




Windows on Computing

New Initiatives at Los Alamos

David W. Forslund, Charles A. Slocomb, and Ira A. Agins

No aspect of technology is changing more rapidly than the field of computing and information systems. It is among the fastest growing and most competitive arenas in our global economy. Each year, more of the items we use contain tiny microprocessors—silicon chips on which are etched hundreds or thousands or millions of electronic circuit elements. Those computer chips direct various operations and adjustments—automatic braking in cars; automatic focusing in cameras; automatic data collection in cash registers; automatic message-taking by answering machines; automatic operation of washers, dryers, and other appliances; automatic production of goods in manufacturing plants; the list could go on and on. Those inconspicuous devices that perform micro-scale computing are profoundly shaping our lives and our culture.



Opening illustration: Elements of high-performance computing: Left, the CM-5 Connection Machine, the most powerful massively parallel supercomputer at Los Alamos; center, the foyer in the Laboratory Data Communications Center (LDCC); upper right, digitized images of Washington, D.C. from a wavelet-based multiresolution database developed at Los Alamos; and lower right, a portion of a “metacomputer,” a 950,000-transistor special-purpose chip for analyzing the behavior of digital circuits. The chip is being developed by a team of graduate students at the University of Michigan.

More visible and no less important are the ways microprocessors are changing the way we communicate with each other and even the kinds of tasks we do. Industries such as desktop publishing, electronic mail, multimedia systems, and financial accounting systems have been created by the ubiquitous microprocessor. It is nothing short of the engine of the digital revolution.

The computer chip was invented in 1958 when Jack Kilby figured out how to fabricate several transistors on a single-crystal silicon substrate and thereby created the integrated circuit. Since then more and more transistors have been integrated on a single chip. By the early 1980s the technology of VLSI, or very-large scale integration (hundreds of thousands of transistors on

a chip), had led to dramatic reductions in the cost of producing powerful microprocessors and large memory units. As a result affordable personal computers and powerful workstations have become commonplace in science, in business, and in the home.

New microprocessors continue to be incorporated into various products at an increasing rate; development cycles are down to months rather than years as the current generation of processors are used to aid in the design and manufacture of the next generation. Because of their economies of scale, off-the-shelf microprocessors are expanding the use of micro- and medium-scale computing in business and in the home. They are also motivating changes in the design

History of Computers at Los Alamos

1943–45

Desktop calculators and punched-card accounting machines are used as calculating tools in the Manhattan Project.

1945

ENIAC, the world's first large-scale electronic computer, is completed at the University of Pennsylvania. Its “shake-down” calculation is the “Los Alamos problem,” a calculation needed for the design of thermonuclear weapons.

1949

IBM's first Card Programmable Calculators are installed at the Laboratory.

1952

MANIAC is built at the Laboratory under the direction of Nick Metropolis. It is the first computer designed from the start according to John von Neumann's stored-program ideas.

1953

The Laboratory gets serial number 2 of the IBM 701. This “Defense Calculator” is approximately equal in power to the MANIAC.

1955

The MANIAC II project, a computer featuring floating-point arithmetic, is started. The Laboratory begins working closely with computer manufacturers to ensure that its future computing needs will be satisfied.



MANIAC II

1956

MANIAC II is completed. The Laboratory installs serial number 1 of the the IBM 704, which has about the same power as MANIAC II. From this point on, the Laboratory acquires supercomputers from industry.

Late 1950s

The Laboratory and IBM enter into a joint project to build STRETCH, a computer based on transistors rather than vacuum tubes, to meet the needs of the nuclear-weapons program.

1961

STRETCH is completed and is about thirty-five times as powerful as the IBM 704. IBM used much of the technology developed for STRETCH in its computers for years afterward.

1966

The first on-line mass-storage system with a capacity of over 10^{12} bits, the IBM 1360 Photo Store, is installed at the Laboratory. Control Data Corporation introduces the first “pipelined” computer, the CDC 6600, designed by Seymour Cray. The Laboratory buys a few.

1971

The Laboratory buys its first CDC 7600, the successor to the 6600. These machines are the main supercomputers in use at the Laboratory during much of the 1970s.

1972

Cray Research, Inc is founded. The Laboratory consults on the design of the Cray-1.

1975

Laboratory scientists design and build a high-speed network that uses 50-megabit-per-second channels.

of large-scale scientific computers. The focus has shifted from single, very fast processors to very fast networks that allow hundreds to thousands of microprocessors to cooperate on a single problem. Large-scale computation is a critical technology in scientific research, in major industries, and in the maintenance of national security. It is also the area of computing in which Los Alamos has played a major role.

The microprocessor has opened up the possibility of continual increases in the power of supercomputers through the architecture of the MPP, the massively parallel processor that can consist of thousands of off-the-shelf microprocessors. In 1989, seeing the potential of that new technology for ad-

ressing the “Grand Challenge” computational problems in science and engineering, Los Alamos set up the Advanced Computing Laboratory as a kind of proving ground for testing MPPs on real problems. Ironically, just as their enormous potential is being clearly demonstrated at Los Alamos and elsewhere, economic forces stemming from reduced federal budgets and slow acceptance into the commercial marketplace are threatening the viability of the supercomputing industry.

As a leader in scientific computing, Los Alamos National Laboratory has always understood the importance of supercomputing for maintaining national security and economic strength. At this critical juncture the Laboratory

plans to continue working with the supercomputing industry and to help expand the contributions of computer modeling and simulation to all areas of society. Here, we will briefly review a few of our past contributions to the high end of computing, outline some new initiatives in large-scale parallel computing, and then introduce a relatively new area of involvement, our support of the information revolution and the National Information Infrastructure initiative.

Since the days of the Manhattan Project, Los Alamos has been a driver of and a major participant in the development of large-scale scientific computation. It was here that Nick Metropolis directed the construction of MANIAC I

1976

Serial number 1 of the Cray-1 is delivered to the Laboratory.

1977

A Common File System, composed of IBM mass-storage components, is installed and provides storage for all central and remote Laboratory computer systems.



The Cray-1

1980

The Laboratory begins its parallel-processing efforts.

1981

An early parallel processor (PuPS) is fabricated at the Laboratory but never completed.

1983

Denelcor's HEP, an early commercially available parallel processor, is installed, as is the first of five Cray X-MP computers.

1985

The Ultra-High-Speed Graphics Project is started. It pioneers animation as a visualization tool and requires gigabit-per-second communication capacity. A massively parallel (128-node) Intel computer is installed.

1987

The need for higher communication capacity is answered by the development of the High-Performance Parallel Interface (HIPPI), an 800-megabit/second channel, which becomes an ANSI standard.

1988

The Laboratory obtains the first of its six Cray Y-MP computers. It also installs, studies, and evaluates a number of massively parallel computers. The Advanced Computing Laboratory (ACL) is established.

1989

The ACL purchases the CM-2 Connection Machine from Thinking Machines. It has 65,536 parallel processors.

1990

A test device for HIPPI ports is transferred to industry. The Laboratory, the Jet Propulsion Laboratory, and the San Diego Supercomputer start the Casa Gigabit Test Project

1991

The Laboratory transfers to industry the HIPPI framebuffer, an important component for visualization of complex images.

1992

A 1024-processor Thinking Machines CM-5, the most powerful computer at the time, is installed at the ACL.

1994

A massively parallel Cray T3D is installed at the ACL for use in collaborations with industry.



The Cray T3D

and II. Maniac I (1952) was among the first general-purpose digital computers to realize von Neumann's concept of a stored-program computer—one that could go from step to step in a computation by using a set of instructions that was stored electronically in its own memory in the same way that data are stored. (In contrast, the ENIAC (1945), the very first General-purpose electronic computer, had to be programmed mechanically by inserting cables into a plugboard.) Maniac II (1956) embodied another major advancement in computer design, the ability to perform floating-point arithmetic—the kind that automatically keeps track of the position of the decimal point. The peak speed of MANIAC II was 10,000 arithmetic operations per second, an impressive factor of 1000 higher than that of the electromechanical accounting machines used to perform numerical calculations at Los Alamos during the Manhattan Project (see Figure 1).

The main function of those accounting machines and very early electronic computers was to simulate through numerical computation the extreme physical conditions and complex nonlinear processes that occur within a nuclear weapon. The continued role of computer simulation as the primary tool for designing and predicting the performance and safety of nuclear weapons has been a major driver behind the development of increasingly powerful computers. The larger the computer, the greater the complexity that could be simulated accurately. That is still true today as the largest supercomputers are being used to simulate realistically very complex phenomena including the interaction of the oceans and the atmosphere on global scales, the basis of bulk properties of materials in the motions of individual atoms, the flow of oil and gas through the porous rock of underground reservoirs, the dynamics of the




Figure 1. The Growth of Computing Power

This plot shows the number of operations per second versus year for the fastest available "supercomputers". Different shapes are used to distinguish serial, vector, and parallel architectures. All computers up until and including the Cray-1 were single-processor machines. The dashed line is a projection of supercomputing speed through 1990. That projection was published in a 1983 issue of *Los Alamos Science*—before massively parallel processors became practical.

internal combustion engine, the behavior of biological macromolecules as pharmaceuticals, and so forth.

After 1956 and MANIAC II the Laboratory stopped building computers and instead relied on close working relationships with IBM and other vendors to ensure that the industry would be able to supply the necessary computing power. A particularly important collaborative effort was STRETCH, a project with IBM started in 1956 to design and build the fastest computer possible with existing technology. For the first time the new semiconductor transistor would be used in computer design. Transistors have a much faster response time and are much more reliable than the traditional vacuum-tube elements of digital circuits. The STRETCH, or IBM 7030, computer was made from 150,000 transistors and was delivered to Los Alamos in 1961. It was about thirty-five times faster than the IBM 704, a commercial vacuum-tube machine similar to the MANIAC II.

In the early 1970s Los Alamos became a consultant to Seymour Cray on the design of the Cray-1, the first successful vector computer. Vector architecture increases computational speed by enabling the computer to perform many machine instructions at once on linear data arrays (vectors). In contrast, computers with traditional serial architecture perform one machine instruction at a time on individual pieces of data (see "How Computers Work" in this volume).

Los Alamos was not only a consultant on the design of the Cray but was also the first purchaser of that innovative hardware. The delivery of the first Cray computer to Los Alamos in 1976 might be said to mark the beginning of the modern era of high-performance computing. The Cray-1 supercomputer had a speed of tens of megaflops (one megaflop equals a million floating-point

operations per second) and a memory capacity of 4 megabytes.

In all these developments at the high end of computing, the Laboratory has taken part in the struggles associated with bringing new technology into the marketplace, testing new machines, developing the operating systems, and developing new computational techniques,

The Laboratory has taken part in the struggles associated with bringing new supercomputing technology into the marketplace, testing new machines, developing the operating systems, and developing new computational techniques, or algorithms, to make full use of the potential computing power presented by each new supercomputer. That type of involvement continues today.

Los Alamos was one of the first institutions to demonstrate the general usefulness of the CM-2 Connection Machine, a massively parallel computer built by Thinking Machines Corporation.

or algorithms, to make full use of the potential computing power presented by each new supercomputer. That type of involvement continues today. The Laboratory was one of the first institutions to demonstrate the general usefulness of

the CM-2 Connection Machine, a massively parallel computer built by Thinking Machines Corporation. The CM-2 was originally designed to investigate problems in artificial intelligence, and when it arrived in Los Alamos in 1988 to be tried out on problems related to atmospheric effects and other hydrodynamic problems of interest in the weapons program, it did not contain the special processors needed for efficient computation of floating-point arithmetic. Laboratory scientists responded by working in collaboration with Thinking Machines on the installation and testing of the appropriate processing units and went on to develop the first successful algorithms for performing hydrodynamic problems on the Connection Machine's parallel architecture.

The question of whether parallel architectures would be a practical, cost-effective path to ever-increasing computer power has been hanging in the balance since the early 1980s. By then it was apparent that the speed of a single vector processor was limited by the physical size and density of the elements in the integrated circuits and the time required for propagation of electronic signals across the machine.

To increase the speed or overall performance of supercomputers by more than the factor associated with the speed-up in the individual processors, one could use parallel systems, for example, a number of Cray-1-type vector processors working together or massively parallel systems in which thousands of less powerful processors communicate with each other. In the early 1980s the latter approach appeared to be intractable.

At the same time that supercomputer manufacturers were worrying about how to improve on the standards of performance set by the Cray-1, an effort was begun to apply VLSI technolo-

gy to the design of a “Cray on a chip.” That effort to produce very powerful microprocessors was motivated in the early 1980s by the rapid growth of and standardization in the PC and workstation marketplace. As microprocessor technology advanced and commercial applications increased, production costs of microprocessors decreased by orders of magnitude and it became possible to produce very high-performance workstations with price-to-performance ratios much lower than those associated with conventional vector supercomputers. The workstations produced by Sun, Silicon Graphics, IBM, and Hewlett-Packard can sustain speeds comparable to the Cray-1 (see Figure 2). Such performance is sufficient for a wide range of scientific problems, and many researchers are choosing to adapt their computing to high-end workstations and workstations networked together to form workstation clusters. The microprocessor revolution has led to the workstation revolution in scientific computing.

Supercomputer manufacturers have also tried to build on the cost savings and technology advances afforded by the microprocessor revolution. Thinking Machines, Intel, Cray Research, and others have been designing MPPs, each containing hundreds or thousands of off-the-shelf microprocessors of the kind found in high-end workstations. Those processors are connected in parallel through various network architectures and are meant to work simultaneously on different parts of a single large problem. As indicated in Figure 2, MPPs hold the promise of increasing performance by factors of thousands because many of their designs are scalable, that is, their computational speed increases in proportion to the number of processors in the machine. One of the largest scalable MPPs in use is at




Figure 2. Supercomputers Today and Tomorrow

The speeds of workstations, vector supercomputers and massively parallel supercomputers are compared by plotting the number of processors versus the speed per processor measured in number of floating-point operations per second (flops). The solid diagonal lines are lines of constant total speed. The maximum theoretical speed of the 1024-node CM-5 Connection Machine is 131 gigaflops. Supercomputer manufacturers are hoping to make MPPs with teraflop speeds available within the next few years. Increases in speed will probably be achieved by increasing both the number of processors in an MPP and the speed per processor.

the Los Alamos Advanced Computing Laboratory—a CM-5 Connection Machine from Thinking Machines containing 1024 SPARC microprocessors. This machine has a theoretical peak speed of 130 gigaflops, more than a factor of 1000 over the Cray-1; some of

the applications (for example, the molecular dynamics calculation in “State-of-the-Art Parallel Computing”) run on the machine have already achieved 40 per cent of that theoretical maximum.

For over fifteen years the Laboratory has had much success with convention-

al vector supercomputers, particularly the Cray X-MP and Cray Y-MP. Los Alamos scientists have grown accustomed to “vectorizing” their computer programs, particularly hydrodynamics codes, and have been quite successful at designing codes that run very efficiently on the vector machines. Nevertheless, during the last five years the Laboratory has been strongly supporting the shift to parallel computing. It is our belief that the best hope of achieving the highest performance at the lowest cost lies with the massively parallel approach. In 1989 the Laboratory purchased its first Connection Machine, a CM-2, and in 1992 the Laboratory acquired the CM-5 mentioned above.

We have also been collaborating with manufacturers of high-end workstations on another approach to parallel computing, the development of workstation clusters. A cluster consists of many stand-alone workstations connected in a network that combines the computing power and memory capacity of the members of the cluster. We are helping to develop the networks and the software necessary to make such a cluster appear to the user as a single computational resource rather than a collection of individual workstations. Clusters are likely to come into greater and greater use because they provide scalable computing power at an excellent price-to-performance ratio starting from a small number of processors. At the higher performance end of computing, where hundreds of processors are involved, clusters do not necessarily compete well with vector or massively parallel supercomputers because the management of the processors becomes very complex and the scalability of the interconnections becomes more difficult to achieve.

Figure 3 shows two workstation clusters at Los Alamos. One, the IBM cluster consisting of eight IBM

RS/6000 560 workstations, is now in general use. The other, which is still in development, consists of eight Hewlett-Packard 735 workstations, each of which is faster at scalar calculations than is a single processor of a Cray Y-MP. We are now working closely with Hewlett-Packard on the software for that new cluster.

Achieving high performance on a massively parallel machine or a workstation cluster requires the use of entirely different programming models. The most flexible approach is the MIMD, or multiple-instruction, multiple-data, model in which different operations are performed simultaneously by different processors on different data. The challenge is to achieve communication and coordination among the processors without losing too much computational time. “State-of-the-Art Parallel Computing” presents a good look at what is involved in achieving high performance on MPPs.

The rich array of programming possibilities offered by parallel architectures has brought back into use many algorithms that had been ignored during the fifteen years or so that vector supercomputers dominated the scientific marketplace. Today the most active area in scientific programming and the one that will have a long-range impact on high-performance computing is the development of algorithms for parallel architectures. The Laboratory is a leader in algorithm development, and this volume presents a few outstanding examples of new parallel algorithms.

The Laboratory has also taken a leadership role in creating an advanced computing environment needed to achieve sustained high performance on the new MPPs and to store and view the enormous amounts of data generated by those machines. The focal point for research on and development and implementation of the elements of the



Figure 3. Workstation Clusters

Workstation clusters are becoming more popular for scientific computing because they have excellent price/performance ratios and can be upgraded as new microprocessors come on the market. The eight RS/6000 workstations in the IBM cluster at Los Alamos (top) are soon to be upgraded from model 560 to model 590. We are collaborating with Hewlett-Packard in the development of software for the very high-performance HP 735 workstation cluster recently assembled here (bottom). This cluster is being outfitted with HIPPI, the networking interface developed at the Laboratory for very high data transmission rates (see “HIPPI—The First Standard for High-Performance Networking”).

new computing environment is the Advanced Computing Laboratory set up at Los Alamos in 1989. The goal of the ACL is to handle the most challenging, computationally intensive problems in science and technology, the so-called Grand Challenge problems. Our 1024-processor CM-5, with its enormous speed and very large memory capacity, is the centerpiece of the Advanced Computing Laboratory. The ACL also houses advanced storage facilities developed at Los Alamos to rapidly store and retrieve the terabytes of data generated by running Grand Challenge problems on the CM-5. A special "HIPPI" network, developed at Los Alamos specifically to handle very high rates of data transmission, connects the CM-5 to the advanced storage facilities and vice versa. The HIPPI protocol for supercomputer networks has since become an industry standard (see "HIPPI").

Five DOE Grand Challenge problems are being investigated at the ACL: global climate modeling, multiphase flow, new materials technology, quantum chromodynamics, and the tokamak fusion reactor. Other very interesting simulations (see "Experimental Cosmology and the Puzzle of Large-scale Structures") are being performed on our CM-5 and have demonstrated the great potential of MPPs for scientific and engineering research.

As we make parallel computing work for Grand Challenge and other problems of fundamental interest, we are in a good position to help industry take advantage of the new computing opportunities. Our work with Mobil Corporation on modeling multiphase flow through porous media (see "Toward Improved Reservoir Flow Performance") uses an intrinsically parallel, Los Alamos-developed algorithm known as the lattice-Boltzmann method to model the flow of oil and water at

the scale of the individual pores in oil-bearing rock. The success of this collaboration provided a working example on which to base ACTI, the Advanced Computing Technology Initiative for

Today the most active area in scientific programming and the one that will have a long-range impact on high-performance computing is the development of algorithms for parallel architectures.

The Laboratory is a leader in algorithm development for the new parallel machines. Two outstanding examples described in this volume are the fast tree code for many-body problems, developed initially to trace the formation of structure in the early universe, and the lattice-Boltzmann method, an intrinsically parallel approach to the solution of complex multiphase flow problems of interest to the oil and gas industry.

establishing research collaborations between the oil and gas industry and the DOE national laboratories.

Another new collaborative effort, led by Los Alamos National Laboratory, Lawrence Livermore National Laboratory, and Cray Research, Inc., is specifically designed to transfer high-performance parallel-processing technology

to U.S. industry with the goal of increasing industrial competitiveness. The \$52 million collaborative program is under the auspices of the DOE's High Performance Parallel Processor program and will involve fifteen industrial partners. Over seventy scientists will be involved in creating computer algorithms for massively parallel machines that are of direct use in simulating complex industrial processes and their environmental impact.

In addition, two networked 128-processor Cray T3D MPPs, one at Los Alamos and one at Livermore, will be used in the industrial collaborations. They will be the first government-owned machines to be focused primarily on solving industrial problems.

The program includes fifteen projects in the fields of environmental modeling, petroleum exploration, materials design, advanced manufacturing, and new MPP systems software. Los Alamos will be involved in seven of the projects, two of which are devoted to developing general software tools and diagnostics that cut across particular applications and address important issues of portability of software from one computing platform to another. Those projects in which Los Alamos will participate are listed in the box "Collaborations with Industry on Parallel Computing."

So far we have concentrated on the Laboratory's initiatives in high-performance computing. But the digital revolution and the culture of the Internet has led us into a second major area—the National Information Infrastructure (NII) initiative. The goal of this federal initiative is to build the electronic superhighways as well as to develop the software and the hardware needed to bring to all groups in the population and all areas of the economy the benefits of the information age.

One feature of the information age is




A digitized fingerprint image consisting of 768 by 768 8-bit pixels.



The corresponding data-compressed image resulting from application of the WSQ algorithm. The compression ratio is 21:1.



Detail of the center image after WSQ compression.



Detail of a digitized fingerprint image.



Detail of the center image after "JPEG" compression.

Figure 4. Wavelet Compression of Fingerprint Data

The FBI has over 29 million active cards in their criminal-fingerprint files. These cards are now being digitized at a spatial resolution of 500 pixels per inch and a gray-scale resolution of 8-bits. Each card yields about 10 megabytes of data, and the potential size of the entire database is thousands of terabytes. The FBI came to the Laboratory for help in organizing the data. The Laboratory's Computer Research and Applications Group collaborated with the FBI to develop the Wavelet/Scalar Quantization (WSQ) algorithm, which has been made a public standard for archival-quality compression of fingerprint images. (The algorithm involves a discrete wavelet transform decomposition into 64 frequency sub-bands followed by adaptive uniform scalar quantization and Huffman coding.) WSQ compression introduces some distortion in the image. The figures demonstrate, however, that the important features of the fingerprint, including branches and ends of ridges, are preserved. In contrast, the international "JPEG" image-compression standard, based on a cosine transform, is applied not to the original image as a whole but rather to square blocks into which the image has been divided. The image resulting from "JPEG" data-compression shows artifacts of this blocking procedure.



the rapid accumulation of very large sets of digital data—in the gigabyte and terabyte range. Such data sets are generated in a wide variety of contexts: large-scale scientific computation, medical procedures such as MRI and CAT scans, environmental surveillance and geospatial exploration by satellites, financial transactions, consumer profiling, law enforcement, and so on. The abilities to “mine” the data—to analyze them for meaningful correlations—and to compress the data for rapid communication and ease of storage are of increasing interest. Also needed are intuitive user interfaces for manipulation of those very large and complex data sets.

The Laboratory has initiated several data-mining projects that use sophisticated mathematical tools to solve practical problems of data analysis, storage, and transmission. One project has resulted in a new national standard, created in collaboration with the FBI, for compressing digital fingerprint data with little loss of information.

Figure 4 shows an original and a compressed fingerprint image. It also compares the method adopted by the FBI, known as the Wavelet/Scalar Quantization (WSQ) algorithm, with a traditional data-compression method. The compressed images will be transmitted between local law-enforcement agencies and the FBI to assist in retrieving past records of suspected criminals. Because the data have been

compressed by a factor of 20, the images can be transmitted in minutes rather than hours.

The WSQ algorithm transforms each image into a superposition of overlapping wavelets, localized functions that vanish outside a short domain (see Figure 5) in contrast to the sine and cosine functions of the standard Fourier transform, which oscillate without end. The discrete wavelet transform includes wavelets on many length scales and automatically produces a multiresolution representation of the image. Thus an image can be retrieved at whatever resolution is appropriate for a particular application.

The algorithm developed for the fingerprint data is also being used to create a multiresolution database for the efficient storage and retrieval of very large images. Aerial photographs of the Washington, D.C. area, supplied by the United States Geological Survey (USGS), were first digitized. Through the discrete wavelet transform, the many separate digital images were put together into a continuous image so that no seams are visible. The resulting multiresolution database is illustrated in Figure 6, which shows the area around the Lincoln Memorial in Washington, D.C. at seven decreasing levels of resolution. At the coarsest resolution (64 meters/pixel) the entire D.C. area can be displayed on a workstation monitor. At the finest resolution (1




Figure 5. Mother Wavelets for the FBI Fingerprint Image Compression Standard

Shown here are the mother wavelets used in the WSQ algorithm for fingerprint data compression. The mother wavelets are translated and dilated to form the set of basis functions used to decompose and reconstruct the original image. The top wavelet is used for image decomposition and the bottom wavelet is used for image reconstruction.

meter/pixel) the user is able to distinguish features as small as an automobile. The USGS has an interest in making such data available for the




Figure 6. A Database of Seamless Multiresolution Images

The area around the Lincoln Memorial in Washington, D.C., is shown at seven resolutions. The digitized images were retrieved from a database constructed by using the discrete wavelet transform.

whole United States and disseminating the data on CD-ROM as well as on the electronic superhighways.

The multiresolution database project is part of a larger effort called Sunrise. The Sunrise project is a unique attempt to create an integrated software and hardware system that can handle many diverse applications of the kind envisioned for the nation's information superhighways—from access to very large databases, to secure electronic commercial transactions, to efficient communication of medical information in a national health-care system. The Sunrise strategy is to use a set of such diverse applications as a starting point for the development of software solutions, the elements of which are general enough to be used in many other applications.

The collaboration with radiologists at the National Jewish Center for Immunology and Respiratory Medicine on developing tools for telemedicine is illustrative of the Sunrise approach. The tools include a multimedia data management system that will display and analyze medical images, manage patient records, provide easy data entry, and facilitate generation of medical reports. The system is also designed to provide transparent access to medical information located anywhere on the information superhighway. Tools for interactive collaboration among physicians, efficient data compression, transmission, and storage, remote data storage

and retrieval, and automated data analysis are also being developed in the Sunrise approach to telemedicine (see "Concept Extraction" for a discussion of the quantitative analysis of LAM disease).

The Laboratory is going through dynamic change and nowhere is the change more visible than in the area of scientific computing, communications, and information systems. Because of the revolution in electronic communications, many people are doing things they never thought possible and in ways that could not have been anticipated (see "@XXX.lanl.gov" for a look into the future of research communication). Los Alamos is no longer focused solely on the high end of scientific computing for national security and basic research. We have become heavily involved in bringing advances in computing and information systems to all members of our Laboratory, to business and industry, and to the general public. We also expect that during the latter half of this decade advanced computer modeling and simulation will make increasingly direct and significant contributions to society. ■

Acknowledgements

We are grateful to the following people for their help with this article: Jonathan N. Bradley, Christopher M. Brislawn, John H. Cerutti, Salman Habib, Michael A. Riepe, David H. Sharp, Pablo Tamayo, and Bruce R. Wienke.

Further Reading

Yuefan Deng, James Glimm, and David H. Sharp. 1992. Perspectives on Parallel Computing. *Daedalus*, winter 1992 issue. Also published in *A New Era in Computation*, edited by N. Metropolis and Gian-Carlo Rota. MIT Press 1993.

N. Metropolis, J. Howlett, and Gian-Carlo Rota. 1980. *A History of Computing in the Twentieth Century*. Academic Press, Inc.

David W. Forslund is a Laboratory Fellow and Deputy Director of the Advanced Computing Laboratory. He is a theoretical plasma physicist who has contributed to controlled fusion and space plasma research and now specializes in distributed high-performance computing and NII technologies.

Charles A. Slocomb is the Deputy Division Director of the Computing, Information, and Communications Division. His primary interest is the future of high-performance computing and its application to scientific problems.

Ira A. Agins is a specialist staff member in the Computing, Information, and Communications Division. His avocation is computer history and particularly the history of computing at Los Alamos.

Collaborations with Industry on Parallel Computing

Bruce R. Wienke

The Computational Testbed for Industry (CTI) was established at the Laboratory in 1991 to provide U.S. industry with access to the computing environment at our Advanced Computing Laboratory and to the technical expertise of Los Alamos scientists and engineers. During this past year the CTI was designated officially as a Department of Energy User Facility. That designation affords us greater flexibility in establishing and implementing collaborative agreements with industry. The number of collaborations has been increasing steadily and will soon total about thirty. The seven projects described here are being established at the CTI through the new cooperative agreement between the DOE and Cray Research, Inc. under the auspices of the DOE's High Performance Parallel Processor program. The projects focus on developing scientific and commercial software for massively parallel processing.

Portability Tools for Massively Parallel Applications Development

Partners: Cray Research, Inc.; Thinking Machines Corporation

Goals: At present, software developed for one vendor's massively parallel computer system is not portable, that is, able to be run on other vendors' computers. The lack of portable programs has slowed the development of applications for every kind of massively parallel computer and the adoption of such computers by industry. This project will work toward removing that barrier by creating common programming conventions for massively parallel machines.

Massively-Parallel-Processing Performance-Measurement and Enhancement Tools

Partner: Cray Research, Inc.

Goals: Create a set of software tools to improve analysis of the system-level performance of massively parallel systems, to maximize their operating efficiency, and enhance the design of future systems. Plans include using this sophisticated automated toolkit to enhance the performance of applications developed in other projects under the cooperative agreement between the Department of Energy and Cray Research, Inc.

Lithology Characterization for Remediation of Underground Pollution

Partner: Schlumberger-Doll Research

Goals: Develop three-dimensional modeling software to cut the costs of characterizing and cleaning up underground environmental contamination. The software is intended for use by the petroleum and environmental industries on the next generation of massively parallel supercomputers.

Development of a General Reservoir Simulation for Massively Parallel Computers

Partners: Amoco Production Company; Cray Research, Inc.

Goals: Oil and gas exploration requires simulations of flow at several million points in reservoirs. Programmers have produced well-developed reservoir simulations for multi-processor vector supercomputers but not for massively parallel systems, so exploiting the potential of massively parallel computers is a high priority. The goal of this project is to adapt Amoco's field-tested reservoir-simulation software so that it performs efficiently on the massively parallel Cray T3D. The resulting program, which will allow much better management of reservoirs, will be made available to the entire petroleum industry.

Materials Modeling

Partner: Biosym Technologies Incorporated

Goals: In the competitive global marketplace for advanced materials, the traditional experimental approach to designing new materials needs to be complemented by materials modeling using high-performance computers. This project is aimed at creating powerful new visual-modeling software tools to improve casting and welding processes and to calculate the fracture properties of new materials designs, including composites.



At left: The Cray T3D was delivered to our Advanced Computing Laboratory in June 1994. It is a 128-processor machine that will be used primarily for collaborative research with industry. Below: Scientists and engineers from Cray Research, Inc. and Los Alamos at a dedication of the new machine.



Application of the Los Alamos National Laboratory Hydrocode Library (CFDLIB) to Problems in Oil Refining, Waste Remediation, Chemical Manufacturing, and Manufacturing Technology

Partners: Exxon Research and Engineering Company; General Motors Power Train Group; Rocket Research Company; Cray Research, Inc.

Goals: The speed and memory size of massively parallel systems will make it possible for U.S. industry to accurately model and improve the efficiency of chemical reactions that involve substances in more than one phase (solid, liquid, or gas). The project with Exxon will advance the simulation of multiphase reactors, which are heavily used in hydrocarbon refining, chemical manufacturing, gas conversion, and coal liquefaction and conversion. The goal of the General Motors project is to improve analysis of important foundry processes. One of the Rocket Research projects is aimed at improving small electric rockets used in satellite stations and has potential applications to microelectronics manufacturing and to neutralizing wastes in flue-gas emissions. Another Rocket Research project involves analysis of the performance, safety, and environmental impact of propellants used in automotive air bags and in fire-suppression systems of aircraft and other mass-transportation vehicles.

Massively Parallel Implicit Hydrodynamics on Dynamic Unstructured Grids with Applications to Chemically Reacting Flows and Groundwater Pollution Assessment and Remediation

Partners: Berea Incorporated; Cray Research, Inc.

Goals: Develop advanced software models to help U.S. industry better address problems involving combustion, pollution, and the treatment of contaminated groundwater and surface waters. These models could also be applied to designing engines, extracting more oil and gas from fields that have been drilled over, and assessing the structural integrity of buildings after a severe fire.