A Decade of Network Development

Leonard Kleinrock *

University of California

ABSTRACT The network technologies developed over the past decade, due largely to incredible advances in integrated chip production, have set the stage for an enormous DP revolution the impact of which will soon be felt in the business community and in the home and consumer markets. We have been able to identify those fundamental principles which render this new technology so cost effective. End-to-end communication services are now coming into place which will pass data, voice, video, fax, graphics, etc. from an individual terminal in an office or a living room across vast distances to remote processing facilities. The technological issues have been addressed, and in this paper we describe some of the successes and some of the remaining problems to be solved. As engineers, we cannot restrict our efforts to the purely technical (and "nice") problems, but rather we must accept the further responsibility of familiarizing ourselves with the overall environment in which our "products" must perform; if we fail in this regard, the world of commerce, trade and business will ignore our best efforts as irrelevant to its needs.

1. The Big Picture and the Real Issues

The age of modern telecommunications began more than a decade ago with the rise of the large-scale packet-switching networks. The first such network was the U.S. Defense Department's ARPANET which became operational late in 1969 as a four-node network [10]. In the decade following that event, we have seen enormous progress in packet communications in a world where data is beginning to dominate.

However, in all this time we have gained little real understanding of these systems. Sure, we know how to move data and we even know how to design "reasonably good" systems. But we really don't understand what distributed communications is all about; we have no idea what the issues of concurrency in processing and in communications actually involve; and we have precious few measures which provide us a meaningful metric of real system performance.

Moreover, in these ten years, as a technology we have made relatively little impact on the real world. Fortunately, that will not remain the case for long. The stage is now set for an enormous penetration in the world of business, commerce, government, education, home services, etc. This creates a challenge, an opportunity and a burden upon us, the system designers, in terms of providing systems which not only perform well according to *our* measures but also which perform well according to the needs of industry and the complex world of the end user.

---

*Leonard Kleinrock, Computer Science Department, University of California, Los Angeles, California 90024. Material based on a keynote speech at 7th Data Communications Symposium, Oct. 27, 1981, Mexico City, Mexico. This research was supported by the Advanced Research Projects Agency of the Department of Defense under Contract MDA 903-77-C-0272.

0276-0037/82/0101-11$3.85/0

Most of us recognize what has been the single most significant development which has acted as the prime mover for this revolution in telecommunications. It is simply the fantastic *microchip technology* and the continued phenomenal improvement in the price-to-performance ratio of these silicon devices. In Figure 1 we see a view of the interplay among technology, processing, communications, systems, end user tools, and the revolution (which has already begun). At the top of this chain we see the *integrated chip* from which we branch off and show its impact in two areas: processing and communications. In the field of *processing*, the rise of personal computers is due mostly to the large scale integrated (LSI) technology which has produced computers on a chip and highly efficient memory systems to go along with them. This also includes printing peripherals whose costs have begun to drop (but which still form a sizable investment for small machines). Following that, in increasing size of processors, we come to the pervasive small business machines which have been based on the eight-bit chips (Z80A, INTEL 8080, etc.) many of which are evolving to systems using 16-bit chips (Z8000, INTEL 8086, Motorola 68000) and soon we will likely see 32-bit small business machines. A number of competing operating systems have already entered the foray including CP/M, UNIX, OASIS, etc. From there we encounter the minicomputers which have grown in power enough to compete with the maxis (next on the list) and then on to the giant supercomputers. All these systems have enjoyed a decreasing price-to-performance ratio due largely to the integrated chip revolution.

The other branch leading from the chip technology is that in *communications* (the subject near and dear to the readers of this journal). Here we have seen a large variety of significant developments over the last decade, most of them concentrated in the last three to five years. The list in Figure 1 is organized in increasing order of the physical distances over which the communications is required to operate. We begin with the problem of providing communications on the tiny integrated chip itself. This turns out be a growing problem as the dimensions of the logical elements on a chip grow smaller and smaller (now challenging sub-micron dimensions) leading to chips with millions of logical elements. Following this we see the recent proliferation of a large number of local network systems providing communications among terminals, computers, peripherals, etc. within the environs of, say, a single office building. A large number of network architectures (perhaps too many) have been implemented including loops and rings, buses, stars, etc. (the most common of the bus architectures is perhaps the well-known ETHERNET). In addition we have seen the private automatic branch exchange (PABX) evolve into a computerized branch exchange (CBX) which has come to compete with the ETHERNET-like systems for in-building communication. Beyond that we have seen very significant improvements in modems (on a chip), modem-sharing units, and smart multiplexers (notably the statistical multiplexer): all these developments are, again, principally based upon the advancement in chip technology. In providing building-to-building and within-city communications, we encounter the problem known as local data distribution networks which tend to be based on some form of radio communication. In the case of mobile terminals, an entire technology of packet radio communications has arisen with its attendant and difficult problems to be described below. The ARPANET spawned a worldwide development of terrestrially-based networks whose backbone communications typically have been broadband telephone channels forming intercity distributed networks across nationwide dimensions. Beyond that we have seen the emergence of satellite services (both point-to-point and broadcast) to provide cost-effective communications over long distances at wide bandwidths and typically crossing hostile terrain such as oceans. All these communication developments are due largely to the advancement in chip technology. (And chips are made of silicon. And silicon is made from the sand of the desert. And wouldn't you know who owns most of the sand!)

The communication components just described have been lashed together to provide a variety of large scale communication *systems* in the last decade. For example, a large number of value added networks (VANs)

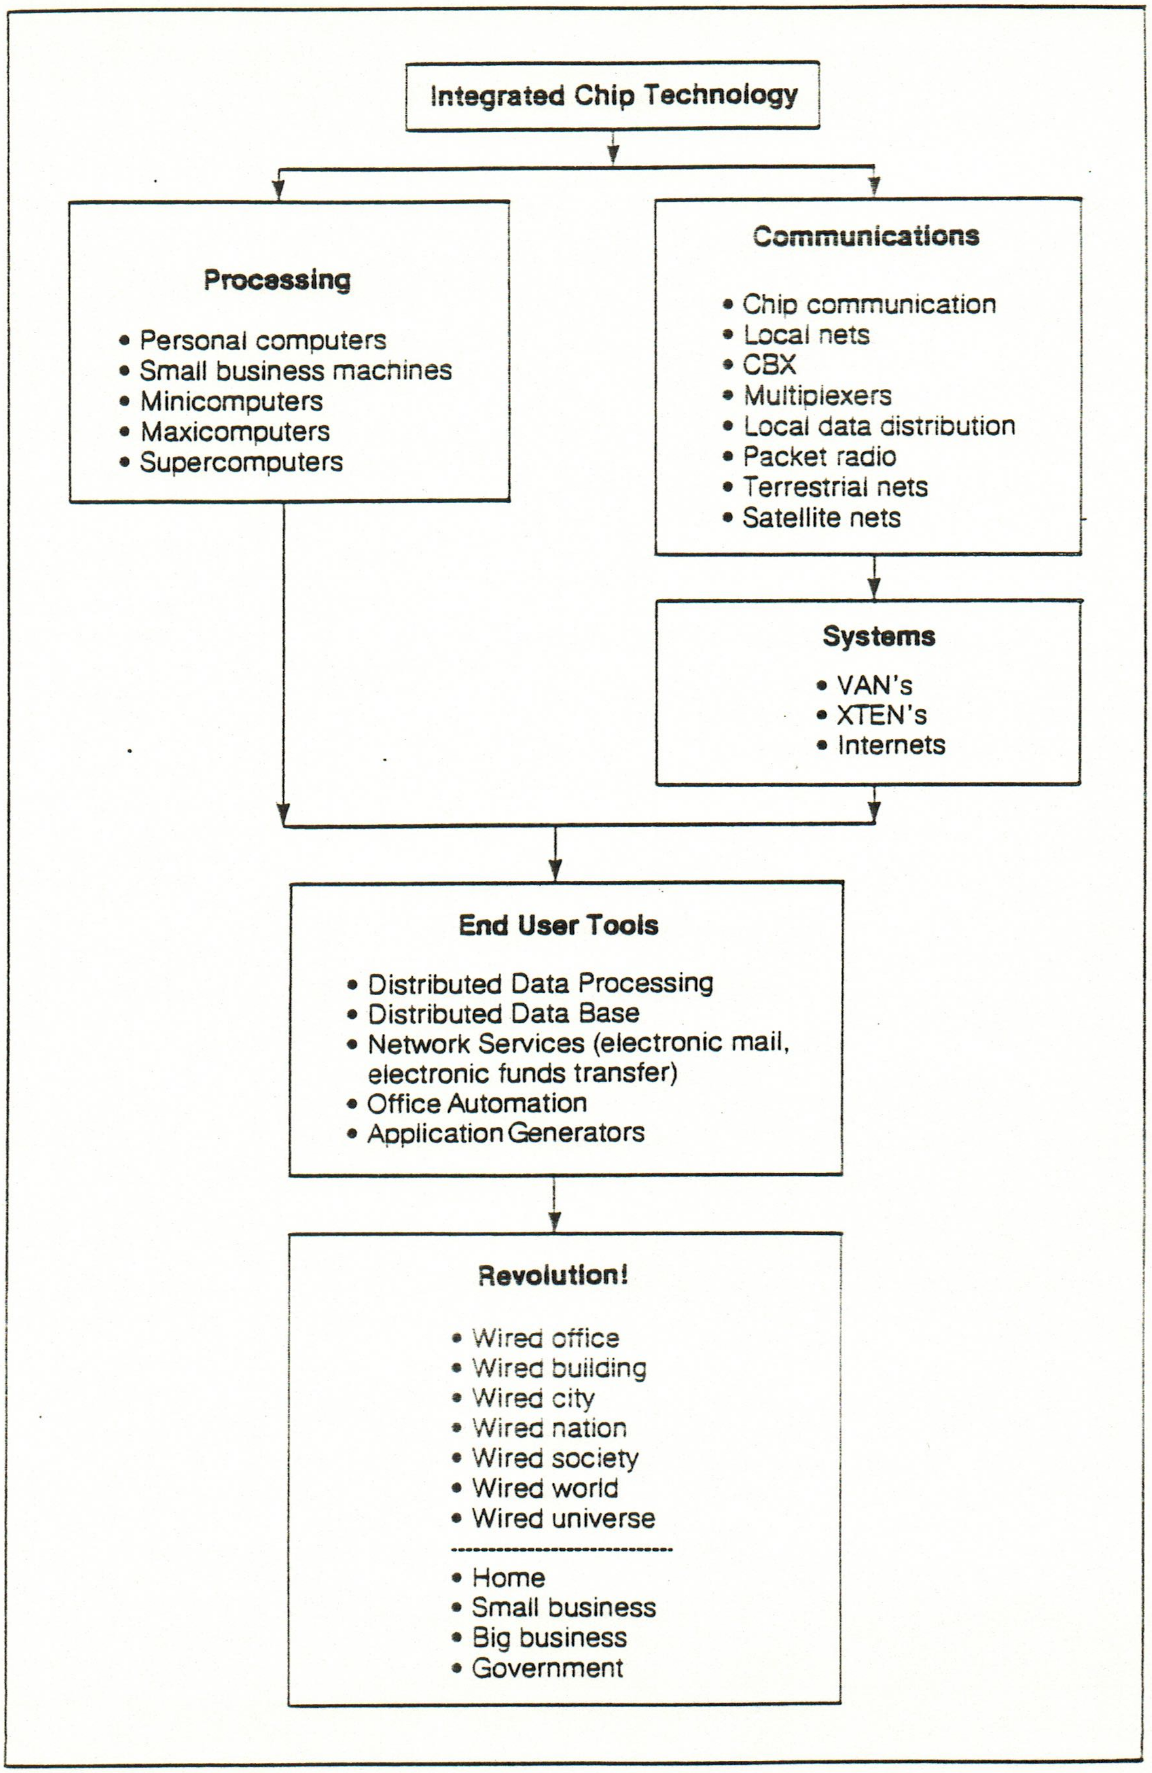

Fig. 1. **The Structure of the Revolution.**

have come into being across the world (TELENET, TYMNET, TRANSPAC, DATAPAC, PSS, EURONET, etc.). Moreover, a number of end-to-end specialized communication systems services have been described (such as Xerox's XTEN system, which, to our great misfortune, was never brought to the marketplace due to economic considerations). Systems similar to XTEN which provide end-to-end services from a terminal in an office, through an in-building local network, to a rooftop antenna, via a local data distribution system across the city, to an earth station feeding a satellite network and back down again through the hierarchy into a remote processing device are sure to emerge in the near future (for example, Isacomm's network and MACOMNET are already discussing such systems which are basically a kind of reincarnation of the XTEN concept). All of this implies the network interconnection of the various components we have described above and this leads to an entire new set of considerations involving internetting.

Referring once again to Figure 1, we see that the processing branch and the communications branch merge once again to provide a collection of irresistible *end user tools*. Among these we list: distributed data processing, distributed data bases, network services (including things such as electronic mail and electronic funds transfer), office automation, application generators, etc.

These end user tools will be packaged together to bring about the *revolution* which is already underway. We will progress from the wired office to the wired building to the wired city to the wired nation to the wired society to the wired world and to the wired universe (if we allow it) [16]. The impact here will extend from the individual personal computer in the home (which will be connected via a communication service to the rest of the world) on to the small business machines (an enormous market) to big business systems to government information systems, etc. The world of telecommunications is upon us! Few people within the world of technology and most people outside that world have any idea of the magnitude of the revolution which is underway. A decade from now when we look back to the early 80's, we will wonder why it was that so few people recognized that the revolution had already begun! It is bound to change almost everything we do in the world of science, commerce, education, business, government, management, work style, recreation, etc. The key is the ability to provide telecommunication networks which tie together all the processing elements, thereby providing a large variety of important services to the end user.

We, as technicians and experts in the field of telecommunications, run a serious danger of losing sight of the big picture in which our technology participates. This is a critical point. The end user really has very little interest in whether an arrival process is Poisson or some other complex stochastic process. He cares not if independence assumptions make sense or if algorithms are known to be optimum. Why should he worry if routing tables in a packet network are updated every half second or every half minute? These are issues with which you and I are concerned. But we are concerned with them in order to provide a product—a product which is useful out there in the real world (a world which understands nothing about our specialities). Their concerns revolve around measures which we seldom incorporate in our analysis and design. Among these measures we might include some which we have thought about but for which we have been able to give little advice, namely: reliability, availability, maintainability, security, privacy, expandability, single vendor lock-in, backup procedures, terminals supported, protocols supported, safety, interconnectability, transparency, software availability, documentation, etc. In fact, were you or I to go out and purchase a computer tomorrow, I suspect that if the application were at all worthwhile, our biggest concern would be *backup* of our important records, followed perhaps by cost, then by the details of the maintenance contract, etc. Way down the list we would finally find those measures with which all of us are so enamored, namely, throughput, response time, queue lengths, efficiency, stability, optimality, etc. This is not to say that our measures are not important (for indeed they are very important); but we must recognize that these other issues are the ones with which the end user is concerned, and too often we forget that these issues must be addressed in our technological evaluations. It is not enough for us to become excellent analysts—we must understand the environment in which our systems are to function, and it is too often that our technical experts fail in that regard. How many of us really understand the key characteristics of floppy disks and of the recent Winchester cartridges? What is the real market for 56 kilobit per second access from one's terminals and hosts? Of what use is a 500 millisecond response time if the CPU overheats and fails? Do you have any idea how distributed processing and distributed data bases really fit into the practical world of data processing and telecommunications? These broader issues come with experience and exposure, and it is up to you as experts to gain that experience and exposure in addition to developing the expertise of your trade.

2. State of the Art in Modeling and Design

In this section we will briefly review some of the key developments which have taken place over the past decade in our area of expertise (not withstanding the warnings at the end of the previous section) and discuss where further work must be done.

A. Resource Sharing. In the field of modeling and design we continually seek efficient ways to allocate processing and communication resources to a collection of demands which require access to these resources. The difficulty in this allocation procedure comes from the fact that these demands almost never tell you ahead of time when they want the resource, they seldom tell you how long they will need the resource, most of the time they never want the darned thing anyway, but when they do occasionally want it, they want it immediately! Such demands are extremely difficult to handle when the resources they want are scarce and precious. This, then, is the underlying problem in many of our considerations [8].

Indeed the problem involves a tradeoff whereby on the one hand we wish to give good service to the demands, but on the other hand we would like efficient use of the resources. At one extreme we could make a permanent assignment of resources to demands (i.e., let the demands permanently own a given resource). This provides terrific service but leads to inefficient resource use when the demands are bursty (in the sense that they only need the occasional use of the resource). In a heavy traffic situation where the demands do need the resource an overwhelming fraction of the time, then a permanent assignment is not especially bad and we may be willing to accept it. At the other extreme, we might choose to give no permanent assignments at all and allow the demands to grab the resources as they need them. In a light traffic situation where there are plenty of resources and few demands, this seems to be an acceptable procedure since it would be foolish to impose access restrictions in a plentiful environment. However, as the demand load begins to increase, this uncoordinated access scheme leads to conflict situations whereby more than one demand is attempting to access the same resource at the same time. This is known as the problem of *multiaccess*. The question is—how can we resolve such conflicts?

Conflict resolution is a key concept in telecommunications. More than one demand requires the same resource at the same time. What are the possible ways of resolving this conflict? Let us list a few of the more common ones:

1. *Queueing:* In this situation, *exactly* one demand gets access to the resource and the others are forced to join a queue and wait their turn. It matters not if we introduce priority queueing or preemption or some other interesting wrinkle on the queueing structure; the key idea here is that exactly *one gets service while all the others wait*. This is a civilized way to proceed and hopefully one which is commonly seen in the real world.
2. *Splitting:* In this case, the capacity of the resource is partitioned into a number of pieces, one piece given to each of the active demands. That is, *all get served simultaneously at some fractional rate*. The splitting need not be equal although it usually is and one well-known implementa-

tion is frequency division multiplexing (FDM).

3. *Blocking:* In blocking, *exactly one demand gets served and the others are refused service* and asked to leave (i.e., they are "blocked" or "lost"). The only demands which get served are those which arrive when the resource is idle.
4. *Smashing:* In this case, *when more than one demand requests the use of a resource then no one gets served*. This is the most disastrous of all the resolution schemes, and yet it does occur in such technologies as packet radio and local bus architectures using contention.

Of course it is possible to form hybrid mixtures of the above schemes. For example, one could invent an access scheme in which the first two customers were allowed to split the resource and the next 15 would be allowed to queue up and wait for the resource, and all the rest would be blocked from any use at all.

The first and second conflict resolution schemes mentioned above (queueing and splitting) are the more civilized. Let us discuss the gains to be had from these two schemes a bit further. The most important observation regarding queueing and/or splitting is that the resource being shared is not permanently assigned to a demand, but rather its full capacity is *dynamically shared* among those demands currently requiring access to the resource; this is a great improvement over the permanent assignment mentioned above and provides the key to the solution of our problem of satisfying unpredictable bursty demands.

We have two important resource sharing principles at work here. The first principle is the *law of large numbers* which simply states that a large population of demands will present a total load to a system of resources which is equal to the sum of the *average* requirements of each individual demand, and further this total load is a highly predictable quantity. This should be compared to the requirement in a permanently assigned environment in which the total resource capacity required is equal to the sum of the individual *peak* demands rather than the sum of the *average* demands. This smoothing effect of the law of large numbers is key to the resolution of resource assignment in a bursty demand environment. The second resource sharing principle has to do with *scaling*. The scaling result simply says that if we start with a certain demand for throughput and a certain total resource capacity, and if we then increase that throughput by a factor and increase the capacity by the same factor, then the response time in satisfying these demands will be reduced by that same factor. This scaling effect is rather important and can be stated in a number of ways [8]. The net result of these two resource sharing principles says that, if possible, we should take large capacity resources and allow large demand populations to share these resources dynamically. In so doing we get the gains from the smoothing effect of the law of large numbers and the gains from the scaling effect as well.

Now how can we apply these ideas to our telecommunications problem? The answer is, in principle, quite straightforward, namely, to gather together a large collection of demands for communication requirements and offer this population a wideband channel which they must share dynamically on a demand basis. Perhaps the best example of such an implementation is that of statistical multiplexing whereby a collection of terminals generates data to be transmitted over a shared wideband channel; their data requests are placed in a queue of other requests waiting for use of the channel in a first-come-first-served fashion. Indeed a queue is a wonderful device (it is a demand multiplexer) which provides the full bandwidth of the resource to one demand at a time. Each terminal transmits its data quickly at the full channel speed and then immediately releases the channel for use by some other terminal. This provides a "burst" communications mode, and it takes advantage of both of our resource sharing principles.

There are four major technologies which have arisen over this past decade of telecommunications progress and they form the components of our overall communications system. Properly designed, they can be made to operate in the burst mode just described. Chronologically, we first saw terrestrial networks rise in the form of packet-switching systems (e.g., ARPANET, TELENET, TYMNET, DATAPAC, TRANSPAC, etc.); the service here tended to be over nationwide distances.

Following that, satellite communications developed as a technology for data communications and was used to span long distances at wide bandwidths, thereby serving the role of connecting different terrestrial networks and also acting as support for the backbone communications of a single terrestrial network. Next we saw the introduction of packet radio technology to serve as local access to networks. Lastly we have seen an enormous rise in the use of local network systems to provide communications over distances typically less than a few miles and usually serving office buildings, universities, etc.; this tends to be communications on a cable or on a fiberoptic channel which connects directly to the terminals and devices of our networks. Each of these four technologies (terrestrial nets, satellite nets, packet radio systems, and local network systems) takes advantage of our two resource sharing principles in a variety of ways. In the following portions of this section we shall discuss these developments and point out where the open problems are from an analytic and design point of view.

B. Terrestrial Nets. Using networks such as the ARPANET as typical examples of implementations of a packet-switching network, we are all aware that progress in modeling, analyzing and designing these networks has been quite successful over this past decade. We have an excellent model for delay in these networks as given by $T = \Sigma_i (\lambda_i/\gamma) T_i$ [10]. Here T is the mean response time of the network, λ_i is the message traffic on the i^{th} channel in the network, T_i is the mean response time of the i^{th} channel (queueing plus transmission), and γ is the total network throughput. This equation is exact and quite general, requiring essentially no statistical assumptions. The particular form for T_i does depend upon the statistical assumptions made [9], and it is here where the analysis bogs down; however, there are excellent approximations which yield results for T which are quite acceptable for the state of the art. (Some particular network configurations do admit an exact solution, but they are quite limited in scope and will not be discussed further.)

So far as the more difficult areas of routing and flow control are concerned, we find that the situation here is not quite as good. First, in the area of routing procedures, we find that there have been a fair number of implementations as, for example, in the ARPANET (using flooding), in TYMNET (using centralized path setups), in TELENET (using call-by-call routing) and in SNA (using explicit path routing from table lookup). The analysis here has been poor to fair in the sense that dynamic routing procedures are difficult to analyze, but static analysis has been possible and provides a good approximation. The issues of looping, stability, transient response, etc. have not really yielded very well to analysis, although progress is being made continually on this front [3]. Static analysis using methods such as the flow deviation method in a centralized fashion [2], or attempts to implement flow deviation in a distributed environment [4] have been fairly effective. We still do find problems with routing procedures as, for example, in the recent article [17] which described yet another system crash due to a malfunction in procedures related to the routing function.

Flow control is perhaps one of the least understood control procedures in terrestrial networks (and in fact in all networks). Here again, a number of implementations may be found in the industry using various throttles such as stop-start procedures, buffer classes, window limitations, isarithmic schemes, etc. The analysis here is abysmal, and very little has been done which is of much use. This is an open area for investigation. Some high level models, however, have yielded to analysis including the work on power [14] and on certain other schemes [7]. The difficulty is that the dynamic nature of flow control is extremely difficult to analyze. In fact, it seems fair to say that flow control has the following cantankerous properties: you need it, it's tough to design, it's nearly impossible to analyze, and it's almost sure to cause you trouble (i.e., create deadlocks, degradations, and other lovely castastrophes). There is great opportunity here for further creative research!

As far as topological design of networks is concerned, here again some satisfactory heuristic design procedures have been in use for many years. The exact solution to the topological design problem is likely never to be found due to its inherent computational and probabilistic complexity, but this is no real impediment at this point since our current design procedures really seem to be quite efficient. However, a breakthrough here would be most welcome.

C. Satellite Networks. Packet switching over a satellite communication channel represents a perfect example of a multiaccess broadcast distributed communication channel. (It is multiaccess in the sense that more than one user may wish to access it at the same time. It is broadcast in the sense that many users hear the transmission in progress. It is distributed in the sense that users are not located in one centralized control facility.) In this case a number of earth stations distributed over a large geographical region attempt to share the common capacity of the satellite channel; the satellite simply transponds back to earth all that it receives in a broad beam covering a large portion of the earth. The question here is how one should access the channel in an efficient way. A number of such access schemes have been suggested [11], and experiments have been conducted as well [6, 12]. The principle feature of these channels is the enormous propagation delay (on the order of a quarter of a second roundtrip for geosynchronous satellites), and this strongly affects the class of access algorithms which make sense. Experiments have shown that packet broadcasting on a satellite channel can be effective not only for data but also for voice communications and for teleconferencing. Commercial offerings of point-to-point satellite channels for data are already available through the commercial organizations (e.g., SBS). The analysis of the access schemes here has progressed reasonably well. The concerns are ones of throughput, response time and stability. The stability question is quite important in these distributed channels and has been addressed at some length in the literature [5]. There is room for further research in this area, and we have seen a number of papers appearing in the recent literature on this subject.

D. Packet Radio. Here we provide the use of a common radio channel to a collection of communicating devices in a relatively small geographical area (on the order of a few tens of miles). The problem is similar to that of satellites, namely, we have a multiaccess broadcast distributed communication channel which places a burden on the access scheme, but in this case the propagation delay is tiny (tens of microseconds) and we can take advantage of this in our channel access protocols. In addition to the access scheme, however, a new class of extremely difficult problems arises, namely, that of routing, path finding, control and management of packet radio networks. In the case of mobile radio terminals, the problems are enormous and have so far met with little analytic success. Even in the stationary case, the problem of multi-hop communications (whereby one does not communicate with one's final destination in a single transmission but rather uses a relay of other packet radios to cross the distance between the source and the destination) is unmanageable. Only very approximate analyses for the multi-hop packet radio systems have been documented [13] and considerable work must yet be done in this area. The Department of Defense Advanced Research Projects Agency has experimented here as well as in the satellite case. A number of testbeds are already in place and experiments are currently going on. Most of the design difficulty focuses around the particular protocols to be used at the access level, at the routing level, at the flow control level and at the management level. A number of extremely difficult problems remain to be handled in this area. As mentioned earlier, a large number of access protocols have been documented and more are appearing every month. Here is a great opportunity for creative work in protocols, analysis and control.

E. Local Networks. Local networks, which provide communications among devices within a building or within a collection of buildings, is one of the newest of the various telecommunications areas. We have seen a veritable explosion in numbers of implementations and products being offered to satisfy the needs of in-building communications. Two access schemes have emerged as the most serious contenders for the Institute of Electrical and Electronics Engineers (IEEE) standard for local network access. These are: Carrier Sense Multiple Access with Collision Detection

(CSMA/CD) and the Token Ring. Analysis and simulation of these schemes are available in the literature [1], and the behavior is relatively well understood. New schemes appear each month in the literature of the field, and one wonders if the standard is not being forced upon us too early in this rapidly evolving technology. However, the access scheme is not the only issue of concern. There are a number of other critical issues which must be addressed. In Table 1, a list is given of considerations one must make in evaluating any local network. The list is divided into two parts, the left half of which describes those technical features which are typically amenable to analysis and to quantitative comparison. These are the measures we often see discussed in the technical literature. The right half of the list represents other (softer) issues and measures which one seldom sees addressed in the technical literature but which are of major concern in the commercial trade literature. It is this set of other issues which tends to determine which architecture is adopted by a given corporation or institution. Most of the readers of this journal are probably far more interested in the left half of the list rather than the right half; this is unfortunate. It is unfortunate because we find ourselves at a kind of dichotomy with the industry; they are concerned with the soft issues, and we keep feeding them results about the technical issues. At some point we will have to cross each other's boundaries and address each other's issues in a serious way if we are to make a significant impact.

Table 1: Local Network Issues

Technical	Other
Speed/Capacity	Cost
Topology	Reliability
Response Time	Installation
Efficiency	Expansion/Contraction
Channel Access	Standards
Length/Spacing	Protocol Interface
Number of Stations	Taps
Addressing	Voice/Data/Fax/Video
Acknowledgements	Security/Privacy
Medium for Transmission	Noise Immunity
	Safety
	Qualified Maintenance
	Transparency
	User Migration
	Vendor Lock-in
	Back-up Facilities
	Can Stations work if net is down
	Availability of Electronic Mail, File Servers, and Print Servers
	Error Checking
	Simple Access to Files
	Regulation
	Charging Policy
	Operations Staff
	Non-technical User Environment Needs

One of the interesting technical issues to discuss is the choice of medium for communication itself. Originally local network access was handled by twisted copper pairs through the use of telephone-supplied equipment, often using private automatic branch exchanges (PABX). More recently we have seen the PABX develop into a computerized branch exchange (CBX) still using twisted pairs. However the use of coaxial cable has recently emerged as a very attractive alternative to twisted pairs [15]. After all, the single half-inch diameter coaxial cable can replace more than 1500 twisted pairs. Baseband coaxial cables can provide data rates up to a few tens of megabits/second (MBPS) in half-duplex mode. Broadband can provide over 100 MBPS in full-duplex mode and it can be implemented using commercially available CATV components. Beyond coaxial cables we are looking at the remarkable technology of fiberoptic channels. These channels are made up of bundles of glass fibers, each a few thousandths of an inch thick and 100 times lighter than copper wire or coaxial cable. They can run four miles without a repeater, and have bandwidths measured in thousands of MBPS. Here we have a technology in which semiconductor lasers convert electronic signals into pulses of infrared light which are sent down the glass fiber and then converted back to electronic signals by photodetectors. Once again silicon (in the form of silica to make very pure glass) has come to the rescue! The fiberoptic channel, in fact, represents an enormous opportunity in providing huge bandwidths in an electonically secure environment (secure from simple tapping, from electromagnetic interference and from crosstalk).

A large number of firms are currently implementing various versions of the ETHERNET protocol (CSMA/CD), and these are already available in the market. For less than $1000 per interface one can connect a terminal to a local network and provide wideband communications among the devices within a given building.

The major question here seems to be one not so much of access protocols for data (these are reasonably well understood) but rather the need for integrated services which permit voice, data, facsimile, video and graphics all to be served by the same communication system. The needs of this integrated service have introduced some serious problems of access and control. Broadband coax using CATV components is being considered as one solution to this problem. The CBX also appears attractive for integrated services. Much research needs to be done in this area.

In all of the systems discussed above (ranging from terrestrial networks to satellite nets to packet radio nets and to local network architectures) we have seen that there is an opportunity whereby we can take advantage of our resource sharing principles in dynamically providing wideband communications among a large number of demands. A number of technical issues remain to be resolved as mentioned above (principally in the areas of routing, flow control, and access control), but some of the major problems lie beyond the analytic area and enter into the softer areas described earlier. Elevating our sights to much higher-level issues, we mention simply the enormously complex issues of tariffs, standards, politics, sociology, trade barriers, etc. Issues such as those will determine the direction in which telecommunications makes its real progress. Political considerations could easily strangle the effective use of telecommunications on a worldwide basis, and we must be wary of such tendencies. Of course, these are issues with which the technician is seldom concerned but issues of which he must be aware even if he contributes little to their solution.

3. The Real World and the Coming Revolution

As mentioned earlier, the telecommunication engineers and analysts must not close their eyes to the nontechnical telecommunications issues. One must be concerned with markets and market penetration. For example, which products should be developed at this point in time, and can they be appropriately marketed? What are the needs of the commercial community, and how can they best be met with our wonderful collection of telecommunication services? Indeed, what should those services be, and how should they be charged for? Should they be regulated, or should they be open to free competition? Who should be concerned with security and safety of the data? What agency or organization should be concerned with tying this collection of heterogeneous networks which span the world into a global network which functions in a cooperative fashion?

Answers to these questions will be forthcoming in one fashion or another. The demand is clearly here. Home services will be the first place where we will feel the impact. Television sets will become two-way terminals to the world's databases and information resources. Children will likely receive significant fractions of their education at a computer terminal in a highly interactive mode. Electronic funds transfer is a virtual certainty as is electronic mail. These are the services that one is likely to see in the home and also moving out into the business sector. At the moment the impact on business is being felt more strongly and has come in the form of software packages for finance, order entry, word processing, etc. These are real business needs which have been well served by the rapid expansion of data processing capabilites. Future systems are likely to include speech input/output word-processing machines (which implies speech recognition, understanding, editing, ect.), nearly paperless offices, large screen executive work stations, fourth generation languages for rapid application development, etc. Meanwhile, the software is abysmal and usually poorly supported and rapidly changing and poorly documented and slow! (The recent growth of the nonprocedural languages appears to be one way out of the software stranglehold; these languages include FOCUS, RAMIS, MAPPER, etc.) Such is the price of passing through a rapid transition into a new era. But that era is upon us, and few people are really aware of its enormity. As technicians we have a clear responsibility to provide the technical support for the telecommunications revolution (or. if you will. information revolution). but we have a second responsibility to keep aware of the real needs and demands as they exist out there in the commercial world and to fold that understanding into our own designs so as to provide a service which is both technically and functionally responsive to those needs.

References

1. Bux. Werner. "Local-Area Subnetworks: A Performance Comparison." *IEEE Transactions on Communications*. Vol. COM-29. No. 10. October 1981. pp. 1465-1473.

2. Fratta. L.. M. Gerla and L. Kleinrock. "The Flow Deviation Method: An Approach to Store-and-Forward Communication Network Design." *Networks*. Vol. 3. No. 2. 1973. pp. 97-133.

3. Gafni. E. M. and D. P. Bertsekas. "Distributed Algorithms for Generating Loop-Free Routes in Networks with Frequently Changing Topology." *IEEE Transactions on Communications*. Vol. COM-29. No. 1. January 1981. pp. 11-17.

4. Gallager. R. G.. "A Minimum Delay Routing Algorithm Using Distributed Computation." *IEEE Transactions on Communications*. Vol. COM-25. No. 1. January 1977. pp. 73-84.

5. Gerla. M. and L. Kleinrock. "Closed Loop Stability Controls for S-ALOHA Satellite Communications." *Proceedings of the Fifth Data Communications Symposium*. Snowbird. Utah. September 1977. pp. 2-10 to 2-19.

6. Gerla. M.. L. Nelson and L. Kleinrock. "Packet Satellite Multiple Access: Models and Measurements." *Conference Record. National Telecommunications Conference*. Los Angeles. California. December 1977. pp. 12.2-1 to 12.2-8.

7. Gerla. M. and L. Kleinrock. "Flow Control: A Comparative Survey." *IEEE Transactions on Communications*. Vol. COM-28. April 1980. pp. 553-574.

8. Kleinrock. L.. "Resource Allocation in Computer Systems and Computer Communications Networks." *Information Processing* 74. Proceedings of IFIP Congress 74. Stockholm. Sweden. August 1974. North Holland Publishing Company. Amsterdam. 1974. pp. 11-18.

9. Kleinrock. L.. *Queueing Systems. Volume I: Theory*. Wiley Interscience. New York. 1975.

10. Kleinrock. L.. *Queueing Systems. Volume II: Computer Applications*. Wiley Interscience. New York. 1976.

11. Kleinrock. L.. "Performance of Distributed Multi-Access Computer-Communication Systems. *Information Processing* 77. Proceedings of IFIP Congress 77. Toronto. Canada. August 1977. North Holland. Amsterdam. 1977. pp. 547-552.

12. Kleinrock. L. and M. Gerla. "On the Measured Performance of Packet Satellite Access Schemes." *Proceedings of the Fourth International Conference on Computer Communication*. Kyoto. Japan. September 1978. pp. 535-542.

13. Kleinrock. L. and J. Silvester. "Optimum Transmission Radii for Packet Radio Networks or Why Six is a Magic Number." *Conference Record. National Telecommunications Conference*. Birmingham. Alabama. December 1978. pp. 4.3.2 to 4.3.5.

14. Kleinrock. L.. "Power and Deterministic Rules of Thumb for Probabilistic Problems in Computer Communications." *Conference Record. International Conference on Communications*. Boston. Massachusetts. June 1979. pp. 43.1.1 to 43.1.10.

15. Levy. W. E. and M. Rothberg. "Coaxial Cable Finds a Home." *Mini-Micro Systems*. Vol. 14. No. 3. March 1981. pp. 99-108.

16. Martin. J. *The Wired Society*. Prentice-Hall. Inc.. Englewood Cliffs. New Jersey. 1978.

17. Rosen. E. "Vulnerabilities of Network Control Protocols: An Example." *Computer Communication Review*. Vol. 11. No. 3. July 1981. pp. 10-16.

About the Author

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

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 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 120 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.

He is a Guggenheim Fellow. an IEEE Fellow. and has received various outstanding teacher and best paper awards. including the 1976 Lanchester prize for the outstanding paper in operations research. and the ICC '78 prize-winning paper award. He was recently elected to the National Academy of Engineering in recognition of his pioneering research contributions and educational leadership in the field of computer communications networks.