

The Computer as a Communication Device

In a few years, men will be able to communicate more effectively through a machine than face to face.

That is a rather startling thing to say, but it is our conclusion. As if in confirmation of it, we participated a few weeks ago in a technical meeting held through a computer. In two days, the group accomplished with the aid of a computer what normally might have taken a week.

We shall talk more about the mechanics of the meeting later; it is sufficient to note here that we were all in the same room. But for all the communicating we did directly across that room, we could have been thousands of miles apart and communicated just as effectively-as people-over the distance.

Our emphasis on people is deliberate. A communications engineer thinks of communicating as transferring information from one point to another in codes and signals.

But to communicate is more than to send and to receive. Do two tape recorders communicate when they play to each other and record from each other? Not really-not in our sense. We believe that communicators have to do something nontrivial with the information they send and receive. And we believe that we are entering a technological age in which we will be able to interact with the richness of living information—not merely in the passive way that we have become accustomed to using books and libraries, but as active participants in an ongoing process, bringing something to it through our interaction with it, and not simply receiving something from it by our connection to it.

To the people who telephone an airline flight operations information service, the tape recorder that answers seems more than a passive depository. It is an often-updated model of a changing situation—a synthesis of information collected, analyzed, evaluated, and assembled to represent a situation or process in an organized way.

Still there is not much direct interaction with the airline information service; the tape recording is not changed by the customer's call. We want to emphasize something beyond its one-way transfer: the increasing significance of the jointly constructive, the mutually reinforcing aspect of communication—the part that transcends “now we both know a fact that only one of us knew before.” When minds interact, new ideas emerge. We want to talk about the creative aspect of communication.

Creative, interactive communication requires a plastic or moldable medium that can be modeled, a dynamic medium in which premises will flow into consequences, and above all a common medium that can be contributed to and experimented with by all.

Such a medium is at hand—the programmed digital computer. Its presence can change the nature and value of communication even more profoundly than did the printing press and the picture tube, for, as we shall show, a well-programmed computer can provide direct access both to informational resources and to the *processes* for making use of the resources,

Communication: a comparison of models

To understand how and why the computer can have such an effect on communication, we must examine the idea of modeling-in a computer and with the aid of a computer. For modeling, we believe, is basic and central to communication. Any communication between people about the same thing is a common revelatory experience about informational models of that thing. Each model is a conceptual structure of abstractions formulated initially in the mind of one of the persons who would communicate, and if the concepts in the mind of one would-be communicator are very different from those in the mind of another, there is no common model and no communication.

By far the most numerous, most sophisticated, and most important models are those that reside in men's minds. In richness, plasticity, facility, and economy, the mental model has no peer, but, in other respects, it has shortcomings. It will not stand still for careful study. It cannot be made to repeat a run. No one knows just how it works. It serves its owner's hopes more faithfully than it serves reason. It has access only to the information stored in one man's head. It can be observed and manipulated only by one person.

Society rightly distrusts the modeling done by a single mind. Society demands consensus, agreement, at least majority. Fundamentally, this amounts to the requirement that individual models be compared and brought into some degree of accord. The requirement is for communication, which we now define concisely as "cooperative modeling" —cooperation in the construction, maintenance, and use of a model.

How can we be sure that we are modeling cooperatively, that we are communicating, unless we can compare models?

When people communicate face to face, they externalize their models so they can be sure they are talking about the same thing. Even such a simple externalized model as a flow diagram or an outline—because it

can be seen by all the communicators—serves as a focus for discussion. It changes the nature of communication: When communicators have no such common framework, they merely make speeches *at* each other; but when they have a manipulable model before them, they utter a few words, point, sketch, nod, or object.


The dynamics of such communication are so model-centered as to suggest an important conclusion: Perhaps the reason present-day two-way telecommunication falls so far short of face-to-face communication is simply that it fails to provide facilities for externalizing models. Is it really seeing the expression in the other's eye that makes the face-to-face conference so much more productive than the telephone conference call, or is it being able to create and modify external models?

The project meeting as a model

In a technical project meeting, one can see going on, in fairly clear relief, the modeling process that we contend constitutes communication. Nearly every reader can recall a meeting held during the formulative phase of a project. Each member of the project brings to such a meeting a somewhat different mental model of the common undertaking—its purposes, its goals, its plans, its progress, and its status. Each of these models interrelates the past, present, and future states of affairs of (1) himself; (2) the group he represents; (3) his boss; (4) the project.

Many of the primary data the participants bring to the meeting are in undigested and uncorrelated form. To each participant, his own collections of data are interesting and important in and of themselves. And they are more than files of facts and recurring reports. They are strongly influenced by insight, subjective feelings, and educated guesses. Thus, each individual's data are reflected in his mental model. Getting his colleagues to incorporate his data into their models is the essence of the communications task.

Suppose you could see the models in the minds of two would-be communicators at this meeting. You could tell, by observing their models, whether or not communication was taking place. If, at the outset, their two models were similar in structure but different simply in the values of certain parameters, then communication would cause convergence toward a common pattern. That is the easiest and most frequent kind of communication.



When mental models are dissimilar, the achievement of communication might be signaled by changes in the structure of one of the models, or both of them.

If the two mental models were structurally dissimilar, then the achievement of communication would be signaled by structural changes in one of the models or in both of them. We might conclude that one of the communicating parties was having insights or trying out new hypotheses in order to begin to understand the other—or that both were restructuring their mental models to achieve commonality.

The meeting of many interacting minds is a more complicated process. Suggestions and recommendations may be elicited from all sides. The interplay may produce, not just a solution to a problem, but a new set of rules for solving problems. That, of course, is the essence of creative interaction. The process of maintaining a current model has within it a set of changing or changeable rules for the processing and disposition of information.

The project meeting we have just described is representative of a broad class of human endeavor which may be described as creative informational activity. Let us differentiate this from another class which we will call informational housekeeping. The latter is what computers today are used for in the main; they process payroll checks, keep track of bank balances, calculate

orbits of space vehicles, control repetitive machine processes, and maintain varieties of debit and credit lists. Mostly they have *not* been used to make coherent pictures of not well understood situations.

We referred earlier to a meeting in which the participants interacted with each other through a computer. That meeting was organized by Doug Engelbart of Stanford Research Institute and was actually a progress-review conference for a specific project. The subject under discussion was rich in detail and broad enough in scope that no one of the attendees, not even the host, could know all the information pertaining to this particular project.


Face to face through a computer

Tables were arranged to form a square work area with five on a side. The center of the area contained six television monitors which displayed the alphanumeric output of a computer located elsewhere in the building but remotely controlled from a keyboard and a set of electronic pointer controllers called “mice.” Any participant in the meeting could move a near-by mouse, and thus control the movements of a tracking pointer on the TV screen for all other participants to see.


Each person working on the project had prepared a topical outline of his particular presentation for the meeting, and his outline appeared on the screens as he talked—providing a broad view of his own model. Many of the outline statements contained the names of particular reference files which the speaker could recall from the computer to appear in detail on the screens, for, from the beginning of the project, its participants had put their work into the computer system’s files.

So the meeting began much like any other meeting in the sense that there was an overall list of agenda and that each speaker had brought with him (figuratively in his briefcase but really within the computer) the material he would be talking about.

The computer system was a significant aid in exploring the depth and breadth of the material. More detailed information could be displayed when facts had to be pinpointed; more global information could be displayed to answer questions of relevance and interrelationship. A future version of this system will make it possible for each participant, on his own TV screen, to thumb through the speaker’s files as the speaker talks—and thus check out incidental questions without interrupting the presentation for substantiation.



At a project meeting held through a computer, you can thumb through the speaker's primary data without interrupting him to substantiate or explain.



A communication system should make a positive contribution to the discovery and arousal of interests.

Obviously, collections of primary data can get too large to digest. There comes a time when the complexity of a communications process exceeds the available resources and the capability to cope with it; and at that point one has to simplify and draw conclusions.

It is frightening to realize how early and drastically one does simplify, how prematurely one does conclude, even when the stakes are high and when the transmission facilities and information resources are extraordinary. Deep modeling to communicate—to understand—requires a huge investment. Perhaps even governments cannot afford it yet.

But someday governments may not be able *not* to afford it. For, while we have been talking about the communication process as a cooperative modeling effort in a mutual environment, there is also an aspect of communication with or about an uncooperative opponent. As nearly as we can judge from reports of recent international crises, out of the hundreds of alternatives that confronted the decision makers at each decision point or ply in the “game,” on the average only a few, and never more than a few dozen could be considered, and only a few branches of the game could be explored deeper than two or three such plies before action had to be taken. Each side was busy trying to model what the other side might be up to—but modeling takes time, and the pressure of events forces simplification even when it is dangerous.

Whether we attempt to communicate across a division of interests, or whether we engage in a cooperative effort, it is clear that we need to be able to model faster and to greater depth. The importance of improving decision-making processes—not only in government, but throughout business and the professions—is so great as to warrant every effort.

The computer—switch or interactor?

As we see it, group decision-making is simply the active, executive, effect-producing aspect of the kind of communication we are discussing. We have commented that one must oversimplify. We have tried to say why one must oversimplify. But we should not oversimplify the main point of this article. We can say with genuine and strong conviction that a particular form of digital computer organization, with its programs and its data, constitutes the dynamic, moldable medium that can revolutionize the art of modeling and that in so doing can improve the effectiveness of communication among people so much as perhaps to revolutionize that also.

But we must associate with that statement at once the qualification that

the computer alone can make no contribution that will help us, and that the computer with the programs and the data that it has today can do little more than suggest a direction and provide a few germinal examples. Emphatically we do *not* say: “Buy a computer and your communication problems will be solved.”

What we do say is that we, together with many colleagues who have had the experience of working on-line and interactively with computers, have already sensed more responsiveness and facilitation and “power” than we had hoped for, considering the inappropriateness of present machines and the primitiveness of their software. Many of us are therefore confident (some of us to the point of religious zeal) that truly significant achievements, which will markedly improve our effectiveness in communication, now are on the horizon.

Many communications engineers, too, are presently excited about the application of digital computers to communication. However, the function they want computers to implement is the switching function. Computers will either switch the communication lines, connecting them together in required configurations, or switch (the technical term is “store and forward”) messages.

The switching function is important but it is not the one we have in mind when we say that the computer can revolutionize communication. We are stressing the modeling function, not the switching function. Until now, the communications engineer has not felt it within his province to facilitate the modeling function, to make an interactive, cooperative modeling facility. Information transmission and information processing have always been carried out separately and have become separately institutionalized. There are strong intellectual and social benefits to be realized by the melding of these two technologies. There are also, however, powerful legal and administrative obstacles in the way of any such melding.

Distributed intellectual resources

We have seen the beginnings of communication through a computer—communication among people at consoles located in the same room or on the same university campus or even at distantly separated laboratories of the same research and development organization. This kind of communication—through a single multiaccess computer with the aid of telephone lines—is beginning to foster cooperation and promote coherence more effectively than do present arrangements for sharing computer programs by exchanging

magnetic tapes by messenger or mail. Computer programs are very important because they transcend mere “data”—they include procedures and processes for structuring and manipulating data. These are the main resources we can now concentrate and share with the aid of the tools and techniques of computers and communication, but they are only a part of the whole that we can learn to concentrate and share. The whole includes raw data, digested data, data about the location of data—and documents —and most especially models.

To appreciate the importance of the new computer-aided communication can have, one must consider the dynamics of “critical mass,” as it applies to cooperation in creative endeavor. Take any problem worthy of the name, and you find only a few people who can contribute effectively to its solution. Those people must be brought into close intellectual partnership so that their ideas can come into contact with one another. But bring these people together physically in one place to form a team, and you have trouble, for the most creative people are often not the best team players, and there are not enough top positions in a single organization to keep them all happy. Let them go their separate ways, and each creates his own empire, large or small, and devotes more time to the role of emperor than to the role of problem solver. The principals still get together at meetings. They still visit one another. But the time scale of their communication stretches out, and the correlations among mental models degenerate between meetings so that it may take a year to do a week’s communicating. There has to be some way of facilitating communication among people without bringing them together in one place.

A single multiaccess computer would fill the bill if expense were no object, but there is no way, with a single computer and individual communication lines to several geographically separated consoles, to avoid paying an unwarrantedly large bill for transmission. Part of the economic difficulty lies in our present communications system. When a computer is used interactively from a typewriter console, the signals transmitted between the console and the computer are intermittent and not very frequent. They do not require continuous access to a telephone channel; a good part of the time they do not even require the full information rate of such a channel. The difficulty is that the common carriers do not provide the kind of service one would like to have—a service that would let one have ad lib access to a channel for short intervals and not be charged when one is not using the channel.

It seems likely that a store-and-forward (i.e., store-for-just-a-moment-

and-forward-right-away) message service would be best for this purpose, whereas the common carriers offer, instead, service that sets up a channel for one's individual use for a period not shorter than one minute.

The problem is further complicated because interaction with a computer via a fast and flexible graphic display, which is for most purposes far superior to interaction through a slow-printing typewriter, requires markedly higher information rates. Not necessarily more information, but the same amount in faster bursts—more difficult to handle efficiently with the conventional common-carrier facilities.

It is perhaps not surprising that there are incompatibilities between the requirements of computer systems and the services supplied by the common carriers, for most of the common-carrier services were developed in support of voice rather than digital communication. Nevertheless, the incompatibilities are frustrating. It appears that the best and quickest way to overcome them—and to move forward the development of interactive *communities* of geographically separated people—is to set up an experimental network of multiaccess computers. Computers would concentrate and interleave the concurrent, intermittent messages of many users and their programs so as to utilize wide-band transmission channels continuously and efficiently, with marked reduction in overall cost.

Computer and information networks

The concept of computers connected to computers is not new. Computer manufacturers have successfully installed and maintained interconnected computers for some years now. But the computers in most instances are from families of machines compatible in both software and hardware, and they are in the same location. More important, the interconnected computers are not interactive, general-purpose, multiaccess machines of the type described by David [1] and Licklider [2]. Although more interactive multiaccess computer systems are being delivered now, and although more groups plan to be using these systems within the next year, there are at present perhaps only as few as half a dozen interactive multiaccess computer *communities*.

These communities are socio-technical pioneers, in several ways out ahead of the rest of the computer world: What makes them so? First, some of their members are computer scientists and engineers who understand the concept of man-computer interaction and the technology of interactive multiaccess systems. Second, others of their members are creative people in other fields

and disciplines who recognize the usefulness and who sense the impact of interactive multiaccess computing upon their work. Third, the communities have large multiaccess computers and have learned to use them. And, fourth, their efforts are regenerative.

In the half-dozen communities, the computer systems research and development and the development of substantive applications mutually support each other. They are producing large and growing resources of programs, data, and know-how. But we have seen only the beginning. There is much more programming and data collection—and much more learning how to cooperate—to be done before the full potential of the concept can be realized.

Obviously, multiaccess systems must be developed interactively. The systems being built must remain flexible and open-ended throughout the process of development, which is evolutionary.

Such systems cannot be developed in small ways on small machines. They require large, multiaccess computers, which are necessarily complex. Indeed, the sonic barrier in the development of such systems is complexity.

These new computer systems we are describing differ from other computer systems advertised with the same labels: interactive, time-sharing, multiaccess. They differ by having a greater degree of open-endedness, by rendering more services, and above all by providing facilities that foster a working sense of community among their users. The commercially available time-sharing services do not yet offer the power and flexibility of software resources—the “general purposeness”—of the interactive multiaccess systems of the System Development Corporation in Santa Monica, the University of California at Berkeley, Massachusetts Institute of Technology in Cambridge and Lexington, Mass.—which have been collectively serving about a thousand people for several years.


The thousand people include many of the leaders of the ongoing revolution in the computer world. For over a year they have been preparing for the transition to a radically new organization of hardware and software, designed to support many more simultaneous users than the current systems, and to offer them—through new languages, new file-handling systems, and new graphic displays—the fast, smooth interaction required for truly effective man-computer partnership.

Experience has shown the importance of making the response time short and the conversation free and easy. We think those attributes will be almost as important for a network of computers as for a single computer.

Today the on-line communities are separated from one another function-

ally as well as geographically. Each member can look only to the processing, storage and software capability of the facility upon which his community is centered. But now the move is on to interconnect the separate communities and thereby transform them into, let us call it, a supercommunity. The hope is that interconnection will make available to all the members of all the communities the programs and data resources of the entire supercommunity. First, let us indicate how these communities can be interconnected; then we shall describe one hypothetical person's interaction with this network, of interconnected computers.

What's in a Node?




* Operating System
Graphic display
Interpreter
User Programs
Files, etc.

Message processing

The hardware of a multiaccess computer system includes one or more central processors, several kinds of memory—core, disks, drums, and tapes—and many consoles for the simultaneous on-line users. Different users can work simultaneously on diverse tasks. The software of such a system includes supervisory programs (which control the whole operation), system programs for interpretation of the user's commands, the handling of his files, and graphical or alphanumeric display of information to him (which permit people not skilled in the machine's language to use the system effectively), and programs and data created by the users themselves. The collection of people, hardware, and software—the multiaccess computer together with its local community of users—will become a node in a geographically distributed computer network. Let us assume for a moment that such a network has been formed.

For each node there is a small, general-purpose computer which we shall call a "message processor." The message processors of all the nodes are interconnected to form a fast store-and-forward network. The large multiaccess computer at each node is connected directly to the message processor there. Through the network of message processors, therefore, all the large computers can communicate with one another. And through them, all the members of the supercommunity can communicate—with other people, with programs, with data, or with selected combinations of those resources. The message processors, being all alike, introduce an element of uniformity into an otherwise grossly nonuniform situation, for they facilitate both hardware and software compatibility among diverse and poorly compatible computers. The links among the message processors are transmission and high-speed digital switching facilities provided by common carrier. This allows the linking of the message processors to be reconfigured in response to demand.




All nodes can be interconnected via their message processors

A message can be thought of as a short sequence of "bits" flowing through the network from one multiaccess computer to another. It consists of two types of information: control and data. Control information guides the transmission of data from source to destination. In present transmission systems, errors are too frequent for many computer applications. However, through the use of error detection and correction or retransmission procedures in the message processors, messages can be delivered to their destinations intact even though many of their "bits" were mutilated at one point or another along the way. In short, the message processors function in the system as traffic directors, controllers, and correctors.


Today, programs created at one installation on a given manufacturer's computer are generally not of much value to users of a different manufacturer's computer at another installation. After learning (with difficulty) of a distant program's existence, one has to get it, understand it, and recode it for his own computer. The cost is comparable to the cost of preparing a new program from scratch, which is, in fact, what most programmers usually do. On a national scale, the annual cost is enormous. Within a network of interactive, multiaccess computer systems, on the other hand, a person at one node will have access to programs running at other nodes, even though those programs were written in different languages for different computers.

The feasibility of using programs at remote locations has been shown by the successful linking of the AN/FSQ-32 computer at Systems Development Corporation in Santa Monica, Calif., with the TX-2 computer across the continent at the Lincoln Laboratory in Lexington, Mass. A person at a TX-2 graphic console can make use of a unique list-processing program at SDC, which would be prohibitively expensive to translate for use on the TX-2. A network of 14 such diverse computers, all of which will be capable of sharing one another's resources, is now being planned by the Defense Department's Advanced Research Projects Agency, and its contractors.


The system's way of managing data is crucial to the user who works in interaction with many other people. It should put generally useful data, if not subject to control of access, into public files. Each user, however, should have complete control over his personal files. He should define and distribute the "keys" to each such file, exercising his option to exclude all others from any kind of access to it; or to permit anyone to "read" but not modify or execute it; or to permit selected individuals or groups to execute but not read it; and so on—with as much detailed specification or as much aggregation as he likes. The system should provide for group and organizational files within its overall information base.



One Message processor can be the messenger between two other message processors



A given pair of nodes may exchange several independent messages for simultaneous users with different interests




Interactive communication consists of short spurts of dialog

At least one of the new multiaccess systems will exhibit such features. In several of the research centers we have mentioned, security and privacy of information are subjects of active concern; they are beginning to get the attention they deserve.

In a multiaccess system, the number of consoles permitted to use the computer simultaneously depends upon the load placed on the computer by the users' jobs, and may be varied automatically as the load changes. Large general-purpose multiaccess systems operating today can typically support 20 to 30 simultaneous users. Some of these users may work with low-level "assembly" languages while others use higher-level "compiler" or "interpreter" languages. Concurrently, others may use data management and graphical systems. And so on.

But back to our hypothetical user. He seats himself at his console, which may be a terminal keyboard plus a relatively slow printer, a sophisticated graphical console, or any one of several intermediate devices. He dials his local computer and "logs in" by presenting his name, problem number, and password to the monitor program. He calls for either a public program, one of his own programs, or a colleague's program that he has permission to use. The monitor links him to it, and he then communicates with that program.

When the user (or the program) needs service from a program at another



. . . filibustering destroys communication.

node in the network, he (or it) requests the service by specifying the location of the appropriate computer and the identity of the program required. If necessary, he uses computerized directories to determine those data. The request is translated by one or more of the message processors into the precise language required by the remote computer's monitor. Now the user (or his local program) and the remote program can interchange information. When the information transfer is complete, the user (or his local program) dismisses the remote computer, again with the aid of the message processors. In a commercial system, the remote processor would at this point record cost information for use in billing.

Who can afford it?

The mention of billing brings up an important matter. Computers and long-distance calls have "expensive" images. One of the standard reactions to the idea of "on-line communities" is: "It sounds great, but who can afford it?"

In considering that question, let us do a little arithmetic. The main elements of the cost of computer-facilitated communication, over and above the salaries of the communicators, are the cost of the consoles, processing, storage, transmission, and supporting software. In each category, there is a

wide range of possible costs, depending in part upon the sophistication of the equipment, programs, or services employed and in part upon whether they are custom-made or mass-produced.

Making rough estimates of the hourly component costs per user, we arrived at the following: \$1 for a console, \$5 for one man's share of the services of a processor, 70 cents for storage, \$3 for transmission via line leased from a common carrier, and \$1 for software support—a total cost of just less than \$11 per communicator hour.

The only obviously untenable assumption underlying that result, we believe, is the assumption that one's console and the personal files would be used 160 hours per month. All the other items are assumed to be shared with others, and experience indicates that time-sharing leads on the average to somewhat greater utilization than the 160 hours per month that we assumed. Note, however, that the console and the personal files are items used also in individual problem solving and decision making. Surely those activities, taken together with communication, would occupy at least 25% of the working hours of the on-line executive, scientist or engineer. If we cut the duty factor of the console and files to one quarter of 160 hours per month, the estimated total cost comes to \$16 per hour.

Let us assume that our \$16/hr interactive computer link is set up between Boston, Mass., and Washington, D.C. Is \$16/hr affordable? Compare it first with the cost of ordinary telephone communication: Even if you take advantage of the lower charge per minute for long calls, it is less than the daytime direct-dial station-to-station toll. Compare it with the cost of travel: If one flies from Boston to Washington in the morning and back in the evening, he can have eight working hours in the capital city in return for about \$64 in air and taxi fares plus the spending of four of his early morning and evening hours en route. If those four hours are worth \$16 each, then the bill for the eight hours in Washington is \$128—again \$16 per hour. Or look at it still another way: If computer-aided communication doubled the effectiveness of a man paid \$16 per hour then, according to our estimate, it would be worth what it cost if it could be bought right now. Thus we have some basis for arguing that computer-aided communication is economically feasible. But we must admit that the figure of \$16 per hour sounds high, and we do not want to let our discussion depend upon it.

Fortunately, we do not have to, for the system we envision cannot be bought at this moment. The time scale provides a basis for genuine optimism about the cost picture. It will take two years, at least, to bring the first interactive computer networks up to a significant level of experimental

activity. Operational systems might reach critical size in as little as six years if everyone got onto the bandwagon, but there is little point in making cost estimates for a nearer date. So let us take six years as the target.

In the computer field, the cost of a unit of processing and the cost of a unit of storage have been dropping for two decades at the rate of 50% or more every two years. In six years, there is time for at least three such drops, which cut a dollar down to 12 1/2 cents. Three halvings would take the cost of processing, now \$5 per hour on our assumptions, down to less than 65 cents per hour.

Such advances in capability, accompanied by reduction in cost, lead us to expect that computer facilitation will be affordable before many people are ready to take advantage of it. The only areas that cause us concern are consoles and transmission.

In the console field, there is plenty of competition; many firms have entered the console sweepstakes, and more are entering every month. Lack of competition is not the problem. The problem is the problem of the chicken and the egg—in the factory and in the market. If a few companies would take the plunge into mass manufacture, then the cost of a satisfactory console would drop enough to open up a mass market. If large on-line communities were already in being, their mass market would attract mass manufacture. But at present there is neither mass manufacture nor a mass market, and consequently there is no low-cost console suitable for interactive on-line communication.

In the field of transmission, the difficulty may be lack of competition. At any rate, the cost of transmission is not falling nearly as fast as the cost of processing and storage. Nor is it falling nearly as fast as we think it should fall. Even the advent of satellites has affected the cost picture by less than a factor of two. That fact does not cause immediate distress because (unless the distance is very great) transmission cost is not now the dominant cost. But, at the rate things are going, in six years it will be the dominant cost. That prospect concerns us greatly and is the strongest damper to our hopes for near-term realization of operationally significant interactive networks and significant on-line communities.

On-line interactive communities

But let us be optimistic. What will on-line interactive communities be like? In most fields they will consist of geographically separated members, sometimes grouped in small clusters and sometimes working individually. They

will be communities not of common location, but of *common interest*. In each field, the overall community of interest will be large enough to support a comprehensive system of field-oriented programs and data.

In each geographical sector, the total number of users—summed over all the fields of interest—will be large enough to support extensive general-purpose information processing and storage facilities. All of these will be interconnected by telecommunications channels. The whole will constitute a labile network of networks—ever-changing in both content and configuration.

What will go on inside? Eventually, every informational transaction of sufficient consequence to warrant the cost. Each secretary's typewriter, each data-gathering instrument, conceivably each dictation microphone, will feed into the network.

You will not send a letter or a telegram; you will simply identify the people whose files should be linked to yours and the parts to which they should be linked—and perhaps specify a coefficient of urgency. You will seldom make a telephone call; you will ask the network to link your consoles together,

You will seldom make a purely business trip, because linking consoles will be so much more efficient. When you do visit another person with the object of intellectual communication, you and he will sit at a two-place console and interact as much through it as face to face. If our extrapolation from Doug Engelbart's meeting proves correct, you will spend much more time in computer-facilitated teleconferences and much less en route to meetings.

A very important part of each man's interaction with his on-line community will be mediated by his OLIVER. The acronym OLIVER honors Oliver Selfridge, originator of the concept. An OLIVER is, or will be when there is one, an "on-line interactive vicarious expediter and responder," a complex of computer programs and data that resides within the network and acts on behalf of its principal, taking care of many minor matters that do not require his personal attention and buffering him from the demanding world. "You are describing a secretary," you will say. But no! Secretaries will have OLIVERS.

At your command, your OLIVER will take notes (or refrain from taking notes) on what you do, what you read, what you buy and where you buy it. It will know who your friends are, your mere acquaintances. It will know your value structure, who is prestigious in your eyes, for whom you will do what



Your computer will know who is prestigious in your eyes and buffer you from a demanding world.

with what priority, and who can have access to which of your personal files. It will know your organization's rules pertaining to proprietary information and the government's rules relating to security classification.

Some parts of your OLIVER program will be common with parts of other people's OLIVERS; other parts will be custom-made for you, or by you, or will have developed idiosyncrasies through "learning" based on its experience in your service.

Available within the network will be functions and services to which you subscribe on a regular basis and others that you call for when you need them. In the former group will be investment guidance, tax counseling, selective dissemination of information in your field of specialization, announcement of cultural, sport, and entertainment events that fit your interests, etc. In the latter group will be dictionaries, encyclopedias, indexes, catalogues, editing programs, teaching programs, testing programs, programming systems, data bases, and—most important—communication, display, and modeling programs.

All these will be—at some late date in the history of networking—systematized and coherent; you will be able to get along in one basic language up to the point at which you choose a specialized language for its power or terseness.

When people do their informational work “at the console” and “through the network,” telecommunication will be as natural an extension of individual work as face-to-face communication is now. The impact of that fact, and of the marked facilitation of the communicative process, will be very great—both on the individual and on society.

First, life will be happier for the on-line individual because the people with whom one interacts most strongly will be selected more by commonality of interests and goals than by accidents of proximity. Second, communication will be more effective and productive, and therefore more enjoyable. Third, much communication and interaction will be with programs and programmed models, which will be (a) highly responsive, (b) supplementary to one’s own capabilities, rather than competitive, and (c) capable of representing progressively more complex ideas without necessarily displaying all the levels of their structure at the same time—and which will therefore be both challenging and rewarding. And, fourth, there will be plenty of opportunity for everyone (who can afford a console) to find his calling, for the whole world of information, with all its fields and disciplines, will be open to him—with programs ready to guide him or to help him explore.

For the society, the impact will be good or bad, depending mainly on the question: Will “to be on line” be a privilege or a right? If only a favored segment of the population gets a chance to enjoy the advantage of “intelligence amplification,” the network may exaggerate the discontinuity in the spectrum of intellectual opportunity.

On the other hand, if the network idea should prove to do for education what a few have envisioned in hope, if not in concrete detailed plan, and if all minds should prove to be responsive, surely the boon to humankind would be beyond measure.

Unemployment would disappear from the face of the earth forever, for consider the magnitude of the task of adapting the network’s software to all the new generations of computer, coming closer and closer upon the heels of their predecessors until the entire population of the world is caught up in an infinite crescendo of on-line interactive debugging.