# CRAY CHANNELS

SUMMER 1992 A CRAY RESEARCH, INC. PUBLICATION


Announcing the large-memory solution: the CRAY Y-MP M90 series

# CRAY CHANNELS

## In this issue

University research provides invaluable input to industry and business—input that can contribute significantly to a nation's economic success. Supercomputers have become standard tools in research and engineering, and it is no accident that universities make up a substantial component of the customer base for supercomputers. Cray Research's general-purpose supercomputers can be found in universities worldwide, supporting a range of research projects in the physical and social sciences, medicine, and agriculture. These computing systems are ideally suited to serve the broad-based research environments typical of universities.

This issue of CRAY CHANNELS presents some examples of university-based research conducted on Cray Research systems. This small sample is not a representative, let alone a definitive, sample, but it underscores the unique contribution universities can make to the scientific enterprise. Also included in this issue is an interview with Kenneth G. Wilson, Nobel laureate in physics and author of the Grand Challenges of science and engineering.

Universities are able to make their unique contributions best when researchers have access to the most advanced computing tools available. Cray Research is committed to producing these tools, and together with its university customers helps ensure that future generations of scientists and engineers are prepared to tackle the technical challenges that lie ahead.

## Features

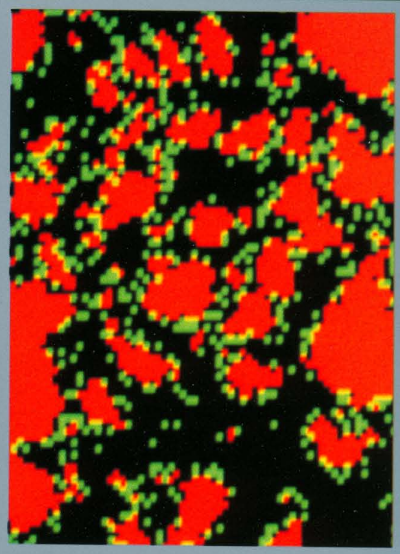

12



## Departments




CRAY CHANNELS is a quarterly publication of the Cray Research, Inc. Marketing Communications Department, Tina M. Bonetti, Director. It is intended for users of Cray Research computer systems and others interested in the company and its products. Please mail feature story ideas, news items and Gallery submissions to CRAY CHANNELS at Cray Research, Inc., 1440 Northland Drive, Mendota Heights, Minnesota 55120. Subscription inquiries and address changes should be sent Attention Department D.

Volume 14, Number 3

**Editorial staff**
Ken Jopp, editor
Mary Miller, associate editor
Heidi Hagen, contributing editor
Molly Sachse, contributing editor

**Design and production**
Barbara Cahlander, typesetter, production artist
Eric Hanson, graphic designer
James Morgan, art director
Cynthia Rykken, graphic designer

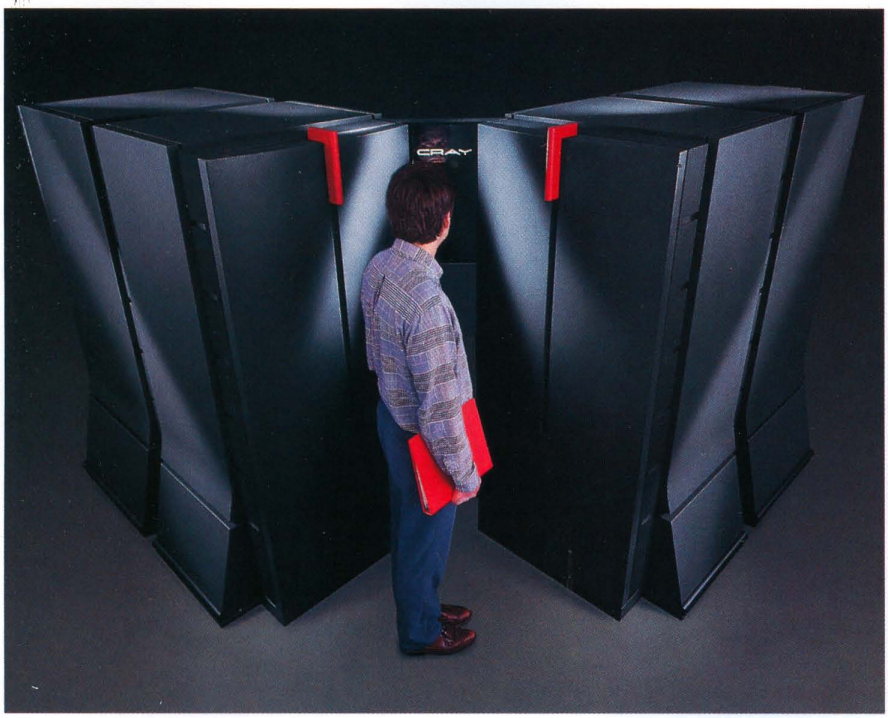


## The large-memory solution:

# Introducing the CRAY Y-MP M90 series of supercomputers

Solving the most challenging computational problems requires the largest computer memories, along with the fastest processors. To meet this need, Cray Research has expanded the CRAY Y-MP family of supercomputers with a series of large-memory systems, the CRAY Y-MP M90 series. These systems combine exceptional functionality with the proven capabilities of Cray Research's UNICOS operating system and provide a high-performance alternative to fast-memory systems when memory size is critical. They can be used as highly responsive interactive tools for application developers and as production throughput engines for the most demanding computational challenges.

The CRAY Y-MP M90 systems offer users an impressive array of productivity advantages:

- ☐ Superior performance on large jobs. Cray Research has optimized important third-party applications to take advantage of the large memories of the CRAY Y-MP M90 systems. The systems minimize solution times by enabling users to run large problems without using peripheral storage media.
- ☐ Improved performance on large interactive workloads. The CRAY Y-MP M90 systems provide optimal performance, increased job size capacity, and the ability to support large numbers of interactive jobs.

**Scattered electrical fields around a double sphere. Such computations are used to obtain radar cross section information and to validate software.**

**Flowfield around an F-18 aircraft showing pressure contours and flow vorticity. The computations required 100 million 64-bit words of memory for the solution at 1.25 million grid points.**

The CRAY Y-MP M90 series; from left to right: the CRAY Y-MP M98, CRAY Y-MP M94, and CRAY Y-MP M92 systems.

- ☐ Improved I/O performance through central memory disk data cache (*ldcache*). By using a portion of the CRAY Y-MP M90 systems' central memory as a disk data cache, users can improve performance on I/O-intensive workloads.
- ☐ Ease of programming. Without the need to write out-of-memory solvers, users can program the CRAY Y-MP M90 systems easily to perform multiple iterations on large jobs in production and research environments.
- ☐ High memory bandwidth. With four memory ports per CPU, the CRAY Y-MP M90 systems feature up to 17.1 Gbytes/sec of memory bandwidth.
- ☐ Easy access to all areas of memory. Unlike the distributed memories and global cache memories used in some systems, any portion of the central memory of the CRAY Y-MP M90 systems can be accessed easily and immediately by any processor. As a result, the central memories of the CRAY Y-MP M90 systems run at full speed on operations, such as global gather, without the performance degradation experienced by distributed and global cache memories.

## The world's largest central memory capacity — the CRAY Y-MP M98 system

The CRAY Y-MP M98 system provides an unprecedented level of computing power, memory capacity, and throughput. It combines a proven, balanced architecture with up to eight CPUs and the latest 4 Mbit and 16 Mbit DRAM memory technology. The CRAY Y-MP M98 system provides access to up to 32 Gbytes (4 Gwords) of central memory. Individual jobs can be as large as 16 Gbytes.

## The CRAY Y-MP M94 system — large central memory capacity in a single cabinet

The CRAY Y-MP M94 system is a powerful, general purpose supercomputer with two to four CPUs, a large central memory, and an I/O subsystem in a single cabinet. By combining its proven, balanced architecture with high-density 4 Mbit or 16 Mbit DRAM memory technology, the CRAY Y-MP M94 system provides 4 or 16 Gbytes (512 or 2048 Mwords) of central memory capacity.

## Large memory within reach — the CRAY Y-MP M92 system

The CRAY Y-MP M92 system offers the advantages of large-scale supercomputers at a cost of ownership that brings it within reach of a broad range of users. With one or two CPUs and as much as 8 Gbytes (1 Gword) of central memory, the CRAY Y-MP M92 system provides a cost-effective solution for users with large memory workloads. The CRAY Y-MP M92 system offers significant performance and price/performance advantages over traditional mainframe and departmental systems. The system can be air cooled and does not require a raised floor, special site plumbing, or motor generator set. As a result, it costs less to install, operate, and maintain. ▬

### CRAY Y-MP M90 system configurations

| Model | CPUs | Central memory size (Gbytes) | I/O clusters | Optional SSD size (Gbytes) |
|---|---|---|---|---|
| CRAY Y-MP M92 | 1 or 2 | 1, 2, or 8 | 1 - 2 | 4 |
| CRAY Y-MP M94 | 2 - 4 | 4 or 16 | 1 - 4 | 4 |
| CRAY Y-MP M98 | 4 - 8 | 8 or 32 | 1 - 8 | 16 |

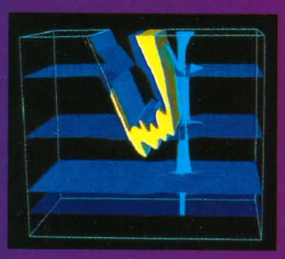

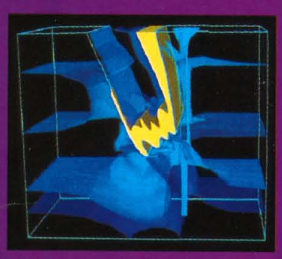

Gold mine model showing effect of ore removal on mine safety. The blue vertical shaft represents the mine shaft; the yellow section represents the home stake formation gold-bearing body.

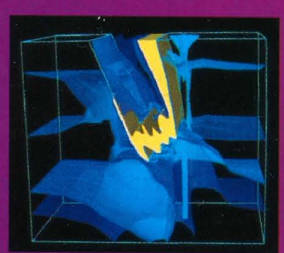

# Computational science and Grand Challenges

## An interview with Kenneth G. Wilson

*Kenneth G. Wilson received the Nobel Prize in physics in 1982 for his contributions to the understanding of phase transitions in matter. He presently is Hazel C. Youngberg Trustees Distinguished Professor of physics at The Ohio State University. From 1985 to 1988 he directed the Center for Theory and Simulation in Science and Engineering at Cornell University, one of five university-based supercomputing centers established by the National Science Foundation (NSF).*

*Wilson coined the term* Grand Challenges *in 1987 to characterize the most pressing and difficult problems in computational science and engineering. These problems include the numerical prediction of global climate change, the mapping of the human genome, and the engineering of ultraminiaturized electronic components. Grand Challenges has been adopted by the Bush administration as the theme for its High-Performance Computing and Communications Program, a federal initiative to fund the development of high-performance computers and networks.*

*Wilson also is co-principal investigator of Project Discovery, an NSF initiative in education reform, a topic about which Wilson is co-authoring a book. In this interview, he shares his thoughts on computational science, the Grand Challenges, and the relationship of universities to both.*

Photo courtesy The Ohio State University

**CRAY CHANNELS:** *Since the development of supercomputers, computational science has been called the third branch of science, joining the experimental and theoretical branches. Do you think this status is deserved?*

**Kenneth Wilson:** I think so. I started thinking along these lines when I noticed that at meetings physicists doing computational work would gather in a corner to discuss the latest workstations and how to deal with algorithms, and so on. The other physicists would come around and quickly get bored and leave. A third branch of science was taking root, and it has been growing ever since. The label seems appropriate also because computational scientists can talk to each other across disciplines, because they share many of the same problems, problems which are different from those shared by theorists or experimentalists. So it's a very practical statement, to call computation the third branch of science.

**CC:** *What has computational science contributed to the research effort that is unique relative to the other branches?*

**KW**: By using computers, scientists are able to tackle problems with more degrees of freedom than they could with a pencil-and-paper approach. Isaac Newton studied the motions of discrete planets, which functionally involved fairly simple orbits depending only on one variable being defined. With computers we are increasingly able to deal effectively with two, three, or more independent variables. As long as you're dealing with functions of multiple independent variables, you need computers to store the information about each function, because it's going to be more than you can remember, and to manipulate the equations that define those functions. Then you need computer graphics to see what they look like. For these reasons computers are having a tremendous impact in fluid dynamics, for example, and weather prediction.

**CC**: *The academic computing community and the supercomputer industry have a longstanding relationship. How would you assess their shared history? What has each contributed to the other?*

**KW**: Some universities already had acquired supercomputers, such as the CDC 6600, in the 1960s. This helped the manufacturers prosper and introduced large numbers of students to scientific computing. But in the late 1970s and early 1980s, funds for computer acquisitions were drying up, and the schools essentially were passed over for the first generation of Cray Research systems. These supercomputers were found largely in the national research labs. As a result, the new levels of computing power were restricted to a narrow focus of applications, those that served the mission of the labs. Meanwhile, there was an explosion of computational activity in the commercial sector, so computer manufacturers were driven toward a more conservative approach to product development to meet commercial needs. When the universities came back into the picture in the mid-1980s, they broadened the base of investigation for supercomputing. I am pleased to see that a realignment of universities and computer companies is taking place. It is important that the most innovative products are in the hands of faculty and students. This obviously is good for science, but the industry also gets a lot of valuable feedback, from the stresses faculty and students place on its products, that it would not get if it served only the commercial markets or the national research labs.

**CC**: *What happened to bring universities back into the picture?*

**KW**: Enough people at universities and at the national labs—and at the NSF—saw that universities had to get back into the picture. They were too important a part of the basic research equation to be left out. And at about this time, global economic competitiveness became a pressing public issue. The competitiveness issue in particular probably was enough to convince Congress and the public to foster computational science in the universities. As a result, the NSF established five university-based supercomputer centers in 1985 and, in so doing, set the stage for a revival of computational science among university researchers.

**CC**: *How did the Grand Challenges idea grow out of this environment?*

## What are the Grand Challenges?

*A definitive list of Grand Challenge problems would be impossible to draft because each area of science and engineering potentially poses Grand Challenges. However, short lists intended to convey the scope of the idea typically include*

- ❑ *High-resolution weather forecasting*
- ❑ *Pollution studies that include cross-pollutant interactions*
- ❑ *Global atmosphere-ocean-biosphere modeling*
- ❑ *Genetic sequencing*
- ❑ *The rational design of new anticancer and anti-AIDS drugs*
- ❑ *The design of new materials, including recording media and high-temperature superconductors*
- ❑ *The aerodynamic design of aerospace vehicles, including the U.S. National Aerospaceplane*
- ❑ *Ignition and combustion modeling in automotive engine design*
- ❑ *Microelectronic design, including the design of quantum switching devices*

*The spirit of the Grand Challenges encourages scientists and engineers to think big. Solving the problems, however, is more than an academic exercise. It is critical to addressing global economics, health, defense, and other vital concerns.*

**KW:** I saw that the researchers tackling the largest problems were trying to do things far beyond the scope of what was possible, even with the available supercomputers. This goes back to the issue of the number of independent variables you want to address. Each independent variable you add to a problem increases the amount of computing power needed enormously. It is not a smooth continuum; each new variable is a big jump in complexity. In quantum chemistry, for example, scientists in the 1930s were trying to perform simple atomic and molecular calculations, but the complexity they faced was so great that they really only turned the corner around 1980. This means they went from just trying to get an answer, any answer, to addressing the margin of error of the answer. Prior to 1980, there were very few quantum chemistry problems about which you could even discuss the margin of error, because for the most part the researchers were struggling just to get any kind of meaningful answer at all. Today, if you calculate very simple molecules, you are concerned about errors of a few percent regarding the configuration, correlation, and energy of the system. This is a meaningful advancement in the state of the art. But even now if you work with molecules that have more than a few atoms, you are back in a situation where the margin of error is uncertain. I think Grand Challenges captures the flavor of these types of problems.

**CC**: *What role do academic researchers play in solving the Grand Challenges?*

**KW**: I see academic researchers spearheading the fundamental breakthroughs and the very broad long-term investigations, trying to bring things down to first principles. Of course people at the national labs also are doing this kind of research, but the breadth you need for some of these problems is only going to come from the universities.

**CC**: *What is the payoff in pursuing the Grand Challenges?*

**KW**: One general payoff is in maintaining a more balanced program of research in the various subdisciplines of science. If you pulled the plug on all the computers involved in atmospheric science, climate change, and so on, for example, the research effort in these areas would largely die, because there would be no way to handle or interpret the data. And these are areas of important practical concern. The primary benefit here is the ability to analyze problems of immediate practical interest. Then again, with something like quantum chromodynamics, the results of the research may not be immediately useful to commercial interests, but one hopes that the effort to use the most advanced computers will act as a technology driver on both hardware and system software.

**CC**: *How do you respond to critics who say that the money spent on so-called big science would better be spent on health care or other social programs?*

**KW**: The problem is that without a growing economy, there will be less and less money available to address social problems, and I believe the economy is not going to grow enough in the next century without basic research investments now.

**CC**: *Since leaving the Cornell Theory Center, you have become an outspoken advocate for education reform. What is your role in this effort?*

**KW**: Education reforms since the 1980s, and even before, largely have been piecemeal efforts in individual disciplines or at individual institutions. These efforts tend to be short-lived because they depend on short-term funding, and there are not enough of them to make a difference on the whole. In the 1980s, there was a movement to get improvements through regulations, especially at the state level, which seems to have succeeded only in complicating the educational bureaucracies. Now an effort is emerging at the federal level focused on trying to specify outcomes, but the NSF already has realized it needs to do more than focus just on outcomes or provide funding for a few years for individual projects. The NSF invited states to submit proposals to reform their math and science education programs. I joined the Ohio group that was drafting a proposal. Our idea was to set up training centers for teachers, and rather than give one- or two-day workshops, have a permanent staff that stayed in touch with teachers to provide long term support. This initiative, Project Discovery, is being funded by the NSF and is now in the start-up phase.

**CC**: *Give us your short answer to the question, "What's wrong with our educational system?"*

**KW**: The main problem is cultural; we don't take teacher education seriously. We have a set of cultural expectations about what goes on in the classroom; for example, the idea that teachers would study teaching is not a part of our educational culture. In Japan, teachers spend less than four hours per day in the classroom; part of the rest of their work day is spent in professional training. They are continually learning to be better teachers. Our culture has not supported that kind of arrangement.

This is a critically important issue and relates to our economic competitiveness. The education of the work force is now a major factor affecting our competitiveness. Initiatives like Project Discovery are coming out of the NSF because the high-level policymakers are painfully aware of the need for better educated professionals and workforce entrants.

Economic competitiveness requires improving math and science education, in particular, for the entire population so that workers can handle the problems that arise in today's jobs. Progress on the Grand Challenges will provide crucial input to U.S. industry in areas such as pollution control and the design of new materials, but for these inputs to show up in new products, U.S. industry needs an adequate supply of skilled scientists, engineers, technicians, managers, marketers, and sales personnel. They also should not have to divert scarce R&D dollars to support remedial training of inadequately educated workers.

# Universities and Cray Research
## A partnership for success

*John Champine, Eric Pitcher, and Chuck Swanson, Cray Research, Inc.*

Of all the tools invented in the past 3000 years in the quest for knowledge and understanding, the supercomputer is one of the most significant. Nowhere has this been more evident than at the leading universities of the world, where there has been an explosive growth in the use of supercomputers to support research in most scientific and many social disciplines. Over the five-year period ending in 1991, the number of Cray Research systems installed in the university market has grown threefold, from 16 systems at the end of 1986 to 48 systems at year-end 1991. The university market currently accounts for 16 percent of the company's installed base. This is the third largest industry segment within Cray Research's diverse customer population, and it represents tens of thousands of scientists, engineers, and teachers who rely on the power of Cray Research supercomputers.

There are many reasons for this dramatic growth in the university market. Some of the growth is due to initiatives started inside the company in 1984, when the company's university programs were begun. At the same time, the U.S. government recognized the need for greater access to supercomputers and laid the foundation for the national communication networks and supercomputing centers we have today.

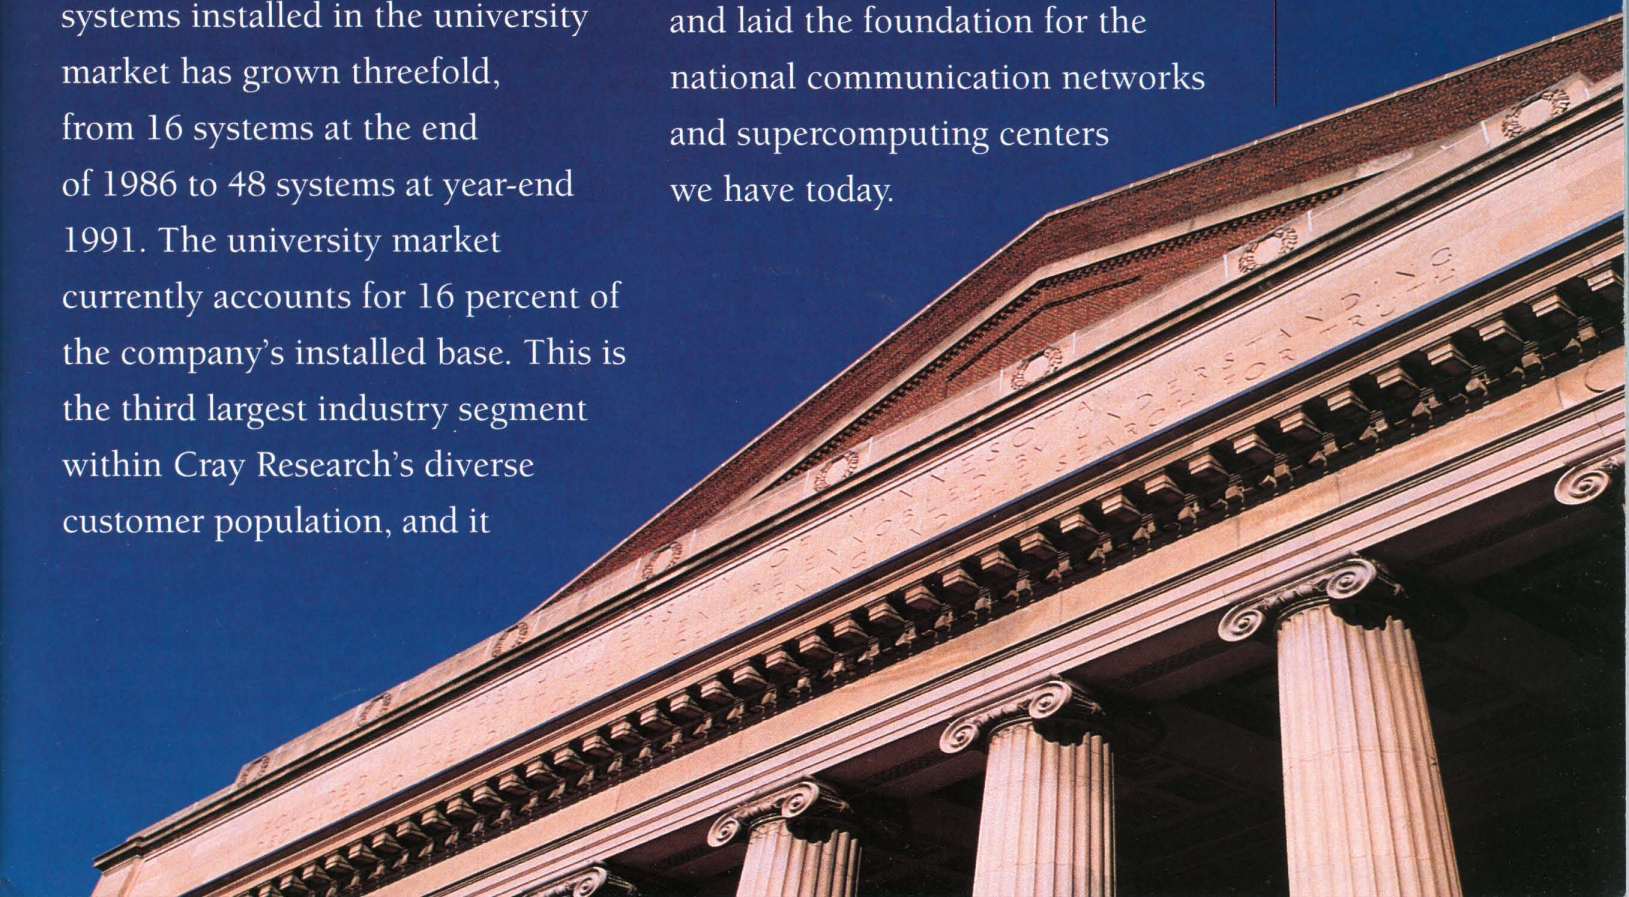

Cray Research's university program was designed to achieve the following objectives:

- ☐ Establish a research and development program through which the company could sponsor leading-edge research at universities that acquire Cray Research systems
- ☐ Act as a catalyst for technology transfer within the computational science arena
- ☐ Promote the expertise of Cray Research in supercomputing applications

To better understand how relationships have evolved with the university community, it is important to understand the market dynamics, including the concerns common to all universities as well as those unique to each.

## The market

Universities have three missions: teaching, research, and service. Supercomputers primarily support research, including the training of graduate students, but they also support the service mission in the form of industrial collaboration between university and industry researchers. Supercomputers also increasingly support the teaching mission as they become incorporated into classroom instruction at both the college and precollege levels.

The importance of supercomputers in advancing basic research at universities is widely recognized. According to the 1989 report published by the U.S.-based National Association of State Universities and Land Grant Colleges (NASULGC), *Supercomputing for the 1990s: A Shared Responsibility,* "Critically important ... is a national commitment to a computing environment that supports and enhances research with the most advanced and most powerful computers of the day—supercomputers." The NASULGC report recommends that this commitment take the form of federally supported national supercomputer centers, as well as regional centers receiving some mixture of federal, regional, and local university support.

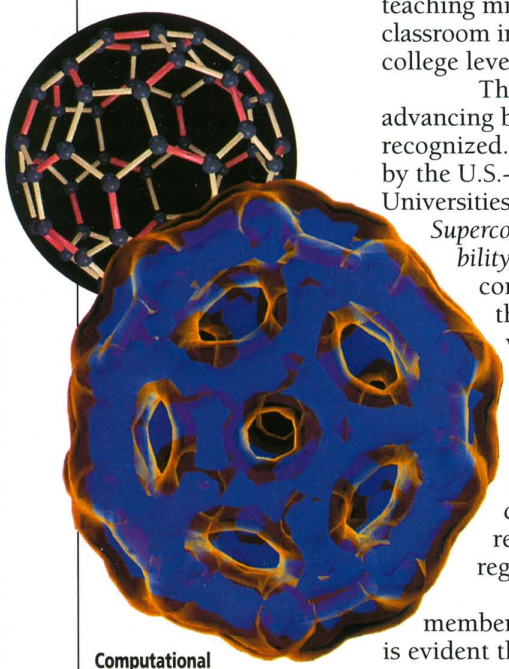

**Computational models of the $C_{60}$ molecule (Buckminsterfullerene), courtesy North Carolina Supercomputing Center.**

Although NASULGC draws its membership primarily from U.S. universities, it is evident that a similar commitment to academic supercomputing exists also in Europe. The establishment of a number of European centers in the late 1970s and early 1980s and the fact that many U.S. researchers were traveling to these centers to access high performance computers provided added impetus for the National Science Foundation's (NSF) creation of the national supercomputer centers in the United States.

## Market segmentation

The university market is characterized by three segments:

- ☐ National centers that provide access to academic researchers through nationwide networks. Cray Research currently has 16 systems at national centers. Examples include three NSF centers in the United States: the Pittsburgh Supercomputing Center (PSC), the National Center for Supercomputing Applications (NCSA), and the San Diego Supercomputer Center (SDSC); as well as national centers in Europe and Asia, such as the Centre for Scientific Computing, Finland (VTKK), the Centre de Calcul Vectoriel pour la Recherche (CCVR) in France, the Science and Engineering Research Council (SERC) in the United Kingdom, and Stichting Academisch Rekencentrum Amsterdam (SARA) in The Netherlands.
- ☐ Regional/state centers that provide network access to researchers within a political region. The size of the systems varies from medium to large, depending on the level of funding. Currently 18 Cray Research systems operate at regional centers. Examples include the Regionales Rechenzentrum der Universität Stuttgart (RUS), in Germany, and the North Carolina Supercomputing Center (NCSC), the Ohio Supercomputer Center (OSC), and the Minnesota Supercomputer Center (MSC) in the United States.
- ☐ Local centers that serve the researchers of the university that purchased the system. There are 14 Cray Research systems at local centers. Examples in this category include the Université Libre de Bruxelles (ULB/VUB) in Belgium, the Aichi Institute of Technology in Japan, the University of Rochester in New York, and the University of California at Berkeley in the United States.

## Technologies and applications

The time-honored model for research is to formulate a hypothesis, develop or extend a computer code, carry out the calculation, and analyze and publish the results. The majority of academic users view supercomputers as the technology necessary to carry out computational research. This situation carries with it a number of implications that set the university market apart from other supercomputing markets.

The applications that run on academic systems are diverse. The subjects addressed include water flow in plant vessels, atomic collisions, prediction of volcanic eruptions, electron mobility in transistors, cross correlation of health records, boundary-layer turbulence, optimal scheduling of transport trucks, black hole simulations, combustion analysis, simulation of the common-cold virus, convection during solidification of alloys, climate change due to deforestation, earthquake analysis of earthen dams, canned-food sterilization, computer vision, mud flow, wave loads on marine structures, and quantum chromodynamics.

As the preceding list suggests, most of the codes that run on academic systems are home grown. Because most of the major academic supercomputer centers around the world have Cray Research systems, the ease with which software can be shared by researchers internationally leverages

the advancements made at one center to benefit all centers.

## The Cray Research University R&D Grant Program

Cray Research sponsors a research and development grant program to promote computational research in science and engineering at universities that operate Cray Research computer systems. The objectives of the program are:

- ☐ To complement Cray Research's internal R&D efforts in vector and parallel algorithms and applications development
- ☐ To educate students at all levels in the use of supercomputers
- ☐ To transfer university research results to industrial users of supercomputers
- ☐ To provide financial support to academic centers

Before the beginning of a grant year, each applicant prepares a proposal describing the research planned for the coming year. The proposals undergo a peer review at Cray Research. Considerations made during the review include the quality of science and engineering involved, requirements for a supercomputer (including a plan for vector and parallel processing), and the potential usefulness of the results to industrial scientists and engineers. Grants are awarded annually and are renewable. Cray Research sometimes assigns a technical contact for a project. This person can help guide the technical specification of the project definition and maintain periodic contact with the researcher. The Cray Research contact may provide relevant technical information and inform the researcher of code development, documentation guidelines, and other useful information, such as implementation conventions used within Cray Research.

A major benefit of the grant program comes from the collaborations that develop between university researchers and other Cray Research computer users. Research reports and other deliverables are disseminated by Cray Research to a wide audience within the company and externally to customers and prospects. This dissemination facilitates the transfer of technology from universities to industrial users of supercomputers.

## Types of research projects

Projects supported by Cray Research at universities comprise two categories—computational science projects and computer science projects. Computational science projects comprise the largest number and focus primarily on the development of algorithms and applications to advance understanding in a particular scientific area. They also are generally performed at the university by the principal investigator(s) and colleagues. Computer science projects deal with "systems" issues, such as operating systems, networking, and user productivity tools. Because this work often complements the work done in Cray Research's software division, the program administrators try to involve one or more persons from Cray Research in these projects.

## Computational science projects

The grant program's effectiveness is demonstrated by the quality of the research produced and the number of researchers educated in computational science methods. For example, Gary Pope and Kamy Sepehrnoori from the Department of Petroleum Engineering, University of Texas at Austin, have developed UTCHEM, an oil reservoir modeling program. UTCHEM is a compositional simulator that can handle various fluids with differing flow, thermodynamic, and chemical properties. Simulators such as UTCHEM are critical for efficient extraction of oil. The ability to handle chemistry makes UTCHEM useful for modeling chemical flooding, an enhanced oil recovery method in which a chemical detergent, or surfactant, is injected into a reservoir to help release trapped oil. A polymer also is injected to improve the flow properties of the resulting fluid mixture, and water injection is included as a driving mechanism. The large number of possible fluids, the complexity of the chemical processes, and the need for a large number of grid blocks require enormous computing power to run field-size simulations with adequate detail. With help from Gene Shiles of Cray Research and Amin Pashapour of Digital Equipment Corporation, UTCHEM was optimized for the CRAY Y-MP system and ran faster than one GFLOPS using eight CPUs.

Using the CRAY Y-MP system at the Pittsburgh Supercomputer Center, John Rosenberg of the University of Pittsburgh is studying the interaction of the enzyme EcoR1 endonuclease with a DNA molecule. The structure of the DNA-EcoR1 endonuclease crystal was determined using the modeling program X-PLOR that combines x-ray diffraction data with a molecular dynamics algorithm. The work included dynamic simulations of DNA that addressed important problems raised by the crystallographic work, and for development of new Monte Carlo methods applicable to biological simulations. This ongoing research is providing new insight into how proteins recognize binding sites in DNA. Rosenberg and his colleagues [Shankar Kumar (University of Pittsburgh), Robert Swendsen (Carnegie-Mellon University), and Peter Kollman (University of California, San Francisco)] won the 1991 Forefronts of Large Scale Computation award for this pioneering work.

The grant program is international, sponsoring research at many European and Asian universities. Recently five graduate students from the Catalonian Research Foundation in Barcelona, Spain, spent several months at Cray Research in Minnesota. They conducted research with assistance from Cray Research personnel in crash simulation, metal forming, maritime hydrodynamics, organometallic chemistry, soil science, and astrophysics.

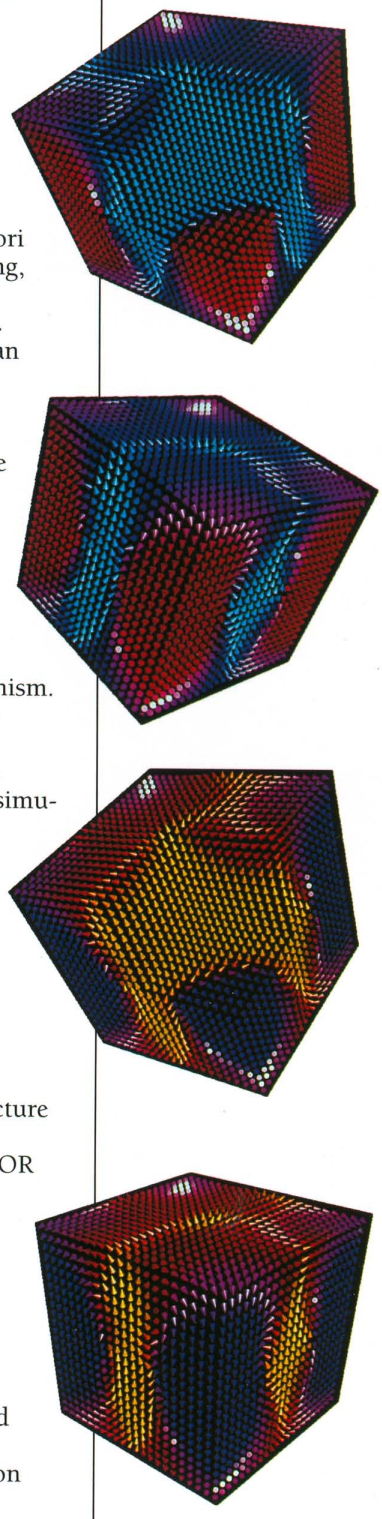

**Models showing magnetic transformations in the iron oxide magnetite, courtesy Ontario Center for Large Scale Computation.**

## Computer science projects

Over the past several years a number of university partners have proposed and implemented several projects in the area of systems software. The National Center for Supercomputing Applications (NCSA) implemented and enhanced many new tools for graphical output of scientific data to workstations and personal computers. They also prototyped an online system for displaying documents and guides to help users more easily program the Cray Research supercomputer. At the Ohio Supercomputer Center (OSC), the development of an advanced graphical display software package called apE (animation production environment) was supported by the R&D grant program for several years. At the central supercomputing facility for the Netherlands, SARA, several projects are under way to investigate effective storage management for the vast quantities of data generated daily on the CRAY Y-MP system.

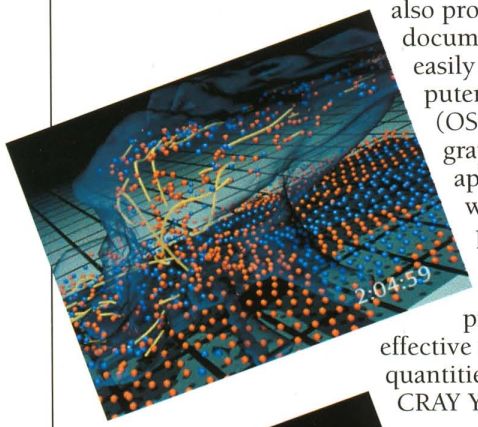



**Severe storm model showing upward (orange) and downward (blue) moving air, courtesy National Center for Supercomputing Applications.**

Two sets of computer science projects, one completed and one just starting, deserve elaboration. The projects are at the San Diego Supercomputer Center (SDSC) and the Pittsburgh Supercomputing Center (PSC).

SDSC is one of five National Science Foundation supercomputer centers. It is a joint project of the University of California at San Diego and General Atomics. In 1989 in conjunction with the acquisition of the CRAY Y-MP8/864 system, Reagan Moore, director of programming and software services at SDSC, proposed six collaborative projects in the area of systems software.

From Moore's viewpoint, the most important projects in support of SDSC users were:

- ☐ The development of a job mix scheduler to allow optimization of throughput on the CRAY Y-MP system
- ☐ The development of improved interactive support by the creation of a second swapper for those jobs
- ☐ The development of resource management including CPU quotas to control multiple projects

Work progressed over the next two and one-half years on these projects. Those applicable to the operating system area (Dynamic Adjustment and Tuning of UNICOS, Resource Control for Principal Investigators, and CPU Time Limit Quota Enforcement) were tackled immediately because they could help boost the high level of throughput at the San Diego site significantly. Some of the ideas developed at SDSC were adopted by the UNICOS development group at Cray Research and have become part of the standard product now being delivered to customers. The file transfer software also was enhanced within the first year at SDSC to provide additional facilities to CRAY Y-MP system users. This included queued file transfer between the CRAY Y-MP system and DECnet, transfers among CRAY Y-MP users, and queued file transfers to remote sites, including directing data to disks and printers. The project, Distributed Transaction-Style Editor, became a true research project insofar as a theory was postulated, and experiments were conducted to prove or reject the theory. The results showed that the technique of distributing the editor worked very well for small files; however, as file sizes grew, although there were significant savings in user time and charges, the amount of data transferred on the network grew substantially. This limited the usefulness of the technique in large file management.

The final project, Autotasking in a Batch Environment, was completed in 1991, and a report was made to the Cray Users Group meeting. This project had three phases:

- ☐ Demonstration of the impact of the Autotasking software feature through the use of the NASA "Kernels" benchmark suite
- ☐ Analysis of the results based on the tuning parameterization developed in the Dynamic Adjustment/Tuning of UNICOS project
- ☐ Throughput improvement by retuning UNICOS to enhance the performance of the Autotasking feature

The results of this study showed that UNICOS 6.1 gave higher speedup and lower overhead than did UNICOS 5.1. It also showed that when running these programs using the Autotasking feature in the highly loaded system at SDSC, the system may have less overhead and still provide reasonable speedups if less than the maximum number of available CPUs is requested for the specific user job.

A second customer with whom Cray Research is cooperating on computer science projects is the Pittsburgh Supercomputing Center (PSC), a joint project of the University of Pittsburgh, Carnegie-Mellon University, and Westinghouse Electric Corporation. In November 1991, PSC and Cray Research reached an agreement to install a CRAY Y-MP C90 system in October 1992. As part of those agreements, PSC committed to working closely with Cray Research to develop

- ☐ Easy and efficient ways for users to access their files from a supercomputer
- ☐ Methods and tools to make it easier for users to run programs in a heterogeneous computing environment—a system linking the CRAY Y-MP C90 system with other distributed computers, such as massively parallel systems

The first project, ease of file access, builds on research performed at PSC for the past several years with the Andrew File System (AFS). AFS was developed at Carnegie-Mellon to make file transfer transparent to computer users at workstations, on a local server or system, or on computer systems located thousands of miles away. PSC has designed and implemented ways to vastly improve the efficiency and performance of AFS. When Cray Research

started to investigate the implementation of the newest derivative of AFS, called the Distributed File System (DFS), it created a natural opportunity for PSC to add its research to that effort. DFS is a component in the Open Software Foundation's Distributed Computing Environment, which has generated great interest in the user community.

In the field of distributed processing, PSC has been doing leading-edge work in "heterogeneous computing." This term refers to applying various computer architectures to a single problem or application. Researchers at PSC had developed several algorithms to reduce the time necessary to run an application, by dividing it up to run on the part of the total computer system which is best adapted for each portion of the problem. One such example is the "MaxSegs" program, which divides a gene and protein matching problem into a broad analysis filtering job that runs on a massively parallel architecture and a fine-grained analysis job that runs on the CRAY Y-MP system. The results are a program that runs five to ten times faster than when both parts are run on either of the two computers. With this kind of experience, PSC and Cray Research will work together in the following areas:

- ☐ PSC will act as a test bed for Cray Research tools and libraries for distributed and heterogeneous computing. PSC researchers will provide feedback and suggestions for improvements.
- ☐ PSC will inform Cray Research of progress made at the center and of new ideas generated for distributed and heterogeneous computing.
- ☐ PSC will host annual conferences on distributed and heterogeneous computing, highlighting parallel vector/scalar and massively parallel system integration successes.
- ☐ PSC will investigate data transmission issues because distributed computing depends heavily on the speed and efficiency of the communication network that connects the various resources. The center will perform analysis and make recommendations in the areas of data format and conversion algorithms, with a view toward reducing latency and increasing bandwidth.
- ☐ PSC will perform analysis and make recommendations in the following areas of distributed computing:
  - Issues related to the portability of programs across computer systems
  - The role of the operating system and the features necessary to support distributed heterogeneous and massively parallel-specific processing
  - Partitioning and distribution of functions to appropriate resources

In all cases, PSC and Cray Research will work together in a research and development partnership. "We are pleased to have this opportunity to work with Cray Research, one of our key vendors, whose hardware and software products have enabled so much good scientific work to be done at PSC," said Michael Levine, co-director, with Ralph Roskies, of the Pittsburgh Supercomputer Center.

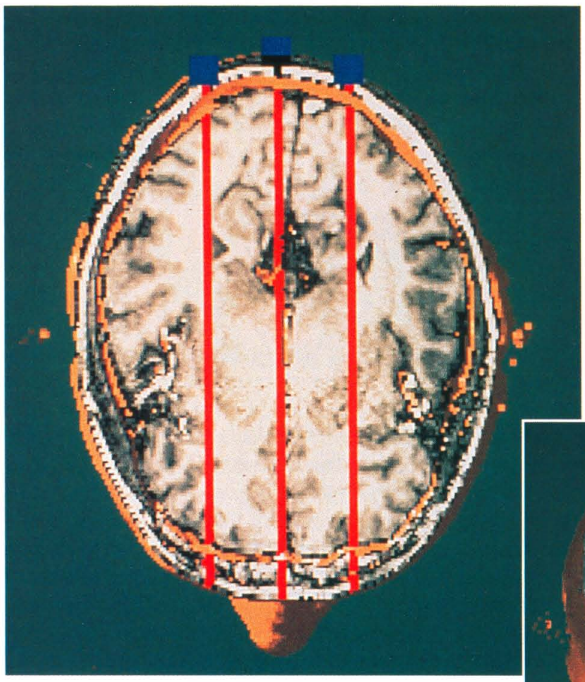

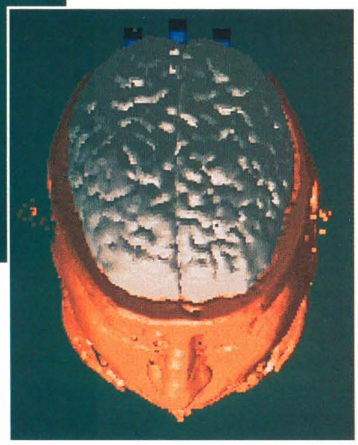

**Three-dimensional brain visualization models, courtesy Ohio Supercomputer Center.**

## Summary

The opportunities created in high performance computing through the partnership of Cray Research and universities have advanced dramatically our understanding of science, nature, and mankind. In just eight years, the number of universities and academic centers with Cray Research supercomputers has grown from three to nearly fifty. The number of researchers and scientists solving problems on supercomputers has grown from hundreds to tens of thousands. The Cray Research University Grant Program and the Computer Science Projects Program have played important roles in accelerating progress in computational science and systems research. The benefits of these partnerships will be felt far into the future—making the world safer and more productive. ▬


***About the authors***

*John Champine is the university and government research and development program manager in Cray Research's software division. He coordinates projects with universities and government research laboratories that complement Cray Research's internal system software design and programming initiatives.*

*Eric Pitcher is director of university marketing for Cray Research. Prior to joining Cray Research in 1987, he conducted global climate research using general circulation models at the University of Miami and the National Center for Atmospheric Research. He received a Ph.D. degree in meteorology from the University of Michigan in 1974.*

*Chuck Swanson is manager of technology transfer in the university programs department of Cray Research. He received a bachelor's degree in physics from the University of Minnesota in 1965 and master's degrees in mathematics from the University of Wisconsin in 1968 and in computer science from the University of Minnesota in 1991.*

# Computer modeling of cement-based materials

*Dale P. Bentz and Edward J. Garboczi*
*National Institute of Standards and Technology*
*Gaithersburg, Maryland*

Cement-based materials are among the most widely used construction materials in the world. Their importance is highlighted by the recent establishment of research initiatives by the National Science Foundation and the Strategic Highway Research Program. The National Science Foundation Science and Technology Center for Advanced Cement-Based Materials (ACBM) is a consortium of five institutions: Northwestern University, the University of Illinois at Champaign-Urbana, Purdue University, the University of Michigan, and the National Institute of Standards and Technology (NIST). The center's goal is to increase the understanding of cement-based materials to establish the scientific basis for developing new, high-performance classes of these materials.

One key goal is to understand the fundamental relationships between the microstructure of a material and its macroscopic properties such as stiffness, strength, transport coefficients, and durability. Without such understanding, materials can be improved only through trial and error. Because cement-based materials are highly random in their internal structure, computer simulation provides the only real hope for discovering relationships between microstructure and macroscopic properties.

Through the ACBM, access to a CRAY-2 system at the Massachusetts Institute of Technology was provided by Cray Research, Inc. This supercomputer was used to study, via large-scale simulation, the microstructure of the interfaces between cement paste and aggregates (sand and rocks), the two major components of concrete. These interfacial zones are often the weak link of the concrete composite with respect to mechanical properties and durability. Understanding their microstructure is critical to designing improved materials.

## Chemistry of cement

The primary cement used today is portland cement, invented in 1824 in England. By firing a mixture of lime and siliceous clay and grinding the product into a fine powder (average particle size of about 1 μm), cement is produced. This cement powder is mixed with water to form cement paste. The cement paste is transformed via chemical reaction between the cement and the water, from a viscous suspension into a rigid load-bearing solid, capable of supporting thousands of pounds per square inch in compression. Concrete is the composite material made from coarse and fine aggregates bound together by the cement paste matrix. To simulate the microstructural development of cement paste, one must consider the physical and chemical transformations that occur as the cement powder reacts with water.

Although cement powder consists of a variety of minerals, the major component is tricalcium silicate, abbreviated $C_3S$ ($C \equiv CaO$, $S \equiv SiO_2$, and $H \equiv H_2O$). A simplified view of how cement paste hardens is as follows: when the cement comes in contact with water, material dissolves off the particle surfaces, forming ions in solution. These ions diffuse throughout the water-filled pore space between the original cement particles and react with the water to form hydration products. Reaction products of two topological types are formed. An amorphous gel product is deposited on the surfaces of the original cement particles while crystalline products nucleate and grow in the water-filled pore

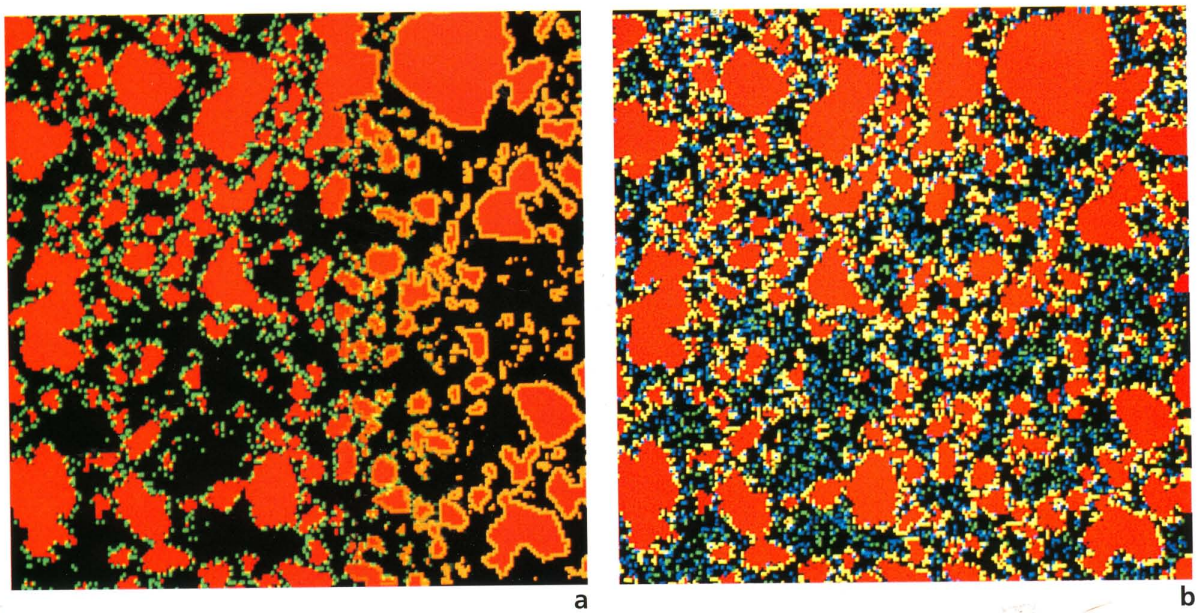

a b

Figure 1. Digital-image-based cement paste microstructural model illustrating (a) initial dissolution of cement particles, (b) diffusion/reaction steps, (c) microstructure after 20 percent of the cement volume has been consumed by hydration, and (d) microstructure after 75 percent hydration.

space. Because water is consumed in these reactions and the lower-density hydration products occupy a larger volume than the cement from which they are formed, porosity decreases as the solid phases grow and connect, eventually producing a rigid solid in a percolation process.[1]

## Simulation of microstructure development during hydration

At NIST, the approach taken to model this microstructural development is based on digital images. A given volume of cement particles randomly dispersed in water is discretized into a three-dimensional simple cubic lattice of cubical pixels, each assigned to a single phase of the cement paste microstructure. Typically, unit cells are 100 to 200 pixels on a side, with periodic boundary conditions used at the outer surfaces. For hydrating $C_3S$, relevant material phases are $C_3S$, water-filled pore space, crystalline calcium hydroxide (CH), and the amorphous calcium silicate hydrate gel (C-S-H), whose chemical composition is approximately $C_{1.7}SH_{4.0}$. Because the reaction and volume stoichiometry of the hydration process are known,[2] microstructure development can be simulated accurately. Because each pixel is approximately 1 μm on a side, the model does not operate at the molecular or ion level, but rather at a subparticle level that is defined by the image resolution. Therefore, the rules of the model are based more on the physical mechanisms of microstructural development than on the precise chemical mechanisms.

The hydration model developed at NIST is based on a cyclic process consisting of dissolution, diffusion, and reaction steps. The model is a type of complex cellular automaton and has much in common with the diffusion-limited aggregation models described by Schaefer.[3] In the dissolution step, all cement ($C_3S$) pixels in contact with water are given a chance to dissolve. The probability of dissolution is based on the local curvature at a surface $C_3S$ pixel, so that sharp edges will dissolve faster than flat surfaces, and smaller cement particles will dissolve faster than larger particles.

When a cement pixel dissolves, it is removed from its location in the microstructure. Diffusing species are created and assigned an initial location within the water-filled pore space of the unit cell. The correct number of diffusing C-S-H and CH species is created to maintain the overall reaction and volume stoichiometry for the hydration reaction. These diffusing species execute random walks, from pixel to pixel, within the available pore space until reaction occurs. The diffusing C-S-H species react by deposition on the original cement or previously deposited C-S-H surfaces. The diffusing CH species react by a nucleation and growth mechanism within the pore space. Thus, for each random step taken by a diffusing CH species, there is a small but finite probability that it may nucleate a CH crystal, upon whose surface other diffusing CH species may deposit.

When all the diffusing species generated in a given cycle have reacted, a new cycle is begun by performing the dissolution process. The reaction is ultimately self-limiting, as the remaining cement surfaces become totally surrounded by hydration products so that further hydration is not possible. Typically, however, 80 to 90 percent of the original cement can be consumed based on these simple rules. This amount of hydration would correspond to real cement specimens one to two years after the original mixing date.

Figure 1a shows a two-dimensional image of actual cement particles, obtained using a scanning electron microscope, with the dissolution phase of the first hydration cycle being performed from left to right. The cement particles are red, the water-filled pore space is black, the diffusing species being generated are green, and the cement surface pixels eligible for dissolution are orange-yellow. In Figure 1b, the dissolution phase is complete and the diffusing species (green: C-S-H; blue: CH) are executing their random walks within the pore space. C-S-H gel (yellow) is seen to be depositing on the surfaces of the original cement particles. Figures 1c and 1d show the same system after 20 percent and 75 percent of the cement has hydrated, respectively. In these images, blue CH crystals have

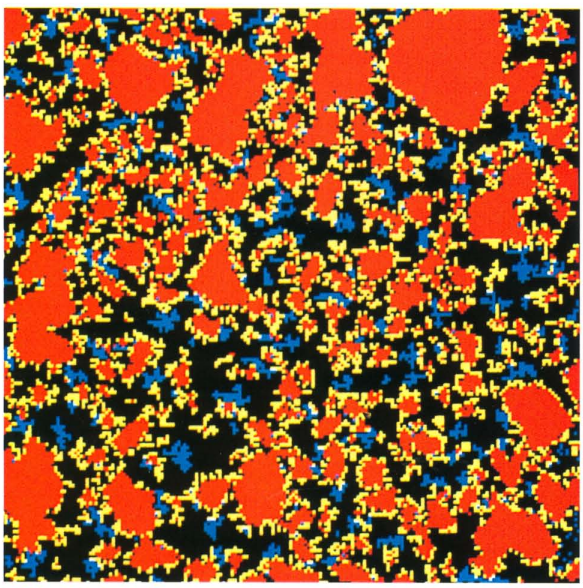

c

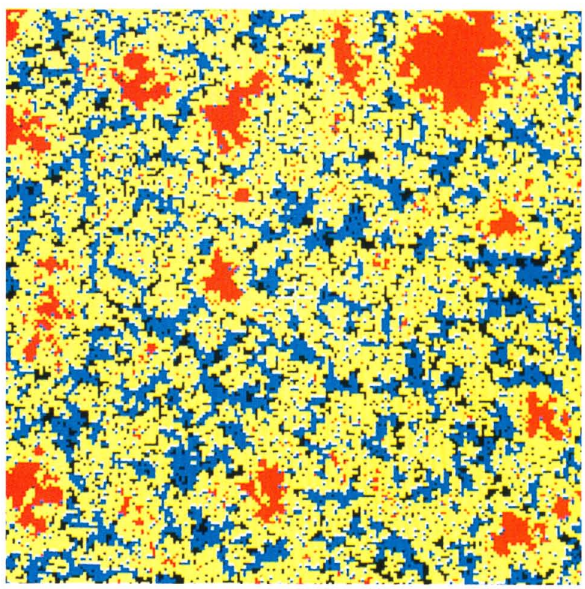

d

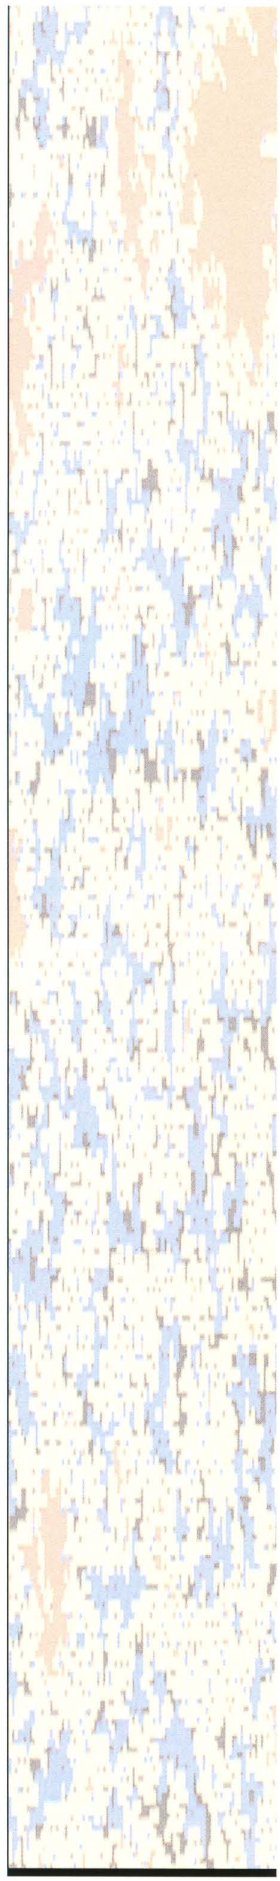

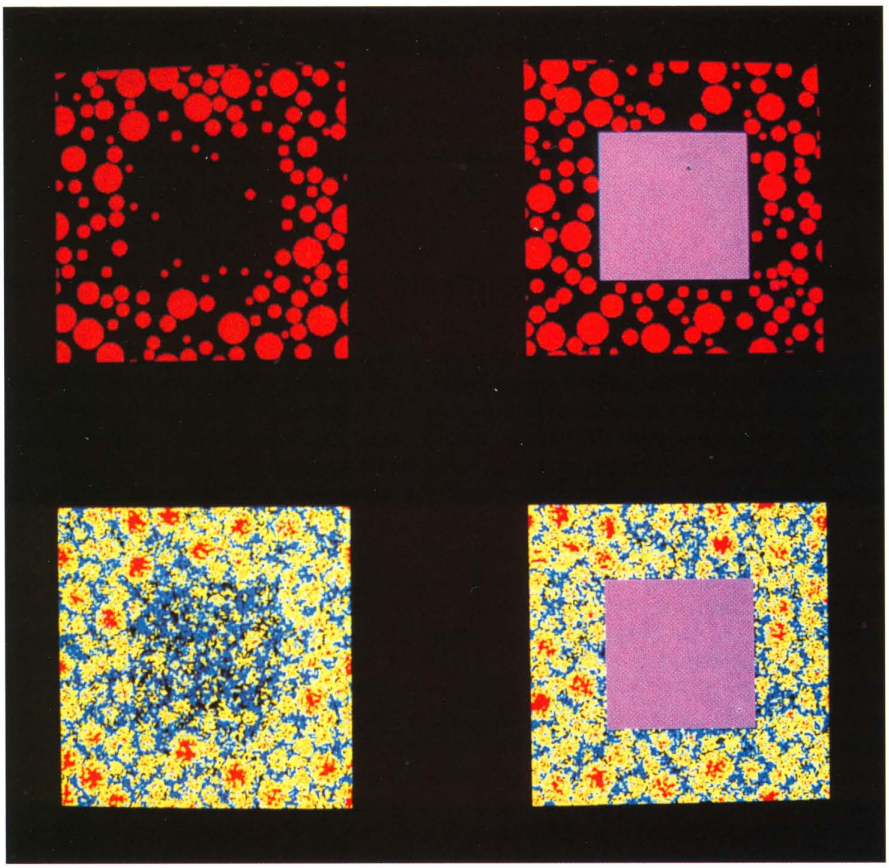

Figure 2. Surface slices (left) and mid-slices (right) for simulated microstructure of normal, neat paste concrete. The top two images show the cement particles before hydration, and the bottom images show the model after 77 percent of the cement has been consumed by hydration.

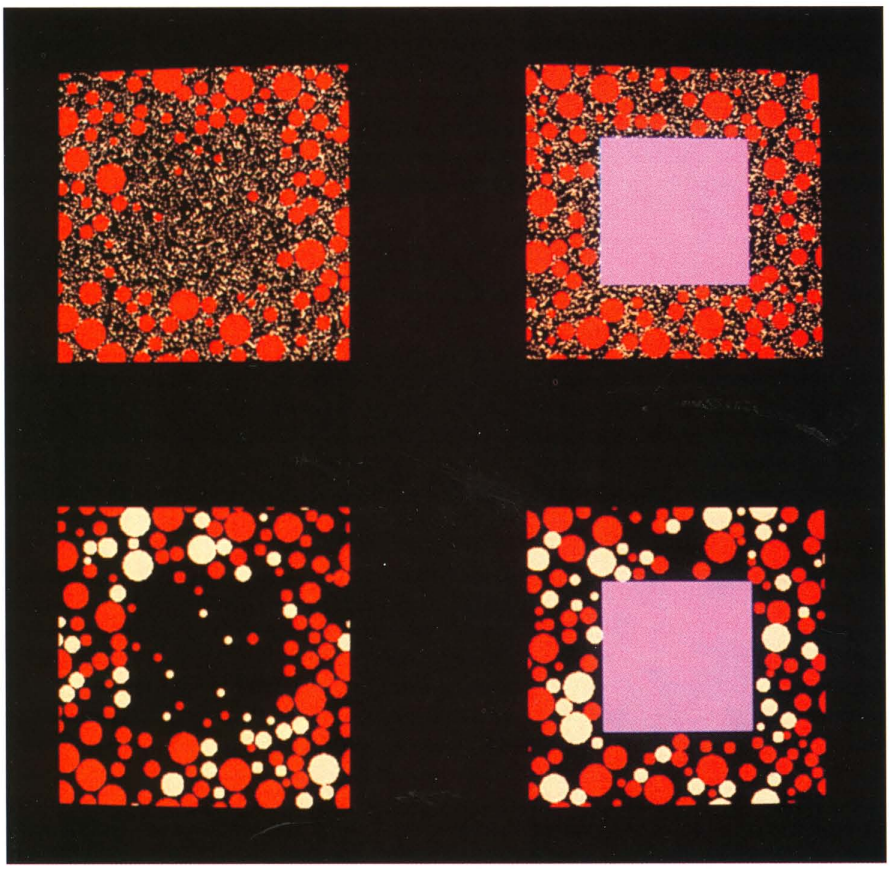

Figure 3. Surface slices (left) and mid-slices (right) for initial particle packings for small (top) and large (bottom) mineral admixture particles, prior to hydration.

nucleated and grown within the black water-filled pore space. In Figure 1d, very little pore space remains, and the remaining cement particles are almost totally surrounded by reaction product, so very little further hydration will occur. Microstructures generated using this model have been shown to compare favorably to those of real specimens.

## Simulation of interfacial zones in concrete

The microstructure of the cement paste formed in the immediate vicinity of an aggregate surface is different from that produced in the bulk paste, far away from any aggregate surface. In general, the interfacial zone extends up to 50 μm from an aggregate, and consists of a region containing less cement, higher porosity, and larger pores than the bulk paste, and often containing large crystals of CH, which have weak cleavage planes and are quite soluble. These interfacial zones are therefore often the weak link in the composite concrete system, both for mechanical strength and durability.

To extend the microstructural model to investigate this phenomenon, a cubic aggregate is placed in the middle of the three-dimensional unit cell, before any cement particles have been placed. In concrete, the smallest aggregate particles (sand) are at least several times larger than the typical cement particle. To satisfy this criterion and still have sufficient room left in the unit cell to randomly place an adequate number of cement particles, the approximate minimum unit cell size required to execute these simulations in three dimensions is 200 by 200 by 200 or 8,000,000 pixels. The aggregate is a 100 by 100 by 100 pixel cube, while the cement particles are digitized spheres with diameters of 11 and 21 pixels. Because of this large system size and the number of operations that must be performed to obtain a completely hydrated microstructure, large-memory supercomputers are necessary.

## Mineral admixtures

Research has been conducted on ways to enhance interfacial zone microstructure and produce stronger, more durable concrete. One of the most successful methods to date has been the incorporation of condensed silica fume (CSF) into the cement paste component of concrete. Silica fume, classified as a mineral admixture, is a very fine powder with particles ten to one hundred times smaller than those of cement powder. It reacts with CH produced during hydration to form pozzolanic or secondary C-S-H. Preliminary simulations suggested that the interfacial zone microstructure in conventional concrete was less than optimal due to the inefficient packing of cement particles around the aggregate.[4] Because silica fume is much finer than cement, it might be expected to pack more efficiently around aggregate particles. Within the model, the well-dispersed silica fume is represented by one-pixel (1 μm diameter) particles.

Two important parameters of mineral admixtures are size and reactivity. The current attainable ranges of these parameters are bracketed by silica fume, which is small and highly reactive,

and fly ash (fine particles collected from coal-fired power plants), which is about the same size as cement and only about 25 percent as reactive as silica fume, because it contains less silica. "Reactivity" is defined here as the volume of CH that can be consumed per unit volume of mineral admixture, producing new CSH via a pozzolanic reaction.

Using the CRAY-2 supercomputer, a model study was conducted in which microstructures were generated using the model outlined above and quantified for each of the six concrete systems listed in Table 1. All mineral admixtures were added by replacing 20 percent of the cement, by mass, with an equal mass of mineral admixture. The simulated images were analyzed by measuring the phase fractions of all phases ($C_3S$, C-S-H, pozzolanic C-S-H, CH, and porosity) as a function of distance from the aggregate surface. "Pozzolanic CSH" means CSH that has been produced by the reaction of CH and the mineral admixture. These phase fractions were averaged over concentric cubic shells going out from the aggregate surface. To gain further insight into the results, two-dimensional color images of slices through the three-dimensional volume were produced. Images taken right through the middle of the aggregate (mid-slices) and images taken just above the aggregate surface (surface slices) were obtained and sent to a color graphics workstation via an Ethernet connection.

## Results of interfacial zone simulations

Figure 2 shows the two-dimensional images obtained for the neat paste concrete, with no mineral admixtures, both before any hydration had occurred, and after 77 percent of the cement had been consumed. The colors are the same as in Figure 1, with the aggregate being the magenta square. The surface slices on the left side of Figure 2 clearly show the interfacial zone to be a region of high porosity, high CH content, low $C_3S$ content, and low C-S-H content. This phenomenon is due both to the particle packing effect mentioned earlier, and a one-sided growth effect. Near the aggregate, the pore space is being filled with hydration products growing in from one side only, unlike the bulk paste where hydration products are growing inward from all directions. With Figure 2 as our baseline, we can proceed to assess the effects of the various mineral admixtures on interfacial zone microstructure.

Because the original arrangement of the cement and mineral admixtures near the aggregate surface influences later microstructural development, the small one-pixel and large (same size as cement) mineral admixtures will differ in their effectiveness in enhancing interfacial zone microstructure. Figure 3 shows the initial particle packings around the aggregate for the systems containing small and large tan mineral admixture (filler) particles at 20 percent mass replacement for cement. The ability of the smaller admixture particles to pack more closely to the aggregate surface is clearly evident in the surface slice images on the left side of Figure 3. The mineral admixture content present in the interfacial zone also is important due to its pozzolanic reaction with the CH forming in this region. This can be seen in Figure 4, which shows

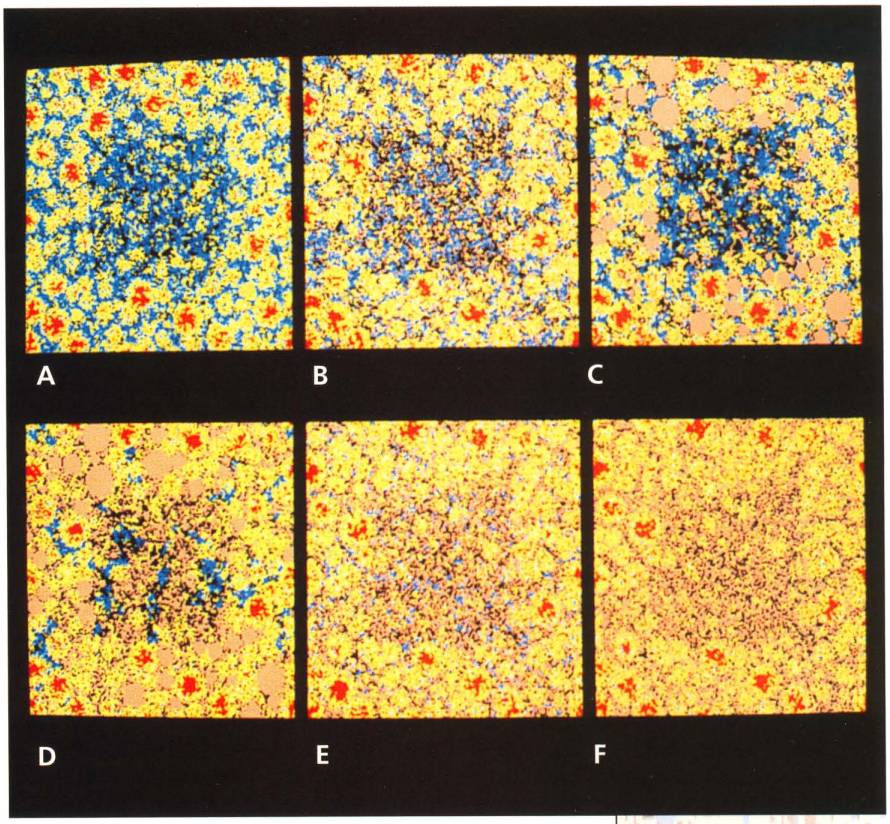


Figure 4. Surface slices for each of the systems listed in Table 1 after 77 percent of the cement has been consumed by hydration.

the final 77 percent hydrated surface slices for each of the six systems listed in Table 1. In Figure 4, the mineral admixtures and pozzolanic C-S-H which they produce are both tan in color. The amount of blue CH remaining in the interfacial zone varies greatly for the different mineral admixtures.

Figure 4 shows that the different mineral admixtures modify the interfacial zone microstructure to varying degrees. The 20 percent small silica fume (System F) is especially effective. In this system, the interfacial zone microstructure is nearly as dense as that of the bulk paste, and the pores remaining in the interfacial zone are much smaller than those found in the interfacial zone of the neat paste concrete. Quantitative results for these six

Table 1. Systems investigated in interfacial zone study.

| System | Mineral admixture size | Mineral admixture reactivity | Physical system |
|---|---|---|---|
| A | — | — | Neat paste concrete |
| B | 1 pixel | 0.47 | Fine fly ash |
| C | Same as cement (11 and 21 pixels) | 0.47 | Fly ash |
| D | Same as cement (11 and 21 pixels) | 2.08 | Agglomerated silica fume |
| E | 1 pixel | 1.04 | Reactive fine fly ash |
| F | 1 pixel | 2.08 | Condensed silica fume |

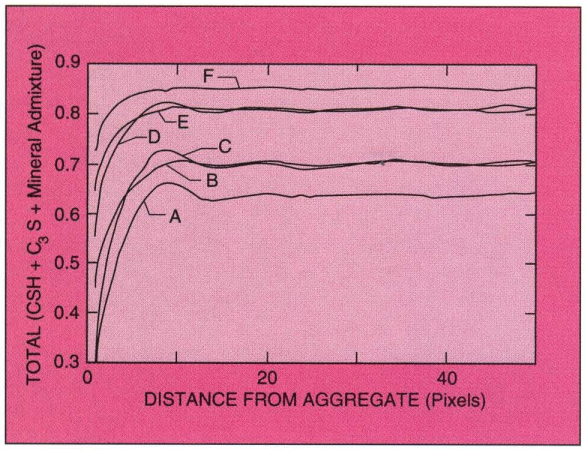


**Figure 5. Total load-bearing phases ($C_3S$ + C-S-H + mineral admixture) versus distance from aggregate surface for each of the systems listed in Table 1.**

systems are given in Figure 5, which plots the total of the ($C_3S$ + C-S-H + mineral admixture) phase fractions as a function of distance from the aggregate surface. These three phases have been combined on the generally accepted assumption that they are the major load-bearing phases in the cement paste matrix. Here again, System F, the small silica fume, is seen to be superior to all other systems as the distribution of this combined phase is nearly homogeneous throughout the concrete. However, if the silica fume tends to agglomerate, as in System D, the interfacial zone microstructure will be significantly less dense than that of the bulk paste. Fly ash is seen to be inferior to silica fume in performance, although System E, based on small reactive fly ash, is slightly more effective than the agglomerated silica fume system. Experimental evidence in support of these results has been reported based on the separation of fly ash into size classifications and the subsequent observation that the finest fly ash provided the greatest strength enhancement in concrete.[5]

## Discussion

The digital-image-based microstructural model has provided insights into the causes of the microstructure formed in the interfacial zones in concrete and the improvements in density and uniformity that may be achieved using mineral admixtures. Also, the predictions of the model have been shown to agree well with measurements carried out on real concrete specimens.[6]

Because digital-image-based models are stringent in both their memory and CPU time requirements, supercomputers will continue to be needed in this area of research. After this study was completed, NIST purchased a CRAY Y-MP computer system. Over the past year, the authors and their colleagues have used this system extensively for simulating cement paste microstructures and determining transport properties such as diffusion coefficients, fluid permeabilities, and elastic moduli. In general, the microstructure development codes are mostly not vectorizable, but the transport property computation codes are highly vectorizable, consisting primarily of different kinds of relaxation routines. Furthermore, this digital-image-based approach to random materials has been adapted to simulating the sintering of ceramic powders. An interfacial zone problem, analogous to that found in concrete, arises in the manufacture of fiber-reinforced ceramic composites, due to the inefficient packing of ceramic particles around the fibers.

Digital-image microstructural models enable the materials modeler to:

- ☐ Deal with random shapes and structures
- ☐ Realistically simulate physical processes such as dissolution, diffusion, and surface reaction
- ☐ Use finite difference techniques to compute almost any property of interest due to the underlying lattice structure of the image

As computer technology continues to evolve, the simulation of larger systems with higher resolution images will become possible, allowing microstructural models of materials to come even closer to their physical counterparts. ▬


### *Acknowledgments*

*The authors thank the National Science Foundation Science and Technology Center for Advanced Cement-Based Materials for partial financial support and Cray Research for donating computer time on a CRAY-2 supercomputer.*


### *About the authors*


*Dale P. Bentz is a chemical engineer in the Building Materials Division at the National Institute of Standards and Technology. His research interest currently is computer modeling of the microstructure and performance properties of building materials including cement and concrete.*

*Edward J. Garboczi is a physicist in the Building Materials Division at the National Institute of Standards and Technology and a participating member of the National Science Foundation Science and Technology Center for Advanced Cement-Based Materials. His current research interests include quantitatively relating microstructure and physical properties of cement-based materials by using the percolation theory and fundamental computer simulations.*


### *References*


1. Bentz, D. P. and E. J. Garboczi, "Percolation of Phases in a Three-Dimensional Cement Paste Microstructural Model," *Cement and Concrete Research*, Vol. 21, p. 325, 1991.
2. Garboczi, E. J. and D. P. Bentz, "Fundamental Computer Simulation Models for Cement-Based Materials," *Materials Science of Concrete II*, American Ceramic Society, pp. 249-278, 1991.
3. Schaefer, D. W., "Fractal Models and the Structure of Materials," *Materials Research Society Bulletin*, Vol. 13, No. 2, pp. 22-27, 1988.
4. Garboczi, E. J. and D. P. Bentz, "Digital Simulation of the Aggregate-Cement Paste Interfacial Zone in Concrete," *Journal of Materials Research*, Vol. 6, No.1, pp. 196-201, 1991.
5. Berry, E. E., R. T. Hemmings, W. S. Langley, and G. G. Carrette, "Beneficiated Fly Ash: Hydration, Microstructure, and Strength Development in Portland Cement Systems," *Fly Ash, Silica Fume, and Natural Pozzolans in Concrete*, Vol. 1, SP-114, American Concrete Institute, Detroit, pp. 241-273, 1989.
6. Bentz, D. P., E. J. Garboczi, and P. E. Stutzman, "Computer Modelling of the Interfacial Zone in Concrete," *Proceedings of the RILEM Conference on Interfaces in Cementitious Composites*, France, 1992.

# Supercomputer applications in livestock breeding

*Julius van der Werf, Wageningen Agricultural University, The Netherlands*
*Jaap Hollenberg, SARA, Amsterdam, The Netherlands*

Animal breeding has become a sophisticated business run by academically trained specialists. These specialists combine statistical techniques and quantitative genetics to rank and select the genetically superior animals on an international basis. Dairy cattle breeding in particular has developed into a multimillion dollar international business.

American and European bulls, for example, are not only used for breeding on a national basis, but their offspring also multiply rapidly in other countries through the export of sperm or of sons as breeding bulls. The genetic harvest of specific bulls, such as the current Dutch top bull "Sunny Boy," is approaching a value of $25 million. Considering that his progeny includes several hundreds of thousands of daughters—or one fifth of the entire Dutch dairy population—the genetic and economic impact of top bulls like Sunny Boy is substantial. Errors in selecting bulls for breeding therefore can have enormous economic consequences internationally.

To optimize international breeding programs, scientists at the Wageningen Agricultural

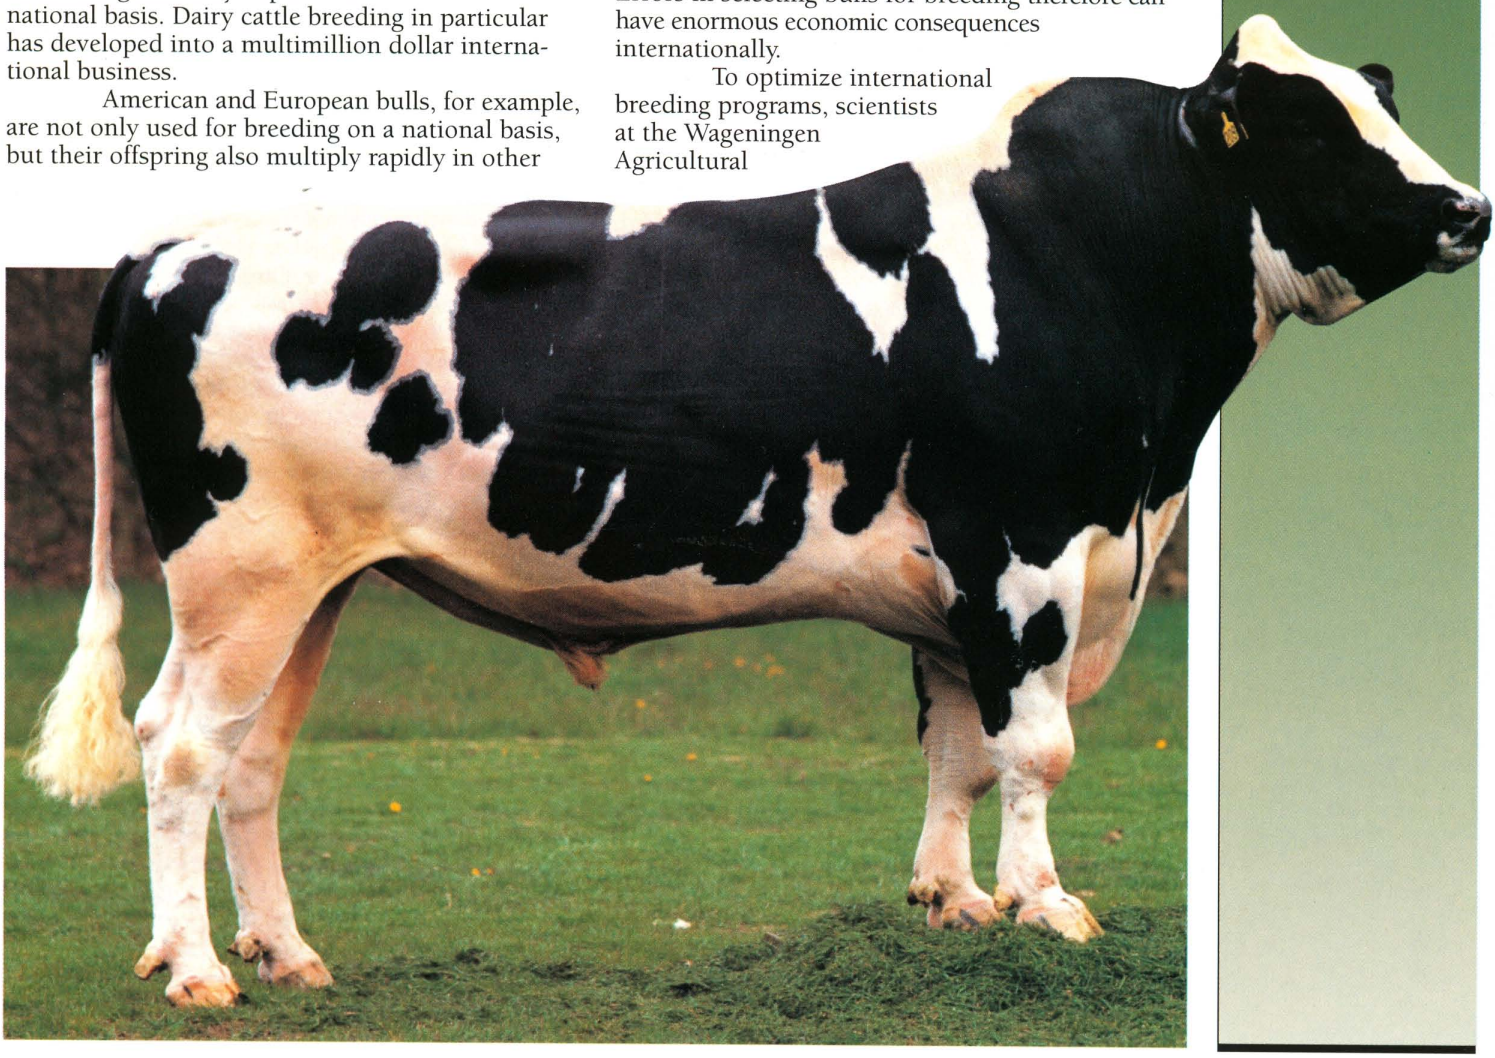

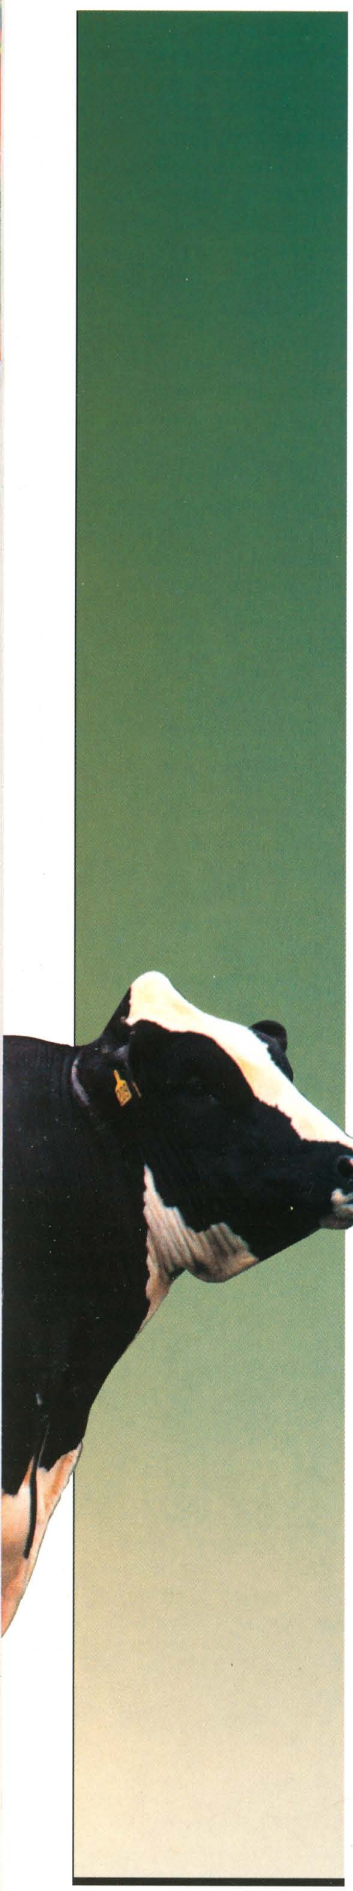

University in The Netherlands are developing computational models for animal breeding, using the CRAY Y-MP4/464 supercomputer at SARA (Stichting Academisch Rekencentrum Amsterdam), the joint computing center of the Amsterdam universities and the Dutch national supercomputer center. This new breed of supercomputer application is fast becoming the tool of choice for international researchers in the animal breeding industry.

## Evaluation of breeding values

Genetic evaluation, including breeding value and genetic parameters such as heritabilities and genetic variation, provides the basis for ranking and selecting dairy animals. This evaluation is based on phenotypic observations about the animal and its family members. A statistical model, which is a mixed linear model, describes the observed performance as a function of systematic environmental effects, herds and season of production, and the effect of the performing animal's potential plus a random residual error term. Animal breeding values are assumed to be randomly distributed genetic effects, with a variance equal to the additive genetic variance of the population under study. Variances and heritabilities are assumed known. The breeding values are solved from the so-called mixed model equations, a system very similar to least-squares equations but with a matrix of additive genetic relationships between all animals added to the diagonal block of animal coefficients. Such breeding values are called best linear unbiased predictions.

Previous systems of equations included the effect of sires, and the total system of equations was equal only to the number of sires used in recent years in a particular area. Most evaluations were national, and the number of animals evaluated was limited to between 10,000 and 100,000. The availability of advanced computer technology has made it possible to use an animal model that includes an equation for each animal, which accounts for its genetic relationship with parents, siblings, and offspring. This extended model is much more accurate in accounting for selection and nonrandom mating of sires and dams. The effects of selection and nonrandom mating would particularly bias the genetic evaluations of top-breeding animals.

In The Netherlands, genetic evaluation of the dairy population is performed twice a year using the animal model. Presently, 5.5 million animals are evaluated, and 11 million equations are solved, requiring the computational power of the CRAY Y-MP4/464 supercomputer.

The system of equations is represented as $Ms = r$, where $M$ is a large coefficient matrix, $s$ is the solution vector, and $r$ is the vector with observations accumulated for each effect. An iterative scheme is used for solving; hence

$$M^*s_{k+1} = r\text{-}(M\text{-}M^*)s_k$$

and certain algorithms coincide with the appropriate choice of approximation matrix $M^*$. For example, for a second order Jacobi iteration, $M^*$ is equal to the diagonal of $M$, and for a Gauss-Seidel iteration, $M^*$ is equal to the lower triangular of $M$. Using a method in which not all coefficients would be stored (thereby saving memory) proved to be a particularly useful strategy. Here only solutions $s$, diagonals $d$, and accumulated corrections on $r$ would be stored while reading through the data and making those corrections. Even this method requires about 72 Mbytes of storage, and on most mainframes the process would be constrained by time-consuming I/O. A supercomputer, such as the Cray Research system at SARA, which can easily provide a working memory of about 250 Mbytes (32 Mwords), is able to store not only the working three vectors, but also all pedigree and performance data.

The advantage of using a supercomputer instead of a mainframe computer is not just processing speed. The supercomputer's larger memory provides additional time savings for solving large systems of equations. Compared with a typical mainframe computer, the supercomputer achieves an estimated 20-fold speed gain, and the gain from reducing I/O is at least as great. It is estimated that the genetic evaluation model for Dutch dairy cows will take less than half a CPU hour on the Cray Research system compared to 100 hours on a mainframe.

## Estimation of genetic parameters

The genetic parameters estimated most often in animal breeding research are the heritability of metric traits, genetic variation, and correlations between traits. The research method generally accepted as most appropriate is the restricted maximum likelihood (REML) procedure. The model is the same as the one used to estimate breeding values, except that genetic parameters have to be estimated rather than assumed. A very attractive feature of REML is its flexibility. It can handle a model with all additive genetic relationships between all animals in a population. Such a model accounts for all selection that has been carried out over generations and basically yields an estimate of genetic variance before selection began. It should be noted that most standard packages for variance component estimation do not handle such models with genetic relationships.

Finding the maximum likelihood of a linear animal model for location and dispersion parameters could in principle be done by several maximization routines, if the data is available. An expectation maximizing algorithm often is used on first derivatives of the likelihood function, and a Fisher's Scoring algorithm would be able to use expectations of second derivatives. However, expressions from such derivatives involve traces of the inverse of the coefficient matrix, which can become very large when an animal model is analyzed. A derivative free algorithm recently has been proposed as an alternative. An array of the coefficient matrix of the mixed model equations is needed, augmented by the right-hand sides of the equations and the total sums of squares of the data. The expressions necessary to evaluate a likelihood can be obtained by a Gaussian elimination on this array, which is a matrix of order equal to the number of animals in the analysis. Special

sparse matrix techniques are particularly useful for this work.

The size of the dataset is limited by the size of the array that can be handled. For example, current examples could use data on 10,000 animals, which implies one million coefficients for each percent of fill of the coefficient matrix. The maximum of the likelihood is obtained in an iterative process in which the likelihood has to be evaluated many times. Obviously, computational speed and total memory requirements are the limiting factors for the size of the datasets to be analyzed. In addition, more complex models could be used, such as multivariate models with observations on several traits, or models accounting not only for additive genetic effects but also for dominant genetic effects or maternal effects. In current applications, the supercomputer has a performance rate 30 times greater than that of a mainframe computer; however, improved programming would increase efficiency further. In addition, large linear equation systems are well suited for parallel processing.

## Animal genotyping

The techniques described previously assume a genetic model in which each trait is determined by a large number of unlinked loci. Individual loci are not distinguished in the model, but their accumulated effect is assumed to behave as a normally distributed variable. There are some examples of quantitative traits that are determined by the large effect of a single locus, for example certain fertility genes in sheep and possibly in pigs. Therefore, models are introduced that have such single gene effects. Another important reason to introduce such models is the rapid development of molecular techniques and the description of specific gene locations. Quantitative models including such "marker genes" could determine their association with important characteristics. In fact, similar models also have been used to analyze medical data in human genetics. Datasets on animals are larger, however, because of their more complicated family structures.

Effects and frequencies of single genes are determined by maximum likelihood techniques. The likelihood of an individual's genotype given its record can be determined. If an individual has three possible genotypes (MM, Mn and nn, respectively), then $m$ independent individuals have $3m$ possible combinations of genotypes. This quickly adds up, because for $m$ = 1000, and for one picosecond for a floating-point operation, one would need many millennia to evaluate a likelihood. Fortunately, the genetic structure of the data often allows a large reduction of possible combinations. However, models of this type are computationally prohibitive unless high-performance computing power is available.

## Stochastic simulation

Two types of simulation studies are used in animal breeding research. In one, properties of certain estimation methods are tested with simulated data. This has proven particularly useful when studying the effect of analyzing selected data. In the other, stochastic simulation has been used for studies related to the optimization of breeding programs. A breeding program aims for a certain maximum genetic change of economically important traits in a population and evaluates annual genetic progress, inbreeding, and variance in response. A breeding program typically must determine the number of animals to be selected, the maximum number of offspring (inbreeding) they should produce, the appropriate time to be selected (generation interval), and the information required to select them.

Deterministic simulations generally are quite difficult because the program must account for complex family structures, overlapping generations, effects of order statistics, and decrease of genetic variation, for example. Because the system studied is stochastic by nature, and many interdependencies arise due to the family structures in a population, stochastic simulation studies have proven very useful. In a study to determine the effect of heterogeneous variances on selection efficiency, we simulated a dairy cattle breeding program for a 25-year period with 10,000 animals annually, for which 20 replicates of each of eight alternatives were required. Such a study would take about 160 CPU hours on a mainframe, while the Cray Research supercomputer could do the job in about 13 hours. The increased computing power allows a comparison of a greater number of alternatives and of more replicates in such studies.

## Summary

Animal breeding researchers have been aware for some time that quantitative genetic research in animal breeding applications requires the large memory and computing power of a supercomputer. The breeding applications developed at SARA are in the early development stage but show great potential for increased accuracy and sophistication as expanded computing capacity becomes available. Because of the large economic stakes involved in research areas such as animal breeding, the growth of the Dutch national supercomputing capacity at SARA is strongly supported by government and the Dutch scientific community. ■

***Acknowledgment***

*The authors thank the Dutch National Computing Facilities Foundation (NCF) for its technical and financial support.*

***About the authors***

*Julius van der Werf graduated from Wageningen Agricultural University in 1984 with a degree in animal science and agricultural economics. He received a Ph.D. degree in Wageningen in 1990 for his research on estimating genetic parameters in dairy cattle populations. Currently, van der Werf is assistant professor at the Department of Animal Breeding at Wageningen Agricultural University, The Netherlands.*

*Jaap Hollenberg graduated with a mathematics degree from the University of Amsterdam. Currently, he heads a group at SARA supporting supercomputer users at Dutch universities.*

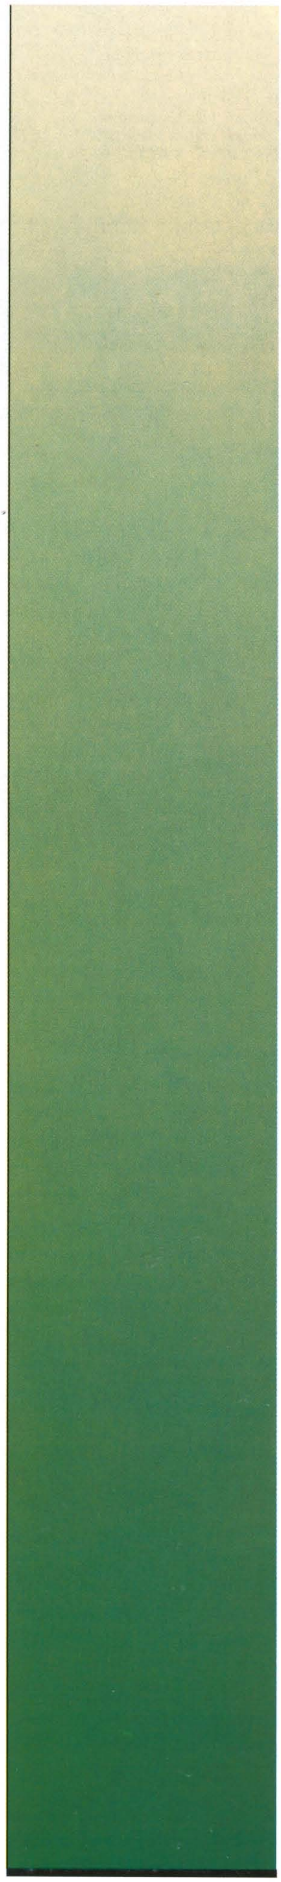

# Supercomputers power discoveries in nursing

*Sharon L. Meintz*
*University of Nevada Las Vegas*

***A new supercomputing application is emerging*** at the Department of Nursing, University of Nevada Las Vegas (UNLV). For the first time, nursing researchers have access to a Cray Research supercomputer to analyze large datasets of nursing and health data collected by government and nongovernment agencies. Although some of these datasets have been analyzed in isolation, they never have been studied broadly enough to identify and predict health trends in human populations—predictions necessary to help nurses provide better preventative and prescriptive health care.

The driving force behind this supercomputer application is the Project for Nursing and Health Data Research (PNHDR), which was created to evaluate the supercomputing resources needed to analyze databases in the nursing and health professions, and to identify any barriers to this application. The project is supported by Cray Research, the National Supercomputer Center for Energy and the Environment, and the National Center for Health Statistics National Data Tapes Program.

## The supercomputer requirement

Most nursing researchers use personal computers for word processing and research analysis. Unless these researchers are especially motivated, however, their computers are not networked to mainframes able to analyze datasets large enough to reveal health trends.

In May 1991, the PNHDR received access to massive nursing and health datasets from government and nongovernment sources. In many cases, this data had been archived without analysis. The computer requirements necessary to process the data are a major obstacle to analyzing such large datasets. For example, the quality and effectiveness data collected for Medicare in just one year documents 10 million cases equalling 15 reels of data, or 1800 megabytes. To complete a five-year comparative or trend analysis, computing capability for 75 reels, or 9000 megabytes, of data is required. Analyzing trends over the length of an average human life would require computing resources of several orders of magnitude greater, and many agencies possessing this data do not have access to the necessary computing resources.

After PNHDR received access to these datasets, researchers used a statistical software package, SPSS-X by SPSS Inc. SPSS-X was ported to a CRAY Y-MP2/216 supercomputer at the National Supercomputer Center for Energy and Environment at UNLV to study selected datasets. Researchers used a multilevel approach to assess the relationships of data in various formats (numeric, text, auditory, and image) to functions (generating,

processing, storing, and transmitting). Although the initial tests indicated that a supercomputer application was required, several barriers to this solution also became evident.

## Barriers to supercomputing

Most nursing researchers need new skills to use supercomputing applications, even with the assistance of engineers and computer scientists. Additionally, many of the available statistical software packages for supercomputers require programming, a skill not currently taught in nursing curricula, and very few nurse scientists understand the components or architecture of supercomputers. Even though they may be skilled in the use of a personal computer and its role in secondary data analysis on small datasets, they need further knowledge to analyze large datasets with supercomputers.

Version 2.0 of the SPSS-X software was chosen because most nurses are familiar with it. However, it does not perform trend analysis, an essential component for nursing and health data analysis. Development of SPSS-X variables for rectangular datasets with between 200 and 1000 variables, over 466 columns and more than 25,000 cases is time-consuming, but it is necessary because these variables currently are not available in a ready-to-use format. Lack of a comprehensive social science statistical analysis package for supercomputers is a significant barrier to the use of supercomputing for nursing and health data analysis.

Additionally, because descriptive statistics are affected by large sample sizes, researchers must establish a percentage index to indicate the sample variability and the number of observations on which the mean is computed.

The supercomputer was able to run large data analyses in seconds or minutes instead of the hours or days it would take using lesser computing resources; however, the hardcopy output of even a simple analysis required about three reams of paper. Because nursing researchers at UNLV have limited access to workstations, they share data with others by transferring hard copy or using campus network systems even though broader network access is available. As this research and development project proceeds, nursing researchers must acquire the same computer resources used by scientists and engineers, and they must become well connected to important networks. Scientific visualization must be among these resources.

## Conclusions

Although there are many barriers to analyzing the massive datasets with supercomputers, it is the only current resource that can support this nursing and health application. The PNHDR's goal is to address the outstanding issues, making the supercomputing solution easily accessible to nursing researchers. As identified by current PNHDR research, ease of access will not be possible until:

- ☐ A new statistical software program is created to analyze trends in massive datasets
- ☐ All interfaces required by nursing researchers are transparent and include a visual postprocessing package
- ☐ SPSS-X variable lists for large datasets are developed
- ☐ Nursing researchers are teamed with the traditional users of supercomputers—scientists, engineers, and computer scientists

When all of these components are in place, the massive amounts of health data that have been stored will be available to researchers to identify and predict health trends. Armed with this knowledge, nurses and all health professionals can better prepare for future patient needs.

## Current research

The PHNDR has identified about 250 gigabytes of health data that will be available to study within the year. To study this data, a new statistical methodology has been proposed: Global Analysis Tera Exploratory Statistics (GATES). This methodology should aid deductive reasoning and enable researchers to reach conclusions regarding specific cases from the study of large databases. In addition, preliminary work has begun on a specialized user interface called VuSTAT. VuSTAT is expected to be similar to the interface of Cray Research's UniChem visualization product, allowing easy interaction with health and nursing data.

## Future directions

The prospect of supercomputing in the nursing profession has generated three new conceptual branches of nursing science:

- ☐ Nurmetrics, which applies statistical techniques to the testing, estimation, and quantification of nursing theories and solutions to nursing problems
- ☐ Computational nursing, which uses mathematical/computer modeling and simulation systems to apply existing theory and numerical methods to new solutions for nursing problems or the development of new computational methods
- ☐ Nursing informatics, which applies the principles of computer science and informatics to understanding the relationship between information form and function for solving problems in nursing administration, education, practice, and research

These new directions in nursing will form the foundation for new discoveries powered by supercomputing. ▬


***About the author***

*Sharon L. Meintz is an associate professor for the Department of Nursing, College of Health Sciences at the University of Nevada Las Vegas and is the director of the Project for Nursing and Health Data Research. She holds M.S. degrees in nursing administration from the University of Colorado and in health education from Memphis State University. She received a Ph.D. degree in education from the University of Denver.*

# Serial and parallel performance on large matrix systems

*Rachid Ababou, Southwest Research Institute, San Antonio, Texas*
*Lynn W. Gelhar, Massachusetts Institute of Technology, Cambridge, Massachusetts*
*Chris Hempel, Cray Research, Inc.*

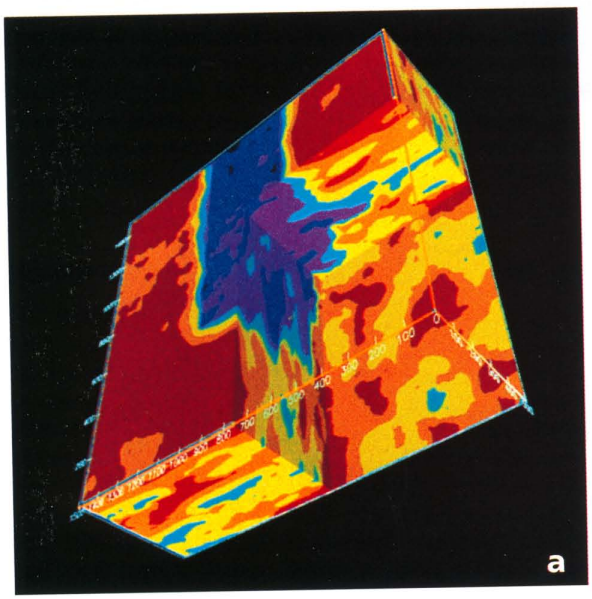


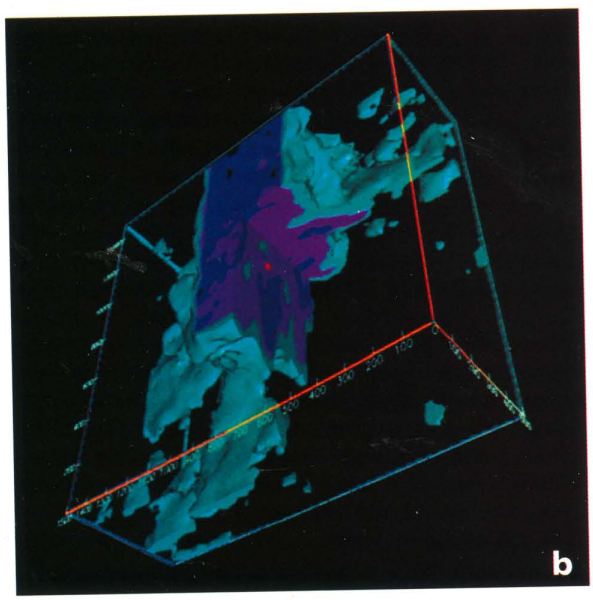


Figure 1. Three-dimensional moisture plume infiltrating from a strip source in a randomly heterogeneous soil (300,000 nodes). Colors represent moisture contents, with purple-blue hues for "wet" and yellow-brown hues for "dry."

Subsurface flow processes are inherently three-dimensional and heterogeneous over many scales.[1,2] Although these characteristics must be taken into account to model fluid flow and contaminant transport accurately, doing so puts heavy constraints on numerical models.

To investigate computational issues associated with flow modeling, we used the general-purpose numerical code BIGFLOW to model in detail large three-dimensional flow systems in unsaturated, partially saturated, and saturated heterogeneous geologic media. The code, based on an implicit finite difference discretization, initially was developed at Massachusetts Institute of Technology (MIT) as a research tool to investigate stochastic flow processes.[1] It is being used to investigate a high-level nuclear waste geologic repository for the U.S. Nuclear Regulatory Commission.[2]

Given the large-scale flow systems to be modeled here, the matrix solver constitutes essentially the "computational kernel" of the code. For linear flow problems in particular, heterogeneous systems on the order of several millions of equations have been solved in moderate times on CRAY-2 and CRAY Y-MP supercomputers using sparse iterative solvers. We focus here on the computational performance of the code for such large problems, based on a highly vectorized, sparse, preconditioned conjugate gradient solver. Computational performance is assessed through analyses of convergence rates, CPU times, and speed-ups due to coarse-grained parallelization.

## BIGFLOW code and three-dimensional applications

The BIGFLOW code solves linear and nonlinear porous media flow equations based on Darcy's law, appropriately generalized to account for three-dimensional, and possibly random, heterogeneity. The code, written in ANSI Standard Fortran 77, is free of machine-dependent directives and is portable without modification to a variety of computer platforms. An implicit finite difference scheme is used for discretization. Optionally, a modified Picard scheme is used to linearize unsaturated flow equations (outer iterations). A preconditioned iterative method is used to solve the resulting symmetric matrix systems (inner iterations).

Iterative matrix solvers used extensively thus far include the Strongly Implicit Procedure (SIP)[3] and the Diagonally Scaled Conjugate Gradient (DSCG) investigated below. (Note that symmetric diagonal scaling is used.) The solution modules were coded especially to take advantage of sparsity and symmetry of the finite difference systems. An auxiliary data processor was developed for interactive entry and analysis of three-dimensional numerical datasets.

Figure 1 shows three-dimensional views of a transient moisture plume during strip-source infiltration (10 days). The plume is viewed from above, first cut away (Figure 1a) and then peeled off (Figure 1b). Colors represent moisture contents, with purple-blue hues for "wet" and yellow-brown hues for "dry." The nonlinear unsaturated soil properties, discretized on 300,000 grid points, were assumed to be random fields with statistically

anisotropic structure, and the flow conditions were similar to an experiment undertaken by P. J. Wierenga and coworkers at the Las Cruces Trench Site.[2] Figure 1 reveals a significant lateral spreading of moisture through marginally wet "fingers" perched over drier zones and confirms theoretical findings on moisture-dependent anisotropy in stratified soils. This particular simulation was conducted on a CRAY-2 supercomputer using the SIP solver.[1]

Figure 2 shows three-dimensional views of hydraulic equipotential surfaces in the case of steady flow in a randomly heterogeneous, isotropic, saturated porous medium. The grid size for this problem was $101^3$ or 1 million nodes. A number of similar problems have been solved with BIGFLOW using either the SIP solver or, more recently, the DSCG solver. The numerical solutions have been exploited to study effective conductivity of random heterogeneous media,[3] to simulate stochastic contaminant transport, and to analyze resulting concentration fields.[4]

## Theoretical estimates of performance

For a wide class of iterative solvers that includes CG (Conjugate Gradient) and DSCG, the number of iterations required to decrease the error by six orders of magnitude, for example, is known to be proportional approximately to the square root of the condition number of the coefficient matrix. In the case of implicit seven-point-centered finite difference systems, the condition number for steady flow is typically $O(n^2)$, where $n$ represents the unidirectional size of the grid along its largest side.[1] For each iteration, the computational work or number of operations is proportional to $N$, the multidimensional number of nodes. Multiplying by the estimated number of iterations yields a total work on the order $O(N^p)$ with exponent $p = 4/3$ for a three-dimensional cubic grid; $p = 3/2$ for a two-dimensional square grid and; $p = 2$ for a one-dimensional grid.

These "order of magnitude" estimates indicate the relation between computer time and problem size. However, they have several shortcomings. First, the convergence rate estimate does not indicate the influence of conductivity heterogeneity and spatial structure. Second, it is only a worst case estimate obtained from an approximate error upper bound. Third, this estimate must break down as the number of iterations approaches the number of equations ($N$), since the CG method gives the exact solution in no more than $N$ iterations (within machine precision). Finally, the assumption that computational work per iteration is proportional to grid size ($N$) does not take into account possible nonproportional speed-ups due to vector and parallel processing.

## Observed convergence rates

Figure 3 shows the number of DSCG iterations versus unidirectional grid size for the case of constant conductivity. In this special case, diagonal scaling has no effect and the DSCG solver is equivalent to the straight CG solver. The numerical grids used in this series of tests were cubic lattices

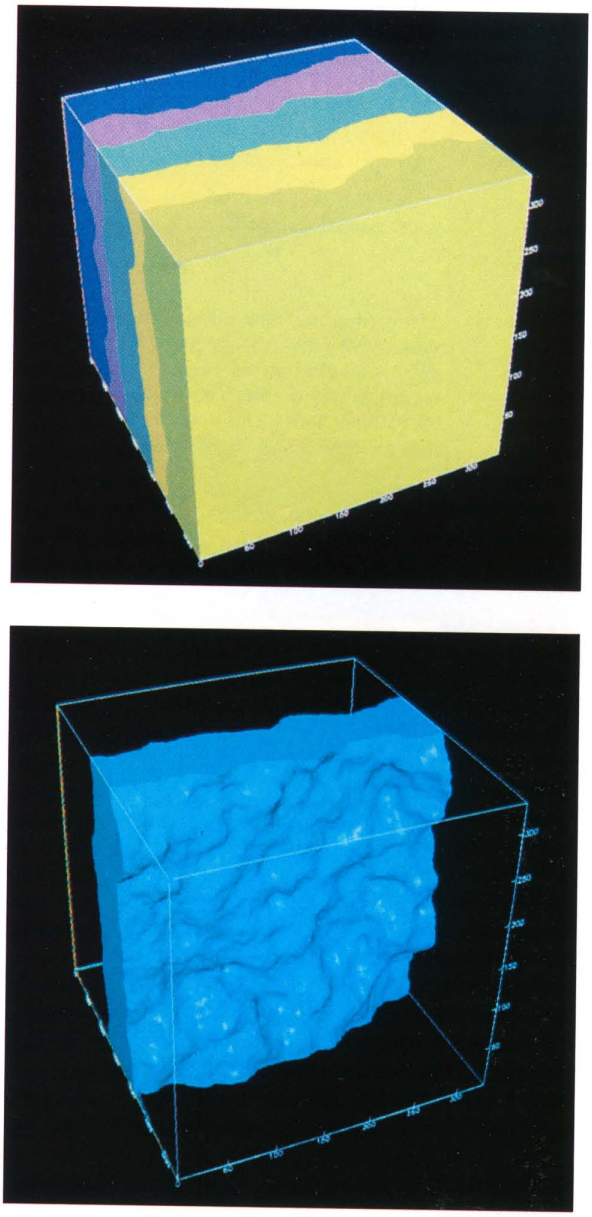

Figure 2. Three-dimensional hydraulic equipotential surfaces for groundwater flow in a randomly heterogeneous, isotropic, saturated porous medium (1 million nodes). Full view (top) and cut-off view of one equipotential slice (bottom).

ranging from 83 to 1283 nodes with the largest grid having over two million nodes. The number of iterations ($I$) was defined as that required to decrease the $L_\infty$-norm of error by six orders of magnitude. The approximately linear increase of $I(n)$, with respect to unidirectional grid size, is in agreement with theory. These tests were implemented on a CRAY Y-MP system with help from G. Wittmeyer.

The convergence behavior of the DSCG solver is, however, more complicated than implied by Figure 3 alone. Figure 4 shows several curves depicting the $L_\infty$-norm of error versus iteration count for test problems with different degrees of heterogeneity, grid sizes, and log-conductivity structures (see Table 1). Grid sizes range from a few thousand nodes up to 7.6 million nodes. The degree of heterogeneity is represented by $\sigma$, the standard deviation of log-conductivity ($\ln K$).

Figure 4 shows the singular behavior exhibited by curve #600B. Initially, convergence is slow, as expected, due to the large unidirectional size of the grid. However, after over 700 iterations, the error drops quickly to machine precision and

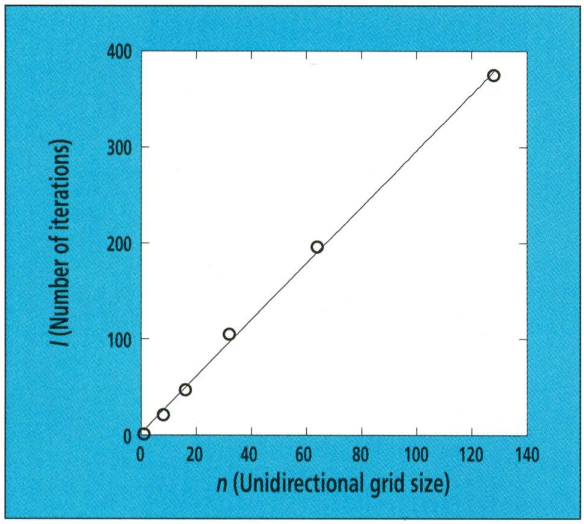


Figure 3. Number of DSCG iterations ($I$) versus unidirectional grid size ($n$) in the case of constant conductivity ($\sigma = 0$).

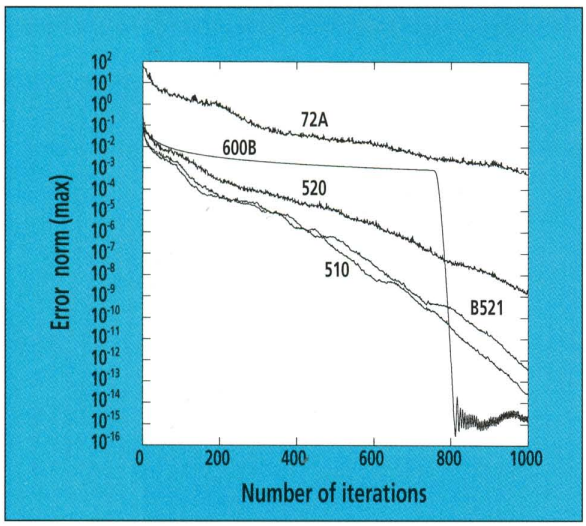


Figure 4. Convergence of DSCG iterations ($L_\infty$-norm of error vs. iteration count) for the test problems described in Table 1.

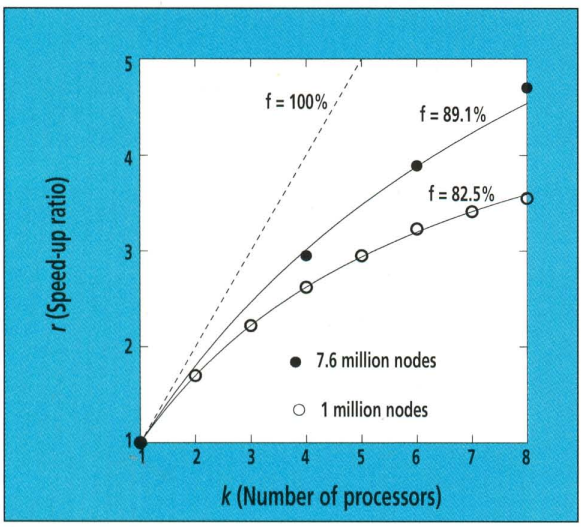


Figure 5. Speed-up curves: parallel/serial speed-up ratio ($r$) vs. number of CRAY Y-MP8 processors running concurrently ($k$).

cannot decrease further. This behavior is in agreement with the fact that the CG method always yields the exact solution, within machine precision, after $N$ iterations at most ($N$ is not very large for this problem). On the other hand, comparing curves labeled #510 ($\sigma = 1$) and #520 ($\sigma = \sqrt{3}$) indicates the influence of degree of heterogeneity (slower convergence). Comparing curves #520 (binary distribution) and #B521 (Gaussian distribution) demonstrates the equally important influence of spatial structure. Comparing curves #520 (1 million nodes) and #72A (7.6 million nodes) shows the influence of grid size (slower convergence).

## Serial and parallel timings

The performance of the DSCG-based code, expressed in CPU seconds, depends on grid size and number of iterations, and on machine-dependent additive and multiplicative factors. Timings can be expressed approximately in the form: $T(I, N) = (a\,I + b)\,N$; where $T$ is the total CPU time (seconds), $a$ represents specific iterative work (seconds/iteration/million nodes), and $b$ represents work spent outside the iterative solution process, or overhead (seconds/million nodes). As before, $I$ is the number of iterations, and $N$ is the number of nodes in multidimensional space. Note that $I$ may be a preselected number of iterations, or alternatively, the number of iterations to decrease the error by a certain amount (say, six orders of magnitude). In view of the theoretical estimate given earlier and confirmed experimentally, the number of iterations is proportional to $N^{1/3}$ for cubic grids.

## Serial timings

For the DSCG-based code running serially on CRAY-2 and CRAY Y-MP machines, we found empirically:

$a_{(\text{serial CRAY-2})} = 0.48$ seconds per iteration per million nodes

$a_{(\text{serial CRAY Y-MP})} \approx 0.20$ seconds per iteration per million nodes

These constants were obtained from timings of several large test problems with randomly heterogeneous conductivities, most involving cubic grids. The dependency analyzer and optimizer fpp, a feature of the CF77 Fortran compiling system, was used on both systems; the "aggressive" optimization option was used on the CRAY Y-MP system. Note that the CRAY Y-MP system is faster than the CRAY-2 system by a factor of about 2.5 for these types of problems (in serial mode).

The serial CRAY-2 system timings were analyzed in detail using the flowtrace utility included in the UNICOS operating system. It was found that the dependency analyzer decreased $a$ by just a few percent. That is, almost all inner loops vectorized with or without fpp. The overhead constant $b$ was found to be sensitive to I/Os; with unformatted I/Os, this constant was found to be:

$b_{(\text{serial CRAY-2})} = 28$ seconds per million nodes, compared to 116 seconds per million nodes with formatted I/Os.

## Parallel timings

Coarse-grained parallelization was studied by allowing the BIGFLOW code to run concurrently on $k$ processors of the CRAY Y-MP8 system in dedicated mode. Again, the DSCG solver was used to solve random conductivity problems involving one to several million grid points. The BIGFLOW source code was not modified for multiprocessing. Instead,

| Test (#) | Ln(K) distribution | Grid size $N = n_1 \times n_2 \times n_3$ |
|---|---|---|
| 600B | Constant $\sigma = 0$ | $N = 1001 \times 5 \times 5$ |
| 510 | Gaussian isotropic $\sigma = 1$ | $N = 101 \times 101 \times 101$ |
| 520 | Gaussian isotropic $\sigma = \sqrt{3}$ | $N = 101 \times 101 \times 101$ |
| 521 | Binary isotropic $\sigma = \sqrt{3}$ | $N = 101 \times 101 \times 101$ |
| 72A | Gaussian anisotropic $\sigma = \sqrt{3}$ | $N = 178 \times 120 \times 357$ |

Table 1. Summary of test problems

we let the Autotasking feature of the CF77 Fortran compiling system perform the necessary code modifications and enhancements (we also used a compiler option to in-line the CG solver module). Estimates of speed-ups ($r$) and of parallelizable fractions of code ($f$) were obtained by comparing cumulated CPU times to wallclock times and by applying Amdahl's law.

Figure 5 depicts two speed-up curves $r(k)$ obtained for a 1 million node test problem (lower curve) and for a 7.6 million node test problem (upper curve). These test problems correspond to #510 and #72A in Table 1. Figure 5 shows both actual speed-ups (circles), and analytical curves $r(k)$ (solid lines). These analytical curves were obtained after evaluation of the parallelizable fraction $f$ from Amdahl's law. The dashed straight line represents the ideal case $f = 100$ percent, corresponding to a fully parallelizable code.

Using the Autotasking capability and timing utilities, we found $f = 82.5$ percent for the 1 million node problem (#510), and $f = 89.1$ percent for the 7.6 million node problem (#72A). The corresponding speed-up ratios for eight processors are 3.59 and 4.53 respectively. The sensitivity of speed-up curve to grid size (and grid geometry) may be due to trade-offs between vector processing and multiprocessing. The largest problem, with 7.6 million grid points, executed at about 750 MFLOPS (wallclock). Having recently identified certain ambiguities in the DSCG solver and the norm calculation modules, we expect to achieve faster rates, possibly well over 1 GFLOPS, by simple modifications of these modules.

Applying a speed-up ratio of approximately 3.5 to 4.5 to the serial CRAY Y-MP timings given earlier, we have now, for three-dimensional heterogeneous test problems of one million nodes or more:

$a_{(\text{parallel CRAY Y-MP8})} \approx 0.04$ to 0.06 seconds per iteration per million nodes.

## Conclusions

The tests presented in this study demonstrate the feasibility of unusually large simulations of heterogeneous flow systems with the general-purpose BIGFLOW code. Given that the code was ported without prior modification to a CRAY Y-MP8 supercomputer, the observed performance and speed-ups due to parallel processing using the Autotasking feature are significant and encouraging. Similar analyses will be conducted for large nonlinear systems arising from models of unsaturated flow in complex geologic media. These new flow simulations will help produce the velocity fields needed for modeling radionuclide transport over thousands of years in a variably saturated, high-level nuclear waste geologic repository. ■


### *Acknowledgments*

*Work by the first author was funded by the U.S. Nuclear Regulatory Commission (FIN#B664); however, the views expressed in this paper are entirely those of the authors. NASA-Ames Research Center provided free computer time on its CRAY-2 system (NAS grant 1990/1991). Most tests on the CRAY Y-MP system were conducted at Cray Research. Volumetric color graphics were developed using the IVM package from Dynamic Graphics. The first author also thanks A. C. Bagtzoglou, G. W. Wittmeyer, T. J. Nicholson, and B. Sagar for their help.*


### *About the authors*


*Rachid Ababou is a senior research scientist at Southwest Research Institute, Center for Nuclear Waste Regulatory Analyses, San Antonio, Texas. His research in stochastic flow and transport phenomena is being used in performance assessment studies of the proposed national high-level waste geologic repository located at Yucca Mountain, Nevada. He received a doctor of engineering degree in fluid mechanics from the University of Grenoble, France, and a Ph.D. degree in civil engineering from the Massachusetts Institute of Technology.*

*Lynn W. Gelhar is a professor of civil engineering at the Massachusetts Institute of Technology where he teaches groundwater hydrology. A leading authority in the new field of stochastic subsurface hydrology, he has received the American Geophysical Union's Horton Award and the O. E. Meinzer Award from the Geological Society of America. Gelhar is preparing a textbook,* Stochastic Subsurface Hydrology, *to be published by Prentice-Hall, Inc.*

*Chris Hempel is an applications analyst with Cray Research in Austin, Texas. He received a bachelor's degree in mathematics from the University of Texas at Austin in 1972 and has worked as a systems analyst at the University of Texas System Cancer Center and at Control Data Corporation. Hempel spent five years providing on-site support at the University of Texas System Center for High Performance Computing.*


### *References*


1. Ababou, R., "Three-Dimensional Flow in Random Porous Media," Ph.D. thesis, Department of Civil Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts, 1988.
2. Ababou, R., *Approaches to Large Scale Unsaturated Flow in Heterogeneous, Stratified, and Fractured Geologic Media,* U.S. Nuclear Regulatory Commission Report NUREG/CR-5743, 1991.
3. Ababou, R., D. McLaughlin, L. W. Gelhar, and A. F. B. Tompson, "Numerical Simulation of Three Dimensional Saturated Flow in Randomly Heterogeneous Porous Media," *Transport in Porous Media,* Vol. 4, pp. 549-565, 1989.
4. Tompson, A. F. B. and L. W. Gelhar, "Numerical Simulation of Solute Transport in Three-Dimensional, Randomly Heterogeneous Porous Media," *Water Resources Research,* Vol. 26, No. 10, pp. 2541-2562, 1990.

# Improved financial and operational forecasting with large-scale reservoir models

*W. Culham, REGA Corporation, Inc., Houston, Texas*
*R. H. Deskin, D. D. Handyside, O. K. Karaoguz, and K.-M. G. Li, British Petroleum Exploration, Inc., Houston, Texas*

Accurately modeling a complex reservoir to predict potential gas and oil production requires computer resources that often are unavailable to engineers. Out of necessity, petroleum engineers have become proficient at circumventing computer resource limitations. Engineers at British Petroleum Exploration, Inc. (BPX) wanted to compare the accuracy of reservoir simulation in constrained conditions, where engineers use a "divide and conquer" approach, with the accuracy of a fine-scale simulation using the CRAY Y-MP 8E/8256 system. The results of this study have formed the basis for operational and financial decisions at BPX.

Accurately simulating a reservoir requires massive amounts of complex data. The reservoir rock is a complex heterogeneous container of fluids that requires a significant diversity of scale in order to model accurately. Heterogeneities range from pore level scale (< 10 microns) to thousands of feet for variations in formation lithologies and geologic structures within a reservoir. At all scales between these two extremes, the heterogeneities must be incorporated into a reservoir description used by a simulator to model fluid flow properly. In addition, an extensive time scale variation is needed to understand detailed fluid flow physics and long term reservoir performance. Relevant time scales range from minutes or hours for single well problems to years (10 to 50) for full field performance predictions.

VIP-EXECUTIVE, a Western Atlas software product, was customized by BPX[1] to provide a parallel processing capability and was used in the case study problem. VIP-EXECUTIVE is a general purpose, highly vectorized simulation product.

## Case study

As is typical of other reservoir studies undertaken by the petroleum industry, this case study required predicting the ultimate recovery from the reservoir, the rate of production, and the impact of operating decisions on recovery and economics. A specific operating question concerned well spacing and determining the optimum number of wells to maximize profit.

The geology of the case study was complex, involving channel fill sequences of hydrocarbon productive sands mixed with nonproductive shales. The channels were complex in their spatial distribution, giving rise to sand continuity and connectivity issues that strongly influence production/recovery

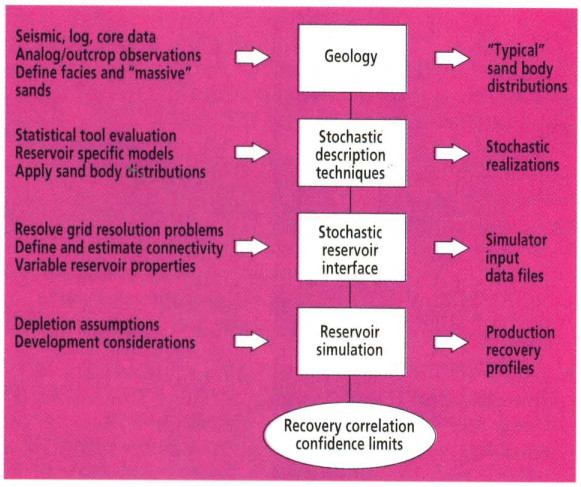


Figure 1 (left). Work description of reservoir study.

Figure 2 (right). Schematic illustrating partitioning used to subdivide study area.

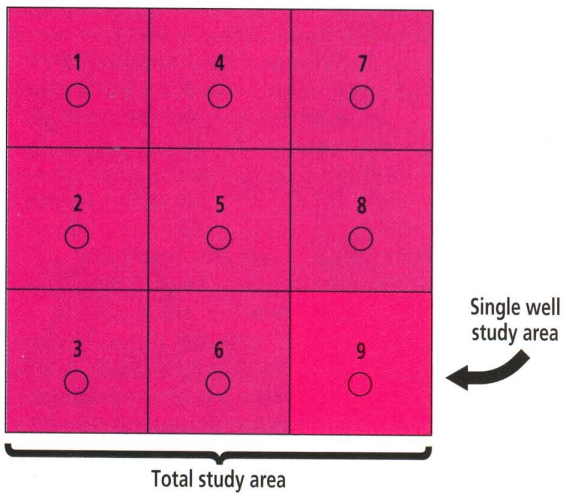

behavior. Seismic data helped define the ratio of pay to nonpay (net to gross ratio) and total formation thickness, but not the actual sand body distributions. Well data provided statistical information on net to gross ratio and the vertical distribution of the sands. Other data from analog reservoirs and outcrop studies helped to develop statistics on sand body widths and lengths.

Because of the complexity of the producing sands and limited data on their distribution, stochastic scenarios of the reservoir were used to develop input data for the simulator.[2] Figure 1 schematically illustrates the sequence of general steps used for the study. Once the reservoir description, fluid properties, and general operating conditions had been defined and provided to the simulator, recovery profile predictions were made for each scenario.

Preliminary reservoir description studies, using the stochastic scenarios, showed that the spatial attributes (connectivity) of the sand bodies could not be retained in grid systems with less than about 200,000 grid blocks over the entire study area. This number of blocks was too large, however, to initialize and run a simulation model with BPX's CRAY X-MP EA/464 supercomputer. Therefore, the study area was partitioned into nine areas, as illustrated in Figure 2, to assess production and recovery performance for individual wells on a 160-acre spacing. Each study area was discretized with a 16 by 30 by 50 layer grid system (24,000 total grid blocks). Corresponding block length, width, and thickness dimensions were 165, 88, and 2.9 feet, respectively. A typical distribution of sand bodies for the total study area is shown in Figure 3, with the front quarter of the model removed for clarity. Figure 4 shows the stochastic sand

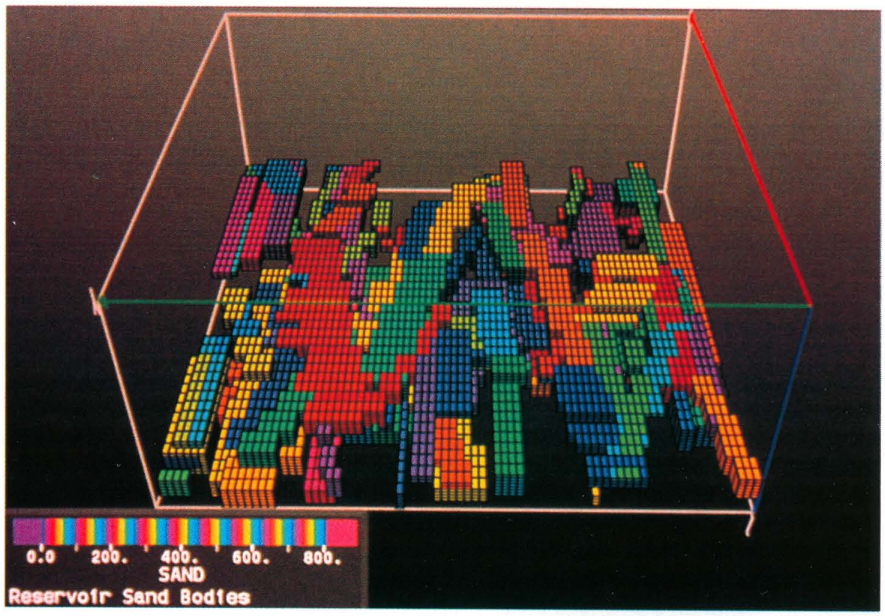


Figure 3. One realization of stochastic sand body distributions for total study area.

Figure 4. Areal view of one realization of stochastic sand body distribution near base of structure.

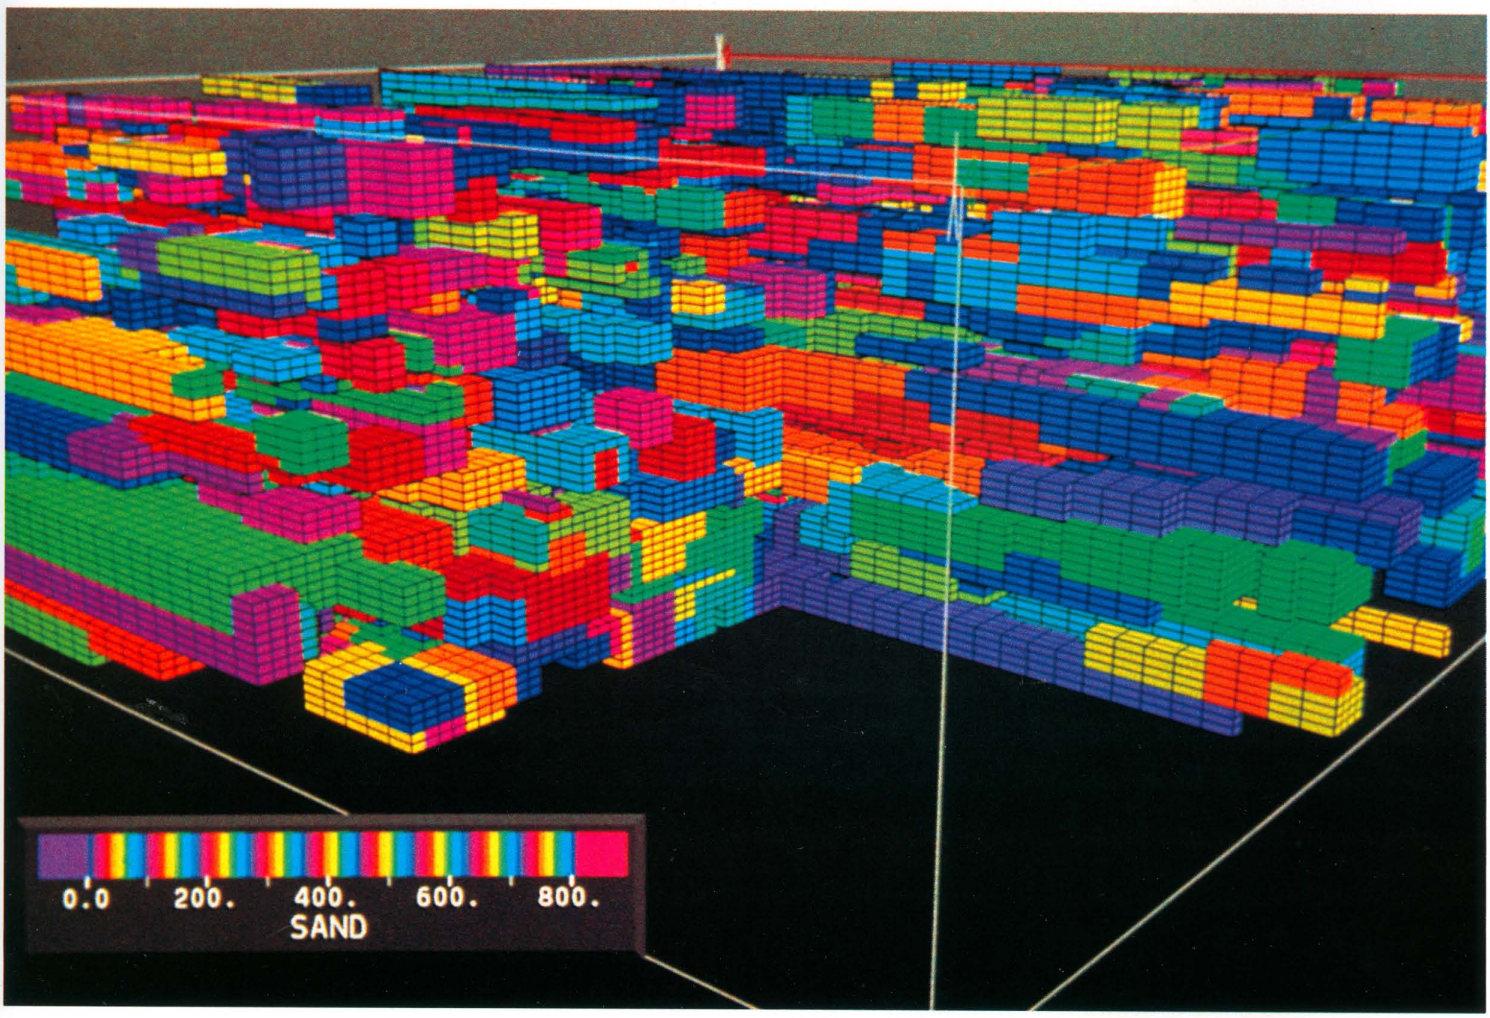

| Spacing (acres/well) | 320 | 160 | 40 |
|---|---|---|---|
| Oil recovery (% of oil in place) | 21.0 | 22.0 | 23.0 |

Table 1. Ten-year recovery predictions for entire study area simulated on the CRAY X-MP EA/464 system.

distribution at a depth near the base of the structure. The single production well for each of the partitioned study areas was located at the center of the Cartesian models and was completed across all sand layers present at this location. Next, each model was run under primary depletion for a period of ten years. Results from the nine separate models were combined to generate the production and recovery profiles for the entire study area. This combination or scale-up procedure is one method of circumventing computing resource limitations. Similar approaches were used to generate results for 320-acre and 80-acre well spacing cases.

The resulting ten-year recoveries, for the entire study area as a function of well spacing with a specified average net to gross ratio, are summarized in Table 1.

Engineers were concerned, however, about the impact of assumed no-flow boundary conditions in using the partitioning/combination approach to calculate total recoveries. The no-flow boundary conditions impose an artificial truncation of sand units that extend beyond the modeled areas. This results in a low estimate of the effective connectivity for the total study area. For example, sand bodies not connected to the well within one model area may be connected to other well locations in adjacent model areas. This additional connectivity would not be accounted for in the partitioning/combination approach. Another concern was the impact of adjustments to the stochastically generated reservoir description required in the process of gridding with some 200,000 blocks. These impacts were assessed through a side study using a CRAY Y-MP 8E/8256 system rather than the CRAY X-MP EA/464 system used in other parts of the study. The expanded memory and additional processors were necessary to model the entire study area at one time, and the parallel processing capabilities of the VIP EXECUTIVE software were used to exploit the multiple CPUs to minimize elapsed time. (One CPU was used for in-house runs on the CRAY X-MP system.)

To model the entire area and ensure an unaltered reservoir description for the simulation model, 1,500,000 grid blocks were used, likely constituting the largest applied reservoir simulation study conducted in the petroleum industry to date. Two separate studies were run, both simulating the total area. One model included internal boundaries (Figure 2 ) to replicate the nine study

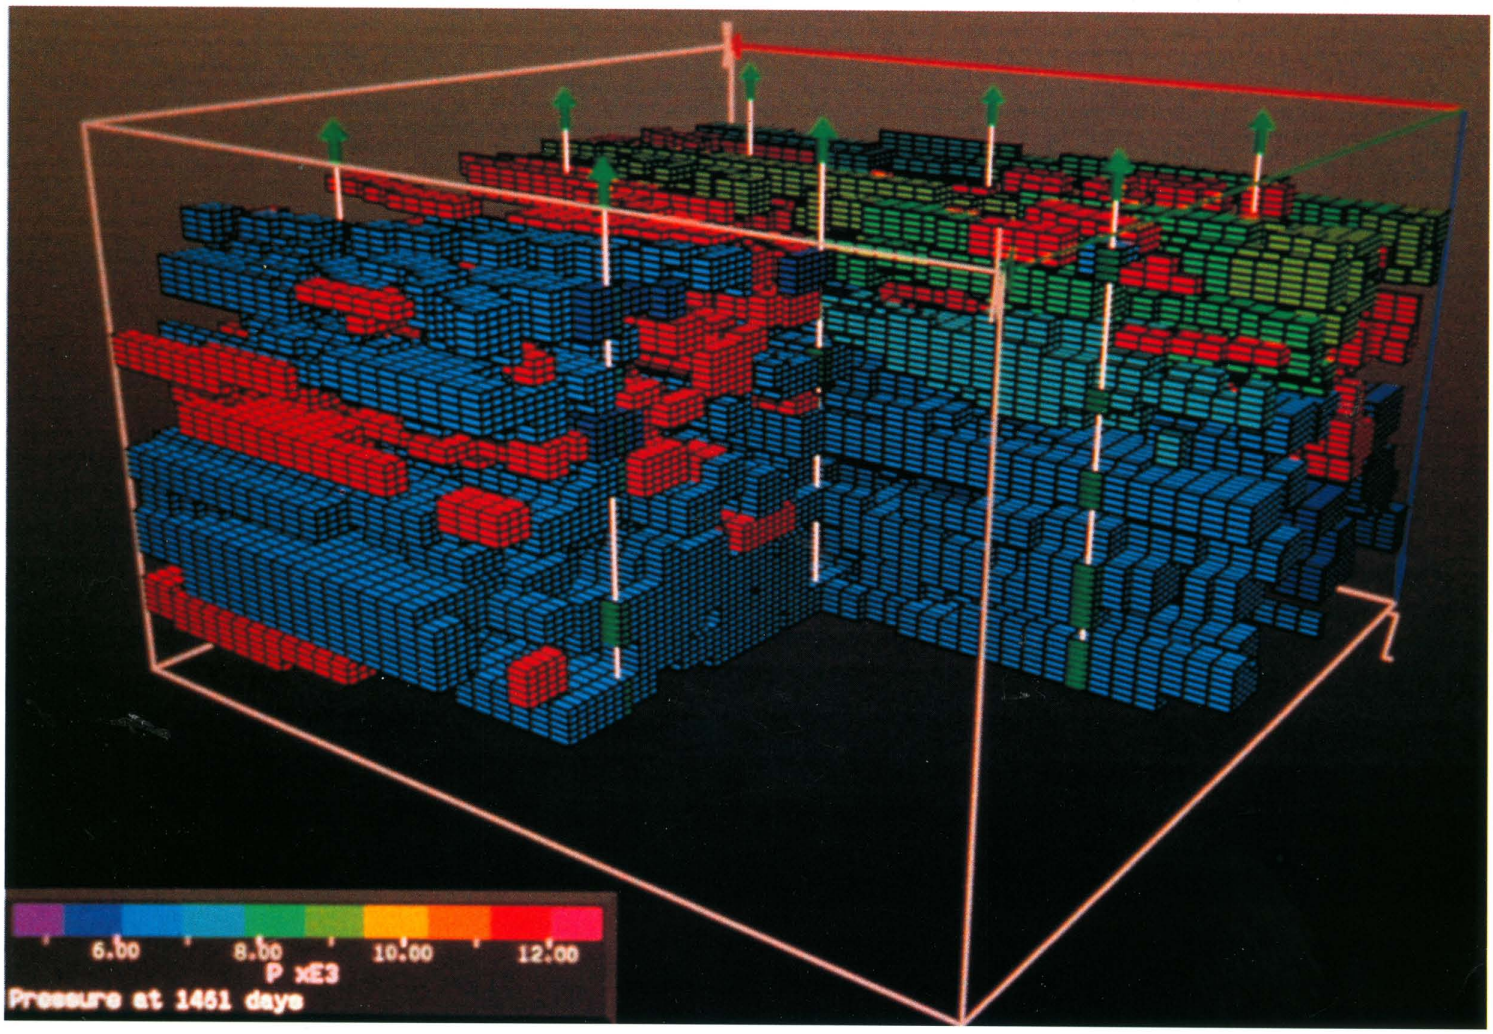


Figure 5. Pressure distribution at 1400 days.

| Run number | Number of grid blocks | Model configuration | Oil recovery (% of original oil in place) |
|---|---|---|---|
| 1 to 9 | 24,000 (each run) | Nine 160-acre partial models | 22.0 |
| 10 | 1,500,000 | Internal boundary present | 23.0 |
| 11 | 1,500,000 | No internal boundary | 28.0 |

Table 2. Results for 160-acre spacing simulations run on the CRAY Y-MP 8E/8256 system.

segments but to retain the fine-scale gridding influence. In this case, the internal boundaries were "no-flow barriers" to duplicate the same boundary conditions used in the nine smaller scale models. In the second study, the internal boundaries were removed. This approach allowed the effect of boundary conditions and reservoir description modifications to be assessed. An example of the pressure distribution across the sand bodies for the combined study area is illustrated in Figure 5 (with the front quarter of the model volume removed for clarity). This figure demonstrates the degree of depletion and connectivity (as inferred from the depletion level) within the sand bodies at a particular time. Results for the 160-acre spacing runs are listed in Table 2.

The results indicate that the adjustment of the original stochastic reservoir description data to conform to the simulation gridding imposed by the 24,000 cell models had only a minimal impact on recovery (sum of runs 1 to 9 versus run 10). However, when the influence of boundary conditions was eliminated by removing internal boundaries (run 11), an increase of 22 to 28 percent in predicted recovery occurred. Thus, a 27 percent (relative to smaller model case) predicted increase in recovery could be justified simply by more rigorous modeling. This recovery increase translates into 40 to 50 million barrels of additional predicted oil recovery. It improved the rate of return by two percentage points on a one-billion-dollar project, substantially improving the economic prospects of the project.

## Computational aspects

The three-dimensional, three-phase, compositional simulator VIP-EXECUTIVE is optimized to run on a single CPU of a Cray Research system. Significant modifications to the simulator were made to use the multiprocessing capability of the CRAY X-MP and CRAY Y-MP systems. The Autotasking feature of the CF77 Fortran compiling system was used to parallelize the simulator. For simple vector loops, the dependence analyzer (fpp) can be used to parallelize the program automatically. For more complicated DO loops, fpp often failed to recognize parallelism or transformed the vector loops into scalar loops. For these cases, the Autotasking directives were put in manually. For the phase equilibrium and well routines, the scalar DO loops contain many levels of nested subroutine calls and branching. This type of code is very difficult to vectorize, but it is ideal for multitasking. Because fpp cannot perform data dependency analysis across subroutines, these loops were parallelized manually. (A detailed description of this work can be found in Reference 1.)

The 1.5 million grid block stochastic model run on the CRAY Y-MP 8E/8256 supercomputer required 230 Mwords of main memory. A ten-year simulation run took 6.5 hours of wall clock time with all eight CPUs working on the problem. The parallel speedup of the program was measured by running a portion of the simulation twice–once with one processor and once with eight processors. The parallel speedups of various segments of the program are shown in Figure 6. The overall program speedup is slightly less than six (subroutine MAIN in Figure 6). Neglecting the multitasking overhead, Amdahl's law implies that about 95 percent of the code was parallelized. Using a speedup factor of six, the estimated wall clock time for the original single CPU version of VIP-EXECUTIVE is about 39 hours (speedup x wall time with eight CPUs). Based on BPX's benchmark results of other smaller reservoir models, the estimated run times of this problem in a dedicated environment on the CRAY X-MP supercomputer (one CPU), IBM550, and HP750 workstations are 59, 234, and 429 hours, respectively. The IBM and HP RISC workstation run-time estimates are actually optimistic. Because neither manufacturer makes a workstation with 256 Mwords of memory, running a model this large on a workstation with a limited amount of

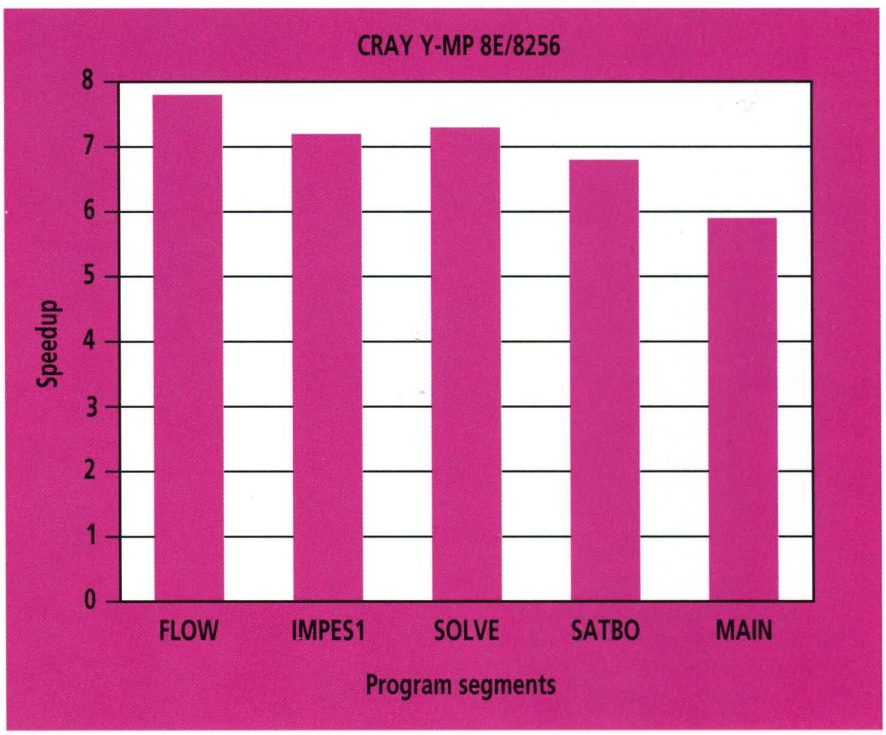


Figure 6. Parallel speedup of VIP software across eight CPUs.

# Reservoir modeling—big problems, powerful solutions

BPX engineers illustrated the advantage of a large-memory system for using simulation to more accurately predict a reservoir's production potential. Their case study is typical of the challenges petroleum engineers face when attempting to accurately model the complexities of a large-area reservoir. Some of these challenges are described below.

*The ability to represent geologic description in sufficient detail to account for its influences on fluid flow and subsequent recovery.* Because the actual reservoir is sparsely sampled (few wells in a large volume), little is known about the reservoir properties between wells. While this uncertainty can be reduced with geologic mapping, well testing, and seismic data, many scenarios may satisfy the requirements of the known data. This has led reservoir engineers and geoscientists to employ geostatistics to develop reasonable reservoir descriptions for simulation purposes.[2] Some screening techniques can be used to reduce the number of scenarios and subsequent simulations to generate recovery predictions. However, even with a reduced set of reservoir descriptions, multiple solutions are often circumvented, and a single case is used as the deterministic performance prediction.

*The ability to represent available geologic detail in coarsely gridded models.* Available detailed reservoir information for a sufficiently fine grid simulation is often not used because of limited computer memory and/or CPU time.

*The ability to utilize detailed fluid composition data.* Depending on the type of reservoir fluid and the recovery process, fluid composition data must be detailed. This in turn may require 20 or more components to represent complex phase equilibrium behavior realistically. Because of constrained computing resources, simplified fluid characterizations are employed, resulting in a compromised process physics representation.

*The ability to achieve numerical accuracy when using discrete numerical procedures to solve continuous systems.* Reservoir fluid flow, which can include one fluid displacing another, can be described as "shock-like." Sufficient grid blocks must be used to control the numerical dispersion of these fluids. Additionally, because many processes are strongly influenced by fluid compositions, cell size used in the simulator (which strongly influences mixture composition) must be discretized finely enough to produce accurate results. Finally, time-steps are bound by grid block size and fluid movement (Courant-like stability limits). These numerical requirements are frequently compromised because of computer resource limitations.

Thus, the reservoir engineer's challenge is to address conflicting requirements simultaneously, including sufficiently detailed, multiple scenarios of the reservoir description, fluid property details, and numerical accuracy. These requirements all are constrained by available computer resources. To partially address these challenges, model sizes have increased in the last two decades from a few hundred grid blocks and 1 or 2 unknowns per grid point, to studies ranging from 10,000 to 100,000 grid blocks and 1 to 20 unknowns per grid point. Within BPX, most studies fall in the range of 10,000 to 50,000 grid blocks with less than 10 unknowns per grid point. Using BPX's CRAY X-MP 4EA/432 system to solve these typical problems requires less than one hour (single CPU) to 150 hours (single CPU) for a complete run. However, modeling restrictions still exist even for this scale of simulation.

## Working with constrained computer resources

To compensate for the lack of necessary computer resources, most BPX simulations, like many in the petroleum industry, employ procedures to reduce system memory and overall computer time. One popular procedure to reduce the number of grid blocks and thus system memory and overall CPU time, is to use pseudo properties (pseudo relative permeabilities, pseudo capillary pressures, and pseudo components).[3] Pseudo properties are correlations developed from fine-scale simulation runs carried out on segments of the entire model. For special application cases and for strict usage conditions, pseudos can provide a close approximation to fine scale simulation results. However they are frequently used under circumstances that invalidate their use. Another procedure used under constrained computer resources is to represent complex displacement processes through the use of pseudo components and simple mixing models.[4] Yet another "work around" is to use two-dimensional models or a partial field model from a representative section of the total reservoir model, approximate true boundary conditions, and attempt to scale up these partial model results to predict total system performance.

## The role of big memory and MPP systems

The current capabilities of the expanded memory Cray Research systems as well as massively parallel processing systems can provide a simple, direct, and cost effective means of quickly meeting technical requirements of reservoir simulation studies. Additionally, the authors project that existing and/or proposed MPP systems (2000 CPUs) can be used to solve full field reservoir problems with as many as 40 to 50 million grid blocks. Models with grid blocks ranging from 2 to 40 million have been identified via sensitivity runs using partial field models as a requirement to accurately solve some of BPX's field scale simulation problems. Models of this magnitude would provide a way to properly represent small scale physics and a detailed reservoir description on an entire field basis.

Using the computational complexity definition proposed by Worlton,[5] the above scale of problems has a complexity of roughly $10^{16}$. With appropriately modified code, it is projected these problems would be solved with less than 100 hours of dedicated time on projected MPP resources. Thus, projected CPU time is comparable or even less than that used for the current scale of problems on BPX's CRAY X-MP system.

The advantages of fine-scale reservoir modeling made possible by large memory systems and/or MPP include:

- ☐ Reduced numerical accuracy problems
- ☐ Accurate representation of the fluid flow/phase behavior
- ☐ Ability to use and accurately reflect existing reservoir description data
- ☐ Reduction in engineer time devoted to conducting studies by using the circumventive approaches, such as the use of pseudo properties
- ☐ Elimination of unverified approximate approaches

Once the petroleum industry can capitalize on existing and future large memory or MPP systems, engineers can return to basic procedures for more accurate predictions of a reservoir's producing potential.

main memory would undoubtedly cause the system to page frequently, resulting in a significant degradation in performance.

## Engineering resources

Time required to develop customized software to use coarse reservoir descriptions in the simulator and engineering time needed to prepare and submit runs, postprocess simulation data, and analyze results for the subject study was roughly 0.6 engineer years (approximately 1000 hours). If access to a CRAY Y-MP 8E/8256 system had been available for the entire study, the time spent customizing software; developing the appropriate interface programs between reservoir description software and the simulator; and preparing, submitting, postprocessing, analyzing and combining separate simulation runs would have been reduced by more than 60 percent. Obviously, through the use of more powerful computers, study methods can be simplified and the engineering time associated with the study reduced significantly. The results of this human resource utilization analysis are similar to results of other studies conducted at BPX. The authors estimate that 50 percent of engineering time can be saved through more direct modeling methods and finer-scale numerical studies. With over 100 engineers involved in reservoir studies within BPX, this represents a 50-engineer-per-year reduction and/or productivity improvement. This further translates into improved staff productivity in excess of $7,000,000 per year.

## Conclusions

When operating under constrained computing resource conditions, petroleum engineers must partition the area to be studied into segments, forcing them to make assumptions about boundary conditions and reservoir description data. In the study carried out at BPX, this problem simplification resulted in conservative projections. The ability to conduct the study without computing resource constraints provided more reliable projections of reservoir performance, increasing predicted recovery by 27 percent and significantly improving the financial assessment. Additionally, engineering time required to customize the problem to meet computing resource constraints could be eliminated and the actual study time reduced, because fewer models would need to be run and analyzed. This would have reduced staff time on the project by 60 percent, significantly improving staff productivity. Computer resources also were used more effectively, as the customized parallel processing option in the VIP-EXECUTIVE software allowed maximum utilization of the computer (eight CPUs) and an acceptable throughput time.

In general, the development of multiprocessor computers and modification of software to exploit all processors will allow the solution of currently intractable reservoir problems. This improvement will simplify the industry's current study methodology, leading to improved productivity and more reliable reservoir performance forecasts. In many cases, the magnitude of the investment in hydrocarbon reservoir development is such that the reduced risk in decision making will more than offset separate and expanded capability system costs needed to improve reliability. ▬

***Acknowledgments***

*The authors wish to thank BPX for permission to publish this article, and Andy Austin of BPX for Figures 3 to 5, which were created with FOTO, from Cognivision, Inc.*

*The article presented here was condensed from the authors' paper,* Improved Financial Decisions on Oil and Gas Reservoir Exploitation with Supercomputers. *The original paper is available through Cray Research by calling Dept. D, 612-683/7215, or from the authors.*

***About the authors***

*W. E. Culham is currently president of REGA Corporation, Inc., an international petrotechnology provision company based in Houston, Texas. At the time of the writing of this article, he managed the Reservoir Technology Division of BPX's Technology Center in Houston, Texas. He has a Ph.D. degree in petroleum engineering from The Pennsylvania State University.*

*R. H. Deskin is a reservoir engineer with BPX's offshore USA business unit in Houston, Texas. She has M.B.A. and chemical engineering degrees from Rice University in Houston, Texas.*

*D. D. Handyside is a staff reservoir engineer with BPX's Technology Center in Houston, Texas. He has a B.Sc. degree in chemical engineering from the University of Alberta, Canada.*

*O. K. Karaoguz is a reservoir engineer with BPX's Technology Center in Houston, Texas. He has a M.Sc. degree in petroleum engineering from Stanford University, California.*

*K.-M. G. Li is a staff reservoir engineer with BPX's Technology Center in Houston, Texas. He has a Ph.D. degree in mechanical engineering from the University of California, Berkeley.*

***References***

1. Li, K-M. G., "Multi-tasking of a Reservoir Simulator on the CRAY X-MP," *Proceedings of the Fifth SIAM Conference on Parallel Processing for Scientific Computing,* March 1991.
2. Haldorsen, H. H. and E. Damsleth, "Stochastic Modelling," *Journal of Petroleum Technology,* p. 404, April 1990.
3. Stone, H. L., "Rigorous Black Oil Pseudo Functions," paper 21207, *Proceedings of the Eleventh SPE Symposium on Reservoir Simulation,* Anaheim, February 17-20, 1991.
4. Todd, M. R. and W. J. Longstaff, "The Development, Testing and Application of a Numerical Simulator for Predicting Miscible Flood Performance, *Journal of Petroleum Technology*, p. 874, July 1977.
5. Worlton, Jack, "Workstations and Supercomputers: A Partnership for the 1990s," *Worlton and Associates Technical Report,* No. 36, January 1991.

# CORPORATE REGISTER

## Cray Research receives first commercial order for CRAY Y-MP C90 system

**Ford Motor Company** ordered a CRAY Y-MP C90 system to be installed at the automaker's Engineering Computer Center in Dearborn, Michigan. The new system, which will replace a CRAY Y-MP supercomputer that Ford installed in 1989, will be used for structural analysis, crash simulation, computational fluid dynamics and other problems related to automotive design and engineering. Ford was the first U.S. automotive company to purchase a CRAY Y-MP system. The CRAY Y-MP C90 system is the follow-on to that system and continues Ford's commitment to using leading-edge computational solutions for computer-aided engineering.

**Lawrence Livermore National Laboratory (LLNL)** has signed an agreement to purchase a CRAY Y-MP C90 supercomputer for delivery to LLNL's National Energy Research Supercomputer Center (NERSC) in 1992. LLNL's order for the CRAY Y-MP C90 system follows the cancellation of a system order from Cray Computer Corporation of Colorado Springs, Colorado.

The Republic of China (Taiwan) **Central Weather Bureau (CWB)** has ordered a CRAY Y-MP 8I system and a CRAY Y-MP EL system. The CRAY Y-MP 8I system, an upgrade replacement for a Control Data Corporation CYBER 205 system, will be used by the Taipei-based weather bureau for daily weather forecasting. The CRAY Y-MP EL system, Cray Research's entry level supercomputer, will be used as a file server. The CRAY Y-MP 8I system also will operate the EMPRESS relational database management system with a unique binary large object (BLOB) attribute. This attribute is ideally suited for storing data that has no easily defined structure, such as satellite data, analyzed fields, forecast fields, and satellite images. The use of EMPRESS on CWB's new system is the first instance of a Cray Research weather customer running a standard database application package on a Cray Research system.

**The Institut Mediterranéen de Technologie (IMT)**, a leading teaching and research institution in Marseilles, France, ordered and installed a CRAY Y-MP 2E system for general academic and industrial research, primarily in the area of computational mechanics applications with emphasis on aeronautics, astronautics, and ocean engineering.

The **Max Planck Society** has installed a CRAY Y-MP 4E supercomputer at its Institute for Plasma Physics in Garching, Germany, a suburb of Munich. The institute is one of the world's premier high energy physics research laboratories. The CRAY Y-MP 4E system, the third Cray Research supercomputer installed at the institute, replaces a CRAY X-MP system. It will be used primarily for research in plasma physics, astrophysics, and solid state physics. The system also will be used by other Max Planck Institutes in Germany for a variety of research applications.

The **Australian Bureau of Meteorology,** Australia's weather agency, has ordered and installed a CRAY Y-MP 2E supercomputer at its computer center in Melbourne. The new system joins a CRAY X-MP system, which was the bureau's first Cray Research system and was installed in July 1990. The CRAY Y-MP 2E system, which will be used primarily for operational weather forecasting for Australia, will be networked to the CRAY X-MP system. The Australian Bureau of Meteorology operates one of the major processing centers within the World Meteorological Organization's Weather Watch program and is the leading center of its kind located in the Southern Hemisphere.

The **Swiss Federal Institute of Technology (EPFL)** located in Ecublens, Switzerland, was the first customer worldwide to install a CRAY Y-MP EL system, Cray Research's entry level supercomputer. EPFL has a nationwide network of approximately 3000 computer systems including workstations, servers, mainframes, and Cray Research supercomputers. EPFL also ordered a CRAY Y-MP 4E system that will operate as a central compute server. The CRAY Y-MP 4E supercomputer will be networked via the 800-Mbit/s HIPPI interconnection to the CRAY Y-MP EL system, which will act as a high-speed file server for the central computers and all the departmental compute servers on the network.

Three other European organizations also have ordered CRAY Y-MP EL systems: **The Defence Research Agency, Maritime Division, Admiralty Research Establishment (ARE)** Haslar, a UK Ministry of Defense (MOD) site located at Gosport in Hampshire; **Lotus Engineering**, an automotive consultant in the UK that is part of Group Lotus, the manufacturer of high-performance sports cars; and **Ecole Centrale de Paris (ECP)**, an engineering school in Paris, France.

ARE researches all aspects of Naval Operations, with a special emphasis on advanced defense technology. The new CRAY Y-MP EL system will be used for computational fluid dynamics applications in ship and submarine hydrodynamics. Lotus Engineering will use its CRAY Y-MP EL system to support the facilities and expertise offered by its Design Analysis Group in Hethel, England. The ECP system will be used by the school's Mechanical Engineering Laboratory, which works to optimize codes and applications for the French automotive, geophysical, civil engineering, and other industries.

## Sun Microsystems and Cray Research to cooperate on development of hardware and software environment

Sun Microsystems, Inc., and Cray Research announced in January that the firms will cooperate in the creation of a hardware and software environment that will allow Cray Research systems and Sun workstations to work together more easily and efficiently.

"There is a very fast growing installed base for SPARC systems, with more than 3800 third-party hardware and software products," said John A. Rollwagen, chairman and CEO of Cray Research. "We're committed to the concept of network supercomputing so that individual users can easily access the power of our supercomputers." Cray Research has joined SPARC International, an organization of more than 200 SPARC vendors that supports binary compatibility among products through the SPARC Compliance Definition (SCD).

"Cray Research has the ability to build the ultimate accelerator for SPARC applications," said Scott McNealy, Sun chairman and CEO, "and by choosing SPARC, Cray Research can take advantage of the benefits of the SPARC community. SPARC users have access to a wide range of compatible systems at many price points from laptops to supercomputers."

## Cray Research and Digital Equipment Corporation reach worldwide marketing agreement

Cray Research and Digital Equipment Corporation announced in January a comprehensive agreement that grants Digital worldwide rights to market, sell, and distribute the CRAY Y-MP EL supercomputer system. The CRAY Y-MP EL air-cooled, entry level supercomputer is fully compatible with Cray Research's entire CRAY Y-MP supercomputer product line and integrates easily into Digital's open distributed computing environment. Under terms of the agreement, Digital will sell the system in conjunction with Digital's own high-performance technical computing (HPTC) products: the VAX vector systems and the DECmpp line of massively parallel systems.

The agreement provides for an orderly marketing and sales transition. For the first six months of the agreement — through June 30, 1992 — both Digital and Cray Research will market and sell CRAY Y-MP EL systems. The agreement does not apply in certain limited cases, such as CRAY Y-MP EL systems sold in Japan and to certain government customers. Starting July 1, 1992, with those limited exceptions, customers will purchase CRAY Y-MP EL systems exclusively through Digital sales representatives. Cray Research will offer back-up sales support and applications expertise for the CRAY Y-MP EL system. Officials of both companies said they consider 1992 a building year for this relationship and that they expect the benefits of the agreement, including sales volume increases, to become more evident in 1993 and succeeding years.

## Cray Research completes acquisition of selected FPS assets, creates wholly owned subsidiary

In December, Cray Research completed its acquisition of selected assets of Floating Point Systems (FPS), producers of a minisupercomputer based on SPARC microprocessor technology. Cray Research has formed a wholly owned subsidiary, Cray Research Superservers, Inc., to market and support current products and to explore technologies for creating software compatibility between the company's supercomputers and SPARC-based systems. Cray Research acquired rights to FPS technologies, system and parts inventories, and equipment to support manufacturing and development efforts. Selected customer and vendor contracts also were assumed as part of the transaction.

## Cray Research announces first graduate fellowships in computational chemistry

In December, Cray Research, in collaboration with the American Chemical Society's (ACS) Division of Physical Chemistry, awarded three graduate fellowships in computational chemistry for the 1991-92 academic year. The Cray Research Fellowship in Computational Chemistry was established to encourage doctoral degree studies in computational chemistry; the fellowships are to be awarded annually on a competitive basis to graduate students in the research and dissertation stage of their doctoral programs.

The award focuses on innovative developments and applications of methods in theoretical chemistry or related disciplines, with emphasis on large-scale numerical simulations on supercomputer systems. Recipients of the 1991-92 Cray Research Fellowships in computational chemistry are

- ☐ Zhihong Deng, Department of Chemistry, University of Pennsylvania
- ☐ Scott M. Le Grand, Department of Chemistry, Pennsylvania State University
- ☐ Randall Q. Snurr, Department of Chemical Engineering, University of California, Berkeley

"These fellowships allow Cray Research to recognize and support some of the finest emerging talent in computational chemistry," said Robert A. Eades, manager of Cray Research's chemical research and engineering group. "We are delighted to be helping these talented graduate students continue their important and innovative research work."

Funded by the Cray Research Foundation, the fellowships provide a cash stipend of $5000 for each student per school year as supplementary financial support. In addition, the students were recognized personally by representatives of Cray Research and the ACS Subdivision of Theoretical Chemistry in April at the 1992 ACS Spring National Meeting in San Francisco.

Applicants interested in the 1992-93 Cray Research Fellowship in Computational Chemistry should submit applications no later than June 30, to Dr. John Tully, AT&T Bell Laboratories, Room 1D346, 600 Mountain Avenue, Murray Hill, NJ 07974. Applications should not exceed five typewritten pages and must include a definition of the proposed research project and an explanation of its scientific importance, the proposed use of the grant money, a description of the importance of high performance scientific computing to the proposed research, and two letters of recommendation.

# APPLICATIONS UPDATE

## IMSL Parallel Libraries available on Cray Research systems

The IMSL Libraries, a collection of mathematical and statistical Fortran subroutines from IMSL, Inc., has been parallelized to take maximum advantage of Cray Research parallel processing technology. The vectorized IMSL Libraries product is installed on more Cray Research systems than any other third-party product. The new IMSL Parallel Libraries product runs on CRAY X-MP and CRAY Y-MP systems under the UNICOS operating system.

The Parallel Libraries provides a significant performance boost to the standard IMSL Fortran Libraries by including fine- and coarse-grain parallelism.

Elements of fine-grain parallelism at the DO-loop level improve performance particularly with linear systems, fast Fourier transforms, and matrix/vector computations. The fine-grain elements allow loop iterations to be parceled among a system's CPUs. The libraries also are linked to the Cray Research level 1, 2, and 3 basic linear algebra subroutines (BLAS) for even better performance. The entire library set of over 900 mathematical and statistical routines can take advantage of fine-grain parallelism.

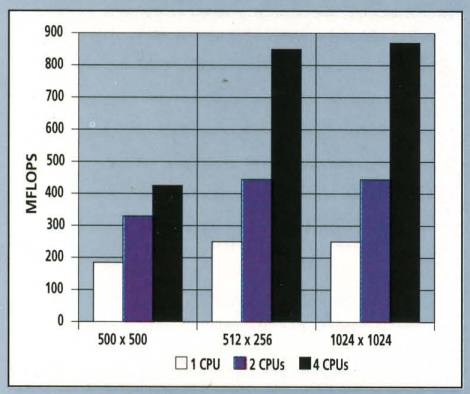


**Performance comparison of a fast Fourier transform (FFT) on a CRAY Y-MP 8/832 system using one, two, and four processors.**

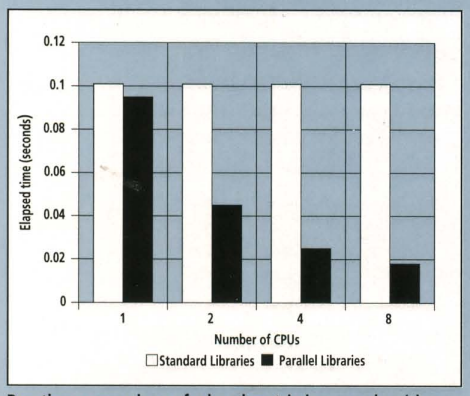


**Run time comparison of a band matrix inverse algorithm (matrix order = 200) on a CRAY Y-MP8/832 system using the IMSL Standard and Parallel Libraries.**

Coarse-grain parallelism lets users execute IMSL subroutines within a parallel section. The product automatically uses either the maximum available number of processors or a number preset by the user. Users can expect near linear or even greater performance increases from the combination of fine- and coarse-grain parallelism.

The IMSL Parallel Libraries product is based on the algorithms in version 2.0 of the vectorized IMSL Libraries. The MATH/Library features routines in the following chapters: linear integration and differentiation, differential equations, transforms, nonlinear equations, optimization, matrix/vector operations, sparse matrices, and more. The STAT/Library includes basic statistics, regression, correlation, analysis of variance, categorical and discrete data analysis, nonparametric statistics, time series analysis and forecasting, discriminant analysis, probability distribution functions, and random numbers.

For more information on using the IMSL Parallel Libraries on Cray Research systems, contact David Barret, IMSL, Inc., 14141 Southwest Freeway, Suite 3000, Sugar Land, Texas, 77478; telephone: 713/279-1019.

## INGRES Intelligent Database 6.3 now available on Cray Research systems

A full implementation of the INGRES product set now is available for use on Cray Research supercomputers, teaming the flexibility of the INGRES RDBMS with the speed and power of Cray Research supercomputers. This implementation includes the INGRES Intelligent Database release 6.3 and the following modules:

- ☐ INGRES/SQL and INGRES/QUEL, two query languages, offer users a variety of options including the flexibility of interactive use and the capacity to embed queries in the C, Fortran, and Ada programming languages.
- ☐ INGRES/MENU makes INGRES subsystems easy to use with a visually oriented, forms-based interface.
- ☐ INGRES/QUERY interface makes data access easy for users without knowledge of the SQL or QUEL query languages.
- ☐ INGRES/APPLICATIONS accelerates development of forms-based INGRES 4GL applications.
- ☐ INGRES/REPORTS generates quality reports using the forms interface or the high-powered report writer language.
- ☐ INGRES/FORMS helps create forms-based terminal display applications with a visual-forms editor (VIFRED).

For corporate-wide data management, INGRES networking tools also are available. INGRES/NET links Cray Research computer systems to client/server environments for easy access to data stored on remote systems, while INGRES/Gateways allow seamless access to information contained in databases other than INGRES, such as IBM's DB2, SQL/DS, and IMS, and Digital Equipment's Rdb and RMS.

INGRES databases combined with the power of Cray Research systems enhance network supercomputing, enabling users at any node in a network to access data in any database supported by INGRES in the network, process it with the speed and power of a Cray Research supercomputer, and view the results at their desks.

For more information on using INGRES 6.3 with Cray Research systems, contact Ingres Corp. at 1/800/4INGRES or Denny Olson, Cray Research, Inc., 655E Lone Oak Drive, Eagan, MN 55121; telephone: 612/683-3677; email: denny@cray.com.

# USER NEWS

## Prime time at Harwell Laboratory

Adrian Powell entered AEA Technology's Harwell Laboratory in England on February 19, 1992, unaware of the discovery that had occurred overnight. As he reviewed the results of a routine preventative maintenance test run on the laboratory's CRAY-2 supercomputer, the Cray Research UK systems analyst saw this message: 756839 prime 68618 sec Wed Feb 19 12:36:58 1992 sn2008. The world's largest known Mersenne prime number was found—and the 32nd perfect number was close at hand.

The new 227,832-digit prime number is $2^{756,839} - 1$, or two multiplied by itself 756,839 times minus one. It is named after Father Marin Mersenne, a 17th century French monk who spent years investigating numbers of the form $2^p-1$ where $p$ is prime. The previous largest known Mersenne prime, $2^{216,091} -1$, was discovered in 1985 during a similar test of a Cray Research system at Chevron Information and Technology Company in Houston, Texas.

The program used to find Mersenne prime numbers is based on an algorithm called the Lucas-Lehmer test. Because it is very computationally intense, the Lucas-Lehmer test is used as a system diagnostic tool. For this program, supercomputers are hundreds of times faster than even the most powerful conventional mainframes or engineering workstations. Cray Research scientists Paul Gage and David Slowinski have optimized the Lucas-Lehmer test to create the version used to find the new Mersenne prime at Harwell.

"What's exciting about this discovery is that we were able to use the products and tools that the people at Cray Research build—the world's best compilers, libraries, microtasking tools, hardware, and performance analyzers—and put them together to create the fastest Lucas-Lehmer test available," said Gage. "It's what I call making your own luck, which is a major component of this discovery."

Every known Mersenne prime is associated with a perfect number: if $2^p-1$ is prime, then $(2^p-1) \times 2^{(p-1)}$ is perfect. In the Harwell discovery, $(2^{756,839}-1) \times (2^{756,838})$ is the perfect number, with 455,663 digits. A perfect number is equal to the sum of all its factors except itself. For example, 28 is perfect because 14+2+7+4+1=28. Of the 32 perfect numbers now known, five were discovered on Cray Research systems over the past 13 years.

The test of the new Mersenne prime took nearly 19 hours on one central processing unit (CPU) of Harwell's four-CPU CRAY-2 system. The number was checked in under 20 minutes on a 16-processor CRAY Y-MP C90 system at Cray Research's Chippewa Falls, Wisconsin, facilities. With a completely different program on another CRAY-2 supercomputer, the discovery was verified independently by Richard Crandall, chief scientist at NeXT Computers, with David Smitley and Josh Doenias. The most powerful conventional mainframes or engineering workstations would take weeks to run such a program—and check just one Mersenne number. There are 41,452 Mersenne numbers between 216,091 and 756,829.

AEA Technology's Harwell Laboratory is one of the largest contract research and development laboratories in Europe with work programs in many nuclear and non-nuclear fields, including complex chemical, physical, and structural studies using high performance computer systems.

## To V or not to V—and when

In February 1992, Fluent Inc. volunteered to model the airflow around a ski jumper using the conventional parallel ski style versus the tips-out, or V, formation being used by many jumpers in Olympic competition. Fluent used RAMPANT, its latest computational fluid dynamics (CFD) software product, and Cray Research ran the mathematical model on a CRAY Y-MP8 supercomputer to provide visualization of the effects of using the V style. Cray Research's Multipurpose Graphic System (MPGS) was used to examine the results. Fluent donated the work on behalf of the U.S. Olympic Team to CBS color commentator Jeff Hastings.

Once skiers are airborne, the entire glide path, including speed and distance, is determined by the airflow around them and the pressure it exerts on the body. Understanding these airflow patterns can provide excellent insight into the benefits of the parallel style versus the V formation. For example, at takeoff it is important for skiers to assume a position that minimizes drag force. Further along the glide path, skiers want to maximize the lift force and "fly."

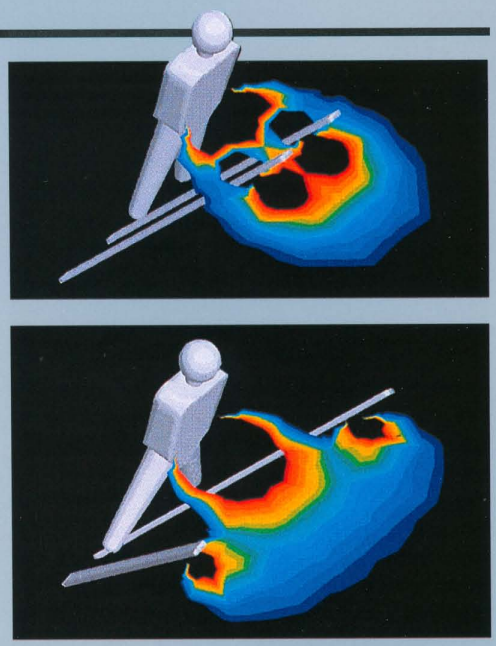

Parallel style ski jump position (top) and V formation (bottom). The difference in pressure, or "sail" area, in the two figures suggests that the V style can mean a 10 percent increase in jump distance, a prediction borne out by Olympic class jumpers.

The modeling results confirmed that the V style produces a more favorable airflow around the skier and will help the skier fly further toward the end of the glide path. If used immediately after takeoff, however, the V style slows down the skier, producing shorter jumps. It is important that skiers assume the V position late in the jump.

Fluent president Bart Patel said, "RAMPANT's ability to model complete geometries along with the speed and visualization capabilities of a CRAY Y-MP8 system provide a new dimension in CFD modeling."

RAMPANT is used to solve compressible and incompressible fluid flow and heat transfer in complex geometries. It uses triangular grids in two dimensions and tetrahedral grids in three dimensions. This grid topology significantly reduces the time and effort required to generate grids for complex geometries. RAMPANT also features a solution-adaptive grid capability that allows local refinement of the mesh to accurately capture flow features such as shock waves and shear layers.

# CRAY USER GROUP

## CUG reports

*Users of Cray Research computer systems established the Cray User Group (CUG) in 1977 to provide a forum for the exchange of ideas related to Cray Research systems and their applications. The group holds two general meetings each year. Its first meeting of 1992 was held April 6-10 in Berlin and was hosted by the Konrad-Zuse-Zentrum für Informationstechnik Berlin (ZIB). Below, CUG president Karen Scheaffer, of the Sandia National Laboratory, offers her comments on the meeting and other CUG-related business.*

"The future isn't what it used to be."
—*Arthur Clarke*

Change was in the air at the spring CUG conference. Many of the attendees' sites are experiencing change in one form or another, such as shrinking budgets, the cracking of the centrally supported "glass house," or corporate consolidation. Cray Research itself has not been immune to change. In recent months we have heard of the following changes: the agreement with Digital Equipment Corporation to sell CRAY Y-MP EL systems, the creation of Cray Research Superservers and the introduction of the CRAY S-MP system, and the licensing of Digital's Alpha microprocessor for use in the Cray Research MPP system.

Because the city in which we met has experienced a great deal of change itself, it was a perfect place to consider how CUG might adapt to and take advantage of industry changes. At the Board of Directors meeting, the Board concluded that all customers purchasing Cray Research computers will be invited to become members of CUG, including CRAY Y-MP EL customers served by Digital and Cray Research Superservers customers. This conclusion is consistent with the CUG bylaws.

The Board is aware of the potential impact of this decision, which likely will cause our membership to increase drastically from the current 193 sites. To this end, the Board began the process of developing a vision for CUG. This process will continue at the summer Board meeting along with a re-evaluation of our financial strategy.

The CUG Board cannot chart a course for the organization without your input, comments, suggestions, and course modifications. To ensure that CUG will continue to be a viable user group for you, the Board of Directors and Advisory Council need your help and input now. Your suggestions are welcome—whether orally, written, or by email—at meetings or informally. The Board has set up an electronic mail reflector (cugbod@sandia.llnl.gov) to allow you to send suggestions and concerns between conferences.

I also believe that we can make change positive by focusing our attention outward. In our industry we have watched the creation of various vendor and user consortia, such as OSF, UI, Sparc International, and ACE. The delineation between vendors is no longer clear. By having an outward focus—networking with other user groups and representing our members in external organizations such as POSIX—CUG will enhance its usefulness.

The CUG conferences are shining examples of volunteer effort, and this conference was no exception. I want to personally thank our host Jürgen Gottschewski, ZIB, and his volunteer team for their time and effort in organizing and staffing the conference.

With the CUG elections just around the corner, we are looking for more volunteers. Mary Zosel, CUG past president, is chairing the nomination committee. She is seeking people to run for president, vice president, and treasurer, as well as two regional directors: one for North, South, and Central America; and another for Asia, Australia, and New Zealand. If you are interested in running for any of these positions, please contact Mary.

On behalf of the CUG membership, the Board submits formal requests to Cray Research. Two of the four outstanding requests were closed in Berlin: the licensing of Cray Research products on non-Cray Research hardware and a review of the decision to support version 5.1 of the UNICOS operating system. Cray Research's plan to license its products on non-Cray Research hardware was presented in Berlin, and the issue of reviewing UNICOS 5.1 support no longer exists because fewer than 20 percent of the sites are running it.

Two open requests remain: I/O support for sites that do not have a model E IOS and a formal mechanism to monitor sites during an upgrade to a released product. The I/O support request is still being studied, and Cray Research's customer service organization is developing a process to monitor upgrades of released products.

CUG provides many useful and cost-saving services for your site. These benefits and services include:

- ☐ The unicos-l electronic mail reflector, created and maintained by Jerry Burkman of UC Berkeley, which can be used to discuss software and hardware issues relating to Cray Research computers
- ☐ CUGDUS (Cray User Group Directory of User Submitted Software), a directory of free software provided by CUG members
- ☐ The CUG institutional representative to POSIX, who represents the CUG viewpoint at POSIX meetings and in POSIX ballots
- ☐ The CUG newsletters and conference proceedings
- ☐ The CUG semiannual technical conferences. The CUG conferences provide a forum where attendees can learn about new Cray Research hardware and software and other advances in high performance computing, discuss issues with other Cray Research sites, and provide input to Cray Research.

Cray Research is listening to its customers, in particular to CUG members. On the last day of the conference I attended a session on CRInform, an on-line information and problem reporting service. CRInform was developed in response to the input from CUG members of the CASPR database. If you haven't signed up for the service, please do so and, most importantly, give Cray Research your input.

CUG conferences provide an excellent opportunity to give your input to Cray Research. The next conference will be held September 14-18, 1992, in Washington, D.C. and will be hosted by Howard Weinberger from Technology Applications, Inc. Please help us celebrate Cray Research's 20th anniversary at the 30th CUG conference. I look forward to seeing you there.

# GALLERY

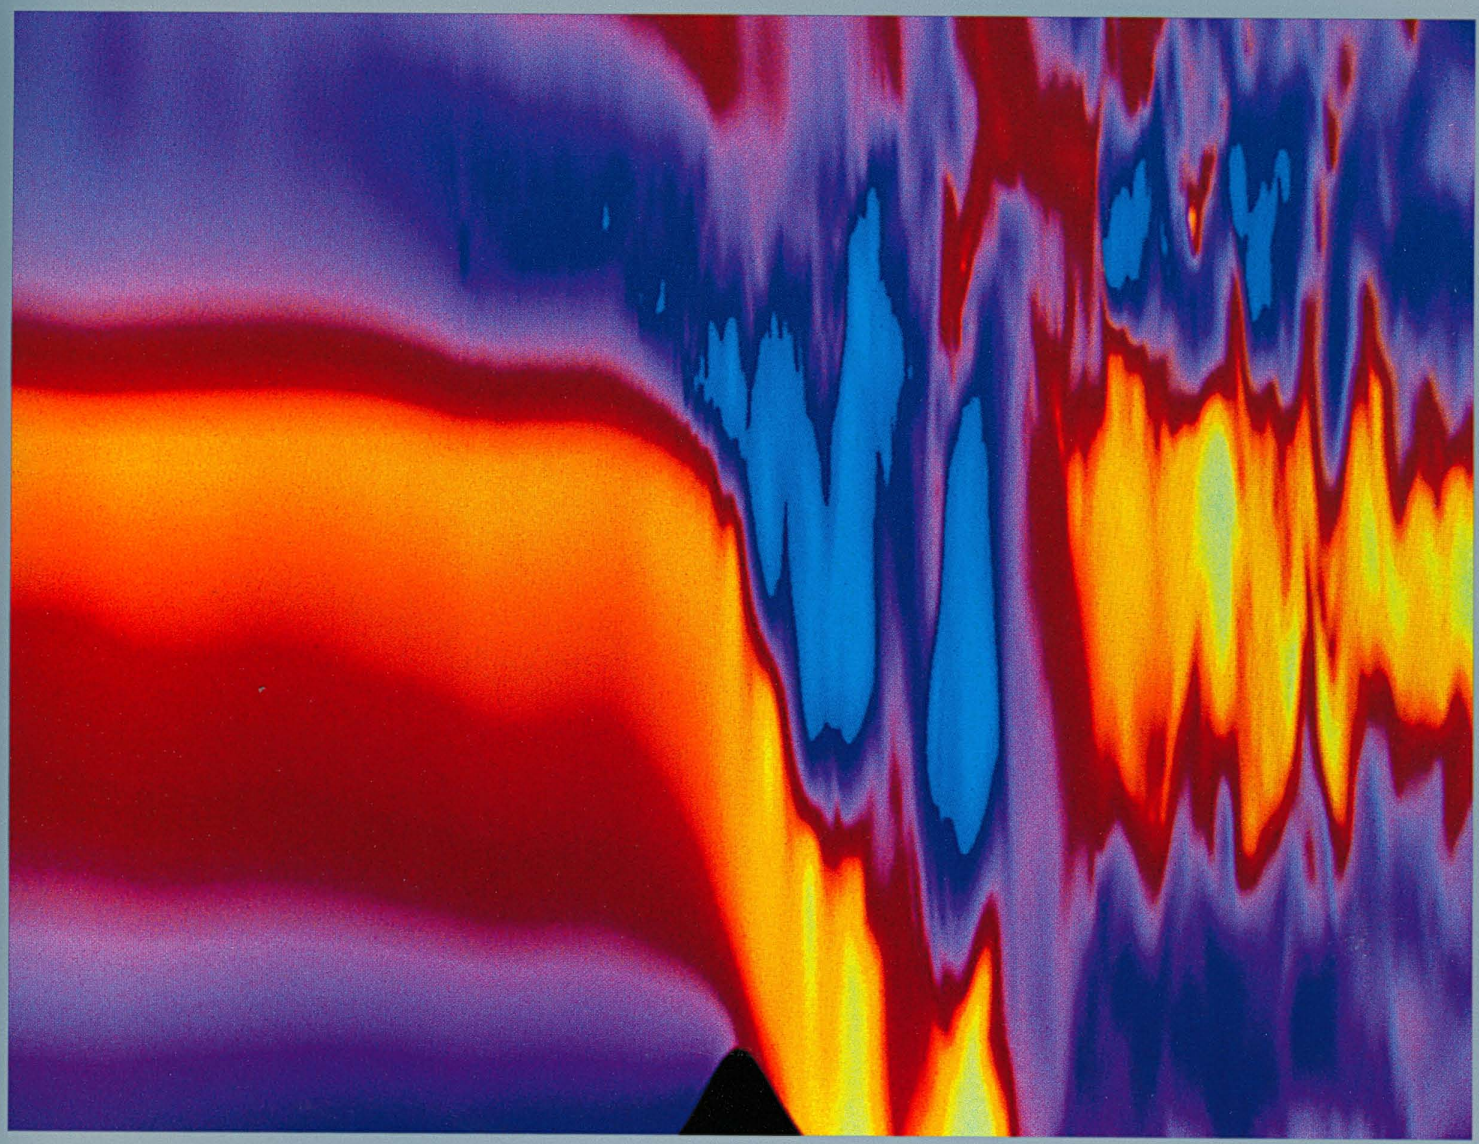

Fluid flow model showing internal waves breaking, and the chaotic state that results, when a flowfield encounters a critical topographic feature. The calculations were performed at the Ontario Centre for Large Scale Computation on a CRAY X-MP/24 system. This flow behavior characterizes that of the Chinook wind storms of Western Canada, for which the Rocky Mountain system is the critical topographic feature. Image courtesy John Scinocca and William Peltier.