A CRAY RESEARCH, INC. PUBLICATION

CRAY CHANNELS

Fall 1986

Announcing new models, improved performance for the CRAY X-MP series

FEATURE ARTICLES:

DEPARTMENTS:

IN THIS ISSUE

The popular image of oil gushers, spraying the sky black while workers celebrate beneath, says little about actual oil recovery operations. Efficient recovery requires skilled engineers using the latest technologies. Among the new technologies to be found in the reservoir engineer's toolbox is the reservoir simulator. Simulators enable engineers to model various recovery strategies on a computer before selecting one to use in the field. However, to run models in a reasonable time, the largest and most detailed simulators require powerful computer hardware. Many of the world's major petroleum companies rely on Cray computer systems to meet this need for high-performance computing.

This issue of CRAY CHANNELS presents articles related to underground fluid flow modeling in reservoir simulation and aquifer decontamination. This issue also announces the latest price/performance improvements in the CRAY X-MP series of computer systems. We'll take a look behind the scenes in Chippewa Falls, Wisconsin, where efforts are concentrated to monitor and enhance the reliability of Cray systems — with impressive results. And we celebrate the tenth anniversary of the first Cray system installation at Los Alamos National Laboratory.

When one considers geological complexities, such as fractures and rock composition, variables of multi-phase fluid flow in multiple dimensions, and the effects of varying temperatures and pressures, it is no wonder accurate reservoir simulation is an immense computational chore. It is also little wonder that Cray computer systems are increasingly being called upon to accommodate the petroleum industry's computing needs.

On the cover is a CRAY X-MP printed circuit board about to receive an electroplating bath. The bath solution selectively deposits copper on the board, forming the board's circuit lines. Following the electroplating bath, boards receive a similar bath that deposits solder over the copper circuit lines, creating an etch barrier and providing protection from contaminants.

CRAY CHANNELS

A Cray Research, Inc. publication

Fall 1986

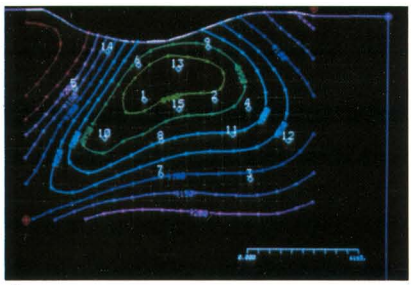

4

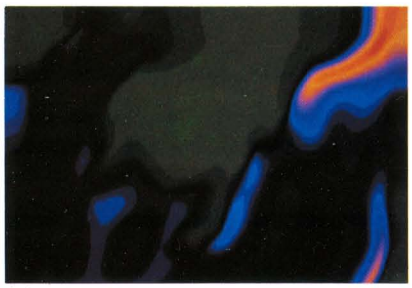

14

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

Volume 8, Number 3
Fall 1986

Editorial Staff
Tina Bonetti
Ken Jopp
Carol Smith
John Swenson

Design
Eric Hanson
Cindy Rykken

FEATURES

DEPARTMENTS

Announcing the enhanced CRAY X-MP series of computer systems

Improved performance, new models mark a continued evolution of the series

Product development at Cray Research doesn't stop when a new computer system goes into production. The company has an ongoing commitment to enhancing the performance and reliability of every product in its line. At the same time, Cray Research is committed to pursuing new technologies, manufacturing techniques, and materials that reduce the cost of that performance and provide better value to the supercomputer customer.

The CRAY X-MP series of computer systems is evidence of that parallel commitment. First introduced in 1982, the series has been enhanced for the second time in just over four years. Cray Research is introducing two new CRAY X-MP models as well.

Faster clock

In 1982, the CRAY X-MP series featured a 9.5-nanosecond (nsec) CPU clock cycle time, and performance three to five times that of the original CRAY-1 computer system. At the time, the CRAY X-MP was the most powerful general-purpose computer system available.

The enhanced CRAY X-MP series now offers even faster processing speeds with the introduction of an 8.5-nsec clock. The faster clock period increases scalar, vector, and memory speeds. That means an 11.7 percent performance improvement over previous CRAY X-MP systems and, on the largest CRAY X-MP models, performance 6 to 12 times that of the original CRAY-1 computer system.

More value for users

CRAY X-MP systems have set a standard for performance and price for more than four years. Thanks to lower manufacturing costs for CRAY X-MP one- and two-processor systems, Cray Research is able to pass on the savings to customers. Now the enhanced performance of the new CRAY X-MP one- and two- processor systems is available at even lower prices. When performance goes up and price goes down, the result is significantly better value.

New models

When the CRAY X-MP series was introduced in 1982, only two dual-processor models were available. The series was expanded over time to include entry-level single-processor models and powerful four-processor systems with memory size options to fit the needs of a growing diversity of users. Cray Research is now happy to announce enhanced two-

and four-processor models with new memory options, bringing the benefits of multiprocessing to a broader range of scientific and engineering users.

The new CRAY X-MP/44 computer system combines four processors with four million words of ECL bipolar memory. Maximum memory bandwidth is 16 times that of the CRAY-1 computer system. Central memory is arranged in 32 interleaved banks and has a bank cycle time of 34 nsec. As with all four-processor models, the CRAY X-MP/44 mainframe is a 12-column 270° arc chassis with the same electrical requirements as the CRAY-1 computer system.

The new CRAY X-MP/22 computer system features two processors sharing two million words of MOS central memory arranged in 16 interleaved banks. Maximum memory bandwidth of the CRAY X-MP/2 systems is four times that of the CRAY-1 and memory bank cycle time is 68 nsec. All CRAY X-MP/2 models consist of eight vertical columns arranged in a 180° arc.

Now totaling 11 models, the CRAY X-MP series offers users the widest range of supercomputer configuration options available. From a single-processor system with one million words of central memory to a top-end system with four processors and 16 million words of memory, the CRAY X-MP series provides an alternative to meet the needs of even the most demanding supercomputer user.

Balanced performance

From the outset, CRAY X-MP system design has been carefully balanced to deliver optimum overall performance. Fast long and short vector processing is balanced with high-speed scalar processing and supported by powerful input/output capabilities.

The flexible CRAY X-MP multiprocessor configurations allow users to employ multiprogramming, multiprocessing, and multitasking techniques. The multiple-processor architecture can be used to process many different jobs simultaneously for greater system throughput. It can also be used to apply two or more processors to a single job for better program turnaround time.

Each CRAY X-MP CPU offers gather/scatter and compressed index vector instructions. These instructions allow for the vectorized processing of randomly organized data, which previously was performed by scalar processing.

The enhanced CRAY X-MP systems support the same range of I/O technology as previous systems. The I/O Subsystem (IOS) acts as a data distribution point for the mainframe, handling I/O for a variety of front-end computer systems and peripherals. It includes two, three, or four interconnected I/O processors, each with its own local memory and a common buffer memory. Cray Research's DD-49 disk drive offers 1200-Mbyte capacity and a transfer rate of 9.6 Mbyte/sec at the user job level. When combined with the data handling and buffering capability of the IOS, the DD-49 disk drive provides unsurpassed I/O performance to complement the power of the enhanced CRAY X-MP mainframe.

The optional SSD solid-state storage device provides up to 128 million words of very fast random-access secondary MOS memory. The SSD is recognized by the mainframe as conventional disk storage. It offers significant potential for performance improvement on I/O bound applications, and thus allows users to solve problems that would be impractical with traditional disk I/O.

Fiber optic link

CRAY X-MP systems are designed to fit easily within a user's existing computer environment. Hardware interfaces enable easy connection between the IOS and a wide variety of front-end computers and workstations. CRAY X-MP systems support networking and Cray software interface support provides a logical connection to a multitude of other vendors' systems.

Now a new fiber optic link further enhances the possibilities available to user computing facilities. The new 3-Mbyte/sec link allows a front-end connection to a CRAY X-MP system to span up to one kilometer (.621 miles) with complete electrical separation from the Cray system.

Software continuity

All software written for the CRAY X-MP series runs on the enhanced hardware. In fact, software developed for CRAY-1 systems runs on all models of the CRAY X-MP series, protecting users' software investments. That portability extends to the two Cray proprietary operating systems, COS and UNICOS. UNICOS, based on AT&T's UNIX System V, is available either as a standalone operating system or as a guest operating system running concurrently with COS. Cray Research also offers automatic-vectorizing ANSI 78 FORTRAN compilers, extensive FORTRAN and scientific library routines, program and dataset management utilities, debugging aids, C and Pascal compilers, a powerful Cray assembler (CAL), and a wealth of third-party and public-domain application codes.

The operating systems, Cray FORTRAN compilers, and library programs allow users to take advantage of the vectorizing, multiprocessing, and multitasking features of the CRAY X-MP systems. Major application codes are offered for the CRAY X-MP series in fields such as computational fluid dynamics, structural analysis, mechanical engineering, nuclear safety, circuit design, seismic processing, image processing, molecular modeling, and artificial intelligence. This abundance of available software gives practical value to the power of the CRAY X-MP series.

The commitment continues

"Once again Cray Research has demonstrated its commitment to product advancement and performance leadership in supercomputers," said John Rollwagen, chairman of Cray Research. "We also continue to lead the industry in supercomputer value. With these performance improvements and price reductions, our customers realize a fourfold improvement in the price/performance of CRAY X-MP systems over the original CRAY-1 system introduced in 1976. We are pleased to make this ongoing commitment to our customers." □

An overview of reservoir simulation

improving petroleum engineering decisions through advanced computer hardware and software technology

James C. Erdle, J.S. Nolen & Associates, Inc., Houston, Texas

Reservoir simulation is a computational tool reservoir engineers use to model the physical processes involved in oil and gas recovery. This article describes how advanced computer hardware and software technology are helping the petroleum industry to:

- ☐ Increase the ability of reservoir simulators to handle problems of more realistic size and complexity
- ☐ Improve the efficiency, accuracy, and certainty with which reservoir management problems can be solved using reservoir simulation
- ☐ Reduce the cost, and thus increase the value, of reservoir simulation
- ☐ Expand the availability and usability of reservoir simulation technology to non-specialist management and engineering personnel working in the operations offices of oil companies of all sizes

The seven steps of reservoir simulation

The process of reservoir simulation can be broken down into seven major steps as shown in the table on the opposite page. The first four steps deal with simulator selection and/or development. The last three steps involve applying the simulator to create a representative model of the petroleum reservoir under consideration, thereby improving reservoir management decisions.

Hydrocarbon recovery processes

Petroleum reservoirs contain hydrocarbon mixtures in various physical states, or phases, and are typically accompanied by water (usually saline). The hydrocarbon mixtures consist of naturally occurring hydrocarbon molecules referred to as *components* of the overall hydrocarbon mixture. Impurities such as nitrogen, carbon dioxide, and the extremely toxic hydrogen sulfide can also be present. The extent to which each component exists in the three possible states, or phases (gaseous, liquid, or solid), at any point in the reservoir, in the wellbore plumbing system, and in the surface facilities, depends on the pressure and temperature to which the specific fluid system under consideration is exposed. The physics of phase behavior is described mathematically within a reservoir simulator to varying degrees of completeness to model the fluid recovery processes with sufficient accuracy.

Petroleum reservoirs are accumulations (usually underground) of hydrocarbons, water, and impurities within the pore spaces of rocks (usually sedimentary). A common misconception is the idea of underground tanks containing oil and gas. Such tanks only exist in the case of the strategic petroleum reserves, which are typically contained in caverns created by the solution mining of salt domes. Two major characteristics must be exhibited by underground rock formations to permit the formation of a hydrocarbon accumulation or reservoir:

- ☐ Porous and permeable rock must exist to facilitate both the migration and accumulation of fluids
- ☐ A trapping mechanism must be present to permit the initial accumulation of fluids, then the subsequent loss or leakage of these fluids to other horizons

The shape and physical characteristics of these accumulations, such as porosity, permeability, and compressibility

are highly non-uniform, necessitating an approach to fluid recovery simulation that can accommodate heterogeneous reservoir descriptions. The modern numerical, finite-difference-based reservoir simulators are especially useful for this purpose. Earlier approaches to reservoir simulation, including the material balance approach (zero-dimensional) and the Buckley-Leverett immiscible displacement approach (one-dimensional) suffer from the inability to model two- and three-dimensional variations in the physical characteristics of both the reservoir rock and fluid. This is not to say that full three-dimensional modeling capability is always desirable or even necessary. The best approach to reservoir simulation is to use the simplest simulator that provides sufficiently accurate results.

The fluid recovery mechanisms at work during the unassisted (primary) production phase of a petroleum reservoir's life cycle are fluid expansion, fluid displacement, gravity drainage, and capillary imbibition. Fluid expansion occurs as fluid is withdrawn from the reservoir, causing the pressure exerted by the fluids on the rock pore spaces to decline. This holds true if the pore spaces do not decrease in volume with fluid withdrawal, a phenomenon which, in a few rare cases, has resulted in surface subsidence. Fluid displacement occurs when water displaces oil or gas from underneath or around the hydrocarbon accumulation, and when gas displaces oil or water from a gas cap that is present above the liquid phases initially or that develops during the process of fluid recovery. Gravity drainage occurs as a result of the density differences between water and the gaseous and liquid hydrocarbon phases. Finally, capillary imbibition causes the displacement of hydrocarbon phases from the pore spaces of the rock as a result of the natural tendency of water to be imbibed preferentially into those same pore spaces, assuming the rock wetting phase is water and not hydrocarbon, as is usually the case.

Petroleum engineers have devised many hydrocarbon recovery aids that vary in process complexity. Such aids range from water injection, which relies on displacement and pressure maintenance, to thermal recovery methods. The latter rely on viscosity reduction via heating, distillation of intermediate hydrocarbon components from the liquid phase to the more mobile gaseous phase, and cracking of the oil phase with associated distillation effects. Coats coined the term *oil mobilization* to describe an additional hydrocarbon recovery mechanism associated with both composition-related (phase behavior) and manmade recovery improvement processes related to enhanced recovery. A wide range of chemical injection-based enhanced recovery methods can also be associated with the oil mobilization recovery mechanism. These include polymer flooding, surfactant flooding, and caustic (alkaline) flooding. Carbon dioxide and nitrogen injection methods may benefit from oil mobilization effects if properly designed.

Mathematical representation of the physical processes

All reservoir simulators evolve from second order, nonlinear partial differential equations that result when conservation of momentum, mass, and heat (if included) equations and transport law equations are developed and combined for an elemental volume of reservoir rock. The more complex the process to be simulated, the more equations that are needed to produce the final set of partial differential equations for each elemental volume comprising the reservoir system. A fundamental choice to be made in selecting or developing a reservoir simulator is whether to model the movement of each component (or lumped pseudo-components) in the overall fluid system, or to model only the movement of the major fluid states or phases. The latter approach is known as *black-oil* simulation, and typically accounts for gas dissolving in the oil phase only, that is, no gas dissolving in water or oil/water phase interchange is allowed. All early simulators incorporated the black-oil component modeling approach.

The individual component modeling approach known as *compositional* reservoir simulation accounts for interphase mass transfer between all fluid components (or pseudo-components). Many real-world petroleum reservoirs contain fluids that exhibit significant interphase mass transfer, thus making black-oil simulators inaccurate for these

The seven steps of reservoir simulation

Steps 1-4: develop or select a simulator

1. Describe the physical processes to be simulated.

2. Develop a mathematical representation of the physical processes.

3. Develop a numerical solution to the mathematical representation (this step has evolved from the initial use of scaled-down physical laboratory models and the subsequent use of analog computer models based on analytical solutions to the fundamental mathematical representation).

4. Develop a digital computer program that incorporates the numerical solution.

Steps 5-7: use the selected/developed reservoir simulator to perform reservoir simulation to improve reservoir management decisions.

5. Build a model of the petroleum reservoir under consideration (by supplying a physical description of the reservoir to the reservoir simulator selected/developed above).

6. Fine-tune the physical description of the reservoir supplied to the simulator by history-matching past reservoir performance (driving the simulator with the measured fluid recovery schedule and comparing measured reservoir and individual well performance variables against those calculated by the simulator).

7. Predict the engineering and economic impact of alternative reservoir management schemes, such as well locations, production rates, and secondary and enhanced recovery projects.

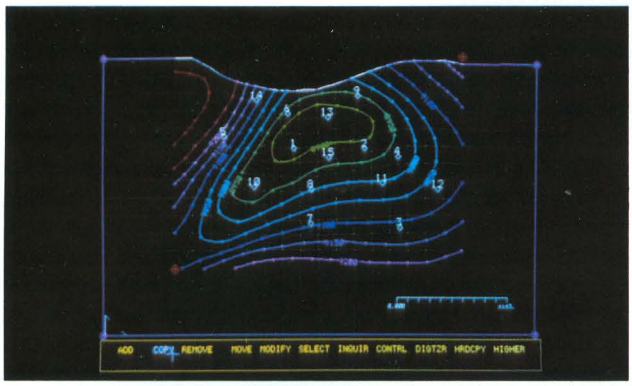

Figure 1. A top-of-structure contour map.

simulation situations. Classic examples of situations requiring compositional simulators include:

- Primary production from volatile oil or gas condensate reservoirs
- Dry or enriched gas injection into black-oil reservoirs to vaporize lighter liquid hydrocarbon components into the more mobile gaseous phase or achieve miscibility between the reservoir fluids and the injected gas
- Carbon dioxide injection to reduce oil viscosity and increase fluid expansion by swelling the oil phase

In addition, most of the enhanced recovery reservoir simulators require the compositional approach to modeling phase behavior. The number of partial differential equations that must be solved at each elemental volume or cell can be substantially greater in compositional simulation than in black-oil simulation. Fortunately, the development and continued refinement of supercomputers, such as CRAY X-MP computer systems, has made field-wide compositional reservoir simulation a reality in many situations.

Numerical solution of the mathematical representation

The partial differential equations describing the recovery of fluids from a petroleum reservoir are solved by finite difference techniques in most modern reservoir simulators used for field-wide applications. The finite difference approach is conducive to developing an intuitive "feel" for the ways in which the simulator deals with a real-world petroleum reservoir problem, even if one has no understanding of the mathematical principles involved. This is because finite difference-based simulators model fluid movement between imaginary "cells" in a reservoir and their immediate neighbors (two, four, or six neighbors for one-, two-, or three-dimensional orthogonal grid systems). Part of the process of using a reservoir simulator involves "gridding the reservoir," or overlaying a one-, two-, or three-dimensional orthogonal or curvilinear grid network throughout the portion of the reservoir in which fluid movement modeling is desired (Figures 1 and 2).

Three main types of finite difference formulations are currently in use. Each has a certain range of application, depending on the size and complexity of the problem. Large-scale reservoir simulation studies using upwards of 15,000 to 20,000 grid blocks or cells are becoming more common and usually are run using the IMPES (Implicit Pressure-Explicit Time) formulation. Studies involving large pressure and fluid saturation gradients, such as "coning" studies using an R-Z coordinate system geometry, require the fully implicit formulation, the most robust formulation available. A third formulation, called the sequential formulation, is used for many intermediate applications requiring a blend of robustness and large problem-handling capability.

Regardless of the formulation chosen to solve a particular problem, a large sparse and, unfortunately, nonsymmetric system of linearized equations must be solved via matrix inversion techniques to move the reservoir simulator from one point in time to the next. The use of corner point geometry descriptions (to model pinchouts and sloping faults, as in Figure 3), dual porosity systems, and multiple layer well completions has led to the use of iterative matrix inversion methods with conjugate gradient residual constraints and preconditioning to produce nearly symmetric, partially vectorizable systems of linear equations (coefficient matrices).

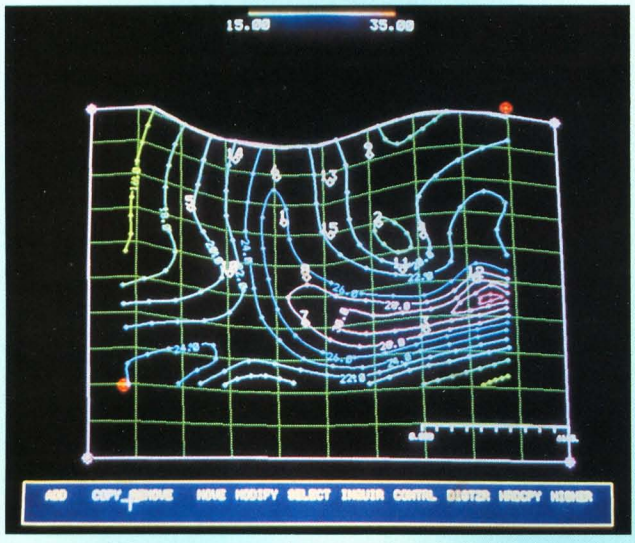

Figure 2. The same map shown in Figure 1 with superimposed grid system.

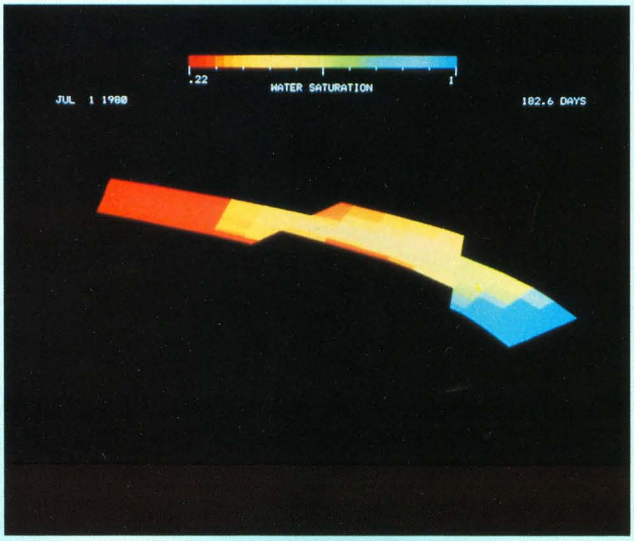

Figure 3. Graphic showing use of corner point geometry to model a sloping fault.

The simulator

Three significant new advances in the computer implementation of reservoir simulator equations have led to an increased use of reservoir simulation technology. These advances are:

- ☐ Vectorized code to take advantage of the vector processing capability of advanced supercomputers
- ☐ Modularized simulator components that allow quick comparisons of alternative model descriptions and/or recovery methods
- ☐ Workstation software that permits input data preparation and output data analysis in menu-driven format, with excellent color graphics display capability, even in a microcomputer environment

J.S. Nolen & Associates pioneered the development of vectorized coding for reservoir simulators with the VIP (Vectorized, Implicit Program) black-oil simulator introduced in 1980. J.S. Nolen has also recently introduced a vectorized and modularized family of reservoir simulators called the VIP-Executive Series. This new line of simulators was developed in response to the petroleum industry's demands for simulators that are easy to use and maintain. With VIP-Executive, a user can alternate between a black-oil approach and a fully compositional approach during model construction sensitivity studies to quickly determine a sufficiently accurate method. In addition, dual porosity behavior can be modeled in either the black-oil or compositional modes with very little additional input.

Another important development in reservoir simulation is menu-driven microcomputer-based data input and output processing software that reduces the need to learn complex keyword commands and to plot hardcopy of input and output data for checking or analysis. The petroleum engineer or consultant who has never had access to reservoir simulation technology before can now use this technique by:

- ☐ Leasing or purchasing the workstation input/output data processing software for a PC
- ☐ Entering a reservoir data set following the menu-driven prompts displayed in plain English by the workstation input processing software
- ☐ Transmitting the input data set to a service bureau (either electronically via a modem or mechanically by sending a floppy disk or tape in the mail) to submit a run on the desired simulator
- ☐ Downloading the simulator output data to a PC for display and analysis using the workstation output processing software
- ☐ Conferring with reservoir simulation experts about unresolvable questions of procedure, including input data preparation, simulator selection, and history-matching

Building a reservoir model

Preparation of reservoir simulator input data has long been considered one of the most time-consuming and tedious phases of the reservoir simulation process. The advent of reservoir simulation input/output processing software for workstations and mainframe terminals has dramatically reduced the time and manpower required for the data input phase of a simulation study. Constraints and upper/lower boundary range values for input parameters also reduce the risk of wasted runs.

History-matching

Reservoir simulation traditionally has included the step of comparing state variables calculated by the simulator with field observations (measurements) of those same state variables at equivalent points in time. Typically, the most commonly "known" state variables are well pressures and rates, as opposed to saturation and pressure distributions in the interwell portions of the reservoir. The reservoir simulator workstation output data processor makes comparison of measured and computed well pressures and fluid production ratios fast and accurate.

Prediction of future reservoir behavior

Modularized reservoir simulator packages coupled with a reservoir simulator workstation input/output data processor, make comparison of alternative recovery schemes a much faster procedure than previously thought possible. Many reservoir and well management decisions are improved through the use of reservoir simulation. Examples include:

- ☐ Development (infill) well drilling patterns and spacing
- ☐ Well completion intervals and zones
- ☐ Field-wide optimization of well stimulation treatment type and sizes
- ☐ Recovery assistance method feasibility studies and design optimization
- ☐ Production allocation in multiple interest reservoir ownership situations
- ☐ Production/injection rate limitations to minimize completion damage and/or reservoir damage
- ☐ Many other engineering and management problem-solving activities

Conclusions

Reservoir simulation is a complex task that is being made more economical, reliable, and accessible to the petroleum industry through advances in computer hardware and software technology. The development of user-friendly software interfaces that prompt user actions in an almost "expert-like" manner are pushing state-of-the-art technology to the levels where it is needed most: the operating offices of oil companies of all sizes. ☐

About the author

James C. Erdle is a well systems analysis consultant and manager of training at J.S. Nolen & Associates. He received his BS degree in petroleum and natural gas engineering from Penn State University in 1971 and his Ph.D. degree in the same discipline from Penn State in 1974. Prior to joining J.S. Nolen & Associates in 1985, Erdle was rocky mountain division operations manager for Flopetrol Johnston/Schlumberger in Denver, Colorado.

Vectorization of a reservoir simulator

Leslie N. Smith, Ganesan Subramanian, John A. Trangenstein, Donald W. Peaceman, and William J. Silliman
Exxon Production Research Company, Houston, Texas

Reservoir simulation is a sophisticated mathematical science that plays a key role in the development of oil and gas resources. Present-generation vectorized simulators are more than an order of magnitude more efficient than previous-generation scalar simulators. This increase in efficiency is possible because the key components of a simulator are amenable to vectorization.

On a worldwide basis, reservoir simulators are used by more than 1000 engineers to make decisions involving investments worth billions of dollars. Simulators have gained widespread use because they can solve problems that are intractable by any other means. No other method exists to describe the flow of oil, gas, and water in multiple dimensions in a real reservoir.

Modern day simulators depend on effective use of state-of-the-art computing equipment such as vector computers. Such computer systems are needed because the numerical solution of the nonlinear algebraic and partial differential equations describing flow in porous media often lead to systems of equations involving tens of thousands of unknowns. Solution of such systems is practical only on vector machines.

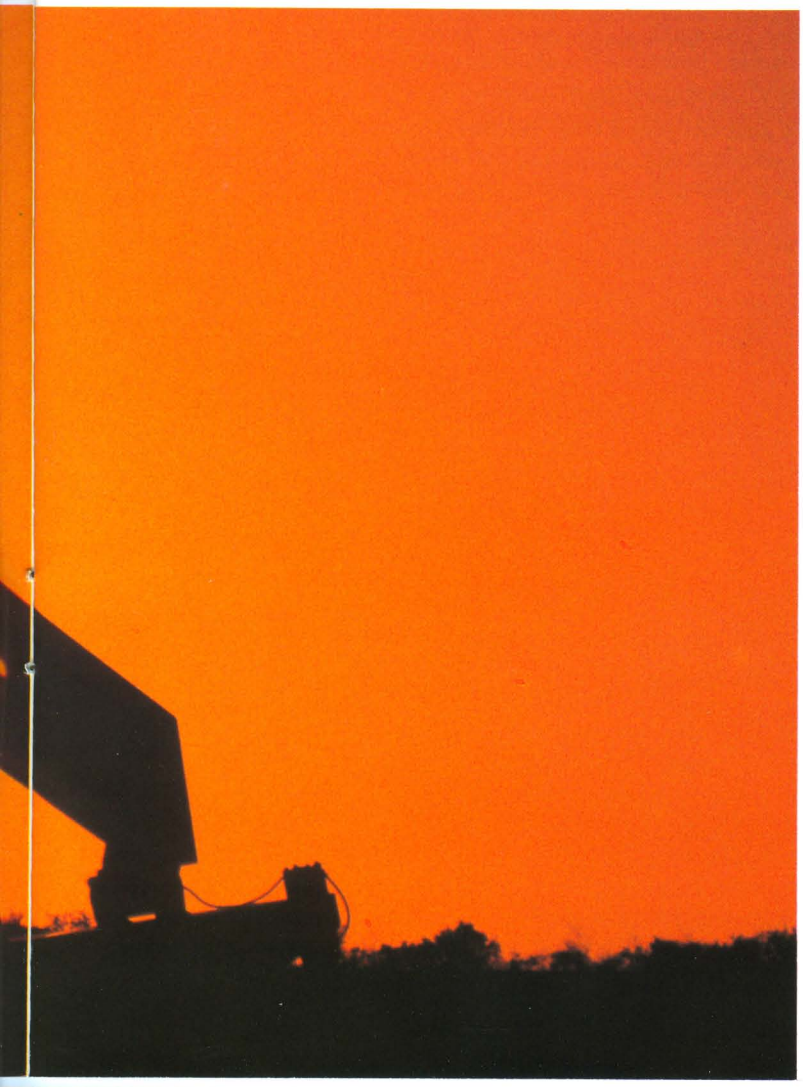

The authors have designed and developed a reservoir simulation program[1], MARS, that is unique because vectorization was a major objective from the beginning of its development process. As a result, the program operates very efficiently on a vector computer such as a Cray computer system. In fact, it operates more than an order of magnitude faster than the previous generation of simulators running on traditional scalar equipment. This article describes the basic structure of the MARS reservoir simulator and some of the techniques used to achieve a high degree of vectorization.

Design of a reservoir simulator

When using a simulator, a grid is superimposed on the reservoir, dividing it into a large number of cells, or gridblocks. Ideally, cells should be small enough so that reservoir and fluid properties can be assumed to be constant within each cell. Models may be one-, two-, or three-dimensional as indicated in Figure 1. With current vector technology, three-dimensional models can reasonably have tens of thousands of cells.

A simulator predicts how the properties and contents of each cell evolve over the life of the field. This is done by taking a series of timesteps: repeatedly taking the state of the reservoir at some given time T_n and predicting the change over a period of time ΔT. Within each so-called timestep several tasks must be performed. The most important of these are the solution of the partial differential equations (PDEs) governing fluid flow, the evaluation of fluid properties, and the calculation of well rates (for example, how much oil or gas is produced).

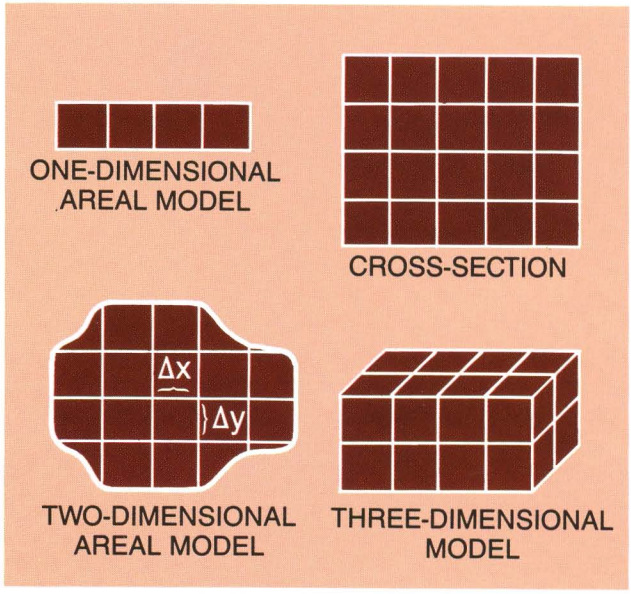

Figure 1. One- and two-dimensional areal, two-dimensional cross-sectional, and three-dimensional reservoir simulation models.

Fluid flow equations

In the most common type of simulator, the black-oil simulator, the reservoir is assumed to be made up of three phases (hydrocarbon liquid, hydrocarbon vapor, and water) with limited mass transfer between the phases. This relatively simple model is sufficient for modeling about 80 percent of all petroleum reservoirs when oil is recovered by fluid expansion or by displacement with water and gas.

The partial differential equations describing flow in porous media derive from mass conservation relations and Darcy's Law for flow in porous media.[2,3] The independent variables for each cell are the pressure (p) and the volume fractions (called saturations, $\underline{s}$) of the gas and water phases. The equations are highly nonlinear and are mixtures of parabolic and hyperbolic partial differential equations.

A "pressure" equation can be generated by combining the conservation equations and using Darcy's Law. This gives:

$$c(p,\underline{s})\frac{\partial p}{\partial t} = Q_p(p,\underline{s}) - \sum_{m=1}^{3} V_{fm}(p,\underline{s})\,\nabla\cdot\left[\sum_{r=1}^{3} X_{rm}(p,\underline{s})\,\nabla p\right] \quad (1)$$

where Q_p is a source/sink term and $c(p, \underline{s})$, $V_{fm}(p, \underline{s})$, and $X_{rm}(p, \underline{s})$ are fluid properties (compressibility, partial molar volumes, and component densities, respectively). This equation is parabolic, although it behaves nearly elliptically because compressibility is normally very small.

The mass balance equations form the basis for the "saturation" equations. They are generally hyperbolic in nature and take the form:

$$\frac{\partial \underline{s}}{\partial t} = \underline{Q}_s(p,\underline{s}) - \sum_{m=1}^{3} \underline{V}_m(p,\underline{s})\,\nabla\cdot\left[u(p)\sum_{r=1}^{3} \underline{F}_r(p,\underline{s})\right] \quad (2)$$

Here $\underline{V}_m$ is a fluid property (phase partial molar volume), $\underline{Q}_s$ is the source/sink term, and $u(p)$ is the total fluid velocity.

These PDEs can be solved in many ways. The MARS code uses the sequential semi-implicit procedure of Spillette, Hillestad, and Stone.[4] In this case two sets of discrete equations are solved at each timestep. The first step is the solution of the pressure equation (1) with the saturation dependence lagged. Because the equation set turns out to be parabolic, it is relatively difficult to solve. However, it is also relatively small, consisting of a single equation per gridblock. These pressures are used to compute the velocity distribution over all phases ($u(p)$), which is then used in the saturation equations (2) to compute the saturations at the end of the timesteps. The saturation equations are essentially hyperbolic and the solution of the discretized set is relatively easy. Additional details of the procedure are given elsewhere[5] and are an extension of other investigators' efforts.[4,6]

In MARS, and in simulators throughout most of the industry, both sets of equations are spatially discretized with standard block-centered finite differences for second derivatives combined with some form of upstream weighting for the first derivatives. Backward differences are normally used in time. Nonlinearities are removed by Taylor series extrapolations about the current values. Once the discrete set of equations is generated, it is solved by any of a wide variety of direct methods (for example, sparse matrix[7] and banded Gaussian elimination[8] algorithms) or iterative methods (for example, line SOR with correction,[9,10,11] nested factorization preconditioning with conjugate gradient iteration,[12] and strongly implicit procedure.[13,14]) In a black-oil simulator, the flow calculations typically consume 50 to 70 percent of the total computation time (Figure 2). As a result, they represent an important component in vectorization.

In enhanced oil recovery (EOR) simulation, thermodynamic and chemical relationships play a more critical role, making the simple black-oil model inadequate. EOR simulation must account for individual species, such as methane, ethane, carbon dioxide, and others, and a mass balance for each species must be substituted for equations (2). In this case, simulation cost is usually dominated by the property evaluation.

Property evaluation

In each timestep a large number of thermodynamic properties, such as densities and viscosities, and their derivatives with respect to primary variables must be calculated. In

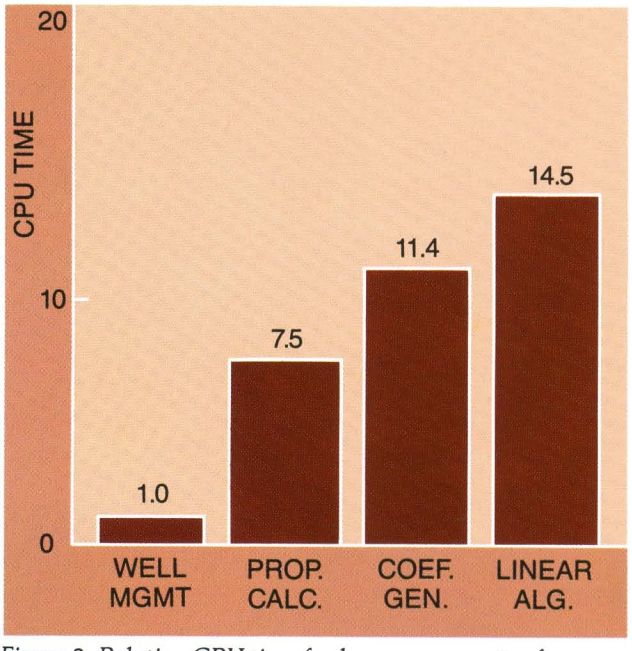

Figure 2. Relative CPU time for key components of a reservoir simulator. Times are for unvectorized code and normalized by the time for the well management routine.

the black-oil case, the thermodynamics is relatively simple and these properties are generally a function of pressure alone. This means that the properties can be expressed in terms of tabular functions, and property evaluations can be accomplished via table look-up. As demonstrated (Figure 2), the property evaluation can consume as much as 30 percent of the overall CPU time.

In the more complicated EOR case, an equation of state such as the Peng Robinson Equation[15] is used to model the thermodynamics. This requires the solution of a highly nonlinear set of equations for each cell in each timestep for the phase compositions. The equations are solved using either a successive substitution or an unconstrained minimization algorithm.[16] If not vectorized, this property evaluation step completely dominates the solution time (88 percent of total CPU time in one test case).

Well management
Simulation models primarily the behavior of the reservoir alone. However, that behavior is influenced strongly by the production strategies and surface facilities through the wells, which might be viewed as variable boundary conditions. As a result, an additional complexity is introduced because the fluid flow equations are not solved in isolation from some model of surface facilities and activities. The logic involved in meeting these specifications is extremely complex and not at all amenable to vectorization. Fortunately, this part of the code does not usually dominate the computations.

Vectorization issues

In developing a simulator, we carefully considered each component as a possible candidate for vectorization. The issues were different in each case, as were the results.

Fluid flow equations
The computational work for the flow equations can be divided into two parts: the calculation of the discretized equations for pressure and saturation (coefficient generation) and the solution of the resulting equation sets (linear algebra).

The coefficient generation routines are the most straightforward to vectorize. Two issues are involved. First, some code is dependent on the type of cell; for example, the user will frequently "key-out" or exclude certain cells from a model. It is best to do extra calculations, if necessary, to preserve vectorization. In particular, we compute coefficients for all cells, including those keyed-out, in a vectorized mode and then zero the appropriate coefficients in a post-processing step. A second issue involved code that is dependent on flow direction, for example, upstream flow. Again, efficiency is improved by performing extra calculations. We do calculations for both upstream and downstream directions and then multiply the contribution from the downstream block by zero. By using this basic philosophy — do extra computations to preserve vectorization — efficiency increased fivefold (Figure 3).

In analyzing the linear algebra routines, traditional measures of efficiency were entirely inadequate. For example, band algorithms were far more efficient than sparse algorithms even though their operation count (number of multiplies, divides, adds, and subtracts required for a solution) was much higher. The reason is that band algorithms vectorize well while the sparse ones do not. With iterative techniques, the key issue appears to be the avoidance of recursion, which inhibits vectorization. Techniques that work best seem to have essentially no recursion (for example, Line Successive Over Relaxation — LSOR — with correction) or have minimized it to a major degree.[17] As with the coefficient generation algorithms, a substantial increase in efficiency was attained (Figure 3).

Property evaluation
In black-oil simulation, two complicating factors occur in property evaluations. First, a large number of table look-ups are required, and these look-ups are computationally expensive when the fluid property tables have inconstant increments. Second, blocks that are saturated (blocks where free gas is present) need to be treated differently than blocks that are undersaturated (blocks where no free gas is present). The first issue was addressed by use of a Cray Assembly Language (CAL) routine. This increased efficiency by over 50 percent. In addressing the second problem, we found significant payoff through the use of gather/scatter techniques. (Note: the authors conducted the

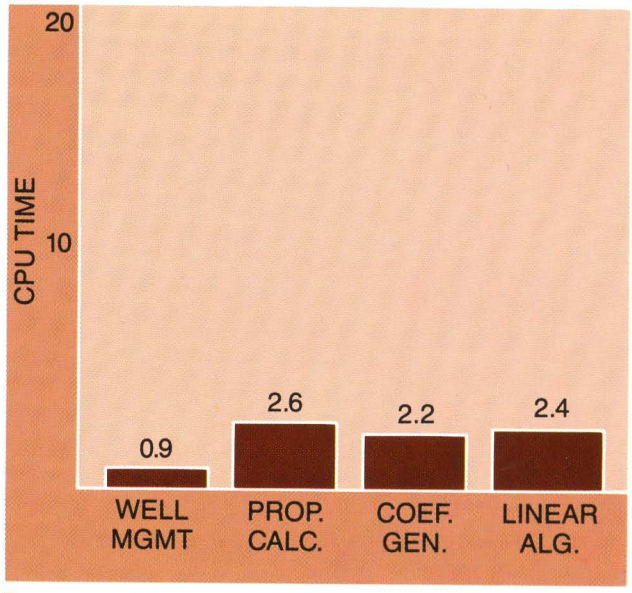

Figure 3. Relative CPU time for key components of a reservoir simulator. Times are for vectorized code. Normalization is the same as in Figure 2.

work reported in this article on a CRAY-1/S system that does not include hardware gather/scatter. Hardware gather/scatter is now included in CRAY X-MP systems -- Ed.) That is, saturated gridblocks are gathered into one group and processed alone. The remaining undersaturated gridblocks are then gathered and processed separately. Finally, the computed properties are scattered back to the original gridblock order (where saturated and undersaturated blocks are mixed). Overall, use of gather/scatters led to a threefold reduction in CPU time.

In the EOR case, the simulator must solve a nonlinear set of equations for each block and then compute a large number of derivatives based on the solutions. Further, each gridblock requires a large amount of storage to execute the property calculations. Straightforward vectorization techniques would require the simulator to store all the information needed for all gridblocks and then solve nonlinear equations for each gridblock in a vectorized loop. This approach would require significantly more storage than that available on the four-million-word CRAY-1/S computer system used by the authors. However, this storage limitation can be addressed by a technique known as "strip-mining." Using this technique, the simulator stores all the information for 64 gridblocks at a time (64 is the optimal vector length for the Cray system) and then solves for all of them simultaneously. It then processes the next 64 gridblocks, and so on. This "strip-mining" technique provides optimal vector performance while minimizing the amount of storage required.

In terms of actually solving the nonlinear equations, we found it best to take a fixed number of iterations for each grid block in a "strip" and then check the entire strip for convergence. We then process further any unconverged blocks in the strip. In general, these vectorization techniques have led to a sixfold to sevenfold increase in efficiency for EOR models. Figure 4 shows the results of strip-mining on 1000-gridblock and 3000-gridblock EOR models.

Well management

We found little to vectorize in the well management code. The basic difficulties stem from the extensive logic involved, the number of special well models — each of which requires special treatment — and the extensive sorting of data required. All of this means that there is little hope of vectorization on current machines with current software. Although we did not put a large effort into vectorization — we were just careful as we went along not to destroy vectorization whenever possible — we did nonetheless see a 10 percent increase in speed (Figure 3).

Remaining challenges

A number of challenges remain. With the vectorization that we have done, we find that unvectorized code is becoming a dominant cost factor. If we want to achieve another substantial increase in speed, we must address this issue.

Another challenge is reflected in the ratio of wall clock time to CPU time. Figure 5 shows the wall clock/CPU ratio for a range of typical simulation models. Not all of the runs were made on an unloaded Cray system, so system wait time was a factor in some cases. While the largest models may only require a few hours of CPU time, they may tie up the computer for ten times as long. This time is spent

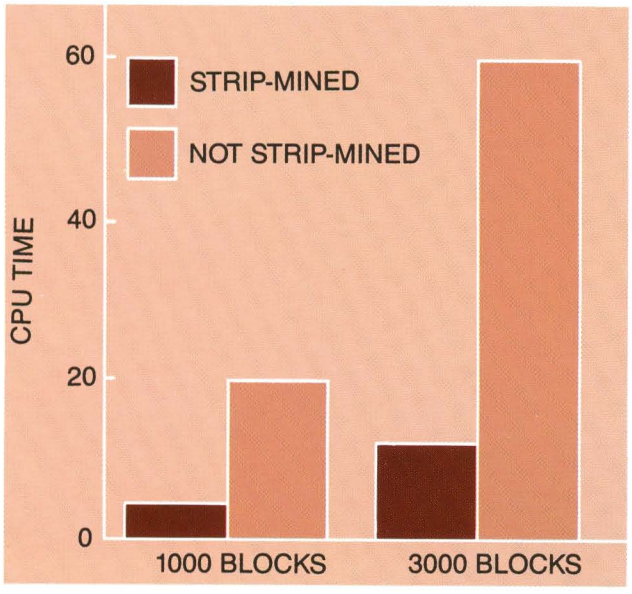

Figure 4. Reduction of execution time through use of strip-mining vectorized technique.

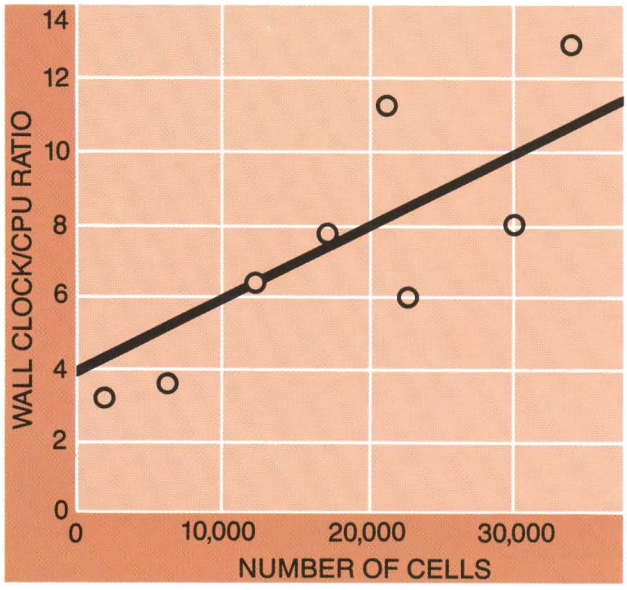

Figure 5. Ratio of the wall clock time to CPU time for a range of simulation codes.

transferring data from disk to memory and back. To more efficiently use vector computers, this factor must be reduced.

As hardware improves, researchers seeking to optimize these codes will need to know how to take advantage of the new technologies. An obvious example is the newer parallel processing machines. Experience dictates that the linear algebra area will be very sensitive to hardware changes and very amenable to exploiting hardware advances.

Finally, the trend in simulation is towards larger and more complex models, particularly in the EOR area. State-of-the-art techniques are not yet adequate to solve all these problems in a reasonable time, and intensive efforts are underway to improve our abilities.

Conclusion

Effective reservoir simulation depends on the use of vector computers, but effective use of such computers requires careful attention to algorithms and computer science techniques.

Conventional methods of evaluating algorithms are frequently misleading when applied to code running on vector computers. It is usually much more important to evaluate the vectorization potential of an algorithm than it is to examine the amount of work it does.

It is possible to develop a reservoir simulation program that operates very efficiently on a Cray computer system. It turns out that most parts of the code are adaptable to vectorization. In our work we achieved a threefold to fivefold increase in simulator efficiency when compared to an unvectorized version of the code. This increase translates into a factor of 13 to 18 when compared to a comparable simulator running on state-of-the-art scalar computers. Overall, vector computing has increased the upper bound for simulation models by more than an order of magnitude.

With vector machines and highly vectorized code, it is possible to adequately model many of the conventional recovery processes important to the oil industry. However, further improvements are needed to successfully model the more difficult EOR situations or situations where the reservoir matrix has structural complexities beyond the scope of present-day simulators. □

Acknowledgment

Modified and reprinted with permission from **Supercomputers: Algorithms, Architectures, and Scientific Computation**, *edited by F. A. Madsen and T. Tajima.*

About the authors

The authors were involved in the development of reservoir simulators at Exxon Production Research Company at the time this article was prepared. Leslie N. Smith has a Ph.D. in chemistry from the University of Illinois at Urbana-Champaign. John A. Trangenstein has a Ph.D. in mathematics from Cornell University and Donald W. Peaceman has a Sc.D. from the Massachusetts Institute of Technology. Both Ganesan Subramanian and William J. Silliman received their Ph.D. degrees in chemical engineering from the University of Minnesota.

References

1. Kendall, R. P., et al., SPE 11483, presented at the Middle East Oil Technical Conference of the Society of Petroleum Engineers, March 1983.

2. Peaceman, D. W., *Fundamentals of Numerical Reservoir Simulation*, Elsevier, 1977.

3. Aziz, K., and A. Settari, *Petroleum Reservoir Simulation*, Applied Science Publishers, 1979.

4. Spilette, A. G., et al., SPE 4542, presented at the SPE 48th Annual Meeting, September 1973.

5. Watts, J. W., SPE 12244, presented at the Seventh SPE Symposium on Reservoir Simulation, November 1983.

6. Acs, G., et al., SPE 10515, presented at the Sixth SPE Symposium on Reservoir Simulation, February 1982.

7. Einsenstat, *Sparse Matrix Proceedings 1978*, SIAM, 1979.

8. Dongarra, J. J., et al., *LINPACK User's Guide*, SIAM, 1979.

9. Watts, J. W., *Society of Petroleum Engineers Journal*, pp. 47-51, March 1971.

10. Nolen, J. S., et al., SPE 7675, presented at the Fifth SPE Symposium on Reservoir Simulation, February 1979.

11. Breitenbach, E. A., et al., SPE 2021, presented at the SPE Symposium on Reservoir Simulation, April 1968.

12. Appleyard, J. R., et al., SPE 12264, presented at the Seventh SPE Symposium on Reservoir Simulation, November 1983.

13. Stone, H. L., *SIAM Journal of Numerical Analysis*, 5, pp. 530-580, 1968.

14. Weinstein, H. G., et al., *Society of Petroleum Engineers Journal*, 10, pp. 99-110, 1970.

15. Peng, D. Y., et al., *Industrial and Engineering Chemistry Fundamentals*, 15, pp. 59-64, 1979.

16. Trangenstein, J.A., SPE 13520, presented at the Eighth SPE Symposium on Reservoir Simulation, February 1985.

17. Wallis, J. R., presented at SEG/SIAM/SPE Conference on Mathematical and Computational Methods in Seismic Exploration and Reservoir Modeling, January 1985.

Contamination of aquifers from sources such as improperly disposed toxic waste, leaking storage tanks, and seepage from polluted streams and ponds has met with increasing public concern. Microbial biodegradation — the decomposition of contaminants by microorganisms — is an important process that can render harmless certain aquifer contaminants such as hydrocarbons. Microbial biodegradation is a natural process that can be accelerated to protect a potable water supply.

Decontamination is physically and chemically complex, involving transport and interaction of hydrocarbons, microbes, and oxygen as well as water movement within the aquifer. Numerical simulation is a necessary predictive tool for understanding and designing microbial decontamination techniques that are physically and economically feasible.

Public concern is understandable considering the health risks that result from ingesting significant amounts of certain substances. For example, U.S. Environmental Protection Agency (EPA) studies indicate that small amounts of gasoline in drinking water result in nausea, dizziness, and headaches, and can cause cancer.[1] A recent EPA study indicates that up to 35 percent of the 800,000 gasoline storage tanks in the U.S. are leaking and that gasoline from perhaps 55 percent of the leaking tanks reaches the water table.[2]

Simulation of hydrocarbon biodegradation in groundwater

Clint N. Dawson and Mary F. Wheeler, Rice University, Houston, Texas
Tung M. Nguyen and Stephen W. Poole, Cray Research, Inc.

Inadequate technology is often the cause of aquifer contamination; waste sites thought to be safe in the past are now known to be leaking. An old creosote pit in Conroe, Texas, is an example of a situation where commonly accepted disposal practices of the 1950s and 1960s have resulted in aquifer contamination.[3] In addition, some ponds and streams have been surreptitiously used as dumping grounds, leaving communities "holding the bag." A stagnant pond northeast of Houston is now a dangerous pool of toxic waste caused by the improper disposal of petrochemicals. No one admits responsibility.

Federal and local governments have responded to the problem with tougher regulations to prevent aquifer contamination. In addition, awards to plaintiffs affected by contaminated drinking water have increased in recent years. Thus, legal and economic incentives for using biodegradation to reduce contamination are significant and are steadily increasing.

In this article we focus our attention on modeling microbial biodegradation. Support for this research has been partially funded by the National Science Foundation (grant DMS85-02320) and Cray Research, Inc.

Description of biodegradation of hydrocarbons

Hydrocarbons such as gasoline are among the most common groundwater contaminants and have been observed to persist in aquifers for many years. However, many aquifers contain microorganisms capable of biodegrading these contaminants. That they do so only at a very slow rate is thought to be due to the small amount of oxygen naturally available in most aquifers. Past microbiological studies have shown that oxygen is required for hydrocarbon biodegradation.[4]

The apparent ability of some microorganisms to consume hydrocarbons in the presence of oxygen has led to several attempts to reclaim contaminated groundwater by the injection of water containing oxygen into the aquifer. Borden and Bedient developed a model describing the flow of oxygen, hydrocarbon, and microbes through the aquifer.[4] Using data from an abandoned toxic waste site near Houston, they have conducted simplified simulation studies in one space variable to test the validity of their model. Standard finite difference techniques were employed. While their ability to model the microbial biodegradation process has met with some success, it has also been hampered by a lack of mathematical tools for simulating in situ biodegradation. In particular, extensive field application requires a numerical method capable of dealing with a non-homogeneous aquifer in multiple dimensions and advection-dominated flow fields.[3,5]

Description of the mathematical model

Borden and Bedient have formulated ordinary differential equations describing the chemical reactions that occur between oxygen, hydrocarbon, and microbes in the biodegradation process.[4] Basically this system of equations expresses the rate at which microbes multiply in the presence of oxygen and hydrocarbon, and the rates at which oxygen and hydrocarbon are consumed. Furthermore, the equations include a rate at which microbes decay in areas where either oxygen or hydrocarbon is missing. The rate at which microbes multiply is essentially exponential. The rates at which oxygen and hydrocarbons are consumed are not exponential but are highly nonlinear. When these reaction equations are combined with the advection-dispersion equations for a solute undergoing linear instantaneous adsorption and with Darcy's Law for an incompressible velocity field, a system of partial differential equations describing the flow of the three single-phase components through the aquifer is obtained.

Darcy's Law and incompressibility combine into an elliptic partial differential equation for pressure and Darcy velocities. This velocity field is created by injecting oxygen in water into the aquifer and producing water containing oxygen, hydrocarbon, and microbes. The injecting and producing are done by wells extending from the surface.

The advection-dispersion-reaction equations describing the flow of the three components through the aquifer are based on the idea of conservation of mass. Advection is controlled by the Darcy velocities. The dispersion terms include the effects of molecular diffusion and velocity-dependent longitudinal and transverse dispersion. These terms describe the "spread" of fronts during diffusion.

Terms representing the injection and production process are also included in the flow equations. Wells are treated as point sources and sinks; that is, they are represented by Dirac delta functions of space. The wells are assumed to inject and produce at a constant rate. Furthermore, we specify an injected concentration of oxygen and produce the same concentrations of oxygen, hydrocarbon, and microbes as are present at the producing wells.

The flow equations for microbes and hydrocarbon are complicated by the inclusion of retardation factors. In the case of microbes, retardation represents, under some simplifying assumptions, the tendency of microbes to grow as colonies attached to the rock formation in the aquifer. This affinity of microorganisms for solid surfaces inhibits the flow of the total population. Some hydrocarbons (tar for example) can also exhibit an affinity for solid surfaces; hence it is necessary to include a retardation factor in the hydrocarbon flow equation.

Description of the numerical method

Simulating the flow of the three components through the aquifer is hampered by several major difficulties. These difficulties are problems in stability and accuracy of numerical methods and problems in the cost and amount of time involved in the computation.

The major difficulty from a numerical analysis point of view is that the flow is advection-dominated. For these flow problems, it is well known that standard finite difference and finite element methods are plagued with numerical instabilities. Moreover, accurate velocities are essential in solving these transport-dominated problems. Hence, a numerical method capable of resolving sharp fronts and calculating higher-order accurate velocities is needed.

From a computational point of view, a major difficulty is that the time scales for reactions and advection-dispersion are quite different; the reactions occur much faster than advection. To solve these nonlinear equations simultaneously with small timesteps would be very costly.

We have avoided such difficulties by solving the problem in a sequential fashion, using a numerical method which is particularly suited to each piece of the computation. First, accurate Darcy velocities and pressure are calculated using a higher-order mixed finite element procedure.[6] A time-splitting algorithm is then employed for treating advection-dispersion reactions. Namely, the reaction terms are separated from the advection-dispersion terms.

Advection-dispersion is handled using a finite element modified method of characteristics. This method, first formulated by Douglas and Russell, combines the time derivatives and the advection term as a directional derivative.[7] In other words, this procedure involves time-stepping along the characteristics, allowing one to use large accurate timesteps. Application of the mixed finite element and the modified method of characteristics to the miscible displacement problem in porous media can be found in references 8 and 9.

Subsequently, the reaction terms become a system of nonlinear ordinary differential equations that can be handled by any stiff ordinary differential equations solver. For purposes of vectorization, we have found that the second order Runge-Kutta explicit method is adequate.

In our time-splitting algorithm, we first perform one advection-dispersion time step, then calculate reactions for many small timesteps. The approximation to the solution at the last timestep is then used as initial data for the next advection-dispersion step.[10]

This time-splitting sequential approach is ideal for parallel computation. The advection-dispersion step can be solved simultaneously on three different processors. Furthermore, the reaction ordinary differential equations can be done in parallel over as many processors as one desires.

Even with these simplifications, the magnitude of the simulation is quite large. For a two-dimensional problem with 61 grid points in each direction, we evaluate and store velocities at 29,040 points. Computing velocities requires solving a linear system of equations with 14,400 unknowns. Handling the reactions alone involves solving 11,163 ordinary differential equations simultaneously for several timesteps. Furthermore, each advection-dispersion step requires solving three linear systems, each with 3721 unknowns. A simulation modeling months or even years requires large amounts of memory and CPU time. Supercomputers such as CRAY X-MP systems make these large-scale simulations possible.

Numerical results

While the governing equations are applicable to a full three-dimensional analysis, we have simplified our assumptions by considering the two-dimensional problem associated with taking an areal view of the aquifer. Gravitational effects can then be ignored.

In our computations, we assumed the aquifer was square, with an injection well in one corner and a production well in the opposite corner. We assumed background concentrations of hydrocarbon and microbes in the aquifer, and injected oxygen in water into the aquifer. We initialized the problem by allowing the oxygen to flow out radially from the injection well with no reactions for 1/100 of one day.

In the cases run so far, we have examined the effects of varying the physical dispersion levels and the retardation factors, both for microbes and hydrocarbons. We have also simulated flow in a nonhomogeneous aquifer.

Test cases were run on a CRAY X-MP/48 computer system. Figures 1-3 represent contour mapping of the concentration levels of oxygen, hydrocarbon, and microbes in the aquifer. The injection well is in the lower left corner and the production well is in the upper right corner of the aquifer. Results were generated on a rectangular, nonuniform 61 x 61 grid. Postprocessing was performed on a VAX 11/750.

Generation of the color contours was a multi-stage process. First, the nodal values for each concentration component were normalized to be in the range zero to one. Bilinear

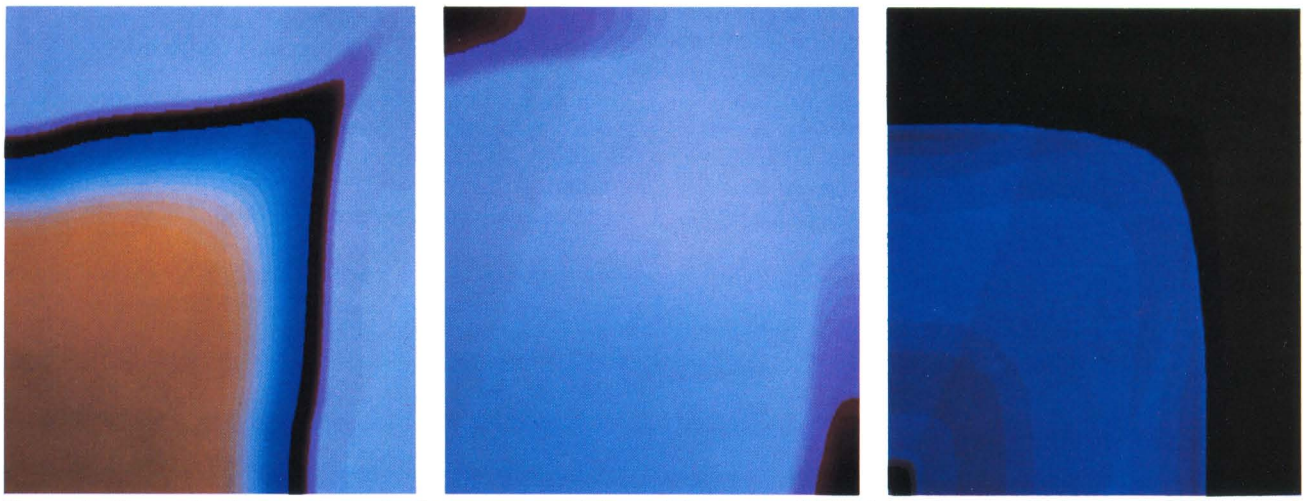

Figure 1. Concentration contours for base case at 60 days: (a) oxygen, (b) hydrocarbon, and (c) microbes.

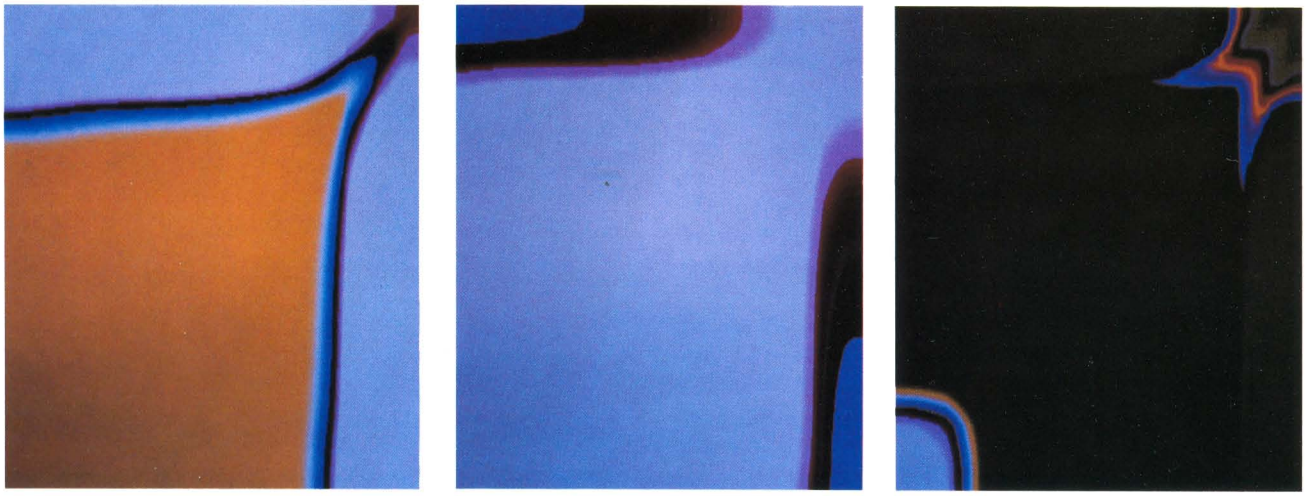

Figure 2. Concentration contours for reduced dispersion case at 60 days: (a) oxygen, (b) hydrocarbon, and (c) microbes.

interpolation was used to define functional values at arbitrary points. One hundred colors were defined as multiple piecewise ramps (step functions). Results were then displayed on a Lexidata Lex-90 using a modified version of a graphics software package developed by Earth Technology Consultants. The final step involved photographing the screen display.

For maximum visibility, output was displayed using the entire screen area. This led to a slight distortion of the results in the vertical direction; in particular, the contours in Figures 1 and 2 are not symmetric because the underlying region is rectangular and not square as modeled. Shadowing effects were caused by timing problems in photographing the screen.

We took for our base case a homogeneous aquifer with a retardation factor of 100 for microbes and 1 for hydrocarbon. This retardation factor for hydrocarbon would be correct for such chemicals as gasoline, since gasoline and water have roughly the same consistency. We also assumed a fairly significant amount of physical dispersion. Figure 1 shows the results of the simulation at 60 days.

The results mimic the type of qualitative behavior one would expect: microbes grow in the presence of oxygen and hydrocarbon and consume the hydrocarbon and some oxygen in the process. Note that by day 60, nearly all the hydrocarbon has been removed from the aquifer except for a small amount still in the upper left and lower right corners (Figure 1b). Also notice the effect of the retardation factor on the flow of the microbes (Figure 1c). By day 60, the microbe front lags significantly behind the oxygen front; that is, the microbes are resistant to flow, as expected.

Figure 2 shows the results of lowering the physical dispersion levels. Pictured are concentration levels at 60 days. Note how much sharper the oxygen front is (Figure 2a). Reduced dispersion results in less interaction between oxygen and hydrocarbon, hence a fairly significant amount of hydrocarbon remains in the aquifer (Figure 2b). The behavior of the microbes near the production well is also quite interesting (Figure 2c). This area is where most of the interaction between oxygen and hydrocarbon is occurring.

As mentioned before, we also tested our method in a nonhomogeneous aquifer, that is, an aquifer where rock permeability varies as a function of space. We generated permeabilities using a program developed by Wilson.[11] Permeabilities were computed on a uniform grid, then interpolated onto our nonuniform computational grid. Figure 3 represents the base case at 60 days in a nonhomogeneous aquifer. Note how the fingering caused by

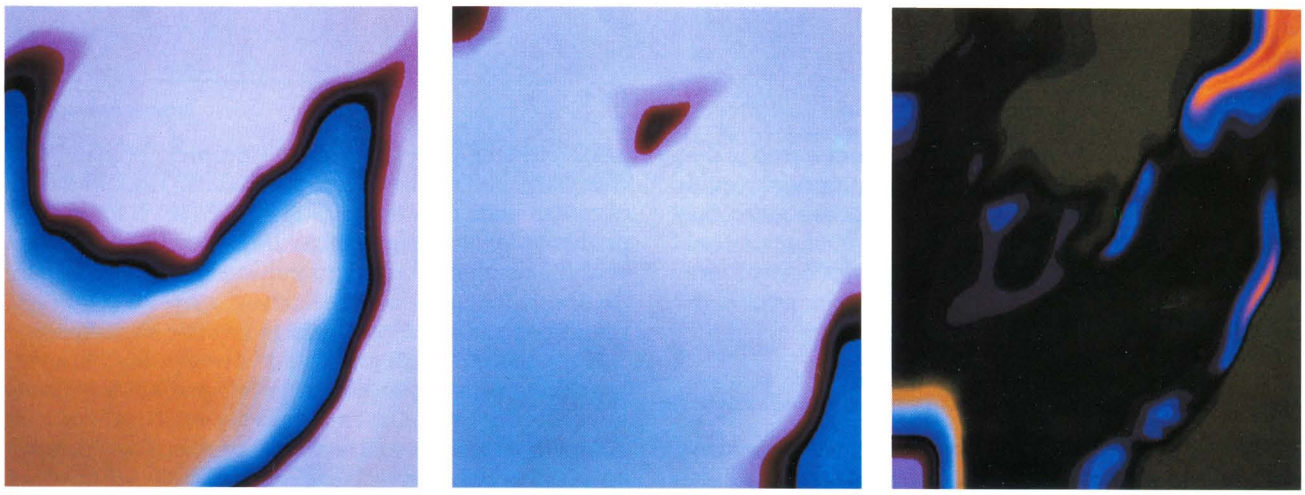

Figure 3. Concentration contours for heterogeneous aquifer at 60 days: (a) oxygen, (b) hydrocarbon, and (c) microbes.

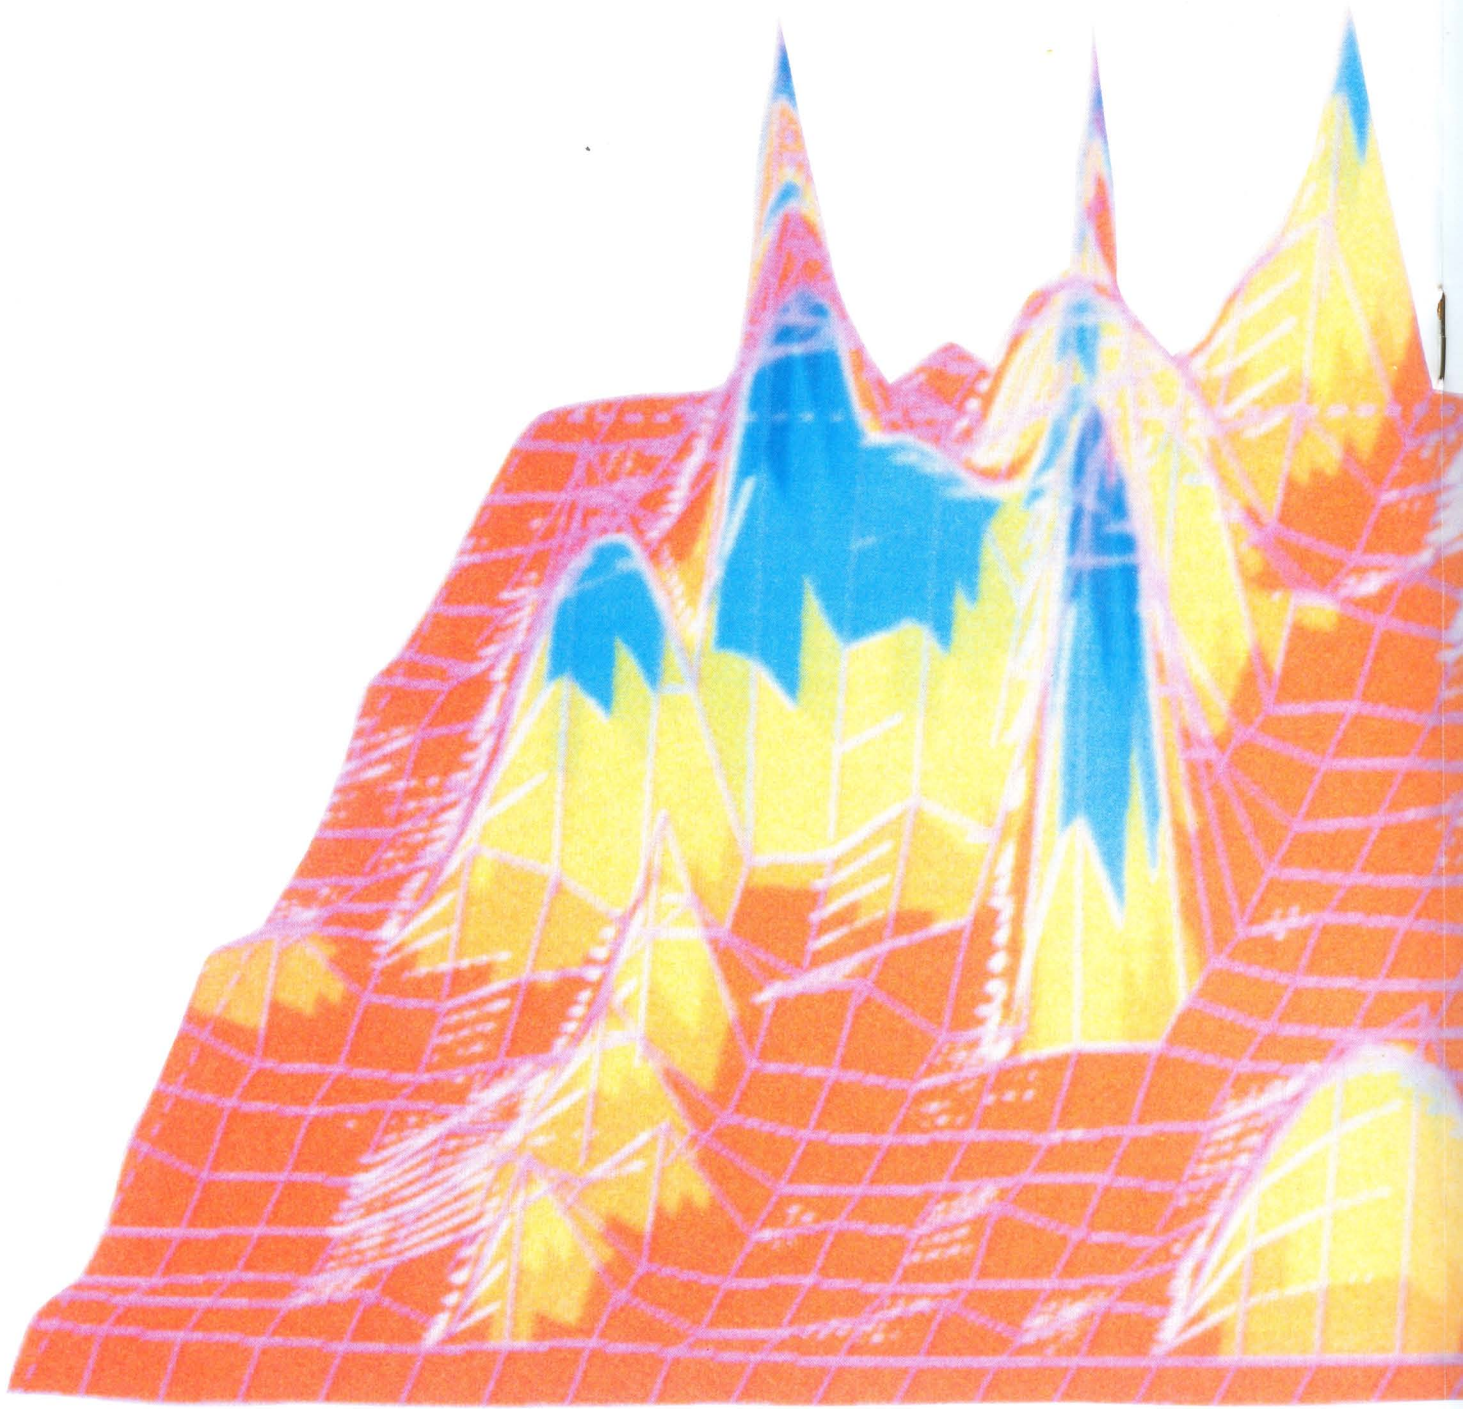

Figure 4. Three-dimensional perspective of permeability function.

the nonhomogeneity results in a less efficient sweep of the hydrocarbon. This is a fairly important feature and would seem to indicate that the less homogeneous the aquifer, the less effective a decontaminant microbial biodegradation may be. Figure 4 is an isometric projection of the permeability function. The peaks indicate zones of high permeability.

Conclusion

Groundwater contamination is an important environmental problem posing a serious threat to the quality of drinking water. The study of decontamination is a relatively new research area involving very complex physics and chemistry, particularly multi-phase flow. The simulation of contaminant flow through the subsurface is a necessary predictive tool to design mitigation methods that will prevent the deterioration of groundwater.

We have formulated a parallel time-splitting algorithm that is accurate and robust and that can be extended to more physically realistic processes involving many more components and mass transfer between components. Simulating these complicated flow problems requires the use of supercomputers such as Cray computer systems and versatile color graphics workstations. □

Acknowledgment

Special thanks to R. C. Borden of North Carolina State University and Peter P. Van Bemmel of Earth Technology Consultants for their help in preparing this article.

About the authors

Clint N. Dawson is a graduate student in the mathematical sciences department at Rice University. He received an M.S. in mathematics from Texas Tech University in 1984. His research interests include numerical simulation of advection-flow problems.

Mary F. Wheeler received her Ph.D. in mathematics from Rice University in 1971. She is a professor in mathematical sciences at Rice. Wheeler's research interests center on numerical simulation of partial differential equations and scientific computation with emphasis on advection-diffusion-reaction and domain decomposition for diffusive problems.

Tung M. Nguyen is a member of a new product development group at Cray Research. He has been employed by Cray Research for over three years, having previously worked in the company's Petroleum Region in Houston. Nguyen received his Ph.D. in mathematics from the University of Houston in 1981.

Stephen W. Poole also is a member of a new product development group at Cray Research. He joined Cray Research in April of this year, having previously worked at ETEC. Poole is pursuing a graduate degree in mathematics at the University of Houston.

References

1. *Houston Post*, "Leaking underground tanks costly," July 6, 1986, Section A, p. 9.

2. *Houston Chronicle*, "EPA study says underground tanks threatening nation's drinking water," June 25, 1986, Section 1, p. 3.

3. Borden, R. C., "Influence of Adsorption and Oxygen Limited Biodegradation on the Transport and Fate of a Creosote Plume: Field Methods and Simulation Techniques," Ph.D. Thesis, Rice University, 1986.

4. Borden, R. C. and P. B. Bedient, "Transport of Dissolved Hydrocarbons Influenced by Reaeration and Oxygen Limited Biodegradation 1. Theoretical Development," to appear in *Water Resources Research.*

5. Borden, R. C., P. B. Bedient, M. D. Lee, C. H. Ward and J. T. Wilson, "Transport of Dissolved Hydrocarbons Influenced by Reaeration and Oxygen Limited Biodegradation 2. Field Application," to appear in *Water Resources Research.*

6. Ewing, R. E. and M. F. Wheeler, "Computational Aspects of Mixed Finite Element Methods," *Numerical Methods for Scientific Computing*, R. S. Stepleman ed., North Holland Publishing Co., 1983, pp. 163-172.

7. Douglas, J., Jr. and T. F. Russell, "Numerical Methods for Convection-dominated Diffusion Problems based on Combining the Method of Characteristics with Finite Element or Finite Difference Procedures," *SIAM Journal of Numerical Analysis*, 19, 1982, pp. 871-885.

8. Ewing, R. E., T. F. Russell and M. F. Wheeler, "Simulation of Miscible Displacement Using Mixed Methods and a Modified Method of Characteristics," SPE 12241, presented at the Seventh SPE Symposium on Reservoir Simulation, San Francisco, California, November 1983, pp. 71-82.

9. Russell, T. F. and M. F. Wheeler, "Finite Element and Finite Difference Methods for Continuous Flows in Porous Media," in *Mathematics of Reservoir Simulation*, R. E. Ewing ed., SIAM Publications, Philadelphia (1983) pp. 35-106.

10. Dawson, C. N., M. F. Wheeler and R. C. Borden, "Numerical Simulation of Microbial Biodegradation of Hydrocarbons in Groundwater," proceedings of the Sixth International Symposium on Finite Element Methods in Flow Problems, Antibes, France, 1986, pp. 353-358.

11. Wilson, J. L., Turning Bands Code, New Mexico Institute of Mining and Technology.

GAMBLING ON THE NEW FRONTIER

Cray Research installed its first customer system just ten years ago. The events that led to that installation are an intriguing story of risk, negotiation, and compromise. For Los Alamos National Laboratory, their CRAY-1 computer system was but one first in a long history of frontier-breaking. For Cray Research, that first installation was a do-or-die maneuver that would mark the beginning of a new era for computing or the end of yet one more bright idea that just didn't work.

An experienced customer

Los Alamos is no stranger to computing. As early as 1945, at about the time of the Trinity Test (the first atomic explosion), Los Alamos began to use the ENIAC (Electronic Numerical Integrater and Calculator) at the University of Pennsylvania. ENIAC could perform in minutes what took days on the Lab's mechanical calculating machines. In 1952, the Lab built MANIAC (for Mathematical Analyzer, Numerical Integrater, and Computer), a machine based on a design pioneered by John von Neumann. In spite of its name, MANIAC was a major step in computer development because with it, computer technology fully progressed from electromechanical to electronic operation.

Whenever possible, Los Alamos has taken advantage of the computer industry's most powerful offerings. That philosophy led to the procurement of Seymour Cray's CDC 6600 system in 1966 and his next machine, the CDC 7600, in 1970. Early in the 1970s, however, principals at the Lab were already searching for even more computer power. "We are always trying to project the Lab's future needs," said Bill Spack, associate C (Computing and Communications) Division leader at Los Alamos. "We stay in touch with the major computer manufacturers, telling them how big our needs are and looking at what they have in development that might satisfy our requirements. Based on that kind of information, we anticipated that a next generation system (a significant step in power above the CDC 7600) would be available sometime in 1977. So we built that assumption into our budget planning for the year," Spack said.

In the fall of 1975, a group from Los Alamos making a round of vendor visits called on Seymour Cray in his small Chippewa Falls, Wisconsin laboratory. "We knew of Seymour Cray from his work at Control Data," said Jack Worlton, Los Alamos Lab Fellow. "One of the problems for someone like Seymour starting his own company is being accused of taking proprietary work from his old employer. For whatever reason, Seymour decided to do something he had never done before: build a vector machine." Much to their surprise, the Los Alamos group learned that Cray planned to have a working prototype of his CRAY-1 computer system, complete with vector capabilities, by March of the following year.

"Try before you buy"

Cray Research knew that if its first installation was to be successful, the customer would need considerable computer expertise. "We had to be very careful where that first machine went," said John Rollwagen, now chairman of Cray Research. "At the time, we thought there might be as many as a dozen potential customers with the historical experience necessary to handle our new computer." Los Alamos wanted the new Cray system, and the Lab had the kind of expertise that made it a good fit for the first installation.

"I had been an advocate of vector processing for some time," said Jack Worlton. "I had even talked with CDC regarding vector capabilities for their Star 100 system. We had a contract for a Star, but the specifications changed significantly by the time it was ready for delivery. Then along comes Seymour with his new CRAY-1, which was much advanced over the Star 100. It appeared to be just the ticket for us."

In the months following that visit to Chippewa Falls in the fall of 1975, representatives of Los Alamos, Cray Research, and the Energy Research and Development Administration (ERDA, the predecessor to the Department of Energy and the governing body that approved Los Alamos procurements) spent a great deal of time trying to work out the details that would allow the CRAY-1 system to be installed at the Lab. Los Alamos had budgeted for a new system in 1977, not 1976, and because the Lab was trying to procure a new system "out of cycle," the process became complicated. In addition, because of the problems of the CDC Star 100, which Lawrence Livermore National Laboratory had experienced firsthand, there was some apprehension over vector processing and the "next generation" of computers.

The delays in dealing with the government were frustrating to the young company, and even more frustrating to its primary founder. "Seymour finally broke the logjam," said John Rollwagen. "He offered to loan the system to Los Alamos, under certain restrictions, for six months. If the system met their requirements, they would pay us. If it didn't live up to expectations, they would send the machine back." Despite the low-risk aspect of Seymour Cray's proposition, government procurement procedures were still very clear on the requirement that competitive bids be solicited for any new computer system. To comply with those guidelines, and in the interest of fairness to other

possible vendors, the "demonstration and evaluation" option was opened up to any vendors that were interested in the arrangement.

By March of 1976, the details were worked out and Cray Research was able to ship and install its first computer system for a six-month demonstration and evaluation. It was decided that at the end of the six-month evaluation period, a competitive procurement would be conducted. Los Alamos was able to assume a minimum of risk and keep well within government guidelines. Cray Research, on the other hand, had basically laid its last card on the table. "Cray was putting everything it had on a truck and shipping it to the mountains of New Mexico," Bill Spack said. Indeed, aside from some lab equipment and office furniture, the CRAY-1 system, serial number 1, was the company's only real asset. "We had used up our venture capital, so we needed to find a home for serial number 1 very quickly" Rollwagen remembers. "It was a critical time. If the machine hadn't performed, Cray Research wouldn't have continued as a company."

Guinea pigs and new frontiers

The first CRAY-1 computer system was delivered without software. Dean Roush, now vice president of engineering for Cray Research, handled the site planning and mechanical installation along with Jack Williams, another early Cray employee. "We had our problems," Roush said. "We ran into copper shavings, problems with refrigeration lines, things like that. But considering it was our first machine, we didn't run into anything too serious. We weren't as nervous about the installation as we were prior to that time — when it wasn't certain whether or not we would find a customer for our first product."

Once the machine was installed, Lab staff and Cray Research employees worked together to make it functional. They had six months to get the kind of results that would justify keeping the system, and both parties had a stake in making the effort successful. The Lab worked with a rudimentary operating system called BOS (benchmark operating system). A FORTRAN compiler for the system didn't exist at the time, so Los Alamos modified a compiler on the CDC 7600 that generated Cray code. The benchmark results were enough to evaluate the new system fairly, and Los Alamos conducted its competitive procurement with Cray Research and two other vendors participating. Cray Research won the bid. "We took the position in the Lab that performance was all-important," Bill Spack said. "We knew we could utilize all the power a manufacturer could provide." Serial number 1 went "on revenue" in October of 1976, infusing new life into Cray Research and beginning a relationship that has resulted in seven Cray system installations at Los Alamos over the past decade.

Compatible philosophies, mutual benefits

Los Alamos has continued to go its own direction with software. Norm Morse, the Lab's current C Division leader, was charged early on with the responsibility of porting an operating system to the Lab's CRAY-1. The operating system that emerged was the Cray Time Sharing System (CTSS), a successor to the Livermore Time Sharing System developed at the Lawrence Livermore National Laboratory and the National Magnetic Fusion Energy Computer Center. Los Alamos also developed the Common File System, a mass storage program for Cray computer systems. Both CTSS and the Common File System are now in the public domain, the latter currently being used by more than 15 user sites.

In addition to the pioneering work the Lab has accomplished and shared with others in the supercomputer community, Los Alamos provides valuable input for Cray Research product development. "Back then (in 1976) one million words was a great deal of memory," Norm Morse recalls. "We suggested to Seymour that he needed SECDED (single-bit error correction, double-bit error detection) in main memory. Serial number 3 was the first machine to get SECDED, and that went to NCAR (the National Center for Atmospheric Research, another early Cray Research customer). We traded serial number 1 for number 4, and we still have that machine today." (Serial number 2 was being built without SECDED and was never completed.) Los Alamos was the first site to receive an SSD storage device, and now is also evaluating the first release of CFT77, Cray Research's latest generation FORTRAN compiler.

In many ways, the success of the relationship between Los Alamos and Cray Research can be attributed to the guiding philosophies of both organizations. The mission of Cray Research is to build the world's most powerful general-purpose computer systems. The mission of Los Alamos is to accomplish research for national security and other advancements, which requires the use of the most sophisticated tools available. "The dominant application for Los Alamos is and always has been nuclear weapon design," said Worlton. "The events being studied in nuclear research — the short time scale and extreme pressure and temperatures — make it unsatisfactory to test in a laboratory. At the same time, field testing is time consuming, expensive, and politically restrictive. Computer modeling has always been complementary to nuclear research. Over the years, field testing has been reduced significantly by the ability to perform computer modeling," Worlton said. "It is important to the mission of the Laboratory that we have very powerful computers, and critical to continuing our work that we constantly have more advanced computers," said Bill Spack. "We want to buy the fastest, most effective machines available for our workload," added Norm Morse. "The goal of our division is to maximize the productivity of our 8000 users. Ultimately, we would like to provide them with the kind of power that allows mathematical simulation with 100 percent certainty that they are simulating nature."

Los Alamos has been tackling the most challenging problems at the frontiers of science for more than 40 years. Cray Research has had a part in the last decade of that adventure. According to Norm Morse, a new Cray system at the Lab is fully utilized within two weeks of installation; Los Alamos still needs more power. As long as that need continues, and as long as customers like Los Alamos continue to provide valuable feedback, Cray systems will no doubt continue to find a home in the mountains of New Mexico. □

The virtue of reliability

Whether shopping for a new wristwatch or a used car, every consumer gives some thought to product reliability. But when considering the purchase of a multimillion-dollar computer system, reliability is a central concern — and rightly so. Cray Research recognizes the importance of providing customers with reliable products and so continuously upgrades and expands its reliability assurance measures. In this way the company hopes to ensure that customers enjoy the fullest benefit from their Cray computer systems.

Cray Research's recent efforts promise to provide even greater product reliability. Improvements to date have been dramatic. Cray computer hardware built, installed, and accepted between 1981 and 1985 showed a sixfold improvement in reliability, as measured by mean-time-to-interrupt. This means that, on average, Cray Research's 1985 hardware ran six times longer between interruptions than the company's 1981 hardware. Such a significant improvement is all the more noteworthy considering the greater complexity of multiprocessor CRAY X-MP computer systems compared with 1981's CRAY-1 uniprocessor computer systems.

"Everyone here plays a part in maintaining and improving the dependability of our products," comments Gary Shorrel, Cray's Research's head of reliability engineering. "People working in engineering, development, manufacturing, testing, and field engineering all contribute to the progress we've seen."

An ounce of prevention

Cray Research has upgraded and initiated many reliability improvement measures, ranging from electrostatic discharge controls and updated training courses for manufacturing personnel, to rigorous testing procedures using the latest equipment available. The newly implemented and expanded test procedures have been key ingredients in the effort.

All Cray Research products are tested thoroughly during manufacturing and prior to shipment. Testing concentrates on two levels of computer operation: that of the individual components, integrated circuits (ICs) and modules, and that of the fully assembled computer system. ICs used in Cray computer systems are subjected to a 48-hour burn-in cycle at 125°C with electricity flowing through the circuits. This accelerated life test simulates about 20 weeks of continuous operation. All ICs must pass through this reliability screen before finding their way into any Cray Research product. In addition, random samples of widely used ICs, as well as all ICs from new vendors and new ICs from familiar vendors, are subjected to 1000-hour burn-in cycles as an added safeguard.

The burn-in tests are valuable filters that catch ICs likely to fail during a system's normal operation. "ICs typically fail within one of two time frames," explains Lou Saye, director of engineering and advanced systems at Cray Research. "Some fail during the first few weeks of use because of process problems or material defects. The ones that don't fail right away tend to operate stably for long periods, up to ten years depending on the part, before exhibiting symptoms of wear. Burn-in tests force out those that would fail early, so they don't end up in our products. The

tests also allow us to provide feedback to our vendors, who have been very responsive to our comments. We have seen definite improvements in the quality of the parts we receive from vendors."

Cray Research's insistence on 100 percent burn-in of all ICs is not a standard industry practice. "Most companies use sample tests, and within that sample they establish an acceptable defect allowance," notes Al Grossmeier, Cray Research's head component engineer. "In fact, many manufacturers' sample allowances exceed our total test allowance." The extensive testing procedures are dictated by the complexity of Cray Research's products. Rather than the hundreds of components used in conventional computers, Cray computer systems use thousands. "It's a costly process and has required a significant investment in IC testers, IC handlers, burn-in equipment, and other tools," Grossmeier says, "but the investment has paid off in system quality and reliability."

At the macro end of the testing spectrum is system testing, which can reveal problems undetected by the testing of individual components and subassemblies. And some aspects of system performance, such as timing, can only be tested after a system has been completely assembled. System testing usually takes from 14 to 18 weeks, depending on the system's size. Typically, four to six weeks of the testing period are used for diagnostic and operating system tests. The final system reliability test consists of an extended run of error-free execution. Cray Research's system testing department uses an extensive suite of application programs to thoroughly exercise all aspects of hardware performance, including disks and tapes when appropriate, before certifying a system fit for shipment.

Testing of components and systems is a crucial part of Cray Research's quality assurance program. Not only does it identify and eliminate failure-prone parts but it also provides valuable feedback to designers and engineers, enabling them to spot potential problems early and stop them from becoming real problems.

Lifelong monitoring

To monitor the performance of Cray computer systems in the field and to help direct some of Cray Research's reliability improvement efforts, the company created a computer database in 1982 to store and analyze information from customer sites. Customer engineers at each site send weekly status reports on each of the site's Cray computer systems to their respective regional offices, which then forward the reports to the central database in Chippewa Falls, Wisconsin. Cray Research personnel use the database for trend analyses to help them identify and prioritize hardware reliability issues. This information is also referred to by designers and engineers when considering modifications to existing product designs.

Each Cray product model must adhere to a reliability standard formulated by Cray Research's technical operations division. The status of any system that breaches the standard is reviewed during a weekly division meeting. "Monitoring systems this way helps us to decide whether a problem represents a trend or an anomaly," explains Stuart Drayton, director of field support in the technical operations division. "Depending on our conclusions, we might consider whether the diagnostics and staff training at the site are meeting the site's needs, or whether to send additional expertise from Chippewa Falls. We also relay diagnostic and troubleshooting information to our people in the field."

Drayton adds that the reliability standard set for each product model is intended to be aggressive without being unrealistic. "We want to challenge our designers and engineers," he says. "We have enough history in the database to make sure the standards are realistic, but once a system achieves its reliability standard, we set a higher target."

Cray Research's hardware designers and engineers work toward clear goals: fast processing, fast I/O, and large memories. Nonetheless, they must design and engineer systems that are maintainable. "We have been very conscientious about building in maintainability," Drayton notes. "The design of the I/O Subsystem (IOS) Model C, for example, addresses problems identified in the previous IOS model."

Now well established, the customer database provides a bounty of information that Cray Research personnel can use in various ways to address reliability issues. And it will be updated and expanded as necessary. "We conduct lifelong monitoring of all installed systems," comments Gary Shorrel. "This, combined with the enhanced testing procedures and other reliability programs, can only mean greater reliability in the future."

Cray Research knows that it benefits along with its customers by upgrading its quality assurance measures. Customers enjoy an increase in productivity and fewer headaches. Cray Research benefits from the greater confidence customers place in its products. The results of the company's reliability improvement measures are already being appreciated. Although Los Alamos National Laboratory is precluded from endorsing commercial products, staff members say they are pleased with their Cray computer systems. "We've seen some very noticeable improvements in overall reliability of the Cray systems, particularly the mainframes," comments Fred Montoya, associate group leader for computer operations at Los Alamos National Laboratory. "When problems do come up, the response and support from Cray is always excellent." Los Alamos National Laboratory currently has four CRAY-1 computer systems, one CRAY X-MP/24 computer system, and two CRAY X-MP/48 systems. Encouraged by such responses, Cray Research is taking steps to ensure even greater product reliability. Says Gary Shorrel, "We've seen some major improvements, and everyone takes pride in that, but we know our job isn't done yet." □

Where are all the cycles spent?

Herbert Cornelius, Cray Research GmbH, Munich, West Germany

Since the first Cray supercomputer was delivered to a customer site, Cray Research has been committed to maintaining and improving both hardware and software performance. Performance optimization utilities and tools are invaluable in helping users obtain maximum performance on Cray supercomputers. Tremendous performance gains can be realized by doing some very simple things to fine tune a program's performance. But before those things can be done, one must be able to see what is going on with a program during its execution. Once understood, the user can then go to work. The performance tools described here provide the insight needed to understand what is happening with — and what can be done to improve — program performance.

Vectorization, parallelism, scalar performance, and I/O operations are among the most important factors affecting overall program performance. Various tools provide Cray supercomputer users with information that can result in reduction of CPU time or elapsed time of programs. Other tools provide comprehensive analyses that can lead to program improvement. Many tools are contained in the Cray operating systems, COS and UNICOS, and in the Cray FORTRAN compilers, CFT and CFT77. Other performance tools are found in the BENCHLIB library, which is a collection of highly optimized useful routines for solving various basic problems.

Accessing powerful hardware

The first step in optimizing code is to compile the program with one of the Cray auto-vectorizing FORTRAN compilers, CFT or CFT77. The compilers provide convenient and efficient access to the powerful hardware, especially its vector capabilities. Graphic output showing DO-loop nesting and vectorization for each program unit can be obtained, and any loop that does not vectorize is flagged and the reason for nonvectorization is indicated. At the end of every program unit a table as shown in Figure 1 gives the user an overview of all DO-loops and their vectorization. After this step one can begin analyzing the program's performance on the Cray system.

Identifying CPU-intensive subroutines

To optimize a program, the user must identify heavily used program units and must understand the program's CPU behavior. To do this, a user may call upon FLOWTRACE, one of the most widely used optimization utilities for Cray supercomputers.

FLOWTRACE analyzes program execution to show information about the overall execution of a program or a subroutine. It performs an automatic runtime program analysis that produces a chart at the end of program execution showing the accumulated CPU time, the number of calls, and the total percentage of the time used by each program unit. A dynamic calling tree is displayed. FLOWTRACE is easily invoked by selecting the compiler option ON=F.

Further optimization — identifying major loops

After identifying the most time-consuming subroutines and functions with FLOWTRACE, the BENCHLIB utility SPY can be used to identify the loops and segments of code that are using the most time within each subroutine. SPY offers a fine-grain look at specified program units and provides a direct graphic readout that is particularly useful. With this information the user can concentrate on code that is most seriously degrading performance.

SPY provides an execution profile of a user program. It samples instruction addresses during the program run and produces a report at the end of execution. The report gives the frequency and the percentage of time that the sam-

```
 1                          SUBROUTINE LOOPS(M,N,X,Y,Z)
 2                          REAL X(N),Y(N),Z(N)
 3                          COMMON A(100,100),B(100,100),C(100)
 4  : . . . . . .           DO  40 L = 1, M
 5  :  V . . . . .          DO  20 I = 1, N
 6  :  V                       Z(I) = X(I)*Y(I)
 7  :  V                       IF ( Z(I) .LT. 0.0 ) Z(I) = 0.0
 8  :  V Ur-                   DO  10 J = 1, 3
 9  :  V Ur                       A(I,J) = B(I,J)/C(I)
10  :  V Ur-              10     CONTINUE
11  :  V                       A(I,4) = EXP(B(I,4))-C(I)
12  :  V . . . . .        20  CONTINUE
13  :     Vc-                 DO  30 I = M+1,N
14  :     Vc                     X(I) = A(1,I)*X(I-M)-Z(I)
15  :     Vc-               30  CONTINUE
16  :                         C(L) = A(L,1)+B(L,1)
17  : . . . . . .          40  CONTINUE
18                          RETURN
19                          END
```

Table of loops encountered

LABEL	INDEX	FROM	TO	ADDRESS	LENGTH	PROPERTIES
40	L	4	17	14a	122	OUTER LOOP
20	I	5	12	25c	34	V – VECTOR
10	J	8	10		INLINE	Ur – UNROLLED
30	I	13	15	77d	12	Vc – CONDITIONAL VECTOR

Figure 1. Output from CFT 1.15 with the LOOPMARK parameter.

pled addresses were contained in each user-defined address "bucket." (A bucket may include a subroutine, loop, block, or smaller segment of code.) The report also shows the percentage of time that each bucket consumes of the surrounding routine and the overall program. An advantage of SPY is its ability to sample system library routines linked with the user program. In addition, SPY does not require any changes to the user program. An example of a SPY report is given in Figure 2.

FLOP TRACE and PERF TRACE

The utilities FLOP TRACE and PERF TRACE are also part of BENCHLIB. These two utilities are available only on CRAY X-MP systems since they make use of the hardware performance monitor feature that is specific to CRAY X-MP systems.

FLOP TRACE gives detailed information obtained directly from the hardware about floating point operations, memory accesses, operation rates, and ratios indicating the balance and optimization of the program. An analysis of the FLOP TRACE information allows the user to understand the CPU behavior of the program, how the compiler has mapped the floating point operations and memory accesses onto the hardware, and what kind of further optimization may be needed. To invoke FLOP TRACE, the user specifies the CFT FLOWTRACE option ON=F when compiling the program and tells the loader to load the FLOP TRACE routines. At the end of program execution FLOP TRACE produces a table containing information about the hardware utilization of each program unit as well as the entire program. The table is sorted by the amount of CPU time consumed.

In addition to FLOP TRACE, four PERF TRACE utilities monitor a certain group of hardware events during program execution. One event group can be selected for each program run. The monitored events include hold issue conditions, memory references, vector operations, and floating point operations. Like FLOP TRACE, each PERF TRACE utility generates a detailed analysis report about every program unit encountered during the program run, and shows all the collected data for the specific monitor group.

ROUTINE	LABEL	ADDRESS	HITS	%PRG	%SUB	SECONDS	CUM%	
SAMPLE	SAMPLE	00000201	2	0.0	0.0	0.001	0.0	:
SAMPLE	10A	00000416	10	0.2	0.2	0.005	0.3	:
SAMPLE	10C	00000427	494	10.9	10.9	0.247	11.1	:*****
SAMPLE	10D	00000436	48	1.1	1.1	0.024	12.2	:*
SAMPLE	20A	00000452	18	0.4	0.4	0.009	12.6	:
SAMPLE	20C	00000464	369	8.1	8.1	0.184	20.7	:****
SAMPLE	20D	00000473	538	11.9	11.9	0.269	32.6	:******
SAMPLE	30A	00000510	10	0.2	0.2	0.005	32.8	:
SAMPLE	30C	00000523	837	18.4	18.4	0.418	51.2	:*********
SAMPLE	30D	00000532	476	10.5	10.5	0.238	61.7	:*****
SAMPLE	40C	00000563	880	19.4	19.4	0.440	81.1	:**********
SAMPLE	40D	00000573	833	18.4	18.4	0.416	99.5	:*********

Figure 2. A sample SPY report.

Using PERF TRACE allows the user to understand the behavior of the program on different parts of the hardware and to see how it can be further optimized. The usage of PERF TRACE is similar to FLOP TRACE, except that different BENCHLIB routines are loaded. Neither FLOP TRACE nor any PERF TRACE utilities influence the execution of the user program.

FLOP TRACE and PERF TRACE are both described in greater detail in CRAY CHANNELS Winter 1986, pp. 18-19. Technical information can be found in the BENCHLIB documentation.

Multitasking

When running a multitasked program, it is sometimes difficult to understand the runtime behavior of all tasks running concurrently. To aid in visualizing this parallel activity, the COS 1.15 multitasking history buffer may be used to monitor all active tasks. This utility traces and records all multitasking library calls and all changes of state within the multitasking library scheduler. A chronological report of the activities of all CPUs is made available. With this information the user can obtain detailed insight into the behavior of the application and identify possible performance bottlenecks.

The COS JCL command MTDUMP generates chronological lists of user-selectable events occurring in the multitasking library. Changes in the states of events or locks show the interactions between all tasks. Most of the produced reports are presented graphically to give the user a better overview. For a detailed discussion refer to the *Multitasking User Guide*, SN-0222. Sample MTDUMP output is shown in Figure 3.

When running a program with microtasking, (the low-level parallelism approach) the COS JCL utility PERFMON from BENCHLIB can be used to obtain information about synchronization overhead. PERFMON, which also uses the Hardware Performance Monitor, tracks the direct use of semaphore bits used in synchronization activities. This information tells the user how well his program is load balanced with respect to parallelization. PERFMON is described in the BENCHLIB documentation.

Microtasking is further discussed in CRAY CHANNELS, Summer 1986, pp. 24-27 and the *Multitasking User Guide.*

Maximizing I/O performance

Because programs with significant I/O operations may require special attention, Cray Research has developed several tools for optimizing I/O. For instance, datasets with dominating I/O operations are easily identified using the COS JCL command OPTION, STAT=ON. Program output shows all I/O transfers involving each accessed dataset, including its size, its I/O wait time, and the number of I/O requests performed on the dataset. The use of these datasets can be optimized either by using faster devices such as the SSD storage device or buffer memory. The user also has the option of using other I/O handling techniques such as disk-striping, asynchronous, raw, or queued I/O to exploit the maximum transfer rates.

	Internal Task Identifier			
Time	1	2	3	4
0	$			
4363	*	$		
359183	e.	!		
360179	.	*		
360640	!	P *		
362333	*	*		
362522	*	C *		
363092	P*	*		
363438	C*	*		
1535105	*	(*		
1535481	*	)*		
1535837	*	$		
1556783	*		$	
1560909	*		*	$
16354735	*		(*	*
16355097	*		)*	*
16355452	*		$	*
16415279	*			$
20645823	$			

Key to symbols used

$ begin task/complete task	(* set lock
* running	)* clear lock
! ready	T* TSKWAIT, no wait
. waiting for lock, event or task	E* EVWAIT, no wait
l. begin wait for lock	P* post event
e. begin wait for event	C* clear event
t. begin wait for task	

Figure 3. Sample MTDUMP output.

Conclusion

Various tools are available on Cray computers to provide the user with information about program performance. Much of this information describes runtime behavior of a program, allowing the user to optimize the program for vectorization, parallelism and input/output operations. The information is easy to obtain since no program changes are required. The knowledge and understanding of how to use these tools enables users to get the most out of Cray systems.

Documentation availability

Those interested in ordering BENCHLIB and its documentation should contact Cray Research, Inc., Applications Department, 1333 Northland Drive, Mendota Heights, MN 55120 or their nearest Cray Research regional office. All other documentation referenced in this article can be obtained by contacting: Dennis Abraham, Distribution Center, Cray Research, Inc., 1440 Northland Drive, Mendota Heights, MN 55120 or by contacting the nearest Cray Research sales office. □

About the author

Herbert Cornelius has been an applications and user consultant for Cray Research GmbH in West Germany since joining the company in 1983. Formerly, he was with the University of Karlsruhe conducting research in the area of applied mathematics. He received his Ph.D. in numerical mathematics in 1981 from the Technical University of Berlin.

CORPORATE REGISTER

Worldwide market grows

Cray Research, Inc., expanded its international customer base this summer with several orders from customers outside the United States.

In June, Cray Research announced it had received an order for the first Cray computer system to be installed in Norway. A CRAY X-MP/22 system purchased by the Norwegian research foundation SINTEF is scheduled for installation at the Universty of Trondheim in Norway early in the first quarter of 1987, pending export license approval. SINTEF, the Foundation for Scientific and Industrial Research, will operate the system on behalf of a consortium composed of SINTEF, the Norwegian Institute of Technology, the Norwegian Applied Science and Pure Science Research Councils, and the Norsk Hydro and Statoil companies. The system will also be available for use by the Norwegian research community.

In June Cray Research also announced the order of a CRAY X-MP/12 computer system from INFOCIP of France. INFOCIP is a company formed by IFP (Institut Francais du Petrole) and CISI (Compagnie International de Services en Informatique) for the purpose of acquiring and operating the Cray system. The system will be installed in the fourth quarter of 1986 at IFP facilities at Rueil Malmasion near Paris, pending export license approval.

Cray Research announced in July that a CRAY X-MP/22 supercomputer with an SSD storage device is scheduled for installation at the University of Toronto in Canada during the third quarter of 1986. The system will be used by Ontario universities and commercial users for various research projects. "Our aspiration is to upgrade this facility to the next-generation Cray supercomputer in three or four years and then with more powerful equipment in the 1990s," said David Nowlan, vice president of research for the University of Toronto.

In addition to receiving international orders, Cray Research recently announced that it installed a CRAY X-MP/44 system for Grumman Data Systems Corporation at the George C. Marshall Space Flight Center in Huntsville, Alabama. The system is an integral part of the Engineering Analysis and Data System (EADS) provided by Grumman at the Marshall center. EADS is used for engineering and scientific data acquisition, reduction, display, and analysis in support of space programs.

Cray Research also announced that a CRAY X-MP/14 computer system was installed in July for Conoco, Inc., at the company's Ponca City, Oklahoma facility. Conoco purchased the system for use in advanced seismic data processing and reservoir modeling to support the company's petroleum exploration and production efforts.

Cray Research offers new station releases

Cray Research recently announced new versions of three station software services. The new station releases include features that further enhance the connectivity of Cray computer systems to user front-end systems.

Apollo DOMAIN Station

Release 2.01 of the Apollo DOMAIN Station provides Apollo DOMAIN users with continued easy access to the

power and performance of the CRAY X-MP and CRAY-1 computer systems running Cray Research's COS operating system. New features provided in this release include an icon-oriented user interface, the ability to dispose graphics output directly from a COS job to a DOMAIN node, and installation exits that allow local software modification. The Apollo Station connects COS version 1.12 or later with version 9.0 or later of Apollo's AEGIS operating system through an NSC HYPERchannel (and an NSC HYPERbus, if required). The HYPERchannel connection consists of an IKON 10077-NSC controller in the Apollo DOMAIN Server Processor (DSP) or peripheral node attachment (PNA). The IKON Board is connected to an NSC A400 adapter, which in turn connects to the HYPERchannel trunk.

IBM MVS Station

Release 2.01 of the MVS Station greatly enhances the usability of Cray computer systems running Cray Research's COS operating system. It provides interactive and multiple CPU capabilities for IBM MVS/TSO users. Release 2 of the MVS Station runs on MVS/XA and MVS/370 systems with either JES2 or JES3. It provides the following new features:

- ☐ Support for COS interactive version 1.14 and higher. Within the interactive environment, the user can issue COS control statements and execute programs. The interactive user can transfer datasets between the Cray computer system and the MVS system in the same way Cray batch jobs are transferred.
- ☐ Multiple CPU support. By means of ACF/VTAM application programs, TSO users anywhere in a multiple CPU environment can use the Station's facilities as though they were running on the processor running the Station.
- ☐ Simplified support for RACF. Support for RACF version 1.6 and later releases is simplified, no longer requiring explicit checking and definition of RACF profiles.

VAX/VMS Station

Release 3.03 of the popular VAX/VMS Station is now available. This release includes the following station enhancements:

- ☐ Support for UNICOS, the new Cray operating system
- ☐ On-line driver diagnostics for testing the integrity of FEI-VAX-to-Cray links
- ☐ The new SHOW NETWORK command, which gives detailed information on transfer activity to and from remote nodes
- ☐ A version of Cray interactive that is decoupled from the rest of the station, giving users better system performance
- ☐ A new version of the Station Installation and Maintenance Reference Manual, which gives more detailed instruction for sites installing multiple stations, as well as advice for tuning stations
- ☐ Additional performance improvements and problem-solving aids

This release of the station continues to provide a sophisticated interface to Cray computer systems for DECnet users. From remote nodes, users can submit and monitor jobs, receive job status, and transfer datasets to and from a Cray system as if their VAX systems were directly connected to it.

Version 3.03 of the VAX/VMS Station links VAX systems running VMS with CRAY X-MP or CRAY-1 systems running COS. It also links to CRAY-2 or CRAY X-MP systems running UNICOS and CRAY X-MP systems running UNICOS in Guest Operating System (GOS) mode. The minimum version levels of the participating operating systems are 4.0 for VMS, 1.13 for COS, and 2.0 for UNICOS.

Cray Research opens three new sales offices

Cray Research has recently opened three new U.S. sales offices. The offices will serve both government and commercial accounts.

A new office in Sunnyvale, California, will serve the Western Region's South Bay District. The district includes much of the Silicon Valley area southeast of San Francisco. The office is located at 894 Ross Drive, Suite 203, Sunnyvale, California, 94086; telephone: (408) 745-6466.

Another new sales office is located in Colorado Springs, Colorado. This Central Region office will serve the area's customers from 5085 List Drive, Suite 104, Colorado Springs, Colorado, 80919; telephone: (303) 599-0355.

Cray Research customers in the state of Ohio will be served from a new sales office in Cincinnati. Ohio is part of the Eastern Region's Great Lakes District. The new office is located at 10200 Alliance Road, Suite 230, Cincinnati, Ohio, 45242; telephone: (513) 891-6690.

Cray Research to sponsor international automotive conference

Zurich, Switzerland will be the site of Cray Research's International Conference on Supercomputer Applications in the Automotive Industry, to be held October 7-9, 1986. Representatives from automobile manufacturers, software developers, research firms, and academic institutions will offer presentations on a variety of topics, including computational fluid dynamics applications, structural analysis, crash simulation, computer-aided engineering, and engine combustion modeling.

A panel discussion will address the question "Crashworthiness Simulation: A Myth or Reality?" Computational and software requirements in automotive engineering will be the subject of a round table discussion.

The conference will be held at the Hotel Zuerich in the heart of Zurich. According to tourist information, early fall is a time of mild weather and numerous cultural activities in the city. The conference will also include several social events for participants and their spouses. A registration fee of $500 (U.S.) is required for the conference. Interested individuals should contact Linda Yetzer, Cray Research, 1333 Northland Drive, Mendota Heights, MN, 55120; telephone: (612) 681-3649.

APPLICATIONS IN DEPTH

Reservoir simulators from SSI available on Cray systems

Scientific Software-Intercomp (SSI) of Denver, Colorado, has made available several reservoir modeling software packages for use on Cray computer systems. Among the packages currently offered on Cray systems are a black-oil simulator, a compositional reservoir simulator, a simulator for thermally enhanced oil recovery operations, a dual porosity simulator for naturally fractured reservoirs, a chemical flooding simulator, and a new pre- and post-processing package.

SimBest

SimBest is a black-oil reservoir simulator designed to model virtually all types of reservoir heterogeneities and production peculiarities that occur in black-oil reservoirs (those in which fluid properties depend only on pressure). The package can efficiently solve problems in which one, two, or three fluid phases are flowing in one, two, or three dimensions. Major applications of SimBest include:

- ☐ Natural depletion of oil and gas reservoirs
- ☐ Waterflooding, including pattern flooding
- ☐ Pressure maintenance by gas and/or water injection
- ☐ Recovery efficiency for various operating schemes and optimization of development strategy
- ☐ Water and gas coning studies
- ☐ Unitization and drainage studies
- ☐ Gas and liquids storage in aquifers and depleted hydrocarbon reservoirs

SimBest can accommodate any number of cells, wells, completions per well, layers, or timesteps. Additional features include the ability to process input data and output results in metric units or U.S. equivalents, true three-dimensional geometry, and the incorporation of ESPIDO, a state-of-the-art linear solver.

COMP III

COMP III is a compositional reservoir simulator for analyzing and predicting the component-by-component behavior of oil or retrograde gas reservoirs with complex reservoir fluid systems and displacement mechanisms. It can be used to model the performance of an entire field, or a simple, single-well cross-section, or any configuration in between. COMP III will numerically simulate one-, two-, or three-phase, multicomponent flow in hydrocarbon reservoirs. There is no program limit on the number of components permitted.

Typically, COMP III is used to simulate the component-by-component behavior of hydrocarbon components from methane (C_1) through hexane (C_6) with the C_{7+} fraction split into two or more pseudocomponents. The simulator also

accounts for the behavior and interaction of reservoir impurities such as hydrogen sulfide, nitrogen, and carbon dioxide. Rigorous treatment is given to the interphase component mass transfer between the oil and gas phases, as well as the phase densities and viscosities. Equation of state techniques coupled with table look-up K-values assure accurate, stable results with maximum computational efficiency.

COMP III addresses a range of problems, including:

- Natural depletion of volatile oil or gas condensate reservoirs
- Revaporization of reservoir liquids with lean gas
- Miscible flooding with LPG, enriched gas or solvent
- Gas cycling of gas condensate reservoirs
- CO_2 flooding of oil reservoirs
- Gas injection into volatile oil reservoirs
- N_2 and flue gas projects

COMP III also incorporates the ESPIDO linear solver.

THERM

THERM is a three-dimensional numerical reservoir simulator developed for the design and analysis of thermally enhanced oil recovery operations. It can simulate both steamflood and in situ combustion processes. In the steamflood mode, applications include hot waterflooding, steam drive, and cyclic steam simulation processes. In the combustion mode, THERM simulates wet and dry, forward and reverse combustion processes. The model is applicable to problems ranging from single-component geothermal reservoirs to multicomponent in situ combustion operations that may include any number of oxidation and cracking reactions.

THERM model describes mass transport by Darcy flow, incorporating gravity, viscous, and capillary forces. Heat transport includes convection and conduction mechanisms within the reservoir and conductive heat loss to the overburden and underlying strata. The model allows any number of chemical reactions with reactants, products, and stoichiometry specified through input data.

The primary oil recovery mechanism reflected in the model calculations is viscous displacement enhanced by reduced oil viscosity at elevated temperatures. An additional recovery mechanism represented in the model is solvent extraction, which results from the stripping of lighter ends from the oil and their travel to and condensation at the leading edge of the steam bank. This oil distillation may be represented by any number of hydrocarbon components of varying volatility.

KAPPA

KAPPA is a fully implicit dual porosity simulator designed for efficient simulation of naturally fractured reservoirs. The simulator uses two grid systems, one for the fractures and one for the matrix. It includes a five-component formulation: three pseudo-hydrocarbon components, inert gas, and water.

Unlike most simulators, which require a single pass through the computer (after the data are acceptable), KAPPA employs a two-pass scheme. First the initialization module is run through the computer. The preprocessor checks only the initialization data. Once it is acceptable, a restart file is created that contains all the arrays and information needed for the execution of time steps. The second pass involves the run module. After the run-time data are checked by the preprocessor and deemed error-free, the simulation proceeds. The advantage of the two-pass approach is that, on subsequent runs, it is not necessary to pass through the initialization phase. Unlike most simulators, the run module calls in the initialization restart file for each new pass.

Although designed to solve difficult dual porosity (fractured) reservoir problems, KAPPA will gracefully degrade to a single porosity problem. This means that it can solve any degree of fracture density variation in the reservoir from no fracturing to dense fracturing within the same mode. Fracture-to-fracture, matrix-to-matrix, and matrix-to-fracture fluid flow are all handled efficiently. Special imbibition and drainage relative permeability and capillary pressure curves can be utilized.

CHEMFLOOD

CHEMFLOOD simulates processes by which liquids are injected into a reservoir to reduce oil-water interfacial tension or to improve mobility control. CHEMFLOOD models micellar/polymer, caustic, and surfactant flooding. Two- or three-phase flow may be simulated in one, two, or three dimensions. A unique feature of the model is the specification of phase behavior either by ternary equilibrium diagrams or K-values. The former has the option to include alcohol or other co-surfactants as a fourth volumetric component in a pseudo-quaternary representation, in addition to the oil-brine-surfactant system. The model calculates oil recovery as a function of various reservoir and operating parameters, such as:

- Pattern type and size
- Injection/production well completion intervals
- Vertical and areal conformance
- Salinity gradients between reservoir brine and injected fluids
- Chemical adsorption as a function of rock type (clay content)

CHEMFLOOD uses an alphanumeric keyword-structured, free-field data format and includes auxiliary programs for generating three-phase relative permeability data tables and ternary equilibrium data tables for input directly into the simulator.

SimEase

SimEase is an interactive input/output processor for reservoir simulators that can be used to:

- Prepare the input data
- Analyze the integrity of all input data
- Automatically format all of the input data into an input data set for the simulator
- Automatically generate the job control language
- Direct the execution of the simulation on the workstation or a remote computer and check the progress of an ongoing simulation

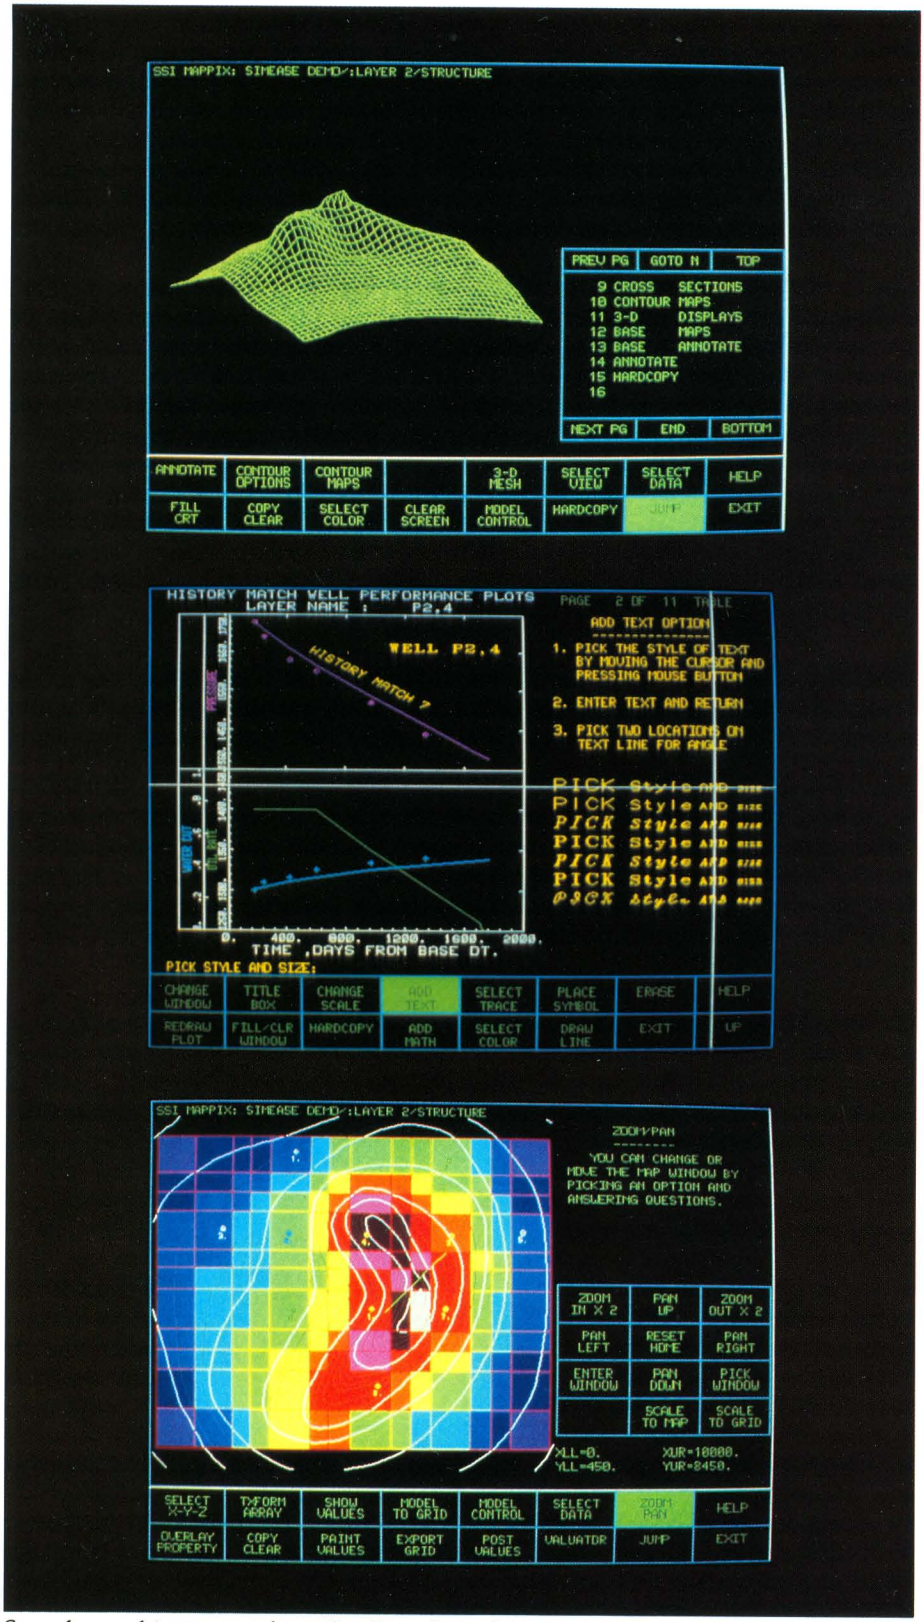

Sample graphic outputs from SimEase. Top, three-dimensional view of user-digitized structure map. Middle, typical well performance plot with user-defined annotation. Bottom, color-coded simulation grid of structure with overlay of structure contours.

- ☐ Analyze the results efficiently in a variety of forms: full-color maps, plots, and tabular reports that can be used for case comparison of history and prediction runs
- ☐ Store and retrieve data via a unique data handler and data pool that minimizes data transfer and storage

SimEase communicates directly with SSI's black-oil simulator, SimBest. It also can be tailored to process input and output data for other reservoir simulators. For more information regarding the use of any of SSI's reservoir modeling packages on Cray computer systems, contact Jim L. Duckworth, Scientific Software-Intercomp, Inc., 1801 California Street, Third Floor, Denver, Colorado, 80202-2699; telephone: (303) 292-1111.

ECLIPSE reservoir simulator available on Cray systems

ECLIPSE 100 is a comprehensive three-dimensional, three-phase, fully implicit, extended black-oil simulator available on Cray systems. It is distributed and supported by Exploration Consultants Ltd. (ECL). Features available with ECLIPSE include:

- ☐ Fully implicit and IMPES options for solving the fluid conservation equations
- ☐ Nested factorization, a fast iterative technique for solving large sets of linear equations
- ☐ Dual porosity/permeability option for modeling highly fractured reservoirs
- ☐ Tracer tracking and API tracking options to determine the movement of "marked" fluid elements and to model the mixing of different types of oil
- ☐ Relative permeability and capillary pressure hysteresis
- ☐ A comprehensive range of well controls and limits
- ☐ Ability to model crossflow and comingling in wells
- ☐ Ability to model highly deviated and horizontal wells
- ☐ Group production and injection controls can be applied at any level in a multi-level grouping hierarchy

- Radial or cartesian grid geometry
- Flexible corner point geometry option to model complex reservoir structures which may include sloping faults
- A range of pre- and post-processors to assist the engineer in preparing the simulation data and analyzing the results

ECLIPSE 100 is undergoing continual development and improvement, and the June 1986 release includes several new facilities. A Carter-Tracy aquifer option is available as an alternative to the existing Fetkovich and numerical aquifer options. The comprehensive well management facilities have been extended still further to include operational constraints on the drilling of new wells (for example, a maximum drilling rate and limits on the number of open wells on each platform). ECLIPSE can open wells automatically in sequence, subject to these constraints.

The ECLIPSE 200 program is also available from ECL. ECLIPSE 200 is a chemical flood and miscible flood extension to ECLIPSE 100 with a fully vectorizing solver.

Using ECLIPSE 200 in the polymer flood application, the water viscosity is treated as a function of the local polymer concentration and, optionally, as a function of the local brine concentration. A separate equation is solved to update the polymer concentration at the end of each time step.

In a physical reservoir, some injected polymer is retained by the reservoir rock. The concentration of a polymer slug is therefore reduced as it advances through the reservoir. Polymer adsorption is modeled using an adsorption isotherm. A simple model of desorption (including the effect of a residual polymer adsorption) is available. In addition, the reduction in rock permeability due to retention of polymer in pore throats can be modeled.

Also available from ECL is ECLIPSE 300, which is a fully-compositional and black-oil simulator with user-controlled options to optimize run-time on vector, parallel, and scalar processors. For more information on ECLIPSE on Cray computer systems, contact Exploration Consultants Ltd., Highlands Farm, Greys Road, Henley-on-Thames, Oxon RG9 4PS, England; telephone: 491-575989.

Reservoir simulators from ERC

ERC Energy Resource Consultants Ltd. has made the PORES and SCORPIO reservoir simulation packages available for use on Cray computer systems.

PORES

The PORES reservoir simulation package is used extensively by the UK Department of Energy to help assess field development plans and to monitor fields that are in production. The PORES black-oil model uses a finite difference method to perform one-, two-, or three-dimensional simulations of oil and gas reservoirs. A new release of the program, PORES 4, has been announced with facilities to model gas systems using a fourth equation. Systems involving any combination of oil, gas, and water are modeled on cartesian or radial grids by using fully-implicit time-differencing in the discretization of the partial differential equations describing the mass conservation of each phase. The nonlinear equations that arise at each timestep are solved using the Newton-Raphson linearization procedure; the resulting linear equations are solved simultaneously by either iterative (nested factorization) or direct methods.

Extensive options exist for representing the characteristics of hydrocarbon reservoirs. PORES has a comprehensive well model that incorporates controls on individual wells and groups of wells, a large module for aquifer modeling including both analytic and numerical models, and a highly developed special (non-neighbor) connection facility. The special connection option can be used for modeling geological faults, aquifers, local grid refinement, dual porosity/permeability systems, and asymmetric radial coning studies. PORES also contains a dynamic pseudo function option and a module for controlling the numerical dispersion associated with the discretization of the equations of motion.

The PORES package consists of eight programs, the primary two being the data initialization and simulation programs. Both are dimensioned dynamically according to problem size by a small preprocessor. This preprocessor also generates all the job control statements necessary to execute a run. The PORES interactive graphics package provides a full-color display and interrogation facility for the results of any simulation. Utility programs for maintaining the PORES source code, generating well bore hydraulics tables, digitizing geologic maps, and summarizing well performance figures complete the package.

SCORPIO

SCORPIO is an advanced three-dimensional reservoir simulator for modeling chemical flood processes. The principal physical and chemical phenomena that SCORPIO models include:

- Temperature distribution resulting from convection in the reservoir fluids and conduction between the reservoir fluids and the rock formation
- Salinity changes that result from convective transport, dispersive mixing, and cation exchange processes
- Non-Newtonian rheology compensating for temperature, salinity, flow rate, and polymer concentration
- Adsorption of each component as a function of fluid composition enabling the modeling of chromatographic separation, cation exchange, and residual resistance
- Definition of reaction rates and stoichiometry for components within a fluid phase allowing for time-dependent chemical and biological degradation of components and for the rate of cross-linking gel formation

SCORPIO is designed to be used for applications such as surfactant flood experiments and field projects, tracer tests, studies of chemical degradation on polymer flood performance, and prediction of in situ gel formation.

The code is run-time dimensioned, enabling storage to be minimized on all major arrays. For example, SCORPIO occupies only 1.3 million words on Cray systems for 2000 grid blocks and 10 chemical components. Machine-specific routines have been used to promote vectorization on Cray systems.

For more information on using PORES or SCORPIO with Cray computer systems, contact Dr. Joe King at ERC Energy Resource Consultants Ltd., 15 Welbeck Street, London W1M 7PF, England; telephone: 019-352315, telex: 24360, Fax: 01 486 3098.

Product monitors Cray system performance

To reap the most number-crunching per day, week, or month from a Cray system, data center managers need to allocate jobs in the manner that most efficiently uses the available CPUs. To assist in this management task, Phoenix Numeric Inc. created the TRACKER, a microcomputer-based system that monitors, reduces, and manages critical performance and utilization data from Cray systems.

Cray-resident software analyzes and delivers performance and utilization data from the Cray system to the TRACKER via the front-end host machine. The TRACKER receives information about CPU performance, job class statistics, memory utilization, I/O characteristics, and job scheduling statistics. Additional performance and utilization data are accurately modeled in the TRACKER software to validate and supplement the information collected from the Cray system itself. Data can be collected at any set time intervals.

The TRACKER presents data to the user through advanced graphics software. A series of nine different displays can be viewed separately or in specified groups. These displays monitor CPU utilization, memory utilization, job roll activity, user I/O, total I/O usage, job class status, jobs in system, job residency, and job residency by priority. Three spare screens are also available for site-specific

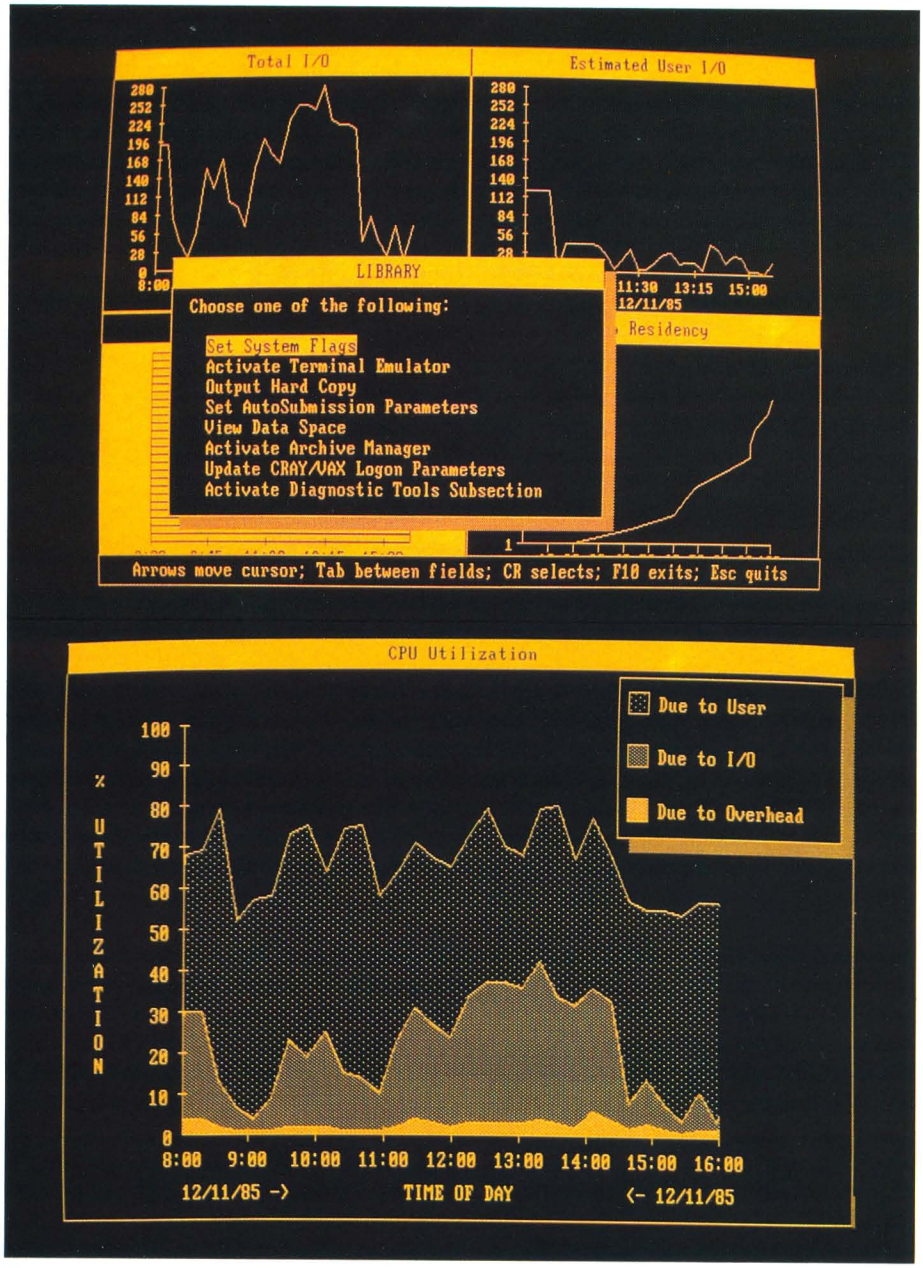

Sample displays from Phoenix Numeric's TRACKER showing library and CPU utilization graph.

displays. The system can recall and analyze user-selected data over any specified time period and can display or correlate information relevant to system utilization and optimization.

The TRACKER is implemented on an IBM AT microcomputer and interfaced directly to a front-end computer through either an RS-232 or an IBM-dedicated synchronous cable. The TRACKER has advanced graphics display capability with a resolution of 720 x 348 pixels. It provides 72 million bytes of on-line storage expandable to 370 million bytes, and utilizes a high performance 60-million-byte tape spooler. Also included is a 160 CPS dot matrix printer and a custom-designed cabinet to accommodate the hardware. For more information on the TRACKER, contact Phoenix Numeric Inc., 9682 Via Excelencia, San Diego, CA 92126; telephone: (619) 549-2929.

USER NEWS

Reservoir modeling: back to the basics

Rumors and investment scams have come and gone, but Minnesota has yet to yield any oil. Nonetheless, oil recovery is on the minds of some researchers at the University of Minnesota's department of chemical engineering and materials science. Pedro Toledo, a graduate student in the department, is using the CRAY-1 computer system at the university to investigate the physical mechanisms that trap oil inside reservoir rock. Such knowledge, he believes, will lead to more accurate reservoir models.

"My goal is to improve understanding of the basic physics at the pore level in reservoir rock," Toledo explains. "The existing industrial reservoir models are macroscopic; they tend to ignore the micro phenomena. I am interested in what goes on inside the pore spaces where the oil actually is trapped. If we can understand this, we may be able to better simulate the processes involved in industrial recovery operations."

The overwhelming complexity of a physical reservoir requires omitting certain factors when designing a reservoir simulator. The factors omitted typically involve fluid flow and surface chemistry in the reservoir's microstructure. These phenomena are difficult to model at the micro level because of the complexities of reservoir topology, geometry, and composition. Topological complexity refers to the intricate shape of the network of interconnections among pore spaces. Geometrical complexity refers to the irregular shapes of the pore spaces themselves. These factors combine to create a situation of such complexity that simulations must be significantly simplified.

The geometry of reservoir rock at the microscopic level can be approximately determined through direct microscopic analysis. Determining topology is much more difficult, however, because the topology of reservoir rock is three-dimensional. Determining topology requires laboratory analysis to interpolate between two-dimensional rock samples. However, even perfect knowledge of geometry and topology would be of limited value because, using current computer hardware, the partial and ordinary differential equations used to describe fluid flow and surface chemistry are manageable only when applied to situations characterized by simple geometry and topology. To compensate for irregular geometry and topology, the equations used in industrial reservoir simulators are augmented by the insertion of phenomenological parameters gathered from experimental studies of fluid flow and displacement in rock samples. These laboratory studies, however, are slow and expensive. As an alternative, Toledo is using the CRAY-1 system at the university to devise computational models that will find the needed phenomenological parameters faster and more economically than current experimental practices.

"Relying on laboratory data is not only slow and expensive, but also can be misleading," Toledo explains. "Laboratory experiments tend to be poor trend predictors. Because the tests are slow and expensive, researchers tend to gather only a limited number of data points for a given sample, so interpolating between the points is very speculative. The computer makes it cost-effective to run many trials to gather more data points and thus makes trend analysis more reliable."

Specifically, semi-empirical parameters gained from laboratory studies fail to reflect the complete history of the saturation process, particularly hysteresis, a shortcoming that leads to conflicts in the solutions of the relevant equations. Hysteresis must be accounted for because recovery operations can create situations where oil pushes water, such as in localized reverse flows and during enhanced recovery operations that follow waterflooding. Again, because of the relative ease of computational modeling, many process histories can be calculated, leading to the determination of more accurate parameters.

"I want to develop a theory of microscopic processes for the prediction of phenomenological parameters, relative permeability in particular," Toledo adds. "I want to be able to tell people in the industry that, for example, in three-phase flow, which is particularly difficult to study experimentally, the relative permeability of water is independent of the other fluids."

Toledo's computer modeling research using the CRAY-1 system relies on several

theoretical tools, including percolation theory, fractals, and resistor network analysis. Network analysis assumes that a network of interconnected pore spaces in reservoir rock can be modeled as if it were a network of interconnected electrical resistors. The flow of fluids through the variously shaped pore spaces is regarded as analogous to the flow of electricity through various resistances, with smaller pore spaces exerting a greater resistance to flow. Pore sizes are varied in the model according to a statistical distribution. The model might step through the displacement process one pore at a time or simultaneously through many pores. The model might continue until all remaining oil is trapped in pore spaces surrounded by those containing water, and is thereby immobilized.

By conducting such studies, and by tracking the advancing front of one fluid as it displaces another under varying conditions of pressure, Toledo hopes to be able to determine computationally the parameters that could otherwise be determined only by laboratory tests. Doing so will require many trials for each process, and Toledo plans to look at all possible histories for three-phase flow, making the speed and memory size of the Cray system a valuable time-saving asset. Although rock samples must still be analyzed physically to get an idea of a reservoir's microstructure, once that idea is obtained, Toledo's methods should provide needed correlations between relative permeability and saturation. This will be a major step toward improving simulator accuracy.

Graphics provides 4-D perspective

Plotting points on a plane of x and y coordinates is a common way to display two-dimensional data. Graphing three-dimensional data is a somewhat greater challenge, requiring the addition of a third coordinate to each plotted point. Representing three-dimensional data that change over time adds considerable difficulty, requiring the display of four dimensions. But this challenge was successfully met by Kendall Preston of the department of electrical and computer engineering at Carnegie-Mellon University in Pittsburgh, Pennsylvania, and the expert graphics staff at Digital Productions, the science and technology division of Omnibus Simulations, Inc. Using custom-written software and Digital Productions' CRAY X-MP computer system, a 3000-frame film was made showing the evolution of polyhedra in a four-dimensional cellular automaton. The data shown graphically in the film were generated by Preston, who has been studying cellular automata for several years and has co-authored a recent book on the subject with Michael Duff of University College, London.

Cellular automata consist of relatively simple central processing units (CPUs) connected to their nearest neighbors in a planar array and running in parallel. John von Neumann, generally regarded as the father of the single-CPU computer, also conceived the multi-CPU cellular automaton. Such arrays have been realized in hardware, early examples being the ILLIAC IV and CLIP computers in the United States and Great Britain, respectively.

Such computing systems are particularly useful for image-processing applications. But their two-dimensional nature limits the application to the processing of black and white images. Applying the same processing techniques to color imagery requires a three-dimensional cellular automaton. The first such device, named TRO, has been built at Carnegie-Mellon University. Using cellular automata to process color imagery that changes over time requires a four-dimensional array. Such a device has yet to be built, but Preston has emulated one in software. Displaying the four-dimensional data produced by the emulator proved to be a challenge, however, which is where Digital Productions and its CRAY X-MP computer system came in. Thanks to a National Science Foundation (NSF) grant that enabled Preston to work with Digital Productions, a Cray computer-generated film was created showing the four-dimensional cellular automata in action. "Each surface contained approximately 100,000 polygons, so using the Cray system was really the only practical way

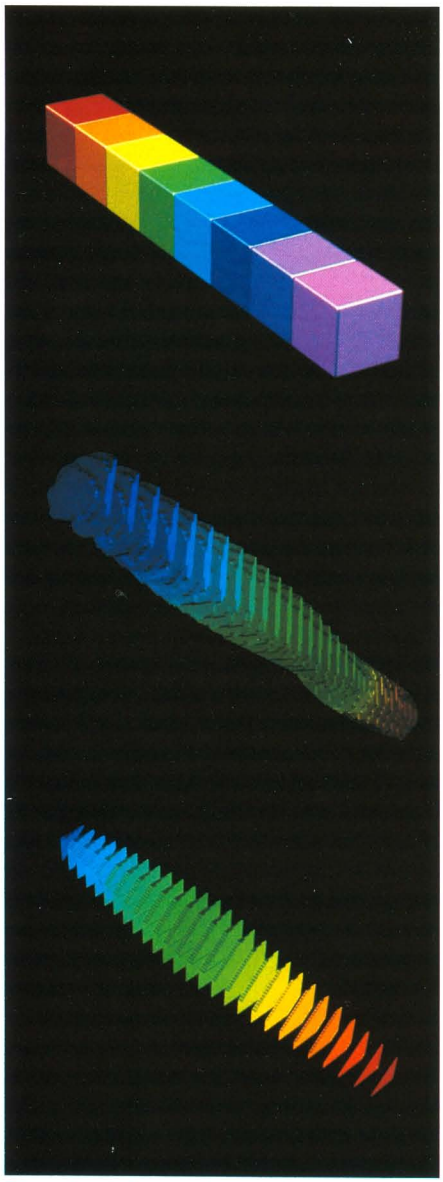

Top, shaded graphic display of a four-dimensional hypercube consisting of eight three-dimensional cubes generated by a color graphics recorder interfaced to the CRAY X-MP computer system at Digital Productions. Middle, multi-dimensional data augmented using the Carnegie-Mellon cellular automaton emulator and displayed in a four-dimensional space consisting of 64 three-dimensional automata each constructed of 262,144 computers. Bottom, four-dimensional hyperzone formed by propagating a single green element sixteen times to form a four-dimensional polyhedron in space. Images courtesy Kendall Preston, adjunct professor, Carnegie-Mellon University.

to compute the images," said Craig Upson, a former producer at Digital Productions, now working at the National Center for Supercomputer Applications at the University of Illinois at Urbana-Champaign.

Processing color images that change over time is perhaps the most obvious application of four-dimensional cellular automata, but four-dimensional data of any sort are good candidates for such treatment. Reflection seismology, for example, may be a suitable application because the data have four coordinates: *x*, *y*, and *z* spatial coordinates, with time as the fourth.

Medical applications of four-dimensional cellular automata have already been explored. Preston is currently working with a medical team to process and plot four-dimensional diagnostic data. The medical data are not four-dimensional in a geometric sense, but comprise four independent measurements taken of cancer patients. The cellular automata emulator has proved an efficient way to calculate the diagnostic sub-type to which each patient belongs. Preston has requested the NSF to extend his present grant to pursue development of four-dimensional graphics for cancer data, again using the Cray system at Digital Productions.

Preston's current plans include not only translating the four-dimensional data into graphic form, but also using the CRAY X-MP system at the Pittsburgh Supercomputing Center to run the emulator itself. The center is administered by a consortium comprising Carnegie-Mellon University, the University of Pittsburgh, and Westinghouse Electric Corporation.

"The software emulator is useful for processing data when the input is relatively slow," Preston said. "But for fast input, processing in four dimensions is practical only if run through actual four-dimensional cellular automata hardware — which doesn't exist as yet. The speed-up could be as great as seven orders of magnitude." If interest in the project grows, Preston said he will pursue such hardware development. In the meantime, much of his energies will go into optimizing the code to run on the Pittsburgh CRAY X-MP system as part of a project to benchmark the code on several types of computing systems.

CRAY BLITZ still reigns supreme

In the game of computer chess, newer isn't always better. Bob Hyatt and Harry Nelson, developers of CRAY BLITZ, learned just that at the World Computer Chess Championships held in Cologne, West Germany, in June. CRAY BLITZ held off a field of 22 hungry challengers to retain its title as the world's computer chess champion.

CRAY BLITZ has held its title since 1983. Many of the challengers were special purpose machines — computers developed solely to play chess — that have only appeared on the computer chess scene in the past year. Other challengers were programs running on networks of general-purpose systems strung together for greater processing power. In contrast, CRAY BLITZ is a FORTRAN program running on the CRAY X-MP/48 computer system, one of the world's most powerful general-purpose systems.

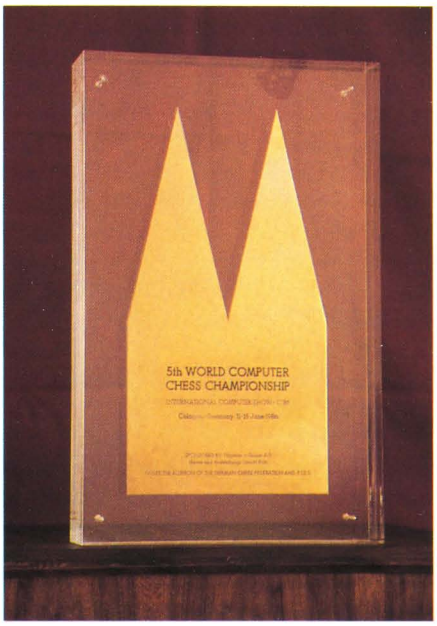

The chess-playing computer program CRAY BLITZ was awarded this trophy after winning the fifth world computer chess championship in Cologne, Germany.

"When we went to our first world championship in New York in 1983, we were the bully on the block," said Hyatt, a researcher at the University of Mississippi. "But now, everyone wants our big general-purpose machine to win. Most of the people who follow these events lean more toward programming than engineering, so they tend to pull for CRAY BLITZ over a black box."

According to Hyatt, that difference is a big part of what keeps CRAY BLITZ ahead of its competition. No current competitor, including CRAY BLITZ, is complete. Human chess masters can still beat the best of them. But while Hyatt and Nelson can change or add lines in a program to improve CRAY BLITZ's playing ability, other competitors must modify hardware.

Although CRAY BLITZ is the reigning champion, Hyatt and Nelson put considerable work into the program recently and expect their development to continue. "When we first started running on a dual-processor CRAY X-MP system in 1983, we knew that we were breaking some new ground multitasking under the Cray operating system, COS," Hyatt said. "Here we were, making our first try at parallel processing just two months before the world championships." Needless to say, CRAY BLITZ walked away with the winner's trophy that year. In the years since, Hyatt and Nelson have improved the program's chess knowledge steadily. In 1984, they developed a new algorithm to exploit the four processors of the new CRAY X-MP/48 system.

What's next for CRAY BLITZ? The ACM (U.S.) championships are coming up later this fall, and Hyatt and Nelson plan to be there with their program. They will also continue to develop CRAY BLITZ. "More Cray processors are on the way, and there's plenty that can still be done with the program," Hyatt said. "In a way, the whole point of the thing is to develop a program that will beat the best human chess master. So that's our goal."

WORLDWIDE BUSINESS CENTERS

Domestic

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

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

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

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

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

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

International

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

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

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

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

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

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

CRAY RESEARCH, INC.