*C-Division Annual Review
and Operating Plan*

August 1990

Norman R. Morse

Contents

Abstract	vii
Section 1: Introduction	1
High-Performance Computing	2
Major Technical Accomplishments	4
C-Division Strategic Plan	9
Section 2: Human Resources	11
C-Division Organization	12
People Profiles	16
Developing Human Resources	27
Awards and Contests	28
Section 3: Services	31
Overview of C Division	32
Software Support	37
Communications Support	40
Direct Support	42
Section 4: Facilities and Capabilities	47
Los Alamos Integrated Communications System	48
Laboratory Data Communications Center	49
Supercomputer Upgrades	50
Central Networking Testbed	51
Network Modernization Project (NET90)	52
Computer Graphics Laboratory	55
Section 5: Leadership	57
Research Projects	58
Advanced Computing Laboratory	70
Benchmarking	72
Collaborations	74
Distributed Computing	76
Conferences and Workshops	82
Publications	86
References	89
Appendix A: Resources	91
Appendix B: Glossary	99

High-Performance Computing

A concerted effort is required for the U.S. to maintain leadership in computing. We have lost the lead in computer chip manufacturing, and not only are other nations beginning to produce very impressive supercomputers themselves, the industrial base in other nations (in particular, Japan) is beginning to apply supercomputing to their needs more extensively than is being done in the United States.

The Office of Science and Technology Policy, which is headed by the President's science adviser, Allan Bromley, has proposed a national high-performance computing (HPC) initiative to help strengthen this nation's eroding lead in computing—from aspects of producing more powerful computing tools (hardware and software) to the actual application of these tools by industry for nationally important problems.

The DOE, and Los Alamos in particular, should play a major role in this initiative.

Many problems of technical and societal importance are far too complex to be addressed adequately even with the advanced computational capabilities now available to scientists and engineers. But with concentrated and focused efforts, we will be able to continue the remarkable technological developments that started this revolution we call the computing age. We are currently at the doorstep of dramatic technological improvements that promise to change the way we do science and business, in fact the way we go about our everyday lives; but there is intense interna-

tional competition for control of the capabilities key to this revolution.

The Laboratory has been a leader in computing since the 1940s, when computers were first used to support the nuclear weapons program. Through the 1970s the Laboratory's primary computing resource was the Central Computing Facility (CCF). This facility was able to satisfy all of the computing needs of the Laboratory's scientists, engineers, and administrators. Today, technological advances have brought significant computing resources to the desktop, and the centralized model of computing is outdated and must be updated.

A task force was formed to investigate HPC at the Laboratory. The task force was called the Computing and Communications Initiative, and its goals were to evaluate the current state of computing at the Laboratory and

Today's distributed computing environment includes powerful workstations, a medium-speed network, departmental class machines, and supercomputers.

We believe that the high-performance computing environment of the future will be massively parallel, with workstations that are more powerful than today's supercomputers. High-speed, reliable networks will transmit data at 1.6 Gbits.

predict a desired future state. The current state of computing is comprised of a distributed environment similar to that shown in the figure at left. The important equipment in the environment includes powerful workstations in clients' offices, a medium-speed network (50 Mbytes) with mass storage output (among other capabilities), departmental class machines (predominantly VAX computers), and supercomputers (predominantly Crays).

Predicting the future computing environment at the Laboratory is a dangerous business. It will be influenced by our institutional goals, technological advancements, Laboratory organization, and funding level. The figure above depicts a future computing environment called the HPC model. This model reflects a shift to an essentially bimodal distribution of workstations providing the general-purpose computing capability and massively parallel supercomputers providing the top-end capability. The two computing components will require very high performance networks to integrate the capabilities into a usable whole.

The workstations in the HPC model will be more powerful than the current supercomputers and will provide "general-purpose" computing and an interface to the massively parallel supercomputers. We believe that

the departmental class computing requirement will be satisfied by a combination of workstation and supercomputer capabilities.

However, because of the mass market nature and relatively low price of the workstations, they, by themselves, will not make the Laboratory unique or more competitive.

The HPC supercomputers will be massively parallel, and the software will be predominantly UNIX, with a graphical interface. A preeminent supercomputer capability will make us unique and allow us to approach scientific problems that others cannot. However, to enhance the scientific reputation of the Laboratory, this resource must be widely available and used by a large segment of Laboratory scientists.

High-speed, reliable networking is needed to support integration of workstations and supercomputers and to allow communication among the scientific community. The HPC network will be based on HIPPI channels and will run at 1.6 Gbits. A fiber-optics plant with supporting communications gear is necessary. Work is needed in high-speed network protocols, and software must be written to integrate the parts of the system.

We foresee that supercomputer suppliers will continue to develop more and more powerful machines, extending to the trillion-

operation-per-second range (compared with the billion range of today) and beyond. These machines will be critical in addressing not only the so-called Grand Challenge problems but other important technical problems as well. High-speed networks and user-friendly interfaces will extend this powerful computing technology into offices and laboratories, and ultimately these capabilities will enter the mass market as the brains of a powerful multimedia center available to the home.

The evolution from distributed to high-performance computing environments will probably take about 5 years. We in the Computing and Communications Division are gearing up to assist our clients in this evolution, and specific projects—already under way—that move us closer to that environment are described throughout this report.

Major Technical Accomplishments

The major accomplishments included here describe some of the areas where we directed resources in FY 1989. These and other projects required the skill, cooperation, and significant effort of people throughout the division.

LDCC

Our most visible accomplishment last year was the opening of the Laboratory's Data Communications Center (LDCC), which was completed on schedule and \$670,000 under budget! Important dignitaries joined the enthusiastic C-Division staff at dedication ceremonies in August. Following

the ribbon-cutting, everyone was allowed inside for a tour of the 95,000 square foot facility, which will be home for the Laboratory's telephone system and up to 10 future supercomputers. Employees had an opportunity to bring their families through the new building the following day, when a "Family Day" activity was held before the facility went behind the secure fence. Soon after, employees began moving into this much-needed space. More details about the LDCC are included in the Facilities and Capabilities Section of this report.

Important dignitaries who played instrumental roles in getting the LDCC funded, designed, and built enjoyed late-afternoon dedication ceremonies before touring the building. From left to right are Sig Hecker, Laboratory Director; Bruce Twining, area manager of the Department of Energy's Albuquerque operations office; Norman Morse, C-Division Leader; Senator Pete Domenici; and Nicholas Nagy, technical manager for the LDCC. In a speech before the ribbon cutting, Morse praised the Albuquerque architectural firm of Flatow, Moore, Brian, Shaffer, McCabe, Inc., and said the building's design may become the prototype for future supercomputing centers.

Research

Computational science has become critical to the performance of the Laboratory's mission, from nuclear weapons design to basic research. Nationally, the importance of scientific computing has also become well recognized. In order to meet the grand challenges, a revolution is taking place in advanced computation and in the environment in which it must take place. Parallel computing, ultra-high-speed networks, visualization, and distributed supercomputing are just some of the technology-based issues that must be successfully dealt with to realize the potential that we now see clearly. These technological opportunities must be woven into the fabric of computational science, here at the Laboratory and at a national level. Los Alamos is involved in several of the most ambitious research programs in computational science. Among these projects are the Advanced Computing Laboratory (ACL); Computational Modeling on the Connection Machine from Thinking Machines Corporation; the NSF Science and Technology Center for Research on Parallel Computation; the DARPA/NSF/DOE Research Internet Gateway Project; the Gigabit National Testbed; the DOE/Energy Research Applied Mathematical Sciences program; and the development of the High-Performance Parallel Interface (now called the HIPPI, it was formerly the HSC, high-speed channel). Details about these and other projects are included throughout this document.

The establishment of the ACL here at the Laboratory is a great source of pride and enthusiasm not only for C Division but also for others throughout the computational science community. The formation of the ACL follows a distinguished history at Los Alamos of research into advanced computing issues such as parallel

computation, visualization, and numerical algorithms. The ACL is intended to provide both the facilities (e.g., computing hardware, communications, space) and the environment to catalyze interdisciplinary research among computer scientists, applied mathematicians, scientists, and engineers. The goals of the ACL are to facilitate Laboratory research in advanced computing, to provide a focal point for applications-driven, interdisciplinary research; to provide an open environment for Laboratory/external interactions and technology transfer; and to provide a bridge to other centers of excellence.

The ACL currently occupies temporary quarters on the first floor of Building SM-123 while awaiting the completion of a permanent facility, a two-story building now under construction. Because the ACL includes innovative and experimental architectures, it encourages researchers and vendors to work on problems that might not be attempted in a production environment. Use of experimental architectures is risky because they may never become commercially successful and hardware and software support could disappear. Yet, if the value of such machines is proven, they potentially could migrate into a production environment where they would enjoy more widespread use.

Despite the risks of innovative and experimental architectures, researchers at Los Alamos are successfully using them in the ACL to solve scientifically interesting problems. Ever since Thinking Machines Corporation delivered a Connection Machine 2 to the ACL in January 1989, this machine has attracted a number of Los Alamos scientists and engineers into research in parallel computation. The peak rate of the CM-2 is roughly an order of magnitude greater than that of the

Los Alamos benchmarks, working through the ACL, did early performance tests of the IBM RISC/6000 workstations announced in February 1990. The report published by the team has become a definitive standard regarding the machine [1]. The ACL is continuing a beta agreement with IBM to further evaluate these products.

CRAY Y-MP (though it is much more difficult to achieve). This computer has 65,536 processors, 2 Gbytes of memory, four framebuffers and monitors, and a 20-Gbyte data storage system. Early in 1990 Thinking Machines upgraded this machine to provide 64-bit floating-point arithmetic; a prototype HIPPI interface is now installed.

Applications work on the CM-2 is in a wide variety of areas: fluid dynamics, quantum chromodynamics, condensed matter physics, plasma physics, transport phenomena, neural nets, and image processing. Of particular interest is work in fluid dynamics utilizing free Lagrangian methods to approximate the hydrodynamics. Problems solved using these techniques, developed at Los Alamos for use on the Connection Machine, exceed the capabilities of the CRAY Y-MPs currently installed in the CCF. These calculations have been so successful that we were able to negotiate with Thinking Machines for a second CM-2; it is now located in the CCF's Secure partition where

it supports the production environment.

Networking

Security, reliability, and improved performance remain priority projects for the ICN and the networks connected to it. Many of the computer networking projects discussed in last year's plan have been completed. All of the ICN front-end switches have been replaced with new hardware and software. All of the XNET-to-ICN interfaces have been replaced with new hardware and software, which has more than doubled the throughput of the XNET systems. (XNET, the Extended Network Access System, provides ICN access for remote computer sites.) Most of the DECnet routers have been replaced with new higher-speed and larger-memory routers, which will migrate to become ISO routers over the next few years. Additional XNET gateways have been added—providing hardware redundancy in all partitions; work is under way to make use of this redundancy in ways more transparent to the end user. Additional

tuning and enhancement continue in each of the above projects.

C Division has a policy of "buy don't build" where vendor products can satisfy the needs of the Laboratory. In order that vendor products will be usable in the Laboratory environment, we attempt to influence the direction of vendors in several ways. One of the most crucial needs we see in the near future is very high speed packet switches and data links.

The Laboratory spearheaded the development of the HIPPI, which operates at 800 Mbit/s. The HIPPI, which is becoming an ANSI standard, is being implemented by many vendors including Cray Research, IBM, Thinking Machines, Network Systems Corp., Ultra Network Technologies, Tektronix, Ancor Communications, Aptec Computer Systems, Digital Equipment Corp., IPITEK, and Maximum Strategy, Inc.

The HIPPI is used in a Los Alamos-developed network architecture called CP* (named for Cross Point Star). This architecture is based on high-speed, nonblocking crossbar switches. The switches were designed by Los Alamos and we are encouraging commercial development of this capability through a collaboration with Digital Equipment Corporation. We have completed the development and checkout of four HIPPI interface cards and two 16-channel CP* switches. We plan to have a testbed in operation using this equipment during FY 1990.

The performance of the XNET gateways has been improved so far by about a factor of 2 because of improved data links to the ICN. Higher speed local area networks of the future (such as the 100-Mbit/s FDDI) will provide an improved speed match to the new XNET-to-ICN data links and should result in another factor of 2 to 5 improvement in performance.

Computer security is a vital element in the Los Alamos com-

puting environment, and the computer network is a crucial element in providing the appropriate security. We are investigating security architectures and protocols that will provide the appropriate protections in the fully distributed computing environment of the future and have proposed changes to some ICN protocols to begin to use these approaches. In addition, we are building online, real-time tools to attempt to detect attacks on the security of the ICN. We have developed a prototype intrusion detection system and are in the process of implementing the functionality of the prototype in the production ICN. The prototype system is based on profiles of user and network component behavior.

We have made about 50 changes to the ICN switch code over the past year in order to improve both performance and reliability. In addition, we have made major advances in understanding the performance and reliability of the ICN. As a result of a study of the core of the network, we have changed all Cray FT links to "four-word mode," increasing throughput by about 40%.

Likewise, we are studying the reliability of the ICN and are studying methods for providing alternate route and improved alternate node capability within the ICN. System failure rates for SYNCs have dropped from 1 or 2 per day to less than 1 per month. System failure rates for FTs have also dropped to less than 1 per month average. This implies a failure rate of 1 per year per FT or SYNC and includes software as well as hardware failures.

The ICN continues to evolve rapidly; during FY 1989, 56 changes were made to the ICN configuration in 24 separate releases. On average, one or more changes were made to the ICN configuration every two weeks.

The extended network (the network of computers outside the CCF) grew at a very rapid pace: the number of distributed processors grew from about 900 to 1450 over the past year.

Network Modeling and Measurements

We developed a discrete-event simulation model of the critical data movement components of the ICN, specifically the Cray supercomputer file transport modules, the network's file transport switching nodes, and CFS file handling and storage modules. The model is helping us predict the effects of new hardware, protocol changes, configuration changes, and software enhancements as well as identify potential bottlenecks in the data movement paths. Tuning of the model to more accurately simulate the components will continue.

We measured data movement through the high activity areas of the ICN, specifically data moving to and from the Cray supercomputers, the Common File System (CFS), the Print and Graphics Express Station (PAGES), and XNET. These performance measurements are helping us understand data movement through the network, which will help us identify potential network bottlenecks and determine trends and rates of growth. This increased understanding will help us predict future network requirements and assist in planning for the future.

We began measuring file transmission rates for UNICOS, the UNIX operating system for the Cray computers, and are attempting to identify the components that need tuning to increase performance.

Wideband Communications Network

In addition to providing computer networks for Los Alamos, C Division is the lead

Laboratory in developing the Wideband Communications Network (WBCN) for the Department of Energy. This network will provide secure, classified computer communication between sites in the Nuclear Weapons Complex. During FY 1989 Phase 2, a major release of the WBCN software, was put in production across the complex. During FY 1990 the software to allow secure file transfers between the Common File Systems at each site was completed, approved, and brought into production at selected sites for testing.

New Computing Capabilities

Many people throughout the division worked together to install two new CRAY Y-MP/832 supercomputing systems, each with eight processors and 32 Mwords of memory, along with two Solid-state Storage Devices (SSDs). One of these machines was later upgraded to 128 Mwords of memory and relocated to the LDCC. Software, including CFTLIB, the system library that runs on the Crays, was converted, enhanced, and upgraded to run on the new computers. The effort expended was well worthwhile, and the transition to the Y-MPs went smoothly for both the users and for C Division. In addition, we installed an IOS on one of the Crays for UNICOS, and we installed an additional X-MP/416.

Operating Systems

We converted CTSS, the Cray Timesharing System, to run on the CRAY Y-MPs, which included the following enhancements: (1) upgrading the IOS channel driver to the latest release for Cray Research, Inc., (2) installing a DD40 disk driver, (3) enhancing the system to use 32 Mwords of main memory, (4) modifying the system tables to handle the increased load, and (5) enhancing security.

UNICOS

In the spring of 1989 we made the decision to run the UNIX operating system on our future supercomputers. At the same time we began a 2-year project to convert our existing Cray supercomputers to UNICOS, the Cray Research variant of UNIX. In order to involve the user community in an active way in planning for this project, a Transition Advisory Group (TAG) was formed with representatives from the major user divisions, C-Division staff, and personnel from Cray Research. This group developed and published a detailed UNICOS Migration Workplan to guide the conversion process. The TAG continues to meet on a periodic basis to monitor the project and discuss its evolution.

The project has been under way for over a year. Significant progress has been made toward providing a production quality UNICOS environment to which our users can migrate. We plan to begin the migration in the fall of 1990 in the Open and National Security partitions. A CRAY Y-MP8/264 running UNICOS will be placed in each partition, and users will begin porting their codes. Spring 1991 we plan to install a UNICOS machine in the Secure partition. This will allow our classified code users to begin their migration.

During the past year we developed a UNICOS version of CFTLIB, Los Alamos's run-time Fortran library. This library bridges many CTSS CFTLIB functionalities to UNICOS capabilities. It also supports Fortran extensions normally not available on UNICOS, such as user interfaces for interrupt handling, "soft" CPU time limits, real-time process status inquiry, byte-addressable random synchronous and asynchronous I/O packages, and standardized error handling.

To help improve the friendliness of UNIX, and to put more

power and flexibility into the user's hands, we have developed and implemented a "next generation" shell called xCL, which will allow users to be even more productive in our new UNIX environment. Among the unprecedented capabilities offered by xCL are a custom programming interface (C-like for C programmers and Fortran-like for Fortran programmers); a built-in, interactive, symbolic script debugger; and drastically improved performance over the existing UNIX shells.

The major graphics libraries such as CGS, as well as the graphics postprocessors, are now available on UNICOS. All of the math libraries that C Division currently supports have been ported to UNICOS. The widely used text editor, FRED, is also running on UNICOS.

In the operating systems area, we have installed UNICOS Version 5.1. We are running it in secure mode in anticipation of using UNICOS in the Secure partition. We have made modifications to UNICOS to utilize our local accounting algorithms, and we have integrated the UNICOS NQS batch processing into the Los Alamos FOCUS Production Control System. In parallel with all of this activity, our consulting staff is coming up to speed on UNIX. In addition, we are preparing migration documentation and training classes in order to help users make the transition.

Work on a new Los Alamos debugger (called LDB) for supercomputing codes is proceeding on schedule. Phase 1, which consists of an important set of basic debugging commands, is complete. Phase 2, which is an implementation of an X Windows interface, IGP graphics, and additional commands, is under way and scheduled for completion in 1991. In addition to its debugging chores, LDB will serve as a vehicle for research into the

debugging of parallel processed algorithms.

User-Visible Software Changes

During this past year we released software that generates and translates raster graphics images. This software complements our vector graphics capabilities. We have adopted the international metafile standard, the Computer Graphics Metafile (CGM), as the vehicle for transmitting raster data. A user can generate raster images directly from a program by using a subroutine library. Users may also take images generated on raster terminals such as the Sun, Iris, or Stardent, and translate these into CGMs. Our central output node, PAGES, will accept CGMs and generate either film or paper hardcopy output.

Users of Sun workstations who are accustomed to a Tektronix terminal environment and have programs that generate Tektronix commands have a Tektronix emulator that allows them to continue working in that mode. This software works with the SunView window system and is being converted to work in the X Windows environment.

Common File System

CFS is a centralized data storage system for the Los Alamos ICN, where it has been in use since May 1979. Performance and reliability of CFS was further improved. The control processor hardware was upgraded, additional disks were acquired to store more of the active files online, and additional tape equipment was obtained. CFS software modifications were made to improve workstation file transmission performance and reliability. A reporting capability was implemented to provide detailed information about CFS workload, performance, and reliability.

A comprehensive study of Los Alamos computer data storage needs in the 1990s was conducted. Surveys and interviews supplied much of the data to support the recommendations that were made, and pros and cons of CFS were reviewed. Numerous existing and planned file systems were also reviewed. The study concluded that CFS is a valuable resource with many strengths that should be preserved for the future. None of the other systems reviewed would perform as well as CFS in the Los Alamos environment. It was recommended that CFS be revamped for the new Laboratory environment of high-performance workstations and supercomputers running UNIX.

With the commercialization of CFS by General Atomics, 25 computing sites are now using the CFS software.

PAGES

We have enhanced the efficiency of the PAGES operations by integrating new equipment. This includes black-and-white printers for document processing (text and graphics) as well as standard ASCII output. These medium-speed printers (92 pages per minute) offer online stapling of

jobs up to 50 pages and a significant reduction in maintenance costs and downtime. Two printers are in operation now, and plans call for purchasing two more.

Raster services using the CGM on 35-mm and color paper output were put in production and have proved to be important to many users.

As part of our charter to keep up with the latest technology, we have

- Evaluated a small raster film recorder, one of a new breed of film recorders that we expect will provide color film services at much higher quality and a fraction of the maintenance cost. We are purchasing a pilot unit based on the promising results of our evaluation.
- Purchased a Canon color copier/printer that offers very high quality color images on plain paper and transparency material. This device, which upgrades the capability for color output, is intended to replace the low-quality color printers in operation today.

Two new high-speed Kodak printers have been integrated into PAGES. Plans include purchasing two more.

C-Division Strategic Plan

A strategic plan is a document that describes the purpose of an organization and the future it wishes to attain. Key elements of a strategic plan include

- Credo—the organization's philosophy or uppermost beliefs.
- Mission—its purpose or reason for being.
- Goals—the "desired futures" an organization wishes to achieve.
- Objectives—measures whose achievement moves the organization closer to its goals.
- Projects—well-specified tasks whose achievement will accomplish the objectives.

Credo

The management of C Division believes we have primary responsibilities in the following areas:

- Pursuit of excellence
- Value of the individual
- Commitment to service
- Commitment to scientific leadership.

Mission

The mission of the Computing and Communications Division is to maximize the productivity of our clients and ourselves through leadership in computing and communications. We will accom-

plish this by balancing our efforts in service, research, and development. We recognize the Laboratory's responsibility as a national resource and shall lead in areas crucial to our problem-solving obligations. Key aspects of our vision include

- An environment that fosters individual initiative and creativity.
- A world-class supercomputing environment supporting a geographically dispersed scientific and engineering community.
- A single integrated communications utility for the entire Laboratory.

Goals

Our goals are in the following four areas: human resources, service, facilities and capabilities, and leadership. Each goal is defined below.

Human Resources Goal

Attract and retain the highest quality staff by valuing their creativity, expertise, and dedica-

tion; nurturing their technical growth; and creating an environment that fosters individual participation in the development of new ideas and in the decision-making process.

Service Goal

Provide services and capabilities that maximize productivity by emphasizing the importance of supporting clients, accepting ownership of problems, and involving clients in projects.

Facilities and Capabilities Goal

Provide world-class scientific computing, network, and communications facilities.

Leadership Goal

Maintain national leadership in computational science through outstanding research and development focusing on advanced computational environments and applied mathematics.

C-Division Organization

A brief description of the division office and the groups is given below. Although the title of each group describes its major function, many projects span group boundaries. There are 350 people who staff the division in its division office and nine groups.

Division Office

Serving as a focal point for the division, the C-Division Office provides managerial and technical guidance to the division's various groups. Reporting to the division leader and two deputy division leaders are a technical and administrative support staff, as well as three program managers. This management team is responsible for coordinating major activities within the division and ensuring the success and continuity of major projects.

Computer Operations

The Computer Operations Group (C-1) monitors and controls the ICN resources in the CCF and LDCC and provides computer operations services 24 hours a day, 7 days a week with a staff of 57 employees. The ICN is available for computing the year round except for Christmas Day and for facility maintenance periods of short duration that normally occur during the annual Laboratory closing that takes place between Christmas and New Year's.

The Computer Operations Group's major responsibilities include the reliable and timely delivery of computing services including oversight of hardware and facilities projects, operator training and development, output generation and distribution, monitor and control of worker

computers and major ICN resources, CCF and LDCC security, telephone dispatching, computer systems status information, and computer vendor liaison. The Computer Operations Group responds to user inquiries and problems and makes every effort to resolve user problems at the operations level.

Computer Documentation

The Computer Documentation Group (C-2), with a staff of 20, provides computing documentation and computing education and supports document processing systems. They develop and maintain user documentation for locally developed computing capabilities and to supplement vendor documentation, and they provide the monthly *Computing and Communications Division News* [2] and *ICN Change Bulletin* [3] to keep all computing users informed about computing topics of current interest. They distribute computing information in printed and online forms.

The staff of C-2 develops video- and computer-based courses. They select courses, including vendor self-paced courses and satellite courses and conferences from major universities around the country. Courses and conferences are broadcast over LABNET (the Los Alamos Broadband Network) and made available in a variety of media. C-2 supports text formatters, such as TeX and TROFF, merging of text and graphics, and local printers as tools in the ICN document processing system. C-2 also offers services on a Form-B basis, including high-quality computer

graphics, computer-based document and data preparation, file conversion, and optical character reading.

Computer Research

During 1989 the Computer Research and Applications Group was reorganized into two smaller, more manageable groups. Under the new organization, Group C-3 is now the Computer Research Group. Staff from the Applied Research Section of Group C-10 joined the existing research sections in C-3 to form the new group. The Computer Research Group performs research applicable to problems of interest to the Laboratory and conducts performance analysis of computer systems of current and potential interest to the Laboratory.

Current efforts include general computational applied mathematics, signal processing, mathematical algorithms and libraries, experimental computing technology, computer performance measurement, expert systems, neural networks, speech recognition, advanced computing architectures, user tools, languages, pattern recognition, and image processing. The group, which has a staff of 35, extensively collaborates with universities, other laboratories, and industrial partners.

Communications

The Communications Group (C-4) is responsible for Laboratory communications. With a staff of 54, Group C-4 designs, builds, installs, and maintains communications systems of all kinds for use throughout the Laboratory and between the Laboratory and external sites. Group C-4 provides Laboratory voice, data, and video communications. A few of their responsibilities include designing and implementing the Los Alamos Integrated Communications System (LAICS); upgrading and maintaining transmission

C-Division organization.

systems including LABNET, outside cable plant, and radio; coordinating communications vendor activities for the Laboratory; administering and automating the communications recharge system; and designing, implementing, and maintaining communications systems in support of distributed computing facilities and external clients.

Computer Network Engineering

The design, development, and support of the ICN's architecture are the primary responsibilities of the Computer Network Engineering Group (C-5). Members of the group, which employs 42 people, design, develop, and maintain hardware and software systems that provide the foundation of the ICN. These systems include high-speed parallel interfaces, high-

speed packet switches, terminal concentrators, gateways to external networks, and authentication and access control servers.

In addition, the group develops specialized systems to support high-speed graphics, provides oversight of the security aspects of the ICN and leads projects for local area network development. The group also does research and advanced development in high-speed networking systems.

Los Alamos is the lead laboratory in developing the WBCN for the DOE. Group C-5 is responsible for development of the network and the VAX host system for this project.

Computer Graphics

The Computer Graphics Group (C-6) provides computer graphics software and systems support for the ICN. With a staff of 18, the group designs, develops, and integrates computer graphics software libraries and utilities for ICN systems and supports PAGES, the ICN's output node. While these systems provide the foundation of computer graphics support in the group, other developments in computing have expanded the group's role in other computer graphics-related areas.

The group has become more involved in the uses of scientific visualization systems for helping users solve their problems. Scientific visualization capabilities, from personal computers to high-end graphical supercomputers in the ACL, are being used to develop and demonstrate solutions to user problems.

Integration of new technology is critical for the Laboratory's future use of computer graphics. Group members are involved with technologies including desktop gray-scale and color scanning, manipulation, and output to high-quality printers; the use of recording video for visualization presentation purposes; multimedia environments including graphics, images, voice, and video; and use of networked window systems for distributed computing environments.

The group's major projects include development and support of computer graphics software, development and integration of

raster graphics capabilities, continued development and enhancement of the PAGES system, evaluation of computer graphics software and related technology, and the development and use of scientific visualization capabilities.

Computer Systems

The Computer Systems Group (C-8) provides systems-level support for the ICN by evaluating, developing, enhancing, and maintaining supercomputer operating systems, archival storage systems (CFS), scientific workstation systems and their networking environments, VMS and UNIX general-purpose computing systems, and production control capabilities. C-8's staff of 36 ensures that users have a stable, reliable, and productive computing environment by providing system development and system management support for CTSS, UNICOS, UNIX, VMS, and NOS systems located in the CCF.

As part of this responsibility, C-8 installs new services in the operating systems, handles changes and upgrades to network protocols from the operating system end, and makes any operating system changes that are required for interaction with other network components. The group is responsible for maintaining and enhancing CFS so it continues to operate at a high level of performance, reliability, availability, and user satisfaction and is responsive to changing user needs. The group is also responsible for managing, developing, and

maintaining the FOCUS production system.

C-8 leads projects in researching and evaluating potential systems for the future in several areas including distributed computing, operating systems, user interfaces, and network performance.

Computer Applications

The Computer Applications Group (C-9) is a new group created in 1989 as part of a reorganization of the old Computer Research and Applications Group. With a staff of 29, the Computer Applications Group provides direct computing support services to the Laboratory. About two-thirds of the group are programmers, software engineers, and database engineers; the other third are targeted for local area network administration.

All of the people within the group work for clients around the Laboratory through "rental" agreements. Specializations within the group include user software interfaces design and advanced information management systems.

Also within the group is the Advanced Information Management Systems (AIMS) Laboratory, which is focused toward creating, developing, and applying tools that help manage large amounts of multimedia information.

Computer User Services

In C Division we often speak of interfaces—communication interfaces, system interfaces, and network interfaces—but from the user's perspective, the most important interface is the human interface (those points at which the user routinely interacts with C-Division software and staff). Group C-10's staff of 30 is responsible for developing and maintaining much of this human interface.

We support a major part of the programming interface employed by users to develop applications programs like the large design codes used at Los Alamos. This interface includes compilers, loaders, run-time libraries, debuggers, utilities, and shells. This software forms the "virtual" supercomputer with which most users interact.

The Consulting Office answers questions and resolves ICN-related problems. The user community requires access to quality educational services; C-10

provides hands-on classroom experience through locally developed classes and vendor classes presented on-site.

Our day-to-day contact with users often leads to fruitful collaborative research projects, including debugging parallel processed code, developing the run-time library CFTLIB and the UNIX shell xCL, the 3-Dimensional Sound Project, and exploring the possible link between supercomputer memory errors and violent astrophysical phenomena.

To help improve the user friendliness of UNIX, and to put more power and flexibility into the user's hands, we have developed and implemented a "next generation" shell called xCL,

which will allow users to be even more productive in our new UNIX environment.

Work on a new Los Alamos debugger (LDB) for supercomputing codes continues on schedule. Phase 1, an important set of basic debugging commands, is complete. Phase 2, an implementation of the X Windows interface, IGP graphics, and additional commands, is under way and scheduled for completion in 1991.

People Profiles

One of the best ways to know an organization is through its people, and C Division is fortunate to have many talented people on its staff. To provide a view of

the division from the staff's perspective, we profiled 12 employees about their work, challenges, and accomplishments.

John Blaylock

John Blaylock leads the Programming Environments Section in the Computer User Services Group. John's section is responsible for developing and maintaining the runtime software and utilities for the Laboratory's supercomputers. John explains, "Most users do not deal with the machine hardware and operating system directly. Instead, users work with a virtual image of the supercomputer created by the operating system utilities and the runtime software."

In addition to the internal Laboratory users, the Programming Environments Section supports many of the external users, such as the Defense Nuclear Agency (DNA), the Nuclear Weapons Complex Network (NWCNET), and the General Dynamics contractors. "I think of all of our users as customers. I like the business connotation because we must satisfy our

customers in order to succeed," says John, who is adept at discerning these needs because he was once a customer of C-Division's supercomputing services himself.

Before coming to Los Alamos, John was manager of the Applications Programming Group at the Florida State University Computing Center. He was hired by the Laboratory in 1979 and worked for 9 years in the Thermonuclear Design (X-2) and Computational Physics (X-7) groups on the development of large design codes. John considers his work on these supercomputing applications to be one of his major accomplishments. In June 1988 John began working in C Division to help provide services to the Laboratory's many supercomputing users.

John also is the coordinator for the UNICOS Migration Project. As coordinator, he serves as Chairman of the UNICOS

Workplan Group, a small group working out the day-to-day details of the migration to UNICOS, and of the Transition Advisory Group, which is a relatively large group with representatives from all of the major user groups for determining the issues and concerns for the change to UNICOS. John likes these challenges, "For me, this is an ideal responsibility because of my experience both as a user and as a software developer." John adds, "We are very enthusiastic about this project and what it portends for the future of our user community."

What John finds most challenging is the effort to manage his time. "But I have never been surrounded by more talented, motivated, and hard-working people than I am in C Division. The work environment is rewarding and conducive to doing good work. I characterize it as being family-oriented because we have frank and open interactions among all of the groups. Even when you lose a technical battle, you feel that your side was heard and considered, and now everyone will pull together to implement whatever was decided."

Marge Devaney

Marge Devaney began her career at the Laboratory in October 1951. She had just graduated from the University of Denver and did not realize what adventures awaited her in Los Alamos. She was a member of the group in the Theoretical Division that was responsible for building and programming the Maniac I, which started running in March 1952. "The Maniac had twice the memory of most other machines in the world—a grand total of 5 kbytes. We programmed everything in absolute machine language and did runtime editing by shorting circuits with clip leads and reading lights on the machine," says Marge who is delighted about the advances that programmers now take for granted.

Marge's group was one of the founding groups of the new C Division. "I worked on weapons programs for 10 years, on editors for over 10 years, and then on operating systems and terminal interfaces for about 5 years. For the past 12 years, I have been working on the Common File System, or CFS, in the Data Storage Systems Section of the Computer Systems Group," Marge says.

Marge is a designer and programmer for CFS, the centralized file management system used for storing and transferring data

among the many machines in the ICN at Los Alamos. CFS has over 4 million files with almost 14 terabytes of data and transfers 100 gigabytes of data per day. "It has been fun to be involved with a system that is an important and heavily used part of the computer network," says Marge.

"The CFS growth rate for this coming year is predicted to be over 6 terabytes of data—almost half as much in 1 year as we

accumulated in the previous 10 years," Marge explains. The challenge will be to improve, reshape, and modify CFS to respond to the changes in the computing environment as more and more computing power is added to the network, as the computers served by CFS vary ever more widely from workstations to supercomputers, as the computer center moves into the UNIX/TCP/IP world, and as graphical user interfaces come into demand. "It should be an interesting year," Marge adds.

CFS has been installed in several of the United States supercomputing centers, the medium-range weather center in England, and quite a few DOE sites. CFS is now being marketed under the name DataTree by a division of General Atomics called Distributed Computing Solutions, or DISCOS. "I wish them every success," says Marge, "and hope we can work together to improve CFS."

Jim Frybarger

Jim Frybarger was part of Control Data Corporation's (CDC's) maintenance organization for about 20 years. Since 1972 he worked at the Laboratory representing CDC, first as a computer technician and later as the engineer in charge of maintenance for the CDC machines. When the

Cray Research computers began replacing the CDC machines at Los Alamos, Jim went to work for C Division. He was hired by the Computer Operations Group as the Laboratory's liaison with the supercomputer vendors. Jim explains, "The Operations Group works closely with Engineering,

the Computer Operating Systems Group, and the vendors to install new equipment, to schedule maintenance, and to diagnose and remedy any hardware problems."

Site preparation is a major part of the installation of the large computers. "The site prep includes not only constructing the building or area, but also determining where and how the machine is going to sit on the floor; how much power, plumbing, and chilled water it needs; how these will be supplied to the machine; and how the machine will be linked to other equipment and resources in the ICN," says Jim.

In the new LDCC, each Cray supercomputer will be supported by 70 tons of refrigerator-type compressors housed on the floor below the computer room. Fifty tons are for the central processing unit and the input/output control units, and 20 tons are for the Solid-state Storage Device.

Jim and his co-workers are currently discussing how to coordinate the upgrade of the Laboratory's CRAY X-MP computers to the latest CRAY Y-MP computers. "For any type of upgrade, maintenance, or relocation, we try not to interrupt the users. We keep our activities as transparent as possible. This philosophy has become increasingly important as the number of users outside the Laboratory has grown," Jim explains.

Jim looked forward to the installation of the first supercomputers into the LDCC. "A lot of effort was put into the design of this building so that the support equipment for the computers would be exactly where it is needed and the environmental conditions are just right," says Jim. This means that the temperature and humidity are stable, and the chilled water and power are readily available with full backup capability. "The installations should be easier and less expen-

sive, and the equipment should actually run better and be more reliable," Jim anticipates.

Jim has learned from his many years of experience that they cannot predict the next challenge.

"Everything changes so rapidly that you simply cannot assume anything. We just have to keep current with the latest technology."

Judy Gonzales

What, your communications equipment isn't working? Who ya' gonna call?

Chances are that you will first explain your problem to Judy Gonzales. Judy is on the front line, so to speak, of the Communications Group (C-4). It is Judy's responsibility to transfer users to the correct person for help. Judy explains, "Our group handles all of the Laboratory's voice and data communications problems. I often ask questions to help analyze the problem so that I know where to direct the caller. This requires patience because people are often frustrated when their equipment is not working properly."

To handle more than 100 calls per day, Judy and her co-worker must know about the many group projects. "We get so many calls about so many different things," says Judy. "Because we also support external customers, our engineers and technicians must be

ready to travel on moment's notice to places such as Washington D.C. We cannot always plan their travel in advance but must fit it into our already hectic schedule." Judy also prepares the group's bimonthly Project Status Report. "It takes a lot of concentration to put the report together and to edit it for a consistent writing style."

Judy was promoted recently to administrative support specialist. She assists Group OS-8 in enrolling people into the palm reader database at the entrance to the LDCC. Palm readers allow people to enter and leave the secure area through an unmanned turnstile; once the individual's hand print and badge information are entered into a database, that individual can use the turnstile to enter and

leave the area, even during weekends and evenings, when a guard normally isn't on duty. Judy's work area is located at the new reception area in the LDCC. She also processes the authorizations, escort papers, and badge exchanges for access to the LDCC. Judy is the person to check with if you would like to schedule the new Al McKnight Conference Room, which is becoming very popular for conferences because of its advanced presentation equipment and unique room design.

Making a to-do-list enables Judy to keep track of all the various work given to her throughout the day, not letting anything fall through the cracks. "I write down each job and mark

it off when it is completed, which makes me feel like I'm making progress. I do like the pace of my work because I find that a little stress helps me to prioritize and organize."

Judy enjoys the people in her group because they are positive thinkers and a great group of people to work with. She thoroughly enjoys working with a Macintosh computer. "Everyone was so helpful when I got the computer. They suggested the software I needed and helped me learn to use it." Judy has taken classes for the Macintosh at the Training Center and at UNM-Los Alamos and feels that the opportunity to take classes is another reason she enjoys her job.

Flavio Gurule

When Flavio Gurule joined the Laboratory 16 years ago, he worked at the Van de Graaff accelerator in the Physics Division, where he helped run tests on the neutron's time of flight along with other experiments. After receiving his MS degree in electrical engineering and computer science, he worked on an environmental research project in the Systems Analysis and Assessment Division, which included collecting samples of air pollutants in several states. To analyze the data, he had to write several computer programs to display the results on cartographic maps

produced using the supercomputers in the Laboratory's CCF.

Flavio was hired by C Division in 1982. Except for a 2-year period in the Analysis and Assessment Division, during which he worked on a Central American project, he has continued to work in the C-Division Office. In his position as a staff member, Flavio is involved in the procurement of computing equipment for the Laboratory. Flavio also participates in writing the many reports required by the DOE and the C-Division Office. One important annual report that he works on is the *Information*

Technology Resources (ITR) Long-Range Plan, which provides quantitative projections of computing equipment for the Laboratory.

Flavio is most amazed at how fast computer hardware and software are evolving. He remembers when there were no desktop computers, and all the programming had to be done on the large mainframe computers. "At that time, we did our calculations using a calculator or else wrote a computer program to solve that very specific task. Sometimes I used to wonder whether the effort spent writing the program was worth the end result of the task," says Flavio. Today, he has a personal computer on his desk with standard, commercial software that can quickly and easily do many of the same tasks that he used to have to program himself.

Marilyn Gurule

Marilyn Gurule has been the Communications Group (C-4) Support Section Leader for the last 3 years, since coming to Los Alamos from Sandia National Laboratories 5 years ago. For the first 2 years, Marilyn provided accounting and budget support for Group C-4. After earning an MBA from the University of New Mexico in 1987, she soon became leader for the Support Section. Marilyn has a dual major in Management Information Systems and Human Resource Management.

As leader of the Support Section, Marilyn is responsible for the MIS and applications programming support in the group, the \$17M communication budget and the Communications and Data Equipment Recharge Systems. As a member of the Los Alamos Integrated Communications System (LAICS) evaluation committee, she is responsible for the selection and implementation of the LAICS System Management software and hardware. This system will support, monitor, and control the LAICS integrated digital voice and data communications network.

"Before divestiture of AT&T in the early 1980s, the telephone company was responsible for the entire communications path, from the phones at the customer's

location to the switches at the central office. Now, however, LANL as a customer is responsible for the purchase and maintenance of all user telephone equipment and inside building wire. This means that all user equipment and services must be tracked so that the group can provide information to help determine the nature of various communications problems and to answer user questions. We also verify and allocate costs from several different communications vendors, including US West Communications and US West Information Systems," explains Marilyn.

As the responsibilities of Group C-4 have increased, the group's database requirements have grown dramatically. "We support communications-related databases for the group, as well as an order-tracking database for US West

Information Systems, the XNET Hostmaster database for C Division, the Trouble Ticket application for Network Control, and the Infotron Network Traffic Application for the Defense Nuclear Agency," says Marilyn.

The advent of LAICS presents Marilyn with many opportunities and challenges. "LAICS will require Group C-4 to act as the Laboratory's telephone company, and the System Management Subsystem (SMS) will support the group in providing the end-to-end communication services to the entire Laboratory. I am especially concerned about the loading of existing data into the SMS database. I want to ensure that the data are as complete and accurate as possible and that this information can be accommodated to the design of the SMS database. We have a lot to learn about this fundamentally new communications network with new service offerings and a new network design," explains Marilyn, who is looking forward to meeting the challenges of this rapidly changing communications environment.

Kathy Holian

Kathy Holian is a staff member in C-Division's newest group, the Computer Applications Group, C-9. About two-thirds of the group are programmers, and many of those are "on loan" to groups and divisions throughout the Laboratory who pay for their services on a Form B basis. When the old Computer Research and Applications Group (C-3) was reorganized in 1989, part of the group became Computer Research and kept the designator C-3, and part of the group, the Applications Section, became Group C-9.

Since her graduation from the University of California, Berkeley, with a bachelor's degree in physics, Kathy has worked in physics research. Her experience with computers in college was minimal. In her first job at Livermore, doing equation-of-state research, she used a tutorial to teach herself Fortran in about a week. Her graphics output was displayed on a monitor connected to a Silent 700, and the calculations were performed on the CDC 7600 computer and the STAR computer, one of the first vectorized computers.

In late 1982 Kathy was hired by the Laboratory's T Division where she continued equation-of-state research. "I thought I had taken a step back to the prehistoric days in terms of equipment," Kathy explains. "Some people in T

Division (which never had very much money) only had Silent 700s *without* monitors!"

Kathy continued working in equation-of-state research until 1988, when she read an article in the Laboratory's weekly *Newsbulletin* about a service provided by C Division informally referred to as "rent a programmer." This appealed to Kathy—she always enjoyed programming and was ready for a change—so she went to talk with Ann Hayes, then group leader of C-3. "I can't say enough good things about Ann Hayes," Kathy says sincerely. "I was really impressed with the unprompted praise members of the group expressed about Ann when I first met them."

Kathy was hired by C Division and right away was "rented" to X Division, where she has since been working on a hydrodynamics code called MESA. Hydrodynamics codes are used to simulate the flow of fluids and solids, using the continuum equations of motion. MESA is used primarily to model the penetration of armor by projectiles. Because MESA is a three-dimensional (3-D) code, the use of graphics is an integral part of interpreting the data. Only recently have 3-D codes become practical, partially because today's supercomputers are capable of performing the calculations and partially because workstations

available today provide the means for visualizing data with color and motion. "Trying to see what is happening in 3-D is challenging," explains Kathy. "One way to look at data (for example, a specific variable such as density) is to color-code it so you can differentiate regions of high density from those of low density; another way is to make movies, which provide much more information than stills because they allow you to see the evolution of a computer simulation as a function of time."

Kathy enjoys publishing her work in technical journals and giving talks at professional meetings. With access to today's tools, such as a Sun workstation and a Macintosh, she can easily write papers and generate viewgraphs. In addition, Kathy is truly impressed with the change in computing equipment and software available to Lab staff. "Seven years is not a long time, but the evolution that I've seen at this Laboratory has been really impressive," she says.

Kathy also understands the importance of maintaining the central supercomputing facility. "There are many significant physics problems to be solved now that require large-scale computing and that may impact the survival of the human race in a way that has never before happened," she says. "Global warming and the continued pollution of our environment are a couple of examples. Private industry will not, and perhaps cannot, be expected to pay for this type of research. The government needs to take more of a lead, and we at Los Alamos need to promote our computing center as an invaluable national scientific resource."

Shortly before this report went to press, Kathy learned that she was a recipient of the Laboratory's Distinguished Performance award for 1990. Details of the award will be included in next year's report.

Rich McLean

Rich McLean is C-Division's Program Manager for External Computing. External computing revenues are expected to be a growing percentage of the total CCF budget, contributing, during FY1990, approximately \$14 million of C-Division's total annual budget of almost \$60 million. The Defense Nuclear Agency (DNA) is the largest external client and, with more than a hundred contractor sites using the CCF, will contribute around \$10 million in FY1990.

For many years the CCF was used exclusively by Laboratory staff in support of Laboratory programs; as remote computing became more feasible, however, external clients wanted to compute here, especially those involved in cooperative projects with the Laboratory. Generally, external clients tend to work on research that complements the Laboratory mission; thus their computing needs are similar to the needs of the Laboratory staff. In addition to contributing to the CCF and thereby enhancing its economies of scale, external clients broaden its base of support and provide funding for projects of mutual interest to both the external clients and C Division.

Rich's work as an advocate of external clients began some years ago in California. A graduate of

the University of California, Berkeley, Rich worked as a computational physicist at Lawrence Livermore National Laboratory from 1955 until 1971, when he started working at home building stock and bond portfolio models for fund managers. These were the early days of remote computing, and Rich was developing software to run on main-frame computers from his home study. He used low-speed lines for dial-in access to computers in New York, Palo Alto, and other locations (his phone bill averaged \$1000 per month!). Between 1973 and 1985, Rich worked in southern California for various defense contractors that relied on remote computing to satisfy their computational requirements. In 1982 he began to focus on issues related to remote computing services, particularly those in support of the DNA RDT&E programs. During the course of that work, he made the decision to come to Los Alamos.

Rich regards his job in C Division as having two main parts: project management and external program development. When he first arrived at the Laboratory, Rich spent most of his time on project management, concentrating on adapting the C-Division services to the needs of DNA and educating DNA users

on ways of effectively using the CCF. These were turbulent times while DNA and C Division were "getting to know each other." It is a testimony to both parties that now the DNA contractors, who do work for a variety of government agencies, prefer coming to Los Alamos for their computing in support of their non-DNA contracts.

More recently, Rich spends most of his time on external program development. With the apparent lessening of tensions in the world today, and the tremendous pressures on the defense budget, some of our long-standing programs are in jeopardy; with change, however, come new opportunities, and Rich is trying to ensure that these are not missed by C Division.

Jim Moore

Jim Moore came to the Laboratory in 1967, intending to work here "for just a couple of years." He had grown up in Santa Fe, worked in California, attended college and graduate school at New Mexico State University, and wanted to live in northern New Mexico. Although Jim's master's degree was in mathematics, his interest was computing, and because interesting work in computing was being conducted at the Laboratory, he accepted an offer here. He started his career at Los Alamos worrying about operating systems. Over the years his experience included work with compilers, graphics, applications programming, and parallel processing research.

Part of Jim's work involves administration; for 12 years he has been a section leader in Group C-3, and he is now the acting interim deputy group leader of that group. Group C-3 went through a recent reorganization, with part of the group splitting off and becoming the Applications Group, C-9. The Research Section of the group remained in C-3, and some people involved in research in C-10 were moved to C-3. The group has not yet stabilized as a result of that reorganization, and Jim agreed to serve as acting interim deputy group leader in an effort to provide stability and

continuity for the group until permanent leaders take the helm. The current acting group leader is Bob Malone, a person for whom Jim has praise. "Bob can walk into a situation and pick up on what's going on very quickly," explains Jim.

Jim's primary research interest is parallel processing; research in this field has been a part of the group's activity for the past 10 years. Parallel processing involves applications, architecture, languages, and a wide spectrum of other activities. "Parallel processing is probably the most dynamic part of computing," explains Jim. "There are lots of possibilities to investigate in the 1990s." The work of C-3 is partially to evaluate machines, yet they try to influence manufacturers and work with people during the early stages of the design cycle. In reality, though, C-3 usually doesn't get involved in a project until the vendor has an installed base and is delivering hardware; by that time they must influence the next generation machine.

The Laboratory is regarded nationwide as a strong leader in the field of advanced computing, and some computing companies want to locate a machine here because it can be mutually beneficial and their credibility often is

enhanced by the association. "Not very many machines get in here," explains Jim. "The company has to be willing to build something that the Laboratory wants. The Connection Machine is an example of a successful machine."

During the 1990s, Jim expects to see several large-scale machines. "Cray may well come out with a somewhat different model, and Thinking Machines will probably build something fairly different," Jim predicts. "The Japanese are essentially following in the footsteps of Cray—building the fastest processors but not very many of them." The area that has a real chance of succeeding in massively parallel, general-purpose computing, Jim believes, is dataflow. MIT, Sandia, Motorola, IBM, the Japanese, and, of course, staff in C-3 are proponents of dataflow, in which one can implicitly express massive amounts of parallelism, the major problem is to "throttle down" to a realizable machine. "The methods you use on Crays to program 8 processors are not the same as the methods for programming 10 thousand processors," says Jim. "MIT's goal is to write programs as easily as you do for sequential machines and give up nothing in efficiency." Problems such as these are challenging, yet solutions appear to be at hand. "Olaf Lubeck, a staff member in C-3, was the first one around here to get interested in dataflow," explains Jim. "He dragged the rest of us kicking and screaming into it. On first glance, it is very difficult to believe that it really is possible to do these things efficiently, but the more we learn about it, the more excited we become."

John Morrison

John Morrison, leader of C-Division's Network Engineering Group, first came to the Laboratory in 1971 to pursue a career in operating systems development. After graduating from the University of Illinois with degrees in electrical engineering and computer science, he worked for Control Data Corporation in Los Angeles. His interest in moving to a more rural area prompted his move to New Mexico.

In his 19 years at the Laboratory, John has seen major changes in the way computing is done and has participated in these transitions. The first transition was the one from a batch environment to a timesharing one. "In 1971, the CCF was largely empty," John explains. "People would walk in the CCF with their card decks, hand them to the computer operator, and walk to the output station to wait for their output. Most of the computer users worked near the CCF; there was little remote access to CCF computing and certainly no networking."

Users wanted a timesharing environment and remote access to the CCF. A contentious issue was over timesharing of the CDC 7600s (the supercomputer of the 1970s). Some claimed that timesharing of supercomputers was an inefficient use of scarce computing resources and others claimed that

the gains in productivity were worth the cost. The Laboratory, after much debate, chose to timeshare the CDC 7600s. This decision and the growth of timesharing on former supercomputers, the CDC 6600s, resulted in the development of a remote access network to CCF resources, a network of supercomputers inside the CCF and centralized network resources such as the Common File System.

When the first Cray computer came along in 1976, some of the same arguments came up. "We believed that we would have to develop our own operating system for the Cray, and there were debates about whether it should be a batch system that fit into a network and allowed distributed computing or a timesharing system," explains John. "There was a lot of turmoil at the time and a lot of people still weren't convinced that timesharing was the best thing to do on a supercomputer." John led a group whose mission was to develop DEMOS, an operating system for the Cray. DEMOS was used for approximately a year before the Laboratory's senior managers decided to timeshare the Cray supercomputers.

The second transition, which C Division is experiencing now, is the growth of distributed computing and the corresponding growth

in networking. "During the seventies and early eighties, we concentrated on networking inside the CCF," John describes. "Now we see the focus shifting to Lab-wide networking." This shift is from terminal-based timesharing on supercomputers to distributed computing on small, very powerful desktop computers connected to supercomputers with high-speed networks. This transition will result in C Division shifting its traditional emphasis on central services to an increased emphasis on distributed services.

For the last 4 years John has been the leader of the Computer Network Engineering Group and has been closely involved with the group's high-speed networking work. This work continues to be very exciting. The HIPPI standard is proving to be quite successful, and Los Alamos will be participating in a major government research program on high-speed networking. This recently announced program will result in a gigabit per second testbed network being established between the California Institute of Technology, the Jet Propulsion Laboratory, the San Diego Supercomputer Center, and Los Alamos National Laboratory over the next 3 years.

John is looking forward to the decade of the nineties with the increased growth of networking technology and changing character of computing.

Jan Sander

For the first 7 years of her career at the Laboratory, Jan Sander worked as a casual secretary in IS Division and then part time as a word processor and section secretary in Q Division. Thursday was her day off, and each week she joined 12 other women, all members of the "Over-the-Hill Hiking" group, for a day of hiking in the spring and fall and cross country skiing in the winter. As time went on, it appeared that the opportunities for advancement and more challenging work went to those who work full time, so Jan decided to look for a full-time position in the Laboratory. Soon after, in 1984, she joined C-2, the Computer Documentation Group, as its group secretary.

During her first year in C-2, Jan was assigned a major project—compiling and editing the *Information Technology Resources (ITR) Long-Range Plan*. This plan, which projects expenditures for major hardware and software acquisitions over a 5-year period, is prepared annually for DOE officials in Washington and Albuquerque. "The experience I brought from my years of preparing LA-series reports in Q Division really helped prepare me for this task," explains Jan. As a member of C-2, she had access to writers and editors who helped her develop the necessary skills to publish an annual report.

One of Jan's favorite duties is the monthly preparation of the Form B reports. Form B is a Laboratory-wide method of recharging for services. "Maybe I should have been an accountant," says Jan, "because I really enjoy making everything balance each month." When the Excel spreadsheet program became available for the Macintosh, Jan converted the Form B reports to Excel, a task she particularly enjoyed. "It's fun to link worksheets and see numbers from the first worksheet automatically transfer to the second worksheet based on the formulas that I've put in the cells. As our needs change, I enjoy the challenge of revising the worksheets to accommodate them."

In addition to the usual duties that all group secretaries handle, Jan was asked to take responsibility for local arrangements for the ACM Conference on Document Processing Systems, held in Santa Fe in December 1988. Work on that began a year before the conference, and Jan was responsible for the myriad of tasks associated with planning a conference of 250 attendees. The conference was a big success, and Jan enjoyed it so much that she's working with a coworker on the Slides Committee for SIGGRAPH '90.

As her editorial skills increased from several years of work on the *ITR Plan*, Jan became a resource for the writers in C-2 who prepared the *Computing and Communications Division News*, a monthly publication that is distributed to about 5,000 people. At first she edited individual articles, then she was given entire sections of the publication to edit. Over time her involvement increased until she was asked to assume responsibility

for the entire publication. "I like working on the *News* because of the people I meet and because each article is a challenge. It's been a great learning experience and has helped me in doing the *ITR Plan*," Jan explains.

Jan's writing and editing skills continued to grow, and she took over full responsibility for the *News*. It was time to make the move from being group secretary, and Jan is now a member of C-2's Writing Section. She still prepares the monthly Form B reports because she enjoys the task, and publishing the annual *ITR Plan* is still her responsibility. She is always searching for ways to improve her work and make her job more fun. Future projects include modifying the Form B worksheets and evaluating whether it's beneficial to convert the *ITR Plan*, currently formatted with TROFF, to a Macintosh-based formatting system and preparing Laboratory brochures. Jan clearly enjoys her work, and when she reflects on her decision to go full time, she sees that it afforded her the opportunities she hoped it would. Her only regret is that she can no longer join the "Over-the-Hill" hikers on their Thursday outings.

Becky Trujillo

Becky Trujillo is a system manager for PAGES, the central output node of the ICN. Her job is to work with a small team of development and operations personnel to keep the PAGES system up and running 24 hours a day. PAGES is a complex system of vendor-supplied and locally developed hardware and software that produces output in many formats for users of the ICN.

Becky has been a member of C Division since 1974, where she initially was a word processor. Her career in system management started in the late seventies when new software was purchased for a PDP-11/34 that was used exclusively by word processors in C-Division's Group C-2. Becky was in charge of software upgrades for that machine, and she developed writeups and helped train people in the use of the new software.

In the early eighties, as text formatting systems were being more widely used throughout the Laboratory, Becky provided Lab-wide consulting for NROFF and TROFF. She also wrote macros, mostly for C-Division use, that formatted manuals so they looked the same whether viewed online or as hardcopy.

A life-long learner, Becky has always had an interest in increasing her knowledge and broadening her skills. During her days as a word processor, her knowledge of writing macros was largely self-taught. She frequently attends courses at UNM-LA and the Community College in Espanola, not in pursuit of a particular degree but for the challenge of learning something new. She also takes advantage of in-house training offered in C Division, such as UNIX and Sun seminars.

In 1984 Group C-6 needed a system manager for the VAX in PAGES. Becky was ready for the challenge and opportunities she felt this career move would

provide. Currently, PAGES consists of three VAX/785s, a network of Sun workstations, and 24 output devices. PAGES is an ever-changing environment, as equipment is constantly upgraded to produce higher quality output more quickly. When new devices are added, Becky and the others on the PAGES team work to make those changes transparent to the users so the user interaction with PAGES remains the same but the overall service level improves.

It turns out that challenge and opportunity were only two of the rewards of her present position; another is the sense of community she feels with the people she works with on the PAGES Team. "I have the world's best mentor in Lynn Jones," says Becky, who describes the team as "family." Settled into her new office in the Laboratory's new Data Communications Center, Becky is surrounded by a few photographs of her children and grandchildren, and her extended family of coworkers reside in offices on both sides of hers. For Becky Trujillo, working in C Division and for the Laboratory has clearly offered not only challenge and opportunity but also rewards and satisfaction.

Developing Human Resources

C Division places strong emphasis on developing its personnel resources. As stated in our strategic plan, we strive to attract and retain the highest quality staff, nurture the staff's technical growth, and create an environment that fosters individual participation in the development of new ideas and in the decision-making process.

An example of one such effort that brings together all of these objectives is our active participation in the Laboratory's student programs. In the last year, 55 students have been associated with the division through the programs listed to the right.

These students have contributed to the division through projects and assignments in areas such as the following:

- Graph theory
- Combinatorics research
- Neural networks
- Dynamical behaviors of classifier systems

- Voice recognition
- Cray UNICOS graph reducer
- Interactive graphics interface
- Workstation and Cray support
- UNICOS migration performance
- CFS utilities
- IOS utility

Clearly, not only does the program provide an opportunity for students to obtain real-life, hands-on experience, but the students contribute their special

skills and talents to the division's productivity. Moreover, the division staff benefits from the students' enthusiasm, insight, and ideas.

To make the most of student efforts, a C-3 staff member, Bonnie Yantis, coordinates student selection, space assignments, and equipment needs. As students arrive, each is assigned a mentor whose goal is to make the student's experience one that is mutually beneficial. To monitor the program's effectiveness, students complete questionnaires assessing their work experience, and mentors evaluate student contributions. This collective input helps us continue to improve the program.

Graduate Research Assistantship (GRA)	22
Undergraduate Student Employment	16
Postdocs	6
SERS Students	4
HBCU	2
High School Cooperative Education	2
STEP	1
National Physical Science Consortium	1
Women in Science	1
Total	55

C Division demonstrated its continuing commitment to excellence in education and research by integrating the two in programs for students. The students pictured here were employed during the summer of 1989. From left, back row, are Kathleen Kelly, Eugenia Lee, Toni Patterson, Jay Waller, Jennifer Fox, Steve Anderson, and Tony Brewster; from left, front row, are Simon Levy, David Huelsbeck, Shawn Rivera, and Jeffrey Hamm es.

Awards and Contests

Technical Communicators Earn Awards

Nineteen Laboratory employees have been recognized by their peers for excellence in technical communication. Of those 19, five were from C Division. The annual competition, sponsored by the Society for Technical Communication, recognizes excellence in the documents produced by technical communicators and recognizes the professionalism they exhibit individually. The local sponsor is the New Mexico Kachina Chapter of STC, which includes all of New Mexico and El Paso, Texas.

Marilyn Foster, C-2, won high honors with her entry *Designing Effective Presentation Visuals*, which, after it won an Award of Excellence (second place) at the regional competition, went on to the international competition, where it won an Award of Achievement (fourth place).

Awards of Distinguished Technical Communication (first place) went to Marion Cohen, C-2, for *Input/Output System (IOS) Reference*, and to Margot Casey, C-2, and Kyle Wheeler, C-2, for *NAVOTTS Information System Users Guide and System Administrators Manual*.

Awards of Excellence (second place) went to Chris Siciliano, C-2, for *CFILES Reference* and to Marilyn Foster for *LaTeX Letter Reference* and for *Designing Effective Presentation Visuals*.

An Award of Merit (third place) went to Kyle Wheeler, C-2, for the *C-Division Annual Review and Operating Plan, February 1989*.

The Society for Technical Communication is the world's largest professional society for technical communicators with 13,000 members and 121 chapters and branches.

International Recognition for Computer Graphics

Melvin Prueitt, C-6, has earned a reputation as a distinguished member of the computer graphics community. His work is seen internationally as evidenced by a short sampling of his recent accomplishments:

- Cover and interior pictures in DOE's publication *Technology '88*.
- DOE poster "High-Temperature Superconductivity," prepared for National Science and Technology Week, 1989.
- Five pictures included in the December 1989 issue of *National Geographic Magazine*.
- Two pictures in the SIGGRAPH 89-90 calendar.
- Image called "Contemplating Liftoff" used by SIGGRAPH in their international brochure and their daily newspaper during their conference.
- Two pictures in a window display at the U.S. Embassy in Vienna.
- Pictures of population statistics included in a publication by World Resources Institute.
- Poster for the Brown Cancer Center.

ACM Award

Ann Solem, C-2, won the Don Morrison Best Presentation Award for her talk entitled "A Training Program for Scientific Supercomputing Users" at the spring meeting of the Rio Grande Chapter of the Association for Computing Machinery (ACM) in April 1989. The award is named for Don Morrison from the University of New Mexico, a man known for his contributions to the local chapter of the ACM and for his encouragement of computer science students.

C-Division technical writers brought home awards from a recent publications competition. Shown from left are Kyle Wheeler, Margot Casey, Marion Cohen, Chris Siciliano, and Marilyn Foster, all of C-2.

FASTCAT

Chris Siciliano, C-2, won a Laboratory-wide contest when she submitted the name FASTCAT for the Lab library's online library catalog. FASTCAT (For Automated Search of The CATalog) allows people to search the library catalog from computer terminals in their offices.

Second Annual Computer Graphics Contest

As mentioned above, C Division employs award winners, but it also provides an opportunity each year for any Laboratory computer graphics user to win an award by entering the C-Division-sponsored computer graphics contest. This year, eight winning images were selected out of 58 entries. The contest is held to recognize the many computer graphics users at the Laboratory and to provide a means for others to view and appreciate their work.

Computer graphics are used at the Laboratory to display scientific data, to model objects, and to present results. Many of the winning images this year are examples of the visualization of scientific data.

Entries were generated on a variety of hardware platforms including the Connection Machine, IBM PC, Macintosh, Silicon Graphics, Sun, and VAX. Some of the software packages used to create the images were Adobe Illustrator, Alias, CGS, DISSPLA, KIVA, Pixel Paint, Presenter, Raytrace, and RSCGI.

This year's award winners are Alex Pang, MEE-10; Richard Andrae, N-6; Mel Prueitt, C-6; John Baumgardner, T-3; Peggy Fairfield, T-3; Cindy Boone, IS-12; Geoff Reeves, SST-8; and Allen Mathews, M-4. Their winning entries will be on display in the CCF corridor until they are replaced by the next contest winners.

In 1989 Ira Agins (left) of the Computing and Communications Support Section and James Hodges (right) of MAT-6 won a Laboratory Distinguished Performance Award in recognition of their teamwork that led to the acquisition of the Laboratory's first Connection Machine and the two CRAY Y-MP computers. Timely acquisition of computing resources demands a thorough understanding of technical requirements, long-range planning, DOE regulations regarding computer acquisitions, and Laboratory procurement, financial, and property management procedures. Ira and Jim have demonstrated an exceptional degree of expertise, creativity, and dedication in all of these areas.

This is a snapshot in a time history simulation of thermal convection in the lunar mantle. Blue denotes cooler temperatures and red hotter ones. The prominent harmonic degree-two pattern correlates with the observed large degree-two non-hydrostatic component in the present lunar figure. This work is the effort of John Baumgardner, Group T-3; it was done on a Tektronix graphics workstation using TERRA software.

Overview of C Division

The network at Los Alamos is one of the most powerful scientific computing facilities in the world. ICN users have available for their use a multitude of computing capabilities. The CCF houses an array of supercomputers and associated network servers. It is the responsibility of C Division to meet the requirements of a growing and rapidly changing technology in the computing and communications area and to provide leading-edge computing and communications capabilities.

In addition to the hardware and software that we provide, we must offer opportunities for users to become more productive by offering a variety of support services including consulting, education, and documentation. Providing these tools, however, requires first and foremost a philosophy about managing such a complex and rapidly changing environment, so we will begin our discussion with a description of that philosophy.

Management Policies and Practices

To correctly support the community that we serve, C Division must excel in customer support and research and development. On the one hand, we must provide high-quality service in the face of complexity, rapid change, and rapid growth. On the other hand, we must look toward the future to anticipate how the computing requirements of the Laboratory will evolve and how we can prepare for that future state. Over the past decade or more, C Division has evolved a number of policies and practices that have served us well in

meeting this challenge. These include

- Working closely with users
- Incorporating the work of others
- Economy of scale
- Levels of support
- Change control
- Commonality/portability
- Centralization/automation
- Research
- Planning

Working Closely with Users

The computing resources in the Los Alamos computing facility must satisfy a wide variety of user requirements. The success of all our major projects depends on interactions with the user community, so we encourage users to participate in the formulation of all major projects. Because of the nature of the Laboratory mission, a large fraction of ICN resources are consumed by a small fraction of users. These users have large computational needs, and their needs provide much of the impetus for advancing ICN capabilities. The services needed today by the users with the most sophisticated requirements will be the services required by all users tomorrow. By working closely with these users, we can focus on the needs of key clients and at the same time develop technology for general use.

Incorporating the Work of Others

A major goal of C Division is to provide the Laboratory with a first-rate computing facility. That goal requires that we provide users with capacity and functionality in all areas of scientific computing and communications.

To realize this goal in an environment that is limited in personnel and other resources, we must complement our development activities with the developments achieved by others. This implies that we adopt an "open system philosophy" in all of our development efforts by using standards whenever possible and appropriate, by considering the short- and long-term issues of support and evolution before we begin any project, and by choosing hardware and software directions that are consistent with industry trends as a primary decision criterion.

In general, we buy rather than build if a product is commercially available, and we may choose a less technically sophisticated product if we believe it has potential for spinoff products that we will need in the future. The open system philosophy does not imply that we will no longer develop software. There are many circumstances in which we must maintain control over hardware and software systems that are too important to leave to others or in which the local requirements of our facility and user community are unique and are not yet supplied by others.

Economy of Scale

Once a mainframe computer has been integrated into the ICN, additional mainframes of the same type can be integrated with relatively modest increases in cost. The resulting economy of scale has enabled us to provide extraordinarily cost-effective services. Consequently, whenever a new class of computer becomes commercially available, we select a standard from it through competitive procurement and then continue acquiring that standard until a new class of computer becomes available. This gives us the full benefit of both standardization and economy of scale.

People access the computing network using terminals and workstations, many of which are connected to local area networks. There are four partitions in which they may work: the Secure, the National Security, the Administrative, and the Open partitions. From any partition users have access to central services such as file storage (CFS) and high-speed output devices (PAGES).

Levels of Support

The variety of hardware and software potentially usable in our network is virtually unlimited, but our resources for supporting hardware and software are definitely limited. To manage this situation, we developed the notion of "levels of support" for most of the software that is available to users on the host computers in the ICN and the software that runs on distributed processors. A list of the supported software, together with the levels of support, is given in the *Catalog of ICN Resources* [4].

Change Control

Networks and computing systems can be excellent examples of unstable systems in that minor changes within them can lead to dramatic perturbations in their behavior. In order to provide our users with a stable production environment, preserve security, and evolve the ICN in an orderly fashion, C Division has instituted ICN change control procedures. All component changes that are user visible are subject to change control procedures. Of course, critical bug fixes and hardware repair are done as required.

Our change control procedures reflect the fact that incremental testing is a cost-effective and safe means of implementing modifications. Consequently, certain network resources are designated for this purpose. Further, all changes that the users will notice and their associated schedules are announced in the monthly *ICN Change Bulletin*.

Commonality/Portability

We are evolving into an environment in which multiple computers will be increasingly

used in the software development process. That is, as a distributed supercomputing environment evolves, our clients will use intelligent workstations as their primary interface to the supercomputing environment because of the superior tools that are available on workstations. This will require that we support a more transparent interface between systems and that some commonality exist between the distributed systems and the central systems. We have been approaching this goal for some time with distributed processors, common libraries, and common interfaces to central services. We hope this idea will come to maturity in our evolving distributed environment.

Centralization/Automation

A salient characteristic of the ICN is that it provides a core of powerful and expensive hardware and software to users throughout the Laboratory and the country. We provide centralized systems that are too expensive to replicate at each user site or that require special site preparations. We operate this extensive centralized facility with a minimal staff by using automated capabilities whenever possible. Examples of this automation are our network services—CFS, PAGES, and FOCUS. This results in faster, more reliable service for our clients, and C-Division's operation costs are reduced. We are continually looking for other ways to automate the operations of our central facility.

Research

The objective of research is simply the acquisition of new information. It involves investigations to determine what can and cannot be done (demonstrating that something cannot be done, although perhaps dissatisfying, is successful research). Given the pace at which high-performance

computing is changing—for example, parallel computing, high-speed networks, and the needs of our users for additional capabilities—research is essential if we expect to remain a leading-edge computing organization and provide the high-performance computing environment that is required by Laboratory programs. This research effort should have a substantial applied component, consistent with our responsibility to prepare the Laboratory for the future in computational science. Currently, only about 10% of C-Division personnel are involved in research and development efforts. We believe that our efforts in this area must continue to grow so that we can be adequately prepared for the future.

Planning

While C Division has prepared an annual plan for a number of years, it wasn't until 1988 that we included much detail about our strategic planning efforts [5]. We believe that our strategic planning efforts help us keep our goals in

focus in a time of new opportunities and increasing demands on our limited resources.

Major ICN Components

The remainder of this section describes the main components of the ICN: the worker computers, the central servers for file storage and output, and the growing use of distributed systems. Information about CAE/CAD/CAM is also included.

Worker Computers

Major computing resources in the ICN are the Cray supercomputers, known locally as worker computers. As of August 1, 1990, there are three Cray computers (17 CRAY-1 equivalents) in the Open partition, one Cray computer (7.5 CRAY-1 equivalents) in the National Security partition, and two Cray computers (20 CRAY-1 equivalents) in the Secure partition. The number of CPUs, memory size, Solid-state Storage Device (SSD) size, and disk space on these machines varies with the different models.

Worker computers in the ICN include a CRAY X-MP/24 (Machine 1), a CRAY X-MP/48 (Machine 2), and a CRAY X-MP/416 (Machine 5) in the Open partition; a CRAY X-MP/416 (Machine 4) in the National Security Partition; and a CRAY Y-MP/832 (Machine 7) and a CRAY Y-MP/8128 (Machine 8) in the Secure partition. Machine 5 is owned by DNA, the Defense Nuclear Agency.

Network Operations Center

The ICN has experienced rapid growth over the past few years, particularly the packet switched service; therefore, it now requires more than a part-time effort to operate. In response to this rapid growth, C Division has established a Network Operations Center (NOC), which is really a consolidation and formalization of many efforts already within C Division. As such, the mission of the NOC is to focus on operational issues in the context of the packet switched service provided within the ICN. The ICN packet switched service has two aspects:

- First, there are local area networks (LANs) within the Laboratory that individual groups and divisions own and operate. Most user workstations are connected to these LANs, and these LANs are managed by a local network manager from that group or division.
- Second, there is the common Laboratory-wide network that connects LANs and worker computers. This common network (known as the Los Alamos Internet) is the regius of the NOC.

The NOC's charter is as follows:

- Maintenance and troubleshooting of the ICN packet switched service (Internet)
- Configuration management and quality control of network components
- Planning for future technologies and growth
- Network security centralization
- Network training
- Network Information Center Hostmaster assignment of host addresses for TCP/IP and DECnet network information dissemination
- Postmaster mail management
- Network consulting for network managers
- Network management and monitoring tools development

The components of the ICN packet switched service that NOC oversees are:

- The worker switch network, which routes SIMP packets at 50 Mbits over a point-to-point topology. This network connects worker computers, CFS, and PAGES with multi-partition security.
- The five Ethernet backbones (Open, Administrative, National Security, Secure Unclassified, and Secure Classified), which route TCP/IP and DECnet packets over a 10-Mbit broadcast network. These networks are gateway (firewall) connected to the worker switch network and they provide single-partition security.
- The LABNET packet switched network, which routes TCP/IP packets over the LABNET cable system at 5 Mbits. This network is connected to the Open Ethernet backbone.
- The connections to external networks, which are:
 - Energy Research, ESNET, TCP/IP, and DECnet;
 - National Science Foundation, NSFnet, TCP/IP;
 - NASA, SPAN, DECnet;
 - New Mexico TECNET, TCP/IP, and DECnet;
 - Many smaller external networks

CFS

CFS is a large, online, centralized file-storage and file-management system for the ICN. It uses an IBM 3090 and an IBM 4381 as control processors. The storage systems include three IBM 3380 disk systems and an IBM 3480 tape storage system. CFS is able to serve a diversity of machines, from supercomputers to personal computers, and a diversity of users, including interactive humans, batch programs, operating systems, and other network servers.

CFS capabilities are provided through a standard set of requests that allows users to store and retrieve files and to build and manipulate tree structures that assist in the management of their

files. CFS is organized as a hierarchical storage system in that active files are stored on disk, and very large files and archival files are stored on 3480 tapes. Files are moved between disk and tape by a migration program based on file activity and file size. The result is a cost-effective system that provides both fast access and large data storage capacity.

CFS Storage. After a decade of use, CFS storage exceeded four million files and 100 trillion bits of data. The daily activity is 22,000 file accesses, which amounts to 77 Gbytes of data transferred. Fifty percent of the files are less than 52,000 bytes, the average size of a stored file is 3.0 Mbytes, and the average size of an accessed file is 3.5 Mbytes.

In October 1989 CFS contained 4.0 million files comprising 12,083 Gbytes of data; 94% of this data was stored offline. Of the 22,000 file accesses per day, about 95% are to disk, whereas about 5% are to offline storage.

CFS Performance. The top priority is to continue improving the throughput, availability, and security of CFS. During second-quarter FY1990 we will install a new CFS control processor. During fourth-quarter FY1990 we will increase the maximum CFS file size and will increase the CFS program size to allow more requests to be executed in parallel.

CFS User Interface. Scheduled to be implemented during fourth-quarter FY1990 as part of the new user interface are the data-compression option, IN/OUT/LOG files, password suppression, and a wildcard capability.

New CFS Initiatives. The following areas will be investigated:

- High-speed data transfer. The goal will be to substantially increase the file transfer rates.
- Workstation storage. The goal will be to make CFS capabilities more available and usable for the workstation environment.

CFS performance is shown in terms of gigabytes stored and number of files stored since its origination in June 1979.

- File transparency. The goal will be to provide users with access to CFS files and directories through the local file system commands and facilities of the machine they are computing on.

PAGES

PAGES is a subsystem of the ICN to which any worker computer, distributed processor, or workstation connected to the ICN can submit jobs to be processed. Choices of output media from PAGES include 8.5 x 11-inch paper, 11-inch and 36-inch roll paper, 105-mm fiche, 35-mm positive film, 16-mm positive film, and cut-sheet film up to 'E' size. CGS metafiles, raster subset CGM files, standard text files, standard document files, and Gerber photoplotter files are all accepted by the system.

PAGES is controlled by several general-purpose processors (approximately two VAX 8600 equivalents) and drives seven film recorders, five high-speed printer/plotters, three electrostatic plotters, two color printer/plotters, and one photoplotter. Output volume exceeds 7 million images per month in the form of 6 million fiche images, 1 million cut-

sheet paper images, 15 thousand roll paper images, 45 thousand slides, and 94 thousand movie images.

We are continuing to work toward a better client interface, a graphical operator interface, and a system that can be distributed across many systems transparently for continued growth. We are also looking at new services such as raster, CGM standard metafiles, and direct PostScript. In addition, we are doing some fundamental redesign at the job-queuing and graphics-imaging model level. We have negotiated a license with Adobe Systems to facilitate support of direct PostScript, CGM, and raster data on PAGES.

FOCUS

FOCUS is a component of the ICN that provides a variety of network services to the network hosts. Services now include: PROD (production or batch workload manager), INFORM (status of ICN nodes), GATHER (host performance measurement data), FONE (interactive laboratory telephone directory), ACS (account control system), ICNHELP (help on worker

utilities), DST (interface to manage dedicated system time), CRAY-LOG (CTSS system operator messages), and more. The concept of a central network server with a generic connection to network hosts continues to be very useful.

The need for managing the production or batch workload has become even more critical with the addition of CRAY Y-MPs. FOCUS currently directs the flow of over 440 charge hours of batch workload in an average night shift of 14.5 hours.

All external users of CTSS are now subject to ACS checks for correctness of charge codes.

We upgraded hardware and operating system software in July to increase I/O throughput and to increase reliability. We are now well positioned for the near term to provide FOCUS network services to a growing ICN.

In FY 1990 we will provide a UNICOS connection to FOCUS to include the batch workload features of FOCUS in the UNICOS environment.

CAE/CAD/CAM

C-Division CAE/CAD/CAM activities have been sharply reduced. Most application activities have been migrated to distributed systems at the user sites, and the ICEM/DDN consulting activities have been moved to the Control Data Corporation support center. The CAE/CAD/CAM NOS machines have been phased out of the CCF.

Software Support

C Division develops, enhances, and maintains operating system software running on the major resource computers in the ICN along with related languages, libraries, system and graphics utilities, and document processing software.

Operating Systems

CTSS continues to be the primary supercomputer operating system at Los Alamos. A large-scale effort is under way to migrate all CTSS supercomputing systems to UNICOS in 1 to 2 years. The decision to provide a UNIX-based environment on Laboratory supercomputers was made in May 1989. Several major factors contributed to this decision:

- A UNIX-based environment provides a good base for distributed computing development, which is important to our clients;
- UNIX-based environments are widely used and vendor supplied, so we are able to leverage our resources;
- C-Division will be better able to make enhanced versions of vendor software available on a timely basis, which will provide our clients with advanced computing capabilities more quickly.

CTSS Software

We will continue to maintain and tune CTSS during the next year. The CTSS message buffers will be upgraded to allow for possible large network messages and data streaming to a graphics terminal. The CTSS file transport module has been modified and tuned for the CRAY Y-MP. We will experiment with the file transport between CFS and CTSS over the 100-Mbyte Cray channels.

UNICOS Software

A comprehensive CTSS to UNICOS migration plan was developed, and implementation of the plan is well underway. We are currently running the UNICOS operating system on a CRAY X-MP/24 in the Open partition, and plans call for UNICOS to be installed on CRAY Y-MPs in the near future. We will continue to install the latest releases of UNICOS and move toward the environment expected and required by supercomputer users. We will integrate the workstation into the UNICOS environment to provide distributed computing capabilities required by our clients.

NOS Software

NOS is the timesharing operating system supplied by Control Data Corporation (CDC). Machine N, the only remaining NOS system, will continue to run NOS for an undetermined period to meet ADP-Division requirements.

VMS Software

VMS is the primary operating system supported by Digital Equipment Corporation (DEC) for the VAX family of computers. The VAX was originally selected as our standard distributed processor. Our support philosophy for VMS is to provide the latest vendor software with a minimum of local modifications. We maintain VMS software on FOCUS and OFVAX, which are used for special applications, and on CCVAX and DNAVAX, which are used as general-purpose workers.

The Laboratory's electronic mail system was converted from DECmail to the ALL-IN-1 office

management system, and hardware and operating system software were upgraded to improve electronic mail service. In addition to electronic mail, ALL-IN-1 provides an online dictionary, thesaurus, telephone directory, and time-management system.

The Laboratory's ALL-IN-1 mail system is supported by the OFVAX cluster. Currently, there are about 3,000 validated users of this system, and we are delivering an average of 23,000 memos per month. About 2,500 memos are being stored on CFS each week.

The ACS now verifies open CTSS charge codes at the user level rather than at the organizational level. FOCUS will also maintain an up-to-date roster of users and their appropriate organizations. If a user moves from C Division to X Division, his bank-account code changes.

UNIX Software

UNIX is an operating system licensed by AT&T and distributed by various vendors. It runs on a diverse set of computers. ULTRIX (LEC's version of UNIX) runs only on VAXs. ICN users access UNIX VAXs for general-purpose computing, word processing, electronic mail, software development, and artificial intelligence research. UNIX software provides an important connection to several national networks, which facilitates our communication with other institutions and government organizations. UNIX or ULTRIX software also supports gateways to the open VMS network, supports PAGES, and provides auxiliary functions for CFS.

The ULTRIX Mail Connection was installed as the UNIX bridge to ALL-IN-1 during the electronic mail conversion.

Languages

Fortran will continue to be the most heavily used programming language at Los Alamos. Fortran 77 is our standard language on all operating systems, and users are requested to use it. In FY 1988 we continued our joint project with Cray Research to ensure compatibility among the products that will become the Y-MP/832 Fortran run-time environment, including the new Cray products CFT77 and CAL2, the Los Alamos loading and debugging tools, and the run-time support libraries. This project has involved staff from all groups within C Division, several Cray internal organizations, as well as a number of users. Extensive effort has been expended to ensure the successful preparation of the C-Division software that provides the user interface and run-time execution environment for the new machines. It is expected that this effort will continue through the coming year as we integrate the Y-MPs into the ICN.

The C language is also very important within the Laboratory not only on VAX machines running UNIX but also on workstations and the Crays. This language is also supported on CTSS and is being heavily used in our development of software for a distributed computing environment.

In addition to Fortran 77 and C, there are a number of other languages in use at Los Alamos. These specialized languages are not available on all systems, however, and their level of support varies from Level 1 to Level 4. They include general-purpose languages (Pascal, BASIC, Prolog, APL, PL/1, LRLTRAN, and Ada), command languages (FCL, CTL, COSMOS, C-shell, and Bourne-shell), and special-purpose languages (APT, MACSYMA, and SIMSCRIPT).

Libraries

C Division supports run-time, mathematics, and graphics libraries. These are designed to support user codes in a wide variety of applications.

Run-Time Libraries

Run-time, or system, libraries support the Fortran and C compilers. They also provide a number of useful extensions to these languages. C Division supports one major run-time library system on CTSS: CFTLIB and its support libraries primarily for Fortran support.

CFTLIB provides the basis for the current Fortran run-time environment on CTSS. This extensive library has been in use at Los Alamos for many years and has evolved to meet the needs of our users very well. During the past year we began a major enhancement effort for this library and its support libraries in preparation for the Y-MP/832 machines. This effort will continue into the coming year and provide efficient run-time support for our users on both the CRAY X-MP architectures and on the Y-MP machines.

Mathematics Libraries

Essential to the writing of large scientific programs is a reliable and robust mathematical software library. This library (called CLAMS) is written primarily in Fortran 77, consists of approximately 1000 routines, and is available on all major computers. CLAMS is accompanied by online and printed documentation.

Other available libraries are CALMATH, a set of specialized routines written in Fortran and CAL (Cray assembly language) for high performance on the Cray architectures; CRAYMATH, Cray's library of intrinsic functions; and IMSL (International Mathematical and Statistical Library), a proprietary library of mathematical and statistical routines.

Online documentation programs, with searches based upon routine name, category, or keyword phrases, exist for both the CLAMS and CALMATH libraries.

Graphics Libraries

The Common Graphics System (CGS) library continues to be the device-independent, system-independent foundation of our Integrated Graphics System here at Los Alamos. The other graphics libraries, CGSHIGH, DISSPLA, GKS, NCAR, and the SC-4020 simulation library, work in conjunction with CGS. Users are able to view their graphical output from these libraries interactively at their terminals or store their graphics in a CGS metafile for later viewing. They also may ship this metafile to a PAGES output device. CGS has been enhanced to include an index color model and to allow users of the direct RGB model to load the lookup table as new colors are encountered.

Much effort has gone into helping define a subset of the emerging standard, the Computer Graphics Interface (CGI), which is supported at all the SLATEC sites. This common foundation has allowed us to share software among the SLATEC sites. Our raster library is built on this standard, and the output file is the raster subset of the international standard Computer Graphics Metafile (CGM). We have a friendly user version of this library and the accompanying utilities being tested by our users. With this system, a user can generate a raster image or a CGS metafile, convert it to a CGM or generate the CGM directly, and ship the CGM to a PAGES film recorder or a laser printer. A user can create a raster image on a MacII or a Sun workstation, transfer the image between these two devices to modify it, and then get a color slide of the image.

Graphics Utilities

The graphics utilities at Los Alamos come in two types: utilities to generate graphical images and utilities to work with a file that contains graphical images. MAPPER is a utility that allows users to generate graphics. The other type of utilities, postprocessors, take a CGS metafile, which is a device-independent, system-independent file of graphical data, and allow users to manipulate that file. The PAGES postprocessors, PFILM, PESF, PPEN, and PRPP, enable users to ship their metafile to a central output device. Other postprocessors, PTEKT and PPS, allow users to view their graphics on a local device. PSCAN and PCOPY are postprocessors that allow users to edit their metafiles and write new metafiles from existing metafiles. Should the user wish to view his metafile on a device not supported by the graphics group, PXXLIB is a library of subroutines that will read the metafile and provide the means for the user to output the graphical information to a device of his choosing.

PostScript files have become a de facto standard in the world of desktop publishing and text and graphics integration. We have a translator, PPS, that converts a CGS metafile into a PostScript file. Also, the raster subset of the CGM, an international standard, is the vehicle we support to transmit raster data. We have a two-step process that will translate from the CGS metafile to this format. A new postprocessor, PGCGM, will allow users to send the CGM file to PAGES. We are convinced that providing means to integrate these standard files into our system will give us and our users leverage in this rapidly changing environment.

System Utilities

System utilities include the tools that users need to write, maintain, execute, and debug programs. These utilities also include tools for communication between workstations, between terminals or workstations and the ICN, and between machines at Los Alamos and machines elsewhere.

We will continue to develop and install our family of command languages. With these powerful and syntactically coherent languages, users can customize their interactive computing environment, and they can write command files to automate repetitious work. These languages rest upon a common foundation of commands and operations. They differ among themselves in that each one models its syntax on a particular programming language. One of these languages—Fortran command language, or FCL—has been working on CTSS machines for over a year. A C command language, along with a command shell, is being developed for UNIX-like systems.

A mouse-driven window interface for CFS is under development, and we are investigating designs of window interfaces for other major system utilities.

We are improving ways to send information between sites in the Nuclear Weapons Complex (NWC). We have developed the NWCfile, a file that contains data and information about the file's attributes, transmission, and creator. These files can be shipped efficiently from one NWC site to the other. The NWCfile Tool Set is a group of subroutines for creating and manipulating NWCfiles. The Tool Set is now available on the IBM/MVS and VAX/VMS systems and will be ported to UNIX. In addition, the subroutines in the Tool Set will be packaged in a set of utilities so non-programmers can use them easily.

The Input/Output System is a facility in the ICN that can read and write tapes and cartridges so that data can be transmitted between the ICN and other computer centers. To make this transmission possible, utility routines are necessary to convert various tape and cartridge formats to and from formats that C Division supports. We have developed converter routines on CTSS for the following text files: CTSS text, NOS text, 7-bit ASCII, EBCDIC, and STXT. Routines for converting between CTSS and NOS binary formats are in progress; converters for various VMS formats are planned in addition to making these routines available on VMS.

Document Processing

The ICN Distributed Document Processing System provides tools for ICN users from any ICN-standard terminal or workstation. The system supports both local and PAGES unclassified output from two common text formatters, TeX and TROFF, and also supports merging of text and CGS or PostScript graphics using TeX or TROFF. The TeX and LaTeX environment is available for VAX/VMS, VAX/UNIX, Sun and Apollo workstations, and IBM PCs; the TROFF environment is available for UNIX.

Supported printers include local PostScript printers and central PAGES output. Laboratory-standard letters can be produced using LaTeX letter style, and Laboratory-standard memos can be produced using LaTeX memo style, TROFF memo macros, or the MEMOMGR utility. Support for classified TeX output through PAGES, input of scanned images, and higher quality output devices are other projects being worked on.

Communications Support

C Division provides Laboratory voice, data, and video communications through a combination of Laboratory and vendor resources. The communications services provided include telephone, radio, point-to-point data communications, local and wide area networks, and closed circuit television. Communications support functions and some of the projects that are necessary to continue providing this support in FY 1990 are described below.

Communications Services Center

The C-4 Communications Services Center provides a central point of contact through which clients may request communications services, verify the status of existing service requests, or inquire about the types and availability of communications services. The Services Center allocates communications resources, recommends system upgrades, configures telecommunications systems, issues all user service requests, coordinates vendor activities, provides user training, and conducts quality assurance audits.

Network Control

The C-4 Network Control Team is the central point of contact for ICN communications problems. Available from 6 a.m. to 6 p.m. on Laboratory working days, we diagnose and repair data communications problems. In addition, we repair faulty circuits, perform preventive maintenance on CCF data communications equipment to improve reliability, and operate secure communications circuits. Our goals are to improve reliabil-

ity and reduce the time required to repair faulty circuits.

Support

The C-4 Support Section is the central point of contact regarding the costs of all communications services and is responsible for C-4's management information system. We answer client questions about communications costs and recharge procedures and monitor communications tariffs to ensure Laboratory costs are as low as possible. Formulating and monitoring the communications budget is our responsibility. In addition, we propose recharge rates and produce monthly billings to recover communications costs. Working with each C-4 section and the local communications vendors, we provide programming and systems support.

Remote Access Networks

Remote access networks use standard protocols and communications circuits to interconnect users and resources. A user on such a network could access any resource or user connected to the network. These networks are critical to research and development because they provide access to computing and information resources that are geographically dispersed. They are especially important to support collaboration between researchers at the Laboratory and around the world.

Off-Site Communications

Off-site communications are point-to-point circuits connecting our external clients to the ICN. C Division provides computing and communications services to a large number of DOE and DoD affiliates. Associated users are distributed throughout the United States. In FY 1990 we will expand off-site communications support to a large concentration of clients on the east coast.

External networks that tie into the ICN are growing in number and in use.

Transmission Facilities

Transmission facilities are provided in support of all Laboratory communications requirements. This Laboratory-wide resource is comprised of radio, twisted-pair copper wire, coaxial and fiber-optic cables, and their supporting structures. Supporting structures include radio towers, conduit systems, and telephone poles.

Radio Systems

Various types of radio communications systems are used to provide instructional television from UNM, secure radio communications for the Laboratory's security force, radio paging, general-purpose two-way radio communications, and back-up telephone service for the Emergency Operations Center. In FY 1990 a trunked radio system will be beta tested.

Los Alamos Broadband Network

LABNET is a Laboratory-wide, frequency-division multiplexed, broadband coaxial and fiber-optic

system that provides data transmission at rates up to 1.544 Mbit/s, local area network backbone transmission at rates of 5.0 Mbit/s, and full-motion video transmission.

Outside Cable Plant

Outside cable plant consists of twisted-pair copper wire and fiber-optic cabling interconnecting buildings throughout the Laboratory's 43 square miles. These facilities are undergoing continual expansion.

Duct Bank

Duct banks constitute a major portion of the Laboratory's outside cable plant supporting structure. Duct banks are below-ground, concrete-encased conduit used to house communications cabling. All of the proposed duct banks involved General Plant Projects funding, typically requiring 3 to 5 years to plan, fund, and construct.

LDCC Communications

Construction of the LDCC was completed in August 1989. The building supports expansion of the ICN, the CCF, and LAICS, which is the Laboratory's new, integrated communication system.

The LABNET head-end was relocated from SM-28 to the LDCC. Moving the head-end involved the rerouting of trunk cables and the relocation of head-end equipment (processors, translators, modulators, and test equipment) to the first floor. LABNET will be available in each office for video or data applications.

The encryption devices were relocated from the SM-132 secure vault to the LDCC secure vault. All local and off-site secure circuits to the CCF requiring encryption were transferred to the new vault. The vault in SM-132 was decommissioned.

Users at sites in a variety of locations use a national communications network to compute in the ICN.

Direct Support

Service is the key to helping our clients get the most out of the many computing capabilities of the ICN, and we are committed to providing the best possible services to meet client needs. Services include around-the-clock computer operations, person-to-person consulting on a wide range of topics, development and distribution of computer documentation in both online and printed forms, and presentation of state-of-the-art computing courses and educational materials. These are a few of the services that help make computing and communications capabilities easier to use.

Other services are provided on a Form-B basis. These include mathematical, database, and applications programming; communications system design, implementation, and maintenance; local area network administration; and document preparation, data entry, and design and generation of presentation graphics.

Operations

Keeping the computers operating is primarily the responsibility of Group C-1; communications operations are provided by Group C-4.

Computing Operations

Computer operations services are provided to all users of the CCF on an around-the-clock basis. The services include monitoring and controlling ICN host computers, assisting users, and distributing computer-generated output. We resolve computer hardware and software problems quickly and accurately. We schedule and announce computing interruptions such as hardware preventive

maintenance, software development, and facilities maintenance activities. We gather, compile, and track hardware performance data for ICN host computers and publish monthly reports for management and the user community. Computer operations personnel are involved, at least in part, in computing activities relative to security, hardware planning and acquisition, site planning and installation, and facilities management. We are often involved in inter-group projects designed to enhance operational efficiency and user productivity. Group C-1, often the first point of contact for users of the CCF, strives to make a good first impression for C Division, and we are frequently called upon to give ICN overviews and tours of the CCF.

Communications Operations

The Communications Group provides Laboratory voice, data, and video communications services. We provide these services through a combination of Laboratory and vendor resources. Communications networks and projects are described in the previous section.

Consulting

Although the consulting function provided by Group C-10 is most visible to ICN users, C Division also provides numerical consulting, communications consulting, and procurement consulting related to data processing and word processing equipment. In addition, we provide database consulting, and the Dispatchers Desk, located near the CCF, is the focal point for many

user services. More details about each of these consulting services are provided below.

Consulting Office

Group C-10's Consulting Office provides consulting services to ICN users. We answer questions on a wide variety of topics, including the use of programming languages, system libraries, graphics libraries, utilities, command languages, controllers, electronic mail networks, and external networks. We also assist users with code debugging. We advise users about what software is available for solving particular problems and where more information can be obtained. We provide feedback to software developers, documentation writers, and the Computing Education Program on those aspects of system usage with which users have difficulty. We teach classes on locally developed software, utility and library routines, operating system revisions, and using the ICN. Where appropriate, we write or modify user programs and utilities. We review and generate documentation in conjunction with Group C-2. When software bugs are reported, we inform the appropriate software maintainer. When problems with the network arise, we help trace the problem until a satisfactory solution is achieved.

- Consulting service is provided from 8 a.m. to 5 p.m., including lunch, on normal workdays, except Mondays between 1 p.m. and 2 p.m. when we have our weekly section meeting. The Consulting Office phone number is (505) 667-5745.

- For non-urgent problems, the Consulting Office can be reached by electronic mail using any of these addresses:

consult@lanl.gov from UNIX or Machine Beta, "consult@lanl.gov" from VMS, or CONSULT@LANL from INFORM on the To: line.

Numerical Consulting

A major activity of many scientists at the Laboratory is the mathematical modeling of physical processes. Often these models take the form of differential or integral equations and are sufficiently complicated to require numerical solution. The Computer Research Group (C-3) provides consulting and research services designed to increase the effectiveness of these models. This involves collaboration between Laboratory groups skilled in modeling and Group C-3, which is skilled in problem-solving techniques, both analytic and computational. We also provide consulting on mathematical software in the general areas of special functions, linear and nonlinear algebra, interpolation, approximation, quadrature, ordinary and partial differential equations, and integral equations. This software is made available to users as the CLAMS Library and as part of a new auxiliary library designed to contain more advanced and sometimes experimental algorithms than those available in the basic CLAMS set. Our phone number in C-3 is (505) 667-7028.

Communications Consulting

The Communications Group provides consulting on the design, implementation, and maintenance of communications systems to the extent desired by each organization. Our phone number in C-4 is (505) 667-7057.

Procurement Consulting and Review

It is mandatory that the Laboratory comply with federal regulations aimed at assuring all automated data processing equipment, word processing equipment, and related equipment will be procured as economically as possible and still perform as required. The regulations are complicated, they change frequently, and they

require interpretation. The C-Division Support Section staff has been delegated responsibility for coordinating Laboratory purchase requests concerned with the acquisition of these computing resources. This is accomplished by guiding the requesters through the maze of federal regulations, interacting with MAT Division, and corresponding directly with DOE. Our phone number in CDSS is (505) 667-6171.

Database Consulting

Group C-10 provides consulting and vendor liaison support for the Laboratory-standard database management systems, S2K (NOS) and INGRES (VMS). Since the vendor has announced that S2K will no longer be supported on CDC computers, a large part of our current effort consists of advising and aiding users in selecting alternate systems for their applications. Our phone number in C-10 is (505) 667-7877.

CCF Dispatchers Desk

The CCF Dispatchers Desk serves as the focal point for many

user services activities. The Dispatchers Desk is the point of contact for users who need computer operator assistance, machine status or general information about C Division, or help locating output. Dispatchers Desk personnel are responsible for routing incoming telephone calls to the appropriate C-1 personnel and for registering official visitors to the CCF. Additionally, we provide users with information about CCF hardware and facilities activities and operate a magnetic tape library. The Dispatchers Desk is located adjacent to the southwest corner of the CCF in TA-3, Building SM-132, and the hours of operation are from 7:30 a.m. until 5 p.m. Monday through Friday. The telephone number is (505) 667-4584.

Computer Documentation

Computer documentation, both locally produced and vendor supplied, is available on a wide range of topics, including ICN systems, scientific and personal

The CCF Concourse, located adjacent to the CCF and near the Dispatchers Desk, is where people pick up their unclassified computer output.

workstations, software applications, graphics, and mathematics. Local documentation about ICN capabilities supplements vendor documentation and makes the complex computing network easier to use. The *Computing and Communications Division News* and the *ICN Change Bulletin* are published monthly to inform users about computing news and software and hardware changes. These publications provide information about all ICN computing systems.

All users have online access to vendor-supplied HELP information, local ICNHELP quick reference information, the CLAMDOC library, and the C2DOC library of local documentation.

We continue efforts to enhance the online presentation of up-to-date information about ICN software and hardware changes.

Users who also require printed documentation may order it electronically by using the Electronic Card Catalog and Ordering System (ECCOS). Both local and remote users who have access to the ICN can use ECCOS to browse through the Computing Information Center (CIC) documentation catalog, search for documents by subject, and order documents from the convenience of their offices. The documents they order are sent to them through the mail.

Manuals, local documents, and vendor publications are available at the CIC. The CIC operates out of two locations: an open area at Pueblo Complex and a secure area at TA-3.

Education

C-Division Computing Education Services offer a wide variety of educational opportunities to ICN users. Local and vendor computing courses are developed or purchased to help both new and experienced computer users do their computing more efficiently.

The emphasis for FY 1989 was placed on increasing course offerings for users. The TA-3 classroom continued to support classes requiring specific terminals or workstations, including Sun, Tektronix, Macintosh, and DEC. Seminars and courses about current topics in computing and communications were received via satellite. As a result of the increase in popularity of microcomputers and workstations, many related courses were added to the library of self-paced courses. Video courses with hands-on labs were offered in the Pueblo Classroom.

We initiated and implemented the first Los Alamos Workshop on Scientific Supercomputing. We plan to continue the workshop when space becomes available.

We will be moving into the UNICOS training arena as soon as user community needs are established. We are also exploring cooperative efforts with vendors for specific subjects related directly to their hardware, i.e., Thinking Machines, Hewlett-Packard, Sun, etc.

New courses are advertised in the *Computing and Communications Division News*. The *Computing Education Catalog* [6] contains a complete listing of self-paced courses; an online version of this catalog is on the INFORM system.

Live Computing Courses

Live computing courses are offered to provide interactive learning opportunities in a classroom environment. Courses are presented in response to user requests; at the request of special interest groups who need a centralized site for ICN-related and distributed training; and whenever significant changes to ICN systems, libraries, and utilities are planned.

A lecture room and a classroom at TA-3 are used for live courses and seminars. The TA-3 class-

room accommodates 16 people for hands-on training. It can be configured for local area networks using Ethernet connections or communication ports to Open and Secure, Unclassified and Secure Unclassified partitions of the ICN. The large-screen projector in this room has capabilities for videotape and most terminal/workstations displays.

The Pueblo Classroom, located in an open area at Pueblo Complex, is equipped with a large-screen video projector and port access capabilities to the ICN. The classroom is used for live classes, group video classes, seminars, and meetings.

Self-Paced Computing Courses

Self-paced courses are offered in a wide variety of media including videotape, audiotape, and computer disk. A library of IBM PC and Macintosh courses is available for users to check out and use at their convenience.

LABNET-Broadcast Courses

Courses, seminars, and special satellite programs are broadcast to viewing areas throughout the Laboratory on LABNET, the Laboratory's broadband network. Viewers may select from a variety of broadcasts scheduled Monday through Friday from 8 a.m. to 5 p.m. on LABNET Channels 9 and 10.

In FY 1989 we increased the broadcasts of satellite teleconferences and seminars to provide users with the most up-to-date technical information. In FY 1990 we plan to further broaden the selection and increase the number of satellite broadcasts, video course offerings, and live broadcasts available on LABNET.

Materials for courses and special broadcasts are distributed to users upon request. LABNET installations are requested through the Communications Group.

Special satellite teleconferences and regularly scheduled courses are announced in the *Computing and Communications Division News*, via flyers, in the *Laboratory Newsbulletin*, on Update, and on the INFORM Bulletin Board.

Users Areas

Computing facilities are provided for computer users at TA-3, SM-200, for Q-cleared users; and at Pueblo Complex for all users. These facilities include terminals and workstations connected to the ICN computing systems. The TA-3 area also has secure terminals for people who need to do classified computing. The Pueblo Complex area is conveniently located in the Computing Learning Center near the Computing Information Center.

Custom Services

C Division provides programming and communications services tailored to suit each client's needs. In addition, we provide documentation support services, including computer-generated graphics, document preparation services, and data preparation services. Details are provided below.

Programming Services

C Division offers custom programming services through the Computer Applications Group (C-9). The purpose of these services is to support any of the Laboratory's groups or programs that have programming requirements, but do not require full, long-term support. This service, offered through "rental" agreements between Group C-9 and the requesting group, may vary in both length of project and level of effort. The average job lasts about 4 months at a half-time level, but jobs can range from a few hours to up to 2 years, and from intermittent support to full-time involvement.

The skills to work on a wide variety of machines, operating systems, languages, and areas of interest are cultivated and maintained within the applications personnel. For example, some of the recent applications jobs have been performed in the areas of computer graphics, weapons code development, 3-D visualization programs for pin shot data, several applications of a fourth-generation language database, a database for chemical information for weapons tests, code vectorization, documentation graphics programs, INGRES database development, analysis-of-risk software, VAX management, code conversion, and a prototype of an advanced information management system for a branch of the office of the Secretary of the Navy. To support this variety and quality of operation, training is encouraged for the personnel. Our goal is to provide high-quality service on an as-needed basis for the widest variety of computing needs.

Advanced Information Management Systems

Within the Computer Applications Group a team of software engineers is developing hardware and software tools to implement computer-based systems that access, understand, and manipulate large amounts of information in different formats (textual, graphical, video, etc.). In particular, support is being given to several federal agencies for the prototyping and proof of concept for information management and presentation systems that address areas that deal with national security issues and interests. This work and other Laboratory assignments will concern the areas of client/server communications, advanced information navigation, integration of databases with hypermedia and expert system shells, the presentation of information using web, network, and

hierarchical structures, and the use of parallel machines.

Local Area Network Administration Services

The Computer Applications Group provides, through "rental" arrangements, administration support for local area networks. This is a complete service that includes not just continual, personal, on-site support of heterogeneous network hardware and software, but also workstation consulting, diagnosing, and debugging of system problems, software updates, system backups, system security, additions to or reconfigurations of networks (with Groups C-4 and C-5), and user orientation. Through this service, supported users benefit from the expertise of other C-Division, Laboratory, and vendor personnel.

Communications Custom Services

The Communications Group designs, implements, and maintains communications systems for users to the extent desired by each organization. We strive to design and implement systems that not only meet current user requirements but also are extensible to meet future requirements and can interface with other communications networks.

Computer Documentation Support Services

Computer Documentation Support includes computer graphics, documentation preparation, and data preparation services for C-Division and Laboratory personnel. We work closely with programmers throughout C Division debugging and evaluating new software and hardware before it is released to the user community.

High-quality computer graphics are prepared for use in publications and presentations. Computer graphics utilities available

on the ICN computers and on microcomputers are used to generate slides and viewgraphs that range from simple to highly technical and complex diagrams. Using a film processor, we produce high-quality color or black-and-white transparencies and prints from a variety of media.

In FY 1989 the staff of groups C-2 and C-6 worked closely to implement the Computer Graphics Laboratory (CGL). As Laboratory staff become aware of the capabilities within the CGL, demand for its services increases. Because of the wide interest in the many new capabilities of the CGL, it is highlighted in a feature article in Section 4, *Facilities and Capabilities*.

In FY 1991 we plan to expand our knowledge about graphics software for microcomputers, and we will continue to keep abreast of advances in high-technology output devices.

Publication-quality technical reports and computer documentation that can be read online or produced as hardcopy are prepared using computer-based formatting tools. In FY 1990 we acquired more knowledge about microcomputers for generating specialized documents.

In FY 1991 we plan to expand our knowledge of desktop publishing software for microcomputers and keep abreast of the advances of new hardware and software for converting and scanning documents.

We enter classified and unclassified text files and computer programs on mainframe computers using the CTSS, UNICOS,

UNIX, and VMS operating systems. Data can also be entered into INGRES and System 2000 databases, DECate spreadsheets, and IBM PC and Macintosh spreadsheets, databases, and text files.

With access to workstations on LANs and microcomputers, and with a continued awareness of state-of-the-art computing capabilities, we are able to accommodate most customers with faster turnaround time.

FACILITIES AND CAPABILITIES

Our strategic planning goal for facilities and capabilities is to provide world-class scientific computing, network, and communications facilities. In previous sections we described how important it is to recruit and retain the best possible people and to maximize our clients' productivity by providing needed services. What we are probably best known for, though, is the facilities that we make available to the users.

We are involved in a number of projects that will help us move

closer to our goal. Our network is a source of great pride; we have a reputation for incorporating state-of-the-art equipment into our complex, powerful network, and we do it in such a way that users of the computing and communications equipment are rarely impacted adversely. We are in the process of upgrading the supercomputers, and our efforts continue to enhance our network for the 1990s. These projects are described in more detail in this section.

Los Alamos Integrated Communications System

Our communications planning is guided by one central theme: provide an effective, secure, and economical communications utility that helps maximize the productivity of Laboratory employees in accomplishing the Laboratory's mission. The utility serves as a foundation of a hierarchy of communications services that span a wide range of capability and reliability. It is based on *demonstrated*, state-of-the-art technology and is characterized by extremely high reliability (1 hour of downtime in 10 years), universal connectivity, integrated voice and data services supporting

circuit-switched data rates up to 64 kbit/s and packet-switched rates up to 192 kbit/s, extensibility, closed-circuit television, and high-bandwidth (several gigabits per second) trunks. On top of this utility we are building a data communications network having two levels. The first level emphasizes reliability and availability; the second level emphasizes speed.

LAICS is a two-phase, 6-year project that will provide the communications utility described above. The project phases are (1) service continuity and (2) system design and implementation.

Phase 1 is complete and assures the continuity of telephone services following the AT&T divestiture. Phase 2 comprises writing technical specifications for transmission, switching, teleservices, and system management subsystems and for vendor services to support LAICS following its implementation.

Draft specifications for LAICS were written in FY 1987. In FY 1988 vendor and Laboratory comments on the specifications were solicited and incorporated into the specifications. Major milestones for LAICS include the following actions:

- Evaluate proposals by August 1990
- Award contract by December 1990
- Begin LAICS implementation by March 1991
- Complete LAICS implementation by February 1992

LAICS communications nodes.

Laboratory Data Communications Center

The Laboratory Data Communications Center is at long last a reality! The dedication ceremony was held on August 25, 1989, and C Division began to occupy the space shortly thereafter. At the dedication ceremony, Senator Pete Domenici noted the importance of computing and communications to the country and that he expected the national laboratories to play a leading role in keeping the United States at the forefront of these technologies.

A facility like the LDCC takes a very long time to become a reality. C Division began the quest for this building in 1982 when it proposed the LDCC as a construction project. Construction funding at the Laboratory is not abundant,

and each year the proposals are reviewed and prioritized. In 1984 the LDCC was the Laboratory's highest priority construction project, and the process began to receive construction authorization from DOE. Initial funding was granted in FY 1986, with subsequent funding in FY 1987 and '88. Senator Jeff Bingaman was among the dignitaries "breaking ground" for the facility on September 2, 1987. Almost 2 years later, the building was dedicated.

The LDCC is a three-story building with large open areas similar to those in the CCF. The floor space is approximately 95,000 square feet. It is located north of the CCF across Coronado Avenue. An enclosed passageway

across Coronado Avenue connects the LDCC to the CCF.

The first floor of the building houses the Laboratory's integrated data- and voice-communications equipment (part of the LAICS project). This includes a computer-based Private Automated Branch Exchange (PABX) and associated test equipment. The first floor also houses building support equipment including electrical substations and switching equipment, water chillers and pumps, air conditions, and an uninterruptible power supply.

The second floor of the LDCC provides space to expand the network so that it is capable of servicing our growing user community. It is the home of much of the equipment that is required for the NET90 project. It also contains modems, multiplexers, line-driver amplifiers, dialup equipment, jackfields, wiring frames, and diagnostic test facilities.

The third floor of the building will house many of the supercomputers in the ICN. The first supercomputer to be installed in the building was a CRAY Y-MP in May 1990.

The second and the third floors could be considered to be an extension of the present CCF. Supporting laboratories and office space are located on each of the floors.

At the dedication ceremony for the LDCC, Laboratory Director Sig Hecker and C-Division Leader Norm Morse shared the dais with United States Senator Pete Domenici (standing), who called on the Los Alamos scientific community to continue its pursuit of the types of research that will enable the United States to maintain a competitive edge in the world market.

Supercomputer Upgrades

The Laboratory has important supercomputer applications that require much larger main memories than the 32 million words available on our Y-MP8/832 systems. As authorized by the Laboratory's Senior Management Group, C Division has responded to this requirement with a five-part acquisition that quadruples the available memory on all supercomputer systems in the ICN.

The first phase increased the memory of the ACL's massively parallel CM-2 system from one-half to 2 Gbytes and acquired a second identical machine for the Secure partition. The other four acquisition phases provided upgrades to four larger memory CRAY Y-MP systems, all of which are expandable to 2-Gbyte (256-

Mword), 8-processor configurations.

Two of these phases provided CRAY Y-MP8/8128 systems, each with 8 processors and 128 Mwords (1 Gbyte) of main memory, replacing the first two Y-MP8/832 systems we acquired in 1988 for the Secure partition of the ICN. These systems, delivered in June and September of this year, are the largest supercomputer systems currently available from Cray Research. The Secure partition now consists of two systems running CTSS, and a 64K processor CM-2 system with 2 Gbytes of main memory. Early in 1990 a CRAY X-MP/48 will be installed in the Secure partition as a UNICOS transition machine.

Two other phases of the acquisition retire our CRAY X-MP

systems in the Open and National Security partitions of the ICN.

Those systems will be replaced with two CRAY Y-MP8/264 UNICOS systems in October and November. Both systems will be initially configured with two processors and 64 Mwords of main memory, and both will be capable of being fully configured to 8-processor, 256-Mword configurations to meet future requirements.

CRAY X-MP system replacements include the X-MP/416 running CTSS in the National Security partition and the X-MP/24 running UNICOS and the X-MP/48 running CTSS in the Open partition. Machine 5, the DNA X-MP/416, is currently slated to remain in the ICN, and the X-MP/24 will leave as soon as the first Y-MP8/264 arrives. The remaining X-MP/416 and X-MP/48 machines will remain for a period of time as UNICOS transition machines.

The supercomputer upgrade not only provides more CRAY-T equivalents in the ICN, it also provides a massively parallel capability in the Secure partition. The proposed configurations are fully expandable, so while they quadruple our present large-problem capability, they afford us future opportunities to expand as well.

Central Networking Testbed

Two C-Division groups, the Communications Group and the Network Engineering Group, are working together on a joint project called the Central Networking Testbed (CNT). Their mission is to evaluate networking systems for conformance to minimum requirements for proper operation within the Laboratory networking environment, and to share this information within the Laboratory to network planners and implementors. Networking systems refer to networking hardware and software technologies, protocols, and C-Division central networking services. A major goal of the testbed is to provide direction for incorporating commercial products and open standards with the very specialized requirements of the ICN, creating a seamless and transparent networking environment.

During the course of the project, we expect to evaluate all ICN networking systems for minimum integration and interoperability requirements before they are put into production. The intent is to provide developers with feedback that prevents gross faults from appearing. The people working on the CNT do not have any regulatory or planning role; they operate in an advisory capacity to C-Division implementors and planners, as well as to others around the Laboratory. In addition, they act as a focus for integration requirements and perceived needs in networking systems by fostering discussion necessary to identify minimum networking requirements. Member of the CNT project will work with the network planners and implementors to define these

requirements and will be responsible for synthesizing them to meet evaluation needs.

Some current projects include:

- High-speed physical and data link technologies such as FDDI, HIPPI, and token bus alternatives for the current broadband network
- Multiprotocol, high-performance routers
- Local and remote Ethernet bridges
- Terminal servers with and without customized security
- Serial link networking technologies

Summary and comprehensive reports will be made available to validated users of the ICN electronically, and some version of the summary report will be published in the *Computing and Communications Division News*.

CNT project suggestions and comments are welcome from any interested party.

This diagram represents one possible configuration of central supercomputers, network routers, and general-purpose computers, including workstations and personal computers. The CNT tests the interdependencies and intricacies of each network protocol layer.

Network Modernization Project (NET90)

The Network Modernization Project (NET90) was chartered to develop a long-term plan for modernizing the ICN. As part of this long-term plan, the project defined a set of goals that would provide a framework for future developments in the network.

Before getting into details of NET90 itself, a description of the network architecture and protocols would be appropriate.

ICN Protocols—Today and Tomorrow

Over the next few years C Division plans to move away from its Los Alamos-developed network architecture and protocols. Current plans are to support only the Transport Control Protocol/Internet Protocol (TCP/IP) and Open Systems Interconnect (OSI) protocols in the ICN and to use, as much as possible, only commercially developed communications hardware and software. This will mean better and faster access to C-Division services using standard, vendor-supplied networking software on workstations or PCs.

Background

From its inception in the mid 1970s, the ICN architecture was designed for a packet-switched network. The original protocol suite was locally designed and implemented primarily to accommodate our security requirements. These Simple Intermachine Protocol (SIMP) based protocols include

- SIMP, which provides a basic datagram service;
- KCC, which provides a virtual terminal service; and
- AFT, which provides a file transport service.

In the late 1970s the VAX was selected as the Laboratory-standard distributed processor. Many organizations were using this system to do local computing, and we needed a way for them to access CCF services. XNET software was developed. This software took the ICN SIMP protocols and applications and used them on top of DECnet. This system used point-to-point communication links connected to the XNET gateways. This has evolved to using Ethernet and high-speed communication links to form a large DECnet wide-area network.

In the mid 1980s, the move to a workstation environment began. Sun and Apollo workstations were selected as the Laboratory-standard scientific workstations. The SIMP-based software was moved to the UNIX platform, and TCP/IP was used as the transport protocol between the workstation and the XNET gateway. The workstation environment has grown tremendously, and many different UNIX platforms and IBM PCs are now running the SIMP software. The Los Alamos TCP/IP network is now tied to the major national internets, including MILnet, NFSnet, Bitnet, and ESnet.

Although the SIMP-based protocols are widely used at Los Alamos, they provide limited functionality, require local support, and are difficult to port to new environments. They tie us to the applications that Los Alamos develops and prevent us from using the applications the rest of the world develops. We intend over the next few years to move our SIMP-based network services

(CFS, PAGES, etc.) first to TCP/IP, and then to OSI.

New Directions

Today we live in a world where networks are widely used. There are several nonofficial U.S. or international standard protocol suites available today; TCP/IP and DECnet are the most widespread in the scientific community. These protocols are available on a wide range of platforms. OSI is the official national and international standard protocol suite. This new protocol, similar to IP, is not yet widely available. These standard protocols provide or will provide support for new distributed computing capabilities such as X-Windows, network file systems, and remote procedure calls that are important for our customers' current and future needs.

TCP/IP

TCP/IP was developed by the Department of Defense ARPANET project. This protocol is available for almost all computer systems being sold today and is widely used in scientific and research fields. The specifications for the implementation of TCP/IP are available to anyone who requests them. This provides an open-system concept where all vendors can compete, which ultimately means a more universal product that can interoperate over many network computer platforms.

OSI

The International Standards Organization has developed the specifications for the OSI family of protocols. These protocols are an international standard and have been adopted by governments around the world. The U.S. government has also directed its agencies to use these protocols for future networking requirements and has developed a specification called GOSIP that describes which

protocols will be used. DOE is working on a GOSIP implementation plan. When this plan is released, we will know better how it will affect Los Alamos.

DECnet

DECnet Phase V will use an OSI transport, and we are currently making plans to route this traffic. DECnet, a protocol suite developed by Digital Equipment Corporation, was the first non-local protocol supported in the ICN. DECnet is used for communication with VMS distributed processors and other systems that communicate with DECnet. The specifications were published by DEC and are non-proprietary. Other vendors have implemented DECnet for compatibility with DEC systems, but DECnet is not as widely used as TCP/IP. DEC has announced that over the next few years it will migrate DECnet to an OSI protocol base.

TCP/IP for VMS

With the large growth in TCP/IP networking, C Division decided to purchase a site license for a TCP/IP package for VMS. C Division negotiated with SRI to provide a site license for the Multinet package that allows VMS systems to run TCP/IP. This license was purchased with funds provided by a number of Laboratory organizations. Other VMS TCP/IP implementations are also available from other vendors, including Digital Equipment and Wollongong. We strongly encourage our DECnet customers to consider TCP/IP, and we have made available the Multinet distribution on CFS for use on Los Alamos systems. We are sorry that we cannot make this available for non-Laboratory systems, but we are negotiating for reduced prices for systems outside the Laboratory. We encourage users to install this software on their VMS hosts and become familiar with it.

Summary of Short-Term Plans

Over the next several months, you will see announcements of changes to network service software as we set the stage for the move to TCP/IP and OSI. We will try to provide equivalent functionality where possible and will make the migration as straightforward as possible. We hope to provide users with more robust solutions to their future networking needs through open, interoperable network protocols and services.

NET90

With a basic understanding of the network architecture and protocols, we can now turn our attention back to the goals of the modernization project. These goals are briefly stated as

- Improving performance
- Improving reliability, availability, and maintainability
- Enhancing security in a distributed environment
- Moving toward an open network architecture

Performance

A major goal of the network modernization work is to upgrade the performance of the ICN. This requirement is driven by the increasing computational power of supercomputers in the ICN, the increasing requirements for high-speed graphics, and the increasing use of distributed computing. We anticipate a hierarchy of performance requirements for the ICN. High-speed graphics applications will be at the top of the performance hierarchy, followed by local area networks (LANs) and terminal communications. Design considerations for the network include the tradeoffs between high performance and reliability. We believe that a majority of the clients served by the network value reliability over performance.

Major projects for FY 1990 that address performance in the ICN are

- Explore the use of FDDI in the ICN. During 1990 a new LAN technology called FDDI will emerge. FDDI is a 100-Mbit/s fiber-optic token ring standard that is in the final stages of completion by ANSI. A number of vendors, including DEC, IBM, Hewlett-Packard, and Sun, have announced products for FDDI, and the first hardware components are being shipped. At first, the technology will be too expensive to replace Ethernet, but it will complement Ethernet in high-performance applications. Over the next few years, FDDI will bring the type of performance we now have between supercomputers and CFS to the desktop workstation.

- Implement higher performance network interfaces. To support the increasing demands of the workstations user and new network technology such as FDDI, access to supercomputers and other CCF services must be improved. Our XNET gateways will be upgraded to handle data transfers between workstations and supercomputers at several times the current data rates.

- Develop a network testbed with gigabit per second communication links and crossbar switches. This testbed, called the CP* testbed, will utilize the emerging ANSI standard HIPPI to connect supercomputers into a network using HIPPI-based crossbar switches. A 32 by 32 crossbar switch can provide an aggregate bandwidth of over 50 Gbit/s. The testbed will be used to develop applications that can utilize the tremendous bandwidths of such a network. It will also be used to evaluate other HIPPI equipment such as the HIPPI-based components being developed by Network Systems, IBM, Thinking Machines, and Cray Research. This evaluation will help us design the next generation network architecture for the ICN.

- Develop a national gigabit network. We are working closely with vendors and other research organizations to develop a very high speed prototype communication network to link Los Alamos with CalTech, Jet Propulsion Laboratory, and the San Diego Supercomputer Center (see the section on collaborations for more details).

Reliability, Availability, and Maintainability

We recognize that the ICN has become a critical resource for many Laboratory organizations. Failures in the ICN that result in denial of service seriously affect the ability of many organizations to function. Improving reliability, availability, and maintainability is an important goal in all our projects. During the upcoming year the projects specifically aimed at this area are:

- Merge the AAC/KCC functionality. During FY 1990 we plan on moving the AAC functionality to the KCC. This will mean the when a user presses carriage return on his terminal, he will be immediately connected to KCC. From there, he can login to a CTSS system or select another service such as OFVAX. By eliminating the intervening system (the AAC), we are providing users with a simpler, more reliable interface to ICN services.
- Added "heartbeat" to FT links. The "heartbeat" is simply a test message containing a prescribed pattern sent by an FT or SYNC at regular intervals. It provides a first-level, online, and continuous diagnostic to verify correct operation of the high-speed parallel interfaces, cabling, and FT/SYNC hardware. As such, it provides error detection of intermittent hardware problems that are now difficult to detect and correct. The test message, after looping through the worker computer, is checked to insure

that the pattern sent is identical to the pattern received.

- Set up the Network Operations Center (NOC). To provide a focal point in managing the ICN and its associated LANs around the Laboratory, the NOC was established. It provides a variety of network management services including network troubleshooting and maintenance; host address assignment and database; network email address coordination and assignment; network traffic statistics monitoring, analysis, and reporting; tracking and coordination of software and hardware changes on all ICN hosts; and assimilation and dissemination of network and security-related information through a central point. The NOC provides on-call support for the ICN 24 hours a day, 7 days a week.

Security in a Distributed Environment

Another major goal is to maintain security in a distributed environment. The FY 1990 projects in this area are:

- Implement an improved ICN authentication scheme. We will be developing an improved user authentication scheme to strengthen our ability to identify users, protect the ICN, and make it easier for the ICN users to access a number of CCF services without having to retype their ICN password for each service.
- Intrusion detection. The goal of this project is to aid the system security personnel in the detection of penetration and abuse of the network. It is being developed on the hypothesis that ICN security violations can be detected by monitoring network audit records for both abnormal and invalid patterns of usage. The long-range goal is to provide real-time detection of attacks, automated or otherwise, on the ICN.

Open Architecture

A reality in our environment is that our systems are and will continue to be heterogeneous, but they must be able to communicate easily and efficiently. In the past, we have developed our own protocol suite and implemented it across this heterogeneous environment. We do not believe this is the appropriate strategy to obtain interoperability in the future. In addition, our protocols do not provide the architectural framework to support the distributed environment that we envision for the future.

These reasons suggest that we replace our own protocol suite with industry standard protocols. An objective to achieve to this goal is to support DECnet and TCP/IP protocols today, continue development of TCP/IP on hosts in the CCF, and begin planning and developing a migration path to the ISO protocols as they become available. This recognizes the leverage that we will obtain by use of the DoD protocols and DECnet in supporting scientific workstations and other commercial systems.

At the same time, we recognize that vendors supporting these protocols are moving toward the ISO suite of protocols. All implementations of TCP/IP throughout the ICN must be done in such a way as to allow the integration of ISO protocols as painlessly as possible.

A specific project to meet this goal is:

- Secure IP router. A router to the ICN is being developed that will allow workstations to directly access CCF services (Crays, CFS, PAGES, etc.) using standard TCP/IP protocols. This router is being developed as a prototype to develop the appropriate security mechanism necessary to protect the CCF and is described in more detail in the section on distributed computing (see Section 5, *Leadership*).

Computer Graphics Laboratory

The Computer Graphics Laboratory (CGL) was established last year to provide users with access to specialized computer graphics systems and capabilities. It also provides users with access to advanced graphics capabilities more quickly. In short, the CGL's purpose is to help users with any computer graphics problem.

The CGL contains a wide variety of hardware, including scientific workstations, gray-scale and color scanners, color printers, high-quality laser printers, Macintosh graphics systems, film recorders and video recorders, as well as the software to run these devices.

The following services are available in the CGL, but by no means is the list complete. People with specialized graphics problems are urged to call the CGL for

advice and consultation. Many problems that once were "impossible" have been solved through the use of CGL equipment, service, and personnel.

- Animation of scientific visualization problems for previewing or presentation purposes
- Video creation of computer graphics from Cray supercomputers and scientific workstations
- Video editing and audio dubbing of computer graphics video recording
- Creation and manipulation of charts and diagrams
- Macintosh to PAGES film output
- Graphics file formats translation
- Color paper and transparency output

- Gray-scale and color scanning and editing
- Direct PostScript output to black-and-white and color devices
- Digitizing
- 600-dpi output from Varityper printers
- Desktop computer graphics applications

The latest technology print and graphics hardware and software will be evaluated in the CGL for possible installation on PAGES. One such item that was evaluated in the CGL and later moved to PAGES was the QMS ColorScript color printer.

Additional services are already planned for the CGL. These include:

- High-quality color output from the Canon color laser printer
- 4 X 5 and 8 X 10 film output
- Frame by frame video recording system to allow real-time playback of scientific animations
- High-quality typeset output

As the CGL grows we will identify other equipment that should be included and provide for working in all partitions of the ICN. In addition, as technology changes, so will equipment brought in for evaluation.

The CGL benefits many people. The user community benefits because they have access to additional resources that might otherwise be months or years away from production status. Users with sophisticated requirements have access to high-quality input and output devices and state-of-the-art software for producing their documents and computer graphics. Users are able to take advantage of the latest technology without having to buy equipment that is too expensive to replicate at each user's site.

The use of the CGL continues to grow as people find out about the services provided. Together with its customers, C Division will determine what equipment and services are beneficial and make them available in the CGL.

Equipment in the CGL provides new opportunities for viewing and editing text and graphics images. Metafiles can be viewed in raster format via the framebuffer on a high-resolution monitor and later converted to a television format that can be recorded onto 3/4-inch tape for editing. After editing, the images can be transferred to VHS tape, which can be shown at talks where many people can view the simulations.

Research Projects

Research is necessary to maintain and enhance our position as a world leader in the fields of computing and communications. Many research projects are under way in C Division, and we cannot begin to describe them all or go into any significant detail about any one project within a publication such as this. Therefore, we have selected a few projects to highlight this year. Technical papers provide much more detail about these and other projects, and a list of recent publications concludes this section of this report.

MediaView, A Multimedia Digital Publication System

Digital publications of many forms are finally a reality. The abundance of high-resolution workstation displays and low cost of mass storage systems enable the realization of an electronic reading environment that far exceeds the capabilities of traditional printed publications. For example, digital publications can include not only text, line art, and still images, but also sound, video sequences, and computer-produced animations. Cast in digital form, the mathematical content of a document can be symbolically and numerically manipulated. Thus, one can experiment with the mathematics, derive new results, and simulate different situations with different parameters. Computer-generated images can be explored by moving the eye-point, changing the lighting conditions, or using the underlying model with one's own algorithms. An animated scientific visualization can be incorporated in a digital publication and

examined in situ or mailed electronically to a colleague for independent study. Clearly, the possibilities for interacting with information are exciting and boundless. Imagine, for example, having such access to the entire Laboratory library collection. Scientists would immediately enjoy dramatically new and efficient ways of working.

While the opportunities for digital publication are great, there are many associated problems. Among these are developing centralized document databases, distributing documents to a workstation, and defining an appropriate user interface to the publication. Research into these and other areas has fundamental relevance to all Laboratory programs. Successful resolution of these research issues can result in a multimedia digital publication system that would profoundly affect the way all staff members interact with information and each other. By acquiring expertise in this area, Los Alamos will be at the forefront of this exciting new field. Specifically, the Laboratory would be well positioned to take the lead in realizing the first digital scientific library.

A brief description of this interpersonal multimedia visualization system follows.

The Concept

Imagine a document, rich in textual and graphical content, which speaks, shows animated sequences, and helps you manipulate mathematical expressions. In an instant you can find any word or passage, excerpt text and figures for incorporation in other

documents, and add voice, textual, or graphical annotations non-destructively. This "document" is not, of course, the usual kind; it is entirely digital and resides on a sophisticated personal computer. Such a document might look like the image on the following page.

MediaView, developed at Los Alamos National Laboratory, is the system that makes it possible to produce a document like this. MediaView is a computer program that provides a generic infrastructure for authoring and interacting with multimedia documents. With MediaView, documents can be developed for professional or personal use, technical or administrative. Because MediaView provides such a rich and flexible environment, professional applications range from physics to law, and administrative uses span the spectrum from CEO to secretary.

How It Works

Taken individually, the components of multimedia technology appearing in MediaView are not new. There are already several programs that manipulate text and graphics, including high-resolution digital color images. A few systems also provide for embedded voice messages, but most are experimental and require expensive add-on hardware. Animation, too, can be produced by video add-in cards and video-disk players. Finally, textual annotation is even available through the Post-It metaphor in some word processing systems. What is not available, however, is a completely integrated system, hardware and software, that combines all of these features. This is one of the most important attributes that MediaView brings to digital publishing, but there are many more.

One example application of MediaView featured here is a collection of technical papers

An example of what a digital document might look like.

from proceedings of recent SIGGRAPH conferences. These papers offer an especially rich and challenging testbed for MediaView because of their extensive mathematical and graphical content. Installed in MediaView, this collection of papers has the following features:

- Keyword retrieval of any word or phrase in the volume and random access to that part of the document. Any collection of text can be copied and then pasted into a text editor or word processing program.
- Figures that accompany text can be viewed and, if desired, extracted for processing by another program. This capability applies

to simple drawings or 24-bit/pixel full color images.

- Equations that appear in a MediaView document have a backing format that is compatible

By double clicking on an equation, its semantics are conveyed to Mathematica, where the user can perform a variety of symbolic and numerical operations. This image was produced using mathematics from the SIGGRAPH'88 paper "Reconstruction Filters in Computer Graphics," by Mitchell and Neftci.

with the Mathematica language. Thus, by clicking on an equation, its semantics are conveyed to Mathematica, where the user can perform a variety of symbolic and numerical operations (see figure).

- While perusing a paper the reader may activate an icon that indicates an animated sequence accompanies the text. A window will open and the animation will be displayed there. Controls attached to the window allow single stepping and selective freezing and extraction of frames. The example shown on the following page is from the SIGGRAPH'89 paper *Analytical Methods for Dynamic Simulation of Non-Penetrating Rigid Bodies* by David Baraff.
- Another icon style indicates that an audio insert is available for listening, e.g., part of the question-and-answer session. The sound fragment can be played back through a built-in speaker or it can be captured by a sound editor for incorporation in another multimedia document.
- The reader may wish to make notes in the "margin" of the document. These could be voice annotations or in the form of textual or graphical "stick-on" notes, which are logically attached to a position in the paper. Subsequent readers can, if permitted, hear and see those comments. Thus, a teacher who is using MediaView in a course can

Another possibility for analysis and visualization is to perform a simulation wholly within MediaView and use its dataset animation facility to view the result. Once the simulation is complete, a window will open and the animation will be displayed there. Controls attached to the window allow single stepping and selective freezing and extraction of frames.

expand on points in the paper in situ.

Other Examples

Other uses for MediaView range from a customized multimedia newspaper to an interactive children's book. In addition, MediaView can accommodate maps, CAD diagrams, and a custom component, which can be any multimedia entity that is not standard to MediaView but is required for a particular application. Finally, anything that MediaView can display can be printed or saved as a PostScript file, which is how all the images included here were captured.

High-Performance Parallel Interface

Los Alamos has been leading an American National Standards Institute (ANSI) effort to establish a "standard" High-Performance Parallel Interface (HIPPI) capable of transporting data at 800 Mbit/s to 1.6 Gbit/s. The speed range of the 800-Mbit/s HIPPI is 80 times that of Ethernet and 16 times that of HYPERchannel, today's common high-speed networks. We plan to use the HIPPI as the basic building block to network the Lab's future supercomputers and to provide movie-mode high-speed graphics to the users.

This effort has paid rich dividends with many vendors, including IBM, DEC, and Cray, supporting the standard and developing hardware. By using a common standard, we can interconnect equipment from different vendors without the expense of special adapter boxes. The definition phase of the standard has been completed, and the approval phase is nearing completion. It is estimated that the HIPPI

will become an official standard in late 1990.

Additional work is being done to develop some of the networking protocols that will run on HIPPI links as well as a way to serialize the data for transmission on fiber optics. Both of these efforts are also being done under ANSI committees in order to define standards for their implementation.

High-Speed Networking

The Laboratory's ICN has been supporting high-speed networking over the past years with in-house-designed links between supercomputers that operate at data rates up to 50 Mbit/s. The switches in the network that control the data flow on these links are currently high-performance minicomputers that have I/O ports that support the link data rates given above. At the same time, the network provides high-speed data links to some users at 300 kbit/s for graphics applications along with three 30-Mbit/s links for special-purpose frame buffering.

There is a demand for higher speed links, both within the network and to users' offices to support various programmatic efforts. The goal for the network is 1-Gbit/s links. With the ongoing HIPPI effort, this goal is within reach. The HIPPI will provide links of 800 to 1600 Mbit/s. The HIPPI, however, only addresses computer channels and data links; it does not address networking.

To provide the remaining components to a network solution, we have developed a new network switch called CP*. A CP* uses the HIPPI as its interface and provides a crossbar switching capability for up to 16 HIPPI connections. A Crossbar Switch (CBS) implies multiple connections operating simultaneously, something a computer (bus-based)

switch cannot do. Meanwhile, Digital Equipment Corporation has developed a distributed protocol processor, again using the HIPPI for interfacing, that, together with the CBS, performs the full CP* function. The DEC system is called the Crossbar Interface (CBI) and essentially off-loads most of the network protocol overhead from the host computers. Connecting multiple CP* systems together provides connectability for a large number of hosts and users. Such a system is being called the Multiple Crossbar Network (MCN).

We envision the HIPPI and MCN network as an addendum to our present network, providing higher speed transmission for the services that can take advantage of it. As computers and workstations become faster in the future, more and more units will be interconnected using these new networking components. Our goal is to provide a growth path for the Los Alamos ICN. By stressing the use of the HIPPI as a standard, we believe that Los Alamos' needs can be almost completely satisfied by commercial equipment, which is in concert with our management philosophy of buying equipment rather than building it.

Our goals for the near future are to put together a CP* testbed using the CBS and CBIs and interconnecting a variety of equipment from IBM, Cray, Thinking Machines, Sun, Tektronix, and others. We are planning for some form of this crossbar network to be integrated into our current network by 1991 to provide users with the increased capabilities.

Applied Mathematics

A fundamental activity throughout the Laboratory is the modeling of physical phenomena. Such models proceed in stages from an understanding of the underlying processes to a repre-

sentation of these processes by equations. Analytical and computational techniques are then applied in a complementary fashion. Analytical techniques yield valuable insights into the solutions of these equations, while detailed information over realistic parameter ranges very often requires that these equations be solved approximately on the computer. The understanding of how to effectively use analytical and computational techniques is one of the domains of applied mathematics and requires expertise in computer science, mathematical physics (chemistry, economics, etc.), and mathematics.

In order to achieve optimal use of computer resources in numerical modeling and to provide the complementary expertise in analytic techniques, C Division supports research and consulting in applied mathematics. The primary foci of this ongoing program are (1) the approximate solution of differential equations; (2) mathematical analysis of models of complex processes; (3) advanced computational algorithms; and (4) collaborations with Laboratory scientists and a distinguished group of visitors. Below we describe just a few efforts in these areas over the past year and indicate the direction of future work.

First Kind Integral Equations

Integral equations of the first kind (FKIs) are an important class of ill-posed problems arising in the data interpretation of a number of experiments. Over the last several years, we have developed a theory that quantifies the ill-posedness of many of the problems, with and without data error. This work is applicable both to the numerical solution of data recovery problems (e.g., tomography) and perhaps to the design of experiments themselves. Active areas in the solution of first kind integral equations include:

the development of new algorithms for solution, parallelization of old algorithms, and new applications in collaboration with others.

Currently, extremely efficient and accurate methods for solving FKIs with noise do not work well (a) in the presence of constraints or (b) on parallel machines. On the other hand, several standard iterative procedures that converge very slowly can handle constraints and probably can be made parallel. The need to solve very large problems with constraints, e.g., in 3-D tomography, necessitates development of new algorithms whose computational complexity, accuracy, robustness, and parallelization are optimal. In addition, new applications of FKIs continually arise. Most recently, we have employed FKIs in novel ways (see the following text on learning theory).

Study of Adaptive and Composite Mesh Techniques for the Solution of Partial Differential Equations (PDEs)

The solution of PDEs with complex localized behavior and in complex geometries can benefit from improved techniques in both adaptive and interactive finite difference grid generation. In many physical problems, such as fluid flow, complex behavior is localized, and so improvements in efficiency and accuracy of computations can be achieved by making local refinements or clusterings of grid points. We have started to study methods for the solution-adaptive generation of both hierarchical refined meshes and evolving curvilinear meshes to accomplish this local refinement.

In the evolving curvilinear mesh approach, an auxiliary PDE or system of PDEs is derived that describes the motion of mesh-points. This PDE can be derived from the requirement that the integral of some weight function related to an error estimate for the

solution be approximately the same for all mesh cells. A new approach is the derivation of a metric tensor weight function, which basically tells what the length of each mesh edge should be. This technique applies more geometric constraints to the mesh than previous approaches, and promises to help avoid the problem of very distorted grids.

For problems in which complex 3-D geometries occur, we have developed the composite overlapping grid approach for solving PDEs. A current research topic is how to best choose appropriate component grids in order to make up the overall 3-D composite grid. Another issue that arises with composite grids when used for compressible flow calculations is the question of conservative interpolation between component meshes. We are in the process of addressing this issue and believe that progress can be made in obtaining a solution.

Mathematical Analysis of Models of Complex Physical Processes

The mathematical modeling of complex physical processes by deriving and analyzing equations that describe the processes is an important area of our effort. In order to develop effective and useful models, a basic scientific understanding of the combined effects of the relevant physical processes is required. This understanding can be used to develop a mathematical model using PDEs or integro-differential equations (IDEs). The model can then be analyzed using asymptotic techniques and by developing numerical methods to approximate and solve the model equations. Models can provide a framework within which the relative importance of various contributing physical effects can be assessed. This often provides information that can be used for engineering design studies, or to

gain information on how to control the physical processes of interest.

An example of a physical process that we have been modeling and continue to explore is the penetration of a solvent into a viscoelastic medium. We have derived PDE and IDE models of nonlinear diffusion type for this situation and are analyzing them both asymptotically and numerically. An unusual situation that arises with many nonlinear diffusion problems is that information propagates with a finite speed rather than the infinite speed that characterizes linear diffusion problems. In addition to the physical models, we have proposed a number of model nonlinear diffusion equations in which this effect can be studied and understood more carefully.

PDEs of reaction-diffusion type arise in many interesting physical situations. An example is the modeling of hazardous chemical contamination of soil and groundwater. By understanding the processes involved we can develop models that can be used to understand how to evolve cleanup strategies for hazardous chemical waste sites. We have extensive experience in the development of models of reaction-diffusion type, particularly those that describe flow in underground aquifers. We expect that the analysis of these models will require singular perturbation techniques, and numerical modeling that will require front-tracking, and adaptive and composite grid techniques.

Geometric Numerical Integration

Many dynamical systems possess symmetries (geometric and dynamic) and/or invariant geometric structures (symplectic, Poisson, volume preserving, etc.) that separate them from the generic dynamical system. Classical analysts always tried to respect

these symmetries or structures; however, traditional numerical analysis of differential equations has typically ignored them. Our work presently consists of developing numerical integration algorithms for such systems. Applications consist of classical Hamiltonian systems, the Vlasov-Poisson and Vlasov-Maxwell equations of plasma physics, the equations of ideal compressible fluids, rigid body motion, coupled rigid body motion, two-fluid electrodynamics, and magnetohydrodynamics (MHD).

Work on symplectic integration for classical Hamiltonian systems clearly demonstrated that the canonical structure of these equations must be preserved when doing numerical integrations. The modern terminology for this canonical structure is called symplectic structure. Recent results show that many important equations of mathematical physics are Hamiltonian with respect to a Poisson structure.

The most elementary Poisson structure is the Lie-Poisson structure, which is generated from a configuration space which is a Lie group, and is relevant to the Vlasov-Poisson equation of plasma physics, the compressible fluid equations, and rigid body motion. We have succeeded in developing a class of Lie-Poisson integration algorithms and implemented them on a fully 3-D moment description of the Vlasov-Poisson equation.

Tackling this problem required the general development of Lie-Poisson integration algorithms and the solution of the truncation problem. Truncation was required for computations since the phase space of the original problem was infinite dimensional. Solution of this problem required the notion of a symplectic groupoid and suggests an approach to the appropriate Hamiltonian description, truncation, and

integration of the multi-blob problem. Namely, color two halves of the Vlasov distribution red and green, determine the red and green moments, and their appropriate evolution. We are presently working on this problem.

Learning Theory

We have been investigating learning theory and its relationship to interpolation. We can easily explain what we mean by learning theory by considering the one-dimensional case. We assume that there will be given a sequence of (possibly noisy) values

$$g(x_0), g(x_1), \dots, g(x_n)$$

of a function g . After $g(x_n)$ is given, we are given x_{n+1} and are expected to predict $g(x_{n+1})$. Learning is said to occur if we do "better" with increasing n . A learning theorem should be specific about this word "better." For example, neural net theory usually appeals to theorems that state that a particular algorithm will produce errors which go to zero.

We hope to produce simpler and faster algorithms (than neural nets) which lead to theorems which provide more information than mere convergence. In the past, we have proved theorems of this type about linear approximation. Unfortunately, the only way we knew to use linear approximation to approximate nonlinear functions is by using a truncated series of the form

$$\sum_{n=0}^N a_n \phi_n(x).$$

The difficulty with this lies in having to identify what functions to use and how many (N) to use. Recently, we recognized that integral equations provide a clue to fast, accurate approximation by

a class of algorithms we call Linear Nonlinear Learning (LNL). We are not only able to prove convergence of the algorithms, but also can provide a strong bound for the errors made by the LNL algorithms in terms of an interpolating function. We have just started to carry out numerical experiments employing LNL in multidimensional settings. Results have been comparable to neural net experiments on the same data, but our algorithm runs 100 to 1000 times faster. We have also extended the theory to allow approximation within certain prespecified error tolerances. We have also tested this version of the algorithm against neural nets, again with comparable results.

Our plan for the future involves expanding our experiments to a much wider class of data. Current experiments just under way involve data with thousands of dimensions. We also plan to clarify the relationship between these algorithms and other iterative methods for solving integral equations and linear interpolation. In particular, conjugate gradient may provide a further increase in speed. In high dimensions, some algorithms in the class are vastly superior to others. We have some notion about why this happens, but we would like to analyze this more thoroughly. Finally, we will try to find out if there is any reason to continue using neural nets in the face of these apparently superior algorithms.

Laboratory Collaborations

Members of the applied mathematics effort are directly involved as collaborators in many scientific application areas. Many of these application areas have significant potential to enhance scientific computing throughout the Laboratory. In addition, we have promoted interdisciplinary team efforts in scientific computing. These include interactions with

groups throughout the Laboratory, e.g., X-6 (multigrid methods for neutron transport problems, optimization), EES-3 (calculation of rock porosity), EES-1 (finite element approximation of a geophysical problem), C-6 and IAO (image compression), T-10 (the human genome project), MEE-13 (coupled fluid and structure codes), and P-14 (design of instruments). We plan to continue and expand our collaborations within the Laboratory.

Supercomputers as Particle Detectors

Since the mid 1970s supercomputers with very large memories have been commercially available. There are now hundreds of such computers having millions of words of memory. The operating systems on these machines are sophisticated enough to keep track of various hardware malfunctions. In particular, individual bit failures in memory can be detected and recorded as a function of time. There has been some interest in circuit failures in computers as a result of cosmic rays, especially for computers aboard spacecraft. There exists the possibility that cosmic ray showers could also affect earth-based computers. Sources of ionizing radiation other than cosmic rays could also cause small-scale hardware errors.

With the number of supercomputers now in existence, it may be possible to establish whether any memory errors are linked to cosmic rays by tracking bit failures as a function of time for a variety of computers. Any periodic fluctuations on a diurnal or sidereal time scale would be of great interest. Connection of a larger than normal number of bit failures with a particular cosmic event, such as supernova 1987a, would also be significant.

At present we are undertaking a very informal preliminary investigation of the feasibility of match-

ing bit failures with cosmic events. We are currently gathering information on what data are available on bit failures on Cray computers and what research has been done on cosmic ray interactions with memory circuits.

Parallel Processing

This year we worked on applications of the Connection Machine to real physical problems, continued our investigation of the dataflow model of computation, and developed simulation tools for the cycle prefix (Faber-Moore) networks we discovered last year.

Quantum Chromodynamics

Quantum Chromodynamics (QCD) is a natural application for high-performance, massively parallel computers, since it requires large lattices, uses only nearest-neighbor communications, and performs a large number of floating-point operations between communications. We collaborated with workers in the field (Rajan Gupta, Laboratory Group T-8; Greg Kilcup and Gerry Guralnik, Brown University; and Clive Baillie, Caltech) on implementation of state-of-the-art QCD calculations on current high-performance machines such as Thinking Machine's CM-2. This project investigated new algorithms, both for physics and for parallel implementations of the physics, and obtained specific physics results from the CM-2. This work was based on existing *Lisp codes that do both Metropolis and over-relaxed updates of pure-gauge QCD on the CM-2. Our investigations included: (1) stability of current algorithms on 32-bit machines; (2) extraction of the glueball spectrum and string tension from pure-gauge configurations on a large lattice; and (3) generation and analysis of full QCD configurations using dynamical Wilson fermions.

Work on implementing QCD on the Connection Machine has been

going on for over a year. Much ground-laying work was done by R. Brickner (C-3) on the Intel iPSC and FPS T-Series machines in early 1988, in collaboration with Gerry Guralnik and Greg Kilcup (Brown University). In May 1988 we visited Thinking Machines Corporation and successfully implemented a *Lisp version of a Fortran QCD production code. This code appeared to run in the near-GFLOPS range on a 64K machine, and we had hopes of significantly improving performance by using specially written low-level math routines. During summer 1988 we converted further Fortran routines to *Lisp, so that by the fall we had a complete pure gauge, Metropolis/over-relaxed update code running, comprising over 3000 lines of *Lisp. Thinking Machines also released Version 5.0 of their software in the fall, which required significant re-writing of the code to make it compatible.

In late fall, Rajan Gupta, Clive Baillie, and Ralph Brickner wrote a new QCD code based on the Fortran code running at MFE and using techniques developed in implementing the pure gauge code on the Connection Machine. This state-of-the-art code embodies dynamical Wilson fermions, uses an advanced inversion algorithm (pre-conditioned minimal residual/conjugate gradient), and employs the projection operator formalism to reduce the spin degrees of freedom necessary to carry out the innermost calculations. With the 5.0 release of Thinking Machines' software, we estimate the code runs at around 1 GFLOP sustained on a 64K CM2. The final debugging was carried out in early 1989, and with the arrival of the Los Alamos 64K machine, production runs began under the direction of the "Los Alamos QCD Collaboration" (R. Gupta and R. Brickner; Greg Kilcup; Clive Baillie; Steve Sharpe, University of Washington;

and Apoorva Patel, CERN). The code is currently running in production mode at Los Alamos, Thinking Machines, Argonne National Laboratory, and NPAC (Northeastern Parallel Architectures Center, Syracuse, New York).

In September 1989 the collaboration presented physics results and a description of the Connection Machine code in multiple invited and contributed talks at the LATTICE 89 International Lattice Gauge Theory Conference in Capri, Italy. Work performed on the Connection Machine with this code comprised a significant contribution to the international QCD effort. Results include data on the hadron spectrum, and strong interaction effects on a variety of weak interaction processes.

Dataflow Research

We have continued the effort begun last year in research on the dataflow model of parallel computation in collaboration with the MIT dataflow effort.

The purpose of the collaboration with the MIT dataflow group was to begin resource management studies in the dataflow execution model using their functional language Id. Functional languages have historically used single assignment semantics resulting in unbounded memory requirements. The point of the current effort is to find places where heap memory and other resources (such as stack space) can be reused while at the same time the high-level functional semantics can be maintained. Manual mechanisms have been developed to free the memory associated with structured variables in the Id language through the use of language annotations. We are now experimenting with them in our Id applications.

The first study focused on the Laboratory benchmark Gamteb, a

Monte Carlo photon transport algorithm. The algorithm written in *Id* is generally recursive and displays immense amounts of parallelism. However, a 32-particle problem uses about 16K of heap memory and needs stack space for 600 simultaneous procedure invocations. These resources are linear with the size of the problem, so it is unrealistic to run a typical problem with no memory management. However, with the annotations to reuse memory and to bound procedure invocations, we found that the Monte Carlo problem could be run with 6K heap space and 100 procedure invocations. The really astonishing result is that these required resources are independent of the problem size. That is, we believe that we can now run a 40000-particle Gamteb problem with about the same amount of storage.

The next stage of this study will involve different algorithms and also inventing compiler analysis that will ensure the safety of the annotations (that is, annotations that lead to non-deterministic execution will be flagged). In the longer term we hope to provide sufficient data that the compiler writers will be able to do without annotations in most cases of interest.

Simulation of Advanced Networks

Two network simulators have been developed during the past year. The original motivation for this was to quantitatively compare the performance of the newly invented "cycle-prefix" (Faber-Moore) networks with commonly used networks, such as hypercubes, meshes, and multi-stage networks. However, the network simulation tools can also be integrated with other tools to investigate multiprocessor design issues. The dataflow/network simulations described at the end are an example.

The two simulators treat different classes of networks. One applies to "static topology" (ST) networks, in which the sources of messages (referred to as "processing elements" or PEs) and sinks ("memories" for data writes, or PEs for data reads) reside at each node of the network. Examples are cycle-prefix, hypercubes, 2-D and 3-D meshes. The other simulator treats multi-stage (MS) networks in which the sources reside at one side, the sinks reside at the other, and the switches in the middle of the network are unencumbered by sources and sinks. For all network types, the user need only specify a few parameters that define the network, and the simulation code automatically generates the interconnection pattern and data structures. Both simulators assumed packet switching, with a full packet passed in a single switch "cycle." The internal structure of the switch is ignored; at present messages are not pipelined in the switch. Simulated workloads are statistically generated: the user specifies the probability that a message will be injected at each source node during each network cycle. Destinations nodes are selected randomly. Work is in progress to provide more realistic workloads (see the discussion below regarding dataflow). Typical diagnostics are acceptance probability, network throughput, network loading, averages and distributions of message delivery time (age), and number of collisions per message.

The implementations of ST networks that we have been studying use desperation ("hot potato") routing. This is possible because multiple paths exist between any two nodes in ST networks. No buffers are provided to hold messages at nodes, so every message must be routed somewhere on every cycle. Two options are used to arbitrate

collisions between messages requesting the same network channel: (a) the older message or (b) the message closer to its destination is given priority. The highest priority message is routed along the shortest path to its destination. Lower priority messages are assigned channels either at random or according to their second and lesser choices.

The MS network simulator currently treats buffered "delta" networks, a large class of shuffle-exchange networks. Delta networks have a unique path from any source to any destination, so buffers are needed between stages to hold messages that cannot proceed because of collisions. On each network cycle, one message at each input port of every switch can be processed if no collisions occur; conflicting requests are arbitrated in a round-robin fashion. The MS simulator is newer and presently has fewer options than the ST simulator.

As a step toward more realistic workloads, we have linked the MS simulator with the GITA dataflow interpreter (both are written in C). GITA can now pass all read and write requests to "I-structure" arrays (shared data structures that reside in global memory of the dataflow machine) to the network simulator. The simulator moves them through the multi-stage network to their destinations ("memory" to write data or to request reading of data; the "originating PE" for data that has been read), accounting for delays that are due to collisions. GITA is notified when the message has been delivered. The injection rates and destination distribution of messages from GITA to the network are obtained from real algorithms executed by GITA, rather than simple statistical samples. In turn, GITA receives a distribution of message delivery times that is more realistic than the constant latency it otherwise assumes. Preliminary results on


```

FUNCTION AINTEG(F,A,B,EPSLON)
EXTERNAL F
YSUM = (F(A)+F(B))/2.0
H = B - A
N = 1
APPROX = YSUM * H
100 CONTINUE
    HNEW = H/2.0
    X = A + HNEW
    DO 200 I = 1,N
        YSUM = YSUM + F(X)
        X = X + H
200 CONTINUE
    APXNEW = YSUM * HNEW
    IF(ABS(APXNEW-APPROX).LE.EPSLON) THEN
        AINTEG = APXNEW
        RETURN
    ENDIF
    H = HNEW
    N = N*2
    APPROX = APXNEW
GO TO 100
END

```

```

Integrate f a b epsilon =
  within epsilon (approxs ((f a + f b) / 2) (b - a) 1)
  where approxs ysum h n =
    (ysum * h) : approxs (fold1 (+) ysum ys) h' (n*2)
    where ys = take n (map f (iterate ((+) h) (a+h')))
    h' = h / 2

```

Two computer codes for the mathematical operation of finding the area under a curve (integration). The Haskell code (right) is shorter, takes less time to write, and is much closer to mathematical notation than the Fortran code (left).

the impact of network delays on available parallelism have been obtained for a Monte Carlo transport code written by Olaf Lubeck (C-3) in the dataflow language Id.

Programming Languages Research

Research continues on functional programming languages, with emphasis on methods for efficient parallel execution of functional language programs. By contrast with imperative languages such as Fortran, functional languages are more concise, expressive, and elegant. Moreover, methods for parallel execution are implicit in the language itself.

For the past 3 years we have participated on an international committee designing the functional programming language Haskell (named for Haskell B. Curry, a logician who contributed much to the study of the logical

theories underlying these languages). Concurrent with this design effort, we have been working with Yale University on an implementation of Haskell targeted for our experimental graph reduction engines. Work has also continued on the reduction engines themselves, including extensive improvements in Cray code generation.

Now that the Haskell definition is complete, we are turning more attention to the application of functional programming to scientific problems. In connection with this aspect of our research, a book is being written on functional programming in Haskell. Intended for a reader with experience in conventional programming and some mathematical maturity, the book will include realistic application programs in several areas. One area of particular emphasis is numerical and scientific programming.

Future areas of research include alternative approaches to compila-

tion of functional languages, expression and implementation of functional I/O, and interfaces to imperative languages.

Natural Language Research

Speech and speaker recognition with neural networks: The speech recognition project pursues the goal of recognizing natural, fluent speech by learning the gestural motions of the speech articulators that produced the speech. We have used this method to recognize digits, and a patent application has been filed on the process. Data have been collected from x-ray records of fluent speech, and used as input to a neural network that learns the relationship between the speech sounds and the articulator motions that produced them. Work is also under way to use this process as the basis for speaker identification and speech data compression.

The neural network technology developed for this project has

proved to be generally useful in the analysis of vector-valued time series, both for classification and signal processing applications.

Speech synthesis with default hierarchy learning paradigm: The default hierarchy learning paradigm is a powerful learning procedure that learns general rules together with sequences of increasingly exceptional cases. In application, it first looks for the exceptional cases, and then defaults to increasingly general cases. We are applying this paradigm to the notoriously difficult problem of learning to pronounce English aloud from written text. Currently, our results are superior to those of other automatic learning procedures (including neural networks and related methods) and nearly as accurate (90.4% correct phonemes) as human-crafted algorithms (93.6% correct phonemes). We hope that with additional tuning and training we will actually surpass human-crafted algorithms in accuracy.

Our recent results have highlighted two characteristics of the default hierarchy learning paradigm, namely, its efficiency (a four- to five-fold reduction in the number of rules required to achieve the same accuracy as a flat rule set), and the need to access general rules in applying the hierarchy.

We are currently applying the algorithms we have developed to the automated pronunciation of names—a practical application with the potential for technology transfer.

Learning Systems Research

During the past year we began work on a joint project with EES-5 on the analysis of ground motion data collected at the Nevada Test Site during underground nuclear tests. With suitable analysis, this non-intrusive technique can be used to estimate shot yields and thus could be an important capability for nuclear weapons test treaty verification. We are employing a variety of linear and nonlinear models and comparing the accuracy of the results as well as their stability under cross-validation.

Five parameters have been derived from the ground motion acceleration data taken in the spall region. These parameters were used as inputs to multilinear regression studies, neural networks, and the recently developed Linear Nonlinear Learning Method of Mycielski and Faber. The nonlinear models examined thus far have not produced results that are superior to linear models.

Previous work has concentrated on higher yield events. During the coming year we will begin to focus our analysis on lower yield events.

Expert Systems Research

A portable chemical identification and safety system combining the use of artificial intelligence, data acquisition, and database techniques is being developed in conjunction with MEE-3. The system consists of a Compaq 386 interfaced to a Neotronics gas detector and an airborne vapor

monitor (AVM) based on ion mobility spectroscopy. Oxygen level, explosive limits, and the presence of an array of chemicals are monitored continuously.

The software system used DESQVIEW to allow multitasking in DOS. There are five main parts to this system: the user interface, data acquisition, alarming, reporting, and an expert system. The user interface allows access to all interactive tasks, such as reports and diagnostics. There are two data acquisition tasks, one associated with the AVM and the other with the Neotronics. These tasks keep track of the gases and chemicals encountered and whether they are in or out of alarm. The alarming task manages the presentation of the alarms to the user. The reports are entered and stored using INGRES. The first four tasks have been developed but remain to be integrated. The expert system will help a user in determining what chemical process is being performed based on the presence of chemicals detected.

As a safety system, short-term exposure limits, time-weighted averages, and ceiling levels are tracked for a set of chemicals. These limits correspond to chemical concentrations over time. Two types of alarms are issued based on these limits. A YELLOW alert indicates that some hazard exists but is not life threatening. A RED alert indicates that a hazard is life threatening.

This system will be completed and delivered to the sponsor later this fiscal year.

Scientific Visualization Projects

In the past C Division often worked mostly on its own or with vendors to develop tools that were used by ICN customers. Recently, C Division has become more involved in solving user problems by collaborating with users and helping them solve their problems.

A couple of examples of this collaborative approach to problem solving are the work in scientific visualization we've been doing with one of our external clients, the Defense Nuclear Agency (DNA), and with an internal Laboratory organization, the Institute of Geophysics and Planetary Physics (IGPP).

The overall goal of these collaborative efforts is to work with the users to identify and develop the best possible computing tools for their applications, and, where appropriate, incorporate those applications in the ICN.

The DNA/LANL Scientific Visualization Project: Providing Visualization Tools for Remote Users

Scientific visualization is a relatively new term used to

describe the inter-disciplinary use of computer graphics to understand scientific data. Scientists generate massive amounts of data and require effective methods to visualize the data's significance. Advances in computer graphics in recent years provide such methods, although in many cases they are not readily available to the scientific community.

Recognizing this, we have begun a collaborative Scientific Visualization Project with DNA to identify, develop, and integrate tools for the creation, transmission, and display of graphical data. Project goals are to identify, evaluate, and develop visualization tools on supercomputers; to evaluate a range of workstations for visualization; to identify several visualization system configurations; and to prototype a remote distributed computing environment. This will be accomplished by working with application programs to gain a better understanding of the scientist's problems and environment. Capabilities developed will benefit local users as well as remote DNA users, who depend heavily on the supercomputers at Los Alamos, yet require low-cost

ways of bringing the necessary visual aspects of graphics work within closer reach.

Early efforts in the project have provided a way to view and animate graphics from Cray supercomputers on relatively low-cost Macintosh and Sun workstations. This capability, extremely useful to remote users, is a foundation that will support all further project developments. The project's use of standard file formats for graphics allows this.

The current focus of the project involves working with scientific data sets, displaying them in many different ways. Volume rendering, raster graphics, innovative uses of color, and other techniques are being examined. Software packages include PV-Wave, TAAC, Vision, Dore, and AVS. Graphics hardware includes a Sun workstation with a TAAC graphics accelerator, an Iris personal workstation, and an AVS workstation. In collaboration with scientists, various configurations of this software and hardware are used to determine the best methods for displaying research results.

Before technology provided scientists with huge supercomputers and monitors for scientific visualization, a common method of visualizing data was the contour plot (image at left). Both of the above images were generated from the same data. Although the one shown in color provides considerably more detail than the contour plot, it doesn't show what happens over time. On some of today's monitors, a scientist can view the sequence as a movie, so color and time provide two additional tools for interpreting changes in the data.

The IGPP Scientific Visualization Project: Providing Visualization Tools for Laboratory Users

In collaboration with the IGPP, we have worked to develop new methods for visualizing global climate change data. IGPP has massive amounts of data gathered over a period of many years. This data needs to be visualized to get a better understanding of global climate change. The purpose of C Division's involvement was to demonstrate scientific visualization techniques in a prototype fashion using IGPP data. Scientists are interested in many aspects of the recorded data, such as surface temperature, pressure data at several geopotential heights, and cloud cover development.

C Division took one year's worth of data from the climate simulation model and developed prototype visualizations using capabilities on several computing and graphics platforms. One full year's worth of data was animated, showing scientists the trends over that period in global surface temperatures. This was done using software on the Cray and animating raster images on the Macintosh and Sun workstations. The graphics generated on the Stardent high-end graphical workstation were particularly exciting for the scientists because they could interact with the image. The data was mapped onto a globe where the scientists could then manipulate the globe to look at various geographic areas. They were also able to view an animation sequence showing several months worth of data. Video recordings were made for later IGPP presentations.

Global climate model data are especially well suited to visualization on a two-dimensional color continuum. This image, from work done on the IGPP project, shows surface temperatures of the globe in July. Animation of this allowed the scientist to view an entire year's worth of data. By watching the changing colors, where blue indicates the coldest region and red indicates the warmest regions, unusual or unique temperature regions are readily recognizable. Images from different time periods can be easily extracted and compared. The original data were generated on a Cray computer and can be viewed on a variety of workstations, in this case, a color Macintosh. Scientists can also view this data in three dimensions in the ACL or on the Crays (see the separate article on the following page about the ACL).

Three-Dimensional Sound

Although a lot of attention is focused on visualization of scientific data as a method of interpretation, a new and exciting area of research explores the possibilities of conveying information by computer-generated sound. This expression includes a new approach to sound generation in that sound is to be given the degrees of freedom of three dimensions and contain the essence of the physical function that it represents. The project, which is a collaborative effort

between staff in C Division and X Division, requires that we learn how to assign sounds to mathematical functions that convey the essence of the details that we intend to convey, and then manipulate such voices in space so that the spatial nature of the function is equally comprehended.

Important components of the project include the UNIX operating system, the C programming language, the NeXT computer, and a Yamaha S77 synthesizer.

Advanced Computing Laboratory

The ACL was instituted to facilitate Laboratory research in advanced computing in an open environment. Established in 1989, the ACL has moved into temporary quarters on the first floor of Building SM-123 while awaiting the completion of a permanent facility, a two-story building now under construction. The goals of the ACL include serving as a focal point for the solution of large, interdisciplinary problems. These typically require vast amounts of computing power, as well as state-of-the-art visualization capabilities and rapid, efficient movement of large amounts of data. Areas of current interest include applied mathematics, chemistry, computer science and engineering, networking, and experimental technology. An extensive, active visitor and student program has been initiated to advance these interests.

An important aspect of the ACL is working with vendors of computing equipment to provide input into future offerings while having access to the latest technology for solution of problems. Collaborations with several companies are ongoing; others are under negotiation. Computers in the ACL are often "here today, gone tomorrow," as we cycle models through on beta test agreements for evaluation. Current computers housed in the ACL are the Thinking Machines Corporation 64K-processor CM-2, a 4-processor Ardent Titan, a dual-headed FPS 350X (Stellar GS2000), an FPS model 500 computer, two IBM RISC/6000 workstations, a Silicon Graphics 4D/380 VGXB, and a Silicon Graphics 4D/25TC, a personal IRIS workstation. An additional

ACL computer, the Convex 220, is located as an open computer in the LDCC.

The advent of the color workstation with accelerated graphics has opened the door to the investigation of new ways for scientists to rapidly view their data. However, a lack of appropriate tools has made this process both time consuming and difficult. A different approach is to use an intuitive and easy-to-use system with reusable subprograms or modules. Using the Application Visualization System (AVS) developed by Stellar (now Stardent), the ACL has begun research using this modular approach. In one application, we have worked with IGPP and EES Division to produce a visualization of their climate model. The data were originally produced during an extensive simulation run on a Cray. We provided multiple methods for viewing this data: a surface temperature color map on a globe (see image below) and a lofted pressure surface (geopotential height) painted using the surface temperature color map. We are involved in providing similar capability on other graphics workstations in the ACL.

However, we are moving rapidly toward a distributed computational environment in which visualization workstations are not standalone resources. We are proposing, in collaboration with C-6, to investigate effective distribution of graphical computation and rendering for scientific visualization throughout a highly networked environment and the use of intelligent agents to assist the study of large data sets. We anticipate that this project will interact strongly with several CM-2 applications and software projects and the HIPPI/CP* testbed. This work will require high-speed VME interface with the CM-2 and a suitable graphics supercomputer (for example,

The image shown here was generated on an FPS 350 (aka, Stellar GS1000) graphics workstation in the users area of the ACL. The images of the globe show surface temperature variations from a climate model; the blue indicates the coolest region and the red indicates warm regions. The scientist views the images on the workstation's monitor and can visualize trends in temperature over a period of time by watching the colors change as a result of the dynamics of weather. Using this 3-D workstation, the scientist can rotate the globe and view it from different perspectives. A similar image from the same work is shown in a separate article in this section (see page 69).

Stardent or Silicon Graphics, Inc.) and a HIPPI application interface allowing transportation of both graphical images and large application data sets.

Staff of the ACL and CGL are developing video capabilities for the ACL and the rest of the Laboratory. The capabilities of the CGL are described in Section 4, *Facilities and Capabilities*. In addition, the CGL will provide a mobile recording station that will tap into the RGB lines on a color terminal to provide off-the-screen taping. The ACL, without duplicating the goals of or equipment in the CGL, will provide its users with video capability. The ACL will obtain a worm video-disk, a scan converter, and a video camera. Through equipment loan, the ACL will also have several scan converters, beta-cam tape decks, and controllers. The video capability for the ACL will allow the CM-2 frame-buffer output to be broadcast over LABNET to users' offices.

Users of the ACL will be able to make 'quick-and-dirty' video recordings of their work. These recordings are not intended for presentation but rather for analysis of data. Should users require presentation-quality animations, the worm video-disk can be used in conjunction with the CGL to produce high-quality video recordings. These two video capabilities complement each other because users will be able to obtain

The computer room of the ACL temporary facility houses a variety of experimental, advanced computers. Shown in the foreground are two graphics supercomputers, the FPS 350 and the Ardent Titan; in the background is the Thinking Machines CM-2, a massively parallel computer used for a variety of large applications. An ACL staff member monitors activity on the system console of the CM-2. With the help of ACL staff, researchers are encouraged to try their applications on any of the ACL computers.

animations of their graphics for both analysis and presentation.

We have several efforts under way in the area of networking. The ACL will serve as the focal point for HIPPI testing at the Laboratory, and many of the projects will center on the availability of a variety of computers in the ACL. The ACL will use the LANL HIPPI tester initially to evaluate the HIPPI channels on each of the hosts that are brought into the Laboratory. We will also be an active participant in the CASA Gigabit Testbed, described elsewhere in this report.

The ACL fosters participation in the development of computational science on a national level, specifically through the DOE Applied Mathematical Sciences Program and the NSF Science and Technology Center for Research on Parallel Computation.

People interested in using the resources of the ACL are encouraged to do so; call the ACL office at 665-4705 for information. Those requesting large amounts of computing time are required to send in a proposal stating their goals and desired computing resources.

**Andy White, Director
Advanced Computing
Laboratory**

The ACL is the culmination of the work of many, many people over a period of years. However, it is really the brain child of its director, Andy White, whose hard work and dedicated effort made it become a reality. Andy, who is also C-Division's Deputy Division Leader for Research, has envisioned a facility such as the ACL since he was a mathematician and Deputy Group Leader of Group C-3. At that time, much of the research into parallel processing

and advanced architectures was being conducted in C-3, and Andy felt that a more visible, Laboratory-wide service was needed. As computational science has emerged as a national issue, it became more apparent that the Laboratory needed and would be supportive of such a capability. Andy's vision and hard work has paid off; he is a dynamic leader, and the ACL is an exciting place to work!

Benchmarking

The Laboratory's leadership in the application of supercomputers to scientific problems depends heavily on its ability to acquire the most powerful computing machinery available. Benchmarking is the process by which the performance of a given machine on computational problems of interest to the Laboratory is determined. The diversity of current and future supercomputer architectures is such that benchmarking can no longer be done simply by looking at the speed of the device technology used in a computer. The overall performance of a machine is based on a balance between elements that often run at widely differing individual speeds; a good example is vector vs. non-vector performance. It is important to measure the contribution of these components separately because if a code cannot make complete use of the faster component, the performance of the code will be determined by the slower.

Because most of the Laboratory's production codes are too big to be ported easily, are classified, or are not written in portable Fortran, we maintain our own set of benchmark codes that represent algorithms of current importance to the Laboratory. We continually update this benchmark set as computing needs change. As part of this update procedure, we have begun to evaluate larger codes for addition to the set. These larger codes are needed in order for us to properly assess newer machines with multiple processors and greatly increased memory sizes. We currently have two unclassified production codes, plus sections of

newly developed codes under evaluation.

Recent Results

During the past year we continued our assessment of Cray Research supercomputers by carefully monitoring the development of the CRAY Y-MP and CRAY-2 lines. We have evaluated several of the larger memory CRAY Y-MPs as well as a faster memory CRAY-2. The Y-MPs show improved memory performance over the X-MPs due to better design in both the network and the memory. Table 1 shows results from some of these tests.

We also continued our efforts at benchmarking the class of machines known as minisupercomputers by testing the newest release from Convex Computer Company, the C240; a Floating Point Systems FPS500; and an early version of a new VAX from Digital Equipment

Corporation, the VAX9000. Two of these machines are multiple-processor vector machines with large amounts of memory, and the other (FPS) is a single-processor vector machine. Table 2 shows the results of the benchmark tests on these machines plus results from an interesting machine no longer being made, the Multiflow TRACE14.

Because scientific workstations have become an important part of the Laboratory's computing program, we began an effort to evaluate some of these workstations. Our recent tests have included workstations from IBM (RISC SYSTEM/6000), Sun Microsystems (Sun 3/60, 3/470, 4/110, 4/260, and three SPARCstations), and Digital Equipment (DECstation 3100). The performance of some of the machines is close to that of the minis mentioned above. The results from benchmarks on three of these machines are shown in Table 3.

We have monitored closely the performance improvements of the ACL's Connection Machine model CM-2 as a result of software evolution. During the 1989-1990

Table 1. Comparison of Single-Processor Benchmark Execution Times (in Seconds) CRAY Y-MP8 and CRAY-2

Code	CRAY Y-MP8/832	CRAY Y-MP8/864	CRAY Y-MP8/8128	CRAY-2 (Static)
FFT	2.8	2.6	3.0	3.7
GAMTEB	3.1	3.0	5.3	3.8
SCALGAM	59.0	-	-	68.6
PHCTON	91.1	87.4	100.7	99.4
LSS	4.3	4.3	6.0	7.4
MATRIX	23.0	22.3	29.6	38.9 ^a
INTMC	9.1	8.7	34.6	17.4
HYDRO	11.3	10.6	34.4	14.8
WAVE	70.3	61.6	81.1	78.4
ESN	14.6	12.1	13.5	16.7
VGAM	-	2.8	-	4.0
MCNP ^b	6.6	6.4	-	-

^aRequires odd dimension for arrays

^b1000 source particles

time frame, the CM Fortran compiler became stable enough for production runs, the *Lisp compiler reached maturity, and the C* compiler was scrapped because of poor performance. At the same time, a new programming system (CMIS) was delivered to allow very close control of the hardware at a very low level.

In support of a collaboration doing Quantum Chromodynamics (QCD) on the Connection Machine, we have developed highly optimized CMIS versions of the heart of all QCD calculations, an on-processor complex matrix multiply. This has been accompanied by development of faster between-processor communications routines, which currently deliver four times the performance of the standard communications software for the QCD application.

In Table 4 we compare performance for the matrix multiply on both Cray and CM systems. Note the Cray figures are for a single processor doing 64-bit arithmetic, while the Connection Machine figures are for a 64K CM-2 doing 32-bit arithmetic versions of the floating-point unit available on the machine.

Table 2. Comparison of Single-Processor Benchmark Execution Times (in Seconds) on Minisupercomputers

Code	Convex C-240	VAX 9000	Multiflow Trace 14/300	FPS 500
FFT	14.7	15.7	17.5	30.1
GAMTEB	10.7	6.5	12.1	14.3
SCALGAM	242.1	121.7	403.5	224.2
PHOTON	338.2	174.1	553.5	310.3
LSS	23.5	22.9	20.7	29.7
MATRIX	140.0	-	259.0	170.0
INTMC	31.1	32.1	26.8	38.8
HYDRO	74.6	145.9	206.7	79.2
WAVE	410.2	272.5	816.0	388.3
ESN	41.9	22.9	56.7	66.8
VGAM	4.4	-	16.4	4.1
MCNP ^a	2525.9	-	-	-

^a60000 source particles

Table 3. Comparison of Benchmark Execution Times (in Seconds) on Workstations

Code	IBM RS/6000 40 ns	VAX DECstation 3100	Sun SPARCstation 4/330
FFT	11.5	68.2	86.3
GAMTEB	8.1	24.2	35.0
SCALGAM	163.4	366.4	674.1
PHOTON	221.3	523.9	945.9
LSS	20.1	119.7	162.0
MATRIX	147.0	1160.0	1140.0
INTMC	18.6	27.0	34.1
HYDRO	200.8	523.3	664.0
ESN	18.0	51.2	43.9
VGAM	7.0	21.0	27.5

^a60000 source particles

Table 4. Matrix Multiply Performance of the Cray and CM Systems

Machine/Compiler	GFLOPS
CRAY X-MP/CFT77 (1 CPU)	0.13
CRAY Y-MP/CFT77 (1 CPU)	0.21
CM-2/CM Fortran	1.87
CM-2/CM *Lisp	1.95
CM-2/CMIS (WTL3132)	8.50
CM-2/CMIS (WTL3164)	10.20

Collaborations

Another part of developing and maintaining our role as a world leader in computing is our collaborations with vendors. It is through these collaborations that we are able to research and develop the new technologies described throughout this publication. Our interactions and collaborations with external organizations tightly couple the Laboratory's work in advanced computing with U.S. competitiveness and with DOE technology transfer initiatives.

IBM

One aspect of the joint study effort with IBM is the implementation of the HIPPI ANSI-standard channel specification onto an IBM 3090 computer. IBM was the first computer vendor to implement this channel, and Laboratory and IBM personnel have been working together to develop user applications and networking techniques that will utilize the 800-Mbit/s data rate capability this channel possesses. Initial plans call for this machine to be connected into a CP* testbed, where the HIPPI link will be used to provide CFS type of access to other systems in the testbed. HIPPI appears destined to be an industry standard, with more than 90 computer vendors showing interest.

Cray Research

The Laboratory has been successfully collaborating with Cray Research since before Cray's first computer was installed in the ICN in 1976. Recent collaboration includes participation on an advisory panel sponsored by Cray Research to review Cray's enhancements and developments of

the UNICOS operating system to ensure that they are consistent with requirements of the various participating installations. The panel is also responsible for suggesting new directions for UNICOS.

Cray Research has also implemented the HIPPI channel onto its computers. Two Cray machines in the CCF currently have this channel. These channels will provide 800-Mbit/s data rates for networking and user applications. Initially they will be connected into a CP* testbed.

Thinking Machines Corporation

As part of a joint study with Thinking Machines, we are helping them understand the behavior of their innovative new architecture in an environment where real supercomputer problems are solved.

A HIPPI interface for the CM-2, which will provide 800 Mbit/s data rates for networking and user applications, will be delivered this summer (1990). The CM-2 in the ACL is also planned to be connected into the CP* testbed (see below).

Sun Microsystems

Workstations from Sun Microsystems are one of the Laboratory-standard workstations. This has resulted in Sun workstations being the predominant scientific workstation at the Laboratory, and, as a result, we have several collaborations with Sun. Several groups in C Division have been "beta-test" sites for new software and hardware from Sun Microsystems. This beta testing has a direct benefit to the Laboratory in that

we get a chance to identify and fix problems with these systems in our environment.

Two collaborations with Sun are in the graphics area. The first is our participation in the beta testing of the window system software called NeWS. These window system and user interface software systems will play a critical part in making the interaction with our computer systems easier for the user. Our participation in this testing program provides us an excellent opportunity to influence and remain at the forefront of this rapidly evolving area of distributed computing.

Sun Microsystems is also providing two prototype HIPPI interfaces to their TAAC-1 graphics accelerator. This will be the first workstation in an office that will allow us to interface to the supercomputer at extremely high data rates (100 Mbyte/s) provided by HIPPI. The workstation will be connected to the CP* testbed via the HIPPI channel and will be used to explore several issues in distributed graphics and high-speed networking from a supercomputer. X Division is a key participant in this project.

NeXT Incorporated

NeXT Incorporated was founded by Steve Jobs, one of the inventors of the Apple computer. In the fall of 1988 NeXT announced their first computer system, a desktop workstation with an innovative user interface based upon the use of the Display PostScript page description language. We have had discussions with NeXT since the summer of 1987, and these discussions have resulted in a joint development effort with NeXT to extend this user interface to interoperate with remote computer systems such as our Cray supercomputers. Discussion continues with NeXT regarding our needs for personal desktop workstations for our

users, especially in the area of graphics.

A club exists at the Laboratory for NeXT computer enthusiasts. Information about the NeXT club is available from Dwight Barrus of Group C-6, (505) 667-8870.

Gigabit Project

Private businesses, universities, and government agencies are teaming up to develop a high-speed computer network that would enable the transmission of 1 billion bits of computer data per second via fiber-optic cables.

This project, funded by a \$15 million grant from the National Science Foundation and the Pentagon and more than \$100 million from the computer and communications industries, will demonstrate supercomputing using wide-area, high-speed (800-Mbit/s) networks. Several technological challenges confront the project, such as the need to develop a new class of computers to keep up with the new "data highways."

Other recent cooperative technological projects have failed because of a lack of money or enthusiasm among government and industry, yet in Japan, government and large corporations

have already begun work on developing data highways.

This may be the first comprehensive effort in the United States to advance the technologically critical area of computer networks. Dozens of corporations will join with several universities and government laboratories in this effort.

This 3-year effort will establish five testbed networks. Los Alamos, Caltech, the Jet Propulsion Laboratory, and the San Diego Supercomputer Center will be the major participants in one of five testbeds. The testbed called CASA will focus on distributed large projects across a high-speed, wide-area network.

Three important scientific problems will be studied by harnessing an 800-Mbit/s network computer and data resources at the four institutions. Los Alamos will contribute the computational resources of the ACL's Connection Machine, provide expertise in network integration, research high-speed interfaces, and collaborate on applications development.

NSF Science and Technology Center for Research on Parallel Computation

Los Alamos is collaborating with Rice University (lead university), Caltech, and Argonne to develop the science and technology base to make parallel computation truly usable.

Other Collaborations

We have other collaborations with major computer vendors. These include Apple Computer, XEROX, QMS, and PARALLAX. University collaborations span some 30 universities covering a broad range of computational science topics. In addition, members of C-Division staff are associated with a variety of committees and advisory boards; these are described in another section.

Distributed Computing

Over the past decade, there has been extensive growth in the populations and capabilities of small computers, primarily personal computers and scientific workstations, but also including VAX-class computer systems. The extraordinary growth in the numbers of these machines is reflected in the graph below. While this graph shows the cumulative numbers of all computer systems at the Laboratory, the principal growth is in small computer systems.

Technology developments have driven the extensive growth in computer systems at the Laboratory. As mini- and microcomputers became more available, Laboratory staff found them convenient for many applications. Small computer systems, from workstations to minisupercomputers, provide many functions that are not available, not appropriate,

and not easily provided by central supercomputer systems. They have proven to increase users' productivity because of the combination of a rich software environment with a good user interface. They are used for a wide variety of tasks including program development, small- to medium-sized calculations, data collection and analysis, and document and presentation preparation.

The utility of many of these systems depends on the availability of networks for communicating with disk and print servers as well as communication among computers across the Laboratory and throughout the world. With the growth in the population of computer systems at the Laboratory, there has been a corresponding growth in the number of local area networks.

The number of computer systems at the Laboratory, including personal computers, scientific workstations, VAX and VAX-class computers, has grown extensively in the past decade. In recognition of this growth, C Division has a Distributed Computing Project whose goals are to support distributed computing and the ICN.

Distributed Processors

Although there were already many standalone minicomputers doing various data management and data acquisition tasks, in 1976 Los Alamos entered the era of geographically distributed, general-purpose computing with the procurement of the VAX computers. These were acquired by user organizations and communicated with the central facility using point-to-point communications and the DECnet suite of data communications protocols. An illustration on the next page shows when the minicomputers arrived at the Laboratory.

There are a total of 361 minicomputers. It is interesting to note that only 6 minicomputers were purchased in FY 1989. Very possibly this is the result of the significant increases in the capabilities of the scientific workstations.

There are currently 269 DECnet hosts distributed among the partitions as follows: Open, 172; Administrative, 26; Secure Unclassified, 34; and Secure Classified, 27.

Workstations

In 1982 and 1984 C Division conducted procurements for desktop workstations. The first of these procurements was for the least expensive machines called "personal workstations," and the next was for more expensive and more capable machines called "scientific workstations." These machines became very popular at the Laboratory because of their user-friendly software (spreadsheets, word processors, and drawing tools), their responsive computing capabilities, and their low cost. They have proved useful to a broad spectrum of the Laboratory population, and, as a consequence, their populations have exploded (see illustrations).

Advances in semiconductor technology have enabled divisions and groups to acquire greater

Minicomputer acquisitions at the Laboratory. It is interesting to note that only six minicomputers were purchased in FY 1989; this is possibly the result of significant increases in the capabilities of scientific workstations.

computing power today than was typically available to institutions only a few years ago. The table shows the decrease in cost for vaguely similar computing capabilities.

The computing power that workstations bring to scientists' desks has changed the way many people do their computing and has increased their productivity. However, there are costs associated with these capabilities. These include software maintenance and services, hardware maintenance and upgrades, and manpower support costs.

Networks

Networking is an essential part of high-performance computing. The concept of network services (such as file servers and print servers) for collections of supercomputers has been a reality at Los Alamos since the early 1970s. With the advent of workstations and distributed processors, many applications are now network intensive; that is, they require robust, reliable, and fast networking just to work correctly. Examples of network-intensive computing applications are mail, network windows, network file systems, and most graphics applications that require a high-speed computational aspect beyond the workstation.

ICN Core Network

Supercomputers, large mainframes, the CFS, PAGES, FOCUS, various access points, and numerous gateways are interconnected using an internally developed 50-Mbit/s High-Speed Parallel Interface. This core of systems communicates using locally developed protocols and a security model unique to Los Alamos that are closely related and allow the ICN to operate as the only approved multiple security level computer network in the nation.

While providing various advantages, these locally developed solutions are difficult to change in a timely fashion, and widely available commercial solutions to computer network problems fit into this framework only after large expenditures of effort on the part of C Division, if at all.

ICN Backbone Network

Outside the core of the ICN, gateways connect not only 60 local area networks based on the Internet protocols and the DECnet protocols, but also numerous nationwide and international networks of various capability and using several different protocols, to the centrally provided services. These networks and their connections to the gateways are based on commercially available communications protocols and security models built into the operating systems.

While the core security model distrusts the hosts and places almost all security assurance in the network, these commercial systems distrust the network and place almost all security functions

Growth of personal workstations at the Laboratory.

Growth of scientific workstations at the Laboratory.

in the hosts. The speed of these connections range from Ethernet, with a signaling rate of 10 Mbit/s (and a usable throughput of 1 to 5 Mbit/s depending strongly on the capability of the hosts involved and their interfaces), down to 9.6-kbit/s asynchronous links.

Also connected to the backbone are a few hosts requiring very high performance and both DECnet and IP routers. An IP router is connected from the backbone to a shared 5-Mbit/s channel on the Laboratory's broadband cable system (known as LABNET) and about 20 other routers connect this shared LABNET channel to local area networks at sites throughout the Laboratory.

Other hosts are connected to backbone routers through point-to-point links (including some dialup links) ranging in speed from T1 (1.544 Mbit/s) to 9.6 kbit/s. A fraction of these connections use bandwidth from LABNET through multiplexors that provide up to T1 speed serial links.

Within the next year the backbones will become 100-Mbit/s Fiber Distributed Data Interface (FDDI) token rings. Each gateway is connected to a backbone Ethernet.

One important example of computers as communications tools is electronic mail. Mail applications are important for collaborations within a nationally distributed scientific community. The Laboratory uses national and international networks for mail exchange. The Bitnet network, for example, is used by LAMPF and others to communicate with colleagues throughout the United States and in Europe. The ARPAnet/Milnet network has been in active use by the Laboratory for several years as a communications medium between Laboratory scientists and scientists across the country. Today, we send and receive about 6000 messages/day from the various national and international networks to which we are connected.

Distributed Computing Project

Because of the explosive growth of distributed computing at the Laboratory, the Distributed Computing Project was formed to help focus the appropriate C-Division resources on the development and support of a distributed supercomputing environment at Los Alamos. The goals of this effort are:

- To provide an open system computing structure within the Laboratory that is consistent; that allows ICN users to use the latest available computing technology that best meets their needs; that enables application programs to be readily shared; and that allows widespread use of distributed applications.
- To provide a security-partitioned, fast, packet-switched network across the Laboratory that is based on the TCP/IP suite of protocols and enables LANs of workstations to communicate with one another and with the CCF. Open partition LANs should be able to communicate with networks external to the Laboratory.
- To provide a network file system that includes the CFS.
- To provide users with a responsive graphical interface to supercomputers.
- To provide users with tools that will make them more productive in their software development and production computing activities,

Costs Declines for Comparable Computing Speeds Over 15 Years

Computer	\$M	Year
CRAY-1	10	1976
Convex	1	1988
IBM RS6000	0.15	1990

e.g., distributed editors and debuggers.

- To provide a centralized support environment so that the expertise required to set up and efficiently manage these complicated environments need not be replicated throughout the Laboratory.

Achieving these goals requires the development of a significant communications infrastructure that provides a reliable and secure network; a body of software that will transparently support distributed applications; a support structure to install and maintain the extensive hardware and software required to implement this environment; and, finally, an education and consulting activity to help new users become acquainted and work easily with this new approach to computing.

We have made much progress toward our goals in the past several years. We have built open and secure capabilities, based on the TCP/IP suite of protocols, to communicate with 40 Open networks, three Secure Unclassified networks, and one Secure Classified network with a total of 1600 workstations. These workstations currently use TCP/IP to deliver messages to gateways and use local protocols and software to access target CCF services. This locally developed software includes a CCF services authentication, a virtual terminal capability, a file movement capability, and access to CFS and PAGES. The software to support this capability runs in virtually the same form on numerous workstation platforms including Sun and Apollo workstations and IBM PCs.

The ICN environment software has been standardized to support both UNIX BSD and UNIX System V workstations. Non-standard systems are not recommended and will require user modifications.

We also support routing between networks, a domain

name server, and electronic mail, and we will be working on network management tools in the near future. In addition, the Open partition workstations are connected through the network to the MILnet, which allows users to communicate using electronic mail with colleagues across the country. We also have links in the Open partition to NSFnet, WESTnet, and New Mexico TECHnet, all of which provide IP access to research centers around the country.

This progress is significant; however, to achieve our goals we need to position ourselves to use commercially available software that will implement the basics of the distributed environment described above. Since all of the major players in this field use the TCP/IP suite of communications protocols as a basis for their distributed computing software, we have decided to do the same. Implementation of these protocols in our environment has required that we carefully consider the implications on security and performance. Several current projects, described in some detail below, are addressing these issues.

Secure IP Router

Our current secure network functions are based on a set of locally developed protocols called SIMP. However, SIMP has one major drawback: it is a proprietary protocol that is unique to Los Alamos, so no commercial applications (such as remote procedure calls) use it. The commercial software needs an end-to-end transport capability like TCP. TCP uses a packet-routing protocol called IP to move packets from source to destination, but to use IP at Los Alamos we need to implement access control for security. Therefore, we are developing an IP router that will allow us to support an end-to-end IP connection from a user's workstation to

certain CCF services such as supercomputers.

An IP router provides both a physical connection between two distinct networks and an IP datagram packet-switching service between the networks. In addition to this basic functionality, our environment requires that this router provide secure access to a target network, audit logging of incoming or outgoing packets from this target network, and tools for correcting and monitoring potential communication problems.

The Secure IP Router (SIPR) is available for testing and debugging in the Open partition. Discussions have been started with various workstation vendors to obtain the means for modifying their respective kernels.

IP Security

The Laboratory would like to use as much vendor-supplied and vendor-maintained software as possible, and in the area of IP routers we would like to influence vendors to provide us the software needed for our secure environment. We are participating in a cooperative effort with Sandia and LLNL to define IP security requirements for commercially available IP routers. Currently, the IP security fields are not adequate to meet the needs of the DOE laboratories. Even the defined security fields are not implemented by all commercial IP routers. This effort will encourage vendors to implement our stated requirements in their future products.

Other Major Projects

Other major infrastructure projects are planned for the next year. Details are given below.

Communications Backbone IP Routers

As the ICN has grown to support multiple heterogeneous

LANs, we have been faced with the problems of conflicting networking software implementations running on the various workstations. This has caused disruptions to service because of broadcast storms. We have decided to protect the various networks from these disruptions by isolating them with IP routers. We are evaluating new IP routers based on their performance, IP and ICMP functionality, dynamic routing protocols support, software configuration, hardware/physical layer support and maintenance, link layer support, troubleshooting and debugging tools, network management tools, and vendor directions and support.

We have ordered IP routers from three vendors for evaluation. The evaluation will be completed by the end of September 1990, and we will take the appropriate procurement action then.

In addition, C Division has established a Central Networking Testbed (CNT) whose mission is to evaluate networking systems for conformance to minimum requirements for proper operation within the Los Alamos networking environment. Details about the CNT are included in the section of this report entitled *Facilities and Capabilities*.

Communications Backbone Upgrade

The existing LABNET communications backbone is a broadband cable television technology that supports multiple 6-MHz video channels. The advantage of this technology is that it can support users over a very wide area. The disadvantage for the long term is that the data rates are limited to approximately 3 Mbit/s per channel. The total capacity of the broadband if all channels were used for data is about 150 Mbit/s; however, only about half of the total bandwidth is used for data. As the data requirements continue

to expand, we will need to increase the capacity of the backbone to support data rates many times this level. In the next year we will be researching various ways to accommodate these dramatically increasing data rates.

User Support Activities

C Division will be involved in various aspects of workstation support for the user community. We are currently involved in the design and installation of LANs, and we have recently expanded that involvement to include help with overall system design, system management, and more explicit help to users in the connection of their proposed LANs to the Los Alamos internet. This effort will involve C-Division personnel and vendor personnel where appropriate.

Sun Microsystems plans to provide additional support to users on a contract basis for system/network management. This will be managed through C Division to ensure consistency across Sun networks. We are pursuing a similar arrangement to help with networked PCs. We are contracting with Dominguez Associates for that effort.

We are also beginning to investigate ways to educate users about this new computing environment. The C-Division classroom has been equipped with a LAN of workstations. This set-up provides live, hands-on courses related to distributed computing and the use of workstations. We will use on-site classes as one mechanism for educating users; other ways will be used as we discover them.

New Initiatives

During the next year, work will continue on new initiatives in the areas of user interfaces, CTS, and distributed applications.

User Interface Team

This new workstation-based environment will add a dimension to the interface that we will be able to present to users. Currently, windows are available on workstations, but when used with supercomputers or remote hosts, they simply implement the virtual terminal inside the window. We are interested in using the emerging standard window systems to present an application-level window on workstations.

The User Interface Team was formed with people from C-6, C-8, C-10, X-5, X-6, and X-7. The team investigated the distribution of one of X-Division's design codes between a workstation and a Cray. The team will use X11 as the windowing environment initially, but they also plan to study NeWS. A prototype system that showed functionally how the distribution will work was developed, and a report was published in the *Computing and Communications Division News*. Motif has been obtained, and we are investigating its use at Los Alamos.

CFS Plans

We are continuing our efforts to understand the role that CFS should play in the workstation environment. We will be re-researching how NFS might be made a part of the CFS environment and what the role of TCP/IP will be in the context of CFS.

Network Computing System

Network Computing System (NCS) was installed on Apollos, Suns, VAXs running UNIX BSD4.3, and Cray computers. Direct TCP/IP access to CTSS from workstations was made available through Machine P. C, X11, and NCS libraries were made available on an experimental basis under CTSS. An NCS demonstration program runs on a workstation and distributes different parts of a problem to a Cray running UNICOS, a Cray running CTSS, a VAX running UNIX BSD4.3, a

Sun, and an Apollo. Also, it is possible for the workstation to have multiple X11 windows to multiple machines. NCS is being installed to run under VMS.

Distributed Applications

The infrastructure will be firmly in place in the next year so that we will be able to begin serious work on distributing applications that require the use of supercomputers (for number crunching) and workstations (for responsive graphical interfaces). This development of a "distributed application" will define the tools that need development. Tools we have

begun work on include a graphical CFS interface for the Sun workstations and a workstation-based debugger based on the CTSS DDT debugger. Another area that needs to be considered is the current Tektronix graphical interface being used by a majority of the supercomputer users; this may not be the best solution in a distributed environment.

Conferences and Workshops

C-Division staff participation in conferences and workshops contributes significantly to the exchange and development of ideas. Members of the staff regularly present papers at international conferences; titles of many of those papers are provided at the end of this section. During the last year, C-Division participated in many conferences and workshops that are particularly noteworthy.

SIGGRAPH

The 1989 SIGGRAPH conference was held in Boston, Massachusetts, July 31–August 4, 1989. This annual ACM conference is the major national forum for the presentation of state-of-the-art technologies in computer graphics. SIGGRAPH has become the premier computer graphics conference in the world. C-Division staff are active in this conference and some served on the conference committee:

- Dick Phillips, C-6, published the first SIGGRAPH panel proceedings in November 1988. This proceedings was sent only to those who had registered for the conference. In addition, the following year, he served on the SIGGRAPH '89 Courses Committee, taught the Introduction to Computer Graphics seminar, and served as the chair of two panel sessions.
- Bob Judd, C-6 Group Leader, served as Panels Chair for SIGGRAPH '89. The SIGGRAPH '89 panel proceedings, published by Bob, was the first to be published as an issue of *Computer Graphics*, the SIGGRAPH quarterly. This proceedings greatly improved quality and was mailed to every member of SIGGRAPH as

well as to those who attended the technical program of the conference.

- Ted Reed, C-6, was on the Panels Committee; he chaired a session and helped present a one-day tutorial on computer graphics.
- Diana Tuggle, C-2, was the Speaker Slides Chair for SIGGRAPH '89. In this capacity, Diana was responsible for coordinating speaker slides and providing on-site graphics services.

The 1990 SIGGRAPH conference was held in Dallas, Texas, August 6–10. For many weeks before the conference C-Division staff actively worked behind the scenes to ensure its success.

- Ray Elliott, C-DO, was treasurer of the conference.
- Diana Tuggle, C-2, was chairman of the Slides Committee. In addition to her work on the conferences, Diana serves an ongoing position on the board of SIGGRAPH as visual aids editor.
- Jan Sander, C-2, served on the Slides Committee as its administrative assistant.

On a local level, Ray Elliott, C-DO, has served as Chairman of the ACM Rio Grande Chapter of SIGGRAPH since June 1989. Other C-Division members have also participated in this local SIGGRAPH chapter, either by serving on the committee or by presenting talks.

Supercomputing '89

Los Alamos continues to be very active in the supercomputing conferences. The most recent conference, *Supercomputing '89*, was held on November 13–17, 1989. The conference brought together supercomputing system researchers, designers, managers,

and computational scientists and engineers to report advances and experiences, state needs, suggest future directions, and exchange information.

- Ray Elliott, C-DO, served as Financial Chairman for the conference. He is actively involved with the 1990 conference to be held in New York City.
- Ted Reed, C-6, presented a one-day tutorial on computer graphics.
- Judy Deaguero, C-2, provided an on-site graphics service at this conference.
- Many other C-Division people actively participated in this conference by presenting talks and demonstrating supercomputing applications.

The *Supercomputing '91* Conference is coming to New Mexico. It will be held in the new Albuquerque Convention Center on November 18–22, 1991. Ray Elliott, C-DO, is chairing the Conference Committee. Dolores Roybal, C-DO, and Ann Hayes, C-DO, are on the conference committee.

Instrumentation Workshop

A "Workshop on Parallel Computer Systems: Performance Instrumentation and Visualization," co-sponsored by the Laboratory (Group C-3) and the Santa Fe Institute, was held May 8–10, 1989, in Santa Fe, New Mexico, in response to the tremendous enthusiasm generated by the first workshop, "Workshop on Instrumentation for Future Parallel Computing Systems" held in 1988. The goals of the workshop were to clarify features of performance instrumentation for current and future parallel computer systems and to determine which visualization techniques and tools would be appropriate for performance measurement of parallel computer systems.

The availability of integrated basic hardware and software for

the collection of performance data, coupled with the tools to enable graphical observation of the measurements, not only would teach us how to program such systems more effectively but also would lead to the design of more effective systems. Performance evaluation of parallel computer systems is a complicated process and a lively field of current research. This process is much more than mere benchmarking of the performance of parallel computers; it also involves an essential feedback mechanism for the design and implementation of hardware and software for the complex supercomputers of tomorrow.

Effective evaluation of these systems based on their performance will require new kinds of data. In addition, the complexity and quantity of data from performance measurements are so overwhelming that new techniques are needed to allow efficient and timely analysis to take place. One such technique is visualization, a method of computing that transforms symbolic data into geometric form, thus allowing researchers to observe their simulations and measurements. Ironically, as the computational power and performance of high-end supercomputing increases, it becomes harder to comprehend the complex results of the computational models. To that end, visualization has become a powerful, almost indispensable mechanism for the scientific user. Once the primary province of the entertainment industry, visualization is becoming a mainstay in the computer scientist's toolbox.

To accomplish the two goals of the workshop, selected designers, researchers, and users were brought together to exchange ideas on the role of visualization in the design of methods and tools for measuring the performance of tomorrow's supercomputers. The workshop began with two key-

note speakers, one from the performance measurement field and one from the visualization field. Daniel Reed, from the University of Illinois, noted that the lack of tools to observe the operation and performance of parallel architectures limits the user's ability to effectively optimize application and system performance. Performance data collection, analysis, and visualization tools are needed to manage the complexity and quantity of performance data. Furthermore, these tools must be integrated with the machine hardware, the system software, and the applications support software. He went on to describe the integrated performance environment he and his associates are developing for the Intel iPSC/2 hypercube.

In his talk, "Scientific Visualization or Just Computer Graphics?," Michael Wozny of Rensselaer Polytechnic Institute discussed model building tools and visualization tools. The model building tools allow a level of abstraction of a particular problem, and the visualization tools allow the visualization mapping of this level of abstraction. Simple visualization tools provide for a lot of insight into complex problems. The main difference between computer graphics and scientific visualization is the emphasis on the user interface in the latter. The visualization tool must be interactive and capable of separating a phenomena into its individual components.

Participants at the workshop learned of various research projects in both the areas of performance evaluation and scientific visualization through presented papers and demonstrations. Working groups, organized to address specific current problems related to the workshop theme, met for three days. Topics covered in the working groups were: (1) monitoring the perfor-

mance of parallel programs, (2) graphical aids to constructing parallel programs, (3) standards in performance instrumentation and visualization, and (4) dataflow and hybrid dataflow architectures.

The book, "Instrumentation for Future Parallel Computing Systems," published from last year's workshop, was such a success that Addison-Wesley has asked us again to publish the proceedings of the 1989 workshop as a book in the ACM "Frontiers in Computer Science" series. The proceedings will include the papers and descriptions of the demonstrations presented at the workshop as well as summaries from the working groups.

Once again the response to the workshop was very enthusiastic. We are currently planning a workshop for 1991 that will focus on the performance instrumentation and analysis issues for current and future parallel computing systems.

National Computer Graphics Association (NCGA) Conference

The NCGA conference, the premier exhibition of computer-graphics vendors, is attended by over 25,000 professionals from all over the world. Included in the four-day conference are vendor exhibits and presentations, technical papers, a video theatre, demonstrations of computer-graphics standards in action, and more. The focus of this annual conference is on the application of computer graphics in the workplace, rather than on current research activities.

The NCGA '89 Conference was held April 17-20, 1989, and the NCGA '90 Conference was held March 19-22, 1990. Ted Reed, C-6, participated in both NCGA conferences by chairing a session and helping present a one-day tutorial on computer graphics.

Gigabit Testbed Workshop

The Los Alamos Gigabit Testbed Workshop, sponsored by the Los Alamos National Laboratory, was held in Santa Fe on January 30-31, 1990. Representatives from industry, universities, and government laboratories were invited to participate in this workshop and to hear presentations about the Los Alamos plans to implement HIPPI channels and switches in a network environment. C-Division staff participated in this workshop, which provided a forum for discussion and initiation of efforts to develop a network testbed.

Scientific Supercomputing Workshop

A new program designed to provide advanced education for supercomputer use in intensive research projects has been organized by C-2. The program is aimed at scientists and engineers in computer-intensive research fields. Four to eight participants are accepted for each workshop, balancing research topics in sciences, engineering, and computer science. Each full-time, three-month workshop immerses participants in diverse aspects of supercomputing. Participants bring to the workshop their own research programs and, with help and support from Laboratory staff, adapt these programs to different supercomputers. Participants take part in frequent discussions with local and visiting faculty on topics relevant to scientific computing and to supercomputing programming. C Division provides each participant with a terminal and computer accounts on different machines; participants are expected to concentrate on the computer aspects of their own research and to present the results of their work at the end of the workshop.

Fall 1989 workshop participants and their research topics were:

- Robert Boland, Computer User Services, C-10. Research Area: Computational mathematics and mathematical software libraries
- Kuok-Mee Ling, Plasma Theory and Computation, CTR-6. Research Area: Plasma physics and controlled fusion research
- Laura Olvey, Thermonuclear Applications, X-2. Research Area: Radiation transport
- Rick Smith, Applied Nuclear Science, T-2. Research Area: Nuclear theory
- Eugene Symbalist, Atmospheric Sciences, SST-7. Research Area: Atmospheric science, hydrodynamics with radiation and non-LTE chemistry models

Workshop on Online Documentation in the Supercomputing Environment

Online documentation in the supercomputing environment got

the attention it deserves in a workshop conducted in Colorado Springs on April 3-6, 1990, and hosted by C-2 and the National Center for Atmospheric Research (NCAR). Thirty participants came from most of the leading supercomputing sites, including the national laboratories, National Institute of Standards and Technology (formerly National Bureau of Standards), the National Science Foundation supercomputing centers, university and regional centers, and Cray Research and IBM. Participants included system implementors, technical writers, consultants, users, and managers.

Evaluations and observations by participants indicated that the workshop was successful at bringing together computing people and technical writers to focus on problems and progress in online documentation.

- Marvin Harlow, X-7, and people from Cray Research said that this workshop was the most

Pictured are the graduates of C-Division's first workshop on scientific supercomputing, an intensive, no-fee program that immersed participants in diverse aspects of supercomputing and how they could apply them to their work. From left: Eugene Symbalist, SST-7; Rick Smith, T-2; Ann Solem, group leader of C-2, who coordinated the program; Robert Boland, C-10; Pat Malone, C-2, who also helped coordinate the program; Kuok-Mee Ling, CTR-6; and Laura Olvey, X-2.

valuable conference they had ever attended.

- With the advent of UNIX on supercomputers, for the first time in recent memory most supercomputing sites are running the same operating system. Thus, the opportunities for collaboration have never been greater.
- With the entry of Cray Research into the world of interactive operating systems, Cray is rapidly understanding that online documentation is now a necessary component of their product set. Cray Research commissioned the National Center for Supercomputing Applications at the University of Illinois to re-implement the DOCUMENT online documentation system (developed by the National Magnetic Fusion Energy Computer Center in LRLTRAN) in C for UNIX. Cray Research's DOCVIEW product, which may be released in UNICOS 6.0, may contribute to unprecedented collaboration between the vendor and customers—customers are excited about this product, and they want to combine local databases of documentation with vendor databases.
- Cray Research representatives at the workshop repeatedly expressed their desire to learn what the customer needs for online documentation.

This workshop came about as a result of discussions at the 1988 Super! meeting at Cornell, the 1988 Online Documentation Working Meeting at San Diego, the 1984-1989 Forums on Computer Systems for Documentation, and the Supercomputing '88 and Supercomputing '89 conferences. The purpose of this workshop was to bring together people of diverse backgrounds to actually work to understand the subject and to lay some tangible foundations for making progress.

This was a true workshop. All participants developed definitions and concepts before the meeting and presented this material as

poster sessions. To assure that everyone knew about existing system capabilities, 14 online documentation systems were demonstrated by dialing up supercomputing sites coast to coast. Then participants broke into teams to brainstorm and develop visions of online documentation. Participants are collecting their contributions into a working paper that will capture the facts and ideas and enable further progress in the development of online documentation; Los Alamos and NCAR are coordinating the working paper. Meanwhile, the steering committee, composed of representatives from all interested sites, will determine how to develop the many good ideas of ways to further expand knowledge, understanding, and collaboration.

International Standards Organization

Ted Reed, C-6, serves as a member of the International Standards Organization (ISO). In July 1989 he was the United States delegate to the Computer Graphics Interface Rappourteur Group meetings to resolve issues arising from the second draft proposal letter ballot. Ted filled in as chairman of the meeting and as US Head of Delegation at various times during the meeting. In October 1989, Ted was acting chairman of the CGI Rappourteur Group meeting that advanced the CGI from ISO Draft Proposed Standard status to ISO Draft International Standard status.

National Committees

C-Division staff regularly serve on and contribute to a wide range of scientific associations. Several members serve on editorial boards for journals or advisory committees. A few of the contributors are listed below.

- Pat Malone, C-2, is a member of the National Advisory Committee

to the Executive Director of the Association for Media-Based Continuing Education for Engineers.

- Jerry Melendez, C-8, is a member of the National Institute of Standards in Technology Panel for Computing and Applied Mathematics.
- Ron Pfaff and Jeff Brown, C-10, serve on Cray Research's C Advisory Board.
- Mark Roschke, C-10, serves on the NEC Super-UX Advisory Committee.
- Ann Solem, C-2, is a member of the Cray Research's Customer Advisory Board on Documentation and Training. Special interests of this board and Cray Research include online documentation and self-paced training.
- John Morrison, C-5, serves on Cray Research's Networking Advisory Board.
- Don Tolmie, C-5, is vice-chairman of the ANSI X3T9.3 Task Group and technical editor of the HIPPI documents.

Publications

C-Division employees are regular contributors to professional journals, conference proceedings, and other publications. This list includes work that was published between January 1989 and spring 1990 and is available to the public. Laboratory technical reports are available through the Laboratory's report library; handouts and internal reports are available from the authors.

Included is information about the group the author was in at the time the work was published; because of reorganization, authors might now be in different groups. In addition, in a few cases where work was prepared by multiple authors, the name of the C-Division author might not be included in this list; the name of the author is on the work and all publications listed include the work of at least one C-Division employee.

This section has been printed in smaller type to save space.

Morse, Norman R. (C-DO), "C-Division Annual Review and Operating Plan, February 1989," Technical Report, LA-11562-MS, Feb. 1989.

Koskela, Rebecca J. (C-8), Simmons, Margaret L. (C-3), and Wasserman, Harvey J. (C-3), "Performance Characterization of the Convex C-240 Computer System," Technical Report, LA-11769-MS, Jan. 1990.

Thomas, Timothy R. (C-10) and Brewster, Tony L. (C-10), "Experiments in Finding Neural Network Weights," Technical Report, LA-11772-MS, Jan. 1990.

Roussel-Dupre, Robert A. (SST-7) and Kelley, Thomas A. (C-3), "Transionospheric Propagation Code (TIPC)," Technical Report, LA-11799-MS, Feb. 1990.

Simmons, Margaret L. (C-3) and Wasserman, Harvey J. (C-3), "Los

Alamos Experiences with the IBM RS/6000 Workstations," Technical Report, LA-11831-MS, Mar. 1990.

Ottenstein, Karl J. (C-3), "PDG Testbed," Computer Code, LA-CC-89-0014, Jul. 1989.

Mniszewski, Susan Marie (C-10), Peter, Eva A. (C-10), and Lloyd, Stephen C. (C-4), "Local Area Network Diagnostic Assistant Expert System," Conference Paper, LA-UR-88-3827, Mar. 1989.

Tolmie, Donald E. (C-5), "High-Speed Channel (HSC) Standard," Conference Paper, LA-UR-88-3903, Feb. 1989.

Yantis, Bonnie (C-10), Rich, David (C-10), and Michelsen, Randy (C-10), "Paragraph-RTS—A Parallel Reduction-Based Run-Time System," Conference Paper, LA-UR-88-4110, Aug. 1989.

Griffin, James Harper (C-3), Hiromoto, Robert E. (C-3), and Wasserman, Harvey J. (C-3), "Parallel Debugging System User's Guide," Unclassified Handout, LA-UR-89-0157, Jan. 1989.

Faber, Vance (C-3), Manteuffel, Thomas A. (C-3), and Parter, Seymour V. (non LANL), "On the Theory of Equivalent Operators and Application to the Numerical Solution of Uniformly Elliptic Partial Differential Equations," Journal Article, LA-UR-89-0313, Jan. 1989.

Hiromoto, Robert E. (C-3), "Exploring Parallel Algorithms Having No Serial Analogues," Conference Paper, LA-UR-89-0314, Jan. 1989.

Brown, David L. (C-3), "On Composite Overlapping Grids," Conference Paper, LA-UR-89-0315, Apr. 1989.

Brown, David L. (C-3), "Sysode—A Test Code for Linear, Stiff, Two-Point Boundary Value Problems," Technical Report, LA-UR-89-0316, Jan. 1989.

Hiromoto, Robert E. (C-3) and Wienke, Bruce R. (X-4), "Chaotic, Single-Pass Inner-Iteration Scheme for the Discrete Ordinate Method SN," Technical Report, LA-UR-89-0359, Apr. 1989.

Zahrt, John D. (X-5) and Wing, G. Milton (C-3), "Time-Dependent Radiative Transfer in a Spherical Vacuum," Journal Article, LA-UR-89-0545, Feb. 1989.

Wendroff, Burton (T-7) and White, Andrew B. (C-DO), "Supraconvergent Scheme for Nonlinear Hyperbolic Systems," Journal Article, LA-UR-89-0777, Mar. 1989.

Hasinger, G. (non LANL), Priedhorsky, William C. (ESS-9), and Middleditch, John (C-3), "Third X-Ray Spectral Mode in Scorpius X-1—Similarities to the Other Bright Bulge Sources," Journal Article, LA-UR-89-1140, Mar. 1989.

Burks, Christian (T-10), Cinkosky, Michael J. (T-10), Gilna, Paul (T-10), et al., "Genbank—Current Status and Future Directions," Journal Article, LA-UR-89-1154, Mar. 1989.

Kubat-Martin, Kimberly A. (INC-4), Kubas, Gregory J. (INC-4), Khalsa, G. Rattan (INC-4), et al., "Molecular Hydrogen Complexes—Structural Characterization of W(CO)3(PR3)2(ETA Super 2-H2) (R = C5H9, C6H11) and MO(CO)3(P[C6H11]3)2(ETA Super 2-IIE)," Journal Article, LA-UR-89-1180, Apr. 1989.

Brown, David L. (C-3), Chesshire, Geoffrey (non LANL), and Henshaw, Willilam D. (non LANL), "Getting Started with CMPCGRD—Introductory User's Guide and Reference Manual," Handout, LA-UR-89-1294, Apr. 1989.

Phillips, Richard L. (C-6), "Distributed Visualization at Los Alamos National Laboratory," Journal Article, LA-UR-89-1531, May 1989.

Hoebelheinrich, Randall L. (C-5), "Multiple Crossbar Network—Integrated Supercomputing Framework," Conference Paper, LA-UR-89-1627, Nov. 1989.

Griffin, James Harper (C-3) and Hiromoto, Robert E. (C-3), "Experience with an Experimental Debugger for Parallel Programs," Journal Article, LA-UR-89-1702, May 1989.

Pfaff, Ronald T. (C-10), Sanders, Claudia (C-2), and Cohen, Marion (C-2), "Software Tools for Controlling, Distributing, and Installing Software at Los Alamos National Laboratory," Conference Paper, LA-UR-89-1829, Nov. 1989.

Coleman, Judith B. (HSE-1), Pfaff, Ronald T. (C-10), and Cohen, Marion F. (C-2), "Change Control of Software at LANL [Los Alamos National

Laboratory], "Conference Paper, LA-UR-89-1832, Apr. 1989.

Mandell, David A. (X-3), Adams, Thomas F. (X-3), and Holian, Kathleen S. (C-3), et al., "Mesa—A 3-D Computer Code for Armor/Anti-Armor Applications," Conference Paper, LA-UR-89-1851, Oct. 1989.

Sydoriak, Stephen J. (C-2), "LaTeX Memos and Letters," Conference Paper, LA-UR-89-2125, Aug. 1989.

Holian, Kathleen S. (C-3), "Hydrodynamics Code Calculations of Debris Clouds Produced by Ball-Plate Impacts," Conference Paper, LA-UR-89-2309 Dec. 1989.

Hagan, Patrick S. (T-7) and Cox, Robert W. (C-3), "Productivity Enhancement in Hydrofractured Oil Reservoirs," Journal Article, LA-UR-89-2380, Jul. 1989.

Cohen, D. S. (non LANL) and White, Andrew B. (C-DO), "Sharp Fronts Due to Diffusion and Viscoelastic Relaxation in Polymers," Journal Article, LA-UR-89-2428, Jul. 1989.

Patel, Nisheeth R. (non LANL), Sturek, W. B. (non LANL), and Hiromoto, Robert E. (C-3), "Parallel Numerical Simulation for Supersonic Flows Using Zonal Overlapped Grids and Local Time Steps for Common and Distributed Memory Multiprocessors," Conference Paper, LA-UR-89-2429, Sept. 1989.

Brown, David L. (C-3), Chesshire, Geoffrey (non LANL), and Henshaw, William D. (non LANL), "Composite Grid Data," Handout, LA-UR-89-2475, Jul. 1989.

Tolmie, Donald E. (C-5), "High-Speed Channel (HSC) and Ultra-Speed Networks," Conference Paper, LA-UR-89-2560, Oct. 1989.

Burns, Patrick J. (non LANL), Christon, Mark (non LANL), Schweitzer, Roland (non LANL), et al., "Vectorization of Monte Carlo Particle Transport—An Architectural Study Using the LANL Benchmark "GAMTEB", " Conference Paper, LA-UR-89-2560, Nov. 1989.

Chorn, Granville E. (C-8), Koskela, Rebecca J. (C-8), and Wilkins, Sally F. (C-8), "Network Load Characterization: 1989," Handout, LA-UR-89-2614, Aug. 1989.

Campbell, John R. (C-3) and McGavran, Lauren P. (C-3), "Integrated Distributed Processing Interface for Supercomputers and Workstations," Conference Paper, LA-UR-89-2706, Sept. 1989.

Gunzburger, Max (non LANL), Meir, Annon J. (non LANL), and Peterson, Janet S. (C-3), "On the Existence, Uniqueness and Finite Element Approximation of Solutions of the Equations of Stationary, Incompressible MHD," Journal Article, LA-UR-89-2708, Aug. 1989.

Sanders, Claudia (C-2), "Bulletin: An Authoring and Change Control Tool," Conference Paper, LA-UR-89-2722, May 1990.

Kristian, J. (non LANL), Pennypacker, C. R. (non LANL), Middleditch, John (C-3), et al., "Submillisecond Optical Pulsar in Supernova 1987A," Journal Article, LA-UR-89-2820, Aug. 1989.

Middleditch, John (C-3), "Adventures in Fourier Space: Hunting the Millisecond Pulsar," Journal Article, LA-UR-89-2821, Aug. 1989.

Hammes, Jeffery P. (C-10), Yantis, Bonnie (C-10), Michelsen, Randy (C-10), et al., "Efficient Unicos Run-Time System for Functional Programs on a Cray X-MP," Conference Paper, LA-UR-89-2974, Nov. 1989.

Rich, David O. (C-10) and Michelsen, Randy E. (C-10), "Writing Parallel, Discrete-Event Simulations in Modsim—Insight and Experience," Conference Paper, LA-UR-89-3104, Sep. 1989.

Ottenstein, Karl J. (C-3), "Simplified Framework for Reduction in Strength," Journal Article, LA-UR-89-3138, Sept. 1989.

Hoebelheinrich, Randall L. (C-5) and Thomsen, Richard G. (C-5), "Multiple Crossbar Network—A Switched High-Speed Local Network," Conference Paper, LA-UR-89-3155, Oct. 1989.

Hoebelheinrich, Randall L. (C-5) and Thomsen, Richard G. (C-5), "Multiple Crossbar Network: Integrated Supercomputing Framework," Conference Paper, LA-UR-89-3156, Nov. 1989.

Baillie, Clive (non LANL), Brickner, Ralph G. (C-3), Gupta, Rajan (T-DO), et al., "QCD [Quantum Chromodynamics] with Dynamical Fermions on the Connection Machine," Conference Paper, LA-UR-89-3311, Nov. 1989.

Saltzman, Jeffrey S. (C-3), "Unsplit 3-D Upwind Method for Hyperbolic Conservation Laws," Journal Article, LA-UR-89-3442, Oct. 1989.

Scovel, J. Clint (C-3), "Symplectic Numerical Integration of Hamiltonian

Systems," Conference Paper, LA-UR-89-3465, Jun. 1989.

Locke, Timothy M. (C-6) and Miranda, Ann S. (non LANL), "OMNI Development of an Engineering Product Data Management System," Conference Paper, LA-UR-89-3474, Aug. 1990.

Tenbrink, Stephen C. (C-5), "Is FDDI Fast Enough for High Speed Graphics?" Journal Article, LA-UR-89-3496, Oct. 1989.

Thudt, Michael E. (P-11), Burczyk, Leonard (non LANL), Dalesio, Leo R. (AT-8), et al., "Los Alamos Accelerator Control System," Conference Paper, LA-UR-89-3571, Oct. 1989.

Faber, Vance (C-3), "A, B, C Bits—A Scheme for Eliminating Livelock in Desperation Routing Networks," Handout, LA-UR-89-3640, Oct. 1989.

Holm, Darryl D. (T-7), Lysenko, Walter P. (AT-6), Scovel, J. Clint (C-3), "Moment Invariants for the Vlasov Equation," Journal Article, LA-UR-89-3652, Oct. 1989.

Ottenstein, Karl J. (C-3), Ballance, Robert A. (non LANL), and Mac Cabe, Arthur B. (non LANL), "Gated Single-Assignment Form: Dataflow Interpretation for Imperative Languages," Conference Paper, LA-UR-89-3654, Nov. 1989.

Christman, Ronald D. (C-3), Cook, Danny P. (C-8), and Mercier, Christina W. (C-8), "Data Storage for the 1990s," Handout, LA-UR-89-3697, Nov. 1989.

Brickner, Ralph G. (C-3) and Baillie, Clive (non LANL), "Pure Gauge QCD on the Connection Machine," Journal Article, LA-UR-89-3739, Nov. 1989.

Cox, Robert W. (C-3), "Stress-Assisted Diffusion—A Free Boundary Problem," Journal Article, LA-UR-89-4088, Dec. 1989.

Kao, Chih-Yue Jim (EES-5), Glatzmaier, Gary Andrew (EES-5), Malone, Robert C. (C-3), et al., "Global Three-Dimensional Simulations of Ozone Depletion Under Post-War Conditions," Journal Article, LA-UR-89-4127, Dec. 1989.

Miller, Janice N. (C-3) and Chen, I-Li (HSE-8), "Terrain-Preprocessor Programming Effort," Handout, LA-UR-89-4146, Dec. 1989.

Ottenstein, Karl J. (C-3), "Unified Approach to Program Transformation," Handout, LA-UR-89-4213, Dec. 1989.

Christman, Ronald D. (C-3), Cook, Danny P. (C-8), and Mercier, Christina W. (C-8), "Re-Engineering the Los

Alamos Common File System," Conference Paper, LA-UR-90-0032, Jan. 1990.

Collins, Maurice W. (C-8), Mercier, Christina W. (C-8), and Stup, Thomas D. (C-8), "Mass Storage System Advances at Los Alamos," Conference Paper, LA-UR-90-0061, Jan. 1990.

Mueller, F. M. (CMS), Fowler, Clarence M. (M-6), and Freeman, Bruce L. (M-6), et al., "Theory of the De Haas-Van Alphen Effect in Pulsed Magnetic Fields," Journal Article, LA-UR-90-0130, Jan. 1990.

Phillips, Richard L. (C-6) and Beach, Richard J. (non LANL), "Image Manipulation in the SIGGRAPH Interactive Proceedings," Conference Paper, LA-UR-90-0540, Feb. 1990.

Smith, Suzanne T. (N-4) and Jalbert, M. L. (C-3), "LAVA/CIS Version 2.0: A Software System for Vulnerability and Risk Assessment," Conference Paper, LA-UR-90-0596, Feb. 1990.

Bucher, Ingrid Y. (C-3), "Access Contention in Multiprocessor Memories, Part I: Queuing Models,"

Conference Paper, LA-UR-90-0605, Feb. 1990.

Steiman-Cameron, Thomas Y. (non LANL), Imamura, James N. (SST-9), Kristian, J. (non LANL), et al., "190 S Periodicity in the Optical Emission of the Enigmatic X-Ray Source GX 339-4," Journal Article, LA-UR-90-1208, Apr. 1990.

Forest, M. G. (non LANL) and Wang, Qi (C-3), "Change-of-Type Behavior in Viscoelastic Slender Jet Models," Journal Article, LA-UR-90-1209, Apr. 1990.

Imamura, James N. (SST-9), Kristian, Jerome (non LANL), Steiman-Cameron, Thomas Y. (non LANL), et al., "Eight Second Optical Quasi-Periodic Oscillations in GX 339-4," Journal Article, LA-UR-90-1211, Apr. 1990.

Hemberger, Philip H. (CLS-1), Alarid, J. E. (CLS-1), Cameron, Douglas (CLS-1), et al., "Transportable Turnkey Gas Chromatograph/ION Trap Detector for Field Analysis of Environmental Samples," Journal Article, LA-UR-90-1437, Apr. 1990.

Ballance, Robert A. (non LANL), Mac Cabe, Arthur B. (non LANL), and Ottenstein, Karl J. (C-3), "Program Dependence Web: A Representation Supporting Control-, Data-, and Demand-Driven Interpretation of Imperative Programs," Conference Paper, LA-UR-90-1474, May 1990.

Figure A-1. During FY 1989, 351 FTEs were available to C Division, 345 within our own division, and 6 from other Laboratory organizations that provided support to our recharge programs. FTE levels are expected to remain relatively constant in FY 1990.

Figure A-2. About 85% of the 351 total FTEs available to C Division in FY 1989 worked on the computing and communications recharge program. In FY 1990, increases in other reimbursable programs are expected to offset decreases in computing recharge and defense programs.

Figure A-3. During FY 1989, C Division received a total of \$86.3M in operating funds. About 85% of these revenues were from the computing and communications recharge programs. Estimated FY 1990 operating funds are \$88.5M.

Figure A-4. During FY 1989, total operating expenses were \$86.3M (\$0.9M of which supported other Laboratory organizations that provided support to our recharge programs). FY 1990 expenses are estimated to be \$88.5M.

Figure A-5. Defense programs continue to be the largest computing client, contributing 60% of FY 1989 revenue.

Figure A-6. The FY 1990 CTSS rate of \$120 per hour is one-fourth that charged in FY 1981.

Table A-1: FY 1990 Rates for Computing Services^a

Worker Machine Services:	
CTSS (recorded hour)	\$120.00 ^b
UNICOS (recorded hour)	120.00 ^b
NOS (Cyber 825) (recorded hour)	240.00 ^b
NOS (Cyber 855) (recorded hour)	270.00 ^b
Tape mount (each)	4.85 ^b
ICN Services:	
ICN access (month)	82.00
Connect (hour)	2.40 ^b
Network transmission (Mword)	2.50 ^b
CFS Services:	
CFS access (each)	0.15 ^b
CFS storage (Mword/month)	0.32
Tape storage (reel/month)	3.15
File count (file/month)	0.05
PAGES Services:	
Media:	
Paper - Line printer (sheet)	0.13 ^c
Graphics/document (sheet)	
Black and white	0.25
Color	0.50
Rolled (sq. ft.)	0.20
Film - 16 mm (frame)	0.07
35 mm (frame)	0.15
105 mm (fiche)	2.35
Photoplotter (sheet)	5.25
Processing:	
Photoplotter services	
Setup (job)	10.00
Device (hour)	35.00
All other PAGES services (Mbyte)	2.30
Other Services:	
VMS/UNIX (month)	35.00

^aA 30% surcharge is added for external users other than those at DOE Integrated Contractor Facilities

^bCharge is subject to shift weights (day = 2.0, night = 1.0, weekend = 0.5)

^cPages are printed on both sides

Table A-2. FY 1990 Communication Services Rates

FTS	0.64/min ^a
FX (ALB. Exchange)	0.19/min ^a
WATS (Intrastate)	0.25/min ^a
Nevada Tie Line	0.19/min ^a
USW Labor	80.00/hr ^a
Radio	11.00/mo
Channel - under 4800	23.00/mo
Channel - 4800 - 19.2K	47.00/mo
Channel - 56K	75.00/mo
Channel - 256K	150.00/mo
Channel - T-1	300.00/mo
Channel - TV Trans.	600.00/mo
LABNET - Video Port	30.00/mo
LAN Gateways	300.00/mo
Message Center - Individual	18.00/mo
Message Center - Group	117.00/mo
Crypto Maintenance	350.00/mo

^aPlus communications surcharge 24.7%

Table A-3. Los Alamos National Laboratory Major ADP Acquisitions (\$k)

Line Item	FY 1990		FY 1991		FY 1992		FY 1993		FY 1994		FY 1995		FY 1996	
	Cap	Exp	Cap	Exp	Cap	Exp	Cap	Exp	Cap	Exp	Cap	Exp	Cap	Exp
CRAY X-MP (4th)		4560		760										
SC*(88)	10595	794												
SC (89)		1879	2577	342										
SC (90)		962		7384	11657	5480	13241	808						
SC (91)														
SC (92)				1306		2612	2612	435	3499	435				
						2730	5460	910	7212	910				
SC (93)														
SC (94)							5324	10648	10648	1774	14116	14116	1774	
SC (95)														
SC (96)													10648	5324
MS ** Upgrade	1855		2300		1200		2065		2550		3100		2850	
OS*** Upgrade			1200						1620					
Network Upgrade			1623		1443		1440		1530		2125		2460	
Total	12450	8195	12900	9792	14300	10822	16746	14204	16361	17317	19341	17746	19426	17746

*CRAY supercomputer.

**MS = mass storage.

***OS = output station.

Glossary

ACL
Advanced Computing Laboratory.

ACM
Association for Computing Machinery.

ACS
Account Control System.

Administrative partition
That portion of the Integrated Computing Network primarily devoted to processing sensitive unclassified data or data subject to a privacy act.

ADPE
Automatic data processing equipment.

AIMS
Advanced Information Management Systems.

ANSI
American National Standards Institute.

Apollo
The brand name of a scientific workstation made by the Apollo Computer Company.

ARPANET
The experimental research and development network that serves a large number of nongovernment research centers. See DDN.

ASCII
American Standard Code for Information Interchange.

AT&T
American Telephone and Telegraph Company.

AVM
Airborne vapor monitor.

AVS
Application Visualization System.

backbone
The primary communications systems used to interconnect various subnetworks or local area networks. For example, the broadband cable we call LABNET is used to interconnect various "open" local area networks, so we could call LABNET a backbone.

bit/s
Bits per second.

CAD/CAM
Computer-Aided Design/Computer Aided Manufacturing.

CBI
Crossbar Interface.

CBS
crossbar switch.

CCF
Central Computing Facility. The physical location of the main Los Alamos computing center.

CDC
Control Data Corporation.

C Division
Computing and Communications Division.

CFS
Common File System.

CFT
Cray Fortran compiler provided by Cray Research, Inc.

CFTLIB
Cray Fortran I/O Library.

CGI
Computer Graphics Interface.

CGL
Computer Graphics Laboratory.

CGM
Computer Graphics Metafile. An ANSI standard.

CGS
Common Graphics System. A basic system of computer graphics primitives that is available for all supported graphics devices and operating environments in the ICN.

CIC
Computing Information Center.

CIM
Computer Integrated Manufacturing.

CLAMDOC
Online documentation for the CLAMS library.

CLAMS
Common Los Alamos Mathematical Software, a library of mathematical routines based on the SLATEC Common Math Library.

CM-2
Connection Machine, manufactured by Thinking Machines, Corporation.

CMPGRD
A package used to generate composite overlapping finite-difference grids.

CNT
Central Networking Testbed.

COS
Cray operating system.

CP*
Cross Point Star.

CRI
Cray Research, Inc.

CTL
A command language controller developed at Los Alamos for CTSS.

CTSS
Cray Timesharing System.

DARPA

Defense Advanced Research Projects Agency.

DDN

Defense Data Network (formerly ARPANET).

DDS

Digital Dataphone Service.

DEC

Digital Equipment Corporation. The manufacturer of the DEC VAX-11 family of minicomputers embedded in the Integrated Computing Network to perform dedicated switching, control, and monitor functions. VAXs are host computers and distributed processors.

DECnet

Digital Equipment Corporation's standard intermachine protocol that C Division uses between XNET and distributed processors.

dedicated

A term used to indicate that a computing resource is reserved for a specific purpose.

DES

Data Encryption Standard.

desktop workstation

A small computer system that is able to be placed in a normal office environment. These systems typically have the following characteristics: a multiprocess interactive operating system, a bit-mapped display, a window-based user interface, and strong network communications software.

DISSPLA

Graphics system that supports high-quality plotting and text capabilities.

distributed computing

The integration of several computer processes running on different computers across a network to accomplish a task. This is distinct from remote computing in that it is dependent on several processes or programs being orchestrated into a unit to achieve the desired result.

distributed processor (DP)

A computer usually located at a user's site that is used for computing purposes such as text processing, highly interactive graphics, and less intensive calculations. Distributed processors are connected to the ICN through the XNET gateways, and C Division provides network support so they can access ICN worker computers and services.

DNA

Defense Nuclear Agency.

DoD

Department of Defense.

DOE

Department of Energy.

DP

Distributed processor (see definition).

DST

Dedicated system time.

ECCOS

Electronic Card Catalog and Ordering System.

Ethernet

A low-level, baseband, local area data communications network developed by Xerox and supported by DEC and Intel, among others.

FASTCAT

For Automated Search of the CAtalog.

FCL

Fortran command language.

FDDI

A new LAN technology; Fiber Distributed Data Interface token rings.

FKIs

Integral equations of the first kind.

FOCUS

Facility for Operations Control and Utilization Statistics.

Form B

A method of recharging for Laboratory services.

FPS

Floating Point Systems.

FT

File transport. A Gould 67 mini-computer that is used as a high-speed packet switch in the ICN.

FTE

Full-time equivalent. One FTE is equivalent to one full-time Laboratory employee.

FTS

Federal Telecommunications System.

FY

Fiscal year.

Gbit

The abbreviation for gigabit, which is one billion bits.

Gbyte

The abbreviation for gigabyte, which is one billion bytes.

GFLOPS

Trillion floating-point operations per second.

GKS

Graphical Kernel System, a proposed international graphics standard.

GRA

Graduate Research Assistantship.

HBCU

Historically Black Colleges and Universities.

HIPPI

High-performance parallel interface.

host computer

A computer in the ICN that executes user programs.

HSC

High-speed channel.

HSPI

High-speed parallel interface.

IBM

International Business Machines Corporation.

ICN

Integrated Computing Network. The Laboratory's main computer network consisting of host computers, file storage, network services, local and remote terminals, data communication interfaces, and distributed processors.

ICNHELP

Online help utility.

IDEs

integro-differential equations.

IGPP

Institute of Geophysics and Planetary Physics.

IMSL

International Mathematical and Statistics Library.

INFORM

Information Management System, a Laboratory-wide system to support management information services activities.

IOS

Disk controller for the Crays.

Input/Output System

A system that provides the ICN with the capability of exporting and importing files using various types of magnetic tape.

IPC

Interprocess communication.

ISO

International Standard Organization.

ISRD

Institutional Supporting Research and Development.

ITR

Information Technology Resources.

kbit/s

Kilobits per second.

KCC

Keyboard Communications Concentrator.

LABNET

Los Alamos Broadband Network.

LAICS

Los Alamos Integrated Communications System.

LAN

Local area network.

LANL

Los Alamos National Laboratory.

LDCC

Laboratory Data Communications Center.

LIBC

Cray C Library.

LLNL

Lawrence Livermore National Laboratory.

LNL

Linear Nonlinear Learning.

MAPPER

A locally developed utility used to prepare computer graphics.

Mbit/s

Million bits per second.

MCN

Multiple Crossbar Network.

metafile

A device-independent, computer-independent representation of graphical data, the data interface to PAGES for graphics.

MFLOPS

Million floating-point operations per second.

MHD

magnetohydrodynamics.

MILNET

A packet data network operated by the Department of Defense.

MS

Multistage networks.

MVS

Multiple Virtual System. An IBM operating system for medium and large computers. Used on CFS.

Mword

Million words.

National Security partition

The fourth partition in the ICN.

NCAR

National Center for Atmospheric Research.

NCS

Network Computing System.

NET90

Network Modernization Project.

NFS

Network File Server.

NOC

Network Operations Center.

NOS

Network Operating System.

NQS

Network Queuing System.

NSA

National Security Agency.

NSC

Network Security Controller.

NSF

National Science Foundation.

NWC

Nuclear Weapons Complex.

NWCNET

Nuclear Weapons Complex Network.

Open partition

That portion of the ICN that is available to anyone who has a legitimate need to compute at Los

Alamos. Only unclassified computing is permitted in the Open partition. Users need not have a Q-clearance.

PABX

Private Automated Branch Exchange.

PAGES

Print and Graphics Express Station. The node in the Integrated Computing Network that is used to produce microfilm, microfiche, electrostatic printer/plotter output, and high-speed printer output.

PC

Personal computer.

PDE

Partial differential equation.

PE

Processing elements.

PMS

Password Management System.

port switch

A port contention device used to pass characters to host computers that have character-oriented systems or to KCCs for line-oriented systems.

PROD

Production or batch workload manager.

Pueblo Complex

The location of some C-Division user services; accessible by all computer users in Los Alamos.

QCD

Quantum chromodynamics code for use in parallel processing research.

RFP

Request for Purchase.

RGB

Red, green, blue.

RPC

Remote procedure call.

SDI

Strategic Defense Initiatives.

Secure partition

That portion of the ICN that may contain classified data. Classified computing in the ICN may be done only on computers in this partition.

SERS

Science and Engineering Research Semester.

SI

Simple Interface. A collection of communication tools to allow one to modify a batch-oriented program running on a Cray so that it has a window interface on a UNIX workstation.

SIGGRAPH

Special Interest Group on Computer Graphics sponsored by the Association for Computing Machinery.

SIMP

A set of locally developed network protocols.

SIPR

Secure IP router.

SLATEC

Sandia, Los Alamos, Air Force Weapons Laboratory Technical Exchange Committee. Also includes LLNL, NMFECC, and Sandia Laboratories/Livermore.

SSD

Solid-state Storage Device, a high-speed memory device used with the CRAY X-MP and Y-MP computers.

ST

Static topology.

STC

Society for Technical Communication.

STEP

Skills Training Employment Program.

SYNC

Synchronous Communications Concentrator. A DEC PDP-11 minicomputer that is used as a packet switch in the ICN.

TA-3

Abbreviation for the Laboratory's Technical Areas. C Division is located in TA-3.

TAG

Transition Advisory Group, formed to help ease the transition to UNICOS.

TCP/IP

Transport Control Protocol/Internet Protocol.

T Division

Theoretical Division.

TEKES

Technical Security Countermeasures Expert Systems.

Telenet

A commercial service network that provides circuit connections from one location to another.

terminal network

Allows any user to access any worker computer provided the user has authorization to do so.

terminal passthrough

An XNET capability that will enable a terminal that is physically attached to a remote computer to be used for interaction with the host computers.

testbed

A collection of equipment configured for experimental or testing purposes. A testbed is normally not used for "production."

TeX

A typesetting system (or a document formatting language) that is used to prepare high-quality technical manuscripts. It is particularly suited for mathematical and tabular work.

TMC

Thinking Machines Corporation, manufacturer of the Connection Machine.

TROFF

A text formatting program that produces letter-quality or typeset output on UNIX systems.

UNICOS

An operating system developed by Cray Research for use on Cray computers. It is based on AT&T System V UNIX.

UNIX

An operating system developed by Bell Laboratories and used on some DEC PDP-11 computers and VAX computers in the Laboratory. UNIX is a trademark of Bell Laboratories.

UNM

University of New Mexico.

VAX

Virtual Address Extension Computer. An upward extension of the DEC PDP-11 family of minicomputers.

VMS

Virtual Memory System. The interactive operating system on many VAX computers at Los Alamos. VMS

was developed by Digital Equipment Corporation.

WAVE

A particle code in the benchmark set.

WBCN

Wideband Communications Network.

WP

Word processing.

WPE

Word processing equipment.

WX Division

Design Engineering Division.

X Division

Applied Theoretical Physics Division.

XNET

Extended Network Access System. Provides ICN access for remote computer sites.

X3

Computing and communications services are recharged to the user organization via the X3 program code.

This report has been reproduced directly from
the best available copy.

Available to DOE and DOE contractors from
the Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831
prices available from
(615) 576-8401, FTS 626-8401

Available to the public from
the National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Rd.
Springfield, VA 22161

Microfiche A01

NTIS		NTIS		NTIS		NTIS	
Page Range	Price Code	Page Range	Price Code	Page Range	Price Code	Page Range	Price Code
001-025	A02	151-175	A08	301-325	A14	451-475	A20
026-050	A03	176-200	A09	326-350	A15	476-500	A21
051-075	A04	201-225	A10	351-375	A16	501-525	A22
076-100	A05	226-250	A11	376-400	A17	526-550	A23
101-125	A06	251-275	A12	401-425	A18	551-575	A24
126-150	A07	276-300	A13	426-450	A19	576-600	A25
						601-up*	A99

*Contact NTIS for a price quote.

END

DATE FILMED

12 / 18 / 90