LA-UR-77-2026

TITLE: EXPERIENCE WITH THE CRAY-1 COMPUTER AT THE LOS ALAMOS SCIENTIFIC LABORATORY

AUTHOR(S): Fred W. Dorr

SUBMITTED TO: Electric Power Research Institute

los alamos scientific laboratory
of the University of California
LOS ALAMOS, NEW MEXICO 87545

Form No. 836
St. No. 2629
1/75

UNITED STATES
ENERGY RESEARCH AND
DEVELOPMENT ADMINISTRATION
CONTRACT W-7405-ENG. 36

EXPERIENCE WITH THE CRAY-1 COMPUTER AT THE LOS ALAMOS SCIENTIFIC LABORATORY

Fred W. Dorr

Computer Science and Services Division
Los Alamos Scientific Laboratory
Los Alamos, New Mexico 87545

INTRODUCTION AND SUMMARY

The Cray Research Inc. CRAY-1 computer (Serial No. 1) has been installed at the Los Alamos Scientific Laboratory (LASL) since April 1, 1976. This report describes our experience with the CRAY-1 since that time. The first six months of the period were spent in a detailed evaluation of the hardware. The machine successfully passed the performance thresholds that had been established, and on October 1, 1976, it became available for productive use. Our initial efforts were concentrated on achieving an early production capability for several large codes. In order to achieve this objective, we established a link between the CRAY-1 and a CDC 7600 running the Livermore Time-Sharing System (LTSS). This link became operational in December 1976 and security approval was received in January 1977. The first major production code became operational in May 1977, and production usage has increased since that time. We are currently developing an operating system for the CRAY-1 that is tailored to the LASL computing environment. We are also in the process of procuring several smaller-scale computers with extensive software facilities that will support the CRAY-1.

EVALUATION

LASL began a performance evaluation of the CRAY-1 computer on April 1, 1976, in order to determine whether the CRAY-1 should be considered for procurement by the Laboratory. The evaluation was conducted by LASL for the U.S. Energy Research and Development Administration (ERDA) with technical consultation provided by FEDSIM, a computer performance consulting branch of the U.S. Air Force. LASL and ERDA wished to verify that the CRAY-1 could meet performance requirements for the LASL workload in the areas of speed and reliability.

This work was supported by the U.S. Energy Research and Development Administration.

Performance threshold criteria were formally established in three areas: scalar processing speed, vector processing speed, and reliability. The evaluation team faced several serious problems. First, there was very little software for the machine; for example, no FORTRAN compiler was operational at that time. Second, only limited I/O capabilities were available at the beginning of the evaluation period, so it was difficult to run programs on the CRAY-1. Finally, Cray Research Inc. had agreed to provide the machine at no expense to the government for a limited period of six months, which provided an inflexible deadline for the completion of the evaluation.

Scalar Processing Speed

The threshold criterion for scalar performance was that the CRAY-1 should be at least twice as fast as a CDC 7600 on the designated LASL workload. This criterion was established in order to estimate a "worst case" level of performance for a supercomputer that had both scalar processing and vector processing modes. The degree to which the vector processing features of a computer can be used depends on the inherent level of parallelism in a particular code, and upon the effectiveness of recognition of this parallelism by special compilers. Since both of these factors could only be estimated for the CRAY-1, a special scalar performance speed test was adopted for the evaluation.

The scalar processing speed test was the most time-consuming phase of the evaluation. We adopted a sampling technique for determining whether the CRAY-1 met the threshold criterion, since the absence of a FORTRAN compiler prevented running entire code segments. The most time-consuming LASL production codes were identified by a workload analysis, and small samples were drawn from these codes for implementation in assembly language on both the CRAY-1 and the CDC 7600. The speed ratios of the 23 implemented assembly language kernels are shown in Figure 1. The sampling techniques were selected to provide us with 90% confidence that the CRAY-1 would perform scalar processing at least twice as fast as the CDC 7600, assuming that assembly language coding was performed equivalently on each machine. The CRAY-1 passed this performance test.

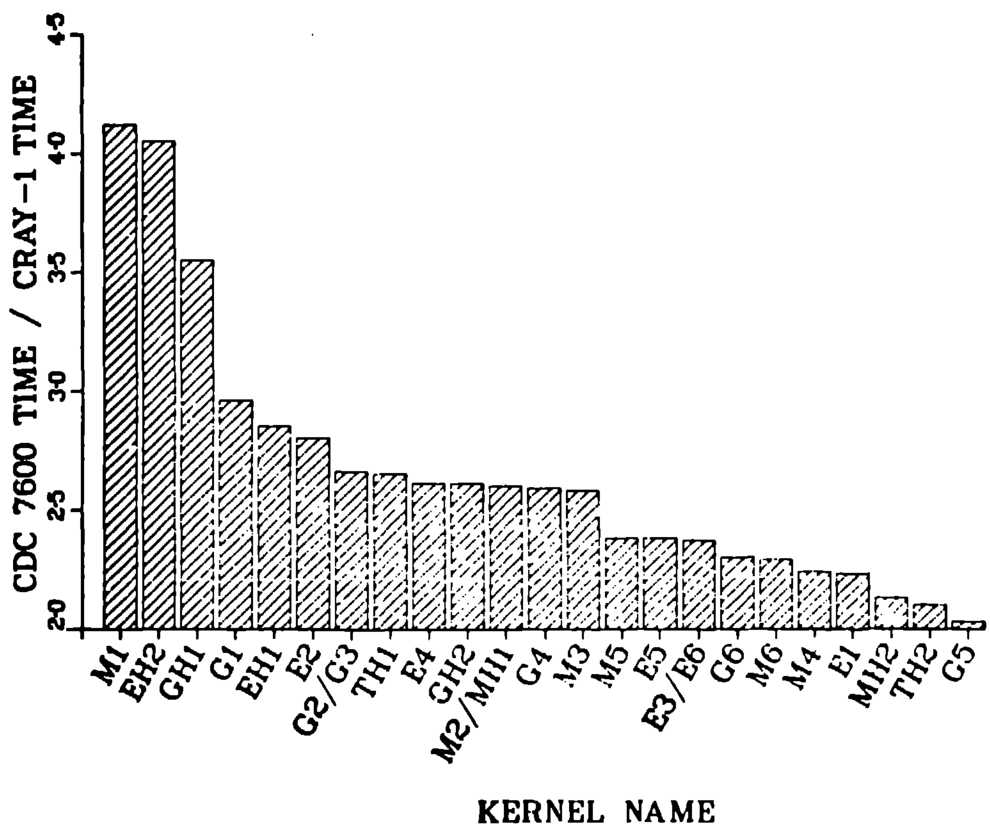

Figure 1. CRAY-1 Scalar Performance

Vector Processing Speed

The threshold performance criteria were that the CRAY-1 should perform operations in vector mode at speeds 3, 4, and 5 times faster than a CDC 7600 for vectors of lengths 20, 100, and 500, respectively. The machine was tested on a set of small vector programs for the various vector lengths. The average speed ratios of five vector functions for the appropriate vector lengths were 3.30, 4.50, and 5.12, respectively. Thus, the CRAY-1 passed the vector processing speed threshold criteria.

Reliability

The LASL reliability requirements specified a minimum level of availability (80%), a maximum mean-time-to-repair (one hour), and a minimum mean-time-to-failure (four hours). We adopted a conservative and demanding reliability testing procedure by running a test program for over 600 hours on the CRAY-1 during the six-month evaluation period. The reliability test program was designed to exercise the machine at resource utilization levels thought to be several times higher than a production workload. The CRAY-1 passed all of the reliability criteria. The best performance data for a particular period of 20 contiguous days were 95% availability, 0.38 hour mean-time-to-repair, and 7.08 hours mean-time-to-failure. Partly as a result of the LASL reliability evaluation, Cray Research announced that error correcting memory would be available as an option on future versions of the CRAY-1, beginning with Serial No. 3.

EARLY PRODUCTION

The critical shortage of LASL computing capability required that the CRAY-1 be made productive as quickly as possible. The mode of operation used during the evaluation period was not adequate to support a full production environment. During the evaluation, the codes that executed on the CRAY-1 were written in assembly language and assembled on the Data General ECLIPSE Maintenance Control Unit (MCU). The binary output from the assembler was then shipped to the CRAY-1 over the CRAY-1/MCU link, loaded into memory, and executed. At the completion of execution, the user examined memory locations and registers to obtain the results of the calculation. The obvious bottlenecks were the MCU link (which could be used by only one user at a time), the lack of a FORTRAN compiler, and the lack of adequate supporting software. In order to achieve the early production requirements at LASL, we needed more efficient and effective access to the CRAY-1, a FORTRAN compiler and I/O library, a mathematics library, and additional utility software.

In order to provide a more effective means of access to the CRAY-1 as quickly as possible, a link was developed between the CRAY-1 and a CDC 7600 running LTSS. The LTSS system was chosen to provide this initial capability because the pre-processing and post-processing programs for production code kernels were already running on LTSS.

The CRAY-1/LTSS link was developed from existing components of the LASL Integrated Computer Network (ICN). A PDP-11/40 using a modified version of existing network software was used to link the CRAY-1 to a CDC 7600. The 7600 to PDP-11 portion of the link was duplicated from existing hardware and software. The CRAY-1 to PDP-11 portion of the link involved the development of new hardware and software. The first jobs were run across the link in late November

1976, the first user codes went across in December 1976, and the CRAY-1/LTSS link was generally available to users by January 1977.

This link provided a capability for LTSS users to submit jobs to the CRAY-1 and to retrieve output from those jobs on LTSS. Utility software was written to convert between the file formats of the CRAY-1 and LTSS.

In order to ensure that a FORTRAN compiler would be available for our early production requirements, we established an effort to develop a FORTRAN capability for the CRAY-1. This approach was based upon cross-compilation: an existing FORTRAN compiler running on another computer system was modified to generate Cray Assembly Language code for the CRAY-1. This technique has provided an early FORTRAN capability for this new supercomputer in a primitive software environment. By the first quarter of 1977, the cross-compiler was successfully running on LTSS and scalar optimization efforts were initiated. The steps in compiling and executing a FORTRAN program on the CRAY-1 during this period were the following:

- Prepare the FORTRAN program, input data, and control statements on the LTSS machine.*
- Compile the FORTRAN program on LTSS using the cross-compiler.
- Send the Cray Assembly Language program, input data, and control statements across the link to the CRAY-1.
- Assemble the program and execute it on the CRAY-1.
- Send the output across the link to LTSS.
- Examine the output on LTSS.

In addition to the LTSS link and cross-compiler, an I/O library, mathematics library, support utilities, accounting procedures, and security software were developed for the CRAY-1 in order to meet early production requirements. The history of utilization of the CRAY-1 for production computing from January 1977 through August 1977 is shown in Figure 2. The unusually high level of use in June and July is due to the special demands of several LASL programs during that period. Figure 3 illustrates the current configuration of the CRAY-1 in the LASL Integrated Computer Network. Note that a PDP-11/70 running the UNIX operating system is used as a support processor in the development of software for the CRAY-1.

---

*CRAY-1 jobs are queued on three LTSS machines, but only one 7600 can communicate with the CRAY-1 at a given time.

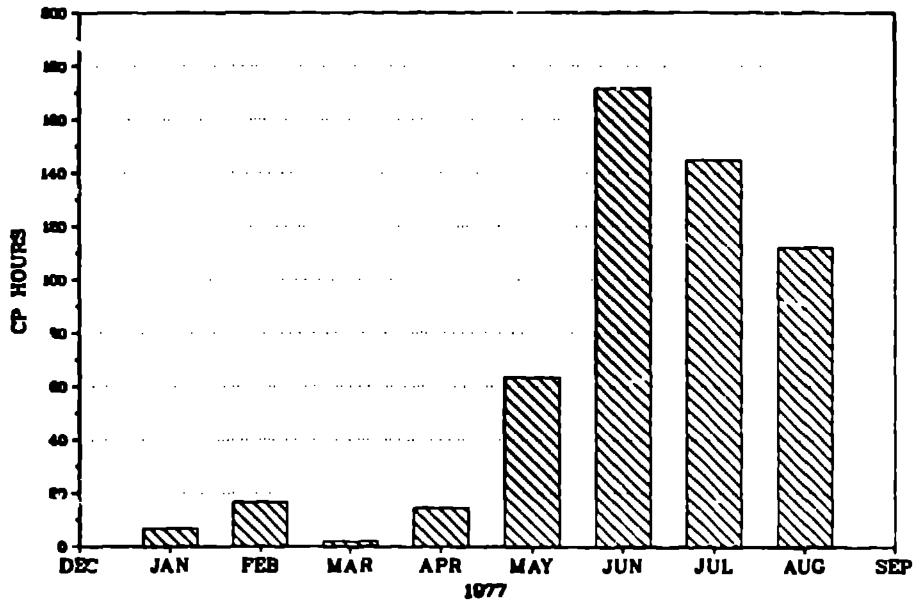

Figure 2. CRAY-1 Production Utilization

SOFTWARE DEVELOPMENT AND FUTURE PLANS

The planned status of the LASL Integrated Computer Network during 1979 is pictured in Figure 4. We are currently developing a new operating system for the CRAY-1, which we call DEMOS. We expect this operating system to provide a foundation for long-term, efficient use of the CRAY-1 at LASL. We are developing our own operating system rather than using vendor-supplied software because of the following factors: the vendor-supplied software was not available in time to meet our early production requirements; our users require an operating system tailored to their specialized computational needs for running very large codes; the LASL computing environment requires multi-level security; the CRAY-1 must be incorporated into the existing LASL Integrated Computer Network; and the CRAY-1 operating system must include the capability of communicating with an interactive graphics system through the network at very high data rates.

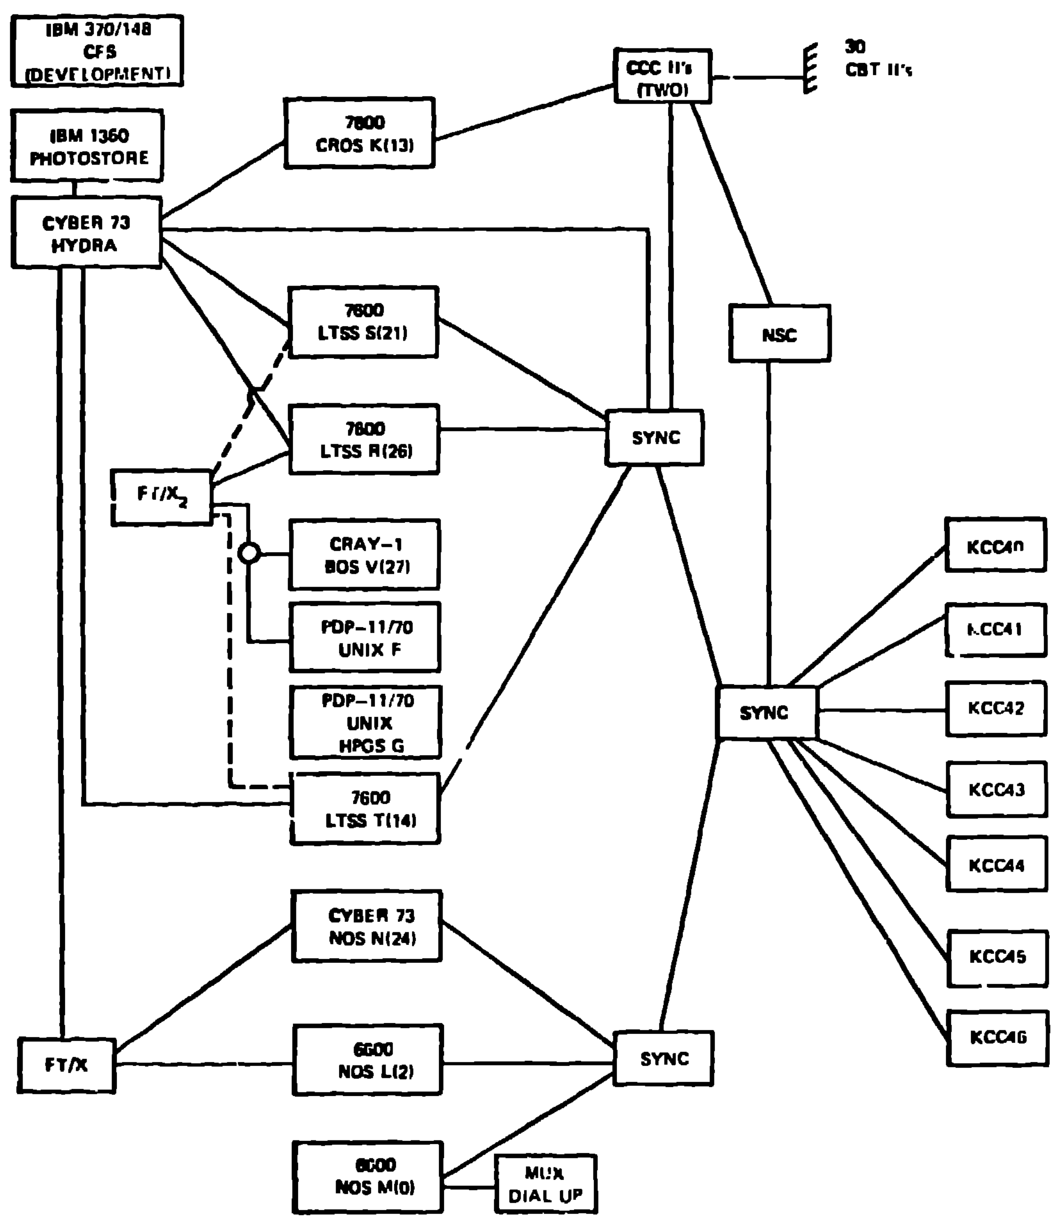

Figure ?. Current ICN Configuration

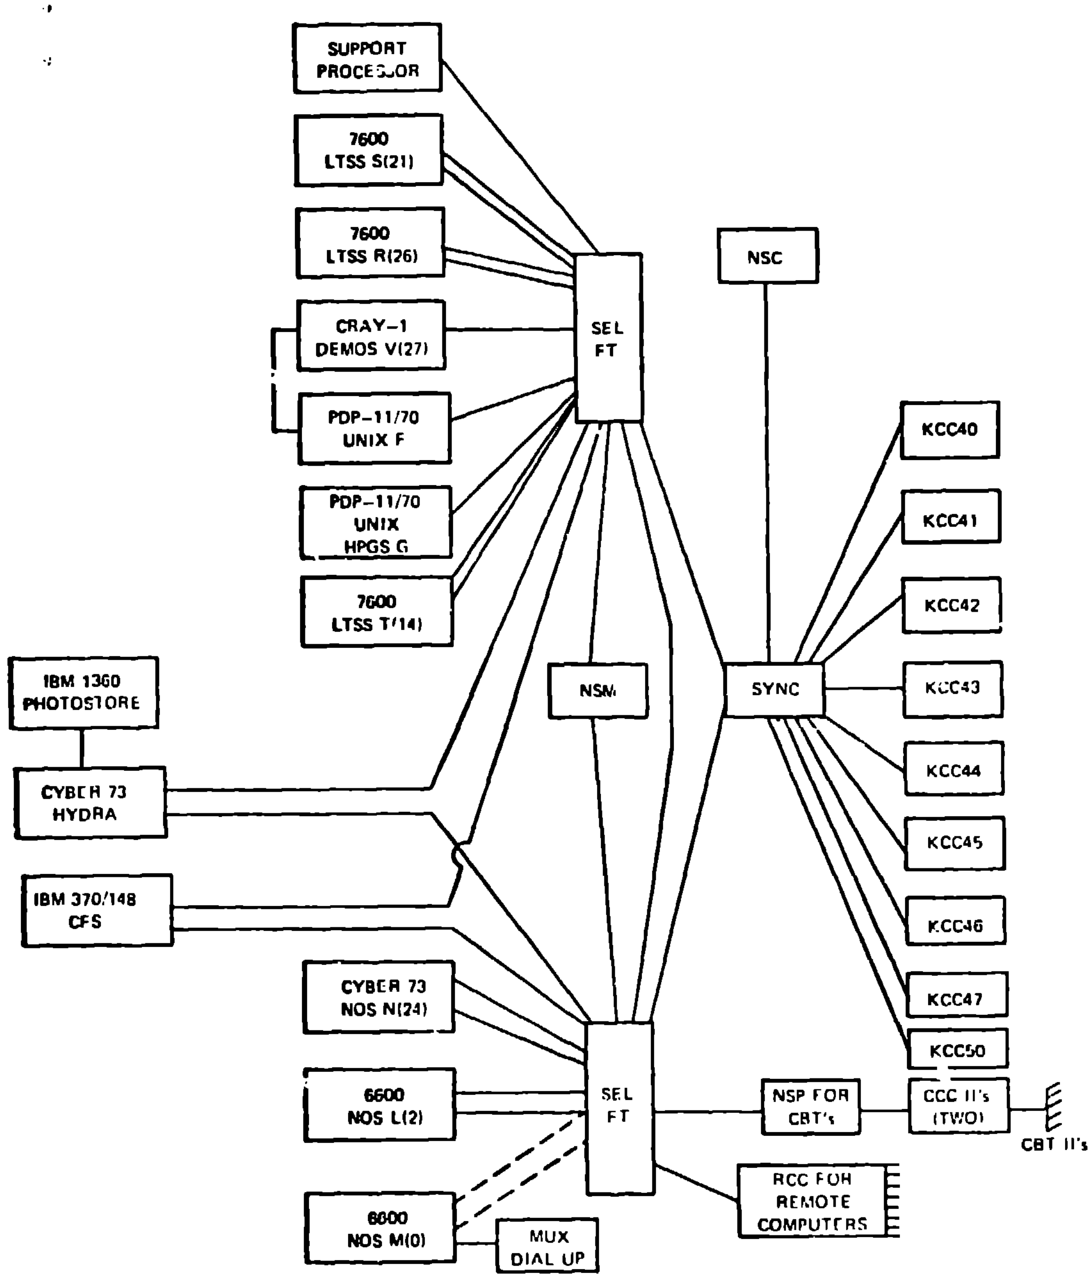

Figure 4. ICN Plans (1979)
(Pending Security Approval)

The first major components of the DEMOS operating system will be available in the first quarter of 1978, with additional components following at about six-month intervals. The major items in this development are briefly described.

- Resident File System--First Quarter 1978. The file system will provide a hierarchical directory structure compatible with the directory structure being developed for the LASL Common File System. It will include efficient buffered I/O for user programs; it will provide a resident file storage facility, which should make the CRAY-1 convenient for LASL users; and it should greatly reduce the network link traffic required to support the machine.
- Multitasking System--Second Quarter 1978. The multitasking structure will enable users to represent the complexities of a job at the command language level so that a great deal of program-to-program I/O can be avoided. This will provide for multiprogramming of user jobs in order to facilitate more efficient use of the CRAY-1, in addition to improving job turnaround.
- Network Interface--Fourth Quarter 1978. The full network interface will allow the CRAY-1 to be accessed from any of the other worker computers in the CCF network, rather than just the current selected front-end machines. The network interface will include the system utilities necessary for users to direct jobs and files to appropriate worker computers. It will also include support for mass storage systems, especially the Common File System.

The High Performance Graphics System (HPGS) consists of a PDP-11/70 with the UNIX operating system and an Evans and Sutherland Picture 2 System. It is being designed to provide a capability for highly interactive graphics involving a production code in execution on the CRAY-1. The target data transfer rate between the CRAY-1 and the HPGS (through the network) is one megabit per second. The user should be able to monitor the progress of a production code and interrupt it in order to do local processing on the PDP-11/70 (for example, rezoning a mesh if the code encounters difficulty during the course of the problem solution). The hardware for the HPGS has been delivered and the software is under development. The initial operational capability is scheduled for the summer of 1978.

We believe that certain portions of jobs that are currently run on supercomputers can be performed more cost-effectively on smaller-scale machines. Examples are the interactive text editing required to set up the input for a large production job, and the interactive graphical analysis of the output of a production job. We use the term "Support Processors" to denote the medium-scale computer systems that are being acquired to provide LASL computer users with a cost-effective alternative for the interactive work required to support supercomputers. The CRAY-1 will be supported by this type of machine in the sense that Support Processors will provide peripherals and software not directly available to users of the CRAY-1. For example, processing of magnetic tapes, printing files, and editing files for CRAY-1 programs can be done on Support Processors. In the future, we may even be able to provide a capability for interactive debugging for the CRAY-1 through Support Processors.

The primary reason for using Support Processors for some of the interactive tasks associated with supercomputers is that smaller computers are becoming increasingly cost-effective for certain kinds of tasks. In addition, distributed computing in this form gives users more control of their computing facilities; it may give them better response time for certain types of jobs; and it can lower the data communications cost. Some types of jobs that have to be executed during the day in the LASL CCF require most of the resources of a very large, fast machine such as a CDC 7600. However, we believe there are also jobs run during the day in the CCF that can be done more cost-effectively on smaller time-sharing machines. Our long-range goal is to provide those services in the CCF that we can do more cost-effectively than the general user community (specifically, production computing using large machines such as the 7600 and the CRAY-1, large mass storage facilities such as the Common File System, and certain types of graphics capabilities). At the same time, we will try to provide access for remote computers to these CCF services. With this communication capability available, users can determine the resources they need for specific jobs and can perform some tasks on smaller machines and other tasks on the large CCF computers.

ACKNOWLEDGMENT

Material provided by Forest Baskett, Tom Keller, and John Morrison has been incorporated in this report.