DOCUMENT RESUME

ED 340 361   IR 015 356

TITLE   High-Performance Computing: Industry Uses of
                     Supercomputers and High-Speed Networks. Report to
                     Congressional Requesters.
INSTITUTION   General Accounting Office, Washington, DC.
                     Information Management and Technology Div.
REPORT NO   GAO/IMTEC-91-58
PUB DATE   Jul 91
NOTE   32p.; For related documents, see ED 327 219, ED 328
                     244, ED 329 226, and IR 015 355.
AVAILABLE FROM   U.S. General Accounting Office, P.O. Box 6015,
                     Gaithersburg, MD 20877 (first five copies: free,
                     additional copies: $2.00 each).
PUB TYPE   Information Analyses (070) -- Reports -
                     Research/Technical (143)

EDRS PRICE   MF01/PC02 Plus Postage.
DESCRIPTORS   Aerospace Industry; *Business Communication;
                     *Computer Networks; Computer Simulation; Computer
                     Software; Cost Effectiveness; *Industry;
                     Postsecondary Education; *Problem Solving; Surveys;
                     *Technological Advancement; Use Studies
IDENTIFIERS   *High Performance Computing; *Supercomputers

ABSTRACT
            This report was prepared in response to a request for
information on supercomputers and high-speed networks from the Senate
Committee on Commerce, Science, and Transportation, and the House
Committee on Science, Space, and Technology. The following
information was requested: (1) examples of how various industries are
using supercomputers to improve products, reduce costs, save time,
and provide other benefits; (2) barriers preventing the increased use
of supercomputers; and (3) examples of how certain industries are
using and benefitting from high-speed networks. The review of
supercomputers examined five industries--oil, aerospace, automobile,
chemical, and pharmaceutical. These industries are known for using
supercomputers to solve complex problems for which solutions might
otherwise be unattainable. Appendices II through V provide detailed
accounts--drawn from 24 companies contacted within these
industries--of how supercomputers are used to improve products and
provide other benefits. Information was also obtained from 10
companies in the oil, automobile, and computer industries on their
use of high-speed computer networks. Appendices I, VI, and VII,
respectively, provide additional information on the objectives,
scope, and methodology of the review, and identify the companies
contacted to assess their use of supercomputers and high-speed
networks. (DB)

July 1991

HIGH-PERFORMANCE COMPUTING

Industry Uses of Supercomputers and High-Speed Networks

IR015356

ERIC

GAO

United States
General Accounting Office
Washington, D.C. 20548

Information Management and
Technology Division

B-244488

July 30, 1991

The Honorable Ernest F. Hollings
Chairman, Senate Committee on Commerce,
Science, and Transportation

The Honorable Al Gore
Chairman, Subcommittee on Science,
Technology, and Space
Senate Committee on Commerce, Science,
and Transportation

The Honorable George E. Brown, Jr.
Chairman, House Committee on Science,
Space, and Technology

The Honorable Robert S. Walker
Ranking Minority Member
House Committee on Science, Space,
and Technology

The Honorable Tim Valentine
Chairman, Subcommittee on Technology
and Competitiveness
House Committee on Science, Space,
and Technology

The Honorable Tom Lewis
Ranking Minority Member
Subcommittee on Technology
and Competitiveness
House Committee on Science, Space,
and Technology

This report responds to your October 2, 1990, and March 11, 1991,
requests for information on supercomputers and high-speed networks.
You specifically asked that we

- provide examples of how various industries are using supercomputers
  to improve products, reduce costs, save time, and provide other benefits;
- identify barriers preventing the increased use of supercomputers; and
- provide examples of how certain industries are using and benefitting
  from high-speed networks.

As agreed with the Senate Committee on Commerce, Science, and Transportation, and Subcommittee on Science, Technology, and Space, our review of supercomputers examined five industries—oil, aerospace, automobile, and chemical and pharmaceutical. These industries are known for using supercomputers to solve complex problems for which solutions might otherwise be unattainable. Appendixes II through V provide detailed accounts—drawn from 24 companies contacted within these industries—of how supercomputers are used to improve products and provide other benefits. We also obtained information from 10 companies in the oil, automobile, and computer industries concerning how they are using and benefitting from high-speed computer networks. We did not verify the accuracy of the examples of benefits provided by the various companies. Appendixes I, VI, and VII, respectively, provide additional information on the objectives, scope, and methodology of our review, and identify the companies examined to assess uses of supercomputers and high-speed networks.

Results in Brief

Supercomputers contribute significantly to the oil, automobile, aerospace, and chemical and pharmaceutical industries' ability to solve complex problems. They enable companies within these industries to design new and better products in less time, and to simulate product tests that would have been impossible without spending months developing and experimenting with expensive product models. Some companies have attributed significant cost savings to the use of supercomputers. For example, although exact figures were not always available, representatives of some automobile and aerospace companies estimated that millions of dollars have been saved on specific models or vehicle parts because of reduced manufacturing or testing costs. In addition, one oil company representative estimated that over the last 10 years, supercomputer use has resulted in increased production of oil worth between $5 billion and $10 billion from two of the largest U.S. oil fields.

Despite widespread use of supercomputers for certain applications, representatives of these companies told us that several key barriers currently hinder their greater use. These barriers include (1) the high cost of supercomputers, (2) a lack of application software, (3) the cultural resistance to the shift from physical experiments to an increased reliance on computational experiments, and (4) a lack of supercomputing education and training.

4

High-speed networks contribute to improved productivity by enabling industries to more efficiently share information and resources and collaborate on product development over distances. Companies within the oil, automobile, and computer industries, for example, rely on high-speed networks to transfer large graphics and data files, and access computers worldwide. In many cases these companies must use high-speed networks (as opposed to those operating at lower speeds) because (1) the applications in use require high transmission speeds, for example to provide instant images for interactive videoconferencing, (2) high volumes of traffic in one or more applications are being transmitted, or (3) fast response is needed for such applications as data base queries. Several companies reported that they would not be able to develop products in a timely manner without high-speed networks.

Background

A supercomputer, by its most basic definition, is the most powerful computer available at a given time. Current supercomputers, costing fro..t about $1 million to $30 million, are capable of performing billions of calculations each second. Computations requiring hours or days on conventional computers may be accomplished in a few minutes or seconds on a superco.nputer.

Although the term supercomputer does not refer to a particular design or type of computer, supercomputers generally use vector or parallel processing. With vector processing, a supercomputer lines up billions of calculations and then uses one or several large processors to perform these calculations. In parallel processing, many smaller processors work on multiple parts of a program concurrently. The trend in supercomputer design is to add more processors to achieve greater performance. Massively parallel supercomputers consisting of between 1,000 and 64,000 processors now exist.

The unique computational power of supercomputers makes it possible to solve critical scientific and engineering problems that cannot be dealt with satisfactorily by theoretical, analytical, or experimental means. Scientists and engineers in many fields—including aerospace, petroleum exploration, automobile design and testing, chemistry, materials science, and electronics—emphasize the value of supercomputers in solving complex problems. Much of this work involves the use of workstations for scientific visualization—a technique allowing researchers to convert masses of data into three-dimensional images of objects or systems under study. These images enable researchers to comprehend more

readily what data reveal and facilitate the understanding of problems by different types of scientists and engineers.

While supercomputers are still relatively limited in use, the number of supercomputers has risen dramatically in the last decade. In the early 1980s, most of the 20 to 30 supercomputers in existence were operated by government agencies for such purposes as weapons research and weather modeling. Today, about 280 supercomputers[1] are in use world-wide. The government (including defense-related industry) remains the largest user, although private industry has been the fastest growing user segment for the past few years, and is projected to remain so.

A high-speed network is generally defined as a network operating at speeds of T1—1.544 million bits per second—or higher. Prior to 1977, high-speed networks were employed exclusively by the telephone companies. By the early 1980s, however, these services had become widely available to commercial customers.

Today, thousands of high-speed networks exist, fueled by demands for a variety of applications, such as electronic mail, data file transfer, and distributed data base access. Networks operating at T1-speeds are common and provide sufficient capability to meet most application needs. However, there is a growing demand for higher-speed networks, such as those operating at T3-speeds (45 million bits per second) or greater, to transmit multiple low-speed applications to many users at the same time. In addition, many industries now look to such networks as a means of transmitting more advanced applications that result from the use of supercomputers and other sophisticated technologies. The growth of T1 and T3 lines is expected to be great, according to Northern Business Information/Datapro, a research company and industry analyst, which projected that revenues for T1 and T3 will increase three-fold between 1990 and 1994.

Industries Benefit From Supercomputing

Supercomputers provide the five selected industries with the ability to develop new and better products more quickly. Although most companies within these industries could not provide precise figures to quantify the extent of gains realized, nearly all believed that supercomputers have enabled them to perform previously impossible tasks, or achieve

---

[1]This figure includes only high-end supercomputers such as those man. factured by Cray Research, Inc. Including International Business Machines (IBM) mainframes with vector facilities would about double this number.

6

significant cost reductions or time savings. Moreover, industry representatives believed that greater benefits would be realized in the future, as these companies move toward using more powerful supercomputers with thousands of processors. Of the 21 companies that commented on the issue, 19 said that they will be using massively parallel supercomputers to a greater extent in the future.[2] Details on each industry's use of supercomputers are in appendixes II through V.

The Oil Industry

As an early user of supercomputers, the oil industry has realized substantial benefits from supercomputer applications. By using two key applications for processing seismic data[3] and simulating reservoirs, oil companies have improved their ability to determine the location of reservoirs and to maximize recovery of oil and gas from those reservoirs. This ability has become increasingly important because of the low probability of discovering large oil fields in the continental U.S. New oil fields are often small and located in harsh environments, making exploration and production difficult. Several industry representatives estimated that the use of supercomputers reduces the number of dry wells drilled (at a cost of $.5 million to over $50 million per well) by about 10 percent. In addition, an Atlantic Richfield Company (ARCO) representative estimated that supercomputer use has led to increased oil production worth billions of dollars at two large fields.

The Aerospace Industry

Engineers and researchers in the aerospace industry have used supercomputers since the early 1980s to design, develop, and test aerospace vehicles and related components. Supercomputers, for example, have enabled engineers to analyze aircraft structural composition for design flaws and to simulate their performance in wind tunnels. This ability is important because wind tunnels are expensive to build and maintain, and cannot reliably detect certain airflow phenomena. Simulation permits a reduction in physical model testing, and substantial savings in time and money. As a major user of supercomputers, McDonnell Douglas estimates that supercomputer simulations saved about a year in the design and testing of its new C-17 military aircraft.

---

[2]Three out of 24 representatives did not comment on the issue for proprietary reasons.

[3]Seismic data reveal characteristics about the earth and are gathered using sound recording devices to measure the speed that vibrations travel through the earth.

The Automobile Industry

Since 1984, automobile manufacturers have increasingly relied on supercomputers to design vehicles that are safer, lighter, more economical, and better built. By the late 1980s, the world's 12 largest automobile companies had acquired supercomputers. A primary supercomputer application—crash analysis—is used to simulate how vehicle structures collapse on impact and how fast passengers move forward. These simulations provide more precise engineering information than was possible from physically crashing pre-prototype vehicles. They also reduce the number of vehicles required for these tests by about 20 to 30 percent. Consequently, companies have been able to save millions of dollars annually. According to General Motors Corporation representatives, for example, supercomputers enabled the company to crash 100 fewer vehicles when developing some of its 1992 models, than it did in 1987. Each test vehicle costs from $50,000 to $750,000, depending on whether a production vehicle or prototype is used.

The Chemical and Pharmaceutical Industries

Supercomputers also play a growing role in the chemical and pharmaceutical industries, although their use is still in its infancy. From computer-assisted molecular design to synthetic materials research, these companies increasingly rely on supercomputers to study critical design parameters and more quickly and accurately interpret and refine experimental results. Industry representatives told us that the use of supercomputers will result in new discoveries that otherwise may not have been possible. Du Pont, for example, is developing replacements for chlorofluorocarbons, compounds used as coolants for air conditioners, that are thought to contribute to the depletion of ozone in the atmosphere. In designing a new process to produce substitute compounds, Du Pont is using a supercomputer to make certain calculations needed for this process. These calculations, on a supercomputer, require a few days at a cost of between $2,000 to $5,000. Previously, however, such tests cost about $50,000 and required up to 3 months to conduct.

Barriers Impede Greater Use of Supercomputers

Although supercomputers have yielded highly visible contributions in the selected industries, representatives told us that many aspects of supercomputer use remain untapped, because of the following significant barriers.

High cost: Currently, supercomputers cost between $1 million and $30 million, not including the cost of software development, maintenance, or trained staff.

Lack of software: While the evolution of software for vector supercomputers has accelerated over the past decade, little reliable software has been developed for parallel supercomputers. This is in part due to the lack of software tools for developing new parallel software and converting vector software so that it can be used on massively parallel supercompt:ters.

Cultural resistance: Many companies or industries, particularly the chemical and pharmaceutical industries, rely more heavily on physical experimentation than necessary, according to representatives. Many scientists and managers see the use of computational science as a dramatic break with past practice, and such a major shift in research methodology is difficult to accept.

Lack of supercomputer training and education: Before 1985, university students and professors performed little of their research on supercomputers. Thus, for many years industry hired students from universities who did not bring supercomputing skills and experience to their jobs. According to Du Pont and Eli Lilly representatives, universities are still not providing a sufficient number of students skilled in the use of supercomputers. A Ford Motor Company representative also noted that there is a scarcity of trained staff in computational fluid dynamics, an important application to the automobile industry. Currently, formal supercomputer education is primarily limited to the National Science Foundation (NSF) university supercomputer centers.

Industry Uses of High-Speed Networks

Like supercomputers, high-speed networks are making valuable contributions to many industries. Companies in the oil, automobile, and computer industries, for example, increasingly rely on high-speed networks to share resources and provide various types of person-to-person communications. Many oil company representatives, in particular, reported network traffic increases, ranging from 10 to more than 100-fold over the past 5 years. Many companies thought that significant benefits—including monetary savings, reduced time-to-market, and improved product quality—have resulted from their use of high-speed networks.

The companies we contacted primarily use high-speed networks operating at T1-speeds (1.544 million bits per second). A significantly smaller, although growing, number of companies also use higher-speed T3 networks of 45 million bits per second. These networks generally consist of private lines, leased exclusively for each company's use,

although many are connected to outside commercial and private networks.

Companies we contacted use high-speed networks for a variety of reasons. In some cases, these networks are used for individual applications that require high transmission speeds, such as interactive videoconferencing. Most companies also used these networks as a more cost-effective way of transmitting large volumes of aggregated traffic from lower-speed applications. These applications include voice communication, remote computer access, and electronic mail.

Landmark Graphics Corporation, a company that develops seismic data processing software for oil exploration, for example, uses an extensive T1 network to support a variety of applications. This network supports up to four voice lines (at 64,000 bits per second each), while providing electronic mail access to hundreds of network users. This network also allows users across the country to work simultaneously on the development of the same software by accessing and sharing files via high performance workstations, and to routinely transfer voluminous files to backup the file system. A Landmark representative said that network use has provided more coordinated and consistent control of product development among the company's different offices, and ultimately, a shortened product development life-cycle.

The Amoco Corporation uses a high-speed network to transmit very large (100 million bit to 1 billion bit) files between its foreign and domestic sites. The files contain large volumes of data such as images of sections of the earth, which measure about 400 square miles wide by 3 miles deep. These data are critical to improving Amoco's ability to locate oil reservoirs. Because of the volume of the data, Amoco representatives said it would be impossible to transmit these files to each work site without high-speed networks. If they did not have the networks, the data would have to be duplicated at each site, resulting in higher costs. Moreover, according to an Amoco representative, access to the supercomputer via the high-speed network enabled them to make a major oil discovery—the details of which are proprietary.

Within the automobile industry, a General Motors (GM) Corporation representative reported that high-speed networks primarily benefit them by reducing costs and increasing productivity. For example, the network permits resource sharing, reducing duplicate hardware and software purchases. One group reported saving $90,000 by using university software over the network, rather than purchasing it. Another group

reported that it did not have to buy a parallel supercomputer because it accessed one at a university via the network. In addition, a corporate networking group projected a $2.3 million cost avoidance for 1991 because the use of a high-speed network enabled them to make large data and graphics files more readily available to remote sites.

We discussed the information in this report with industry representatives and experts, and incorporated their comments where appropriate. Our work was performed between October 1990 and May 1991.

As agreed with your office, unless you publicly announce the contents of this report earlier, we plan no further distribution until 30 days from the date of this letter. We will then send copies to interested congressional committees and others upon request. Please contact me at (202) 275-3195 if you have any questions concerning this report. The major contributors to this report are listed in appendix VIII.

Jack L. Brock, Jr.
Director
Government Information
    and Financial Management

Contents

12

Appendix VIII
Major Contributors to
This Report

29

Abbreviations

ARCO	Atlantic Richfield Company
GAO	General Accounting Office
GM	General Motors Corporation
IBM	International Business Machines Corporation
IMTEC	Information Management and Technology Division
NASA	National Aeronautics and Space Administration
NSF	National Science Foundation

Objectives, Scope, and Methodology

At the request of the Senate Subcommittee on Science, Technology, and Space, the Senate Committee on Commerce, Science, and Transportation; the House Subcommittee on Technology and Competitiveness; and the House Committee on Science, Space, and Technology; we reviewed various industries' use of supercomputers and high-speed networks. The purpose of our review was to (1) illustrate how the automobile, aerospace, petroleum, and chemical and pharmaceutical industries are using supercomputers to improve products, reduce costs, save time, or provide other benefits; (2) describe barriers that inhibit the increased use of supercomputers; and (3) provide examples of how certain industries use high-speed networks and their associated benefits.

To illustrate how industries are using and benefitting from supercomputers and identify barriers to their increased use, we interviewed managers, scientists, and engineers from the 24 companies listed in appendix VI. We selected these companies on the basis of recommendations from various experts knowledgeable about industrial supercomputer use. Most of the companies we selected are Fortune 500 companies, largely because of the resources required to purchase, maintain, and use supercomputers.

We also interviewed and obtained background information on supercomputers and on industry applications and future trends from industry analysts and consultants, hardware vendors, and government officials. The industry analysts and consultants included those from Research Consortium, Inc., Dataquest, The Superperformance Computing Service, Gartner Group, Inc., and the Institute for Supercomputing Research Recruit Co., Ltd. The hardware vendors included Cray Research, Inc., International Business Machines Corporation, Thinking Machines Corporation, and Silicon Graphics, Inc. The government officials included those from the Office of Science and Technology Policy, International Trade Administration, Department of Commerce, Lawrence Livermore National Laboratory, Department of Energy, and National Aeronautics and Space Administration Ames Research Center, National Science Foundation (NSF), and NSF supercomputer centers—San Diego Supercomputer Center, National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign, Cornell Theory Center, and Pittsburgh Supercomputing Center. We also interviewed and obtained documents from representatives of the Institute of Electrical and Electronics Engineers, Inc., and the American Petroleum Institute.

To assess how industries use high-speed computer networks, we collected information from companies in various industries and procured

15

the services of Texas Internet Consulting, a firm that designs and implements networks for large corporations. This firm subsequently interviewed the ten computer, automobile, and oil companies listed in appendix VII to determine how these companies use and benefit from high-speed networks. These companies had been selected based on the recommendation of experts as being frequent users of high-speed networks. Texas Internet Consulting also provided background information on high-speed networks and related applications.

We discussed the information in this report with scientists, engineers, and other experts from 14 oil, aerospace, automobile, chemical and pharmaceutical, and computer companies and have incorporated their views as appropriate. However, we did not verify the validity or accuracy of the examples of dollar savings and productivity improvements provided to us by the various companies. In some cases, we were unable to obtain such information because it was considered proprietary and could not be released. In other cases, company representatives said they had not performed the extensive analysis necessary to quantify such benefits.

Our review was conducted from October 1990 to May 1991 primarily in Washington, D.C., and other locations listed in appendix VI.

The Oil Industry

Many large oil companies worldwide—including American, British, French, and Middle Eastern companies—use supercomputers to better determine the location of oil and gas reservoirs, and to maximize the output from these reservoirs. The oil industry is among the early users of supercomputers, with ARCO being the first company to purchase a Cray supercomputer, in 1980, to model the largest oil reservoir in the U.S.—Prudhoe Bay, Alaska. Since that time, the oil industry has invested hundreds of millions of dollars in developing software for supercomputer applications. In addition, the industry uses off-the-shelf, consultant-developed and university-developed software.

The oil industry is now looking at massively parallel processing to achieve the next-order-of-magnitude improvement in speed. The industry has many large computational problems—such as modeling entire reservoirs in greater detail—that cannot be practically attempted on today's supercomputers. Many companies are experimenting with parallel supercomputers and are converting vector software to parallel software. According to an ARCO representative, such conversion is time-consuming and expensive because programs are extremely large. Despite industry movement toward parallel supercomputers, most representatives thought that vector supercomputers would continue to be used to a great extent. The oil industry uses two key supercomputer applications—seismic data processing and reservoir simulation—to aid in oil and gas exploration and production. Most representatives from the eight companies we contacted said that supercomputers have greatly helped them reduce costs. They also stated that supercomputers have greatly enabled them to (1) perform previously impossible tasks and (2) improve the quality and timeliness of their simulations.

Seismic Data Processing

Seismic data processing is used to produce images of subsurface geology through calculations involving large volumes of seismic data. Analysis of these images increases the probability of determining the location of oil reservoirs. This is important because the vast majority of holes drilled (ranging in cost from $.5 million to over $50 million) are dry and this is the most reliable way of determining the presence of oil or gas without drilling. While two representatives mentioned that seismic data processing had been critical in making significant oil and gas discoveries, the details on these finds are proprietary.

Although most representatives indicated difficulties in estimating the economic benefits derived from using supercomputers for seismic data processing, various industry representatives said that it reduces the

number of dry wells drilled by about 10 percent, saving hundreds of millions of dollars over the last 5 years. Chevron Corporation representatives stated that it reduces drilling of dry off-shore wells by about 50 percent. Drilling costs for the oil industry are significant, totaling over $200 billion over the past decade.

Using a supercomputer for seismic data processing also permits a faster payoff for millions of dollars in property, equipment, data collection, and other up-front investments. An ARCO representative estimated that it saves about a half million dollars per development well (a well located in a known oil field) because oil is located and recovered more quickly.

Representatives said it would be impractical to run their seismic data processing applications on less powerful computers. They said that supercomputers increased the speed of results from 3 to 100 times that of other computers, depending on the amount of data being processed. This permits them to run more iterations of a seismic image, which in turn improves the image, and the quality of their drilling decisions.

In addition, supercomputers enable these companies to perform more complex seismic data processing tasks than would be possible on other computers. Depth migration and modeling reservoirs located around salt domes are two such tasks that require massive computations. Depth migration is a process that removes some of the distortion of seismic images caused by layers of rocks acting as lenses. Companies are continually improving this method to provide more accurate images, particularly of areas where rocks are layered, such as in the foothills of the Rocky Mountains.

Another recently developed process that can only be run on supercomputers—producing accurate images through salt—is being used by several companies to improve their ability to discover oil in areas with salt domes. This process is important because salt domes are likely traps for oil. Yet, it is difficult to get a clear view of what lies under salt domes because salt absorbs sound waves and distorts seismic images. In 1987, Oryx Energy Company used a Cray X-MP to better understand the shape of salt domes. While they had previously thought the domes were cylindrical, supercomputer simulations revealed that they were smaller and mushroom-shaped. Understanding their shape enabled Oryx to change their drilling path and discover oil in a well in the Mississippi Salt Basin.

Reservoir Simulation

Reservoir simulation, an important tool used to increase the amount of oil and gas that can be extracted from a reservoir, allows engineers to "experiment" on a field by trying out various recovery methods and sizes and types of facilities (i.e., equipment used in extraction). By analyzing the alternatives, the most cost-beneficial development methods for a field can be identified before the actual production work is undertaken. The importance of reservoir simulation derives from the need for long-term optimization of recovery of the world's limited fuel resources.

Although representatives said that it is difficult to estimate the economic benefits derived from supercomputing in reservoir simulation, various representatives said that it permits increased oil yield, which one representative estimated to be worth $10 million annually to their company. One Amoco Production Company representative estimated that the increased yield due to reservoir simulation was about 25 percent for some wells. Simulation also reduces the risk of losing the oil in a reservoir during production—if the oil in a well is produced too fast, it may move to another area of the reservoir and be lost beyond recovery. Using supercomputers for reservoir simulation also reduces the amount of money spent on unnecessary recovery methods and facilities and equipment. These methods and facilities are expensive—recovery methods cost from about $1 to $25 per barrel, while surface facilities and equipment for a very large field can cost billions of dollars. Investment in the wrong recovery methods and facilities can break a company.

Supercomputers have enabled the industry to perform reservoir simulations of entire large fields, previously considered too impractical or costly. Full-field models of Prudhoe Bay and Kuparuk, Alaska—the largest and third largest on-shore fields in the U.S.—were used to improve recovery methods. An ARCO representative estimated, conservatively, that reservoir simulation used for these oil fields has resulted in about a 500 million barrel increase in production, worth between $5 billion and $10 billion.

The Aerospace Industry

Aerospace companies were among the earliest users of supercomputers, which are regarded as essential tools in the design, development, and testing of aircraft, missiles, and other aerospace vehicles and components. Furthermore, representatives generally consider supercomputing critical to competing for advanced military and other contracts. Of the six American companies we contacted, four began using vector supercomputers in the early 1980s, and by the late 1980s all had at least one supercomputer.

Most of the aerospace software applications were initially developed in federal government laboratories. For example, NASTRAN, one of the early software programs for analyzing the structure and shape of an object, was developed through research sponsored by the National Aeronautics and Space Administration (NASA) in the mid-1960s. Supercomputer software is commercially available for some aerospace applications, but companies also use software programs developed by government facilities, and develop or modify software for proprietary use.

Since 1987, five of the six companies have begun experimenting with massively parallel supercomputers to determine how to best use them and to develop application software. One of the companies has a massively parallel supercomputer dedicated to classified military projects. Company representatives indicated that in the future, their companies expect to use massively parallel supercomputers to a greater extent, primarily for computational fluid dynamics—an application that permits engineers to simulate wind tunnels. However, the lack of software is currently limiting use of this technology.

The major supercomputer applications used by the aerospace industry are computational fluid dynamics, structural analysis, and computational electromagnetics. Most companies we contacted reported that these applications allowed them to perform previously impossible tasks, improve the quality of products, and reduce the time required to make products commercially available. For example, a Lockheed Aerospace representative said that supercomputers allow engineers to perform analyses in areas of fluid dynamics, in which it is impossible to conduct physical tests. A McDonnell Douglas representative stated that supercomputers help engineers improve product quality by permitting them to perform analyses and simulations much faster and over a broader range of parameters.

Computational Fluid Dynamics

Computational fluid dynamics models enable aerospace engineers to simulate the flow of air and fluid around and through proposed structures and components. For instance, engineers can simulate the performance of aircraft and other aerospace vehicles in a wind tunnel. In addition, the performance of other components, such as engines and electrical systems, are simulated using this application. Prior to supercomputers, engineers primarily evaluated aerodynamic designs by testing physical models in wind tunnels and assessing the aircraft during actual flight tests. However, by using a supercomputer, engineers can determine the effectiveness of an aerodynamic design and make modifications before constructing physical models.

Boeing Aircraft Company, for example, used computational fluid dynamics extensively in the design of its newest commercial airplane, the Boeing 777.[1] The company's chief researcher in aerodynamics attributes a faster and more efficient design of this airplane to using computational fluid dynamics on a supercomputer. Those engineers designing the shape of the airplane and those studying its aerodynamics were able to work on the same model simultaneously prior to constructing physical models. Before the use of a supercomputer, these two efforts were done sequentially, relying on physical models and actual wind tunnel tests, which added months to the process.

Boeing was also able to make a computer-generated model of the Boeing 777, rather than a physical model, available to potential customers during the design process. As a result, customer requests for changes were more easily accommodated. For example, one request made during the design phase caused Boeing to redesign a fold in the airplane's wings so that the wing tips would raise, allowing the airplane to fit through existing airport gates. Using computational fluid dynamics on a Cray Y-MP supercomputer, Boeing engineers quickly evaluated three potential designs and identified the configuration with the right balance of aerodynamic efficiency, structural weight, and cost. Boeing's confidence in the results is evidenced by the fact that it committed the wing design to production before conducting physical wind tunnel tests.

In addition, Lockheed Aeronautical Systems Company used computational fluid dynamics on a Cray X-MP supercomputer to develop a computer model of the Advanced Tactical Fighter for the U.S. Air Force. In

---

[1] Boeing announced the development of the 777 aircraft in 1990 and plans to make its first delivery in 1995.

2()

about 5 hours of processing time, Lockheed simulated the fighter's performance using a full-vehicle computer model to obtain aerodynamic data that is normally only provided through actual wind tunnel testing. By using the supercomputer simulations, Lockheed shortened the development phase of the prototype fighter by several months and reduced the amount of traditional wind tunnel testing by 80 hours. The latter resulted in savings of about a half million dollars.

McDonnell Douglas found that using computational fluid dynamics on a supercomputer to simulate the performance of the C-17 military aircraft reduced design and wind tunnel testing time. Engineers used this application to analyze 76 potential wing configurations and identify the best three designs. Thus, they only had to conduct traditional wind tunnel tests on the final three designs. This complete process required 9 months and 350 hours of wind tunnel testing. In contrast, designing the wings of the DC-10 commercial airplane without using this supercomputer application, took 2 years and required 1,200 hours of wind tunnel testing on more than 50 designs.

Structural Analysis

Aerospace companies also use supercomputers to analyze the structure of aircraft and other aerospace vehicles and components to determine optimum weight, shape, and composition. Structural analysis applications simulate the stresses and strains on a given object that result from applied pressure or loads. The structure of an object, such as an aircraft, is first defined as a grid of elements representing its shape and composition. Engineers can then use the model to ascertain the effects of pressure and loads on each element of the structure (e.g., metal fatigue). The supercomputer permits the analysis to be done quicker and more often than performing physical tests on actual structures.

McDonnell Douglas engineers, for example, uncovered and corrected a helicopter fan design flaw in 2 days on their Cray X-MP supercomputer, after spending 6 weeks trying to uncover the problem using other computers. The design flaw concerned the level and location of stress experienced within the helicopter fan hub when spinning. (A fan hub looks similar to a wheel on an automobile and generates air flow inside the tail of a helicopter.) Using structural analysis, engineers were able to identify the fan hub's high-stress areas and make the necessary design modifications.

Computational Electromagnetics

Supercomputers have also been used to simulate the electromagnetic characteristics of military aircraft to make them more difficult to detect using radar. The supercomputer models the chemical composition of the aircraft's outer surface, the geometry of the aircraft, and the reflections of electromagnetic waves off the aircraft's outer surface.

Lockheed Aeronautical Systems Company, for example reduced the radar signature (the size of an image appearing on a radar display) of a low-observable, stealth-like military aircraft, using computational electromagnetics on Cray X-MP and Cray Y-MP supercomputers. This reduced the signature substantially beyond what had been previously obtained using other means. As a result, Lockheed was able to construct fewer physical models and conduct fewer physical electromagnetic tests at a savings of $4.4 million and $1.5 million, respectively.

Automobile manufacturers have been using supercomputers increasingly since 1984 as a design tool to make cars safer, lighter, more economical, and better built, with significant time and dollar savings. By 1989, each of the 12 largest automobile companies worldwide had acquired one or more Cray supercomputers. These supercomputers are enabling automotive engineers to create increasingly sophisticated and realistic simulation models to design and test future vehicles. It would be impractical to perform many of these simulations, such as large three-dimensional models, on anything less powerful than a supercomputer.

Although commercial supercomputer software is available for most automotive applications, much of this software originated from federally-supported research. For example, KIVA, a computer program developed at Los Alamos National Laboratory, is used to model, in three dimensions, the interactions of air and liquids flowing through an engine. NASTRAN, developed under the sponsorship of NASA, is used to analyze various automobile structures. DYNA3D, developed at Lawrence Livermore National Laboratory, is used to simulate car crashes. Commercial vendors and the automobile companies have also modified these applications to meet specific needs of the automobile industry.

Although no companies reported owning massively parallel supercomputers, some of these companies were exploring the potential of this new technology. Some company representatives indicated that massively parallel supercomputers will be important in order to process larger and more complex models, particularly for crash analysis and computational fluid dynamics.

The primary applications used by the automobile industry are automobile crash analysis, structural analysis, and computational fluid dynamics. The five companies we contacted—both American and Japanese—reported that these applications have had a moderate to great impact on improving product quality, reducing costs, and performing previously intractable tasks. According to a Ford Motor Company representative, the intense competitiveness of the industry has made it necessary to use a supercomputer. A Nissan Motor Company representative also stated that Nissan would not consider developing automobiles without a supercomputer.

23

Automobile Crash Analysis

Crash analysis is one of the primary applications for supercomputers in the automotive industry because it provides an alternative to physically crashing test vehicles. The primary advantage of modeling a crash on a supercomputer is that more data about the crash is produced at a cost and within a time frame that would be otherwise impossible to achieve. For example, crash simulations can show how the vehicle structure will collapse during a crash, how fast the driver and passengers will move forward during impact, and when air bag sensors will be activated.

Initially, vehicle crash simulations involved modeling a half-car frontal crash. Today's models have become more complex, involving full vehicle crashes, two full vehicle crashes, and side impact crashes. Using these models, companies can then compare many more designs and optimize a vehicle's structure for weight, stiffness, and strength before a full proto-type vehicle is crash tested.

Physical crash tests require prototype vehicles, which are time-con-suming and expensive to manufacture. According to one automobile company, each test costs between $50,000 and $750,000, depending on whether a production vehicle or a prototype is used. In addition, a pro-totype vehicle can take as long as 8 months to manufacture. In contrast, full car crash simulations can be processed on a supercomputer in less than 20 hours.

According to General Motors Corporation (GM) representatives, crash testing prototype vehicles has been substantially reduced as a result of supercomputing. In testing passenger restraint devices, such as seat belts, during the development of some of its 1992 automobiles, GM crashed about 100 fewer vehicles than it did in 1987. According to company representatives, the reduction was largely due to crash simulations on a Cray supercomputer.

In addition, the Saturn Corporation, a subsidiary of GM, estimated that they performed over 100 vehicle crash simulations on a supercomputer between 1986 and 1990. Data obtained from these simulations were then used to modify the design of the vehicles. As a result, Saturn reported a savings of more than $2 million in development and test costs.

Structural Analysis

Automobile companies use structural analysis models to simulate the physical structure of an automobile, including parts and components, to improve functionality and reduce manufacturing costs. The simulations help engineers design stronger parts that are lighter and less susceptible

to problems caused by vibration and stress. Lighter parts also contribute to improvements in fuel economy.

Chrysler Corporation improved the design of a small car's body structure and a convertible's floor structure by modeling them on a Cray X-MP supercomputer. The stiffness of the small car body structure was improved by 10 percent, while its weight was reduced by 45 pounds. The improvement in body stiffness on the small car made the vehicle easier to handle and gave it a better ride. The floor of the convertible was lengthened by 8 percent, yet reduced in weight by 9 pounds. Chrysler estimates that it will save about $3.9 million annually in reduced raw materials to manufacture both vehicles.

Chrysler was also able to eliminate the need for a dash board bracket in a new minivan by modeling the dash board structure on a Cray X-MP supercomputer. The model showed that the remaining dash boa !
brackets were sufficient to hold the dash board structure in place. Chrysler estimates that the elimination of the $2 bracket will save about $940,000 annually in the cost of parts to manufacture the vehicle.

In addition, Chrysler improved the design of a new 2.0 liter engine by modeling its structure on a Cray X-MP supercomputer. The stiffness, or rigidity, of the engine structure was increased by 3 percent, thus contributing to a smoother running engine. In addition, its overall weight was reduced by 9 pounds. Chrysler estimates that it will save about $1.8 million annually in reduced raw materials to manufacture the engine.

Computational Fluid Dynamics

Computational fluid dynamic models are used to simulate the flow of air or liquids around or through automobile parts or structures. Specifically, models simulate the exterior flow of air around a vehicle, the flow of air within the vehicle, and the flow of air and liquids within the engine, cooling system, and air conditioning system.

At Chrysler Corporation, for example, engineers used a Cray X-MP supercomputer to simulate the flow of liquids through the cooling system of one of its vehicles. This analysis helped determine which design was most efficient in cooling the vehicle's engine. Thus, the design of the system was optimized to improve its cooling efficiency while reducing the number of needed parts. Chrysler estimates that this reduction will reduce manufacturing costs by about 5 percent, saving about $1.6 million annually.

The Chemical and Pharmaceutical Industries

The chemical and pharmaceutical industries are the fastest-growing group of industrial supercomputer users, although their use is still in its infancy. The impetus for these industries to begin using supercomputers was a combination of recent advances in theoretical chemistry and three-dimensional visualization techniques. Du Pont became the first company in these industries to buy a vector supercomputer in 1986. Since then, several other American and Japanese companies have begun using supercomputers. Two companies—Dow Chemical Company and Eli Lilly & Company—are now experimenting with massively parallel supercomputers. Lilly is using a supercomputer at the National Center for Supercomputing Applications—a NSF supercomputer center at the University of Illinois—to develop code to be used on massively parallel supercomputers.

Some of the important chemical software applications used today were derived from applications developed at government laboratories, such as NASA Ames Research Center. However, little application software has been available until very recently, because the chemical industry only recently began using supercomputers. In order to develop sophisticated application software on a reduced, shared-cost basis, Cray Research, Inc., other vendors, chemical companies, and government laboratories have formed a chemical software consortium. Member companies include Lilly, Monsanto Company, Exxon Research and Engineering, Du Pont, and 3M Corporation.

The major supercomputer applications in the chemical and pharmaceutical industries include molecular modeling, structural analysis, and computational fluid dynamics, and are used in basic research, product development, manufacturing process design, manufacturing plant design, environmental impact assessment, and waste disposal. Thus, supercomputing affects many stages of the industry's product lines. These product lines number in the thousands and include industrial chemicals, polymers,[1] and biological materials for agriculture and medicine. All company representatives reported that the use of supercomputers greatly helps them perform previously intractable tasks—such as the study of complex molecules. Several representatives also said that supercomputers helped them reduce cost and time to market, improve quality, and develop a greater variety of products.

---

[1] A chemical compound or mixture of compounds containing repeated structural units of the same original molecules, such as Nylon.

26

Molecular Modeling

Molecular simulations enable scientists to study the molecular proper-ti s of chemical compounds used in the development of drugs and other products. One of the keys to understanding molecules lies in gaining a clear appreciation of their three-dimensional shape. Unlike the rigid "ball and stick" models that scientists built in the past, the atomic positions in a molecule are constantly changing. Using supercomputers in conjunction with workstations, scientists are able to construct images, such as those of large, complex human proteins and enzymes. Scientists can then rotate these images to gain clues on biological activity and reactions to various drug candidates.

The use of molecular modeling has also been important, from an economic perspective, in the development of new drugs. A Du Pont scientist estimated that about 30,000 compounds are synthesized—at a cost of about $5,000 per synthesis and initial screening—for every new drug that is developed. As such, as much as $150 million can be invested in discovering a drug, even before clinical testing, federal government approval, and manufacturing and development costs are added. By making this drug discovery process more "rational," with less trial and error, a Du Pont representative estimates that millions of dollars can be saved.

Du Pont is currently developing replacements for chlorofluorocarbons, compounds used as coolants for refrigerators and air conditioners, and as cleansing agents for electronic parts. These compounds are being phased out because they are thought to contribute to the depletion of ozone in the atmosphere. Du Pont is designing a new process to produce substitute compounds safely and cost-effectively. These substitutes will be more reactive in the atmosphere and will decompose faster. Du Pont is using a supercomputer to calculate the thermodynamic data needed for developing this process. These calculations can be completed by the supercomputer in a matter of days, at an approximate cost of $2,000 to $5,000. Previously, such tests were conducted in a laboratory, and required up to 3 months to conduct, at a cost of about $50,000. Both the cost and time required for such traditional methods would have substantially limited the amount of testing that is now being done.

Structural Analysis

Computer-based structural analysis techniques are used in the chemical and pharmaceutical industries to determine stress and durability of both products and processing equipment. For example, Du Pont engineers used structural analysis to optimize the design of a mold used in a new

process for manufacturing Corian[2] sinks. This enabled them to quickly determine the wall thicknesses, ribbing, and reinforcements necessary to withstand the molding pressure while minimizing its weight. Thus, this application enabled them to develop the manufacturing process 6 to 12 months sooner than would have been possible using a series of proto-type molds. This saved development costs and increased revenue by getting the new process in operation faster.

Computational Fluid Dynamics

The chemical and pharmaceutical industries use computational fluid dynamics to model products and product performance, and to aid in the design of manufacturing processes and plants. These models predict how fluids flow through equipment and can simulate processes, such as mixing, drying, cooling, and separation. For example, Du Pont has successfully used this application on a Cray Y-MP supercomputer to model manufacturing processes for sheet products, such as film and plastic wrappers. These products, particularly X-ray film, must be uniform in thickness to meet the required quality standard. Thus, computational fluid dynamics provides Du Pont with an additional design tool to improve sheet product thickness uniformity, without increasing manufacturing time.

Du Pont also models the manufacturing process for Nylon, one of its major products. Computational fluid dynamics solutions to these models accurately quantify the flow through the processing equipment and locate where clogging is most likely to occur. Clogging is a major problem in Nylon plants, causing periodic shutdowns. By modeling the process, Du Pont has been able to reduce the number of shutdowns.

---

[2]Corian is a registered trademark of E.I. du Pont de Nemours and Company.

Companies Interviewed Regarding Supercomputer Use

Aerospace Companies	The Boeing Company, Seattle, Washington General Dynamics Corporation, Fort Worth, Texas Grumman Corporation, Bethpage, New York Lockheed Corporation, Calabasas, California McDonnell Douglas Corporation, Hazelwood, Missouri Pratt & Whitney, United Technologies Corporation, East Hartford, Connecticut
Automobile Companies	Chrysler Corporation, Highland Park, Michigan Ford Motor Company, Dearborn, Michigan General Motors Corporation, Warren, Michigan Honda Motor Company, Ltd., Tokyo, Japan Nissan Motor Company, Ltd., Tokyo, Japan
Chemical and Pharmaceutical Companies	Dow Chemical Company, Champaign, Illinois E.I. du Pont de Nemours and Company, Wilmington, Delaware Eli Lilly & Company, Indianapolis, Indiana Merck & Company, Rahway, New Jersey Monsanto Company, St. Louis, Missouri
Petroleum Companies	Amoco Production Company, Tulsa, Oklahoma Atlantic Richfield Company (ARCO), Plano, Texas British Petroleum (BP), Houston, Texas Chevron Corporation, Houston, Texas Exxon Production Research Company, Houston, Texas Oryx Energy Company, Dallas, Texas Phillips Petroleum Company (Phillips 66), Bartlesville, Oklahoma Shell Oil Company, Houston, Texas

Companies Interviewed Regarding High-Speed Network Use

Amoco Production Company, Houston, Texas
Apple Computer, Inc., Cupertino, California
General Motors Research Corporation, Warren, Michigan
Hewlett-Packard Company, Cupertino, California
Intel Corporation, San Jose, California
International Business Machines Corporation, Austin, Texas
Landmark Graphics Corporation, Houston, Texas
Schlumberger Well Services, Houston, Texas
Sun Microsystems, Inc., Milpitas, California
Tandem Computers, Inc., Cupertino, California

3 9

Major Contributors to This Report

Information Management and Technology Division, Washington, D.C.

Linda D. Koontz, Assistant Director
Valerie C. Melvin, Assignment Manager
Beverly A. Peterson, Evaluator-in-Charge
Nancy M. Kamita, Computer Scientist

Los Angeles Regional Office

Allan Roberts, Assistant Director
Ambrose A. McGraw, Regional Assignment Manager
Benjamin H. Mannen, Senior Evaluator
Shawnalynn R. Smith, Staff Evaluator

San Francisco Regional Office

Frank Graves, Regional Assignment Manager
Don Porteous, Staff Evaluator

3 1

Ordering Information

The first five copies of each GAO report are free. Additional
copies are $2 each. Orders should be sent to the following
address, accompanied by a check or money order made out to
the Superintendent of Documents, when necessary. Orders for
100 or more copies to be mailed to a single address are
discounted 25 percent.

U.S. General Accounting Office
P. O. Box 6015
Gaithersburg, MD 20877

Orders may also be placed by calling (202) 275-6241.

BEST COPY AVAILABLE