BNL--30609

DE82 006935

MASTER

The Traditional Computing Center as a
Modern Network Node*

S. Heller, A. M. Peskin
Brookhaven National Laboratory
Upton, NY 11973

November, 1981

DISCLAIMER

This book was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

* The submitted manuscript has been authored under contract DE-AC02-76CH00016 with the U. S. Department of Energy.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency Thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

The Traditional Computing Center as a Modern Network Node

Abstract

There is an obvious trend toward decentralization of computing power from the traditional, large computing center. Even so there remains a generous, but changing role for such centers to play. Their capabilities would then be complementary to smaller, individualized facilities, so the user would benefit greatly from a general purpose, local network on which the large center represented a node.

There is no network currently available that exhibits all the attributes of the ideal local for this environment. It can be approached, however, by combining several diverse products as network segments, which are interconnected via processor gateways. This is in fact the strategy being followed at Brookhaven National Laboratory, which has a computing environment typical of a large class of institutions.

The attributes of the ideal network are presented. A brief discussion of the current state-of-the-art in networking is then given. Finally, the particulars of the Brookhaven implementation are offered as a case history.

THE TRADITIONAL COMPUTING CENTER AS A MODERN NETWORK NODE*

by S. Heller and A. Peskin
Brookhaven National Laboratory

I. Motivation

It is generally recognized that there is a trend toward decentralization of computing power from the traditional large computing center. Even so, where such centers exist, there remains a generous, albeit changing role for them to play. There will still be a need to provide general purpose terminal-based interactive computing, and processing of extremely large compute-bound codes. There need be a repository for equipment so specialized or so costly that it must be shared among organizational sub-units. There should be a single center of expertise, so that duplication of staff and effort is minimized. And finally, resource sharing is best accomplished with local computer networking and this implies an organization to develop, manage and maintain the network.

The ideal network for a large organization such as a university or research laboratory might be envisioned to possess the following attributes:

- It would intermix a large number of diverse "host" computers through well defined hardware and software interfaces and protocols.
- Terminals connected to any host would be accessible to all hosts through the network, and furthermore host-to-host and terminal-to-terminal communication would be supported.

- Process-to-process and file transfer communication would be supported.
- The minimum host should be extremely modest, perhaps a microprocessor based terminal.
- Network reliability must be greater than 99% with no single point of failure.
- There should be a choice of bandwidth ranging from 300 cps to 50 mhz, and there should be software compatability among the classes of service.

A network with such attributes exists, for the time being, only in the minds of industry planners and standards committee writers. All currently available commercial products fall short in one or more respects. Those installations with enormous need and/or resources have endeavored to develop their own customized networks with mixed success. For the remainder of organizations, there are two important constraints. The first is the cost of entering this brave new world compared to that of computing at an existing, less cost-effective computing center. The second problem is confusion; developments in networking are proceeding at such a tremendous rate that it is difficult to make a decision of direction that is likely to stand the test of time.

Realistic cost requirements tend to force planners into the design space of commercially

available products, rather than trying to mount a major development effort. And while no one network product can meet all the requirements an installation is likely to have, a collection of network segments, interconnected through gateway computers, provide almost the same desired effects. The central computing facility then becomes a network node for shared facilities, and a repository for the gateway computers as well.

The example of the Central Scientific Computing Facility (CSCF) at Brookhaven National Laboratory provides a typical case history. It contains very large computers (Control Data 6600s and 7600) well past their prime. It exists in an environment of rather autonomous scientific departments, most of which own and operate several mini and mid-size computers, usually dedicated to a specific task. For various reasons, one minicomputer vendor seems to dominate in terms of numbers (in this case DEC) and this is again not atypical of a campus or research facility environment. The computing center staff and users alike look forward to an updating of their antiquated service model through an infusion of new equipment based on the network architecture of Figure 1, which shows the network segments as concentric arcs and the gateways as radial spokes. The elements of this architecture are discussed in section III.

II. General Discussion of Networks and Gateways

One convenient way to categorize networks is according to the bandwidth capabilities. Generally, a medium which can accommodate higher speed data transmission

MASSNET

Figure 1

Network Architecture

also requires more expensive equipment to interface computer systems to the network. The most common media are cable, coax or possibly optical, and the simplest topology is a line or a bus on which hang the communication interface units (CIU's) and from them computer system hosts. We will consider mainly this geometric model in our discussion. Many variations exist for host arrangement and connection, but this topology together with its associated "carrier sense multiple access" (CSMA) network is emerging as the most popular and robust of networks - both from the viewpoint of hardware and software at the physical and link level protocols.

With this in mind, four categories of networks can be distinguished:

- 1) Very high speed local network
- 2) High speed local network
- 3) Medium speed local/global network
- 4) Slower speed global network.

Table 1 summarizes the salient characteristics of these four types. The table omits the cost of media which can be considerable depending on the type chosen and the distance to be covered.

Table 1

Network Type	Typical to max. line speed	Typical to max. distance	Typical CIU Costs	Typical software available
1) Very high speed local	50 Mb/, sec. up to 200 Mb/sec.	1000 feet, up to 10,000 feet	\$35,000 or higher	level 1&2 in ISO standards (physical + link levels), possibly more with homogeneous set of hosts, and operating systems.
2) High speed local	1 to 5 Mb/, sec. but up to 10 Mb/sec.	2500 feet, up to 15,000 feet	\$1,000 to \$5,000	"
3) Medium speed	50 Kb/sec. to 1 Mb/sec	cross country, 3,000 mi.	\$1,000 to \$4,000	"
4) Slower speed	5 to 50 Kb/sec.	cross country & overseas 10,000 mi.	\$1,000 to \$3,000	"

Examples of types 1) to 4) are currently available commercially and many are in operation. Most use packet switching techniques. An illustration of type 1) is Network Systems Corp.'s Hyperchannel [1], usually consisting of a coax cable (but also supporting a fiber optic cable or a micro-wave link), and CIU's called adapters (each specially devised for a particular vendor's computer) having a nominal speed on the cable of from 40 to 50 Mb/sec. It is a CSMA network. Type 2) is well represented by Xerox's Ethernet [2], probably the best known of the CSMA/CD (with Collision Detection) networks. An example of type 3) which departs from the simple bus structure is DECNET [3]. The best example of type 4) is the ARPANET packet switching network. From the beginning ARPANET included specifications for file transfers amongst dissimilar hosts.

Most modern network implementations attempt to follow the layering of protocol levels as defined in the ISO model [4]. As indicated in table 1, most network suppliers offer lower level protocols but very few attempt higher level implementations, especially among dissimilar hosts, that is, a "heterogeneous" environment. Most of the available networks encompassing the higher levels are of the "homogeneous" type that is all the nodes of the network are the same computer and operating system, or at least similar.

Apart from the difficulty a network may have handling heterogeneous hosts, the two most glaring shortcomings that one notices in networking are:

- A) Difficulty in internetting or connecting together already existing networks.
- B) The absence or lack of agreed upon standardization of higher level protocols (level 5, 6, and 7 of the ISO model).

The solution to A) via gateways is discussed below. For B) a vendor may typically develop higher level protocols for the homogeneous case, which will frequently contain quirks peculiar to that vendor's operating system or data structures. Where an effort has been made to solve the heterogeneous case, given there is no standard defined in this area, something new has to be done to link together with another network having its own but different higher level protocols. Thus problems A) and B) are somewhat related. It would seem that an effort towards level protocol standardization of products for levels 5, 6, 7 is needed for applications spanning networks to readily proceed. Such an effort has already been made below the user level (e.g. RS-232, X.25, CSMA standards, etc.).

A gateway is a device or system that converts one or more protocols from one network to another. It can be thought of as a switch or interface device interconnecting networks and allowing communication from one network to another. It can be implemented in varying ways. The translation from one protocol in one network to another in another network can be tightly coupled, that is, as close to a one-to-one mapping as possible.

Such a tight coupling can be very difficult to implement since it may mean altering the protocols in each of the networks, at least on the gateway node. For packet switching networks a level 4 (transport layer) has been defined which standardizes the level at which this tight coupling should be done. TCP (Transmission Control Protocol) [5], the Dept. of Defense standard to be implemented on ARPANET in 1983, is an example of a protocol at this level.

Of course if one wishes to interconnect Network 1 with Network 2 both must have the transport layer implemented and with the same protocol. Even when these conditions are satisfied, the user still may not have all that is needed to accomplish the application -- the higher level protocols must exist on both Network 1 and Network 2 which are compatible with each other, if not identical. Frequently, however, some means are found to get the bits from one host to another and then each particular pair of communicating users agree upon their own private arrangements to effectively accomplish the application layers. While this solves the immediate problem for a particular pair of users, it does nothing for others and is in the long run a very wasteful way of doing business.

An alternative to introducing common protocols at level 4 and above is to forego the generality of complete internetting. Instead, one may translate an entire set of protocols in one network as a functional application

to an entire set of protocols in the other having the equivalent functional application without altering or adding to the protocols of either network. This could be called a "loose coupling" gateway. For example, suppose we want a file transfer from a host on Network 1 to a host on Network 2 where Network 1 has its file transfer protocol = FTP_1 , and Network 2 has its file transfer protocol = FTP_2 ($FTP_1 \neq FTP_2$). First we would transfer the file from the host on Network 1 to the gateway, using FTP_1 , then using FTP_2 transfer the file from the gateway to the host on Network 2. The gateway acts very much like a message switching node in the loose coupling while it is like a packet switching node in the tight coupling.

The loose coupling implementation is much simpler but less general in its capabilities than is afforded by a tight coupling gateway. For example, process to process communication is possible under tight coupling but not generally feasible in a loose coupling. Also the tight coupling, on paper at least, looks like it should be more efficient. For example, a file need not be stored twice (once in the gateway and once in the receiving host). The gateway in a loose coupling would generally need rotating mass storage, while a tight coupling gateway could all be handled in memory (but flow control problems may make that impractical). However, the simplicity of the loose coupling is very desirable. Moreover, in a loose coupling situation the

the set of application functions translated can be tailored or selected according to the needs at hand.

III. Putting It All Together and the Prototype at Brookhaven National Laboratory

Brookhaven National Laboratory (BNL) is faced with the problem of providing access across networks and to a wide variety of users both on site locally and at geographically distant places. There is a high speed (12Mb/sec) local network called BROOKNET which has been in use for 13 years and now must be replaced with a more modern network [6]. The problem of setting up all four types of networks is a formidable one especially with a limited staff. The approach taken at BNL is to put together as many commercially available parts as possible while providing minor development and support for certain functional capabilities that are not available in the marketplace. Heavy reliance is placed upon configuration and network simulation analysis to provide tools for coherent management of the various facilities involved.

For the type 1) network BNL has contracted with the Masstor Corporation to provide a network called MASSNET. It uses the Network Systems Corporation Hyperchannel medium. It will initially have two CDC 6600s (NOS/BE operating system) and a Masstor/Fujitsu MC5 processor (IBM OS/VS2 system) as nodes [7]. The MASSNET control hardware includes a network adapter, disk storage for file staging, and keyboard/display equipment for operator

control. The MASSNET software consists of four components: the access methods (NETIO and NETCOM), the Network Manager (NETMGR), and the Network Administrator (NETADM). Together these provide the following capabilities:

1. Programmable process-to-process and file transfer communications
2. Multi-vendor networking
3. Network resource allocation
4. Connections that can survive failure of any host or application program
5. Global network security control
6. Remote job/process start capability from any host

There are plans to add a variety of other computers to MASSNET including a VAX 11/780 (VMS system) and a PDP-11/44 (UNIX) system. One of the principal applications of MASSNET is to provide access from other hosts to a large mass storage facility; the MC5 computer provides storage management and access control of an M860 tape cartridge system having an initial capacity of 5×10^{11} bits.

DECNET is chosen as the type 3) network since so many of the on-site computers are from DEC. DECNET offers a somewhat larger menu of facilities than MASSNET but at lower speeds. It offers task to task communications, network file transfer, downline systems and task loading, network command terminals, and network resource

sharing capabilities using its own set of protocols. In its latest phase III implementation it includes communication with adjacent and/or non-adjacent nodes. Eventually the vendor plans to integrate X.25 standard interfaces into its product. ARPANET on the 11/44 is our type 4) network and will connect to the 6600s via MASSNET.

The VAX 11/780 is the loosely coupled gateway between DECNET and MASSNET. A File Transfer Protocol (FTP) is the principal application to be provided. Commands will be implemented under the VMS operating system to allow a user (on DECNET connected to the VAX) to send or get files resident on a MASSNET host. Auxiliary functions such as MASSNET file directory information, network and host status (on MASSNET) will also be provided. In ARPANET parlance, for a first approach, MASSNET hosts will look like FTP servers to DECNET hosts, and DECNET hosts will behave as FTP users. Fully symmetric capabilities would be developed if sufficient need for it is expressed by the network users. In a similar fashion, the PDP 11/44 would be both an ARPANET node and a MASSNET node and would serve as a gateway between the two for the FTP.

The type 2) network has not yet been selected. Potentially it may have the largest number of nodes since it is envisioned to connect various personal computers, mini's and micro's, to that network and by means of a gateway, to MASSNET as well. While many companies offer hardware configurations with up to the link level software, the only higher level protocols available are designed to run under

a single operating system on at least a homogeneous class of hosts. The cost of the CIU for this network is a factor to be considered since it must be relatively inexpensive to be attractive to its large potential community of users.

Simulation played a large role in the design of this configuration, particularly in the multiple main-frame MASSNET segment. Simulation studies were used to size the system, perform the cost-performance tradeoff analysis and predict the ultimate performance over various load assumptions. Most of these analyses were done with a GPSS simulation model implementation, with the studies focusing on the user's-view of file access and retrieval. These studies also produced interesting results on the stochastic distribution of files in an "open shop" multi-user, multi-processor environment, which in turn suggested new optimization methods for file staging in that environment [8].

VI. Conclusion

It is hoped that this networking arrangement can be eventually upgraded to or replaced by a single robust computer network. An ETHERNET-like linkage, for example, would appear most promising. Until such a unified system is readily available, however, the "gatewayed segment" approach described above appears to afford the best opportunities for revitalizing a traditional facility.

References

1. Christensen, G. B.,: Hyperchannel Data Trunk Contention, Proceeding of the Minneapolis Conference on Local Networks, University of Minnesota, October, 1977. (See also various Network Systems Corporation publications such as Course 400: Host Software Considerations, or NSC publication No. A01-0000-02 on Systems Description: Hyperchannel Network Adapters.)
2. Metcalfe, R. M. and Boggs, D. R.: Ethernet: Distributed Packet Switching for Local Computer Networks, Comm. ACM, 19:7, July, 1976, pp. 395-404.
3. Introduction to DECNET, Digital Equipment Corporation, (order No. AA-J055A-TK), January, 1980.
4. Desjardins, R.: ISO/ANSI Reference Model of Open Systems Interconnection, Proceedings, Trends & Applications: 1980 Computer Network Protocols, NBS Conference, May 29, 1980. (See also reference Model of Open Systems Interconnections ISO/TC97/SC16/N2227.)
5. Transmission Control Protocol, Defense Advanced Research Projects Agency, Internet Program Protocol Specification, September, 1981.
6. Campbell, G., Fuchel, K., Padwa, S. L., BROOKNET-2: A Guide to the Brookhaven Online Computer Network, BNL report #50404, January, 1974.
7. Masstor Systems Corporation: Concepts and Facilities for MASSNET and SVSS for a CDC NOS/BE Host, September, 1980.
8. Peskin, A. M., and Yuen, K.: A Model for the Computer Aided Design of Local Computer Networks, Summer Computer Simulation Conference, July 15-17, 1981, Washington, DC.