the Background Processors, Common Memory and peripheral controllers. System communications occurs through four high-speed synchronous data channels.

The maximum number of I/O devices possible in a CRAY-2 configuration is 40 (a maximum of ten devices per communication channel is possible). A typical configuration might include nine disk drives (typically DD-29s) and one adapter per channel, for a total of 36 disk drives and four adapters in the maximum configuration. More than four adapters may be configured, but at the expense of disk storage.

UNIX is a trademark of AT&T Bell Laboratories.

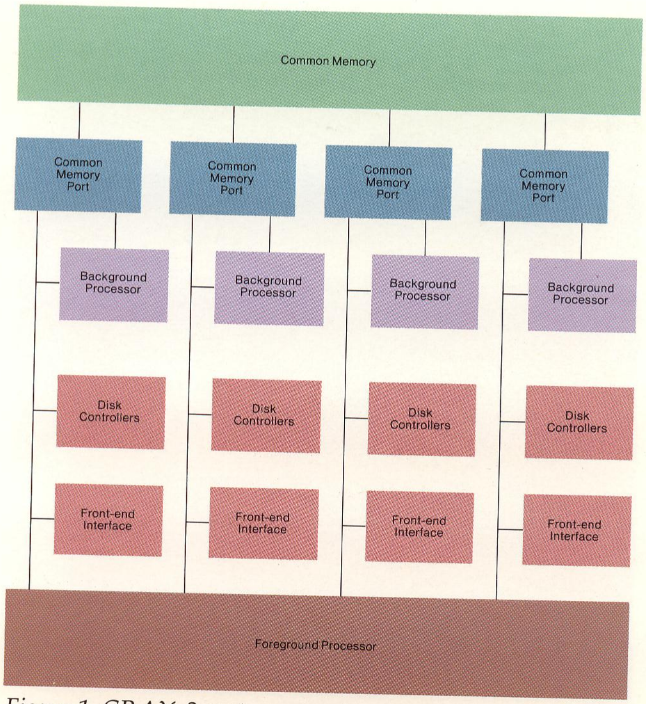

Figure 1. CRAY-2 system overview.

Physical characteristics

The CRAY-2 mainframe is elegant in appearance as well as in architecture. The memory, computer logic and DC power supplies are integrated into a compact mainframe composed of 14 vertical columns arranged in a 300° arc. The upper part of each column contains a stack of 24 modules and the lower part contains power supplies for the system. Total cabinet height, including the power supplies, is 45 inches, and the diameter of the mainframe is 53 inches. Thus, the "footprint" of the mainframe is a mere 16 square feet of floor space. The mainframe weighs 5500 pounds, including 2000 pounds of coolant.

The CRAY-2's pluggable modules are three-dimensional structures; eight printed circuit boards form the module. Circuit interconnections are made in all three dimensions. Each module measures 1 x 4 x 8 inches, weighs 2 pounds, consists of approximately 40% integrated circuits by volume and consumes 300 to 500 watts of power.

Effective cooling techniques are central to the design of high-speed computational systems. More densely packed components result in shorter signal paths, thus contributing to higher speeds. Traditionally, the tradeoff has been lower reliability due to increased operating temperatures, but with the advent of the CRAY-2 this is no longer a limitation. The liquid immersion cooling technology used by the CRAY-2 is a breakthrough in the design of cooling systems for large-scale computers. It places the cooling medium in direct contact with the components to be cooled, thus efficiently reducing and stabilizing the operating temperature and increasing system reliability.

The CRAY-2 mainframe operates in a cabinet filled with a colorless, odorless, inert fluorocarbon fluid. The fluid is nontoxic and nonflammable, and has high dielectric (insulating) properties. It also has high thermal stability and outstanding heat transfer properties. The coolant flows through the module circuit boards and is in direct contact with the integrated circuit packages and power supplies.

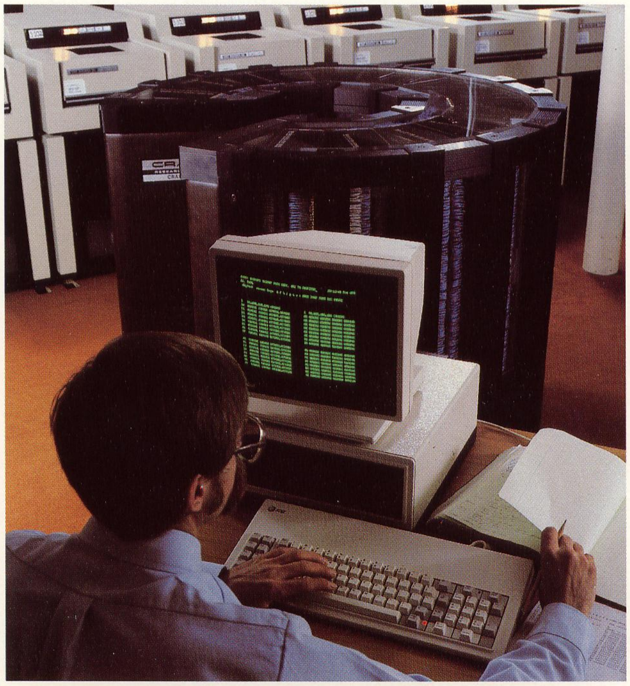
Development work is done with the CRAY-2.

CRAY-2 reliability and maintenance

The CRAY-2's immersion cooling technology contributes to its high reliability. All components rapidly dissipate heat to the fluid, thus preventing high chip temperatures. In fact, chip temperatures on the CRAY-2 are substantially lower than those achieved by other types of cooling and result in significantly reduced chip failure rates. Efficient heat dissipation also prevents destructive thermal shocks that might result from large temperature differentials and fluctuations. Additionally, a fifteen-to-one decrease in module count per CPU from the CRAY-1 and a ten-to-one reduction in memory module count enhance failure isolation, producing a corresponding increase in maintenance efficiency.

If a module should fail, effective and timely maintenance is a routine operation. Diagnostic software quickly isolates the problem to the failing module. The immersion fluid is quickly pumped into the reservoir adjacent to the mainframe. The front panel is easily removed for ready access to the module, which can then be replaced. The front panel is then reinstalled and the fluid quickly returned to the mainframe. The entire operation requires only a few minutes. Once the system is restarted, further diagnosis and repair of the module can occur on-site at the maintenance station.

CRAY-2 software

Cray Research has made a major commitment to the development of a comprehensive and useful user environment through an aggressive software development program. The CRAY-2 Computer System comes with state-of-the-art software including an operating system based on AT&T UNIX System V, an automatic vectorizing FORTRAN compiler, a comprehensive set of utilities and libraries and a C language compiler. The choice of an operating system based on UNIX provides the CRAY-2 user with a well-defined program development environment joined with the advanced computational power of the CRAY-2. The user accesses the power of the system through the proven automatic vectorizing standard-FORTRAN compiler and library routines.

The CRAY-2 Operating System contains a kernel and a large, diverse set of utilities and library programs. The kernel is procedure-oriented, encompassing many processes that dynamically share a common data area used to control the operation of the system. The system is oriented towards an interactive environment with a hierarchical file structure, which features directories, user ownership and file protection/privacy. The kernel of the CRAY-2 Operating System has been substantially enhanced in the areas of asynchronous I/O processing and in the efficient use of very large data files. Other significant enhancements include support for multiprocessing and user multitasking. A batch processing capability is provided for efficient use of the system by large, long-running jobs. The operating system supports high-level languages (including C and FORTRAN) and the mechanism to deliver a common operating system environment across a variety of interconnected computer systems. It provides the user with the ultimate in computational performance from the CRAY-2.

The CRAY-2 FORTRAN compiler, CFT Version 2, is based on CFT, the highly successful CRAY-1 compiler that was the first in the industry to automatically vectorize programs. CFT Version 2 automatically vectorizes inner DO-loops, provides normal program optimization and exploits many of the unique features of the CRAY-2 architecture, all without sacrificing high compilation rates. The compiler and FORTRAN library offer current Cray customers a high level of source code compatibility by making available the same FORTRAN extensions, compiler directives and library interfaces available on other Cray Research products. The FORTRAN library and a library of highly optimized scientific subroutines enable the user to take maximum advantage of the architecture of the hardware. The I/O library provides the FORTRAN user with convenient and efficient use of external devices at maximum data rates for large files.

In conjunction with vectorization and large memory support, a flexible multitasking capability on the CRAY-2 provides a major performance step in large-scale scientific computing. The user interface to the CRAY-2 multitasking capability is a set of FORTRAN-callable library routines compatible with similar routines available on other Cray products.

The C programming language is a high-level language used extensively in the creation of the CRAY-2 Operating System and the majority of the utility programs that comprise the system. It is available on processors ranging from microcomputers to mainframe computers and now to CRAY computers. C is useful for a wide range of applications and system-oriented programs. The availability of C complements the scientific orientation of FORTRAN.

A useful and appropriate set of software tools assists both interactive and batch users in the efficient use of the system. Operational support facilities enable proper management of the system.

The CRAY-2 Assembler, CAL Version 2, provides a powerful macro assembly language that allows users to take advantage of all CRAY-2 instructions, while using an instruction syntax and macro capability highly compatible with the CRAY-1 assembler.

Conclusion

The CRAY-2 Computer System represents a major advance in large-scale computing. The combination of four high-speed Background Processors, a high-speed Local Memory, a large Common Memory, extremely powerful I/O and a comprehensive software product offers unsurpassed performance for today's supercomputer user. With its balanced architecture and large Common Memory, the CRAY-2 offers users dramatically increased throughput rates. In conventional memory-limited computer systems, I/O wait times for large problems that use out-of-memory storage run into hours. On the CRAY-2, however, problems previously considered large-scale become medium- or even small-scale. And problems considered unsolvable or too costly become not only solvable, but economically feasible as well.

The CRAY-2 computer system . . . setting the standard for the next generation of supercomputers. □

Summary of CRAY-2 features

Hardware:

- ☐ 256 million words of directly addressable Common Memory
- ☐ 4.1 nsec clock cycle time
- ☐ Integral Foreground Processor and four independent Background Processors
- ☐ Powerful I/O
- ☐ Very high reliability
- ☐ Liquid immersion cooling
- ☐ Three-dimensional modules

Software:

- ☐ The CRAY-2 Operating System, based on the proven UNIX System V and enhanced to fit the large-scale scientific computer environment
- ☐ CFT Version 2, a vectorizing and optimizing FORTRAN compiler
- ☐ A FORTRAN standard mathematical and I/O subroutine library
- ☐ A scientific subroutine library optimized for the CRAY-2
- ☐ A multitasking library that allows user partitioning of an application into concurrently executing tasks
- ☐ A wide variety of system utilities to support the needs of interactive and batch processing
- ☐ A C language compiler that supports the needs of system software written in C
- ☐ CAL Version 2, the CRAY macro assembler, which provides access to all CRAY-2 instructions

Computer simulation and nuclear safety

The nearly 80 nuclear power plants operating in the United States today contribute more than 13 percent of our national electric power supply. Another 50 or so are currently under construction, so that by the year 2000, nuclear power will provide upward of 20 percent of all the electricity used in the United States. As these figures indicate, nuclear power has become an integral part of our national energy grid.

Internationally, nuclear power also has an important role. France and Japan have made it a central part of their national energy policies and both have built numerous nuclear reactors. In addition, nuclear reactors are producing electricity in the United Kingdom, Canada, West Germany, Belgium, Sweden, Korea and other industrial nations.

In spite of its practical success, however, the nuclear energy industry remains embroiled in controversy, at least in the United States. In the mind of much of the public, the risks surrounding nuclear power outweigh its value as an energy source. Concerns about power plant safety have prompted the U.S. Nuclear Regulatory Commission (NRC) to issue nearly 2000 regulatory guides, letters, bulletins, orders, notices and standards since 1982. Among current regulations is the requirement that proposed power plant designs adhere to specific criteria for emergency core cooling system (ECCS) performance. Utilities must demonstrate that for a given accident the ECCS performance meets the Federal regulations before they can obtain a license for new plant construction.

"The analysis assumptions described in Appendix K of 10 CFR 50 (the Federal regulation specifying analysis assumptions that must be used in ECCS analyses) have really become the basis for power plant simulations run on computers," explained Brian Sheron, chief of the Reactor Systems Branch, Division of Systems Integration of the NRC's Office of Nuclear Reactor Regulation. "The NRC has developed its own codes and approved others for simulating power plants based on the regulation. Typically, vendors run their proposed designs through the simulations, and those results are submitted to the NRC as proof that the plant satisfies the criteria."

Today, computer simulation is the most important tool available for assessing nuclear power plant safety. Simulations play a central role in the licensing process, the development of power plant designs and basic research. By modeling a power plant's internal processes on a computer, scientists can study the performance of the plant's safety features and predict the outcome of foreseeable failures. Since no one would ever want to subject a full-scale nuclear power plant to severe accident conditions, and because the cost of full-scale experiments for all of the possible accident paths is prohibitive, the nuclear power industry relies more heavily on theoretical analysis of design and safety features than any other high-technology industry. The computer codes used to perform these analyses are tested against experimental data from scaled test facilities in the United States and abroad.

Supercomputers are required to run complex reactor simulation codes in a practical time frame. For example, CRAY systems play a central role in nuclear power plant safety analyses conducted at the Los Alamos National Laboratory (LANL) in Los Alamos, New Mexico. The laboratory developed a program called the Transient Reactor Analysis Code (TRAC), which is now run on its CRAY computers. A second major nuclear power plant simulation code, the Reactor Excursion and Leak Analysis Program (RELAP), has been developed at the Idaho National Energy Laboratory (INEL). The latest version, RELAP5, has recently been converted for use on CRAY systems.

Reactor primer

Although a modern nuclear power plant is a very complex system designed to exacting specifications, a nuclear reactor, by itself, is a relatively simple device. In 1942, Enrico Fermi and his colleagues built a crude reactor on the first try. By placing pieces of natural uranium in a stack of graphite blocks, they achieved a self-sustained and controlled nuclear fission chain reaction, thus demonstrating the potential for generating a large amount of usable energy.

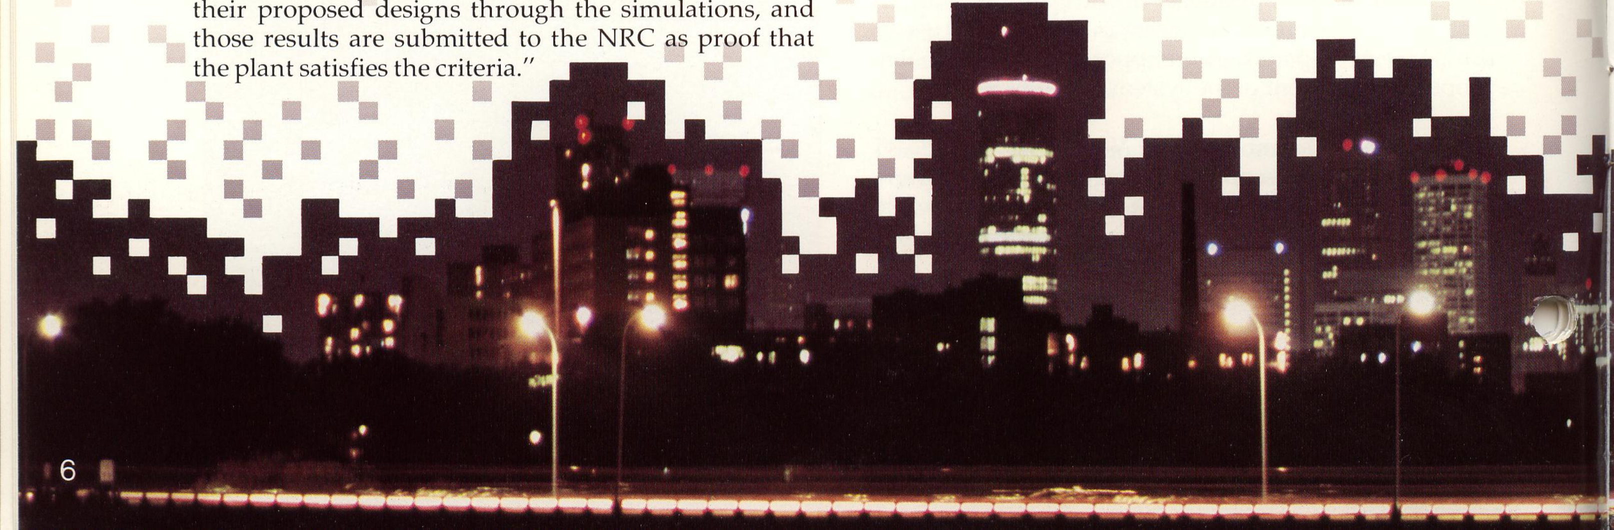

The energy-producing process is nuclear fission, in which an atomic nucleus absorbs a neutron and breaks apart into several fragments, typically two smaller nuclei, two or more neutrons and gamma rays. The resulting neutrons can themselves initiate fission of other nuclei and so begin a chain reaction. Sustaining this chain reaction, however, requires a sufficiently large mass of fuel. With too little fuel, too many neutrons escape and the chain reaction stops.

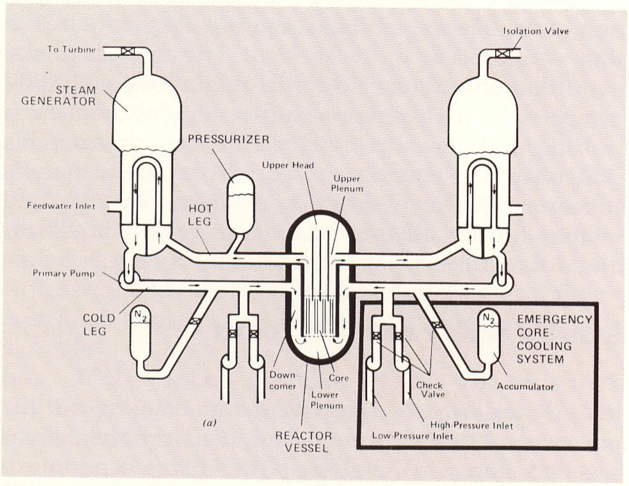

Figure 1. Coolant flow pattern through the primary system of a pressurized-water reactor.

Several methods are commonly used to control the fission reaction rate in nuclear reactors. Older plants typically rely on control rods that are moved in and out of the reactor core where the fission reaction is taking place. These rods contain materials such as boron or cadmium that readily absorb neutrons without undergoing fission, thus removing them from further participation in the ongoing reaction. Newer pressurized-water reactor plants, however, typically control the fission reaction rate with boron dissolved directly in the primary coolant. As the coolant bathes the reactor core (Figure 1), free neutrons are exposed to the neutron-absorbing boron. The efficiency of the fission reaction in reactor cores is maximized by slowing down neutrons with graphite or water, and/or reflecting them back into the fuel region.

Modern nuclear power plants include fail-safe systems for the rapid insertion of control rods into the core to halt the chain reaction altogether under emergency conditions. This is referred to as a reactor scram. In addition, some reaction rate control occurs spontaneously due to temperature changes in the core. In light water reactors, for example, changes in the core's material properties tend to shut down the chain reaction as the core temperature rises. This self-regulation provides a nuclear reactor with a certain amount of inherent stability and safety.

Most of the energy released by nuclear fission is the kinetic energy of the lighter nuclei that are formed when the heavy nuclei split. When these nuclei collide with neighboring fuel nuclei, their kinetic energy is converted to heat, which then transfers from the fuel to a liquid or gas coolant pumped through the reactor core. Since the rate of heat transfer to the coolant must equal the rate of energy production in the core to prevent the core from overheating, ensuring appropriate coolant flow is crucial to the proper and safe functioning of a reactor. Heat transferred to the coolant/moderator and from the coolant to a secondary coolant loop is used to produce steam for generating electric power.

Power plant safety

Power plant safety considerations center on two potential trouble spots: the long-lived decay heat (heat generated by the nuclear breakdown of radioactive material) and the radioactivity of the fission products. Provision for the removal of decay heat is a critical requirement for all power reactors. Excess decay heat can damage the core and potentially lead to the release of radioactive materials into the biosphere.

To make sure the radioactive products are contained, power plants incorporate four distinct barriers: ceramic (uranium dioxide) fuel pellets, fuel-rod cladding, a primary system boundary and the containment building itself.

The *ceramic fuel pellets* have an exceptionally high melting point (3040° Kelvin, or about 5010° Fahrenheit) and chemical stability that prevent escape of most fission products except in extreme accident conditions. The trace of fission products that normally escapes is confined by the second barrier, the *cladding* surrounding the fuel pellets. The cladding typically is made of a zirconium-based metal that has low neutronic absorption properties and high strength, low permeability to fission product gases and good thermal response properties. If the core temperature were to increase to the point where the cladding failed, the fission products would be contained by the third barrier, the *primary system boundary*. In spite of the inherent strength of the thick vessels and piping in this boundary, spontaneous small and large breaks can occur, thus initiating loss-of-coolant accidents.

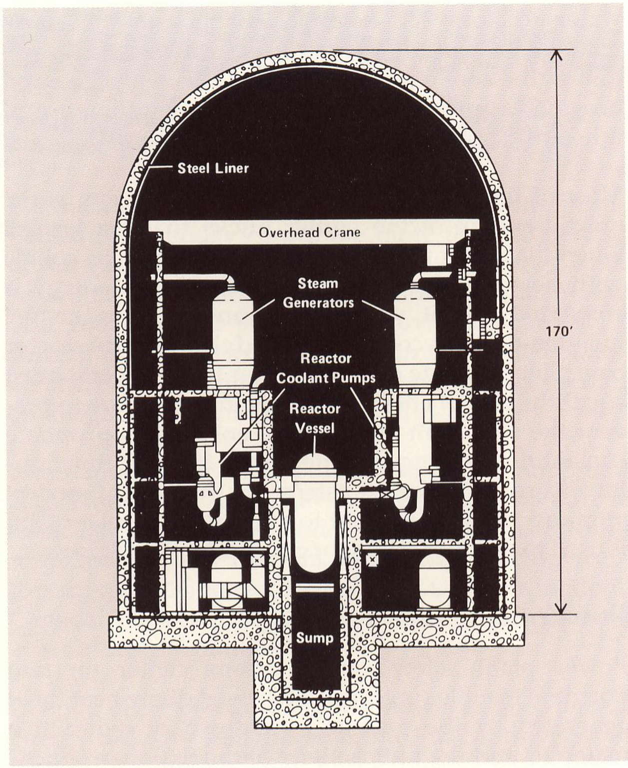

Figure 2. Cross section of a typical containment building for a pressurized-water reactor housing the entire primary system, the pressure control system, ventilation equipment and part of the emergency core-cooling system.

The *reactor containment building* (Figure 2), is the fourth and final barrier to fission product release. It generally consists of a steel liner surrounded by a structural concrete shell. The fourfold barrier combination prevents leaks and can withstand substantial internal overpressure, as well as external impacts caused by tornadoes, external explosions or aircraft crashes.

Safety analysis

The safety analyst's job is to determine, for any postulated accident, whether the maze of barriers stays intact and whether radioactive materials stay contained. But the maze is complex and changing during an accident. The locations and sizes of barrier failures, the release paths and the transport mechanisms all depend on temperature and pressure. The analyst must start from the beginning and predict fluctuating thermal and physical conditions throughout the entire accident. The analysis typically involves a sophisticated computer model that breaks down the system into many cells and audits the mass, temperature and velocity of the materials in each cell.

The most widely used power plant safety analysis codes integrate models describing physical processes such as two-phase fluid dynamics, heat transfer, reactor kinetics, control systems, structural processes (especially fuel) and component dynamics (pumps, valves, etc.) With such codes the analyst can predict the physical state, including pressures, temperatures, flow rates, coolant distribution and power in a reactor system or component in both time and space. The credibility of reactor analysis codes rests on their continued assessment against experimental data.

The analysis usually begins with the reactor running smoothly at full power. Then something is assumed to go wrong — a pump fails or a coolant-carrying pipe breaks — and the computer calculation follows the changes in water and steam flow rates and in system temperatures and pressures. Reactor scram and injection of emergency cooling water are also simulated as they would occur in the accident. The analysis tracks, over time, the system's thermal hydraulics, including compressible two-phase steam-water flow, an engineering and computational problem of considerable difficulty. The equations used in computer analysis codes assume conservation of mass, energy and momentum for all the materials in each of the hundreds of cells in a typical calculation.

Reactor safety analysis codes

Two primary safety analysis codes, TRAC and RELAP, are used by the NRC for its ongoing nuclear power plant safety research projects. They and derivative codes are also used by power plant vendors for design research. TRAC and RELAP are "best estimate" codes intended to simulate power plant dynamics as accurately as possible.

Development of these codes was commissioned by the NRC and carried out at LANL and INEL, respectively. "RELAP was developed first, and at the time we considered a one-dimensional code to be adequate," said Sam Bassett, director of the Division of Accident Evaluation at the NRC's Office of Research. "But work to make it more accurate showed us we had to include three-dimensional effects like cross flow and radial flow, as well as asymmetric system effects. So we initiated a *de novo* attempt at a three-dimensional code at Los Alamos, TRAC. In the long run, one of the programs may prove to be better than the other, but for the time being, the NRC supports both — and both are heavily used around the world."

RELAP and TRAC are based on similar physical models. Both solve the Navier Stokes equations for the void (gas), fluid and boron flow fields. They also calculate conductive, convective and radiative heat transfer — a highly nonlinear problem. Each code consists of a large set of subprograms that can be assembled to simulate the thermohydrodynamics that might occur during any specified transient in any realistic reactor design. (A transient is a change, over time, in the physical properties of the internal reactor system.) There are subprograms for the reactor components — the core, pipes, pressurizer, valves, steam generators, pumps and accumulators. Other subprograms simulate the physical processes such as steam-water fluid dynamics, heat generation in the core and heat transfer between the two phases of the coolant (steam and water) and between the coolant and the solid structures. When assembled into a large systems code and run on a high-speed computer like the CRAY, these subprograms numerically simulate the complete course of reactor transients, most notably the loss-of-coolant accident.

The modular designs of TRAC and RELAP make them flexible enough to be suitable for studying many types of transients. By joining the modules (subprograms) in a meaningful way, the user can simulate a wide range of phenomena, from a simple blowdown to a multiple-failure transient. The user need supply only the problem geometry and boundary conditions.

The codes produce an extraordinary amount of information during the course of a calculation. At each step and for each mesh cell, they provide values for the following variables: fluid pressure, void fraction, temperatures and velocities of the two coolant phases (for vessel cells, the velocities are vector quantities) and temperatures of solid materials, such as the cladding. Other variables, such as mass and momentum fluxes and fluid density, can be obtained from these basic variables.

"In the early 1970s, when these codes were first being developed, they included some very conservative assumptions," said Dr. Richard Lee, a program manager at the Reactor Systems Research Branch at the NRC. "But since then, we've learned so much about reactors, heat transfer in particular, that the newer versions are much more realistic. And because they're more realistic, their predictive capabilities are much better."

"We have more faith in our methods today," added Dr. Paul Turinsky, head of the department of nuclear engineering at North Carolina State University. "The codes are qualified against data gathered from large-scale experiments on power plant models conducted by the NRC. Considerable conservatism needed to be built into the early codes because of their limited accuracy, to assure adequate safety margins. The newer, more accurate, codes give us a better quantitative understanding of the real margins of safety we have to work with."

Computer simulations are used extensively by the NRC to understand plant behavior and to study generic problems that affect more than one plant. "We use TRAC and RELAP in our Incident Response Center where we simulate and analyze the location of an accident in a plant and how it's progressing," explained NRC's Sheron. "We use the results to provide utilities with evaluation and guidance, although the ultimate responsibility for any accident is with the operator/owner. We also might make recommendations about whether or not to evacuate an area and how much area to evacuate, if we were asked to do so."

Licensing applications

Along with research addressing general issues of plant design and operation, computer simulations are used extensively in the process of licensing new plant construction. Analyses are usually performed by the vendor selling the plant to the utility company, although occasionally they will be run by an independent contractor or by the utility itself. Computer analysis results are also used by utilities when applying for an amendment to their license, as when a plant is reloaded with new fuel. Furthermore, when the NRC issues regulations that require new analyses for plants already licensed, utilities must resubmit their results to show that the new regulations are met.

Typically, licensing codes are proprietary to the power plant vendor and the analysis is included in the vendor's contract with the utility. Westinghouse, the only nuclear power plant vendor that currently owns CRAY systems, is in the process of transferring its proprietary licensing codes to the CRAY. "The proprietary codes we use for licensing currently run on a CDC 7600, but we're starting to vectorize them now and are in the process of moving them over to the CRAY for use in the future," said Jack Olhoeft, a technical assistant in Westinghouse's Nuclear Safety Department.

For licensing codes to meet NRC regulations, they must include conservativisms that guarantee large safety margins. "Best estimate" codes such as TRAC and RELAP do not include extra conservativisms, but are intended to produce the most accurate results possible. Nonetheless, best estimate codes play an accessory role in power plant licensing. "Along with verifying compliance with the regulations, we want to know what's actually happening in a plant," said Cliff Davis, a senior engineering specialist at INEL. "The conservatisms in the licensing codes may cause them to respond differently than the plants. While it's not a required part of the licensing process, it's not uncommon that the NRC asks one of the national labs to do an independent audit of the results submitted to the agency. We conduct these audits using best estimate codes. Some people have even begun to think that it might be better to do best estimate analyses, with uncertainty bands applied to the results, for licensing purposes."

The suggestion that best estimate codes be allowed for licensing analyses is receiving serious attention within the NRC. The Office of Nuclear Regulatory Research currently is drafting an amendment to Appendix K that would qualify best estimate codes for licensing submissions. The office hopes to put out a public notice of the proposed changes later this year, although the exact timing depends on the speed of the NRC's internal review process, according to Louis Shotkin, chief of the Reactor Systems Research Branch, Division of Accident Evaluation, at the NRC's Office of Research.

"The existing Appendix K requirements continue to be adequate and we intend to grandfather their use," said Shotkin. "However, today we have computational tools that are much better than those we had when the regulations were written. In addition, codes like TRAC and RELAP have been assessed against a considerable amount of experimental data. We think best estimate codes should be acceptable for evaluating both small and large-break LOCAs, because now we're able to quantify the uncertainty in the codes and it's acceptable."

RELAP and TRAC are today the most widely used best estimate codes. Westinghouse researches power plant design using a TRAC-derivative running on its CRAY systems and both TRAC and RELAP are used for research and licensing around the world. Current work on these codes is described below.

TRAC

Originally written in the mid-1970s to run on a CDC 7600, TRAC was successively refined and eventually put up on the CRAY. "We reached the point where there was not enough memory available on the 7600 for many of the large calculations," explained John Mahaffy, former associate leader of the safety code development group at Los Alamos, now working at Pennsylvania State University. "We put it on the CRAY because we had access to CRAYs at Los Alamos and we needed the larger memory. The code has a huge number of arithmetical operations for each time step and the calculations require 64-bit arithmetic. With the CRAY's one million words of memory, we could run code on it and get around disk I/O altogether."

"We're optimizing TRAC for the CRAY, and that project is just about completed," Mahaffy added. "Pacific Sierra Research Corporation has done some timing studies and we figure we'll get a speedup of slightly over two. The work in the code is widely spread. It's not like there are a few DO loops to vectorize — about 10,000 lines of code must be reworked. We're also restructuring the database so we can feed in large vectors." The vectorized version of TRAC will be incorporated into the Nuclear Plant Analyzer, an interactive program under development that will make power plant simulations easier to analyze.

Along with the CRAY's large memory, its speed offers a distinct advantage for power plant simulations. The CRAY's speed presents the potential for real time, or better than real time, simulations using TRAC. "This is what we're looking at in the future. Simulations running at real time or better give us the possibility of direct guidance to plant operators," said Mahaffy. "For example, during a slow transient, like Three Mile Island where they had a couple of hours to react, we might be able to run several simulations to see what the best response would be or what caused the accident."

RELAP

The development of power plant safety analysis codes began nearly twenty years ago at INEL with the RELAPSE code. RELAPSE has since evolved into RELAP5, the latest version of the code to come out of INEL. While TRAC and RELAP5 are very similar, they differ in that TRAC includes a three-dimensional core model, while RELAP5 is entirely one- and pseudo-two-dimensional.

A conversion effort to get RELAP5/MOD2 running on CRAY computers was completed in early 1985. "As the code grew, we ran into memory constraints on the machines for which it had originally been developed," explained Richard Wagner, principal engineer with the Code Development and Analysis Program, Water Reactor Systems Research, at INEL. "For example, the severe core damage simulation used to be a separate program, but we incorporated it into RELAP5 because it was more accurate to include feedback effects among the different plant components. We had to segment the expanded code, which created a lot of disk activity and had quite an impact on the code's performance. So our first motivation to convert for the CRAY was memory, but we appreciate its speed too."

Recently, RELAP5/MOD1 was multitasked experimentally to run on the four-CPU CRAY X-MP/48, with significant performance improvements. An entire accident sequence was run (1000 seconds of simulated time with as good as round-off accuracy) for a four-inch small cold-leg break pressurized water reactor accident using a licensing model. With about 70 percent of the code multitasked, it showed performance improvements of 1.5 to 2.5 over the original code.

"A small part of that speedup was due to new algorithm developments suitable for parallel machine architectures," explained Henry Makowitz, the CRAY analyst responsible for the development work on RELAP5/MOD1. "But the real significance of our results was that we were able to obtain first principles faster than real time simulations using the original numerics of the code for the full length of the accident sequence. This has never been done before for a first-principles calculation for the entirety of such a severe accident sequence. However, the speedup we're able to get is limited by the sparse matrix solver, which cannot yet be multitasked." Makowitz is currently working with North Carolina State University's Department of Nuclear Engineering to develop parallel sparse matrix solvers.

A system analysis conducted by Makowitz (Figure 3) indicates that up to 95 percent of new power plant safety analysis codes could be multitasked if parallel techniques and philosophy were adopted for code development. A nodal approach to building the matrices,

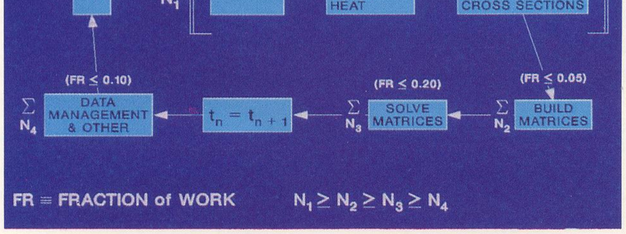

Figure 3. Levels of parallelism proposed for future code development. Such an approach could lead to significantly faster than real time simulations.

followed by a parallel sparse matrix solution for an implicit finite difference formulation, will probably optimize performance for vectorizing concurrent multiprocessing machines such as the CRAY X-MP/48. "Designing new codes based on multitasking could result in simulations five to ten times faster than real time," predicted Makowitz.

Makowitz also is working with the University of Illinois' Nuclear Engineering Department on the development of artificial intelligence codes for computer assessment of slow transients. Such codes, coupled with the capability of significantly faster than real time simulation, present safety analysts with the possibility of running several simulations during the course of an actual transient. Analysts would receive intelligent suggestions for intervention strategies based on probable events from the computer. "Such a capability could play a central role in projects such as the Nuclear Plant Analyzer," noted Makowitz.

The Nuclear Plant Analyzer

To improve the ease of power plant simulation, the NRC has commissioned the development of a computer-software interface for executing TRAC and RELAP5. The interface, called the Nuclear Plant Analyzer (NPA), is designed to use advanced supercomputers, long-distance data communications and a remote workstation terminal with interactive computer graphics. An NPA to drive TRAC has been developed at Los Alamos and a separate one to drive RELAP5 has been developed at INEL, both using common guidelines.

The NPA provides TRAC and RELAP5 users with a tool that can significantly reduce the time and effort required to analyze power plant transients. Weeks to months of human effort are required to prepare high quality input data, to execute TRAC or RELAP5 using that data and to interpret the results of the calculation. The NPA is designed to automate most of this procedure and to provide interactive capability to the user during the calculation. Computed results are presented in graphics displays as the calculation proceeds. Sample graphic output is shown in Figure 4. Control of the plant, as defined by the input data, can be overridden at any time during the calculation by hardware-adjustment commands issued by the NPA user. The NPA handles all interaction with the computing environment. This allows the user's attention to be devoted fully to the transient event being analyzed.

"No code experience is needed to run the TRAC and RELAP5 codes using the NPA," explained Los Alamos' Mahaffy. "With the NPA, analysts only need to understand the complex thermal-hydraulic phenomena occurring in power-plant transients." Testing the NPA interactive coupling for each code was recently completed at LANL and INEL. "The next step will be combining these two NPA versions using the best features of each and incorporating software to access the nuclear plant database," Mahaffy said.

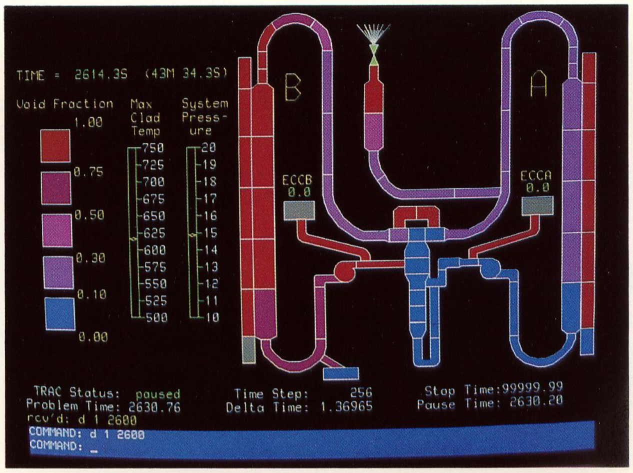

```
 TRAC Status:  paused           Time Step:       256          Stop Time:99999.99
Problem Time: 2630.76         Delta Time:  1.36965          Pause Time: 2630.20
rcv'd: d 1 2600
COMMAND: d 1 2600
COMMAND: _
```

Figure 4. Sample graphic output from the TRAC NPA. Any variable in the code, such as void fraction, cladding temperature or pressure, can be color coded and mapped onto the noding diagram.

Conclusion

Combining supercomputer power with codes like TRAC, RELAP5 and the Nuclear Plant Analyzer helps to assure maximum safety in nuclear power plant design and operation. The extreme complexity of these codes and the desire for faster than real time simulations make supercomputers an essential part of power plant safety analysis. Licensing codes may also find themselves running on supercomputers more and more if new plant construction picks up significantly. But whatever the future of the nuclear power industry, detailed computer simulations will continue to play a central role in ensuring nuclear reactor safety. □

References

Johnsen, G. W., "Reactor Safety Computer Code Development at INEL," *Report to the Department of Energy*, January 1985.

Makowitz, H., "Numerical Experiments in Concurrent Multiprocessing with the RELAP5 Nuclear Reactor Systems Code," *Proceedings of the Eleventh International Meeting on Advances in Nuclear Engineering Computational Methods*, Knoxville, Tennessee, April 9-11, 1985.

Makowitz, H., "Numerical Experiments in Concurrent Multiprocessing with Thermohydraulic Nuclear Reactor Systems Codes," *proceedings of the Eighth International Conference on Structural Mechanics in Reactor Technology (SMIRT)*, Brussels, Belgium, August 19-23, 1985.

Stevenson, M. G. and Jackson, J. F., "A Primer on Reactor Safety Analysis," *Los Alamos Science*, Vol. 2, No. 2, 1981.

Vigil, J. C. and Pryor, R. J., "Accident Simulation with TRAC," ibid.

Weaver, L.E., "The Outlook for Nuclear Power: Is it still an Energy Option in the United States?," *Nuclear News*, August 1984.

Multitasking at Cray

Luann Nelson and Peter Rigsbee
Cray Research, Inc.

High-speed performance. It is coveted by all involved in large-scale industrial and scientific computing. Traditionally, performance increases have come from speedups in the CPU, but this is becoming more difficult. Designers have therefore moved on to another method for improving total system performance: the coupling of multiple processors into a single system. This is an approach taken by Cray Research — in fact, most computers being delivered by the company today are multiprocessors. The CRAY X-MP Series includes one-, two-, and four-processor models; meanwhile, the CRAY-2 is comprised of four central processors built into a single mainframe.

The performance gains promised by multiple-CPU systems can be achieved in two ways. The first approach treats each CPU independently and schedules separate jobs to run on each of them. As a result, total system throughput is a multiple of that achievable with a single processor. The second approach allows a single program to use more than one processor. When multiple processors execute portions of a single program concurrently, program turnaround and thus, wall-clock time, should be reduced. Cray Research calls this multitasking.

Some definitions

Before going much further, several related terms require clarification. The definitions presented here are those used by Cray Research.

- *Multiprocessing* is a property of the hardware in which two or more CPUs are available. For example, the CRAY X-MP/2, the X-MP/48 and the CRAY-2 are all multiprocessors.

- *Multiprogramming* is a property of the operating system allowing overlapping and interleaving execution of more than one program or job at a time. System resources, such as processors, I/O devices and central memory can be shared efficiently among many jobs. When one program performs I/O, another uses the CPU for calculations. Multiprogramming is supported by most computer vendors for all but the simplest single-user computers.

- *Multitasking* is the structuring of a program into two or more parallel instruction streams that execute concurrently. These instruction streams are called tasks. Multitasking is a software construct and pertains only to the manner in which an application program is organized. Performance gains are realized when the "multitasked" application is run on a multiprocessing system. If the tasks execute at the same time on different processors, overall job turnaround, wall-clock time, can be reduced. Cray Research distinguishes between two types of multitasking, *macrotasking* and *microtasking*, both of which will be explained later in this article.

Cray has developed software to allow users to implement multitasking on multiprocessor systems. This article examines what multitasking is, when it should be considered and how it can be used.

Applications for multitasking

In a general sense, almost all applications run on CRAY computers are candidates for multitasking. The primary use of a CRAY system is to model physical processes such as weather, nuclear reactions and acoustical images. Weather phenomena occur simultaneously; the particles in a nuclear plasma all move at the same time and their interactions are resolved instantaneously. It would be ideal if application programs could be structured to execute as real-life phenomena do.

However, scientists are forced to describe and model physical processes in sequential terms on single processor computers. A period of time called a time-step is chosen, and instantaneous interactions during that period are approximated by a sequential series of computations. With many computer systems, this is the only option. The computer executes instructions sequentially, and therefore, the applications running on it must be sequential.

CRAY computers offer two significant approaches in which parallelism inherent in an application can be exploited to improve performance. First, vector registers in the hardware allow certain computations to be carried out in parallel on up to 64 independent elements at a time. Vectorization is the method by which a code is modified by the compiler to use this hardware feature. Second, multiple processors allow certain computations to be carried out in parallel on two or more processors. Multitasking is the method used to modify program code to use multiple processors.

Vectorization yields performance gains by overlapping CPU cycles. Thus, a gain is realized by using the CPU more efficiently. Multitasking yields performance gains by splitting a portion of a program among two or more processors. Therefore, a gain is achieved in wall-clock time.

Vectorization and multitasking can be used on the same application; the overall result is a performance improvement that is a combination of the improvements from the two separate methods. While the Cray FORTRAN Compiler automatically modifies a program to allow vectorization, the user must explicitly start up the multitasking process by calling for the library routines provided by Cray Research.

Work has been done on a number of application programs to allow them to make use of Cray's multitasking capabilities. They include:

- ☐ Medium-range weather forecasting
- ☐ Short-term spectral weather forecasting
- ☐ Particle-in-cell simulation
- ☐ Monte Carlo gamma ray transport
- ☐ Seismic three-dimensional migration
- ☐ Nuclear reactor simulation (RELAP5)
- ☐ Circuit design (SPICE)
- ☐ Computer chess (CRAY BLITZ)

Now, multitasking is implemented at the subroutine level; research is being done to develop multitasking at the DO-loop level (called microtasking).

Multitasking considerations

The implementation of multitasking by Cray is best-suited for a particular class of applications that can be characterized as:

- ☐ Written in FORTRAN
- ☐ Long running or frequently running
- ☐ Using most or all of memory
- ☐ Requiring a dedicated environment for any reason

The decision whether or not to modify a program to use multitasking is influenced by these four attributes.

Most major applications run on CRAY supercomputers are written in FORTRAN. Thus, the Cray approach has been to support multitasking through FORTRAN, keeping enhancements to a minimum.

The length of a job is important because there is a programmer cost to multitask a job. Saving 1.5 hours of a six-hour job is better than saving three minutes of a six-minute job that will run only once. (Of course, if the six-minute job is run a thousand times a day, multitasking may very well pay off.)

Multitasking introduces some system overhead that is not present in a traditional, sequential application. The amount of computation that each task performs must be weighed against the overhead introduced by creating that task. If the amount of computation in a portion of a program is too low, the performance gain from the parallelism will be exceeded by the overhead of creating and communicating among the tasks. The longer a task can execute, especially without need to communicate with other tasks, the greater the performance increase from multitasking.

Multitasking is also well-suited for jobs that execute in a dedicated environment. This can take two forms: the job that runs dedicated during off hours, or the job that requires most or all of memory during production hours. Multitasking dedicated jobs permits them to finish sooner or permits the scientist to solve a more complex or detailed problem. Multitasking memory-bound jobs, using otherwise idle processors, improves system throughput since jobs will move through the system faster.

Multitasking implementations

Cray Research distinguishes between two types of multitasking. The first is *macrotasking,* which organizes a program into large sections involving significant amounts of work that can be executed in parallel. This form of multitasking is currently supported by a set of standard library routines. The second type of multitasking has been termed *microtasking,* which partitions work at a very low level, typically alternate iterations of a FORTRAN DO-loop. Microtasking is currently in an experimental stage, but appears to offer significant performance enhancements for programs that cannot be partitioned into larger pieces.

The initial Cray implementation of multitasking centers on the development of a set of library subroutines that uses a basic set of primitive multitasking functions. The user is asked to analyze the application and insert library calls as required to create and coordinate or synchronize task execution at the level of the user's subroutines.

Analyzing the application code is key to producing a successful multitasked application. Precisely defining and adding the necessary communication and synchronization mechanisms between parallel tasks is important. After a program is developed, debugged and vectorized, one must look at the program's flow to locate sections that have the potential to run in parallel.

The next step involves looking at variables used in the tasks, and deciding which must be shared between tasks and which must be private to a single task. Variables that will be shared may require protection. Cray Research has developed and is enhancing various tools to aid the user in these analyses.

Code modifications are made after the analysis is completed. Modifications involve inserting library subroutine calls to use the multitasking support. The major library subroutines:

- ☐ Create tasks and wait for their completion
- ☐ Test and signal software EVENTs, which synchronize the execution of tasks
- ☐ Control software LOCKs, which allow protection of resources and variables shared by tasks

Other code changes are required to ensure that variables are properly treated as private or shared. Three basic options exist:

- ☐ COMMON variables can be shared by all tasks and are allocated to static locations in memory
- ☐ TASK COMMON variables can be shared by, and only by, the subroutines included in a task and are allocated dynamically at runtime
- ☐ Other variables are local to a task and are (usually) allocated dynamically at runtime

A major advantage of the current implementation is that it is system independent. The same FORTRAN interface is supported on CRAY-1, CRAY X-MP and CRAY-2 systems, so applications developed today for one CRAY computer can be easily transported to another CRAY model.

Future development and research

Cray is investigating alternative machine-dependent implementations that allow the minimum task size to be reduced. CRAY X-MP hardware includes features that allow rapid task switching and communication. Currently, these features are available to the user for experimentation through assembly language. (They are not available on the CRAY-2.)

In addition, Cray is investigating automatic code partitioning by the new FORTRAN compiler. Automatic partitioning, in conjunction with an efficient multitasking implementation, will allow users to multitask some loops within an application with no more effort than is required to vectorize loops today. The performance gains that can be achieved with automatic partitioning will depend on the sizes of the loops selected and their importance within the application. Generally, however, performance gains will be smaller than could be achieved by the analysis described earlier. But the small effort required by the user will frequently make this worthwhile.

Microtasking development continues. Later this year, multitasking will be implemented in Pascal in the same manner as with FORTRAN. Other work being done by Cray focuses on additional and enhanced tools and products for use in analyzing and debugging multitasked applications.

This article covers only a few of the concepts and issues related to multitasking, and only provides a glimpse into the features and software available from Cray Research. A Multitasking User Guide (SN-0222) describing analysis and implementation of applications for CRAY X-MP computers is available. It may be purchased for $2.50 by ordering from: Dennis Abraham, Cray Research, Inc., 1440 Northland Drive, Mendota Heights, MN 55120. Future issues of CRAY CHANNELS will feature additional articles about multitasking. ☐

Acknowledgement

CRAY CHANNELS thanks the authors, who are members of the Cray Research Software Development Department, for their work in preparing this article. Additional thanks are extended to John Larson of the CRAY X-MP Development Group in Chippewa Falls, WI for the extensive background information he provided.

Theoretical and actual multitasking speedups

What kind of performance gains should be expected from multitasking an application on a CRAY or any other processor? Research has been done in this area, and it turns out that the maximum theoretical gain is a function of only two variables: the fraction of the application that is multitasked, and the number of processors. An equation known as Amdahl's Law predicts the speedup based on these variables:

$$S(P,f) = \frac{T1}{Ts + Tm}$$

$$= \frac{T1}{T1 * ((1-f) + (f/P))}$$

$$= \frac{1}{(1-f) + (f/P)}$$

where:

$T1$ = original execution time (no multitasking)

f = fraction of time multitasked

$Ts = (1-f) * T1$ = time to execute sequential part

$Tm = (f/P) * T1$ = time to execute multitasked part

P = number of processors

For a given, fixed number of processors, a plot of speedup versus 'f' produces the Amdahl's Law Curve as illustrated at right.

There are two key messages from this. First, significant speedups are not possible unless significant portions of the code are multitasked. Second, for a fixed percentage of multitasked execution time, the speedup does not increase as fast as the number of processors. The general belief that with P processors, speedups should be P is a myth.

How close can a user come to obtaining these theoretical gains? Take Amdahl's Law and add one additional term:

$$S(P,f) = \frac{T1}{Ts + Tm + OH}$$

where:

OH = overhead time because of the multitasking

The other terms are as defined above

The theoretical gains are for those cases in which the overhead is zero. Overhead is a combination of factors:

- ☐ Time to execute the multitasking subroutines
- ☐ Time to dynamically allocate local variables
- ☐ Time lost to one task while another uses a common resource.

For most multitasked applications on the CRAY X-MP, overhead is typically less than two percent of the total application time.

Amdahl's Law Curve

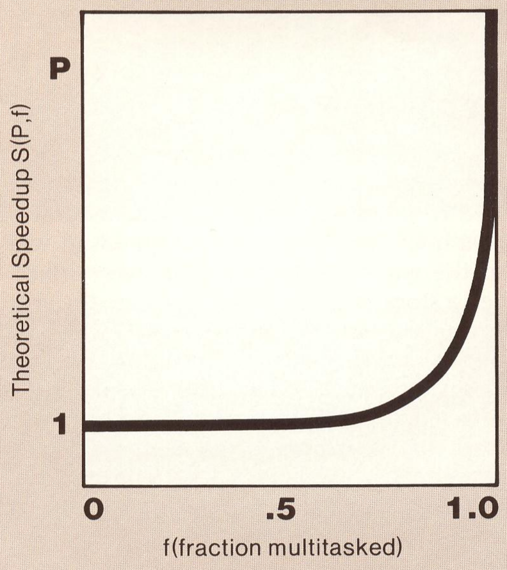

Migration of seismic source records

Don E. Larson
Geoquest International, Inc.

In the search for valuable mineral deposits buried beneath the Earth's surface, reflection seismology is one of our most powerful tools. Most often used to explore for oil and gas, the method is also used to search for other minerals and even for water.

Although seismic data was collected and analyzed before computers were available, the industry has been very quick to exploit each new advance in computer technology. When digital computers became generally available, the industry responded by developing digital recording equipment. This, combined with the introduction of the common depth point (CDP) method in exploration, resulted in an enormous increase in both the quantity and quality of seismic data.

In the 1970s, wave theory was introduced into seismic data processing. Wave theory methods are very demanding of computer resources and thus have been applied only after much data reduction has been performed using other less expensive techniques.

The advent of supercomputers such as the CRAY has now made it possible to use wave equation methods prior to data reduction. The ability to do such processing enables us to develop interpretable images from seismic data in geologic settings where it had previously been impossible.

Seismic exploration

Seismic exploration involves exciting the surface of the earth using either explosive or vibratory sources. The induced vibrations propagate downward and outward through the earth's crust in expanding waves much like the waves that radiate on the surface of water when a stone is thrown into it. Where the acoustic and/or elastic properties of the buried rock layers change, this radiating energy is partitioned and some of the energy continues to radiate downward but some is reflected and radiates back toward the surface. By recording this returning energy at the surface and analyzing the elapsed time between the excitation of the surface and the arrival of the returning energy, the geophysicist can construct an image or model of the geologic structure below the earth's surface.

Although the acoustic energy propagates through the Earth's crust in waves, it is much simpler to think of it as traveling along rays that obey the rules of geometric optics. This is how it will be treated here.

The amount of energy that returns to the earth's surface is relatively small and the instruments needed to record it are necessarily quite sensitive. Unfortunately, such instruments record many undesirable noises as well as the returning reflected waves. Some of these noises arise from sources other than our excitation of the earth. Others result from wave propagation modes that do not carry any information useful to our exploration objective.

A powerful weapon we can use against noise is redundancy. If our recording of data from a physical experiment is contaminated by random noise, we can repeat the experiment many times and sum the results. The repeatable parts of the recording, the signal, will be reinforced in the sum but the random part, the noise, will not. The result is an improvement in the signal to noise ratio of our data.

While the simple redundance of repeating the same experiment many times will discriminate against random noises, it does not offer much help with systematic noises like surface waves and reverberations (multiples). A further disadvantage of simply repeating the same experiment over and over is that it takes a lot of time. Both of these problems are attacked by the common depth point (CDP) method.

Common depth point (CDP)

In the *common depth point method,* we combine the results of several separate physical experiments to improve the signal to noise ratio of the data. This method differs from simple redundancy because the experiments used differ from one another.

Figure 1 illustrates the CDP method. If we assume that the surface of the earth is (locally) flat and that all rock interfaces below the surface are also flat, and if we

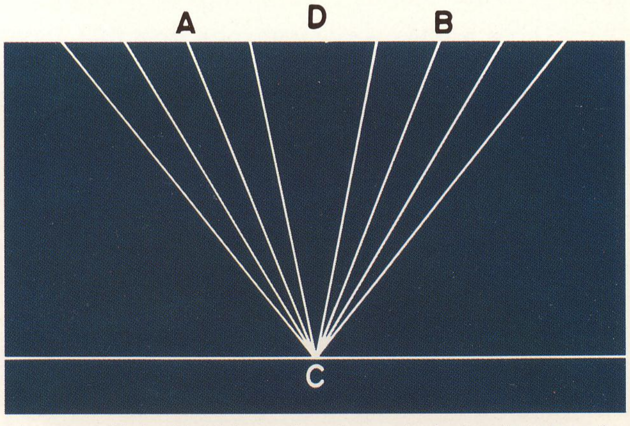

Figure 1. If the Earth's surface is excited at point A and a recording is made at point B, it will record a reflection from C, directly below D, the midpoint between A and B.

excite the surface of the earth at point A and record reflected energy at point B, the point of reflection on the interface shown will be at point C, directly below the midpoint between A and B. As the figure shows, we can make many experiments with different source and receiver locations and as long as the midpoint between the source and receiver is the same, the reflection point on the rock interface will remain the same.

The elapsed travel time from source to receiver is different for each pair shown because the travel distance is different. It is obvious that if we sum traces having different travel times, we will not reinforce the events at all. However, for each of the traces, knowing the arrival time of the reflection event and the distance from source to receiver, we need only the propagating velocity of sound in the earth and the Pythagorean Theorem to calculate the vertical travel time from the midpoint D to the reflection point C and back. Expressed mathematically:

$$T_0^2 = T^2 - (X/V)^2$$

The vertical travel distance is the same for all traces that have the same (common) midpoint between their source and receiver locations. Thus, if we correct all the travel times to the vertical travel time, we can then sum the traces and the reflection events will reinforce one another.

The piece of information that we do not have for this correction is the velocity of sound in the earth. A process called velocity analysis performs the time correction and summing for a suite of different velocities and the velocity for which the sum has the best signal to noise ratio is selected as the correct velocity. The process of summing the traces is called stacking and consequently the velocity selected is called the stacking velocity. Where there are many reflecting rock interfaces there is not a single best stacking velocity but a set of velocities that vary with vertical travel time.

When all the traces in a common depth point set have been corrected to vertical travel time and stacked into a single trace, that trace is used as if the source and receiver had both been located at the common midpoint of all the source/receiver pairs in the set. By performing this travel time correction and stacking at many CDP locations along the seismic line we produce a seismic cross section such as that shown in Figure 2a.

Field records

While the common depth point method deals with sets of seismic traces that have a common midpoint location, it is far more efficient to gather data using a single source location and many receiver locations. For each receiver in this configuration the location of the midpoint between source and receiver is different and these traces are not suitable for stacking.

By moving the whole configuration of source and receivers along by an appropriate distance, we can record another set of traces for which most of the source/receiver midpoints coincide with the midpoints from the previous position, but the source-to-receiver distance for each midpoint is different. By doing this repeatedly, it is just a sorting task to gather sets of traces with the same midpoint but different source-to-receiver distances. Such sets, called common depth point gathers, are the sets of traces upon which the velocity analysis and stacking will be performed.

Migration

In petroleum exploration the objective is to locate "traps" in the subsurface where hydrocarbons could accumulate. In certain geologic settings it is even possible that seismic data will contain so-called "direct hydrocarbon indicators" which are usually variations in reflection strength caused by the presence of hydrocarbons in the rock pore spaces.

One problem faced by the geophysicist when using reflection seismic data to construct a geologic model is that the seismic recordings measure elapsed travel time while the interpreted model must describe depth. Figure 2 illustrates one facet of this problem in a simple way. While the detectors used to record the returning energy are generally omni-directional, the trace that represents the returning energy is displayed as deflections about a vertical base line beneath the detector location. As a result, the time image of the anticline, shown in Figure 2a, appears much broader than the actual anticline as shown in Figure 2b. The

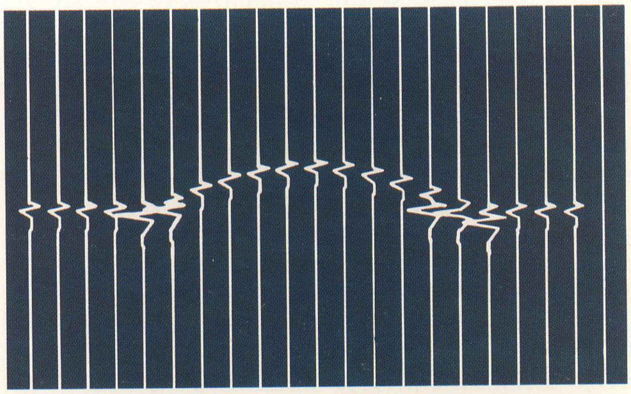

Figure 2a. The apparent width of the anticline on the seismic section is greater than its actual width.

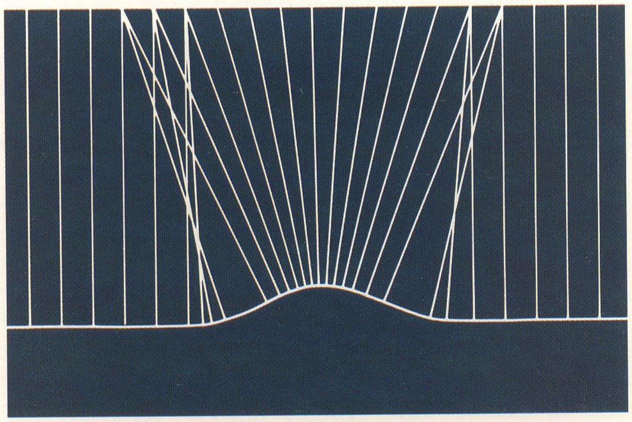

Figure 2b. Rays radiating from the top of an anticline illustrate the cause of one type of distortion that effects the seismic image.

rays radiating from the anticline to the surface in the model represent the minimum travel time path from each surface location to the reflector and back.

Not only is the time image of the anticline distorted, but the image of anything below it will also be distorted because the travel direction of energy passing through the anticlinal interface will be changed by that interface. One can think of each rock layer as a lens that may distort our view of everything beneath it if it is not flat, uniformly thick and homogeneous in its acoustic properties. Geology meeting those requirements is unlikely to contain hydrocarbons because it lacks a trapping mechanism.

Migration is a process by which distortions can be removed from the seismic data. The earliest migration methods were performed on the desktop using drafting tools. Methods were developed later to migrate seismic data using analog computers, and when digital computers came into common use for processing seismic data, it wasn't long before there were programs written to migrate digital data. In the early 1970s, wave equation migration algorithms were introduced and soon became the preferred methods.

These migration methods were all capable of removing the type of distortion inherent in the manner of recording and displaying seismic data, as illustrated in Figure 2. But they were not able to remove the distortion caused by the fact that each rock layer acts like a lens that affects the view of everything below it. It wasn't until the late 1970s that programs were developed that could properly account for and remove the distortion caused by this lensing effect. Such migration programs have been given the name *depth migration.* By default, the older methods that cannot remove the lensing effects are called *time migrations.*

Depth migration demands much more computer time than does time migration. This has caused some resistance to its widespread use. A less obvious perceived obstacle to the use of depth migration is that to properly remove the distortion caused by overlying geologic structure, we must be able to describe that structure accurately. This has led to the complaint that "to use a depth migration program you have to know the result before you can migrate the data."

The answer to the first complaint lies in the demonstrably superior results. The second complaint provides its own answer and actually points out one of the primary advantages of depth migration. If an interpreter describes a geologic model and uses it to migrate the seismic data, the result will tell him if the model was correct. If the model was correct, then the migrated result will conform to the model; if the model was not correct, the migrated result will provide direct guidance for making changes to the model. Thus, depth migration tests the geophysical interpreter's model of the geology and provides guidance for improving the model. Migration then becomes an iterative process that can be used to develop a model of the geologic cross section.

The idea of using a depth migration program iteratively when we think of the computer resources required is worrisome, but the speed of CRAY supercomputers makes the whole idea quite practical. Experience using a CRAY-1/S shows that depth migrations that take several hours of elapsed time on conventional mainframe computers take only minutes on the CRAY.

CDP assumptions violated

Just as the lens-like behavior of the rock layers can distort the geophysicist's view of the deeper layers, it also violates the basic CDP assumption that all the layers are flat and homogeneous.

Figure 3 shows the same anticlinal structure as in Figure 2 except an interface is added below the anticline and the rays representing a CDP gather are shown. We can see that, although there is a common midpoint on the surface, the reflection points are anything but common. In such a situation the CDP stack will probably produce a weak image at best and that image will not represent a reflection at a single location but a kind of average over an area. Such an image cannot be expected to produce high resolution data needed to derive a good geologic model.

Unfortunately, those parts of the Earth's crust that have some of the best potential for containing hydrocarbon traps also violate the CDP assumption. For those areas we need another way to process the data to develop an accurate image of the geology.

Depth migration before stack

Experience with depth migration of stacked data has demonstrated that lensing effects can be removed. It has now been demonstrated that depth migration of seismic data before stacking can prevent the degradation of the stacked image by that same lensing.

The algorithms used for this depth migration before stack are based on wave theory, which describes physically realizable experiments. It follows then, that we must use them to migrate the common source records that are recordings of actual experiments used to produce the seismic data.

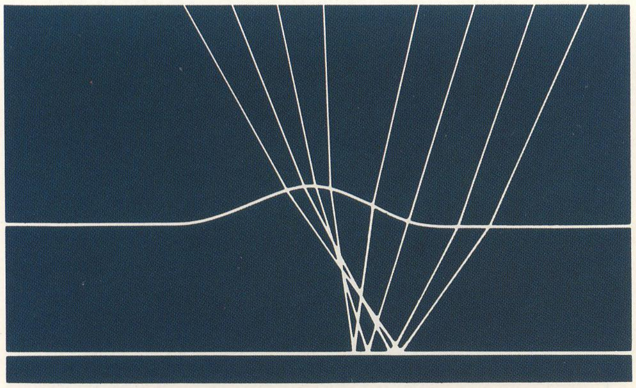

Figure 3. Even though the source/receiver pairs all have a common midpoint, the reflection points are different because the anticline acts as a distorting lens.

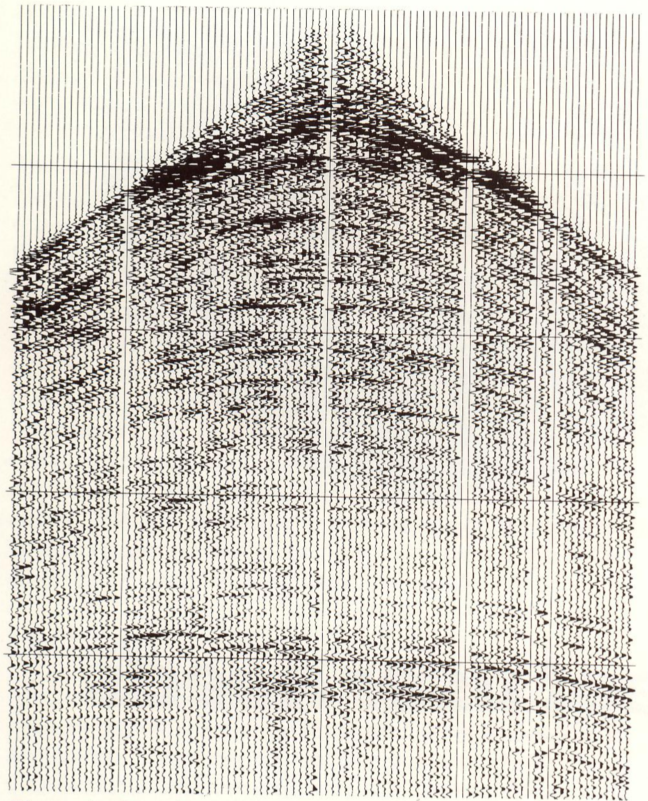

Figure 4a. Sample seismic record before depth migration.

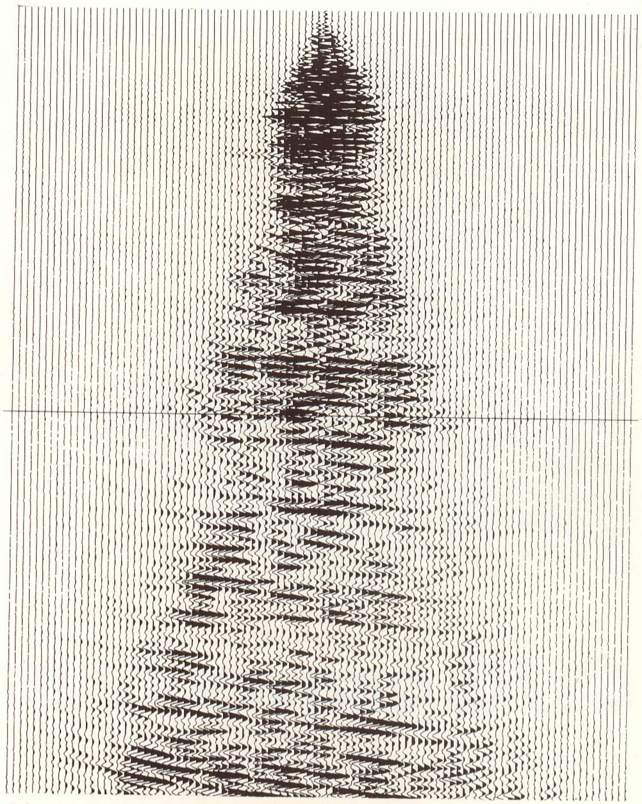

Figure 4b. Same record after depth migration.

When we depth migrate a common source seismic record, the result is an image that looks like a geologic cross section. If the velocity model used in the migration was correct, the migrated image will confirm this; if it was not correct then the migrated image will provide guidance for correcting the model. This result is much like that obtained by depth migrating stacked seismic data But we obtain additional leverage from the fact that the images obtained from depth migration of several records that cover overlapping parts of the geologic cross section must match one another as well as the velocity model.

Where overlapping migrated records do not match one another, the amount of mismatch can be measured and used to correct the velocity model. Thus, the depth migration of common source records can be used iteratively to determine the propagation velocity within the rock layers and the location of the layers.

Figure 4a shows a common source seismic record before it was depth migrated. The same record after migration appears in Figure 4b. It contained 96 traces, each six seconds long and sampled at an interval of four milliseconds. The total number of samples in the record is 144,000. The migration algorithm used is recursive and required 200 major iterations through the dataset. Each major iteration contained a mixture of ten recursive and nonrecursive iterations. When run on a single processor of a CRAY X-MP/48, this record required 51 CPU seconds to migrate.

Several seismic source records can be migrated independently using this method. Among other things, this means that one can make use of the multiprocessing capabilities of the CRAY X-MP simply by dividing the records to be migrated into separate data streams that can be processed simultaneously.

Applications

The most obvious use for depth migration of common source seismic records is to develop an image of the geologic cross section in prospective areas where the assumptions of the CDP method are violated. There are many exploration provinces in the world where this is the case and there has been hitherto no good way to explore them seismically.

Other possibilities exist however. For example, there may be situations where it is impossible or uneconomical to record the many separate experiments necessary to have enough data to form common depth point gathers. Such situations might include vertical seismic profiling in a well bore or the use of reflection seismics in petroleum reservoir engineering.

Conclusion

The introduction of wave theory based methods of processing seismic data has led to a great improvement in our ability to process and interpret that data. The great computer resource requirements of these methods has limited their use to stacked seismic data.

The availability of the CRAY computers has made it practical for us to begin using these methods to process data before stacking. These developments will enable us to use reflection seismology in areas where it has heretofore not been effective. □

Cray User Group

CUG grows to meet user needs

People with common interests naturally meet to share notes, compare ideas and swap yarns. But if one's interest is supercomputing, one's colleagues are scattered across the globe. Bringing this far-flung community together is the mission of CUG, the Cray User Group.

CUG is an international association organized and run by Cray customers. Begun in 1977 as a kind of ad hoc support network for CRAY computer users, CUG has grown from a handful to over 60 member installations today. The latest step in its maturation was an extensive reorganization last fall. As a result, the organization's bylaws have been rewritten and its steering committee replaced by an elected board of directors. Current plans call for CUG's formal incorporation later this year.

The group's reorganization was forced by its own rapid growth. "When there were still only 20 or 30 member sites, we were able to keep things pretty informal," said Bob Price, manager of Engineering Computer Systems at Westinghouse and CUG treasurer. "But with our current membership, we've had to make it more structured just to keep things organized." Membership is granted to all sites with a CRAY computer installed or with a contract for an installation. Each member site is granted at least one voting delegate to CUG, but any number of bona fide users of a member's services can attend CUG activities as participants. (Sites with three to five CRAYs are eligible for two delegates. Those with six or more are eligible for three delegates).

Membership in CUG doesn't bring the usual benefits such as monthly newsletters or discounts, however. "The main benefit to members is the exchange of information the network provides," said Price. Virtually all CUG members agree.

"For me, it's a chance to find out what new installations are up to and to talk with other operations managers about what it's like to run a site with a CRAY," said Gary Jensen, operations manager at the National Center for Atmospheric Research (NCAR) and chair of CUG's operations workshops. "As a user, I think anyone who stays away from the CUG meetings is really missing out."

The early years

CUG began informally in an Albuquerque, New Mexico hotel room. A small group of interested users from NCAR, Los Alamos National Laboratory and the National Magnetic Fusion Energy Computer Center, along with a few Cray representatives, got together to discuss their interests and concerns. "The main point was to share common experiences," said Gene Schumacher, group head of systems programming at NCAR and regular attendee of the early CUG meetings. "There were not too many of us (CRAY users) at the time. We wanted to see if we had similar concerns and to let Cray Research know about them."

"I saw the meetings primarily as a forum to discuss software development," said George Grenander, a Cray marketing employee during CUG's early years, now a consultant to the company. "At that time there was little software to run on the CRAY, and what existed was newly developed. The customers wanted to have some input into the company's software development and, at the same time, they wanted to work together and share ideas — to see what the other sites were up to so that they wouldn't waste their time reinventing the wheel. It was a kind of balancing act for us — trying to work with the existing customers and doing software development to attract new customers."

The early informal meetings became regular events. New installation sites were invited to join and were often asked to host the meetings. During the course of the last few years, the meetings have evolved into large semi-annual conferences, typically with every third conference hosted by a European site. These meetings have changed not only in their size, but also in their orientation. "Initially the highest levels of customer management were involved," said Grenander. "But now CUG involves more technical people, typically operations, applications or software development managers, or team members." Gene Schumacher concurred, "When it started out it was oriented towards systems programming, but now CUG has grown to include more operations and applications."

CUG today

CUG conferences give CRAY users an opportunity to present papers, exchange ideas and experiences and to learn from company representatives something about Cray's status and current research. "A real high point is when Cray Research gets up there and opens up about what's going on," said Gary Jensen. "It's refreshing to hear a company talk that way instead of just reciting some 'Corporate Plan'."

The conferences typically feature guest speakers and a large menu of technical workshops on topics such as operations, COS, CTSS, networking, graphics, languages and multitasking. "Networking is a particularly hot topic now," noted Helene Kulsrud, research staff member at the Institute for Defense Analyses and CUG vice president. "Most users originally front-ended their CRAYs, but now they're doing more networking. There's also a lot of interest in multiprocessing and multitasking."

The speakers and workshops cover all levels, Kulsrud added. "In the parallel sessions, for example, we have a number of workshops that address 'how to' questions, like how to optimize code or how to report statistical results to management. There are also technical areas that involve the mathematics used in various applications. We're starting to include tutorials on more specific topics, such as how to schedule jobs efficiently."

CUG has also begun inviting speakers who are not necessarily from member sites or who are not CRAY users themselves. "We recently had some people come from Bell Labs to talk about UNIX, and they were not using the Lab's CRAY," said Kulsrud.

"We also had a speaker from Network Systems Corporation," added Karen Friedman, writer/editor for NCAR's systems programming department and CUG secretary. "It's important for us to see what other vendors are doing that relates to CRAY use."

All CRAY users are invited to submit technical papers for review by the conference program committee. Accepted papers are included in the conference proceedings, often in digest form, and authors given the opportunity to present their papers. "We try to set up a program for two meetings ahead, though that's pretty ambitious," said Friedman. "We do have to plan things four months ahead though, so papers should be submitted for the next meeting as soon as possible after the last one."

The next CUG conference is scheduled for the fall of 1985. The Canadian Meteorological Centre in Dorval, Quebec will host. "The theme for the meeting will be performance evaluation for the CRAY X-MP Series, so papers on this topic will be particularly welcome," said Kulsrud. The Dorval conference will also include the election of new officers. The conference after that is scheduled for Seattle (in spring 1986), with Boeing Computer Services hosting.

Although CUG is an independent organization, it has vigorous supporters within Cray Research. "As far as I'm concerned, the CUG conference is the premiere supercomputer conference," said Dave Sadler, Cray corporate interface to CUG. "Ninety percent of the supercomputer expertise in the world is represented there. The most recent conference had over 200 participants."

While CUG members acknowledge that Cray Research representatives add to the interest and value of the conferences, Cray has also benefited from CUG's input. "Among other things, they keep us honest," said Sadler. "And they do a good job of pointing out areas for the company to address."

"In the early days, we brought reassurance back to Cray management that we had happy customers," added Grenander. "We also learned from customers how fast the applications were running on our machines and we learned about and incorporated some of the customers' programming techniques. For example, ideas developed during early work on Fast Fourier Transforms at NCAR and ECMWF (European Centre for Medium-range Weather Forecasts) were eventually incorporated into some of our scientific software packages."

CUG's effort to foster the exchange of information among Cray users is a valuable service to this specialized community. As CUG matures into a professionally managed corporation, it will become even better able to serve the needs of CRAY users. "It's becoming much more professional," said Vice President Kulsrud. "We're getting more and more technical papers and working to improve the quality of the printed proceedings. At the same time we're trying to keep the flavor informal so that people can get to know everyone they want to know and exchange information freely." Although CUG is an autonomous organization, Cray Research supports its efforts to build cohesiveness and self-awareness among Cray computer users. □

CORPORATE REGISTER

Cray ships 100th system

A significant corporate milestone was passsed in March with the shipping of the 100th Cray supercomputer. In 1976 Cray estimated that the entire market for supercomputers would only be 80 to 100 systems. Nine years later, some experts are pegging the supercomputer market at about 1000 computers. Who knows how long it will be before Cray ships its 200th system? The way the worldwide supercomputer population is growing, chances are good we won't have to wait another nine years.

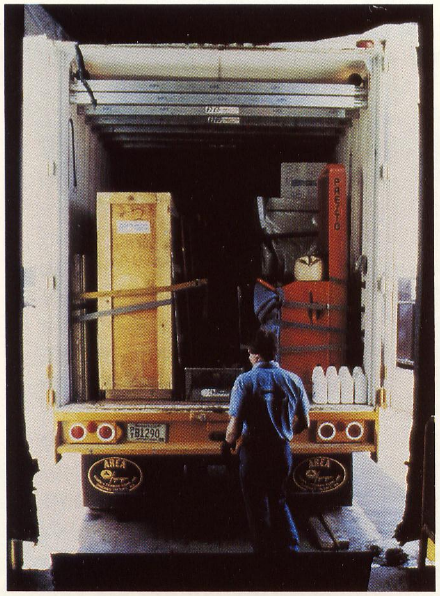

Cray's 100th supercomputer heads out from Chippewa Falls.

Japanese auto maker to get CRAY

At the company's annual meeting in May, Cray Chairman John Rollwagen told shareholders that Nissan Motor Company, Ltd. of Japan has announced its intention to obtain a CRAY supercomputer. The order, for a CRAY X-MP/11, is subject to negotiation of a final contract and obtaining an export license, Rollwagen said. The system currently is scheduled for installation in the first quarter of 1986. This will be the fifth CRAY system installed in Japan. The others include systems recently installed at Toshiba and Nippon Telephone and Telegraph.

Cray announces international orders

In March, Cray announced that the Abu Dhabi National Oil Company (ADNOC) ordered a CRAY X-MP/14 computer system. The system, which will be purchased, is to be installed in the fourth quarter of 1985 at ADNOC's headquarters in Abu Dhabi, subject to U.S. export license approval. The system will be used for oil reservoir engineering. ADNOC and its subsidiary companies operate large oil fields that require advanced management techniques supported by the Cray supercomputer.

Adam Opel AG, an automotive company located in Ruesselsheim, West Germany recently ordered a CRAY-1 S/1000 computer system. The system will be used for automobile development, primarily in the areas of structural analysis, aerodynamics, kinematics and CAD/CAM. Installation is scheduled for mid-1985.

Shell UK Limited of London has ordered a CRAY X-MP/14 computer for installation during the third quarter of 1985. The new system will replace Shell's existing CRAY-1 and will be used by the operating companies within the Royal Dutch/Shell Group for reservoir simulation studies and other technical applications.

NSF funding presents academic research options

On February 25, 1985, the National Science Foundation awarded grants of about $200 million to establish supercomputing centers at four U.S. universities. Two of the four centers plan to install Cray computer systems in the near future. Three grants will be spread over the next five years (one will be spread over three) and each of the four recipients will receive between $7 million and $13 million in 1985.

One center will be located at the University of Illinois in Urbana-Champaign. Illinois plans to install a CRAY X-MP/24 in the third quarter of this year and upgrade to an X-MP/48 in the future.

The University of California, San Diego, was also awarded a grant to establish a center that will be operated by personnel of GA Technologies. The center was proposed by GA Technologies with the support of 18 academic and research institutions across the U.S. A CRAY X-MP/48 will be installed at year-end, and the center will be fully operational in early 1986.

Contracts for the University of California, San Diego, and University of Illinois Cray systems are currently being negotiated.

The NSF also awarded grants for the establishment of supercomputer centers at Cornell University and Princeton University. The grant to Cornell will extend for three years. The university will install an IBM 308x with add-on vector units from Floating oint Systems; Princeton will install a Control Data CYBER 205, with plans to install an ETA GF10 when available.

The NSF initiative is part of the Federal Government's supercomputing development plan. In 1983, the White House Office of Science and Technology Policy asked the Federal Coordinating Committee for Science, Engineering and Technology to investigate supercomputer R&D issues, based on the results of the Lax Panel Report, scientists' activities and concerns about America's ability to remain internationally competitive in many areas of basic research. Specifically, the committee was asked to recommend ways to improve the nation's scientific and engineering researchers access to high-performance systems, to recommend a research and development plan that would lead to much faster computers and software and to coordinate long-range research issues.

n spite of increasing requirements for computing within academia, few of today's supercomputers have been installed at colleges and universities in the U.S. In the late 1960s the National Science Foundation had a program to install large (at the time) scientific processors for use in academic research programs. The NSF effort ended in about 1970 and until this year there had not been a comparable program. Without that program, the installation of state-of-the-art scientific computers at U.S. colleges and universities was halted for about a decade. In fact, no CDC 7600s were ever installed in U.S. universities. However, several European colleges and universities did install 7600s and more recently, CRAY-class systems. In recent years, many U.S. academicians have had to look to European university computing centers to access supercomputers.

The new government program addresses two distinct areas:

- ☐ Access — providing access to supercomputing for university researchers, and
- ☐ Research — developing R&D programs to address the theory, design and software for future supercomputer systems. Several Federal agencies are participating in the program including the Department of Energy and the Department of Defense Advanced Research Projects Agency (DARPA).

The NSF program addresses the former segment of the program. John Connolly of the NSF explained, "We're not interested in designing new computers. Our main interest is in computers as a tool to solve scientific and engineering problems and to train students." It is expected that several additional university supercomputer contracts will be awarded over the next two years.

Cray Research is very pleased with the National Science Foundation program for several reasons. Certainly, the company is pleased to have the opportunity to install two additional systems this year. But more importantly, over the long-term, the availability of supercomputers in the university arena gives scientists the opportunity to conduct research that was not possible before — perhaps not even thought of before. The results of that research are sure to benefit both government and industry.

Recognizing that fact, Cray Research is sponsoring a corporate grant program that will fund selected research and development proposals submitted by universities. The company is most interested in R&D projects relating to the development and conversion of applications for use on CRAY systems. Among the application areas given special consideration are chemistry, medical science, computer science, electronics and many types of engineering. Other projects that Cray expects to consider will focus on algorithm development, work on selected system related software such as compilers, libraries, networking and operating systems, and other research projects.

As university students and faculty begin to conduct research in a variety of disciplines on supercomputers, new applications programs will emerge. And the experience gained with supercomputing at the university level will help ease the current shortage of supercomputing expertise in industry today.

Currently, Cray and the University of California, San Diego supercomputing center are discussing research and development work in computational chemistry, bioengineering and graphics. Likewise, the University of Illinois and Cray are discussing R&D work in parallel processing and operating system related topics.

Cray is pleased to be a part of this national program that promises to further many areas of technology, and is looking forward to working with universities on several levels to help ensure the success of their installations as the supercomputing projects get underway.

New Cray facilities evidence manufacturing, support demands

"Sometimes it seems like we're in the business of building buildings and not supercomputers!" was the exasperated comment from Chippewa Falls. One can certainly appreciate the sentiment. Between now and the end of 1985, employees at Cray's Mendota Heights and Chippewa Falls operations will live through about ten moves and expansions. Of course, it's not really surprising; moving departments to new locations seems to occur regularly at Cray as the company continues to grow.

The moves and expansions taking place involve primarily divisions dedicated to manufacturing systems and supporting customers. It stands to reason — the larger the Cray computer population, the greater the customer support requirements.

Some of the groups involved in the upcoming moves include software and customer training, technical publications support, CRAY-2 production, CRAY X-MP manufacturing, printed circuit board development and the distribution center.

A new 49,000 square foot facility is under construction at Riverside for CRAY-2 production. At the main Chippewa Falls complex, the manufacturing division is moving into a 100,000 square foot building and is completing construction on a new 26,000 square foot facility. Both of these buildings are located in the same industrial park as the main complex.

In total, manufacturing production and administration space is being increased by more than 180,000 square feet. Expansion activities will be completed by fall of this year, and things should settle down a bit. But by then, company growth may dictate yet another round of building and moving.

Pictured above is the first CRAY-2 building. Construction on a second CRAY-2 facility nears completion.

CRAY family grows in MH

Probably most visitors to Cray's Mendota Heights facility over the past few months have had a chance to sneak a glimpse of the CRAY-2 installed in the curtained computer room. So the information in this article may be a little old to some. But for those who haven't ventured to Minnesota and scrutinized the new family member, we thought we'd tell you a little about the CRAY-2 hidden behind closed doors.

Late in August last year, the first prototype CRAY-2 was moved from Chippewa Falls to Mendota Heights for continuing software development. The system contains one processor and four million words of central memory, and is commonly known as a CRAY-2 quad. (A full production model CRAY-2 includes four processors and 256 million words.)

A few special arrangements had to be handled in order to install the CRAY-2 in Mendota Heights, but by and large, it was an easy addition. Before installation, a water chiller was installed in the mechanical room for the heat exchanger. In addition, a three-inch dam was built under the floor around the perimeter of the system in case of a leak. If the liquid were to leak, a vacuum would suck it up from the dam and the fluid would then be shipped back to manufacturing for filtering and refinement before being used again in the system. Once the CRAY-2 quad was onsite, it was up and running within 24 hours — not bad for a prototype.

For some time before the quad was moved to Mendota Heights, software had been running on the system in Chippewa Falls. Since joining its Mendota brethren, the CRAY-2 has seen more and more activity. A good portion of the software development work is done on the CRAY X-MP, then shipped to the CRAY-2 via NSC HYPERchannel connections for testing and refinement. A VAX 11/750, available via HYPERchannel, is used for software development. In addition, an AT&T personal computer connected to the CRAY-2 functions as the operator console and four DD-29s are linked to the mainframe I/O system.

Software development has continued at a rapid pace. With the installation of the first CRAY-2 at a customer site, it is expected that development activity will continue to heat up. The fact that Software Development has had easy access to the system for over a year is evidence of Cray's commitment to strong software support for the new generation of supercomputers.

APPLICATIONS IN DEPTH

RELAP5 converted for CRAY use

The nuclear reactor safety analysis code RELAP5 was recently converted to run on CRAY machines. The latest version, RELAP5/MOD2, can execute on CRAY-1 and CRAY X-MP computers operating under COS 1.13 or 1.14, or CTSS.

RELAP5 was developed to describe the behavior of light water reactor transients. It can be used to analyze large and small break loss-of-coolant accidents, operational transients, transients in which the entire secondary system must be modeled and system behavior simulation up to the point of core damage. The controls, turbine, generator, condenser and feedwater system can be included.

The program's hydrodynamic model is integrated in time using a semi-implicit finite difference scheme which is stable for time steps less than the material Courant limit. Heat transfer processes are modeled by means of "heat structures" in which a transient heat conduction solution is used with a variety of boundary conditions including convective heat transfer to fluid control volumes. The heat structures can be used to model nuclear fuel pins, steam generator tube wall and piping system boundaries with environmental heat loss. The reactor kinetics model is a point formulation and includes moderator, doppler and boron concentration feedback. Reactor controls are silmulated by means of control components such as summers, function generators, integrators, differentiators, delay lines, lead/lags and a rotating shaft for coupling of turbines, pumps, generators and motors.

The code uses a five-equation, two-phase flow hydrodynamic model consisting of two phase continuity equations, the two phase momentum equations and an overall energy equation augmented by the requirement that one of the phases is assumed to be saturated. In this model, only two interphase constitutive relations are required — those for interphase drag and interphase mass exchange. Models are included for abrupt area changes, choking, mass transfer, interphase drag, wall friction and branching. In addition to reactor applications, the program can be applied to transient analysis of other thermal-hydraulic systems with water as the fluid.

RELAP5 is available to National Regulatory Commission-approved users and is supported by the Idaho National Engineering Laboratory. Those interested in additional information about RELAP5 on the CRAY should contact Richard Wagner, Reactor Systems Technology Division, Idaho National Engineering Laboratory, P.O. Box 1625, Idaho Falls, ID 83415. Acquisitions must be approved through the NRC through Dr. Richard Lee, at the Reactor Systems Research Branch, Division of Accident Evaluation, U.S. Nuclear Regulatory Commission, Washington, D.C. 20555.

ACSL updated

Advanced Continuous Simulation Language (ACSL) was developed for the purpose of modeling systems described by time-dependent, non-linear, differential equations and/or transfer functions. Typical applications include control design, chemical process representation, missile and aircraft simulation and fluid flow and heat analysis.

ACSL is now available in a version that runs on CRAY X-MP systems under COS 1.14. ACSL features include on-line interaction through any terminal supported by the monitor, FORTRAN interface capability and choice of integration routines (including Gears stiff algorithm). There are no limits on problem size since dynamic tables expand to the full extent of available memory. ACSL follows standards established by the Continuous System Simulation Language Technical Committee in 1967. Array capabilities are complemented by vector and matrix integration operators. Simple structures with default options are included for the novice and advanced facilities for the experienced user.

APPLICATIONS IN DEPTH

For more information on ACSL on the CRAY contact Mitchell and Gauthier Associates, Inc., P.O. Box 685, Concord, MA 01742; telephone: (617) 369-5115

Math and statistics library available on the CRAY

Many CRAY users need a comprehensive on-line selection of subroutines for mathematical and statistical applications. The IMSL Library, a collection of 540 mathematical and statistical subroutines for scientific problem-solving, is designed to meet that need.

The IMSL Library addresses mathematical problems including:

- ☐ Differential equations, quadrature and differentiation
- ☐ Eigensystem analysis
- ☐ Interpolation, approximation and smoothing
- ☐ Linear algebraic equations
- ☐ Vector/matrix arithmetic
- ☐ Nonlinear equations
- ☐ Optimization
- ☐ Linear programming and transforms.

The Library's statistical selection includes subroutines in areas including:

- ☐ Basic statistics
- ☐ Regression analysis
- ☐ Analysis of variance
- ☐ Nonparametric statistics
- ☐ Time series, forecasting and econometrics
- ☐ Observation structure and multivariate statistics
- ☐ Categorized data analysis and sampling.

Also included are some general-purpose subroutines, as well as subprograms for evaluating special functions and for generation and testing of random numbers.

The IMSL Library currently is used in widely divergent fields. In geological research it has found application in processing data about water occurrence, availability and quality and hydrologic processes. A leading biotechnology firm uses the Library's differential equation solvers to study the optimal conditions for the growth of microorganisms and the production of proteins. In addition, a major U.S. automobile manufacturer has recently integrated the IMSL Library into a geometric modeling system for designing industrial parts on a CRAY computer.

The IMSL Library is complete and well documented. For more information about the Library, contact IMSL, NBC Building, 7500 Bellaire Boulevard, Houston, Texas 77035; telephone: (713) 772-1927 or (800) 222-IMSL; telex: 79-1923 IMSL INC HOU.

Cray hosts science and engineering symposium

Seymour Cray shared his views on the CRAY-2 and CRAY-3 with symposium attendees.

Automotive, aerospace and petroleum engineers, astrophysicists, chemists, and even a chess champ convened in Minneapolis this spring for a Cray Research-sponsored symposium, "Science, Engineering and the CRAY." The three-day symposium treated participants to a smorgasbord of information on the latest developments in supercomputer applications. Experts in disciplines ranging from combustion chemistry and speech analysis to weather forecasting and nuclear energy discussed new code development and optimization techniques in their respective fields. Speakers shared their insights and experiences with over 350 scientists and engineers from industry, education and government. Representing the international supercomputing community were researchers from Canada, England, Finland, France, Italy, Japan, the Netherlands, Norway, Saudi Arabia, Sweden, Switzerland and West Germany.

Along with the formal presentations, the symposium gave participants ample opportunity to meet in small groups to discuss their particular interests. "We were especially pleased to see how much interaction there was between the commercial and university people," noted Derek Robb, manager, market development for Cray Research. During the second day's presentations, Seymour Cray addressed a standing room only audience to discuss the CRAY-2 supercomputer and his work on a follow-on system, the CRAY-3.

"The CRAY-2's 256 million words of directly addressable memory is a two-orders-of-magnitude increase in memory size over the original CRAY-1," noted Cray. "That's a shocking number, because never in my experience have we taken a basic factor in our computer design and talked about a one-order-of-magnitude increase, let alone two." Cray continued, "In the 1970s, everyone was talking about the computing business plateauing — that there were some real limitations — and we were making progress of factors of four each generation. Now, from the CRAY-1 to the CRAY-2, we're taking order-of-magnitude steps or greater, both in compute power and in memory size."

The symposium concluded with tours of the CRAY X-MP manufacturing and testing facilities in Chippewa Falls, Wisconsin, and graphics and benchmarking demonstrations at Cray's Mendota Heights, Minnesota facility. This was the second science and engineering symposium hosted by Cray Research in the past three years, with attendance at this year's meeting nearly double that of the previous one. "Not only did we see a dramatic growth in the number of attendees," said Robb, "but we saw a lot of interest in some newly developed applications, things like computational chemistry, particle physics and artificial intelligence."

Symposium proceedings are available from the Cray Research Distribution Center, Mendota Heights, Minnesota. A nominal fee will be charged to recover printing costs. For information, contact Dennis Abraham at (612)681-3091.


People from diverse fields attended the symposium to learn about supercomputer applications and programming techniques.

Following is an alphabetical listing of this year's symposium presentations:

Anderson, M. Paul, Ford Motor Company, *"Computational Requirements for Power Train Analytical Simulation"*

Berkhout, A.J., Delft University, *"Prestack Seismic Migration in Three Dimensions"*

Booth, Mike, Cray Research, Inc., *"Microtasking on the CRAY X-MP"*

Burridge, David, ECMWF, *"Global Weather Forecasting"*

Centrella, Joan, Drexel University, *"Studying the Universe on a CRAY"*

Cray, Seymour, *"The CRAY-2, CRAY-3 and Beyond"*

Dixon, David, DuPont, *"Large-Scale Modeling of Molecular Systems at DuPont"*

Douglass, Robert, Los Alamos National Laboratory, *"Artificial Intelligence Applications on Supercomputers"*

Fredriksson, Billy, Saab-Scania AB, *"Large-Scale Linear and Nonlinear Structural Analysis and Optimization"*

Hibbitt, David, Hibbitt, Karlsson & Sorenson, *"Non-linear Finite Element Problems"*

Hsiung, Chris, Cray Research, Inc., *"The Role of Multiprocessing and High-speed Computing"*

Hunten, Keith, General Dynamics, Fort Worth Division, *"Trials of Using CRAYs for Structural Analysis"*

Hyatt, Robert, University of Southern Mississippi, *"CRAY BLITZ — 1984 Chess Champion"*

Kollman, Peter, University of California, San Francisco, *"Simulations of Complex Molecules — Computational Requirements for the 80s"*

Komornicki, Andrew, Polyatomics Research Institute, *"GRADSCF: A Research Tool in Computational Chemistry"*

Levinson, Steve, Bell Laboratories, *"Speech Analysis Applications"*

Martin, Werner, Adam OPEL AG, *"Automotive Engineering on Supercomputers"*

McDonald, Alvis, Mobil, *"Vector Computer Applications in Reservoir Simulation"*

Meintjes, Keith, General Motors Research Laboratories, *"Toward More Realistic Computations of Reacting Flows"*

Moriarty, Kevin, Dalhousie University, *"Efficient Multitasking of the SU(3) Lattice Gauge Theory Algorithm on the CRAY X-MP"*

Patterson, Stuart, Jr., TCS Inc., *"The Technical Computation Facility"*

Peterson, Vic, NASA-Ames Research Center, *"Supercomputers and Computational Fluid Dynamics"*

Pople, John A., Carnegie-Mellon University, *"Gaussian 82 on the CRAY"*

Schmidt, Wolfgang, Dornier Aircraft, *"Euler Applications and Fine Mesh Aerodynamic Calculations on the CRAY X-MP"*

Schwenke, David, and Truhlar, Donald, University of Minnesota, *"Large-Scale Calculations on the Quantum Mechanical Description of Energy Transfer in Molecular Collisions"*

Smarr, Larry, University of Illinois, *"Black Holes in Living Color"*

Smith, Alvy Ray, Lucasfilm, *"The Making of Andre and Wally B"*

Stanisforth, Andrew, Canadian Meteorological Center, *"Numerical Weather Forecasting in the Canadian Weather Service"*

Vanderplaats, Garrett, University of California, Santa Barbara, *"Optimization Techniques for Design: The State of the Art"*

Wheeler, Mary, Rice University, *"Modeling of Miscible Displacement in Porous Media"*

Wolf, Malcolm, Systems Designers International, *"Synthetic Aperture Radar Processing on a Supercomputer"*

Woodruff, Susan, Los Alamos National Laboratory, *"Faster-Than-Real-Time Simulation of Nuclear Reactor Accidents"*

USER NEWS

Model for peptide bond analyzed

Biotechnologists are developing and refining ways to make a variety of biological substances, typically proteins or protein-based substances that occur naturally. They are also using site-specific mutagenesis to produce proteins with altered functions. Many observers believe that the next major step in biotechnology will be the engineering of synthetic protein molecules, such as hormones and enzymes, that are tailor-made to perform specific functions. But development of this technology — protein engineering — requires much basic research to unveil the chemical basis of protein structure.

Protein structure and dynamics are being studied theoretically in the group of Professor Martin Karplus at Harvard University's department of chemistry. The group is using a program for the building, energy minimization and dynamic simulation of proteins. The essential element of such a program is the empirical energy function that determines how the energy of a protein changes in different conformations. Since hydrogen bonds are one of the essential elements of protein structure, it is important to treat them accurately in the energy function. Walter Reiher, a graduate student in the Karplus group, has been studying hydrogen bonding between peptide groups and water, using the formamide molecule as a peptide model. Reiher is performing the calculations on a CRAY X-MP/48 using the program Gaussian 82 developed by Professor John Pople and his collaborators at Carnegie-Mellon University.

Modeling water-formamide hydrogen bonding is important in refining mathematical models of proteins because, in the solutions where proteins naturally occur, many hydrogen bonds exist between water molecules and the atoms of a protein. Reiher has also simulated formamide-formamide hydrogen bonding as a model for hydrogen bonding between peptide groups, which is important in determining the secondary structure of proteins. Since water-water hydrogen bonding has been studied experimentally, Reiher has studied hydrogen bonding between water molecules as a reference for the other models. These calculations provide information about the relative strengths of the hydrogen bonds that can form between water and peptide groups. Knowledge of the energetics of these hydrogen bonds is essential for exploring whether water is more likely to hydrogen bond with itself or with a protein.

"Formamide is the smallest molecule containing a peptide group," explained Reiher. "My objective is to calculate the energy cost of stretching and bending the hydrogen bonds in the molecule. This has to be done with high accuracy so that the results can be used to improve the mathematical description of proteins." Using several different approximations and very large basis sets (sets of functions for describing electron distributions), Reiher has been able to perform detailed and reliable calculations of water-formamide, formamide-formamide and water-water hydrogen bonding.

"The basic calculation is the Hartree-Fock energy evaluation," he explained. "I also performed Moller-Plesset and configuration interaction calculations, which are improvements over the Hartree-Fock method. These options are included in Gaussian 82. They give the molecular wave function, which describes the distribution of electrons in the molecules and gives the energy of hydrogen bonds between the molecules."

The level of accuracy that these calculations can achieve is determined in part by the size of the basis set. While larger basis sets yield more accurate results, they also significantly increase computation time. Although accurate water-water calculations can be performed on a VAX-11/780, the time required for extremely accurate water calculations — or for calculations of larger molecular systems using cruder approximations — can be prohibitive even on conventional mainframe computers.

"The use of the CRAY has allowed me to carry out some very large calculations," said Reiher. "The results provide useful references to compare with results from smaller calculations. I have found an empirical way to scale the results from the smaller calculations to give results that are in good agreement with the large calculations I carried out on the CRAY. No study has been published with calculations performed at the level of detail I achieved using the X-MP/48. The study would not have been possible without it."

Reiher has also run benchmark comparisons of Gaussian 82 on the CRAY X-MP/48 and on a VAX 11/780. Identical jobs were run using basis sets of various sizes and different approximations. The larger calculations ran significantly faster on the CRAY. "In general, the Hartree-Fock calculations ran with relatively smaller increases in speed, while the more detailed calculations were considerably faster," said Reiher. "For example, the integral transformation portion of the Moller-Plesset calculations of the formamide-water system would have required several days of CPU time on the VAX, whereas on the CRAY it took just under 23 minutes of CPU time. Based on these results I decided not to attempt formamide-formamide calculations on the VAX because of their size."

"We saw that the speed advantage of the CRAY over the VAX increases with the size of the molecular system being modeled," Reiher added. "Speed enhancements seen with both of these factors — the detail and the size of the calculations — are encouraging because the kinds of jobs we want to run on supercomputers are very detailed calculations of small molecular systems, or less detailed calculations of larger systems."

When the I/O required for the largest calculations became impractical using disk storage, Reiher transferred his files to the CRAY Solid-state Storage Device (SSD). "Using disks, I/O time was about the same as CPU time," he noted. "But the SSD brought I/O time nearly down to zero. We saw throughput increase by about a factor of two with the SSD."

The modeling research has two major implications, said Reiher. It demonstrates the ability to scale up small calculations to mimic larger, more accurate ones. And data now available will be used to reformulate the mathematical model of proteins.

The research formed the central part of Reiher's doctoral thesis in theoretical chemistry, but Reiher said he sees his work in the larger context of protein engineering. "My work really is a precursor to another type of calculation, molecular mechanics, which is a much more approximate way to characterize larger systems," he said. "It's hoped that through refining our mathematical model of proteins it will eventually be possible to design modified proteins for specific functions and determine binding sites for particular drugs. Other students in our group are already using the new potentials based on my calculations to predict changes in protein structure resulting from single-site mutations."

Not in Kansas anymore

Tornadoes are a recurring source of awe and tragedy in many areas of the United States. Common as they are, however, no one quite knows why they occur. But computer simulations recently executed on a CRAY-1 at the National Center for Atmospheric Research have taken scientists a step closer to understanding the birth of tornadoes. By combining CRAY power with interactive graphic displays, scientists Joseph Klemp and Richard Rotunno have discovered some significant features about the early stages of tornado formation.

Klemp and Rotunno's research involves modeling the structure of convective clouds which produce many of the local severe weather events we experience, such as thunderstorms, hailstorms and tornadoes. A convective cloud may cover an area many tens of kilometers square, soaring thousands of meters into the atmosphere. Consequently, portions of it are exceedingly difficult to observe directly in detail. But where direct observation may not be feasible, investigations can be carried out computationally.

Klemp and Rotunno used an extremely detailed model in their research that can generate as many as one billion bits of data in a single run. "This problem only became feasible in the last decade or so," said Rotunno. "The first cloud models had extremely low resolution, but the

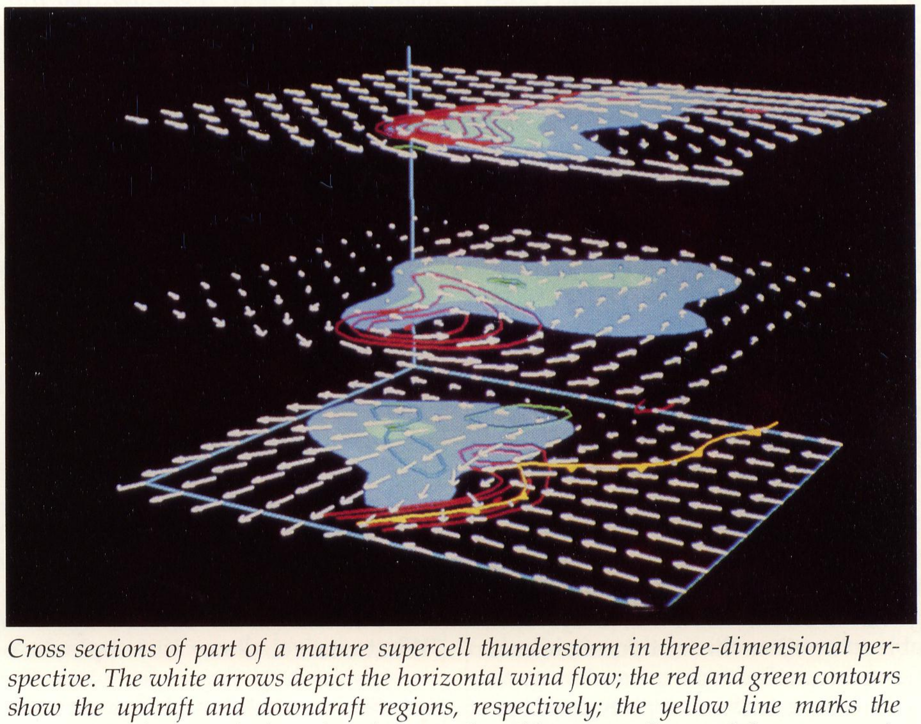

Cross sections of part of a mature supercell thunderstorm in three-dimensional perspective. The white arrows depict the horizontal wind flow; the red and green contours show the updraft and downdraft regions, respectively; the yellow line marks the storm's gust front — the boundary between the cold storm outflow and the warm environmental air; and the blue areas mark the precipitation fields, with the lighter blue indicating heavier precipitation. The indentation visible in the red and green contours (rain field) of the middle level, called a hook echo, is frequently observed by radar in supercell storms and indicates that a tornado may form.

field has really come of age since the CRAY went on-line."

"Essentially, the CRAY is necessary because we handle such large amounts of data," said Klemp. "The computation time itself depends on the number of grid points used, which varies considerably, but typically our runs are about 30 percent faster than real time."

The model Klemp and Rotunno used was originally developed by Klemp and Robert Wilhelmson of the University of Illinois, and is unusual in that its motion equations account for sound wave motion. "Sound wave motion is not believed to be important in convective processes, so most cloud modelers filter out the sound waves," said Klemp. "But that simplification forces them to solve Poisson's equation for pressure." To accommodate sound waves, which limit the size of the time steps that can be used, the model treats the sound waves in smaller time steps, separately from the rest of the code. "This way we achieve numerical efficiency in the model, and it allows us to compute all other processes using larger time steps that are more appropriate for those processes," said Klemp.

While it is not possible to test the cloud model's accuracy by studying the behavior of clouds in a laboratory, comparisons with real clouds are possible using Doppler radar. The radar is used to track the motion of water droplets inside storm clouds and these results are then used to assess the performance of the model.

Using the model, Klemp and Rotunno have made several discoveries about how convective storms generate tornadoes. First, they were able to monitor the motion of air as it flowed into the stormcell and observe how strong rotation at low levels of the cloud was produced. At low levels, air approaching a convective cloud or storm must pass along the fringe of the cold downdraft flowing out of the cloud. The meeting of markedly warmer and colder air generates strong rotation about a horizontal axis. As it is swept into the updraft region of the storm, its axis of rotation tilts into the vertical and becomes further concentrated. This concentrated vertical vorticity marks the transition of a storm into its tornadic phase.

"The significance of this discovery is that we now know low-level rotation originates in a different way than middle-level rotation, which arises from environmental shear," said Klemp.

Another important outcome of the scientists' efforts has been the development of a consistent theory for wall clouds, which project below the main body of a massive storm cloud and often indicate that a tornado is forming. Rotunno and Klemp were able to trace the parcel of air that forms the wall cloud. They found that the air feeding the wall cloud originates in the upper, colder portions of the cloud, revealing that the localized lowering of the cloud base — the wall cloud — occurs because a portion of the air feeding into the storm is much colder than its surrounding environment. This feature was not noticed in two-dimensional cross-sections but only became apparent in three-dimensional visualizations of the model data.

In spite of these advances, no theory yet exists to explain every phase of tornado formation. "We still don't know how the axisymmetric structure of a tornado fits in," said Klemp. "Most tornado research is based on axisymmetric studies, but storm cells are highly non-axisymmetric. We're still filling in the gaps between these two structures."

"What we've done is supplied a number of hypotheses for observationalists to test," added Rotunno, "such as the importance of the cold-air boundary, which was first reported by observation. More analyses of observed storms are needed to see how accurate our theories are."

Astronomer sees spots

Jupiter's Red Spot has intrigued astronomers since Galileo first aimed his telescope at the giant planet. Nowhere else in our solar system has a spot like Jupiter's been found and, so far, astronomers have been unable to explain Jupiter's. But now a Harvard astronomer, Philip Marcus, believes he has found a relatively simple explanation. Marcus has been running simulations of the jovian atmosphere on a CRAY-1 at the National Center for Atmospheric Research and has been able to generate Red Spots in the simulated atmosphere. Marcus says his results indicate that just two physical forces are responsible for the spot on Jupiter: wind shear and the Coriolis force.

Jupiter's Red Spot is a kind of super-hurricane — a swirling mass of jovian atmosphere that is uncannily stable. The spot has persisted since it was first reported by Galileo 300 years ago. It is also relatively stationary and is confined to an atmospheric band running the circumference of the planet just south of the equator.

"Researchers have taken two main approaches in explaining the Red Spot," said Marcus. "One has been to treat it as an isolated object. There was a theory that it was due to circulation around a mountaintop, for example. The second approach is more meteorological and considers the context of the entire planet."

Using the second approach, Marcus has determined that only two physical forces operate on a fast enough time scale to account for the Red Spot. One is wind shear, the sideways force generated by two gas streams flowing past each other in opposite directions. The other is the Coriolis force, which pushes the atmosphere on a rotating planet perpendicular to the planet's direction of rotation.

Computationally, Marcus' method is unusual for this type of problem While the majority of researchers modeling fluid flow use the finite dif-

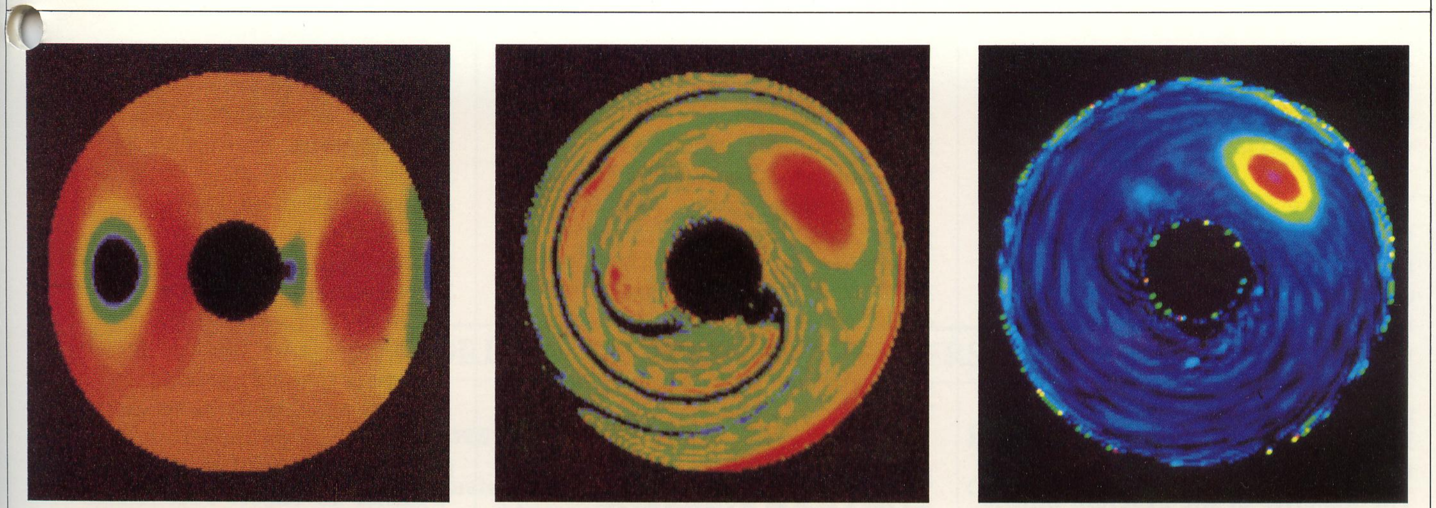

Three frames from a movie of a red spot simulation experiment. The region shown is between 20 and 26 degrees south latitude as projected onto a two-dimensional surface. The East-West wind moves clockwise near the outer edge of the disk and counter-clockwise near its inner edge. Initial starting conditions (left) show an anticyclonic spot (red patch) and a cyclonic spot (blue patch) of equal magnitude. The range of colors — red, orange, yellow, green, blue — represent different speeds of rotation, from the most anticylonic to the most cyclonic. After two jovian days (middle), the red spot adjusts and becomes slightly elongated while the blue spot is twisted into a spiral and small-scale chaos. Four jovian days later (right), only the red spot and chaos remain.

ference method, which divides a flow zone into a discrete number of grid points, Marcus defines each variable (velocity, temperature, pressure, etc.) as the sum of a finite Fourier series.

Each Fourier series used in a calculation can include up to 50,000 terms. "The program might step forward 300,000 terms in each time step and it might take 1000 to 2000 time steps for things to settle down," said Marcus. "Handling that many terms really requires a supercomputer. With the CRAY-1, we can do a single run in about five minutes."

"In the simulations I've run, there have been only two final results," Marcus added. "Either there are no stable vortices or there is one and only one anticyclonic (rotating in the direction opposite the planet's) circulating spot."

Like the simulated spots, Jupiter's Red Spot is singular and anticyclonic. When Marcus gave the simulation a spot rotating in the wrong direction, the spot either broke up or changed direction. When a random number generator bombarded the simulation with many mini-spots, they eventually coalesced into one big spot. "We added three-dimensional effects like convection and still got the same thing," said Marcus. "All of the numerous simulations we've run indicate that there can be one and only one spot."

In Marcus' simulations, a single stable spot forms within a time equal to about six jovian days. "I don't think the physics that other people have proposed would operate fast enough to play a role in spot formation," said Marcus. "But our physics, just the shear and Coriolis forces, is relatively simple, so if we're right we should be able to reproduce our simulated results in the lab." To that end, Marcus and a colleague are devising a laboratory apparatus incorporating a rotating sloping-bottomed bucket filled with water to simulate the Coriolis force. "The computer is good for initial experiments," said Marcus. "But in the lab we'll see if our physics is really right."

Although Marcus' theory remains to be verified, his work demonstrates the value of high-speed computing to basic research. Scientific advances ultimately rest on experimental proof, but, as Marcus' research into the Red Spot shows, computer simulations can indicate which experimental approaches are most likely to be fruitful.

The 29th is A-OK

About a year and a half ago, the 29th Mersenne Prime was found in record time by the first six-column CRAY X-MP. It's pretty exciting to think that a Mersenne Prime was discovered in only 3904 seconds on the new X-MP, but one must admit there was some luck involved. The 28th Mersenne Prime had been discovered one year earlier on a CRAY computer. Who ever heard of two Mersenne Primes being discovered in a year's time when only 29 have been discovered in the last 2500 years?!

Verifying the discovery was quite another story. It took about nine CPU hours for the old program to complete the verification, but that didn't worry the programmers. In fact, those who worked on verifying the number were really quite pleased. You see, the Lucas-Lehmer test, which is used to test the primality of a Mersenne number involves extremely complex mathematics, clever algorithms and programming, and a lot of compute power. Not only that, but as Mersenne Primes become larger, the time needed to verify them increases commensurately. (See CRAY CHANNELS Vol. 4 No. 1, p. 15.) Those involved in verifying the

29th Mersenne Prime say the task was the easiest of all the prime number verifications they have worked on so far.

The Prime Finder program used to verify the primes has undergone four rewrites to make it faster and more efficient. Dave Slowinski, one person involved in the Cray effort to discover new Mersenne Primes, explained, "The kernel of the Lucas-Lehmer test goes through a loop that squares large integers. The multiplication of the large integers must be verified. The majority of the program time is spent squaring and verifying them. The multiplication, called convolution, is executed via a very clever method called the Schonhage-Strassen Fast Fourier Transform (FFT) Multiply. It's highly vectorized and performs remarkably well on the X-MP."

Slowinski went on, "We were able to use the optimized FFT routines in SCILIB to implement the Schonhage-Strassen FFT. Earlier versions of the Prime Finder code included hundreds of lines of highly vectorized assembly language. Much of the newest version is written in FORTRAN and is faster than the old. The more efficient program and faster computer mean that the new program executes more than 100 times faster than the 1979 version run on a CRAY-1."

In addition to its importance in finding new Mersenne Primes, the Prime Finder code has a practical function. It uses a method called "residue check" to verify quickly the correctness of each of the thousands of 80,000-digit integer multiplies used to test potential Mersenne Primes. This residue check is executed dozens of times each second and immediately flags a hardware fault. Therefore it is a good system confidence test and is used by Cray Research in system check-out before new computers are approved for shipment. Because so many Cray customers use their systems for similar computations, the program is even more appropriate for system testing.

In any event, the 29th Mersenne Prime is found and verified. It is indeed a big number — 39,751 digits long. If you are interested in seeing the number, a new 29th Mersenne Prime poster is now available. The poster can be ordered from Creative Publications, Order Department, 5005 W. 110th Street, Oak Lawn, IL 60453; telephone: 1-800-624-0822, in Illinois: 1-800-435-5843. Reference catalog number: 93006. The cost is $3.50 plus $2.00 shipping charge.

Efforts continue on FORTRAN standard

It is often said that something good is worth waiting for, and the new ANSI FORTRAN standard is no exception. The standard is a product of the American National Standards Institute (ANSI), a voluntary organization of corporations and universities established to identify standards for many industries. The importance of system compatibility and measurements makes ANSI standards virtual laws for many industries — including the computer industry.

Now, X3J3, the ANSI technical subcommittee dealing with the FORTRAN standard, is compiling FORTRAN 77 revisions and additions into a new version. For several years the committee, composed of about 40 members from all areas of computing, has been developing the new standard that will eventually replace FORTRAN 77.

Dick Hendrickson, Cray's representative to the committee, commented on the status of the new standard, "Most of the technical work is completed. We are now in the process of finalizing the proposed standard for publication next year. But beyond that it will be a year or two before the new standard is approved and adopted." He went on to say, "Most committee members are truly concerned with good FORTRAN for the scientific community, as opposed to promoting features that benefit their cause. The broad range of considerations we have to look at makes it very time consuming to develop standards that will meet everyone's needs. One point we are committed to is ensuring that programs conforming to the existing standard will automatically conform to the new."

Some of the major new technical features include:

- ☐ Vector and array processing syntax that allows most vector and array operations to be done without DO-loops and subscripts
- ☐ A module facility that makes it easier to collect subroutines that share common data
- ☐ A facility that allows programmers to define data types and to define and do operations on those data types
- ☐ Free-form source input.

A release date for the new standard has yet to be set. If one were a gambler, one might wager that the new standard will be adopted in 1988 (The current FORTRAN standar was adopted in 1977. FORTRAN 66 reigned before that.) More optimistic predictions, however, put the date as early as late 1986 or 1987.

Oscar, meet the CRAY

It wasn't even considered for writing, acting, producing or directing, but a CRAY X-MP was acknowledged for its special effects contributions to film at this spring's Academy Awards presentation. The computer helped earn a scientific and engineering award for Digital Productions of Los Angeles, which uses its CRAY to generate images for motion pictures. Simulated scenes of the planet Jupiter that were created on the CRAY appeared in the film "2010" and about 25 minutes of CRAY-generated battle footage appeared in "The Last Starfighter". Digital Productions received the award for its ability to "create motion picture segments entirely from the imaginaton without the physical requirements of sets or props," according to the Academ Award citation. And who said only factory jobs are lost to automation?

WORLDWIDE BUSINESS CENTERS

Domestic

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

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

Central Region Office
5330 Manhattan Circle, Suite F
Boulder, CO 80303
(303)499-3055
Albuquerque
Chicago
Dearborn
Minneapolis

Eastern Region Office
11710 Beltsville Drive, Suite 500
Beltsville, MD 20705
(301) 595-5100
Atlanta
Boston
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
The Atrium, Suite 350
Pleasanton, CA 94566
(415) 463-2800
Los Angeles
Seattle

International

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

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

Cray Research GmbH
Perhamerstrasse 31
8000 Munich 21, West Germany
89-56014-0

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

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

CRAY
RESEARCH, INC.

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