The ARPANET & Computer Networks

Lawrence G. Roberts
NetExpress Inc.

In 1964 only large mainframe computers existed, each with its own separate set of users. If you were lucky the computer was timeshared, but even then you could not go far away since the terminals were hardwired to it or connected by local phone line. Moreover, if you wanted data from another computer you moved it by tape and you could forget wanting software from a different type of computer. Thus, most users were tied by their computer and terminal to a very restricted environment.

Today, in 1985, your terminal could well be a microcomputer networked with a very large, worldwide collection of other computers. You can obtain data and software from all these computers relatively easily with room for improvement) or, where convenient, use the software and data on its home computer by remote access, computer to computer.

This change, which has occurred over the past 20 years, is in part a massive and evolutionary change in computer technology, and in part a modest and revolutionary change in communications technology. The revolution in communications started with an experiment in computer networking, the ARPANET, and grew into a communications revolution called packet switching. Today virtually all the world is linked by packet switched communications service so that any terminal can access almost any computer in the world. This packet switched data network has grown up independent of the telephone network, but over the next 20 years the basic fabric supporting all switched services (data, telephone and video) appear likely to become converted to packet switching, completing the revolution.

History of Network Concepts

Going back to examine the history of computer networks, the first event for me took place in November 1964 at the Second Congress on the Information System Sciences in Hot Springs, Virginia. In informal discussions with J. C. R. Licklider, F. Corbato, and A. Perlis, I concluded that the most important problem in the computer field before us at that time was computer networking; the ability to access one computer from another easily and economically to permit resource sharing. That was a topic in which Licklider was very interested and his enthusiasm infected me. My interest was more toward the networking and communications issues rather than the computer language and compatibility issues which were foremost in Lick's mind. For at least the prior year, Licklider, who was then running the ARPA IPT Office (then called Command & Control Research), had been pursuing the concept of the "Intergalactic Computer Network," trying to define the problems and benefits resulting from computer networking. In any case, that Hot Springs discussion convinced me that I should change my carrier objectives to concentrate on computer networking and the related communications problems.

One year later, in 1965, a second important meeting took place at MIT. Donald Davies from the National Physical Laboratory in the U.K. was at MIT to give a seminar on time sharing. Licklider, Davies, and I discussed networking and the inadequacy of data communication facilities for both time sharing and networking. Davies reports that shortly after this meeting he was struck with the concept that a store and forward system for very short messages (now called packet switching) was the ideal communication system for interactive systems. He wrote about his ideas in a document entitled "Proposal for Development of a National Communication Service for On-Line Data Processing" which envisioned a communications network using trunk lines from 100K bits/sec in speed to 1.5 megabits/sec (T1), message sizes of 128 bytes and a switch which could handle up to 10,000 messages/sec (Historical note: this took 20 years to accomplish). Then in June 1966, Davies wrote a second internal paper, "Proposal for a Digital Communication Network" in which he coined the word **packet**, a small sub part of the message the user wants to send, and also introduced the concept of an "interface computer" to sit between the user equipment and the packet network. His design also included the concept of a Packet Assembler and Disassembler (PAD) to interface character terminals, today a common element of most packet networks.

As a result of distributing his 1965 paper, Donald Davies was given a copy of an

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the
© 1986 ACM-0-89791-176-8-1/86-0051 \$00.75

publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

internal Rand report, "On Distributed Communications" by Paul Baran of the Rand Corporation, which had been written in August 1964 (1). Baran's historical paper also described a short message switching network using T1 trunks and a 128 byte message size but was oriented toward providing extremely reliable communications for secure voice and data in a military environment. In all, there were 11 reports written for the Air Force in the Rand Memorandum group, of which a couple were classified and unfortunately the others were very sparsely published in the scientific press, thus their impact on the actual development of packet switching was mainly supportive, not sparking its development - that happened independently at Rand, NPL and ARPA.

The First Network Experiment

Convinced that computer networking was important, the first task was to set up a test environment to determine where the problems were. Thus, in 1966, I set up two computer networks between Lincoln Laboratory's TX-2 computer and System Development Corporation's Q-32 computer using a 1200 bps dial channel (high speed those days). Each computer was operating in time-sharing mode and permitted any program to dial the other computer, log-in, and run programs much as it would execute a subroutine call. The experiment showed that there was no problem getting the computers to talk to each other and use resources on the other computer - time-sharing operating systems made that easy. The real problem uncovered was that dial communications based on the telephone network were too slow and unreliable to be operationally useful. This work, jointly authored with Tom Marill, was published in the AFIPS FJCC proceedings, Nov. 1966 (2). The lesson learned - **a new data communications network is needed in order to successfully network computers.**

ARPANET Development

The chance to develop and build a major computer network experiment based on radically new communications technology came within a few months. I was asked to take over the responsibility of the ARPA Information Processing Techniques (IPT) office and manage and build its programs. ARPA was sponsoring computer research at leading universities and research labs in the U.S. These projects and their computers provided an ideal environment for an experimental network project; consequently, the ARPANET was planned during 1967 with the aid of these researchers to link these projects' computers together. One task was to develop an computer interface protocol acceptable to all 16 research groups. A second task was to design a new communications network technology to support 35 computers at 16 sites with 500,000 packets/day traffic. The initial plan for the ARPANET was published in October 1967 at the ACM Symposium on Operating System Principles in Gatlinburg Tennessee (3). The reasons given at that time for establishing a computer network were:

- A. **Load Sharing:** Send program and data to remote computer to balance load.
- B. **Message Service:** Electronic mail service (mailbox service).
- C. **Data Sharing:** Remote access to data bases.
- D. **Program Sharing:** Send data, program remote, e.g. Supercomputer.
- E. **Remote Service:** Log-in to remote computer, use its programs and data.

The communications network design was that of the now conventional packet network; Interface Message Processors (IMP's) at each node interconnected by leased telecommunication lines providing a store and forward service on very short messages. The main difference from later packet nets was that the IMP's were located at the computer sites and connected by a short parallel cable rather than a communications line interface.

Also presented at the Gatlinburg Symposium was Donald Davies's first open publication on the NPL packet network concepts presented by Roger Scantlebury, "A Digital Communication Network for Computers Giving Rapid Response at Remote Terminals." (4) It detailed the concept of a high level packet net with high capacity nodal switches and interface computers in front of mainframe computers. This was the first time that either Davies or I knew anything about the work of each other since our 1965 contact. The NPL paper clearly impacted the ARPANET in several ways. The name "packet" was adopted, much higher speed was selected (50 Kilobit/sec vs. 2.4 Kilobit/sec) for internode lines to reduce delay and generally the NPL analysis helped confirm the concept of packet switching.

Another confirmation of the basic concepts came from finally being able to read the Rand reports on distributed communications. The Rand work was very detailed, since it covered the whole network including microwave and one valuable analysis on routing. Their hot-potato routing algorithm was a useful starting point for the ARPANET routing design.

During 1968, a request for proposal was let for the ARPANET packet switching IMP equipment and the operation of the packet network. The RFP was awarded to Bolt Beranek and Newman Inc. of Cambridge, Massachusetts, in January 1969. The RFP specified the general packet switching concept, packet size, and interface protocol so that bidders could not totally change the system concept, to circuit or message switching for example. The two largest computer companies to receive the RFP no bid it because they didn't have mini-computers with which to make an economic bid. BBN bid the Honeywell 516 mini-computer which was ideal for the task in 1969. Significant aspects of the network's internal operation, such as routing, flow control, software design, and network control were developed by a BBN team consisting of Frank Heart, Robert Kahn, Severo Ornstein, William Crowther, and David Walden. By December 1969 four nodes of the net had been installed and were operating effectively.

The first set of detailed papers covering the ARPANET were published in May 1970 at the AFIPS SJCC. (5-9) These papers reported the motivation and economics (5), the detailed design of the IMP (6), the network delay analysis and experience (7), the topological design programs and results (8), and the host-to-host protocol (9). These papers showed the world for the first time that **packet switching works**, that it is **economic** and that it is **reliable** and **virtually error free**. They also provided a complete description of how a working network was designed. As such these papers were the technical and motivational basis for many other network experiments around the world.

The ARPANET utilized minicomputers at every node to be served by the network, interconnected in a fully distributed fashion by 50 KB leased lines. Each minicomputer took blocks of data from the computers and terminals connected to it, subdivided them into 128 byte packets, and added a header specifying destination and source addresses; then, based on a dynamically updated routing table the minicomputer sent the packet over whichever free line was currently the fastest route toward the destination. Upon receiving a packet, the next minicomputer would acknowledge it and repeat the routing process independently. Thus, one important characteristic of the ARPANET was its completely distributed, dynamic routing algorithm on a packet-by-packet basis, based on a continuous evaluation within the network of the least delay paths, considering both line availability and queue lengths.

The technical and operational success of the ARPANET quickly demonstrated to a generally skeptical world that packet switching could be organized to provide an efficient and highly responsive interactive data communications facility. Fears that packets would loop forever and that very large buffer pools would be required were quickly allayed. Since the ARPANET was a public project connecting many major universities and research institutions, the implementation and performance details were widely published. (10, 11, 12, 13, 14). The work of Leonard Kleinrock and his associates at UCLA on the theory and measurement of the ARPANET has been of particular importance in providing a firm theoretical and practical understanding about the performance of packet networks.

The ARPANET was first demonstrated publicly at the first International Conference on Computer Communications (ICCC) in Washington, D.C. in October 1972. Robert Kahn of BBN organized the demonstration installing a complete ARPANET node at the conference hotel, with about 40 active terminals permitting access to dozens of computers all over the U. S. This public demonstration was, for many (if not most) of the ICCC attendees, proof that packet switching really worked. At this time, it was difficult for many experienced professionals to accept the fact that a collection of computers, wideband circuits, and minicomputer switching nodes (equipment totaling well over 100 pieces) could all function together reliably. The ARPANET demonstration lasted for three days and clearly displayed its reliable operation in public. The network provided highly reliable service to thousands of attendees during the entire duration of the conference.

Industry Reception

From the first time I distributed a description of packet switching outside the computer research community (the 1967 paper) until about 1975, the communication industry's reaction was generally negative since this was such a radically different approach. In some of the initial technical speeches I gave, communications professionals reacted with considerable anger and hostility, usually saying I did not know what I was talking about since I did not know all their jargon vocabulary. The most common technical flaw suggested (before the ARPANET was built) was that the buffers would quickly and catastrophically run out. After the ARPANET was operating successfully, their pitch changed to be that packet switching would never be economic without the government subsidy. Paul Baran reported the same reaction to his papers when he presented them; this reaction was the major reason his proposals never moved the military. Donald Davies reported a somewhat less angry response from the British Post Office, more one of mild interest but no serious consideration.

I learned a major lesson from that experience; **people hate to change the basic postulates upon which considerable knowledge has been built**. In the case of packet switching, the first postulate to change was the statistical nature of the traffic - data versus voice. The second was that computing was expensive. Some people find it is impossible to consider such a major jolt to their memory organization - they avoid it with putdowns if possible, if not with anger. Other people are more willing to reconsider, but for everyone it requires considerable effort. Those of us proposing packet switching all came from the computing field and did not need to change lots of prior concepts and knowledge. Many of those in the communications field still have not accepted packet switching. (The fight is heating up again as voice packet switching starts to be considered.)

ARPANET Growth

As soon as the first four nodes were brought up and tested in December 1969 the network grew very rapidly. One year later, in December 1970, the network had grown to 10 nodes and 19 host computers. By April 1971, there were 15 nodes with 23 host computers. The topology was as shown in Fig. 1. By this time, it was clear that connecting terminals directly to the network through a PAD-type device was important. Such a device was designed and built in 1970/1971, and the first Terminal Interface Processor (TIP) was added

The diagram illustrates the ARPANET network structure, showing the connections between various computer systems and IMPs (Interface Message Processors). The network is organized into two main horizontal rows of IMPs, connected by vertical lines. The top row includes IMPs for SRI, UTAH, ILLINOIS, MIT, LINCOLN, and CASE. The bottom row includes IMPs for UCSB, STANFORD, SDC, BBN, HARVARD, and BURROUGHS. Various computer systems are connected to these IMPs, including PDP-10, PDP-11, TX-2, GE-645, 360/75, 360/67, DDP-516, XDS SIGMA-7, 360/91, IBM 1800, 360/65, PDP-1, and ILLIAC IV.

to the network in Aug 1971. This permitted users with no computer to select a computer from all those around the country. In many cases having the user attach his terminal to a TIP and access even his own host(s) through the network was found to be more reliable. This was the start of a trend which today is almost the rule, **workstations should attach to a network, not a computer**. By January 1973, the network had grown to 35 nodes of which 15 were TIP's and was connected to 38 host computers. Fig. 2 shows a map of the network at this point. Network traffic had grown rapidly in 1972, from 100,000 packets/day to 1,000,000 packets/day, exceeding the original estimates I made in 1967. Also in 1973, the first satellite link was added to the network with a TIP in Hawaii; later in the year a pair of TIP's were added in Europe. In September 1973, the network was up to 40 nodes, 45 host computers with internode traffic of 2,900,000 packets/day, had clearly reached a stage of operational stability, heavy usage, and was by any measure a major success. It was at this point that I left ARPA to spread the technology to the commercial world.

The ARPANET has continued to grow since 1973, with 111 host computers in 1977 and more than 400 hosts in all the interconnected networks by 1983. The network has become a utility for both ARPA and the Department of Defense as a whole. Research has continued into internetting and packet radio but little change has occurred in the basic network.

1969 Crossover

In 1974, after it was clear that the technology worked, I finally spent time analyzing the economic trends which made packet switching possible.(15) These trends could have been analyzed in the mid-sixties if someone were to have asked the right questions, but as it often turns out they were examined after the event they could have predicted. Simply, since packet switching requires more computation than circuit switching but saves transmission bandwidth, the trend analysis looked at the cost of computation and the cost of communications and found the situation illustrated in Fig. 3. For packet switching the cost of computation dominated in the early 1960's but in 1969 the cost curves crossed and afterward the cost of communications dominated. The composite cost of packet switching thus fell below the cost of circuit switching also about 1969 and since then the margin of advantage has quickly widened to where it equals the full peak to average transmission ratio for data of about 8 to 1. Subsequently, in a 1982 paper (16) I did a similar analysis for voice and concluded the crossover there was in 1978 and therefore packet voice would have been cheaper. In a somewhat similar fashion to data, the industry may recognize this ten to fifteen years after the event.

FIGURE 3
**PACKET SWITCHING:
COST PERFORMANCE TRENDS**

Impact on Computer Resources

One of the original goals of the ARPANET was resource sharing to optimize computer resources. In 1973 I surveyed the ARPANET users to see what alternative computer facilities they would need if they did not have the network. In a number of cases there were ARPA projects which did not have a computer at all but utilized computer power from resource centers which had developed. They obtained access through a TIP, typically onsite, and used one of a collection of computers often accross the country. They were generally happy not to have to run their own computer and had far more reliable service than if they had been limited to one machine. Also due to the statistics of large groups and the benefit of multiple time zones (from Hawaii to Europe) the number of computers required for these users was far less through consolidation. Another group of users were using the Illiac IV remotely for large numerical problems and would have required their own supercomputer or perhaps had to commute if the network were not there. In other cases individuals needed to use software unique to foreign hardware and thus needed access to other computers besides their own. After finding rational solutions (fiscally sound) to all the ARPA users needs assuming no network, the computing cost was compared to the then actual computer cost and found to be **three times as expensive**. Independent of the details of the survey, it is clear that huge savings were possible (in that era of mainframes) by pooling computing demand. Today when all my computer resources come from micro-computers I couldn't save anything using remote computers for computing power. Supercomputer users still might save but in large part this benifit is mainly historical.

Other Network Development

CYCLADES: In France the interest in computer networks grew quickly during the early 1970's. In 1973 the first hosts were connected to a network called CYCLADES which linked several major computing centers throughout France. (17) The name CYCLADES refers to both the communications subnet and the host computers. The communications subnet by itself is called CIGALE and is a pure **datagram** network, moving only disconnected packets and delivering them in whatever order they arrive without any knowledge or concept of messages, connections or flow control. The ARPANET also operates using datagrams but perhaps the most avid supporter of the concept is the designer of CYCLADES, Louis Pouzin. As with any datagram network, a large part of the communications funtions must be implemented in the host computers; the packet ordering, message formation, flow control, and virtual connections support.

RCP: Another packet network experiment was started in France at about the same time by the French PTT Administration (18). This network, RCP, first became operational in 1974 and was the experimental testbed for the French public network service TRANSPAC. The design of RCP, directed by Remi Despes, differed sharply from that of CIGALE and ARPANET being organized around the concept of the **virtual circuits** rather than datagrams. RCP's character as a prototype public network may have been a strong factor in this difference since a vitrual circuit service is more directly marketable, not requiring substantial modifications to the customers' computer. In any case, RCP pioneered the incorporation of individually flow-controlled virtual circuits into packet switching networks.

Telenet: The success of the ARPANET led Bolt Beranek and Newman (the primary contractor for the ARPANET) to establish a commerical network company, Telenet in late 1972. In October 1973 I joined Telenet as president, and we filed with the FCC to become a carrier and construct a public packet switching network. The FCC approved the request six months later, and Telenet started the world's first public service in August 1975 with 7 nodes. By April 1978 Telenet's network had grown to 187 nodes providing service to 180 host computers and supporting direct terminal access service in 156 cities and interconnections to 14 other countries. Telenet was designed from the start to appear to the user as a **virtual circuit** service with the host interface being implemented over a communications line rather than with a box on site. However, for the first several years Telenet operated a core network based on datagrams copied from the ARPANET but implemented virtual circuits at all interfaces. It wasn't until the complete shift was made to Telenet's TP-4000 packet switch around 1980 that the savings of vitual circuits in the core net could be realized (about 30% for Telenet with a 32 byte average packet size).

EIN: Originally known as COST Project II and later as the European Informatics Network, EIN is a multi-nation funded European computer network (19,20). The project director was Derek Barber of NPL, one of the original investigators of packet switching in the U.K. It was organized in 1971 but due to the difficulties of multi-national funding it did not become operational until 1976.

Other Public Data Networks: Public packet networks were announced in the U.K., France, Canada, and Japan in the 1973-1974 period. In 1977 the Experimental Packet Switched Service (EPSS) in the U.K. and DATAPAC in Canada became operational. Also in 1977 a second public network TYMNET became approved in the U.S. In 1978 TRANSPAC in France and DX-2 in Japan became operational as public networks. By the early 1980's public packet network service was available in most countries in the world as an international service and as a domestic service in many of these countries.

CCITT Standards- X.25

With five, independent, public packet networks under construction in the 1974-1975 period (U.S.-Telenet, Canada, U.K., France, Japan), there was strong incentive for the

nations to agree on a standard user interface to the networks so that host computers would not have unique interfacing jobs in each country. Unlike most standards activities, where there is almost no incentive to compromise and agree, carriers in separate countries can only benefit from the adoption of a standard since it facilitates network interconnection and permits easier user attachment. To this end the parties concerned undertook a major effort, to agree on the host-network interface during 1975. The result was an agreed protocol, CCITT Recommendation X.25, adopted in March 1976.

The X.25 protocol provides for the interleaving of data blocks for up to 4095 virtual circuits on a single full-duplex leased line interface to the network, including all procedures for call setup and disconnection. A significant feature of this interface, from the carriers' point of view, is the inclusion of independent flow control on each virtual circuit (VC); the flow control enables the network (and the user) to protect itself from congestion and overflow under all circumstances without having to slow down or stop more than one call at a time. In networks like the ARPANET and CYCLADES which do not have this capability, the network must depend on the host (or other networks in interconnect cases) to insure that no user submits more data to the network than the network can handle or deliver. The only defense the network has without individual VC flow control is to shut off the entire host (or internet) interface. This, of course, can be disastrous to the other users communicating with the offending host or network.

Another critical aspect of X.25 is that it defines interface standards for both the host-to-network block transfer and the control of individual VC's. In datagram networks the VC interface is situated in the host computer; there can be, therefore, no network standard for labeling, sequencing, and flow controlling VC's.

The March 1976 agreement on X.25 as the technique for public packet networks marked the beginning of the second phase of packet switching: large interconnected public service networks. In the years since X.25 was adopted, many additional packet standards have been agreed on as well. X.28 was adopted as the standard asynchronous terminal interface and X.29, a protocol used with X.25 to specify the packetizing rules for the terminal handler, was adopted as the host control protocol. Also, a standard protocol for interconnecting international networks, X.75 has been adopted.

Packet Switched Voice

Packet networks today typically only utilize trunks and lines up to 64 kilobits/sec in speed. As a result, end to end delay in these networks remains high - over 100 ms. Network delay would however be vastly reduced given high speed trunks (T1-T3) and high speed packet switches (>100,000 packets/sec). Even with a shift to T1 trunks, the network delay would drop to 10 ms plus propagation. With T2 trunks, the delay would not be practically different than with a circuit switch (1-2 ms + propagation). For most data applications this is unnecessary but for voice it is critical. The voice network will often return an echo from analog phones and even 20 ms of delay can make such an echo sound bad. Short packets like 32 bytes (4.5 ms of 56 kb voice) are also necessary to reduce delay. But with high speed trunks and short packets, packetized voice is not only toll quality but also saves 67% of the transmission and switch bandwidth due to silence suppression. Thus in designing tomorrow's central office switch and toll switches for the telecommunications network it is clearly extremely desirable to use packet switches for voice of any rate (56kb, 32kb, 16kb, etc), for data of all speeds and for switched video. Without packet switching the flexibility is vastly reduced and each service must be switched separately. Packet switches of the required speed are both feasible and economic today if restricted to fixed-length, short packets (16 or 32 bytes), all flow control and error checking are done at the network interfaces - not for each link, and a shortened virtual circuit header is used (about 1 byte) with only the link number. AT&T has a prototype switch of 2 million packets/sec in test and designs exist for switches up to 35 million packets/sec (21).

The most interesting facet of high speed switch technology is that for large central office or tandem office switches the peak switch throughput is limited by the speed of the fastest memory for a time switch and similarly for a packet switch but the packet switch is able to switch a whole packet not just one byte like the time switch. Thus in general for 32 byte packets, **packet switches can be made about 32 times the maximum throughput of circuit switches.** This becomes very valuable at high level tandem switching offices to switch traffic between many fiber paths.

Based on the importance of coping with variable bitrate demands for data, voice, and video and the cost savings for integrating all services with one switch it is clear that **the whole telecommunications plant will convert to packet switching in the future.** The date is harder to predict but conversion should occur during the 1990's and be complete by 2000. The basic service will be small packets on virtual circuits and each application will have its own interface units. Data services will use an X.25 interface which does the flow control and error checking for the VC on an end-to-end basis.

Future Computer Networks

If we look forward to the year 2000, based on historical trends computers will be about 500 times cheaper per computation than today. Almost all devices in the home, office or plant will have complex computers incorporated into them. Simple local area network facilities like shielded twisted pair will connect the majority of these together and a gateway will connect this local network to the public telecommunication system over

(hopefully a revised and truly integrated) ISDN facility. Distinctions as to which device is the phone may become difficult because many of the device-computers may use the voice network from time to time for warnings, ordering, and certainly in assisting us to find the party to whom we wish to talk. Phone service today is in great need of automation; we waste large amounts of time non-productively answering and calling to find people. Thus as the telecommunications network brings data and voice together home and office automation will integrate them.

REFERENCES

- (1) P. Baran et al, "On Distributed Communications, Vols. I through XI," RAND Corporation Memos, August 1964.
- (2) T. Marill and L. G. Roberts, "Toward a Cooperative Network of Time-Shared Computers," Proceedings FJCC 1966, pp. 425-431.
- (3) L. G. Roberts, "Multiple Computer Networks and Intercomputer Communication," ACM Symposium on Operating System Principles, October 1967.
- (4) 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," ACM Symposium on Operating System Principles, October 1967.
- (5) L. G. Roberts and B. D. Wessler, "Computer Network Development to Achieve Resource Sharing," Proceedings SJCC 1970, pp. 543-549.
- (6) 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 36, June 1970, pp. 551-567.
- (7) L. Kleinrock, "Analytic and Simulation Methods in Computer Network Design," AFIPS Conference Proceedings May 1970.
- (8) H. Frank, I. T. Frisch, W. Chou, "Topological Considerations in the Design of the ARPA Computer Network," AFIPS Conference Proceedings, May 1970.
- (9) S. Carr, S. Crocker, V. Cerf, "HOST-HOST Communication Protocol in the ARPA Network," AFIPS Conference Proceedings, May 1970.
- (10) 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 40, June 1972, pp. 243-254.
- (11) L. Kleinrock, "Performance Models and Measurement of the ARPA Computer Network," Proceedings of the International Symposium on the Design and Application of Interactive Computer Systems, Grunel University, Uxbridge, England, May 1972.
- (12) L. Kleinrock and W. Naylor, "On Measured Behavior of the ARPA Network," AFIPS Conference Proceedings, Vol. 43, NCC, Chicago, Illinois, May 1974, pp. 767-780.
- (13) L. Kleinrock and H. Opderbeck, "Throughput in the ARPANET -- Protocols and Measurement," Fourth Data Communications Symposium, Quebec City, Canada, October 1975, pp. 6-1 to 6-11.
- (14) H. Frank and W. Chou, "Topological Optimization of Computer Networks," Proceedings of the IEEE, Vol. 60, No. 11, November 1972, pp. 1385-1397.
- (15) L. Roberts, "Data by the Packet," IEEE Spectrum, Vol. 11, pp. 46-51, February 1974.
- (16) L. Roberts, "Packet Switching Economics," presented at L. M. Ericsson Prize Presentation, Stockholm, Sweden, May 1982.
- (17) L. Pouzin, "Presentation and Major Design Aspects of the CYCLADES Network," Third Data Communications Symposium, Tampa, Florida, November 1973, pp. 80-85.
- (18) R. F. Despres, "A Packet Network with Graceful Saturated Operation," Proceedings ICC, Washington, D.C., October 1972, pp. 345-351.
- (19) D. L. A. Barber, "The European Computer Network Project," Proceedings ICC, Washington, D.C., October 1972, pp. 192-200.
- (20) D. L. A. Barber, "A European Informatics Network: Achievement and Prospects," Proceedings ICC, Toronto, Canada, August 1976, pp. 44-50.
- (21) A. Thomas, J. P. Coudreuse, M. Servel, "Asynchronous Time-Division Techniques: An Experimental Packet Network Integrating Videocommunication," Proceedings International Switching Symposium, Florence, Italy, May 1984.