#/139

Special Editorial

Packet Switching Principles*

Reprint from
Journal of Telecommunication Networks
11 Taft Court, Rockville, Maryland 20850

Leonard Kleinrock
University of California

If you attended any of the large computer conferences of the mid-1960s, you surely heard at least one panel heatedly debating the central issue of data communications. On the one side, we heard the representatives from the data processing industry complaining bitterly that no suitable facilities existed for efficient communication of computer-generated data. On the other side, we heard the representatives from the communications industry claiming that each country was essentially a copper mine of interlaced telephone channels which could be used for data communication. However, this telephone network proved to be a completely unsatisfactory solution for computer-communications since the typical use of a computer terminal was to send less than one second's worth of data, and the telephone plant required tens of seconds to set up the connection and (in the United States, for example) there was a minimum duration charge for three minutes! This impasse led to a revolutionary new method for using communication channels which has come to be known as *packet switching*.

Before describing the principles behind packet switching, it is important to elaborate on the nature of computer-generated data. Information processing devices (especially computer terminals) tend to generate data in widely separated bursts. Indeed, computer terminals almost never warn you ahead of time when they need to send data down a communication channel; they seldom tell you how much data they wish to send. Most of the time they send nothing at all, but when they do occasionally want the channel, they want it immediately! This nasty combination produced most of the problems we faced in data communications.

In the mid-1960s, the classical technique for assigning voice channels to telephone conversations was *circuit switching*. With this method, a connected path of channels would be set up for the duration of a "call" in a dedicated fashion. When this technique was applied to data communication, these dedicated communication channels were idle most of the time waiting for occasional bursts of data.

Clearly, what was needed was some method of assigning the communication channels to the terminals only during those few instances when required. This would alleviate the enormous inefficiencies of the classical techniques in this new environment. The solution to the problem required that *intelligence* be installed in the switches which assigned the channels. The cost of this computerized intelligence had been falling dramatically because of the unbelievable progress in microelectronics—a trend which continues to this day. By 1970, the savings in communications as a result of the rapid dynamic channel assignment by these intelligent switches exceeded the cost of the switches themselves, and so the economic forces made a technology such as packet switching unavoidable. In fact, by 1969, a fledgling packet-switching network (the U.S. Defense Department's ARPANET) was already operational and expanding.

Packet switching works as follows. Imagine that a message, such as the text of the previous sentence, must be sent from a source to some remote destination through a network (as shown in Figure 1) using the technology of

---

*Presented at the 1982 L. M. Ericsson Prize Ceremony, Stockholm, Sweden, May 5, 1982.

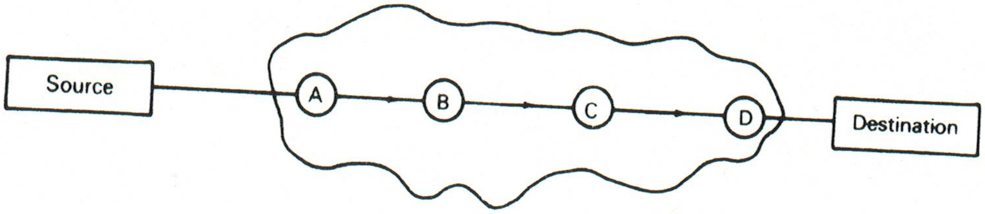

Figure 1. Communication network.

packet switching. The text of that sentence consists of thirty-four characters (including spaces). Now, in packet switching, the text is broken up into segments called *packets*, and these packets are transmitted through the network separately, passing from one node to another node, hop by hop. Each packet has a maximum allowable length which, for our example, we will assume to be thirty-two characters. Thus, for our sentence, we have a message consisting of two packets, one of length thirty-two characters and another of length two characters. In addition to the data characters, each packet carries with it a coded version of the address of the destination. When the first packet is presented to node A, the intelligent switch (in that node) will select the first step in that packet's journey (obviously, it will select node B for this simple example). If the selected channel connecting these two nodes is busy transmitting another packet, then our first packet will wait in a queue until its turn for transmission comes up, at which time transmission begins. When it is received at node B, the process is repeated, and we note that the packet "hops" from node to node through the network using only one channel at a time, possibly queueing at busy channels on its journey. Each packet is treated in this fashion. Thus, many packets of the same message may be in transmission simultaneously, thereby allowing packets to be "pipelined" down the chain.

This procedure differs significantly from circuit switching. In particular, we no longer dedicate any channels ahead of time; they are used only on demand in a dynamic fashion and need be acquired only one at a time instead of in groups. It is this feature of packet switching which provides the efficiency of channel use. The decomposition into packets provides a reduction in transmission time for the entire message through the network because of pipelining. Furthermore, the intelligent switches are constantly sensing conditions to select good paths through the network; this provides considerable improvement in reliability since alternate paths will be used in the case of failure or congestion along primary paths. Since the source and destination are not directly connected (as was the case with circuit switching), it is possible for devices of vastly different speeds to communicate with each other through the network; that is, a packet switching network provides the capability of speed conversion. Thus, with packet switching, we have an efficient, rapid, reliable, and flexible communication system.

With any system as complex as a large-scale packet switching network, it is necessary to develop methods for performance evaluation and for system design. If one creates a mathematical model of packet networks, one finds that an exact analysis of the system behavior is hopeless! Our current tools are far too crude for this purpose. In spite of this, we had to find some method of analysis. To our good fortune, it is possible to introduce certain assumptions which simplify the problem to manageable proportions and which yield excellent tools for approximate system performance evaluation. The same situation exists with regard to network design; we have no adequate methods for least-cost design, but we have developed effective techniques for low-cost design.

It is worthwhile to point out that much of the effectiveness of packet switching comes from the "largeness" of the networks in which we use the technique. Specifically, we have found that two *resource-sharing principles* come into action here. Recall that the "bursty" behavior of the terminals gave rise to many of our communication problems. There exists a principle, known as the "law of large numbers," which basically states that a large number of bursty data sources will collectively behave in a very "smooth" fashion whereby a predictable and steady flow of data will emanate from the group as a whole. (The insurance companies know this—they know almost exactly how many people will die next year. They simply don't know which particular individuals will die, so they "bet" with everybody, and they usually win the bet; that is, the mortality tables are extremely accurate because of the law of large numbers.) The second principle is often referred to as the "economy of scale" principle. This principle states that if we begin with a service system which handles a certain load and which is endowed with a certain capacity for handling that load, then doubling the load and doubling the capacity will cause the response time of that system to improve by a factor of two! In order to see how these principles apply to our networks, let us examine the components of a typical system. In Figure 2 we show a more complete picture of such a network. At its periphery we note the various kinds of terminals and computer facilities. It is the purpose of the "communications sub-network" to provide communications among these various devices, and in this sub-network we see the intelligent switching computers which are connected with the high-speed communication lines. Observe that this sub-network carries traffic from *many* terminals and computers (hence, we expect our first resource-sharing principle, the law of large numbers, to work for us). Since the network is handling so much traffic from this large number of devices, we require *large-capacity channels* and so we expect to reap the benefits from the economy of scale (our second resource-sharing principle). In fact, these two resource-sharing principles apply not only to the expensive communication channel capacity, but they also apply to the two other key network resources, namely, the storage capacity and the processing capacity of the intelligent network computers themselves. In summary, then, it is precisely when large populations dynamically share large capacity resources that we enjoy significant performance efficiencies; packet switching networks are prime examples of such systems.

However, we have an unresolved problem. In Figure 2, we note that "remote" terminals must pass through a small network of their

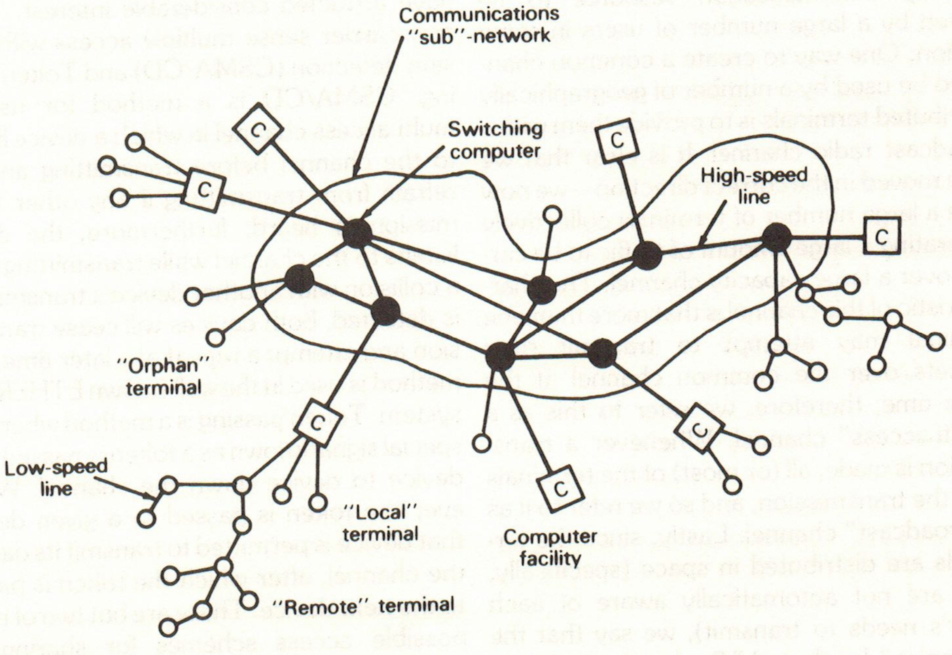

Figure 2. The structure of a computer-communication network.

own before they reach the communication sub-network in which all the efficiencies of large shared systems are to be found. If we concentrate on the channel labelled "Low-speed line" in that figure, we can inquire regarding the efficiency of its use. We see that only *one* terminal can "share" this line (and one is *not* a large number); similarly, since only one terminal's traffic is passing over this channel, then the load is small and so also will be the capacity. We, therefore, conclude that neither of our resource-sharing principles applies to this "local access network." Things are fine deep in the sub-network, but we are in trouble at the periphery of our system. Unfortunately, a significant fraction of the system cost is invested in these local access networks. What shall we do in this case? The consideration of this question ushered in the technology of local access, which has recently become the major focus of computer-communications. In fact, in the form of local area networks, it has become the critical element in the realization of the automated office.

In order to resolve the problem of local access, we once again take a cue from our two resource-sharing principles. These principles tell us that we must provide a large capacity communication resource to be shared by a large number of users in some fashion. One way to create a common channel to be used by a number of geographically distributed terminals is to provide them with a broadcast radio channel. It is clear that we have moved in the correct direction—we now have a large number of terminals collectively generating a large amount of traffic to be carried over a large capacity channel. The characteristic of this channel is that more than one terminal may attempt to transmit their packets over the common channel at the same time; therefore, we refer to this as a "multi-access" channel. Whenever a transmission is made, all (or most) of the terminals hear the transmission, and so we refer to it as a "broadcast" channel. Lastly, since the terminals are distributed in space (specifically, they are not automatically aware of each other's needs to transmit), we say that the channel is "distributed." Such a channel may, therefore, be described as a *multi-access broadcast distributed* channel. In the recent past, considerable effort has gone into the design of such channels in order to find ways to use them efficiently. A number of exciting access methods based on packet switching concepts have been developed and experimented with. Of more interest, perhaps, is the application of these multi-access techniques not to radio but to in-building communication systems commonly known as *local area networks*.

Local area networks are intended to provide high-speed, low-cost communications within a building or between buildings. The typical application is for office automation, namely, to connect together the many information-processing devices that one finds in the modern office; these include terminals, work stations, microcomputers, minicomputers, printers, storage devices, electronic blackboards, displays, etc. We have seen a veritable explosion in the number of implementations and products being offered to satisfy the needs of in-building communications. In most of these systems, we see the basic principles of packet switching applied once again, namely, dynamic access from many bursty devices to a common wideband channel. Among the access schemes, two have attracted considerable interest. These are: carrier sense multiple access with collision detection (CSMA/CD) and Token Passing. CSMA/CD is a method for using a multi-access channel in which a device listens to the channel before transmitting and will refrain from transmitting if any other transmission is heard; furthermore, the device listens to the channel while transmitting and if a collision with another device's transmission is detected, both devices will cease transmission and attempt a repeat at a later time. This method is used in the well-known ETHERNET system. Token passing is a method whereby a special signal known as a token is passed from device to device down the channel. Whenever the token is passed to a given device, that device is permitted to transmit its data on the channel, after which the token is passed to the next device. These are but two of many possible access schemes for sharing the capacity of the channel. Aside from the access method, there is the choice of medium for the communication itself. One common medium is twisted pairs of copper wire; these have been used for the traditional private automatic branch exchange (PABX) for handling voice switching and have recently been upgraded into computerized branch exchanges (CBX), which are capable of efficiently handling data as well as voice. Coaxial cable is an attractive medium as well, both in its baseband and broadband implementations. Beyond coaxial cable, the remarkable technology of optical fiber channels is extremely exciting; with this medium, enormous bandwidths are available in a medium which is small, lightweight, flexible, low loss, immune to electromagnetic interference, immune to high temperature, etc. The major impediment to the widespread use of optical fiber at present is the lack of a cost-effective method for connecting and tapping the cable; however, when the economics become competitive, optical fibers will find an enormous use in local area networks and in point-to-point long-haul communications. These revolutionary developments in local communications all take advantage of the packet switching technology in that they dynamically assign capacity among a large number of bursty devices.

The application of these revolutionary developments in communications does not stop with pure data transmission. The requirement to send video, facsimile, voice, graphics, etc., as well as data, certainly exists in today's automated environment. The technology for these applications is currently coming into place. We have already seen a large number of packet-switching networks spring up around the world. Many local area networks are already in place. Satellite data networks are in operation. Many of these separate networks have already been interconnected, and more are being attached every week. We are entering an era of worldwide access to data and to information processing resources, this access being made available through the sophisticated and cost-effective packet-switching computer networks we have been discussing. The 1970s was the era of computer network development. The 1980s will be the decade of network applications and will provide the penetration of the information revolution into many additional areas of activity.

About the Author

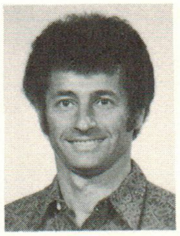

Dr. Leonard Kleinrock received his B.S. degree in electrical engineering from the City College of New York in 1957 (evening session) and his M.S.E.E. and Ph.D.E.E. degrees from the Massachusetts Institute of Technology in 1959 and 1963 respectively. While at M.I.T., he worked at the Research Laboratory for Laboratory for Electronics as well as with the computer research group of Lincoln Laboratory in advanced technology. He joined the faculty at UCLA in 1963. His research interests focus on computer networks, packet radio systems, and local area networks. He has had over 120 papers published and is the author of three books, *Communication Nets: Stochastic Message Flow and Delay*, 1964; *Queueing Systems, Volume I: Theory*, 1975; *Queueing Systems, Volume II: Computer Applications*, 1976 and also the *Solutions Manual for Queueing Systems, Volume I*, 1982. Professor Kleinrock served as the head of the UCLA Computer Science Department Research Laboratory and is a well-known lecturer in the computer industry. He is principal investigator for the ARPA Advanced Teleprocessing Systems contract at UCLA and co-principal investigator for the NSF Advanced Network Environment for Distributed Systems Research Project. He was recently elected to the National Academy of Engineering, is a Guggenheim Fellow, an IEEE Fellow, and serves on the Boards of Governors of various advisory councils in the computer field. He is a member of the Science Advisory Committee for IBM. He has received numerous best paper and teaching awards, including the ICC '78 Prize Winning Paper Award, the 1976 Lanchester Prize for outstanding work in Operations Research, and the Communications Society 1975 Leonard G. Abraham Prize Paper Award. In 1982, as well as having been selected to receive the C.C.N.Y. Townsend Harris Medal, he was co-winner of the L. M. Ericsson Prize, presented by His Majesty, King Carl Gustaf of Sweden, for his outstanding contributions in packet switching technology.