Resource Sharing Computer Networks

Dr. Lawrence G. Roberts

Just as time-shared computer systems have permitted groups of hundreds of individual users to share hardware and software resources with one another, networks connecting dozens of such systems will permit resource sharing between thousands of users. Each system, by virtue of being time-shared, can offer any of its services to another computer system on demand. The most important criterion for the type of network interconnection desired is that any user or program on any of the networked computers can utilize any program or subsystem available on any other computer without having to modify the remote program.

The objective of this program is twofold:
to develop techniques and obtain experience on interconnecting computers in such a way that a very broad class of interactions are possible, and to improve and increase computer research productivity through resource sharing.

By establishing a network tying ARPA- sponsored computer research centers together, both goals are achieved. In fact, the most efficient way to develop the techniques needed for an effective network is by involving the research talent at these centers in prototype activity.

Scientific Environment

Currently there are thousands of computer centers in the country, each of which operates almost completely autonomously. There is some trading of programs between those machines, which are sufficiently similar to allow this, and there is technical communication through publications of technical meetings describing techniques developed. However, since the computer field is growing at such a rapid rate, a more immediate mechanism must be developed if there is to be significant cross-fertilization in sharing between these many centers. Although the same problem exists in many technological areas, the solution is most easily found and implemented by the computer community. If a sufficiently reliable and capable network were established linking these centers, many improvements could be obtained. There network were established linking these centers, many improvements could be obtained. There would be less duplication of large programs and systems, some of which require hundreds of man months of effort. Currently such programs must be reprogrammed for each machine where they are needed even if they are only required occasionally. It is estimated that such duplicative efforts more than double the national costs of creating and maintaining the software. A network will not eliminate all of this duplication but can be used for those functions which are only infrequently called and those which only need to be tested. Further, there are large data files available at individual locations which are not valuable enough to warrant duplication at every computer center, but from which segments could be obtained at any network location. For example, within the ARPA research centers there are files of speech samples, digitized pictures and the semantic definitions of most English words.

Often it is important at a research establishment to test a new language developed at another installation to determine what features should be incorporated into local languages. Currently one either reprograms the language on his local machine or obtains sufficient remote console time to evaluate the language. Although it may be preferable to use the original system via remote consoles, this is often difficult or impossible due to console incompatibility. With an interactive network it is possible to use ones local consoles through the local computer to access the remote system, thus eliminating the need for compatible consoles and at the same time reducing the communications costs by several orders of magnitude.

Another important application of a network is to link specialized computers to general purpose computer centers. ILLIAC IV is an outstanding example of such a specialized machine. With recent improvements in the hardware area, it will become more cost effective to design and construct computers particularly efficient at other specialized tasks (e. g. compiling, list processing and

information retrieval). Making such machines available to all the computer research establishments would significantly increase the capability at these other centers.

Military Environment

The military environment like the scientific environment, includes thousands computers of various vintages and vendors. The traditional staff elements (Personnel, Intelligence, Operations (Command and Control), Logistics and communications) throughout all the Services are various machines with varying degrees of success. With the current fractionation of computer resources in the absence of any technology permitting the interconnection and sharing of these resources, the current situation can only get worse. Those data files and programs which have common utility to many military organizations and installations must be stored, created and maintained separately at each different machine. Military systems interconnected in a distributed interactive network obviate such constraints.
Previous Work

Relatively little work has been done in the past on interactive computer networks and it is mainly with the advent of widespread time-sharing that such nets become feasible. Most previous work has concentrated on either load sharing or message handling goals. Several attempts at load s have been made including ones at Bell Telephone Lab and UCLA. The desire was to improve processor utilization through load equalization, out unfortunately the precise compatibility required Is almost impossible to maintain.

More recently, experiments have been carried out between Lincoln Laboratory and System Development Corporation to test the feasibility of more general computer-computer interaction. This experiment demonstrated the relative ease of modifying time-sharing systems to permit network interactions and provided some statistics on the message lengths encountered. This experience has been added to through the introduction of the 338 display a computer at ARPA tied to the Lincoln system. Although the requests in this communication link are totally in one direction, the form of communication utilized is identical with that expected in network activities and has extended the techniques to include graphic display interactions.

Preliminary Network Planning

In early 1967 preliminary plans for an interactive computer network were discussed with ARPA Information Processing Techniques contractors. Working groups were established to design standardized communications protocol and to specify their network requirements. A preliminary protocol was developed and discussed with interested parties during the summer of 1967.

Network Information Center

In order for people to utilize the envisioned computer network effectively, it was necessary to provide extremely good documentation on what programs and files are available throughout the net. This information should be available online to any individual in the network. It should be possible for him to add new program descriptions, edit previous descriptions, retrieve relevant information based on keyword searches and affix comments to program descriptions, which he has used. To achieve this goal, Stanford Research Institute has agreed to develop such a facility. This is an extension of the capability already achieved at SRI and is already in progress in order that it may become available concurrently with the network.

Communication System

Multi-point, fast response, high capacity, reliable communications are required for an interactive computer network. The traffic between nodes is expected to consist mainly of short digital messages with a wide dispersal of destinations. Initially, message length will vary from one to one thousand characters and with an expected average length of 20 characters. Since a cross country 50 kb

communication line has a delay equivalent to 150 characters, messages must be continuously multiplexed into each line in order to maintain reasonable efficiency. Since the dispersion of destinations is large, messages with different origins and destinations must be concentrated into the same line. This can only be achieved with a store and forward system.

Message delay for on-line interactive work should be well below one second (origin to destination). This cannot be achieved with voice grade communication lines in a store and forward system. However, with 50-kilobit communication lines, the required response speed can be attained. The additional capacity obtained with 50 kb lines is also important, but is not the prime factor dictating the choice of these lines.

After considering the trade-offs associated with the communications subsystem, it was decided to design and build a store and forward net using message processors at each research center interconnected with 50 kb communication lines. Such a distributed communication system will be revolutionary, providing vastly reduced transmission costs, fast response and high reliability. The effect of providing such an efficient communication capability to the computer researchers should be to inspire the development of creative and effective network techniques.