The Convergence of Computing and Telecommunications Systems

David Farber and Paul Baran

A major change in computer communications is taking place. As the cost of communications and computation declines, the number of new services that can be offered will increase and a large number of new industries can be expected to emerge. A battle between the potential suppliers is now under way.

This article considers: (i) the nature of the new services, (ii) their importance to the economy, and (iii) the arguments for regulatory control as proposed by the major potential competing suppliers. The basis for seeking regulation hinges primarily on the use of a communication component provided by a natural monopoly that in the past has been highly regulated. Rapid changes in the technological base make it difficult, if not impossible, to resolve the differences between the three major competing suppliers [American Telephone & Telegraph Company (AT&T), International Business Machines Corporation (IBM), and the smaller computer companies] within the existing communications regulatory framework. Considering computation and communications as distinctly separable functions does not provide the basis to resolve boundary disputes rapidly when regulatory issues are involved. This article examines the underlying technical changes, reflected in the conceptual changes, which now make the regulators' tasks difficult. It suggests that part of the present turmoil and uncertainty is a result of the growing use of digital processes within the communications industry itself—an inevitable evolution of the new technology. An appendix describes the evolution of data transmission technology.

We do not intend to propose solutions to the very difficult problems that this article raises. We do not know of any completely workable solutions. Our purpose is rather to sensitize the reader to the issues and problems that our country will face now and in the future.

Dr. Farber is a professor of information and computer science and electrical engineering, University of California, Irvine 92717. Mr. Baran is president of Cabledata Associates, 701 Welsh Road, Palo Alto, California.

The Nature of the New Services

This article was written by two authors 500 miles apart, using a computer communication system. Each writes on the same sheet of "paper." Each amplifies, modifies, and clarifies the words of the other. Misspellings are corrected, and the order of phrases is changed. Each author takes his turn tidying up the modified manuscript by pressing a few buttons and giving a few commands. This takes care of the nitty gritty in paper writing, gobbling up the blank spaces and even justifying the right margin. Discussions on content and outlook take place with the use of the same computer communications-based message systems. The method used to produce the article you are now reading is not tomorrow. It is today.

Tomorrow, computer communications systems will be the rule for remote collaboration, like the above example. As computer communications systems become more powerful, more humane, more forgiving, and above all, cheaper, they will become ubiquitous. The obvious next question is, "Why bother to retype this manuscript manually, set type, and print this journal on paper?" Once the words in this article have been entered into the computer system, they are in fact "published" if a copy has been stored in a publicly accessible memory portion of the computer communications system. Think of the paper wasted in conventional periodical publications when only a small portion will ever be read.

The cost component of the computer communications cost-demand curve declines each year. Yet improvements of literally several orders of magnitude are possible on the basis of what is known already. Accelerated expansion of demand is a certainty. With increased use of computer communications, changes in our institutions will follow. The example given for publishing is only suggestive of the range of change. The immediate reaction of those in publishing to such change is to feel threatened. Institutions rarely take kindly to the prospect of obsolescence of their present justification.

Importance to the Economy

Our simple illustration of the potential change in publishing could, and most likely will in time, spread across the entire economy. Even today we are an information-processing society. More gross national product is developed in the United States today in activities in the information sector than in the production of tangible goods.

We have chosen a dramatic application of computer communications solely for illustration. More important are the mundane applications: automated hotel reservations, credit checking, real-time financial transactions, access to insurance and medical records, general information retrieval, and real-time inventory control in business. A wide range of services such as these and beyond is coming. Automation of the office is almost a certainty. The only questions are "when," "how," and "by whom?" Electronic transmission of much of the current first-class mail similarly will follow.

Change will be felt widely throughout the economy and, in a fundamental manner, by many institutions. Such change is generally perceived as crisis and rarely accepted with grace. Most people believe that change is something best absorbed by others. The interim period will be a period of inevitable discomfort to those organizations and professions that are affected.

The Argument for Regulation

History of potential suppliers in this information economy. Although computer communications promises to provide huge new markets for potential suppliers, there is uncertainty about who they will be. It is not simply a choice of AT&T, IBM, or a diversity of smaller computer companies. Rather, the question is which supplier will capture the largest, most profitable slice of the pie, and with the minimum degree of the unpleasantness of competition.

Each supplier arrives at its claim of legitimacy to supply the new markets by a different route. The telephone industry supplies communications and is now a regulated monopoly. (The most economical way of providing basic communication channels is on a common shared-resource basis—it is almost unarguably a "natural monopoly.") The IBM claim in this new industry is well-founded. It will provide a significant share of the new market's hardware. It is cash-rich and thus unique in the computer industry in having the financial resources to make the needed capital investments. And it is the winner in the proof-by-adversary test in the free economy. It has been tried and has won in the free marketplace. The third contender, the remainder of the computer industry, feels that IBM has won unfairly and seeks to handicap the giant. (There are few computer companies today that have not brought suit against IBM for alleged violations of anti-competitive statutes.) These smaller companies feel that IBM may exercise monopolistic power in a sector in part deemed not to be a natural monopoly and thus not regulated. This may all seem strange in a free economy. But the regulatory and institutional constraints are so treacherous that the outcome of the debate may hinge more on political adroitness than on market acceptance.

The disintegration of the boundary between suppliers. The basic premise that communications functions and computational functions are separable has been rendered meaningless by a basic and pervasive change in technology. The separate "atoms" of communications and computers form a new "molecule" with properties different from those of their atomic components.

Consider the recent case in which AT&T wished to offer its subscribers the sophisticated model 40 communications terminals made by its wholly owned subsidiary, Teletype Corporation. A series of delays ensued as competitors argued that AT&T was entering the processing business; AT&T said "nonsense, it is just in the communications business." The Common Carrier Bureau of the Federal Communications Commission (FCC) decided against AT&T's position. After a short delay, the commissioners overturned the decision of their own bureau. Indecisiveness and on-again, off-again decisions at this boundary line are common. There is no longer an unambiguous dividing line upon which to reach decisions. The lawyers involved in resolving these arguments are assured of an employment stability lacking in most other sectors of the economy, as the endless suits at the interface between suppliers and the government seem to be careers rather than occasions.

The cause of the disintegration. The regulatory paradigm assumes that communications is communications, and computing is computing. Unlike oil and water, they now mix well—too well to treat separately, once combined. They simply are no longer separable into orthogonal components subject to different legal treatment. Separating communications and computational processes is becoming an almost nonsense activity, like unscrambling an omelet. It is a task better left to philosophers than to lawyers on regulatory commissions.

This has not always been the case. The basic changes have crept upon us gradually. The forces melding these two previously separable domains are several. At a gross level, there is the economics of distributed processing. It may be cheaper in the future to provide independent computing power for those who need it, at their terminal, office, or home. But just having one's own computer is in general not adequate. There is the occasional or possibly frequent need for shared information where, irrespective of cost, access to remote shared data in real-time is mandatory. But, most important, the basic nature of the communications process is that digital bits provide the natural common form for all signal transmission and processing. The digital wave form becomes the universal package with which to send voice, data, facsimile, or anything else. Why this is so, and how it got that way, provides insight into the evolving technology and the increasing difficulty that will be felt by those who try to reach decisions on the premise that communications are separable in computer communications networks. What are the alternatives available for solving our dilemma?

Why aren't microcomputers a perfect substitute for large interconnected computers? There is a school of thought that argues that computer communications is a chimera. There is no need to use the resources of AT&T or IBM. According to this viewpoint, computers are going to be so small and cheap in the future that we won't need to use communications lines to talk with big computers. Is the low-cost mini-, or micro-, or picocomputer the answer? Powerful computers have become cheap and will get cheaper. The explosion of the computer hobby market during the last year and a half has been remarkable. There are now 150 specialty stores selling low-cost computers in hobby kit form. More computer development work is being done by hobbyists today than by computer science researchers funded by public agencies. The hobbyists are not all amateurs. There are some whose vocation is also their avocation. They are intelligent, dedicated persons no longer blocked from entry into computer research by the cost of the equipment. And there are many more now at work (or play?) than the government could possibly support.

The reasons why access to communications capability will always be a necessary adjunct to the small computer are basically unchanged:

1) Whenever it is necessary to make infrequent use of a highly specialized resource, it is cheaper to share (why build a machine to make nails when you can buy the few you need at the corner hardware store?).

2) The output of one data processing task is usually the input for another process.

3) By agglomerating computing power as a result of the interconnection of a number of previously independent systems, applications that are currently not feasible in the data processing area can be handled. These include (i) systems which provide high availability, (ii) systems which can grow modularly as their environments grow (the problem of what to do with a system that supports N users when the $N+1$ user arrives), and (iii) systems which match the physical organization of computing power with the corporate organization, thus relieving organizational strains and stresses.

Distributed processing technology and its fulfillment. All of the above applications can be satisfied by distributed processing, loosely defined as the use of a multiplicity of computational devices for a single processing task. Distributed processing implies that the units that compose the distributed computer can freely intercommunicate. If the distribution is other than local, then remote computer communications facilities are required.

Data-processing begets data-processing. Consider the following example suggesting the nature of distributed processing. A cash register in a grocery, really a counter computer, computes a customer's bill quickly. But it also generates the information needed by the store to order fresh turkeys today for delivery tomorrow. This same information can be used by a hypothetical turkey farmer to determine instantly whether there will be a shortage or a surplus of his product, and whether to hold out for a better price or sell his turkeys now. This information has economic value to his procurement decision. The key feedback signal of a free economy in the past has been the flow of money. More rapid signals can be obtained as an inexpensive by-product of the flow of information from other processes.

Economic examples are not the only use for computer networking and distributed processing. Equally, or more significant, are those that deal with human communications. Just as in the preparation of this article we were able to interchange ideas and manuscripts over the computer communications media interactively, so users of the low-cost computer systems of the future will want to intercommunicate for ideas, pleasure, and business. One has only to look at the meteoric growth of citizen's band radio to realize the new opportunities for communications between people in our society. Microcomputers with a near-zero cost making efficient use of a shared communications resource can open up entirely new forms of human communications.

The digital channel—the evolving choice for communications transmission. Before considering this and other issues in detail, it will be helpful to review some of the characteristics of the digital channel and the reasons why it is becoming desirable even for the communication of basically analog signals. An analog signal is a voltage sent over the usual telephone line representative of (or an analog of) the pressure wave form of the human voice. Why should such a conceptually simple wave form be translated into a series of bits? The reason is that, once converted to digital format, it is possible for the wave form to tolerate a tremendous amount of distortion of its wave shape (the signal) and still be capable of simple reconstruction into a neat distortion-free series of pulses. This allows greater economy in the packing of information in a communications circuit and allows the use of lower quality (and cost) channels. Once in digital format, pulses that have been lost or converted into the wrong symbols are recoverable.

It is relatively easy to carry out many rapid calculations upon the data stream to detect and correct errors. (The price is a small cost for a few extra bits to serve as error detectors and correctors.) The ultimate in such distortion removal is cryptography, where a bit stream is purposely modified by a known-in-advance bit stream (called the key). One "subtracts" one binary wave form from the other and reconstructs the initial wave shape without distortion of the transmitted intelligence. An unauthorized observer sees only a binary stream with a seemingly random chance of either being a "1" or a "0," where "1's" and "0's" are the names of the two states possible in the digital channel. Affordable secure transmission mandates digital transmission. The digital stream simplifies the task of transmission design. Here again, if the signal does not become too distorted, below 50 percent of its initial value, it can always be reconstructed without errors. Digital transmissions, unlike analog transmission, which requires critical adjustments merely to minimize distortion, is thus the ideal technique.

Switching. Communication networks require both transmission and switching. Switching is the process of connecting any network user to any other network. Switching, by its nature, is a digital process. In the 1930's, an operator plugged one circuit into the next of a tandem chain of connections through telephone switching offices. This is a digital process as each switch is either open or closed, a neat binary process. The decision process of laying out switches and relays for telephony was among the earliest practical uses for Boolean or binary logic. The foundation of the logical design process in computer design owes much to the telephone industry. Moreover, the integrated circuit devices which make real-time computer applications feasible are all offsprings of the transistor developed at the Bell Telephone laboratories. The telephone dial itself is a digital device. A series of pulses from the dial acts to set switches.

As computer technology matured to the point where its products were realizable as well as elegant and complex, the quaint electromechanical computers that formed the telephone switching network of the past were replaced by computer circuit logic. This evolution continues as digital transmission is now being combined with digital switching. The two separate processes in time are melding into a single integrated process. Bits going into one end of the network are thus beginning to be treated almost the same way as a bit stream in a computer, albeit with less processing.

What Is the Regulatory Issue?

Who should be the future legitimate supplier of the digital system? As the telephone technology slowly evolves toward an almost all-digital system in the far future, it will or can develop the technology needed to provide many of the new computer communications-based services. For example, AT&T has already developed devices for credit card transactions and sophisticated computer input-output terminals through its Teletype Corporation.

Equally comfortable in the domain of processing and perhaps even transmission is the highly competitive computer industry. During the last decade, the FCC encouraged competition in transmission for other than voice telephone services. New transmission companies and independent "value-added carriers" were authorized. The value-added carriers essentially lease telephone channels from the communications carriers, add processing, and use the leased lines so effectively that they are able to sell the joint product for a profit over and above the cost of the basic regulated channel.

Today, much of the in-fighting involves legal and regulatory actions seeking to block new service offerings and capabilities. It is, in part, an attempt to prevent the opponents from establishing beachheads by precedence. And, in part, it is a way of adding uncertainty to the industry to discourage capital investment. Any new computer communications service offered today by a potential supplier is certain to be opposed by the other major contenders. IBM's plan for a communications satellite network for major business users, which bypasses AT&T Long Lines, faces opposition from both AT&T and computer companies fearful of IBM's assumed long-range intent. AT&T's offering of sophisticated terminals is met with countervailing opposition from the suppliers of computer terminals. The outcome of the competition between the suppliers will be decided, in part, by the staying power, legal resources, and political muscle each allocates to the fray. Little attention is paid to the "public interest." The term defies definition. Is the public interest the interest of the cross-subsidized residential telephone user? Is it the interest of a business which faces a reduced communications bill? Is public interest to be viewed primarily in the short term, irrespective of long-term damage to existing institutions in achieving immediate savings? There is no simple calculus to reach decisions. At present, there is a struggle between those who point out the advantages of competition versus those who hold that better use of limited capital results by virtue of monopoly. It is unclear which view will prevail. Until it becomes clear, resources will be wasted and future economically beneficial services to the nation will be delayed simply by the inability of industries to cooperate.

Summary

Since we do not know of any completely workable solutions to the problems we have considered, we shall settle for a few general observations:

1) The public availability of socially useful computer communications services is and has been held back by legal battles that are now under way between the potential suppliers.

2) No simple resolution of these issues in the near future seems likely in view of the past conceptual separation of computers and communications doctrines.

3) The current policy is to determine whether the nation shall or shall not have certain computer communications services, by the adversary process. In this process, often only the voices of the loudest adversary suppliers are heard.

4) Although there can be no certainty that better alternatives cannot be devised, we believe that such a possibility assumes a higher probability if the key actors come from the technical community sectors more representative of the future consumers.

5) If we are to have the new services that are possible, we need an approach that makes better use of the technologists' dreams and goals rather than have future prospects excessively bound by lawyers paid to preserve the interests of their clients, irrespective of any secondary consequences.

6) We cannot be sanguine about this possibility, as technological statesmanship is too easily corrupted by the same forces that have placed us in this predicament. Furthermore, even if not corrupted, beneficial cooperation can too readily be regarded as simply collusion.

7) Although we do not have any clear answer, we do know that present approaches are not taking us where we want to go very rapidly and that alternative approaches should at least be considered.

Appendix: Advances in Digital Computer Communications Technology

We consider here the technical evolution of digital communications.

Modulation. The first digital transmission circuit was the electrical telegraph, predating the voice telephone by 35 years. Although local telephone lines can be and are used for simple on-off direct-current (d-c) signals, the need to transmit signals beyond the limited distances where a d-c path exists led to the use of tone or alternating-current (a-c) transmission.

The device that converts the binary signal into an a-c signal for transmission is called a "modulator." The device that converts the a-c tone back into a properly shaped digital signal is called a "demodulator." The generally used arrangement, called a "modem" (MOdulator/DEModulator), combines both functions.

The first of the present-day-type modems was developed in the 1950's primarily by researchers at the Lincoln Laboratories of the Massachusetts Institute of Technology for transmitting digital signals between air defense sites over analog telephone lines. The efficacy of these devices has improved steadily since then as a result of improvements from a number of organizations. Today we are reaching perhaps 50 percent of the theoretical transmission rate of the conventional telephone line. Modem costs have been declining steadily as the signal processing within the modem is itself converted to digital processing. This conversion is carried out effectively with the use of large-scale integration semiconductor chip circuits. The telephone plant in existence today provides a widespread access channel with low entry cost for many of today's computer applications.

Bandwidth availability. The figure of merit for a modem is the number of bits per second of digital information that can be transmitted over a conventional telephone channel. Data rates on the order of 300 bits per second (about 30 characters per second) simultaneously in both directions are readily achieved at low cost without electrical connection by means of acoustic coupling over any telephone line. Data rates as high as 9600 bits (one way) (about 960 characters per second) are achievable over most (but not all) voice-grade circuits with relatively expensive state-of-the-art technology.

Transmission carriers lease full-time circuits with greater bandwidths than voice circuits. Millions of bits per second can be transmitted over television bandwidth channels. Naturally, the tariffs for such arrangements are expensive, and the applications for such channels limited.

Multiplexing. Even though broad bandwidths are technically feasible, most applications can be satisfied by low data rates. A local telephone call or circuit rental has historically been inexpensive, whereas long-distance circuits have been expensive. The cost of a short-period terminal connection to a computer on the other side of the country, an expensive business, is equal to a long-distance voice call, albeit few bits may be exchanged.

The answer, of course, is combining many users' signals to share a single expensive long-distance channel. Many telegraph, and later, data signals were combined together on a single line. This process is called multiplexing. One form is called "frequency multiplexing." Here each user is assigned a different frequency, and the individual frequencies are separated at the receiving end by tuned filters. Alternatively, in time division multiplexing each user's bit streams are interleafed sequentially. Fundamental to the notion of multiplexing is the fact that each data user will not be transmitting most of the time. The fastest typist at a keyboard generates only a few tens of bits per second. For the hunt-and-peck typist, the result is only a few bits per second. The telephone channel can readily transmit 4800 bits per second, providing a potential for sharing.

Statistical multiplexing. When there is a cluster of data users who wish to send signals to a distant computer, a single expensive circuit can be used. To pack more users onto the line means taking advantage of the fact that, at any one instant of time, only a portion of the possible users would want to transmit. Moreover, to make the averaging even more effective, advantage can be taken of readily available digital storage to hold data and thus smooth out the peaks. "Time buffering" increases the efficiency of the multiplexing process. So devices were developed that combined signals from many users and elegantly packed them together for transmission over a single line.

Spatial multiplexing or packet switching. Multiplexing is highly effective if all users who want to intercommunicate with one another are in only two locations. If the potential users are spread all over the map, the same statistics can be utilized, although it becomes necessary to identify uniquely where the bits came from and where they are to go, and to add a little housekeeping data to detect damaged packets in transit. To minimize the processing equipment, the bits are formed into consistently organized packages called packets. Such packet switching appears on the ascendancy. Most widespread computer communications networks to be built in the future probably will use packet switching, at least in part.

Packet switching networks are now

being built in the United States, Canada, Europe, and Japan. Interconnection between these networks is planned. Packet switching has several advantages: (i) a very robust structure can be built that readily permits the building of systems whose operation is much more reliable than its elements (communications lines and packet switching centers); (ii) it provides the highest degree of statistical averaging to make most effective use of the basic resource; (iii) the standard format package simplifies full effective interconnection between completely different computer systems and terminals; and (iv) no better alternative is in sight for most user-to-computer and must computer network applications.

Future directions for digital networks. Packet switching, although important, is by no means the end of the development for future computer communications systems. The telephone plant itself is moving on to a mostly digital structure in order to achieve a number of economic advantages. In the process it looks more like a computer communications network but with tremendously greater data-handling capacity. One might argue that the two systems are really the same. However, the minor differences between the two systems are certainly of major significance. An argument can even be raised that the conventional telephone set is nothing more than a "terminal" which generates and accepts commands and which receives and generates data to other such instruments, with the "data" being digitalized voices. A single digital voice channel in today's digital telephone systems carries 64,000 bits per second without a conventional modem. Clearly this is a tremendous capacity. There are future needs that could well use some of these high-capacity trunks of the telephone plant. There is a need in the distributed processing area for large bandwidths. In the case of intercomputer file transfers or access, bandwidths on the order of two to three megabits are appropriate whereas for the interchange of high-fidelity pictorial information 10 to 50 megabits are necessary.

Broad bandwidths are not needed everywhere. The same situation prevails in the local versus long-distance distribution of data as in the local versus toll distribution of voice communication; namely, the cost of local circuits is much less than that of toll circuits. In the case of local distribution of high-speed data, new systems such as the Pierce loop developed at Bell Laboratories, the Irvine ring designed at the University of California, Irvine, and the Ethernet produced by the Xerox Palo Alto Research Center have been proposed and constructed to attack the problem. These systems are marginally applicable to the distribution of such capabilities over large distances. The interfacing of such local data networks to national and international networks represents one of the present edges of computer communications technology.

The Information Economy and Public Policy

Manley R. Irwin and Steven C. Johnson

If a single word summarizes U.S. telecommunication policy, that word is conflict. Telephone markets are vulnerable to competition from specialized carriers, the telephone industry is challenged by domestic satellites, telephone equipment is experiencing increased competition from outside manufacturers; the computer and communications industries are engaged in skirmishes for the lucrative market of data processing.

And this is only the beginning. Both state agencies and the FCC have squared off in a quarrel over their respective regulatory jurisdictions; controversy has erupted as to the proper policy mix between competition and regulation; both government and private antitrust suits have challenged the structure of the telephone carriers; and legislation has been introduced declaring a moratorium on competition in telecommunication.

In this article we attempt to provide the framework in our search for answers to questions concerning the cause of the conflict, the forces that contribute to it, the policy alternatives that are open for future consideration, and their effects on service markets.

U.S. Telecommunications

Conflict surfaces when two concepts, two ideas, or two philosophies find themselves juxtaposed. The telephone carrier industry, its services, its equipment, its organization—indeed the premise of regulation—is today subject to scrutiny and reappraisal. What has been the basis of the U.S. telecommunications industry for the past decades? The answer must begin with an examination of the telephone industry.

Service. Telephone companies have traditionally regarded their mandate as rendering a universal service to all subscribers (*1*, p. 4). Such universality, plain old telephone service, presumed carrier ownership of transmission, switching, terminals, and local loop plant. Local exchange services were priced at flat rates within specific areas and tariffs between customer classes rested on the concept of service value (*1*). Long distance tariffs were determined by the variables of time and distance.

Telephone companies were assigned exclusive service areas. Costs were aggregated or spread among the various types of equipment, vintages of plant, and density of population areas. Inevitably, subscribers in some areas subsidized customers in others.

The industry was capital-intensive. The investment required to establish a large telephone system was deemed to erect prohibitive barriers to entry. The industry was also said to enjoy economies of scale or size that gave monopoly a social advantage over competition (*2*). The results saw the industry evolve into a private monopoly—a development that laid the foundation for government regulation.

Dr. Irwin is a professor of economics in the Whittemore School of Business and Economics, University of New Hampshire, Durham 03824. Mr. Johnson was formerly with the Federal Communications Commission, and is now a staff economist with the National Research Council, Washington, D.C. 20418.