This essay looks at some future effects on computing and telecommunications of some countervailing technology trends. The projected time is the mid-1980s, and the technology trends considered are those of Large Scale Integration (LSI), related storage technologies such as charge coupled devices (CCDs) and magnetic bubbles, optical fiber transmission systems, and satellites. The principal focus is upon the implications of these trends for distributed processing and computer networks.

Potential technology implications for computers and telecommunications in the 1980s

by W. D. Frazer

In this article, we briefly discuss recent developments in several technological areas, both new and old. We then attempt to project—some might say speculate about—the future implications for communications and computing, with particular emphasis on implications for computing in the early and mid-1980s.

Figure 1 Computer logic costs

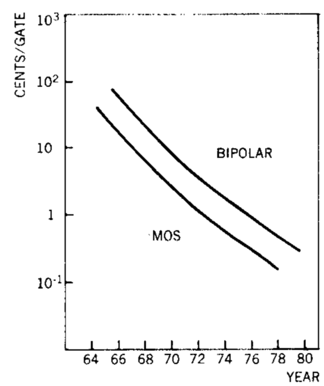

Large Scale Integration and storage technologies

One of the most significant technological developments in recent memory has been the dramatic and accelerating decrease in the cost of electronic circuitry brought about by Large Scale Integration (LSI). The magnitude of past change and a projected future are clearly indicated[1] in Figures 1 and 2 and are discussed as well by Branscomb elsewhere in this issue.[2] We shall focus first on the implications for communications of the new economic ground rules for logic and memory circuits along with related storage technologies such as charge coupled devices (CCDs) and magnetic bubbles.

Figure 2 Cost per bit of memory (Adapted from "Microelectronics" by Robert N. Noyce. Copyright © 1977 by Scientific American, Inc. All rights reserved.)

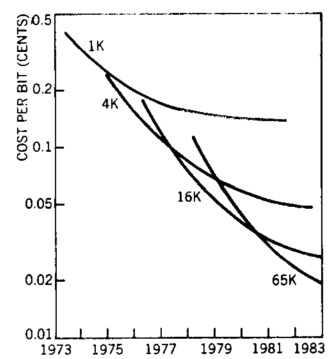

The major consequence for our purposes of the "LSI revolution" in circuit costs depicted in Figures 1 and 2 is a visible trend toward the decentralization of many aspects of data processing function. There are many reasons for this, and it is instructive to examine a few of the most significant ones.

decentralization

The principal motivation to decentralize is unquestionably an economic one, and the economic incentives have several sources. First, as Branscomb has observed,[2] the cost of electronic logic and memory circuitry has decreased in recent years at a rate several times the rate of decrease in communication costs. This argues for the replacement of communications bandwidth with remote processing whenever possible. Once one begins to think in these terms, other favorable factors also enter the picture: The performance requirements imposed upon a small remote data processing system are generally much less stringent than those imposed on a large central host. This is because the diversity of demand for function is usually much greater in a larger central host. Similarly, the system software can generally be simpler and less highly optimized in performance. Less expensive technologies can sometimes be used to advantage in such systems, contributing further to the overall cost reduction.

In large systems, a huge pool of memory can be time shared by many programs through multiprogramming and through the use of virtual storage, orchestrated by a sophisticated and complex operating system. On the other hand, smaller systems, more than larger ones, tend to be dominated both in cost and performance limitations by the amount of processor memory. This is why progress in reducing the cost of memory is so crucial and has such a major effect on decentralized data processing. Until quite recently, it was common for small systems to be limited to memory sizes of a few tens or hundreds of kilobytes. That barrier is now being broken,[3] and as a result one can expect to see a qualitative change in years to come. The recently announced IBM 4331 Processor and System/38 exemplify this trend. The effects of memory size, particularly in small systems with a limited degree of multiprogramming and limited addressing, are highly nonlinear. The time required to invert a matrix or to sort a file of records, for example, is reduced in large increments as critical boundaries of memory size—those necessary to hold and access critical subsets of programs or data—are reached and passed. It is safe to predict that each new advance in LSI memory density technology will quickly be reflected in a capability to decentralize an additional set of applications as well as to decentralize larger versions of applications which are currently running remotely. We shall return later to limits to this trend.

interactive computing

Another aspect of the economic picture frequently stems from the interactive nature of many of the applications users wish to execute. As geographic distance and the volume of data transmitted between locations increase, so does response time. Even at 1979 costs, the tradeoff between response time—i.e., the cost of idle labor—and hardware cost frequently favors the latter. This trend can only accelerate as pressure to increase office and clerical productivity intensifies, for the two costs are moving in opposite directions. Two examples of such applications are data entry and text editing. In both cases, the terminal operator is frequently limited in speed by the application response time. On the other hand, many—if not most—of the terminal interactions have minimal processing requirements. They consist of verifying the correct entry of data, for example, or reworking a single word or line of text. Thus, moving these simpler interactions out to a remote processor, or even into the terminal itself, offers both the potential for productivity gain and increased user satisfaction. The difference between a three-second and a two-second response time for data entry provides potential for up to a thirty-three percent productivity gain. Communication is reserved, in this scenario, for those more complex cases in which centralized processing and/or remote data access or storage are required.

microprocessors

The trend to decentralization does not stop at the application program level. Just as applications have begun to move from centralized sites to the sites of users, so also have communication-related tasks begun to move from processors of whatever size into tiny microprocessor chips. Among the most notable of these in communications are the highly repetitive functions required to manage communication lines at the bit level, including data assembly, error checking and correction, link-level protocol processing, etc. Chips which provide line handling at the level of the Synchronous Data Link Control (SDLC) line protocols of IBM's SNA are commercially available today. The use of such microprocessor chips makes possible a much more modular and cost-effective approach to handling the large numbers of lines necessary for many of the networks of the early 1980s. In addition, by offloading low-level functions from the host to the microprocessor, this type of design makes available more expensive processor cycles, which can then be applied to other tasks.

reliability of remote units

A second major factor contributing to the trend toward decentralization is the growing reliability of remote units. Reliability manifests itself in a combination of product cost and service cost, and contributes in this way to the economics of decentralization. A regular program of frequent preventative maintenance and inspection is normally quite feasible when amortized over all the equipment in a large installation. Such a program may quickly become a dominant cost, depending upon frequency, when applied to a collection of geographically distributed sites with implied travel time costs. The semiconductor-based components of data processing systems have long had intrinsically high reliability—at least in controlled environments. Connector reliability between boxes is significantly lower, however, as is the reliability of electromechanical devices subject to the high accelerations which characterize high speed impact printing, disk actuators, etc. Con-

Figure 3 Storage technology costs (Reprinted from *Electronics*, October 26, 1978. Copyright © McGraw-Hill, Inc. 1978.)

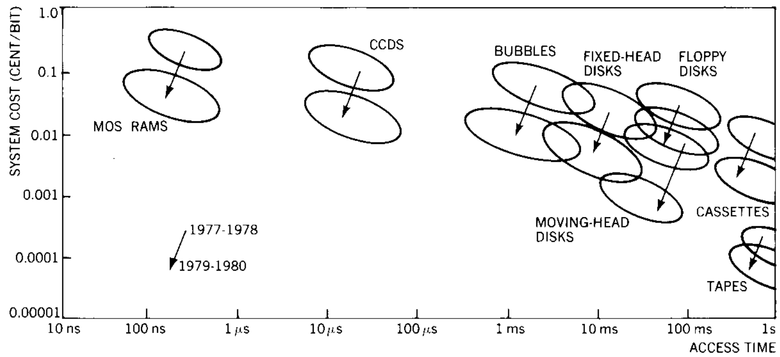

nector reliability has increased dramatically as LSI has replaced box-to-box connections with chip-to-chip connections within a single box and on a single board.

new storage technologies

Fortunately, we have also seen advances both in nonimpact printing technologies and in nonmechanical alternatives to conventional disk storage technologies in the form of, e.g., magnetic bubble storage devices and electronic storage technologies such as charge coupled devices, as seen in Figure 3. Although not presently cost competitive with disk or tape for large volumes of data, such devices may well become competitive in smaller sizes.[4] They also have the native long-term potential for higher reliability in unattended operation. In addition, their intrinsically shorter access time offers the small system designer storage hierarchy tradeoff opportunities not available so far with disk files.

A major reliability benefit to the remote system accrues simply from the availability of cost-effective nonvolatile storage. In the event of software error—or even some kinds of hardware malfunction—one cannot normally determine initially which portions of processor memory may have been "contaminated" by an errant program. A standard approach to this situation, as in the IBM Information Management System (IMS) or the Network Control Program (NCP) today, is to rely on a pristine version of error recovery programs stored on a write-protected nonvolatile storage medium. These programs systematically investigate, test, and reconstruct the data structures necessary to resume normal operation. Although error recovery routines could conceivably be stored in an appropriately protected and segregated portion of processor memory, such a solution would offer no protection against a partial or total loss of power. They would in any event be uneconomical to store in processor memory because of the large number of special cases, given the cost of processor memory and the typically low utilization and large physical volume of such programs. Lacking such a cost-effective nonvolatile storage technology for small systems, one would have to rely on communication with a remote, perhaps central, location to do diagnosis and software recovery. This is what was done frequently in the past in the U.S. Department of Defense network, ARPANET, for example, and in NCP. This approach obviously increases overall communication costs, however, and adds to overall system complexity by complicating local responses to recovery.

The same cost trends could also work to reduce service cost by permitting customer service in some instances. Routine customer service guided by illustrated printed material has long been common for reproduction equipment. It is quite conceivable to imagine diagnostic programs and/or hardware as a part of terminals or small processors which would permit a user to locate and replace malfunctioning subassemblies. Again, cheap nonvolatile storage is the catalyst making this interaction between user and system economically feasible.

personal interaction

The availability of inexpensive on-line storage is crucial also to certain types of applications particularly suited to human interaction. Two of the most important of these are voice and image processing. Both consume substantial amounts of storage space because of the analog nature and intrinsic entropy of the data. A substantial processing capability is generally required as well, but the "naturalness" of the resultant human interface justifies the expenditure. Where feasible, of course, a remote text editing function of the kind previously discussed can economize on transmission costs by transmitting finished documents in character form, rather than as images.

decentralization for flexibility

A third motivation for decentralization is flexibility. The small size and resulting portability of terminals made possible by LSI technology will make it increasingly feasible to bring data processing into homes, vehicles, and other locations which would have been unthinkable several years ago.

In another area, the LSI revolution has brought about the replacement of electromechanical telephone line switching technology with digital control of electronic switches, as in the IBM 3750 Switching System.[5,6] This has made it possible to incorporate vast amounts of new function and flexibility into such systems by replacing former hardware function with programmed function. Customizeable call-forwarding, conference calling and call intercepting are but three examples of added flexibility. Audio text editing, filing and distribution systems based on telephone input are already in experimental use.[7] Clearly, such trends as this will continue at least well into the 1980s. In addition to added function flexibility, there has also been a startling decrease in the physical size of switching equipment from that of the electromechanical behemoths of the past, yielding improvements in physical flexibility as well.

factors that oppose centralization

It is time, perhaps, to examine the other side of the decentralization ledger: The degree to which decentralization can be carried out is highly application dependent. To the extent that an application requires the sharing of data across geographic locations, it will be hard to decentralize. Care must be taken, for example, to insure that the state of a data base remains consistent by synchronizing changes made to it. Although considerable research effort in computer science is being invested in improving locking and synchronization techniques for sharing of distributed resources, one seems to require a significant amount of communication among the participants in the sharing process simply to orchestrate their sharing. At the point where this communication becomes too burdensome, the clear choice is to centralize some or all of the serially reusable resources.

Other significant technological factors also argue for centralization. Two of the most important are storage and operations costs. In the case of the storage of large amounts of data, the cost per bit is and will be affected significantly by economies of scale, as we have seen. On the basis of this projection, it appears that large centralized systems will continue to have a significant large scale storage cost advantage, at least for several years to come. Cheap large scale storage is essential to the feasibility of many sophisticated applications involving facsimile processing, document distribution and image processing. Thus, in many instances these applications will very likely continue to require large centralized processing installations; further, interactive or remote batch applications based on them will increase communication demand, as well as storage demand in remote locations.

The area of operations cost, too, is one where economies of scale frequently come into play. Widespread replication of operations and programming support staffs is costly and inefficient and will also serve as a limiting factor in the trend to decentralize. An unacceptable additional burden of administration is sometimes required simply to keep multiple locations in step with one another, both in operational procedures and in software.

Optical fiber technology

We now explore some of the new technological advances in telecommunications and comment on their implications. Although we have previously postulated that processing costs are decreasing more rapidly than communications costs, the advent of satellite communications and optical fibers among other technologies will contribute to an overall decrease in cost in communications as well.

Optical fiber technology is another area which offers striking new opportunities for communication innovation. Its major communications-relevant advantages are low cost, low loss, low dispersion, small size, and immunity to electrical noise.

The optical fiber transmission medium is a glass fiber acting as a wave guide for a modulated light beam.[8] For high bandwidth applications, the source is typically a laser, but may also be a light-emitting diode or other source. Detection is normally done by means of photo diodes. Because there is no electromagnetic radiation outside the fiber, it must be physically penetrated in order to tap the signal; this provides some degree of security against unauthorized taps. The low cost, generally projected to be less than coaxial cable when volume production is achieved[9-11] (see Figure 4 and Table 1), arises from recent and major advances in manufacturing and materials technology. It is anticipated[11] that with mid-1980s volumes, cabling costs (sheath, structural reinforcement, etc.) will be a much greater component than the cost of the fibers in the cable. The small size[12] is an attribute of substantial importance when dealing with the confined cabling space within a building or plant site, for example, or even within a computer room. A major new degree of flexibility in establishment-wide communication linkage results from low cost combined with very low weight/bandwidth and volume/bandwidth ratios of optical fibers.[13-14] One can conceive the practicality of access to a glass fiber near every electrical outlet in new construction to provide flexibility in moving terminals. In existing construction, it may become cost-effective to replace copper wires *in situ* with much broader bandwidth links while still leaving conduit space to spare.[15]

optical fiber characteristics

The most striking characteristic of optical fibers is bandwidth. Their bandwidth arises principally from the optical frequencies of transmitter and receiver, but is also made possible by their low dispersion and low loss. The low dispersion permits full use of the bandwidth capability by providing an intrinsically low signal distortion, even at high modulation rates. The low loss dramatically lengthens the distance between repeaters, thus limiting the number of instances in which they must be used. Since such repeaters must presently transform signals from optical mode to electrical mode and back again, they must clearly be used only when absolutely necessary. Bandwidths of over 400 Mbps on a single fiber have been achieved successfully in experiments at distances of a few kilometers; one might expect to see systems in regular commercial operation at 400 Mbps with distances of ten kilometers between repeaters (or between source and destination) by 1985.

Figure 4 Approximate capital cost per circuit relationship for a range of optical fiber and metallic cable digital systems (installed cable plus repeaters), based on routes of 50 to 100 km with no multiplexing included (From *Telecommunications*, April 1977; reprinted by permission.)

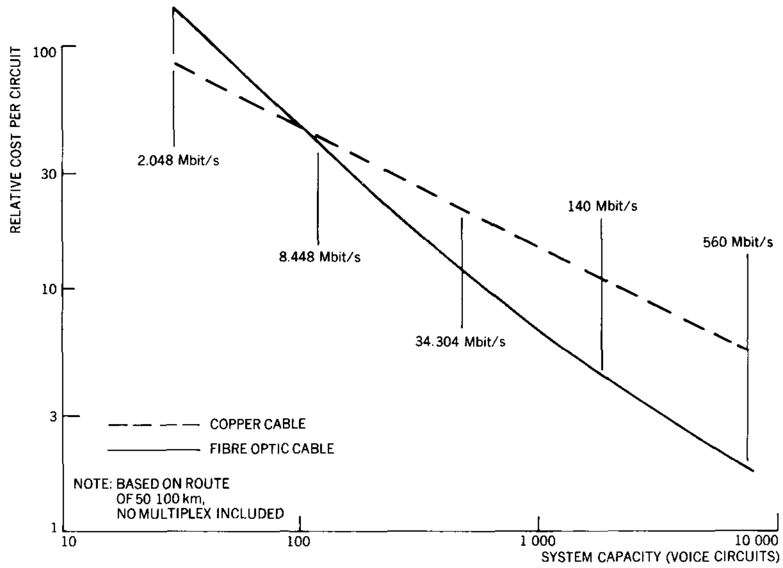

Table 1 Comparison of some existing or proposed metallic cable and possible optical fiber digital systems (From *Telecommunications*, April 1977; reprinted by permission.)

System capacity		*Metallic systems*		*Optical fiber systems*	
Mbits/s	*Voice channels*	*Cable type*	*Repeater spacing* (km)	*Fiber type*	*Repeater spacing* (km)
2 048	30	Voice frequency	2	Step index	12–16
8 448	120	0.6–0.9 mm screened pair	3.4	Graded index	10–14
34 304	480	2.8 mm coaxial cable	2	Graded index	10–12
140	1 920	4.4 mm coaxial cable	2	Mono mode or graded index	8–10
560	7 680	9.5 mm coaxial cable	16.2	Mono mode	
		(4.4 mm coaxial cable	1)		5–6

potential applications

The future implications for communications are significant. A private high bandwidth communications network might easily span a very large plant site or an entire skyscraper complex, providing "local" linkage between any two points. The bandwidth available is sufficient for most current applications, including extensive display support. Historically, there has been a major technical distinction in software support between "local" and "remote" attachment between data processing components based on communication path error rate, bandwidth, and—to a lesser extent—communication path delay. Such distinctions blur significantly with fiber optics. "Local" can now mean "within a few kilometers."

Since a single fiber can suffice to meet the needs in most foreseeable applications, the signal skew problems inherent in multi-line connections are avoided, and extraordinary distances can be spanned by a continuous link, as noted. On the other hand, the convenience and flexibility in signal switching which we now enjoy with electrical signal communications technology is much more difficult to attain with optical fibers. While fiber optic switching technology is still in its infancy, the problem of designing an all-optical cross point switch seems formidable. Electro-optical switches have been built, but they suffer from high losses. Electromechanical switches are the most promising current technology. Even multidropping, the tapping of a single optical signal at several junctions on a fiber link, is costly in signal strength with present technology.

A related area in which recent progress has been made is that of power transmission and control. Power transmission has been pursued in telephony for such purposes as bell ringing,[16] but it could be important in other applications as well. One example of such an application is remote power control for devices and/or instruments, which now appears to be feasible; whether power itself can ever be supplied in significant amounts over an optical link is much more speculative. Control of power via signals on a fiber optics link is currently being pursued actively by the Japanese.[10]

A final word is in order regarding freedom from electrical interference of signals on an optical fiber link. Interference becomes significant whenever exposure to high voltage electrical sources or ground loops becomes a problem—on an open air plant site, for example.[17,18] It has recently been estimated[19] that a fiber optics cable with an appropriate interface is cheaper even today for short runs in such environments than a linkage of point-to-point copper wires protected by lightning arrestors.

To summarize, optical fibers appear to offer significant potential as a communications technology, a potential which we have only begun to understand and exploit. At this point, it appears that their only serious shortcoming stems from the fact that optical signals and transmission modes differ significantly from electrical ones. Transformations are thus required where the two modes must meet, with resulting inefficiencies of one kind or another.

Satellites

Significant recent progress in satellite transmission technology has for the first time made feasible the widespread local use of the vast communications potential of satellites. The advances in electronic circuitry and storage technologies we discussed earlier have contributed in a major way to this progress. So, also, have improved antenna designs resulting from experience in the years since the Apollo space program.

Figure 5 Intelsat shows investment cost improvement of thirty percent per circuit per year

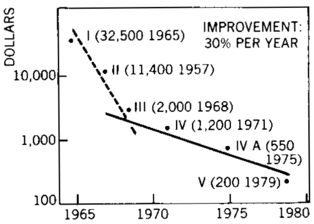

Historically, the major barrier to use of satellite-based communications on a regular and frequent basis for data processing was cost. This cost derived in part from the capital investment represented by an earth station, but also in part from the cost of the terrestrial lines necessary for transmission of data between the earth station and the end user or originator—be it terminal or data processing host. The cost of satellite fabrication and launch, for example, is frequently small in comparison with these costs. There has, however, been a consistent downward trend in investment cost per circuit per year,[20] as shown in Figure 5.

It is now economically feasible to consider a local rooftop or plant site antenna linked, perhaps, by an optical fiber network of the kind discussed in the last section to data processing components within a line radius of a few kilometers and communicating with a satellite in geostationary orbit. The details of satellite technology are spelled out elsewhere;[21-25] we focus here on some data processing implications.

network considerations

To begin with, it is very likely that voice and data will be combined in such a network because of the significant bandwidth available and the need to multiplex signals in order to amortize the cost of the earth stations. Data traffic demands will initially be small by comparison, but the result will likely be broadband digital communications at very low cost indeed by present standards. Note that the satellite connection network can be essentially fully connected (i.e., each point connected to all others) because of the intrinsic broadcast potential of satellites; the satellite antenna can be designed to be uni- or omni-directional as needs dictate. Major rethinking of enterprise communications and organization procedures may be needed in order to capitalize fully on broadcast capability: For example, to what extent can simultaneous on-line availability of text in widespread geographic locations supplant face-to-face discussions? Experience among programmers, both inside IBM and among university users of networks, has indicated that such a facility is extraordinarily useful; whether the same will be true for nonprogrammers remains to be determined.

Another major area to be explored is that of data processing load sharing among geographically remote areas. Such a strategy was unthinkable when communication bandwidths were measured in kilobytes or tens of kilobytes per second, but load sharing becomes a somewhat more realistic possibility at megabyte-per-second rates across multiple time zones. Some interactive applications will still be unsuitable to consider because of the delay inherent in satellite transmission, a subject to which we shall return below. Nevertheless, batch or "complex transaction" applications appear to be suitable candidates to consider for transmission to remote and underutilized processors.

One of the most interesting and relatively unexplored areas of technology is that of incorporating satellite links into computer networks on a widespread basis. The same satellite links which can feed directly onto plant sites have characteristics other than huge bandwidths which distinguish them from terrestrial links. One of these is intrinsic error rate. The small antenna systems which make local reception possible have error rates which will require more redundancy in transmitted data to achieve an acceptable overall error rate, particularly at very high frequencies. Fortunately, this function can probably be provided at reasonable cost through use of "line handling microprocessors" as discussed earlier. Some additional bandwidth must be consumed at higher (12–14 GHz) frequencies in order to overcome atmospheric attenuation at acceptable error rates. A second characteristic is the broadcast capability alluded to above. Some ingenuity will be required to use this capability to the fullest in data processing. One can envision the use of broadcast to initiate searches of data bases or directories in multiple locations, but the technology necessary to support such applications—particularly with respect to a synchronization of changes and recovery—is not yet in hand.

transmission time delay

A more significant characteristic, however, is transmission delay. An end-to-end transmission across a geostationary satellite link takes roughly one-quarter second. A round trip thus consumes a minimum of one-half second, yielding a lower bound of about one-half second response time on satellite-based transactions. This delay, a rather long one by data processing standards for a high speed data link—as distinct from a channel-connected terminal link—also imposes special stress on resource allocation strategies in computer networks: Let us suppose that a collection of, say, two million bytes of data is to be sent from A to B. If the satellite link supports a data rate of two million bytes per second, even allowing overgenerously for fifty percent redundancy will still yield a net rate of one million bytes per second which must be buffered and accepted by B. Now, if B allocates his buffers at the instant he agrees to accept the transmission, they will sit idle for a minimum of one-half second. To complicate matters further, A may not send the data immediately if other demands on A's time intervene. On the other hand, if B tries to optimize use of memory too closely—by trying, for example, to use the buffers for other purposes until the data begins to arrive from A—then B may lose data, and a rather costly retransmission will be required. This example should give some idea of the complexities of the tradeoffs involved.

In summary, satellite-based facilities offer a qualitatively new communications capability which calls for fresh approaches in both data processing system organization and user enterprise procedures to fully exploit.

Some further implications for networks

We have argued here that a growing trend toward decentralization is a natural consequence of the recent dramatic history of LSI, storage technology, fiber optics and satellites. It is reasonable to inquire at this point, "What consequences follow for computer-based networks?" One of the most obvious is a significant growth in the number of distinct end points that a network must accommodate and manage. This will require a new variety of software technology. Directory structures which suffice for networks of a few dozen end points will have to be redesigned to accommodate to an order of magnitude or so in growth.

Current computer networks worldwide have a fairly simple logical connectivity structure in contrast to that which will be required in the future. Telephone networks, by contrast, are richer in connectivity, but lack the responsiveness and resource management sophistication necessary to meet computer network needs: A typical telephone call consumes a few seconds in establishing connectivity and consumes a few kilobits of bandwidth. Recovery is minimal—mainly by user retry. On the other hand, connections to a processor capable of executing, say, ten to twenty million instructions per second and transmitting data at the rate of several megabytes per second must be made and unmade at a much faster rate. New public data networks may be expected to fulfill this need.[26] To minimize disruption at the applications processor, the computer network using these facilities should assume a significant share of responsibility for successful delivery of the data as SNA does today. Ideally, it should also provide substantial flexibility in choosing the level of recovery support desired—from very little to very much. In this way, the application can tailor the service to its needs. Financial transactions will require extensive recovery support, while data-logging may require almost none.

Another subject of intensive investigation is resource management. When a message is accepted by a computer network, such as SNA provides, an implicit commitment is made to supply link capacity to transmit it to its destination and buffer capacity to hold it all along its route. For very large computer networks to be economical, this commitment must furthermore be made dynamically, without first securing and holding these resources. If overcommitment occurs, network congestion, or even resource deadlock, may result; undercommitment is uneconomical. What is required is a distributed control mechanism which can adapt gracefully to changing demand, while achieving a cost-effective average resource utilization. There are several promising proposals in these areas but no clear winners as yet.[25,27-29]

Recovery is another area which must be handled in large computer networks in a qualitatively different way. Here again, it is infeasible to centralize control over the recovery process. Parts of the network will be "healthy" while other parts are "sick," and cooperative procedures and protocols must be found to heal the latter with minimal disruption to the former.

Finally, we need to capitalize more fully on the wealth of data processing capability to be made available to us by new technologies in order to provide simpler and more effective mechanisms for the management and control of such large and complex networks by owners and users.

Summary and conclusions

We have discussed some technological trends and some new technologies and have attempted to forecast some of the implications of each to communications and computing. Some general conclusions which cut across technologies are worth reemphasizing in a more global context:

We are on the verge of an era in which we will have significantly more bandwidth and interconnection flexibility than we have ever had. It is likely than there will be more, initially at least, than we now know how to use to advantage. Similarly, we will have a greatly enhanced capability to distribute processing among a large number of data processing hosts. At present, however, we have no good criteria by which to determine when distributed processing makes good long-term economic sense; frequently it is difficult even to see if distributed processing is currently feasible in a given application situation because of the complexities introduced by sharing data bases and application logic, for example.

Economic pressures and technology trends argue forcefully for trading labor and transport costs for computing and communication costs even today. This will result in more and larger networks which must be supported and managed. The demand for such networks is growing now and will very likely accelerate, fueled by economic factors.

Problems will continue to occur in justifying replacement and displacement of installed technologies in all areas. Justifications can range from economic—as when fiber optics displaces twisted pair in telephone cable ducts—to political—as when the Federal Communications Commission required two-way communications capability on Cable Television links, opening the potential for widespread high bandwidth residential links.

Finally, hybrid blends of computer and communications technology, as in digital telephone transmission and switching, are already feasible and will continue to grow and prosper in appropriate applications environments.

ACKNOWLEDGMENTS

We gratefully acknowledge the comments and criticisms of S. Calo, M. Reiser, M. Sachs, and G. Stierhoff.

CITED REFERENCES AND NOTE

1. W. J. Kleinfelder, *The Impact of Technology on Computing; The Future is Now*, Keynote Address, ACM Region 16 Meeting, Atlanta, GA (April 1978).
2. L. M. Branscomb, "Computing and communications—A perspective of the evolving environment," *IBM Systems Journal* **18,** No. 2, 189-201 (1979, this issue).
3. D. P. Bhandakar, "Dynamic MOS memories: Serial or random access?" *IEEE Computer Society Annual Conference, 16th, San Francisco*, pp. 162–164. See also *Computerworld* **12** (September 25, 1978) p. 6.
4. D. Toombs, "CCD and bubble memories: System implications," *IEEE Spectrum* **15,** No. 5, 36-39 (May 1978).
5. P. J. Schuyten, "Digital dawn for telephony," *The New York Times*, D5 (December 13, 1978).
6. The IBM 3750 Switching System is marketed in Europe.
7. D. Zeheb and S. J. Boies, "The speech filing migration system," *Proceedings of the Fourth International Conference on Computer Comunications,* Kyoto, Japan, September 26-29, 1978, pp. 571-574.
8. R. D. Maurer, "Glass fibers for optical communications," *Proceedings of the IEEE* **61,** No. 4, 452-462 (1973).
9. K. C. Kao and M. E. Collier, "Fiber-optic systems in future telecommunications networks," *Telecommunications* **11,** No. 4, 25-32 (April 1977).
10. *Fiber Optics Handbook and Market Guide* (First Edition), Information Gatekeepers, Inc., Brookline, MA (1978).
11. T. G. Giallorenzi, "Optical communications research and technology: Fiber optics," *Proceedings of the IEEE* **66,** No. 7, 744-780 (July 1978).
12. R. A. Cerny and M. C. Hudson, "Fiber optics for digital systems," *Digital Design* **7,** No. 9, 39-50 (August 1977).
13. E. N. Randall and P. J. Dobson, "Designing fiberoptic cable," *Laser Focus* **14,** No. 4, 56-59 (April 1978).
14. N. Snyderman, "Optical cables: Down to business," *Electronic News* **24,** No. 1210, 54-57 (November 13, 1978).
15. T. C. Cannon, D. L. Pope, and D. D. Sell, "Installation and performance of the Chicago lightwave transmission system," *IEEE Transactions on Communications* **COM-26,** No. 7, 1056-1060 (July 1978).
16. B. C. DeLoach, Jr., R. C. Miller, and S. Kaufman, "Sound alerter powered over optical filter," *Bell System Technical Journal* **57,** No. 9, 3309-3316 (November 1978).
17. "BART using optical-fiber link," *Communications News* **15,** No. 5, 129 (May 1978).

18. J. J. Pan and D. E. Halley, "Fiber-optic links for microwave and satellite communication terminals," *National Telecommunications Conference, Dallas, 1976,* **3,** 37.3-1-37.3-3 (1976).
19. L. A. Price II, personal communication.
20. B. I. Edelson, "Global satellite communications," *Scientific American* **236,** No. 2, 58-73 (February 1977).
21. C. E. Catoe, "Low cost information distribution; New directions for technology developments," *Astronautics and Aeronautics* **16,** No. 2, 68-75 (February 1978).
22. H. L. Van Trees, E. V. Hoversten, and T. P. McGarty, "Communications satellites: Looking to the 1980s," *IEEE Spectrum* **14,** No. 12, 43-51 (December 1977).
23. R. G. Gould and Y. F. Lum, Editors, *Communications Satellite Systems: An Overview of the Technology*, IEEE Press, New York (1976).
24. J. Martin, *Future Developments in Telecommunications,* Prentice-Hall, Inc., Englewood Cliffs, NJ (1977), Chapter 13.
25. J. J. Spilker, *Digital Communications by Satellite*, Prentice-Hall, Inc., Englewood Cliffs, NJ (1977).
26. J. R. Halsey, L. E. Hardy, and L. F. Powning, "Public data networks: Their evolution, interfaces, and status," *IBM Systems Journal* **18,** No. 2, 223-243 (1979, this issue).
27. S. S. Lam and M. Reiser, "Congestion control of store-and-forward networks by input buffer limits—an analysis," *IEEE Transactions on Electronic Computers* (to be published).
28. E. Raubold and J. Haenle, "A method of deadlock-free resource allocation and flow control in packet networks," *Proceedings of the International Council for Computer Communication*, Toronto, 1976, pp. 483-487.
29. M. Reiser, *A Queuing Analysis of Computer Communication Networks with Window Flow Control*, Research Report RC7218, IBM Thomas J. Watson Research Center, Yorktown Heights, NY 10598.

Reprint Order No. G321-5099.