Energy & Technology Review

- The Industrial Computing Initiative
- Artificial Hip Joints: Applying Weapons Expertise to Medical Technology
- Research Highlights

University of California

Lawrence Livermore National Laboratory

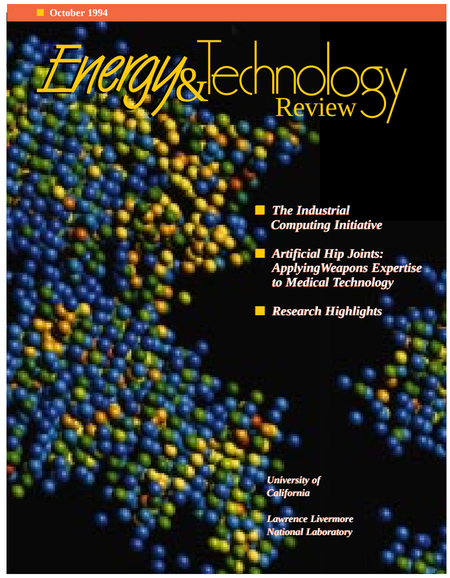

About the Cover

The benefits of the massively parallel computing being explored by the Industrial Computing Initiative (ICI) are great to both the international scientific community and U.S. industry. The U.S. semiconductor industry, for example, will greatly reduce development time and cost in the fabrication of the next generation of semiconductor devices by combining accurate simulations on the atomic scale with modern three-dimensional visualization techniques to model diffusion mechanisms in silicon like the one on our cover this month. For further discussion of the ICI and its recent achievements and potential benefits, turn to the article beginning on p. 1.

About the Journal

The Lawrence Livermore National Laboratory, operated by the University of California for the United States Department of Energy, was established in 1952 to do research on nuclear weapons and magnetic fusion energy. Since then, in response to new national needs, we have added other major programs, including technology transfer, laser science (fusion, isotope separation, materials processing), biology and biotechnology, environmental research and remediation, arms control and nonproliferation, advanced defense technology, and applied energy technology. These programs, in turn, require research in basic scientific disciplines, including chemistry and materials science, computing science and technology, engineering, and physics. The Laboratory also carries out a variety of projects for other federal agencies. *Energy and Technology Review* is published monthly to report on unclassified work in all our programs. Please address any correspondence concerning *Energy and Technology Review* (including name and address changes) to Mail Stop L-3, Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, CA 94551, or telephone (510) 422-4859, or send electronic mail to etr-mail@llnl.gov.

Prepared for **DOE** under contract
No. W-7405-Eng-48

■ October 1994

SCIENTIFIC EDITOR
William A. Bookless

PUBLICATION EDITORS
Lori McElroy, Dean Wheatcraft

CONTRIBUTING EDITORS
Carolin Middleton, Stephanie Shang

WRITERS
Kevin Gleason, Robert D. Kirvel

GRAPHICS COORDINATOR
Paul Harding

DESIGNERS
Paul Harding, George Kitrinos, and Ray Marazzi

ART PRODUCTION
Treva Carey

COMPOSITOR
Louisa Cardoza

PROOFREADER
Catherine M. Williams

Printed in the United States of America
Available from
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Road
Springfield, Virginia 22161

UCRL-52000-94-10
Distribution Category UC-700
October 1994

Energy & Technology Review

The Industrial Computing Initiative

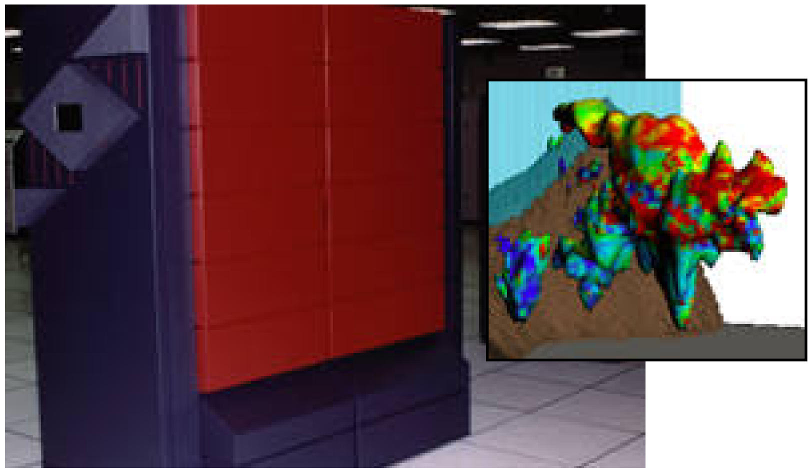

Our three-year, multiparty collaboration is addressing several different problems that have limited more widespread use of massively parallel computing by researchers in government, academia, and industry. The delivery of a set of tools and efficient applications that can be run on different machines will accelerate the use of high-performance parallel processing to increase U.S. industrial productivity and competitiveness.

FOR more than four decades, the strength of LLNL has been based on large-scale facilities and associated science teams working to make bold advances in science and technology. One of the areas in which we have become world renowned for our resident expertise is high-performance computing.

The history of LLNL intersects strongly with the history of computing. We have made important contributions in software, operating systems, scientific applications, and computing techniques. We also have a history of successful partnerships with private industry and other government laboratories.

Today, many computational science projects within the DOE laboratories and U.S. industry are facing a challenge. To move forward, the projects need to add further realism, which will come from augmenting physical effects, resolution, or dimensionality. Increasing realism will expand the demands on a computational resource by orders of magnitude. This demand has driven the movement toward the use of computers with multiple processors. These processors, working together, can rapidly solve a single problem. This concept is called parallel processing.

However, massively parallel computing has not been adopted as the high-end standard in computational research as rapidly as it might have because of the difficulty in creating efficient, high-performance parallel programs. The obstacles do not arise from a single source. Rather, they are due to deficiencies both in the hardware design and in the software programming environment of virtually all massively parallel systems. In addition, each vendor offers a unique architecture and creates unique software products. Such variations, along with other issues described in this article, now

Glossary

Central processing unit (CPU) The part of the computer containing the circuits required to interpret and execute the instructions.

CMOS Complementary metal oxide semiconductor.

Flops Floating-point (arithmetic) operations per second; a commonly used measure of the speed of calculation.

H4P High-Performance Parallel Processing Project.

ICI Industrial Computing Initiative.

Latency The waiting time between the issuing of a read instruction and the receipt of requested data.

Massively parallel processor A parallel processing machine with 100 or more microprocessor-based CPUs.

Microprocessor A single silicon chip on which the arithmetic and logic functions of a computer are placed. A typical microprocessor contained about 35,000 devices in 1982; recent ones contain about 3.5 million.

Node An intersection point in the communication topology of a massively parallel processor. The CRAY T3D has two processors per node, and the communication topology is a three-dimensional torus.

Parallel processor A machine that uses more than one processor running simultaneously to speed up the solution of a computational problem.

Pipelining The computer architect's version of an assembly line. Instructions are overlapped in execution, and a new operation is started every clock cycle even though it takes several clock cycles to complete one operation.

Porting Moving an application code from one computer system so that it runs on another of a different type, for example, from the Thinking Machines CM-5 to the CRAY T3D or vice versa.

Production environment All of the system components needed for a user to develop an application efficiently, debug it, execute it, and assimilate the output. These components include the computer itself along with one or more graphics workstations, high-speed networks, high-performance storage systems, and documentation, consulting, and software tools.

Scalar code An application that is not, or cannot be, vectorized is called scalar. (See vector processor.)

Supercomputer A computer that is among those with the highest speed and largest memory at any given time.

Vector processor A computer with hardware instructions that can each operate on a set of data elements, achieving high speed by streaming the set of elements through the hardware segments in a pipelined fashion. (See pipelining.) A code executing in this mode is called a "vectorized" code.

VLSI Very large-scale integration.

impede the movement of parallel applications across different platforms. Investment in the new technology both by government laboratories and private industry will accelerate as the difficulties are overcome.

The Industrial Computing Initiative (ICI) represents a collaborative effort by major industrial partners and government laboratories to develop applications targeted at parallel computers. Leverage can come from working in common; in particular, the ICI effort is of sufficient scale to allow for a general and more complete assessment of massively parallel technology.

The ICI is one part of a broader project called the High-Performance Parallel Processing Project (H4P). The H4P, valued at $66 million over three years, is funded by the DOE and private industry in a 50-50 cost-shared manner. The ICI portion of the project, in which the Laboratory is playing a key role, is valued at $52 million.

What Has Changed?

The current situation in high-performance computing has been compared to what troubled the railroad industry in the 19th century when different gauges of track prevented the transport of goods on different lines. Now, industry is experiencing the computational equivalent of that situation. To understand how we arrived at this place, some basic definitions and a brief look at the unprecedented growth of computing power over the last few decades will be helpful.

When the first commercial computers came on the market in the 1950s, each user had access to an entire machine—its processor, memory, and storage. In a sense, a user temporarily "owned" a machine while others waited their turn. By the mid 1960s, the idea of timesharing came into being. In this approach, the operating system directs the central processing unit (CPU) to work on several jobs in tandem. Instead of letting the CPU sit idle while a time-consuming step in one job is completed (often the input/output function), the operating system juggles jobs in and out of the CPU to make the most efficient use of a single processor and to give each user a sense of interactive control.

Modern conventional supercomputers provide not only time-shared access but also extremely rapid computational performance. One measure of performance is speed of calculation, usually expressed as Mflops, which is a million floating-point (or arithmetic) operations per second. Most current supercomputers achieve hundreds or sometimes thousands of Mflops by streaming a set of data elements through the hardware segments in a pipelined fashion.

Parallel computing represents a leap forward in the efficient and potentially flexible use of processor resources to solve computational problems. Parallel computing is simply the simultaneous execution of operations, transmission of information, or storage of data through the use of multiple processors. In massively parallel processing (MPP), large numbers of processors are used to attack an even larger set of tasks that together compose the problem to be solved. Whereas current conventional supercomputers still use one or a few powerful CPUs, massively parallel machines can have 100, 1000, or even more microprocessor-based CPUs. There is no fixed boundary between parallelism and massive parallelism, but a machine with more than 100 CPUs is generally called massively parallel.

The transition to massively parallel systems has been driven by the performance revolution in microprocessor technology. When personal computers came on the market some 15 years ago, they sparked a revolution. The "brain" in a personal computer consists of a simple computing device on a single silicon chip, the microprocessor. Early microprocessors were slow by current standards and were useful only for limited applications.

In the past decade, microprocessor technology has improved rapidly in both absolute and relative performance. One measure of absolute performance is the clock rate—the rate at which a processor operates, also called the processor cycle time. For microprocessors, the clock rate improved by a factor of about 50 from 1982 to the present. In contrast, the clock rate for conventional supercomputers improved over the same period by a factor of about 3. Microprocessor clock rates will exceed those of supercomputers within a couple of years.

In terms of relative performance, the story is similar. Figure 1 shows the evolution of performance for microprocessors versus conventional supercomputers.[1] In this graph, the measure of performance is the peak rate for floating-point operations on a common benchmark problem. The two curves show that microprocessors are closing the performance gap with supercomputers and will eventually take the lead. The comparison becomes increasingly sobering for traditionalists when the measure used for comparison is based on cost performance.

In a recent study conducted at LLNL, the performance of computers over the years was measured in terms of dollars per mips (millions of instructions per second).[2] Depending on the type of technology, it takes from one to a few instructions to complete a floating-point operation. Figure 2 shows that a sharp change in the slope of the cost-effectiveness curve occurred in the late 1980s.

We can conclude from such data that parallel machines, leveraging off commodity processors, have an inherent advantage in cost performance. Moreover, distributed-memory MPP machines using microprocessor technology have three other built-in advantages over current vector supercomputers. They offer higher peak speeds (for systems with larger numbers of processors), larger aggregate memory, and, in some cases, abundant software designed for the workstation market. Putting all this information together, the potential advantage for parallel computers appears to be overwhelming. For example, an examination of results from a widely accepted benchmark suite (the NAS Parallel Benchmarks[4]) shows that the latest parallel machines are superior for several, but not all, types of problems. However, the dominance of parallel machines has been slower to emerge and is not yet as overwhelming and complete as might have been expected from optimistic forecasts.

What Is the Problem?

We have seen that massively parallel machines are becoming increasingly cost effective relative to established supercomputers. Moreover, MPP systems can motivate larger, more sophisticated codes, in part, because they support a much larger memory than a conventional vector supercomputer. Why then isn't MPP technology already in wide use in the industrial sector? In a nutshell, the problem is a combination of unresolved technical challenges in both the hardware and software.

The most serious problem facing parallel computers is that each processor has a local memory. When one processor requires information available on another processor's local memory, a communication process must be effected. This represents an extra dimension of complexity—one that is built into a parallel machine—and is largely absent from a sequential monoprocessor design. To complicate matters, the computation and communication hardware components for parallel processing are not designed to cooperate from the outset but, instead, are engineered to cooperate after the microprocessor has been designed (for the workstation). Cooperation between separately designed hardware components makes communication less efficient than would otherwise be the case. This situation is unavoidable because microprocessors are commodities.

During a calculation, every problem is broken down into tasks or fragments. Unless all tasks represent completely independent problems, the processors assigned to a specific task will, at some point, need to have information computed elsewhere. This need for information raises issues concerning communication between the processors, its rate and its latency (how fast does the first word get back?). MPP programmers worry about synchronization between tasks, the layout of data across the separate memories of all the processors, data transfer rates, and processor idle time. In other words, it is easy to design a parallel program that is inefficient either through lack of insight or experience on the part of a programmer or through poorly designed, developed, and tested software and hardware in the target system. Unless a balance exists among processor speed, latencies, communication rates, and input/output capabilities, a system will lack generality. In this case, only a small class of problems (those that

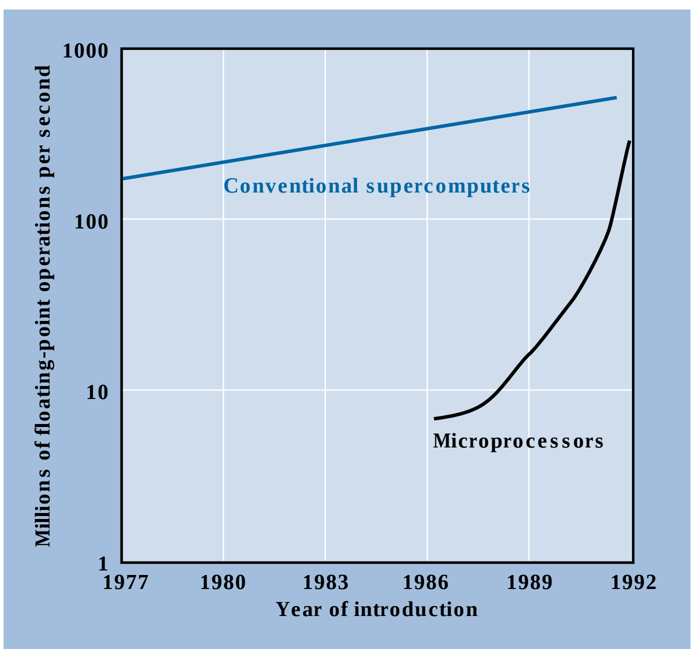

Figure 1. The evolution of performance for microprocessors vs conventional supercomputers central processing units.[1] The supercomputer curve shows a steady, but gradual, increase in performance over the last 15 years. Dramatic improvements in integrated circuit technology are allowing microprocessors to close the performance gap with conventional supercomputers.

do not tickle the system's weak points) runs well on that system.

Unfortunately, as the number of processors in a system increases, the number of tasks must increase as well, or some processors will be idled. With the increase, weaknesses in system design are aggravated and become more obvious, leading to decreased efficiencies, saturation, and, ultimately, defeat. The massively parallel processor lives in the niche of high processor counts, so it is the most vulnerable architecture with respect to the issue of balance.

Because each system must contend with a complicated architecture involving multiple processors, communication hardware, interconnect topology, and input/output processors, a vendor faces a myriad of design choices. As a result, each manufacturer today offers a system that is very different from any other competing system. The divergence of products begins at the hardware level. After layers of software, operating systems, programming models, debuggers, and performance-monitoring tools are added, the result is that codes designed for one system (even if ported with only a few weeks' work) will generally run poorly on another system until they are tuned. After all, the codes were not designed with the idiosyncrasies or weaknesses of some alternative system in mind.

To recap and add to the list of problems, the limitations and perceived weaknesses of MPP machines emanate from:

- Difficulty in programming efficiently and portably, especially for certain applications requiring extensive communication and abundant synchronization.
- The lack of sophisticated and time-tested support tools, such as debuggers and performance-analysis tools.
- Inadequate timesharing support, making it difficult for many users to share a machine efficiently for code development and short debugging runs.
- The lack of refined parallel physics, engineering, and chemistry application codes as well as efficient parallel mathematical libraries.
- The lack of efficient, standardized, and error-proof parallel input/output protocols and capabilities.
- A lag in processor capability. The latest microprocessor frequently

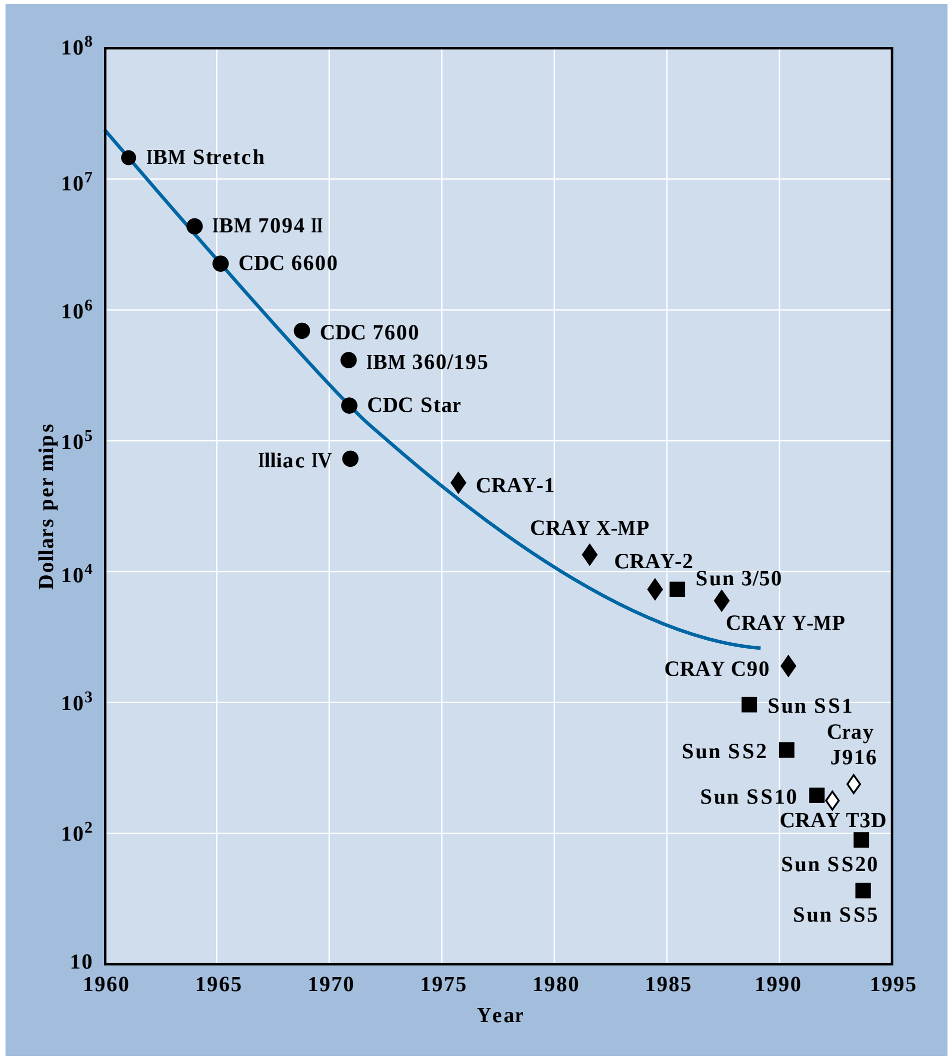

Figure 2. The performance of computers over the years measured in terms of dollars per millions of instructions per second (mips).[2] A sharp change in the slope of the cost-effectiveness curve occurred in the late 1980s. Solid circles are dollars per mips for conventional supercomputers from a study by R. Turn.[3] Solid diamonds are dollars per peak Mflops vs arrival time for supercomputers. (In this case, instructions per second is roughly equivalent to flops per second.) Solid boxes are the microprocessor-powered equipment from Sun Microsystems plotted as dollars per VAX mips as a function of year of availability. Open diamonds are the CMOS VLSI-powered equipment from Cray Research. The T3D is the CRAY distributed-memory massively parallel processor (using a single chip microprocessor) that is being used as part of the current ICI project. The CRAY J916 is a recent, shared-memory, CMOS-technology parallel processor.

does not reside in the latest MPP system because of the design time of the MPP itself.

Despite many years of development in parallel processing, major problems and challenges have effectively separated parallel processing research from the mainstream of supercomputing applications. Industry sees a gulf between what is required in real-world computing and the idealizations that have dominated the field of parallel computing research. Stated in simple terms, those companies that contemplate a move to massively parallel computing face a real dilemma: the risk of pursuing MPP, which entails a significant investment of resources, and the risk of falling behind international competition by not pursuing the most advanced computer technology.

U.S. industries simply do not have efficient applications, and they lack the necessary software tools and languages to facilitate code development in an MPP environment. Developing efficient applications that can realistically simulate complex problems and that have the potential to run on MPP hardware from different vendors is what the ICI is all about.

Figure 3. **The CRAY T3D after its recent installation at LLNL. This massively parallel computer includes a high-performance, 88-Gbyte disk subsystem. The memory of the T3D is logically shared, but physically distributed, within a three-dimensional toroidal array of nodes. All system memory is directly accessible to all 128 processors. Each microprocessor provides 150 Mflops of peak performance. The number of processors and the disk capacity will be doubled in December for the benefit of LLNL and University of California researchers.**

The Partners

The issues being addressed by the ICI demand an unprecedented level of innovation and cooperation from computer manufacturers, computer scientists, engineers, programmers, and participants from the industrial sector. Overall, the project will tap and put to use the expertise of more than 40 scientists from two national laboratories, six specialists in parallelization from Cray Research, two specialists from Thinking Machines Corporation, and at least 17 industrial scientists.

National Laboratories

The two national laboratories, Lawrence Livermore and Los Alamos, bring to the endeavor their expertise in supercomputing applications and MPP enhancements gained from large-scale simulation and modeling of complex phenomena. Through their weapons and supporting science programs, they have developed a unique science base that is necessary to solve large problems. The laboratories also offer extraordinary infrastructure support in both intellectual and physical terms. Thus, the two national laboratories are well positioned to fulfill the project's objectives and the DOE's mandate to transfer new technologies to the private sector.

The importance of computing at these facilities is revealed in the financial commitment to it. At LLNL, for example, some 10% of the nearly billion dollar annual budget is spent on the development of applications and system software, and nearly 10% of all LLNL employees work for the Computation Organization.

Cray Research

Cray Research has played a major role in the supercomputer industry for almost two decades and has a well-established collaborative relationship with the national laboratories. The company has made available for the project two T3D high-performance massively parallel computer systems (Figure 3). One of the two T3Ds is now sited at Livermore's National Energy Research Supercomputer Center (NERSC); the other is sited at Los Alamos' Advanced Computing Laboratory (ACL), a High-Performance Computing Research Center. These machines will eventually be linked together in a prototype distributed MPP computing environment.

The T3D hardware design permits a global address space implemented in the hardware. This design means that it is possible for a programmer to view the memory as shared, not distributed. A programmer must still worry about indiscriminate addressing of memory because an off-processor fetch still takes about six times as long as a local fetch. However, the machine's designers focused on the requirements of balance, and the machine's interprocessor communications network today defines the state of the art.

The T3D hardware sets the stage for significant software improvements. In particular, it permits efficient implementation of all three programming models that are now commonly used (the models are called message passing, data parallel, and work sharing). This feature is important because all three programming models can be implemented in one application. Thus, scientists can choose an appropriate combination of models that best fits the algorithms employed, rather than trying to fit the application to a single model provided by a hardware vendor.

Thinking Machines

Thinking Machines Corporation is a small, privately held company that established itself as a major MPP system supplier over the last 5 years. This company was the first to demonstrate the routine use of tens of thousands of processors on a single program. It was also the first to demonstrate a scalable programming standard for parallel computing. Its advanced massively parallel system, the 1024-node Connection Machine (CM-5) supercomputer system, is currently at the Los Alamos ACL and will be made available to the project team.

Industrial Partners

The industrial partners are major U.S. companies spanning a spectrum of commercial enterprises. Among the partners with whom LLNL is negotiating CRADAs are: AT&T, Alcoa, Boeing, Hughes, Halliburton, Areté, IT Corporation, and Xerox. The issues these companies need to address are ones of national concern. They range from protecting and preserving the environment to developing advanced materials, improving manufacturing techniques, and enhancing the nation's high-performance computing infrastructure. Some 17 scientists from industry will participate in the ICI.

Project Objectives

The central objective of the three-year project is to accelerate the development of parallel computing so that it can serve as an effective competitive tool for U.S. industry. This development will be implemented by a selected set of industrial applications that will have practical use and that are written in portable programming languages.

Whereas the effort addresses many different technical issues, the principal objectives fall into two broad areas:

- Critical applications development to address problems that are of interest to the laboratories and that industry has identified as top priority.
- Accelerated development of integrated production MPP environments, including the development of portability tools and standards that will allow applications to run on hardware from different vendors.

Critical Applications

We have chosen to demonstrate the range and depth of science that is accessible through MPP by developing a few key, representative applications. Some of the highly complex applications that fit this criterion are, for example, modeling pollution transport in the environment (Figure 4), tools to study materials around boreholes for petroleum exploration (Figure 5), advanced materials design (Figures 6 and 7), and advanced manufacturing problems (Figures 8 and 9).

Many of the applications we are addressing had their origins in defense programs. They will now serve a dual use in that enhanced versions of the codes can be moved back into defense work and serve as industrially useful software products as well. The box on pp. 10–11 summarizes the various applications on which we will focus at first. As the collaboration develops, work should expand to include a wider range of applications.

Complex applications codes usually consist of various modules, each of which treats a different aspect of the physics or other work, such as the numerical algorithm used to advance equations towards solution. We now know that each module or algorithm of an application maps best to a particular programming model. Nonetheless, very few applications programs actually use multiple programming models; efficiency suffers as a result of having to adhere to a single programming model because of limitations in hardware or in the software environment. As mentioned earlier, one of the many attributes of the newer machines is an ability to use two or three different programming models (such as message passing, data parallel, or work sharing) in one application. Thus, another aspect of developing efficient applications will be to make efficient use of available programming models.

Integrated Production Environments

A massively parallel computational environment is made up of many components, including graphics workstations, high-speed networks, high-performance storage systems, and software tools, as well as the supercomputers themselves. If a user can efficiently develop an application, debug it, execute it, and assimilate the output, a production environment exists. If a user cannot,

Figure 4. In one ICI effort, LLNL researchers are developing a global atmospheric chemistry model. Highly complex problems, such as environmental modeling, can be best addressed by the new MPP machines now arriving in the marketplace. This work can help guide industry in developing and analyzing products through a better understanding of their possible environmental consequences. The simulation shown here is a model of the distribution of ozone during the month of December at an altitude of about 20 km above Earth's surface. ▶

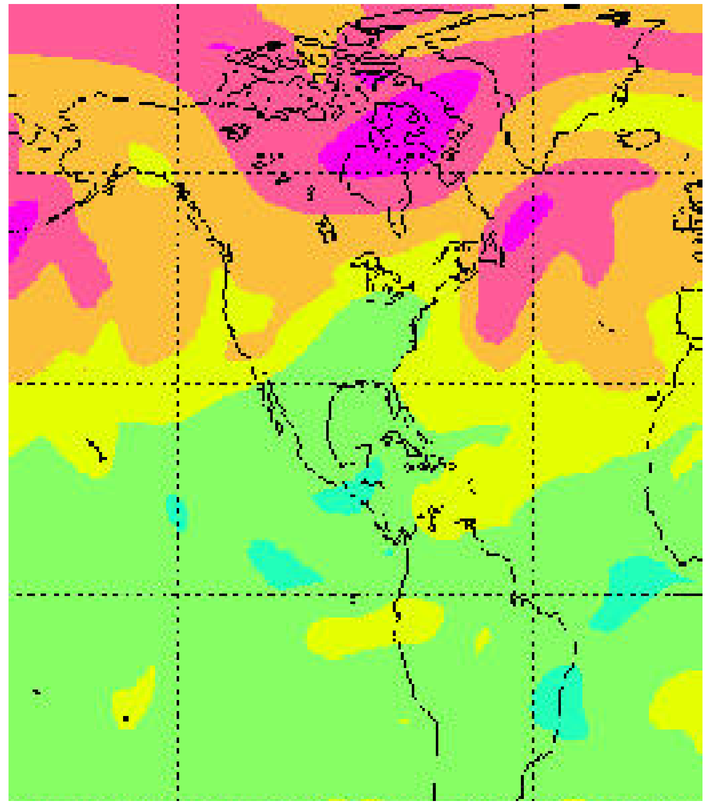

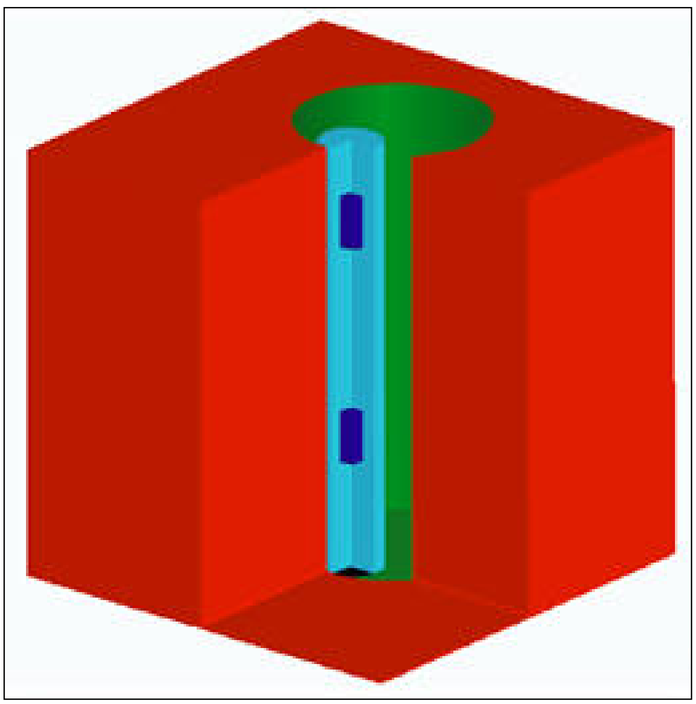

◀ **Figure 5.** We are developing MPP computer codes to calculate the response of nuclear logging tools used to study material surrounding exploratory boreholes in the quest for petroleum. This simulation shows the typical configuration of a borehole logging tool. Neutrons from a point source (black dot) at the bottom of the tool (light blue) are scattered in the borehole (green) and surrounding formation (red) before striking detectors (dark blue) on the tool.

Figure 6. Efficient materials modeling tools in MPP environments will benefit many sectors of the industrial community. Such tools can help to accelerate the synthesis–processing–fabrication–manufacturing cycle and ultimately contribute to the goal of materials by computer design. These illustrations show how electrons (red) interact with (a) defects and (b) impurities in silicon devices during the processing step. By understanding such interactions, we can better control processing temperatures and procedures to increase the efficiency of manufactured silicon devices. ▶

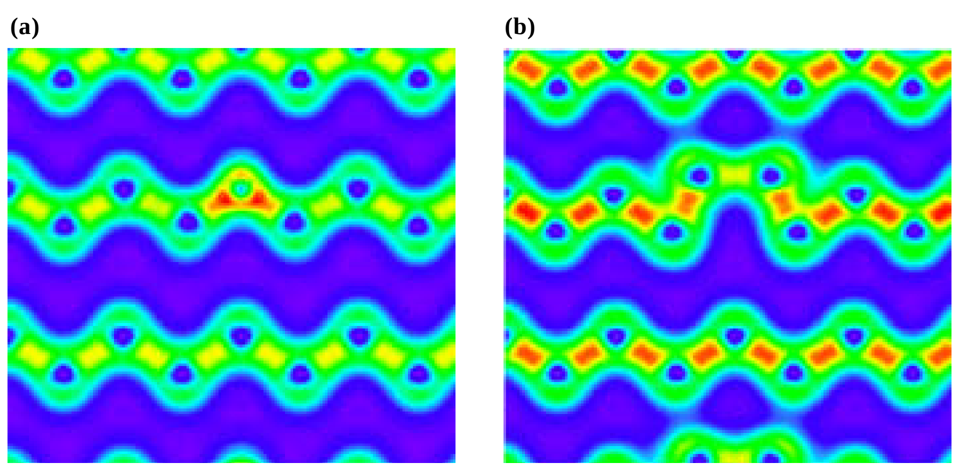

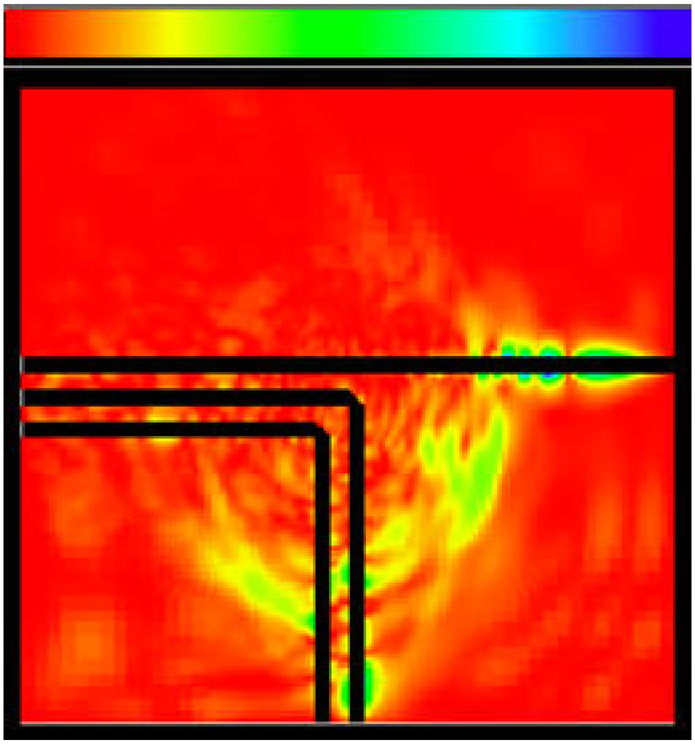

Figure 7. Another ICI project involves microwave components design. In this simulation, the color scale at the top represents the intensity of an electric signal, ranging from weak (red) to very strong (blue). The square below the scale represents a circuit board containing a three-way interconnect. Each microstrip conductor (black lines) is a few tenths of a millimeters in size. A test signal with a pulse width of two trillionths of a second should ideally propagate from left to right along the center conductor and follow the 90-degree corner without dispersion to the adjacent conductors. Instead, the electric fields show cross-coupling to the other conductors, particularly evident on the upper conductor (notice blue on the right-hand side). This type of simulation illustrates the difficulties that can arise in moving to higher-frequency signals that are needed for faster operation of devices.

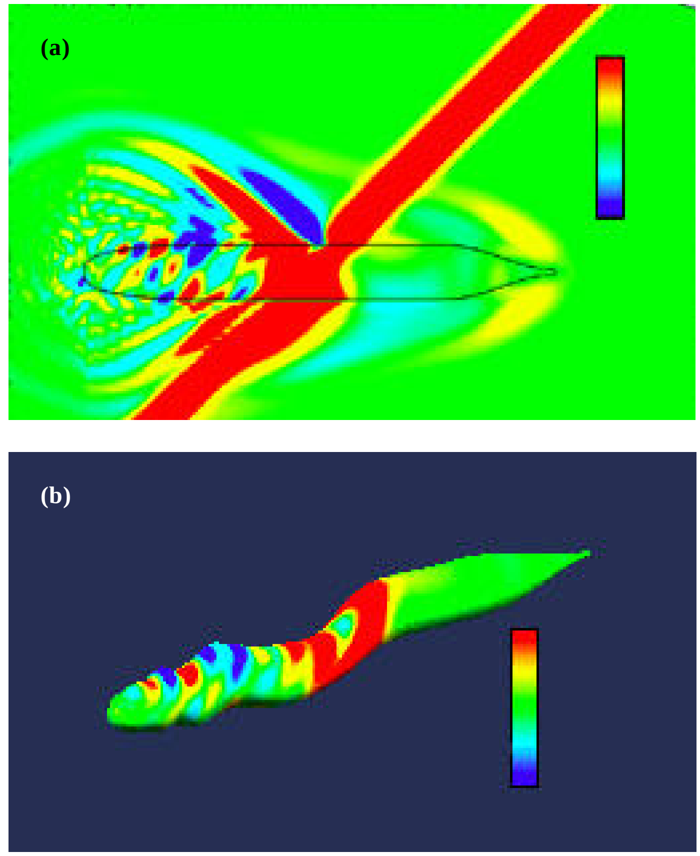

Figure 8. Our work in fluid–structure interactions will create a new simulation capability in the area of structural acoustics by the shipbuilding and aerospace industries. (a) In this snapshot of total pressure, an incident pressure wave (red) is midspan on a simulated elastic body (black outline). The pressure field highlights the presence of both a supersonic structural wave and flexural waves on the elastic body. (b) When the total pressure is mapped onto the surface of the elastic body itself, the exaggerated displacements help to show the presence of the flexural waves on the body.

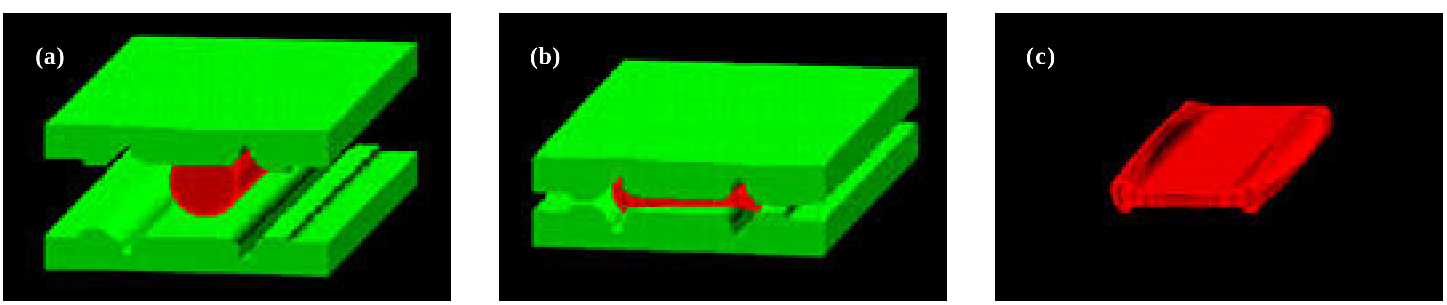

Figure 9. MPP codes for fluid dynamics and structural response will allow the U.S. metals industry to develop improved forging, extrusion, and casting processes. Here a three-dimensional code treats complex flows for a variety of materials. (a) At the start of a forging process, an aluminum rod (red) is positioned in a die (green). The forging process is represented in (b), and the final result is (c) an aluminum billet. Although the simulation represented here was run on a CRAY Y-MP machine, the code is currently running on the new CRAY T3D supercomputer at LLNL.

such an environment does not exist. At present, a production environment for parallel computers does not exist. As sophisticated as they may be, many industrial and DOE scientists do not enjoy working in a developing computational environment.

The ICI project will include a comprehensive effort to build an enhanced infrastructure for large-scale computing. The activities include developing high-speed network connections, a new multi-user scheduling system, code portability standards and techniques, performance enhancement tools, and visualization tools. A collaboration of systems and application developers will ensure that the final product is a fully tested MPP environment that will support the development of applications as well as large-scale production.

First, a mechanism is needed that allows many programmers to work on a machine simultaneously and to develop code efficiently. One answer is a space- and time-shared scheduler. Without a scheduler, it is impossible to halt a long-running program to allow others to use the processors without cancelling the original job. It is impossible to use machine partitions efficiently for code development because no mechanism is available to share the processors.

The DOE national laboratories have developed scheduler software for earlier massively parallel

What We Can Expect from the New MPP Applications

The Industrial Computing Initiative will initially focus on problems that can be best addressed and solved by the low-latency, tightly coupled, large-memory computer systems now arriving on the marketplace. Whereas the applications address a very broad range of topics, they share a common objective: helping the U.S. to compete more effectively in the global marketplace. All of the following topics are of interest to the DOE, the national laboratories, and the manufacturers of MPP machines as well. The benefits associated with each project are explained in the right-hand column. The nine projects involving LLNL are identified as blue type, and each one is described in this article.

Environmental Modeling

Lithography characterization for remediating underground pollution	Will reduce the costs of characterizing and remediating sites with underground contamination.
Unstructured grids for three-dimensional (3D) representation of heterogeneous materials	Will help industry cost-effectively design internal combustion engines, contain pollution, remediate groundwater contamination, and treat contaminated surface water.
Subsurface flow and chemical migration	Will allow industry to characterize the migration and transformation of contaminants in soils at waste sites. A new 3D simulation capability will improve the design and management of engineered remediation strategies.
Global atmospheric chemistry models	Will guide industry in developing and analyzing products through a better understanding of their possible environmental consequences.

Petroleum Applications

General reservoir simulation	Will allow oil exploration and oil service companies to perform simulations at fine grid resolutions.
Nuclear well logging	Will provide computer codes to calculate the response of nuclear logging tools used to study material surrounding exploratory boreholes.

machines—in particular, for the BBN TC-2000 as part of the LLNL Massively Parallel Computing Initiative. A similar, general scheduled environment is planned for the T3D, as shown in Figure 10. Developing a scheduler rapid enough for interactive code development will be complicated by the inability of the T3D processors to context switch (this means that multiple processes cannot remain resident in memory). Therefore, an entire job must be swapped to another system (called the front end) or to disk to allow access to the processors by waiting processes. Although swapping is possible, it is also more time consuming than context switching, so the resulting environment will probably be less than optimal for a programmer interested in expeditious debugging. Nonetheless, much will be learned from the process of developing, implementing, and testing the scheduler. Even with limitations, a scheduled environment will represent a considerable improvement over a strictly space-shared environment. In this part of the project, Cray Research will develop the checkpoint and swapping software, and LLNL will develop a prototype scheduler along with the X-window interface that will allow a user to see which jobs are resident, which ones are waiting, and which processors are being used.

Materials Design	
Materials modeling	Will enable industry to simulate processes, such as casting and welding, that are involved in making new materials.
Advanced materials design	Will make available to industry advanced materials modeling codes. Codes will be ported, optimized, and integrated into Cray Research's UniChem computational chemistry system.
Microwave components design	Will provide state-of-the-art microwave simulation exploiting the latest advances in computing capabilities.
Advanced Manufacturing	
Hydrocode library	Will address specific hydrodynamic simulation problems; make practical the analysis of 3D, reactive, multiphase flows common in U.S. industry; and provide adequate resolution of physical and chemical models.
Fluid–structure interaction	Will create a new time-domain simulation capability for structural acoustics by the shipbuilding and aerospace industries.
Fluid dynamics and structural response	Will allow the entire U.S. metals industry to develop improved forging, extrusion, and casting processes.
Shallow-junction devices	Will allow predictive modeling for the fabrication of next-generation semiconductor devices, reduce development time, and reduce cost.
Tools to Maximize the Use of MPP	
Dynamic time-sharing scheduler	Will create a more user-friendly environment for debugging and production.
MPP performance measurement	Will allow for the more efficient use of high-performance computers through better data collection, analysis, and visualization tools.
Portability tools	Will provide a means for writing portable applications, thus removing a major barrier to industrial use of MPP computers.

Workers at Los Alamos will address the need for portability tools. On conventional computers, users are accustomed to applications that can be moved to (ported) and executed on a variety of computers and workstations. In contrast, portability standards do not yet exist for massively parallel machines. Users must either execute their programs exclusively on one vendor's hardware, or they must maintain separate sources for separate machines. An inability to write portable code and the lack of portability standards and tools have been major barriers to the adoption of massively parallel machines by industry.

A broad objective of the ICI is to establish a common set of conventions for programming massively parallel machines, including common mathematical library calls, common message-passing interfaces, common parallel extensions, and so forth. The basic idea is to ensure that there is a common subset of languages and libraries that are supported on MPP machines. Portable applications will not only benefit the vendor and scientific communities, but they will also remove a major barrier to the industrial use of MPP computers.

Potential Results and Benefits

Benefits to U.S. Industry

The single, most important achievement of the project will be physically realistic simulations using advanced models that run routinely on massively parallel computers. Ultimately, such applications can be used by industry to help the U.S. to compete effectively in the global marketplace. The applications, concepts, and tools that will be transferred to industry will enhance competitiveness by reducing product cycle times, producing higher-quality products, requiring fewer prototypes, and developing more efficient industrial processes.

Another possible gain from the ICI is best understood from the perspective of an industrial partner in the project. All of the industrial partners will have invested significant staff time in developing certain applications. However, from their point of view, the ICI still represents a high-leverage situation. As the codes evolve, industrial participants will be able to judge for themselves whether or not the applications provide a competitive advantage. If they do, then the partners will have employed an optimal strategy to uncover what might otherwise have been an expensive question to answer.

With the considerable cost-shared contribution by Cray Research, the infrastructure for this project is provided by the laboratories at no cost to the other industrial partners. In addition, the laboratories are making available a constellation of computational experts and specialists in parallel computing. Therefore, it is possible for a company to develop a large-scale application and determine its usefulness without the major investment associated with building a local infrastructure—an investment involving tens of millions of dollars. With the development of a major code, a partner is free to decide if an investment in a dedicated machine for the purposes of production is beneficial. In effect, the expensive research and development costs and risks associated with such work will have been reduced to nearly zero.

Over the short term, our applications will address a few, critical industrial problems—ones also of interest to the laboratories—that have been only partially solvable to date. The box on pp. 10–11 summarizes the results we can expect from each application. The benefit of each one is substantial; taken as a whole, the benefits are as broad as the spectrum of individual topics. The payoffs range from increased cost

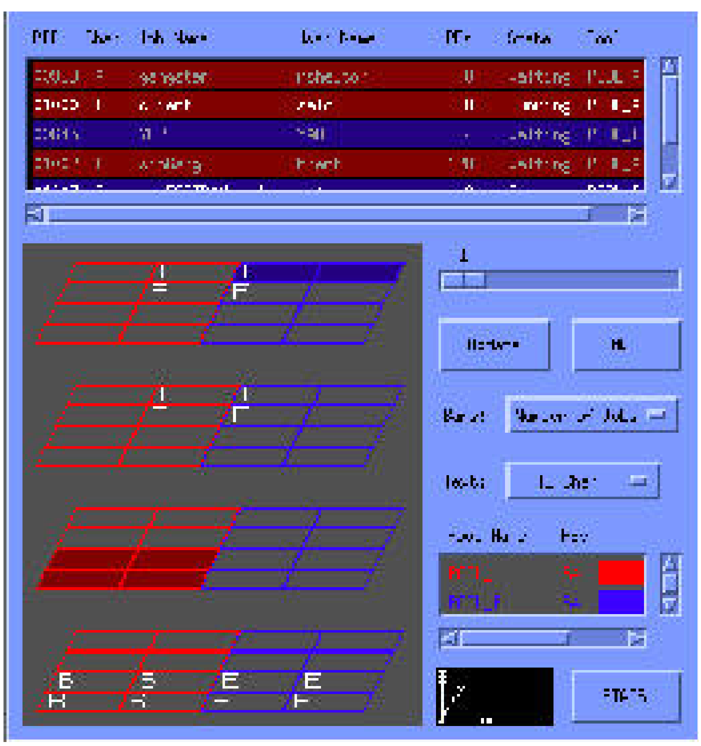

Figure 10. The Livermore Gang Scheduler will provide time-sharing support for massively parallel processor architectures. Such support greatly enhances the development environment on the most advanced systems and allows work to continue on production codes. In the display at the top, the status of each job and which processors are being used can be seen at a glance.

effectiveness in the design stage of a particular product or device to greater realism in three-dimensional (3D) models of chemical transport in soils or through the global atmosphere.

For example, many U.S. industries and the DOE face the enormously complex and expensive task of cleaning up contaminated soils and groundwater at thousands of facilities. Traditional two-dimensional models cannot adequately characterize the variations in subsurface properties that affect how contaminants migrate. In one of our research projects (described in more detail on pp. 24–25), we will develop a sophisticated 3D model to rapidly and accurately simulate fluid flow and chemical transport in heterogeneous, porous soils around waste sites.

Another effort in which LLNL and AT&T researchers will work together combines accurate simulations at the atomic scale with modern 3D visualization techniques to model diffusion mechanisms in silicon (Figure 11). Predictive modeling for the fabrication of next-generation semiconductor devices will reduce development time and cost. This work could have a significant effect on the U.S. semiconductor manufacturing industry.

Benefits to the DOE

The DOE's national laboratories have invested hundreds of millions of dollars in the computational software and hardware needed for their energy and defense missions. The phenomena that are modeled—often through techniques developed at the laboratories themselves—include materials properties, structures, pollution migration, charged-particle transport, shocks and detonation, fluid dynamics, acoustic and elastic wave propagation, combustion, plasma physics, nuclear physics, and electromagnetism.

The ICI serves a dual-use function. Applications originally created specifically for DOE programs will now be transformed into efficient massively parallel applications for use by U.S. industry. The modified and enhanced codes can be simultaneously reintegrated into ongoing DOE programs. For example, the parallelization of existing codes will serve the reliability and safety programs at LLNL. Furthermore, scientists not directly involved in the ICI will have access to the new techniques and machines for their own code development (see the box on p. 14). Such access and the acquisition of skills associated with new computational techniques can facilitate a wide range of ongoing DOE programs, including genetics research with its need for vast computational resources, environmental modeling, and weapons-related work.

To carry out its mission related to the Comprehensive Test Ban, the DOE requires more extensive use of advanced simulations. The ICI gives the DOE one more way to evaluate a possible path to massively parallel computing through cost sharing with computer vendors and other project participants. In addition, the kind of production environment we will develop can serve as a model for an MPP production environment on classified machines.

Benefits to All

A unique aspect of the ICI is that it will create an infrastructure that couples the most advanced MPP resources at the two national laboratories. The prototype interconnection by a high-speed network will allow researchers at either the Livermore or Los Alamos locations to access the resources of both sites using distributed computing techniques. Not only will the prototype environment increase the total computing capability at each location, but it will also give industry the benefit of experimenting

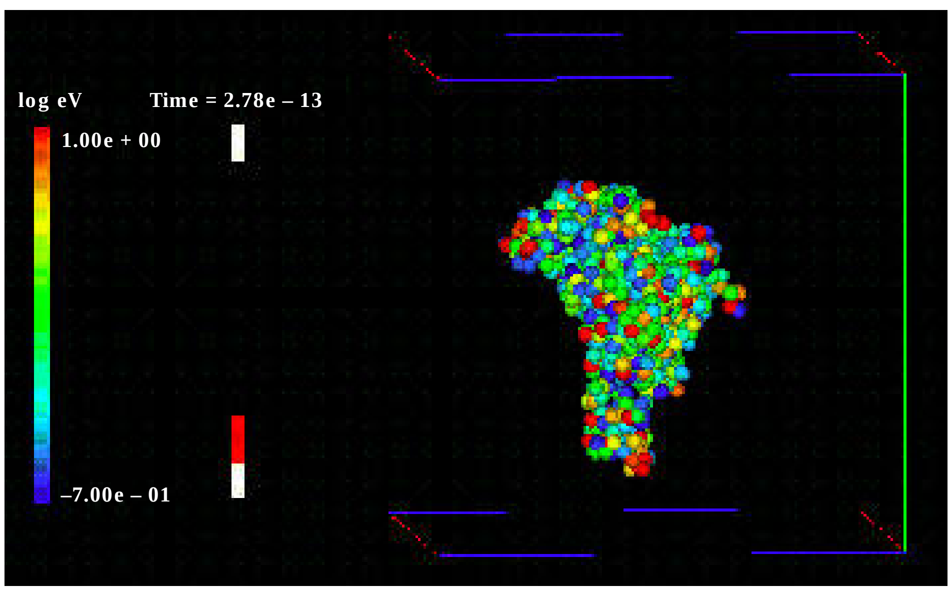

Figure 11. Predictive modeling for the fabrication of next-generation semiconductor devices will reduce development time and cost. Here, a silicon crystal is bombarded by an ion, and color is used to show the energy distribution. When doping crystals, this type of simulation helps us understand how atoms are displaced and how defects arise.

Building on the ICI

The machine resources at LLNL for the Industrial Computing Initiative (namely, the 128-processor CRAY T3D machine with 88 GB of disk) were chosen to accommodate nine different industry projects at LLNL. However, we will have little capacity to accommodate more ICI-like projects. To add worthy projects, such as the grand challenge problems in tokamak physics and climatology mentioned in this article, LLNL has decided, as an institutional project funded through the Director's Office, to double the number of processors and the disk capacity on the T3D. Sixteen processors will also be added to the 48-processor Meiko CS-2 parallel computer sited in the open environment.

The T3D enhancement will be provided in the same cost-shared manner as the first 128 processors, through a special arrangement with Cray Research. The resources available on two parallel computers will allow for ample expansion of ICI-like projects and for access by researchers at LLNL and within the UC system interested in developing high-performance massively parallel research projects. The environment across the two machines will be coordinated to allow, for instance, file sharing through a common file system as well as some common programming models. In this manner, researchers engaged in code development will be able to migrate and port across systems with minimal disruption. Some sharing of the tertiary storage environment at the Open Computing Facility is also planned. The Laboratory is particularly interested in using the enhanced computational resources to develop additional partnerships with industry and in developing collaborations with researchers at the UC campuses.

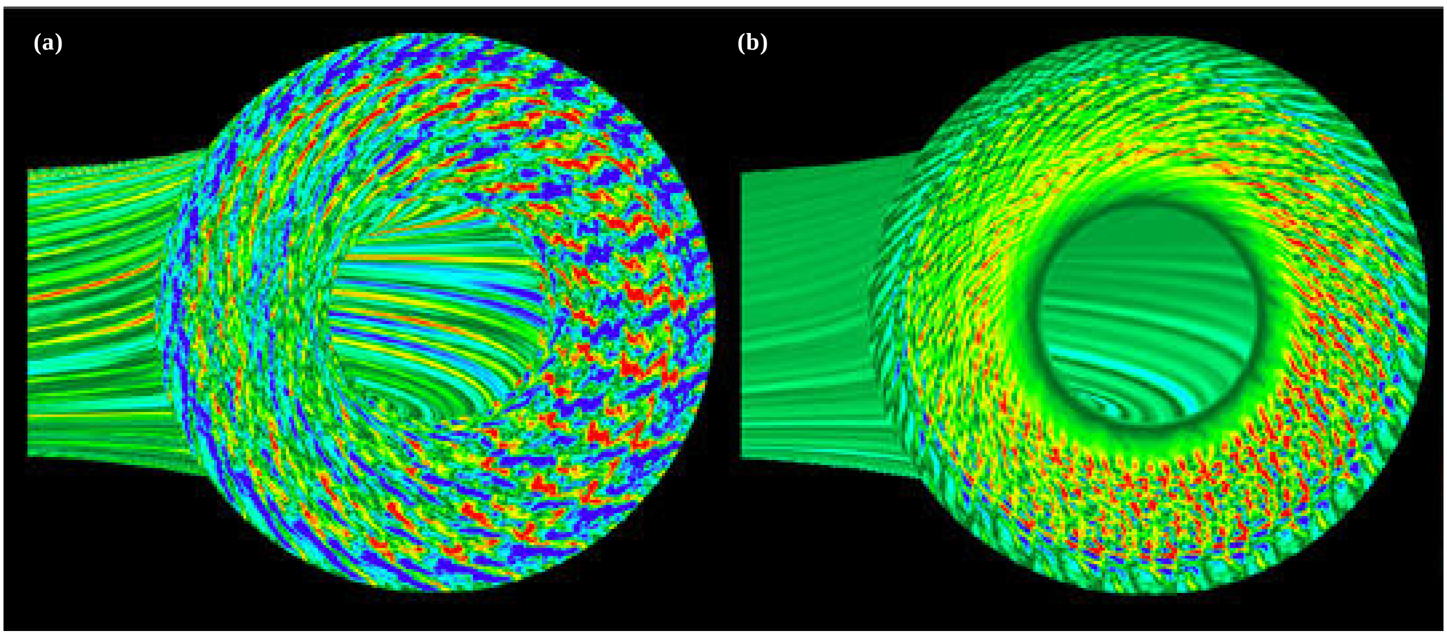

Figure 12. The Numerical Tokamak (NT) Project represents a grand challenge problem in the area of plasma physics. Along with other NT consortium participants, we are developing advanced computational models of tokamak physics using the most powerful high-performance computers. (a) This physical model shows how fusion ignition is prevented by the turbulent mixing of hot core plasma with cooler edge plasma. Red represents positive density perturbation; blue corresponds to negative perturbation. The view shows transient turbulent structures in cross section. (b) Numerical models explore the conditions that can reduce turbulence. Although this work is not part of the ICI, tokamak physics codes and those for other grand challenge problems will benefit from what is learned through the current ICI effort. (This calculation was done on a CM-5 machine at the Los Alamos ACL by LLNL researchers; results are displayed with a parallel ray-tracing technique developed here.)

with a distributed MPP computing environment before making the investment required to implement one. Coupling several dispersed machines to attack a single, but very large, application is of great interest to industry and the nation as a whole because it represents an economic approach to addressing large-scale computational needs.

Learning how to develop efficient parallel codes will also allow us to better address the problems identified as "grand challenges." These problems are of such generality and complexity that they can be addressed only on large parallel computers; standard vector machines do not have sufficient speed and memory, and loosely coupled multicomputers do not have sufficient interprocessor communication efficiency.

An example of a grand challenge problem in the area of plasma physics is the Numerical Tokamak Project. In the 21st century, magnetic fusion could be the source of large amounts of electricity without contributing to global warming or acid rain. The U.S. goal is to have a magnetic fusion Demonstration Power Reactor on line by the year 2025. Now, a multi-institutional, multi-disciplinary collaboration, including researchers at LLNL, is developing some of the most advanced computational models of tokamak fusion devices to identify the most cost-effective design (Figure 12).

Another grand challenge problem involves global climate modeling. As part of a major international effort, LLNL investigators are trying to determine why different global climate models give dramatically different results even when they use the identical set of data. Grand challenge problems, such as the Numerical Tokamak Project, global climate modeling, and others, have great significance in addressing society's problems and needs. Although these particular problems are not a part of the ICI, grand-challenge codes will benefit from what we learn through the ICI.

Summary

The widespread use of massively parallel computing as a scientific and industrial tool has been impeded by technical problems associated with available hardware and software. U.S. industries do not have the applications and necessary software tools and languages to facilitate use of MPP environments. To address this need, the Industrial Computing Initiative, which is part of the High-Performance Parallel Processing Project, is developing a set of efficient applications that can realistically simulate complex problems and are written in a way that allows them to run on MPP hardware from various vendors.

The ICI involves more than 40 scientists from Lawrence Livermore and Los Alamos national laboratories, six specialists in parallelization from Cray Research, two specialists from Thinking Machines Corporation, and nearly a score of industrial scientists. The delivery of a set of efficient parallel applications, serving as guideposts for subsequent work, can help U.S. industry compete more effectively in the global marketplace. The new applications, together with improved methods and tools, can reduce product cycle times, produce higher-quality products, and speed the development of more efficient industrial processes. At the same time, the new computational advances in massively parallel computing technology can be reintegrated into ongoing programs at the national laboratories to serve the missions defined by the DOE and to address national needs.

Key Words: computer performance; grand challenge problems; High-Performance Parallel Processing Project (H4P); Industrial Computing Initiative (ICI); massively parallel processing (MPP); parallel processing.

Notes and References

1. F. Baskett and J. L. Hennessey, "Microprocessors: From Desktops to Supercomputers," *Science* **261**, 865 (1993).
2. This graph, developed by LLNL's Eugene Brooks, is an extension of a projection first made by R. Turn (Ref. 3) 20 years ago and more recently used by Baskett and Hennessey (Ref. 1).
3. R. Turn, *Computers in the 1980s* (Columbia University Press, New York, N.Y., 1974), p. 80.
4. D. Bailey et al., *The NAS Parallel Benchmarks*, NASA Ames Research Center, Moffett Field, CA, RNR Technical Report RNR-94-007 (1994); D. H. Bailey et al., *NAS Parallel Benchmark Results 3-94*, NASA Ames Research Center, Moffett Field, CA, RNR Technical Report RNR-94-006 (1994).

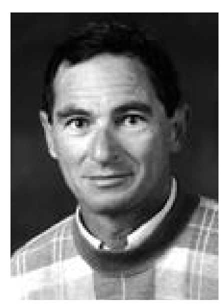

For further information contact Michel G. McCoy (510) 422-4021.

Artificial Hip Joints: Applying Weapons Expertise to Medical Technology

Oxidized zirconium is an excellent material for the artificial femoral heads used in hip replacements, but it cannot be easily shaped by conventional grinding. We drew on our expertise in high-precision machining to show one manufacturer how to apply precision, single-point turning to the manufacture of zirconium femoral heads.

MATERIALS fabricators at LLNL have spent several decades developing sophisticated precision machining techniques, such as single-point diamond turning, in connection with past weapons research. More recently, we are using these skills to assist U.S. manufacturers in meeting technological challenges. This article tells the story of how LLNL machining specialists transferred expertise gained through weapons research to the field of medical technology—particularly the manufacture of improved, innovative prosthetic hip joint replacements.

The Medical Problem

Hip replacement has become increasingly commonplace as the medical technology has improved, as life spans increase, and the population ages. It is most commonly performed to reduce the pain and immobility of arthritic patients over 55 years of age. In 1993, an estimated 129,000 people received total hip replacements in the U.S.[1] Sales of the prostheses alone generated between $600 million and $800 million in 1993.

In a total hip replacement, the pelvic bone must be cut away and a polyethylene cup fitted to receive the artificial femoral head.[2] This head is often attached by an artificial stem to the altered femoral bone. Unlike most other forms of surgery, this procedure is followed by an intensive course of antibiotics and immunosuppressives. Hip replacement surgery is not to be undertaken lightly; yet the nonsurgical alternative (medication) for long-term management of arthritis is very expensive over the long term and may leave the patient debilitated or disabled.

Artificial hip joints, however, undergo a very gradual wear process and have a limited lifespan. The hip joint transfers very large loads. The hip joints carry the weight of the upper body, and the legs transfer the forces of locomotion to the torso. In spite of these conditions, an artificial hip joint can perform well for 15 years or more, depending on type and level of activity. The younger the patient at the time of replacement, the greater the likelihood a repeat operation may have to be considered in a normal lifespan.

Moreover, as the polyethylene cup wears and third body debris is produced, the metal surface of an artificial femoral head may become rougher. This roughness in turn may promote further abrasive wear of the polyethylene cup. (Figure 1 shows an exaggerated schematic of roughness.) Wear produces debris. Although the debris particles are very small, usually less than a micron across, they are not tolerated by the human body. The body's normal cell growth and regeneration decrease in their vicinity.

Eventually, then, the question of another replacement arises, but the decision is complicated by such factors as the age and health of the patient and the cost of the artificial piece and of the operation. There are therefore strong health and economic incentives to lengthen the useful life of artificial hip components.

The Materials Solution

Smith and Nephew Richards, an established firm based in Memphis, Tennessee, with a wide inventory of prosthetic implants and synthetic replacements for bones in the body, is a major supplier of artificial hip joints. Historically, they have offered three artificial femoral heads, one made of cobalt–chrome, the others of a solid ceramic (alumina and zirconia, Al_2O_3 and ZrO_2). The ceramic is extremely wear-resistant and wears away the polyethylene cup much more slowly than the cobalt–chrome head but is more expensive to manufacture. Smith and Nephew Richards' researchers were aware that the metal zirconium can be oxidized to acquire a hard zirconia ceramic surface. They reasoned that if zirconium metal could be shaped to exact dimensions by machining, they could combine the best of both technologies: they could fabricate the femoral head at a relatively lower cost than that of a solid ceramic but achieve the same wear resistance. Zirconia's lower coefficient of friction than cobalt–chrome's for lubricated sliding against polyethylene and thus lower rate of wear on the polyethylene cup is expected to lengthen the useful life of the unit. Even if it lasted only a few years longer, it could mean the difference, for a hip-replacement patient in his or her thirties or forties, between facing one rather than two additional replacement operations in a lifetime.

The Process Problem

Initially, Smith and Nephew Richards tried unsuccessfully to use traditional hardware and conventional techniques to grind the zirconium into the spherical femoral head. Grinding typically works best on hard, brittle materials, but zirconium is a relatively soft, ductile material. Experience has shown that ductile materials are difficult to grind; the material builds up on the grinding wheel and produces a rough, galled surface on the workpiece.

Some of the initial trials on zirconium parts that Smith and Nephew Richards' fabricators attempted to grind were rough and

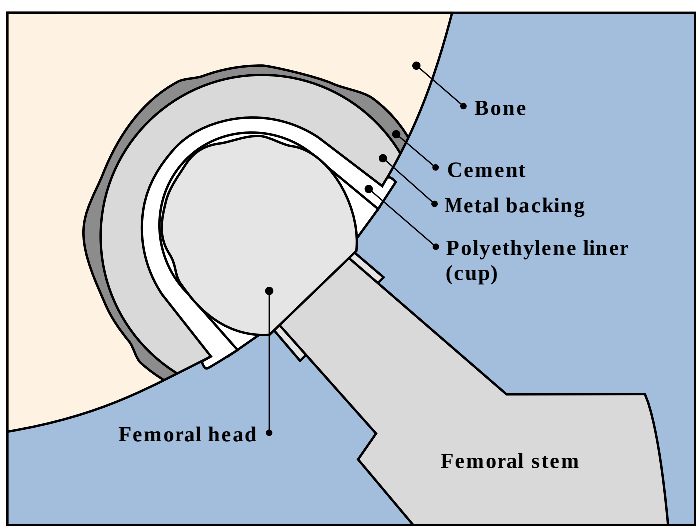

Figure 1. Schematic showing the geometry, close up, of an insufficiently smooth ball (head) in a cup. The roughness is enormously exaggerated to show how the head fails to distribute pressure uniformly, creating pressure points that will become wear points.

out-of-round, and the grinding wheels lasted minutes instead of their normal wear life of hours or days. When consulted, neither the grinding machine manufacturer nor the grinding wheel manufacturer had a solution.

Through hearing about our CRADA (cooperative research and development agreement) with COM (the Consortium of Optical Manufacturers), Smith and Nephew Richards learned that we at LLNL have decades of experience grinding and that we have a comprehensive grinding capability. They described their problems to us, and we visited their plant to better understand the problems and to see first hand what type of equipment they use and products they manufacture. We signed an LLNL research agreement with Smith and Nephew Richards to do the work under the U.S. Department of Energy's National Machine Tool Partnership (an initiative engaging the efforts of several national laboratories to improve the competitiveness of U.S. firms), and set about evaluating their fabrication processes. Our research on zirconium led us to the conclusion that grinding was not the most effective material removal process.

The Process Solution

Since the femoral head is spherical, we used our expertise in precision turning acquired through years of fabricating spherical components for weapons research. In simple terms, precision turning is a point-defined process that draws a single hard, pointed tool across a rotating surface

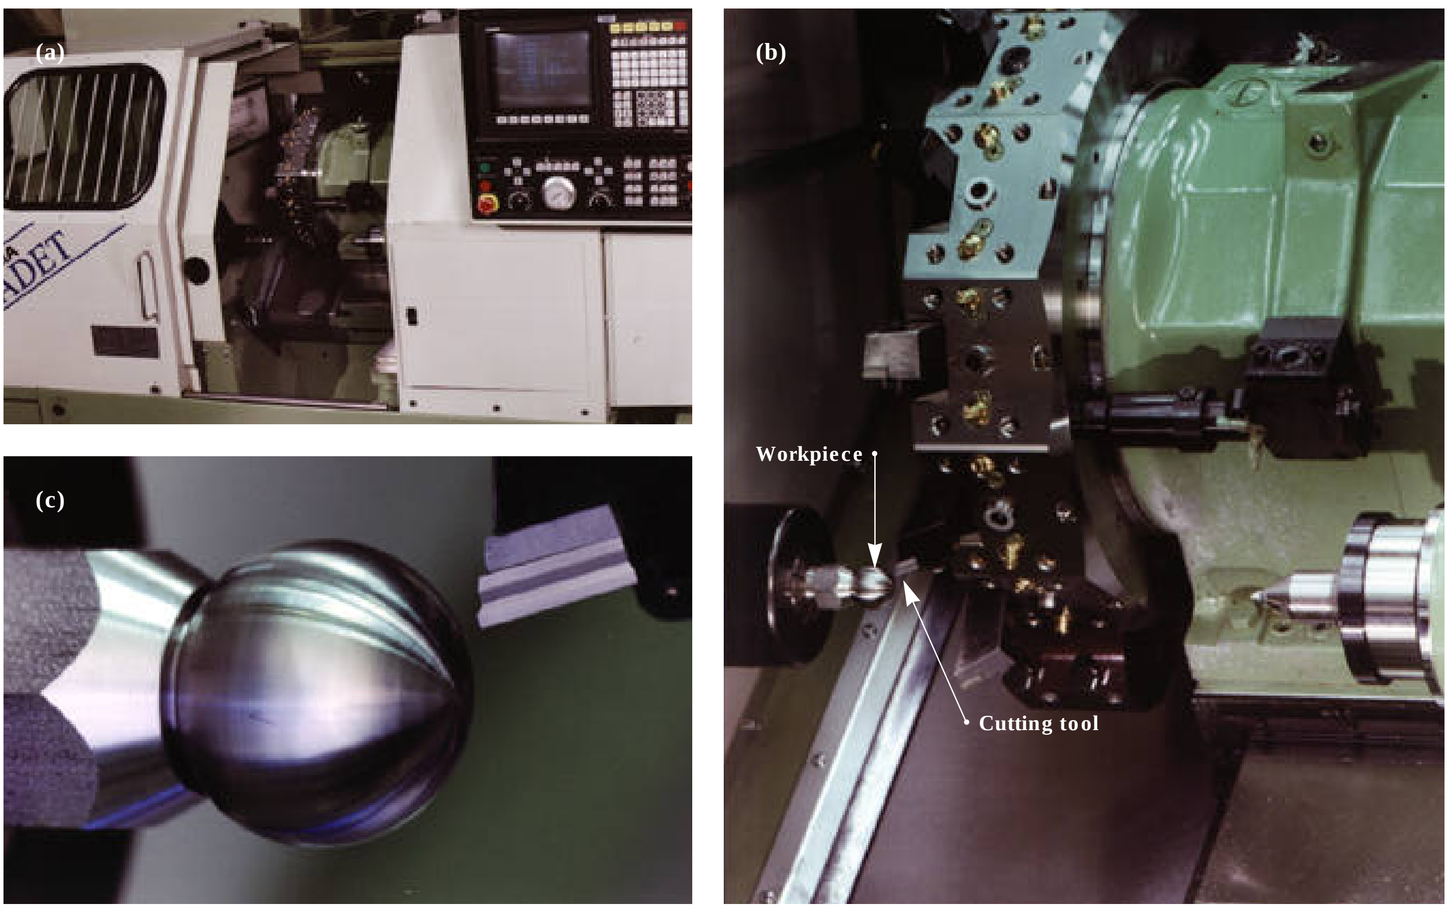

Figure 2. Three views of the lathe, identical to a machine in the Smith and Nephew Richards shops. (a) The safety enclosure has been rolled open to reveal the workbay; also visible are the computer numerical controls. (b) The spherical workpiece, lower left, is dwarfed by the tool carousel, which can hold a dozen cutting tools at a time (five tool holders are visible). (c) A workpiece and a cutting tool. The two are brought into contact by precise control of their relative positions, and the workpiece is rotated against the cutting tool.

in a highly controlled manner. For our weapons work, we have had to achieve contour accuracy to within 25 to 50 nm (1 to 2 millionths of an inch or 1000 times smaller than a human hair). We also developed the compensated tool paths and inspection hardware and processes to produce these very accurate spheres.[3]

The Smith and Nephew Richards' assignment was an excellent opportunity to transfer our knowledge directly from weapons work to an application in industry. Through testing, we confirmed that the single-point turning process would work for their purposes. We requested a list of the machine tools that Smith and Nephew Richards possessed and cross-checked it with the tool inventory at the Lab. We identified one semiprecision machine tool, a CNC (computer numerical control) lathe, that could make the parts and that was present in both the firm's plant and at the Lab. We had to demonstrate not only that single-point turning could machine the ductile zirconium, but also that the lathe could achieve Smith and Nephew Richards' standards of size, contour, and surface finish.

Contour is a very stringent requirement for Smith and Nephew Richards' application. The relationship between the contour of the femoral head and the contour of the polyethylene cup in which it rides greatly determines the rate of wear. Pressure within the joint has a first order effect on wear. Controlling contour errors will provide reduced (uniform) pressure by maximizing the area of contact. Smith and Nephew Richards' dimensional requirements for contour accuracy were on the order of 5 μm. The company wanted the femoral head to have an R_a (roughness average) surface finish of about 0.25 μm *before* it is polished. Polishing then produces about a 25- to 50-nm R_a finish before the head undergoes the oxidation process that creates the ceramic layer of zirconia. (Just as ductile materials, such as zirconium, do not submit well to grinding, they are difficult to polish; they tend to sleek and gall. Accordingly, pre-oxidization polishing is done more lightly than post-oxidation polishing.)

After the zirconium head is polished, it is oxidized at an elevated temperature and forms a zirconia (ZrO_2) layer. (Oxidation turns the outer zirconium layer into zirconia; the process is accelerated by heat and the presence of certain gases.) Oxidation does not make the head smoother, only harder, and makes the resulting zirconia, like its ceramic counterpart, easier to polish to a very smooth surface.

Using their material, we produced femoral heads and other workpieces to their specifications on the CNC lathe (see Figure 2). We used cutting tools that were made or altered to the Laboratory's specifications (such as more exacting standards for tool nose contour) and that have since become readily available to the precision machining industry. We used a two-step turning process. The first step results in a contour that falls short of the client's specifications by a factor of two. The second step encompasses a process that LLNL has refined using inspection feedback to produce a compensated tool path: that is, a tool path that corrects for all repeatable machine errors. Typical machine errors might include nonstraightness and misalignment of machine slideways. The two-step process is the key to producing workpieces with the accuracies down to the limitations of the machine.

Ultraprecision Measurement

The femoral head shown in Figure 3 is one that we fabricated and sent to Smith and Nephew Richards for polishing. They then returned it to us so that we could measure it to be sure the polishing had not degraded the contour. Our equipment is among the world's most capable for measuring contour, size, and surface roughness. Figure 4 shows one of our devices, a small-radius gauge that measures contour by means of a linearly variable differential transformer (LVDT) probe. The accuracy of the small-radius gauge system is 25 nm. Figure 5 shows a contour error plot generated by the small-radius gauge's data acquisition

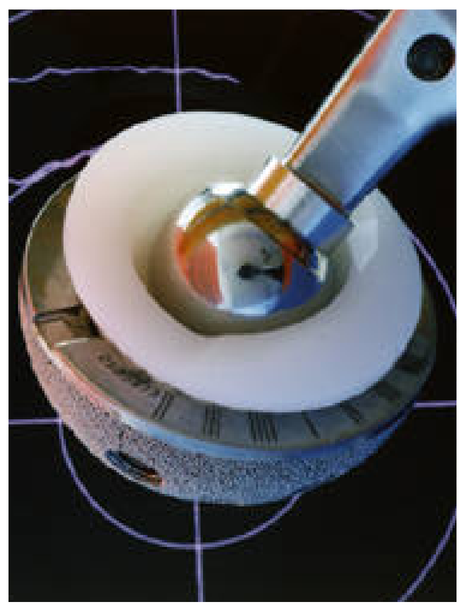

Figure 3. **This zirconium head has been oxidized to acquire a zirconia outer layer, which was then polished. The head has a 28-mm diameter. The stem is a titanium alloy (Ti-6Al-4V).**

computer and a surface-roughness trace generated from a profilometer.

Having sent them back parts that we had measured for size, contour, and surface roughness, we undertook a secondary collaboration with Smith and Nephew Richards through the National Machine Tool Partnership Program. We developed a set of standards for Smith and Nephew Richards; we also performed measurements on ten parts (cobalt-chrome and ceramic femoral heads) that they supplied to us. They used these measured pieces to cross-check their inspection equipment and to improve further the accuracy of their measurements of precise contours and spherical shapes. They then had confidence that they could measure certain aspects of these parts, like

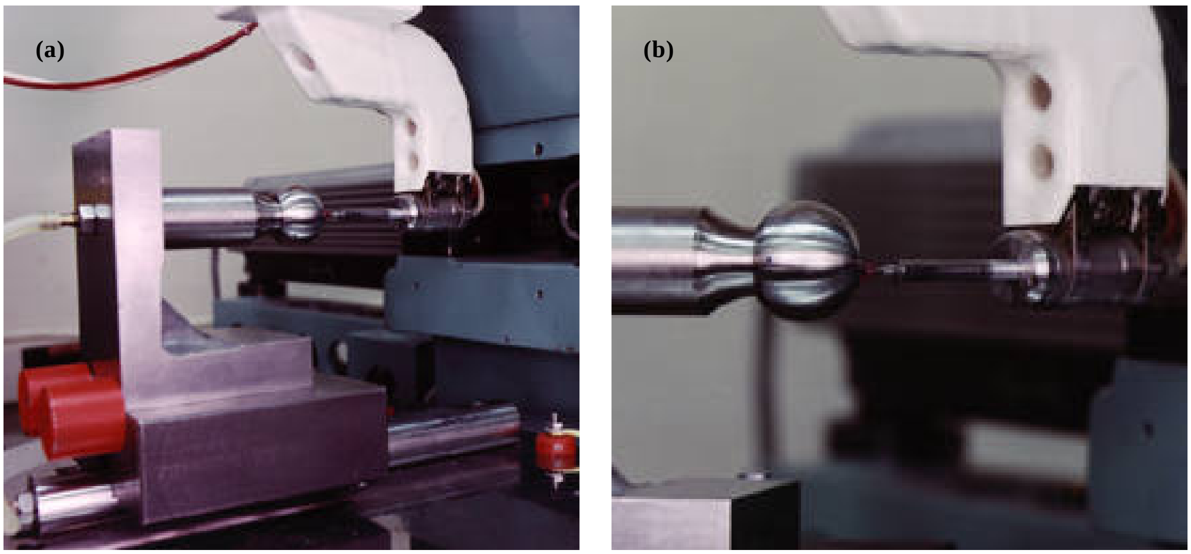

Figure 4. (a) The small-radius gauge, a commercial machine that LLNL upgraded: we added a precision air-bearing spindle, a laser interferometer position-feedback system, and a data-acquisition computer. (b) Closeup of a laser stylus attached to the air-bearing LVDT probe designed and built by LLNL. The system can measure the contour of workpieces to 2.5 nm.

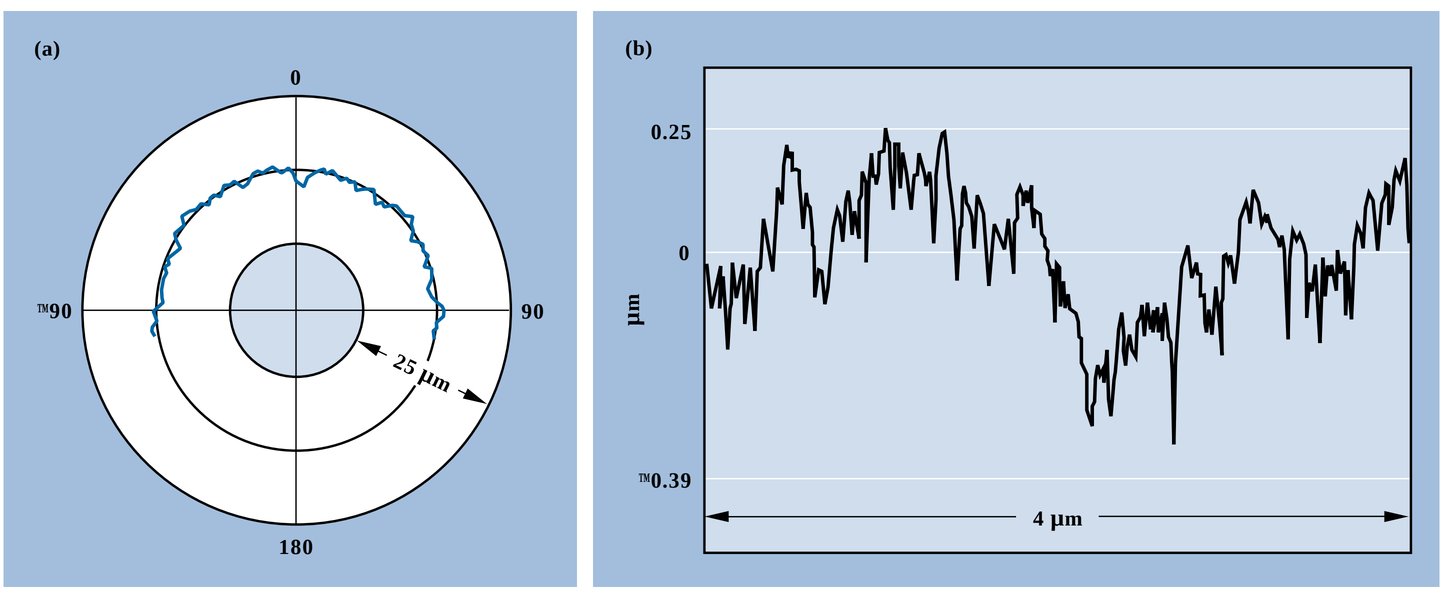

Figure 5. (a) Plot generated by the data-acquisition computer showing contour errors generated in the single-point turning process. (Note that the distance between the inner and outer circles is 25 µm.) (b) A trace showing the roughness of a single-point turned head. The roughness average (R_a) is 0.12 µm. The white line, at 0, represents a theoretical ideal smooth surface.

contour, on their gauge with known certainty. The minor discrepancies between our measurements and theirs quantified the uncertainty of their inspection process.

Several U.S. companies and tens of companies worldwide make femoral heads. Each firm has its own specifications and standards. It must be said that the industry produces workpieces that are, for the most part, extremely accurate spherical components. The specifications for contour and surface roughness are very stringent and have motivated these companies to develop advanced manufacturing efforts. Some pieces, made from solid ceramic, have a contour accuracy of 0.25 μm.

The Deliverable

We presented Smith and Nephew Richards with a “turnkey” package: step-by-step procedures, complete documents on the computer-controlled toolpath program, and an LLNL design for a specialized toolholder that provides fine control over tool height. We qualified our process by producing parts that met their specifications.

Summary

Hip replacement surgery is performed most often to relieve the pain and immobility resulting from arthritis. Although current artificial hip systems work quite well, reducing the slow rate of wear can further improve the success of the operation. An improved, more cost effective bearing surface can be produced by using oxidized zirconium alloy technology. There are health and financial incentives for developing a device that optimizes longevity and affordability.

The materials scientists at Smith and Nephew Richards, an established manufacturer of prosthetic devices, decided that a femoral head made of zirconium would combine the lower production costs of metal heads with the much greater longevity (less wear) of ceramic heads, because zirconium, a metal, can be oxidized to acquire a surface layer of zirconia, a ceramic. Smith and Nephew Richards’ machinists, however, encountered difficulties in grinding zirconium heads to the needed size, contour, and smoothness. Having learned of the Laboratory’s accomplishments in machining to extremely high precision, Smith and Nephew Richards and LLNL entered into a research agreement under the National Machine Tool Partnership to explore solutions to their manufacturing problem.

We at the Lab soon established that grinding was a rather difficult method for removing material compared to turning/cutting. The ideal technique is precision single-point turning, one that we have decades of experience with through weapons work. We located a turning machine at the Lab that is identical to one that Smith and Nephew Richards uses and turned some pieces on it to their specifications. We delivered to Smith and Nephew Richards a process, a design for a custom workpiece holder for high-precision single-point turning, and the software for generating the compensated tool paths. In the course of this work, Smith and Nephew Richards discovered that they could improve upon their methodology for measuring sphericity and size using the expertise of LLNL. By measuring the same pieces at both sites and comparing the results, we quantified the maximum accuracies that they could achieve with their equipment.

Key Words: National Machine Tool Partnership; precision engineering; single-point turning.

Notes and References

1. *Orthopedic Network News* **5**(3), 1-5.
2. B. Friedman and A. Elixhauser, “Increased Use of an Expensive, Elective Procedure: Total Hip Replacements in the 1980s,” *Medical Care 31*(7), 581–599.
3. For a description of precision engineering at LLNL, see the article beginning on p. 16 of the January 1986 *Energy and Technology Review* (UCRL-52000-86-1), and for a more comprehensive description of the philosophy, capabilities, and accomplishments of precision engineering at LLNL to that point, see the entire September 1987 issue of *Energy and Technology Review* (UCRL-52000-87-9).

This project was funded by the U. S. Department of Energy’s National Machine Tool Partnership Project.

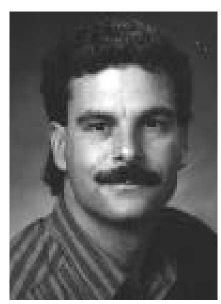

For further information contact Pete J. Davis (510) 422-3494.

KEN Project: Real-World Face Recognition

The silicon retina extracts object information from the unilluminated face area that cannot be gathered by the CCD camera. Left to right, compare the images taken in two-sided illumination by the CCD camera and silicon retina with the pair taken in one-sided illumination by the CCD camera and silicon retina.

THE Institute for Scientific Computing Research at the Laboratory recently developed the real-time, face-recognition system KEN. This system was developed through the integration of existing hardware and software. The name KEN is derived from the verb "to ken = to see, to recognize" and is related to the meaning of the noun in the Shakespeare quote "'tis double death to die in ken of shore."

Handwriting and speech recognition are already part of today's multimedia PCs and personal digital assistants. Object recognition, specifically face recognition, will be a natural extension of this technology. Associative query of image databases (e.g., search of medical image databases) and the interpretation of facial expressions for human-computer interfaces are likely applications. The existence of video-conferencing and the advent of interactive television widen the field of possible applications (e.g., model-based compression techniques, low-bit-rate video conferencing).

In security and surveillance, applications include ID verification at automatic teller machines (ATMs), credit card verification, preselection of crime suspects by matching sketches to a database, and identification of suspects for America's most wanted. The following scenarios are thinkable: An automatic recognition system preselects a few pictures of suspects from a crime scene out of a large database or alerts security personnel if suspicious activity occurs around an ATM or a property. Recognition systems included in valuable properties, such as cars or boats, could be used to identify the owner and prevent unauthorized access. This recognition technology is not limited to faces but can also be applied to other object classes, such as signatures and footprints.

Recognition Process

KEN demonstrates object recognition technology by recognizing human faces in a real-world environment. The demonstration system contains several modular components: The input sensor is either an off-the-shelf charge-coupled device (CCD) camera or an analog silicon retina chip (from Caltech). The silicon retina normalizes contrast locally and offers distinct advantages under strongly varying lighting conditions by extending the dynamic range beyond that of a normal CCD camera. A computer-controlled lens system with autozoom, autoiris, and autofocus and a pan-tilt unit track the possible candidates. A datacube real-time video processing board extracts features and segments the input images. The control software for mechanical components and the classification software run on a workstation.

KEN's software represents faces as labeled graphs. The vertices of these graphs are labeled by feature vectors and their links by distance vectors. To "acquaint" the system with a person during the learning phase, the person's face is extracted as a frontal view and stored as a face graph model in the system's knowledge database. The knowledge base contains all memorized faces in the form of sparse graphs. During recognition, all face models are compared to the unknown face candidates appearing

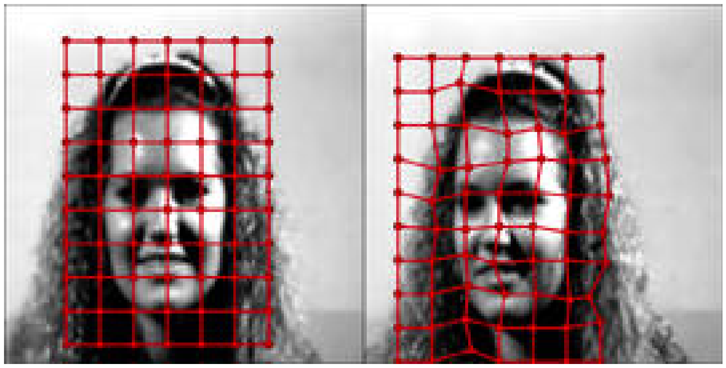

Rubber sheet effect: The labeled graph model (overlay) for the face in the left image is distorted to fit best over the face in the right image. The quality of the fit describes how closely the two face images resemble each other.

in the input. The face models are positioned over a face candidate and normalized in size and orientation by a template match. Subsequently, the graphs are locally distorted by an elastic matching process to fit the input image as closely as possible. The distortion effect is similar to the stretching of an elastic membrane (e.g., a rubber sheet). Cost values are assigned to the distorted face models according to the proximity of the final match to the input. A statistical evaluation of the resulting cost values of all models allows the system either to qualify a match or to reject poor matches and ambiguous results. KEN is currently capable of classifying a face, almost in real time, from a database of 100 faces with a positive identification rate of up to 90% and less than a 1% false positive recognition rate.

Environment

The hardware for KEN is off the shelf with the exception of the silicon retina prototype. The software is called AVision. It is written as a collection of object libraries in C and C++ and runs on major workstations (Sun, SGI, PC) and supercomputers, depending on the availability of C++ compilers. An upcoming parallel implementation in SISAL will help to speed up the development of learning algorithms that require the performance of supercomputers.

Future Directions

We plan an extension of this work to motion sequences of objects in a scene, such as the parts of a face in facial expressions or the body parts of a moving person. This involves constructing a choreographic database to implement a motion recognition system, because motion, such as the distinct trajectories of moving limbs, is used by humans as an additional visual clue to recognition. Motion processing appears to be an important component in a real-world application of recognition technology. Other research is aimed at organizing large databases to improve the efficiency of the elastic matching process. Clustering processes, active data acquisition, and learning algorithms for the weighting of components are under development. The associative query of pictorial databases (e.g., of medical images) is a similar problem that may also be researched.

Acknowledgments

We are grateful to Kawabena Boahen and Carver Mead, Caltech, for providing the prototype of a silicon retina. Joachim Buhmann, Bonn, Germany, contributed active vision concepts, learning algorithms, clustering schemes, and programming effort to the project.

For more information contact Martin Lades (510) 422-0893, hml@llnl.gov.

Modeling Groundwater Flow and Chemical Migration

GROUNDWATER contamination is a major environmental problem throughout the world. In the United States, for instance, numerous governmental and industrial sites require remediation. The Department of Energy is currently cleaning up several of its contaminated facilities, including LLNL.

At LLNL, chemical solvents and petroleum products were dumped onto the ground surface in the 1940s when the present site was a naval air station. Forty years of research and development activities also contributed to the contamination problem. LLNL is obligated to characterize the contamination and clean it up. To this end, various engineered remediation techniques (e.g., pump-and-treat and biofilters) are now being studied, tested, and implemented.

To understand better the efficacy of a given cleanup strategy, as well as to determine the most economical implementation for a specific site, engineers frequently employ mathematical modeling tools to aid in their analysis and design of remediation procedures. Unfortunately, many models are based on unrealistic assumptions about the subsurface media and flow behavior. For example, many models in use today ignore the fact that the subsurface is heterogeneous in composition and in spatial distribution. These simple homogeneous models are incapable of capturing important physical phenomena (such as nonuniform diffusion) that have a significant impact on contaminant migration.

Another flaw of many models is that they employ outdated and inefficient numerical methods, which preclude running realistic simulations on even the largest conventional vector supercomputers.

To address these deficiencies, we are developing ParFlow, a simulator for modeling fluid flow through heterogeneous porous media. To enable detailed simulations of large sites, this code uses the latest numerical methods and high-performance computing technologies.

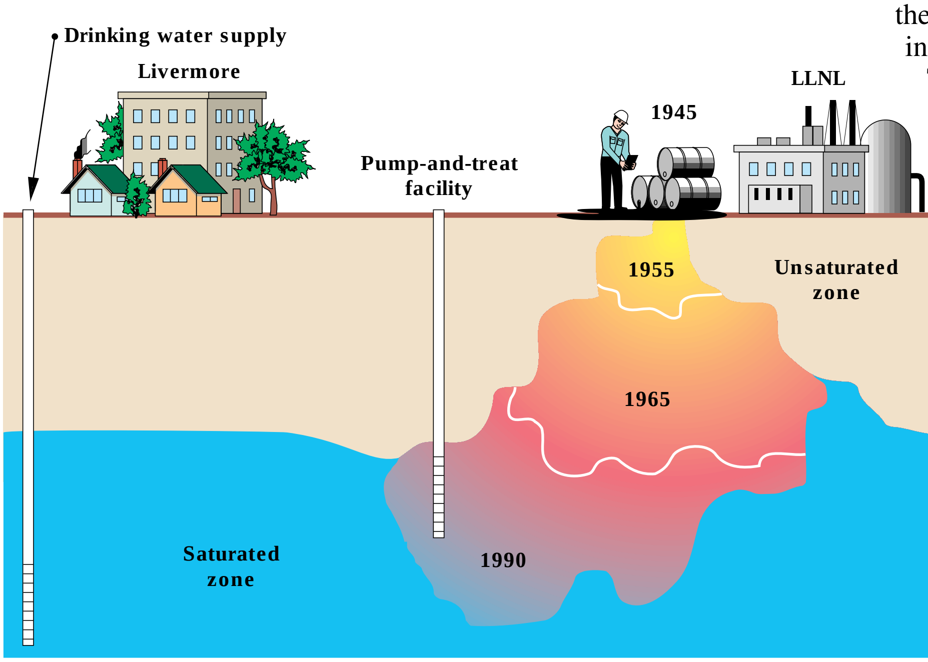

Contaminants have migrated through the unsaturated zone into the more mobile groundwaters. LLNL is designing and executing remediation procedures.

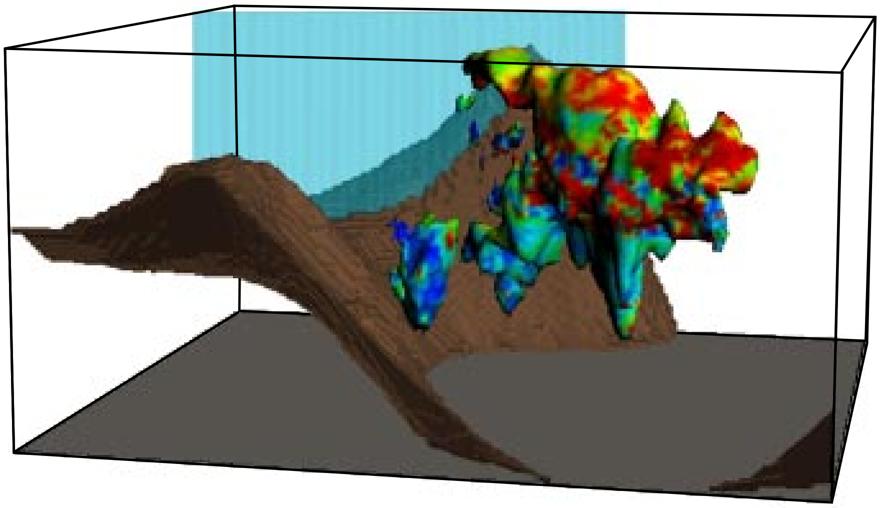

A snapshot in time of contaminant migration through the LLNL subsurface (hypothetical, multilayered, heterogeneous realization). A clay layer is shown in brown. The blue plane represents a fault zone. (The vertical scale is exaggerated 10 times.)

Modeling Flow and Contaminant Migration on MPPs

ParFlow is a portable and parallel code for modeling multiphase fluid flow and multicomponent chemical transport through heterogeneous porous media. This scalable code presently runs on the CRAY T3D, IBM SP-1, and nCUBE/2 massively parallel computers, as well as on a two-processor SGI Onyx and a cluster of Sun Sparcstations. It was used recently to simulate groundwater flow at LLNL (using a hypothetical subsurface model). The size of the physical site (several square kilometers) and the need to resolve the subsurface heterogeneities (to within a few meters) resulted in a grid with more than one million spatial zones.

To obtain the necessary hydraulic conductivity values, we implemented an algorithm that generates a statistically accurate subsurface realization from the given field data. We then discretized the modeling equations and solved the resulting linear system with preconditioned conjugate gradients. The velocity field was then passed to a particle-in-cell code to simulate contaminant migration.

It is important to note that the subsurface heterogeneities give rise to preferential flow channels. These channels, which are not reproducible by homogeneous codes, lead to nonuniform contaminant migration and have a major impact on the efficacy of engineered remediation procedures.

Graphical User Interface

To facilitate code development and maintenance, we are developing XParFlow, an X-Windows (Motif) interface to ParFlow. This program displays the current simulator configuration as a connected graph. By selecting a node in the graph, the user is able to customize the simulator. For example, the user can choose from several linear solvers and preconditioners.

To facilitate problem specification, we are currently using the GMS package from Brigham Young University and the Department of Defense. This new graphical user interface was specifically designed to be a front-end to groundwater modeling codes such as ParFlow.

Multidisciplinary Collaboration

The ParFlow project is an interdisciplinary effort involving scientists from various programs within the Laboratory. LLNL and International Technology Corporation are partnering to commercialize some of this work under the auspices of a cooperative research and development agreement (CRADA).

For further information contact Steven Ashby (510) 423-2462, sfashby@llnl.gov.

Gas and Oil National Information Infrastructure

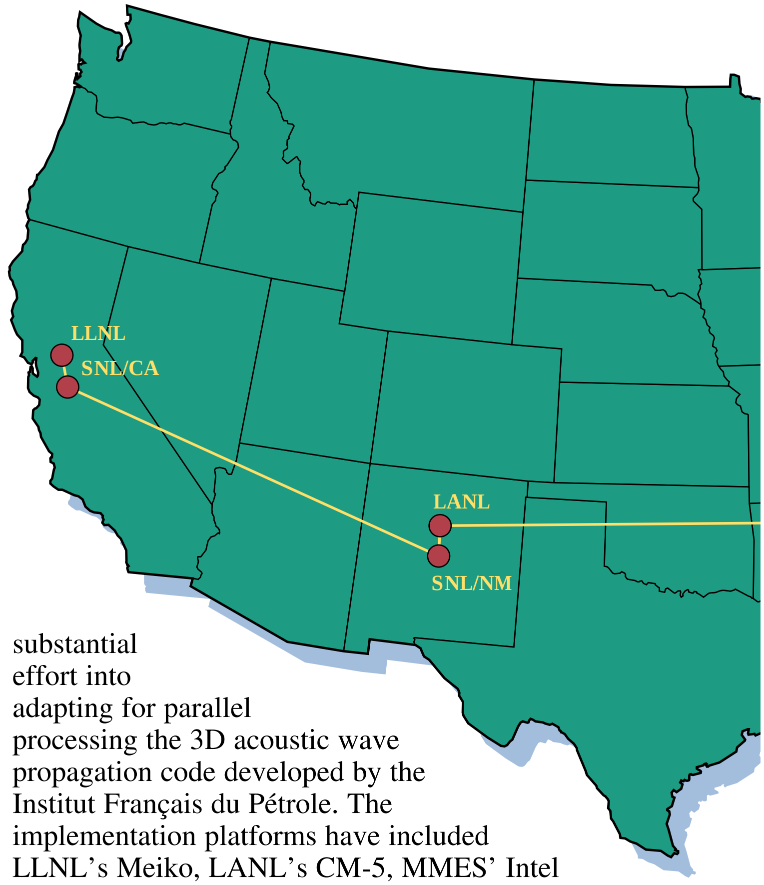

THE U.S. Department of Energy pilot National Information Infrastructure (NII) programs are designed to create computing and communications testbeds with U.S. industries. One such program is the Gas and Oil National Information Infrastructure (GO-NII), which unites four of DOE's national laboratories: Lawrence Livermore National Laboratory (LLNL), Los Alamos National Laboratory (LANL), Martin Marietta Energy Systems (MMES) at Oak Ridge, and Sandia National Laboratories (SNL). GO-NII focuses on networking, high-speed input and output, mass storage, collaborative and productivity tools, scientific data management, data visualization, data security, and distributed computing. Industry can evaluate the feasibility of using new advanced technologies in the GO-NII program while reducing the financial investment and risk of supporting its own research program. Technology and information will be transferred from the national laboratories to all segments of the gas and oil industry.

GO-NII Demonstration

At the Supercomputing '94 conference in Washington, D.C., GO-NII demonstrated the use of an advanced high-speed, long-haul communication system; a new architecture in high-performance storage systems; calculating, searching, browsing, and cataloging of three-dimensional (3D) seismic data and metadata; and productivity tools to enhance remote collaboration among scientists. Although this demonstration focused on the gas and oil industry, the technology can be applied to other U.S. industries.

Synthetic Seismic Dataset

A major undertaking of the GO-NII effort is a project with the Society of Exploration Geophysicists and the European Association of Exploration Geophysicists to design salt and overthrust 3D models and then simulate realistic 3D surveys based on those models. The four GO-NII laboratories are working jointly to generate this 3D synthetic seismic dataset for the verification and validation of seismic processing tools used in the gas and oil industry. These national laboratories have put substantial effort into adapting for parallel processing the 3D acoustic wave propagation code developed by the Institut Français du Pétrole. The implementation platforms have included LLNL's Meiko, LANL's CM-5, MMES' Intel 860 and Paragon, and SNL's Intel Paragon. The laboratories will be applying massively parallel processing at the equivalent of approximately 7000 hours at a billion calculations per second (Gflops) toward this calculation between October 1994 and September 1995.

The salt model and overthrust model synthetic seismic datasets calculated at the four laboratories will be stored at LLNL's Facility for Advanced Scalable Computing Technology (FAST). FAST will provide information and expertise in emerging storage technology, as applied to the use and distribution of these massive datasets. FAST will provide a base on which to build distributed visualization, remote access, and collaborative tools to enable more efficient analysis and processing of seismic information.

At the end of the project, the resulting model, tools, and data will be made publicly available to companies within the gas and oil industry for further commercialization and research.

ARIES Testbed Partnership

The Amoco ARIES testbed, as described in *Network World* and other trade journals, is a partnership between Amoco and about 15 industrial communications and computing vendors to explore the use of asynchronous transfer mode networks in the oil and gas business sector. The four GO-NII laboratories offer their capabilities in network technology, distributed computing, and system integration to partner with the ARIES testbed in

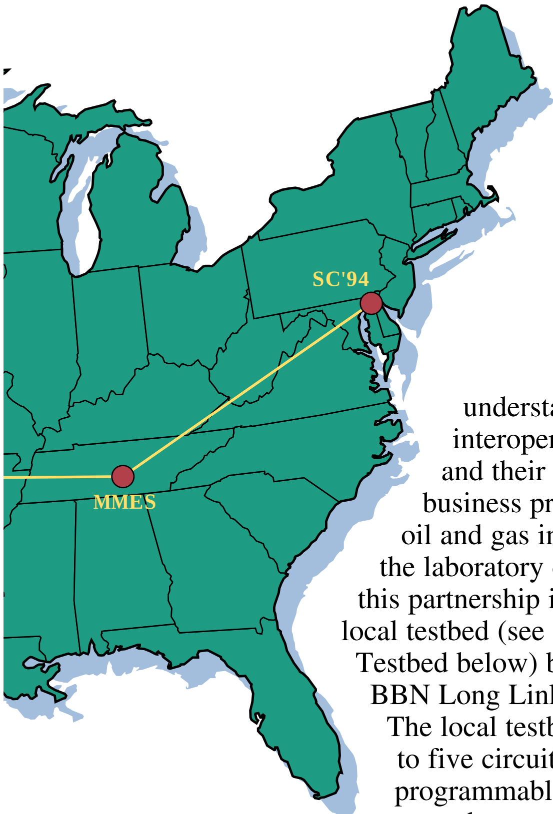

At the IEEE Supercomputing '94 Conference, DOE national laboratory researchers demonstrated long-haul, high-speed communications, high-performance storage, and collaborative tools, all of which can be used in gas and oil industry applications.

understanding interoperability issues and their impact on business processes in the oil and gas industry. Part of the laboratory contribution to this partnership is to establish a local testbed (see GO-NII Testbed below) based on the BBN Long Links Emulator. The local testbed provides up to five circuits of programmable delay and error characteristics that will be used to emulate nationwide networks within a more controlled testing environment, eliminating the need for long-haul communication lines in the early stages of testing and development.

GO-NII Testbed

The GO-NII testbed will provide an environment through which the gas and oil industry can evaluate, benchmark, and access new computing infrastructure technologies of interest to them. The GO-NII testbed will simulate long-haul, high-performance networks with access to supercomputers, workstation clusters, and high-performance mass storage in a transparent distributed computing environment. These include the National Storage Laboratory and High-Performance Storage System environment, massively parallel systems, emerging technologies in parallel input/output, cluster environments, and dynamic software to manage the environment and its resources. The GO-NII testbed will provide integrated access tools to effect the desired seamless environment.

Petroleum Technology Transfer Council

A GO-NII focus area is to work with independent oil and gas producers in making use of advanced computing technologies to increase domestic oil production. Independent of GO-NII, the Petroleum Technology Transfer Council (PTTC) was formed in January 1993 as an *ad hoc* council to address the technology needs of the U.S. oil- and gas-producing community and to identify the best mechanisms for improving the transfer and communication of technology to domestic producers. Ten regional lead organizations (RLOs) representing the oil-producing areas of the United States have been identified, and regional resource centers will be established at these sites. A producer advisory group in each region will help to define the priority technical problems encountered in the region and to determine the needs of the producers in accessing available technologies to resolve them.

The four national laboratories participating in GO-NII are working with the PTTC RLOs to provide expertise, advice, and technology transfer possibilities and thus to help in forming the electronic information system linking the ten RLOs and the approximately 8000 independent oil and gas producers. Potential areas in which the laboratories and RLOs will team up are computing infrastructure for electronic commerce, full spectrum network access, network security, distributed visualization, data integrity and conversion, system integration, electronic storage, compression, and database technologies.

ACTS Testbed

Amoco, the American Petroleum Institute, Naval Research Laboratory, DOE laboratories, and NASA are planning a joint demonstration of satellite communications capabilities as applied to the gas and oil industry. NASA's Advanced Communications Technology Satellite is operating multiple T1 (1.5-megabit-per-second) circuits, with the capability of OC-3 (155-Mb/s) and OC-12 (622-Mb/s) speeds. This capability could have a dramatic impact on the ability of the gas and oil industry to retrieve seismic and reservoir data from remote sites, such as the Gulf of Mexico, Alaska, and foreign locations. A low-speed demonstration of this capability is scheduled for December 1994.

For further information contact Carol Hunter (510) 422-1657, chunter@llnl.gov.

At the IEEE Supercomputing '94 Conference, DOE national laboratory researchers demonstrated long-haul, high-speed communications, high-performance storage, and collaborative tools, all of which can be used in gas and oil industry applications.

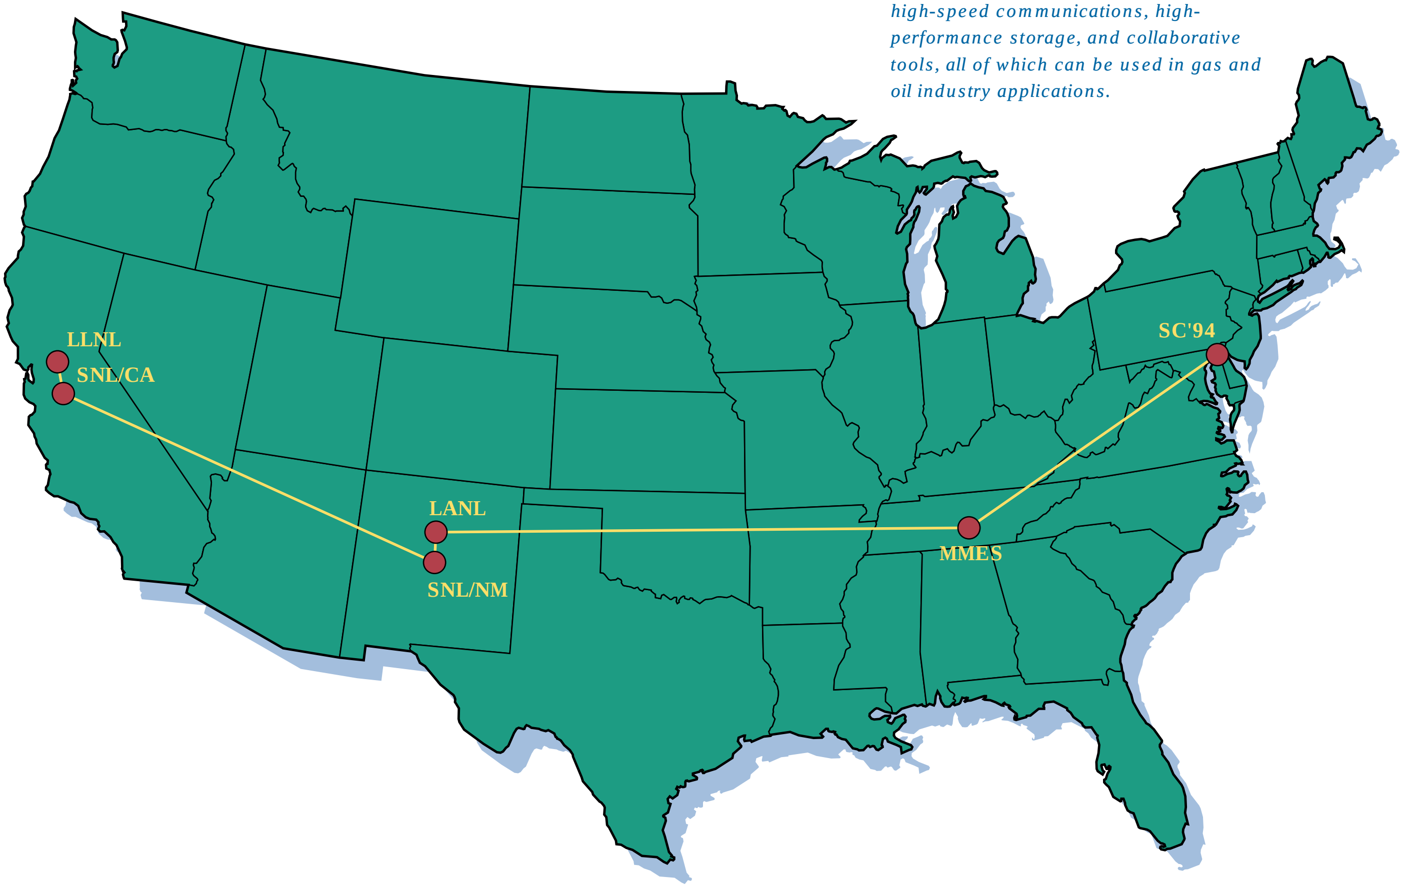

Figure for Gas and Oil National Information Infrastructure Highlight

The Industrial Computing Initiative

The widespread use of massively parallel computing as a scientific and industrial tool has been impeded by technical problems associated with available hardware and software. U.S. industries do not have the applications and necessary software tools and languages to facilitate use of massively parallel computing environments. To address this need, the Industrial Computing Initiative (ICI), which is part of the $66-million High-Performance Parallel Processing Project, is developing a set of efficient applications that can realistically simulate complex problems and are written in a way that allows them to run on massively parallel hardware from various vendors.

The ICI involves more than 40 scientists from LLNL and Los Alamos National Laboratory, six specialists in parallelization from Cray Research, two specialists from Thinking Machines Corporation, and nearly a score of industrial scientists. The delivery of a set of efficient parallel applications, serving as guideposts for subsequent work, can help U.S. industry compete more effectively in the global marketplace. The new applications, together with improved methods and tools, can reduce product cycle times, produce higher-quality products, and speed the development of more efficient industrial processes. At the same time, the new computational advances in massively parallel computing technology can be reintegrated into ongoing programs at the national laboratories to serve the missions defined by the DOE and to address national needs.

Contact: Michel G. McCoy (510) 422-4021.

Artificial Hip Joints: Applying Weapons Expertise to Medical Technology

Materials scientists at Smith and Nephew Richards, a firm that makes artificial hip joints, decided that a femoral head made of zirconium, which can be oxidized to acquire a surface layer of zirconia, a ceramic, would combine the lower production costs of metal heads with the much greater longevity of ceramic heads. Smith and Nephew Richards machinists, however, encountered difficulties in grinding zirconium heads to the needed size, contour, and smoothness. Because of LLNL's expertise in precision engineering, Smith and Nephew Richards and the Lab signed a research agreement under the National Machine Tool Partnership to help them solve their problem.

We soon established that the preferred technique for use with zirconium was not grinding, but single-point turning. We thus had an opportunity to turn decades of experience acquired through weapons work to advances in medical technology. Identifying an appropriate turning machine that the Lab and Smith and Nephew Richards both have, we turned some pieces on it to their criteria. We delivered to Smith and Nephew Richards a turnkey package: a custom workpiece holder for high-precision single-point turning and the software for operating the compensated tool path. During the course of this work, Smith and Nephew Richards and the Laboratory discovered that the Lab's inspection capabilities for measuring the size, contour, and smoothness of the femoral heads were more precise than theirs, and so we undertook and completed a secondary task of helping establish the exact limits of their fabricating and inspection equipment.

Contact: Pete J. Davis (510) 422-3494.