

On Resource Sharing in a Distributed Communication Environment

Leonard Kleinrock

The efficiency of resource sharing provides the cost-effectiveness of packet switching in many of our computer-communications systems.

A revolution is in the making! We are witnessing a growth rate in technological change which is overwhelming. Thanks to enormous advances in data communications and in integrated chip technology, we are in the midst of a computer communication explosion which has already made significant changes in the field of data processing. The early phase of the revolution has passed—we have developed cost-effective data communication systems. Indeed in the last five years we have witnessed the rise of *computer networks* whose function it is to span intercontinental distances and provide communication among computers across nations and across the world. There now exists a large number of national networks which are in the process of interconnecting to each other in such a world network.

These networks have hastened the next phase of the revolution, namely, the widespread acceptance and application of teleprocessing and networking by the business sector of our economy. As this second phase proceeds, we will see a stress placed on our computer networks in two areas. First, in the need for long-haul, wide-band inexpensive communications deep in the backbone of our networks; one answer to this need is the introduction of sophisticated packet **satellite** radio data communication systems. The other environment in which we will see stress is at the periphery of our networks where local access is the major problem. The early

The author is with the Department of Computer Science, University of California, Los Angeles, CA 90024.

generation computer networks did not properly solve the local interconnection problem, namely, how to efficiently provide access from the user at the terminal to the network itself. A potential solution to this problem is the use of **ground radio** packet communications. In this article we describe some of the recent technological advances which have provided solutions to these two problems. Indeed the systems issues involved for both radio communications problems are very similar, although the technological implementations are quite different as we shall see. To begin with, we discuss the general principles of resource sharing which provide the key to the cost-effectiveness of radio packet switching.

RESOURCE SHARING

A privately owned automobile is usually a waste of money! Perhaps 90 percent of the time it is idly parked and not in use. However, its "convenience" is so seductive that few can resist the temptation to own one. When the price of such a poorly utilized device is astronomically high, we do refuse the temptation (how many of us own private jet aircraft?). On the other hand, when the cost is

Cost-effective computer networks have hastened the widespread acceptance and application of teleprocessing and networking by the business sector of our economy.

extremely low, we are obliged to own such resources (we all own idle pencils).

An information processing system consists of many poorly utilized resources. (A resource is simply a device which can perform work for us at a finite rate.) For example, in an information processing system, there is the CPU, the main memory, the disk, the data communication channels, the terminals, the printer, etc. One of the major system advances of the early 1960's was the development of multiaccess time-sharing systems in which computer system resources were made available to a large population of users, each of whom had relatively small demands (i.e., the ratio of their peak demands to their average demands was very high) but who collectively presented a total demand profile which was relatively smooth and of medium to high utilization. This was an example of the advantages to be gained through the smoothing effect of a large population (i.e., the "law of large numbers") [1]. The need for resource sharing is present in many many systems (e.g., the shared use of public jet aircraft among a large population of users).

In computer communication systems [2] we have a great need for sharing expensive resources among a collection of high peak-to-average (i.e., "bursty") users. In Fig. 1, we display the structure of a computer network in which we can identify three kinds of resources:

1) the *terminals* directly available to the user and the *communications resources* required to connect those terminals to their HOST computers or directly into the network (via TIPS in the ARPANET, for example)—this is an expensive portion of the system and it is generally difficult to employ extensive resource sharing here due to the relative sparseness of the data sources;

2) the *HOST machines* themselves which provide the

information processing services—here multiaccess time-sharing provides the mechanism for efficient resource sharing;

3) the *communications subnetwork*, consisting of communication trunks and software switches, whose function it is to provide the data communication service for the exchange of data and control among the other devices.

The HOST machines in 2) above contain hardware and software resources (in the form of application programs and data files) whose sharing comes under the topic of time-sharing; we dwell no further on these resources. Rather, we shall focus attention on those portions of the computer communications system where packet communications has had an important impact. Perhaps the most visible component is that of the communications subnetwork listed in item 3) above. Here packet communications first demonstrated its enormous efficiencies in the form of the ARPANET in the early 1970's (the decade of computer networks) [2]. The communication resources to be shared in this case are *storage capacity* at the nodal switches (the IMP's in the ARPANET), *processing capacity* in the nodal switches, and *communications capacity* of the trunks connecting these switches. Packet switching in this environment has proven to be a major technological breakthrough in providing cost-effective data communications among information processing systems.

As stated earlier, deep in the backbone of such packet-switched networks there is a need for long-haul, high-capacity inexpensive communications, and it is here where we see the second application of packet communications for resource sharing in the form of *satellite packet switching*. The third application may be found in

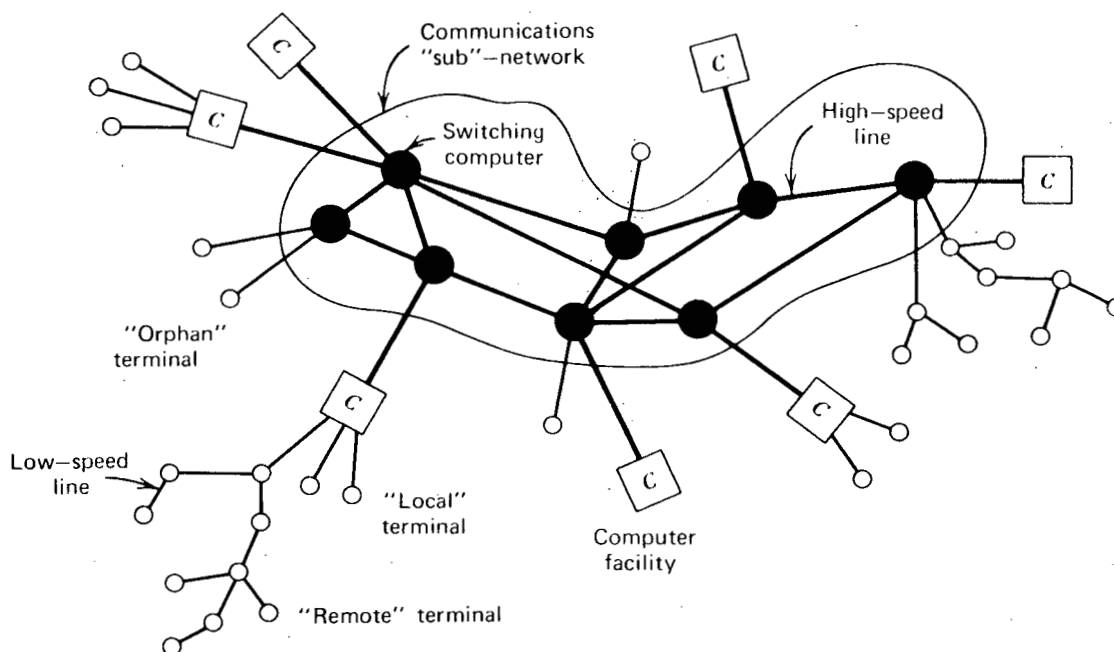


Fig. 1. The structure of a computer communication network.

- ARPA** — Advanced Research Projects Agency of the Department of Defense.
- ARPANET** — The ARPA computer network which interconnects university-computing locations and other research establishments.
- Packet** — A group of binary digits including data and call control signals which is switched as a composite whole. The data, call control signals, and possibly error control information are arranged in a specified format. (CCITT definition).
- Packet Switching** — The transmission of data by means of addressed packets whereby a transmission channel is occupied for the duration of transmission of the packet only. The channel is then available for use by packets being transferred between different data terminal equipment. (CCITT definition).
- IMP** — Interface Message Processor. A device at each node of the ARPANET which performs message switching and interconnects the research computer centers or "hosts" with the high bandwidth leased lines.
- TIP** — Terminal IMP. An IMP with multiplexing and demultiplexing equipment that collects characters from terminals, packages them in the form suitable for processing by the IMP, and sorts out characters destined for particular terminals.
- Contention** — The condition where two or more users may want access to a single resource with the possibility that their demands could come at the same time.

the local access problem stated in item 1) above which also lends itself to the use of packet switching to provide efficient communications resource sharing; this takes the form of the use of a multiaccess broadcast channel in a local environment, commonly known as *ground radio packet switching*. The common element running through all these systems is the application of the smoothing effect of a large population to provide efficient resource sharing, an exquisite example of which is provided by packet communications.

Let us consider two important examples of the effectiveness of resource sharing. Both cases involve the sharing of communication lines. In the first case we consider a voice communication system in which we provide m trunks to serve a population of users attempting to place telephone calls. We assume that a user call is blocked (and hence lost) if all trunks are occupied when he attempts to place his call. The common measure of load on such a system is expressed in *Erlangs* (one Erlang represents the full-time use of a single trunk). Clearly, we cannot handle more than m Erlangs, but if we even

TABLE I
Resource Sharing of Telephone Trunks

Load (Erlangs)	Number of Trunks (m)	Percentage Blocked
2/3	1	40%
8/3	4	17%
32/3	16	3%
128/3	64	0.05%

approach a load of m Erlangs, then due to statistical fluctuations in customer behavior, we know that many calls will be blocked and this is a situation we wish to avoid. The game, therefore, is to design enough trunks into a system to satisfy a given load so that the probability of blocking is small enough to satisfy the users' needs in an economical fashion. The analysis for this classic problem was solved 70 years ago by A. K. Erlang [1] and that solution has the amazing characteristic shown in Table I. Here we see the way which the percentage of blocked calls varies with the number of trunks (m) and the load level. Now for the resource sharing. Suppose we had a population presenting a load of 2/3 Erlangs to a single trunk; we then see that the percentage of blocked calls is an atrocious 40 percent. Obviously, if we had four such populations, each accessing their own single trunk, then each group would experience the same 40 percent blocking. However, if these four groups would simply pool their trunks yielding a total population load of 8/3 Erlangs sharing the set of 4 trunks, then we see that the percentage blocked has now drastically reduced to 17 percent. By pooling 16 times the load (32/3 Erlangs) onto 16 trunks, the percentage blocked drops to 3 percent and if we go further down the table to 64 times the original load (128/3 Erlangs) sharing a pool of 64 trunks, we reduce the percentage blocked to a mere 0.05 percent. This represents a gain of almost three orders of magnitude! Thus, in effect, by the creation of a larger population sharing pooled resources, we have gained enormously in system performance.

By the creation of a larger population sharing pooled resources, we have gained enormously in system performance.

Our second case involves a data communication system in which random blocks of data (i.e., messages) arrive (at a rate of λ msg/s). Let us assume that we have available a single data communication channel of C bits/s to satisfy these demands. When an arriving demand finds the channel busy transmitting some other message, this arriving demand will queue up and wait its turn for service (rather than being blocked, and therefore lost, as in the first case). As before, we can describe the load on the system in terms of a number ρ which is the fraction of time the channel is transmitting data (in this case of a single channel, ρ is exactly equal to the load in Erlangs); as ρ approaches one, the average delay increases without bound. We wish to provide a large enough channel capacity to keep the delay within acceptable bounds. For this case (where a given population presents a load ρ to a channel of capacity C) let us denote the resulting average delay by $T(\lambda, C)$. If we had two populations each presenting a load ρ to their own channel of capacity C , then clearly each group would suffer an average delay of $T(\lambda, C)$. On the other hand, let us consider merging these two data streams, thereby resulting in a

total rate of 2λ msg/s (with each message having the same average length as before) and merging the two capacities into a single channel of capacity $2C$; the resultant delay may be denoted by $T(2\lambda, 2C)$. The load ρ in this merged case would be the same as it was before the merging. However, the startling fact is that average delay in the merged case (the case of resource sharing) is *half* that of the former situation. That is,

$$T(2\lambda, 2C) = \frac{1}{2} T(\lambda, C).$$

In fact if we merged N such groups, then the overall delay will drop by a (sharing) factor of N [3]. Again we see the remarkable advantages of resource sharing. The principle here is that serving a large population with a large shared resource is extremely efficient in terms of performance.

It is this efficiency of resource sharing that provides the cost-effectiveness of packet switching in so many of our computer communication systems. This is accomplished by assigning network resources (channel capacity, storage, logical links, etc.) on a demand basis; these resources are designed to be shared among a large population of users and are therefore endowed with a large capacity for serving these users. It is precisely when large populations share large-capacity resources that we enjoy the performance efficiencies of resource sharing.

MULTIACCESS BROADCAST CHANNELS

Let us now characterize the satellite and ground radio switching systems which we introduced earlier. Such channels may be described as *multiaccess broadcast channels in a distributed environment*. Indeed, the object is to properly share this precious communications resource among a collection of user terminals where the resource is to be used to provide communications among the users. From the previous section we recognize that if the users are bursty, then we should not partition the channel into small pieces (each such piece assigned to a fraction of the user population), but rather, we should look for ways to share the full channel among all users on a demand basis. Immediately we can separate two cases: the first case is that in which all users are within radio range and line-of-sight of each other (we speak of this as a *one-hop* system); the second case is where not all users are within range and sight of each other (in which case we have a *multihop* environment). The adjective *multiaccess* refers to the fact that the many users are trying to share the channel simultaneously in some cooperative fashion. The adjective *broadcast* refers to the fact that each channel can hear the transmission from many or all other terminals.

The important characteristic is described by the adjective *distributed*, which refers to the fact that our terminals are geographically distributed in a way which makes controlling their behavior an issue of importance. The key parameter describing this notion of distributed sources is usually taken to be the ratio a of the propagation delay

(the time it takes electromagnetic energy moving at the speed of light to pass between two separated terminals) to the transmission time of a packet. For example, consider 1000 bit packets transmitting over a channel operating at a speed of 100 kbits/s. The transmission time of a packet is then 10 ms. If the maximum distance between the source and destination is 10 mi then the (speed of light) packet propagation delay is on the order of $54 \mu\text{s}$. (This is a typical example for a ground radio packet-switching system.) Thus the propagation delay constitutes only a very small fraction ($a = 0.005$) of the transmission time of a packet. On the other hand, in a satellite environment, this ratio is more often on the order from 10 to 30; for example, a geostationary satellite introduces a propagation delay on the order of 250–270 ms, and for the 10 ms packet transmission time mentioned above, we would then have a ratio of propagation delay to packet transmission time of roughly $a = 25$.

Now how do we pull off the “resource sharing”? An ever-increasing number of access schemes have recently been described in the published literature [4] which more or less succeed at this; these we describe shortly. Before doing so, however, let us discuss the *price* one must pay for sharing a communication channel in such a distributed environment.

THE UNAVOIDABLE PRICE

As with most contention systems, two factors contribute to a degradation in performance: first, there are the usual queueing effects due to the random nature of the message generation process; second, there is the cost due to the fact that our message sources are geographically distributed. If all the terminals were collocated (i.e., coordination among them was free and instantaneous) then we could form a common queue of the generated message packets and achieve the optimum delay-throughput profile, namely, that of the $M/D/1$ queueing system [1] described later. Unfortunately, coordination is not free and we must expend some effort in organizing our many terminals which are distributed and which independently generate messages. The total capacity we have available is fixed and we are faced with controlling access to this channel from these distributed message sources in which the control information must pass over the same channel which is being controlled (or over a subchannel which is derived from the data channel).

We have a spectrum of choices for introducing this control, ranging from no control at all to dynamic control, and finally to extremely tight static control. For example, we could allow the terminals to access the channel using PURE (i.e., unslotted) ALOHA in which a terminal transmits a packet as soon as it is generated hoping that it will not collide with any other packet transmission; if there is a collision, then all packets involved in that collision are “destroyed” and must be retransmitted later at some randomly chosen time. This *uncontrolled* scheme is extremely simple, involves no control function or hard-

ware, but extracts a price from the system in the form of wasted channel capacity due to collisions. At the other extreme, we could have a very tight *fixed control* as for example in FDMA or TDMA (see next section) where each terminal is permanently assigned a subchannel derived from the original channel. Such a *fixed control* scheme certainly avoids any collisions, but is inefficient for two reasons: first, because terminals tend to be bursty sources and therefore much of their permanently assigned capacity will be wasted due to their high peak-to-average ratio; and second, the response time will be far worse in this channelized case due to the scaling effect which is especially apparent in FDMA (see the second section). A *dynamic control* scheme such as reservation-TDMA, or Roberts' reservation scheme [5] makes use of a reservation subchannel through which terminals place requests for reserved space on the data channel; this system permits dynamic allocation of channel capacity according to a terminal's demand, but requires overhead in order to set up these reservations.

Thus we see that the issue of allocating capacity in a distributed environment is a serious one. In one form or another nature will extract her price! This price will appear in the form of collisions due to poor or no control, wasted capacity due to rigid fixed control, or overhead due to dynamic control. These comments are summarized in Table II below.

TABLE II
The Price for Distributed Sources

	Collisions	Idle Capacity	Overhead
No Control (e.g., ALOHA)	Yes	No	No
Static Control (e.g., FDMA)	No	Yes	No
Dynamic Control (e.g., Reservation Systems)	No	No	Yes

In general, as the number of terminals grows, and as the geographical separation grows, then also grows the price we pay for distribution.

A FAMILY OF MULTIACCESS METHODS [4]

Multiaccess methods for distributed computer communication systems have recently been evaluated. In this section we describe a variety of these suitable for one-hop systems. It is perhaps best to think of all terminals as transmitting fixed-length packets to a central station which is the destination for these transmissions (this is not a necessary assumption since point-to-point communication also fits this model).

We now consider nine random multiaccess broadcast schemes, and for each we give a reference and an extremely concise definition:

PURE (UNSLOTTED) ALOHA [6]: A newly generated packet will be transmitted by its terminal at the instant of its generation; collided packets destroy each other and must be retransmitted.

SLOTTED ALOHA [2], [7]: The same as PURE ALOHA except that new packet transmissions must begin at the next slot point, where time is slotted into lengths equal to a packet transmission time.

CSMA (Carrier Sense Multiple Access) [2], [8]: The same as PURE ALOHA except that a terminal senses (listens to) the channel and can hear the carrier of any other terminal's transmission; if such a carrier is detected, then the terminal refrains from transmitting and follows one of many defined protocols for later attempts.

POLLING [9]: A central controller sends a "polling message" to each terminal in turn; when a terminal is polled, it empties all of its data before indicating its empty buffer condition whereupon the next terminal is polled in sequence.

FDMA (Frequency-Division Multiple Access) [9]: The bandwidth of the channel is divided into M equal subchannels, each reserved for one of the M terminals.

TDMA (Time-Division Multiple Access) [9]: Time is slotted and a periodic sequence of the M integers is defined such that when a terminal's number is assigned to a slot, then that terminal (and only that terminal) may transmit in that slot; typically each terminal is given one out of every M slots.


MSAP (Mini-Slotted Alternating Priority) [10]: A carrier-sense version of polling whereby a polling sequence is defined and when a terminal's buffer is empty, it simply refrains from transmitting; after a time interval equal to the propagation delay, the next terminal in sequence senses the channel idle and proceeds with its transmission, etc. (This is also known as hub go-ahead polling.)

RANDOM URN SCHEME [11]: An optimal distributed control adaptive scheme in which a fraction N/M of the terminals is given permission to transmit (and each will do so if it has anything to transmit). M is the total number of terminals and N is the number (assumed to be known) that wish to transmit.

M/D/1 [1]: The classic first-come-first-serve single-server queueing system with Poisson-distributed arrivals and constant service time equal to a packet transmission time. This is an ideal system which neglects the fact that the terminals are geographically separated.

In Fig. 2, we plot the mean response time of the system (the time from when the packet wants to transmit until it is correctly received at the destination) as a function of the channel load ρ , for a number of these schemes. This figure is for the case of $M = 100$ terminals and shows the relative performance of the various access schemes. We note, for example, that PURE ALOHA gives the best performance at extremely small loads whereas the MSAP scheme seems to give the best performance at high loads

M/D/1 SHOE SHINE 50¢
 SERVICE TIME: 6 MIN.



The M/D/1 Queue

The shoeshine boy above illustrates the simplest kind of waiting-line situation, the M/D/1 queue. This queueing system has Poisson-distributed arrivals at a rate λ customers per second with a constant service time of x seconds. With the Poisson distribution, the probability of k arrivals in a time interval t seconds is given by

$$P(k,t) = e^{-\lambda t} (\lambda t)^k / k!$$

for $k=0,1,2,\dots$. The average response time T (waiting time plus service time) is [1]:

$$T = \frac{\lambda x^2 / 2}{1 - \lambda x} + x.$$

Therefore, for customers arriving for a shine at an average rate of one every 10 min, and a service time of 6 min a shine, on the average a customer can expect to spend a total of 10.5 min at the shoeshine stand. In the packet network application, λ is messages per second and x is equal to a packet transmission time.

(excluding the ideal scheme M/D/1). Indeed a well-designed static control system such as TDMA will also perform very well at high loads. What is important is to find a scheme which adapts its behavior between that of an ALOHA-like scheme at light loads to a static control scheme at heavy loads. Such schemes are beginning to appear in the literature and an example of one is the URN scheme described above. Another example is a scheme known as Scheduled Retransmission Upon Collision (SRUC) which was recently described in [12].

The performance profiles shown in Fig. 2 represent some of the better known access schemes currently available. Many more are being studied and will soon be available in the literature. Again, the idea is to create access schemes which perform well in this multiaccess

broadcast distributed environment. To perform well means to find an efficient way to share the common channel capacity.

APPLICATIONS AND THE FUTURE

Two applications we have mentioned are: wideband satellite systems with an enormous propagation delay and ground radio systems with a tiny propagation delay. It is worthwhile observing that the Advanced Research Projects Agency (ARPA) has been conducting experiments for both of these systems. The first is an experimental satellite network which currently connects three countries across the Atlantic [13]; the measurement and implementation results from this system have been quite




Fig. 2. Delay-throughput profile for multiaccess broadcast schemes (100 terminals). The parameter a is the ratio of propagation delay to packet transmission time.

encouraging and indicate that satellite packet switching is a cost-effective and viable technology. Also, ARPA is conducting an experiment in the Palo Alto, CA area for a ground radio packet-switching environment including mobile terminals [14]. Here too the early measurements indicate that the system is feasible and effective.

It is fair to say that radio packet switching is a new technology about to explode in the applications area. The satellite application deep in the backbone of our computer networks is clear and needs no further justification. Access for local terminals has been a long outstanding problem to teleprocessing system designers in that the cost of the local access portion of the network has been far too high relative to the rest of the communication system. Radio packet switching promises to reduce that cost significantly, and we can expect to see such systems available in the near future; an interesting example is Xerox's proposed XTEN network. Indeed the ground radio applications include such things as communications among moving vehicles (taxicabs, police cars, ambulances, private fleets), communications among aircraft, and indeed communications among any mobile units or any widely distributed units in a sparse environment. We may even see the use of radio packet broadcasting on a tiny scale down at the integrated chip level if cost-effective lasers can be implemented on a chip; laser packet switching among logic elements on a chip may greatly simplify the interconnection and/or prototype problem in chip design.

An exciting application of these radio packet access schemes has been under development recently and we are already beginning to see products and services based upon this new development. The application is to use packet radio access schemes *not* in a radio environment,

but rather on a coaxial cable or other wire-communication media. Indeed the entire technology of loop and ring structures has recently taken advantage of these access schemes. For example, consider the case of a data communication bus to which are attached a number of devices (e.g., a CPU, a disk, a drum, terminals, minicomputers, etc.). Until recently, contention for access to this bus had usually been resolved by a central controller. Clearly, any of the access schemes which we described in the previous section also lend themselves for application to this wire-communication bus. Indeed we are already seeing products based upon this idea in which demand access and random access are used to govern the use of a communication bus; a prime example of this is the ETHERNET [15] developed by Xerox. Here, a 1 km coaxial cable is being used to connect up to 256 devices which transmit data at 3 Mbits/s, using a variation of CSMA. The variation is simply that a device cannot only listen *before* it transmits, but it can also listen *while* it is transmitting; this permits it to detect collisions and then to abort its own transmission in the event of a collision, thereby saving considerable wasted time on the channel. The early indications are that such a scheme works extremely well and we can look forward to many more applications of packet radio access schemes to wire communications. For example, there is no reason why all future aircraft and naval vessels should not be wired up in such a fashion. Furthermore, all office buildings could have a common pipe running throughout the building, attached to which are all terminals requiring access to each other and to a centralized computer located perhaps in the basement. Indeed, not only in-building but in-plant communication should be run this way among a number of buildings at a given site. The applications here are unlimited and in fact, one may expect to see the applications to wire-based communications appearing on the market before the radio schemes are available.


The key to the success in all these developments is simply that large populations sharing large resources provide enormous efficiencies in performance and are to be incorporated whenever possible. The analytic and design problems which remain in studying single-hop and multihop schemes continue to occupy the analysts and designers in computer communications. The applications have not, should not, and will not wait for analytic results as long as a cost and performance savings can be demonstrated. We can expect much future work in this area, both in analysis as well as in development.

REFERENCES

- [1] L. Kleinrock, *Queueing Systems, Vol. 1: Theory*. New York: Wiley, 1975.
- [2] —, *Queueing Systems, Vol. 2: Computer Applications*. New York: Wiley, 1976.
- [3] —, "Resource allocation in computer systems and computer-communication networks," *Proc. IFIP Cong.-74*, pp. 11-18, 1974.
- [4] —, "Performance of distributed multi-access computer-

- communication systems," *Proc. IFIP Cong.-77*, pp. 547-552, 1977.
- [5] L. G. Roberts, "Dynamic allocation of satellite capacity through packet reservation," *AFIPS Conf. Proc. 1973 Nat. Comput. Conf.*, vol. 42, pp. 711-716, 1973.
- [6] N. Abramson, "Packet switching with satellites," *AFIPS Conf. Proc. 1973 Nat. Comput. Conf.*, vol. 42, pp. 695-702, 1973.
- [7] L. Kleinrock and S. Lam, "Packet switching in a slotted satellite channel," *AFIPS Conf. Proc. 1973 Nat. Comput. Conf.*, vol. 42, pp. 703-710, 1973.
- [8] L. Kleinrock and F. Tobagi, "Random access techniques for data transmission over packet-switched radio channels," *AFIPS Conf. Proc. 1974 Nat. Comput. Conf.*, vol. 44, pp. 187-201, 1975.
- [9] J. Martin, *Teleprocessing Network Organization*. Englewood Cliffs, NJ: Prentice-Hall, 1970.
- [10] M. Scholl, "Multiplexing techniques for data transmission over packet-switched radio systems," Ph.D. dissertation, Comput. Sci. Dep., Univ. of California, Los Angeles, Eng. Rep. UCLA-ENG 76123, Dec. 1976.
- [11] L. Kleinrock and Y. Yemini, "An optimal adaptive scheme for multiple access broadcast communication." *ICC'78 Conf. Rec.*, vol. 1, pp. 7.2.1-7.2.5, June 4-7, 1978.
- [12] F. Boronovo and L. Fratta, "SRUC: A technique for packet transmission on multiple access channels," *Proc. of 4th Int. Conf. on Comput. Commun.*, pp. 601-607, 1978.
- [13] I. M. Jacobs, R. Binder, and E. V. Hoversten, "General purpose packet satellite networks," *Proc. IEEE*, vol. 66, pp. 1448-1467, Nov. 1978.
- [14] R. E. Kahn, "The organization of computer resources into

- a packet radio network," *IEEE Trans. Commun.*, vol. COM-25, pp. 169-178, Jan. 1977.
- [15] R. M. Metcalfe and D. R. Boggs, "Ethernet: Distributed packet switching for local computer networks," vol. 19, no. 7, pp. 395-404, July 1976.



Leonard Kleinrock received the B.E.E. degree from City College of New York, New York, NY in 1957 and the M.S.E.E. and Ph.D.E.E. degrees from Massachusetts Institute of Technology, Cambridge, in 1959 and 1963, respectively.

He was an Assistant Engineer at the Photobell Co. in New York from 1951 to 1957 and a Staff Associate at the MIT Lincoln Laboratory from 1957-1963. In 1963 he joined the faculty of the School of Engineering and Applied Science, University of California, Los Angeles, where he is now Professor of Computer Science. His research spans the fields of computer networks, computer systems modeling and analysis, queueing theory and resource sharing, and allocation in general. At UCLA, he directs a large group in advanced teleprocessing systems and computer networks. He is the author of three major books in the field of computer networks: *Communication Nets: Stochastic Message Flow and Delay* (New York: McGraw Hill, 1964; also New York: Dover, 1972); *Queueing Systems, Vol. I: Theory* (New York: Wiley-Interscience, 1975); and *Queueing Systems, Vol. II: Computer Applications* (New York: Wiley-Interscience, 1976). He has published over 85 articles and contributed to several books. He serves as consultant for many domestic and foreign corporations and governments and he is a referee for numerous scholarly publications and a book reviewer for several publishers.

Dr. Kleinrock was awarded a Guggenheim Fellowship for 1971-1972 and is an IEEE Fellow "for contributions in computer-communication networks, queueing theory, time-shared systems, and engineering education." He was the co-winner of the 1976 Lanchester Prize for his book *Queueing Systems, Vol. II: Computer Applications* (J. Wiley, 1976).

Call for Papers

Special Issue on Military Communications

IEEE Transactions on Communications

The IEEE Communications Society is planning a Special Issue of the TRANSACTIONS dedicated to Military Communications Systems to be published in Fall 1980. The issue will contain invited papers on:

- Strategic, Tactical, and Base Communications Systems
- NATO Communications Systems
- Specific Communications, such as AUTOVON, AUTODIN, JTIDS, etc.
- Future Concepts

Papers are solicited in the area of Technology for Military Communications. Specifically we are interested in papers concerned with:

- Transmission
- Switching
- Terminals
- Control
- Security

Prospective authors should prepare a 500-word summary to indicate interest. A complete manuscript must then be prepared in accordance with the "Information for Authors" published in the TRANSACTIONS. Please send five (5) copies of the summary to D. L. Schilling, Department of Electrical Engineering, City College of New York, New York, NY 10031, before July 1, 1979. Five (5) copies of the entire manuscript are due before December 31, 1979.

For additional information call any of the Guest Editors:

- D. L. Schilling (212) 862-3737
- I. Lebow (202) 692-1765
- R. L. Pickholtz (202) 676-6538
- G. Coviello (202) 692-0093