NMFECC Network History

Jim Leighton

The Beginning

Kirby Fong's article in this month's *Buffer*, "Origins of the National Magnetic Fusion Energy Computer Center," details how NMFECC was selected in 1973 to meet the ever growing computing needs of geographically distant researchers in the field of magnetic fusion energy. The establishment of a single, central facility to provide supercomputer support through a data communications network was extremely foresighted. At the time, there was only one major computer network, the ARPAnet, which had been established to do research in the networking area, and no one was using supercomputers on a regular production basis via a data communications network. NMFECC is now providing this service through an associated network known as the MFEnet. This article describes the evolution of that network.

The First RJE Station

NMFECC began operation in July 1974 with a borrowed Control Data Corporation (CDC) 6600 computer. The data communications capability consisted of four dial-up lines, manually answered, each capable of running at ten characters per second. We began construction of a remote job entry (RJE) station based on a Digital Equipment Corporation (DEC) PDP-8 minicomputer, connected to the CDC 6600 through a leased line running at 4800 bits per second. This RJE station, delivered to Princeton Plasma Physics Laboratory (PPPL) in February 1975, was capable of outputing text or graphical data, and also provided a card reader for data input.

We subsequently made a number of improvements to the RJE station: (1) In an effort to reduce plotting time, we devised a data compression scheme to compress graphical data before it was sent to be plotted. This scheme is still in use today. (2) We designed new software and hardware to attach terminals directly to the RJE station. In August 1975, we added four terminals to the PPPL RJE station, each capable of running at 30 characters per second. (3) We replaced the original four dial-up modems on the CDC 6600 with 16 self-answering dial-up modems.

The Terminal Concentrator

In September 1975, we installed a CDC 7600 and made it available to remote users by connecting the 16 dial-up ports of the CDC 6600 so that they could also access the CDC 7600. We built and installed a new PDP-11-based terminal concentrator to provide better terminal access to the CDC 7600. By January 1976, twelve dial-up ports operating at 30 characters per second were available, and we installed one at the new dial-up speed of 120 characters per second at the University of California at Los Angeles (UCLA) and the University of California at Berkeley (UCB). By June 1976, we had increased dial-up access to sixteen 30-characters-per-second ports and four 120-characters-per-second ports. In December 1975, we moved PPPL's RJE station from the CDC 6600 to the CDC 7600.

The Initial Data Communications Plan

Meanwhile, the initial data communications plan began to take form. The hub would be a PDP-11/50 at NMFECC attached by four 50-kilobits-per-second dedicated phone lines and two direct lines to PDP-11/40s, with each PDP-11/40 serving as a front-end communications processor for each of six PDP-10s. We delivered and installed the PDP-10s by October 1975, and installed the 50-kilobits-per-second lines by January 1976. By June, network operation had progressed to the point where a remote PDP-10 user could submit a batch job to the CDC 7600 and have output from the job returned to his PDP-10. By July, file transport between PDP-10s was available. By August, the network included an initial version of software that supported both terminal traffic and file traffic on the same communications link.

In December 1976, we added a seventh PDP-10 (already in existence at Scientific Applications, Inc. (SAI) of La Jolla, California) to the network. It was connected to the PDP-11/40 at General Atomic (GA) in San Diego, California, requiring the PDP-11/40 to do the routing of traffic for the GA and SAI sites. This was the first time more than one user site had been supported on a single communications link.

Reliability

At this point, the framework of a functional data communications network was in place. Although the network was operational, reliability was a major issue and remains today a very large consideration in our network planning.

One step taken to improve reliability was the development of a "watch-dog" timer, installed in each remote PDP-11. If the PDP-11 failed to reset the timer every few seconds, it would force the PDP-11 to restart. This simple device eliminated the necessity of contacting remote site personnel to manually restart a failed PDP-11. Another step taken was to use the Center's second PDP-11/50 as a "hot standby" in case a remote PDP-11 failed. We installed an elaborate bus switch to allow all communications links to be switched from a failed processor to the functioning one.

The Computer Standard Interface Channel

The Center experienced a significant increase in power with the arrival of a CRAY-1 supercomputer, which was operational by May 1978. This machine was about four times faster than the CDC 7600. We connected the CRAY-1 to the CDC 7600, which served as a front-end for connections to the network and the terminal concentrator. This connection scheme was adequate, but limited, and we made plans to connect the CRAY-1 directly to the network and the concentrator.

The engineering group had been busy developing new hardware, called the Computer Standard Interface Channel (CSIC), to facilitate interconnections between all of the different types of computers in use at the Center. In June 1978, the CDC 7600 started using a CSIC channel to the network. Connections to the network and terminal concentrator were completely operational by August 1978. The next step was to develop an interface for the CRAY-1. By June 1979, the CRAY-1 had direct connections to both the network and to the new terminal concentrator. The new terminal concentrator was a combination of new hardware and software designed to replace the initial PDP-11-based concentrator. The new hardware controller was specifically designed to allow use as pairs of units, allowing a failed unit to be replaced in a matter of seconds by throwing a switch.

The RUSS and the NAP

Network planning was now aimed toward providing access for user sites that did not have PDP-10s. Two efforts were under way. We first refined the old RJE concept by using a PDP-11 that supported 8 to 16 local terminals and a printer/plotter. Because the name RJE implied a batch operating environment and we wished to emphasize the timesharing provided by the operating systems on the CDC 7600 and the CRAY-1, we renamed the station Remote User Service Station (RUSS). We directed another effort toward user sites that wanted to connect their local computers to the network. Again based on a PDP-11, we developed the concept of a Network Access Port (NAP). A NAP would be inserted between the remote site computer and the network to simplify the task of connecting it, and also to protect the network from inadvertent or intentional wrong-doing on the part of the user.

In October 1978, we installed the first RUSS using a 4800 bits-per-second dedicated line to Berkeley. Requests for similiar installations at the University of Washington, University of Wisconsin, and New York University were already at hand, and we were receiving additional requests regularly. We scheduled the first NAP installation to connect a VAX computer at the University of Texas to the network. By October 1980, the NAP was operational with file output to the VAX and with interactive terminal traffic, using a 4800 bits-per-second link. File input from the VAX followed a few months later.

The Satellite Communications Link

In June 1980, we selected American Satellite Corporation to provide dual links to replace each 50-kilobits-per-second link. Initially, each link pair carried the same traffic, with the receiver selecting one stream of data or the other by a manual switch. We later enhanced this scheme to allow both links to carry independent data, with the network automatically switching the data to one link if the other was failing. The proposed system required dual seven-meter antennae at the Center and dual five-meter antennae at the four major remote sites.

By April 1981, the first satellite link was online to Princeton, and it operated flawlessly — a tribute to our months of preparation. We didn't tell Princeton that they were on the satellite for several days to avoid biasing their observations. For the most part, they were unable to tell that they were communicating via a satellite link. Subsequently, we installed additional links at Oak Ridge National Laboratory (ORNL), Los Alamos National Laboratory (LANL), and GA. Satellite communications have proved to be a dramatic improvement in reliability.

The Japanese Connection

We held an initial meeting in November 1983 at Nagoya, Japan to discuss possible approaches to connecting the computer center at the Institute of Plasma Physics in Nagoya with NMFECC. Actual work on the link began in late 1984, and by July 1985, the link was available for use. This link now allows any remote authorized user on the MFEnet to log onto the Nagoya host computer. Conversely, authorized Japanese users may log onto or transport files to and from any host computer on the MFEnet.

Expansion to the ER Community

Until mid-1983, NMFECC had provided services almost exclusively to researchers in the field of magnetic fusion energy. In May 1983, the Center was asked to expand its services to include researchers in other fields of energy research, under the Office of Energy Research (ER). We will add a total of thirteen new ER sites to the network in 1985, approximately doubling the number of NMFECC users to an estimated total of 3500. The ER community funded a third CRAY computer at the Center, a CRAY X-MP, which we installed in November 1984, nearly doubling the Center's computer power.

The CRAY-2

Our most recent hardware acquisition is a CRAY-2 computer that contains four central processors and 64 million words of memory. This machine represents as much computer power as the rest of the Center's CRAY computers combined. The CRAY-2 was slightly more difficult to integrate into the network, because it has a new architecture and instruction set, but connections were completed in a matter of weeks.

Where We Are Today

Currently, the MFEnet includes a CRAY-1, a CRAY-1S, a CRAY X-MP, a CRAY-2, a CDC 205 (at Florida State University), the original seven PDP-10s, 22 DEC VAXs, and over 80 PDP-11s serving as NAPs, RUSSs, and communication processors.

We have expanded dial-up service to include 24 ports at 120 characters per second, and 8 at 30 characters per second. We also provide 24 ports for terminal traffic coming in through TYMNET, a commercial packet-switched network.

An electronic mail system is available on all the CRAYs, PDP-10s, and VAXs, allowing users to readily communicate with each other from their computer terminals. An interface to the ARPA Internet allows mail to be exchanged with members of that network as well as BITnet.

The network provides access for over 3000 users at over 138 customer sites. Network traffic averages over 300 million characters of data daily. The interest in links to several European sites is growing and some initial planning is already under way to install these connections.

The network's original purpose was to provide access to and from NMFECC's computers. We are now being asked to provide a general purpose data communications network that will include non-NMFECC users. New satellite links are becoming available that use much smaller earth stations than the current links and are much less expensive to install. This may open the door to providing direct satellite links to user sites that were too small to consider in the past.

The MFEnet is today the model for efficient, reliable, and cost effective network access to supercomputer resources. No other organization in the nation currently provides the same level of remote-access resources. The MFEnet will continue to grow in terms of new user sites and increased bandwidth. New capabilities will continue to be introduced, including some of those mentioned above and almost surely in areas that we have not yet considered.