CONF-9410472--

HIGH-PERFORMANCE COMPUTING IN

SEISMOLOGY

NATIONAL RESEARCH COUNCIL

CONF-9410472-

HIGH-PERFORMANCE COMPUTING IN SEISMOLOGY

Committee on Seismology
Board on Earth Sciences and Resources
Commission on Geosciences, Environment,
and Resources
National Research Council

NATIONAL ACADEMY PRESS
Washington, D.C. 1996

Support for this study was provided by the National Science Foundation, the Air Force Office of Scientific Research, the U.S. Department of Energy, and the U.S. Geological Survey.

Additional copies of this report are available from:
Committee on Seismology
Board on Earth Sciences and Resources
National Research Council
2101 Constitution Avenue, NW-HA 372
Washington, DC 20418

The cover shows a snapshot of simulated strong ground motions (indicated by the dark areas) in the Los Angeles region at 50 seconds following a magnitude 7.75 earthquake on the San Andreas fault. The heavy white line is the Pacific coastline, the thin white lines are major freeways, and the vertical direction is approximately north-northeast. These and similar simulations show that resonating seismic waves in the underlying sedimentary basins increase the strength and duration (and hence the danger) of ground shaking from earthquakes in the Los Angeles area. This is one of the examples of high-performance computing in seismology discussed in the report. (Figure reprinted with permission from Olsen, K., R. J. Archuleta, and J. R. Materese. 1995. *Science,* 270, 1628-1632. Copyright 1995 American Association for the Advancement of Science.)

Printed in the United States of America

COMMITTEE ON SEISMOLOGY

THOMAS H. JORDAN, *Chairman*, Massachusetts Institute of Technology, Cambridge
WALTER J. ARABASZ,+ University of Utah
F. A. DAHLEN,* Princeton University, New Jersey
STEVEN M. DAY, San Diego State University
THOMAS C. HANKS, U.S. Geological Survey, Menlo Park, California
CHARLES A. LANGSTON, Pennsylvania State University, University Park
THORNE LAY, University of California, Santa Cruz
STEWART A. LEVIN, Mobil Exploration & Production Technical Center, Dallas, Texas
STEPHEN D. MALONE, University of Washington, Seattle
DONALD L. PAUL,+ Chevron Petroleum Technology Company, Houston, Texas
JAMES R. RICE, Harvard University, Cambridge, Massachusetts
PAUL G. SOMERVILLE, Woodward-Clyde Consultants, Pasadena, California
ANNE M. TREHU, Oregon State University, Corvallis
JOHN E. VIDALE, University of California, Los Angeles

National Research Council Staff

CHARLES MEADE, Study Director
WILLIAM E. BENSON, Senior Program Officer
KEVIN D. CROWLEY, Senior Program Officer
JUDITH L. ESTEP, Administrative Assistant

* *Term expired December 31, 1995.*
\+ *Term expired December 31, 1994.*

BOARD ON EARTH SCIENCES AND RESOURCES

J. FREEMAN GILBERT, *Chairman,* University of California, San Diego
THURE CERLING, University of Utah, Salt Lake City
MARK P. CLOOS, University of Texas, Austin
JOEL DARMSTADTER, Resources for the Future, Washington, D.C.
KENNETH I. DAUGHERTY, E-Systems, Fairfax, Virginia
WILLIAM R. DICKINSON, University of Arizona, Tucson, *emeritus*
MARCO T. EINAUDI, Stanford University, California
NORMAN H. FOSTER, Independent Petroleum Geologist, Denver, Colorado
CHARLES G. GROAT, University of Texas, El Paso
DONALD C. HANEY, Kentucky Geological Survey, Lexington
SUSAN M. KIDWELL, University of Chicago, Illinois
SUSAN KIEFFER, Kieffer & Woo, Inc., Palgrave, Ontario
PHILIP E. LaMOREAUX, P. E. LaMoreaux and Associates, Inc., Tuscaloosa, Alabama
SUSAN M. LANDON, Thomasson Partner Associates, Denver, Colorado
J. BERNARD MINSTER, Scripps Institution of Oceanography, La Jolla California
ALEXANDRA NAVROTSKY, Princeton University, New Jersey
JILL D. PASTERIS, Washington University, St. Louis, Missouri
EDWARD C. ROY, JR., Trinity University, San Antonio, Texas

National Research Council Staff

CRAIG M. SCHIFFRIES, Staff Director
THOMAS M. USSELMAN, Associate Staff Director
INA B. ALTERMAN, Senior Program Officer
WILLIAM E. BENSON, Senior Program Officer
KEVIN D. CROWLEY, Senior Program Officer
ANNE M. LINN, Senior Program Officer
CHARLES MEADE, Senior Program Officer
LALLY A. ANDERSON, Staff Associate
VERNA J. BOWEN, Administrative Assistant
JENNIFER T. ESTEP, Administrative Assistant
JUDITH ESTEP, Administrative Assistant

COMMISSION ON GEOSCIENCES, ENVIRONMENT, AND RESOURCES

M. GORDON WOLMAN, *Chairman*, The Johns Hopkins University, Baltimore, Maryland
PATRICK R. ATKINS, Aluminum Company of America, Pittsburgh, Pennsylvania
JAMES P. BRUCE, Canadian Climate Program Board, Ottawa, Ontario
WILLIAM L. FISHER, University of Texas, Austin
JERRY F. FRANKLIN, University of Washington, Seattle
GEORGE M. HORNBERGER, University of Virginia, Charlottesville
DEBRA S. KNOPMAN, Progressive Foundation, Washington, D.C.
PERRY L. MCCARTY, Stanford University, California
JUDITH E. MCDOWELL, Woods Hole Oceanographic Institution, Massachusetts
S. GEORGE PHILANDER, Princeton University, New Jersey
RAYMOND A. PRICE, Queen's University at Kingston, Ontario
THOMAS C. SCHELLING, University of Maryland, College Park
ELLEN K. SILBERGELD, University of Maryland Medical School, Baltimore
STEVEN M. STANLEY, The Johns Hopkins University, Baltimore, Maryland
VICTORIA J. TSCHINKEL, Landers and Parsons, Tallahassee, Florida

National Research Council Staff

STEPHEN RATTIEN, Executive Director
STEPHEN D. PARKER, Associate Executive Director
MORGAN GOPNIK, Assistant Executive Director
GREGORY SYMMES, Reports Officer
JAMES MALLORY, Administrative Officer
SANDI FITZPATRICK, Administrative Associate
SUSAN SHERWIN, Project Assistant

The National Academy of Sciences is a private, nonprofit, self-perpetuating society of distinguished scholars engaged in scientific and engineering research, dedicated to the furtherance of science and technology and to their use for the general welfare. Upon the authority of the charter granted to it by the Congress in 1863, the Academy has a mandate that requires it to advise the federal government on scientific and technical matters. Dr. Bruce Alberts is president of the National Academy of Sciences.

The National Academy of Engineering was established in 1964, under the charter of the National Academy of Sciences, as a parallel organization of outstanding engineers. It is autonomous in its administration and in the selection of its members, sharing with the National Academy of Sciences the responsibility for advising the federal government. The National Academy of Engineering also sponsors engineering programs aimed at meeting national needs, encourages education and research, and recognizes the superior achievements of engineers. Dr. William A. Wulf is interim president of the National Academy of Engineering.

The Institute of Medicine was established in 1970 by the National Academy of Sciences to secure the services of eminent members of appropriate professions in the examination of policy matters pertaining to the health of the public. The Institute acts under the responsibility given to the National Academy of Sciences by its congressional charter to be an adviser to the federal government and, upon its own initiative, to identify issues of medical care, research, and education. Dr. Kenneth I. Shine is president of the Institute of Medicine.

The National Research Council was organized by the National Academy of Sciences in 1916 to associate the broad community of science and technology with the Academy's purposes of furthering knowledge and advising the federal government. Functioning in accordance with general policies determined by the Academy, the Council has become the principal operating agency of both the National Academy of Sciences and the National Academy of Engineering in providing services to the government, the public, and the scientific and engineering communities. The Council is administered jointly by both Academies and the Institute of Medicine. Dr. Bruce Alberts and Dr. William A. Wulf are chairman and interim vice-chairman, respectively, of the National Research Council.

Contents

Executive Summary

Seismologists were among the first scientists to exploit the capabilities of advanced computing technology. Thirty years ago oil companies purchased the first "supercomputers" to process and analyze seismic reflection data in the search for petroleum. Shortly afterward, supercomputers were also used for seismological modeling in support of nuclear test detection. The field of seismology and society have both benefited tremendously from these activities. These projects engaged seismologists at the leading edge of technology, spurred the development of new computer technology, promoted a number of spin-off applications with important scientific implications (e.g., reflection seismic imaging of deep geologic structures), enhanced the discovery of energy resources, and supported initial efforts to limit the use of nuclear weapons by international treaties.

Since these early efforts, computer technology has evolved dramatically. Rather than supercomputing, the focus has shifted to the broader concept of high-performance computers matched to a diverse range of applications. The capabilities of high-performance computers may be used to display high-resolution graphics, perform complex model simulations, or archive large volumes of data. Today, high-performance computers are of critical importance to seismology because they are used at all stages of data acquisition, communication, modeling, and analysis. These applications include visualization of three-dimensional seismic images, automated processing of worldwide earthquake records, and complex simulations of earthquake processes.

For these reasons the future strength of seismology will be closely coupled to the degree that the field incorporates and utilizes advances in high-performance computing. This report outlines the opportunities and challenges for this effort, ranging over the seismological subdisciplines of earthquake monitoring and physics, comprehensive test ban treaty verification, petroleum exploration, and global and strong motion[1] seismology.

In its charge from the National Research Council, the Committee on Seismology was requested to identify the major computational challenges in seismology, assess specific research areas where emerging technologies may be decisive, and assess what is needed from technology to meet these computational challenges. As the first step in this process, a workshop on these issues was held at the San Diego Supercomputer Center in October 1994. Based on presentations at the workshop and from its own subsequent deliberations, the committee concludes that high-performance computers present significant opportunities for fundamental breakthroughs in seismological research.

Particularly promising areas for future research include simulations of earthquake processes, advanced modeling techniques for seismic wave propagation, and new methodologies for high-resolution images of the Earth's interior (from the crust to the core). The results of this work will have important implications for studies of geologic processes throughout the Earth and for understanding the causes of earthquakes. These applications also have tremendous practical applications that range from the mitigation of seismic hazards to locating undiscovered petroleum reserves. The importance of these problems provides strong incentives to develop high-performance computing applications throughout the field of seismology.

At present, the federally coordinated High Performance Computing and Communications Initiative (HPCCI) provides funding for a range of high-performance computing activities in seismology. The largest of these is the Advanced Computational Technology Initiative

---

[1] Strong motion seismology focuses on modeling and measuring the intense ground motions close to an earthquake source that are sufficiently large to cause damage to structures.

(ACTI) sponsored by the U.S. Department of Energy. Based on this program, high-performance computing activities in reflection seismology are the largest seismological component in federally funded HPCCI activities, exceeding support for other subdisciplines by more than an order of magnitude. Modeling of seismic wave propagation in three-dimensional sedimentary basins, one of the Grand Challenges of HPCCI, is the second largest.

The committee notes that there will be great challenges to making full use of high-performance computing technology throughout the field. Much of the problem stems from the recent shift toward parallel and distributed computing architectures as replacements for previous vector and serial machines. In principle, this transformation allows massive increases in computer speed and performance; however, it also requires a complete re-engineering of software and algorithms for scientific computing. In effect, the challenge of high-performance computing has shifted from issues of hardware design to problems of writing better software. Thus, taking advantage of advances in high-performance computing will require focused efforts to train and educate seismologists in the use of this technology. To promote these goals, the committee makes the following recommendations.

1. Focused efforts to develop validated documented software for seismological computations should be supported, with special emphasis on scalable[2] algorithms for parallel processors.

There has been little effort to validate or benchmark the software for a wide range of seismological problems. A strong effort in this area could facilitate the growth of a new generation of software for high-performance computing applications and would allow researchers to focus on the more advanced problems associated with developing software for parallel computers. It is critical that these efforts focus on scalable parallel algorithms because this will be the trend of future

---

[2] Scalable algorithms are designed so that their speed "scales" linearly with the number of processors in a parallel computer. The design and implementation of scalable algorithms are key factors in attaining high computational speeds with massively parallel architectures.

software developments in high-performance computing. The committee observes that much of the software development for scalable algorithms is currently feasible in academic environments (outside supercomputer centers) using networks of linked workstations and standardized software for parallel computing. Such an effort would increase the exposure of a wide range of seismologists to high-performance computing technology with relatively small investments in new computing facilities.

2. The education of seismologists in high-performance computing technologies and methodologies should be strengthened.

Breakthroughs in computational seismology will require a detailed understanding of high-performance computing software and hardware. The capabilities of seismologists in this area should be improved through broad educational efforts for researchers at all levels.

3. Collaborations between seismologists and computational scientists and engineers should be strengthened.

High-performance computing challenges in seismology are similar to a wide range of active research issues in computational science (e.g., I/O for large data sets, visualization). Seismology would benefit from increased interdisciplinary collaboration that includes computational scientists and engineers, but the mechanisms currently available to sponsor such collaborations appear to be limited. Activities of this type could be promoted through increased dissemination of research results from computational seismology and by greater participation by seismologists at workshops and conferences on high-performance computing. Also, collaborations between scientific societies could play an important role by facilitating cross-disciplinary forums and workshops for earth scientists and computational scientists and engineers.

4. The infrastructure for archiving, disseminating, and processing large volumes of seismological data should be expanded.

Seismology has entered a new era that is characterized by (1) significant increases in the volumes of recorded seismic data, (2) explosive growth in the number and size of centralized data archives, and (3) "real-time" recording from global seismic networks. Full utilization of these future data streams will require a significant upgrade of the infrastructure for data communication and storage. To facilitate this transition, there should be a sustained effort to exploit the capabilities of new technology for seismological data communication and archiving. Recent developments in real-time seismic monitoring for earthquake hazards and nuclear test detection represent important opportunities in this area. Also, there would be substantial benefits from developing widely accessible computer archives from the records of the thousands of continuously recording seismic stations throughout the world. In this last area, international scientific unions could play a key role in facilitating the development of international data archives.

1

Introduction and Background

Historically, computer performance has been measured by the speed of the fastest machines, and, by this standard, the advances have been remarkable. For the past 50 years, computational speeds have increased by two orders of magnitude per decade (Pfeiffer, 1994). Computer scientists expect this trend to continue for a number of years, but there is a growing realization that qualities other than speed are increasingly important to a wide range of computer applications. These include the capability to display high-resolution graphics, to transmit information over networks, to store large amounts of data on disks and in memory, and to perform many applications in an office setting. This diversified view of computer qualities has led to the new concept of "high-performance computing," which refers to the most powerful computing and communications technology *for a particular task*. Thus, high-performance computers may archive huge data sets, display complex images, or run detailed model simulations. With such a broad perspective, the standards and qualities of high-performance computing are quite fluid because of the rapid changes in the technology and capabilities of computers.

The emerging technology of high-performance computing has dramatically affected the science of seismology. Driven by the enhanced capabilities of computers to collect, analyze, and visualize information, the field is experiencing a huge surge in the volume of primary data streams, and there is growing interest in computational models that require advanced processors to operate on vast amounts of input data. At the same time, there is a growing realization that many of the critical problems in seismology are highly nonlinear, highly complex, or both. Examples of the latter include understanding the nature of earthquakes and providing detailed images of subsurface geologic structures. In this environment it has become clear that the future vitality of seismology will depend on how effectively the field can harness the full capabilities of high-performance computing technology to solve its most important problems.

To this end, the present report by the National Research Council's Committee on Seismology is intended to survey the status and potential of high-performance computing activities in the field of seismology. The motivation for this work is straightforward: taking advantage of future breakthroughs in high-performance computing will require a detailed understanding of the technology and how it can best be applied to seismological problems. This report is based in part on the presentations made at a workshop convened by the committee on October 2-4, 1994, at the San Diego Supercomputer Center. The workshop included briefings by more than two dozen experts representing computational sciences and seismology from the perspectives of academia, the petroleum industry, government research laboratories, and the consulting industry. The workshop participants are listed in Appendix A, and the workshop presentations are listed in Appendix B. The seismological disciplines represented at the workshop included earthquake monitoring and physics, comprehensive test ban treaty verification, petroleum exploration, and strong motion and global seismology.

In its charge from the National Research Council, the Committee on Seismology was requested to hold a workshop to (1) address recent developments in high-performance computing and the objectives and achievements of the High Performance Computing and Communications Initiative; (2) address the nature and scale of current problems in computational seismology; (3) identify the major computational challenges in seismology and assess where emerging technologies may be decisive; and (4) assess what is needed from technology to meet the computational challenges. This report integrates the results of the workshop with findings developed by the Committee on Seismology during subsequent deliberations.

HISTORICAL ROLE OF COMPUTING IN SEISMOLOGY

The field of seismology has advanced dramatically with the growth in computer capabilities over the past 20 years. For example, in the 1970s almost all earthquakes were detected with analog devices that could only record over a limited range of frequencies and amplitudes. Seismograms were recorded on paper, requiring time-consuming hand transcription before any computer analysis could be performed. Given these limitations and the primitive state of computers, only the simplest aspects of earthquake rupture and seismic wave propagation could be determined. By comparison, high-performance computers and communications networks are currently used at all stages of data collection, processing, and analysis in earthquake seismology. As a result, there have been great advances in understanding the causes and consequences of earthquakes (Jordan, 1995). These results, together with comparable advances in other subdisciplines of seismology, have demonstrated that **the field is critically dependent on high-performance computing technology to acquire and transmit data, to carry out simulations and analysis, and to visualize results. For this reason, future breakthroughs in seismology will be closely coupled to advances in high-performance computing and communications technology.**

Historically, seismologists have been pioneers in the use of high-performance computing, spurred by the investments of oil companies in new technologies for the exploration and development of fossil fuel resources. Beginning in the mid-1950s, the industry started using computers to process large volumes of seismic reflection data in an attempt to systematize the discovery of oil. Seismic reflection techniques were capable of providing subsurface images of geologic structures that significantly reduced the financial risk of exploratory drilling. Computational seismology thus provided great economic benefits, and in the 1960s oil companies began to purchase the first supercomputers for their own private use (Bloomquist, 1993). From the 1970s to the present, the continuously expanding efforts by oil companies to collect and process seismic reflection data have stimulated the development of new vector and scalar supercomputers by IBM and Cray, and, in a trend that has continued up to the present, oil companies have purchased prototypes of the most advanced vector and scalar supercomputers (Jordan, 1994). At present, the annual volume of seismic reflection data collected by industry is estimated to be greater than three petabytes (10^{15}), and it is growing at close to 20 percent per year (S. Levin, personal communication, 1995). **For at least the next 10 to 15 years, processing and analyzing this data stream will present a significant computational challenge for seismology that will continue to require the most advanced computational resources.**

Access to high-performance computers by seismologists outside industry has been shaped by federal policy regarding the support and development of the scientific computing infrastructure. To this end, there have been dramatic changes over the past 20 years. In the 1970s National Science Foundation (NSF) terminated its support for campus computer centers and there were no university "supercomputers." Other government agencies (e.g., the Advanced Research Projects Agency [ARPA], predecessors of the U.S. Department of Energy [DOE]) supported supercomputer facilities with moderate networking capabilities largely at the national laboratories. These programs supported classified seismological research on nuclear explosions. University seismologists without ties to the national labs or industry, however, were largely limited to minicomputers, which were extremely primitive machines relative to current technology. These restrictions placed severe constraints on the scope of computational research for seismological problems other than reflection seismic imaging and nuclear test detection.

Beginning in the 1980s, there were two key developments that have had a long-lasting impact on high-performance computing in seismology. First, the development of powerful workstations offered the first opportunity for powerful computing on a laboratory scale. Large numbers of seismologists adopted workstation technology in a trend that continues today. The growth of this technology has allowed academic and government seismologists to work on a wide range of new, computationally intensive problems that would have been impossible in previous decades. This work has included the development of tomographic techniques for resolving elastic heterogeneities through the Earth's mantle and detailed analyses of earthquake rupture mechanisms (Jordan, 1994). **Indeed, a review of the presentations made at the**

committee's workshop indicates that workstations used for a wide range of applications (visualization of supercomputer results, networking, data archiving) are an important component of high-performance computing in seismology. Since the 1980s, workstations have spread rapidly through the scientific community because NSF and other funding agencies supported their purchase. Today, the purchase and upgrading of workstation technology for individual researchers continue to be important components of the nation's policy to support scientific computing. This policy has important implications for the level of high-performance computing in seismology because of the relatively large role of workstations compared to supercomputer centers.

The second important development in the 1980s was the beginning of a large federal effort to build a national infrastructure to support supercomputing at universities. These programs were spurred by several reports pointing to a severe shortage of advanced academic computing facilities which hindered research for a broad range of scientists and engineers. The most influential study to address the problem was a National Science Board report, **Report of the Panel on Large Scale Computing in Science and Engineering**, edited by Peter Lax (Lax, 1982). Its recommendations were revolutionary at the time, and they continue to be an important component of federal policy on high-performance computing. The Lax report specifically called for government action to increase access to regularly upgraded supercomputing facilities via high-bandwidth networks; increase research in computational mathematics, software, and algorithms; train people in scientific computing; and invest in research on new supercomputer systems.

In response to the Lax report, NSF initiated a program to build and support five supercomputer centers for academic and industrial researchers throughout the country. The centers were sited at major universities and were operational by 1986. (Four of the centers remain in service today, but, their future is uncertain.) The impact of the centers was almost immediate. First, they significantly increased the volume of supercomputer CPU hours for scientific and engineering computations in the United States. Second, they signaled the beginning of a major federal investment in a national infrastructure for high-performance computing. In addition to funding the supercomputer centers, NSF began construction of NSFnet (a predecessor of the current Internet) to provide communication between the centers and users at remote sites.

These initiatives to support supercomputing were extremely successful throughout the sciences. Between 1986 and 1990 the total number of users at NSF supercomputer centers increased by more than a factor of five; the allocated CPU hours (accounting for increases in computer performance) jumped by almost an order of magnitude; and there was a large increase in the participation of graduate, college, and high school students (National Research Council, 1995a). Analyses of supercomputer center users showed little geographic bias, indicating that access to supercomputing, facilitated by the evolving Internet, was spreading throughout the science and engineering communities in the United States (National Research Council, 1995b).

During this period of growth, academic and government seismologists were relatively small participants in NSF's supercomputer centers. But there is clear evidence that the field of seismology reaped large benefits from the federal effort to construct a national infrastructure for supercomputing. The most significant of these benefits was the development of the Internet to connect users to the supercomputer centers and to each other. The Internet has been a boon for seismology because it provides a means to obtain records from remote seismographic stations in near real time that can then be compiled into large data sets and transferred over great distances. Easy access to these data sets by individual scientists has in turn spurred a rapid growth in computational seismology for modeling and analysis. The Internet will continue to play an important role in archiving and disseminating seismic data, as well as in supporting emerging technologies such as distributed computing. Seismology has also benefited from the "spillover" effects of supercomputing technology. The features of the advanced computers that were developed for NSF's supercomputer centers were rapidly transferred to workstation technology for large numbers of seismologists (and other scientists) in subsequent years. There is abundant evidence that this trend continues today, as seismologists are now utilizing workstations based on moderately parallel architectures.

Because of the federal government's efforts, high-performance computing methods became an important tool for a large number of scientists and engineers in the United States by the end of the 1980s. In many fields, computational science associated with simulation, modeling, and analysis assumed a comparable role to the traditional methods of theory and experiment. Consequently, computational science (distinct from computer science) has emerged as a new field to address computing problems across a range of disciplines within the physical and biological sciences and engineering. The breadth and significance of this transition demonstrated that government could play a critical role in fostering the growth of new computer technologies. To build on this success, there was considerable interest in coordinating the efforts of federal agencies with regard to high-performance computing by the end of the 1980s. Within the Office of Science and Technology Policy (OSTP), several interagency coordinating committees began to focus on issues of high-performance computing. Through the work of these committees and a program plan published by OSTP in 1989, an initial High Performance Computing and Communications Initiative (HPCCI) was formed. The initiative focused on coordinating the efforts and budgets of the four major agencies involved in computational research: DOE, NSF, National Aeronautics and Space Administration, and the Department of Defense (through the Advanced Research Projects Agency). By 1992 the HPCCI received legislative support with the passage of the High Performance Computing Act. This law authorized a 5-year program to coordinate federal funding for high-performance computing across a broad range of agencies (initially 10 were involved). The goals of the HPCCI as stated in the law were to extend U.S. leadership in high-performance computing and networking; disseminate new technologies to serve the economy, national security, education, health care, and the environment; and spur gains in U.S. productivity and industrial competitiveness.

One of the key mechanisms for achieving these goals has been the focus of resources on select Grand Challenges and National Challenges. Grand Challenges are defined as "fundamental problems in the sciences and engineering with broad economic and scientific impact, whose solutions require the applications of high-performance computing." National Challenges are viewed as "fundamental applications with broad and direct impact on the Nation's competitiveness and the well-being of its citizens." By design, these programs involve a multidisciplinary approach because the scope of the problems is beyond a single research area, scientific discipline, or computer. In 1997 the HPCCI will end its 5-year term. For fiscal year 1995 the budget was $1.1 billion, spread over 12 departments and agencies. Plans are under way to extend the program beyond 1997.

The HPCCI has significantly affected computing activities throughout science and engineering. Each year its accomplishments and goals are described in the "Blue Book" published by the National Coordinating Office for High Performance Computing and Communications within the OSTP (National Coordination Office, 1994). In addition, the HPCCI has been reviewed in three external reports: the National Research Council's **Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure** (National Research Council, 1995a); the NSF Blue Ribbon Panel's report, **From Desktop to Teraflop: Exploiting the U.S. Lead in High Performance Computing** (NSF, 1993), and the U.S. General Accounting Office's **High Performance Computing and Communications: New Program Would Benefit from a More Focused Effort** (U.S. GAO, 1994). Within the scope of the present report and the concerns of computational seismology, the most important impacts of the HPCCI have been (1) a shift of emphasis to the concept of "high-performance" rather than "super" computing, (2) the development of new funding opportunities for computational seismology, and (3) the establishment of parallel computing as the architecture for the next generation of supercomputers. These issues are discussed below.

HIGH-PERFORMANCE COMPUTING

Perhaps the broadest impact of the HPCCI has been the growing emphasis on "high-performance computing" as the driving theme for the development of computational resources in the United States. Within the scientific and engineering communities this has created a recognition that solutions to important problems depend on more than processor speed. Consequently, there is a growing focus on such problems as (1) scalable I/O,[3] (2) the need for new database technologies for massive data archives, (3) the continuing demand for improved workstation performance for visualization, and (4) the importance of improved network communications in the development of distributed computing.

With the emphasis on high-performance computing, the target community for national computing policies has broadened significantly beyond the users of NSF's supercomputer centers. Indeed, the most recent Blue Book cites the development of the Mosaic program and the World Wide Web, used by millions of people, as one of the HPCCI's greatest accomplishments. Within this context, the HPCCI is focused on developing an integrated infrastructure for computing and communications that can accommodate the largest number of users at levels ranging from desktop computers to massively parallel machines.

FUNDING FOR HIGH-PERFORMANCE COMPUTING IN SEISMOLOGY

At present, the HPCCI and related programs provide a range of funding for high-performance activities in seismology. The largest of these is the Advanced Computational Technology Initiative (ACTI) sponsored by DOE. This program is targeted at technology transfer from the national laboratories to promote the competitiveness of domestic oil and natural gas companies. Much of the program's funding is on a cost-sharing basis with industry, and roughly a third of its budget is focused on the development of technology for reflection seismology. In 1995 funding for high-performance computing applications related to processing, inversion, and modeling of reflection seismic data consisted of $9 million from DOE and $12 million from industry. DOE supported work at the national laboratories and at universities with research programs in seismology. Industry contributions supported their own expenses while participating in ACTI. For 1995, $3 million of the ACTI

---

[3] Scalable I/O uses software and hardware designed to scale the input and output speed to the number of processors in a parallel computer.

funds were counted as part of the HPCCI. (Specific ACTI projects are discussed in the next chapter.) Based on the funding for ACTI, **high-performance computing activities related to the processing, modeling, and inversion of reflection seismology are the largest seismological component in federally funded HPCCI activities, exceeding the support for other subdisciplines by more than an order of magnitude.**

The HPCCI provides support to high-performance computing in seismology through a number of smaller programs, including the Grand Challenge program, and the Scalable I/O Initiative and through the NSF's earth sciences program. Of the 16 Grand Challenges, one involves seismology: earthquake ground motion modeling in large basins. This particular project is based at Carnegie Mellon University; however, a number of other computational groups in the United States receive NSF support to carry out similar simulations (see Chapter 2 for examples). **These observations indicate that modeling of seismic wave propagation in three-dimensional sedimentary basins is the second largest seismological component in federally supported high-performance computing programs.** Finally, the Scalable I/O Initiative provides limited funding for work on software and hardware development for applications that require huge amounts (terabytes) of input data. This work funds the benchmarking of codes for parallel processing of reflection seismic data.

PARALLEL COMPUTING

Throughout the 1970s and 1980s, supercomputers were based on scalar or vector architectures that utilized one or a handful of specially designed processors to achieve extremely high speeds. Much of the cost (and speed) of these machines was associated with the effort to design and construct the processors. In the late 1980s it was recognized that supercomputer speeds could be achieved at relatively low cost by arranging large numbers of off-the-shelf workstation processors in a parallel architecture. This approach exploited the difference between the price and performance of workstations relative to supercomputers: a workstation processor might be 10 times slower yet cost a thousand times less than the CPU for a scalar supercomputer. Because the (idealized) speedup for parallel computing is proportional to the number of processors, a cost-effective supercomputer could utilize large numbers of these mass-produced workstation CPUs.

At the beginning of the 1990s, parallel computing was an emerging technology with an uncertain future. Within the field of high-performance computing, there was considerable disagreement about its utility for supercomputer applications, and none of the established supercomputer manufacturers offered parallel models. In this context the HPCCI played a critical role in demonstrating and establishing parallel computing by supporting experimental machines and software development at NSF's supercomputer centers. As a result of HPCCI support, there has been a dramatic change of opinion regarding parallel computing with a corresponding shift in the market for supercomputers. Today, it is widely viewed that scalar and vector supercomputers are close to their last generation.

While the transition to parallel computing has enhanced the performance of supercomputers, it also presents significant problems. This point is emphasized in a recent National Research Council (1995a) report in which it is said that "In parallel computing the fundamental challenge is not building the machines, but learning how to program them" (p. 6). Parallel computers can only attain high speeds if models and algorithms are optimally designed for parallel computation. This approach is not straightforward, however, for nonlinear problems that involve complex or interdependent phenomena and that are intrinsically difficult to divide for parallel computation. As a result, the focus of high-performance computing has shifted away from technological concerns of building faster computers to scientific questions involving new methods of inquiry.

In the past few years there have been changes in the design of parallel computers that have important implications for software and algorithm design. These modifications involve the detailed configuration of memory, CPUs, and communication networks. Optimizing the performance of a parallel computer requires codes that are specially tailored for the hardware design. Because of the proliferation of computer designs and changes in architecture over time, libraries of standardized software for scientific computing on a broad range of platforms are in short supply. Indeed, writing software for parallel supercomputers can often require custom subroutines for such basic operations as matrix multiplication and Fast Fourier Transforms. At present, the trends in hardware (and hence software) development in parallel computing are highly uncertain. **It is clear that future increases in computational speed will rely on a close coupling of advances in the hardware and the software for parallel computing.** Presentations at the committee's workshop indicated that there is considerable interest in developing standardized protocols to support distributed computing over networks of linked workstations. Workshop participants indicated that this technology could flourish with the development of scalable algorithms that will support parallel computing on a wide range of scales.

2

Scientific Challenges

This chapter briefly describes some of the current computational challenges in seismology. These issues span a wide range of disciplines and often have significant economic and social implications, such as the mitigation of seismic hazards, verification of a Comprehensive Test Ban Treaty for nuclear weapons, and increased discovery of economically recoverable petroleum resources. Indeed, it is these rewards that have demonstrated the importance of high-performance computing in seismology and that provide a great incentive for supporting future research in this area.

Consistent with the qualities of high-performance computing, the computational requirements for these problems are variable. Some of the challenges require the fastest and most advanced supercomputers, whereas others utilize the capabilities of workstation technology. Recognizing these inherent differences in scope and difficulty, the purpose of this chapter is to describe the scientific and computational aspects of a wide range of high-performance computing problems in seismology, with emphasis on the potential areas for future breakthroughs.

Computational problems in seismology concern the recording, analysis, and simulation of seismic waves in the Earth. They also involve the analysis of fault rupture, explosions, and the nonlinear response of materials to high-amplitude seismic waves. In this work the underlying physics of linear seismic wave propagation is well understood for simplified hypothetical models. In the Earth, however, seismic waves display a rich complexity because of the extreme heterogeneity of geologic materials and the dynamic nature of earthquake processes. In recent years the greatest computational challenges in seismology have been driven by growing scientific, economic, and social interest in the details of these phenomena. To this end, there has been great interest in high-resolution images of the crust and mantle, dynamic simulations of earthquake rupture, and complete models of earthquake effects (e.g., Su et al., 1994; French, 1992; McLaughlin and Day, 1994; Olsen et al., 1995). Simulating or analyzing these processes with even a moderate resolution in space or time often requires state-of-the-art computing resources.

At present, the major challenges in computational seismology fall into two classes. First, there are problems for which the underlying physics and mathematics are adequately understood but are nevertheless intractable because of computational limitations. For example, seismologists engaged in petroleum exploration and development know how to construct accurate three-dimensional images of subsurface structures using prestack depth migration techniques (French, 1992). There would be tangible economic returns from exploiting this well-understood methodology, but current supercomputers are not fast enough to make the computations feasible except for limited regions of the crust (Abriel, 1994). The problem of computing the ground motion of large sedimentary basins during earthquakes is another example of an important problem in which the physical and mathematical formulations are well understood. But three-dimensional calculations covering the full range of frequencies of interest to structural engineers are currently not possible.

Second, there are research problems in which high-performance computing resources are required to advance scientific understanding of seismic wave generation and propagation. For example, seismologists are confronted with challenging computational problems in efforts to produce earthquake simulations for faults with small slip weakening or frictional state transition slip distances. Another example is the problem of understanding the effects of geometric disorder of fault zones on spatial-temporal patterns of earthquakes. Such understanding may be an important part of future progress in the field of intermediate-term earthquake prediction.

EXAMPLES OF RESEARCH ISSUES

Strong Motion Response of Sedimentary Basins

Sedimentary basins trap and amplify earthquake-induced ground shaking in a complex manner that depends on the details of the basin geometry and the local geology. In effect, basins resonate after an earthquake, and at particular frequencies and orientations, these effects can have devastating consequences. The Michoacan earthquake of 1985 demonstrated this phenomenon: the epicenter of the magnitude 8.1 event was located 300 km from Mexico City, yet resonance within the sediments underlying the city contributed to severe ground shaking. In the United States there are several large population centers situated on similar sedimentary basins with high exposure to earthquake hazards (e.g., Los Angeles, Salt Lake City).

Compared to the high spatial resolution required for this problem, current supercomputers are only able to calculate ground motions from hypothetical earthquakes in simplified models of sedimentary basins (see Figure 2.1, p. 29). These simulations are restricted to the low-frequency components of the seismic energy, and the response of near-surface soils is treated in a highly simplified way. Despite their simplicity, these calculations have revealed important site-specific amplification effects such as out-of-plane scattering, focusing, and resonance. Extending these simulations to shorter periods, incorporating the nonlinear viscoelastic properties of soils, and accounting for realistic basin geometries are crucial for assessing the vulnerability of urban structures to earthquakes. Because of the difficulty of the numerical calculations and the great social importance of this issue, simulations of seismic wave propagation in sedimentary basins have been identified as a Grand Challenge problem within the HPCCI. Future improvements in these models will require increases in supercomputer speeds, new algorithms for parallel computation schemes, and enhanced spatial resolution in models of regional geology and soils.

Earthquake Rupture Processes

A closely coupled problem involves the modeling of earthquake rupture processes in the lithosphere. Although it is widely recognized that earthquakes originate from slip along geologic faults in the Earth's crust, there is little understanding of the processes that control the distribution and occurrence of seismicity. This is a complex issue that spans a huge range of time scales and distances. From historical studies it has become clear that seismic activity may not be uniform over time or space (e.g., Dolan et al., 1995; Rice, 1993). Faults that have produced earthquakes of small magnitude in the recent past may be capable of very different behavior in the future. Understanding this behavior will be extremely difficult because earthquakes are not driven by isolated faults. Rather, they are the result of interactions within complex fault zones that are composed of networks of interacting faults of various sizes. A complete understanding of earthquake hazards requires information about the behavior of these networks of crustal faults. To this end, numerical simulations have become an essential tool.

With current supercomputers, only systems of very limited size can be investigated and only with rather rough approximations to the governing physics (see Figures 2.2 and 2.3, pp. 31 and 33). One of the greatest challenges for this work is the need to treat the details of fracture on a minute scale, together with the large-scale dynamics of slip along the entire fault. For simulations over an entire earthquake cycle (about 100 years for portions of the San Andreas Fault), the time scales must account for both the slow pace of stress accumulation and the rapid nature of fracturing. Because of ambiguities in the physics of earthquakes, many of the model simulations are run on workstations to address a limited subset of the problem (e.g., Ben-Zion and Rice, 1993). Much of this work is limited to two-dimensional approximations of fault zone geometries, although there have been recent attempts to model in three dimensions (Ben-Zion and Rice, 1995). As seismologists gain experience with simple models and future improvements in high-performance computing technology become available, numerical study of the dynamics of complex fault systems promises to be an important computational challenge that will greatly improve our understanding of long-term earthquake potential. This line of inquiry may even provide insights into the possibility of observing short- and intermediate-term precursors in advance of earthquakes.

Computation of Seismic Wave Propagation

The majority of computational seismology focuses on the simulation and analysis of seismic records from natural and artificial sources. The results of this work are used to construct subsurface images for geologic studies and oil exploration and to support worldwide monitoring of nuclear explosions. The scientific issues in these studies are based on three physical principles. First, seismic wave velocities are material properties that vary with the chemical composition, permeability, porosity, and fluid content of rocks. Thus, determinations of subsurface velocity variations provide information about changes in geologic properties. Second, seismic waves are reflected, refracted, or scattered from interfaces between dissimilar seismic wave velocities. Finally, the transmission of seismic energy is strongly influenced by the elastic anisotropy that is common in geologic materials. Together, it is these phenomena that govern the complexity of the signal that follows the first arrival on a seismogram.

There has been great interest in developing efficient algorithms for detailed simulations of seismic wave propagation through realistic geologic media. Such calculations provide a foundation for a wide range of problems in computational seismology, and they might be performed billions of times in a single model. In practice, realistic simulations are difficult because of the significant heterogeneity that characterizes natural rocks and because of the complex wavefields produced by multiple reflection, refraction, and other forward- and back-scattering events. For example, finite-difference simulations of wave propagation in realistic geologic models can involve tens of millions of grid points, gigabytes of computer memory, and gigaflop-hours of CPU time to simulate the propagation from a single seismic source. To alleviate such severe computational requirements, there has been considerable effort to develop alternative approximate solutions to the wave equation (see Figure 2.4, p. 35).

Discrimination Between Earthquakes and Nuclear Explosions

Seismic wave propagation calculations will play an important role in the verification of a nuclear Comprehensive Test Ban Treaty because seismic recording networks provide the primary technology for detecting nuclear explosions on a global basis (see Figure 2.5, p. 36). Because the recorded signal includes both source and path effects, the challenge will be to identify clandestine explosions by comparing computational models of seismic wave propagation through the Earth with observed seismograms. These simulations, performed for a wide range of geologic terranes and seismic sources, will be used to discriminate suspected nuclear explosions from the many natural and man-made events that produce measurable seismograms (e.g., earthquakes, quarry excavation, induced seismicity from oil recovery). Differentiating between such events will require a deep understanding of the physical processes that produce small changes in recorded signals.

Seismic Imaging for Natural Resource Exploration

Wave propagation calculations are also a critical component of seismic reflection imaging techniques for petroleum engineering. These images are developed by recording the signal from artificial seismic sources using dense grids of seismometers. For each seismic shot the network records hundreds of seismograms that are individually sensitive to the characteristics of the path between the source and the receiver. The subsurface geologic images are then refined by fitting portions of calculated seismograms to the measured data. In recent years reflection seismology has focused on new computational algorithms and data acquisition techniques to allow increasing levels of detail and accuracy in the model inversions (French, 1992) (see Figure 2.6, p. 39). The capabilities of massively parallel computers have played an important role in this effort. Simply visualizing the huge amount of information in these models has become an important high-performance computing application in itself (see Figure 2.7, p. 41). Future advances in subsurface imaging will utilize entire seismic waveforms to invert for variations in elastic properties (rather than just reflection properties). On a detailed scale, with high-frequency data, such inversions are only possible today for limited two-dimensional sections (see Figure 2.8, p. 43).

Because these images reduce the financial risk of drilling exploratory wells, there is great economic incentive to develop algorithms and computer hardware to increase the resolution of subsurface seismic models. For this reason the petroleum industry has been a significant commercial user of supercomputing technology for the past 30 years. In the future, processing of reflection seismic data will continue to be one of the largest computational challenges for seismology. For example, the input data from a typical seismic survey currently exceed 1 terabyte. Using these data to refine a three-dimensional model over a 100-km^2 region with a 25-m grid spacing requires approximately 2×10^{15} floating point operations. On a state-of-the-art commercial supercomputer, such a model might run for about 10 days. At present, there is great interest in extending imaging techniques to larger areas; however, the scale is strongly limited by the capabilities of existing computer hardware.

In some respects the computational characteristics of reflection seismic imaging are ideally suited to the qualities of massively parallel supercomputers. Calculating large numbers of synthetic seismograms between all sources and receivers in a large three-dimensional model can easily be parceled to thousands of individual parallel processors. In contrast, however, the data requirements for these models are staggering. Because these models are constrained by terabytes of input data, simply getting the information into the computer and storing the associated variables in memory present a significant computational bottleneck. In large part this is an issue of computer design that is decoupled from the concern of raw processor speed. Problems that involve the transfer of large volumes of data will be limited by the slowest link in the system, which is usually not part of the central processor (e.g., disk I/O, data bus).

Because of the extreme data requirements, processing of seismic reflection data has been identified as one of the activities within the Scalable I/O Initiative of the HPCCI. The purpose of this initiative is to study the I/O needs of I/O-intensive programs on testbed computers and to use the results to guide the development of software and computer languages for I/O-intensive applications.

The only way to calibrate the resolution of seismic images and to assess the accuracy of algorithms for inverting seismic data is to perform forward simulations of the complete wavefields. The computational requirements for this effort are enormous, however, because of the wide range of length scales that control the qualitative and quantitative nature of the data. For applications to oil exploration, the Society for Exploration Geophysics and the European Association of Exploration Geophysicists are striving to simulate a three-dimensional seismic reflection survey across two hypothetical geologic structures (a salt dome and an overthrust belt). This work is supported in part by DOE's Advanced Computational Technology Initiative and involves collaborations with scientists at the national laboratories. By calculating the seismic response of a "known" model, these simulations will allow a benchmark calibration of the different inversion methods are used to infer unknown structures from measured acoustic data. Under the most favorable assumptions, the simulation of a single model survey (over 400 km^2) will consume an aggregate 900 hours of supercomputer time, distributed over four national laboratories, for a total of 7,000 gigaflop hours. Yet this simulation will involve only a limited portion of the acoustic spectrum. A realistic simulation incorporating the full elastic wavefields would require much more computer time. Only through future hardware advances will computer simulations achieve their full potential to guide industry seismologists in the processing and interpretation of three-dimensional reflection surveys.

Imaging the Earth's Plates

Techniques developed by the petroleum industry are being extended to large-scale imaging of assemblages in the Earth's crust (see Figure 2.9, p. 45). Historically, seismic data have been analyzed from the perspective of discrete geologic units (e.g., for a particular petroleum-bearing formation). Recent high-resolution, wide-aperture seismic data, coupled with geologic observations, indicate that the

Earth's crust is heterogeneous at all resolvable scales and that the severity of heterogeneity depends on the tectonic environment (e.g., Lafond and Levander, 1995; Holliger and Levander, 1994). A more precise understanding of crustal structure is fundamental for a wide range of geologic studies and practical applications (e.g., knowledge of subsurface linkages among faults and magmatic systems is important for assessing earthquake and volcanic hazards).

Because analytical elastic wave theory cannot describe wave propagation through heterogeneous geologic materials, numerical simulations provide the only means to predict seismic response in realistic settings. Using numerical integration techniques such as finite differences, complex crustal models require fine discretization (20 to 60 m) in large two-dimensional models (200 × 50 km) to simulate elastic wave propagation in the appropriate bandwidth (10 to 40 Hz). Iterative, two-dimensional, finite-difference modeling is feasible at some institutions using super-scalar or parallel machines. At present, three-dimensional, finite-difference simulations of realistic crustal models are prohibitively time consuming and costly, and no academic institution has the computational resources to simulate large portions of the crust in the appropriate bandwidth. In the future, massively parallel machines with super-vector performance on individual nodes, in combination with advances in algorithms and theoretical research, will be required to address realistic seismic reflection and refraction problems with finite-difference methods.

Whole-Earth Imaging

For imaging applications over much larger scales, seismologists utilize data from earthquakes throughout the world in combination with tomographic techniques to create three-dimensional images of seismic velocities over the Earth's mantle and core (National Research Council, 1993) (see Figure 2.10, p. 47). At present, these models are constrained by travel-time data for body and surface waves and the waveforms from long-period seismograms that are sensitive to the normal modes of the Earth (Su et al., 1994). Given the volume of input data and the inversion techniques, the spatial resolution of these models is still quite coarse (500 to 1,000 km). Future increases in detail will be provided by coupling wave propagation calculations over a wide range of frequencies to new nonlinear inversion techniques. For high-resolution images over the vast scale of the Earth's mantle, these calculations are beyond the capabilities of current supercomputers.

Data Acquisition and Archiving

In recent years there has been a significant increase in the volume of automatically recorded and telemetered seismic data from global and regional networks of seismometers. Used for basic geologic research, earthquake monitoring, and nuclear test verification, and collected by a wide range of government, university, and international organizations, this data stream is greater than 10 gigabytes per day and is growing at an annual rate of approximately 10 percent. At present, the Incorporated Research Institutions for Seismology (IRIS) makes a large fraction of these data available using rapid automated data archiving and retrieval networks. At present, more than 2 terabytes of recorded seismic information can be randomly accessed through the IRIS archives. Although the current monitoring and communication network operates with commercial technology, the issues of automated seismic data acquisition and archiving may evolve to be important computational challenges in the future. In particular, there will be a growing need to link the large number of data archives and recording stations that are spread throughout the world. The rate of this transition will depend on the number and sampling interval of future seismic instruments, the eventual size of the historical archives of worldwide seismic data, and the evolving need for real-time global seismic records.

HIGH-PERFORMANCE COMPUTING EXAMPLES

Three-Dimensional Simulation of Seismic Wave Propagation in a Sedimentary Basin

The disasters of the Michoacan, Mexico (1985) and the Loma Prieta (1985) earthquakes have motivated research into the amplification of seismic ground motions in alluvial basins. Until recently, such predictions have been limited to one- and two-dimensional modeling. With current high-performance computing resources, it is now possible to extend the analysis to fully three-dimensional elastic modeling (Olsen et al., 1995). Although limited to somewhat low-frequency simulations (less than about 1 Hz) because of the extensive computational requirements, the three-dimensional modeling accounts for important site-specific effects such as out-of-plane scattering, focusing, and resonance.

Figure 2.1 shows the time progression of seismic wave propagation in the Los Angeles and San Fernando valleys for a hypothetical M 7.75 earthquake on the San Andreas fault. Hot (cool) colors depict positive (negative) ground velocities in the direction parallel to the fault. The thin lines show the major freeways of the Los Angeles area. The thick line is the coastline. The simulation utilized a model with 23,209,875 cubes of 0.4-km dimension, and it required 23 hours of CPU time on a 512 processor parallel computer.

Such an earthquake would have potentially devastating effects because of its orientation with respect to downtown Los Angeles. The results of these and other simulations for other faults in the region show that large peak motions can occur at great times and distances from an earthquake's epicenter because of resonance effects that amplify the local ground motions in the sedimentary basin. In this simulation such resonance is observed at times between 50 and 100 s after the first seismic waves traverse the basin. Away from the epicenter, strong ground motions are often observed above the deep parts of the basins and near the steep edges. In general these simulations show extreme spatial heterogeneity of ground shaking across the valley, indicating that damage to buildings and infrastructure will be sensitive to location (see Olsen et al., 1995).

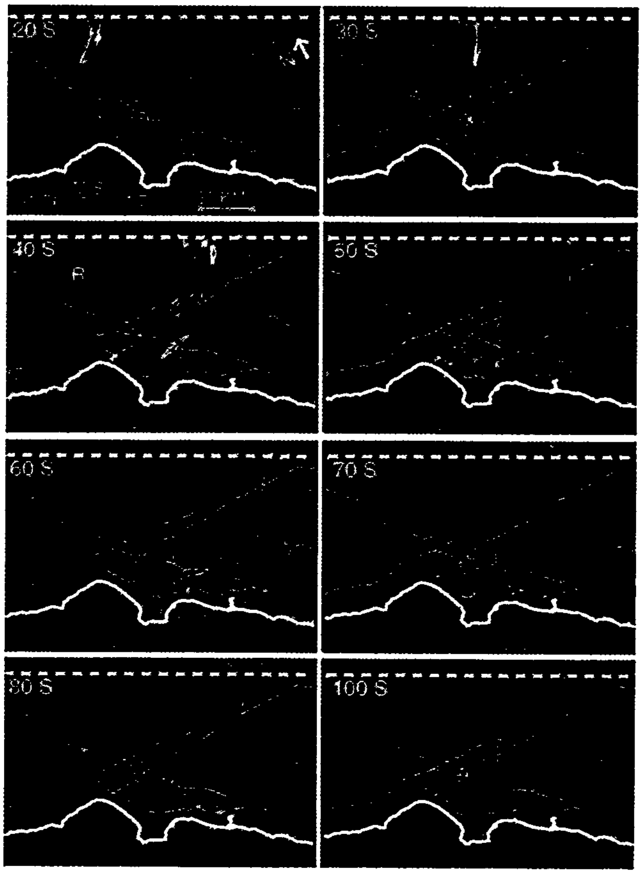

FIGURE 2.1 Simulated seismic wave propagation through the Los Angeles and San Fernando Valleys from a M 7.75 earthquake on the San Andreas fault. (Figure reprinted with permission from Olsen, K., R. J. Archuleta, and J. R. Materese, 1995. *Science,* 270, 1628-1632. Copyright 1995. American Association for the Advancement of Science.)

A Nonlinear Problem: Interacting Earthquake Faults

The dynamics of earthquake fault rupture have been investigated using finite-difference simulations since the 1970s. Most of this work has idealized the fault zone as a single planar surface, and the nonlinear processes have been subsumed into a friction law. However, real earthquake faults are complex, with recognizable bends and offsets, delineating a series of distinct segments along the fault length. Individual fault rupture events (earthquakes) often stop when they encounter segment boundaries. On some occasions, however, ruptures are able to jump across one or more segment boundaries, producing a longer rupture involving multiple noncoplanar fault segments and resulting in a larger earthquake. The Landers, California, earthquake of 1992 is a good example of such a multiple-segment rupture: it jumped over fault segments spaced 1 to 2 km apart.

Recent finite-difference simulations provide an improved understanding of the interactions between offset fault segments during rupture (Harris and Day, 1993). Simulations show that segment offsets of several hundred meters pose a relatively weak barrier to rupture, that distances of 1 to 2 km (as in the Landers earthquake) will delay but not stop an earthquake, and that separations of 3 to 5 km pose a formidable barrier to rupture extension. Furthermore, the amount and sense of offset of the fault trace influence the time evolution of the rupture and therefore the nature of the ground shaking (e.g., Olsen et al., 1995).

As an example of this behavior, Figure 2.2 shows the absolute value of the ground velocity for a simulated earthquake that ruptures two fault segments offset by 1 km. The region shown is 20 by 40 km. Locations of the faults are indicated by the dashed lines in the first panel. The earthquake initiates at the far left and propagates toward the right. The first panel shows the velocity field 2.7 s into the simulation, just before the rupture encounters the offset in the fault trace. In the subsequent panel the rupture stops at the fault segment. After a pause, it then jumps to the right-hand fault segment, and extends farther to the right as movement diminishes on the left-hand fault segment. Two-dimensional simulations such as these can be performed by using workstation technology; however, fully three-dimensional models are beyond the capabilities of current high-performance computing resources.

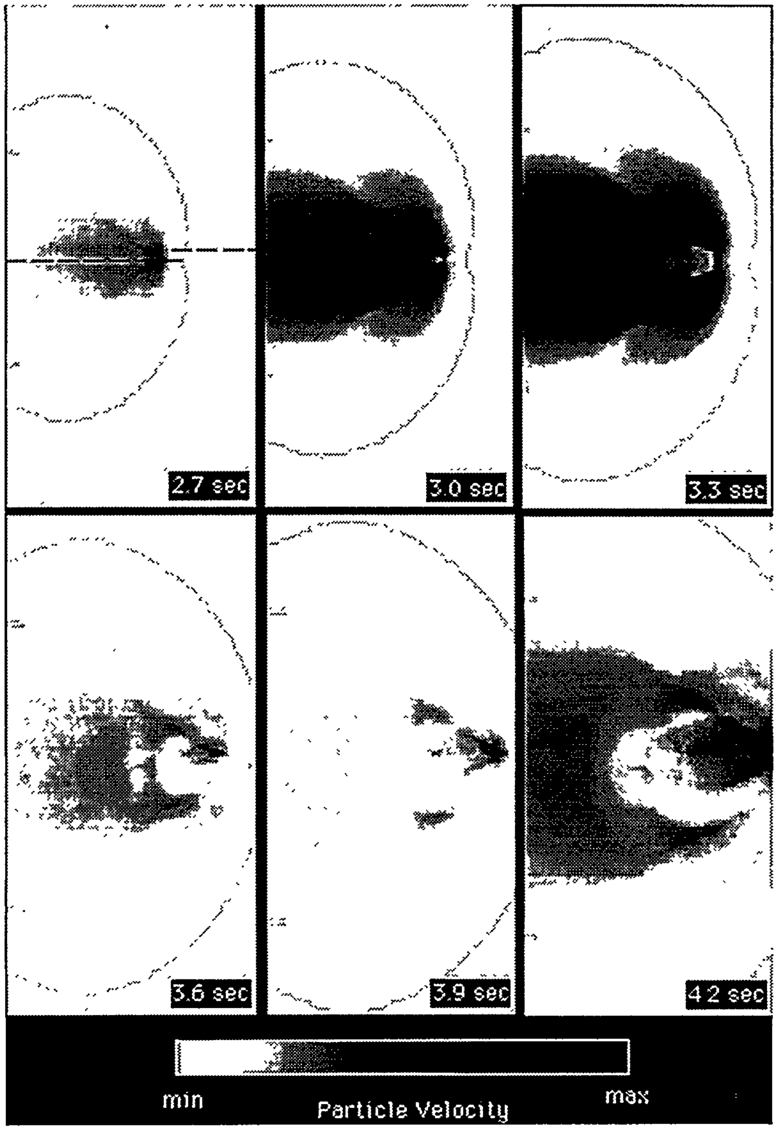

FIGURE 2.2 Time progression of two dimensional earthquake rupture on two offset fault segements. (Figure courtesy of R. Harris and S. Day.)

Elastodynamic Simulations of Earthquake Rupture

What are the response of fault zone materials and the segmented geometry of faults to the mechanics of geologic-scale stressing? How do these processes influence the onset, recurrence, rupture, and propagation modes of earthquakes? Addressing these issues in seismic source theory requires numerical simulations that combine fault and crack dynamics. The purpose of this work is to develop a model for fault instabilities (earthquakes) that incorporates both the large-scale mechanical regimes and the small-scale complex fracture heterogeneities. The challenge is to simulate crustal stressing and seismicity over hundreds of years, during which the fully elastodynamic equations are required for only a few seconds near sporadic earthquakes. For this reason the implementation requires highly nonuniform time steps that vary by more than 10^{10} in some cases (being shortest during the earthquake slip).

Faults in these studies are modeled as surfaces of displacement in a three-dimensional continuum. Modeling is carried out in a fully dynamic context where feasible, with laboratory constitutive relations and geologic boundary conditions. Much of the work on this problem is focused on developing scalable models to simulate components of the earthquake cycle. As greater understanding is gained, the models can be combined for large-scale parallel computation.

An example of this work, Figure 2.3 illustrates the magnitude of slip at different positions and times for a rupture event (earthquake) along a fault surface. In the simulation, accelerating slip begins under quasi-static conditions over a slowly growing zone of nucleation shown at a depth of about 5 km at times between 0 and 3 s. From this region, rupture breaks out into a fully dynamic process. The rupture history is complex, with a first propagation upward to the surface and then a few seconds later a propagation downward toward the base of the seismogenic zone, which induces a step of further slip in the shallow crust. These simulations were performed on a computer workstation; however, the methodology has been scaled to elastodynamic simulations using 512 parallel processors (e.g., Perrin et al., 1995).

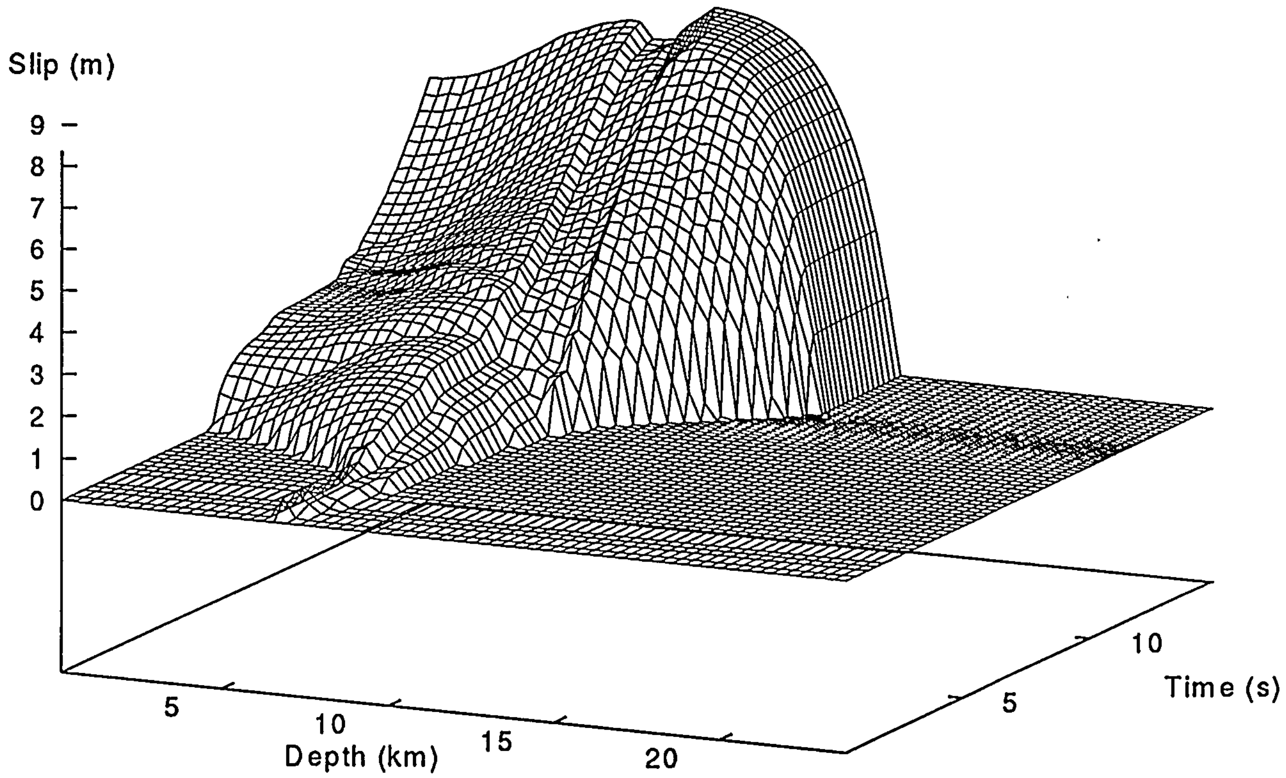

FIGURE 2.3 Dynamic simulation of earthquake rupture on a fault surface. The figure illustrates the amount of slip at depth on the fault during a 15 *s* interval that begins with the initiation of rupture (Figure courtesy of Y. Ben-Zion and J. R. Rice.)

Computational Methods for Simulations of Seismic Wave Propagation

Three-dimensional elastic wave propagation in a heterogeneous medium poses a major computational challenge for seismology. Complete and accurate wave field calculations can only be performed for small models, which precludes their application for oil exploration and imaging of deep geologic structures. To address this problem there has been a great effort to develop computationally efficient methods that account for the most important seismic wave effects. Figure 2.4 illustrates the results for one of these techniques: the "one-way" approximation to the elastic equations that accounts for the forward-scattered energy. The figure shows the time progression (indicated in seconds at the lower right of each frame) for a compressional seismic wave traversing an elastically heterogeneous region (bounded by the two light blue lines in each frame). The box is approximately 4 × 12 km in dimension (the tick marks are spaced at 1 km). At $t = 1.0$ s, the wave enters the heterogeneous media. The colors are proportional to the wave amplitude. Red (blue) colors denote positive (negative) amplitudes. As the wave enters the heterogeneous zone at $t = 1.0$ s, the elastic energy is uniformly distributed. As the wave traverses this zone it become corrugated, and amplitudes vary on the wave front due to focusing and defocusing effects. Converted compression to shear wave energy trails the initial wave front starting at $t = 2.0$ s. These calculations contain most of the transmitted wave effects, and they are thousands of times faster than complete finite-difference methods. Even with this gain in speed, however, only supercomputers can use these techniques to perform realistic calculations for oil exploration applications.

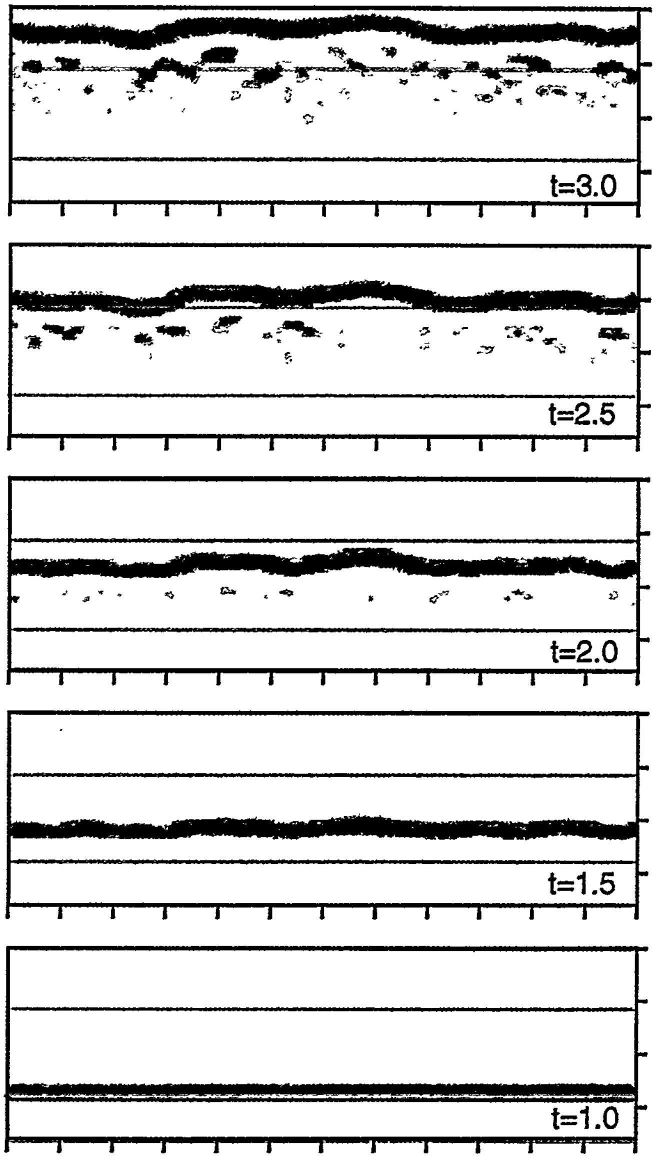

FIGURE 2.4 Seismic wave propagation through an elastic heterogeneity calculated with the "one way" approximation to the elastic wave equations. (Figure courtesy of Xiao-Bi Xie and Ru-Shan Wu.)

Simulations of Subsurface Nuclear Explosions

Figure 2.5 shows the result of a nonlinear finite-difference calculation of the pressure (denoted by the gray scale) and particle velocities (denoted by arrows) at 0.0061 s after a 0.38-kiloton nuclear explosion in an elliptical cavity surrounded by a geologic salt formation. Calculations such as these are used to investigate the seismic detection of possible clandestine violations of a Comprehensive Test Ban Treaty. The seismic waves radiating from an underground nuclear explosion are greatly reduced by detonation in a large cavity. Current research focuses on understanding the relationship between the radiated seismic signal and the details of the cavity design.

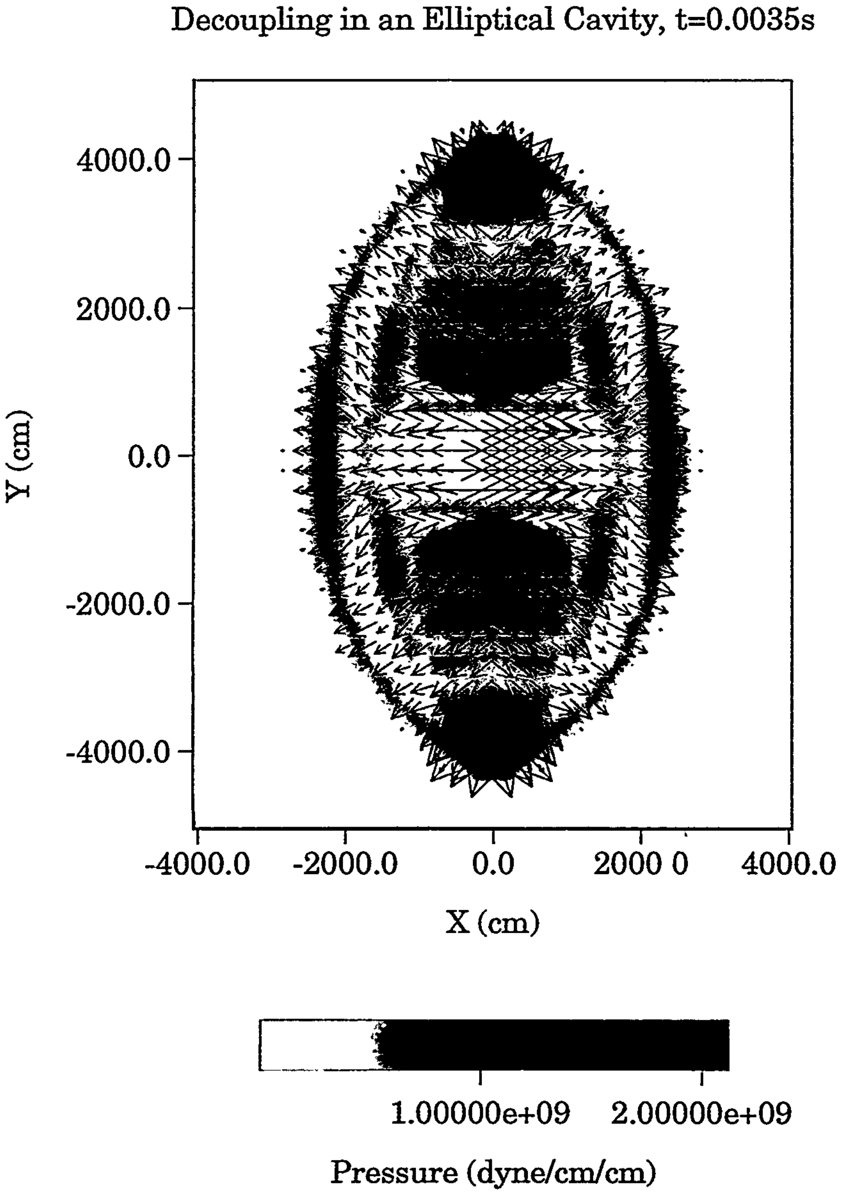

FIGURE 2.5 Finite difference simulation of the stresses induced by a nuclear explosion in an elliptical cavity (Figure courtesy of J. L. Stevens and K. McLaughlin.)

Three-Dimensional Seismic Imaging

With the capabilities of the fastest supercomputers it is now possible to construct three-dimensional seismic reflection images over sizable sedimentary basins. Figure 2.6 shows the results for such an area that is 10 km wide, 2.4 km across, and 5 km deep. The dark (light) colors shows the presence of strong (weak) reflectors of seismic energy. This model was derived from a reflection seismic survey using parallel arrays of seismometers and a two-dimensional grid of different shot points. The model was refined using prestack depth migration, a technique that involves billions of synthetic seismogram calculations to identify subsurface structures that provide the best match to the measured reflection data. These modern imaging techniques have two important implications for petroleum exploration efforts. First, the depth migration techniques provide structurally accurate images of subsurface structures. A comparison of these model results with previous images from fast approximate time migration techniques (see Figure 2.6) shows that many spurious features are removed when the data are appropriately imaged in depth. Second, three-dimensional results dramatically improve the ability of geologists to identify promising structures prior to drilling.

For the future there is considerable interest in extending the scale and resolution of three-dimensional volume images. The use of full wave propagation calculations to fit the entire reflected seismic waveforms should allow a dramatic increase in resolution. At present, these calculations are not feasible, except for small, two-dimensional regions. Imaging deep salt features in the Gulf of Mexico and complex geologic terranes (such as overthrust belts) will also be important challenges. There is also great interest in developing techniques for high-resolution imaging of the elastic properties of the crust (see Figure 2.7, p. 41).

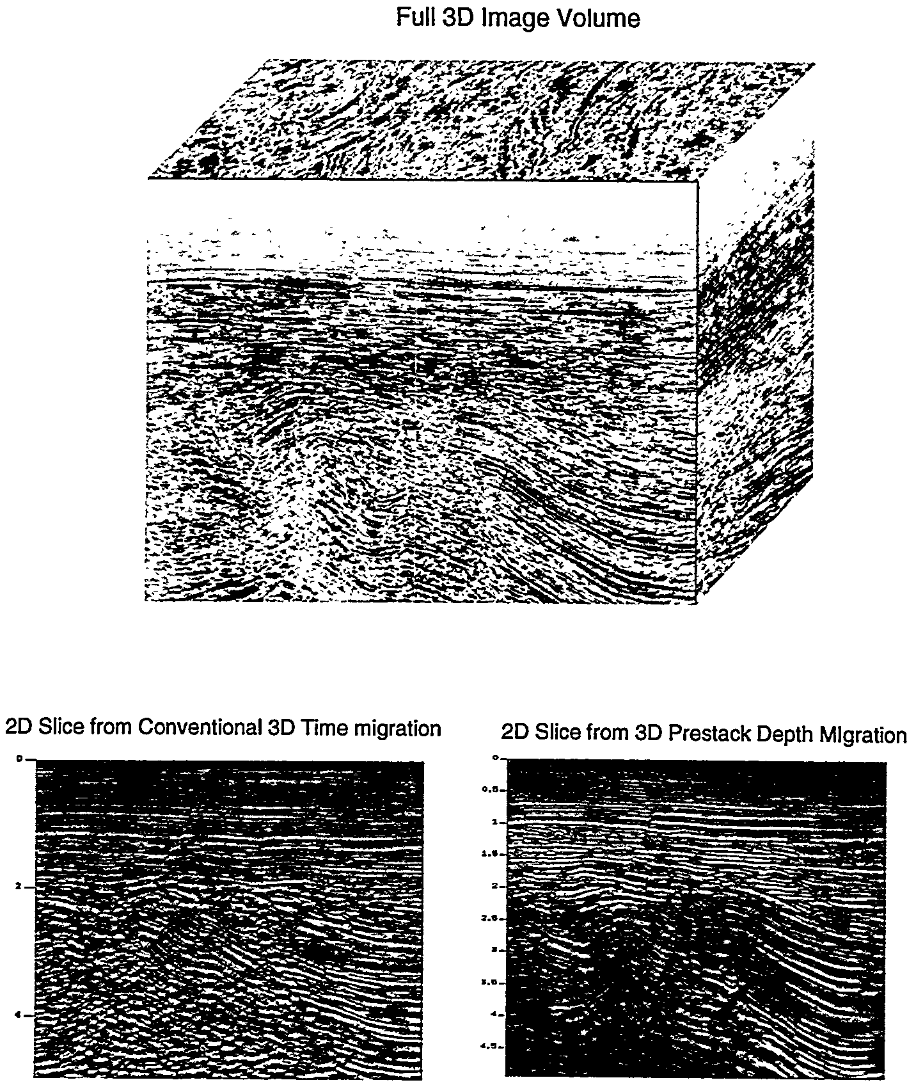

FIGURE 2.6 Three-dimensional seismic reflection image over a region 10 km wide, 2.4 km across, and 5 km deep. The bottom panels compare two dimensional slices of the data derived with two different inversion techniques. The depth migration methods provided the most accurate images. (Figure courtesy of S. Levin.)

Visualization of Subsurface Seismic Images

Modern exploration and development of petroleum reserves depend on three-dimensional seismic imaging to provide information about the nature of the reservoir. A typical modern three-dimensional seismic survey contains a terabyte of data prior to processing. The final processed seismic volume may consist of over 3 billion data points representing the response of the Earth over an area of approximately 300 km^2 to a depth of 6 km and occupying 12 gigabytes of disk space. This volume of data must be interpreted in detail to develop an understanding of the reservoir and to optimally position the development wells.

The most effective way to understand and interpret this volume of data is to visualize it on a computer screen, so that the three-dimensional results can be understood. Simple planar slices through the model are of limited use in this effort (see Figure 2.6, p. 39). Figure 2.7 illustrates such a three-dimensional interpretation over a depth range of approximately 4 km, showing two surfaces that are important for locating development wells in the North Sea gas basin. These surfaces have been selected (and interpreted) from the results of a large, three-dimensional seismic survey. The surrounding material is invisible in this rendering to highlight the location and properties of these geologic units. The colors on the surfaces represent the strength of the reflected seismic waves (red, strong; blue, weak). The lighting emphasizes the faults that cut and displace the surfaces. The locations of these faults play an important role in the transport of hydrocarbons.

These technologies are critical to the economics of developing new reservoirs. Interpreting detailed reservoir surfaces and fault patterns and visualizing the reservoir structures in three dimensions allow a better understanding of reservoir structure and potential flow paths. By identifying drainage blocks and connectivity in a reservoir, a development drilling program can be optimized, which minimizes development costs and maximizes the hydrocarbons recovered. At a current cost of $10 million per development well in the North Sea, the economic impact of this technology can be substantial.

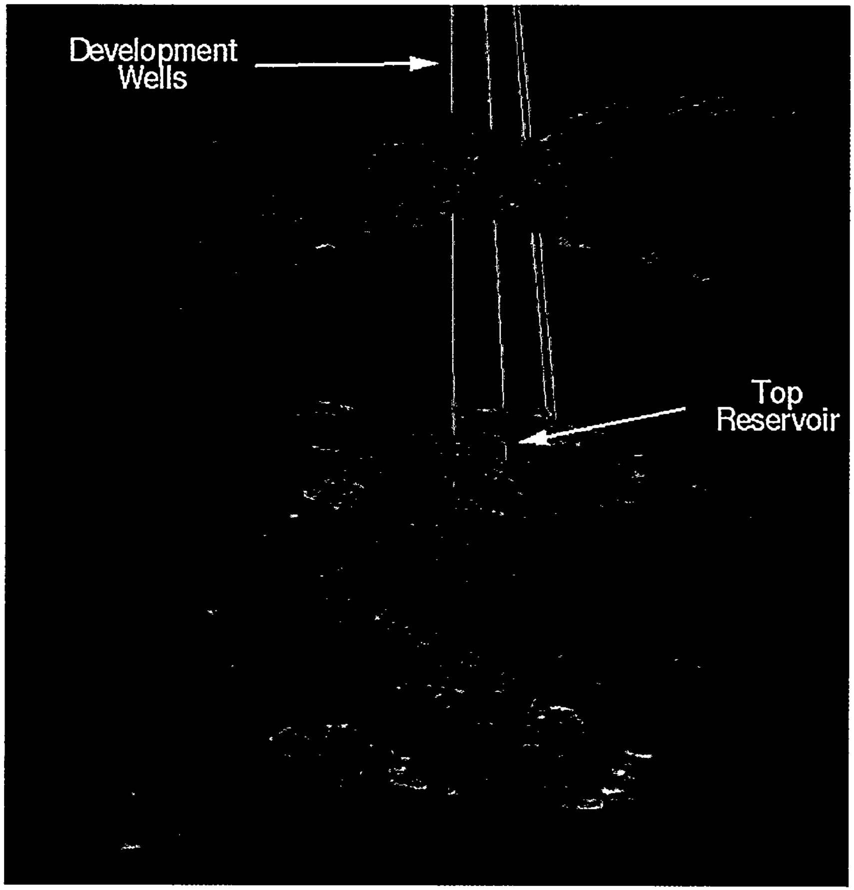

FIGURE 2.7 Three dimensional visualization of the results from a reflection seismic survey. (Figure courtesy of G. Dorn.)

Elastic Travel Time and Waveform Inversion of Cross-Hole Seismic Data

Enhanced oil and gas recovery operations significantly increase the yield of hydrocarbon reservoirs. One of the most promising tools for this effort is the use of cross-well seismic tomography to image local geologic structures and to develop optimal recovery operations at a particular site. Utilizing explosives and vibrators placed down a source well and seismic data recorded from adjacent wells, cross-well tomography determines the three dimensional velocity structure in the interwell medium.

Historically, cross-well tomography has only utilized measurements of the travel time for seismic waves between the source and the receiver. This approach provides a coarse resolution of the interwell structure: with few independent observations, one cannot specify the large number of unknown parameters in a model of the local geology. In an effort to increase the resolution of cross-well tomography, recent analyses have emphasized the use of amplitude data (i.e., the loudness and quietness) as well as travel-time information.

Figure 2.8 illustrates the improvements from these new analytical methods for a single cross-section. In this figure, the wells are located along the vertical edges. The variations in color represent differences in acoustic velocity that can be correlated with changes in rock type. The tomogram on the left is derived from models utilizing only travel-time information. The image on the right uses an inversion for both the travel time and the amplitudes of the acoustic pulses. A comparison of the two images shows that the new techniques enhance the spatial resolution by an order of magnitude and that they resolve previously undetected layering and structures. The computational demands for these improvements, however, are enormous. Today, a typical application involves simulations of thousands of time steps on a grid of 300,000 points, requiring the solution of more than 10^{13} algebraic equations for the inversion of the entire data set. This corresponds to more than a week of computing on a multigigaflop computer. Future applications will require simulations over larger models (i.e., large oil fields), with concomitant demands on computing resources.

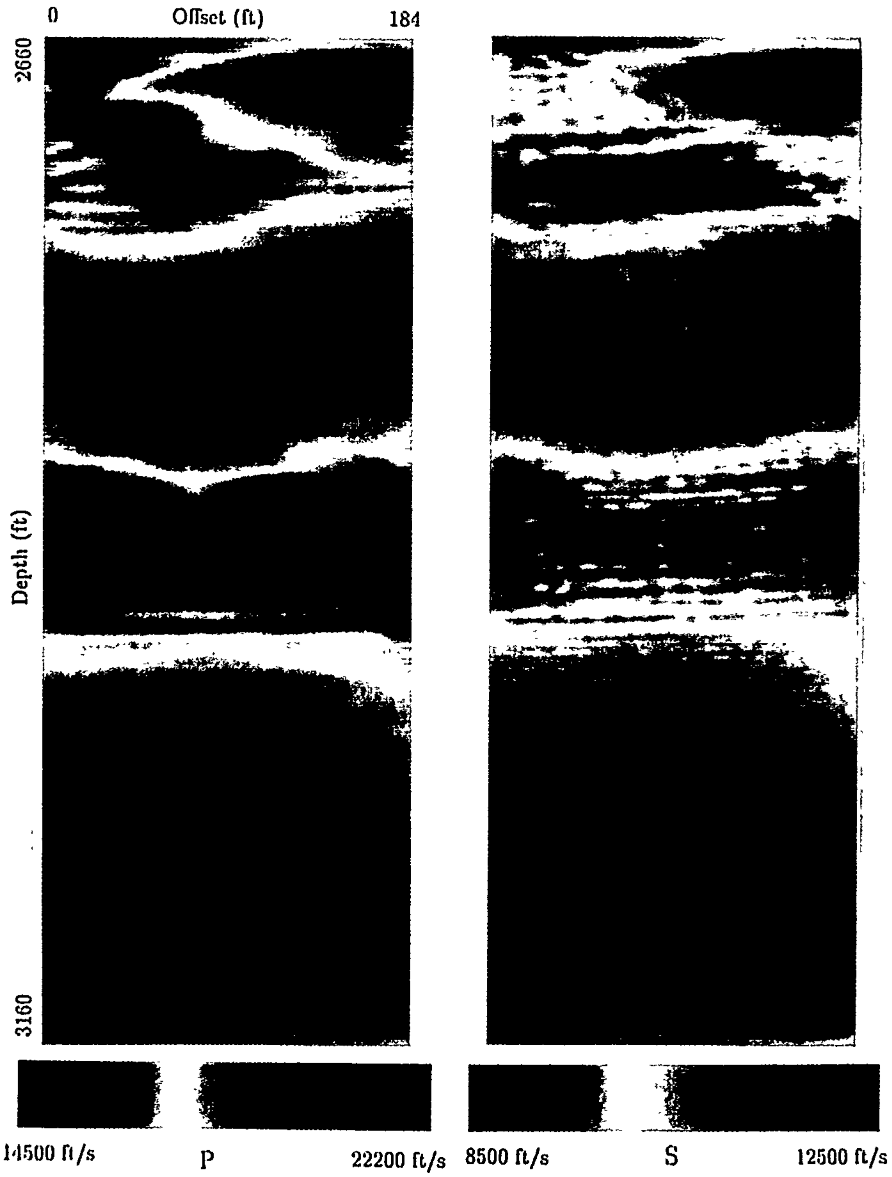

FIGURE 2.8 Comparison of cross-well tomographic images using conventional techniques (left) and full waveform inversions (right). (Figure courtesy of G. Schuster.)

Simulation of the Seismic Response for Realistic Crustal Models

Figure 2.9 shows the results of two-dimensional wave propagation calculations for a model of the Basin and Range in Nevada. The left portion is an example of the heterogeneity of seismic wave velocities over a limited portion of the model. The distribution of seismic velocity fluctuations is indicated schematically. The upper crust (to 11.5 km) has self-affine Gaussian velocity fluctuations. The middle crust and the lower crust have three zones of self-affine binary fluctuations that represent different geologic units. Below 30 km the mantle is uniform and there are no free surface effects. The figure also compares the results of a simulated large-aperture survey across the heterogeneous model with a measured profile across the Basin and Range. In both the simulations and the measurements the times of the first arrivals are sensitive to the average velocity of the section. Most of the subsequent seismic signal results from fluctuations in the seismic velocities resulting from heterogeneities in the geologic units. The general features of the simulated section compare favorably with the data. Improving the match between the model and the observations will require more detailed considerations of scattering from velocity heterogeneities in the crust. These two-dimensional simulations can be performed using extended runs on advanced workstations.

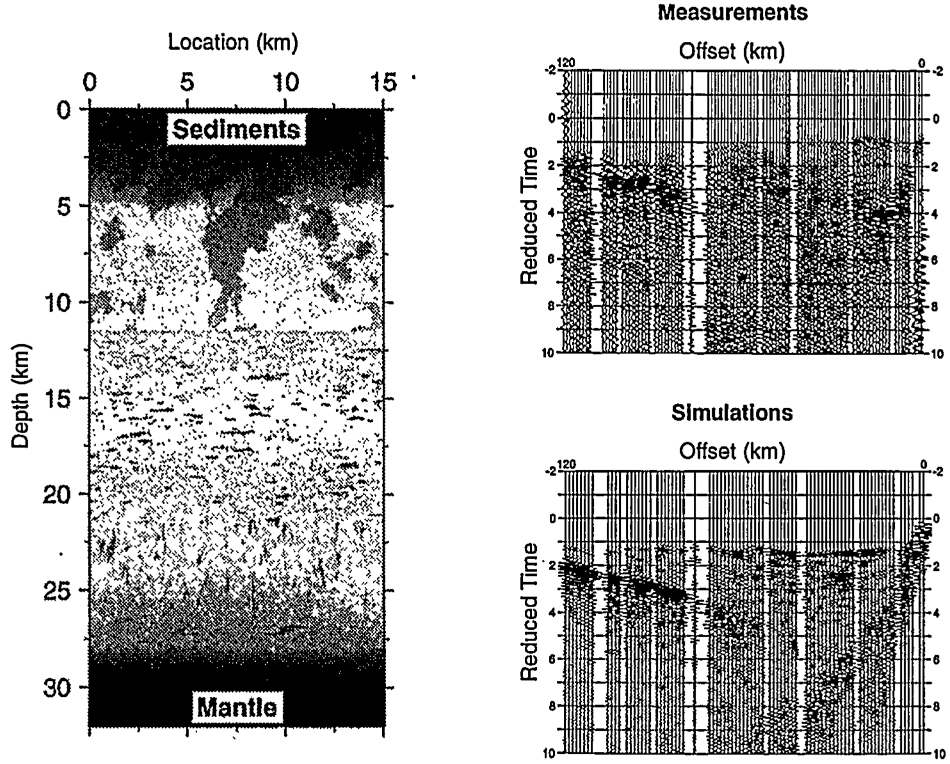

FIGURE 2.9 Simulations and measurements of a wide aperture seismic survey across the Basin and Range in Nevada. The elastic heterogeneity for the model calculations is illustrated schematically (right). (Figure courtesy of A. Levander.)

Tomographic Images of the Earth's Interior

Seismologists use the records from thousands of earthquakes around the world to image elastic heterogeneities throughout the deep interior of the Earth. Figure 2.10 shows the results of such a three-dimensional model based on the analysis of 14,000 travel-time observations and 27,000 long-period seismograms. The model extends from a depth of 1,000 km to the boundary between the lower mantle and the core (2,891 km in depth). The red (blue) shapes are isosurfaces for seismic shear velocities that are 0.3 percent slower (faster) than the spherically averaged value through the Earth at that depth. The viewpoint is at 50,000 km above the Pacific Ocean. This model has been low pass filtered so that only spherical harmonics 1 through 4 are plotted. The results of this and similar models provide important information for constraining tectonic processes deep in the Earth's mantle.

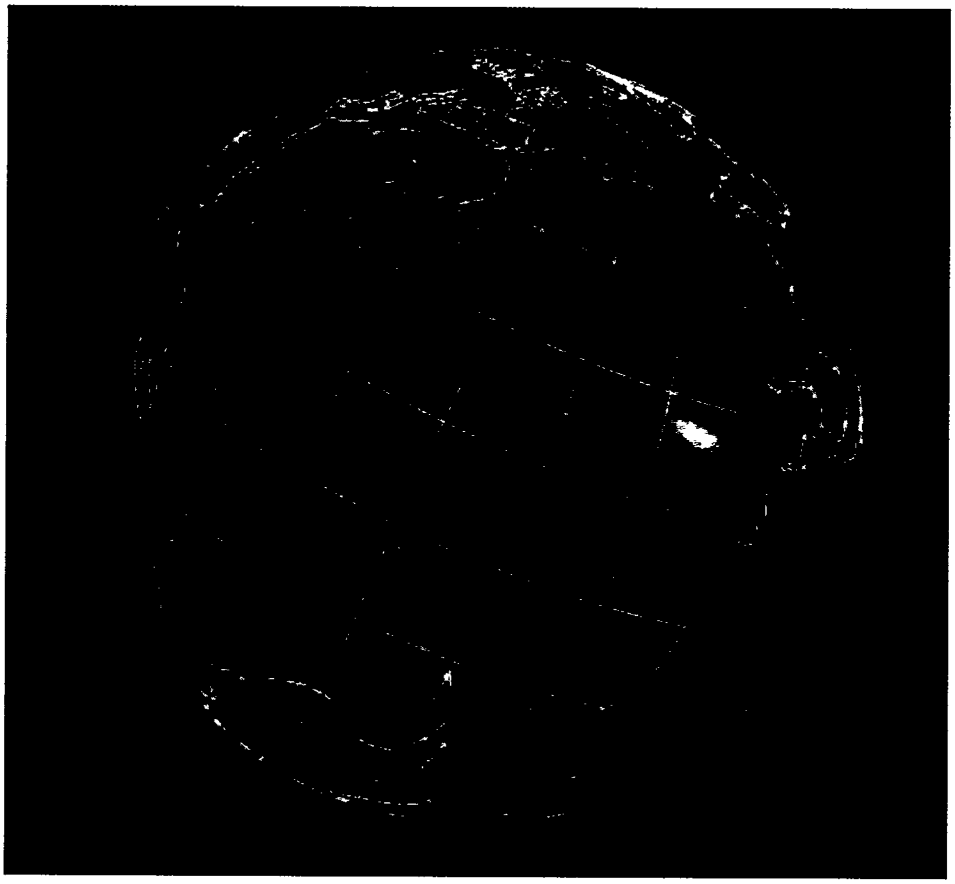

FIGURE 2.10 Three dimensional tomographic image of the Earth's mantle showing isosurfaces of fast (blue) and slow (red) seismic velocities. (Figure courtesy of Wei-jia Su and A. Dziewonski.)

CONCLUSIONS

The scientific, technical, and economic importance the issues discussed here presents a clear agenda for future research in computational seismology. In this way these problems will drive advances in high-performance computing in the field of seismology. There is a broad community that will benefit from this work, including the petroleum industry, research geophysicists, engineers concerned with seismic hazard mitigation, and governments charged with enforcing a comprehensive test ban treaty. These advances may also lead to new applications for seismological research. The recent application of high-resolution seismic imaging of the shallow subsurface for the environmental remediation industry is an example of this activity.

Solutions to these problems will require advances in hardware, software, and scientific understanding, as discussed below.

Hardware

Improvements in processor speeds, particularly for workstations, will continue to be of great value. Historically, increases in computational speed have allowed greater detail in model simulations. Applied to workstations, future advances in processor speeds will be necessary for enhanced graphics and visualization capabilities. There is a need for computational facilities that allow the use of massive data sets (10^{12} to 10^{15} bytes) in model simulations; in communications across global, national, regional, and local networks; and in the development of manageable data archives. These resources would be of immediate benefit to seismic imaging applications. A key component of this effort will involve the development of high-bandwidth networks and I/O capabilities between parallel processors. A related issue involves the need for improved hardware to support the automated collection, dissemination, and archiving of primary seismic data. This information provides the empirical foundation for a broad range of studies, and it is clear that the health and strength of seismology are strongly dependent on these efforts.

Software

There is a great need for algorithms to make full use of the capabilities of massively parallel computers. The demand extends to the most basic level of subroutine programming (e.g., matrix multiplication, Fourier transforms). These issues are key to attaining the highest speeds on the fastest supercomputers. There is also a need to develop efficient solutions to seismological problems (e.g., approximations to full finite-difference solutions). Historically, there is abundant evidence that the impacts of such software advances can equal (or exceed) the impact of hardware improvements.

Scientific Understanding

To make full use of hardware and software advances, there is a need for increased understanding of complex geologic systems. Because of the vast range of scales in seismological problems, simulations are needed to elucidate the interactions among complex processes that give rise to nonlinear, large-scale behavior. Such work has direct applications to developing models of tectonic stressing and earthquake failure. It is also needed to develop an understanding of scattering phenomena throughout the Earth that give rise to much of the complexity in measured seismograms (from natural and artificial sources).

3
Recommendations

Based on presentations at its workshop and subsequent deliberations, the Committee on Seismology makes the following observations and recommendations.

First, it is apparent that the ongoing development of high-performance computing technology will provide continued improvements in workstation and supercomputer performance for seismological applications. The committee recognizes that solutions to a number of problems in seismology are limited by current computer capabilities. Consequently, there will be great rewards from a strategy that strives to exploit the state of the art in high-performance computing resources. Presentations at the committee's workshop described great advances in the speed of network communications, the capabilities of distributed computing, the performance of massively parallel computers, and the technology for archiving extremely large data sets. These presentations indicated that developments in high-performance computing will continue on two fronts: in the development of the hardware and software for large-scale parallel computations and in the development of advanced workstation technology for visualization and laboratory-scale modeling and analysis.

Second, there is considerable uncertainty regarding the shape and direction of these technological developments. In part this reflects the diverse range of applications currently driving the market for high-performance computing (e.g., scientific computing, multimedia, financial data processing). The committee believes that this has important implications because high-performance computing in seismology relies on a broad range of computer capabilities, including fast processor speeds, advanced graphics, large amounts of memory, and high-bandwidth communications. Also, details of software design are often closely coupled to hardware, especially in developing applications for massively parallel computers.

Finally, the workshop presentations clearly described a broad range of seismological problems that will require a focused application of high-performance computing technology. Many of these issues have important scientific, economic, and social implications such as modeling of earthquake processes and imaging of geologic structures. For most of these problems, high-performance computing is required because of the great complexity of seismic phenomena and the large data sets that are needed to constrain geologic models.

The committee believes that these scientific problems articulate an unambiguous agenda for computational seismology for the next 10 to 15 years. Examples of long-range goals for high-performance computing in seismology include:

- development of three-dimensional simulations of seismic wave propagation and ground shaking from complex fault rupture events over regional scales (10^6 km^2) and up to frequencies of 10 to 20 Hz,

- widespread improvements in seismic wave propagation modeling and tomographic inversion techniques to allow accurate, high-resolution imaging of the Earth's interior over a wide range of scales (less than 10 km for global-scale models, less than 10 m for crustal-scale models); and

- construction of communication networks, data storage facilities, and seismic instrumentation to support worldwide access to archived and real-time global seismic data obtained from large numbers of recording sites (10^4 and greater).

Meeting these scientific challenges will require a focused effort to advance the breadth and depth of seismological applications in high-performance computing. To this end the committee makes the following recommendations.

1. Focused efforts to develop validated documented software for seismological computations should be supported, with special emphasis on scalable algorithms for parallel processors.

The committee observes that there has been too little effort to validate or benchmark the software for computational seismology. In many cases researchers start from scratch when developing software for a wide range of problems. Several of the workshop presentations revealed a large degree of "parallel" efforts expended by different research groups on particular computational problems (e.g., wave propagation in sedimentary basins). While recognizing the important educational benefits of software development (see below) and of exploring new numerical algorithms, the committee believes that a strong effort to develop validated algorithms for computational seismology could facilitate the growth of a new generation of software for high-performance computing applications. Such an effort would allow researchers to focus on the more advanced problems in developing software for parallel computers. Activities in this area should include:

- efforts to develop validated routines for a wide range of problems in computational seismology,
- benchmarking activities to identify algorithms that can exploit the benefits of parallel computer architectures,
- maintenance of widely accessible computer archives for software and documentation, and
- dissemination efforts to publicize and document the above efforts.

These activities could be coordinated through periodic workshops on computational seismology. Ideally, these meetings would be held at supercomputer centers or national laboratories where there is abundant expertise and resources for high-performance computing. The committee notes that similar workshops, held on an annual basis, have significantly enhanced computational activities in the field of geodynamics. The goal of these activities should be to support an infrastructure for developing software and expertise for high-performance computing in seismology. The committee observes that much of this work currently relies on the initiative of individual investigators and corporations (i.e., to develop software for particular research questions).

It is critical that the above effort focus on the development of scalable parallel algorithms. Other National Research Council reports, together with the presentations made at the committee's workshop, indicate that this is the trend of future software developments in high-performance computing. Taking advantage of parallel computer architectures will require a new generation of software and algorithms for scientific modeling and analysis. Currently, the technology for parallel computing is in flux, yet it is clear that scalable algorithms will be required to avoid massive reprogramming of software with each incremental change in computer technology. To this end the committee observes that much of the software development for scalable algorithms is currently feasible outside supercomputer centers using networks of linked workstations and standardized software for parallel computing (e.g., MPI[4]). The committee believes that a focused effort to develop scalable algorithms would have important benefits. It would allow extensive benchmarking and testing of new software using widely available computational resources, and it would increase the exposure of a wide range of seismologists to high-performance computing technology with relatively small investments in new computing.

---

[4] MPI (Message Passing Interface) is a freely available software protocol that permits a heterogeneous collection of computers, hooked together by a network, to be used as a single large parallel computer. With MPI (or similar software) large computational problems can be solved by using the aggregate power and memory of many computers.

2. The education of seismologists in high-performance computing technologies and methodologies should be improved.

The workshop presentations demonstrated that accelerated breakthroughs in computational seismology will require widespread understanding of high-performance computing software and hardware. The capabilities of seismologists in this area could be improved through broad educational efforts for researchers at all levels. The computer skills of undergraduate and graduate students in seismology should be strengthened through course work in computational science combined with broad exposure to high-performance computing technology. For postgraduate scientists there should be funding for participation in the ongoing seminars and outreach programs at the National Science Foundation's supercomputer centers. In this area, scientific societies could play a large role by organizing workshops and special sessions at national meetings to focus on the resources for computer education. The goal of these efforts should be to strengthen the role of computer education in developing the research capabilities of seismologists.

3. Collaborations between seismologists and computational scientists and engineers should be increased.

Several of the workshop presentations emphasized the similarity between seismological problems and a wide range of active research issues in computational science (e.g., I/O for large data sets, visualization). At the same time, however, there was only limited evidence of cross-collaboration between seismologists and scientists involved in state-of-the-art high-performance computing. During its deliberations, the committee heard presentations indicating that computational seismology would benefit from broad interdisciplinary collaborations (e.g., with computational scientists and engineers), yet the opportunities and interest in this work are currently limited. Specifically, it is often a challenge to attract computational scientists to work on seismological problems. For this reason the committee believes that an effort to foster collaborations between seismologists and a broad range of computational scientists would be valuable. Such collaborations could be promoted through increased efforts to educate the scientific community regarding the nature and importance of seismological problems and increased attendance by seismologists at workshops and conferences devoted to high-performance computing technology. Again, scientific societies could play a large role by facilitating cross-disciplinary forums and workshops for earth scientists and computational scientists and engineers.

4. The infrastructure for archiving, disseminating, and processing large volumes of seismological data should be improved.

Workshop presentations demonstrated that the seismological community is entering a new era that will be characterized by (1) significant increases in the volumes of recorded seismic data, (2) explosive growth in the number and size of centralized data archives, and (3) "real-time" recording from global seismic networks (e.g., National Research Council, 1991). The workshop presentations also indicated that there have been great advances in the communication and archiving of seismological data. At present, much of this activity utilizes off-the-shelf components (e.g., phone lines, T1 internet connections, gigabyte hard disks) and rudimentary data archiving software. Full utilization of future data streams likely will require a significant upgrade of the infrastructure for data communication and storage. The committee observes that the seismological research community will benefit tremendously from recent improvements in data communication and archiving facilities to support the monitoring of a Comprehensive Test Ban Treaty (see National Research Council, 1995b). Based on this experience, the committee believes there should be a sustained effort to exploit the capabilities of new technology for seismological data communication and archiving. Recent developments in real-time seismic monitoring for earthquake hazards and nuclear test detection represent important opportunities in this area. Also, there would be substantial benefits from developing widely accessible computer archives from the records of the thousands of continuously recording seismic stations throughout the world. In this last area, international scientific unions could play a key role in facilitating the development of international data archives.

References

Abriel, W. 1994. Presentation at National Research Council workshop, "High-Performance Computing in Seismology," San Diego Supercomputer Center, October 2-4.

Ben-Zion, Y., and J. R. Rice. 1993. Earthquake Failure Sequences Along a Cellular Fault Zone in a Three-Dimensional Elastic Solid Containing Asperity and Non-Asperity Regions. *Journal of Geophysical Research, 98,* 14109-14131.

Ben-Zion, Y., and J. R. Rice. 1995. Slip Patterns and Earthquake Populations Along Different Classes of Faults in Elastic Solids. *Journal of Geophysical Research, 100,* 12959-12983.

Bloomquist, M. G. 1993. Presentation before the U.S. House of Representatives Committee on Science, Space, and Technology, Subcommittee on Science, October 26.

Dolan, J. F., K. Sieh, T. K. Rockwell, R. S. Yeats, J. Shaw, J. Suppe, G. Huftile, and E. Gath. 1995. Prospects for Larger or More Frequent Earthquakes in the Los Angeles Metropolitan Region, California. *Science, 267,* 199-205.

French, W. S. 1992. Implications of parallel computation in seismic data processing. *Geophysics: The Leading Edge of Exploration,* June, pp. 22-25.

Harris, R., and S. Day. 1993. Dynamics of Fault Interaction: Parallel Strike-Slip Faults. *Journal of Geophysical Research, 98,* 4461-4472.

Holliger, K., and A. Levander. 1994. Lower Crustal Reflectivity Modeled by Rheological Controls on Mafic Inclusions. *Geology, 22,* 367-370.

Jordan, T. H. 1994. Presentation at National Research Council workshop, "High-Performance Computing in Seismology," San Diego Supercomputer Center, October 2-4.

Jordan, T. H. 1995. Presentation before the U.S. House of Representatives Committee on Science, Subcommittee on Basic Research, October 24.

Lafond, C. F., and A. Levander. 1995. Migration of Wide-Aperture Onshore-Offshore Seismic Data, Central California: Seismic Images of Late Stage Subduction. *Journal of Geophysical Research, 100,* 22231-22242.

Lax, P. (ed.) 1982. *Report of the Panel on Large Scale Computing in Science and Engineering.* U.S. Department of Defense, in cooperation with the National Science Foundation and the U.S. Department of Energy, Washington, D.C.

McLaughlin, K. L., and S. M. Day. 1994. 3D Elastic Finite-Difference Seismic-Wave Simulations. *Computers in Physics, 8,* 656-663.

National Coordination Office for High Performance Computing and Communications, Office of Science and Technology Policy. 1994. *FY 1995 Implementation Plan.* National Coordination Office for High Performance Computing and Communications, Rockville, Md.

National Research Council. 1991. *Real-Time Earthquake Monitoring.* National Academy Press, Washington, D.C.

National Research Council. 1993. *Solid Earth Sciences and Society.* National Academy Press, Washington, D.C.

National Research Council. 1995a. *Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure.* National Academy Press, Washington, D.C.

National Research Council. 1995b. *Seismological Research Requirements for a Comprehensive Test-Ban Monitoring System.* National Academy Press, Washington, D.C.

National Science Foundation, Blue Ribbon Panel on High-Performance Computing. 1993. *From Desktop to Teraflop: Exploiting the U.S. Lead in High-Performance Computing.* National Science Foundation, Arlington, Va.

Olsen, K., R. J. Archuleta, and J. R. Materese. 1995. Three-Dimensional Simulation of a Magnitude 7.75 Earthquake on the San Andreas Fault. *Science, 270*, 1628-1632.

Perrin, G., J. R. Rice, and G. Zheng. 1995. Self-Healing Slip Pulse on a Frictional Surface. *Journal of Mechanics and Physics of Solids, 43*, 1461-1495.

Pfeiffer, W. 1994. Presentation at National Research Council workshop, "High-Performance Computing in Seismology," San Diego Supercomputer Center, October 2-4.

Rice, J. R. 1993. Spatio-Temporal Complexity of Slip on a Fault. *Journal of Geophysical Research, 98,* 9885-9907.

Su, W., R. L. Woodward, and A. M. Dziewonski. 1994. Degree 12 Model of Shear Velocity Heterogeneity in the Mantle. *Journal of Geophysical Research, 99,* 6945-6980.

U.S. General Accounting Office. 1994. *High Performance Computing and Communications: New Program Would Benefit from a More Focused Effort.* U.S. General Accounting Office, Washington. D.C.

Appendix A
Workshop Participants

W. L. Abriel
Chevron U.S.A.
935 Gravier Street - Room 1763
New Orleans, LA 70112

Kevin Anderson
Battelle Pacific Northwest Laboratory
Box 999, M/S K7-34
Richland, WA 99352

Thomas C. Bache
SAIC
10260 Campus Point Dr.
San Diego, CA 92121

Craig Beasley
Western Geophysical Company
Box 2469
Houston, TX 77252-2469

Larry A. Bergman
Jet Propulsion Laboratory
California Institute of Technology
M/S 300-329
4800 Oak Grove Drive
Pasadena, CA 91109

Jacobo Bielak
Dept of Civil & Environmental Engineering
Carnegie Mellon University
Pittsburgh, PA 15213-3890

Biondo Biondi
Geophysics Department
Mitchell Building
Stanford University
Stanford, CA 94305

Hans-Peter Bunge
Institute of Geophysics &
Planetary Physics
Los Alamos National Lab
Mail Stop C305
Los Alamos, NM 87545

Jack Caldwell
Geco-Prakla
1325 S. Dairy Ashford
Houston, TX 77077

Debbie Campbell
Sandia National Laboratories
P.O. Box 5800, MS 0977
Albuquerque, NM 87185

Ted Charrette
Massachusetts Institute of
Technology
MIT Building E34-554
42 Carleton Street
Cambridge, MA 01773

Jon Claerbout
Geophysics Dept.
Stanford University
Stanford, CA 94305-2215

F. A. Dahlen
Dept of Geological and
Geophysical Science
Princeton University
Guyot Hall
Princeton, NJ 08544-1003

Peter Davis
Institute of Geophysics and
Planetary Physics, 0225
University of California,
San Diego
9500 Gilman Drive
La Jolla, CA 92093-0225

Steven M. Day
Department of Geological
Sciences
San Diego State University
San Diego, CA 92182

John A. Dickinson
Exxon Production Research
Company
P.O. Box 2189
Houston, TX 77252-2189

Geoffrey Dorn
ARCO Exploration and
Production Technology
PRC—D2226
2300 Plano Parkway
Plano, TX 75075-8427

Adam M. Dziewonski
Harvard University
Department of Earth &
Planetary Sciences
20 Oxford St.
Cambridge, MA 02138

Paul Earle
Institute of Geophysics and
Planetary Physics
Scripps Institution of
Oceanography
University of California,
San Diego
La Jolla, CA 92093

Ray Ergas
Geophysics Center
Chevron Petroleum Technology
Co.
La Habra, CA 90633-0446

William Farrell
SAIC
10260 Campus Point Drive
San Diego, CA 92121

Stanley M. Flatté
Physics Department
University of California
Santa Cruz, CA 95064

David B. Fogel
Natural Selection, Inc.
1591 Calle de Cinco
La Jolla, CA 92037

J. Freeman Gilbert
Institute of Geophysics and
Planetary Physics
University of California,
San Diego
La Jolla, CA 92093-0225

Peter Goldstein
Lawrence Livermore National
Laboratory
L-205
P.O. Box 808
Livermore, CA 94550

Robert W. Graves
566 El Dorado
Woodward-Clyde Consultants
Pasadena, CA 91101

Jeff Hanson
Institute of Geophysics &
Planetary Physics, 0225
Scripps Institution of
Oceanography
University of California,
San Diego
La Jolla, CA 92093

Alistair Harding
Institute of Geophysics and
Planetary Physics (0225)
Scripps Institution of
Oceanography
University of California,
San Diego
La Jolla, CA 92093-0225

Ruth A. Harris
Geophysics Department
Mitchell Building, Room 361
Stanford University
Stanford, CA 94305

Lisa A. Heizer
CG-234
Department of Geological
Sciences
San Diego State University
San Diego, CA 92182

Clifford Jacobs
National Science Foundation
Division of Atmospheric
Sciences
4201 Wilson Blvd
Arlington, VA 22230

Thomas H. Jordan
Professor and Head
Department of Earth,
Atmospheric and
Planetary Sciences,
54-918
Massachusetts Institute of
Technology
Cambridge, MA 02139

James R. Kamm
Los Alamos National
Laboratory
EES-3, MS F659
P.O. Box 1663
Los Alamos, NM 87545

Randy Katz
ARPA
3701 N. Fairfax Drive
Arlington, VA 22203-1714

Valeri Korneev
Lawrence Berkeley Laboratory
50E, ESD
1 Cyclotron Road
Berkeley, CA 94720

Art Lerner-Lam
Lamont Doherty Earth
Observatory
Palisades, NY 10964

Alan Levander
Geology & Geophysics
MS 126
Rice University
Houston, TX 77005

Stewart A. Levin
Mobil Exploration and
Production Technical
Center
3000 Pegasus Park Drive
Dallas, TX 75247

William C. Luth
Engineering and Geosciences
Division
U.S. Department of Energy
Washington, DC 20585

Joseph Matarese
Massachusetts Institute of
Technology
MIT Bldg., E34-554
42 Carleton Street
Cambridge, MA 01773

Keith L. McLaughlin
S Cubed
PO Box 1620
La Jolla, CA 92038

J. Bernard Minster
Institute of Geophysics and
Planetary Physics (0225)
Scripps Institution of
Oceanography
University of California,
San Diego
La Jolla, CA 92093-0225

David Oglesby
Institute for Crustal Studies
University of California
Santa Barbara, CA 93106

Kim Olsen
Department of Geological
Sciences
University of California at
Santa Barbara
Santa Barbara, CA 93106-9630

John Orcutt
Director, Institute of
Geophysics and Planetary
Physics (0225)
Scripps Institution of
Oceanography
University of California,
San Diego
La Jolla, CA 92093-0225

Jeffrey Orrey
Department of Physics/TAGG
Campus Box 583
University of Colorado
Boulder, CO 80303

Richard Ottolini
Unocal Science and Technology
P.O. Box 68076
Anaheim, CA 92817-8076

Gagan B. Patnaik
President
Advanced Geocomputing
Technologies
4660 La Jolla Village Drive,
Suite 500
San Diego, CA 92122

Donald L. Paul
President
Chevron Petroleum Technology
Company
2811 Hayes Road
Houston, TX 44082

Wayne Pfeiffer
San Diego Supercomputer
Center
P.O. Box 85608
San Diego, CA 92186

Delaine Reiter
ASEC@PL/GPE
Hanscom AFB, MA 01731

James R. Rice
224 Pierce Hall
Division of Applied Sciences
Harvard University
Cambridge, MA 02138

Walter Ritchie
Western Geophysical
10001 Richmond
Houston, TX 77042-4299

Wilmer Rivers
3003 Farm Road
Alexandria, VA 22302

Salvador Rodriguez
CGG American Services, Inc.
2500 Wilcrest Suite 200
Houston, TX 77042-2797

Daniel Rosenblatt
7734 Esterel Drive
La Jolla, CA 92037

Dave Schneider
626 Engineering and Theory Center
Cornell University
Ithaca, NY 14853

Gerard T. Schuster
Geology & Geophysics
University of Utah
Salt Lake City, UT 84112

David Simpson, President
The IRIS Consortium
1616 N. Fort Myer Drive
Suite 1440
Arlington, VA 22209

Paul G. Somerville
Woodward Clyde Consultants
566 El Dorado Street
Suite 100
Pasadena, CA 91101

Ralph Stephen
Senior Scientist
Department of Geology & Geophysics
Woods Hole Oceanographic Institute
Woods Hole, MA 02543

Kris Stewart
San Diego State University
Computational Science
Curriculum Coordinator
San Diego Supercomputer Center
810 Ave De San Clemente
Encinitas, CA 92924

Paul Stolorz
MS 198-219
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109

Brian W. Stump
Los Alamos National Laboratory
P.O. Box 1663, EES-3
MS C335
Los Alamos, NM 87545

Paul Tackley
Earth and Space Sciences Department
University of California
Los Angeles, CA 90024

Fumiko Tajima
Institute for Geophysics
The University of Texas at Austin
8701 N. Mopac Boulevard
Austin, TX 78759-8397

Anne Trehu
Oceanic and Atmospheric Sciences
Oregon State University
Ocean Administration Building
104
Corvallis, OR 97331

Alexie G. Tumarkin
Institute for Crustal Studies
University of California
Santa Barbara, CA 93106

Junho Um
Institute of Geophysics & Planetary Physics
Scripps Institution of Oceanography
University of California, San Diego
La Jolla, CA 92093-0225

David Wald
U.S. Geological Survey
525 South Wilson Avenue
Pasadena, CA 91106

Richard Watson
Lawrence Livermore National Laboratory
P.O. Box 808—L560
Livermore, CA 94550

Nadya Williams
Institute of Geophysics & Planetary Physics
Scripps Institution of Oceanography
University of California, San Diego
9500 Gilman Drive, Dept. 0225
La Jolla, CA 92093-0225

Tony Williams
PO Box 2113
La Habra, CA 90632

Mara Yale
Institute of Geophysics and
Planetary Physics
Scripps Institution of
Oceanography
University of California,
San Diego
La Jolla, CA 92093-0225

Guang Yu
Department of Geological
Sciences
San Diego State University
San Diego, CA 92182

Appendix B
Workshop Presentations

Steven Day	*Welcome*
Thomas Jordan	*Workshop Overview*
Clifford Jacobs	*Overview of HPCC-National Science Foundation's Role*
Randy Katz	*New Directions for HPCC*
Donald Paul	*An Overview of the Advanced Computational Technology Initiative*
Wayne Pfeiffer	*Scalable Parallel Computing*
Richard Watson	*Developments in Mass Storage Technology*
William Farrell	*Database Management*
Peter Bunge	*Large Scale Mantle Dynamics Modeling on Parallel Virtual Machines*
Larry Bergman	*Interactive Visualization Using CASA Meta-Supercomputer*
David Simpson	*The IRIS Data Management System*
Jon Claerbout	*15 Years of Super Computing and Other Problems*
Walter Ritchie	*Computing Environment Demands for Very Large 3D Surveys*
Geoffery Dorn	*Visualization for Petroleum Industry Applications*

Brian Stump	*Data Management and Visualization Tools in Comprehensive Test-Ban Treaty Monitoring*
Thomas Bache	*Fault-Tolerant Information-Driven, Distributed Multi-Processing on Networks of UNIX Workstations*
Stanley Flatte	*Numerical Simulation of 1000-km Acoustic Pulse Propagation in the Ocean Waveguide, Including Comparison with Experiment*
Jacabo Bielak	*3D Finite Element Modeling of Earthquake Ground Motion*
Kim Olsen	*Simulation of 3D Elastic Wave Propagation*
Alan Levander	*Simulation of Seismic Response for Realistic Crustal Models*
Ruth Harris	*Computational Challenges of Simulating Earthquakes on Complex (Geologically Acceptable) Faults*
James Rice	*Problems in Crack and Fault Dynamics*
Richard Ottolini	*Modeling of 3D Seismic Reflection Surveys*
Paul Tackley	*3D Simulations of Mantle Convection: Spectral and Multigrid Methods*
Edmond Charrette	*Seismic Imaging and Inversion on a Parallel Computer*
Stewart Levin	*3D Prestack Depth Migration at Mobil*
Adam Dziewonski	*Global Seismic Tomographic Models*
William Abriel	*Reservoir Imaging in the Petroleum Industry*

John Dickinson	*3D Prestack Depth Migration*
David Fogel	*Evolutionary Programming for Computationally Intensive Problems in Geophysics*
David Wald	*Inversion for Earthquake Rupture Dynamics*
Gerald Schuster	*Elastic Traveltime and Waveform Inversion of Crosshole Seismic Data*