

A Forward Look
Dr. Lawrence G. Roberts
Director
Information Processing Techniques
Advanced Research Projects Agency

COMPUTER-TO-COMPUTER data communications today might be compared with people-to-people communication via telegraph before the day of the telephone. Sending a message by telegraph was so slow that the media could only be used for non-interactive transmission of essential information. As such, its use was limited. The telephone provided an ability for people to interact, thus permitting a whole new range of applications.

Considering people somewhat mechanistically, we might view their use of the telephone as inter-human resource sharing. To solve a problem a man will call those people who have bits of data which he needs and will call on specialists for opinions, thus making use of other human resources. This is achieved because the media is appropriately responsive for human requirements and permits interactive conversation, thus eliminating the need for transmitting excessive detail, much of which may be unnecessary. Also, with an interactive dialog, information does not need to be formatted in a standard way since details can always be clarified if misunderstood. This increase in utility and the many new applications thereby permitted have resulted, as we know, in a vast increase in telephone traffic volume over telegraph traffic levels.

Communication between computers would most likely be affected in an analogous manner if a data communication system were made available which matched the needs of computers as well as the phone matches the needs of humans. Such a system would, of course, have to have different technical parameters such as connection time, data rates and reliability, than those required for voice communication, but if it permitted truly interactive conversations between a large ensemble of computers, the effect should be much the same, in that it would permit remote access to specialized hardware and software resources, joint problem solving and the dynamic retrieval of data from remote files. The analogy with the telephone is just one way of examining the potential impact of substantially improved data communication between computers and the resultant increase in applications and traffic which such a change might bring about.

Intercomputer communication has many quite substantial differences from interpersonal voice communication. Whereas voice conversation is a rather continuous, constant data rate process, communication with computers, either from computer consoles or other computers, requires a burst transmission rate several orders of magnitude higher than the average rate, even during a single conversation. Since there has been very little experience so far with real intercomputer traffic where two programs are talking to each other, it is useful to examine the characteristics of computer console traffic which is both a component of, and is also likely to have the same general parameters as, computer-to-computer traffic.

From statistics on teletypes, graphic consoles and remote batch stations, it appears that the ratio of burst rate to average rate is approximately 100 to one. This means that if a standard communications line is established for a computer conversation, the average utilization of that line will only be about 1 percent and therefore the cost will be ten to 100 times higher than the raw cost of moving the bits. A second characteristic of computer-to-computer communications is that the connect time to establish a conversation must be short enough that the computers or the computer users are not held up unduly when the need to access a special resource is determined. For computers the "connect time" should be considerably less than a second as opposed to the 20 or 30 seconds commonly experienced for voice communications. Third, the maximum data rate required in man-machine interaction must be considered. It is known that for useful comprehension by a human, the peak data rate for graphical material is on the order of 20 Kilobits per second, which suggests the required bandwidth for console-

A Forward Look
Dr. Lawrence G. Roberts
Director
Information Processing Techniques
Advanced Research Projects Agency

to-computer communications. This also suggests at least a minimum for computer-to-computer communications.

Finally, the error rate for intercomputer traffic must be far lower than required for voice communications or computer console traffic since there is usually very little, if any, redundancy inherent in the data. For many applications the error rate must be less than one in 10¹² bits. At the same time, the reliability (up time) of the data communications system must be very high if the user is to depend on remote resources. The cost of a data service providing the characteristics outlined above must be compared with the cost of duplicating the computer resources involved. Very simply, if the monthly cost of adequate communications service exceeds or even approaches the cost of a reasonably well endowed computer installation, it is not likely to be economical to use that communications service rather than duplicate that facility. Arbitrarily setting a threshold at 20 percent of a computer facility cost, it can be predicted that the communications system should not cost more than \$10- K per month per node.

A few years ago no communications system in existence even came close to providing the type of service just described. Therefore, ARPA undertook to develop such a capability so as to make resource sharing between computers possible. The communications system which resulted is utilized in the ARPANET and currently interconnects more than 20 computers at 15 locations around the country. By early 1972, expansion to 25 locations is expected (Figure 1). A delay- engineered message switching system, the ARPANET consists of Interface Message Processors (IMPs) at each node intercommunicating over 50 kilobit-leased communication lines and connected to one or more Host computers at each site. The IMP accepts messages from the Host, breaks them into thousand bit packets and sends each packet toward the destination over whichever communication line is currently optimal. Each IMP in turn checks the error detection code on the packet and, if it checks, routes the packet on to the next node and sends an acknowledgement to the previous node. At the destination, packets are assembled back into a message and delivered to the Host.

In practice, this organization proves to be extremely responsive, delivering short messages anywhere in the country within .1 second and permitting throughput rates for long messages of up to 80 kilobits per second. By adjusting the number of communication lines which are leased, the network can be engineered to have almost any desired overall average capacity between 2 kilobits per node and 60 kilobits per node. Since each communication line is being used for traffic between many pairs of nodes simultaneously, it can be loaded quite efficiently even though the individual Host-to- Host conversations have such a high ratio of burst rate to average rate.

A Forward Look
Dr. Lawrence G. Roberts
Director
Information Processing Techniques
Advanced Research Projects Agency

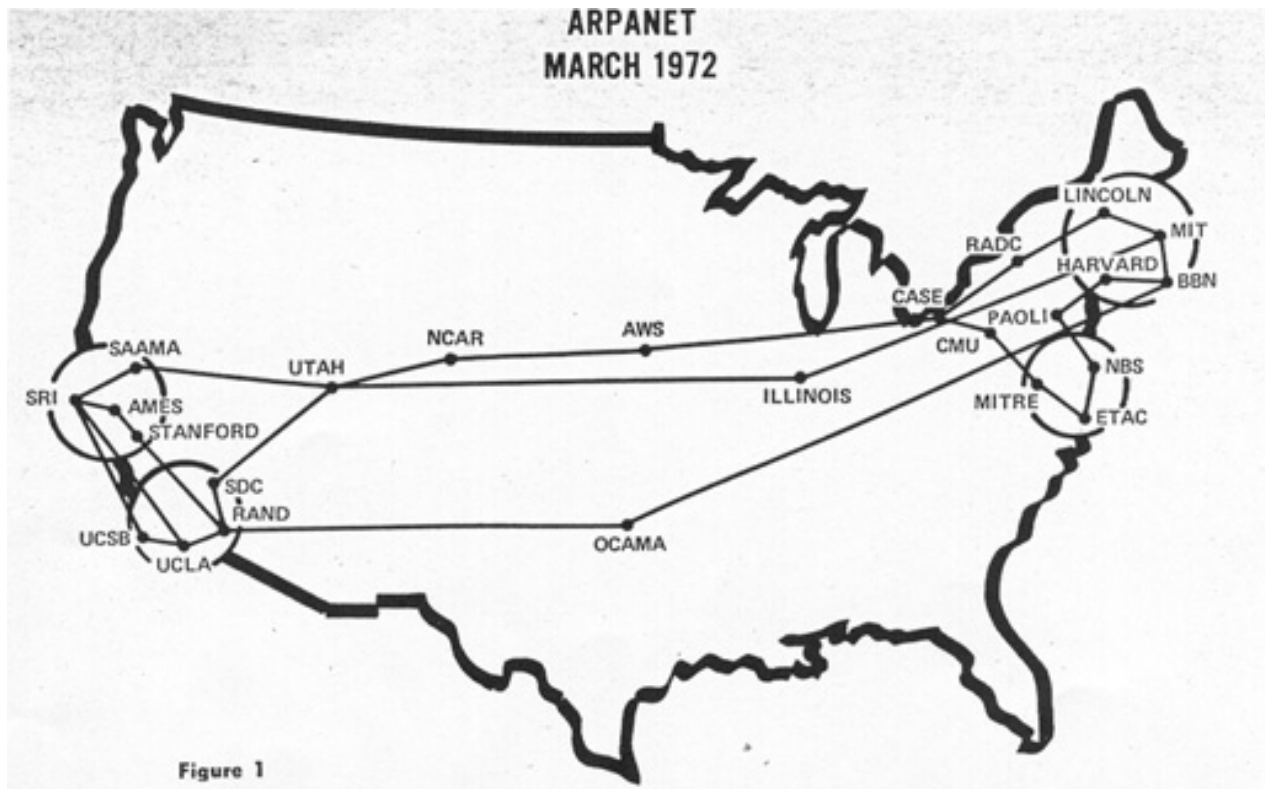


Figure 1

The actual cost of the total network communications system including the cost of IMPs, maintenance and communication lines ranges from \$3K to \$6K per month per node, depending on the overall traffic levels and the facilities required at each node. For new people entering the network, the February 1972 network of 23 nodes is currently estimated to cost \$4.8K/node/month; \$3.1K for an equal share of the communication lines cost and \$1.7K for the lease of a minimal IMP.

If a user wishes to provide direct console access to the network, a Terminal Interface Processor (TIP) would be used. The TIP, which will become available in August 1971, will act both as an IMP and as a simple Host, permitting up to 64 consoles and peripheral devices to intercommunicate with any Host in the network at rates up to 20 kilobits/ sec. Thus the TIP expands the network concept to include nodes without an interactive host of their own, but who wish high bandwidth support for graphic consoles, printers, and large collections of lower speed devices. Use of a TIP increases the cost by \$1.6K/month.

Although an equal share of the communication line cost is currently allocated to each node, this policy will be changed, as soon as feasible, to one of charging only for the bits actually sent from each node. Referring to Figure 2, it can be seen that the cost of the network increases almost linearly with capacity, at least for bandwidths below 16 KB/node. Also, it turns out that the capacity and cost of these distributed networks are remarkably insensitive to the distribution and destination of traffic, the total traffic being the only important parameter. Thus, it is appropriate to charge for traffic initiated at a node, based on the cost of increasing the total capacity of the network by that amount. From Figure 2 it can be seen that this will be 11 ~ / megabit for the ARPANET. However, since the network cannot be expected to be always fully loaded to peak capacity, day and night, it is likely the actual rate will be

A Forward Look
Dr. Lawrence G. Roberts
Director
Information Processing Techniques
Advanced Research Projects Agency

30~/megabit based on an estimated 36 percent average loading. The total cost per node would then be \$1.7K/month plus 3Q~/megabit.

Looking ahead, assuming the broad availability of a data communications service similar to the ARPANET system, it is clear that very significant changes in computer system organization will take place. Some of these changes will occur rapidly-within the first five years- and others will take a decade or more before people fully accept the concepts. Soon after a network with a dozen or more reliable computer services becomes available, many institutions will find it far more economical to obtain their computing services from a selected set of these remote systems, rather than run their own computer center.

For example, take the case of an institution about ready to upgrade its facility. One choice would be to obtain a medium scale, general purpose batch system, admittedly a compromise for their large numerical users and time-sharing users, but the best single system that they can afford. Alternatively, they could buy no new machine and obtain access to several of the systems on the network through a Terminal Interface Processor. This approach permits their large numerical users to use a large "number-cruncher", their statistical and payroll users to access a large scale general purpose system, and their interactive users to have teletype or graphic console access to a good time-sharing system.

Overall, the cost of each service is less than it would have been on a dedicated G-P computer by factors between two and ten. Also, they can buy just the capacity they need and expand smoothly rather than having to pay for an oversize machine for a year or two. The peripherals cost the same in either case, and the network cost is negligible compared to the direct computer cost savings. As added benefits, the computer services they use are probably better and more reliable than they could hope to do themselves, since the services must stay competitive; a wider range of software is available and can be accessed directly without translation or transfer, and as new hardware is introduced which is economically useful, they can transfer jobs to it on a selected and leisurely basis. The direct use of distributed hardware services just described will probably account for most of the initial use of the network. This growth should proceed about uniformly over the next eight years- two computer replacement cycles. Some additional traffic will be introduced by the gradual transfer of current data traffic from other data communication networks to the computer network due to the economy or reliability, but the total quantity of this traffic is minor in comparison to the new traffic generated by the computer resource sharing activity.

A second major application of the computer network is data base sharing-direct retrieval from remote, one of a kind, data bases. Currently, when large data bases or files are needed at several computer centers, duplicates are maintained at each center. This difficult and costly practice can be avoided if the access speed through the network is fast enough so that neither human users nor computer processes are unduly delayed. The ARPANET response speed of one-tenth of a second for a question and three-tenths of a second for a one-page answer is quite acceptable for a human user and for a computer program it is no worse than a slow disc.

To start with, this response appears adequate; however, further experience may indicate a need for faster response in future networks. Data base sharing will not build up as rapidly as hardware service sharing, however, since it represents only an incremental saving for an installation and demands considerable faith in the network. Copying a 109 bit data base monthly might cost \$2,000-less than the minimum network cost and therefore not a prime motivation for joining the network. However, the cost of accessing the data base through the network would cost at most \$300 even if all the data were

A Forward Look
Dr. Lawrence G. Roberts
 Director
 Information Processing Techniques
 Advanced Research Projects Agency

required, providing a considerable cost saving and convenience as long as the network connection had other justification. Of course for very large data bases, such as the 1011 bit weather-climate data base being developed by the Air Weather Service for the ARPANET, the cost of either copying or storing a duplicate would immediately justify network connection.




Figure 2

A Forward Look
Dr. Lawrence G. Roberts
Director
Information Processing Techniques
Advanced Research Projects Agency

In most cases very large data bases would not be developed at all without a network making possible nation-wide access, since the cost would be prohibitive. Data base sharing, therefore, is not likely to grow rapidly until the network is reasonably well established, lagging the service-sharing growth by perhaps two years, but then growing exponentially as everyone requests access to all the information available.

Software sharing, the third major application, is the remote use of software subroutines and packages, programs not available on the users' primary computer due to incompatibility of hardware or languages. An example of this type of activity might be the use by MIT scientists of the Stanford Heuristic Dendral System, a program for determining molecular structure given the mass spectrum. On a computer at MIT the scientists would collect and preprocess the mass spectrum data. Then, much like using a subroutine, the Stanford computer would be called, the data sent and the molecular structure, when determined, sent back. If interaction were required, the MIT scientist would be interrogated much as if he were at Stanford, thus building up the heuristic model based on nation-wide inputs. The MIT computer, upon receiving the response, would proceed locally with the calculations or displays desired.

Software sharing like this will be required if we are to maintain maximum progress as the volume of useful software continues to expand. Since the annual cost for software is already larger than that for hardware and, to some extent, should be cumulative rather than wearing out, the long range importance of software sharing is clearly greater. However, due to human inertia and a strong "not invented here" syndrome associated with software, it is clear that the cross-utilization of software will take years to develop. The buildup of software sharing activity will most likely begin very slowly, growing exponentially, but not become a major factor until the network becomes well established in four to eight years.

Besides hardware, software, and data base sharing, there are many other important network applications, all of which require a large viable network before they become important in their own right. These include teleconferencing, publishing, library services and office paperwork filing and distribution. Ten to twenty years from now these applications may well dominate computer usage and network usage but they are not likely to be important factors for at least five years.

Overall, then, hardware service sharing is likely to be the major factor causing networks to come into existence, since the effective cost of computing can be drastically lowered for only a moderate communications cost. Then, data base sharing will become the dominant force expanding the traffic in three to four years. Software sharing, although very important in the long run, will not become a major factor for four to eight years. The text oriented services, libraries and office work, will then come into their own in ten to twenty years. The whole trend should decrease the importance of the general purpose computers as stand-alone systems, and substantially increase the importance of specialized systems-ones which can provide a specific service at the lowest cost.