The Latency/Bandwidth Tradeoff in Gigabit Networks

Gigabit networks really are different!

Leonard Kleinrock

The bandwidth for data communications has been growing steadily and dramatically over the last twenty years. Some of us remember the early days of data modems which provided access speeds of 10 characters per second (cps) in the late 1960s. When 300 baud speeds became available (providing 30 cps), we thought of it as a major improvement (and it was).

In the mid-70s, as packet-switched networks [1] began to proliferate, we saw the standard set at 64 kb/s trunk speeds;[1] of course, by the time one paid for the software and protocol overhead, we were happy to end up with about 10 kb/s file transfer speeds (by now, the dial-up data modem speeds had reached 2400 bits per second). The killer application which drove the penetration of these X.25 networks was that of transaction processing.

In the 1980s we witnessed the proliferation of T1 channel speeds, providing 1.533 megabit per second (Mb/s) trunk speeds. Private T1 networks exploded in the 1980s because of the cost savings they provided by allowing corporations to integrate their voice and data networks into a single network. This was the killer application for corporate T1 networks. In the scientific community, T1 was introduced toward the end of the 1980s due to the killer applications of e-mail and file transfers; the load from these applications arose very quickly due to the enormous growth in the number of connected users. However, the packet-switched networks still had 64 kb/s backbone speeds due largely to the complex operations the switches were required to carry out; specifically, each switch had to process every packet up to the third layer (the network layer) of the seven-layer OSI architecture.[2]

As we entered the 1990s, we saw a grass roots development in the form of Frame Relay networks [3,4]. These nets offer packet switching at T1 speeds, a significant step above the 64 kb/s packet switching nets of the 1980s. Both hardware and software developments led to these higher speed packet switched networks. On the hardware side, the widespread deployment of fiber optic communication channels by the long-haul carriers was critical. Besides having enormous bandwidths, these fiber optic channels are extremely noise-free, thereby greatly relieving the network of extensive error control. Faster switches have also been developed due to the progress in VLSI technology. However, the communication bandwidth has grown much more rapidly due to fiber optics than has the speed of the switch due to VLSI. Indeed, prior to the fiber optic revolution, the communication link had represented the performance bottleneck, and so one was prepared to waste switch capacity in order to save communication capacity. This took the form of packet switching in which intelligent switches were introduced into our data communication networks in order to dynamically assign the channel bandwidth on a demand basis. However, now that fiber optics has appeared, the communication bandwidth is no longer a constraint; in fact, a reversal in the relative cost of switching and transmission has taken place and has led us to architectures in which the switch has now become the economic as well as the performance bottleneck. Considerable research and development effort is currently under way to produce high-speed packet switches [5].

New protocols which take advantage of these hardware improvements have also been developed. In particular, the ISDN signalling channel (the D channel) uses a streamlined protocol for routing signalling packets (known as the Link Access Protocol for the D channel - LAPD) [6]; indeed, it only processes these packets up to the second layer (the data link layer) of the seven layer model, extracting a minimal amount of network layer information. Frame Relay uses the LAPD protocol for the data channel (rather than just for the signalling channel), thereby achieving much higher transfer speeds than were possible with X.25 packet networks. Thus, by relegating as much function to hardware as possible, by moving function out of the network when possible (e.g., error control on the data packets), and by taking advantage of streamlined packet protocols, Frame Relay is able to achieve packet switching at T1 speeds. The killer application which has been the driving force behind Frame Relay is that of local area network (LAN) interconnection.

Dr. LEONARD KLEINROCK is professor and chairman of computer science, University of California, Los Angeles.

---

[1] *In fact, in 1969, we already had the early ARPANET operating at 50 kb/s out of UCLA.*

0163-6804/92/$03.00 1992© IEEE824-7504

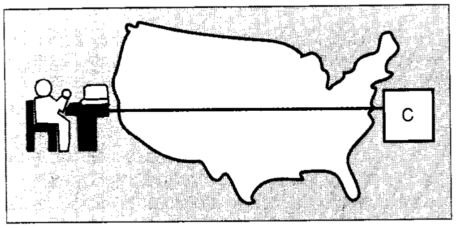

■ **Figure 1.** *Sending a 1-Mb file across the U.S.*

In addition, we have seen some multi-megabit data network plans, announcements and offerings. Among these are the Fiber Distributed Data Interface (FDDI) at 100 Mb/s [7], Switched Multimegabit Data Services (SMDS) at 45 Mb/s [8], the Distributed Queue Dual Bus access protocol at 45 and 150 Mb/s [9], ATM switches and Broadband ISDN [10] at 155 Mb/s up to 2.4 gigabits per second (Gb/s), the High Performance Parallel Interface (HIPPI) at 800 Mb/s, etc. Indeed, the Synchronous Optical Network (SONET) [11] standard has defined speeds for optical systems well into the multigigabit range.

It is clear we are moving headlong into an era of gigabit per second speeds and networks.

The Major Issue: Latency vs. Bandwidth

As we move into the gigabit world, we must ask ourselves if gigabits represent just another step in an evolutionary process of greater bandwidth systems, or, if gigabits are really different? In the opinion of this author, gigabits are indeed different, and the reason for this difference has to do with the effect of the latency due to the speed of light.

Let us begin by examining data communication systems of various types. It turns out that there are a few key parameters of interest in any data network system. These are:

C = Capacity of the network (Mb/s)
b = Number of bits in a data packet
L = Length of the network (miles)

It is simplest to understand these quantities if one thinks of the network simply as a communication link. One can combine these three parameters to form a single critical system parameter, commonly denoted as a, which is defined as:

$$a = 5LC/b \qquad (1)$$

This parameter is the ratio of the latency of the channel (i.e., the time it takes energy to move from one end of the link to the other) to the time it takes to pump one packet into the link. It measures how many packets can be pumped into one end of the link before the first bit appears at the other end [12]. The factor 5 appearing in the equation is simply the approximate number of microseconds it takes light to move one mile.[2] Now, if we calculate this ratio for some common data networks, we find the values shown in Table 1: Note the enormous range for the parameter a. At one extreme, namely, local area networks, it is as small as 0.05, while at the other extreme, namely, a cross-country gigabit fiber optic link, it is as large as 15,000. This is a range of nearly six orders of magnitude for this single parameter!

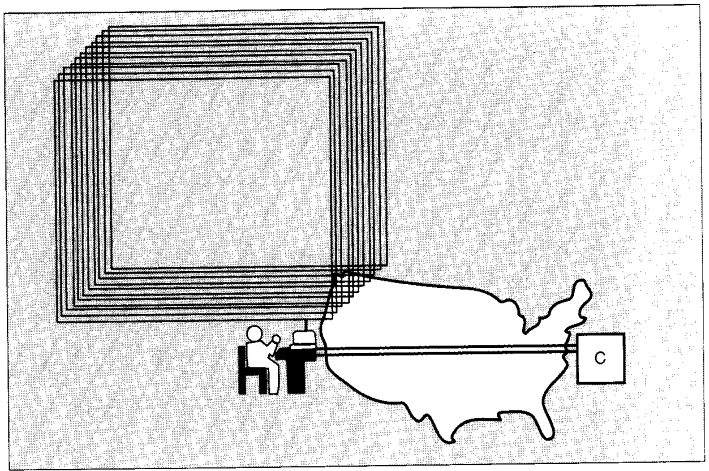

■ **Figure 2.** *Sending a 1-Mb file across the U.S. via an X.25 network.*

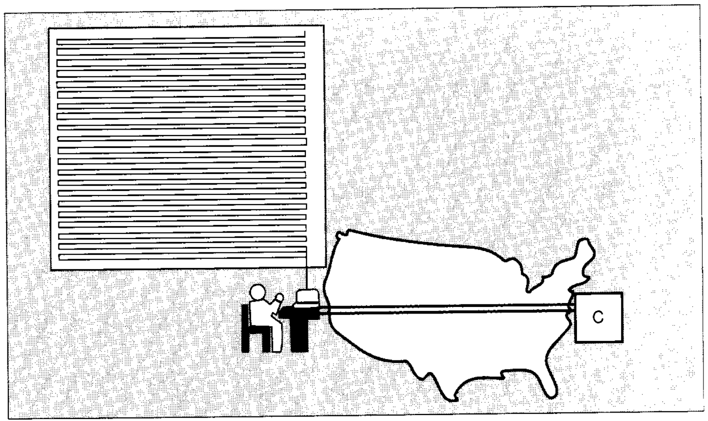

■ **Figure 3.** *Sending a 1-Mb file across the U.S. via a T1 channel.*

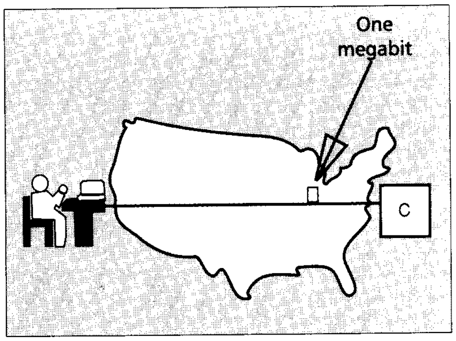

■ **Figure 4.** *Sending a 1-Mb file across the U.S. via a 1.2-Gb link.*

We see that a grows dramatically when we introduce gigabit links. So we naturally must ask ourselves if networks made out of gigabit links are different in some fundamental way from those made out of kilobit or megabit links. There are two cases of interest to consider. First, we have the case that a large number of users are each sharing a small piece of this large bandwidth. In this case it is fairly clear that to each of them, a gigabit network looks no different from today's networks.

However, if we have a few users each sending packets and files at gigabit speeds, then we do

---

[2] *Throughout this paper we make the simplifying assumption that light propagates in a fiber optic channel as quickly as it propagates in free space.*

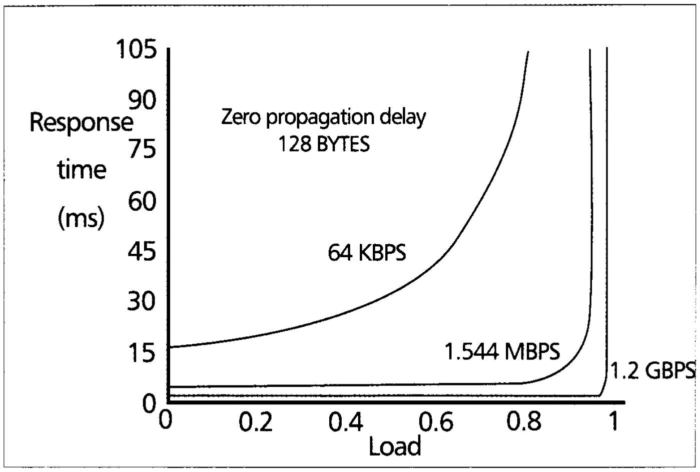

■ **Figure 5.** *Response time vs. system load with no propagation delay.*

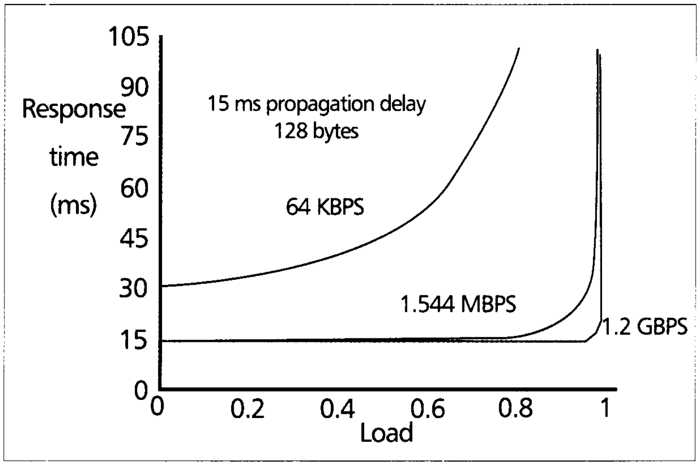

■ **Figure 6.** *Response time vs. system load with a 15 ms propagation delay.*

see a change in behavior and we do run into new problems. At these speeds, *a* gets very large. To see the effect of this change, let us consider the following scenario. Assume we are sitting at a terminal and wish to send one megabit across the United States to some remote computer as shown in Fig. 1.

Now, if the speed of the communication channel we have available is 64 kb/s (as in, say, an X.25 packet network), then, as shown in Fig. 2, the first bit of this transmission will arrive at the East Coast computer after approximately 1000 bits have been pumped into the channel. Thus we see that the channel is buffering roughly 0.001 of the message; that is, there is 1000 times as much data stored in the terminal's buffer as there is in the channel. Clearly, if we had a higher-speed channel, the time to transmit our 1 Mb file could be reduced. That is, we can benefit from more bandwidth.

	Capacity (C) MBPS	Pkt Lngth (b) Bits	Prop Delay (*t*) Microsec	Ratio "a"
Local net	10.00	1000	5	0.05
Wide Area net	0.05	1000	20,000	1.00
Satellite	0.05	1000	250,000	12.50
Fiber link	1,000.00	1000	15,000	15,000.00

■ **Table 1.** *Effect of parameters: propagation delay/Packet Tx time — an enormous variation.*

Thus, let us now increase the speed of the channel and use a T1 channel (1.544 Mb/s). In Fig. 3 we show this new configuration. Now we find that the terminal is buffering roughly 40 times as much data as is the channel. Once again, we see that we can benefit from more bandwidth.

Let us now increase the channel speed to a gigabit channel; in particular, we will assume a 1.2 Gb/s link (the OC-24 SONET offering). This case is shown in Fig. 4 where we see the entire 1 Mb file as a small pulse moving down the channel. Indeed, the pulse occupies roughly only 0.05 of the channel "buffer." It is now clear that more bandwidth is of no use at all in speeding up the transmission of the file; it is the latency of the channel that dominates the time to deliver the file!

Therein lies the fundamental change that comes about with the introduction of gigabit links into nationwide networks. Specifically, we have passed from the regime (of pre-gigabit networking) in which we were capacity limited, to the new regime of being latency limited in the post-gigabit world. Things do indeed change (as we shall see below). The speed of light is the fundamental limitation for file transfer in this regime! And the speed of light is a constant of nature which we have not yet been able to change!

In the considerations above, we assumed that our file was the only traffic on the link. Let us now consider the case of competing traffic with smaller packets. Indeed, let us now assume that we have the classical queueing model of a Poisson stream of arriving messages requesting transmission over a communication link, where each message has a length which is exponentially distributed with a mean of 128 bytes (i.e., a classic M/M/1 queueing system) [13]. If, as usual, we let ρ denote the system utilization factor, then $\rho = \lambda(1024/C)$ where λ is the arrival rate (messages per microsecond) and C is the channel capacity (Mb/s). In this situation, we know that T, the mean response time (milliseconds) of the system (i.e., the mean time from when the message arrives at the tail of the transmit queue until the last bit of the message appears at the output of the channel, including any propagation delay), is given by

$$\mathbf{T} = \frac{1.024/\mathbf{C}}{1-\rho} + \tau \qquad (2)$$

where τ is the propagation delay (i.e., the channel latency) in milliseconds.

Let us ask ourselves if gigabit channels actually help in reducing the mean response time, T. In Fig. 5, we show the mean response time (in millisec) versus the system load ρ for three different channel speeds. In this figure, we assume that the speed of light is infinite, and so $\tau = 0$. The channel speeds we choose are the same as those considered above, namely 64 kb/s, 1.544 Mb/s and 1.2 Gb/s. We note a significant reduction in T when we increase the speed from 64 kb/s to 1.544 Mb/s; thus, the faster T1 channel helps. However, note that when we go from 1.544 Mb/s to 1.2 Gb/s, we see almost no improvement. (The only region in which there is an improvement with gigabits is at extremely high loads, a situation to be avoided for other reasons). As far as response time is concerned, gigabits do not help here!

One might argue that the assumption of zero propagation delay has biased our conclusions. Not so; in Fig. 6 we show the case with a 15-ms propagation delay, (i.e., the propagation delay across the USA) and we see again that gigabits do not help.

We can sharpen our treatment of this latency-versus-bandwidth discussion as follows. Let us assume that we have an M/M/1 model as above, where the messages have an average length equal to b bits. Assume we wish to transmit these files across the United States, as in the earlier figures. Now, as can be seen from Eq. (2), there are two components making up the response time, namely, the queueing-plus-transmission time delay (the first term in the equation) and the propagation delay (τ). In this paper, we have been discussing the relative size of each of these and we referred to regions of bandwidth- limited and latency-limited systems. Let us now make those concepts more precise. We choose to define a sharp boundary between these two regions. In particular, we define this boundary to be the place where the two terms in our equation are exactly equal, namely, where the propagation delay equals the queueing-plus-transmission time delay. From Eq. (2) we see that this occurs when the bandwidth of the channel takes on the following critical value,

$$C_{CRIT} = \frac{1000b}{(1-\rho)\tau} \tag{3}$$

In Fig. 7, we plot this critical value of bandwidth (on a log scale) versus the system load ρ; we have drawn this plot for the case of τ = 15 ms and a message length of one megabit. Above this boundary, the system is latency limited, which means that more bandwidth will have negligible effect in reducing the mean response time, T. Below this boundary, the system is bandwidth limited which means that it can take advantage of more bandwidth to reduce T. Note that for these parameters the system is latency limited over most of the load range when a gigabit channel is used; this means that for these parameters, a gigabit channel is overkill so far as reducing delay is concerned.

We repeat this plot in Fig. 8 for a number of different message sizes. Without labeling the regions, the same comments apply, namely, systems above the curve are latency limited, and below they are bandwidth limited. We note that gigabit channels begin to make sense for message sizes of size 10 megabits or more, but are not helpful for smaller file sizes. This comment about message size refers to the file size that the user application generates; the fact that ATM uses 53-byte cells has little to do with this comment.

Figures 7 and 8 apply to the case of a cross country link (i.e., with a propagation delay of roughly 15 ms). For other than τ = 15 ms, the critical bandwidth which defines the boundary is given from Eq. (3).

Other Issues

We have dealt with the latency-bandwidth tradeoff for gigabit networks in this paper. Of course there are a number of other issues to be addressed in gigabit nets, some of which we choose to mention in this section.

Consider the example from the previous section,

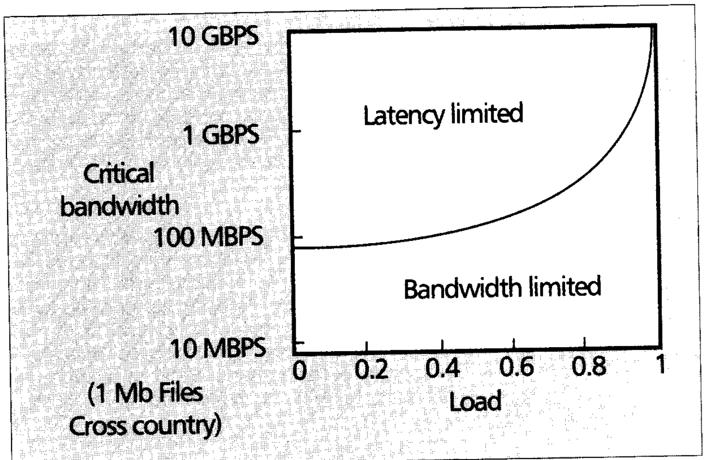

■ **Figure 7.** *Bandwidth vs. system load for a 1-Mb file sent across the U.S.*

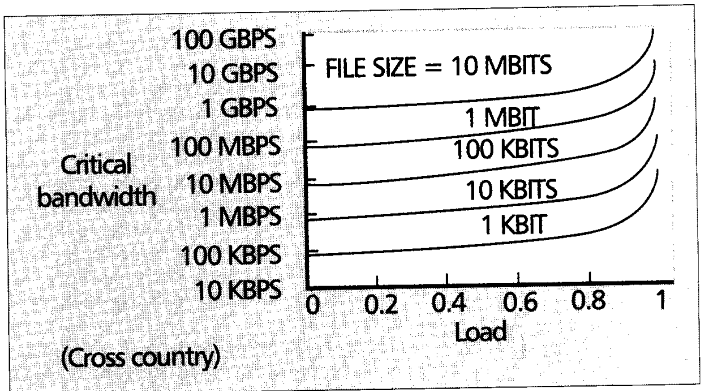

■ **Figure 8.** *Bandwidth vs. system load for files sent across the U.S.*

namely, a gigabit link spanning the United States. Suppose we start transmitting a file a time t=0. Roughly 15 ms later, the first bit will appear across the country. Now suppose that the receiving process decides immediately that it cannot accept this new flow which has begun. By the time the first bit arrives, however, there are roughly 15 million bits already in the pipe heading toward this receiving process! And, by the time a stop signal reaches the source, another 15 million bits will have been launched! It does not take too much imagination to see that we have a problem here. It is basically a congestion control and flow control problem. Clearly, a closed control feedback method of flow control is too sluggish in this environment (due, once again, to latency). Some other forms of control must be incorporated. For example, one could use rate-based flow control in which the user is permitted to transmit at a maximum allowable rate.

Moreover, at the application level, it is important to find ways to hide this latency, in order to get full advantage of the gigabit links and of the high performance processors attached to a gigabit network. One way to hide latency is to use some form of parallelism (or pipelining) such that while one process is waiting for a response, another process, which does not depend upon this response, may proceed with its processing.

Another issue has to do with the maximum attain-

Gigabit networks have forced us to deal with the propagation delay due to the finite speed of light.

able efficiency that one can obtain by taking advantage of statistical multiplexing of bursty sources in a gigabit environment. If we have a large number of small bursty sources, then statistical multiplexing takes exquisite advantage of the Law of Large Numbers [13], and allows one to drive these channels at very high efficiencies. However, if we have a small number of large sources, then the multiplexing does not usually lead to very high efficiencies. This is because statistical smoothing of a small number of sources is not sufficient to bring about the advantages of statistical multiplexing. Furthermore, if we have a large number of nonhomogeneous sources, one must calculate the effective number of such sources in order to calculate the efficiency to be expected from multiplexing [14].

Conclusions

The major conclusion of this paper is to recognize that gigabit networks have forced us to deal with the propagation delay due to the finite speed of light. Fifteen milliseconds to cross the United States is an eternity when we are talking about gigabit links and microsecond transmission times. As we saw earlier, the propagation delay across the USA is forty times smaller than the time required to transmit a 1-Mb file into a T1 link. At a gigabit, the situation is completely reversed, and now the propagation delay is 15 times larger than the time to transmit into the link. We have moved into a new domain in which the considerations are completely reversed. We must rethink a number of issues. For example, the user must pay attention to his file sizes and how latency will affect his applications. The user must try to hide the latency with pipelining and parallelism. Moreover, the system designer must think about the problems of flow control, buffering, and congestion control. Some form of rate-based flow control will help the designer here. He must also design algorithms which make rapid decisions if enormous buffer requirements are to be avoided. The designer cannot depend on global state information being available in a timely fashion; this affects his choice of control algorithms. In many ways, the user will see gigabit networks as being different from megabit networks; the same is true for the designer/implementer.

Much more research must be done before we can claim to have solved many of the problems that this new environment has exposed. We must solve these problems in the near future if we are to enjoy the benefits that fiber optics has given us in the form of enormous bandwidths.

References

[1] Kleinrock, L., *Queuing Systems, Volume II: Computer Applications*. New York, Wiley, 1976.

[2] Tanenbaum, A., *Computer Networks, 2nd Edition*. New Jersey, Prentice Hall, 1988.

[3] Lowe, S.J. "Plugging into frame relay for speed and flexibility", *Data Communications Magazine*, pp. 54-62, April 1990.

[4] Mier, E.E., "New signs of life for packet switching", *Data Communications Magazine*, pp. 90-106, December 1989.

[5] Hui, J. Y. *Switching and Traffic Theory for Integrated Broadband Networks*. Boston, Kluwer, 1990.

[6] Stallings, W., *ISDN: An Introduction*. New York, Macmillan, 1989

[7] Thurber, K. J., "Getting a handle on FDDI", *Data Communications Magazine*, pp. 28-32, June 1989.

[8] "Generic system requirements in support of switched multi-megabit data service," Bellcore Technical Advisory, TA-TSY-000772, Issue 3, October 1989.

[9] Kessley, G. C., "IEEE 802.6 MAN," *LAN Magazine*, pp. 102-16, April 1990.

[10] Handel, R., "Evolution of ISDN towards broadband ISDN," *IEEE Network Magazine*, pp. 7-13, January 1989.

[11] Ballart, R. and Y.C. Ching, "SONET: Now it's the standard optical network," *IEEE Communications Magazine*, pp. 8-15, March 1989.

[12] Kleinrock, L., "Channel efficiency for LANs," *Local Area and Multiple Access Networks*, Computer Science Press, Raymond L. Pickholtz, ed., pp.31-41, 1986.

[13] Kleinrock, L., *Queuing Systems, Volume I: Theory*. New York, Wiley, 1975.

[14] Kleinrock, L., "Performance of distributed multiaccess computer-communication systems", *Proceedings of IFIP Congress 77*, pp. 547-52, August 1977.

Biography

LEONARD KLEINROCK is professor and chairman of computer science, University of California, Los Angeles. He is the principal investigator for the DARPA Parallel Systems Laboratory project at UCLA. Dr. Kleinrock is the founder and CEO of Technology Transfer Institute, a computer/communications seminar, conference, and consulting organization located in Santa Monica, Calif. He is a fellow of the IEEE, member of the National Academy of Engineering, recipient of the Ericsson Prize for Outstanding contributions in packet awitching technology, and recipient of the 12th Marconi International Fellowship Award for pioneering work in the field of computer networks.