High-Performance Computing and Communications

Rick Stevens
Mathematics and Computer Science Division
Argonne National Laboratory
Argonne, IL 60439
e-mail: stevens@mcs.anl.gov

ABSTRACT

This presentation has two parts. The first part discusses the U.S. High-Performance Computing and Communications program—its goals, funding, progress, revisions, and research in high-performance computing systems, advanced software technology and applications, national research and education networking, information infrastructure technology, and basic research and human resources. The second part of the presentation covers specific work conducted under this program at Argonne National Laboratory. Argonne's efforts focus on computational science research, software tool development, and evaluation of experimental computer architectures. In addition, we describe collaborative activities at Argonne in high-performance computing, including an Argonne/IBM project to evaluate and test IBM's newest parallel computers and the Scalable I/O Initiative being spearheaded by the Concurrent Supercomputing Consortium.

OVERVIEW

In 1991 the U.S. Congress passing the High Performance Computing and Communications (HPCC) Act. This bill established a significant increase in federal funding for the four principal agencies engaged in supercomputing research (ARPA, DOE, NASA, and NSF). The legislation specified certain roles for each agency (e.g., networking to NSF, hardware development to ARPA), but much was left to the agencies to determine. Following the establishment of the federal HPCC program, many universities and national laboratories established complementary programs.

This report has two parts. Part 1 is a general overview of the federal HPCC program and some of the projects started since December of 1991. We also discuss the role of the Department of Energy (DOE) national laboratories in the HPCC program. Part 2 is a description of the HPCC program at Argonne National Laboratory. We conclude with comments on the impact of high-performance computing to industry, computer science, and computational science research.

THE FEDERAL HPCC PROGRAM

The federal HPCC program began in 1991 with four components: high-performance computing systems, advanced software technology and algorithms, national research and education network, and basic research and human resources. Late in 1993, a fifth area was added to the HPCC program: Information Infrastructure Technology and Applications (see Table 1).

Table 1. HPCC budgets by agency (millions of U.S. dollars)

Agency	FY1992	FY1993
ARPA	232.2	275.0
NSF	200.9	261.9
DOE	92.3	109.1
NASA	71.2	89.1
HHS/NIH	41.3	44.9
DOC/NOAA	9.8	10.8
EPA	5.0	8.0
DOC/NIST	2.1	4.1
Total	654.8	802.9

Initially, four agencies—ARPA, DOE, NASA, and NSF—received the bulk of the funding and initiated the majority of the programs. Later, other agencies—EPA, NSA, NIST, NIH and NOAA—joined the main four. ARPA (or DARPA at the time) was selected to lead the development of massively parallel computing hardware and at the time was supporting significant development efforts at Intel Supercomputer Systems Division and at Thinking Machines Corporation. NSF was selected to lead the effort to build the NREN and, with ARPA, launched the national Gigabit Testbed Projects that will in 1993-94 demonstrate applications of gigabit networks. DOE was selected to acquire several large-scale prototype massively parallel processing (MPP) systems and make them available for applications and tools projects. NASA was to collaborate with DOE on applications and to formulate new models and methods for aerospace and space science applications. All agencies would continue to support research in parallel software tools and basic research in computer science necessary to provide a new stream of ideas for more applied work.

High-Performance Computing Systems

The federal agencies have supported different stages in the development of high-performance computing systems. ARPA has led the research efforts, whereas all the agencies have been involved in prototyping (by funding purchases for testbeds, etc.). A good example of "partnership development" of high-performance computing systems is the Concurrent Supercomputing Consortium (CSCC), which purchased the Intel Touchstone DELTA System. The DELTA, developed by Intel with ARPA support, was the first general-purpose (MIMD) parallel computer to be significantly faster than a conventional vector supercomputer. The CSCC action has become a model for other groups acquiring parallel computers under the HPCC program.

During the past two years, a number of other MPP systems have been produced and marketed in the United States, including Intel

Paragon, TMC CM-5, Cray Research T3D, IBM SP1, Kendall Square KSR1, and nCUBE nCube-2. Current systems under development include fine-grain parallel machines, machines with support for global shared memory, and machines that combine the latest workstation technologies into an MPP node (including hardware support for multimedia). Scalable I/O architectures remain a significant technical challenge, and integrated single-kernel image-operating systems are an ongoing research activity.

Advanced Software Technology and Algorithms

The advanced software technology and algorithms component of the national HPCC program supports work in four areas: Grand Challenges, software components and tools, computational techniques, and High-Performance Computing Research Centers. In many cases, the work in software components and tools and in computational techniques was already in progress under previously support efforts—primarily, computer science research. The main new elements are Grand Challenges and High-Performance Computing Research Centers.

Grand Challenges combine applications scientists with applied mathematicians and computer scientists to solve large-scale computing problems. The applications areas range from computational fluid dynamics to quantum chromodynamics and global climate modeling. Typical projects are supported at the level of $500K-$1000K per year for three to five years. Three agencies have supported calls for Grand Challenges (DOE, NSF, and NASA). About 35 Grand Challenge projects were under way at the end of FY 1993, with at least an equal number of smaller computational science projects funded at a lesser level.

High-Performance Computing Research Centers (HPC-RCs) have been established at several DOE sites at the NSF Supercomputer Centers and at other supercomputing facilities. HPC-RCs are characterized by the installation of one or more MPPs (current sites have TMC CM-5, Intel Paragon and DELTA, Cray T3D, IBM SP1, Meiko CS-2, KSR KSR1), active research projects in the development of software tools and systems administration tools, and development of systems software. The primary mission of the HPC-RCs is to provide the cycles and services for the Grand Challenge applications teams; however, they are also functioning as a de facto testbed for the large-scale systems they employ; some centers also provide user services and support. In addition, the HPC-RCs are providing a valuable service to the vendor community, enabling them to test and evaluate these new systems in environments that are more friendly than typical commercial installations.

National Education and Research Network

The National Education and Research Network (NREN) is the union of all the federally supported Internets plus the local or regional networks that connect to the backbones. At the end of FY93 the NREN consists of the National Science Foundation's NSFnet, DOE's ESnet, ARPA-supported MILnet, and NASA's SCInet, as well as about a dozen local networks serving the principal metropolitan areas of the country. The fastest links in this integrated Internet are the 45-Mb/s backbone links of the NSFnet; however, selected ESnet connections will be upgraded soon. In addition to the "production" network, the NREN component of the HPCC program supports the development and limited deployment of gigabit networks. Currently, five gigabit testbeds are under development and could form the nucleus of a gigabit NREN in the mid-1990s. The networks are all capable of supporting TCP/IP protocols and have led to the development and widespread use of network resources like Gopher, World Wide Web, and the MBONE desktop video conference systems. It is not uncommon today to be searching the Library of Congress book catalog while watching a scientific conference on the MBONE and computing at a remote supercomputer. Indeed, the rapid success of the NREN is largely driving the follow-on program to the HPCC the National Information Infrastructure Initiative.

Basic Research and Human Resources

To ensure the continued production of computational scientists trained in the use of applying parallel supercomputers to problems in science and engineering, a program was created as part of the HPCC program to fund graduate fellowships and postdoctoral positions at HPCC research centers. This program also provides for limited support of existing basic computer science research programs that the agencies already had in place. Some agencies have used this part of the program to build education projects that expose undergraduate and precollege students to the world of supercomputing and computational science.

Role of the DOE National Laboratories in HPCC

DOE has been involved in supercomputing longer than any other federal agency. It has developed a computational science culture that reflects the easy access to large computers: problems are routinely conceived and solved that would not even be attempted at a university or company. In support of this computational science culture, research programs have grown up to provide numerical methods, software libraries, algorithms, and programming environments that are needed to utilize advanced computers.

The laboratory system, while sometimes competitive, is capable of cooperation to provide new capabilities. Unlike a university, the labs can engage in large interdisciplinary projects without suffering from academic turf battles. Unlike some of the NASA centers the DOE laboratories have expertise in many disciplines locally available. Unlike the NSF supercomputer centers which must serve a large community, the DOE computing centers have the ability to focus resources on a limited number of problems, thus ensuring that progress in made.

But the most significant value that the DOE brings is forty years' experience with high-performance computing. Many of the production codes now being used (serially in most cases) in industry were developed by the national laboratories. This intimate relationship with the applications codes that are at the core of modern industrial supercomputing is DOE's biggest asset. Ironically, it is one that is being underutilized in the current funding scenarios. Many researchers believe that, with adequate resources, DOE could provide three important roles in the HPCC:

1. DOE could develop a new *software infrastructure* for American industry. This software infrastructure will be one of the foundations of the National Information Infrastructure (NII) and would leverage DOE research programs in applied mathematics, computer science, and

computational science.

2. DOE could be a *national testbed* for research computing environments. Specifically, DOE could enable the vendor community to experiment within the DOE environment before widespread deployment in the private sector.

3. DOE could promote widespread utilization of *high-performance visualization and portable computing*, thereby increasing the productivity of the researcher.

Technology Transfer and Support for American Industry

The DOE laboratories have enjoyed close relationships with industry in the past. The HPCC program has placed increased emphasis on "formalizing" industrial relationships. Technology transfer for the HPCC program generally takes two forms: (1) DOE joins with a vendor (hardware or software) in development or evaluation or both; or (2) DOE joins with a "user" industry (a company that is a user of computing technology rather than a developer or vendor of that technology).

DOE has always had strong relationships with vendors since the early days of supercomputing. The largest payoff to the American economy, however, is through the user industry. Under the auspices of the HPCC program, DOE labs are helping technical users adopt computing technology demonstrated at the laboratory. For example, Argonne has worked closely with Motorola to develop a model of a piezoelectric crystal oscillator that is used in cellular phones and computers. This model runs on an MPP system and can solve a problem that Motorola was struggling with on a conventional supercomputer.

Industrial Supercomputing Applications

American industry has been relatively quick to adopt computing tools that can be demonstrated to improve productivity in the short term. The industry has been slower, however, to exploit supercomputers to improve its capability for better design and simulation in such areas as structural mechanics, computational fluid dynamics, chemical process dynamics, molecular modeling, electromagnetics, materials properties, crystallography, computerized tomography, and image compression.

The root cause may be a a lack of awareness about (1) which tools or programs might improve a particular product or process, (2) what kind of integration is available to make the results usable for automated manufacturing, and (3) whether the performance needed by the particular application is achievable with existing hardware and existing budgets.

Independent software vendors are a good vehicle for solving the first two components of this awareness problem, because they are closest to the marketplace. DOE, on the other hand, is likely to make a major contribution in the third area. In this area, industry has neither the financial nor the human resources to undertake in-house development or porting of applications. Similarly, software vendors—often relatively small organizations—cannot afford to invest in research activities that may not pay off for a number of years. Many national laboratories, on the other hand, have had significant experience in developing, maintaining, and porting large applications codes and would be in an ideal position to port popular applications to MPP systems that may provide significant price/performance important to industry.

The challenge is to avoid having to produce a separate version of each application for each machine. One proposal is to use a portable programming environment such as HPF or MPI to produce one version of each application; that version would then be tuned to a specific architecture by the hardware or software vendor. In this way the national laboratories could, within two or three years, provide American industry with portable parallel implementations of critical simulation and design software.

Observations on the U.S. Program

The U.S. high-performance computing scene can boast of several significant accomplishments: (1) establishment of numerous computing consortia, typically a mixture of universities, labs, and industry; (2) installation of more than 100 MPP computers; and (3) porting or re-implementation of diverse scientific applications on parallel computers. Most of these computers and applications have been influenced either *directly* by participation in an HPCC consortium or *indirectly* by use of the resources assembled by the consortia. And although the largest computers are still going to national laboratories, they are no longer being limited to defense or classified work (e.g., at Los Alamos the largest CM-5 is installed in the open use partition).

Another significant trend is evident at the university-based NSF-supported supercomputer centers. While they are still largely based on vector supercomputers (Crays and IBMs). they have begun making the transition to MPPs, with each center focusing on a different vendor: the National Center for Supercomputer Applications, a 512-CPU CM-5; San Diego Supercomputer Center, a 400-CPU Intel Paragon; Cornell National Supercomputing Facility, a 64-CPU IBM SP1; and the Pittsburgh Supercomputer Center, a 256-CPU Cray T3D.

We also note a growing interest in establishing parallel computing standards. Motivated by the large number of message-passing systems that have been developed, scientists have initiated a national effort to standardize the syntax and semantics of message passing into a new standard called Message Passing Interface (an effort loosely coordinated by the NSF Center for Research in Parallel Computation and currently involving more than thirty companies and software developers). A similar effort has led to the parallel language standard High Performance Fortran, which is now being implemented by many vendors.

In summary, the U.S. high-performance computing scene can be described as quite healthy. One note of caution: The pace of development of parallel systems has increased significantly in recent years, and the major vendors are now offering new systems about every 12-18 months. This situation places an enormous strain on software developers and on applications users. Effective exploitation of this new technology is one of the major challenges facing the federal HPCC program.

Information Infrastructure Technology and Applications: The Next Step

The National Information Infrastructure (NII) initiative is a proposed federal program to develop both the technology and applications to take maximum advantage of U.S. computing and networking resources for both federal and private use. The NII will expand the national networking connectivity to include homes, businesses. and schools. Specifically, it will develop large-scale digital libraries; advanced database and multimedia capabilities; and the gigabit networking technology needed for real-time interactive video computing, wide-area distributed computing applications, and the rapid exchange of large volumes of data from coast to coast. The NII will provide to businesses the ability to link databases for manufacturing and design and will allow small companies to combine resources for attacking large problems. It will also provide transparent access to supercomputers for design and engineering simulation.

The DOE laboratories can play a major role in the development and deployment of the NII. As a microcosm of advanced R&D and environmental management, the DOE labs can be an ideal role model for American industry on how to utilize the emerging NII to improve productivity, safety, environmental monitoring and compliance, and education and training.

Five application areas and seven technology areas have been proposed for DOE's NII program. The applications are manufacturing, environment, energy demand management, digital libraries, and education. The technology areas are user interfaces and multimedia, very large databases and database mining, internetworking and the integration of WANs and LANs, enterprise computing, virtual reality environments, and telepresence. These core areas are closely supported by existing work in the DOE funded by a variety of programs and will provide the key technologies for exploitation of NII by DOE.

ARGONNE'S HIGH-PERFORMANCE COMPUTING AND COMMUNICATIONS PROGRAM

As an outgrowth of parallel computing research supported at Argonne during the 1980s, and in response to the federal HPCC program, Argonne started its own High-Performance Computing and Communications program in 1990. The Argonne HPCC program includes research in Grand Challenge applications, development of software tools for parallel computing, evaluation of experimental computer architectures, and collaborative projects with research institutions and industry.

Grand Challenge Applications

Argonne's HPCC program supports five Grand Challenge-class projects: computational chemistry, computational biophysics, superconductivity, global climate modeling, and stellar astrophysics. In each case, applications scientists are collaborating with computer scientists to develop parallel implementations of models and to design new algorithms for exploiting the architecture of MPP systems.

Software Tools for Parallel Computing

At Argonne a number of computer science projects are under way with the objective of producing programming tools and libraries that will improve scientists' ability to develop portable codes for MPPs. Argonne is supporting the development of Fortran M, PCN, p4, Upshot, Chameleon, and BlockSolve and is an active participant in the Message Passing Interface research project. Since scientific visualization is a critical component of analyzing the results of large-scale computations, Argonne also supports projects developing parallel renderers, data management schemes, and high-speed networks to support visualization.

Experimental Computer Architectures

Argonne has established a partnership with IBM's Highly Parallel Supercomputer Systems Laboratory to explore the development of MPP systems based on workstation node technology. Together Argonne and IBM are developing and testing an experimental configuration of IBM's new SP1 scalable parallel supercomputer. The Argonne configuration consists of 128 nodes each with a 62.5 MHz RIOS1 chipset, 128 MB of RAM, 1 GB local disk, and several network connections. This system will be used for applications development and as a testbed for novel operating systems and parallel computing tools. The nodes are packaged in racks of 16 with N+1 redundant power supplies. Each node is connected to other nodes via a high-speed Omega switch with a bandwidth of 40 MB/s and less than 1 usec hardware latency. In addition to the high-speed switch, each node is connected to two ethernets (for booting and for tcp access), and every fourth node is connected via Fiber Channel Standard (FCS) to a set of eight fileservers. These fileservers are large IBM 9 70B servers with FCS and HIPPI interfaces. Four IBM 9570 disk arrays are connected via HIPPI to the set of fileservers. Connected to the HIPPI switch is an automated tape library with 6.4 TB of DD2 and three 19-mm tape drives (see Figure 1).

Collaborative Projects

Central to Argonne's activities in high-performance computing is participation in multi-institutional collaborative projects. Efforts to date have largely been in four areas:

1. *NSF Science and Technology Center for Research in Parallel Computation:* Together with Rice University, Caltech, Los Alamos National Laboratory, Syracuse University, the University of Maryland, and the University of Tennessee, Argonne was involved in the formation of the National Science Foundation Center for Research in Parallel Computation (CRPC). The CRPC conducts research projects in the development of parallel languages (Fortran-based languages and others), methods for parallel solution of linear equations and numerical PDEs, parallel software for solving large scale optimization problems, and tools for automatically compiling parallel programs from sequential codes. The CRPC has also funded the (partial) acquisition of parallel computer systems including the Touchstone DELTA and Paragon, CM-5, KSR, MasPar-2, and SP1.

2. *Concurrent Supercomputing Consortium:* With Caltech, JPL, Purdue, CRPC, LANL, Sandia National Laboratories, Pacific Northwest Laboratory, the National Science Foundation, ARPA, and NASA, Argonne participated in the commissioning of the 512-node Touchstone DELTA system in 1991. Recently this system has been upgraded to a 512-node Intel Paragon. This system has been used by hundreds of researchers to attack problems from molecular modeling to global climate change.

3. *CAVE Automatic Virtual Environment:* The University of Illinois at Chicago, in collaboration with Argonne and the National Center for Supercomputing Applications, has developed a walk-in virtual reality

environment called the CAVE. CAVE has a number of advantages over head-mounted systems: it allows multiple researchers to experience the same virtual environment, it enables users to interact with the virtual environment in a natural way, and it has over four times the resolution of typical displays. CAVE has been used for advanced three-dimensional visualization of biological and chemical data.

4. *Scalable I/O Initiative:* The Scalable I/O Initiative brings together systems software developers, computer vendors, and applications teams to address the problem of developing hardware and software systems to support scalable I/O for high-performance computer systems. Unlike the current state of parallel operating systems where a commonly available software platform exists (e.g., Mach), vendors must work largely from scratch when developing the filesystems and user software needed to support a scalable I/O system. This problem unnecessarily duplicates efforts and and results in multiple solutions to a single problem. The Scalable I/O Initiative will provide a commonly available set of software modules that will serve as the foundation for the next generation of parallel I/O systems. A full-scale testbed will also be provided for the development and evaluation of new systems software for scalable I/O.

IMPACT

One objective of both the federal and local High-Performance Computing and Communications programs is to encourage industrial use of advanced computers. The solution of many important applications (e.g., in structural mechanics, computational fluid dynamics, and materials science) requires the use of the most powerful supercomputers available. However, since industry is often not in a position to undertake inhouse applications development, the U.S. government has emphasized the need for collaborative projects. A recent example of such a collaborative project at Argonne involves a consortium of chemical companies, which are cooperating to develop a new generation of computational chemistry applications codes. The objective is to enable the chemical companies to model chloroflorocarbon substitutes for environmental problems before manufacture. This and similar projects are expected to lead to a high-performance revolution, with widespread interest in and use of MPP systems by industry.

The HPCC program has also profoundly affected U.S. computer science research by providing a focused goal for the more applied work in tools and languages and for architecture development.

Finally, in a more indirect way, the HPCC program has created a new community of young researchers with the goal of developing useful computing systems and tools. These researchers will be developing the computing technology for the next decade. The HPCC program has established the framework and foundation for this work.

ACKNOWLEDGMENT

This work was supported by the Office of Scientific Computing, U.S. Department of Energy, under Contract W-31-109-Eng-38.

Figure 1. ANL's SP1 configuration

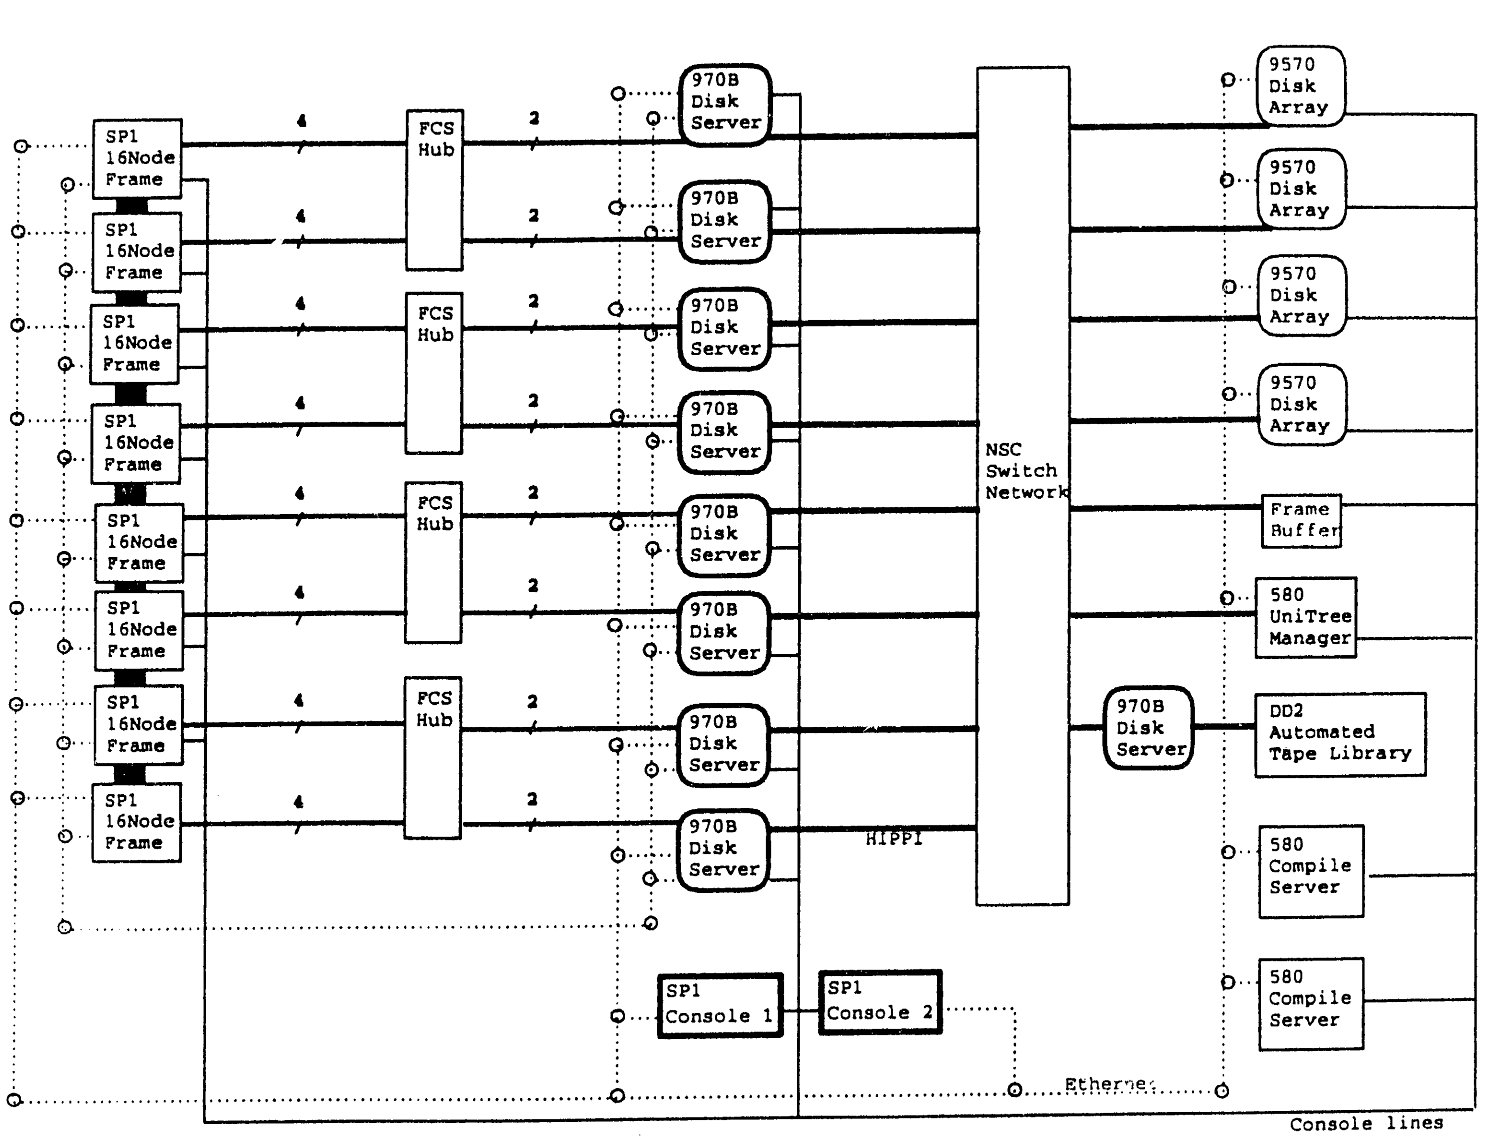

END

DATE FILMED

12/27/93