

AD 748927

ASPECTS OF LARGE-SCALE
RESOURCE SHARING THROUGH NETWORKS
OF COMPUTERS

Eric Harslem
John Heafner

May 1972

D D C
RECEIVED
OCT 8 1972
RECEIVED
B

P-4833

Reproduced by
NATIONAL TECHNICAL
INFORMATION SERVICE
U S Department of Commerce
Springfield VA 22151

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

CONTENTS

	Page
ABSTRACT	v
The Market for Resource Sharing	1
Limitations of the Local Timesharing Service	3
First Order Solutions	4
The Emergence of Networks of Computers	6
The ARPA Network: Developing Network Methodology	8
ARPA Network Properties	8
Advantages of the ARPANET Design	9
ARPANET: Experience and Problems	11
Communications Subnetwork	11
User/Server Resource Sharing	11
Problems	11
Prognosis of Networks of Computers	13
Direction of Technological Advance	13
Future of Networks	14
A Glance at Some Non-Technical Problems	15
REFERENCES	17
BIBLIOGRAPHY	18

ABSTRACT

This paper discusses some advantages of resource sharing through a network of computers as compared to resource sharing through other forms of timesharing services. The ARPA network of computers, sponsored by the Defense Advanced Research Projects Agency of the Department of Defense, is used as an example of large-scale resource sharing in a computer network. This paper discusses the technical and economic aspects of computer networks, touching only briefly on legal and social implications. This paper describes some difficulties encountered in the use of computer networks and possible ways to address these difficulties. The problem areas discussed are those with application to a broad group of users, and with high potential for solution in the near future.

ASPECTS OF LARGE-SCALE RESOURCE SHARING
THROUGH NETWORKS OF COMPUTERS*

Eric Harslem
John Heafner

The Rand Corporation, Santa Monica, California

THE MARKET FOR RESOURCE SHARING

The predominant trend of the 1960s in the data processing services industry was the offering of computer timesharing services. A growing user population is indicative of the continuation of the trend, in the 1970s, toward resource sharing of computer systems on a much larger scale [1,2]. Resource sharing has expanded from the sharing of computer hardware to include (1) sharing of computer hardware and software maintenance, (2) software development, and (3) sharing of communication facilities.

A closer look at the timesharing clientele reveals a rapidly broadening customer base. The users, no longer predominantly from scientific and engineering companies, are now being attracted to timesharing as a solution for business data processing and management information and decisionmaking problems. Timesharing services are moving into educational areas--with the notion of computer access from the home just over the horizon.

* Any views expressed in this paper are those of the authors. They should not be interpreted as reflecting the views of The Rand Corporation or the official opinion or policy of any of its governmental or private research sponsors. Papers are reproduced by The Rand Corporation as a courtesy to members of its staff.

This paper will be presented at the American Management Association's conference on Systems of Computers, Los Angeles, California, 20-22 June 1972.

The development of ARPANET was supported by the Defense Advanced Research Projects Agency.

The changes in customer base indicate the transition from the use of timesharing services for scientific calculation to problems requiring that large volumes of data be exchanged between the user and his program. This transition has been made possible through the expansion of the relationship between the computer technologies and the communications industry. The growing acceptance of remote-access computing is shown by various market surveys of the growth of the communications industry:

- The computer communications market now stands at \$450 million and will reach \$2-6 billion by 1980 [3].
- Growth in peripheral devices must parallel communications growth to provide access to shared resources. *Telecommunications* reports that the \$6 billion computer peripheral market of 1970 should grow to \$11 billion by 1975 and that the \$60 million market for graphics terminals in 1971 will grow to \$450 million in 1980.*
- Creative Strategies, Inc., anticipates a \$2.2 billion revenue from cable television by 1976.† The Institute for the Future, in a recent market analysis on home consumption of information utilities via cable television and picture-phone, projected a \$15-20 billion market by the close of the 1980s.

Given this growing customer base for remote-access computing, it is of interest to examine the motivations for this growth from the small isolated timesharing services of the early 1960s to the larger, wide-area timesharing services of today, and to extend these motivations to a larger scale to demonstrate the need for, and the trend toward, networks of computers.

* *Telecommunications*, February 1971.

† *Telecommunications*, February 1972.

LIMITATIONS OF THE LOCAL TIMESHARING SERVICE

The most prominent configuration for a timesharing service in the early 1960s was a single-processor installation with many dial-up or hard-wired lines to remote-access terminals, see Fig. 1.




Fig. 1--User's terminals

Within its local sphere, this arrangement provides a high degree of resource-sharing for the computing system, while offering users computing power at a much lower rate than individually owned batch systems. However, with growth of the user population and user requirements, this simple structure presents the following:

- The user generally cannot access a large enough set of hardware and software resources for his total needs.
- The local timesharing service has few of the redundancies to provide back-up when key elements of the system fail.
- Geographical separation limits the set of users of such a local service and limits the number of services available to isolated users. This limitation is generally in the form of increased communication charges commensurate with the greater distances.
- Attempts by local timesharing services to enhance hardware and software resources often fail economically because the user population they address is not large enough or diverse enough to support special purpose services.

FIRST ORDER SOLUTIONS

In the era of the local timesharing service, the users developed solutions to the problems just mentioned. These solutions were, and usually are, aesthetically unacceptable and wasteful from an economic standpoint; however, they suffice as a temporary solution.

To expand the software and hardware resources at his disposal, and to provide a degree of back-up, a user gets access to a group of timesharing vendors. If he is in an area dense with these vendors, he has a relatively low communications expense. However, in some cases, using a group of suppliers requires increased communication cost to reach a larger area. In addition, there is no one terminal device that guarantees access to all the services of interest, which means the added expense and inconvenience of multiple terminals. Thus, while increasing the available resources, this technique increases communication and terminal costs and causes the users to learn many different access methods.

At this point, a fairly large step was taken to reduce the problem of high-cost communications. Larger timesharing services evolved (see Fig. 2).

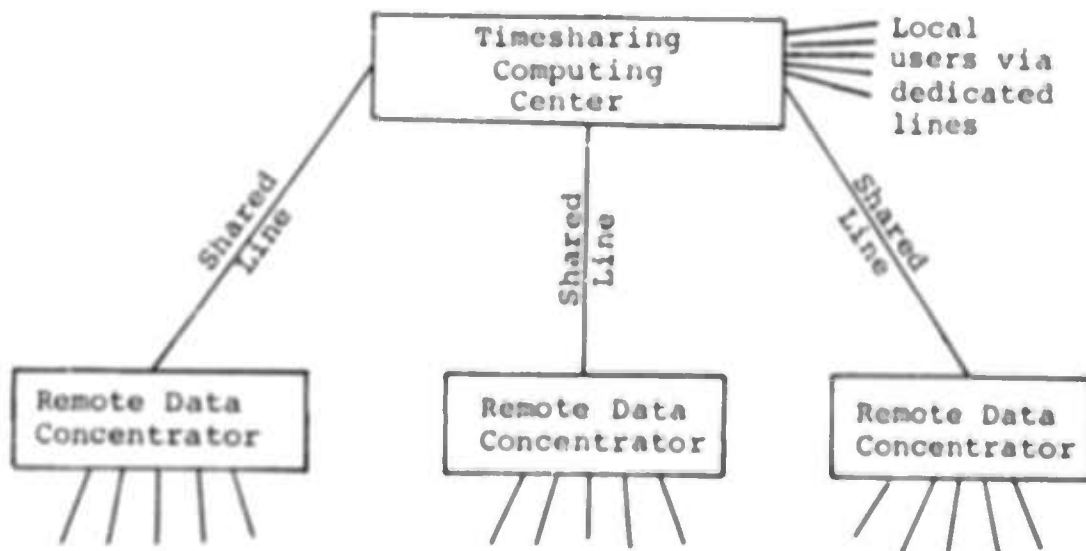


Fig. 2--Data concentration

In this model, the vendor expands his user population by putting data concentrators in areas of high user density. These users are faced only with local communication charges to access the services via the data concentrator and a shared line to the computing center.

This increase in the user community of a timesharing service provides the vendor an economic base to solve several problems previously mentioned:

- The larger revenue allows the vendor to keep more back-up equipment at his installation, providing the user greater reliability.

- The vendor is able to expand his hardware and software resources, using his increased revenue, to provide more diverse services for the user and, in turn, to attract more users.

THE EMERGENCE OF NETWORKS OF COMPUTERS

The two solutions to resource acquisition/sharing problems, mentioned earlier, are, in fact, primitive instances of computer networks. In the first case, the user at his terminal(s), with the capability to access several vendors' services, forms a *centralized network*, with the user at the center and the vendors' services being his "slave" processors. In the second case, the supplier's computing center is *the center of a network of computers*, with the data concentrators (mini-computers) being the "slave" processors.

With the continued growth of timesharing services using the data concentrator/shared line approach to expand the market place, the demands for resources easily exceed the capabilities of a single-processor. Thus, these centers usually expand their processors and peripherals in a homogeneous fashion, i.e., they add another of the same "model" equipment. The expansion serves to increase back-up and reliability.

With the continued growth of such services, it became obvious that there was no need to centralize the processing resources and that communication charges could be further reduced by distributing the centers of computation geographically. This led to the emergence of the more general form of networks of computers, shown in Fig. 3.

There are other benefits in this homogeneous expansion of resources. Data protocol problems are less severe in communicating between like systems, and identical software can be used in each system. However, with these benefits of homogeneity an older problem reappears. The expanding community of users needs access to a greater variety of computing resources. Contrary to the claims of computer hardware vendors, no one "brand" of computing hardware can solve the problems of a large user community in a cost effective manner.

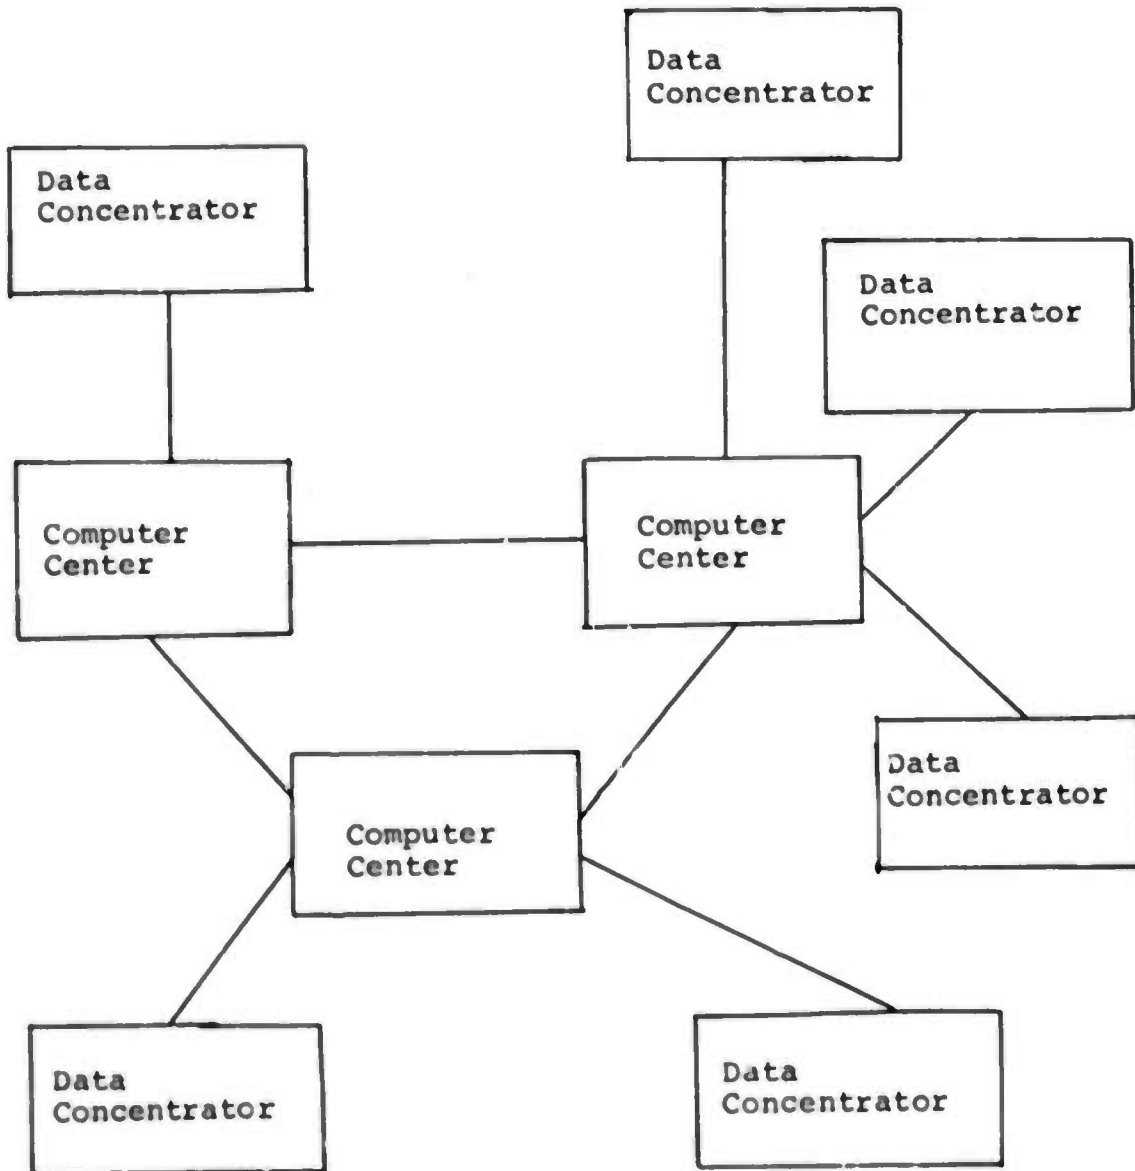


Fig. 3--Distributed computing

Thus, the current evolution and experimentation in resource sharing via networks of computers involves the interconnection of heterogeneous computing resources. We examine the Defense Advanced Research Projects Agency's (ARPA) Network as an example of this work and its associated problems.

THE ARPA NETWORK: DEVELOPING NETWORK METHODOLOGY

ARPA NETWORK PROPERTIES

Studies in the early 1960s [4] indicated the desirability of distributed digital communications networks. The preceding sections have indicated the trend of the timesharing industry toward this concept. One of the more prominent examples of computer networks is the experimental ARPA Network [5], called ARPANET.

The ARPANET* is a nationwide network, currently interconnecting 26 ARPA-sponsored research installations. The network is distributed (rather than centralized) in structure and heterogeneous in content.

As shown in Fig. 4, the ARPANET consists of two major parts: the subnetwork and the host computers. The subnetwork consists of a series of 50 kilobit communications lines and small message processors, called IMPs, at each node. The lines provide redundant communications paths between the nodes. The IMPs are responsible for handling message flow through the subnetwork. The subnetwork operates in a store-and-forward mode [6] with traffic routing governed adaptively (according to load) by the IMPs. The IMPs also handle error checking and retransmission. The subnetwork is capable of reconfiguring, when possible, to circumvent line failures.

The other part of the ARPANET is the set of host computers. At each node, there is one or more host computers. Some of these offer services to the ARPANET users, while

* ARPANET was designed by ARPA and many of its contractors. Most notably, American Telephone and Telegraph designed and implemented the circuits; ARPA provided the "plan"; Network Analysis Corporation selected the topology; Honeywell fabricated the subnet equipment; and Bolt Beranek and Newman, Inc., was responsible for most of the subnet implementation, installation, and checkout.

others only provide access to the Network for local terminals. The magnitude of the host computers varies from the ILLIAC IV to the Terminal IMPs* (TIPs) and PDP-11's.

ADVANTAGES OF THE ARPANET DESIGN

One of the greatest advantages lies in the subnet/host separation. Since the communication, error handling and routing functions are concentrated in the *homogeneous sub-network*, the *heterogeneous* collection of host computers can be interconnected without implementing communications programs in each computer.

Another advantage is the sharing of high-data-rate lines among a large group of users. In fact, the anticipated charge[†] for use of the subnetwork is \$16,500 per year plus free transmission of 4.5 million bits per month. Above this monthly transfer rate, there will be a charge of 30¢ per million bits transferred, independent of distance.

The greatest advantage, however, is the wide array of resources which such a heterogeneous collection of host computers can offer a user, with a relatively low initiation fee. For example, a TIP with communication lines rents for about \$1700/month. Via a TIP, 63 terminal devices can be interfaced at that node. Thus, for about \$30/month over the cost of a terminal, access to the resources of the ARPANET is available.

Previously, we indicated that accessing heterogeneous resources often required different terminals because of inconsistent input/output requirements of timesharing systems. In the ARPANET, each terminal is given access to all terminal applications through a terminal-oriented protocol [7].

* Terminal IMPs are slightly expanded IMPs with a minimal host facility to provide basic terminal access to the Network.

[†] *Datamation*, April 1972.

In summary, the ARPANET solves many problems regarding conventional timesharing systems. It offers at low capital (or monthly) investment rates (1) efficient use of communication facilities, (2) low data transfer charges, (3) a vast array of resources shared over a large population, and (4) single-terminal access to these resources.

ARPANET: EXPERIENCE AND PROBLEMS

COMMUNICATIONS SUBNETWORK

The performance of the subnetwork has exceeded design goals. The average outage of the communication path between two sites (using alternate paths when possible) has been under 0.5 percent. The error detection and retransmission techniques in the IMPs reduce undetected errors to 1 in 10^{12} bits transmitted. At the current network traffic level, this amounts to less than 1 bit per year of use. Currently, the response time for a message through the subnetwork averages 0.3 seconds for a maximum length message. The capacity of the ARPANET is a function of the topology and can be increased by adding higher bandwidth communication lines when and where necessary.

USER/SERVER RESOURCE SHARING

The ARPANET has been productively used since February 1971 when UCSB offered network Remote Job Entry (RJE) to their host IBM 360/75. Subsequently, other RJE facilities have joined the ARPANET at UCLA and UCSD, with the ILLIAC IV complex scheduled to offer services in late 1972. The set of timesharing service centers is considerably larger, embracing several PDP-10s, a Sigma 7, Multics, two 360/67s, and a Burroughs 6500.

Although the ARPANET was first used for remote batch services (which probably still accounts for a large amount of the network data transfer) the predominant use is now the terminal-based timesharing systems.

PROBLEMS

One major problem in turning the ARPANET into a usable entity is the heterogeneity of people associated with the

network sites. There are not only different sets of terminology, but also different sets of operating system techniques and principles. Many working sessions for developing communications protocols were initially hindered by communication gaps and divergent technical philosophies.

User problems are quite similar. As mentioned before, a terminal protocol has been developed to allow user access to many resources from the same terminal. However, the appearance of each installation, in terms of the syntax and semantics of commands, remains different. This area is one with much greater inertia (i.e., the operating system) and will not readily yield to standardization. Studies in this area are currently underway.

PROGNOSIS OF NETWORKS OF COMPUTERS

DIRECTION OF TECHNOLOGICAL ADVANCE

ARPA plans to relinquish the ARPANET when the developmental and experimental phases are complete. The ARPANET and other smaller networks serve as examples of large-scale resource sharing. The advances in technology gained from these network experiments will serve as a basis for larger and more sophisticated networks for both government and private industry.

Exploration of networking methodology, however, will continue for some time in many areas. Several prospective areas for investigation are mentioned below:

- Mass file technology has brought about feasible, rapid-access storage on a large scale (8). Individual ownership of such mass-storage systems is usually economically prohibitive; thus, they are ideal candidates for resource sharing in a network. Access to such systems opens many new possibilities for new development in areas such as management information systems and medical data bank systems, which have been constrained by their data volume in the past. Exploratory work will concentrate on access and organization of large, possibly distributed, data bases.

- Encryption techniques (9) that satisfy exacting governmental requirements for purposes of security will be further refined and applied to transmitted data. Encryption is also of concern to commercial enterprises from the standpoint of privacy of corporate data.

- Current technology allows the extension of networks via transoceanic cable, microwave systems, and satellite. Thus, intercontinental networks with even greater resource

* *Datamation*, April 1972.

sharing and communications facilities are possible. A proposed domestic, nationwide, satellite communication system could be operational by 1974.*

- The current terminal device is a keyboard and a printer. The power of a user's "station" could be enhanced with digitally encoded audio/video adjuncts. Although such additions may be economically infeasible at present, modest cost voice encoder/decoders and compression techniques are under current exploration. Such "stations" open possibilities for such fields as remote conferencing, which have application to management interaction, and social and policy experiments, which require group interaction.

- Microprogrammable subnet components will be employed and tailored to real-time applications involving complex arrangements of analog and digital equipment. With the steady improvement in the price/performance ratio of the minicomputer, such configurations may be attractive to such military and defense areas as radar and sonar research, as well as to industrial process control.

- Fault detecting and self-correcting hardware [10] will be further explored and employed in future networks. In the current ARPANET, for example, line failures are automatically detected and reported and messages are re-routed over other paths. In the event of detected software failures, cooperating programs permit a subnetwork computer to be reloaded from one of its neighbors.

FUTURE OF NETWORKS

Where do we go from here? Among those who foresee an information utility, Dr. George Feeney predicts [11] that

* *Telecommunications*, June 1971.

networks of computers will continue to grow as supplements to conventional timesharing, thus doubling the current timesharing systems' capacities by 1977. He foresees this trend leading to a computer utility where computer resources will be available and marketed in the same manner that electric power and telephone services are marketed now.

The growth in communications and computer technology toward a *computer utility* leads to an interactive computer station for the home that will not compete with necessities such as home appliances, but will compete with luxury items in the same price range.

In the final analysis, however, it appears that neither technology nor economics are significant impediments in the growth of networking. The ARPANET and others have shown the technical feasibility and the economic desirability of networks of computers.

A GLANCE AT SOME NON-TECHNICAL PROBLEMS

The computer utility is feasible, but is it socially desirable? It is feared by some persons that such a utility will widen the communications gap between levels of our society. They anticipate that those on the lower economic scale will not be offered information in a usable form at a cost they can easily bear. That is to say, it is a utility for the rich, not for the poor who are not motivated to make use of the information.

Privacy of information is another area of continuing concern. The large amount of private information now stored in computers is increasing. Existing laws and technical means for the protection of that information against unauthorized access are inadequate. Furthermore, there is little or no legal foundation providing penalties for deliberate or accidental disclosure of private data.

Technically, much work has been done to protect against erroneous computer programs and malfunctioning hardware. The

weakest link appears in the communications facilities, where little work on security has been done in areas outside government use.

There is a need for a basis from which to construct protection mechanisms. Professional licensing standards have been suggested as a foundation around which to build adequate safeguards and develop penalties for violation. Other regulatory needs are being examined by the FCC. The large body of current regulations are based on telephone and telegraph communication. Data transmission has different requirements, a variable bandwidth and high reliability. Users wish to pay for data transferred rather than a connect charge, because the data are typically sent in bursts, leaving the lines idle most of the time. Current regulations hinder small time-sharing companies from concentrating their traffic and thus forming consortia to achieve acceptable unit costs.

Tariffs are not based on the kind of equipment used with the phone lines. There are also constraints on the kind of equipment that can be connected. The FCC is not insensitive to those problems, as evidenced by the Carterphone Decision a few years ago, which permitted the attachment of foreign devices. Computer manufacturers believe that they can offer terminals at lower cost if they are allowed to provide the circuitry to interface with the phone lines. However, there is still the need for pricing structure to charge for line use at a different rate depending upon the equipment involved.

Another bottleneck in the evolving communications technology is the lack of viable rules governing the use of microwave for data transmission.

The FCC appears to be taking steps in the right direction. The problems, however, are enormously complex and there is no question but that the social, political, and legal problems (rather than the technical ones) will delay the coming of the computer utility.

REFERENCES

1. Kolker, R. J., "Time-Sharing Services--Region I," *Modern Data*, January 1971, pp. 42-48.
2. Kolker, R. J., and J. A. Murphy, "Time-Sharing Services--Region II," *Modern Data*, July 1971, pp. 46-57.
3. Poliski, Iris, "Data Communications," *Business Automation*, October 1971, pp. 22-24.
4. Baran, Paul, "On Distributed Communication Network," *IEEE Transactions on Communication Systems*, Vol. CS-12, March 1964.
5. Roberts, L. G., and B. D. Messler, "Computer Network Development to Achieve Resource Sharing," *AFIPS Conference Proceedings*, Vol. 36, 1970, pp. 543-549.
6. Heart, F. E., R. E. Kahn, S. M. Ornstein, W. R. Crowther, and D. C. Walden, "The Interface Message Processor for the ARPA Computer Network," *AFIPS Conference Proceedings*, Vol. 36, 1970, pp. 551-567.
7. Crocker, S. D., J. F. Heafner, R. Metcalfe, and J. Postel, "Function-Oriented Protocols for the ARPA Computer Network," *AFIPS Conference Proceedings*, Vol. 40, 1972.
8. McFarland, K., M. Hashiguchi, "Laser Recording Unit for High Density Permanent Digital Data Storage," *AFIPS Conference Proceedings*, Vol. 33, Part 2, 1968, pp. 1369-1380.
9. Petersen, H. E., and R. Turn, "System Implications of Information Privacy," *AFIPS Conference Proceedings*, Vol. 30, 1967, pp. 291-300.
10. Levitt, K. N., M. W. Green, and J. Goldberg, "A Study of the Data Communication Problems in a Self-Repairable Multiprocessor," *AFIPS Conference Proceedings*, Vol. 32, 1968, pp. 515-527.
11. Feeney, G., "A Three-Stage Theory of Evolution for the Sharing of Computer Power," *Computer Decisions*, November 1971, pp. 42-45.

BIBLIOGRAPHY

1. Gourley, D., "Digital Net Promises Faster and Cheaper Data Communications," *The Communications User*, Nov-Dec 1970, pp. 35-37.
2. Frank, Howard, "Spin Your Data Links into an Optimal Network," *Computer Decisions*, December 1971, pp. 6-11.
3. Hobgood, W. S., "Evaluation of an Interactive-Batch System Network," *IBM Systems Journal*, Vol. 11, No. 1, 1972, pp. 2-15.
4. Marill, Thomas, and L. G. Roberts, "Toward a Cooperative Network of Time-Shared Computers," *AFIPS Conference Proceedings*, Vol. 29, 1966, pp. 425-431.
5. Kleinrock, Leonard, "Analysis and Simulation Methods in Computer Network Design," *AFIPS Conference Proceedings*, Vol. 36, 1970, pp. 569-579.
6. Frank, H., I. T. Frisch, and W. Chou, "Topological Considerations in the Design of the ARPA Computer Network," *AFIPS Conference Proceedings*, Vol. 36, 1970, pp. 581-587.
7. Ornstein, S. M., F. E. Heart, W. R. Crowther, H. K. Rising, S. B. Russell, and A. Michel, "The Terminal IMP for the ARPA Computer Network," *AFIPS Conference Proceedings*, Vol. 40, 1972.
8. Frank, H., R. E. Kahn, and L. Kleinrock, "Computer Communication Network Design--Experience with Theory and Practice," *AFIPS Conference Proceedings*, Vol. 40, 1972.
9. Thomas, R. H., and D. A. Henderson, Jr., "McROSS--A Multi-Computer Programming System," *AFIPS Conference Proceedings*, Vol. 40, 1972.
10. Roberts, L. G., "Extension of Packet Communication Technology to a Hand-Held Personal Terminal," *AFIPS Conference Proceedings*, Vol. 40, 1972.
11. Menkhaus, E. J., "Time Sharing--More Glitter Than Gold," *Business Automation*, Nov. 1970, pp. 36-42.
12. Gaines, E. C., and J. M. Taplin, "The Emergence of National Networks," *Telecommunications*, Dec. 1971, pp. 27-29.

13. Hirsch, Phil, "Multi-Access Computer Networks," *Data-
mation*, June 1970, pp. 153-154.
14. Cerf, V. G., E. P. Harslem, J. F. Heafner, R. M.
Metcalf, and J. E. White, "An Experimental Service
for Adaptable Data Reconfiguration," *IEEE Transactions
on Communication Technology, Special Issue on Computer
Communications*, June 1972.
15. Tynes, L. R., "TYMNET--A Terminal Oriented Communication
Network," *AFIPS Conference Proceedings*, Vol. 38, 1971,
pp. 211-216.
16. Irwin, M. R., "The Telecommunications Equipment Market--
Public Policy and the 1970s," *AFIPS Conference Proceed-
ings*, Vol. 37, 1970, pp. 269-274.
17. Borko, H., "National and International Information
Networks in Science and Technology," *AFIPS Conference
Proceedings*, Vol. 33, Part 2, 1968, pp. 1469-1472.