72c

A snapshot of seven typical networks, showing the variety of work under way and the problems that all have in common

Networks: An Introduction

by David J. Farber

A computer network is an interconnected set of dependent or independent computer systems which communicate with each other in order to share certain resources such as programs or data—and/or for load sharing and reliability reasons. This survey is based on information gathered from the IEEE Computer Society Workshop on Computer Networks (at Lake Arrowhead, September, 1971), the Mitre Corp. report on Computer Networks, the technical literature, and the experience of the author.

This article is intended to give an overview of this expanding field by the examination of seven typical networks. These seven were chosen to show various aspects of the subject, not because they are necessarily the best, or the most advanced, efforts. The key points of comparison among these networks are summarized in Table 1.

After we have described the seven networks we will examine some of the common problems that occur in all networks, such as data conversion, rigidity imposed by the protocols, etc.

The ARPA network

The ARPA (Advanced Research Projects Agency) network (Fig. 1) is a nationwide system designed both to explore network technology and to interconnect and service ARPA-sponsored research centers. The key aim of the system is to allow the accessing of programs, services, and data from any place on the network.

The ARPA network is a distributed network: sites (nodes) on the network are connected to each other either directly or indirectly through intermediate sites. This is to be distinguished from a centralized network where all sites are connected together via one central site. The computers and associated software systems that make up the ARPA network are heterogeneous, not all from the same source.

The network can be broken into two parts. One part consists of the computers which will provide the computational services of the network—the hosts; the other part deals with the function of servicing the communication needs of the network.

The communication section of the ARPA network consists of modified Honeywell DDP-516 computers connected via 9- and 50-kilobit leased telephone lines. The DDP-516 machines are called IMPS (Interface Message Processors). The communication system operates in a message-oriented store-and-forward fashion: a message is stored at intermediate points as it makes its way toward the destination. Each time the message is handed forward correctly, the handing node is freed from any further responsibility for the message. Since it is often necessary to send messages of substantial size, the network breaks long messages into smaller submessages called packets. These packets of about 1,000 bits are independently forwarded through the communications network. A duty assumed by the network, through the IMPS, is to insure that the packets are reassembled into the original message for transmission to the destination host. In addition the IMPS govern routing of messages through the network in order to minimize the transit time of the message and to increase the utilization of the transmission facilities.

Each host computer is equipped with a program called the NCP (Network Control Program). The NCP arranges for connections to be established and terminated between programs on one host and programs on another host and performs other monitoring functions for user programs.

There are currently 23 host machines on the existing ARPA network. These range from a PDP-11 through the ILLIAC IV. The network is managed by the ARPA agency and is technically directed by a steering committee of the Network Working Group, an organization of host representatives who are charged with the technical evolution of the system. In addition, Bolt, Beranek, and Newman (BBN) is charged with maintaining the communications section of the network.

The ARPA network is today the main candidate for becoming a nationwide data network. There is considerable pressure for universities, government agencies, and other organizations to be allowed to join the network. Some of these groups would like to form closed subnets communicating only among themselves, while others desire to join the larger group in hopes of utilizing the services of the existing sites—in particular, the ILLIAC. It is clear that there is a movement toward removing the network from the status of a DOD-sponsored research activity and evolving it into a "commercially" run computer network.

The CYBERNET network

This network is included here as representative of a currently operational commercial network. While its technology is not as sophisticated as that of the ARPA network, it does face up to the practicalities of the real world. It was formed, basically, to connect Control Data Corp.'s existing data centers. The expectations were that by interconnecting the centers they could gain: better reliability by making available an alternate machine in case of local failure; greater throughput by load balancing across machines situated in different time zones; greater manpower utilization by using corporation manpower facilities more effectively, allowing access to others' programs and data bases; and, finally, the convenience of enabling a customer to choose a configuration that is best suited to his problem rather than the one which is best located geographically.

CYBERNET is a distributed network consisting of CDC machines such as 6600s and 3300s linked by wide-band and voice-band lines.

CDC speaks of the 6600s and other similar CDC machines as the primary computing capability of the network and calls them "centroids." It considers the 3300s as the front-end machines and concentrators for the centroids and calls these the nodes of the system. Terminals and satellite computers are supported for interactive and remote job submission operations.

The communications system of CYBERNET utilizes a broad spectrum of switched, leased and satellite communications facilities. It counts heavily on essentially hand-established connections for terminal-to-computer and computer-to-computer links. Thus the network by itself cannot reconfigure itself. Alternative paths do exist in some cases between nodes and centroids but in general a link failure will necessitate human intervention.

CYBERNET is operating as a commercial entity and is offering general computation services to its users.

The DCS

The Distributed Computer System (DCS) shown in Fig. 2, is an experimental computer network being developed and constructed at the Univ. of California at Irvine. Its stated aims are: low cost, reliability, easy addition of new services, a modest startup cost, and low incremental expansion costs. It is primarily intended to service mini and midi scale computers. Its communication architecture is based on a digital communication ring topology utilizing essentially the Bell System T1 technology and fixed-length messages. The computers are interfaced to this circular transmission medium using a fairly sophisticated piece of hardware called a ring interface, not a computer. The main novel feature of the communications protocol is that messages are ad-

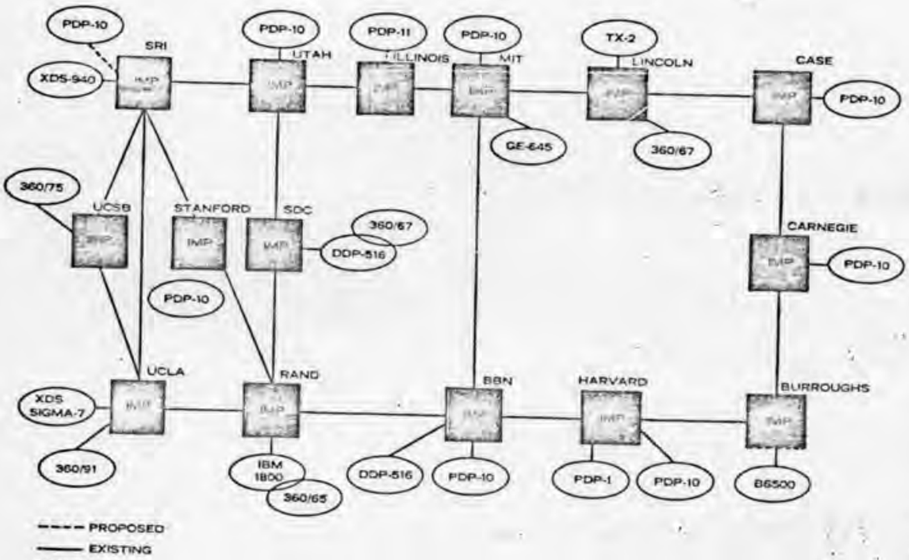

NASA Ames' ILLIAC IV has now been added to the ARPA network

SOURCE: Bolt Beranek and Newman

Fig. 1. ARPA Network Topology, February 1971.

	ARPA	CYBERNET	DCS	MERIT	OCTOPUS	TSS	TUCC
Organization	Distributed	Distributed	Distributed	Distributed	Mixed	Distributed	Central
Composition	Heterogeneous	Heterogeneous	Heterogeneous	Heterogeneous	Heterogeneous	Homogeneous	Homogeneous
Number of nodes	23	36	9	3	10	9	4
Geography of nodes	USA	USA	UC, Irvine	Michigan	LBL	USA	North Carolina
Machine size	Mixed	Large	Mini	Large	Large	360/67	360
Communication interface machines	Honeywell DDP 516	CDC 3300 PPU	Ring Interface	PDP 11	CDC PPU	IBM 2701	IBM 2701
Communication protocol	Message switched	Message switched	Mixed	Message switched	Point to point	Point to point	Point to point
Transmission medium	Leased lines	Leased lines	Twisted pair—coaxial	Telpak	Coaxial	DDD	Telpak
Data rates bps	50,000	100-40,800	2-5 million	2,000	1.5-12 million	2,000, 40,800	100-2,400, 40,800
Transmission mode	Analog	Analog	Digital	Analog	Digital	Analog	Analog
Message format	Variable length	Fixed length	Variable length	Variable length	Variable length	Variable length	Variable length
Message size	8,095 bits	1,024 characters	900 bits	240 characters	1,208 or 3,780,000 bits	8,192 bits	1,000 bytes

Table 1

dressed to the receiver by means of the name of the receiver, not by a location at which that receiver lives. Thus the receiver can be allowed to migrate to other computers without having to inform the transmitter of that fact.

There are three types of ring interface—one to support a computer, which could be a front end machine; one to allow the direct attachment of a terminal to the ring; and one designed to allow the construction of a network of rings. This "ring of rings" operates essentially the same as the basic ring.

The DCS effort plans a distributed data base capability and a set of services for the users. It is not intended to provide a commercially viable system but rather to explore the issues involved in distributed architecture.

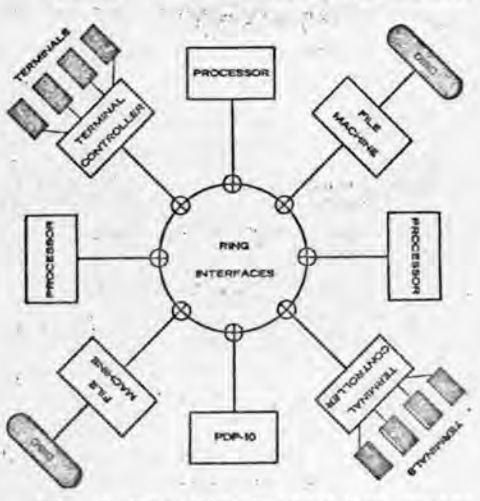

Fig. 2. The Distributed Computer System.

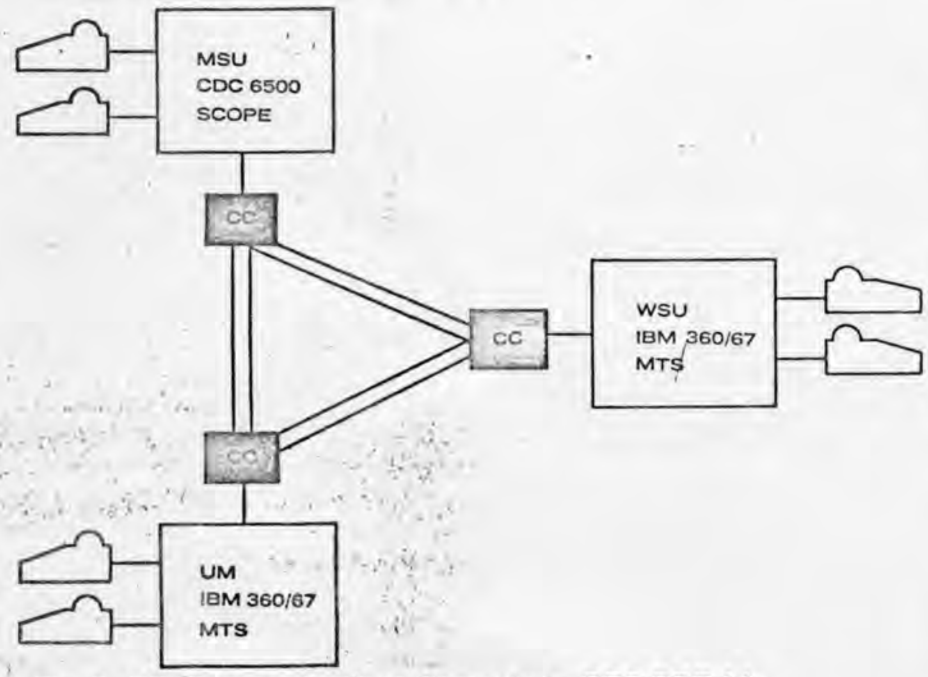

Fig. 3. CC: Communications Computer—DEC PDP-11.

However, it is intended that an operational system be built both to test the ideas and to explore the utilization of the large number of minicomputers that have appeared on the university campuses and at health care centers.

The MERIT network

The MERIT Network, the Michigan Educational Research Information Triad, Inc. (Fig. 3), is a joint cooperative effort between Michigan State Univ., Wayne State Univ., and the Univ. of Michigan. Its stated aim is to create an educational computing network to allow the computing at the member schools to be shared.

The MERIT Network is a distributed network consisting of three nodes. Its computers are heterogeneous. Each host computer is connected to the communications network by means of a modified DEC PDP-11/20. The communication lines interconnecting each site are a group of 2000 bps voice-grade lines.

The communications computer, the PDP-11/20, is capable of providing, through the facilities of a host interface hardware module, a variable-length message transfer from PDP-11/20 main storage to the host core and the communications system. In addition, it allows the host computer to treat its communication computer as several peripheral devices. This simplifies the host software system considerably since it allows the dedication of a pseudo-peripheral device to each user.

The communications computer (CC) is capable of acting as a store-and-forward system. Thus, if a path is destroyed, an alternative path exists via another of the CCs. The MERIT Network employs voice grade dial-up lines allowing the economic savings associated with the Telpak lines of the existing tri-university voice network.

The MERIT personnel feel strongly that networking will have a synergistic effect on the total computing environment. They are seriously facing the management difficulties inherent in the interconnection of educational computers.

The Octopus system

The Octopus system is a heterogeneous network developed at the Lawrence Berkeley Laboratory (formerly the Lawrence Radiation Laboratory) of the Univ. of California. It connects a complement of devices including two CDC 6600s, two CDC 7600s, and eventually will include a CDC STAR. All of these machines, called workers, are operated as time-shared facilities. The laboratory plans to provide for a centralized hierarchical data base and for a wide variety of input-output devices which can view the network as a single resource.

The communications system utilizes a store-and-forward protocol. The workers in the Octopus are interconnected via 12-megabit-capacity hard-wired cables. The system can be considered as two superimposed subnetworks. The first is a File Transport subnet consisting of the workers, a transport control computer, a dual DEC SYSTEM 10 and the file storage. The second network is a Teletype subnet consisting of PDP 8s (each supporting 128 terminals), the workers, and the transport control computer. Notice that the Teletype subnet is a distributed network while the File Transport subnet is a centralized subnet. The dual DEC SYSTEM 10 insures reliability in this centralized subnet. In addition, while the subnets are logically independent, there are cross couples providing redundant paths in the event of failure.

The Octopus network is one of the more elaborate networks currently in operation. It is also one of the few networks which has been designed to handle security materials. One should note, however, that Octopus lies entirely on Lawrence Berkeley Laboratory premises.

The TSS network

The TSS Network (a cooperative venture between IBM and some of its 360/67 customers) was developed as a network of homogeneous computers operating in a distributed fashion. Each of the hosts operating on the TSS Network consists of a 360/67 using the IBM TSS/360 operating system. Some of the nodes have local networks consisting of 360s appearing as devices, not hosts, to the network.

The communications facilities between the 360/67s utilize voice-grade switched lines. These lines are interfaced to the 360/67s by means of IBM 2701s or 2703s. Thus, while this network utilizes off-the-shelf hardware, this hardware, insofar as the communication protocols are concerned, is not programmable. Thus all the programs such as store and forward, error, etc., are resident on the host machines. Indeed, the communications software operates as a user program via an access method. There are plans to attach an IBM 370/145 to act as a communications computer and data base manager in the future. There are also plans and capabilities for utilizing 50,000 bps lines when the demand exists.

The TSS Network is experimental. Since all machines on the network are similar, program and data interchange is available. Both dynamic file access and remote batch are available over this network.

A notable feature of this network is that the host machines are IBM 360s utilizing standard hardware. One could, in theory at least, buy a copy of this network from any IBM salesman.

The TUCC network

The Triangle Universities Computation Center (TUCC) Network is a joint undertaking of the Duke, North Carolina State, and North Carolina Universities. It is an example of a relatively simple, straightforward undertaking in networking. It has been operational since 1966. It is a centralized network of homogeneous machines. At each of the three nodes of the net there are IBM 360/40s or 50s. These 360s do local batch jobs in addition to handling the telecommunications necessitated by the net.

The nodes of the net are connected to the central facility by means of a leased 40,800 bps half duplex line. This line is interfaced to the 360s by means of IBM 2701 Data Adapters.

In addition to the three TUCC nodes, local schools and colleges are serviced by the central computer via a variety of medium and low speed input/output devices.

This network is simple. It uses off-the-shelf hardware with minimal extensions to the basic IBM 360 Operating System software.

Discussion and comparisons

We have briefly scanned the architecture of a number of major attempts at computer networks. There are many others that we have not discussed. Some of these are highly specialized, such as the California Law Enforcement Telecommunications System; some are of very limited applicability. In addition to those being developed in the U.S., there are a number of efforts under way in Canada, Great Britain, France, and other countries. In particular, the network developed at the National Physics Laboratory (NPL), Great Britain, is an early example of experimentation in this field. In general, however, most of the foreign networks are planned to be basically patterned after the ARPA system. There is also an interest in loop networks such as the DCS.

All the networks we have discussed have some common objectives and some common problems. There are a number of services that can be considered as standard offerings in a large number of networks. We will define and discuss these at this point.

Load sharing, the ability to take a given workload and to distribute it among the computers of a network in order to make equal use of the resources of the network, is one of these services. It is offered by the CYBERNET, DCS and TSS networks but is not a basic feature of the ARPA, MERIT, Octopus or TUCC systems. In the case of all these systems, load sharing could be added as a user-supplied feature subject to certain restrictions. All the networks surveyed provide a form of program sharing. That is, they all allow data to be sent to a node at which a desired program is resident. A common problem encountered in heterogeneous networks first appears here. While we don't expect programs written for one brand of computer to run without human change on another brand, we do, perhaps naively, expect data generated by one brand to be transmitted and understood by a program running on another brand of computer. In general this is a difficult problem requiring careful design of message formats and protocols. In addition, in the case of data sharing of files, data conversion services must be available for converting between notations and conventions of the different computers.

There is also a form of sharing sometimes referred to as data sharing. If I have a large data set that I need processed and that data set is on another node of the network, then it might be more economical to send the program to the data rather than vice versa. All systems give this capability in one form or another.

There exists a feature called dynamic file access which is in essence the ability of a program to access a remote data set as if it were local. This allows the program to operate on a distributed data base with no special planning. It is central to the DCS design and also available on the MERIT and TSS networks.

The central area which is common to all networks is the communications and operational protocols, that is, the rules and regulations which define how one is to handle an event and what to do when an error occurs. Most of the design time spent in the construction of computer networks is involved in the formulation and debugging of these protocols. These having been defined, a number of different computers can each be programmed to behave the same, at least with respect to their appearance to each other. Most networks demand that all joining nodes conform to one rigid protocol. In effect, they all decide to talk the same language to each other. In the case of the DCS an attempt is made to support different protocols with respect to establishing contact between programs, machines, etc. The rigidity and complexity of the protocols affects the cost of joining a network.

I will make no attempt to judge which of these networks is better, or more indicative of future directions. As a designer of one of them, I clearly have biases.

As mentioned earlier, the ARPA net is the biggest and best developed network. It is, however, expensive. The IBM 360 networks are less elegant but use off-the-shelf components. Networks such as DCS, Octopus, and TSS are basically research efforts and may show future directions. In all cases, issues such as network economics and management problems transcend in difficulty the technical problems. One is left with a strong feeling that technically networks are here to stay but how we use them in our existing corporation and university structures, and how we pay for their use, are unsolved problems. □

Prof. Farber is on the faculty of the Univ. of California at Irvine and is a founder of Caine, Farber & Gordon, Inc., a Pasadena, Calif., computer systems company. He is a principal investigator of the DCS project, a network research program sponsored by the National Science Foundation. He is co-author of the SNOBOL language series and his previous computer experience includes positions at Xerox Data Systems, RAND, and Bell Telephone Laboratories.