# Implications of the Computer-Communication Partnership

F. E. HEART

Heart, F. E.: Implications of the Computer-Communication Partnership. *Medinfo 1974*. pp. 21–27.

The partnership that has recently grown up between computers and communications systems promises to have a substantial impact upon many areas of society. As computers have become less expensive and more reliable, and bulk communication bandwidth less expensive and more available, interest in computer networks has grown dramatically. The ARPA Network is prototypical and its main characteristics are described.

By allowing large numbers of geographically separated individuals and facilities to share increasingley sophisticated computer resources, and to do so efficiently and reliably, such networks can be expected to: alter the economics of computer use and the quality of computer service; catalyse the availability of specialized computer-based resources and the accessibility of centers of expertise; modify the use of small computer facilities; modify techniques for document preparation; and, in the broadest sense, change some of the traditional forms of human communication.

*F. E. Heart, Bolt Beranek and Newman Inc., 50 Moulton Street, Cambridge, Mass. 02138, U.S.A.*

## I. Introduction

Over the last few years, a strong and significant partnership has grown up between computers and communications systems. On the one hand, computers are being used to effect far-reaching improvements in communications systems; on the other, communications systems are being used to increase and extend the utility of computers. This partnership is an important new tool which will have an impact upon all the areas of society where computers are already in use, and which will also rapidly open up new areas of computer application that were heretofore not appropriate. The present paper attempts to review the recent "network" stages of growth in the computer/communication partnership and to discuss the implications of this new tool, both in a general way and with specific reference to the life sciences.

The association of computers and communications is not, in itself, a recent phenomenon. Communications circuits were attached to computers as early as the 1950's, and during the last decade the remote use of computer services has become a viable industry. Further, computers have been used within communications systems for a long time; for example, in specialized and elaborated military systems very early and in numerous "front-end" communications processors of all kinds. These instances, however, have been one-of-a-kind combinations: held together with special engineering, requiring dedicated leased lines, adapting themselves to communications systems designed for voice; as a result they have generally been difficult to implement, expensive to operate, and often unsatisfactory. To show how this not-very-easy association could become a workable partnership, it is useful to point to certain major lines of development in electronic technology.

Over the twenty-year period since the early 1950's, the electronics industry in general and the computer industry in particular have been in a state of extraordinarily productive turmoil. Two basic trends (of many that might be noted) have been the decreasing cost and increasing reliability of circuits and of computers constructed from those circuits; and the increasing availability and decreasing cost of bulk communications bandwidth. (Very roughly, computer costs per CPU cycle have dropped by a factor of ten every five years [1].) A very recent important factor in the decreasing cost of bandwidth has been the rapid integration of communications satellites into the inventory of the carriers. As an interesting example of inexpensive bulk communication, a 50 kilobit per second satellite data circuit from California to Hawaii was leased at only slightly greater cost than a voice circuit over that path (because, on such satellite circuits, voice is in fact carried as 50-kilobit PCM, and providing the data service actually required less equipment).

The improvement in cost/performance/reliability of circuits and computers has had several important results, including: (1) the development of increasingly larger computational resources. (2) the development of very small computers whose cost and reliability make them sensible for use as flexible switchgear in communications systems and competitive with special-purpose hardware. The availability of economic bulk bandwidth has led to: (1) a desire to share bandwidth in a reliable and efficient manner among users with limited needs; (2) the increasing potential competitiveness of using a distant computational resource rather than obtaining that resource locally.

## II. Computer networks

The mushrooming of developments in computer/communication during the last five years has led to a rapid growth of interest in computer networks. A recent bibliography on Resource-Sharing Computer Networks [2] lists the number of published contributions for the years 1969, 1970, 1971 and 1972 as, respectively, 46, 94, 121, and 144, and this bibliography does not include several extensive series of informal documents. This sudden increase in interest was triggered by the United States' decision, in 1968, to build the ARPA Network. The ARPA Network project focused many diverse interests and needs, and has been prominent to the point that terms such as "packet switching", "Host", "IMP", and "resource sharing" have become part of the technological jargon of the day. The ARPA Network has been widely reported, with many citations in [2]; a few key references are [3, 4, 5, 6, 7, 8], and a recent interesting summary of packet switching technology is given in [9]. Network projects are being undertaken all over the world; a recent document describing many such projects is [10], certain others are described in [11, 12, 13] and the involvement of satellite channels is discussed in [14]. In the United States, commercial carriers based upon the new technology are rapidly evolving [15, 16].

In order to provide background for discussing the implications of computer networks, and since the ARPA Network is prototypical, some characteristics of that network will be briefly summarized below. Fig. 1 is a recent "logical" map of the network, showing Host computers as well as the network nodes (Interface Message Processors or IMPs, and Terminal IMPs or TIPs). By early 1974, internode traffic on this net had reached ∿0.75 billion bits/day, and new nodes were joining the net at about one per month. Most internodal network circuits operate at 50 kilobits per second (Kbs); however, two circuits operate at 230.4 Kbs, the satellite link from the United States to Norway operates at 7.2 Kbs, and the Norway-United Kingdom circuit operates at 9.6 Kbs. Host interfaces to IMPs and TIPs typically operate at bit rates from 50 to 300 Kbs, and most terminals are connected to TIPs at rates under 2400 bits per second.

Important features of the ARPA Network include the following:

*(a)* It is a store-and-forward packe-switching system in which the path for a given message is not established in advance, each message carries an address, and messages typically traverse several nodes in going from source to destination.

*(b)* It is a highly reliable and high speed communication system in which there are typically several possible paths between source and destination, packets are held at each node until successfully received at the succeeding node, and total sourcedestination transit time is usually under 0.5 seconds.

*(c)* The standardized IMPs/TIPs and the adaptive routing of packets make the net robust in the face of circuit or node failure and also permit extremely simple addition of circuits and nodes for expansion or reconfiguration.

*(d)* The arrangements for Host connection facilitate the interconnection of computers from different manufacturers and of one-of-a-kind research computers.

*(e)* The TIP provides access to the net, and in turn to the resource pool of Hosts, for a wide variety of terminal types. By use of a "network virtual terminal" mechanism, each network Host need deal with only that "virtual" terminal, all other terminal idiosyncrasies being dealt with by the TIP.

*(f)* The basic way in which messages use the network capacity make it appropriate to consider transmission costs to be independent of distance, and to consider such costs to be based primarily on quantity of data sent rather than connect time; both major departures from prior communication economics.

*(g)* Because circuit bandwidth may be quite flexibly and efficiently multiplexed, transmission costs are very favorable compared to other competitive approaches; in a proposed commercial ARPA-like net, prices may be only 50 cents per megabit [16].

*(h)* Programs running in the IMPs and TIPs may be remotely debugged, modified or reloaded from a Network Control Center, and a new version of the program can be released to the entire net in less than two hours.

*(i)* The "Host-to-Host protocols"—disciplines by

## ARPA NETWORK, LOGICAL MAP, JANUARY 1974

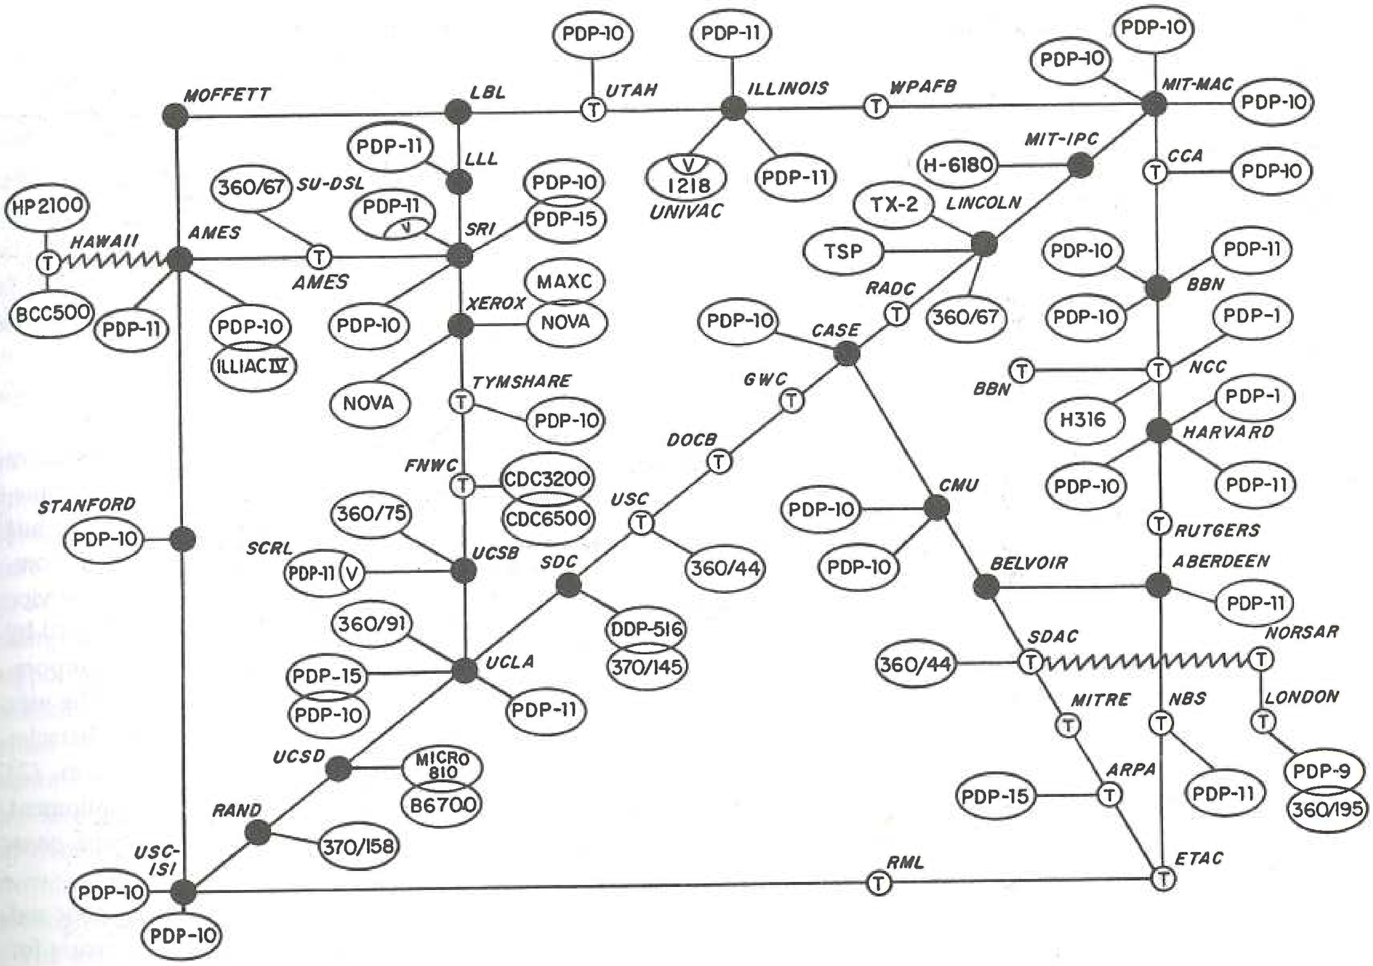


Fig. 1.

which processes in one host communicate with processes in Hosts elsewhere on the network —are largely independent of the communication system, and a given Host grouping may implement and use a network standard protocol or may use a private specialized protocol. Thus the network conveniently allows diverse private "communities of interest" to be built up "on top of" the basic network functions.

As experience has been gained with the ARPA Network, and as the number of people "thinking network" has increased, many modifications to the ARPA Network design have been promulgated. With such a young discipline many areas for improvement exist, and there is also a strong tendency to invent new approaches. On the other hand, as networks appear in many countries, and even multiple nets in the same country, the demand for intercommunication among different nets becomes important; 1974 is in fact a very active year for study of possible standards for international net interworking, and some experiments are already designed.

It is now obvious that packet-switching nets of this general type will be commonplace in a few years, and it is thus important to understand how such nets will affect the use of computers in general, specifically with relation to such important disciplines as the life sciences. In considering the implications of these new tools, a prefatory observation is in order about the difference between bare "possibility" and "convenience":

> It has been "possible" for a long time for an individual to make use of several diverse remote computers; likewise, it has been "possible" to divide a problem among two or more geographically distant computers, or to use many different types of terminals with a computer service, etc., etc. Such arrangements have been so laborious to actually accomplish, however, that they have occurred very rarely. The new computer communication partnership, and the burgeoning computer network techno-

logy have moved a myriad of such activities from *possible* to *convenient*, and now such things can and will really occur on a wide scale in many areas of application.

## III. Implications

It is diffcult to predict which of the many implications of computer networks will turn out to be most important; so many of the potential implications seem capable of profound impact that I cannot claim any particular ordering of the following thoughts:

### *Economy*

Networks are going to permit many straightforward but significant economies. For example, one U.S. university made the following approximate computation:

| | |
|---|---|
| Cost of existing campus computer center | \$1.5 M/yr |
| Cost of equivalent computer service from large machine 100 miles away | \$0.6 M/yr |
| Cost of network service to access remote machine | \$0.1 M/Yr |
| Saving | \$0.8 M/Yr |

This kind of thrift is a considerable incentive to determining the feasibility of employing the remote center; the cost squeeze on universities becomes worse every year and a good deal of pride can be swallowed for three quarters of a million dollars per year. (By way of partially explaining the surprising economy available by such routes, one should note that the total cost of a computer center, including building space, air conditioning, operators, etc. is usually about *three times* the cost of the computer in the center.) Of course, there are real problems – loss of control and potential lack of responsiveness to changing needs, priority conflicts, etc. – but many of these problems have been solved to a considerable degree in real cases. As an example, the University of Illinois has a sizable Research and Development group fully dependent upon the ARPA Network for all its computing. As a second example, Harvard University has discontinued operation of a computer center in favor of obtaining service commercially from a local company, and MIT is also considering such approaches.

This kind of calculation applies equally well to industrial situations. Consider that the General Motors Corporation in the United States operates on the order of 250 computers of many sizes and types [17]. If netting some or all of these machines would lead to the elimination of a mere 10 %, the svings would be enormous. It is not surprising that G. M. has been studying networking very seriously.

At a different level, for groups that need high bandwidth service some of the time, the use of a leased line paid for all the time, or even of dial-up service, is very expensive compared to paying only for bits actually sent.

At a still different level, consider the terminal user who can make a local call to a network access port and then, independent of distance, use the most suitably sized and priced computer on a huge net: that user is able to obtain the kind of service he needs much more economically than heretofore.

### *Competition*

In many industries, it is competition that induces cost-effective performance. A competitive forcing function has been present in the minicomputer market, but has been usually absent in the domain of large computers, corporate/university computer centers, or service bureaus. All too often the user is effectively trapped by the following constraints: (1) The university (or corporation) computer center is the only one to which the user is attached, all his files are there, and political obstacles may exist to his even considering outside services. (2) The user is technically bound by terminal equipment and protocols that are unique to the entrapping computer.

Under such conditions, service may be very poor and very expensive, and there is little economic pressure for improvement. Once connected into a network, classes of users with straight-forward computational problems are suddenly able to move files independent of distance, and to use the best, most cost-effective service on the net. There is immediate competitive pressure for improvement, easy comparison with alternatives, and a quantum increase in efforts toward good documentation, reliability, friendliness, improved rate structures, etc. In the long run, such competition will be fed back to the manufacturers as well, and we can hope for survival of the fittest.

### *Expensive and specialized resources*

Despite a certain pining for low-budget renaissance men, there has been a tendency for science to require expensive facilities. In the computer field, such facilities include very large computers (e. g., ILLIAC IV), unusually large mass storage devices ($10^{12}$ bits or more), specialized input/output devices (e. g., film readers, photocomposition devices), as well as large data bases, or large specialized bodies of software. At present the cost of such facilities is usually too great for any one group to afford, and the facility may quite literally not exist unless its cost can be distributed over the needs of many user groups. Physicists do travel to accelerator sites, and astronomers to observatories, but this kind of sharing of a facility by physically traveling to its location is difficult at best. Network technology provides a convenient, highly cost-effective way to allow many geographically separated groups to share such computer resources, and the network is therefore a major factor in whether the facility can be afforded at all. In the context of the ARPA Network a decision was reached that the ILLIAC IV and a trillion bit store would be used *only* via the network.

The notion of highly specialized resources is clearly interesting in the context of the life sciences. For some time, the U.S. Medlars system has provided medical library services, and more recently some of these services have been available via the Tymshare Network, and via an *ad hoc* connection to the ARPA Netvork. As another example, the U.S. National Institutes of Health is supporting the development and operation of a specialized computational resource called "PROPHET"[18] which is now being used by several geographically separate groups of pharmacology researchers; the system provides computational tools for molecule handling, tabular and graphical data handling, etc., and also provides a potential for collaborative research and pooling of computational tools.

The explosive growth of networks will greatly accelerate the consideration of specialized and expensive resources of all kinds. In the life sciences the possibilities abound, including: large data bases (collections of literature, drug data, public health data), specialized software systems to support research in particular fields (EKG/ECG analysis systems, biological modeling facilities, chemical analysis/synthesis system [19, 20, 21], and powerful bibliographic systems).

More generally, most computer systems, and especially those in the life sciences, embody hardware and software which is inherently difficult to transfer to other settings; the issue of program transferability has been a pernicious millstone around the neck of the computer industry since that industry's inception. There always seems to be some unique factor, and most developments continue or reinforce this waste through duplication. The network approach to resource sharing–obtaining use of someone else's program and system by using it remotely on its home ground–is a very important new way to reverse this trend.

*Centers of expertise*

As a sub-case of specialized resources, it is important to point out that human endeavor is naturally geographically distributed. Scientific groups acquire reputations as centers of expertise in some specialty, and many formal and informal paths exist for such groups to market their talents. Network technology promises another path which may be particularly rapid and effective. Consider a hypothetical group studying drug interactions, collecting drug interaction data, and systematizing that data on a local computer facility; what better way for the rest of the country to obtain accurate and complete information than to access that computer via a network? Or a hypothetical group developing software for sophisticated clinical fluid balance and electrolyte computations; what better way to obtain training or advice than to access the carefully designed public version of that software on its home computer when the need arises?

In a way this new distribution path is also an alternative to the journal literature; the ARPA Network already supports an "on-line journal" and the listings of articles, meetings, calendars, and reports can be very timely indeed.

*Marketing*

The networks of the future will serve as a very general and quite new form of marketing and distribution tool for all kinds of goods and services. Network users will include producers, wholesalers, and retailers (and perhaps thieves and police as well). As a simple but real example from ARPA Network experience, a large university computer center was experiencing difficulty with finances; the network, while connected primarily to serve research projects, was also most welcome and financially rewarding as a way to reach a national market for the services of that center. As a hypothetical example in a clinical setting, a sophisticated menu-planning service might sell services to hospitals and other institution, and take into account the prevailing local market prices of food commodities, specialized diets, etc.

*Modified roles for small machines*

In many laboratory settings, a small computer with analog-to-digital is virtually standard equipment. The storage and displays is virtually standard equipment. The advent of networks has two distinct potential impacts on this arrangement. First, it may be feasible in some instances to reduce such instruments to still smaller "terminal-like" devices and then to access larger machines on the net for mass storage, common software packages, and other assistance. In some cases such a step would be most welcome in reducing the burden of running a computer center. Biomedical laboratories in particular may view a medium size computer installation as quite a diversion, involving funding and engineering staffing that may be hard to locate and retain.

Second, networks will permit an important and novel way to expand the utility of small local computers, by sending programs from a central distribution point for actual running in the small local machine. The ARPA Network has shown that with reliable communication, it is perfectly feasible to load new code into distant machines, and this approach may permit tailoring a small local machine to the needs of the moment. In both approaches, more power is obtainable with a smaller local system.

### *Human-human interaction*

Once the ARPA Network had become operational on a wide-spread basis, a relatively simple-minded idea which, in rudimentary form, had been a "curiosity" at some installations for years, suddenly *and quite unexpectedly* became a very successful fringe benefit. The current form of the system on the ARPA Network is called "network mail", and provides that many individuals have computer "mailboxes" in their favorite Host computer; whenever A desires to communicate a message to B, C, D, and E, he logs into a convenient Host, uses a program which sends mail, lists the addressees with an indication of the addressees' mailboxes, and types the message. No matter where B, C, D, and E might be, they occasionally enter the net at a local node, log into their (mailbox) Host, and as part of the login process are automatically told that a message exists for them. They may then obtain any recent messages, using a program to read their mail.

For significant numbers of people, this system has become, almost overnight, an essential aid to their work. A directory [22] issued in January 1974 included more than 450 individuals with network addresses, and the actual number having addresses and using the facility was probably several thousand; many of them do not let a day go by without checking their network mail. The system has many elaborations, including distribution lists, filing aids, etc., but its primary impact has been to permit communication at any hour of the day or night without even knowing the city, state or country where the addressee might currently be traveling.

As a management tool, or for communicating with busy people on the move, or as an aid to working at odd hours, the system has substantially changed the lives of many individuals. Some people now carry terminals for use from hotel rooms (via acoustic couplers) or even from airport phone booths that happen to be near a power plug.

Other similar systems for human-human interaction are being designed, including various forms of teleconferencing [23], and it is easy to imagine these new approaches to human communication becoming very popular among the academic and industrial communities.

In a broader context, the above is an example of the truly unexpected developments that can come about when a "community of interest" forms that is ten or a hundred times larger than before; and I believe that such spontaneous inventions should be "expected" in the life sciences and elsewhere.

### *Document preparation*

Another kind of human-human interaction can be seen emerging from the already widespread use of systems for on-line text editing and document preparation. The advent of networks potentially provides broad access to such systems, and increases the value of keeping documents on-line for both modification and distribution. The notion of collaborative work on a paper by people who are geographically distant becomes obvious, and the "turnaround" time becomes minutes rather than days for new drafts. Further, there is a strong incentive for the development of high quality final output printing systems, spelling-correction systems, grammar-assistance systems, etc. A number of documents have been produced by "distant" groups of people, using the ARPA Network, and the interaction has apparently been very convenient and stimulating. Within the ARPA community a major effort is now being mounted to pursue these kinds of tools for expediting human interactions.

### *Employment/Training*

The computer industry has had a tradition of providing some limited amount of employment to part-time people, to people who wished to work partially at home, and to people who worked over a terminal and seldom came to the office. However, this kind of job market has always been a rather localized affair, only sensible quite nearby the computer facility and not very practical when terminal use involved toll telephone charges. As a somewhat speculative notion, the advent of networks and of charging policies independent of distance may open up new flexibility in employment for people who live at great distances from a prospective employer. In some cases moves may be avoided or delayed, and greater individual freedom may accrue.

Along a related line of thought, it has always been necessary for people to do a certain amount of "going where the jobs (or schools) are," and individuals whose mobility is limited have often been unable to do so. For example, many kinds of physical handicaps make going to an office difficult, and moving from city to city prohibitive. Network technology may conceivably provide a wider potential job market for such individuals; training centers may be able to train people who are far away, and even groups whose mobility is very poor (prison inmates?) might find an expanding set of options coming in on a wire.

### *Computer research*

Finally, and as is only appropriate, the new partnership of computers and communications has a direct feedback path into the further advance of computer technology. There has not yet been an operating system designed specifically with a network in mind. There has not yet been enough work on programs which allow "self-help" in the use of the program for a distant and naive user.

Study of the proper ways for processes in a Host to communicate with other processes has been given a shot of adrenalin by the *need* to standardize and understand such interactions. Programs that live in several geographically separated computers and interact are being studied, multi-computer operating systems are being studied, input-output efficiency is being carefully considered and optimized etc. The network technology has been a large stimulant to computer technology already, and we have barely scratched the surface.

## References

1. Roberts, L. G.: Data by the packet. *IEEE Spectrum,* pp. 46–51, Feb. 1974.
2. Annotated Bibliography of the Literature on Resource Sharing Computer Networks. U.S. Department of Commerce, National Bureau of Standards, NBS Special Publication 384, SD Catalog No. C13.10:384, Washington, D.C., Sept. 1973.
3. Roberts, L. G. & Wessler, B. D.: Computer network development to achieve resource sharing. AFIPS Conference Proceedings *36:* 543–49, June 1970.
4. Heart, F. E., Kahn, R. E., Ornstein, S. M., Crowther, W. R. & Walden, D. C.: The interface message processor for the ARPA computer network. AFIPS Conference Proceedings *36:* 551–67, June 1970.
5. Ornstein, S. M., Heart, F. E., Crowther, W. R., Rising, H. K., Russell, S. B. & Michel, A.: The terminal IMP for the ARPA computer network. AFIPS Conference Proceedings *40:* 243–54, June 1972.
6. McKenzie, A. A., Cosell, B. P., McQuillan, J. M. & Thrope, M. J.: The network control center for the ARPA network. Proceedings of the First International Conference on Computer Communication, pp. 185–191. Washington, B. C., Oct. 1972.
7. McQuillan, J. M., Crowther, W. R., Cosell, B. P., Walden, D. C. & Heart, F. E.: Improvements in the design and performance of the ARPA network. AFIPS Conference Proceedings *41:* 741–54, Dec. 1972.
8. Heart, F. E., Ornstein, S. M., Crowther, W. R. & Barker, W. B.: A new minicomputer/multiprocessor for the ARPA network. AFIPS Conference Proceedings *42:* 529–37, June 1973.
9. Metcalfe, R.: Packet Communication. MAC TR-114, Project MAC, Massachusetts Institute of Technology, Cambridge, Mass., Dec. 1973.
10. Proceedings of the International Advanced Study Institute on Computer Communication Networks (Sponsored by NATO, held at Univ. of Sussex, Sept. 1973). Noordhoff International Publishing Co., Leyden, The Netherlands, 1974.
11. Pouzin, L.: Presentation and major design aspects of the cyclades computer network. Proceedings of Datacomm 73, pp. 80–85. IEEE Catalog Number 73 CHO828–4C, Nov. 1973.
12. Oizumi, J.: Plans for a Japanese university computer network. Proceedings of the Hawaii International Conference on System Sciences (Computer Network Subconference). Western Periodicals Co., Jan. 1974.
13. Barber, D. L. A.: The European computer network project. First International Conference on Computer Communication, pp. 192–200. ICCC 72 CHO 690–8c, Oct. 1972.
14. Abramson, N.: Packet switching with satellites. AFIPS Conference Proceedings *42:* 695–702, June 1973.
15. Packet Communications Inc.: Application to the FCC for Common Carrier Status. FCC No. P–C–8533, Jan. 1973.
16. Telenet Communications Corporation: Application for a Public Packet Switched Data Communications Network. FCC No. P–C–8750, Oct. 1973.
17. Request for Information, General Motors Corp., Data Communications Task Force, Aug. 30, 1972.
18. Castleman, P. A., Russell, C. H., Webb, F. N. Hollister, C. A., Siegel, J. R., Zdonik, S. R. & Fram, D. M.: The implementation of the PROPHET System. To be presented at the National Computer Conference and Exposition, Chicago, Ill., May 1974.
19. Corey, E. J. & Wipke, W. T.: Computer-assisted design of complex organic syntheses. Science *166:*3902, 178–92, Oct. 1969.
20. Feigenbaum, E. A., Buchanan, B. G. & Lederberg, J.: On generality and problem solving: a case study using the dendral program. *In* Machine Intelligence *6.* Edinburg Univ. Press, 1971.
21. Buchanan, B. G. & Lederberg, J.: The heuristic dendral program for explaining empirical data, pp. 91–100. IFIP Congress 1971, Ljubljana, Aug. 1971.
22. ARPA Directory. NIC 19275, Network Information Center, Menlo park, Calif., Jan. 1974.
23. Several Papers in a Session on teleconferencing. First International Conference on Computer Communication. ICCC 72 CHO 690–8c, Oct. 1972.