


FEBRUARY 1974
IEEE

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

spectrum

More Power by Networking



More power by networking

Whether computers communicate as equals or in a superior-subordinate mode, the outcome is usually positive

Computer networks have evolved for reasons of economics and efficiency to permit information resident in one computer to be shared with others, to enable large numbers of users to have remote access to a computing center, and to control the flow of information from one point to another. Not surprisingly, computer networks have evolved in much the same way as other types of networks in our society such as those in communications, power, and transportation. The evolutionary pattern has been to begin with pairs, which are expanded to chains, which in turn lead to series-parallel arrangements such as trees and loops. The networks are interconnected in a variety of ways—they may be based on superior-subordinate relationships, such as commonly found in task-oriented organizations (e.g., industrial or government networks), or else they may be predicated on a more democratic relationship because of more general goals (e.g., networks for research and education).

In all networks, communication is initiated by one computer. However, the ensuing transactions with respect to the initiator can be as a superior demanding a task be done by a subordinate or simply as a request for help.

C. Gordon Bell Digital Equipment Corp.

Some networks consist of a number of small computers used as an alternative to a single large computer. This is done to ensure that the computer structure is operational at all times. It also can simplify system development if functions are partitioned carefully. Such a network permits greater economy by selecting machines that are the most cost-effective for each part of the problem being solved, and then distributing the processes according to costs and capabilities. Such a network may be under control of either a large or small computer.

A key element in all networks is intercomputer and computer-to-terminal communications. A well-established trend is the use of minicomputers connected to larger machines to handle the communications line control functions, commonly called front-end processors. (The minis replace hardwired, inflexible communications control units.) In addition, the minicomputer may perform a variety of simple data processing tasks on a more cost-effective basis than a large machine. The tasks include editing, calculations, and interpretation of languages such as BASIC. Although such a simple pairwise intercommunication could hardly be called a network, there is a trend to use interconnected computers whose functions are specialized. Another similar trend is to use remote job entry computers as intelligent terminals, connected in the form of a one-level tree (or star).

Some of the tasks performed by minicomputers in controlling communications include transforming and controlling the communications lines and modems, controlling a wide variety of terminals (e.g., performing character code conversion, echoing, error checking, and speed selection), and multiplexing characters from multiple incoming lines into a single outgoing line to a host computer.

Network configurations

There are three basic types of network structures (Fig. 1): directly connected networks, store-and-forward networks, and hybrid networks. The direct-connected structure provides for a direct communications link between the transmitter and receiver. Most of today's networks operate in this manner because it offers the simplicity of a fixed communications pattern. The structure may be a one-level hierarchy (called a one-level tree or star), most often in the form of a central site connected to a number of satellites. Strictly speaking, a tree is not a network, since there are no closed loops for alternative routes among the links. The characteristics of directly connected networks are listed in Table I.

Store-and-forward networks contain separate switching computers, which provide buffering to as-

I. Overview of Network Structures

	Direct connection	Independent store-and-forward network
Cost	Cost small for simple networks	Additional machines needed
Limitations	Not applicable to large nets or long distances. Applicable to limited traffic matrix (e.g., STAR)	
Ease of design	Simple	Functionally independent of host and its operating system, but the net is costly and complex
Reliability	Based on redundancy	Can be made arbitrarily large and therefore reliable
Performance	Fixed performance which can be modified using dial-up to assist in overloads	A network can increase performance with more links
Applicability	For small minis	Large interconnection nets where communication costs are small fraction of overall cost

sure that the networks operate smoothly regardless of the speeds or data formats of the machines involved. This type of network also permits any machine to communicate with any other in the network. The ARPA network (covered later in this article) is a good example of this type of structure and, in the author's opinion, large-scale networks will inevitably follow this approach.

A hybrid configuration is a direct-connected network, but with the intermediate computer nodes providing store-and-forward capabilities. This approach has seen little use.


Considerations in network design

An important element to successful network design is a thorough knowledge of computer components, together with a proper definition of the information processing problem over the life of the network. The most important component parameters are link capacities, error rates, costs, and reliability. Processing capacity also is an important factor and it must be measured, not on the basic hardware, but in terms of the operating system and languages specifying the information processing problem.

The information processing problem should be defined both in terms of data processing needs and the location of hardware. The processing needs should be based on the number of statements, file sizes, and the delay time that can be tolerated for the processing services.

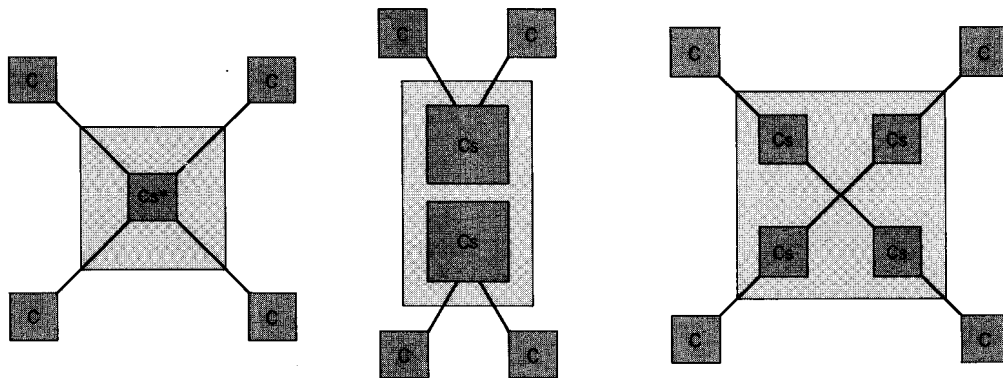
The physical location requirements are based on the economics of whether the processing equipment

[1] Network structures. Direct-connected networks (A) may be configured in a number of different ways, but most often they take the form of a star (a central computer connected to satellites). Store-and-forward networks (B) contain independent switching computers to optimize the data flow and work assignments. In a hybrid configuration, (C), the store-and-forward function is performed by a single, dedicated channel in a directly connected network.

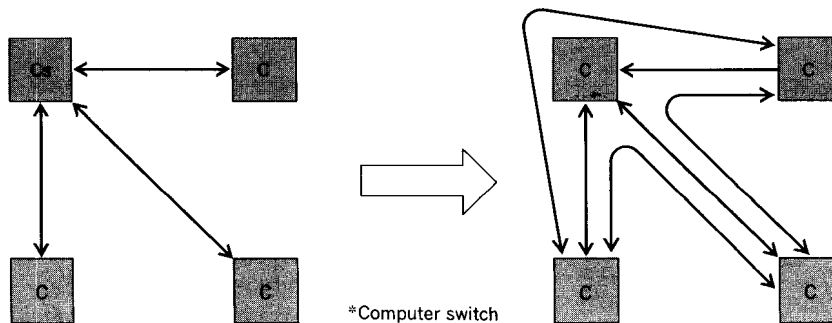


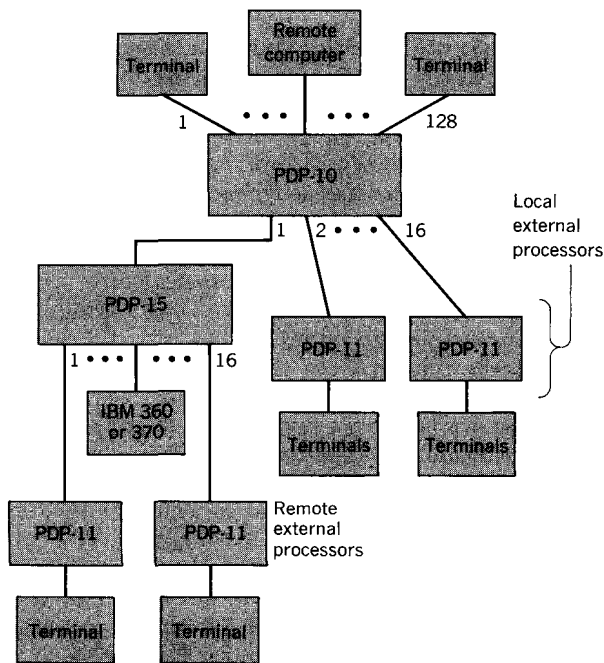
A

B



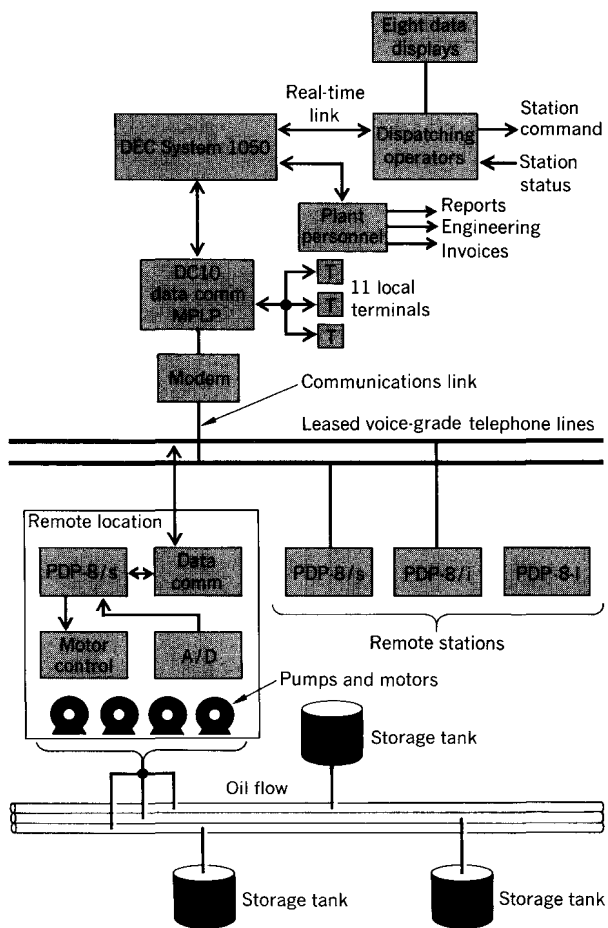
C





[2] Laboratory interconnection programming system is designed for control of diverse computers. At the heart of the system is a PDP-10, which functions as the central computer. Its purpose is to file, to provide central support of satellite computers, and to switch information.

[3] Canadian International pipeline network communicates with minicomputers over voice-grade telephone lines. The minicomputers are located at remote, unattended pumping and oil storage stations along the 1900-mile pipeline.



should be local or remote. For example, suppose it is desired to display the output of a flow-rate meter and there are switches with which an operator can enter the flow-rate limits. Processing of this problem can be handled entirely locally or at a remote computer—or possibly even distributed among several machines (if the information is needed elsewhere).

In designing a network, the criteria that usually apply to components—i.e. cost, performance, and reliability—also apply to networks. Several additional guidelines are:

- Minimizing the number of computers and computer types tends to decrease programming costs.
- Designing a network, where possible, with identical components allows a variety of functions for backup.
- Duplicating functionally specialized computers offers reliability.
- Unless the node is replicated or made into a duplexed tree, tree-structured (star) networks are not satisfactory for most networks that require high reliability.
- Communication links and the protocols that establish error-free transmission are probably the least understood by computer engineers.
- Never build a structure before competing, alternative structures have been considered.

Typical networks in use today

Several networks in use today illustrate the various network approaches. Those covered below are the DEC Laboratory Interconnection Programming System (LIPS), a simple tree structure; the Lawrence Livermore Laboratory system, a directly connected network with some switching (hybrid) capability; and the ARPA network, designed for computer network research.

The DEC LIPS System is a general purpose, PDP-10 computer for controlling a network of diverse computers. Strictly speaking, it is a one-level tree, which permits a number of independent satellite computers to be connected to a large central computer, as shown in Fig. 2. Its overall function is to file, to provide central support of satellite computers, and to switch information among the computers. The network has been implemented at Oak Ridge National Laboratory (IBM 360/91, 360/75, 2780, three Systems Engineering Labs 817/840s, seven PDP-4/7/15s), the Abbott Laboratory, Knolls Atomic Power Laboratory (two Control Data Corp. 6600s, one 7600), and Rolls Royce Engine Testing Laboratory (two English Electric KDF-9s, one IBM 360, eight PDP-8/11s).

The LIPS network has appeared in various laboratories that use a number of large, scientific computers. The functions being performed by the relatively large, general-purpose time-sharing system are:

- Scheduling and loading of the large satellite computers. In several cases (e.g., Oak Ridge and Knolls), the PDP-10 appears to other machines as a remote job entry terminal. In the case of Rolls Royce, where the PDP-10 schedules two KDF-9s, utilization of the latter was increased from 25 percent to 75 percent and the number of supporting operators was reduced from 22 to 10.
- Switching. Since there are several large machines available, a given job can be routed to the machine that can best handle load.

- Editing. Since the files reside in the PDP-10, they can be edited locally, avoiding costly movement of the files and trivial processing of them.
- Interactive processing of small jobs. These tasks are carried out in a manner akin to editing, using either interactive interpreters (e.g., APL) or load-and-go compilers (e.g., BASIC).
- Pooled specialized central facilities. Printing, punching, display, film reading, and specialized plotting are provided by the central facility.

The LIPS approach has also been used in minicomputer networks. It has seen service in laboratories where minicomputer nodes are employed to meet real time and preprocessing demands. These structures are similar to the industrial control systems in which minicomputers control individual, real time processes. The need for such a network is predicated on the poor adaptability of a single, large computer to meet real time response demands; poor reliability of a single site; the high cost of simple, short wordlength calculations required in real time preprocessing; and the physically isolated inputs and outputs associated with a process.

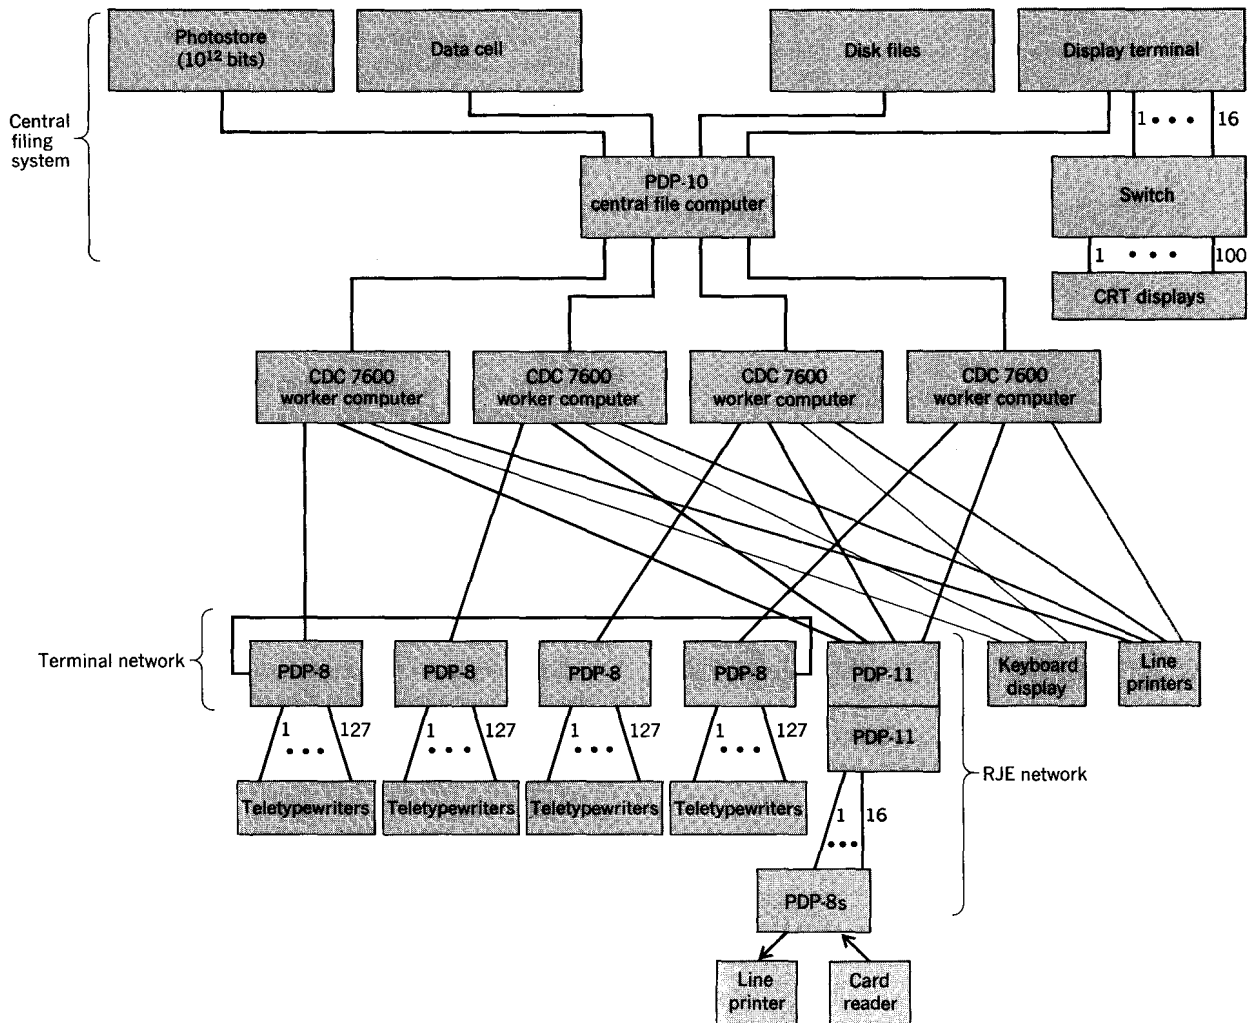
The functions usually performed at a LIPS central

computer site include:

- Management of a central data base. Usually a significant economy results from removing file devices from the local computers and placing them in a single site. Overall, there is less hardware and software for operating systems in the individual machines.
- Centralized editing and language translation (program preparation). Large machines, which contain large primary memories, usually provide the best facilities for minicomputer program preparation.
- Pooling of specialized facilities (e.g., printers, punches, tapes, displays, plotters).
- Significant computations that cannot be handled at the minicomputer nodes. This permits the use of more extensive languages such as COBOL and PL/1, which usually are not supported by minicomputers.

The Canadian Interprovincial Pipeline (IPL) network (Fig. 3) is a one-level tree in which 44 PDP-8 minicomputers are connected to a central PDP-10 computer via several pooled, synchronous, communication links. The minicomputers are located at remote, unattended pumping and oil storage stations along the 1900 miles of the pipeline. Each minicomputer controls several large pumps, and monitors

[4] Lawrence Livermore Laboratory's Octopus network gives several hundred subscribers access to four worker computers and a central file computer. The terminal network accommodates 500 teletypewriters, 16 remote job entry consoles, keyboard-displays, and printers.



pressures, flow rates, gravity, viscosity, power, and other station parameters. The central site also contains a minicomputer to back up the central central computer, but with fewer monitoring functions. The central computer monitors and records the behavior of the pipeline and controls the individual sections. One advantage of such centralized control is that the amount of electrical energy supplied to move the oil can be optimized. Thus, the network has both local and global control.

An alternative approach would be to multiplex all inputs and outputs, and then transmit them to a cen-

tral site for control. This would require that communications lines always be operational, whereas the existing structure carries out local control at each site without frequent intervention from the central control computers.

The Lawrence Livermore Laboratory's Octopus network, started in 1964, appears to be the earliest, and most general, of the general-purpose computer networks, because it provides a wide range of functions with different computers. Several hundred simultaneous on-line users are given access to four worker computers.

At the heart of the network (Fig. 4) are four large worker computers (three CDC 7600s and one CDC 6600) and a central file computer (a DEC dual processor PDP-10) with access to a data cell, disk files, and a 10^{12} -bit IBM photostore (controlled by an IBM 1800 computer).

Hardwired 10-Mb/s links connect the large computers to the file computer. The terminal computers and the large file are connected to the switching computer via high data rate links. The terminal network connects over 500 teletypewriters to the four worker computers via four PDP-8s, and the remote job entry (RJE) network connects up to 16 remote job entry consoles. In addition, a set of keyboard displays and line printers are connected to all the large computers in the system.

The main purpose of the network is to permit any user to have access to any of the large computers so that the network will operate even with a number of machines down.

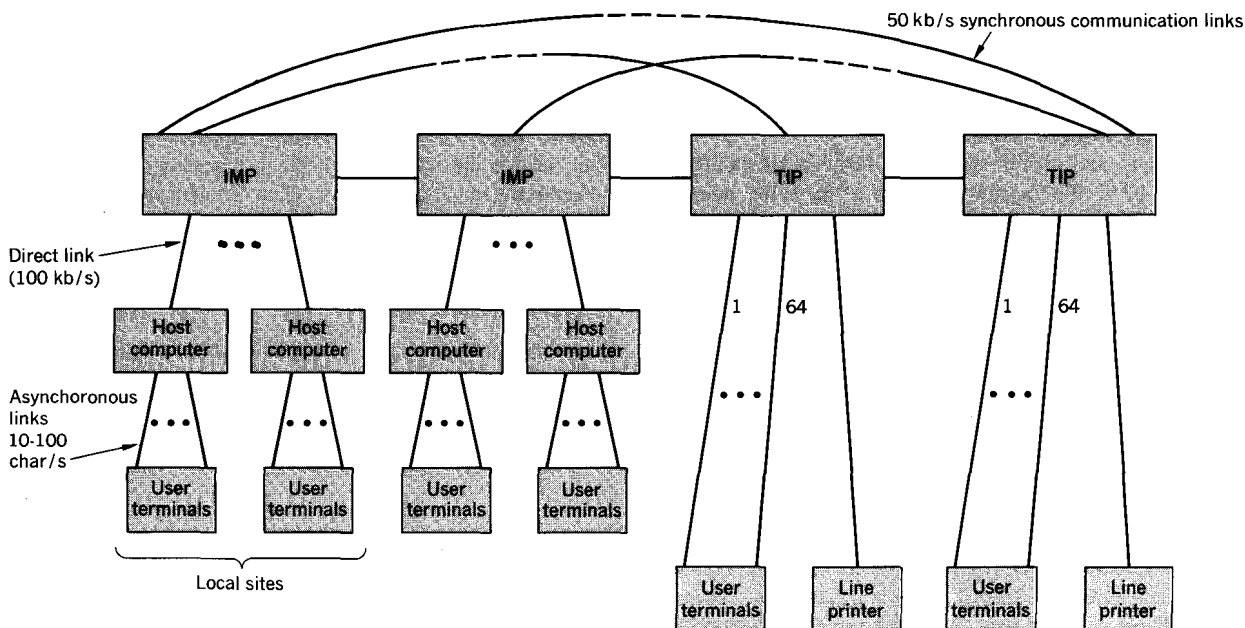
The Advanced Research Projects Agency (ARPA) network was conceived in 1968 and placed in operation at four sites in 1969. As of January, 1973, 34 centers had been connected.

The computer utility

To date, computer networks have been built by special classes of users. However, the next logical step is an information utility for whole "communities," such as businesses, homes, and government departments, which would provide services such as credit card transactions, printed message delivery, news distribution, and library information retrieval. Such a utility network might take advantage of the economy of scale by assigning information processing tasks to machines best able to handle a given processing task or else they could distribute the work load among many machines for optimum efficiency. Computer networks already take advantage of geographic time zones to assign jobs to machines during off-hours.

In the nearer future, the existence of specialized data bases will begin to make program and file sharing desirable because of the excessive time and costs involved in shipping data in physical form and in the costs of maintaining and updating redundant files. For example, we might expect to see specialized networks for libraries as a first step in retrieval networks. Indeed, computer networks may eventually provide the only economical solution to the problem of maintaining and disseminating the information in libraries.

[5] The ARPA network structure consists of large-scale central computers (hosts) and interface computers (IMPs and TIPs). The TIPs were added after the network was initially designed.



For more on networks

Bell, C. G., Habermann, A. N., McCredie, J., Rutledge, R., and Wulf, W., "Computer networks," *Computer*, Oct./Sept. 1971, pp. 14-23.

Cady, G. M., and Luther, G., "Trade-off studies in computer networks," *IEEE Compton*, 1973, pp. 147-150.

Frank, H., and Frisch, I. I., *Communication, Transmission, and Transportation Networks*. Reading, Mass.: Addison-Wesley, 1971.

Ornstein, S. M., Heart, F. E., Crowther, W. R., Rising, H. K., Russell, S. B., and Michael, A., "The terminal IMP for the ARPA computer network," *Proc. AFIPS SJCC*, 1972, vol. 40, pp. 271-280.

Roberts, L. G., and Wessler, B. D., "Computer network development to achieve resource sharing," *Proc. AFIPS SJCC*, 1970, vol. 36, pp. 543-567.

Roberts, L. G., "Network rationale: a 5-year re-evaluation," *IEEE Compton*, 1973, pp. 3-5 and pp. 39-44.

Speers, G. S., "Monitoring control by distributed computing," *Datamation*, July 1973, pp. 47-49.

Wecker, S., "A design for a multiple processor operating environment," *IEEE Compton*, 1973, pp. 143-146.

The purpose of this system is to investigate broadly the use of a computer network; to explore an alternative method of message switching; to provide a wide range of computing facilities to a community of computer and physical scientists, for program and file sharing; and to permit the users to communicate with each other in a variety of ways.

The ARPA network (Fig. 5) is a packet switching, store-and-forward network, whose nodes consist of interface processor computers (IMPs). The computers are linked together by up to four 50-kb/s synchronous communication lines. Messages of up to 8000 bits are transmitted among central (host) computers on a packet basis (up to 1000 bits/packet), with a packet delay of approximately 0.1 second when several IMPs are involved.

A TIP (terminal interface processor) is an IMP without a host computer, but with the capability for up to 64 terminal users to access other hosts. TIPs were added after the initial network was designed, and are used at about one third of the sites.

The hosts are either standard, large-scale computers (e.g., PDP-10s or 360/370s using the same operating systems) or else specialized computers. The latter type includes the ILLIAC-IV at NASA Ames, a large 360, a CDC 7600, a Burroughs B6700, the MIT Multics System, and one or two large 10¹²-bit files. In addition, hosts monitor and report the network activity, provide information about documentation and programs of interest, and hold user "mailboxes" through which individuals communicate with each other.

The network has been operated in the following ways:

- Remote use of computers either from a termination on a host or via a TIP, or else on a batch or interactive basis. Since computers can be used over long distances, the network functions as a broker for computing facilities. In this way, a particular site can choose

not to own and manage a facility, but instead to buy file storage and computation at a variety of other sites. Since ARPA controls the operation of the individual sites, when a site becomes overloaded, the agency can expand the configuration or move some of the users to another site.

- File movement and printing. A user may retain text or programs at one or more sites and then transfer files to particular sites for program execution. Some of the sites have elaborate printing devices for both arbitrary character sets and graphics.

- Personal messages. Users can communicate with each other in various ways. A message for another user may be placed in a "mailbox" located at a host. For example, a message was placed in a mailbox of another user at 11 p.m., requesting latest information on the ARPA network use. By 10 a.m. the next morning, the information was available. This form of communication is faster than a letter or telegram; it is low-cost; and it doesn't require the simultaneous availability of the communicants. Still another advantage of the mailbox is that it permits batch processing of messages by each user.

On the other hand, users can be directly linked via terminals, for both direct two-way and conference real time communication. In the conference mode, a message is broadcast to all attendees.

Finally, terminals may be cross-connected to a single program. In this mode, multiple terminals can have their inputs and outputs connected in parallel. Thus, one user can demonstrate a program by taking several users through a "scenario."

- Machine-to-machine subroutines. In the simplest case, a user has access to a program in a given machine. If the user program requires a program on another machine, the first program calls it in a manner akin to subroutines stored in a single computer.

- Access to a large common data base. For example, a 10¹¹-bit meteorological data base is being prepared to be shared among six research sites.

The ARPA Network is important because it provides multiple links per node, so that nodes and links can fail, but the network still functions. Although it could have been built earlier, it requires high-reliability, high-performance minicomputers; time-shared system modes; and 50-kHz synchronous communication links that became available in the past few years. The network design was carried out in an exemplary fashion with extensive analysis, internal documentation, and technical papers. As it became operational, measurement experiments were conducted to verify the analytical tools it developed. It is clearly the archetype for future large-scale networks.

C. Gordon Bell (F) is vice president of engineering for Digital Equipment Corp., Maynard, Mass., on leave as professor of electrical engineering and computer science at Carnegie-Mellon University, Pittsburgh, Pa. He received the B.S. and M.S. degrees in electrical engineering from the Massachusetts Institute of Technology in 1956 and 1957, respectively. Mr. Bell is the holder of several patents in digital systems, the designer of several computers, and he has published books and papers on computer architecture.

digital

DIGITAL EQUIPMENT CORPORATION, Maynard, Massachusetts, Telephone: (617) 897-5111 • ARIZONA, Phoenix, CALIFORNIA, Sunnyvale, Santa Ana, Los Angeles, Oakland, San Diego and San Francisco (Mountain View) • COLORADO, Englewood • CONNECTICUT, Meriden, Fairfield • DISTRICT OF COLUMBIA, Washington (Lanham, Md.) • FLORIDA, Orlando • GEORGIA, Atlanta • ILLINOIS, Northbrook • INDIANA, Indianapolis • LOUISIANA, Metairie • MASSACHUSETTS, Marlborough and Waltham • MICHIGAN, Ann Arbor and Detroit (Southfield) • MINNESOTA, Minneapolis • MISSOURI, Kansas City and Maryland Heights • NEW JERSEY, Fairfield, Metuchen and Princeton • NEW MEXICO, Albuquerque • NEW YORK, Huntington Station, Manhattan, New York, Syracuse and Rochester • NORTH CAROLINA, Durham/Chapel Hill • OHIO, Cleveland, Dayton and Euclid • OKLAHOMA, Tulsa • OREGON, Portland • PENNSYLVANIA, Bluebell and Pittsburgh • TENNESSEE, Knoxville • TEXAS, Dallas and Houston • UTAH, Salt Lake City • WASHINGTON, Bellevue • WISCONSIN, Milwaukee • ARGENTINA, Buenos Aires • AUSTRALIA, Adelaide, Brisbane, Canberra, Melbourne, Perth and Sydney • AUSTRIA, Vienna • BELGIUM, Brussels • BRAZIL, Rio de Janeiro, Sao Paulo and Porto Alegre • CANADA, Calgary, Montreal, Ottawa, Toronto and Vancouver • CHILE, Santiago • DENMARK, Copenhagen • FINLAND, Helsinki • FRANCE, Grenoble and Paris • GERMANY, Berlin, Cologne, Hannover, Frankfurt, Munich and Stuttgart • INDIA, Bombay • ISRAEL, Tel Aviv • ITALY, Milan and Turin • JAPAN, Tokyo • MEXICO, Mexico City • NETHERLANDS, The Hague • NEW ZEALAND, Auckland • NORWAY, Oslo • PHILIPPINES, Manila • PUERTO RICO, Santurce • SPAIN, Barcelona and Madrid • SWEDEN, Stockholm • SWITZERLAND, Geneva and Zurich • UNITED KINGDOM, Birmingham, Bristol, Edinburgh, London, Manchester and Reading • VENEZUELA, Caracas