Applied Math 254: Computer Networks

1. Introduction to Computer Networks

John M. McQuillan and David C. Walden
January 1975

Table of Contents

1. The Growth of Computer Communications*

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 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. Teleprocessing and remote job entry are now quite conventional techniques. Further, computers have been used within communications systems for a long time; for example, in specialized and elaborate military systems very early and in numerous "front-end" communications processors of all kinds. These real-time applications guarantee rapid response to events in a transaction-oriented system. Finally, we may cite the growth of time-sharing systems, which have also stimulated the development of enhanced I/O capability. 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 tentative association has 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 [Roberts 74].)

The improvement in cost/performance/reliability of circuits and computers has had several important results,

----

*The material in this section is based on [Heart 74].

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.

2. The Growth of Computer Networks

We have noted that 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 technology, 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.

We will now explore the reasons for the growth of networking, its problems, and its implications. We start by defining a computer network as an interconnected group of independent computer systems which communicate with one another and share resources such as programs, data, hardware, and software.

2.1 Reasons for Computer Networks*

2.1.1 Motives for Interconnection

1. Coordination. Where the activities of different organizational units interact strongly, links between their information systems will provide horizontal communication to aid operational coordination. Similarly, within an organization, hierarchical links between operational information systems and those serving higher levels will provide vertical communication for planning and control purposes.

2. Hardware resource sharing or load leveling. The processing resources of the network--CPU's, core, on-line storage, etc.--can be pooled. Dynamic load leveling can be achieved by shipping programs and data between nodes, or by shipping data only and using locally held software. With resource pooling, general growth in demand can be handled without having simultaneously to step up the capacity of each node. The steps in incrementing network capacity are therefore smaller, making for less spare capacity between increments. For example, one U.S. university made the following approximate computation:

cost of existing campus computer center $1.5M/yr

----

*The material in this section is based on [Bernard 73].

cost of equivalent computer service from large machine 100 miles away $0.6M/yr

cost of network service to access remote machine $0.1M/yr

saving $0.8M/yr

One should note that the total cost of a computer center, including building space, air conditioning, operators, etc, is usually about <u>three times</u> the cost of the computer in the center.

3. Data and program sharing. By making the data base and software library of each node accessible to the others, one can eliminate duplication of effort in their construction and maintenance.

4. Enhanced Availability. Pooling of hardware and software resources can also provide redundancy. The facilities of a remote node can be used as back-up in case of local failure.

5. Access to specialized resources. Where a node possesses specialized resources--i.e., specialized hardware or software facilities, or data bases--these can be made available to users at other nodes. Moreover, with economies of scale the enlarged user base created by the network may justify the creation of specialized resources which would not be justified for a single node.

2.1.2 Motives for Decentralization

Clearly, an alternative means of exchanging information between information systems and of pooling resources may be provided by centralizing all data processing functions in a single installation. Centralization provides advantages of scale in hardware costs and in the data processing operation generally. It eliminates the considerable technical and managerial problems of operating a network. Thus, we must examine reasons why the network alternative may be chosen.

1. Scale of demand. Demand for data processing facilities may be so large that a multiple computer network would still utilize the economically optimal size of machine. In this case, there are no potential economies of scale in centralization.

2. Local control. Decentralization allows for local control of data processing and provides responsiveness to local needs. It is arguable whether these results require actual physical localization of the hardware, but in many cases--particularly where several independent organizations

are involved--centralization may be politically impossible.

3. Distribution of risk. A centralized system is totally vulnerable to failures, while decentralization spreads this risk. For similar reasons, a data base may be decentralized to reduce its vulnerability to security breaches.

4. Reduction of communication costs. Where a set of geographically-distributed terminals is to be given access to computer services, the communication costs involved will often be lower when this is done through a set of regional centers rather than a single central installation.

5. Economies of specialization. In terms of pure cost per processor cycle there are undoubted economies of scale in combining several data processing tasks on a large multi-programming machine. However, the overhead costs of providing the required complex general-purpose hardware and software are enormous, and small simple machines without these high overheads in capital and operating costs may provide an attractive alternative. Distribution of functions among such machines also permits choice of each machine to fit the specialized requirements for its particular function.

6. Enhanced adaptability. Distribution of functions among several machines enforces a modularity on the overall system, making it easier to incorporate changes in technology or requirements that affect only one processor or function.

2.2 Problems with Computer Networks

1. Loss of local control. One of the obvious difficulties with using a network to access remote computational resources instead of using local facilities is that there is the danger that the remote users will have less effective control over the computer they are using than if it were local. This may be reflected in matters as diverse as scheduling, maintenance, documentation of changes, and so on.

2. Lack of responsiveness to changing needs. A related danger is that the network will become a barrier between the service supplied by a computer center and the users of that service, so that their needs are only perceived and acted upon slowly if at all. This is a particular problem if the network environment <u>creates</u> the need for changes in the service.

3. Priority conflicts. There may be conflicts between what the managers of the computing facility want to

accomplish and what the remote network users want, and even among the wishes of the many individual users who may be geographically scattered and working on unrelated projects.

4. Performance difficulties. As we will detail in subsequent notes, the presence of a network implies very different performance for the end user than a simple communications circuit. The delay and bandwidth of his communications may be restrictive, and may be highly variable. Further, he may face different kinds of reliability problems due to the many additional components which can fail in the network.

We conclude this brief summary of some of the typical difficulties of network use by noting that most of them can be solved by good network design or by careful network management; they are not intrinsic problems with the computer communications themselves.

2.3 Advantages of Computer Networks*

2.3.1 Competition

In many industries, it is competition that induces cost-effective performance, 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 straightforward 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.

----

*The material in this section is based on [Heart 74].

2.3.2 Expensive and Specialized Resources

There has been a tendency for science and some industries to require expensive facilities. In the computer field, such facilities include very large computers (e.g., ILLIAC IV), unusually large mass storage devices (10 to the 12th 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.

2.3.3 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.

2.3.4 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. Many large university computer centers experience difficulty with finances. A network, which may be used primarily to serve research projects, may turn out to be most welcome and financially rewarding as a way to reach a national market for the services of that center.

2.3.5 Modified Roles for Small Machines

In many laboratory settings, a small computer with analog-to-digital conversion equipment, modest mass 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.

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.

2.3.6 Human-Human Interaction

Once the ARPA Network had become operational on a widespread 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 issued in January 1974 included more than 4500 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.

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.

2.3.7 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.

2.3.8 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 understand and standardize 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.

3. Characteristics of Computer Networks

Having considered some of the reasons people use networks, and some of the effects networks have on them, we turn to the technical aspects of networks. Here we are concerned more with an explanatory catalog of possibilities rather than a detailed look at any one technique.

3.1 Computer Network Organization*

First we look at the possible topological layouts for a network, together with the underlying organization.

3.1.1 Computer Networks and Communication Networks

A network can be defined as an interconnected set of points, or "nodes." The difficulties arise due to different interpretations as to what constitutes a node of the network. Figure 3-1 depicts the simplest form of computer-communications system, in which terminals are linked into a single central computer in a "star" configuration. If terminals and multidrop points, as well as the central computer are regarded as nodes, such a system may be described as a "computer-communications network" or simply a "computer network." However, it does not constitute a computer network in the sense of a network of many computers. In this latter sense, the term "computer network" refers to a system in which several computers are linked together with tasks distributed among them.

----

*The material in this section is based on [Bernard 73].

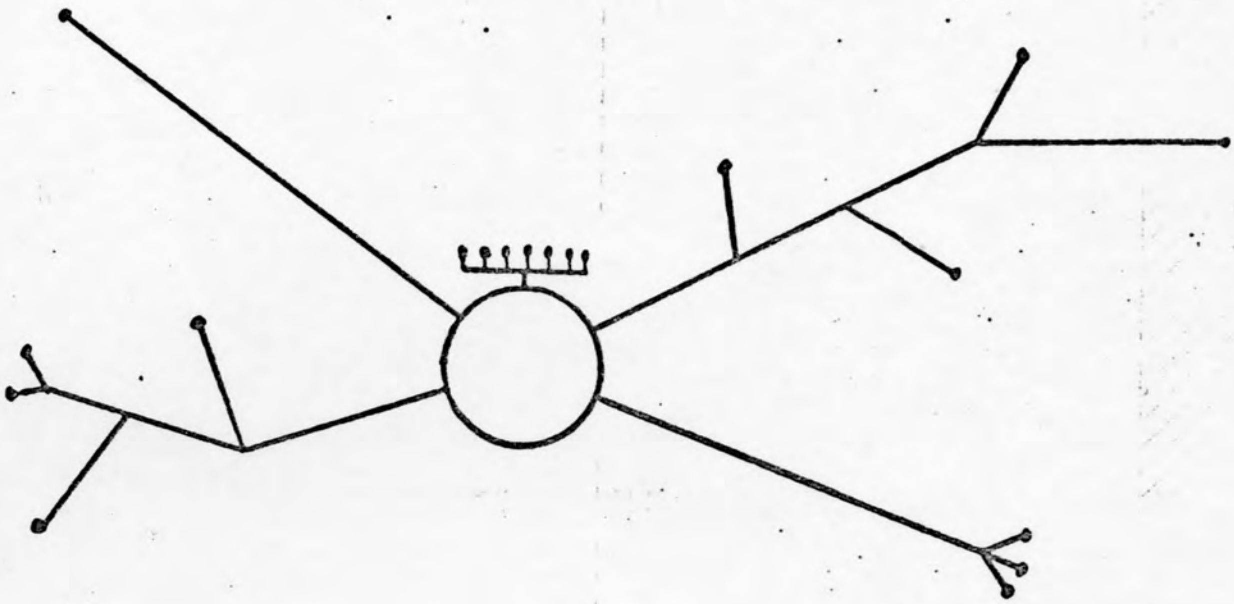

A computer-communication system focused around a central computer installation.

Figure 3-1 Computer Communications System

3.1.2 Centralized and Decentralized Organization

The distribution of functions among the nodes of a computer network may take many forms. It will be useful to draw a distinction initially between two extremes.

Figure 3-2 illustrates one approach. Various functions that would normally be carried out by the central machine are now handled by satellite processors--typically mini-computers. In the example shown, machine (a) is a communications processor, handling message assembly, buffering, translation, transmission, error recovery and polling tasks associated with the communications network; machine (b) is a line concentrator, carrying out similar functions at a remote location; machine (c) is a peripheral processor controlling the data base and carrying out associated file-handling functions; and machine (d) is an "intelligent" terminal with some autonomous processing capability. The underlying communications system has a basic "star" configuration focused around the central installation, and there is an hierarchical relationship between this machine and its satellites. Accordingly, this form is described as a <u>centralized</u> network. A highly centralized network raises few radically new problems beyond the existing technology of computer-communications systems, so that this report is not greatly concerned with this form.

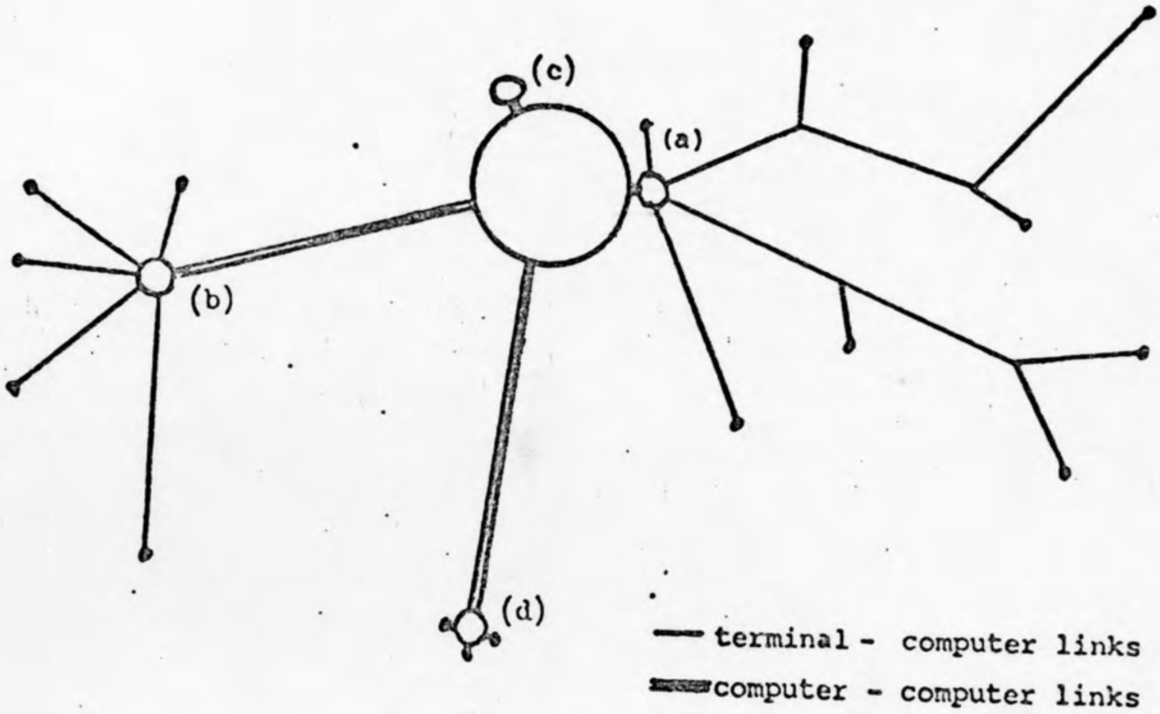

A centralized computer network, formed from a large central processor and several functionally specialized satellites.

Figure 3-2 Centralized Computer Network

Consider by contrast the system depicted in Figure 3-3. Here two completely autonomous computer systems have been linked together to permit data exchange and/or sharing of resources. Again, processing and storage functions are distributed between the nodes of the network, but here the nodes are completely autonomous, cooperate on an equal basis, and there is no central focus to the network. The network is fully decentralized or distributed. The intercomputer link here represents an entirely new level of interconnection in the communications system.

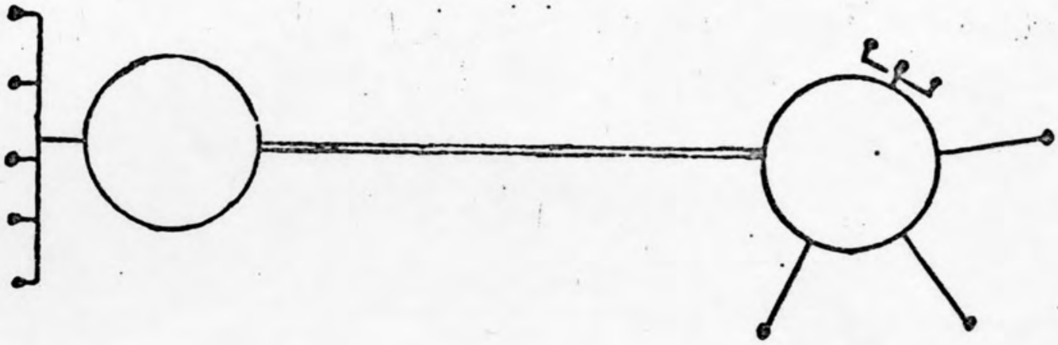

Two autonomous computer systems linked together to permit resource sharing, forming a fully decentralized network.

Figure 3-3 Decentralized or Distributed Computer Network

Full centralization or decentralization represent the extremes of a spectrum of possible organizations, reflecting different degrees of autonomy for the satellites. A common form of semi-centralized system is illustrated in Figure 3-4. The satellites in this network perform most processing themselves locally, but for some purposes they act as terminals of the central installation. For example, the satellites might pass large processing tasks such as mathematical programming jobs up to the central machine. Or they might only support time-sharing work, with all background batch work transmitted through to be processed centrally. There are also forms or organization in which a central machine acts effectively as a slave to its "satellites," performing common functions and perhaps holding a common data base.

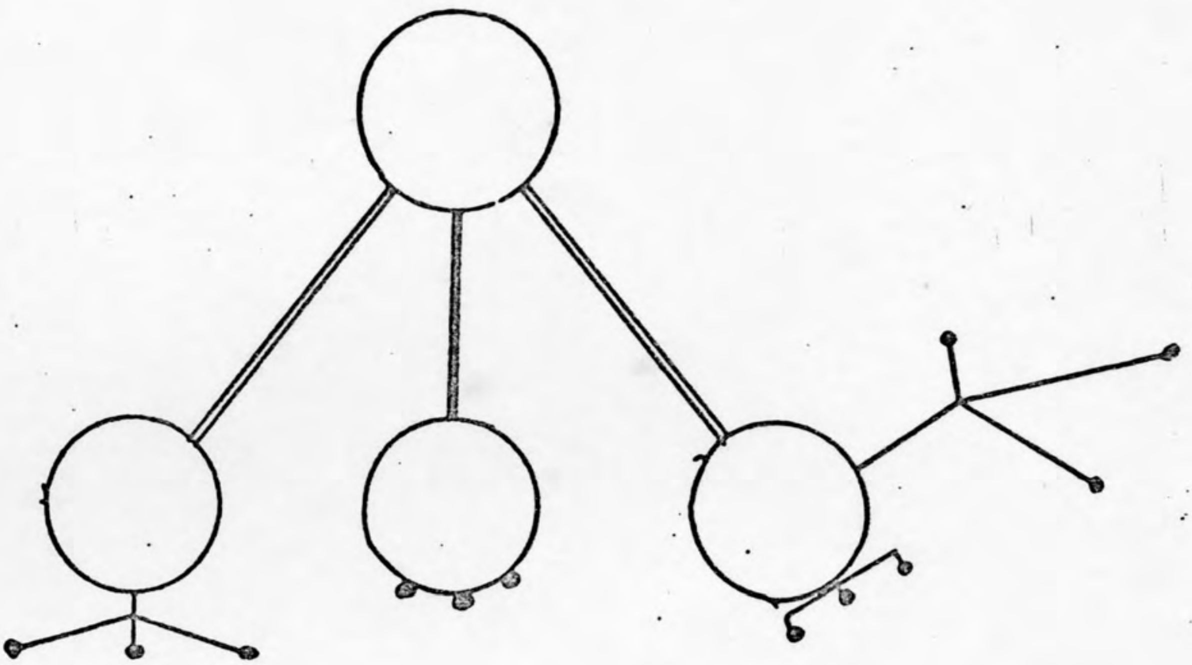

A semi-centralized network in which the satellites have considerable autonomy but act as terminals to the larger central machine for certain functions.

Figure 3-4 Semi-Centralized Computer Network

To summarize some of these concepts in abstract form, we present several kinds of network topologies in Figure 3-5. The three networks in Figure 3-5a represent various specialized approaches to network connectivity. They range from the star to the fully-connected network in terms of the number of lines, but they are all relatively simple in terms of topological layout. The other three networks are abstract examples of more general network geometries, which we have called distributed networks.

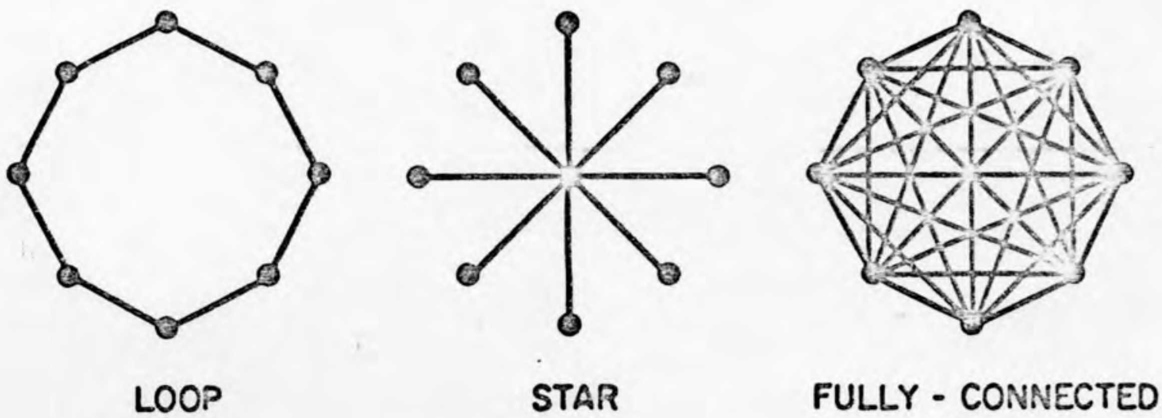

a. SOME NON-DISTRIBUTED NETWORK TOPOLOGIES

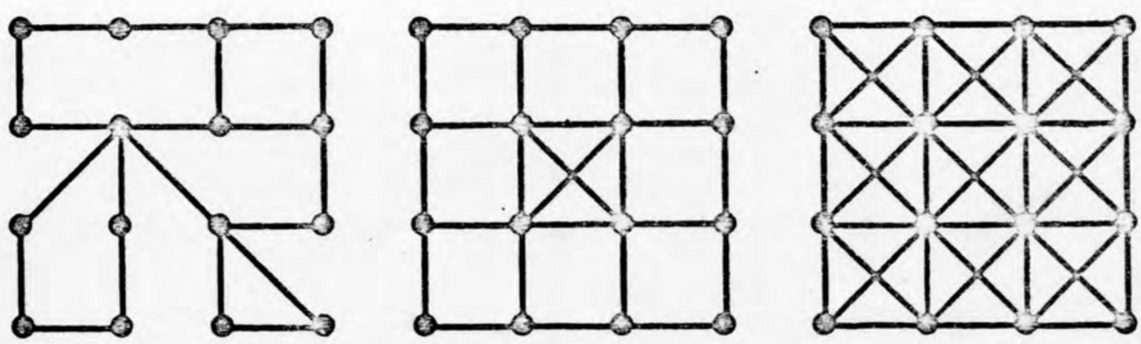

b. SOME DISTRIBUTED NETWORK TOPOLOGIES

Figure 3-5 Some Distributed and Non-distributed Networks

3.1.3 Functional Differentiation

Another important parameter reflecting the organization of the network is the degree of differentiation between the functions of the various nodes. Each processor may perform a highly specific function in the network, or they may all have the same general-purpose capability. The centralized network in Figure 3-2 is functionally quite heterogeneous, while the fully decentralized network of Figure 3-3 is functionally completely homogeneous. The degrees of centralization and of differentiation of function are not necessarily correlated, however; there are fully decentralized networks which are functionally heterogeneous.

3.1.4 Front-End Communications Processors

A second important design decision related to the organization of the network, and to its homogeneity is the choice between using a front-end processor to handle all communications functions and using the main computers for these tasks. The advantages of the front-end approach are many: the smaller machine may be more cost-effective at communications; it is easier to standardize and control the network connections; this approach may lead to more reliable communications. The two main drawbacks are the added cost and added complexity of more computers in the whole system.

3.1.5 General-Purpose Organization

The form of organization employed in a network must be judged not simply by the size and capabilities of the systems involved, but by the functions and functional relationships entailed in the particular use that is made of the network. Where the systems are general-purpose rather than being dedicated to a particular task, the network may be designed to be used for several different purposes. Thus, one can find many distinct organizations superimposed on the same physical network, a common user network.

3.1.6 Hierarchical Networks

Intercomputer networks may also be arranged in hierarchies. For example, the network in Figure 3-6 has a hierarchical structure in which the upper level is a fully decentralized network, itself formed partially from centralized networks. Again, the hierarchical relationships must be judged in terms of functional rather than physical organization: if computer (a) acted as a satellite of computer (b) in Figure 3-6, it would have to be regarded as part of the latter's subnetwork rather than as an independent node of the higher-level network.

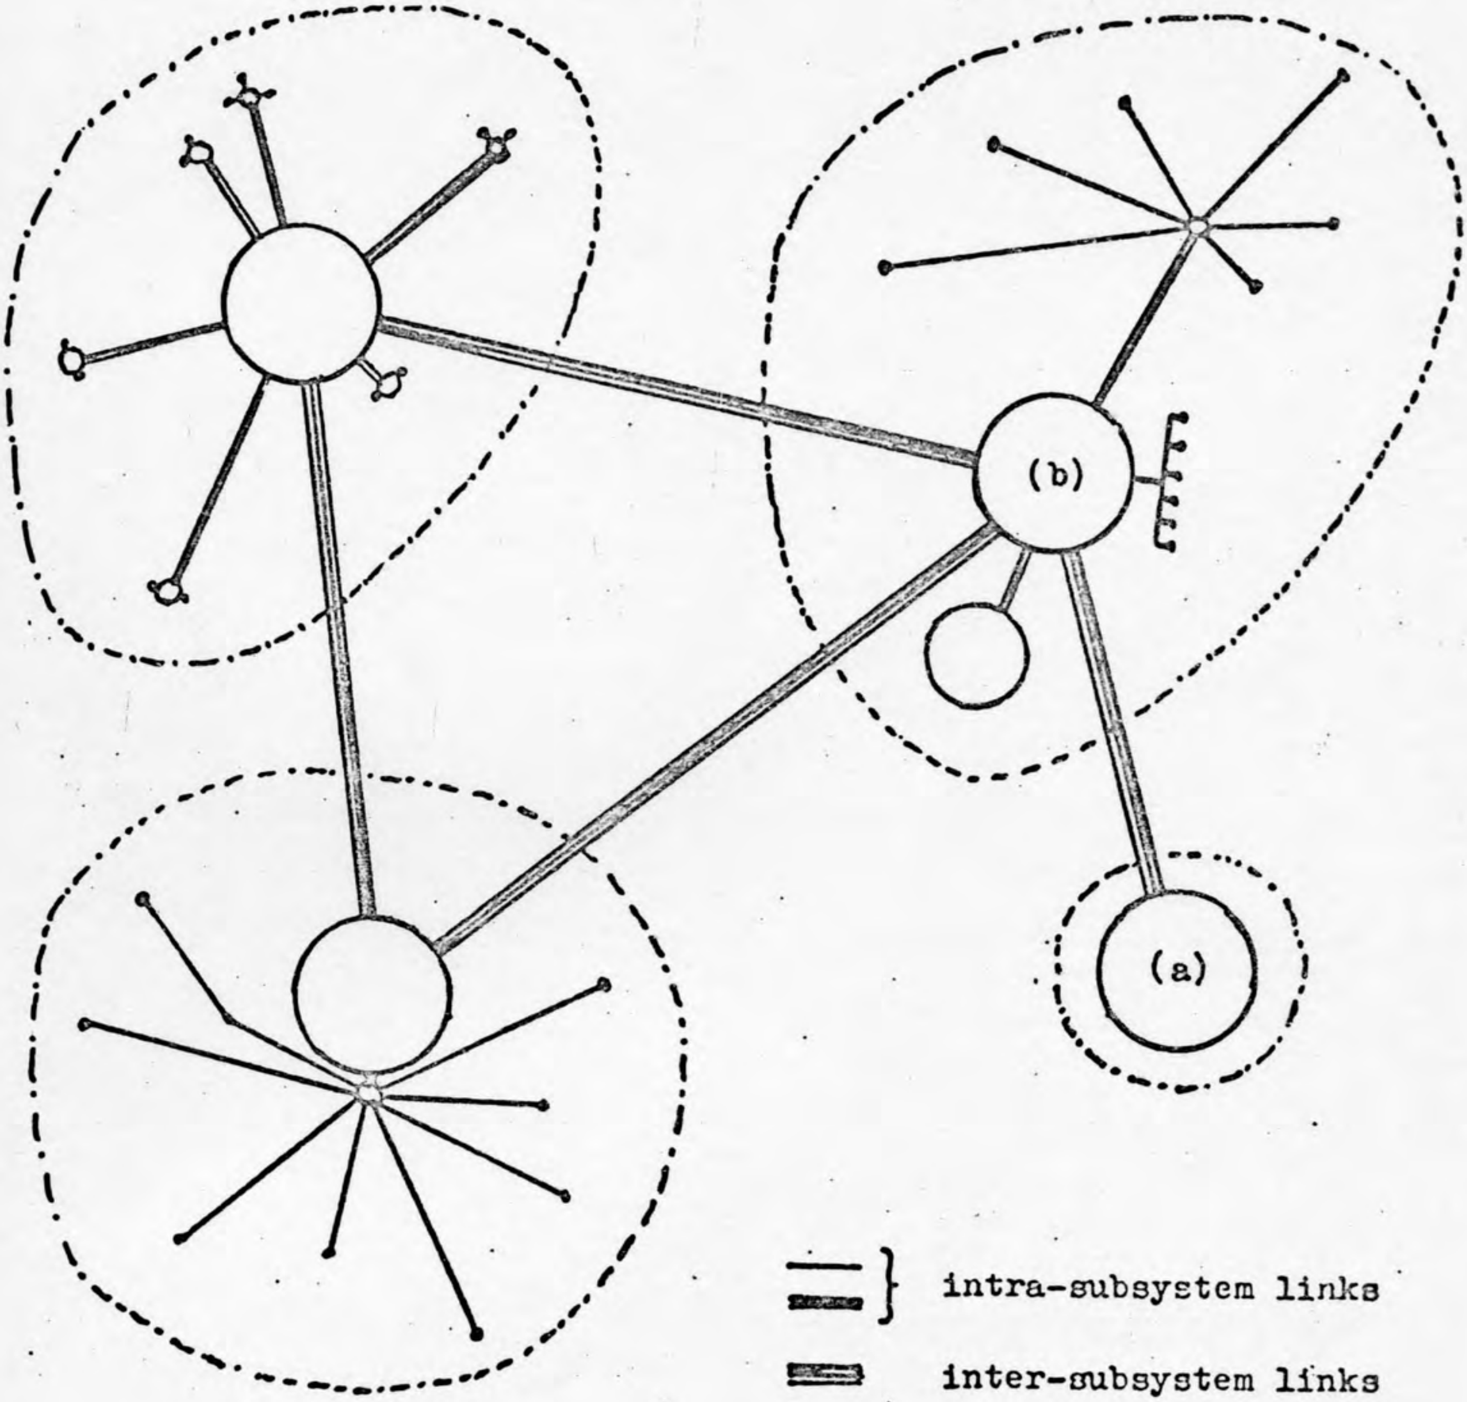

Hierarchical intercomputer network: a fully decentralized network some of whose nodes are the foci of centralized sub-networks.

Figure 3-6 A Hierarchical Computer Network

3.2 Technology

In this section we introduce several aspects of the technology used to implement computer networks. These issues are described in order to provide a framework for comparing the approaches used in different kinds of networks.

3.2.1 Switching Protocol

The major distinction to make in discussing the technology of computer networks is the switching protocol used by the nodes to transfer the information through the network. We distinguish three major categories: circuit switching, message switching, and packet switching.

Circuit Switching

Circuit switching is the traditional method used by the telephone network for the establishment of a voice circuit through many geographically separated switching centers. It refers to the technique of identifying a dedicated path of several communications lines linked together for the purpose of providing a virtual circuit from source to destination. Circuit switching has been adapted to provide a means for linking together computers and terminals in a switched network, primarily in simple applications like terminal-to-terminal communication at a fixed rate corresponding to the circuit rate. A circuit-switching network should not be confused with the use of the common dial-up telephone network for the low-speed communication requirements in many computer applications. A circuit-switching network explicitly manages a number of circuits (which may be dial-up or leased) on the basis of the demands of the users of the network, just as the automatic telephone switching system manages the use of the lines for voice communication.

Message Switching

Message switching is a technique developed in the 1950's for computer networks which is based on the earlier technology for telegram and other torn-tape transmission systems. When a message is sent into the network, the source node waits until it is fully accumulated, and then stores the message on secondary storage. Then, minutes or hours later, the source node attempts to transmit the message over the outgoing communications circuit. After the message is successfullly received at the next node, the process is repeated, perhaps several times, until the message reaches the destination. The nodes attempt to optimize the use of the (generally low-speed) transmission circuits by buffering all messages until the lines are ready to send them, even if it means holding a message for some hours. In this respect, message switching is a direct descendant of those previous technologies for message transfer (telegram, etc) in which human intervention was required for each message, and for which physical storage (bins of torn tape on shelves) was reserved for messages in transit. This type of system is also known as a store-and-forward network because of the action of transmission and acknowledgement of receipt at each step of the network.

Packet Switching

Packet switching is a technique which can be seen as an enhancement of message switching, particularly with respect to the speed with which the communications equipment is switched between one message and the next. In a packet-switching network, no secondary storage is used for long-term message holding, and a message may be sent from the source to the destination in pieces known as packets. In fact, the first packet of the message may be transmitted through the network before subsequent packets have been received at the source. A store-and-forward discipline is followed here as well: each node holds a copy of the packet it transmits until it receives an acknowledgment from the adjacent node. However, this storage is for periods of tens or hundreds of milliseconds only, so that a small amount of primary storage is sufficient for this function.

The most important aspect of packet switching is that it represents the most dynamic switching protocol of the three methods discussed here. Circuit switching dedicates a path to a conversation for minutes at a time. Message switching waits until the circuit is free, and then sends what may be a lengthy stream of data for a message. Packet switching uses a network line long enough to send out a single packet (tens of milliseconds) and then makes a new decision about which traffic to send on the line. This makes it possible for highly interactive traffic to be routed through a packet-switched network quickly, and for the bandwidth of such a network to be used very efficiently if there are a number of sources of traffic with non-uniform input rates.

3.2.2 Transmission Modes

In the previous section, we discussed three techniques for the transmission of messages along the network communications paths made up of many lines. Now we discuss the technology used for transmission over individual network lines.

Digital vs. Analog

Communications facilities have traditionally been analog systems, with carrier and various modulation techniques, in order to transmit voice efficiently. Although computers and other digital devices can communicate over such lines through the use of modems (modulator-demodulators), it is likely that the common carriers will begin to make digital service more readily available in the future.

Synchronous vs. Asynchronous

The transmission of a single bit of information on a communications circuit depends on the receiver knowing exactly when to sample the input signal to tell if it is a one or zero. This can be accomplished in one of two basic ways: either the receiver can sample the signal synchronously, at some fixed rate, and in phase with when the sender asserts the data, or the receiver can rely on a second signal to tell when to sample the input asynchronously, according to when the sender turns on the second, control signal.

Full Duplex vs. Half Duplex

Some communications facilities work in only one direction at a time, and special signals must be transmitted to "turn around" the line; these lines are known as half duplex, and are most common in some terminal applications and in locations where circuit costs are high. The alternative is a full duplex line which can be operated in both directions at the same time.

Data Compression

Some systems use data compression to reduce the total number of bits in transmissions by one of several techniques. Repeated bits or characters can be replaced by a count of the number of occurrences; encoding techniques can reduce the number of bits needed to represent the most commonly used bit patterns, and so on.

3.2.3 Transmission Rates

Generally, the transmission rates available to the designers of computer networks are restricted by the offerings of the common carriers, unless special transmission facilities are developed for the network. Some standard rates are:

2400 baud
4800 baud
9600 baud
50 kilobaud
56 kilobaud
230.4 kilobaud

3.2.4 Transmission Medium

As mentioned above, the designers of a computer network have two basic alternatives in acquiring the communications circuits for the network: use a common carrier offering, or develop a special facility. The common offerings include dedicated leased lines, dial-up lines, digital data service, and so on. There are such services that work within cities, going through relatively old central office equipment, as well as those which use high-speed microwave relay, and also satellite circuits. A local distribution network may be able to obtain much better service in terms of data rate and error rate by the use of special equipment, such as direct-wire connections or coaxial cable transmission facilities.

3.2.5 Message Format

The last technical consideration of a general nature that we mention here is that of message format. Some networks restrict messages to be a fixed length regardless of how much data is sent, while others permit a variable length with a fixed maximum. Others may permit virtually infinite-length messages, though some fragmentation is likely. Some networks deal with a single character set only, or with a limited number of codes, while others permit complete transparency in the binary data transmitted. Although problems of standardization, and interfacing to existing equipment, as well as efficient use of network resources, may dictate such restrictions, clearly the fewer restrictions that are placed on the network user, the easier it is for him to adapt to the network environment.

3.3 Applications of Networks*

A discussion of existing and potential networks will demonstrate the great range of application areas within which the networking trend is discernible. It also serves to introduce a number of specific network examples which provide a context for the general discussions that follow.

3.3.1 Message-Switching Systems

---

*The material in this section is based on [Bernard 73].

An early application of computers was in the automation of the teletype message-switching systems operated by large organizations and telegraphic common carriers, such as Western Union and Cable & Wireless. Special-purpose real time computers replaced the manual "torn-tape" systems previously in use at switching centers. These computers handle storing and forwarding of messages, detection of format and transmission errors, and auxiliary functions, such as accounting, message logging, and broadcasting.

Decentralization of Switching Centers

Message-switching systems based around a single switching center are limited in their capacity and are obviously highly vulnerable to failures. Large systems typically employ a network of interconnected centers; messages are queued and forwarded through a chain of these centers from source to destination, as illustrated in Figure 3-7. Geographic distribution of the centers is encouraged by the fact that in large systems the message-switching load is generally highly localized, so that with multiple centers a majority of the traffic can be handled at the local level.

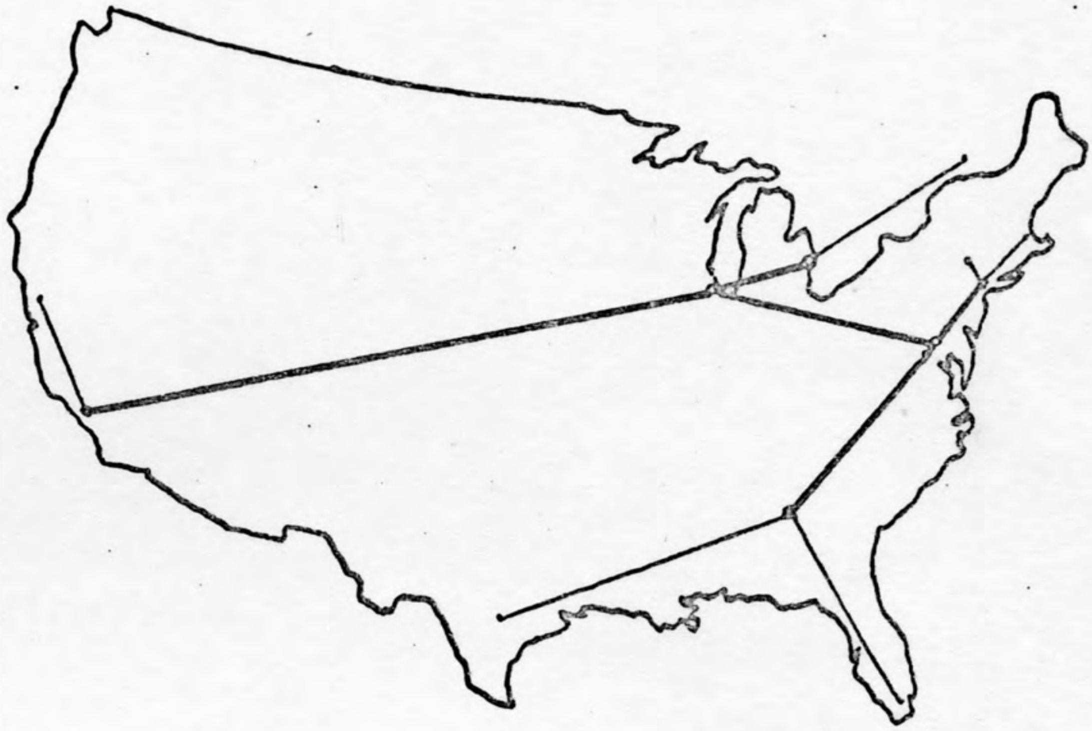

A decentralized message-switching network such as might be set up by a large nation-wide organization.

Figure 3-7 A Decentralized Message Switching Network

Distributed Communications

The vulnerability of a message-switching system to line and node failures can be greatly reduced by over-connecting the network--i.e., by providing multiple routes for transmission between any pair of points (see Figure 3-8). The RAND Corporation carried out a detailed feasibility study of this form of system in the early sixties [Baran 64a]. The principles established in this study have had a strong influence on subsequent designs.

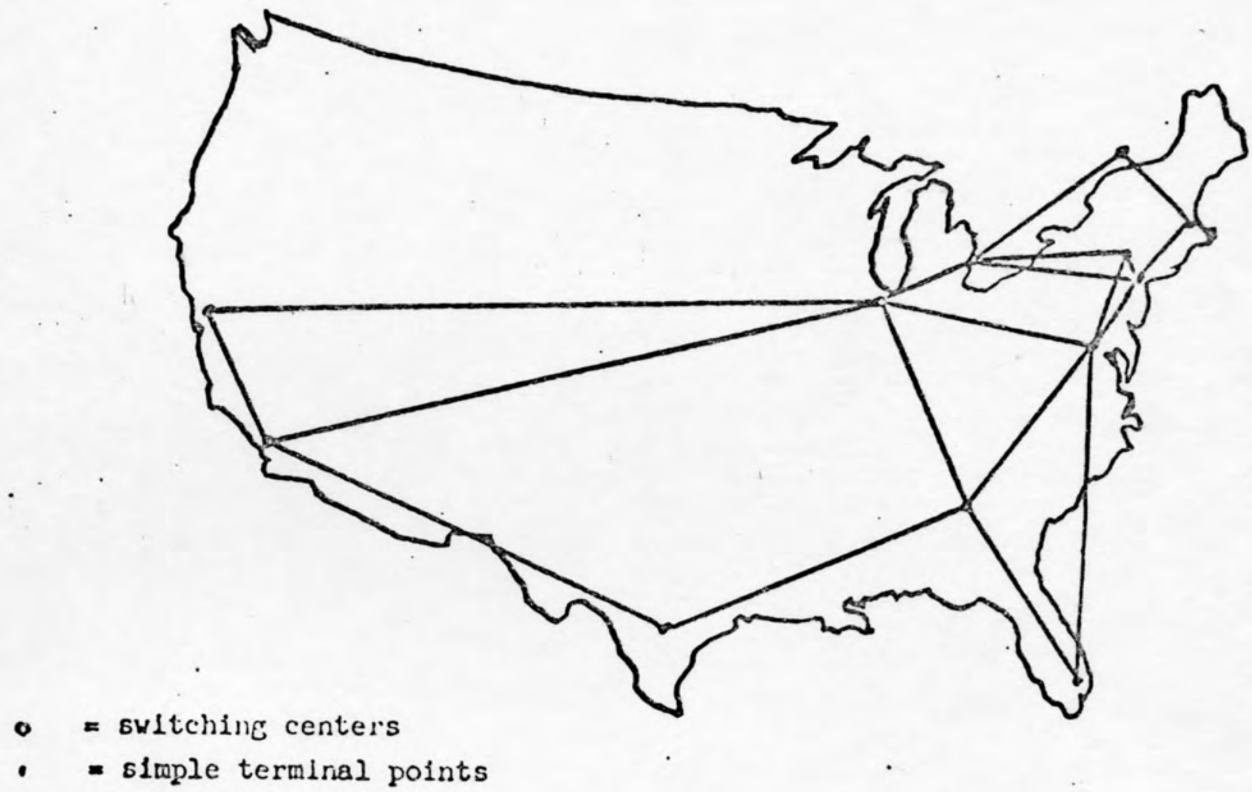

A fully distributed network, employing redundant switching centers and interconnections. The provision of multiple paths between any pair of points minimizes the effects of failure of a node or interconnection.

Figure 3-8 A Distributed Message Switching Network

As examples of major message-switching systems we may cite the common-carrier systems operated by Western Union, ITT, Cable & Wireless. ARINC (Aeronautical Radio Inc.), which is jointly owned by a consortium of U.S. internal airlines, operates a common message-switching system to serve their operational requirements; SITA (Societe Internationale de Telecommunications Aeronautiques) provides similar facilities for the European airlines. U.S. military requirements are handled by a fully integrated world-wide system, AUTODIN. Perhaps the best publicized is the NASA Communications Network (NASCOM) established to interconnect the various ground centers and the space craft for the Apollo program.

3.3.2 Command and Control Systems

The development of revolutionary new weapon systems in the nineteen-fifties and sixties created a need for equally revolutionary changes in the techniques used to plan and control military operations. The speed and power of supersonic bombers and missile systems required much faster response from defense systems; the vast number of weapon systems employed to handle any given situation demanded a far higher degree of coordination between operational units; the volatility of the world situation and the dangers of escalation inherent in any military action required much faster decision-making and tighter control from higher command levels. This demanded facilities for rapid transmission, processing and display of information; the development of real time command and control systems to provide these facilities has been at the forefront of the development of real time systems in general.

The first such system was SAGE, built to coordinate the components of the U.S. air defense system--radar, anti-aircraft missiles, interceptors, etc. It replaced a manual system that had become too slow and localized to cope with nuclear delivery systems. By correlating information from radar scanners, ground stations, and aircraft, SAGE maintained a complete picture of air and ground situations. Operations such as missile firing and guidance were handled automatically, and higher, manual levels of control were supported through real time displays of the air and weapons situations. The system was based on computers at some thirty direction centers, each gathering and maintaining information on a local sector. Each center communicated with those in adjacent sectors to coordinate tracking and interception along the path of an aircraft or missile, and also provided information to a hierarchy of interlinked computers which filtered information upward to the various levels of command.

Since SAGE a number of real time command and control systems has been developed to handle various portions of the defense complex. SAGE has been superseded by a more sophisticated version, NORAD. Other similar systems include the Naval Tactical Data System (NTDS) and the Strategic Air Command's system.

The most ambitious project in this field is the proposed formation of an integrated World-Wide Military Command and Control System (WWMCCS). This involves interconnection, through AUTODIN, of the various existing military command systems, plus certain other military and civil information systems. This would facilitate coordination of activities over the entire defense complex, and provide integrated real time command and control facilities right up to the presidential level.

Apart from the military sphere, similar systems can be expected to emerge wherever there is a need for real time control of complex and geographically or organizationally decentralized activities. Air traffic control constitutes a current example of this.

3.3.3 Travel Services Networks

Real Time Reservation Systems

Real time reservation systems support a wide range of reservation services offered to the general public, including seat reservations for air and rail travel, and booking services for associated hotel and car rental requirements.

The airlines led this field with seat reservation systems such as American Airlines' SABRE, Pan Am's PANAMAC and B.O.A.C.'s BOADICEA. Such a system performs the airline's sales entry and seat inventory control functions in real time. It enables its sales offices to give immediate confirmation of a reservation, as well as automating functions such as waitlisting and checking for confirmations. In the more advanced systems the sales and inventory control functions have been linked to other operational support systems such as accounting and check-in processing, and to higher-level functions such as flight and crew scheduling.

Many hotel chains and consortia have for similar motives set up systems which maintain a central real time data base giving room availability, so that a traveller may reserve a room at any hotel in the group, with instant confirmation. Similarly, the national and international car rental agencies such as Hertz and Avis have centralized systems that provide for booking of a rental car at a remote destination.

Interconnection of Reservation Systems

All of the above services are connected with the travel industry, so that any given customer is quite likely to require several of them at the same time. Thus, a traveller might wish to make reservations for a journey involving flight segments with several different airlines, hotel stays at various stopping points, and perhaps car rental and entertainment tickets as well. By linking several independent systems into a data base sharing network, each participating system becomes able to provide the full range of services offered by the consortium.

Just as the airlines were among the leaders in the development of real time systems, so they have also been pioneers in the networking of these systems. A large proportion of airline business consists of "inter-line" bookings--that is, bookings that an airline passes on to another for flight segments not covered by its own services. The ability to confirm inter-line segments along with its own is clearly important to an airline, particularly one with a limited set of routes, and most airlines have direct inter-line booking arrangements with others whose route structures complement their own. With the airlines concerned using real time reservation systems these arrangements are generally reflected in inter-system communications links over which inter-line reservations can be handled automatically. As an example, we have the following account of the inter-line functions of Eastern Airlines' system:

> "The system automatically generates messages to sell or request seats on other airlines. Eastern maintains availability on 24 other airlines, accommodating more than 90% of the requests we receive. Also, we supply many other airlines with Eastern's flight availability. Messages selling or requesting seats received from the other airlines are automatically processed and replies sent out if necessary. Messages are automatically sent to other airlines advising of schedule changes. A reply of seat confirmation or advice from another airline of a schedule change is automatically processed and placed on special queues alerting the agent to notify the passengers. All such messages (an estimated 37 million in 1969) were previously handled manually." [Datamation, March 1969, p. 32]

Inter-system links have also been set up for other purposes. For example, Eastern Airlines' system maintains a lost-baggage file that is common to all the airlines. The major reservations systems are linked into Eastern's system to provide them with on-line access to the file.

In addition to inter-airline links, many airlines have established similar links with hotel booking and car rental systems, enabling them to offer real time hotel and car reservation services directly through their agents' terminals. This sort of inter-service arrangement is spreading rapidly, again with the links tending to be arranged individually between pairs of operators.

Travel Agent Networks

Travel agents are naturally anxious to obtain the advantages of on-line access to the various real time reservation systems. In recent years several proposals have been developed for network systems that would provide travel agents with such a service. The general form of these proposals is shown in Figure 3-9. Terminals in travel agents' offices are connected with a central installation which is in turn linked to the systems of participating airlines, hotel groups, etc. The central installation acts in effect as another reservation system in the inter-line network, except that it does not offer any reservation inventory itself. It holds availability data obtained from the participating systems. It generates and holds PNR's (passenger name records--the basic data unit associated with a booking) for bookings made from travel agents' terminals, carries out service functions of such a billing and notification of schedule changes, and distributes sales notifications or requests as appropriate to all the operators involved in a given booking.

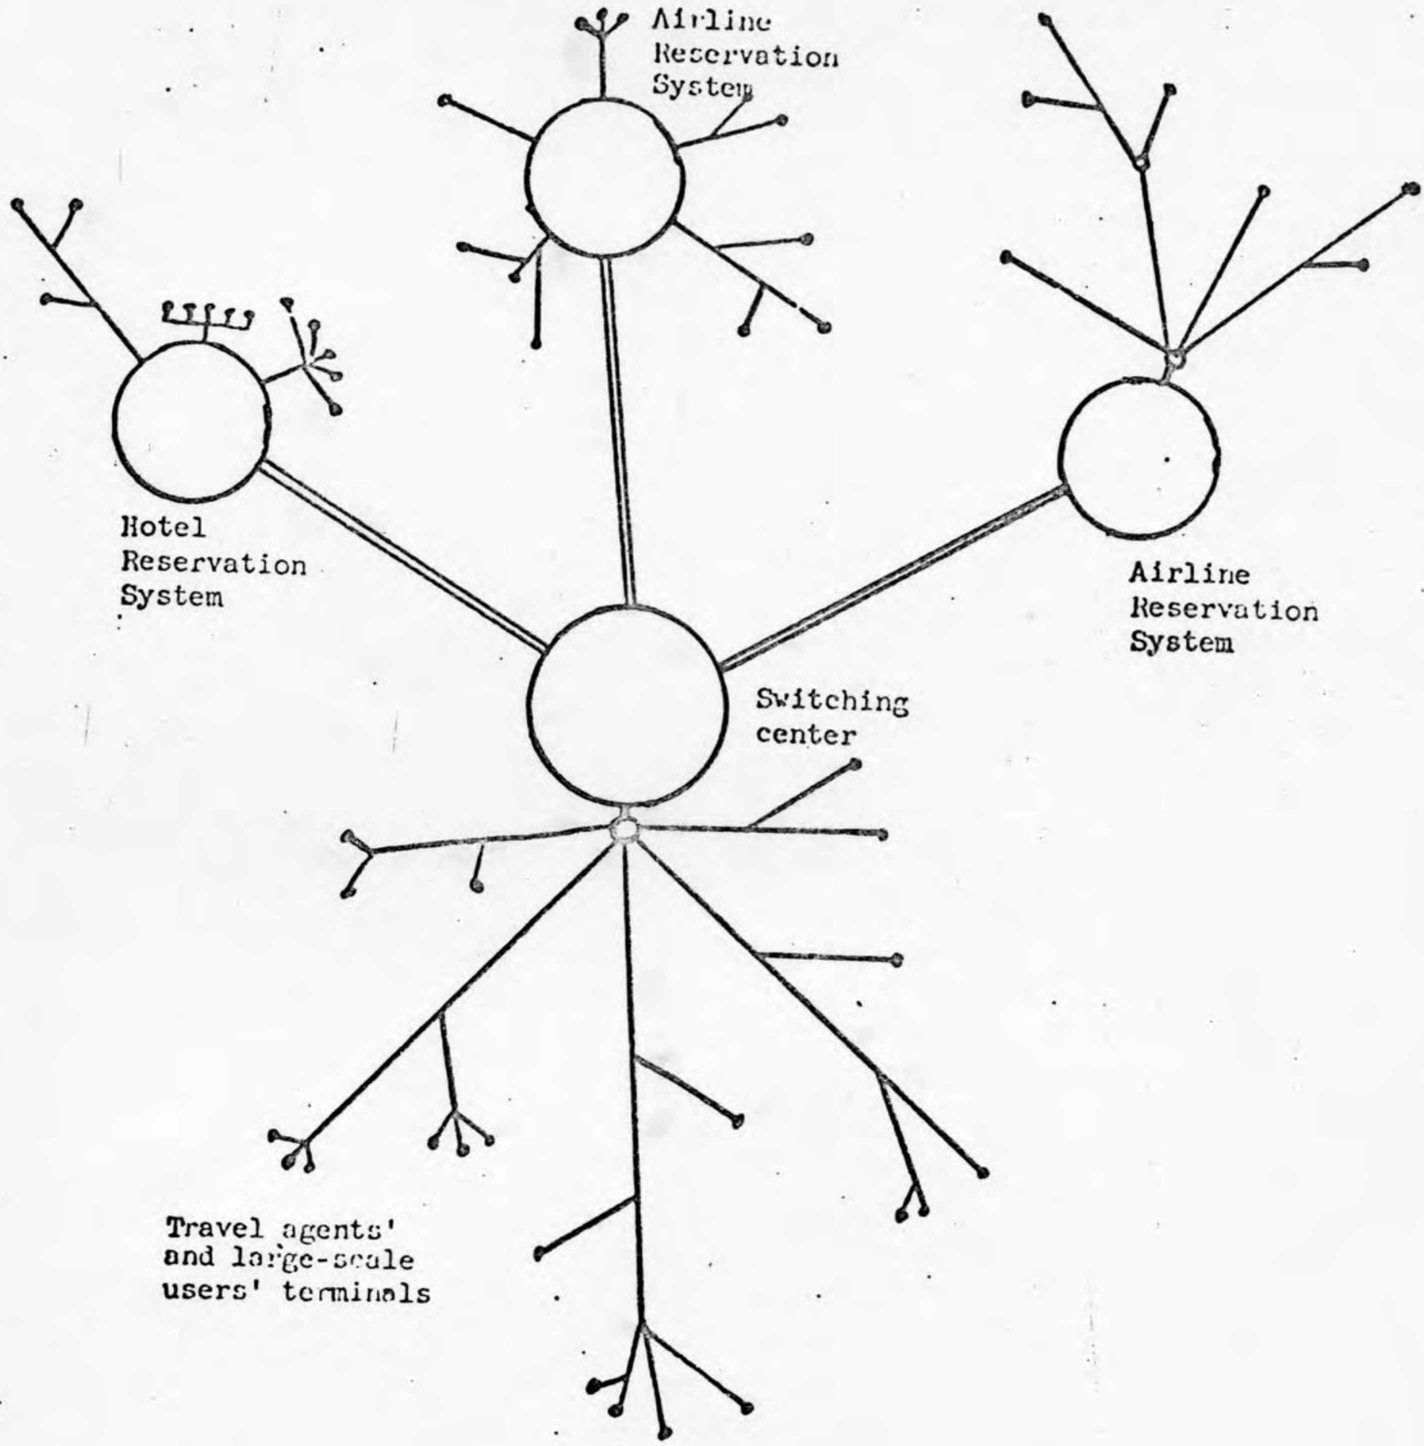

A unified approach to reservation system networking, providing inter-line facilities and common booking facilities for travel agents.

Figure 3-9 A Travel Agent Network

The difficulties facing this kind of potential system are not so much technical as organizational and political. The pairwise links involve straightforward agreements between two cooperating operators, and they require relatively small investments of effort and capital. On the other hand, integrated networks of the kind described above are major projects requiring continuing cooperation and commitment from many independent--and often competing--operators.

3.3.4 Electronic Payments Systems

Electronic payments systems, also referred to under labels such as "paperless exchange," "electronic funds transfer systems," and "the cashless society," are computer-communications based money transfer mechanisms that seem likely to replace cash and checks in many sectors of the economy. The data processing literature abounds with scenarios envisaging fully integrated money handling centers operated by, or working in close conjunction with the banking system. One such example foresees the emergence of a new form of computer-based utility, "SAVE," which would perform, on-line, the functions of present-day clearing houses, credit card companies, finance houses, etc.

Future Developments

The future development of electronic payments systems is difficult to predict, but it obviously will involve a growing network of links between the computers of banks, large corporations, government agencies, institutions such as credit card bureaus and finance houses, and new centers such as the automated clearing houses. What is not clear is the extent to which these links will be based on data communications rather than tape transfer, and will operate on a transaction basis rather than in batch mode. In addition to the economic considerations involved, the problems of providing security and hard-copy backup will severely constrain the technology which can be used. The ultimate form of the network is likely to be determined not so much by what constitutes an optimal solution by technical or economic criteria, but rather by what is feasible given the organizational and political context.

3.3.5 Corporate Information Networks

The typical large corporation operates several data processing centers on a divisional or regional basis. These centers might be linked together for any of the motives outlined earlier: to permit load sharing and provide back-up, to link the systems of headquarters and local operating units access to a centrally held data base, to give headquarters access to operating files, etc.

Forms of Corporate Network

In medium-sized companies a typical arrangement can be described as follows: The organization has a medium-scale computer at a central data center serving the largest plant and also corporate headquarters. Each of the ten plants has a satellite machine, permanently connected by data transmission line to the center. In the larger plants these satellites are small computer systems capable of performing most applications independently. In the smaller plants the satellites are remote batch terminals, used only for transmission of batches of data to the central system for processing.

This semi-centralization permits some local control of, and specialization in, data processing facilities at the local level, while providing all units with the facilities offered by a large, general-purpose machine.

The example above illustrates some reasons why data processing may be deliberately distributed around a network as an alternative to a single central installation; the motives for doing this were discussed more fully earlier.

Functionally Specialized Networks

Some people have recently been advocating an extreme form of distribution, in which data processing functions are split among a network of specialized mini-computers. They say that there should be separately sized systems for headquarters, regional offices, and local installations such as airports. Likewise, the many functional areas should have separate dedicated systems, e.g., accounting, advanced models for simulation, random-access systems as opposed to static displays, etc. The headquarters system should be a command and control system with lateral interfaces to all other parallel systems, and vertical interfaces to remote systems and to functionally dedicated systems.

One major advantage of multiple systems is the opportunity to take advantage of the best available components to meet the unique requirements of each functional or decision level.

Inter-Corporate Networks

Where the operations of two or more firms are closely integrated, we may expect to see links formed between their information systems to aid coordination of their activities. The airline networks described above illustrate this point. Another example is provided by the Blue Cross-Blue Shield group, which consists of autonomous but closely cooperating companies operating in each state. BC-BS has the problem of handling claims made in one state on a policy held in another. The volume of these claims is quite high, since they originate not only from cases where an individual requires treatment while away from his home state, but also from individuals insured under group policies that are held in one state but whose beneficiaries are distributed nation-wide. To deal with these claims, BC-BS seems likely to set up a network of links between the data processing centers in each state.

3.3.6 Government Information Networks

Government uses data processing for much the same purposes as the private sector, and to this extent one may expect network formation to occur for the same broad range of motives. For example, the Army Materiel Command has formed a network primarily for load sharing purposes.

However, a particularly significant activity of government at all levels is the acquisition and distribution of information, both for its own use, and the use of organizations and individuals in the private sector. As one might expect, data sharing provides the major thrust behind networking in the government area. Networks can be expected to occur when there are advantages to be gained from pooling the files maintained by different governmental agencies. They may also be found within a single agency, in cases where there are advantages to distributing a data base around a network of regional centers rather than centralizing it.

National Security Network

As an example of the former case, a network has been formed to link the information retrieval systems operated by various government security agencies. The network provides terminals associated with each information system with on-line access to data-bases held by the others.

Criminal Justice Information Network

A further example is provided by the emerging criminal justice information network. Law enforcement agencies now rely heavily on on-line computer systems for communication and information retrieval. Numerous efforts are now under way to interconnect these systems at the local, state, and national levels, to facilitate exchange of information and sharing of data bases.

Health Information Network

A potential example of a regionally distributed data base occurs in the area of health information. At present, individual medical records are maintained by hospitals, with little coordination and considerable duplication. Similar information-handling problems exist with other health-related information. There would be great advantages to centralizing health records at a regional center, with on-line access provided to local hospitals, etc. Interlinking of such regional centers would probably be justified -- for example, to provide access to individual records held in one region from a location in another, or to allow transmission of data to a central location for statistical analysis.

Distributed vs. Centralized Data Bases

Where government files hold information on individual citizens problems of security and privacy are of central concern. In this context, distribution of files within a network offers certain advantages over a centralized data base.

3.3.7 Utility Networks

A computer utility offers data processing services on a commercial basis. Users are generally linked on-line to the utility which may offer interactive time-sharing or remote batch services, or both. In some cases the utility simply provides the user with raw computer power, supported by a suitable range of languages and other software. In others, it may also offer specialized services aimed at specific applications (e.g. payroll, or fuel oil delivery scheduling), or may provide services to construct software that meets a user's particular needs.

Utility Chains

Many computer utility operators are based on a "chain" of computer centers located at major business centers nation-wide or world-wide. These chains clearly possess major advantages of scale over the individual center: multiple centers allow load-levelling, economies of scale in software development, justification for specialized resources, etc.

Utility Networks

Realization of many of the advantages of operating a chain -- for example, load levelling -- requires communications links between the centers. In addition, any utility has to link the user terminals into its centers, and for a chain this is most satisfactorily done by setting up its own network of lines, concentrators, and switching points. Thus, many of the larger, more advanced chains have embedded their computer centers in a communications network that provides for both terminal-to-center and center-to-center linkage. Examples of such systems are Computer Science Corporation's INFONET [Schwartz 72], G. E. Information Services' network [Hench 72], [Schwartz 72], and CDC's CYBERNET [Luther 72]. The design aims of these systems are broadly similar. The communications system allows a user to be switched into any center, so that he can be offered specialized facilities available at remote centers and provided with back-up in case of local failure. The inter-center links also allow load-levelling. In some cases the user can transmit data files between centers; thus, geographically separated users can interface through the network.

CYBERNET

CYBERNET is a distributed network consisting of CDC machines such as 6600s and 3300s linked by wide-band and voice-band lines (see Figure 3-10). 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 the system. It is a message-switching network, utilizing a broad spectrum of switched. leased, and satellite communications facilities, running at rates from 100 baud to 40 Kbs. It counts heavily on hand-established connections for terminal-to-computer and computer-to-computer links; the network cannot reconfigure itself. Alternative paths do exist in some cases between nodes and centroids, but in general a link failure will necessitate human intervention. Messages in the network are a fixed length, 1024 characters.

Figure 3-10 CYBERNET

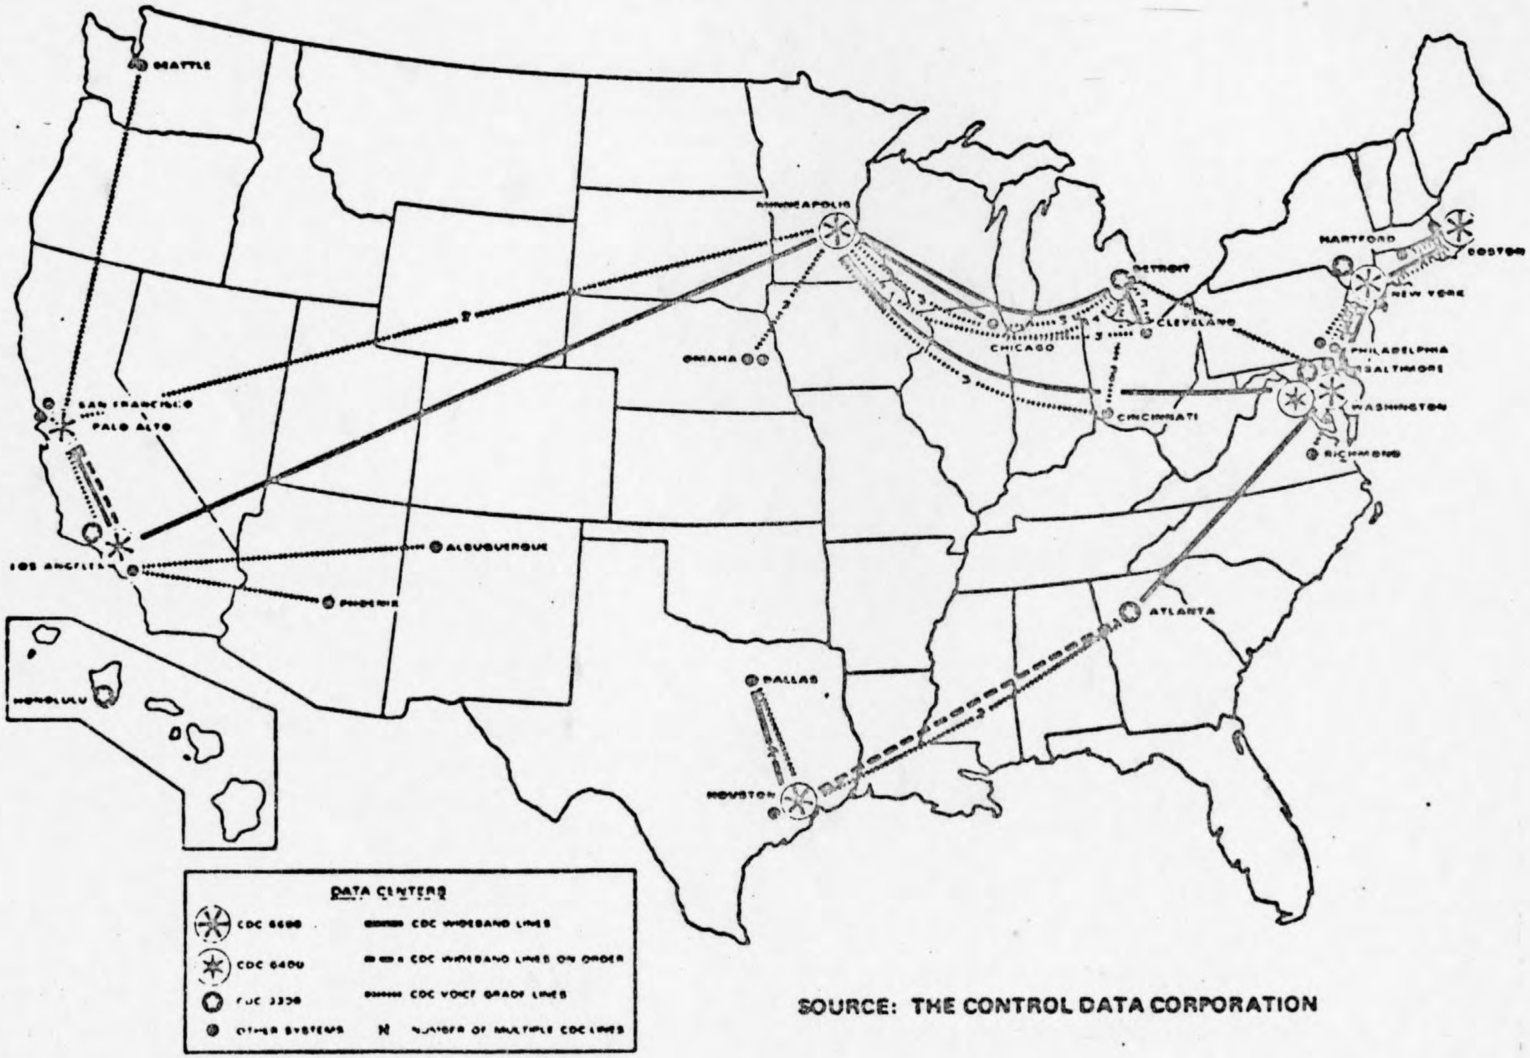

SOURCE: THE CONTROL DATA CORPORATION

The CYBERNET Network

3.3.8 Research Networks*

The final application of computer network technology that we will consider is academic and institutional research.

The MERIT Network

The MERIT Network, the Michigan Educational Research Information Triad, Inc., is a joint cooperative effort between Michigan State University, Wayne State University, and the University of Michigan. The network geometry is shown in Figure 3-11. 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 baud 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.

Some references for the MERIT system are [Herzog 72], [Aupperle 72], [Aupperle 73], and [Cocanower 72].

----

*The material in this section is based on [Farber 72c].

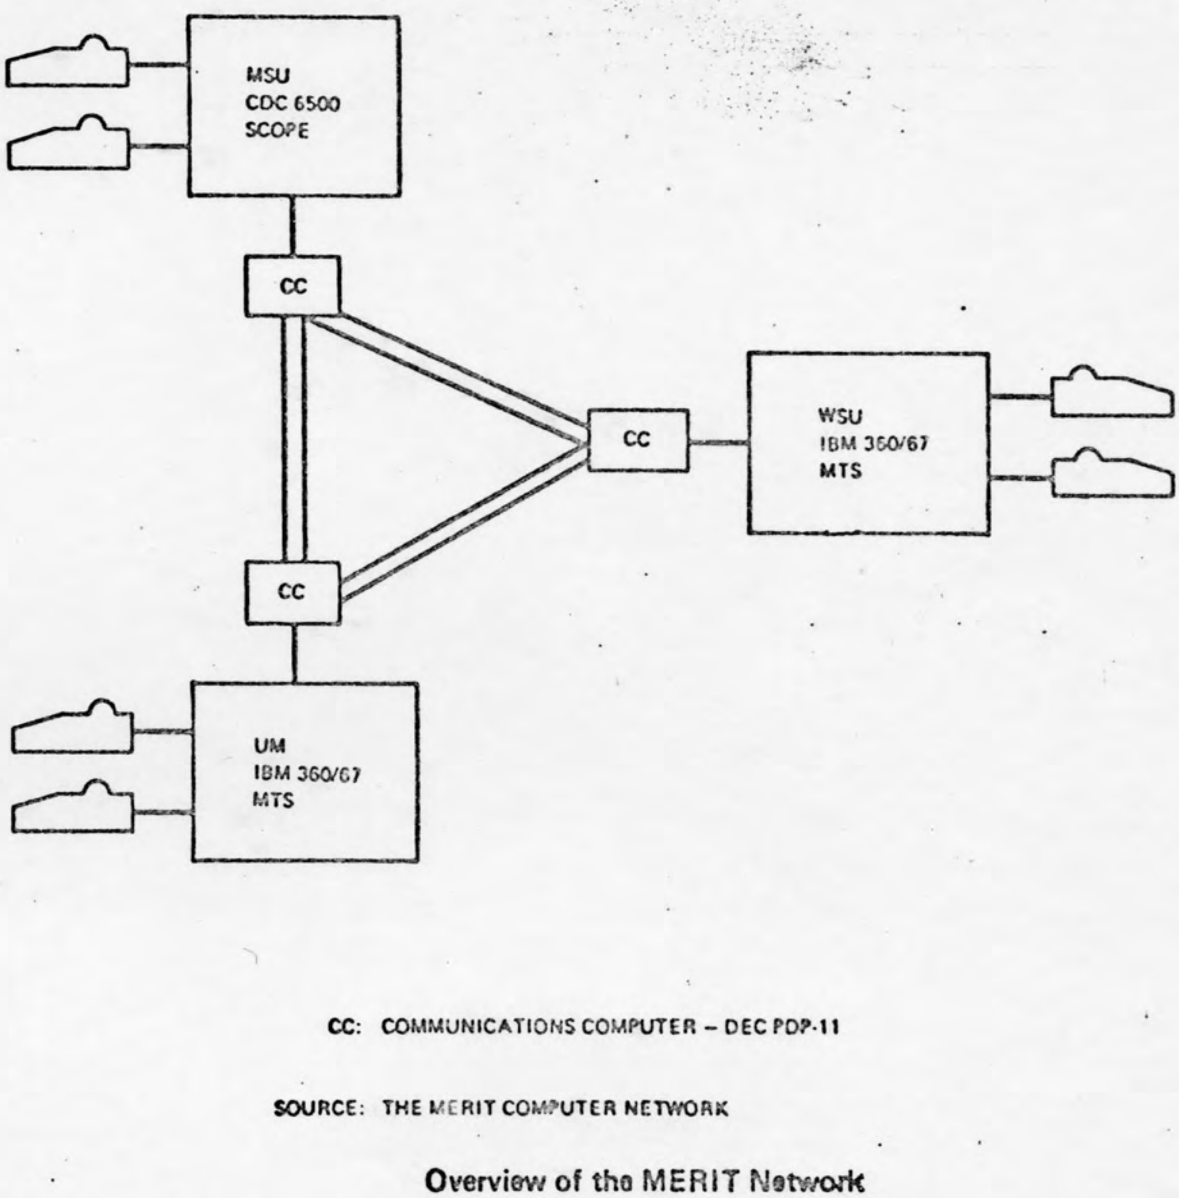

Figure 3-11 The MERIT Network

The OCTOPUS System

The Octopus system is a heterogeneous network developed at the Lawrence Berkeley Laboratory of the University of California. It connects a complement of devices including two CDC 6600s, two CDC 7600s, and a STAR. The network is shown schematically in Figure 3-12. 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-megabaud 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 PDP-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 PDP-10 insures reliability in this centralized subnet. In addition, while the subnets are logically independent, there are cross couplers providing redundant paths in the event of failure.

References to the OCTOPUS network are [Fletcher 73a] and [Fletcher 73b].

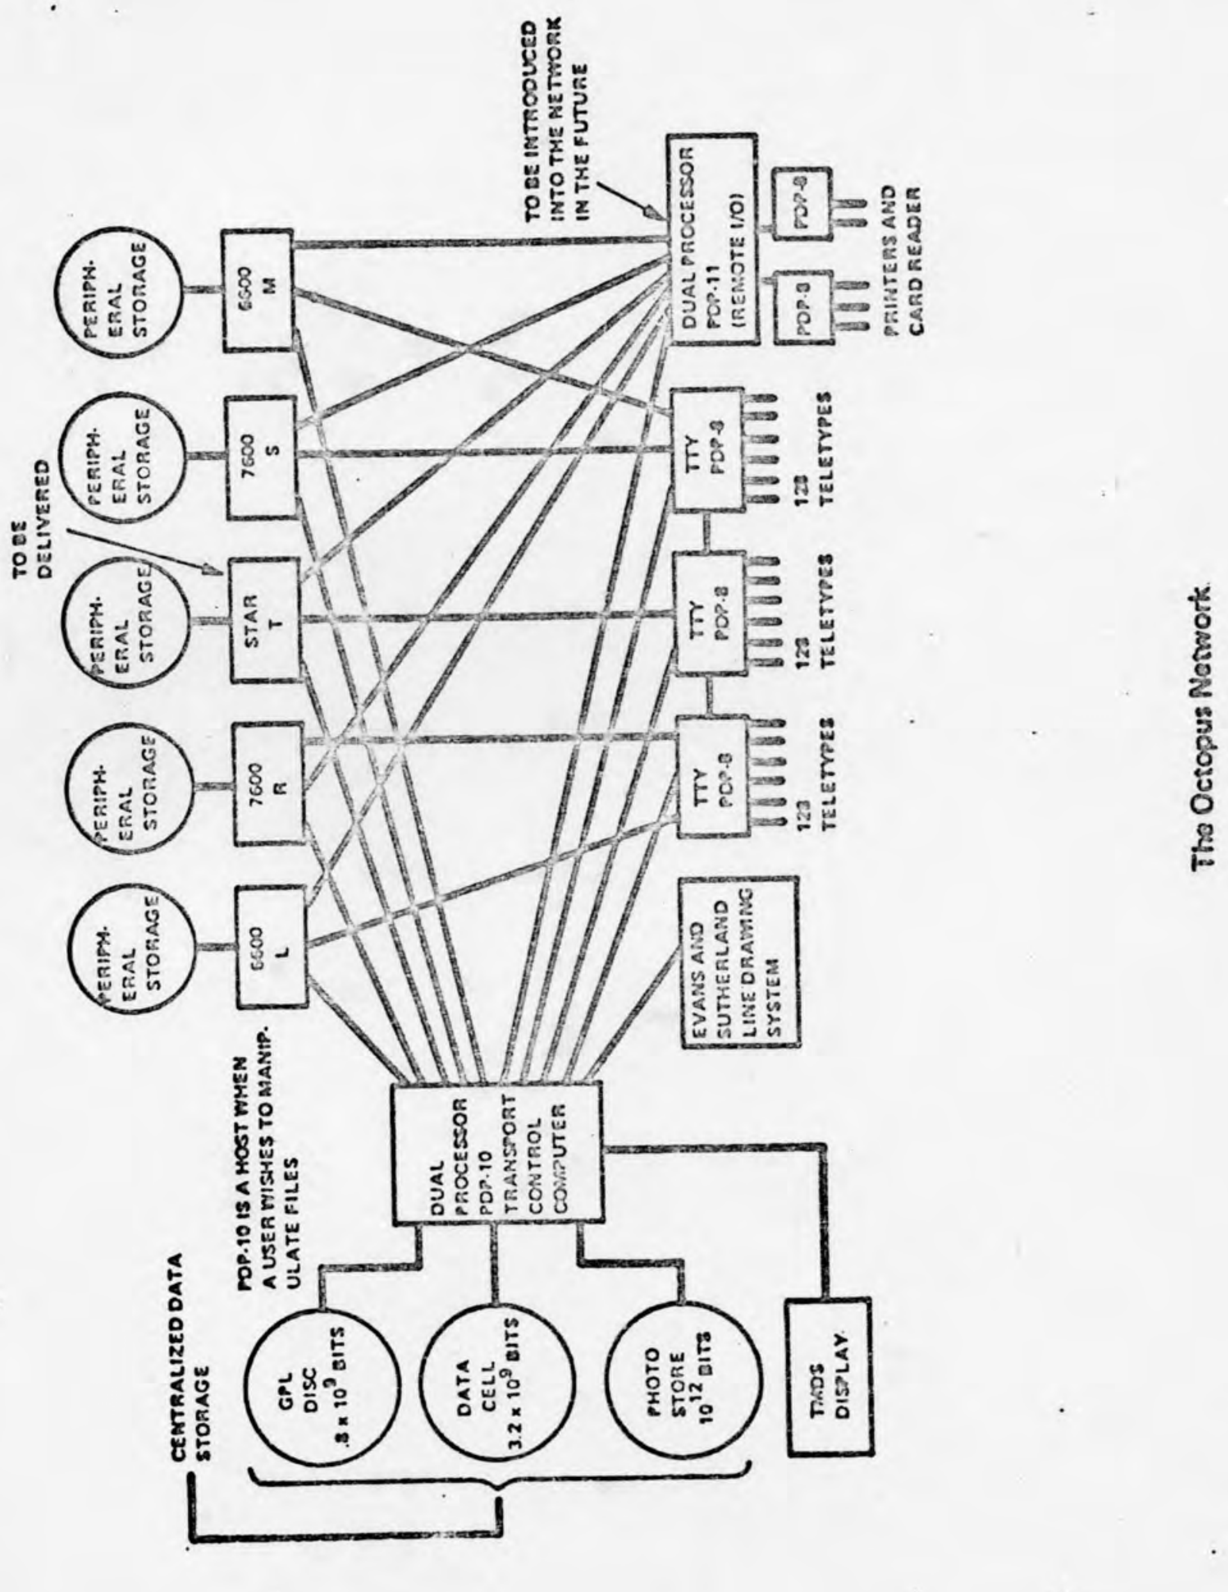

Figure 3-12 The OCTOPUS Network

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, and the network configuration is as shown in Figure 3-13. It is a centralized network of homogeneous machines. At each of three nodes of the network there are IBM 360/40s or 360/50s. 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 baud half duplex line. This line is interfaced to the 360s by means of IBM 2701 Data Adapters. This network is a simple extension to the basic 360 hardware and OS 360 software.

Some references to the TUCC net are [Brooks 68] and [Williams 72].

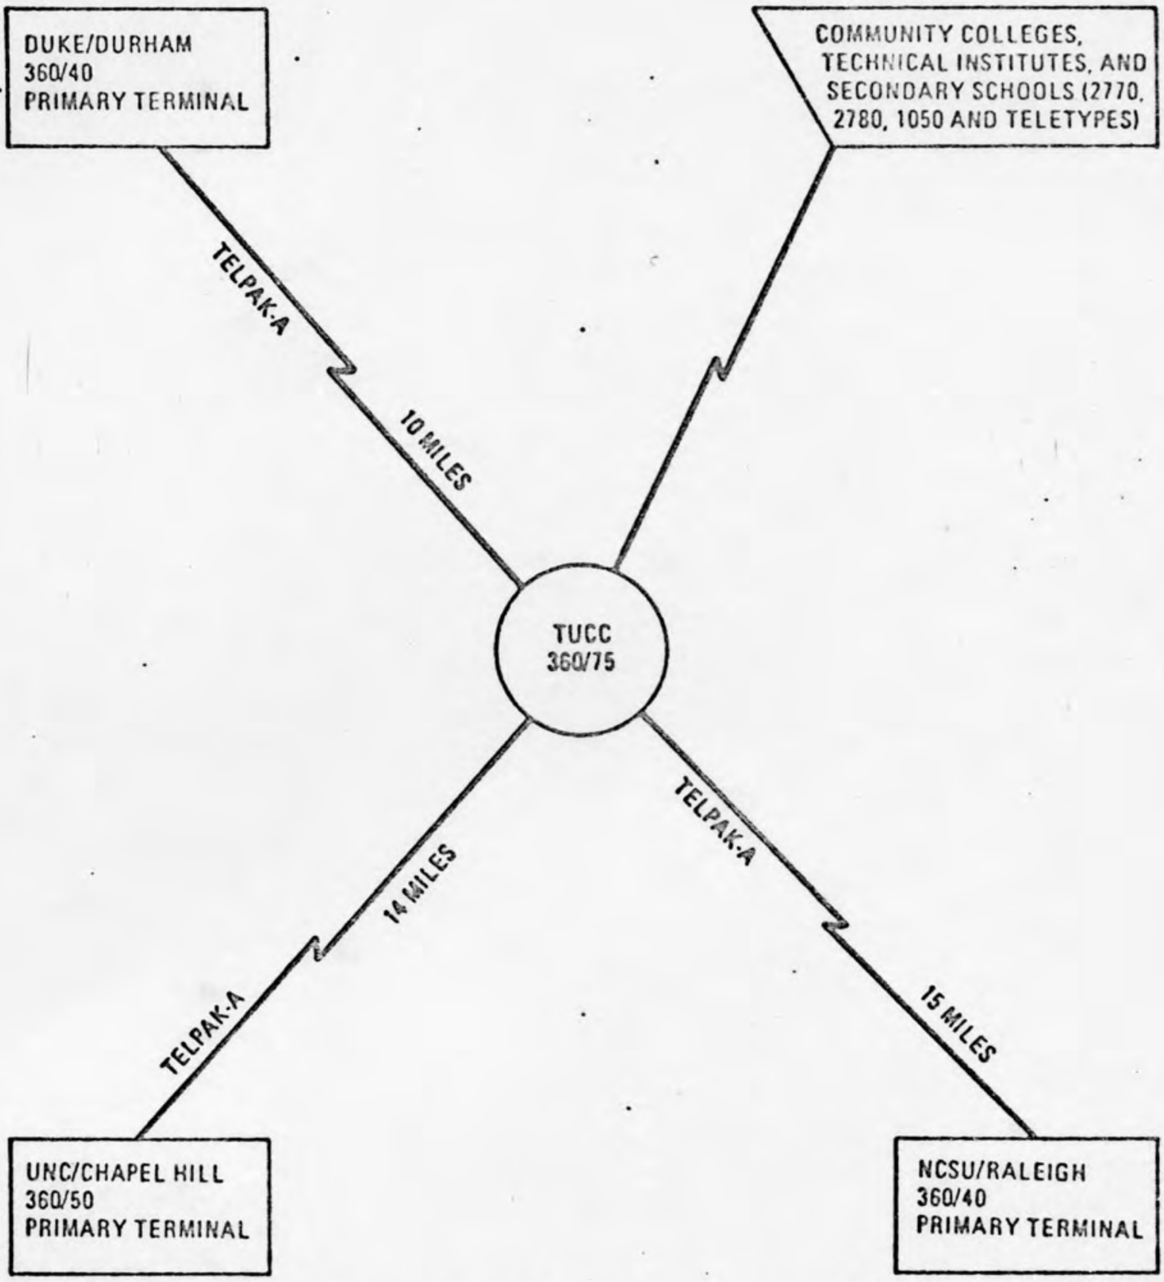

NOTE: IN ADDITION TO THE PRIMARY TERMINAL INSTALLATION AT DUKE, UNC, AND NCSU, EACH CAMPUS HAS AN ARRAY OF TERMINALS INCLUDING 2780, 2741, 1050, 1130 AND TELETYPE TERMINALS DIRECTLY CONNECTED TO THE TUCC 360/75.

An Overview of the TUCC Network

Figure 3-13 The TUCC Network

4. Development of Packet Switching Networks

In this section we survey several networks which have had influence on the development of packet switching.

4.1 Baran

Our survey begins with the extensive paper study done under the leadership of Paul Baran, then of the Rand Corporation, in the early 1960's. Comprising some 11 thick reports, this study had great influence on the later development of packet switching networks. Baran was mostly concerned with improved reliability of distributed networks over centralized networks. In particular, he studied grid networks (such as shown in Figure 4-1) being interested in their survivability and high performance in the face of actual attack. Baran anticipated many of the features later incorporated into actual packet switching networks such as a terminal concentrator device, user-to-user encryption, microwave relays, satellite links, 1000 bit messages, hot potato routing (which is good in the face of massive failures due to the isolated nature of the algorithm), etc. The Baran study demonstrated the feasibility and cost-effectiveness of such networks.

References to the Baran study are the reports themselves [Baran 74a], [Boehm 64], [Smith 64], [Baran 64b], [Baran 64c], [Baran 64d], [Baran 64e], [Baran 64f], [Baran 64g], [Baran 64h], and [Baran 64i]; a brief summary of the Baran study is given in [McQuillan 74a].

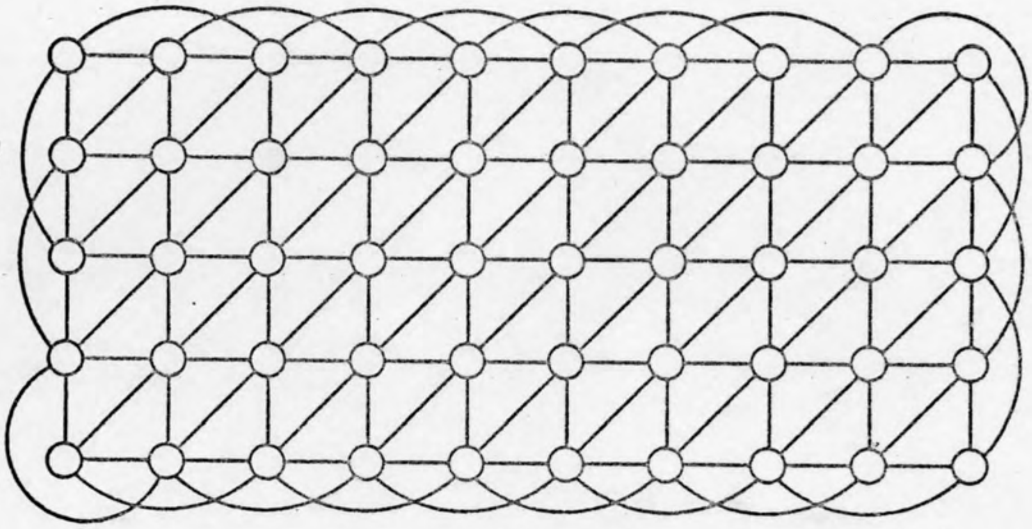

Figure 4-1 Baran Network Design

4.2 Kleinrock

The next significant influence on the development of packet switching networks came in Kleinrock's thesis [Kleinrock 64]. In his thesis, Kleinrock proved convincingly that message delays in very large store-and-forward networks can be made very low and thus put to rest one of the early and persistent objections to networking.

McQuillan's thesis [McQuillan 74a] contains a summarization of this fundamental result of Kleinrock's.

4.3 The National Physical Laboratory Network

The National Physical Laboratory Network (NPL Network) was the next important step in the development of packet switching networks. Largely the brainchild of Donald Davies, it unfortunately never really got off the ground as far as implementation goes. Nonetheless, numerous reports and theoretical studies have resulted which have had great influence.

The NPL Network was to have two parts: a high level network and local distribution centers. These are shown in Figure 4-2 with the high level trunk network shown in the center circle, the local distribution centers shown with the circle-I's and the terminals and computers (little-T's and circle-C's) attached to them.

Great concern was taken with a standard interface to the terminals, and a few such interfaces and one local distribution center were actually constructed. (The standard interface has the disadvantage that no commercially manufactured terminal has such an interface.)

Another main thrust of the NPL Network has been "isarithmic" control of congestion, in which the number of packets in the network is held constant. This will be discussed again later.

One possible high level network is shown in Figure 4-3.

References to the NPL Network may be found in Davies' and Barber's book [Davies 73] and a summary in [McQuillan 74a]; a reference on isarithmic control is [Davies 71].

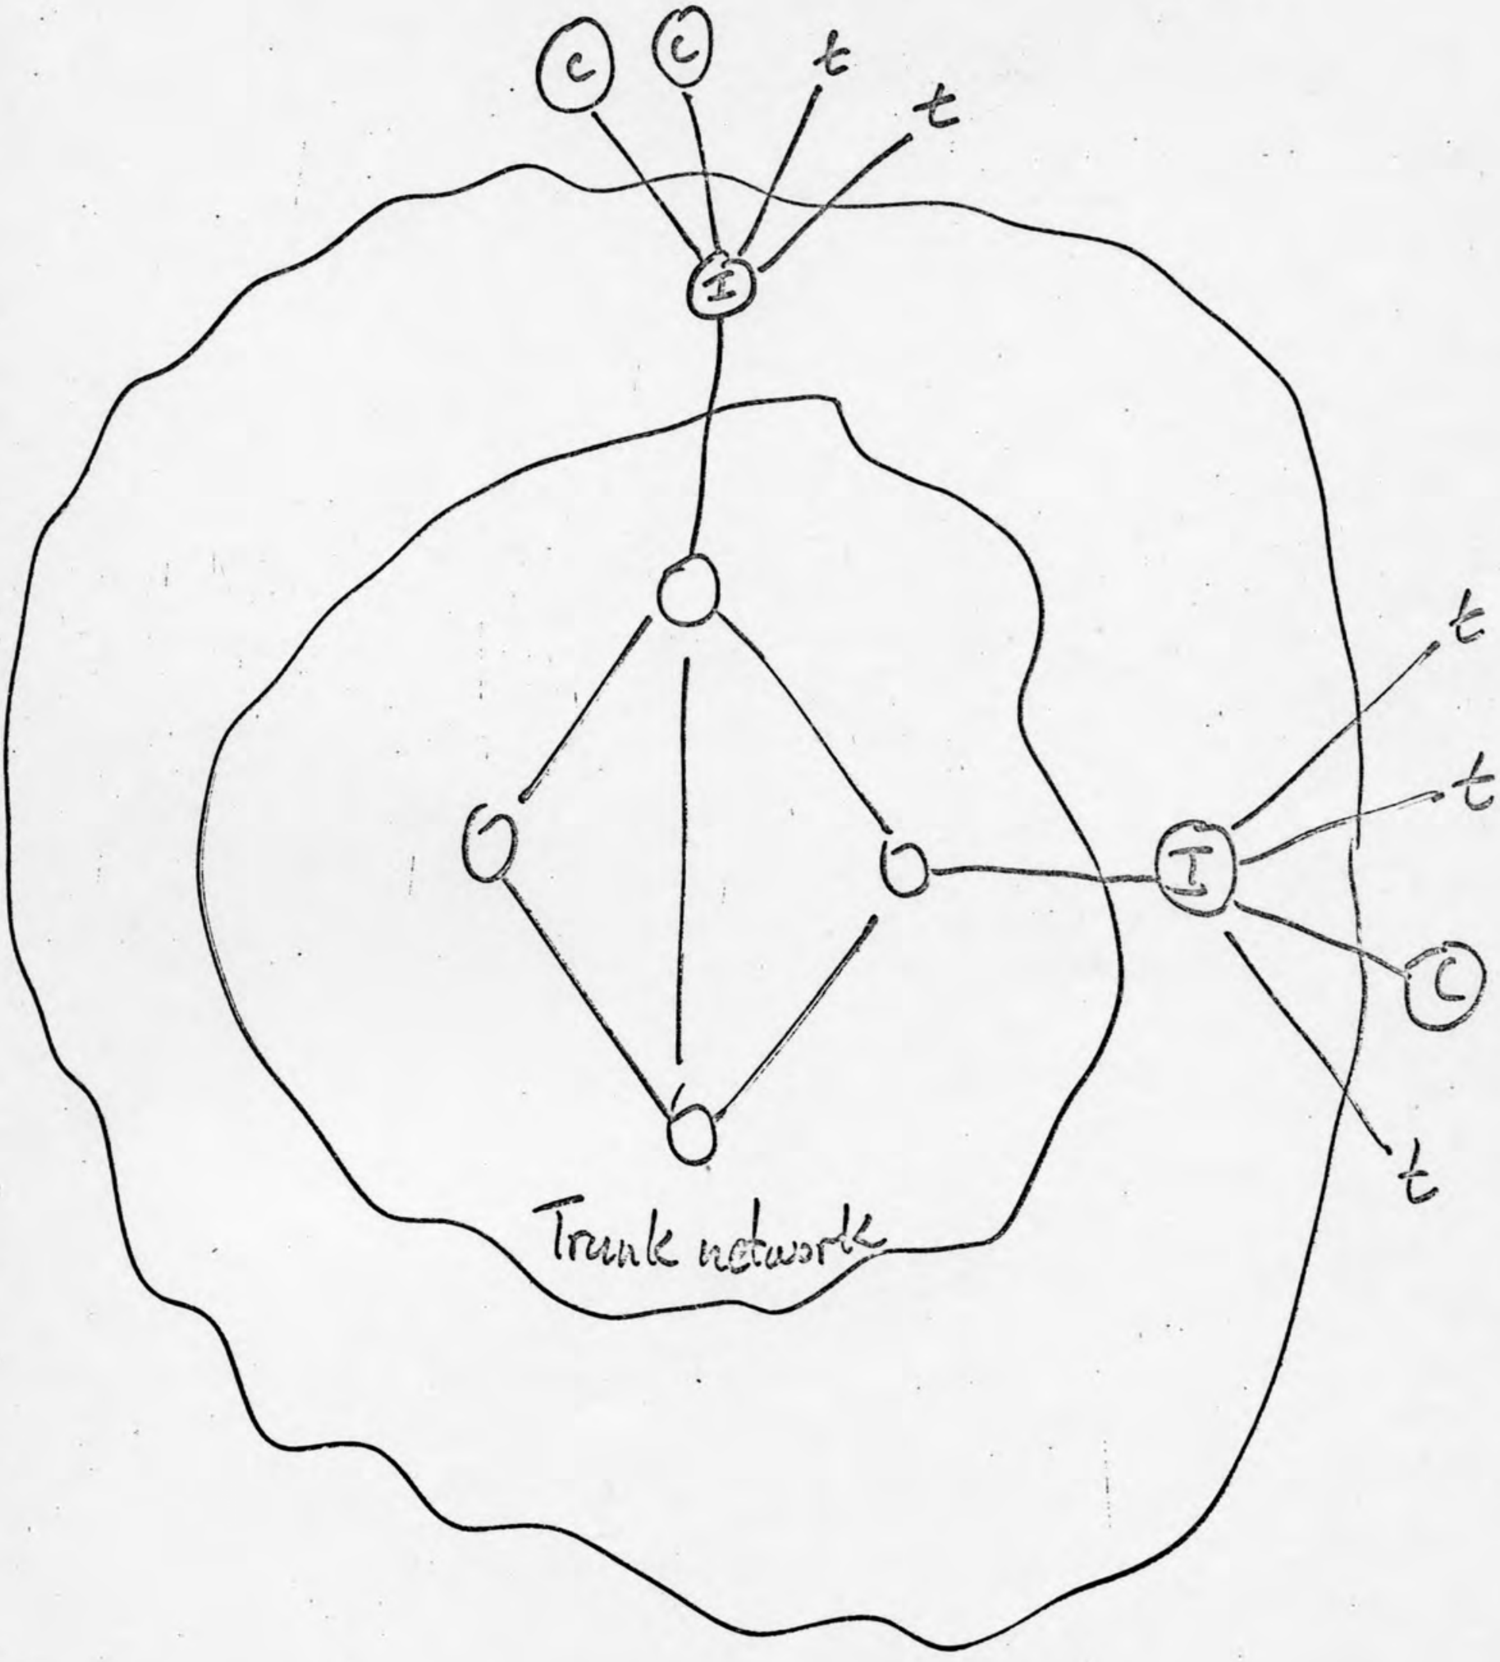

Figure 4-2 Structure of the NPL Network

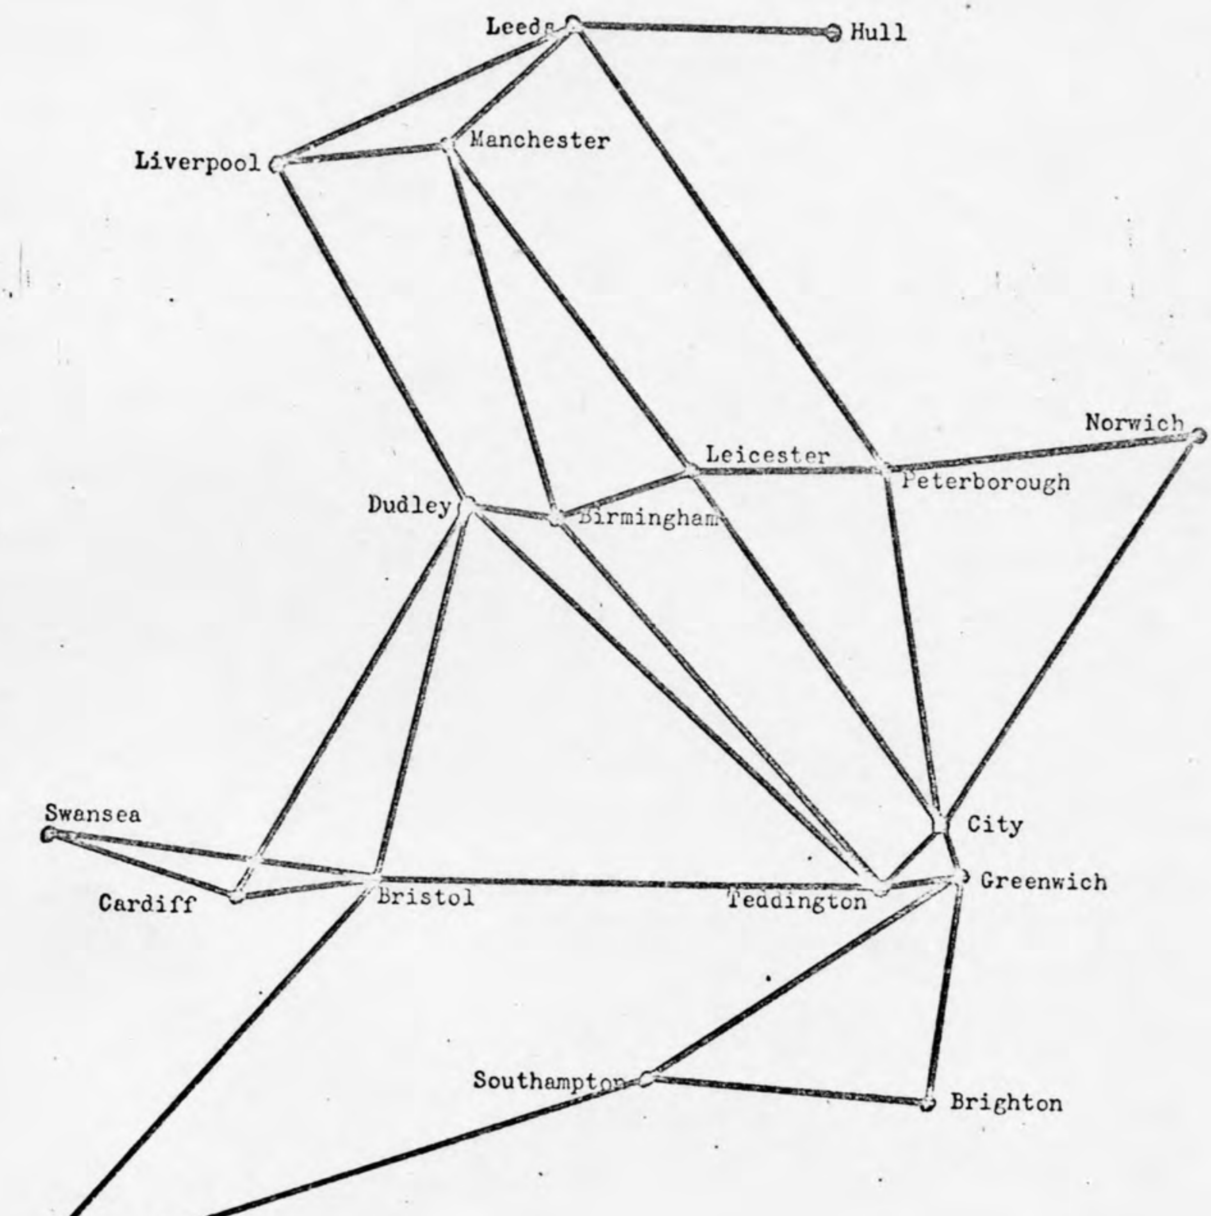

Figure 4-3 Example NPL High Level Network

4.4 The SITA Network

Begun in 1949, the Societe Internationale de Telecommunications Aeronautique (SITA) Network was originally meant to be a low speed message switching network able to aid with ticket sales and to move operational information. It was implemented well in time to be available when the first computer based airline reservation systems came into being.

In 1964 the SITA Network was reappraised and it was decided to adapt it to handle three kinds of traffic: a) single address, inter-computer, rapid response traffic using 5, 6, or 7 bit codes; b) conventional teleprinter traffic using CCITT 5 or 7 codes with single or multiple address traffic up to 4000 characters in length with three grades of delivery (within 2 minutes, within 30 minutes, and within 12 hours); and c) single address traffic using CCITT codes 2 or 5 with the same three delivery priorities as for b). At the time of this decision to upgrade the SITA Network, it had 100 centers and was handling 100 million telegrams per year.

The design for the expanded network involved a high-level network with each center having regional responsibility. Twenty-four hundred baud lines were used with store-and-forward blocks of variable size to get the necessary response time and efficiency. The block size chosen was 256 characters (as most SITA traffic was under 256 characters in length). Strangely, different computers were used at the different regional centers. The high-level network was completed in 1970. In parallel the regional centers were upgraded. The network in 1970 is shown in Figure 4-4.

A little more detail about the design of the SITA Network will be interesting. The blocks were checksummed on the links using a double parity scheme. Acknowledgments and negative acknowledgments were used to indicate the correct and incorrect reception of blocks on the line. In the absence of real traffic, dummy traffic is sent at least every three seconds on the links to test their operability. If three tries doesn't get a block of traffic (or dummy traffic) through, the line is declared dead and only check messages are sent. If a link is determined to be fault, a status message is sent to all centers which causes the routing paths to be changed. If a center is disconnected due to faulty links, low priority messages are buffered while other traffic is discarded.

Source to destination ordering of traffic is done with a destination to source acknowledgment being sent for at least every 16 blocks. Duplicate detection is also provided.

The most stunning thing about the SITA Network is its great similarity to the ARPA Network, which was built without benefit of study of the prior SITA work.

References to the SITA Network are [Brant 72] and the excellent overview in [Davies 73].

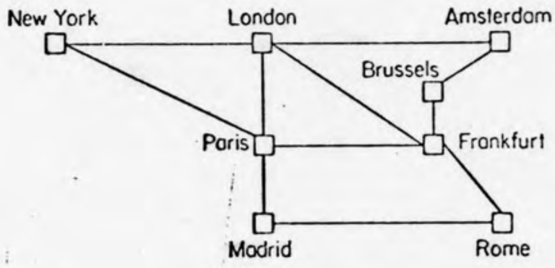

The SITA high-level network, 1970

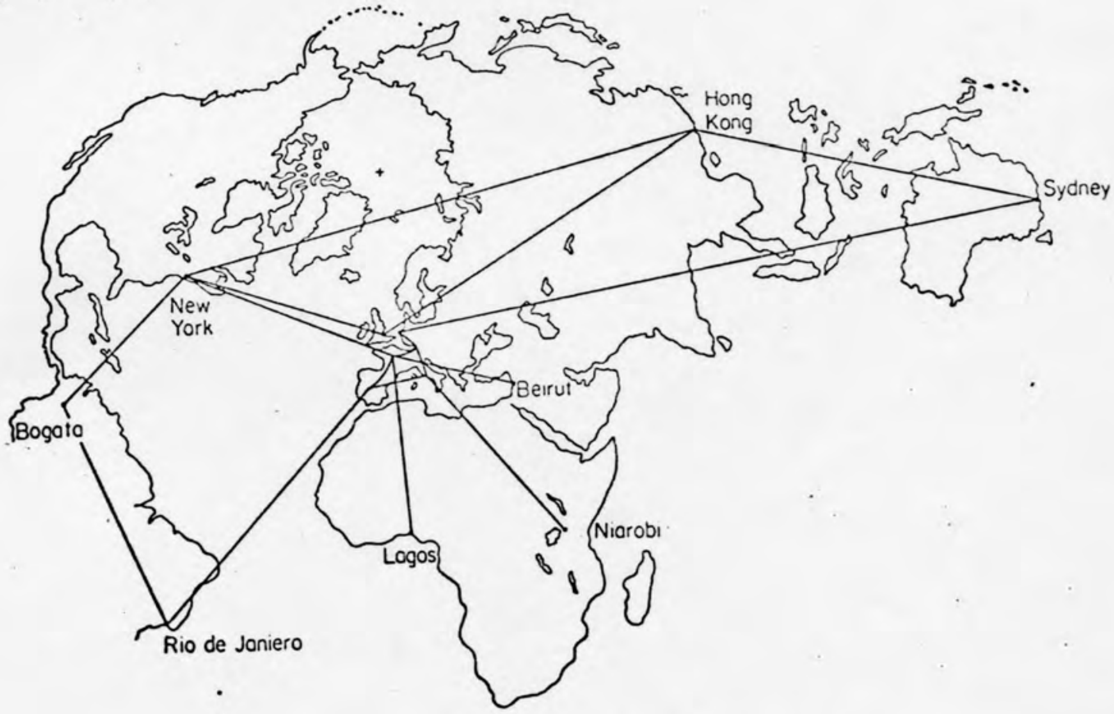

Some planned SITA developments

Figure 4-4 The SITA Network

4.5 The Lincoln to SDC Experiment

In 1966 and 1967 an experiment was performed between the TX-2 computer at MIT Lincoln Laboratory and AFNSQ7 computer at Systems Development Corporation. This experiment is described in [Marill 66]. The experiment's major importance is that it was the immediate predecessor of the very extensive ARPA Network experiment.

4.6 The ARPA Network

The ARPA Network will be discussed very extensively later and we will not describe it further here other than to provide some geographic and topologic maps illustrating the development of the network (see the following Figures) and to note some of the major references [Heart 70], [McQuillan 72], and [Ornstein 72].

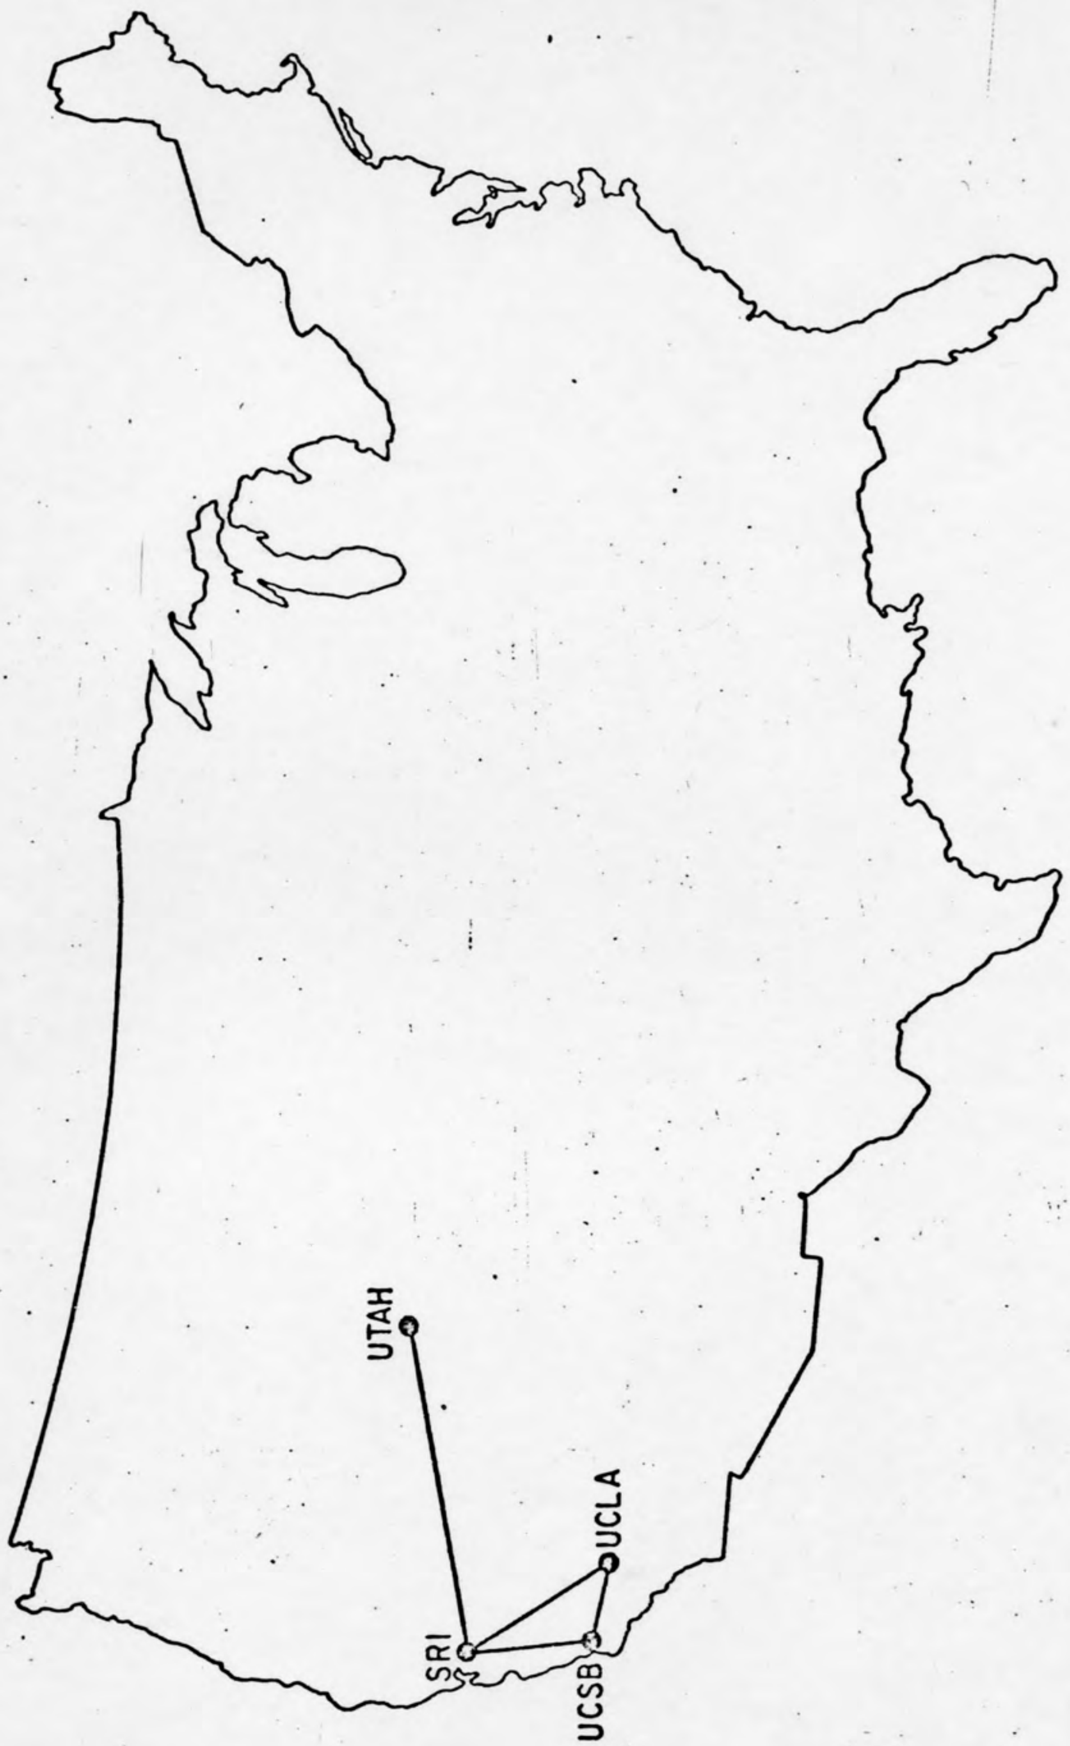

Figure 4-5a Early Geographic Map of the ARPA Network

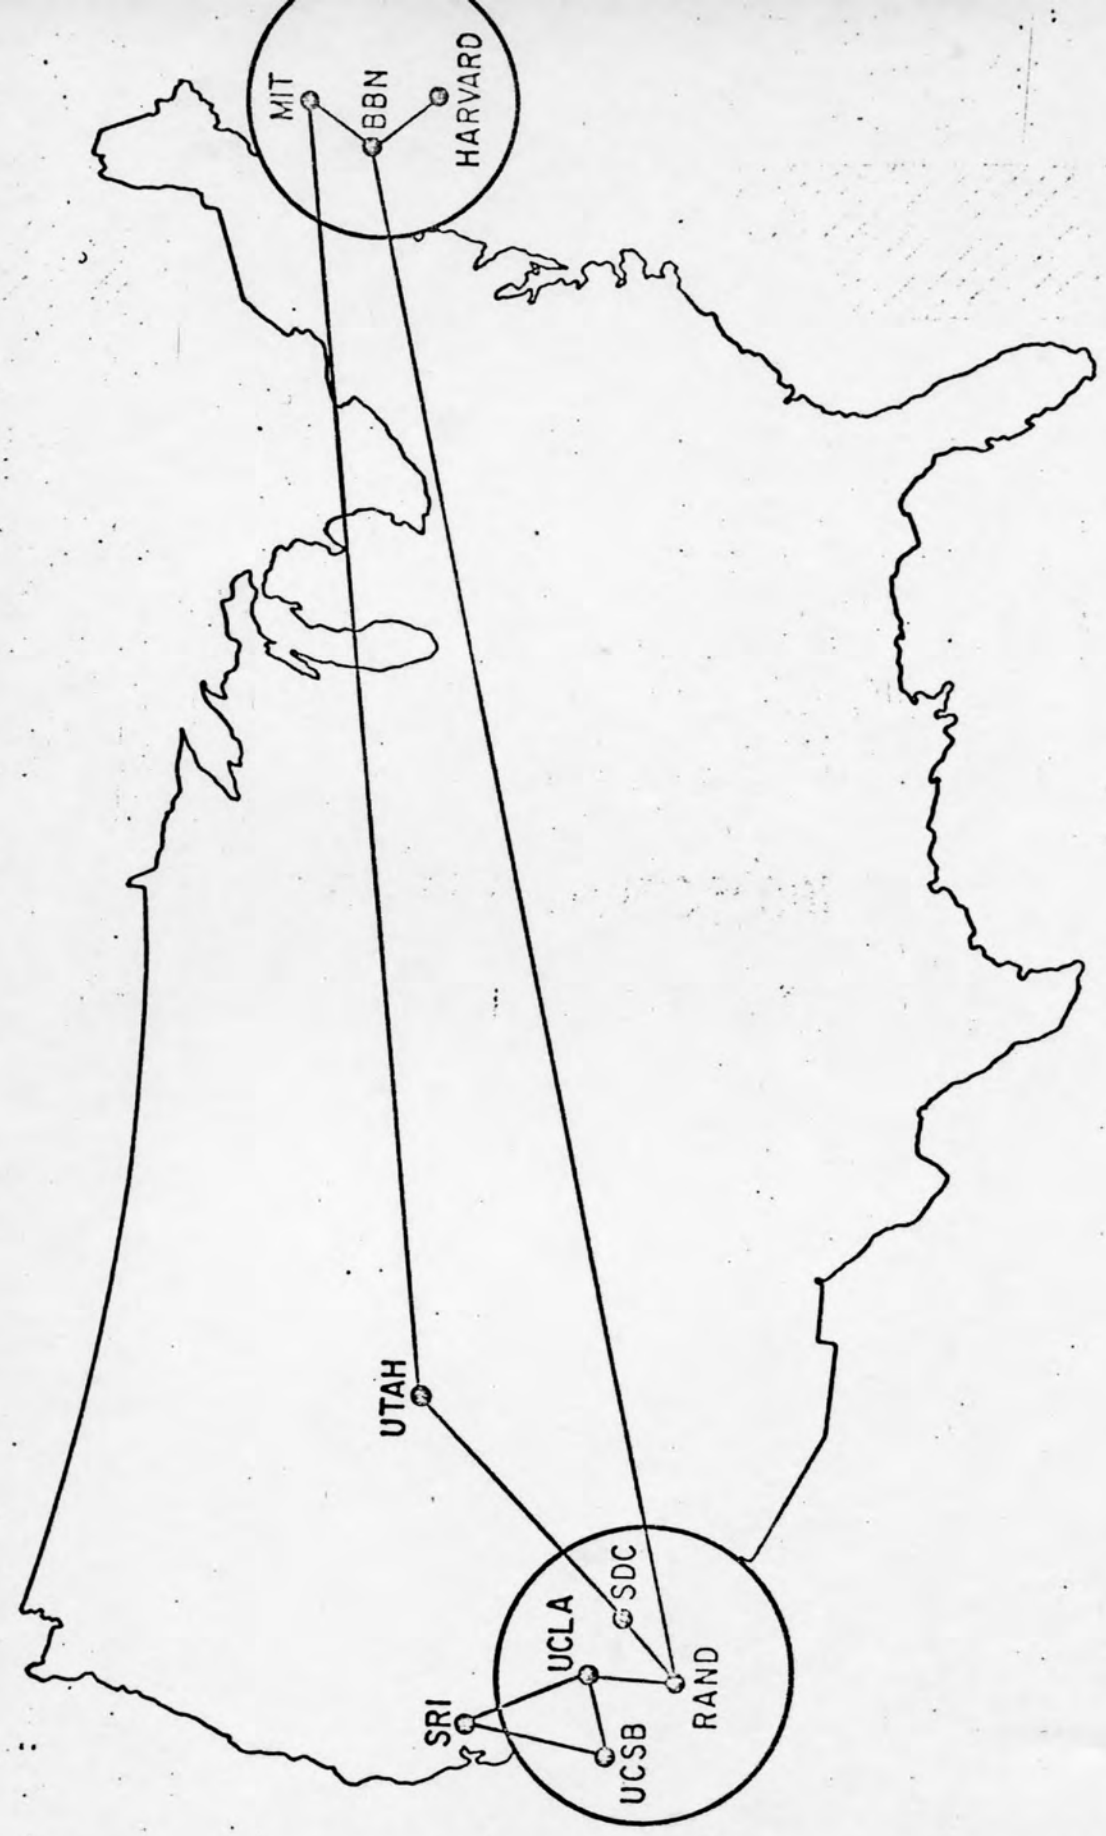

Figure 4-5b Later Geographic Map of the ARPA Network

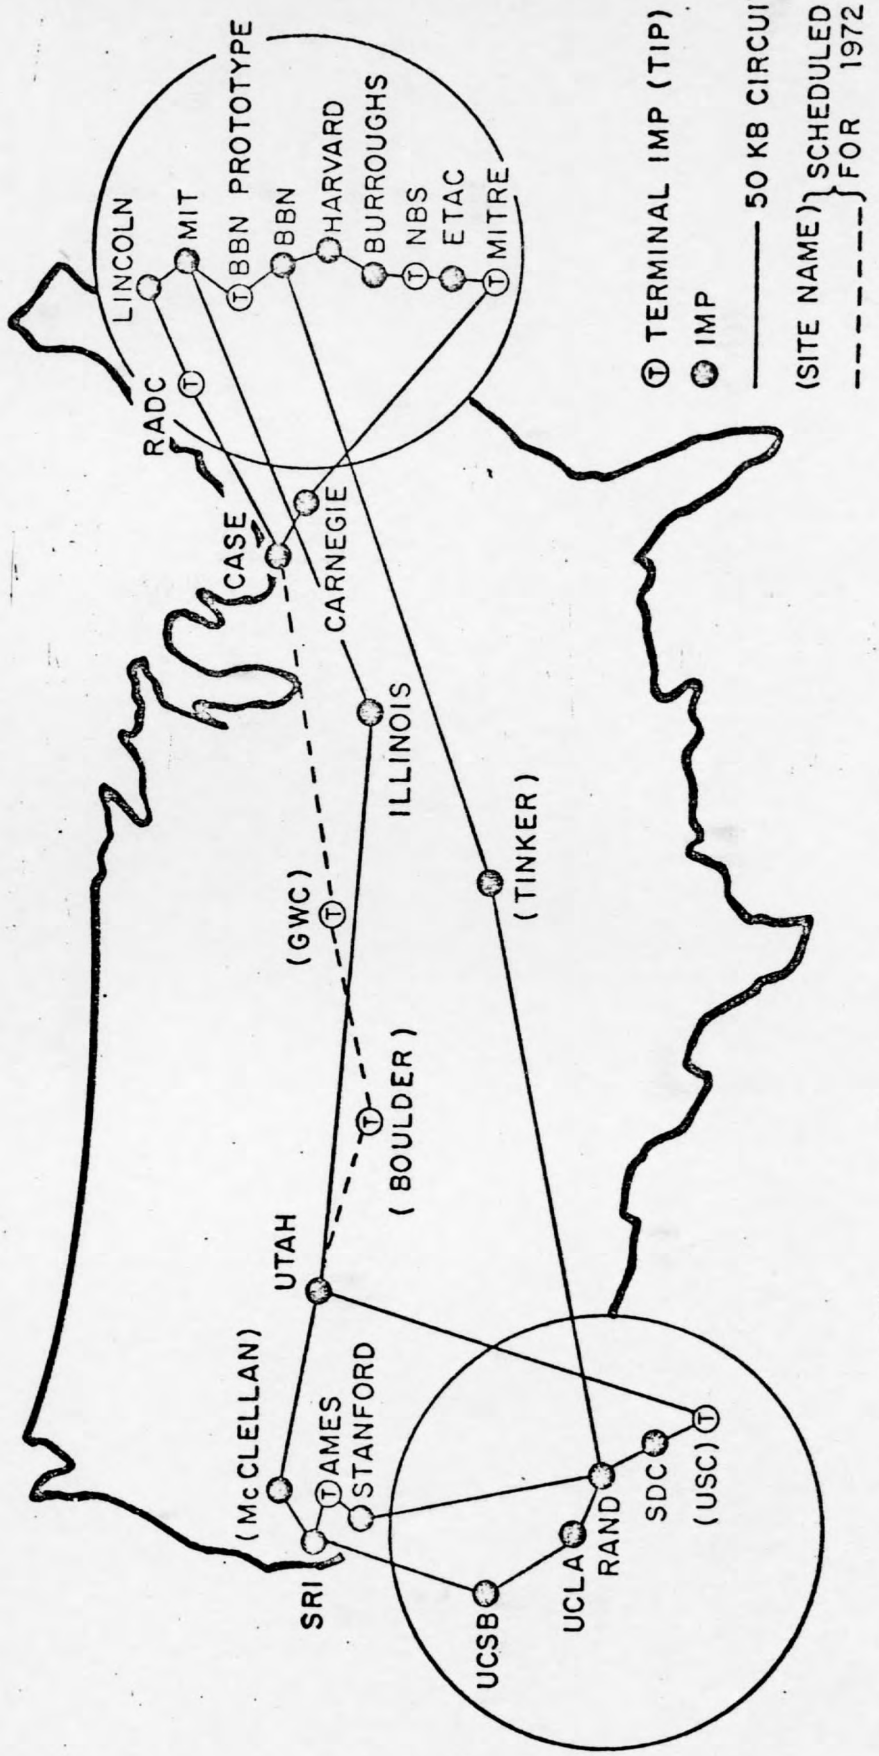

Figure 4-5c Later Geographic Map of the ARPA Network

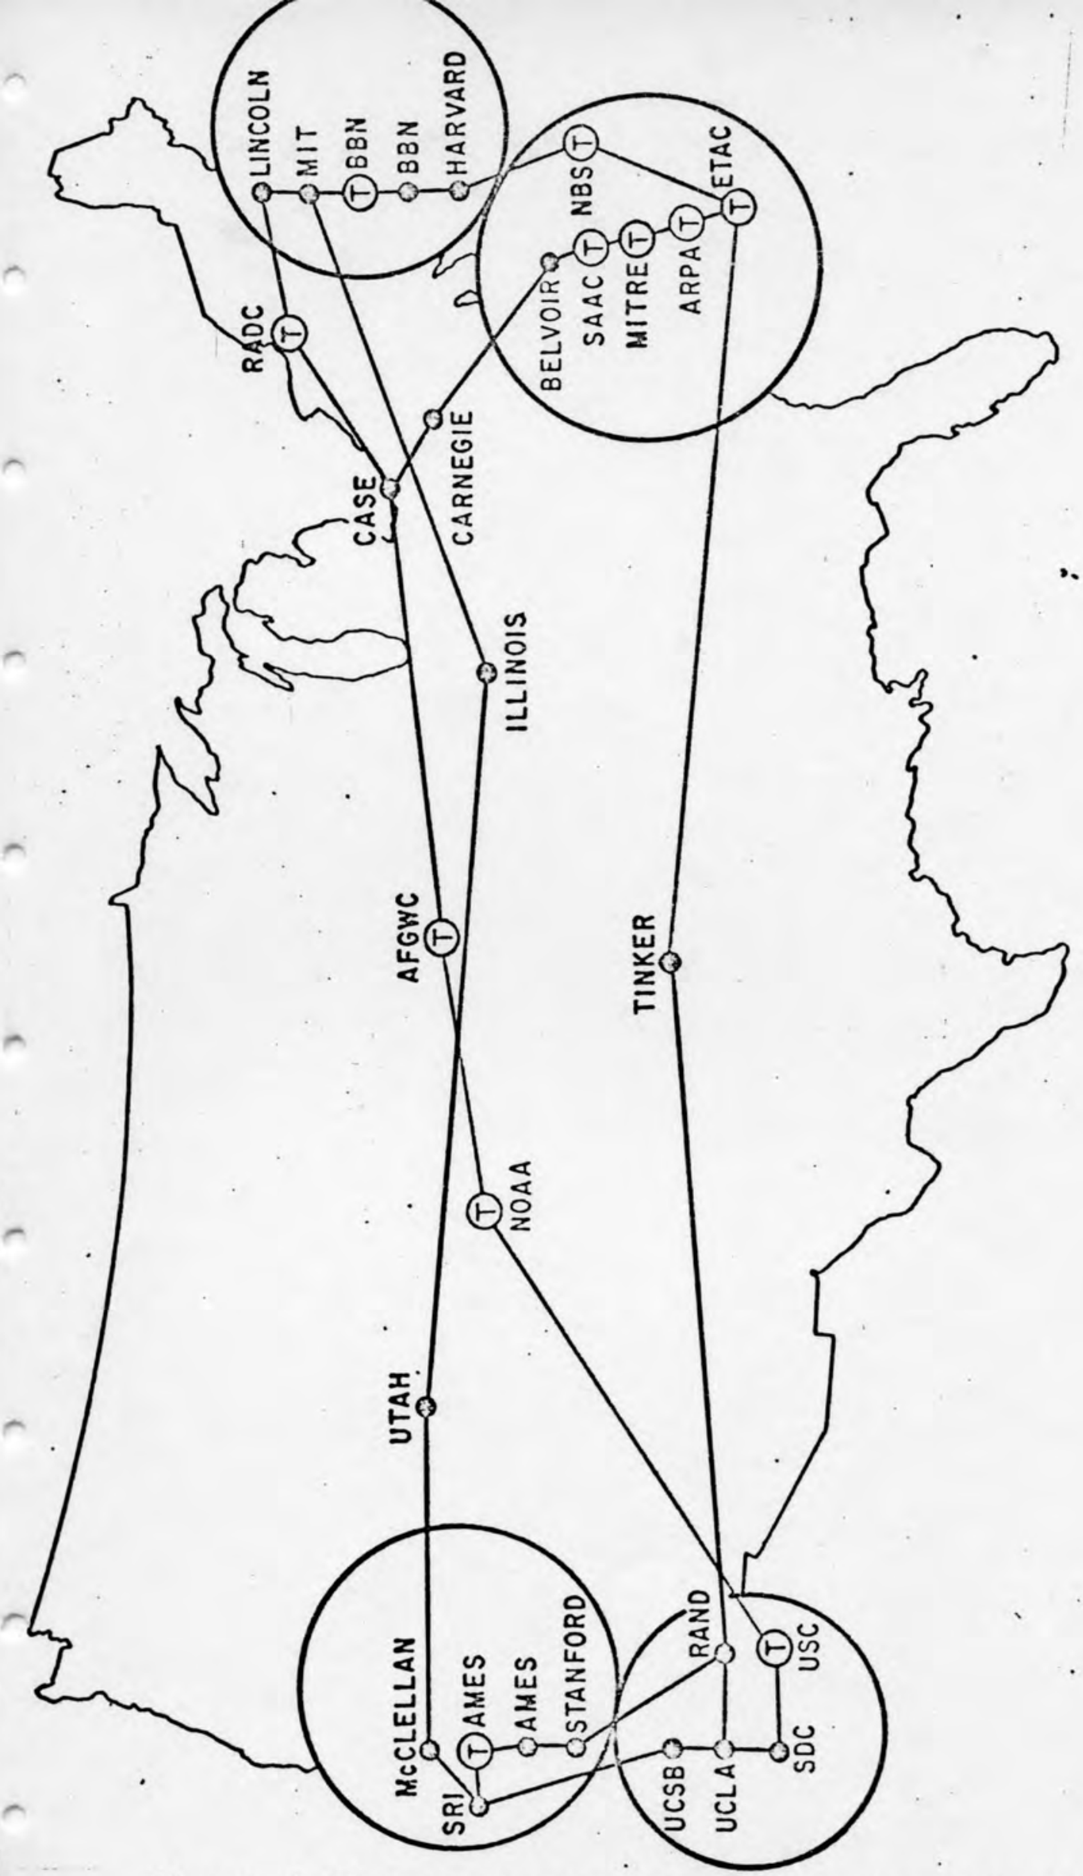

Figure 4-5d Later Geographic Map of the ARPA Network

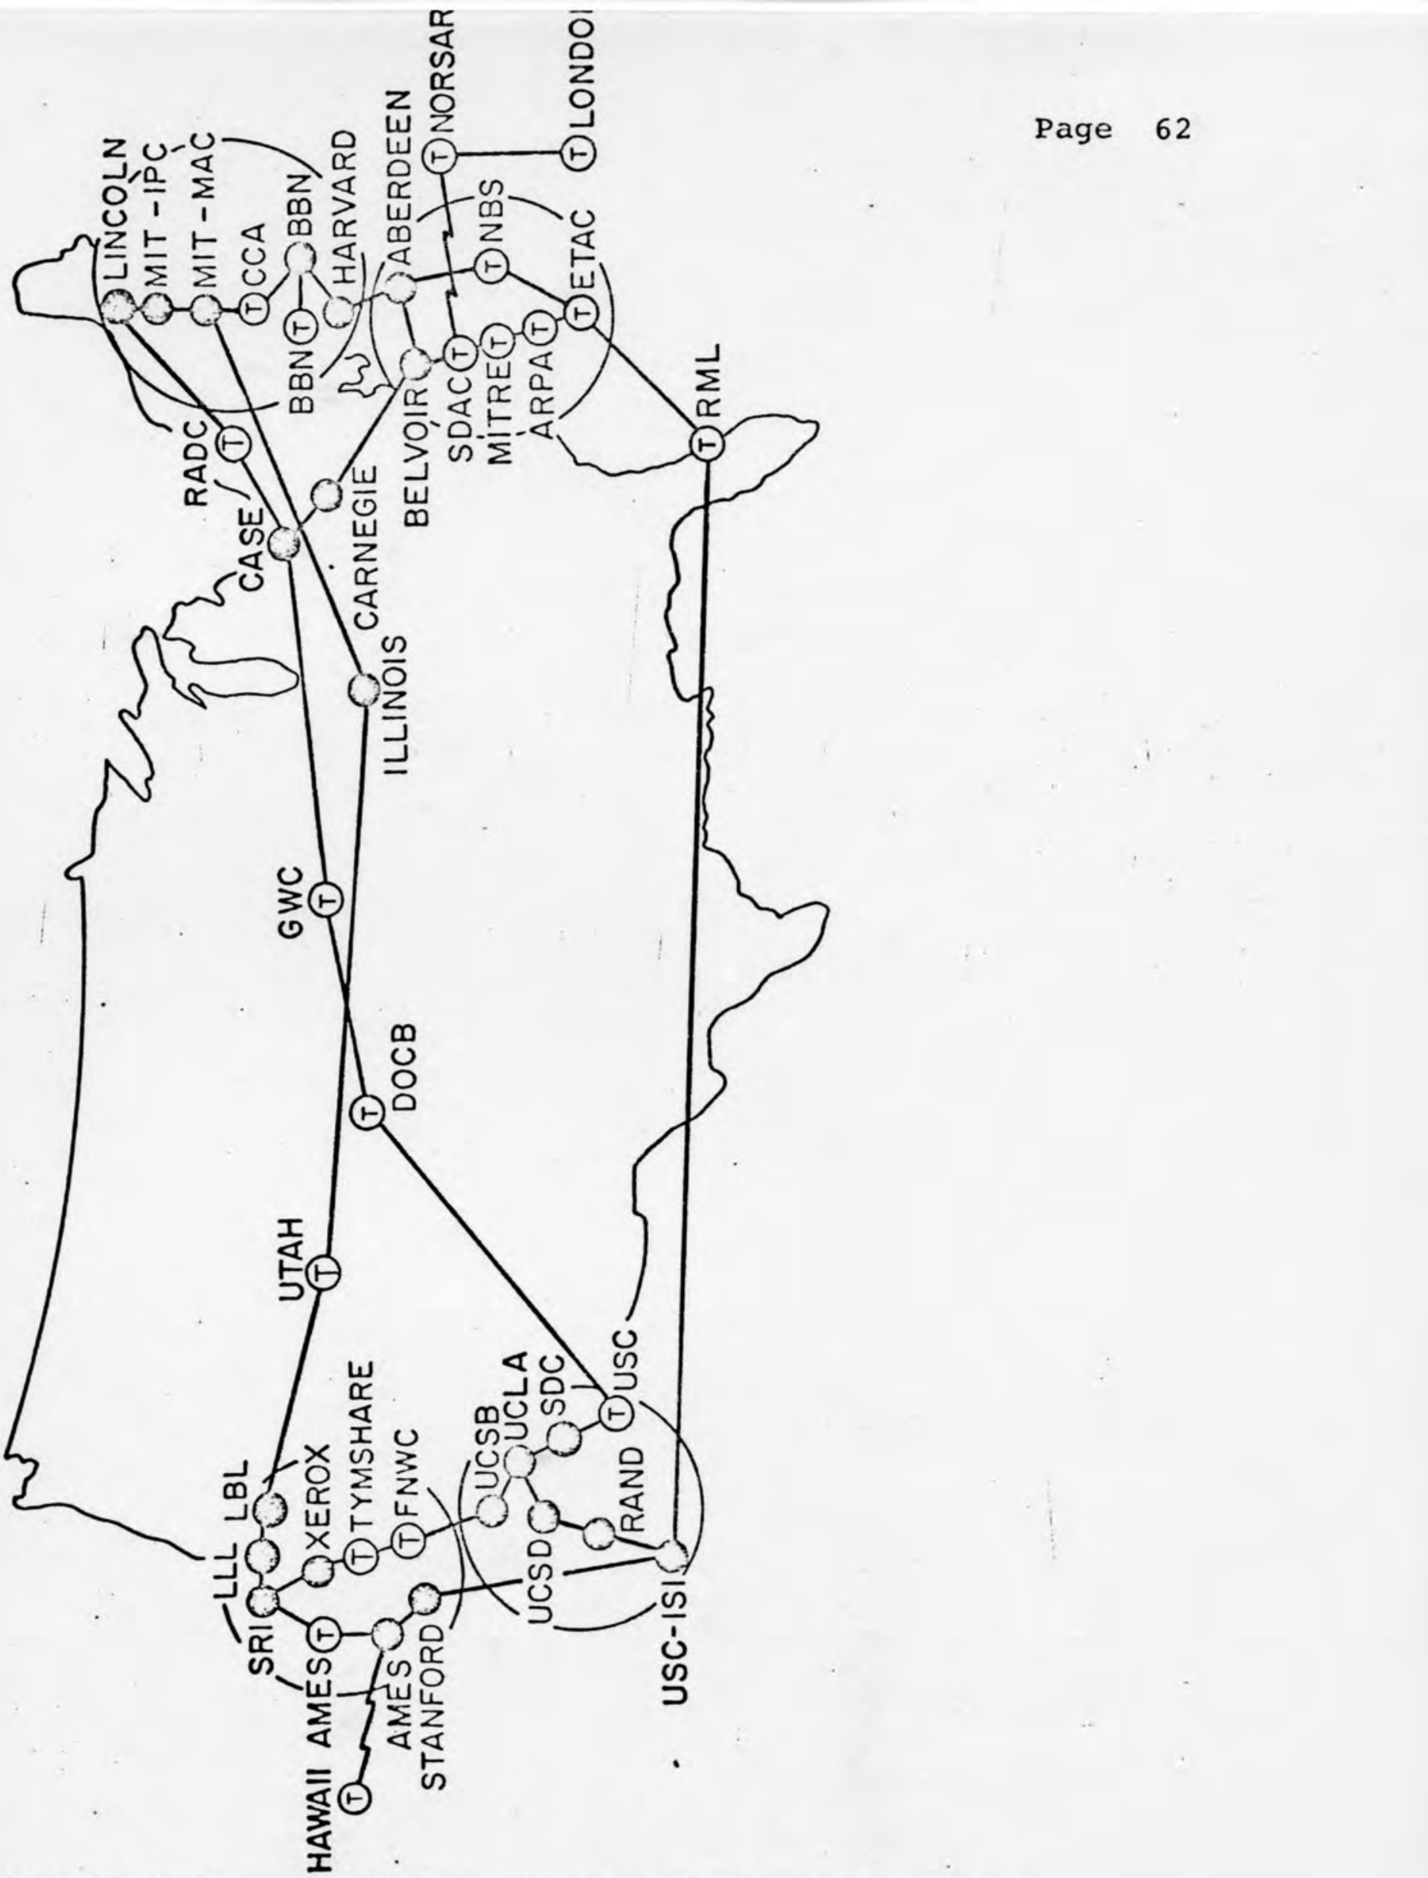

Figure 4-5e Later Geographic Map of the ARPA Network

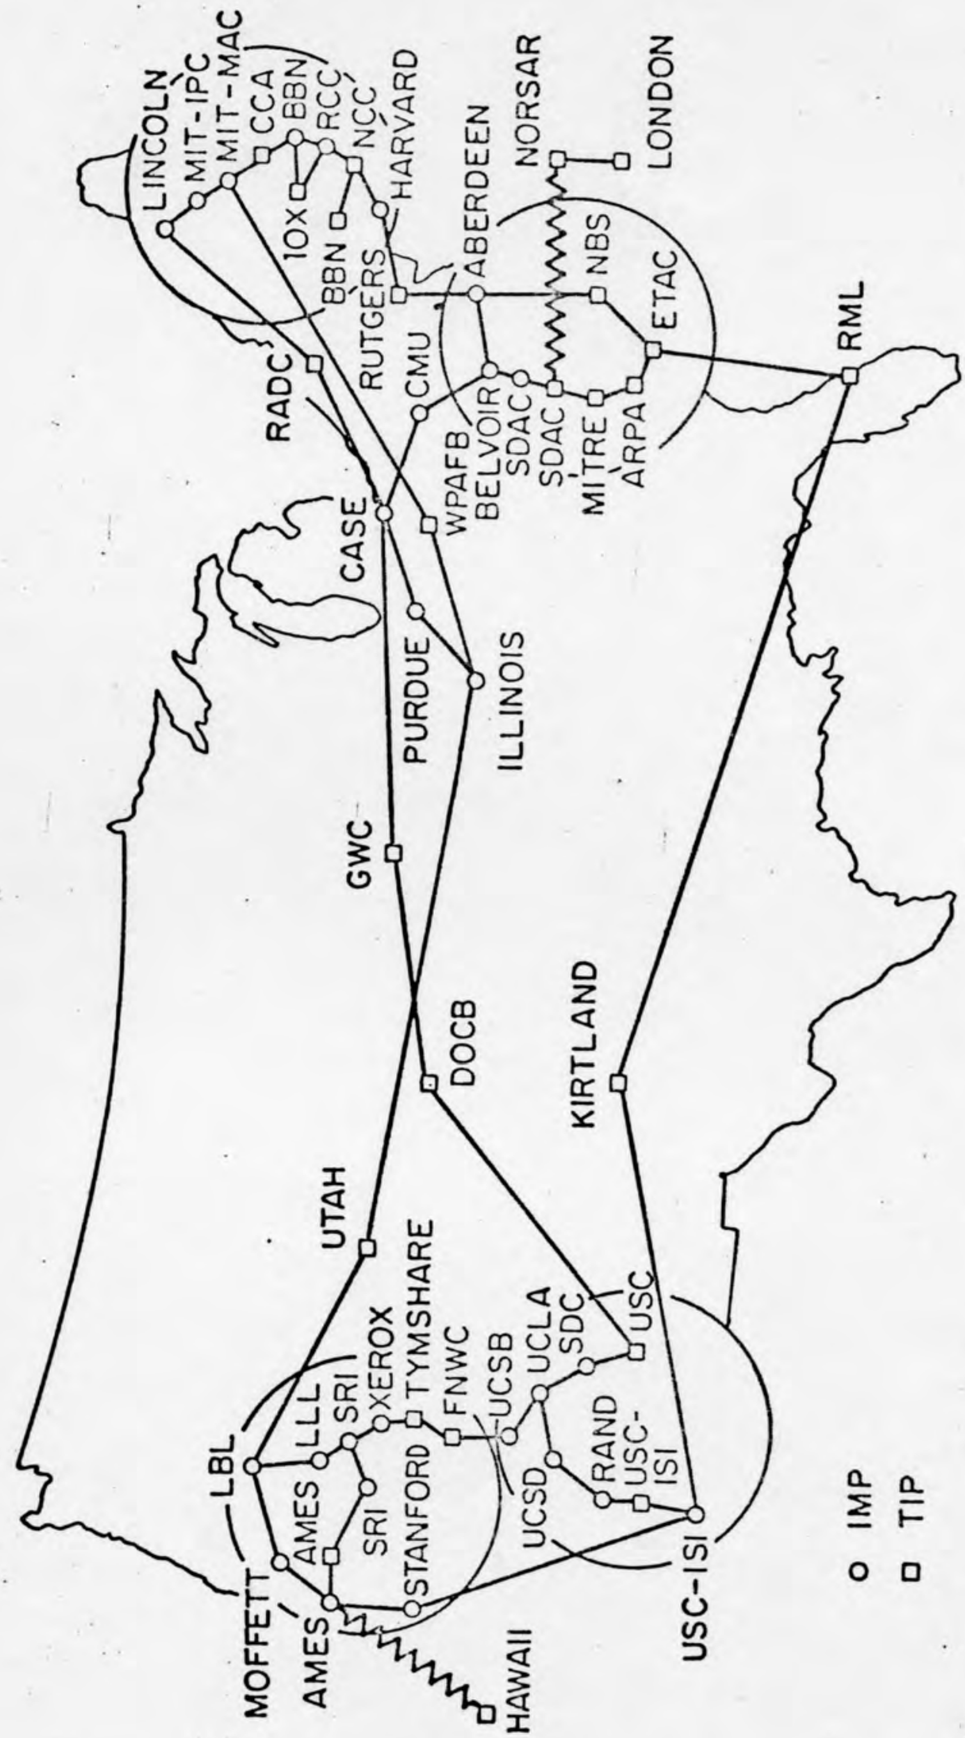

Figure 4-5f Current Geographic Map of the ARPA Network

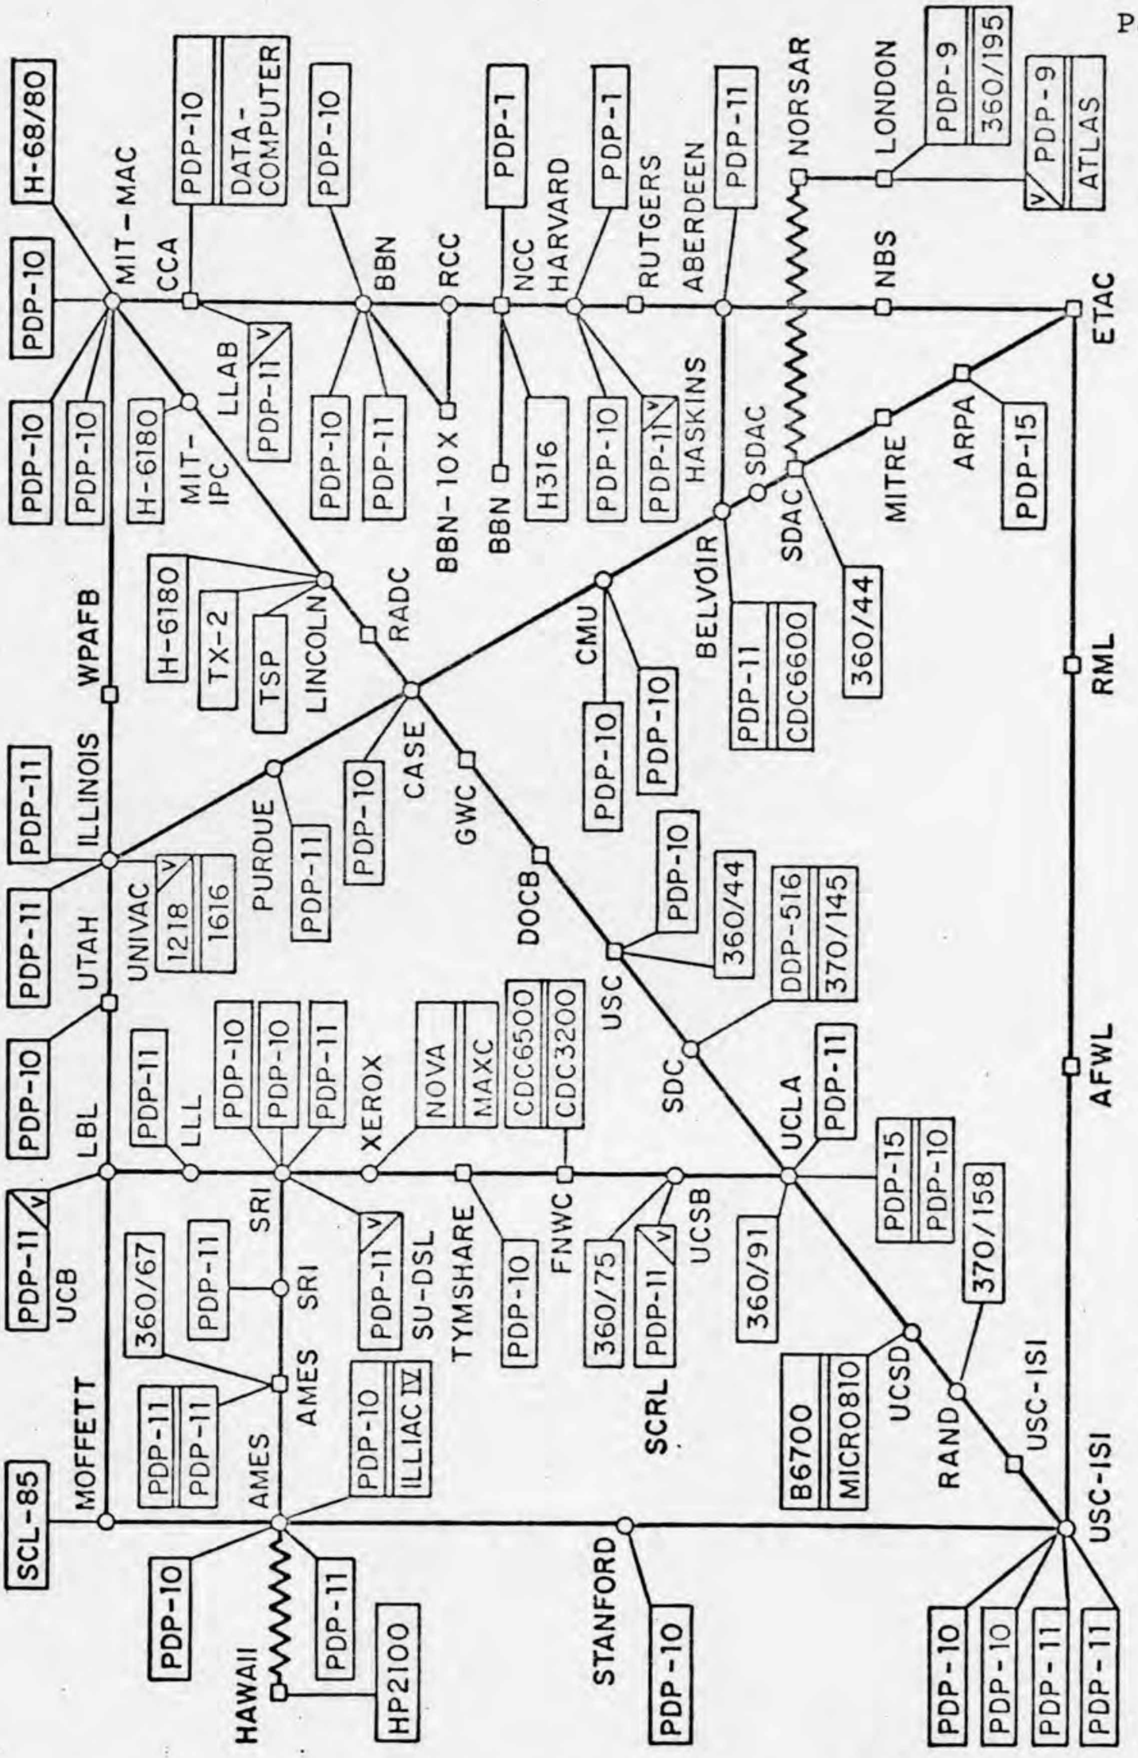

Figure 4-6 Current Logical Map of the ARPA Network

4.7 The ALOHA System

Shortly after the inception of the ARPA Network, the group led by Abramson at the University of Hawaii began work on the ALOHA System. The ALOHA System is essentially a ground-based radio packet switching star. The Menehune takes the place of the ARPA Network IMP as the packet switcher. As shown in Figure 4-7, a number of terminals and mini-computers connected to Terminal Control Units (TCUs) are distributed about. The TCUs consist of a line buffer control unit, a modem, a transceiver, and a UHF antenna. When a line has accumulated in the line buffer unit to be sent to the Menehune, it is burst onto the radio channel without coordination with other TCUs. If the burst is correctly received at the Menehune via the radio receiver and receive modem, it is acknowledged. If it is not correctly received, perhaps due to conflict with a burst from another TCU, the TCU retransmits the burst. Transmissions from the Menehune to the individual TCUs do not conflict. The Menehune is based on a HP2115A computer, commercial transceivers are used, and the radio channel is utilized via 24,000 baud modems.

References to the ALOHA System are the paper by Abramson [Abramson 70] and that by Kuo and Abramson [Kuo 73].

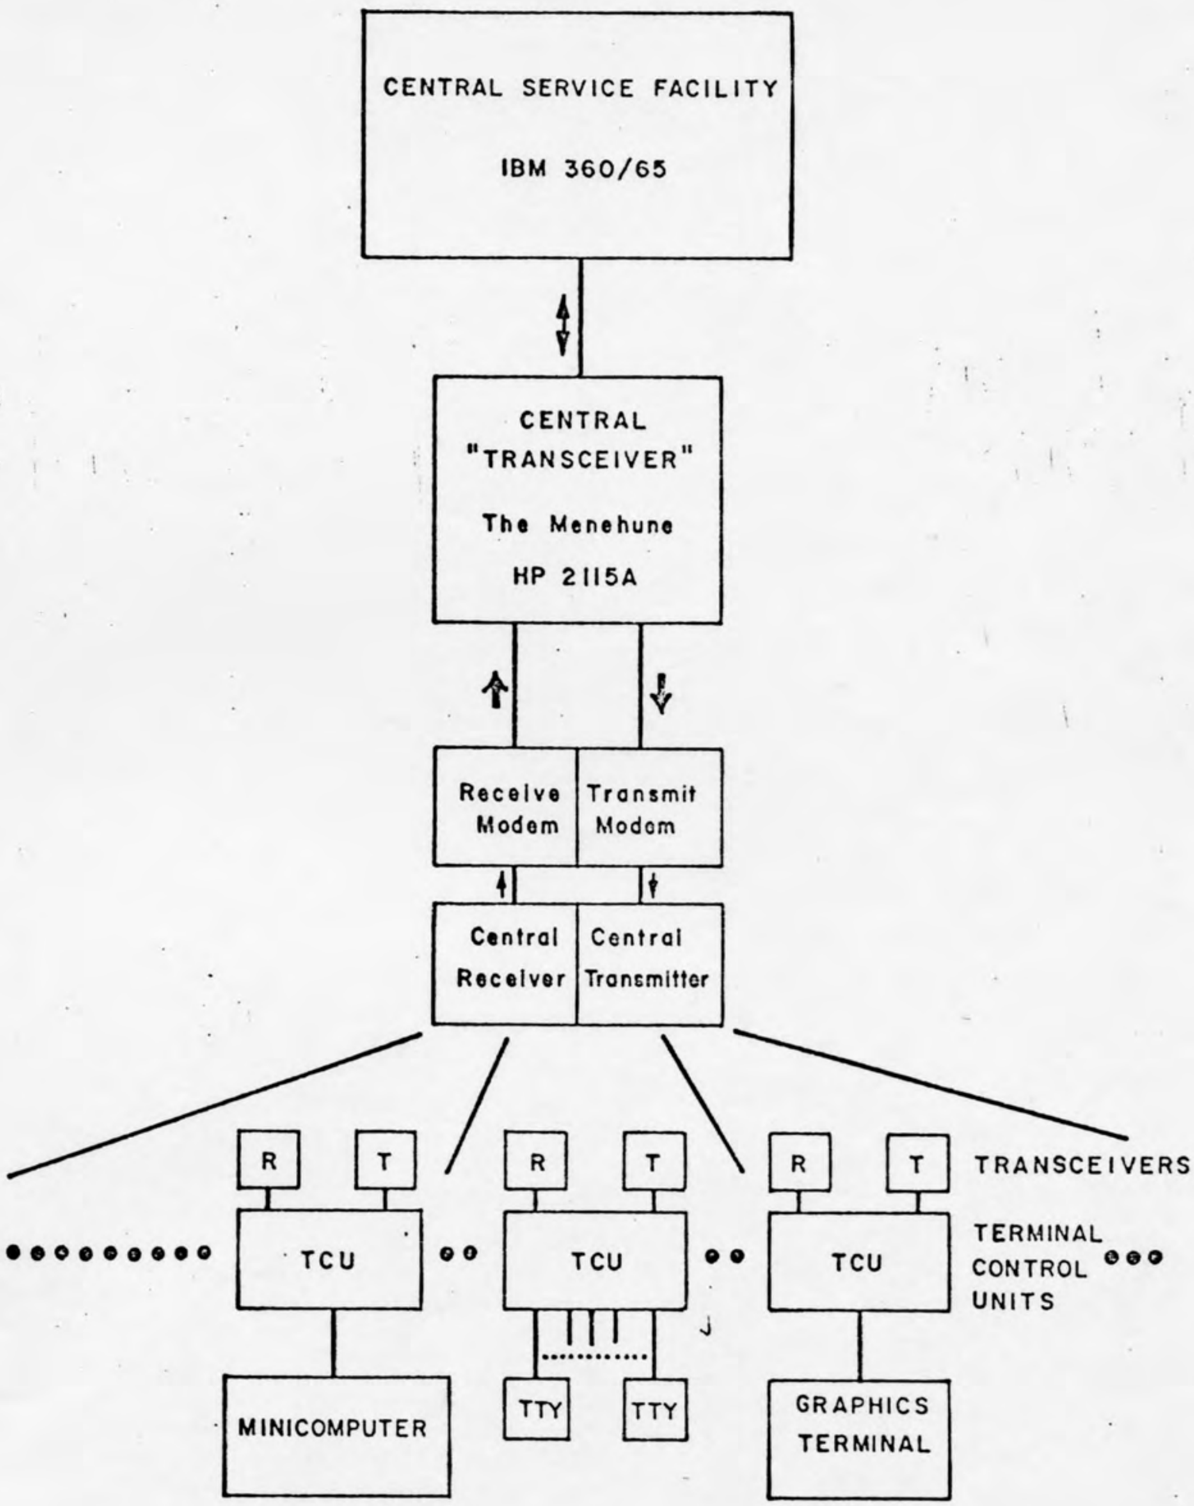

Figure 4-7 The ALOHA System

4.8 The Distributed Computer System

At the University of California at Irvine, Farber has built a packet switching ring network. As shown in Figure 4-8, the system is based on a high speed, one way closed communications loop. A hardware device called a Ring Interface (RI) interfaces minicomputers to the ring network and performs the functions of inserting messages onto the ring, deleting them from the ring, and relaying them around the ring. Each RI has an associative memory listing the unique names of the processes in the local computer. Messages are sent to a process by name. The RI picks off messages for its processes as they come by on the ring and at the same time relays them with the "copy" bit set. Messages not for local processes are just relayed. The originating machine removes messages it originally sent and notes the state of the "copy" bit which acts as an acknowledgment bit.

There are many references to the Distributed Computer Network: two are the papers by Farber and Larsen in the Brooklyn Polytechnical Conference Proceedings [Farber 72a], [Farber 72b]; another is a good summary in [Akkoyunlu 74], others are [Farber 72d] and [Farber 73].

Although slightly out of the current mainline of packet switching, DCS is especially interesting as it calls to our attention the various loop networks; e.g., [Pierce 72].

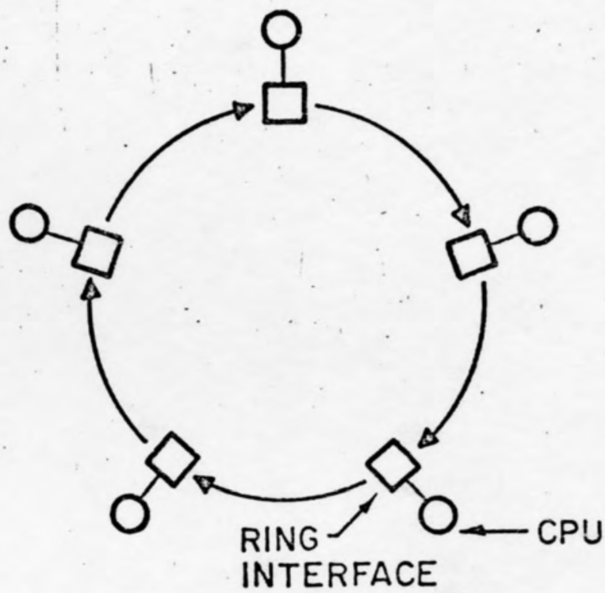

Figure 4-8 The Distributed Computer System

4.9 Packet Broadcast Satellite Communication

A direct outgrowth of the ALOHA System and the ARPA Network is the satellite broadcast system implemented in the Satellite IMP. This is described in [Butterfield 73] and will be discussed in more detail later. Figure 4-9 and Figure 4-10 illustrate the broadcast concept and its inclusion in a wider network.

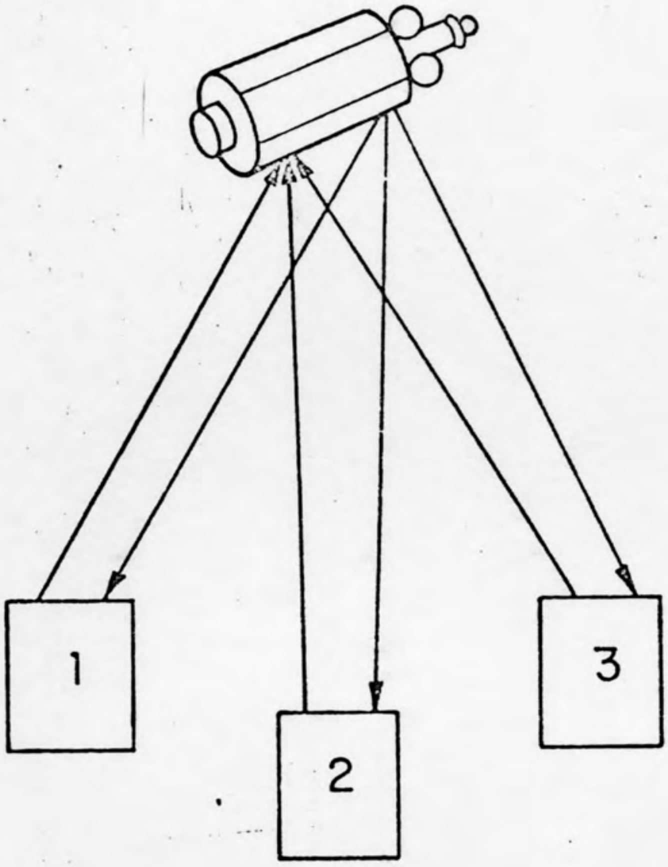

- SAME TRANSMITTING AND RECEIVING FREQUENCIES
- CARRIER ON/OFF CONTROL WITH FAST ACQUISITION
- LONG ROUND TRIP TIME ~1/4 SE

Figure 4-9 The Packet Broadcast Satellite Concept

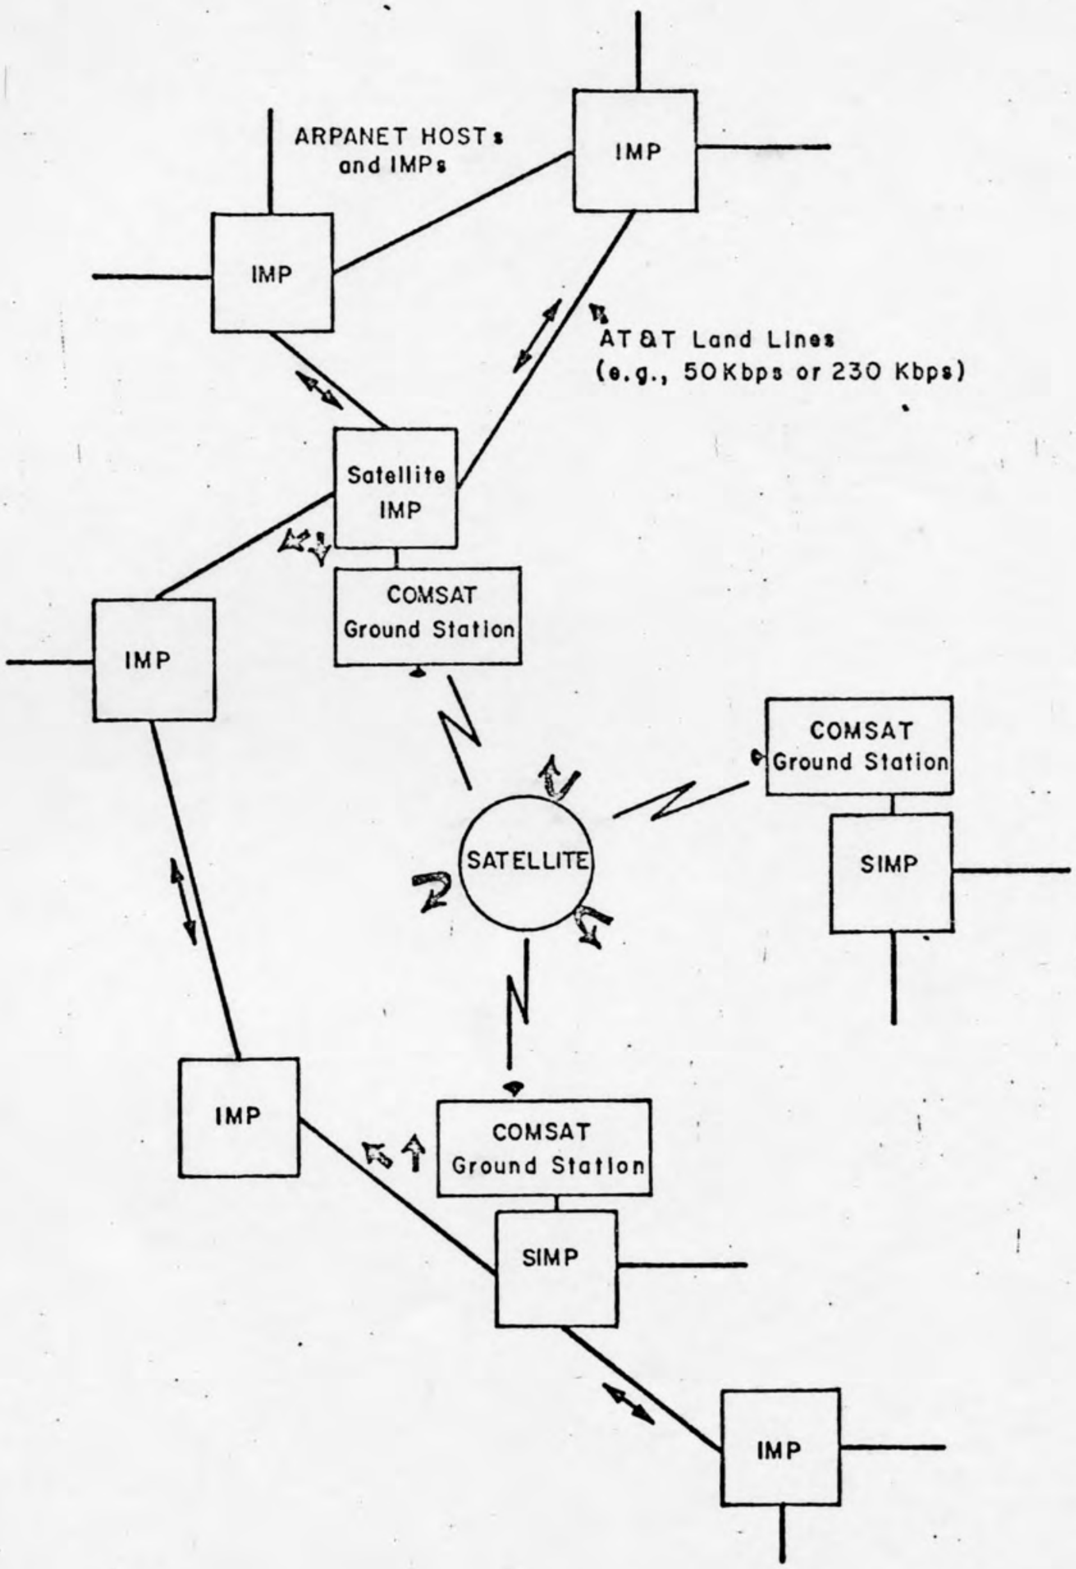

Figure 4-10 Packet Broadcast Satellite in a Wider Network

4.10 Packet Radio

A similar outgrowth of the ALOHA System and the ARPA Network is the ARPA Packet Radio System suggested in [Roberts 72]. Shown in Figure 4-11 which was taken from [Metcalfe 73], we see that the packet radio consists of Packet Radio Stations, Radio Relays, and the Packet Radio Terminals. A terminal is generally within range of several stations and relays, but not within range of all. The system is also operated at low power. Thus, by adding a routing function to the stations which instruct the relays which ones should repeat bursts to a given destination, it is possible to gain much greater capacity out of the system than would be possible if every relay repeated every burst. Part of this process is for the stations with the aid of the relays to detect the movements of the terminals so as to adjust the routing.

The comprehensive set of references for the packet radio system is [Kahn 75], [Abramson 75], [Garrett 75], [Kleinrock 75], [Frank 75], [Fralick 75], and [Burchfiel 75].

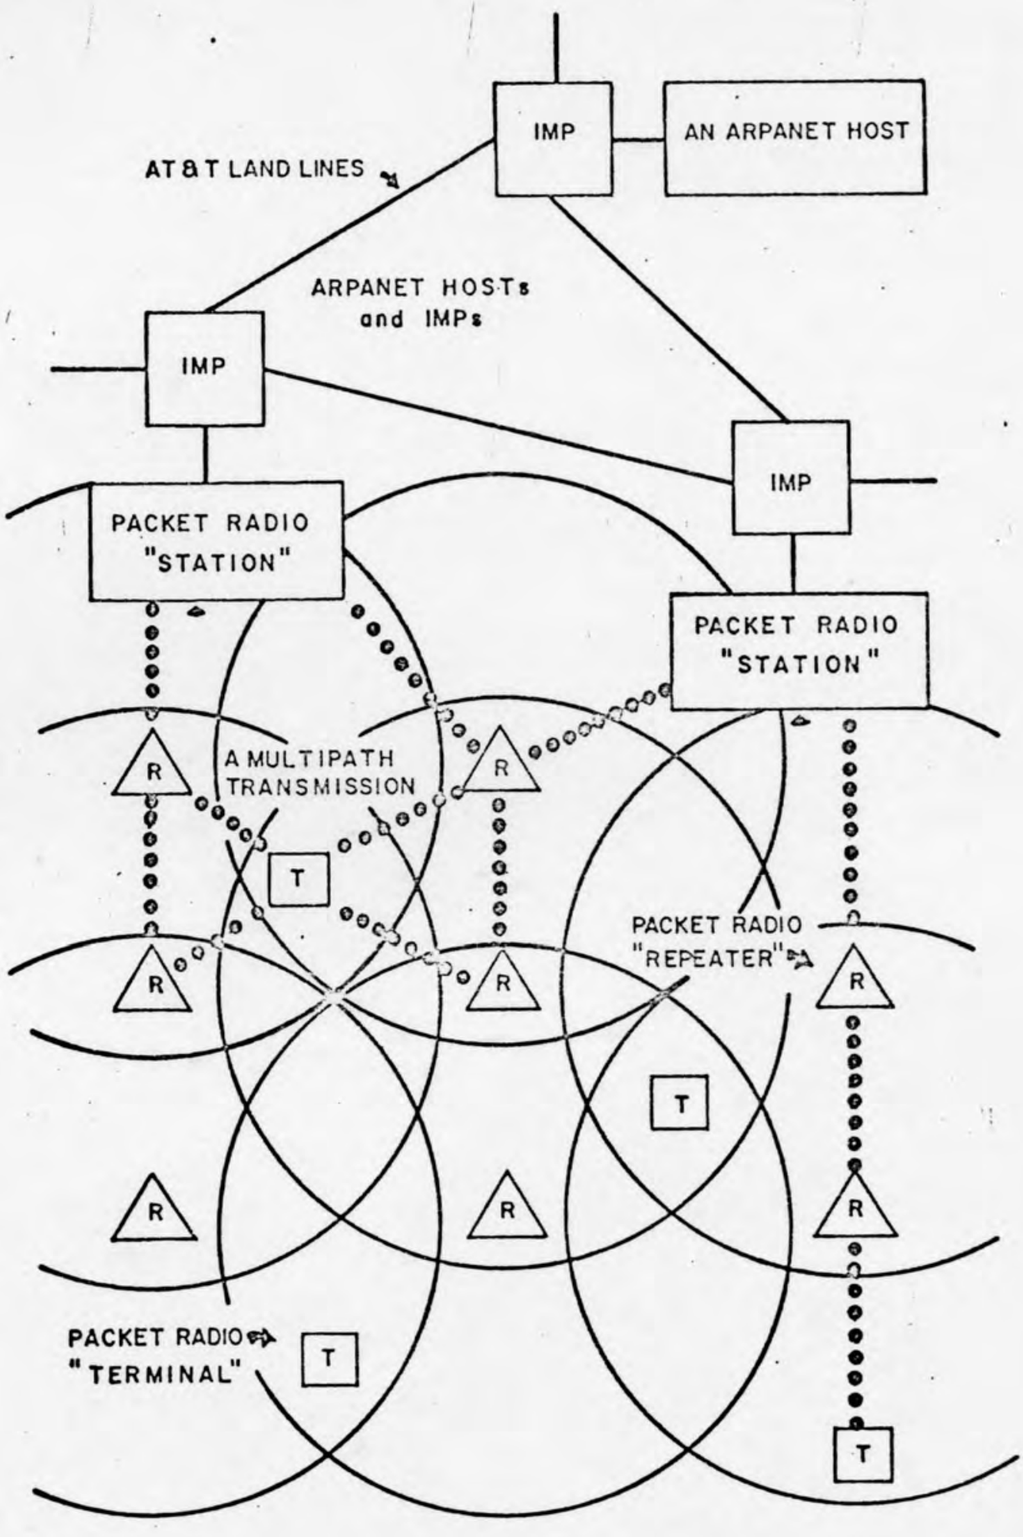

Figure 4-11 Packet Radio

4.11 ARPA Network Possible Extensions.

We merely list a number of the possible extensions to the ARPA network in order of decreasing probability.

a. more nodes in the continental network (six are scheduled and six more are probable).

b. Satellite IMPs in the Atlantic area (at Etam, Goonhilly, Tanum, and possibly Teheran).

c. use of a domestic satellite (two 1.5Mbs Satellite IMPs are under construction and several more are a possibility).

d. a Packet Radio Station will be installed in the Palo Alto area with a connection to the ARPA Network.

e. Satellite IMPs are a possibility in the Pacific area (e.g., in Hawaii, California, Alaska, and the far east).

f. once there is a Satellite IMP in the far east and one in Teheran, use of the Indian Ocean Satellite becomes a possibility.

g. a shipboard Satellite IMP with local distribution via packet radio is a possibility.

h. connection to other networks (e.g., Telenet, Cyclades) is a possibility.

4.12 Packet Communications Incorporated

The first of the so-called Value Added Networks is PCI. A Value Added Network uses the circuits of the already existing carriers with the addition of something, namely packets. Thus, the VANs don't have to condemn people's land, build microwave-towers, etc. They just least existing services. The following three Figures (4-12, 4-13, and 4-14) show PCI's plans. (PSPs are like the ARPA Network IMPs and TAPS are like the terminal portion of ARPA Network TIPs.)

PCIs intention is to give their systems increased reliability over the ARPA Network machines through the selection of better computers.

The best source for information on the PCI network is their FCC filing [PCI 73]; an alternative is the Auerbach evaluation of their service [Auerbach 74].

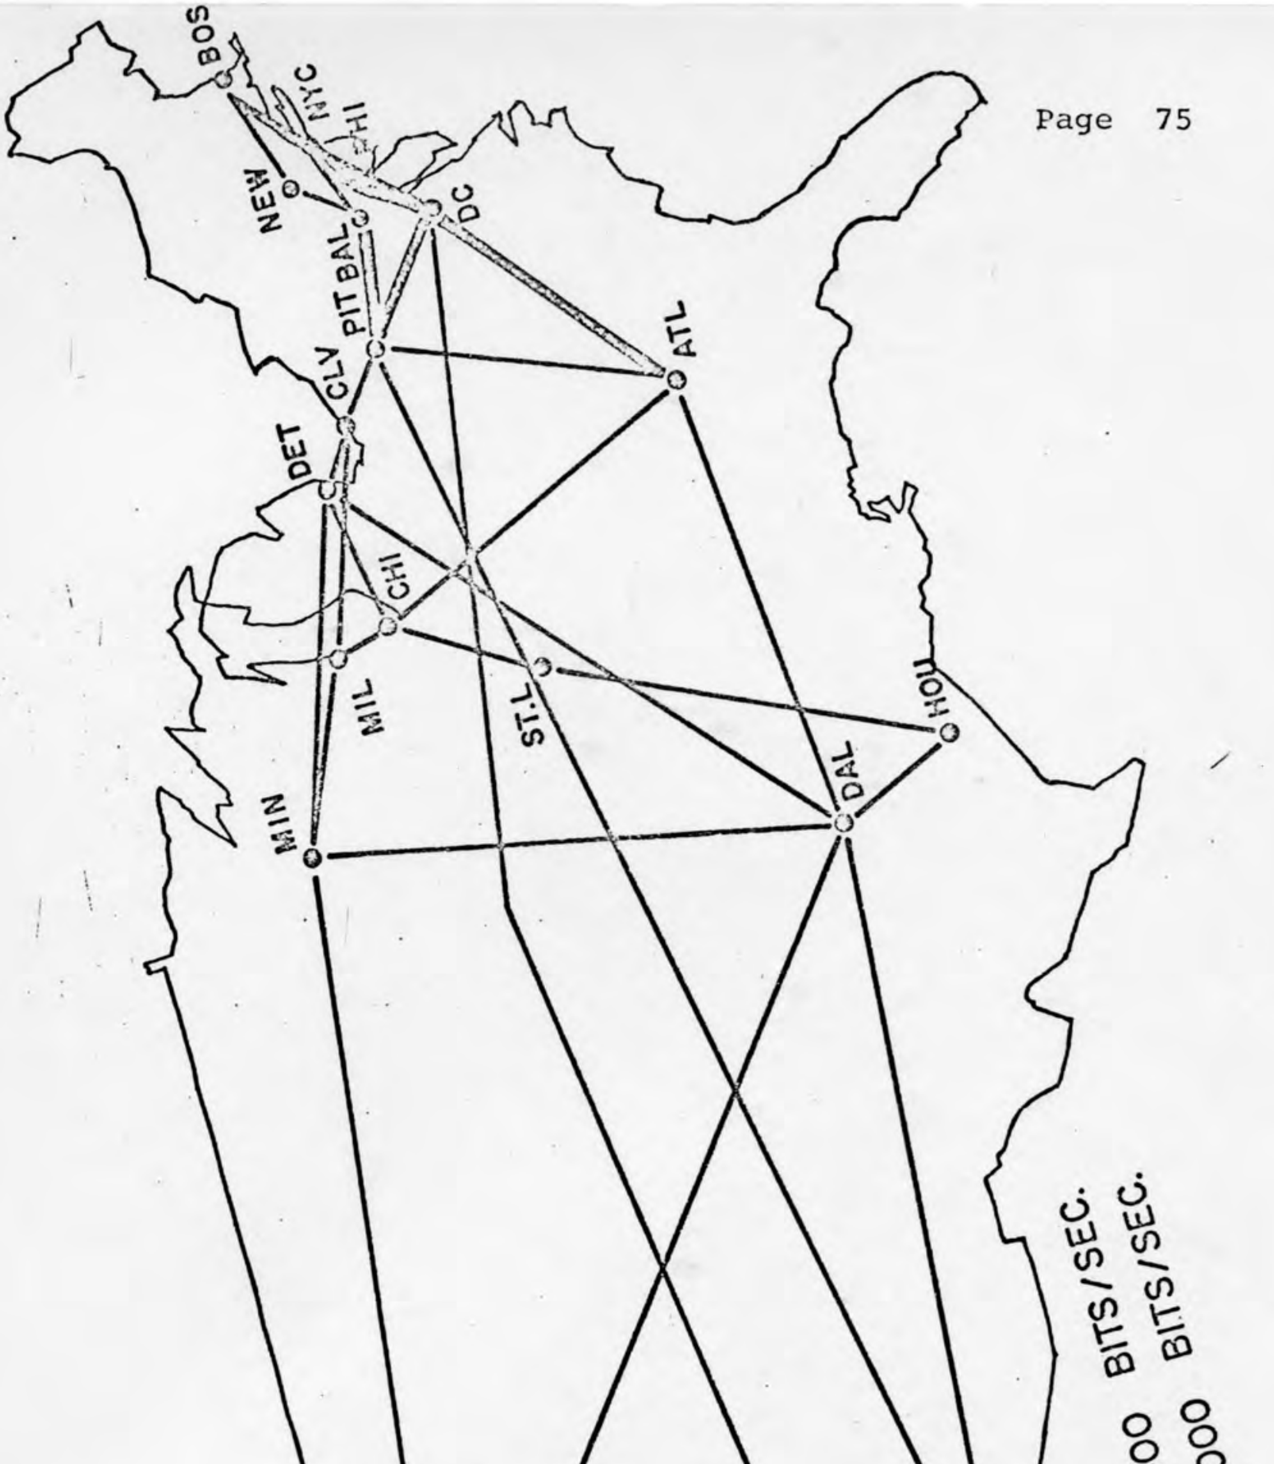

Figure 4-12 PCI Example Network, Circa 1975

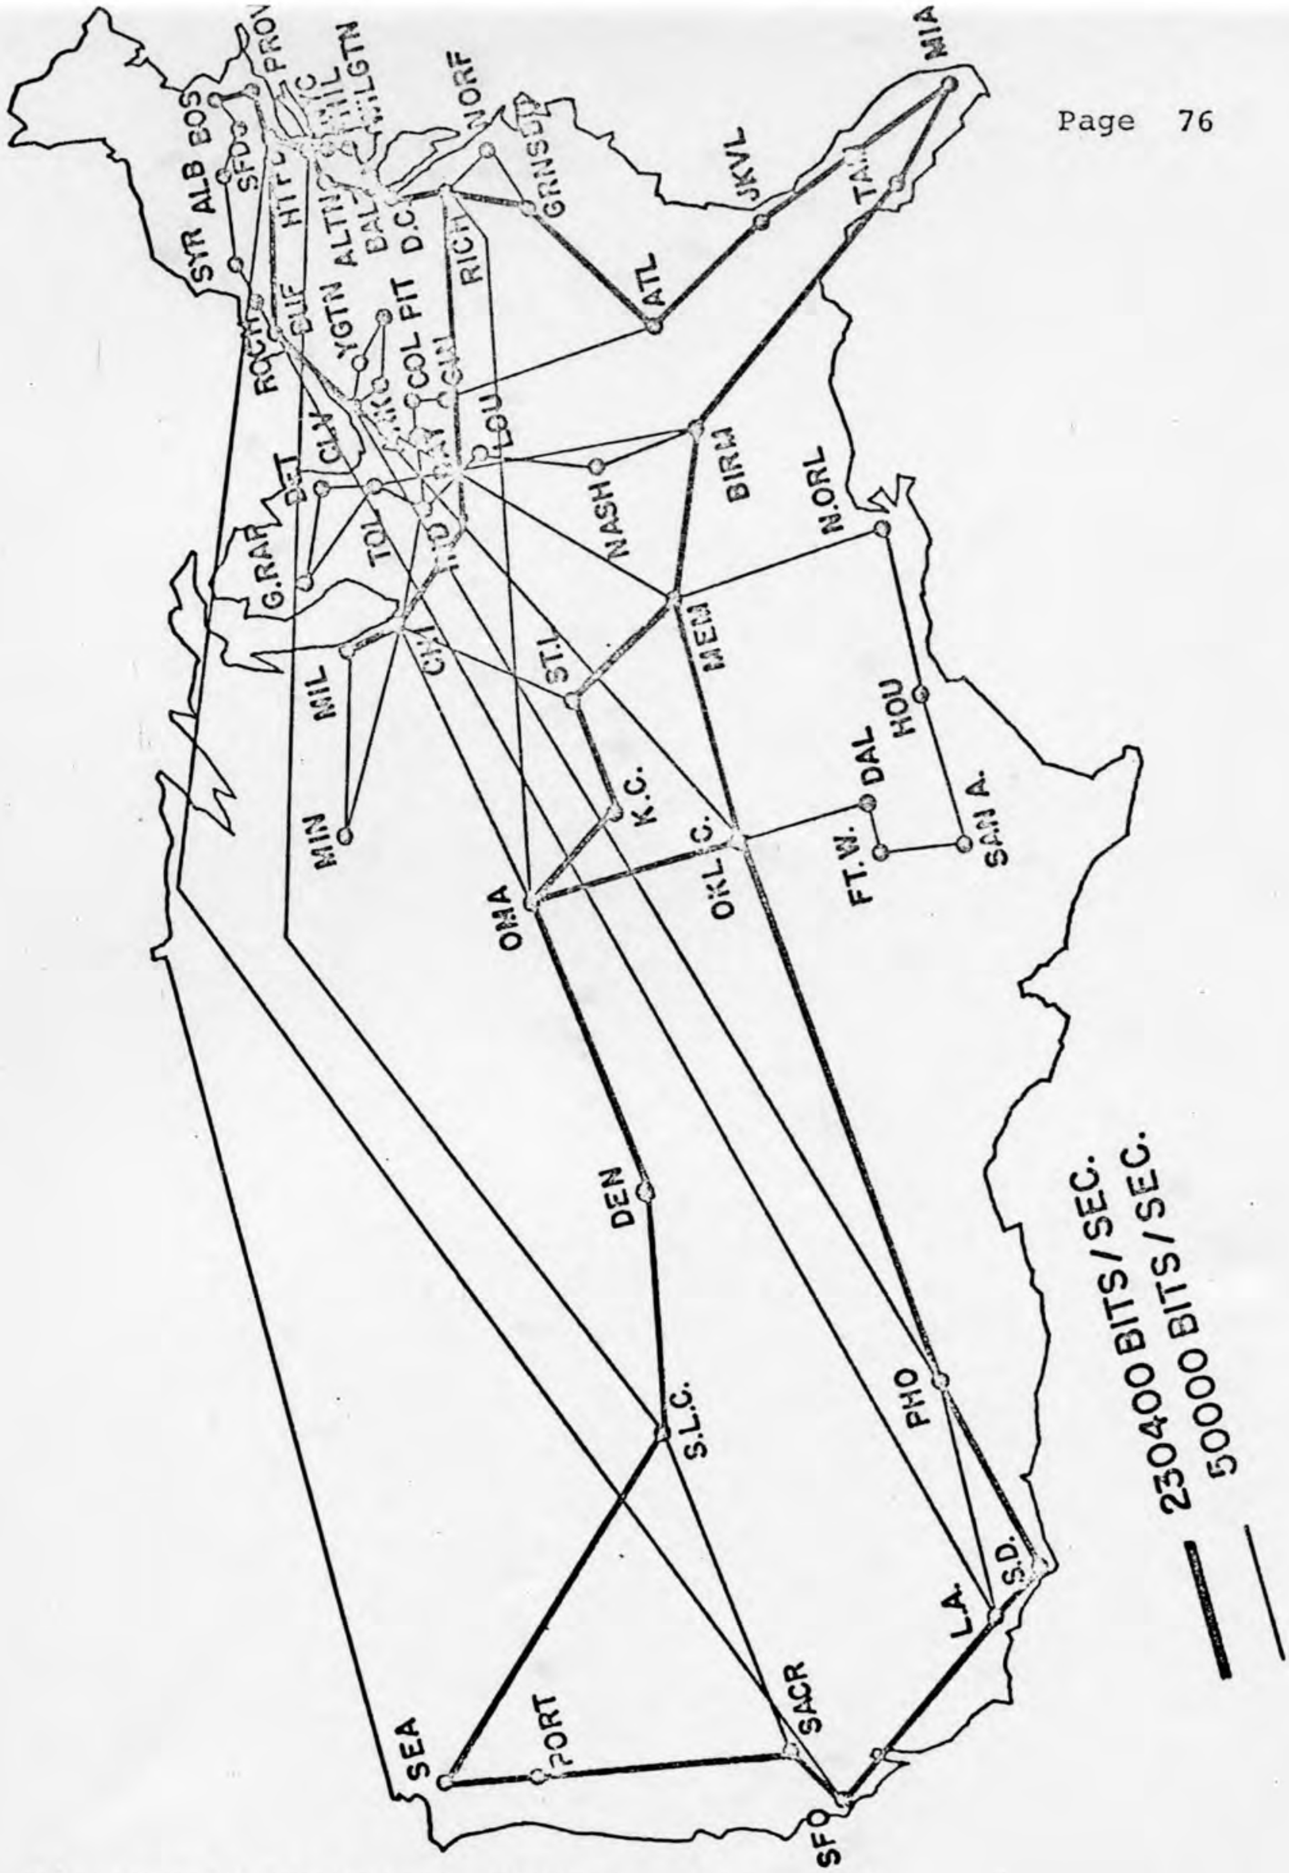

Figure 4-13 PCI Example Network, Circa 1978

Figure 4-14 PCI Example Network, Schematic

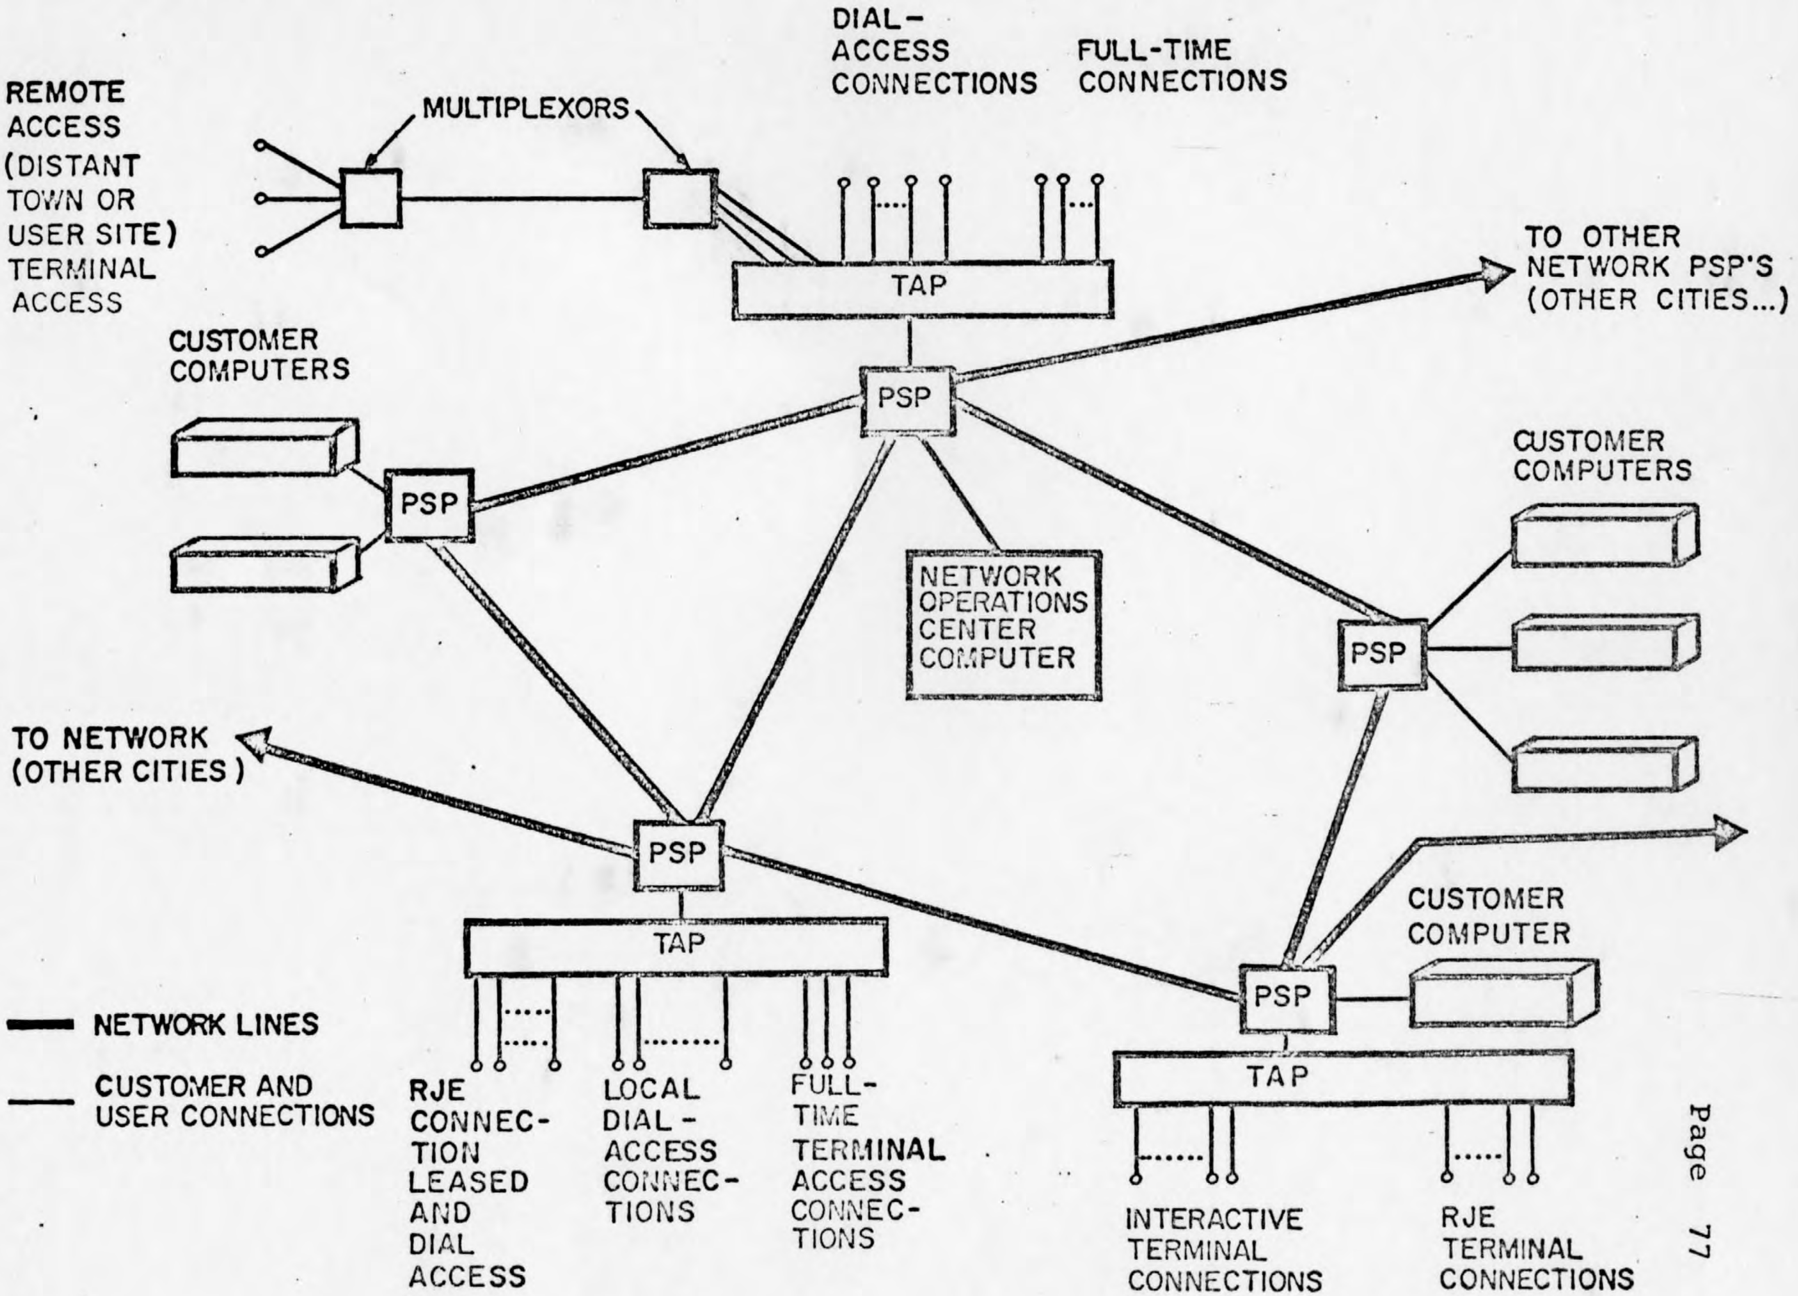

4.13 Telenet Communications Corporation

Telenet is another VAN, set up by BBN, the prime contractor for the ARPA Network, which now holds a minority interest in the company. With L. G. Roberts, the father of the ARPA Network as its president, Telenet is quite closely related to the ARPA Network.

Telenet received its FCC approval in April of 1974. The present goal is for initial service in mid-1975. As shown in Figure 4-15, this initial network will consist of 50Kbs and 9.6Kbs circuits. The network will be fully 3-connected for reliability. The 9.6Kbs lines will be upgraded to 50Kbs as the capacity requirement goes up and rapid expansion to 18 cities is planned as shown in Figure 4-16.

Telenet has chosen their hardware to provide for easy adaption of the ARPA Network IMP and TIP software which is in the public domain.

The ARPA technology will undergo some modification to better fit it to an operational environment.

a. for reliability IMPs and TIPs will be duplexed and will reside in a central office which will be manned and have backup power.

b. a Host interface to IBM computers which is native to them will be provided to avoid changes in the IBM Host software.

c. remote job entry terminals will be accommodated.

d. there will be a security option.

One thousand bit packets will be used with the charge basically independent of distance and based on the amount of data transmitted. This is in sharp distinction to the presently available phone services which are dependent on distance and for which one must pay a fixed lease charge regardless of the data transmitted. The charge which has been suggested is $1.25 per kilo-packet. A modest connect charge is also made: this amounts to $50/month for a leased 9.6Kbs line and $100/month for a leased 50Kbs line; for dial-in $1, $2, or $3.50 is charged depending on the speed of the line, up to 18000 baud, 2400 baud, or 4800 baud.

The best source for further information is the Telenet FCC filing [Telenet 73], although a good summary is provided in the Auerbach evaluation [Auerbach 74].

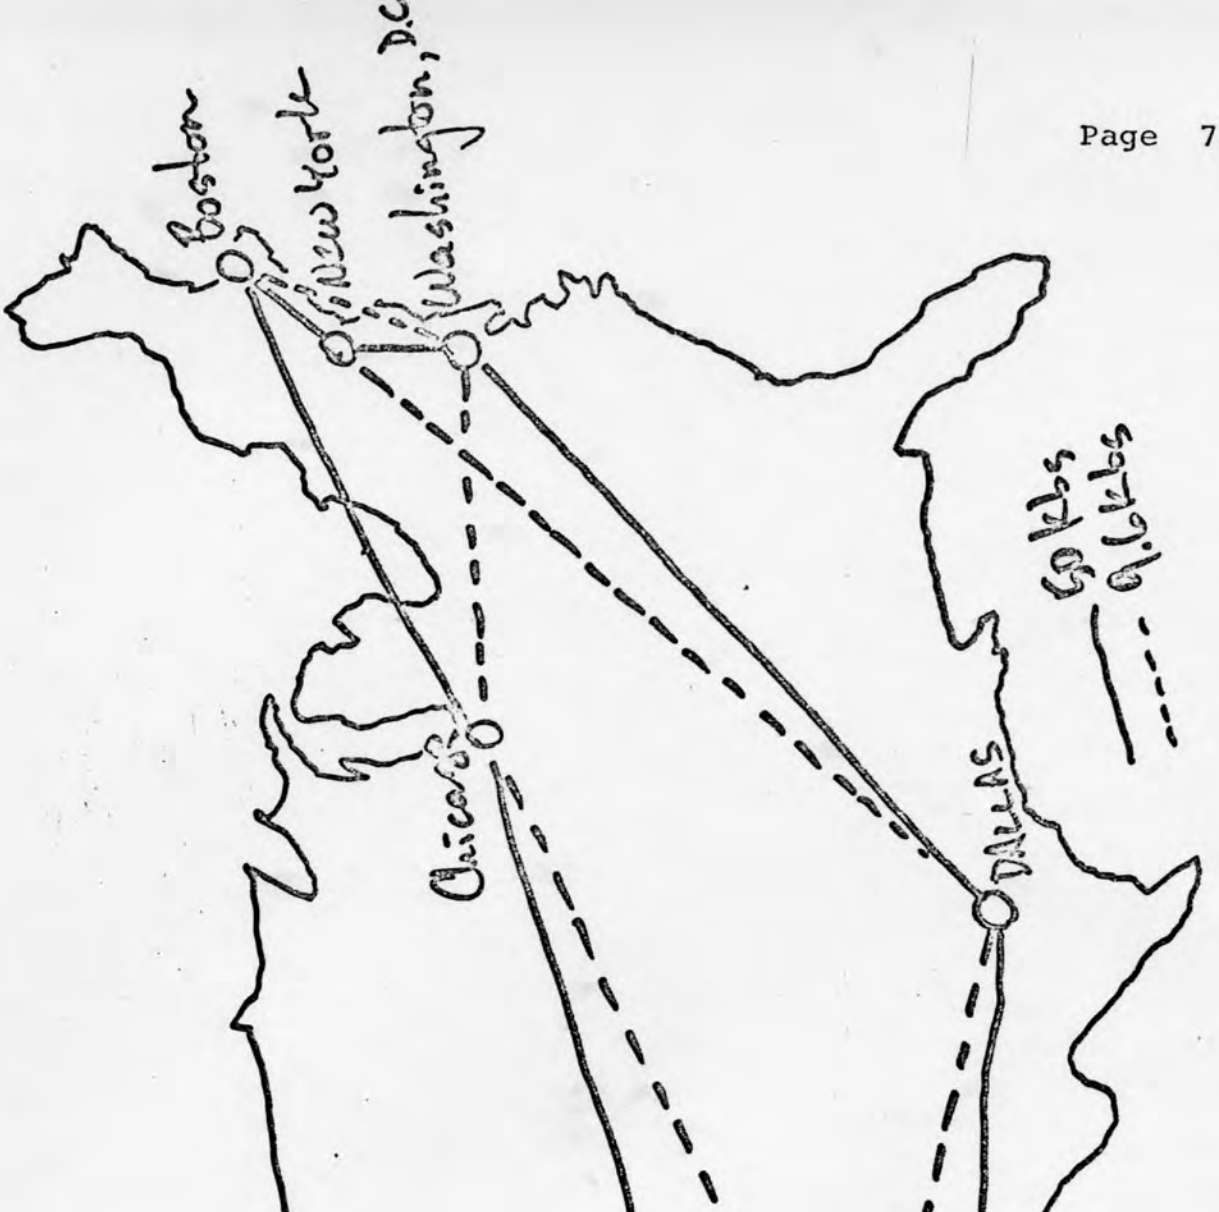

Figure 4-15 Telenet, Tentative Early Configuration

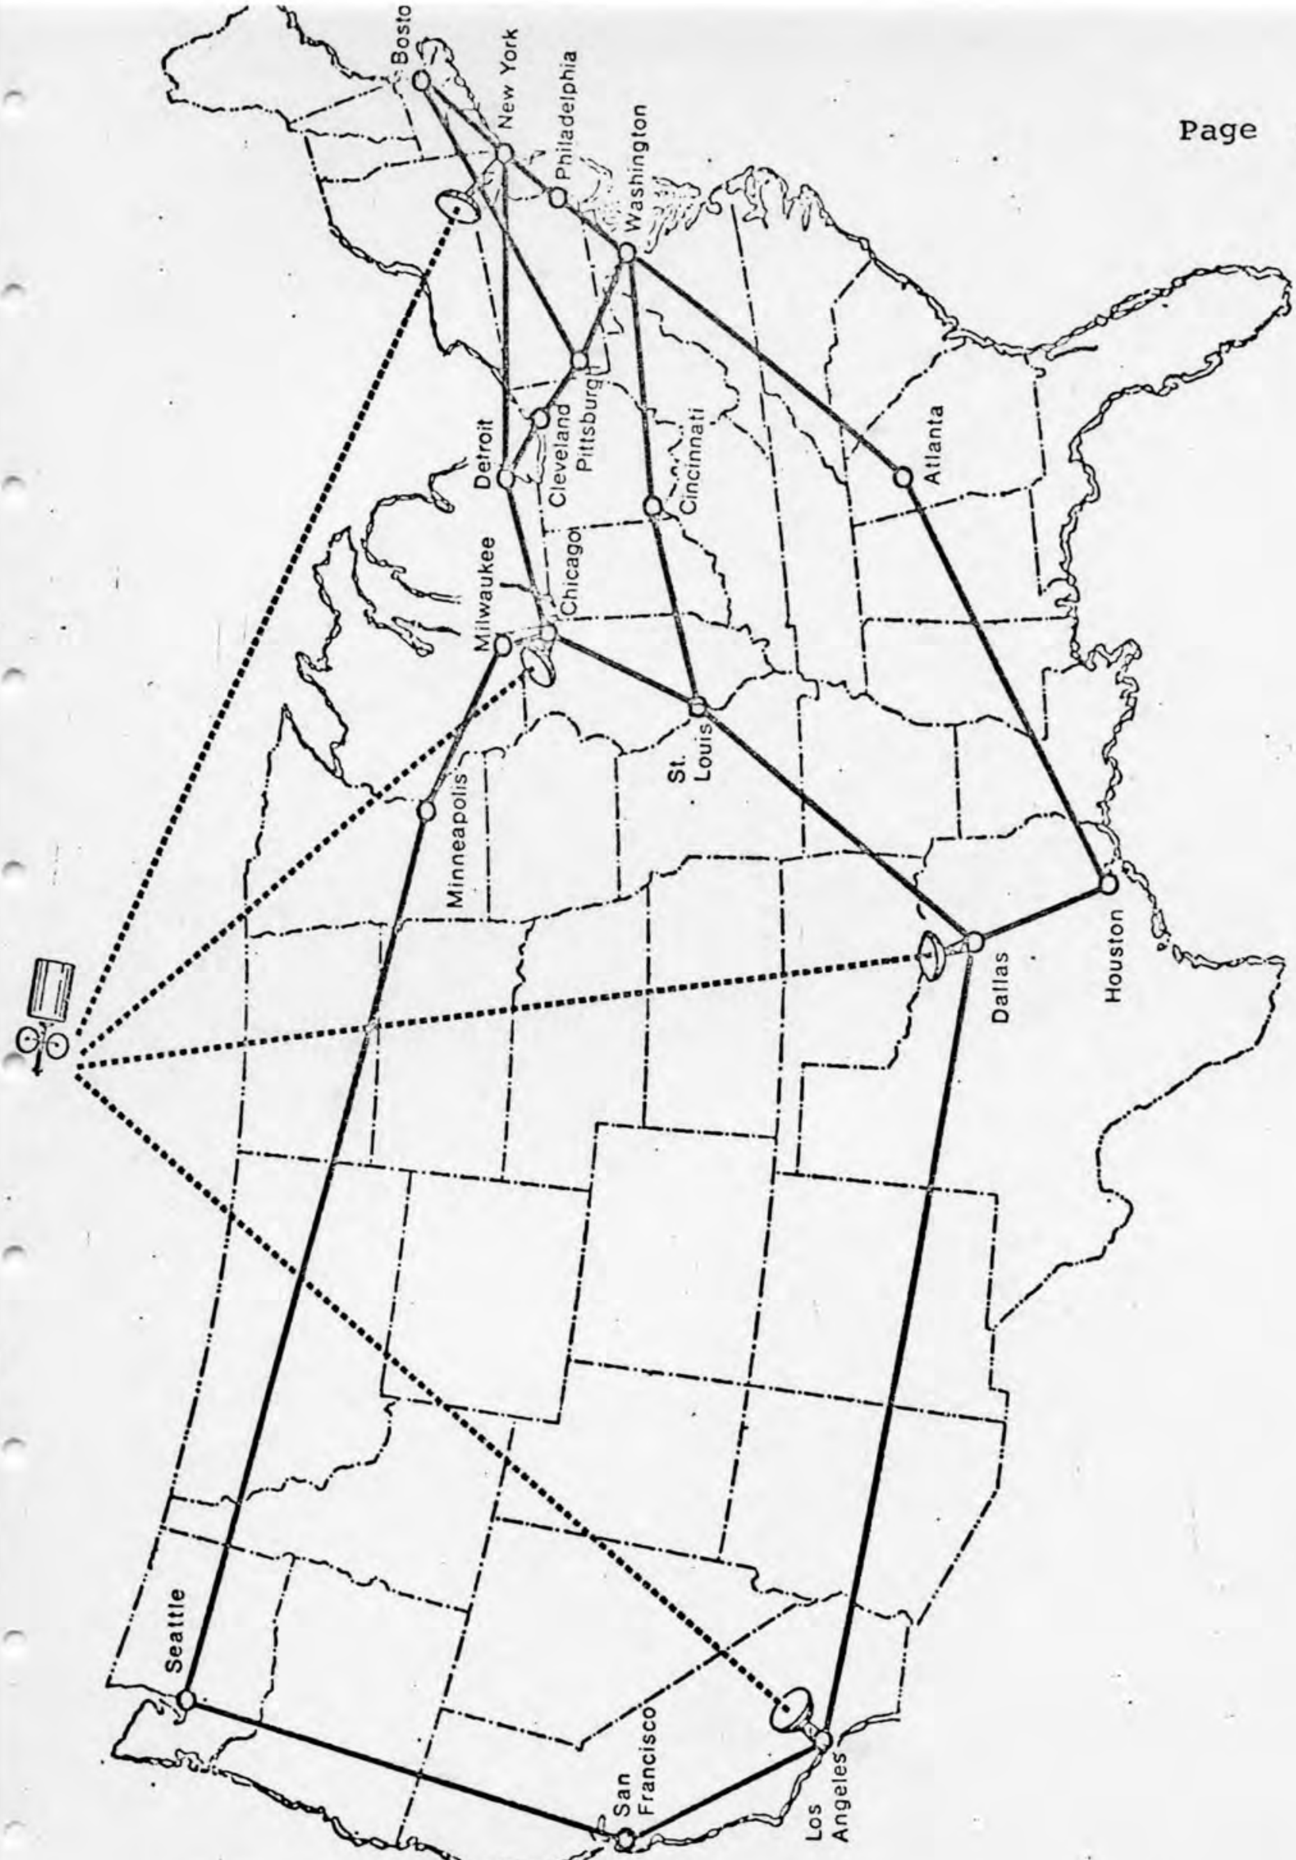

Figure 4-16 Telenet, 18 City Plan

4.14 Cyclades

The Cyclades network is the brainchild of Louis Pouzin of the French government laboratories called IRIA. The network is to have the form shown in Figure 4-17 in the not too distant future. Today, several nodes are operating to the point where communication through them is possible, and a very interesting experimental connection to the NPL network has been made (using the Zimmermann protocol).

Cyclades is based on a packet switching subnetwork called "Cigale". Cigale is an interesting contrast to the ARPA Network. Addressing is by Host rather than by network node. Provision is made for hierarchical addressing. Internal "Transfer stations" (Hosts) in the nodes are available to carry out special functions. The network deals only in packets; there are no messages. No packet ordering is maintained; packets are delivered as they arrive at the destination (the combination of this and Host addressing allows a Host to be simultaneously connected to multiple nodes). Explicit congestion control of an as yet undefined nature will be used. Nodes are allowed to discard packets if necessary. A dynamic routing algorithm will be used, probably based on least delay, but differing from the ARPA Network algorithm in that a packet will not be allowed to reverse direction on a path.

At the Host level Cyclades is again an interesting contrast to the ARPA Network.

a. In Cyclades, "transfer stations" are the equivalent to the NCPs of the ARPA network.

b. Multiple transfer stations are possible per Host and in fact the nodes themselves contain several.

c. Addressing is by a permanent name space of subscribers which has the form "(TS name") : "(local name") and a transient name space of ports.

d. Four types of communication are possible between transfer stations:

   i. letters -- no setup, optional acknowledgment, 2000 bits long, primarily for terminal to Host communication.

   ii. liaisons -- port to port setup, error control, flow control, no ordering, primarily for transfer of random access files which can just be written on the disk in the order they arrive.

iii. connections -- the same as liaisons but including ordering, primarily for transfer of sequential files.

iv. events -- short, out of band, interrupt letters.

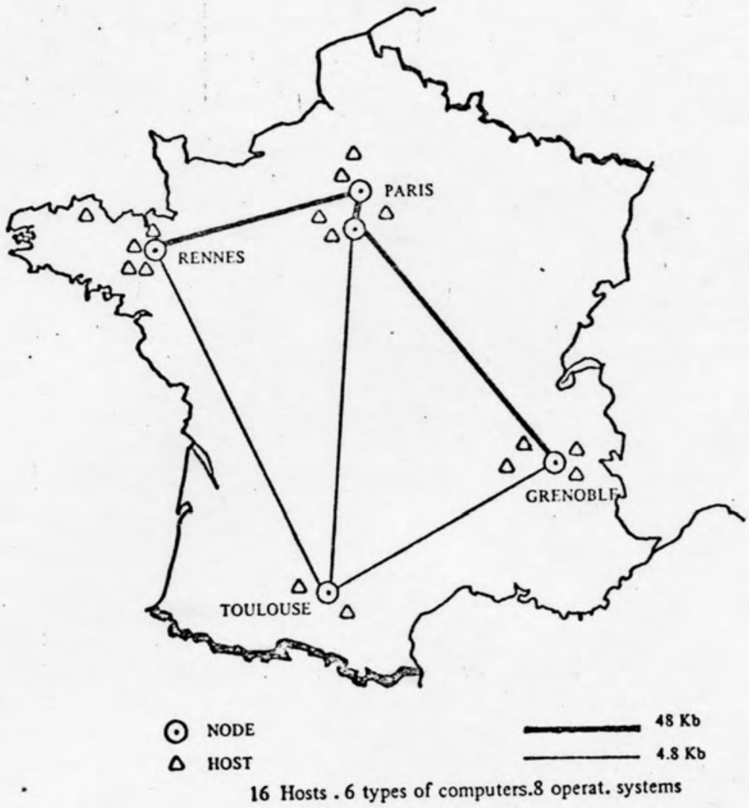

Figure 4-17 Cyclades

4.15 European Informatics Network (EIN)

Construction on EIN has begun (see Figure 4-18). Specifications call for nodes to be operational (at NPL and in Paris) in September 1975. Six to twelve months later, five nodes are to be operational. Possibly later the network may expand to 20 nodes.

The computer choice has been part technical, part economical, and all political. This will possibly result in a computer which is all wrong from the point of view of control and operations as well as one possibly unsuited technically for the job.

EIN is a blend of the ARPA Network and Cyclades. There are no messages (a hobbyhorse of Pouzin). Packets are 2000 bits long (a hobbyhorse of Davies). The nodes do store and forward. All Hosts are connected via communication circuits as are the ARPA Very Distant Hosts. Ordering is optional. There is no storage control: no copy at the source, no reserved space at the destination, and packets are lost if there is no space. Hosts can connect to multiple nodes, but then the ordering option cannot be selected. Adaptive routing based on delay will be used. The inter-node acknowledgment scheme is that used in the ARPA Network. Each node can support 4 lines to other nodes and eight Hosts.

References to EIN include [Barber 72] and [Barber 74].

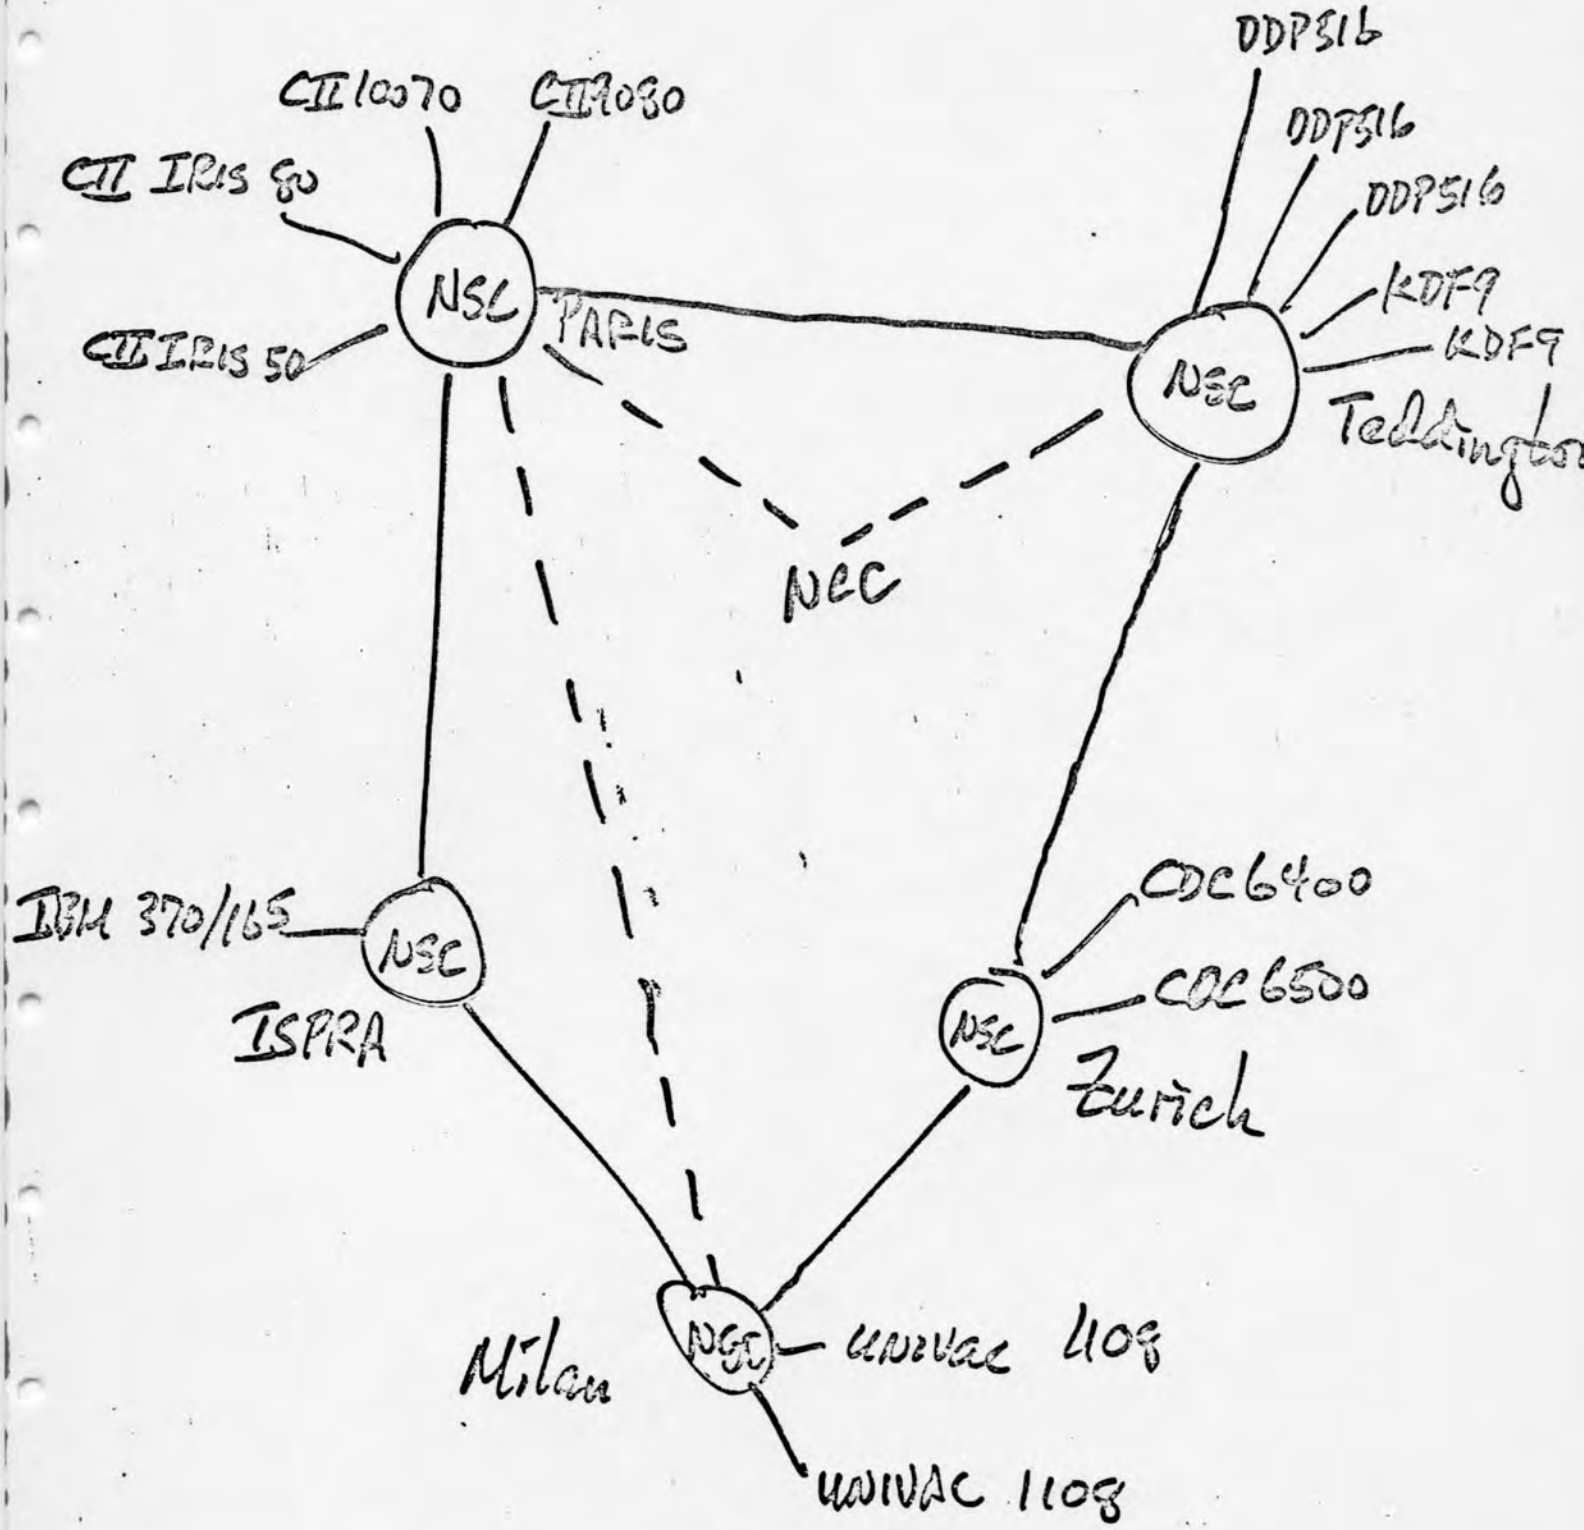

Figure 4-18 European Informatics Network

4.16 Proposed Commercial and Government Networks

The major point to be made at this time about commercial and government packet-switching networks is that the bandwagon seems to be pulling out and everybody is climbing on board. ARPA has sold is customer, the rest of the Defence Department, on packet-switching it appears. The commercial world is beginning to demand packet switching. And the national communications authorities are beginning to take notice despite the entrenchment of the old-line telephony people who appear generally unsold on packet-switching.

To be a bit more specific, industrial outfits the size of General Motors are taking a serious look at packet switching. When we looked at their communication needs several years ago, they had some 300 big! computer systems to connect together, if memory serves. The British Steel Corporation is currently procuring a big packet switching network to connect their computing resources (valued at over 50 million pounds per year) together. We have reason to believe that several of the U.S.'s biggest banks are seriously considering packet switching.

Within the U.S. government, the National Security Agency is building a small packet switching network. SAC is considering a large packet-switching network, to be called SATIN 4. The Defence Communications Agency is considering a large packet-switching network to be called AUTODIN II. We recently saw a request for proposal for a study for the U.S. Navy which hinted at the eventual possibility of a packet-switch on every! Navy ship.

We have already mentioned the U.S. common carriers who are attempting to offer packet-switching service. In other countries, the state equivalents of the U.S. carriers are also becoming active in packet-switching. The British Post Office has its EPSS experimental system. Bell Northern Research is building a packet-switching network. The French PTT (postal, telegraph, and telephone authority, the common abbreviation for any of the foreign governmental communication authorities) is building an experimental packet switching network. Norway built an experimental packet switching network several years ago.

Additionally, on a multi-national note, we have already mentioned EIN, the common market network; CCITT (the U.N.-like confederation of national communication authorities is studying standards for packet-switching; and IFIP (the international computer organization) has a study group considering connecting many separate packet-switching networks together.

We will probably mention a number of the above mentioned packet-switching efforts again later, either for purposes of comparison with our major example, the ARPA Network.

The point that should be remembered is that packet-switching appears to be a big thing, perhaps even more an "in" thing in the 70's than time-sharing was in the 60's. Further, the size of the packet-switching networks that are being discussed are frequently enormous, either in physical size or dollar value and most often both. The latter point, the large amounts of money that will be involved, guarantees great interest in packet-switching on the part of the computer and communications industries. Everybody will want a piece of the action, and they are beginning to fight for "their share."

5. Economics of Packet Switching

We are not particularly expert on the subject of the economics of packet-switching. Therefore, this section will mainly consist of quotes and figures taken from the writings of others who have made statements on the economics of packet-switching. Although the logic of the statements frequently sound reasonable to us, we have not ourselves investigated enough to be able to stand behind these statements and assertions. Nonetheless, we believe that as examples of the kinds of statements that are made about packet-switching economics, the following will be interesting.

Roberts has been a frequent spokesman for the economics of packet-switching. In [Roberts 70] he said the following about cost:

> "To be a useful utility, it was felt that communications costs for the network should be less the 25 percent of the computing costs of the systems connected through the network. This is in contrast to the rising costs of remote access communications which often cost as much as the computing equipment.
>
> "If we examine why communications usually cost so much we find that it is not the communications channel per se, but our inefficient use of them, the switching costs, or the operations cost. To obtain a perspective on the price we commonly pay for communications let us evaluate a few methods. As an example, let us use a distance of 1400 miles since that is the average distance between pairs of nodes in the projected ARPA Network.* A useful measure of communications cost is the cost to move one million bit of information, cents/megabit. It is assumed for leased equipment and data set rental that the usage is eight hours per working day.

----

*The actual average distance between nodes in the ARPA Network is probably considerable shorter than this distance -- McQ and DCW.

Media		
Telegram	$3300.00	For 100 words at 30 bits/wd, daytime
Night Letter	565.00	For 100 words at 30 bits/wd, overnight delivery
Computer Console	374.00	18 baud avg. use[2], 300 baud DDD service line & data sets only
TELEX	204.00	50 baud teletype service
DDD (103A)	22.50	300 baud data sets, DDD daytime service
AUTODIN	8.20	2400 baud message service, full use during working hours
DDD (202)	3.45	2000 baud data sets
Letter	3.30	Airmail, 4 pages, 250 wds/pg, 30 bits/wd
W. U. Broadband	2.03	2400 baud service, full duplex
WATS	1.54	2000 baud, used 8 hrs/working day
Leased Line (201)	.57	2000 baud, commercial, full duplex
Data 50	.47	50 KB dial service, utilized full duplex
Leased Line (303)	.23	50 KB, commercial, full duplex
Mail DEC Tape	.20	2.5 megabit tape, airmail
Mail IBM Tape	.034	100 megabit tape, airmail

Figure 5-1 Cost per Megabit for Various Communications Media 1400 Mile Distance

"Special care has also been taken to minimize the cost of the multiplexor or switch. Previous store and forward systems like DoD's AUTODIN system, have had such complex, expensive switches that over 95 percent of the total communications service cost was for the switches. Other switch services adding to the system's cost, deemed superfluous in a computer network, were: long term message storage, multi-address messages and individual message accounting.

"The final cost criteria was to minimize the communications software development cost required at each node site. If the network software could be generated centrally, not only would the cost be significantly reduced, but also the reliability would be significantly enhanced."

Roberts also did some cost comparison with alternative network communications systems designs in [Roberts 70], including the following analysis of cost vs. delay for the various alternatives considered.

"For the purpose of this comparison the capacity required was set at 500baud to 1KB per node-pair. A minimal buffer for error checking and retransmission at every node is included in the cost of the systems.

"...

"The systems chosen for the comparison were fully interconnected 2.4KB and 19KB leased line systems, Data-50 the dial-up 50KB service, DDD the standard 2KB voice grade dial-up system, Star networks using 19KB and 50KB leased lines into a central switch, and the ARPA Network using 50KB leased lines.

"The graph (see Figure 5-2) shows the cost per megabit versus delay. The rectangle outlines the variation caused by a block size variation of 1 to 10 Kilobits and capacity requirement variation of 500 to 1000 baud. The dial-up systems were used in a way to minimize the line charges while keeping the delay as low as possible. The technique is to dial a system, then transmit the data accumulated during the dial-up (20 seconds for DDD, 30 seconds for Data-50). The dial-up systems are still very expensive and slow as compared with other alternatives. The costs of the ARPA Network are for optimally designed topologies. The 19KB Star was eliminated because the system saturated just below 1KB per node-pair which did not provide adequate growth potential though the cost was comparable to the ARPA Network. For the 50KB Star network the switch is assumed to be an average distance of 1300 miles from every node.

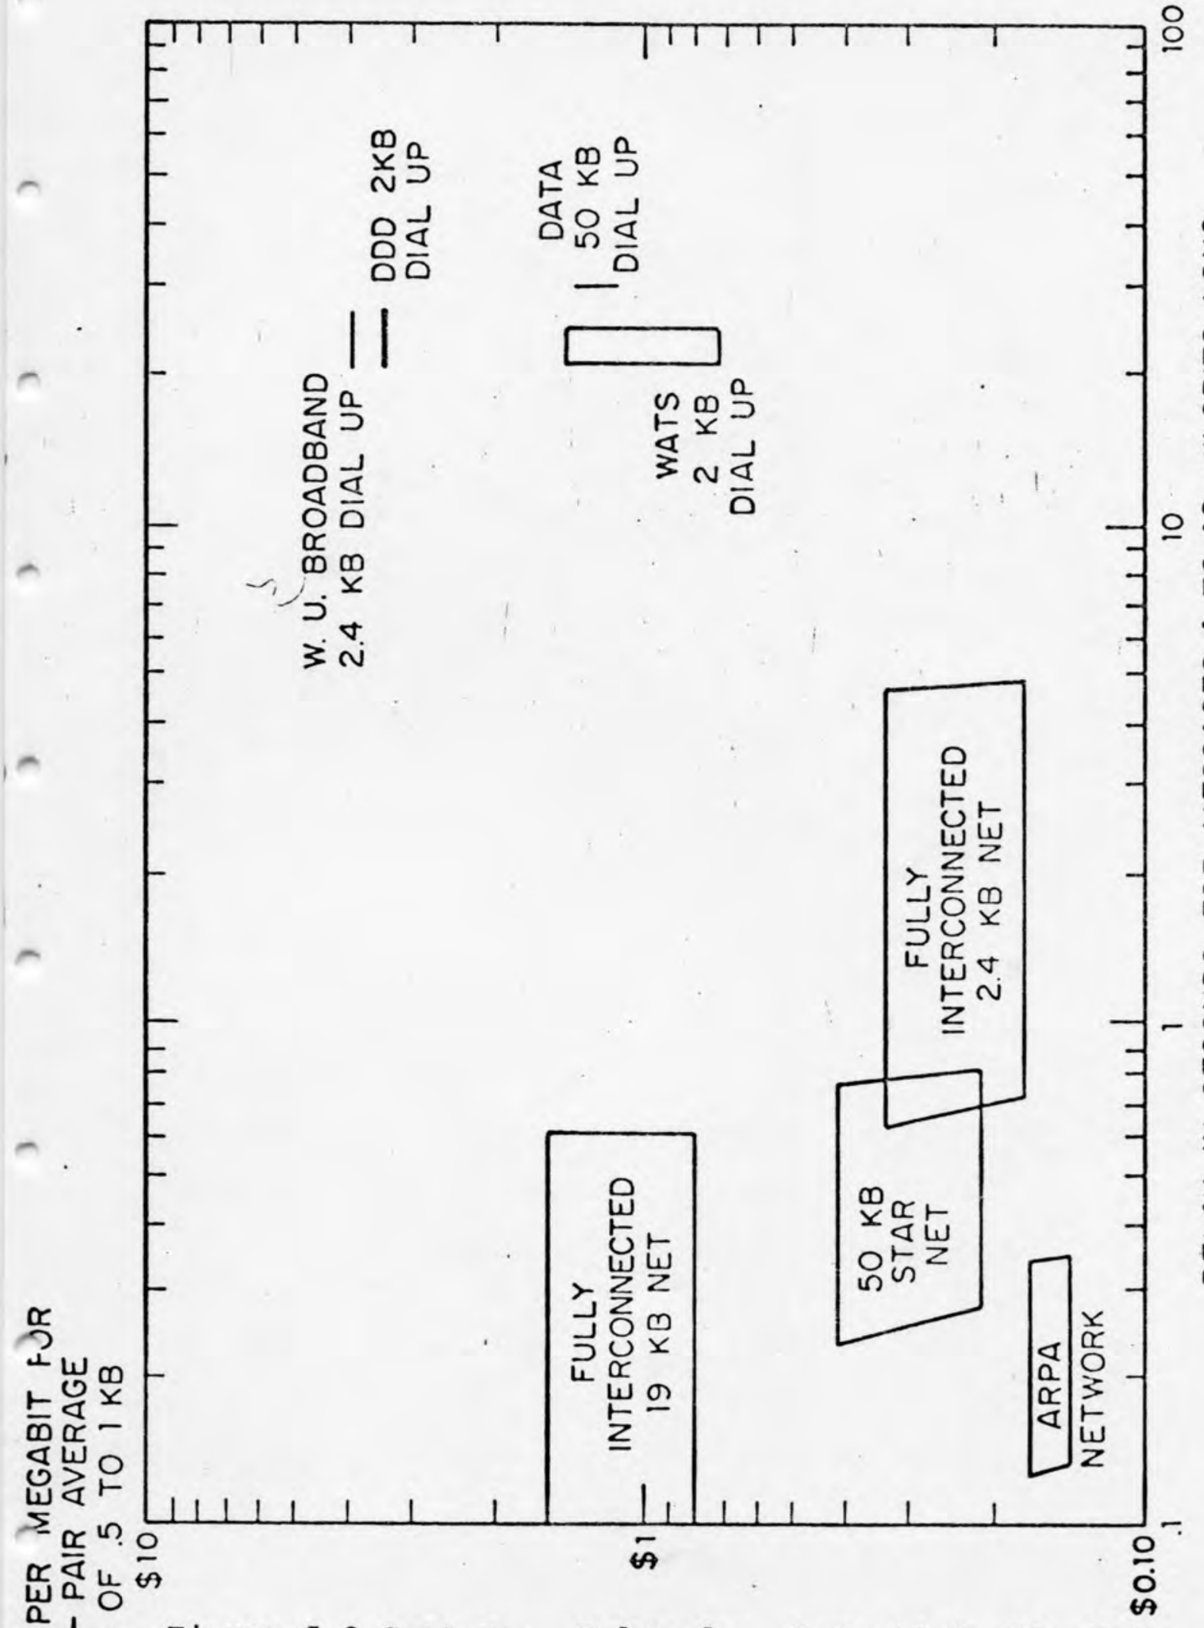

Figure 5-2 Cost vs. Delay for Potential 20 Node Network Designs

In [Roberts 72] can be found an analysis of how tariffs might be set up for a packet-switching service.

> "Our analysis indicated that network costs are surprisingly independent of many factors which one would initially assume to influence tariff rates. For example, we did a number of simulations of various traffic characteristics. We looked at users who only wanted to communicate with nearby machine, communications use that fell off rapidly with distance, as well as uniformly distributed communication anywhere in the country, and asked what the effect was on network costs. In other words if, through a tariff or some other means, the distribution of people's choice were skewed in terms of distance, did it affect our costs in providing that service? The answer was, 'No, it does not affect it at all.' For any kind of traffic distribution in terms of distance, we can design a network equally as good as the ARPA Network Design, with a slightly different topology but the same cost. It does not matter to us whether people send their messages across the country or to their neighbors. It does not make sense to introduce a distance-based tariff.
>
> "Second, we found that the network is independent of the traffic distribution pattern. That is, if most of the traffic comes from a subset of the nodes, and those were communicating in the main with another subset of nodes, and you introduced very strong constraints on the probabilities of communication between other nodes, the network would still work at the same cost per bit. It has the same throughput capability, within about five percent. This holds even with some nodes transmitting ten times as much traffic as others. Hence, it doesn't matter who generates the traffic or where it goes, we only need to look at the total network load and determine what the cost should be for this total amount of traffic. All we need to do is count bits. Actually, in our case, packets are the best thing to count.
>
> "Because we only need to count traffic, and it doesn't matter where it's going, we really don't need to ask a new customer whom he is going to talk to. It really doesn't matter. The distributive nature of the network accepts statistical loads in both time and space very well; their distribution doesn't seem to affect it.
>
> "The third point is that the network is independent of the peak rate. We have analysis of what would happen if a good percentage of the people in the

country wanted to join the ARPA Network but didn't care to pay for fifty-kilobit line interconnection. Suppose they only needed a peak rate which could be achieved by 9.6 kilobit lines, an effective rate of 15 kilobits. In a doubly connected network like the ARPA Network, the effective data rate for long transmissions is about eighty percent of twice the basic line rate. Thus the capacity of a network using fifty kilobit lines is about 80 kilobits, and that of a network with 9.6 kilobit lines is about 15 kilobits. We asked what would happen if half the people on the network decided to use 9.6 kilobits instead of fifty. The result was that, although they got worse service, the cost of the network was the same whether we put in 9.6 kilobit lines for them to fifty kilobit lines for everybody. A substantial number of fifty kilobit lines were needed in any case for the large service centers and other high-bandwidth users. If you have a fifty kilobit line going by a 9.6 kilobit customer you might just as well bring it in to his building rather than run a low-speed alternate line. Analyses made with a random distribution of these factors generally show that there is no variation in cost with varying distribution of peak rate requirements.

"The forth factor is that, within the effective traffic levels, cost is a linear function of traffic volume. Most of you have seen Figure 5-3 before, but it is useful to make this and my final points. Note that for capacities anywhere from zero to twenty kilobits per node average rate, the cost per node depends linearly on the traffic. This means you can charges a fixed price per bit regardless of the actual traffic volume.

"The fifth point is that network cost is largely independent of the network size beyond about forty nodes. The shaded area extends to about a hundred nodes. We have gone beyond that in analysis, but within that area, it's largely independent of the size of the network, once you have spanned the country and paid for that effectively through twenty to forty nodes.

"One primary conclusion can be drawn based on the foregoing analyses. Charging for network use should be purely on a cost per packet basis plus a connection charge. There are no other factors we need consider, and that makes life very simple."

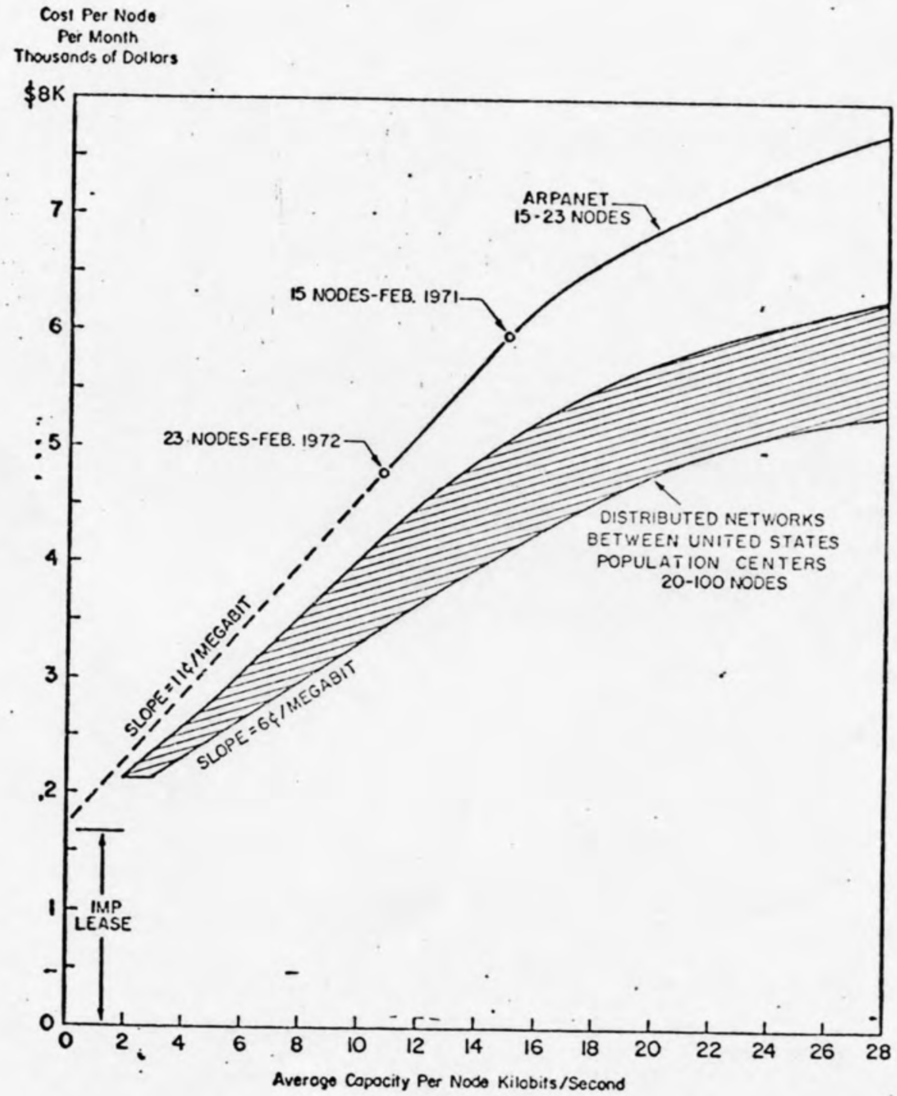

Figure 5-3 ARPA Network Cost vs Capacity

In [Roberts 73a] we find the following discussion of cost effectiveness.

> "In order to fully support computer resource sharing reliably and responsively on a nationwide basis, the communications network must be approximately th size and cost of the current ARPA Network (34 nodes and $2M/yr). Based on recent measurements the network traffic generated by a fully loaded, moderate size, time-shared computer is 720,000 packets per day. A minimal network such as the current ARPA Network has a basic capacity of 10M packets per day. Therefore, to fully utilize the basic capacity of the ARPA Network would require 14 moderate size computers to be fully accessed through the network. Additional capacity can easily be added beyond this point but with no great economy of scale. Figure 5-4 shows this effect for generalized national networks of ARPA Network technology. By utilizing the above measurement of network traffic produced by a host and estimating its rental at $720,000 per year, total network traffic can be related to the total computer resource value accessed via the network (each computer dollar produces 365 packets of traffic). Thus, the relative cost of network communication can be related to overall the annual computer value used. Since there are some fixed costs associated with adding each additional node, the cost-effectiveness depends partially on the number of nodes, but the main effect is produced by the dollar volume of usage. This means at least 10-20 million dollars of computer time usage must be expected before a nationwide network becomes optimally cost--effective. Once this activity level is reached, the main benefit from increased usage is improved reliability and increased peak throughput capability."

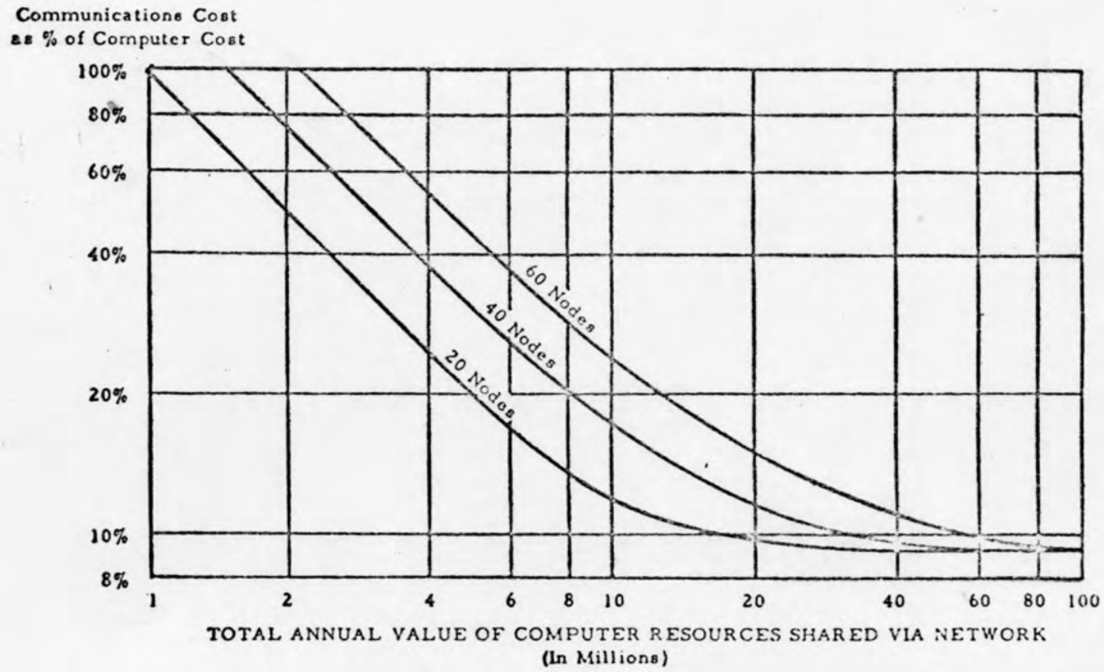

Figure 5-4 Cost - Effectiveness of Nationwide Networks

We have repeated some of the economic arguments that are made in favor of packet-switching by Roberts. Other relevant arguments are made in [Barber 71] and [Dunn 74].

We conclude this section by directing the readers attention to yet another paper by Roberts, [Roberts 74], in which he makes very compelling economic arguments for the concept of transmitting "data by the packet."

References

Abramson 70.
N. Abramson, "THE ALOHA SYSTEM -- Another Alternative for Computer Communications," Proceedings of the AFIPS 1970 FJCC, pp. 281-285.

Abramson 75.
N. Abramson, R. Binder, F. Kuo, Okinaka, and Wax, "ALOHA Packet Broadcasting, a Retrospect," to be presented at the AFIPS 1975 NCC&E, May 1975.

Alsberg 75.
P.A. Alsberg, "Distributed Processing on the ARPA Network -- Measurement of the Cost and Performance Tradeoffs for Numerical Tasks," Proceedings of the Eight Annual Hawaii International Conference of System Science, January 1975.

Akkoyunlu 74.
E.A. Akkoyunlu, A.J. Bernstein, and R.E. Schantz, "Interprocess Communication Facilities for Network Operating Systems," IEEE Computer, Vol. 7, No. 6, June 1974, pp. 46-55.

Aupperle 72.
E.M. Aupperle, "MERIT Computer Network: Hardware Considerations," Proceedings of the Courant Institute Symposium on Computer Networks, New York, 1970, R. Rustin (ed.), Prentice Hall, 1972, pp. 49-64.

Aupperle 73.
E.M. Aupperle, "MERIT Network Re-Examined," Proceedings of the Seventh Annual IEEE Computer Society International Conference, San Francisco, February 1973, pp. 25-30.

Auerbach 74.
Auerbach Publishers Inc., Public Packet Switching Networks, Data Processing Manual No. 3-08-04, 1974.

Baran 64a.
P. Baran, "On Distributed Communications: I. Introduction to Distributed Communications Networks," RAND Corp. Memo RM-3420-PR, August 1964, 37p.

Baran 64b.
P. Baran, "On Distributed Communications: IV. Priority, Precedence, and Overload," RAND Corp. Memo RM-3638-PR, August 1964, 63p.

Baran 64c.
P. Baran, "On Distributed Communications: V. History, Alternative Approaches, and Comparisons," RAND Corp. Memo RM-3097-PR August 1964, 51p.

Baran 64d.
P. Baran, "On Distributed Communications: VI. Mini-Cost Microwave," RAND Corp. Memo RM-3762-PR, August 1964, 101p.

Baran 64e.
P. Baran, "On Distributed Communications: VII. Tentative Engineering Specifications and Preliminary Design for a High-Data-Rate Distributed Network Switching Node," RAND Corp. Memo RM-3763-PR, August 1964, 85p.

Baran 64f.
P. Baran, "On Distributed Communications: VIII. The Multiplexing Station," RAND Corp. Memo RM-3764-PR, August 1964, 103 p.

Baran 64g.
P. Baran, "On Distributed Communications: IX. Security, Secrecy, and Tamper-Free Considerations," RAND Corp. Memo RM-3765-PR, August 1964, 39p.

Baran 64h.
P. Baran, "On Distributed Communications: X. Cost Analysis," RAND Corp. Memo RM-3766-PR, August 1964, 21p.

Baran 64i.
P. Baran, "On Distributed Communications: XI. Summary Overview," RAND Corp. Memo RM-3767-PR, August 1964, 23p.

Baran 64j.
P. Baran, "On Distributed Communication Networks," IEEE Transactions on Communications Systems Vol. CS-12, March 1964, pp. 1-9.

Barber 71.
D.L.A. Barber and R.A. Scantlebury, "Current and Future Computer Network Options," Infotech State of the Art Report 6: Computer Networks, 1971, pp. 365-406.

Barber 72.
D.L.A. Barber, "The European Computer Network," Proceedings of the First International Conference on Computer Communication, October 1972, pp. 192-200.

Barber 74.
D.L.A. Barber (ed.), "A Specification for a European Informatics Network," Co-Operation Europeenne dans le Domaine de la Recherche Scientifique et Technique, January 4, 1974.

Bartlett 69.
K.A. Bartlett, R.A. Scantlebury, and P.T. Wilkinson, "A Note on Reliable Full Duplex Transmission Over Half Duplex Lines," Communications of the ACM, Volume 12, No. 5, May 1969, pp. 260-261.

BBN 74
"Specification of the Interconnection of a Host and an IMP," Bolt Beranek and Newman Report 1822, December 1974 revision.

Belsnes 74
D. Belsnes, "Flow Control in Packet Switching Networks," INWG Note No. 63, October 1974.

Bernard 73.
D. Bernard, "Intercomputer Networks: An Overview and a Bibliography," University of Pennsylvania, May 1973, available from the National Technical Information Service, AD769232.

Boehm 64.
S. Boehm and P. Baran, "On Distributed Communications: II. Digital Simulation of Hot-Potato Routing in a Broadband Distributed Communications Network," RAND Corp. Memo RM-3103-PR, August 1964, 49p.

Boehm 66a.
B. W. Boehm and R. L. Mobley, "Adaptive Routing Techniques for Distributed Communications Systems," RAND Corp. Memo RM-4781-PR, February 1966, 78p.

Boehm 66b.
B. W. Boehm and R. L. Mobley, "A Computer Simulation of Adaptive Routing Techniques for Distributed Communications Systems," RAND Corp. Memo RM-4782-PR, February 1966, 35p.

Boehm 69.
B. W. Boehm and R. L. Mobley, "Adaptive Routing Techniques for Distributed Communications Systems," IEEE Transactions on Communication Technology, Vol. COM-17, No. 3, June 1969, pp. 340-349.

Brant 72.
G.J. Brant and G.J. Chretien, "Methods to Control and Operate a Message-Switching Network," Computer-Communications Network and Teletraffic, Polytechnic Press of the Polytechnical Institute of Brooklyn, Brooklyn, N.Y., 1972.

Brooks 68.
F.P. Brooks, J.K. Ferrell, and T.M. Gallie, "Organizational, Financial, and Political Aspects of a Three-University Computing Center," Proceedings of the IFIPS International Conference, Edinburgh, 1968, pp. 923-927.

Burchfiel 75.
J. Burchfiel, R. Tomlinson, and M. Beeler, "Functions and Structure of a Packet Radio Station, to be presented at the AFIPS 1975 NCC&E, May 1975.

Butterfield 73.
S.C. Butterfield, R.D. Rettberg, and D.C. Walden, "The Satellite IMP for the ARPA Network," Proceedings of the Seventh Annual Hawaii International Conference on System Sciences, Honolulu, Hawaii, January 1974, Computer Nets Supplement, pp. 70-73.

Cabledata 74a.

Cabledata 74b.
Cabledata Associates, Inc., "ARPANET Management Study: New Application Areas," Quarterly Technical Report for Contract F08606-74-C-0043, August 1974.

Calveley 71.
C.E. Calveley, "Cost Considerations of Proposed Computer Networks," Infotech State of the Art Report 6: Computer Networks, pp. 477-492.

Cerf 74a.
V. Cerf and R. Kahn, "A Protocol for Packet Network Inter-Communications," IEEE Transactions on Communications, Vol. COM-22, No. 5, May 1974, pp. 637-648.

Cerf 74b.
V. Cerf, "An Assessment of ARPANET Protocols," RFC 635, NIC 30489, April 1974, a limited number of copies available for the cost of reproduction and handling from INWG, c/o Prof. V. Cerf, Digital Systems Laboratory, Stanford, CA. 94305.

Chou 72.
W. Chou and H. Frank, "Routing Strategies for Computer Network Design," Proceedings of the Microwave Research Institute International Symposium on Computer-Communications Networks and Teletraffic, April 1972, pp. 301-309.

Chou 73a.
W. Chou, H. Frank, R. Van Slyke, "Simulation of Centralized Computer Communications Systems," Proceedings of the ACM Third Data Communications Symposium, Florida, November 1973, pp. 121-130.

Chou 73b.
W. Chou and A. Kershenbaum, "A Unified Algorithm for Designing Multidrop Teleprocessing Networks," Proceedings of the ACM Third Data Communications Symposium, Florida, November 1973, pp. 148-156.

Cocanower 72.
A. Cocanower, "MERIT Computer Network: Software Considerations," Proceedings of the Courant Institute Symposium on Computer Networks, New York, 1970. R. Rustin (ed.), Prentice Hall, 1972, pp. 65-78.

Crowther 73a.
W.R. Crowther, R.D. Rettberg, D.C. Walden, S.M. Ornstein, and F.E. Heart, "A System for Broadcast Communication: Reservation-ALOHA," Proceedings of the Sixth Hawaii International Conference on System Sciences, January 1973, pp. 371-374.

Crowther 73b.
W.R. Crowther, J.M. McQuillan, and D.C. Walden, "Reliability Issues in the ARPA Network," Proceedings of the ACM/IEEE Third Data Communications Symposium, November 1973, pp. 159-160.

Crowther 74.
W.R. Crowther, F.E. Heart, A.A. McKenzie, J.M. McQuillan, and D.C. Walden, "Network Design Issues," BBN Report No. 2918, November 1974, to be available from National Technical Information Service.

Crowther 75.
W.R. Crowther, F.E. Heart, A.A. McKenzie, J.M. McQuillan, and D.C. Walden, "Issues in Packet-Switching Network Design," to be presented at the AFIPS 1975 National Computer Conference, May 1975.

Davies 67.
D.W. Davies, K.A. Bartlett, R.A. Scantlebury, and P.T. Wilkinson, "A Digital Communication Network for Computers Giving Rapid Response at Remote Terminals," Proceedings of the ACM Symposium on Operating Systems Principles, October 1967.

Davies 71.
D.W. Davies, "The Control of Congestion in Packet Switching Networks," Proceedings of the Second ACM/IEEE Symposium on Problems in the Optimization of Data Communications Systems, October 1971, pp. 46-49; also in NPL Com. Sci. T.M. 56, September 1971, 7p.

Davies 73.
D.W. Davies and D.L.A. Barber, Communication Networks for Computers, London: John Wiley and Sons, 1973.

Despres 72.
R.F. Despres, "A Packet Switching Network with Graceful Saturated Operation," Proceedings of the First International Conference on Computer Communication, October 1972, pp. 345-351.

Dunn 74.
D.A. Dunn and M.J. Eric, "The Economics of Packet Switching," Technical Report CN74-17, Engineering-Economic Systems Department, Stanford University.

Farber 72a.
D.J. Farber and K.C. Larson, "The System Architecture of the Distributed Computer System - The Communications System," Proceedings of the Microwave Research Institute International Symposium on Computer-Communications Networks and Teletraffic, April 1972, pp. 21-27.

Farber 72b.
D.J. Farber and K.C. Larson, "The Structure of a Distributed Computing System - Software," Proceedings of the Microwave Research Institute International Symposium on Computer-Communications Networks and Teletraffic, April 1972, pp. 539-545.

Farber 72c.
D.J. Farber, "Networks: An Introduction," Datamation, April 1972, pp. 36-39.

Farber 72d.
D.J. Farber and F.R. Heinrich, "The Structure of a Distributed Computer System - The Distributed File System," Proceedings of the First International Conference on Computer Communication, October 1972, pp. 364-370.

Farber 73.
D.J. Farber, J. Feldman, F.R. Heinrich, M.D. Hopwood, K.C. Larson, D.C. Loomis, and L.A. Rowe, "The Distributed Computing System," Proceedings of the Seventh Annual IEEE Computer Society International Conference, San Francisco, February 1973, pp. 31-34.

Fletcher 73a.
J.G. Fletcher, "OCTOPUS Communications Structure," Proceedings of the Seventh Annual IEEE Computer Society International Conference, San Francisco, February 1973, pp. 21-24.

Fletcher 73b.
J.G. Fletcher, "OCTOPUS Software Security," Proceedings of the Seventh Annual IEEE Computer Society International Conference, San Francisco, February 1973, pp. 61-62.

Floyd 62.
R.W. Floyd, "Algorithm 97, Shortest Path," CACM 5(6), June 1962, p. 345.

Fralick 75.
S. Fralick, D. Brandin, and F. Kuo, "Digital Portable Terminals," to be presented at the AFIPS 1975 NCC&E, May 1975.

Frank 70.
H. Frank, I. T. Frisch, and W. Chou, "Topological Considerations in the Design of the ARPA Computer Network," AFIPS Conference Proceedings, Vol. 36, June 1970, pp. 581-585.

Frank 71a.
H. Frank, I. T. Frisch, R. Van Slyke, and W. S. Chou, "Optimal Design of Centralized Computer Networks," Networks, Vol. 1, no. 1, 1971, pp. 43-58.

Frank 71b.
H. Frank and W. Chou, "Routing in Computer Networks," Networks, Vol. 1, No. 2, 1971, pp. 99-112.

Frank 71c.
H. Frank, "Computer Network Design," IEEE International Convention Record, March 1971, pp. 220-221.

Frank 71d.
H. Frank and W. Chou, "Throughput in Computer-Communications Networks," International State of the Art Report No. 6: Computer Networks, Infotech Information Ltd. Maidenhead, Berkshire, England, 1971, pp. 493-512.

Davies 67.
D.W. Davies, K.A. Bartlett, R.A. Scantlebury, and P.T. Wilkinson, "A Digital Communication Network for Computers Giving Rapid Response at Remote Terminals," Proceedings of the ACM Symposium on Operating Systems Principles, October 1967.

Davies 71.
D.W. Davies, "The Control of Congestion in Packet Switching Networks," Proceedings of the Second ACM/IEEE Symposium on Problems in the Optimization of Data Communications Systems, October 1971, pp. 46-49; also in NPL Com. Sci. T.M. 56, September 1971, 7p.

Davies 73.
D.W. Davies and D.L.A. Barber, Communication Networks for Computers, London: John Wiley and Sons, 1973.

Despres 72.
R.F. Despres, "A Packet Switching Network with Graceful Saturated Operation," Proceedings of the First International Conference on Computer Communication, October 1972, pp. 345-351.

Dunn 74.
D.A. Dunn and M.J. Eric, "The Economics of Packet Switching," Technical Report CN74-17, Engineering-Economic Systems Department, Stanford University.

Farber 72a.
D.J. Farber and K.C. Larson, "The System Architecture of the Distributed Computer System - The Communications System," Proceedings of the Microwave Research Institute International Symposium on Computer-Communications Networks and Teletraffic, April 1972, pp. 21-27.

Farber 72b.
D.J. Farber and K.C. Larson, "The Structure of a Distributed Computing System - Software," Proceedings of the Microwave Research Institute International Symposium on Computer-Communications Networks and Teletraffic, April 1972, pp. 539-545.

Farber 72c.
D.J. Farber, "Networks: An Introduction," Datamation, April 1972, pp. 36-39.

Farber 72d.
D.J. Farber and F.R. Heinrich, "The Structure of a Distributed Computer System - The Distributed File System," Proceedings of the First International Conference on Computer Communication, October 1972, pp. 364-370.

Frank 72a,
H, Frank and I, T, Frisch, "The Design of Large-Scale Networks," Proceedings of the IEEE, Vol, 60, No, 1, January 1972, pp, 6-11,

Frank 72b,
H, Frank, R, E, Kahn, and L, Kleinrock, "Computer Communication Network Design - Experience with Theory and Practice," Proceedings SJCC 1972, pp, 255-270; also in Networks, Vol, 2, No, 2,1972, pp, 135-166; also in Advances in Computer Communications, W, W, Chu (ed,), Artech House Inc,, 1974, pp, 254-269,

Frank 72c,
H, Frank and W, Chou, "Topological Optimization of Computer Networks," Proceedings of the IEEE, Vol, 60, No, 11, November 1972, pp, 1385-1397,

Frank 72d,
H, Frank, "Optimal Design of Computer Networks," Proceedings of the Courant Institute Symposium on Computer Networks, New York, 1970, R, Rustin (Ed,), Prentice-Hall, 1972, pp, 167-183,

Frank 73,
H, Frank, "Providing Reliable Networks with Unreliable Components," ACM Third Data Communications Symposium, Florida, November 1973, pp, 161-164,

Frank 75,
H, Frank, R, Van Slyke, and I, Gitman, "Packet Radio Network Design," to be presented at the AFIPS 1975 NCC&E, May 1975,

Garrett 75,
Garrett and S, Fralick, "A Technology for Packet Radio," to be presented at the AFIPS 1975 NCC&E, May 1975,

Gerla 73,
M, Gerla, "Deterministic and Adaptive Routing Policies in Packet-Switched Computer Networks, Proceedings of the Third ACM Data Communications Symposium, November 1973, pp, 23-28,

Gitman 74,
I, Gitman, R, M, Van Slyke, and H, Frank, "On Splitting Random Accessed Broadcast Communication Channels," Proceedings of the Seventh Hawaii International Conference on System Sciences, January 1974, 5p,

Heart 70.
F.E. Heart, 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, June 1970, pp. 551-567; also in Advances in Computer Communications, W.W. Chu (ed.), Artech House Inc., 1974, pp. 300-316.

Heart 73.
F.E. Heart, S.M. Ornstein, W.R. Crowther, and W.B. Barker, "A New Minicomputer/Multiprocessor for the ARPA Network," AFIPS Conference Proceedings, Vol. 42, June 1973, pp. 529-537; also in Selected Papers: International Advanced Study Institute, Computer Communication Networks, R.L. Grimsdale and F.F. Kuo (eds.), University of Sussex, Brighton, England, September 1973; also in Advances in Computer Communications, W.W. Chu (ed.) Artech House Inc., 1974, pp. 329-337.

Heart 74.
F.E. Heart, "Implications of the Computer-Communication Partnership," to appear in the proceedings of MEDINFO 74, Stockholm, August 1974.

Hench 72.
R.R. Hench and D.F. Foster, "Toward an Inclusive Information Network," AFIPS Conference Proceedings, Vol. 41, October 1972, pp. 1235-1241.

Herzog 72.
B. Herzog, "The MERIT Computer Network," Proceedings of the Courant Institute Symposium on Computer Networks, New York, 1970, R. Rustin (ed.), Prentice Hall, 1972, pp. 45-48.

Kahn 71.
R.E. Kahn and W.R. Crowther, "Flow Control in a Resource-Sharing Computer Network," Proceedings of the Second ACM/IEEE Symposium on Problems in the Optimization of Data Communications Systems, Palo Alto, October 1971, pp. 108-116; also in IEEE Transactions on Communications, Vol. COM-20, No. 3, Part II, June 1972, pp. 539-546; also in Advances in Computer Communications, W. W. Chu (ed.), Artech House Inc., 1974, pp. 230-237.

Kahn 72.
R.E. Kahn, "Resource-Sharing Computer Networks," Proceedings of the IEEE, Vol. 60, No. 11, November 1972, pp. 1397-1407; also in "Advances in Computer Communications," W.W. Chu (ed.), Artech House Inc., 1974, pp. 208-218.

Kahn 75.
R.E. Kahn, "The Organization of Computer Resources into a Packet Radio Network," to be presented at the AFIPS 1975 NCC&E, May 1975.

Kershenbaum 72.
A. Kershenbaum and R. Van Slyke, "Computing Minimum Spanning Trees Efficiently," Proceedings of the ACM Annual Conference, August 1972, pp. 518-527.

Kleinrock 64.
L. Kleinrock, "Communications Nets: Stochastic Message Flow and Delay," Dover Publications, New York, 1964, 209p.

Kleinrock 70.
L. Kleinrock, "Analytic and Simulation Methods in Computer Network Design," AFIPS Conference Proceedings, Vol. 36, June 1970, pp. 569-579.

Kleinrock 71.
L. Kleinrock and G. L. Fultz, "Adaptive Routing Techniques for Store-and-Forward Computer Communication Networks," International Computer State of the Art Report No. 6; Computer Networks, Infotech Information Ltd. Maidenhead, Berkshire, England, 1971, pp. 541-562; also in Proceedings of the International Conference on Communications, Montreal, Quebec, Canada, June 1971, pp. 39.1-39.8.

Kleinrock 74.
L. Kleinrock and W. Naylor, "On Measured Behavior of the ARPA Network," AFIPS Conference Proceedings, Vol. 43, May 1974, pp. 767-780.

Kleinrock 75.
L. Kleinrock and F. Tobagi, "Random Access Techniques for Packet Radio Networks," to be presented at the AFIPS 1975 NCC&E, May 1975.

Kuo 73.
F.F. Kuo and N. Abramson, "Some Advances in Radio Communications for Computers," Proceedings of the Seventh Annual IEEE Computer Society International Conference, San Francisco, February 1973, pp. 57-60.

Luther 72.
W.J. Luther, "Conceptual Bases of CYBERNET," Proceedings of the Courant Institute Symposium on Computer Networks, New York, 1970, R. Rustin (ed.), Prentice Hall, 1972, pp. 111-146.

McKenzie 72.
A.A. McKenzie, B.P. Cosell, J.M. McQuillan, and M.J. Thrope, "The Network Control Center for the ARPA Network," Proceedings of the First International Conference on Computer Communication, Washington, D.C., October 1972, pp. 185-191.

McKenzie 74.
A.A. McKenzie, "Some Computer Network Interconnection Issues," Proceedings of AFIPS National Computer Conference and Exposition, Chicago, Illinois, May 1974, pp. 857-859.

McQuillan 72.
J.M. McQuillan, W.R. Crowther, B.P. Cosell, D.C. Walden, and F.E. Heart, "Improvements in the Design and Performance of the ARPA Network," AFIPS Conference Proceedings, Vol. 41, December 1972, pp. 741-754.

McQuillan 73.
J. M. McQuillan, "Throughput in the ARPA Network - Analysis and Measurement," BBN Report 2491, January 1973, 34p.

McQuillan 74a.
J.M. McQuillan, "Adaptive Routing Algorithms for Distributed Computer Networks," BBN Report No. 2831, May 1974, available from the National Technical Information Service, AD781467.

McQuillan 74b.
J.M. McQuillan, "The Evolution of Message Processing Techniques in the ARPA Network," to appear in International Computer State of the Art Report No. 24: Network Systems and Software, Infotech, Maidenhead, England.

McQuillan 74c.
J.M. McQuillan, "Design Considerations for Routing Algorithms in Computer Networks," Proceedings of the Seventh Annual Hawaii International Conference on System Sciences, Honolulu, Hawaii, January 1974, pp. 22-24.

Marill 66.
T. Marill and L.G. Roberts, Toward a Cooperative Network of Time-Shared Computers," AFIPS Conference Proceedings, Vol. 29, November 1966, pp. 425-431.

Metcalfe 73.
R.M. Metcalfe, "Packet Communication," Massachusetts Institute of Technology Project MAC Report MAC TR-114, December 1973.

Mimno 73.
N.W. Mimno, B.P. Cosell, D.C. Walden, S.C. Butterfield, and J.B. Levin, "Terminal Access to the ARPA Network -- Experience and Improvements," Proceedings of the Seventh Annual IEEE Computer Society International Conference, San Francisco, February 1973, pp. 39-43.

NAC 70a.
NAC First Semiannual Technical Report for the Project, "Analysis and Optimization of Store-and-Forward Computer Networks," June 1970, 62p.

NAC 70b.
NAC Second Semiannual Technical Report for the Project, "Analysis and Optimization of Store-and-Forward Computer Networks," January 1971, 96p.

NAC 71a.
NAC Third Semiannual Technical Report for the Project, "Analysis and Optimization of Store-and-Forward Computer Networks," June 1971, 109p.

NAC 71b.
NAC Fourth Semiannual Technical Report for the Project, "Analysis and Optimization of Store-and-Forward Computer Networks," December 1971, 123p.

NAC 72a.
NAC Fifth Semiannual Technical Report for the Project, "Analysis and Optimization of Store-and-Forward Computer Networks," June 1972, 92p.

NAC 72b.
NAC Final Technical Report for the Project, "Analysis and Optimization of Store-and-Forward Computer Networks," October 1972, 120p.

Ornstein 72.
S.M. Ornstein, F.E. Heart, W.R. Crowther, S.B. Russell, H.K. Rising, and A. Michel, "The Terminal IMP for the ARPA Computer Network," AFIPS Conference Proceedings, Vol. 40, June 1972, pp. 243-254; also in "Advances in Computer Communications," W.W. Chu (ed.), Artech House Inc., 1974, pp. 317-328.

Ornstein 74.
S.M. Ornstein, W.B. Barker, R.D. Bressler, W.R. Crowther, F.E. Heart, M.F. Kraley, A. Michel, and M.J. Thrope, "The BBN Multiprocessor," Proceedings of the Seventh Annual Hawaii International Conference on System Sciences, Honolulu, Hawaii, January 1974, Computer Nets Supplement, pp. 92-95.

Ornstein 75.
S.M. Ornstein, W.R. Crowther, M.F. Kraley, R.D. Bressler, A. Michel, and F.E. Heart, "Pluribus -- A Reliable Multiprocessor," 1975 National Computer Conference and Exposition.

PCI 73.
Packet Communication Inc., "Before the Federal Communications Commission: In the Matter of Packet Communication Inc. Application for authority to institute and operate a packet-switched communication network in the Continental United States," January 23, 1973.

Pierce 72.
J.R. Pierce, "Network for Block Switching of Data," The Bell System Technical Journal, Vol. 51, July-August 1972, pp. 1133-1145.

Pouzin 73.
L. Pouzin, "Presentation and Major Design Aspects of the Cyclades Computer Network," Proceedings of the Third ACM Data Communications Symposium, November 1973, pp. 80-88.

Pouzin 74.
L. Pouzin, "Basic Elements of a Network Data Link Control Procedure (NDLC)," INWG 54, NIC 30375, January 1974, a limited number of copies available for the cost of reproduction and handling from INWG, c/o Prof. V. Cerf, Digital Systems Laboratory, Stanford, CA. 94305.

Roberts 70.
L.G. Roberts and B.D. Wessler, "Computer Network Development to Achieve Resource Sharing", AFIPS Conference Proceedings, Vol. 36, June 1970, pp. 543-549.

Roberts 71.
L.G. Roberts and B.D. Wessler, "The ARPA Network," NIC document 7750, May 1971.

Roberts 72.
L.G. Roberts, "ARPANET: Current Status, Future Plans," EDUCOM Spring Conference, April 1972, pp. 7-12.

Roberts 73a.
L.G. Roberts, "Network Rationale: A 5-Year Reevaluation," Proceedings COMPCON 1973, February 1973, pp. 3-6.

Roberts 73b.
L.G. Roberts, "Dynamic Allocation of Satellite Capacity through Packet Reservation," Proceedings AFIPS 1973 NCC&E, pp. 711-716.

Roberts 74.
L.G. Roberts, "Data by the Packet," IEEE Spectrum, February 1974, pp. 46-51.

Rosner 73.
R.D. Rosner, "A Digital Data Network Concept for the Defense Communications System," Proceedings of the National Telecommunications Conference, Atlanta, November 1973, pp. 22C1-6.

Schwartz 72.
M. Schwartz et al. "Terminal Oriented Computer-Communication Networks," Proceedings of the IEEE, Vol. 61, November 1972, pp. 1408-1422.

Smith 64.
J.W. Smith, "On Distributed Communications: III. Determination of Path-Lengths in a Distributed Network," RAND Corp. Memo RM-3578-PR, August 1964, 91p.

Telenet 73.
Telenet Communications Corporation, "Before the Federal Communications Commission: Application for a Public Packet Switched Data Communications Network," October 9, 1973.

Thomas 72.
R.H. Thomas and D.A. Henderson, "McRoss- A Multi-Computer Programming System," AFIPS Conference Proceedings, Vol. 40, June 1972, pp. 281-293.

Thomas 73.
R.H. Thomas, "A Resource Sharing Executive for the ARPANET," AFIPS Conference Proceedings, Vol. 42, June 1973, pp. 359-367.

Thomas 74.
R.H. Thomas, "ARPANET TENEX - A Step Toward a Network Operating SYstem," reprint of a presentation at 1974 National Computer Conference and Exposition Panel Session, May 1974.

Thomas 75.
R.H. Thomas, "JSYS Traps - A TENEX Mechanism for Encapsulation of User Processes," to be presented at the AFIPS 1975 NCC&E, May 1975.

Tomlinson 74.
R.S. Tomlinson, "Selecting Sequence Numbers," INWG - Protocol Note 2, August 1974, a limited number of copies available for the cost of reproduction and handling from INWG, c/o Prof. V. Cerf, Digital Systems Laboratory, Stanford, CA. 94305.

Van Slyk 71a.
R. Van Slyke and H. Frank, "Network Reliability Analysis: Part I," Networks, Vol. 1, No. 3, 1971, pp. 279-290.

Van Slyk 71b.
R. Van Slyke and H. Frank, "Reliability of Computer-Communication Networks," Proceedings of the 1971 Winter ACM Simulation Conference, December 1971, pp. 71-82.

Walden 72.
D.C. Walden, "A System for Interprocess Communication in a Resource-Sharing Computer Network," Communications of the ACM, Vol. 15, No. 4, April 1972, pp. 221-230; also in "Advances in Computer Communications," W.W. Chu (ed.), Artech House Inc., 1974, pp. 340-349.

Walden 74.
D.C. Walden, "Host-to-Host Protocols," to appear in International Computer State of the Art Report No. 24: Network Systems and Software, Infotech, Maidenhead, England.

Williams 72.
L.H. Williams, "A Functioning Computer Network for Higher Education in North Carolina," AFIPS Conference Proceedings, Vol. 41, October 1972, pp. 899-904.

ARPA NETWORK, LOGICAL MAP, JANUARY 1975

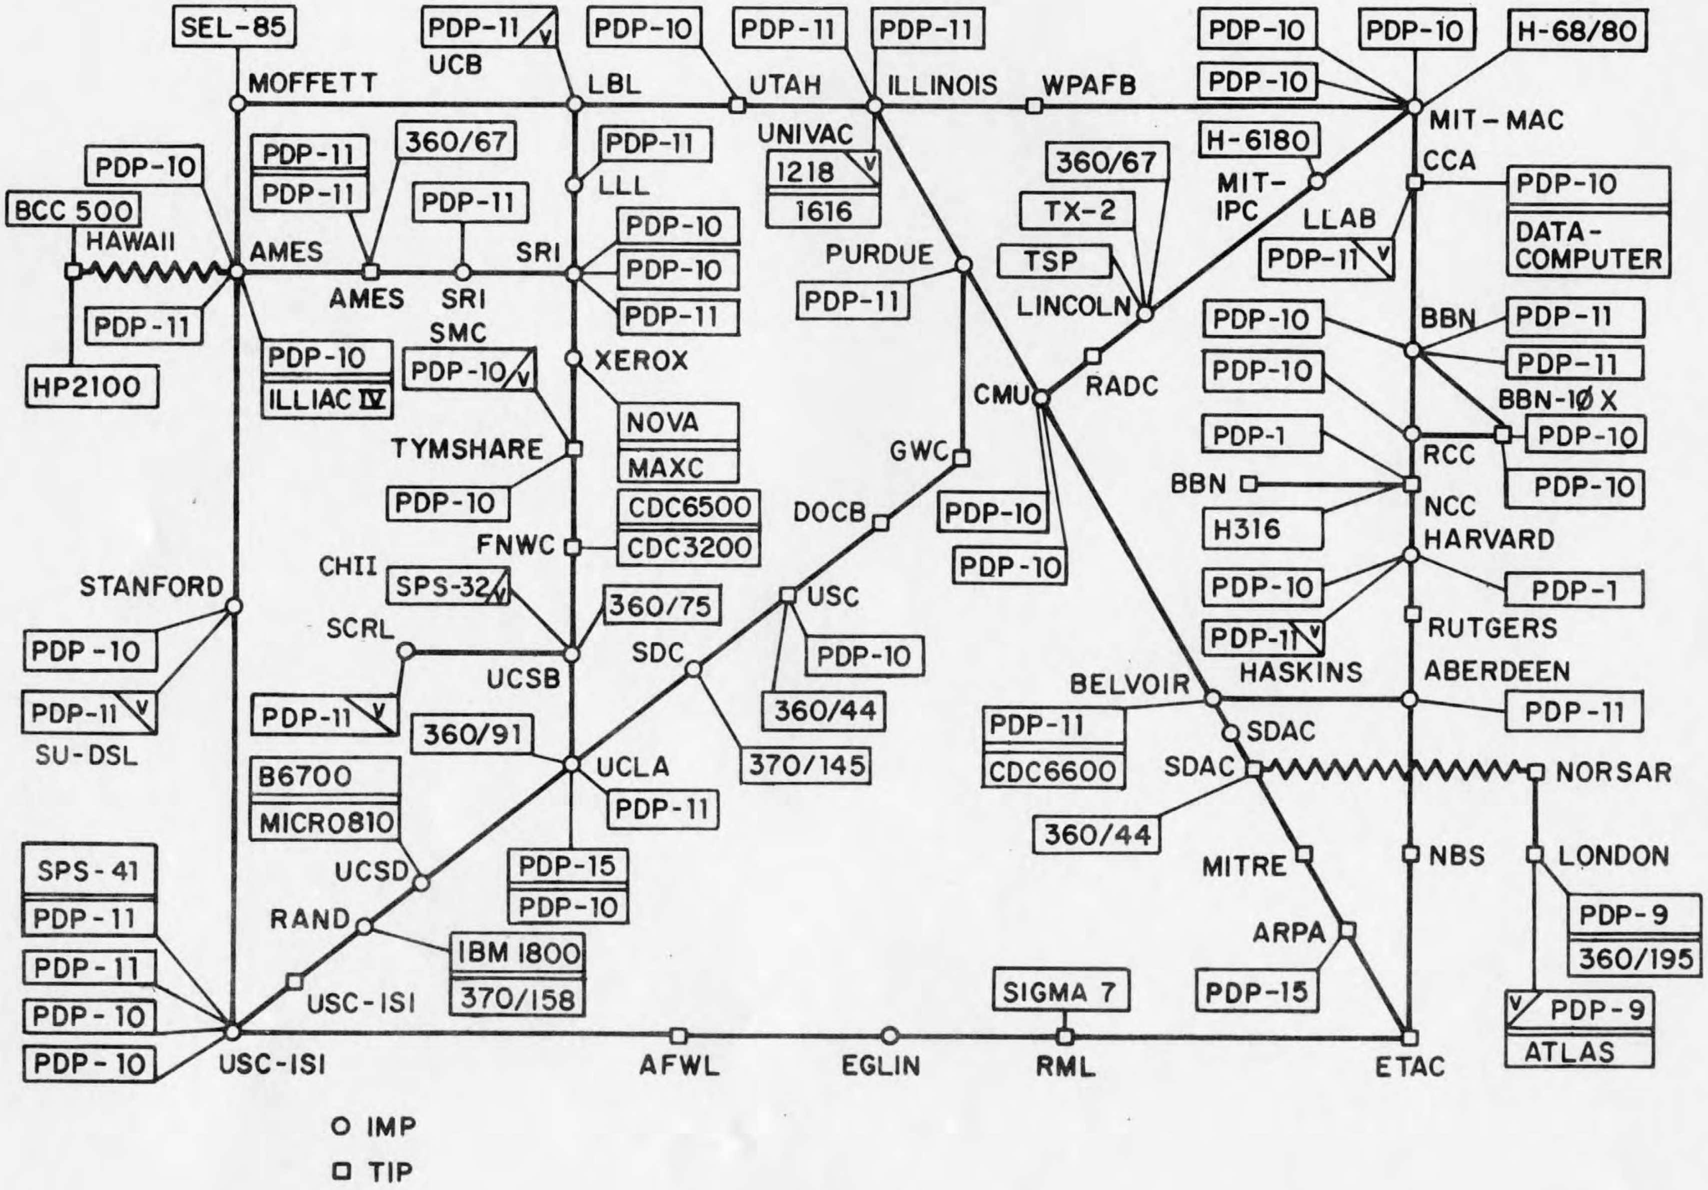

ARPA NETWORK, GEOGRAPHIC MAP

JANUARY 1975