Data by the packet

Because computing costs are now so low, an unusual new concept of data communications is feasible

The most dominant force over the past 20 years in both computer and communications architecture has been the continuous and rapid decrease in the cost of computer hardware. Not only has this electronics revolution affected the capability and design of computer systems, but it has also made possible a radically new concept of data communications. Called *packet switching*, this concept is strongly dependent on the cost of computing since it uses computers to correct transmission errors, to provide high reliability through alternate routing, and to allocate communication bandwidth dynamically on a demand basis rather than as a preassigned bandwidth.

The impact of widespread packet-switched communications networks on the computer field should be momentous. With data communications priced on a distance-independent basis, it should become economically feasible for terminals and computers throughout the country to access efficiently—on demand—a wide variety of computer services. The economics of resource sharing will undoubtedly eliminate the need for many medium-scale data processing centers and permit a considerably higher degree of specialization for the larger service centers. Moreover, packet switching will make possible far more economic realization of a comprehensive electronic funds-transfer system, on-line credit authorization, integrated corporate data networks, as well as nationwide access to all forms of data banks and retrieval services.

All these developments hinge on the availability of low-cost computer communications. Although packet switching economizes on transmission costs by maintaining high line utilization, until recently the cost of computer hardware had continued to be too high to permit practical packet-switching networks. The price of computing has been decreasing far more rapidly than the price of communications, however, and one could have anticipated a crossover point where the cost of using computers to allocate bandwidth became cheaper than the cost of the communications themselves. According to my estimate, this crossover occurred during 1969. As a result, packet switching has now become quite an economic and viable technology, permitting the establishment of large, cost-effective networks nationwide. This development in data communications will in turn have substantial effect on the course of computer technology as preliminarily indicated by results from the United States ARPA (Advanced Research Projects Agency) network research activity.

One of the main reasons for the great interest in computer networks today is the considerable economies that can be achieved through resource sharing. With a computer network that is reliable and responsive enough to permit the full resources of a remote computer to be employed locally, it becomes possible for organizations to consider the elimination or reduction of local computing facilities and the utilization of remote service centers instead. This not only can lead to considerable cost savings, as we shall see, but also has many additional advantages, including increased reliability of a multicomputer service center, improved service from a variety of competing suppliers, and increased flexibility for expansion. The main incentive for an organization changing over, however, is the cost saving. To get a feel for the magnitude of this effect, it is useful to examine the initial experience in an environment—the ARPA Computer Network (ARPANET)—where remote use of computing has been fostered by providing adequate communications facilities.[1-3]

Lawrence G. Roberts
Telenet Communications Corporation

Cost-effective remote-access computing

Although initiated in 1969, the ARPANET was not sufficiently developed to permit useful resource sharing activity until mid-1971. As network reliability and effectiveness proved out, ARPA and its contractors found many cases where newly required computer capacity could effectively be obtained through the network rather than by adding local computing resources. By March 1973, several contractors were making substantial use of the network for a majority or all of their computing resources. Several of the computer centers on the network had grown to become substantial suppliers of computer service, providing not only time-sharing service but also remote batch service. At that time, an accounting was made of the total computer usage obtained through the network, and an estimate was made for each user of the cost of purchasing comparable time on outside computers or leasing the necessary in-house computer facilities to do the same job apparently being done through the network. Table I tabulates this information by user organization, identifying both the cost of computing with and without the network. A total of $2 million a year was being spent on computing resources accessed through the network—resources that would have cost $6 million per year if the network had not existed! This savings of $4 million per year more than offset the $3.5 million annual cost of the network.*

At this stage in its development, the network was only using 20 percent of its capacity; thus, at the

* This cost includes amortization of the system's IMP (interface message processor) minicomputer units.

Data by the packet

Because computing costs are now so low, an unusual new concept of data communications is feasible

The most dominant force over the past 20 years in both computer and communications architecture has been the continuous and rapid decrease in the cost of computer hardware. Not only has this electronics revolution affected the capability and design of computer systems, but it has also made possible a radically new concept of data communications. Called *packet switching*, this concept is strongly dependent on the cost of computing since it uses computers to correct transmission errors, to provide high reliability through alternate routing, and to allocate communication bandwidth dynamically on a demand basis rather than as a preassigned bandwidth.

The impact of widespread packet-switched communications networks on the computer field should be momentous. With data communications priced on a distance-independent basis, it should become economically feasible for terminals and computers throughout the country to access efficiently—on demand—a wide variety of computer services. The economics of resource sharing will undoubtedly eliminate the need for many medium-scale data processing centers and permit a considerably higher degree of specialization for the larger service centers. Moreover, packet switching will make possible far more economic realization of a comprehensive electronic funds-transfer system, on-line credit authorization, integrated corporate data networks, as well as nationwide access to all forms of data banks and retrieval services.

All these developments hinge on the availability of low-cost computer communications. Although packet switching economizes on transmission costs by maintaining high line utilization, until recently the cost of computer hardware had continued to be too high to permit practical packet-switching networks. The price of computing has been decreasing far more rapidly than the price of communications, however, and one could have anticipated a crossover point where the cost of using computers to allocate bandwidth became cheaper than the cost of the communications themselves. According to my estimate, this crossover occurred during 1969. As a result, packet switching has now become quite an economic and viable technology, permitting the establishment of large, cost-effective networks nationwide. This development in data communications will in turn have substantial effect on the course of computer technology as preliminarily indicated by results from the United States ARPA (Advanced Research Projects Agency) network research activity.

One of the main reasons for the great interest in computer networks today is the considerable economies that can be achieved through resource sharing. With a computer network that is reliable and responsive enough to permit the full resources of a remote computer to be employed locally, it becomes possible for organizations to consider the elimination or reduction of local computing facilities and the utilization of remote service centers instead. This not only can lead to considerable cost savings, as we shall see, but also has many additional advantages, including increased reliability of a multicomputer service center, improved service from a variety of competing suppliers, and increased flexibility for expansion. The main incentive for an organization changing over, however, is the cost saving. To get a feel for the magnitude of this effect, it is useful to examine the initial experience in an environment—the ARPA Computer Network (ARPANET)—where remote use of computing has been fostered by providing adequate communications facilities.[1-3]

Cost-effective remote-access computing

Although initiated in 1969, the ARPANET was not sufficiently developed to permit useful resource sharing activity until mid-1971. As network reliability and effectiveness proved out, ARPA and its contractors found many cases where newly required computer capacity could effectively be obtained through the network rather than by adding local computing resources. By March 1973, several contractors were making substantial use of the network for a majority or all of their computing resources. Several of the computer centers on the network had grown to become substantial suppliers of computer service, providing not only time-sharing service but also remote batch service. At that time, an accounting was made of the total computer usage obtained through the network, and an estimate was made for each user of the cost of purchasing comparable time on outside computers or leasing the necessary in-house computer facilities to do the same job apparently being done through the network. Table I tabulates this information by user organization, identifying both the cost of computing with and without the network. A total of $2 million a year was being spent on computing resources accessed through the network—resources that would have cost $6 million per year if the network had not existed! This savings of $4 million per year more than offset the $3.5 million annual cost of the network.*

At this stage in its development, the network was only using 20 percent of its capacity; thus, at the

Lawrence G. Roberts
Telenet Communications Corporation

* This cost includes amortization of the system's IMP (interface message processor) minicomputer units.

same communication cost, the ARPANET could permit five times as much remote computing to be performed. At full traffic load, the network cost would be 35 percent of the actual computing cost, or 11 percent of the original computing cost. Clearly, for a moderate investment in communications, a computing cost reduction factor of three can be achieved through resource sharing.

Examining the ARPA data, it is useful to identify a few of the different categories of network usage and explore the sources of the cost savings. The largest individual user organization was the University of Illinois, which before the network came into full operation leased its own Burroughs 6700 computer to develop programs for the ILLIAC IV computer. After the University of California at San Diego put its large B6700 complex on the network, Illinois was able to terminate its computer lease and obtain the computing support far more economically from the larger installation through the network. Moreover, Illinois required access to the ILLIAC IV computer that it obtained from the NASA Ames Research Center through the network, remote batch service on the IBM 360/91 at U.C.L.A., and time-sharing services from U.S.C. and M.I.T. To obtain this collection of services locally, Illinois would have had to continue leasing the Burroughs computer, purchase time on a local IBM 360, and either purchase a time-sharing machine or obtain service from a commercial vendor. Therefore, not only would it cost the University three times as much to obtain equivalent computing power locally, but it would not have had a convenient way of accessing these machines and transferring data and programs between them.

This case is an excellent example of a user requiring many different types of service, each being most efficiently provided by a large specialized service center. For example, program development and debugging is extremely cost-effective on a time-sharing computer, but the execution of large numerical applications programs can often be done a factor of ten less expensively on a large batch processing machine.

Another type of saving is exemplified by Massachusetts Computer Associates, a small computer software firm. This group would use the equivalent of one third of the capacity of a PDP-10 time-sharing system, if they operated locally on a dedicated basis. The required facilities, however, were provided through the network from the University of Southern California, taking advantage of the fact that peak usage hours on the east and west coasts of the U.S. do not coincide due to the time zone difference. A large number of such users of various sizes can be accommodated by a single service center with considerably greater economy of operation than with a number of independent centers. The consolidation of maintenance, software, and operational personnel in one location considerably reduces costs, as most of the time-sharing companies have found, and nationwide operation increases the machine utilization considerably due to the time zone spread.

Although there are many other combinations of resource sharing represented within the ARPA network community, these two examples illustrate the main sources of cost savings resulting from hardware sharing. Additionally, large but less easily quantifiable

I. Computer resource usage within ARPANET*

User Organization	Activity	Remote Usage, in thousands	Projected Cost for Local Replacement, in thousands
University of Illinois	Parallel processing research	$ 360	$1100
NASA Ames Research Center	Air foil design and ILLIAC	328	570
Rand Corporation	Numerical climate modeling	210	650
Mass. Computer Associates	ILLIAC IV compiler development	151	470
Lawrence Livermore Laboratory	TENSOR code on ILLIAC	94	370
Stanford University	Artificial intelligence research	91	180
Rome Air Development Center	Text manipulation and resource evaluation	81	450
ARPA	On-line management	77	370
Seismic Array Analysis Center	Seismic data processing	76	300
Mitre Corporation	Distributed file network research	60	240
National Bureau of Standards	Network research	58	200
Bolt Beranek & Newman Inc.	TENEX system support	55	80
Xerox Palo Alto Research Center	Computer science research	47	100
University of Southern California Image Processing Lab	Picture processing research	35	70
University of California, Los Angeles	Network measurement	28	90
Systems Control, Inc.	Signal processing research	23	70
University of California, Santa Barbara	Network research	22	70
U.S. Air Force Range Measurements Lab	ARPANET management	17	60
Institute for the Future	Teleconferencing research	13	40
Miscellaneous	Computer research	192	580
	Total	$2018	$6060

*Annual remote computer usage cost is based on March 1973 data.

cost savings are anticipated to be forthcoming in the future from the sharing of data base and software resources.

Computer–communication interdependence

The large savings possible from resource sharing are dependent upon having available an economic means of switched communications to access these resources. This communication facility must be reliable, sufficiently responsive to meet all of the interactive demands of time-sharing users, have sufficient data transmission speed for high-speed printer and display output, be distributed to scattered users throughout the country, and most importantly be sufficiently economic so as not to destroy the cost savings achieved through resource sharing. In the early 1960s, this kind of communications service was not available. It is only with the extremely rapid development of computers themselves that it has been possible to reduce drastically the cost of providing such a communications service to computer data users.

Since computer input/output typically is extremely bursty in nature, requiring peak data transfer rates ten to one hundred times as great as the average data rate, it is necessary to share statistically among many users a communication channel of fixed data transmission capacity in order to achieve reasonable economy. Without doing this, the cost of data communications would be 10 to 100 times as expensive as the raw communication bandwidth, thereby making resource sharing cost-ineffective over even moderate distances.

Packet-switching technology as is used in the ARPANET has been developed to permit the statistical sharing of communication lines by many diverse users. By dividing the data traffic into small addressed packets (1000 bits or less in length), extremely efficient sharing of communication resources can be achieved even at the burst rates required by interactive computer traffic. Since resource-sharing computer usage has only become economically viable with the development of packet switching, it is instructive to look at the cost trends that made this technology possible. To do this, it is necessary to have a model of the average resources utilized in moving data through a packet-switching communications network.

The basic design of a packet-switching network, as exemplified by the ARPANET, consists of a collection of geographically dispersed minicomputers called *interface message processors* (IMPs) interconnected by many 50-kilobit/second (kb/s) leased lines. An IMP accepts traffic from a computer attached to it called a *host,* formats it into packets, and routes it toward its destination over one of the 50-kb lines tied to that IMP. Each IMP in the network receiving a packet examines the header and, making a new routing decision, passes it on towards its destination, possibly through several intermediate IMPs. Thus, a packet proceeds from IMP to IMP in making its way to its destination. The destination IMP collects the packets, reformats them into messages in the proper sequence, and submits them to the destination host computer. Throughout the process, each IMP checks the correctness of each packet by means of both hardware- and software-based error-control techniques. If the packet is received incorrectly due to a transmission error on the line, the IMP does not acknowledge receipt and the preceding IMP must retransmit the packet, perhaps over a different path. Because the network uses high-speed transmission lines and short packets, and all data is stored in high-speed primary memory in the IMPs (as opposed to disk drives and other secondary storage devices) average end-to-end transit delay for a packet is 0.1 second.

This general design of a packet-switching network has proven so successful in the ARPANET that a new component of the communications industry has been formed—"value-added" packet carriers, offering the public packet communications service on a nationwide regulated basis. Two of the four value-added carriers that have applied to the FCC—Telenet Communications Corporation and Packet Communications, Inc.—intend to offer a service patterned directly on the ARPANET. The FCC has responded extremely positively to this new development in the communications field and, indicating its intention to adopt a policy of "liberal entry" in this field, has already approved the first of the value-added carrier applicants.

In a nationwide packet-switching network, there is practically no dependence of cost on distance, and for this reason both Telenet and PCI intend to charge only for the *number* of packets moved, independent of the distance. If we assume that the traffic itself is evenly distributed geographically, the average distance traversed by a packet would be about half the cross-country mileage, or 1200 miles in a network spanning the U.S. The number of *hops* (intermediate leased lines traversed) required on the average to deliver a packet is a logarithmic function of the number of *nodes* in the network and therefore increases very slowly once the network exceeds about 40 nodes. For triply connected networks of 40 to 100 nodes, the average number of hops required is three or four. Since

II. Data communications costs since 1960

Date Introduced	Data Rate	Cost per Million Bits
1960	2.4 kb	$1.00
1963	40.8 kb	$0.42
1964	50.0 kb (Telpak A)	$0.34
1967	50.0 kb (Series 8000)	$0.33
1974-76	56.0 kb (DDS)	$0.11 (est.)

III. CPU time and cost estimates for the ARPA network

Date	Machine	Cost	Total Processing Time per Kilopacket	Cost per Kilopacket
1970	Honeywell 516 IMP	$100k	26.2 seconds	$0.168
1971	Honeywell 316 IMP	$ 45k	35.4 seconds	$0.102
1974	Lockheed SUE IMP	$ 10k*	41.2 seconds	$0.026

*Estimated per processor based on multiprocessor configuration.

commercial packet-switching networks, such as that proposed by Telenet Communications Corporation, will employ "central office" facilities serving many users in a city, it is not likely that the number of central office nodes will exceed 100 for many years to come. Therefore, taking four hops as a reasonable estimate, the average packet would traverse five IMPs and four lines on its course between two locations 1200 miles apart. The average length of these leased lines would be 300 miles, a number that clearly fits the model and also matches the ARPANET experience.

Therefore, two cost factors are involved in moving the average packet: (1) the communications cost for four 300-mile leased lines, and (2) the computation required in five IMPs.

Communications cost trends

Over the past decade, there has been very little change in the actual per-mile price of leased communications channels. However, in the more recent past there have been substantial changes in modem technology, permitting higher-speed data transmission over these lines, and in the near future the American Telephone and Telegraph Co. expects to introduce Digital Data Service (DDS), a totally digital technology at a substantially lower price.

Considering the period 1960–1980, one detects a downward trend in data communications cost. To quantify this, let us examine the cost per million bits of data moved for reasonably high-bandwidth leased lines 300 miles in length (including termination and modem costs). The actual costs tabulated in Table II have been multiplied by four so that they are truly representative of the incremental cost for moving one million bits (or a thousand packets) over a packet-switched network for a total distance of 1200 miles. The calculations are based on the lines being fully utilized in full-duplex mode for eight hours each working day (173 h/month—the same level of utilization assumed in calculating subsequent computer costs). It should be noted that a single leased line 1200 miles long would in general be 20–50 percent less expensive, but such a dedicated line would not permit the multipoint access and usage-based charging provided by a packet network.

Based on this data, if one takes the cost of service for each year between 1960 and 1980 and computes the least-mean-square exponential fit, one finds that the cost of service has been decreasing by 11 percent per year over that period, resulting in a factor-of-10 decrease in communications cost each 22 years. A similar analysis was made of the costs of dial telephone service over a distance of 1200 miles, using the highest data rate modem feasible over dial telephone lines. For the dial case, a comparable rate of decrease in service costs was found, largely due to improvements in modem performance, but with an average cost per megabit eight times greater than the wideband leased line case. Thus, at least over the last decade or so, the cost of data communications service has been decreasing fairly consistently and it is safe to assume that continued advances in digital transmission technology will continue this trend into the future. There is, however, no indication that a remarkable breakthrough is about to occur, drastically decreasing the cost rate of terrestrial transmission service from common carriers.

Computer cost trends

The cost per unit capability of electronics and computer equipment has often been studied, and in all cases the trend was found to be remarkably consistent at about a factor-of-ten decrease in cost every five years. In a 1969 study made by the author[4] analyzing the cost performance of complete computer systems vs. their date of first delivery, it was found that the cost performance was increasing exponentially at a rate of 1.56 per year. This study was based on the least-mean-square best fit for all machines introduced between 1955 and 1969 having at least 32-bit word length. In a later study by Newport and Ryzlak[5] examining the cost performance of communication processors, it was found that the cost performance was increasing at a rate of 1.61 per year. In this study, minicomputers introduced between 1960 and 1972 were also examined.

Given the remarkably consistent rate of increase found by these and other studies it is quite clear that the rate of growth of computer cost performance is well established. In order to fix such a curve absolutely, however, it is necessary to examine the actual CPU time consumed by an IMP to process 1000 packets of information. Very accurate studies have been made by Bolt Beranek and Newman Inc. (BBN)—ARPA's contractor for building and operating the ARPANET—of the CPU time required for the IMPs used within the ARPANET. The time consumed depends somewhat on the length of the packet, but for consistency (equating 1000 packets to one megabit of data) full packets of a thousand bits will be considered. In Table III, the three computers developed for ARPA by BBN to provide IMP service are considered. The date of introduction, the cost, and the total CPU time consumed to move a thousand full packets through five IMPs in the network are listed. The additional processing required to accept and deliver these packets to the host computer is also included. The cost per kilopacket of traffic was found by depreciating this equipment over 50 months and assuming operations and maintenance expense to equal equipment cost; operation was set at eight hours each working day for consistency with our earlier communications cost calculations.

These costs are quite consistent with the previously described growth rate of cost performance with time and form the basis of an actual trend curve for the computational cost of packet transmission. To be consistent with the treatment of communications trends, the dates that should be utilized in considering the IMP processing costs should be delayed a year from the date of introduction to correspond more accurately with their average usage date.

Composite costs for packet switching

The graph shown in Fig. 1 shows the trend curve for the data communications and computational cost associated with moving a million bits or a thousand packets of information over an average distance of 1200 miles within a nationwide network. As can be seen, even with an 11-percent decrease per year, the cost of communications remains relatively flat,

[1] **A comparison of cost performance trends for computing, data communications, and packet-switching systems over a nationwide network. The impact of satellite communications on packet switching is particularly emphasized in the graph on the right, which also demonstrates the cost effectiveness of reducing the regional size of the national communications network via satellites.**

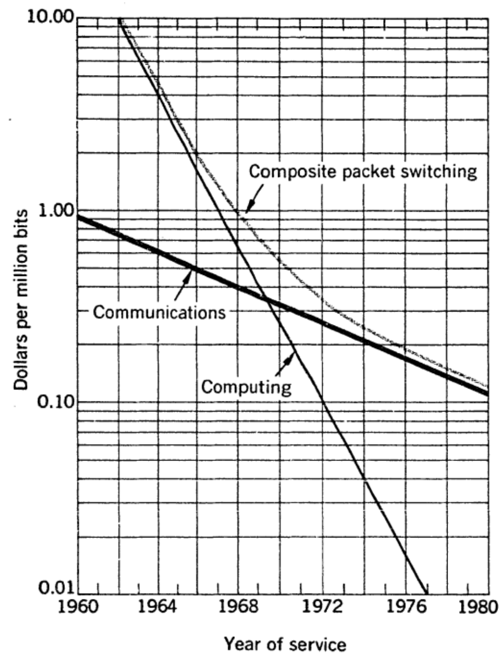

whereas the cost of computing has been decreasing with extreme rapidness; these curves intersect in 1969. Before that year, the cost of the computer power required to provide dynamic allocation of communication resources would have made packet switching extremely expensive to implement and, in fact, the first packet-switched network—the ARPANET—was not initiated until 1969. The colored curve shows the overall cost of communications plus computation associated with packet switching. Since only the incremental cost of moving a kilopacket through the net have been considered here, these results should not be interpreted as indicating the entire cost of operating a communications service, but only the incremental costs associated with moving large quantities of data. Also, since the communications trend in particular has been smoothed considerably, whereas in fact the present cost of 50-kb/s communication lines is still 33 cents/Mb, the resultant cost for any particular year should not be considered as an accurate indicator of precisely what the price ought to be. For example, in Telenet Communications Corporation's publicly filed prospective tariff, incremental packet traffic is priced at 48 cents/kilopacket. This is very close to the total cost of communication (33 cents) and IMPs (10 cents) today but corresponds to 1970 on the graph.

A more general interpretation can be made of Fig. 1. Except for the numerical absolute scale numbers, it is entirely representative of the costs associated with any system that utilizes both communications and computing components in a fixed ratio! If the communications segment is only a small fraction of the overall system, the crossover will be much later in time, but the identical shape will be preserved. Thus, for any data processing system providing remote computing capability, the curve indicates that as time passes the costs will become more and more dominated by communications. This means that it is all the more important for new communications techniques such as packet switching to be introduced in order to reduce this component of the total system cost.

IV. Cost estimates for INTELSAT communication satellites

INTELSAT	Usage Year	Number of Circuits	Lifetime, years	Total Cost	Cost per Circuit per Year
I	1965–67	240	1.5	$ 8.2M	$22 800
II	1967–68	240	3	$ 8.1M	$11 300
III	1968–71	1 200	5	$10.5M	$ 1 800
IV	1971–78	6 000	7	$26.0M	$ 600
V*	1978–85	100 000	10	$28.5M	$ 30

*Estimated.

The future—satellite communications

Although terrestrial communications cost appears to limit the future price of computer-communication service, including packet-switching networks, the situation is rapidly changing with the introduction of domestic satellites. The cost trends for satellite technology are not yet sufficiently well documented for great confidence; however, a preliminary estimate can be made from the international satellite progress over the last decade. Table IV presents the space-segment cost per circuit for the INTELSAT series of satellites.

Applying the least-mean-square exponential fit to this data, the rate of technological improvement in the cost performance for satellites is found to be 40.7 percent per year, or a factor of ten every 6.7 years. This can only be treated as a crude estimate of the cost trend for satellite communication, but since it is quite in keeping with the general cost trend for electronics, it is a quite credible growth rate.

Preliminary estimates of the actual cost of wide bandwidth digital communications service on domestic satellites are in the vicinity of 5 cents/Mb for service in early 1974. As depicted in the right-hand chart of Fig. 1, if the cost trend for domestic satellites follows the initial international trend, the projection of satellite communication cost performance is very close to that of computers.

Satellites can be used extremely efficiently within a packet network when used on a broadcast basis between a large number of ground stations. A detailed study of the techniques for broadcast satellite use and

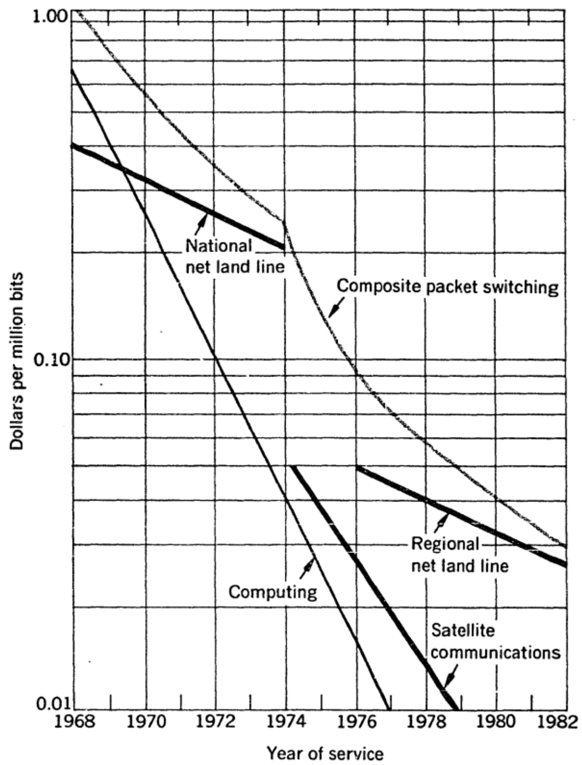

their superior performance over the use of satellites as point-to-point channels has already been described in a previous paper.[6] The cost of long-distance communications between ground stations can be reduced significantly as indicated by the satellite trend in Fig. 1. However, terrestrial leased lines are still required to interconnect the central offices of the packet network within each ground station region. By 1976–78, there should be sufficient domestic satellite stations in operation so that 12 regions can be established within the U.S., thereby reducing the size of the regional packet network to a 200-mile radius. In this case, instead of using four 300-mile leased lines, each packet would traverse four 50-mile leased lines (two in each ground station region) plus the satellite hop.

Thus, the overall cost of packet switching would be the sum of three components: computing, the satellite link, and the 50-mile terrestrial links. The trends of these three areas are plotted in Fig. 1, and the overall cost of packet switching shown to break in 1974 with the introduction of satellite service. Satellites will play an important role in reducing the future cost of packet-switching service and, for this reason, broadcast satellite service has been included as an integral part of the proposed offering by Telenet for its nationwide packet-switching system.

In the more distant future, additional technological improvements are expected to continue to occur in the packet-switching field. When they are introduced, high-frequency (15–30-GHz) satellites will most likely permit direct satellite interconnection of the central offices in a packet-switching network. This will produce an additional break in the cost trend and leave only the local distribution problem. The utilization of cable television systems and packet radio broadcast techniques[7] should permit substantial improvements in the future local distribution scene. Hence, it is entirely possible that the cost trend for packet switching will continue to decrease at almost the same rate as the cost of computation if these new packet communication techniques are vigorously pursued.

Where we stand

Experience with ARPANET has demonstrated that computing service can be obtained remotely through a computer network at one third the cost of a local dedicated system. To obtain this factor-of-three cost reduction, however, one must have available a highly responsive and reliable communications system capable of handling the peak data rate of the burst traffic normally associated with computer usage. In order to provide such an enhanced communications service, it is necessary to embed significant computational and logical capabilities within the network itself.

Until recently, the cost of computer hardware, as compared to that for the raw communication facilities, was too high to permit widespread use of computer intelligence within a data communications system. However, computer hardware costs have been falling so rapidly relative to communication facilities costs that "intelligent" and enhanced data networks have suddenly become a viable technology.

REFERENCES

1. Roberts, L. G., and Wessler, B., "Computer network development to achieve resource sharing," *Proc. AFIPS Spring Joint Computer Conf.*, pp. 543-549, 1970.

2. Heart, F., *et al.*, "The interface message processor for the ARPA computer network," *Conf. AFIPS Spring Joint Computer Conf.*, pp. 551-567, 1970.

3. Kahn, R. E., "Resource-sharing computer communications networks," *Proc. IEEE*, pp. 1397-1407, Nov. 1972.

4. Roberts, L. G., "Data processing technology forecast," internal U.S. Government paper, Apr. 23, 1969.

5. Newport, C. B., and Ryzlak, J., "Communication processors," *Proc. IEEE*, vol. 60, pp. 1321-1332, Nov. 1972.

6. Roberts, L. G., "Dynamic allocation of satellite capacity through packet reservation," *Proc. National Computer Conf.*, pp. 711-716, 1973.

7. Roberts, L. G., "Network rationale: a 5-year reevaluation," *Proc. COMPCON 73*, Feb. 1973.

Lawrence G. Roberts (M) obtained the Ph.D. degree in electrical engineering from M.I.T. in 1963. He is presently president of Telenet Communications Corporation, Washington, D.C., a firm that plans to establish and operate a nationwide packet-switching service as a common carrier under FCC regulation. Formerly director for information processing techniques with the Advanced Research Projects Agency of DOD, Dr. Roberts was principally concerned with new information processing techniques in the areas of computer systems, graphics, artificial intelligence, signal processing, weather modeling, and man-machine interaction. In addition, he was responsible for developing the ARPA computer network. Prior to his government experience, he led the computer software group at M.I.T.'s Lincoln Laboratory and was instrumental in developing a three-dimensional input device (the Lincoln Wand), the APEX time-sharing system, the CORAL list structure system, and several graphics systems. Dr. Roberts has written many professional articles and has lectured throughout the world. He was selected the 1965-66 national lecturer for ACM, of which he is a member; he is also a member of Sigma Xi.